



UNIVERSITY OF LEEDS

This is a repository copy of *Towards a Computed-Aided Diagnosis System in Colonoscopy: Automatic Polyp Segmentation Using Convolution Neural Networks*.

White Rose Research Online URL for this paper:
<http://eprints.whiterose.ac.uk/158057/>

Version: Accepted Version

Article:

Brandao, P, Zisimopoulos, O, Mazomenos, E orcid.org/0000-0003-0357-5996 et al. (9 more authors) (2018) *Towards a Computed-Aided Diagnosis System in Colonoscopy: Automatic Polyp Segmentation Using Convolution Neural Networks*. *Journal of Medical Robotics Research*, 03 (02). 1840002. p. 1840002. ISSN 2424-905X

<https://doi.org/10.1142/s2424905x18400020>

Electronic version of an article published as *Journal of Medical Robotics Research*, 03, 02, 2018, 1840002, www.doi.org/10.1142/s2424905x18400020 © World Scientific Publishing Company <https://www.worldscientific.com/worldscinet/jmrr>

Reuse

Items deposited in White Rose Research Online are protected by copyright, with all rights reserved unless indicated otherwise. They may be downloaded and/or printed for private study, or other acts as permitted by national copyright laws. The publisher or other rights holders may allow further reproduction and re-use of the full text version. This is indicated by the licence information on the White Rose Research Online record for the item.

Takedown

If you consider content in White Rose Research Online to be in breach of UK law, please notify us by emailing eprints@whiterose.ac.uk including the URL of the record and the reason for the withdrawal request.



eprints@whiterose.ac.uk
<https://eprints.whiterose.ac.uk/>

Towards a computed-aided diagnosis system in colonoscopy: Automatic polyp segmentation using convolution neural networks

Patrick Brandao^a, Odysseas Zisimopoulos^a, Evangelos Mazomenos^a, Gastone Ciuti^b, Jorge Bernal^c, Marco Visentini-Scarzanella^d, Arianna Menciassi^b, Paolo Dario^b, Anastasios Koulaouzidis^e, Alberto Arezzo^f, David J Hawkes^a, and Danail Stoyanov^a

^a Centre for Medical Image Computing, University College London, London, UK
E-mail: patrick.brandao.15@ucl.ac.uk

^b The BioRobotics Institute, Scuola Superiore Sant'Anna, Pisa, Italy.

^c Computer Science Department at Universitat Autnoma de Barcelona, Barcelona, Spain

^d Multimedia Laboratory Toshiba Corporate Research and Development Center, Kawasaki, Japan.

^e UEndoscopy Unit, The Royal Infirmary of Edinburgh, Edinburgh, UK

^f Department of Surgical Sciences, University of Turin, Turin, Italy.

Early diagnosis is essential for the successful treatment of bowel cancers including colorectal cancer (CRC) and capsule endoscopic imaging with robotic actuation can be a valuable diagnostic tool when combined with automated image analysis. We present a deep learning rooted detection and segmentation framework for recognizing lesions in colonoscopy and capsule endoscopy images. We restructure established convolution architectures, such as VGG and ResNets, by converting them into fully-connected convolution networks (FCNs), fine-tune them and study their capabilities for polyp segmentation and detection. We additionally use Shape-from-Shading (SfS) to recover depth and provide a richer representation of the tissue's structure in colonoscopy images. Depth is incorporated into our network models as an additional input channel to the RGB information and we demonstrate that the resulting network yields improved performance. Our networks are tested on publicly available datasets and the most accurate segmentation model achieved a mean segmentation IU of 47.78% and 56.95% on the ETIS-Larib and CVC-Colon datasets, respectively. For polyp detection, the top performing models we propose surpass the current state of the art with detection recalls superior to 90% for all datasets tested. To our knowledge, we present the first work to use FCNs for polyp segmentation in addition to proposing a novel combination of SfS and RGB that boosts performance.

Keywords: convolutional neural networks; colonoscopy; computer aided diagnosis.

1. Introduction

Colorectal cancer (CRC) is the most frequent pathology of the gastrointestinal tract and accounts for nearly 10% of all forms of cancer. When the disease reaches an advanced stage, the 5-year survival rate of CRC patients is lower than 7%, but early diagnosis with successful treatment can dramatically increase this figure to more than 90%.¹ Conventional colonoscopy is the reference standard for CRC screening and diagnosis. It provides direct visualization of the inner surface of the colon for acquiring biopsies and performing therapeutic procedures on early stage neoplastic lesions. Wireless and tethered endoscopic capsules and robotic endoscopes have recently been reported

to overcome the pain and discomfort associated with conventional colonoscopy.¹ Despite assisted articulation or the use of wireless capsules, the success of the exam still mostly depends on the operator's skills both for dexterous maneuvering of the camera and for ensuring full exploration of the colon wall.¹ With wireless capsules the problem is more focused on the processing and analysis time of the long and often occluded video signal. As a result, even experienced endoscopists can miss polyps during examination.² To overcome this, computer-aided diagnostic (CAD) systems can be used to support clinical decisions with automated polyp detection algorithms.

Despite significant progress in recent years, CAD polyp detection is still far away from routine clinical use. The

problem is challenging because there is large variety in the size, shape, colour and textures of colon polyps, which alongside the presence of specular reflections, endoluminal folds and blood vessels in colonoscopy recordings, impedes detection accuracy and induces a significant number of false detections. Early CAD polyp detection methods used texture and colour as the main features.^{3,4} More recently, shape information has been incorporated to increase feature complexity and discriminating power and provide improved detection models.^{5,6}

We build onto recent advances in computer vision and incorporate shape information to build a system based on a convolutional neural network (CNN) framework that can accurately detect and segment polyps and high-risk regions in colonoscopy images, while minimizing false detections. We combine state-of-the-art deep learning techniques designed for generic object detection and adapt them to endoscopic images. To our knowledge, this is the first work to fine-tune cutting edge CNN architectures, such as VGG and ResNets, for polyp segmentation. To increase detection accuracy, we also extracted a relative measure of depth from the colonoscopy images, using Shape-from-shading (SfS), and input it as an additional feature to the networks to enhance polyp representation. Our study also includes investigation into the influence of batch normalization in training convergence.

2. Related work

Propelled by large scale datasets such as ImageNet⁷ and Pascal VOC,⁸ deep learning approaches are the current paradigm for many computer vision problems, often surpassing by some way traditional methods for classification, segmentation, detection and tracking. Since the first deep CNN won the Imagenet Large Scale Visual Recognition Competition (ILSVRC) classification challenge⁹ continuous breakthroughs in CNN architectures, training optimization and fast Graphic Processing Units (GPUs)¹⁰ have led to continuous performance improvements. Studying the influence of network depth in large-scale image recognition resulted in the powerful 16 and 19 layer VGG network.¹¹ In the same year, the concept of "inception architecture" was developed¹² and the GoogLeNet won the Imagenet challenge. The current leader in results successfully trained a network (ResNet) 8-times deeper than VGG by using a residual learning framework.¹³ Several CNNs^{16,37,38} were also used in semantic segmentation achieving impressive results. Along with training on RGB data, image depth has also been used to provide a richer scene representation, in detection and segmentation tasks using a conditional random field (CRF) model.¹⁴ Depth can also be used directly as a fourth input channel and achieve improved results comparing to RGB-only training¹⁵ and can improve the performance of the fully convolutional version of VGG by performing training with RGB-D images.^{16,39}

Large scale polyp datasets, such as CVC-ColonDB,¹⁸ that recently became publicly available, facilitated the de-

velopment of deep learning techniques for polyp detection in endoscopic images. Zhu *et al.* fed features extracted through a CNN, into a Support Vector Machine (SVM) to identify lesion areas in the colon.¹⁹ On the other hand, Tajbakhsh *et al.* used a variety of pre-calculated features as inputs in a set of CNNs to classify polyp candidates.²¹ However, these methods are tested in different datasets, which limits comparison between the reported performance. In 2015, the MICCAI sub-challenge on automatic polyp detection provided a set of guidelines and datasets that allowed direct comparison between different methods. The best results were achieved by models learned end-to-end. The OUS participation achieved a detection precision and recall of 69.7% and 63.0%, respectively, in the ETIS-Larib dataset by slightly modifying the AlexNet model. The CUMED submission used a fully convolution network (FCN), initialized from models trained in large scale datasets to achieve the best overall performance, with a 72.3% detection precision and 69.2% recall in the same dataset.²⁰

Despite the clear advantages that CNN offer, defining and training a CNN architecture from the beginning is not a trivial task. The training process typically requires an excessive amount of labeled data, something that is generally difficult to be obtained for medical images. In addition, selecting the architecture that provides the best compromise between convergence speed and inference ability is both challenging and time consuming. Finally, potential issues of convergence and overfitting often require significant tuning of the learning parameters.²² Alternatively, it has been suggested²² that fine-tuning deep-learned architectures for specific tasks (e.g. image detection and segmentation), provides better performance than designing and training a new network model. Such an approach can be applied in tasks involving medical images, despite being substantially different from the natural images used in training the initial network, by using a smaller set of images to re-train (fine-tune) a pre-existing network. We adhere to this strategy and present a number of established CNNs which after appropriate fine-tuning are optimized for detecting and segmenting polyps in colonoscopy recordings.

We are expanding on the work developed in³¹, where we segmented polyps using the same principles used by Long *et al.*¹⁶ to semantically segment objects in natural scene images. First, we exploit the use of residual networks¹³ which significantly outperformed previous models in large scale recognition challenges. We also evaluate the effect of using depth estimations as a extra input feature for the best models. Finally, to guarantee generality, all models are evaluated in an extra publicly available dataset, the CVC-Colon.²⁶

3. Proposed Methods

We develop polyp segmentation CNN models, by utilizing well-known architectures initialized with weights obtained from pre-training, in large image-datasets and then fine-tune them with labeled colonoscopy images from publicly

available databases (ETIS-Larib, CVC-ColonDB etc). The models we explore have optimized architectures and learned feature extractors that are capable of solving complex object detection tasks. For performing polyp segmentation instead of simple detection, we opted to convert the pre-trained CNNs into FCNs by adding deconvolution layers in order to obtain a pixel classification map as the output. Moreover, we applied an efficient SfS technique to recover depth from colonoscopy images and provide it as an additional, to the standard RGB channels, parameter in the formulation of our network models.

3.1. CNN and FCN Basis

Irrespective of the architecture, CNNs always integrate three basic components: convolution, activation function and pooling operation layers. They operate on local inputs, depending only on relative spatial coordinates. An example of the effect of convolution and activation on a colonoscopy image is depicted in Fig. 1.




Fig. 1. The basic CNN operations on a single CNN neuron from the first layer of the FCN-VGG with batch normalization. Image sequence left to right: input image, receptive field, convolution results, normalized image, and ReLU activated image.

Considering $x_{i,j}^k$ as the input data vector at location (i, j) in layer k , the input of the following layer, $x_{i,j}^{k+1}$, is computed by

$$x_{i,j}^{k+1} = f_S^k(x_{si+\delta_i, sj+\delta_j}^k, 0 \leq \delta_i, \delta_j \leq w). \quad (1)$$

where S is the stride or subsample factor and f^k represents the type of operation of the layer k . In classification CNNs, the network ends with one or more fully connected layers that produce non-spatial outputs.¹⁶ A loss function l compares the prediction outputs of the last layer f^K to the desired result y as:

$$l(x^K, y) = \sum_{i,j} l(x_{i,j}^K, y_{i,j}). \quad (2)$$

Using the chain rule, the gradient of the loss is back-propagated throughout the network and the parameters of all layers are updated via stochastic gradient descent (SGD).²³

Traditional CNN architectures, such as AlexNet and VGG, are used for classification problems, which mean that take an input image and output a single classification score for all the possible classes. To obtain pixel wise segmentation, these networks need to be converted to a fully convolutional network. A fully connected layer can be viewed as a convolution layer where the kernel has the same dimensions as the input. By replacing these with convolutions, it is possible to convert traditional classification networks into FCNs that take inputs of any size and output coarse classification maps. While the resulting maps are equivalent to the processing of the individual patches on the original network, the computational cost is highly amortized by the inherent efficiency of convolution. Even though the output maps can yield any size, these are typically reduced by subsampling within the network.¹⁶ To connect these coarse outputs to dense pixels, an interpolation strategy needs to be used.

Convolution is a linear operation, and as such, it can be expressed in a matrix multiplication form. Assuming Ω as the map of size $W \times H$ to be convoluted by the kernel θ of size $W' \times H'$ with a stride S , the convolution operation can be expressed as

$$vec(\psi) = Cvec(\Omega). \quad (3)$$

where $vec(\Omega)$ represents Ω flattened to a WH dimensional vector, $vec(\psi)$ is a vector with size $D = (\frac{W-W'}{S} + 1) \times (\frac{H-H'}{S} + 1)$ and C is a sparse matrix of size $D \times WH$, where the non-zero elements are elements of K . The vector $vec(\psi)$ can be later reshaped to a $(\frac{W-W'}{S} + 1) \times (\frac{H-H'}{S} + 1)$ convoluted map [14]. During CNN training, the loss ψ_l is backward passed to the lower level layers by convolution transpose

$$vec(\Omega_l) = C^T vec(\psi_l). \quad (4)$$

where Ω_l and ψ_l have the same dimensions of Ω and ψ in the forward pass, and Ω_l connectivity pattern is compatible with C by construction.²⁴

If $S > 1$, convolution implements a subsample operation. Intuitively, transpose convolution is a way to upsample the input by a factor of S . Following this principle, by simply reversing forward and backward pass operations, it is possible to implement in-network upsampling. The transpose convolution layer, also known as deconvolution layer, does not need to have a fixed filter (doing bilinear interpolation, for example) but can also be learned and adjusted during training. This provides very fast and effective upsampling used to structure efficient FCNs, capable of achieving state-of-the-art results in semantic segmentation.¹⁶

3.2. Batch normalization

SGD iteratively estimates the global gradient of the loss by using a limited set of samples. Changes in distribution of the inputs hinders convergence, as the parameters of each layer need to adapt to a new distribution. This slows down training by demanding lower learning rates and careful parameter initialization. In deep networks, this effect is intensified because small changes in the parameters are greatly amplified throughout the network, as the inputs of each layer are affected by parameters of all preceding layers.²⁵

To overcome this, we evaluate the incorporation of batch normalization into standard CNN architectures to compensate batch distribution changes. Batch normalization layers perform in-network normalization by linearly transforming each training mini-batch to have zero mean and unit variance. This technique has proved to yield improved results on classification tasks using considerable less training iterations.²⁵ An example of the batch normalization process is illustrated in Fig. 1.

3.3. Shape-from-Shading

The increased detection performance of CNNs by incorporating depth information, motivated us to employ a SfS technique²⁷ to extract depth from colonoscopy images and include it in the formulation of our models. SfS aims to recover the 3D shape of an object by analyzing the illumination variation across the image. Subsequently, SfS is suitable for approximating depth in colonoscopy recordings with a monocular view without requiring stereo or multi-view matching³⁴ and structure from motion estimation.³⁵ While limited to only relative depth, SfS does not require texture assumptions like shape-from-texture³⁶ techniques and is useful for extracting geometric information easily from existing clinical colonoscopy systems.

The majority of SfS approaches²⁷⁻²⁹ assume a light source either coinciding with the optical center or infinitely far away from the scene. These conditions are unrealistic in the case of colonoscopy even though the light source and the camera are both at the tip of the instrument. This is because despite the small distance between the camera and the light, the observed tissue is also very close and highly dependent on small illumination changes. To overcome this limitation, we use a method which approximates the position of the lightsource at the tip of the endoscope and uses the position directly in the SfS problem formulation.³⁰

Assuming the camera positioned at $\mathbf{y} = (a, b, c)$ and a set \mathcal{A} of all image points $\mathbf{x} = (x, y)$, the SfS problem is formulated as a Hamiltonian-Jacobi non-linear Partial Differential Equation (PDE),

$$H(\mathbf{x}, \nabla v) = \frac{1}{\rho} I(\mathbf{x}) \sqrt{(v_x^2 + v_y^2) + J(\mathbf{x}, \nabla v)^2} \cdot Q(\mathbf{x})^{\frac{3}{2}}, \quad (5)$$

where

$$Q(\mathbf{x}) = (x + a)^2 + (y + b)^2 + (f + c)^2, \quad (6)$$

$$J(\mathbf{x}, \nabla v) = \frac{(x + a)v_x + (y + b)v_y + 1}{f + c}, \quad (7)$$

$$\mathbf{I}(x) = \rho \frac{I \cdot \mathbf{n}}{r^2}, \quad (8)$$

$$v = \log d(x). \quad (9)$$

$d(\mathbf{x})$ is the depth of point \mathbf{x} , v_x and v_y are the spatial derivatives of v , ρ is the surface albedo and \mathbf{I} and \mathbf{n} are the light and surface normal vectors, respectively. The Lax-Friedrichs method³² is used to solve the resulting PDE and a specular highlight triangulation method is used to solve for the unknown albedo. An example of depth extraction from a single colonoscopy image can be seen in Fig. 2.




Fig. 2. SfS method employed. (a) image from the CVC-ClinicDB dataset; (b) depth estimation from SfS; (c) 3D surface recovered from SfS depth.

3.4. Proposed architectures for polyp detection

We investigated several state-of-the-art convolution architectures and adapt them through fine-tuning, for polyp detection. Specifically, we test six different architectures: AlexNet,⁹ GoogLeNet,¹² VVG¹¹ and three version of the ResNets architecture with 50, 101 and 152 layers of depth.¹³ The AlexNet and VGG are converted into FCNs (FCN-AlexNet and FCN-VGG) by discarding the two fully connected layers and replacing them with 1×1 convolution layers with the same 4096 dimensions of the fully connected layers. The final scoring layers were also replaced with a 2D, 1×1 convolution to produce the background and polyp pixel classification maps as the output. The conversion of GoogLeNet and ResNets into FCNs only requires the replacement of the scoring layers with a 2D convolution. We also increased the resolution of the output coarse map by discarding the final averaging layer. ResNets already incorporate batch normalization, while to apply this to the remaining networks, we added a regularization operation between every convolution and activation layer, as illustrated in Fig. 1. Every network is finalized with a deconvolution layer with stride $S = 32$ and a kernel of size $W' = H' = 64$, responsible for upsampling the coarse output to a dense scoring map with the same dimensions as the input. Even though the CNNs output a coarse segmentation map, a single deconvolution layer can accurately upsample the blob-like structures of most polyps. We verified that adding extra deconvolution layers from the finer

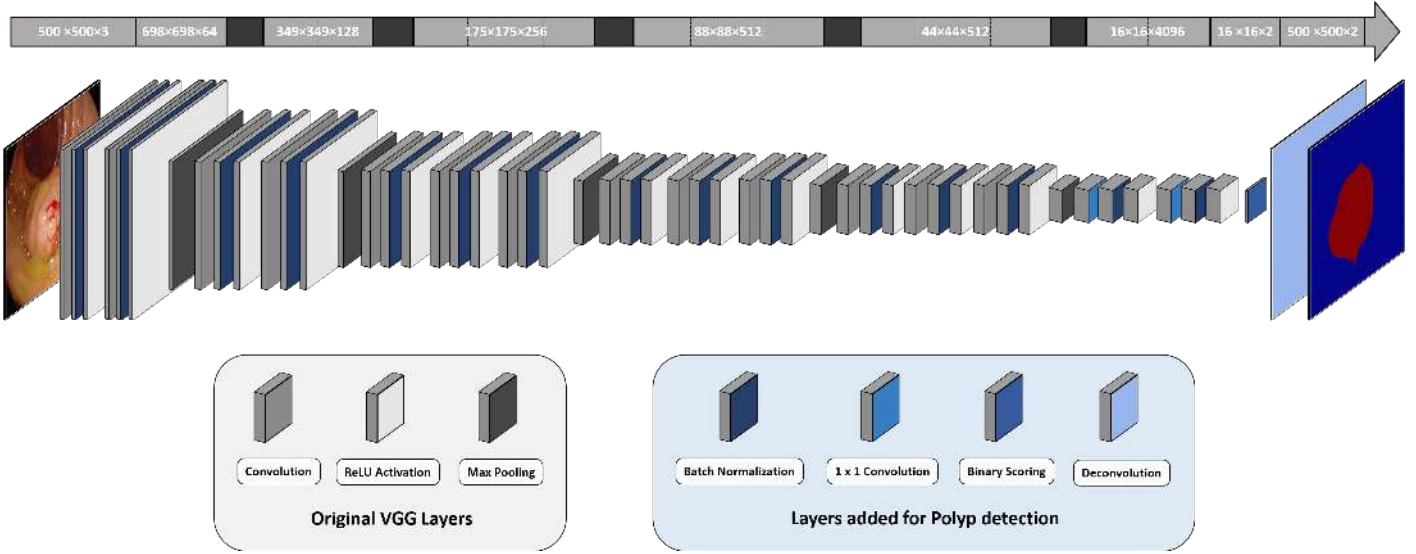


Fig. 3. Illustration of the proposed BN-FCN-VGG architecture with batch normalization for a 500×500 size image. The values on the top array represent the output size of each layer underneath. The fully connected and scoring layers of the original VGG were removed. Grey coloured layers were loaded from the original model while blue coloured layers were added or modified for polyp segmentation.

levels of the models did not improve the results. An example the proposed fully connected version of VGG with batch normalization (BN-FCN-VGG) is illustrated in Fig. 3.

3.5. Implementation and training details

Developed networks were optimized by SGD with a 0.99 momentum and all layers were updated by back-propagation. Classes probabilities are calculated with Softmax function and cross-entropy was used as the loss function and the learning rate for convolution bias was doubled. In GoogLeNet, the two deeper loss function were discarded and only the last one was used for fine-tuning. For FCN-Alexnet and FCN-VGG, the convolution filters were initialized by copying weights from public available models trained on the PASCAL segmentation dataset. Because no trained segmentation models are publicly available, the fully convolutional GoogLeNet (FCN-GoogLeNet) and the three fully convolution ResNets (FCN-ResNet) were initialized by loading classification models trained on the Imagenet dataset. New convolution layers were zero-initialized and the learning rate of scoring layers was increased by a factor of 10. We fine-tune the networks with the highest fixed learning rate that did not cause loss divergence. For FCN-GoogLeNet this corresponds to a learning rate of 10^{-12} , while all other FCNs were optimized with a learning rate of 10^{-10} . Convergence was achieved after 30K iterations for FCN-ResNet-51 and FCN-ResNet-101, 40K for FCN-GoogLeNet, 50K for FCN-VGG and FCN-ResNet-151 and 120K for FCN-AlexNet.

Images were resized to 500×500 and random flip-

ping was used for data augmentation during training. Non residual FCNs were trained with a random single image per batch. All FCN-ResNets were trained with 224×224 patches randomly sampled from the training images. This allowed to increase the batch size even with limited memory resources. The same type of sampling was performed during training of the batch normalization versions (BN-FCNs) of the non residual networks. Batch sizes of 20 were used for BN-FCN-AlexNet, BN-FCN-GoogLeNet and FCN-ResNet-51. Due to memory constrains, smaller batch sizes were used for other FCNs: 16 for FCN-ResNet-101, 8 for FCN-ResNet-151 and 5 for BN-FCN-VGG. When training networks with depth (D-FCN), the Sfs values are concatenated to the RGB channels to create a new 4-channel input. A new channel is added to every first layer convolution filter and it is initialized by averaging the values of the other filter dimensions. The learning rate of this layer is increased by a factor of 10. All models were trained and tested using the Caffe³³ software library in a single NVIDIA Tesla K40 GPU.

4. Experiments and Results

We used the public datasets from the MICCAI 2015 polyp detection challenge.²⁰ For comparison purposes, we divided the dataset for training and testing as it was suggested in the MICCAI challenge guidelines: CVC-CLINIC and ASU-Mayo for training and ETIS-Larib for testing. Furthermore, we also report results from a second public available dataset (CVC-ColonDB).²⁶ The datasets were obtained with different imaging systems and contain manual segmentations of every detected polyp. We specifically used

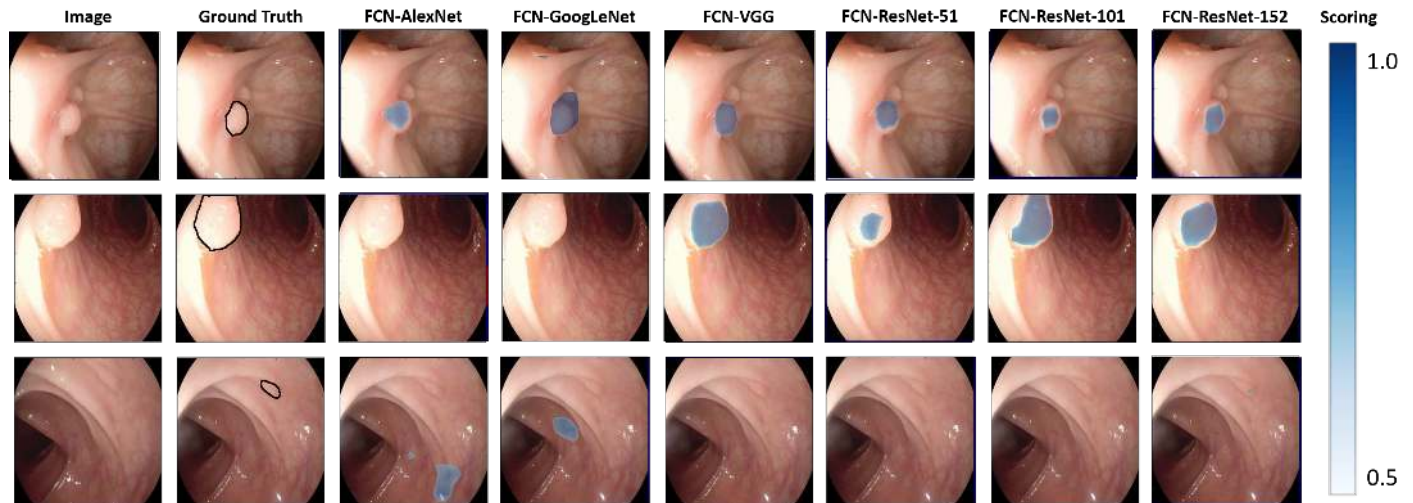


Fig. 4. Example of three different scored segmentations produced by the six proposed FCN networks. The colorbar defines the scoring probability of each pixel to belong to the polyp class. Third image results are best viewed in colour electronically because the FCN-ResNet-152 detection is very small.

Table 1. Segmentation and detection precision (prec) and recall (rec) in % obtained by the proposed FCNs. Mean interception over union (IU) is also presented for segmentation. The best result for each metric is highlighted

	ETIS-Larib					CVC-ColonDB				
	Segmentation			Detection		Segmentation			Detection	
	Prec	Rec	IU	Prec	Rec	Prec	Prec	IU	Prec	Rec
FCN-AlexNet	27.87	35.54	15.7	44.08	63.78	40.3	20.71	15.77	45.29	54.68
FCN-GoogLeNet	25.83	29.82	12.29	41.85	62.76	37.46	12.93	12.71	42.26	45.25
FCN-VGG	70.23	54.2	44.06	73.61	86.31	76.06	60.46	54.01	79.57	86.01
FCN-ResNet-50	55.75	23.43	19.72	73.84	76.53	67.76	25.64	22.74	82.89	82.38
FCN-ResNet-101	63.26	53.88	41.35	75.32	91.66	73.85	50.73	46.23	83.70	88.20
FCN-ResNet-152	65.26	38.24	33.19	79.42	89.75	72.85	50.72	43.28	82.08	93.27

the following grouping of images for training (fine-tuning) and testing:

- CVC-CLINIC: 612 SD training frames with at least one polyp each;
- ETIS-Larib: 196 HD testing frames with at least one polyp each;
- ASU-Mayo: 36 small SD and HD videos sequences, divided into training frames with and without polyps.
- CVC-ColonDB: 379 testing frames from 15 different colonoscopy sequences with at least one polyp each.

In total, the MICCAI challenge training data has 19514 frames from CVC-CLINIC and ASU-Mayo datasets. However, only 4664 of these corresponds to images with polyps. We verified that networks trained with the full dataset did not converge satisfactory in practical time and performance was substantially worst. To avoid this, we fine-tuned all proposed FCNs only with images containing polyps, achieving better performance.

The developed FCNs were formulated to produce

dense pixel-wise polyp segmentations. As such, we report results using three common segmentation evaluation metrics: mean pixel precision, mean pixel recall and interception over union (IU). If a pixel of polyp is correctly classified it is counted as a true positive (TP). Every pixels segmented as polyp that fall outside of a polyp mask counts as a false positive (FP). Finally, Every polyp pixel that has not been detected counts as a false negative (FN). The evaluation metrics are calculated according Equation 10.

$$Prec = \frac{TP}{TP+FP} \quad Rec = \frac{TP}{TP+FN} \quad IU = \frac{TP}{TP+FP+FN}. \quad (10)$$

Since, in the MICCAI challenge results are reported in terms of polyp detection, we also evaluate the polyp detection rate using the metrics advocated by the challenge directives;²⁰ detection precision and recall. If a segmented blob falls within the polyp mask it is counted as a TP. If the detected blob falls outside the ground truth mask it is a FP. Every polyp in the image that has not been detected counts as a FN. Only one TP is considered for polyp, no matter how many detections fall within the polyp mask. Detection precision and detection are calculated with the




Fig. 5. Segmentation comparison obtained by the three non-residual architectures with and without batch normalization. FCN-VGG results are best viewed in colour electronically the detection is very small.

Table 2. Segmentation and detection precision (prec) and recall (rec) in % obtained by the non residual FCNs trained with batch normalization. Mean intersection over union (IU) is also presented for segmentation. Metrics improved by adding BN are highlighted in bold

	ETIS-Larib					CVC-ColonDB				
	Segmentation			Detection		Segmentation			Detection	
	Prec	Rec	IU	Prec	Rec	Prec	Rec	IU	Prec	Rec
BN-FCN-AlexNet	30.05	29.07	17.41	38.95	62.76	46.22	28.78	21.19	43.69	80.87
BN-FCN-GoogLeNet	49.36	23.85	20.36	53.96	63.10	63.87	25.92	23.04	62.56	75.99
BN-FCN-VGG	56.87	66.59	42.32	56.24	94.01	66.8	61.3	47.18	61.57	95.16

same formulas of Equation 10

4.1. Results from RGB data

We first train every FCNs using only RGB data and compare their performance on both testing datasets. Table 1 presents the segmentation and detection results for all proposed network architectures and Fig. 4 illustrates representative examples of polyp segmentation for each network. FCN-ResNet-152 and FCN-ResNet-101 proved to be the best polyp detectors achieving the highest recalls in both databases. In some situations (third row in Fig. 4), FCN-ResNet-152 was the only network capable of correctly detecting the polyp, even with limited segmentation accuracy. Both deep architectures (FCN-ResNet-101, FCN-ResNet-152) proved to be able to learn complex filters capable of 90% detection recall in the testing datasets. FCN-ResNet-50 resulted in less accurate detections than its deeper counterparts, with approximately 10% lower detection recall. These observations indicate that, while more than 50 layers are essential in handling the high complexity of detecting polyps of various sizes, shapes and textures, as indicated by the improved performance of FCN-ResNet-100 against

FCN-ResNet-50, the addition of even more layers in FCN-ResNet-152 does not improve the detection performance.

In the non residual architectures, FCN-VGG outperforms the other FCNs by achieving detection recalls of 86% in both datasets. The simpler FCN-AlexNet successfully detected 63.78% of the ETIS-Larib polyps, and 54.68% of the CVC-ColonDB, and resulted in a considerable amount of false positives, as exemplified by the second and third segmentations of Fig. 4. Finally, the FCN-GoogLeNet produced the worst detection performance of all networks studied. Although, GoogLeNet is a deeper architecture than the other two, this does not necessarily translate to better inference ability, as the network is notoriously hard to optimize. Without the two loss functions that allow better convergence in the deeper modules, the network is not able to completely adapt to the new learning problem.

In terms of the segmentation results, FCN-VGG outperformed all other networks with an IU of 44.06% and 54.01% for ETIS-Larib and CVC-ColonDB datasets, respectively. Subsequently, even though FCN-VGG detects a smaller number of polyps, the overall quality of the segmentation it provides is superior to other networks. An example of this is depicted in the second polyp of Fig. 4. Sim-

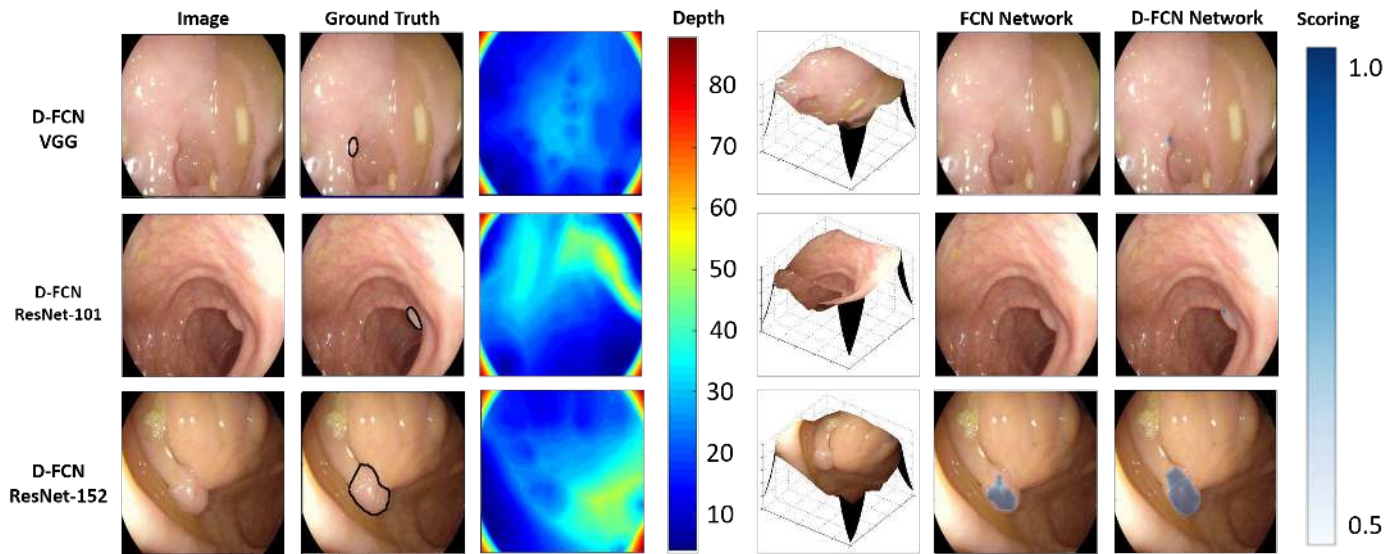


Fig. 6. Comparison between segmentations obtained by the three top-performing architectures trained with and without depth.

Table 3. Segmentation and detection precision (prec) and recall (rec) in % obtained by the three FCNs with the best performance trained with RGB-D data (D-FCNs). Mean interception over union (IU) is also presented for segmentation. Metrics improved by adding depth information are highlighted in bold

	ETIS-Larib						CVC-ColonDB				
	Segmentation			Detection			Segmentation			Detection	
	Prec	Rec	IU	Prec	Rec	Prec	Rec	IU	Prec	Rec	
D-FCN-VGG	68.68	62.16	47.78	73.32	88.01	74.94	68.02	56.95	76.85	91.03	
D-FCN-ResNet-101	55.63	61.11	40.99	70.62	95.83	74.31	58.15	49.65	83.17	90.47	
D-FCN-ResNet-152	67.66	39.78	33.77	77.95	90.2	71.16	47.95	41.32	80.78	92.5	

ilar levels of segmentation quality were achieved by FCN-ResNet-101, which indicates that it learned more general filters than its deeper residual counterpart (FCN-ResNet-151). Finally, similar to the detection results, FCN-AlexNet and FCN-GoogLeNet achieved the worst performance in segmenting the polyps.

As far as we know, our method is the first to produce dense polyp segmentations, which only allow comparison with other algorithms with the use of detection metrics. The current state of the art was set by the top deep learning method in the 2015 MICCAI polyp detection challenge (OUS), which achieved 73.3% detection precision and 69.2% recall in the ETIS-Larib dataset.²⁰ As seen in Table 1, four of our models (FCN-VGG and all three FCN-ResNets) surpass this results, with improvements in precision and increases in recall as high as 20%. The OUS methodology was not made publicly available yet, so direct comparison is not possible. However, the huge difference in accuracy shows how important proven CNN architectures and a good initialization are to achieved a better solution.

4.2. Adding batch normalization

Batch normalization is not implemented in the original AlexNet, VGG and GoogLeNet. We investigate the influence of adding batch normalization in these FCNs and list the results in Table 2. During training of the batch-normalized (BN-FCNs) versions, we decreased the learning rate by a factor of 100 and convergence was achieved after 30K iterations for all BN-FCNs. Due to memory limitations, relatively small batch sizes were used. BN-FCN-AlexNet resulted in a slight increase in IU segmentation and detection recall for ETIS-Larib, while in CVC-ColonDB, every single evaluation metric was improved, especially for detection, where the recall increased by more than 25%. Similar improvements in segmentation IU and detection accuracy, are observed with BN-FCN-GoogLeNet for both datasets. Examples of improved segmentations are illustrated in Fig. 5 for all three non-residual networks. Batch normalization enabled BN-FCN-VGG to increase the amount of polyps detected, with recalls higher than 94% for both datasets. However, this was accompanied with an decrease in precision. The third row in Fig. 5 shows an example of a polyp being misdetected without batch normalization (FCN-VGG) while being successfully recovered

in the batch-normalized version (BN-FCN-VGG).

4.3. Results from RGB-D data

To evaluate the addition of SfS-extracted depth, as an additional feature, we restricted ourselves to the three architectures that achieved the best detection and segmentation results with RGB data; FCN-VGG, FCN-ResNet-101 and FCN-ResNet-152. The results after the inclusion of depth information are listed in Table 3.

The addition of depth information allowed D-FCN-VGG to perform slightly better than its RGB counterpart. Segmentation IU and detection recall improved approximately by 2% for both datasets. Similar increases were verified with D-FCN-ResNet-101, elevating its detection recall to more than 95% for the ETIS-Larib. Fig 6 illustrates three examples where depth information allowed the networks to either detect a polyp that would otherwise miss (D-FCN-VGG, D-FCN-ResNets-101) or improve segmentation accuracy (D-FCN-ResNets-152).

With the current architectures, bigger improvements are hampered by the difficulty of propagating meaningful gradients through the model. This issue is more evident in the FCN-ResNet-152, which had comparable detection performance with and without depth data. Alternative ways to incorporate depth information into the models might facilitate the learning of more meaningful RGB-D feature extractors.

4.4. Computation speed

The inference speed of each network highly depends of the amount of learned parameters and the number of layers. Table 4 lists the average time required for each FCN to segment a single 500×500 image. The addition of batch normalization and the use of depth features slows down inference, as more operations are required to produce the final segmentation maps. The VGG architecture has the highest number of receptive fields, which results in the slowest average inference speed of all networks. On the other end, FCN-AlexNet has the fastest average inference both with (136ms) and without (51ms) batch normalization.

Table 4. Average inference time in milliseconds (ms) for a 500×500 image. If applicable, average inference time is shown , without batch normalization (no BN), with batch normalization (BN) and with the inclusion of depth (Depth)

Networks	Average Inference Speed (ms)		
	no BN	BN	Depth
FCN-AlexNet	51	136	-
FCN-GoogLeNet	60	193	-
FCN-VGG	295	412	536
FCN-ResNet-51	-	164	-
FCN-ResNet-101	-	206	265
FCN-ResNet-152	-	319	387

5. Conclusion

We have presented a deep learning framework for automatically detecting and segmenting polyps in colonoscopy images. This is achieved by taking advantage of very rich representations available in CNNs trained on large databases which we fine-tune to perform polyp detection and adapt them, by converting them to FCNs for achieving segmentation. We compare the networks' ability to accurately detect and segment polyp structures in experiments on publicly available datasets with annotated ground truth. Obtained results suggest that the two deepest residual architectures (ResNet-101, ResNet-152) were able to cope with the complexity of polyp structures and achieve the best detection results. On the other hand, the higher number of receptive fields in the VGG, allowed the network to achieve a better overall segmentation output. These three networks achieved detection recall rates around 90% both in the ETIS-Larib and CVC-ColonDB, considerably surpassing the state-of-the-art in polyp detection. We also introduce relative depth information, derived from SfS as an additional input channel. Results show that including depth can improve polyp representation and lead to increased detection rates and segmentation accuracy. We note thought that it is hard to propagate meaningful gradients through deep networks, thus the directly visible improvement was small. In the future, a modeling strategy that facilitates learning from depth might yield better prediction and inference rates rather than direct SfS embedding. Inference times are still a limitation for the use of CNNs in real-time CAD systems suitable for clinical practice. Strategies to speed up the inference process, such as the use of FFT convolutions, could be used for achieving minimum latency, required in clinical applications. Furthermore, a more detailed analysis of receptive field activation might allow the compression of large scale CNNs for simpler problems, providing even faster inference times.

References

- [1] Ciuti, G., Calio, R., et al., "Frontiers of robotic endoscopic capsules: a review," *Journal of Micro-Bio Robotics* 11(1), 1-18 (2016).
- [2] Leufkens, A. M., Van Oijen, M. G., et al., "Factors influencing the miss rate of polyps in a back-to-back colonoscopy study," *Endoscopy* 44(5), 470-475 (2012).
- [3] Karkanis, S. A., et al. "Computer-aided tumor detection in endoscopic video using color wavelet features." *Inf Technol Biomed* 7(3),141-152 (2003).
- [4] Ameling, S., et al. "Texture-based polyp detection in colonoscopy." *Bildverarbeitung fr die Medizin*, 22(25), 346-350 (2009).
- [5] Tajbakhsh, N., Suryakanth R. G. and Jianming L. "Automated polyp detection in colonoscopy videos using shape and context information." *Trans. Med. Imag.* 35(2),630-644 (2016).

- [6] Bernal, J., Sanchez, F. J., et al., "WM-DOVA maps for accurate polyp highlighting in colonoscopy: Validation vs.saliency maps from physicians," *Comput. Med. Imaging Graph.* 43, 99-111 (2015).
- [7] Russakovsky, O., Deng, J., et al., ImageNet Large Scale Visual Recognition Challenge, *Int. J. Comput. Vision* 115(3), 211252 (2015).
- [8] M. Everingham, S. M. A. Eslami, et al., The Pascal Visual Object Classes Challenge: A Retrospective, *Int. J. Comput. Vis.*, vol. 111, no. 1, 98136 (2014).
- [9] Krizhevsky, A., Ilya S. and Geoffrey E. H.. "Imagenet classification with deep convolutional neural networks." *NIPS*. 25, 1106-1114 (2012).
- [10] Russakovsky, O., et al. "Imagenet large scale visual recognition challenge." *INT J COMPUT VISION*. 115(3), 211-252 (2015).
- [11] Simonyan, K. and Zisserman, A., Very deep convolutional networks for large-scale image recognition, in *Proc. ICLR*, 114 (2015).
- [12] Szegedy, C., Liu, W., et al., Going deeper with convolutions, in *Proc. IEEE CVPR*, 19 (2015).
- [13] K. He, X. Zhang, S. Ren, and J. Sun. Deep residual learning for image recognition. *arXiv preprint arXiv:1522.03385* (2015).
- [14] Silberman, N. and Fergus, R., "Indoor scene segmentation using a structured light sensor", in *Proc. ICCV*, 601-608 (2011).
- [15] Couprie, C., Farabet, C., et al., "Indoor Semantic Segmentation using Depth Information", *arXiv preprint arXiv:1301.3572* (2013).
- [16] Long, J., Shelhamer, E., et al., "Fully convolutional networks for semantic segmentation" in *Proc. IEEE CVPR*, 3431-3440 (2015).
- [17] Bernal, J., Tajkbaksh, N., et al., Comparative Validation of Polyp Detection Methods in Video Colonoscopy: Results from the MICCAI 2015 Endoscopic Vision Challenge, *IEEE Trans. Med. Imag.*, (2016).
- [18] Bernal, J., *et al.* Towards Automatic Polyp Detection with a Polyp Appearance Model *Pattern Recognition*, 45(9), 31663182 (2012).
- [19] Rongsheng, Z., *et al.* "Lesion detection of endoscopy images based on convolutional neural network features" *CISP*, 372-376 (2015).
- [20] Bernal, J., Tajkbaksh, *et al.* "Comparative Validation of Polyp Detection Methods in Video Colonoscopy: Results from the MICCAI 2015 Endoscopic Vision Challenge," *IEEE Trans. Med. Imag.* (2016).
- [21] Tajbaksh, N., Suryakanth R. G., and Jianming L., "Automatic polyp detection in colonoscopy videos using an ensemble of convolutional neural networks." *ISBI*, 79-83 (2015).
- [22] Tajbaksh, N., et al. "Convolutional Neural Networks for Medical Image Analysis: Full Training or Fine Tuning?" *Trans. Med. Imag.* 35(5), 1299-1312 (2016).
- [23] Zeiler, M. D. and Fergus R.. "Visualizing and understanding convolutional networks." in *Proc. ECCV*, 13(1), pp. 818-833 (2014).
- [24] Dumoulin, V. and Visin, F., A guide to convolution arithmetic for deep learning, *arXiv:1603.07285*, 128 (2016).
- [25] Ioffe, S. and Szegedy, C., "Batch Normalization: Accelerating Deep Network Training by Reducing Internal Covariate Shift". *ICML*, 37(1), 448-456 (2015).
- [26] Bernal, B., Sanchez, F. and Vilario, F., Towards Automatic Polyp Detection with a Polyp Appearance Model. *Pattern Recognition*, 45(9), 31663182 (2012).
- [27] Ahmed, A. H. and Farag, A. A., "Shape from Shading Under Various Imaging Conditions", in *Proc. CVPR*, 1-8 (2007).
- [28] A. Ahmed and A. Farag, "Shape from Shading for Hybrid Surfaces", *ICIP*, volume 2, 525-528 (2007).
- [29] Zhang, L., Yip, A.M. *et al.*, "A unified framework for document restoration using inpainting and shape-from-shading", *Pattern Recognition*, 42(11), 2961 - 2978 (2009).
- [30] Visentini-Scarzanella, M., Stoyanov, D. and Yang, G.Z. "Metric depth recovery from monocular images using shape-from-shading and specularities", in *Proc. ICIP*, 25-28 (2012).
- [31] Brandao, P., Mazomenos, E., *et al.*, "Fully convolutional neural networks for polyp segmentation in colonoscopy". *SPIE*, 191349 (1) (2017).
- [32] Kao, C. Y., Osher, S. and Qian, J., "Lax-Friedrichs sweeping scheme for static Hamilton-Jacobi equations". *J. Comput. Phys.*, 196(1), 367-391 (2004).
- [33] Jia, Y., Shelhamer, E., et al., Caffe: Convolutional Architecture for Fast Feature Embedding, *Proc. ACM MM*, 675-678 (2014).
- [34] Montserrat, T., et al., "Depth estimation based on multiview matching with depth/color segmentation and memory efficient belief propagation." in *Proc. IEEE ICIP* (2009).
- [35] Nistr, D. "Preemptive RANSAC for live structure and motion estimation." *Machine Vision and Applications*, 321-329 (2005).
- [36] Forsyth, D. A. "Shape from texture without boundaries" *ECCV*, 225-239 (2002).
- [37] Chen, L., et al. "Semantic image segmentation with deep convolutional nets and fully connected crfs." *ICLR* (2015).
- [38] Noh, H., Seunghoon H. and Bohyung H., "Learning deconvolution network for semantic segmentation" in *Proc. ICCV*, 1520-1528 (2015).
- [39] Hazirbas, C., et al., "FuseNet: Incorporating Depth into Semantic Segmentation via Fusion-based CNN Architecture." *Proc. ACCV*. Vol. 2 (2016).