



Focus U-Net: A novel dual attention-gated CNN for polyp segmentation during colonoscopy

Michael Yeung^{a,b,*}, Evis Sala^{a,c}, Carola-Bibiane Schönlieb^d, Leonardo Rundo^{a,c,**}

^a Department of Radiology, University of Cambridge, Cambridge, CB2 0QQ, United Kingdom

^b School of Clinical Medicine, University of Cambridge, Cambridge, CB2 0SP, United Kingdom

^c Cancer Research UK Cambridge Centre, University of Cambridge, Cambridge, CB2 0RE, United Kingdom

^d Department of Applied Mathematics and Theoretical Physics, University of Cambridge, Cambridge, CB3 0WA, United Kingdom

ARTICLE INFO

Keywords:

Polyp segmentation
Colorectal cancer
Colonoscopy
Computer-aided diagnosis
Focus U-Net
Attention mechanisms
Loss function

ABSTRACT

Background: Colonoscopy remains the gold-standard screening for colorectal cancer. However, significant miss rates for polyps have been reported, particularly when there are multiple small adenomas. This presents an opportunity to leverage computer-aided systems to support clinicians and reduce the number of polyps missed.

Method: In this work we introduce the Focus U-Net, a novel dual attention-gated deep neural network, which combines efficient spatial and channel-based attention into a single Focus Gate module to encourage selective learning of polyp features. The Focus U-Net incorporates several further architectural modifications, including the addition of short-range skip connections and deep supervision. Furthermore, we introduce the Hybrid Focal loss, a new compound loss function based on the Focal loss and Focal Tversky loss, designed to handle class-imbalanced image segmentation. For our experiments, we selected five public datasets containing images of polyps obtained during optical colonoscopy: CVC-ClinicDB, Kvasir-SEG, CVC-ColonDB, ETIS-Larib PolypDB and EndoScene test set. We first perform a series of ablation studies and then evaluate the Focus U-Net on the CVC-ClinicDB and Kvasir-SEG datasets separately, and on a combined dataset of all five public datasets. To evaluate model performance, we use the Dice similarity coefficient (DSC) and Intersection over Union (IoU) metrics.

Results: Our model achieves state-of-the-art results for both CVC-ClinicDB and Kvasir-SEG, with a mean DSC of 0.941 and 0.910, respectively. When evaluated on a combination of five public polyp datasets, our model similarly achieves state-of-the-art results with a mean DSC of 0.878 and mean IoU of 0.809, a 14% and 15% improvement over the previous state-of-the-art results of 0.768 and 0.702, respectively.

Conclusions: This study shows the potential for deep learning to provide fast and accurate polyp segmentation results for use during colonoscopy. The Focus U-Net may be adapted for future use in newer non-invasive colorectal cancer screening and more broadly to other biomedical image segmentation tasks similarly involving class imbalance and requiring efficiency.

1. Introduction

Globally, colorectal cancer (CRC) ranks third in terms of incidence, and second only to lung cancer as a leading cause of cancer death [1]. The absence of specific symptoms in the early stages of disease often results in delays in diagnosis and treatment, with the stage of disease at diagnosis strongly linked to prognosis. In the United States, the 5-year relative survival rate for Stage I colon cancer is 92%, decreasing to 12% in those with Stage IV [2].

In 1988, Vogelstein proposed the adenoma-carcinoma sequence

model for CRC carcinogenesis, describing the transition from benign adenoma to adenocarcinoma with associated well-defined histology at each stage [3]. Importantly, there is a prolonged, identifiable and treatable preclinical phase lasting years prior to malignant transformation [4,5]. As a result, CRC is highly suitable for population level screening, which has been shown to be effective at reducing overall mortality [6,7].

Non-invasive CRC screening tests include stool-based tests, such as the faecal occult blood test, and more recent blood-based tests, such as Epi proColon® (Epigenomics AG, Berlin, Germany). Capsule colon

* Corresponding author. Department of Radiology, University of Cambridge, Box 218, Cambridge Biomedical Campus, Cambridge, CB2 0QQ, United Kingdom.

** Corresponding author. Department of Radiology, University of Cambridge, Box 218, Cambridge Biomedical Campus, Cambridge, CB2 0QQ, United Kingdom.

E-mail addresses: mjyy2@cam.ac.uk (M. Yeung), es220@medschl.cam.ac.uk (E. Sala), cbs31@cam.ac.uk (C.-B. Schönlieb), lr495@cam.ac.uk (L. Rundo).

endoscopy and CT colonography are newer, non-invasive radiological investigations useful for screening high-risk individuals unsuitable for colonoscopy. Invasive options include flexible sigmoidoscopy and colonoscopy, offering direct visualisation and the ability to obtain biopsy specimens for histological analysis. Sigmoidoscopy is limited to cancer in the rectum, sigmoid and descending colon, and colonoscopy remains the gold-standard screening tool for CRC with the highest sensitivity and specificity [8]. However, colonoscopy is associated with significant miss rates for polyp detection, contributed by both patient and polyp-related factors [9–11]. The risk of missing polyps significantly increases in patients with two or more polyps, with higher miss rates for flat or sessile compared to pedunculated or sub-pedunculated polyps and miss rates vary from 2% for adenomas ≥ 10 mm to 26% for adenomas < 5 mm.

The difficulty in detecting polyps during colonoscopy presents an opportunity to incorporate computer-aided systems to reduce polyp miss rates [12]. Polyps may remain hidden from the field of view, for which a real-time Artificial Intelligence (AI) model has been developed to assess the quality of colonoscopy [13]. Alternatively, polyps may enter the field of view but remain undetected by the operator. In this case, polyp segmentation approaches not only aim to detect polyps, but to also accurately delineate the polyp border from surrounding mucosa. Early automated methods to segment polyps relied on hand-crafted feature extraction, using either shape-based [14–16] or texture and colour-based analysis [17,18]. While considerable advancements were made, the accuracy of polyp segmentation remained low with hand-crafted features unable to capture the scale of polyp heterogeneity [19].

This paper is structured as follows. Section 2 outlines the state-of-the-art of polyp segmentation on colonoscopy images. Section 3 describes the architecture of the proposed Focus U-Net. Section 4 describes the analysed datasets and the evaluation metrics used in this study. Section 5 present the experimental results. Finally, Section 6 provides a discussion and concluding remarks.

2. Related work

In recent years, significant improvements have been achieved by adopting automatic methods based on deep learning. The introduction of Fully Convolutional Networks (FCN) enabled Convolutional Neural Network (CNN) architectures to tackle semantic image segmentation tasks [20]. The application of FCNs to polyp segmentation has yielded impressive results [21,22]. Currently, the state-of-the-art approaches are largely based on the U-Net, a modified FCN architecture developed for biomedical image segmentation [23]. The U-Net consists of an encoding network used to capture the image context, followed by a symmetrical decoding network enabling localisation of salient regions. UNet++ extends the U-Net by incorporating a series of nested skip connections, reducing the semantic gap between the features maps of the encoder and decoder networks prior to fusion [24–26]. The ResUNet++ combines residual units with the spatial attention-based Atrous Spatial Pyramidal Pooling (ASPP) and channel attention-based squeeze-and-excitation block [25,26]. Similarly, both attention components are incorporated into the DoubleU-Net, which further leverages transfer learning from the first U-Net to generate features as input into the second network [27]. Despite excellent segmentation results with these models, the large memory and associated long inference time limits use in clinical practice where real-time polyp segmentation is required. Recently, several efficient models with significantly faster inference times, in addition to greater accuracy, have been proposed. Pritising efficiency over performance, PolypSegNet introduces the depth dilated inception module, enabling efficient feature extraction across a range of receptive field sizes [28]. Similarly, ColonSegNet is a light-weight network that includes residual connections and channel attention to achieve real-time polyp segmentation [29]. PraNet uses a two-step process that involves initial localisation of the polyp area, followed by progressive refining of the polyp boundary, resembling the method by which humans identify

polyps [30]. HarDNet-MSEG uses a low memory latency HarDNet68 backbone [31], together with a Cascaded Partial Decoder [32] for fast and accurate polyp segmentation. Progressively normalised self-attention network introduces a self-attention module that incorporates channel-split, query-dependent and normalisation rules to improve computational efficiency [33]. The feedback attention network (FANet) uses a form of hard attention based on an iterative refinement method using Otsu thresholding [34].

In this paper, we introduce a novel attention-gated U-Net architecture, named the Focus U-Net, which uses a new attention module known as the Focus Gate (FG), incorporating both spatial and channel-based attention with a focal parameter to control the degree of background suppression. Using this architecture, we achieve state-of-the-art results across five public polyp segmentation datasets. With an efficient and accurate polyp segmentation algorithm, we provide the latest advancement towards using AI in colonoscopy practice, with the aim of assisting clinicians by increasing polyp detection rates.

3. The proposed focus U-Net architecture

In this section, we introduce the techniques used in the Focus U-Net, beginning with the FG and associated channel and spatial attention modules, followed by explanations of deep supervision and loss function optimisation.

3.1. Overview of the focus U-Net

The architecture of the Focus U-Net is shown in Fig. 1. Similar to the U-Net, the Focus U-Net begins with an encoding network, capturing features relevant to polyps such as edges, texture and colour. The deepest layer of the network contains the richest information relating to image features, at the cost of spatial resolution, and forms the gating signal used as input into the FG. The FG uses the gating signal to refine incoming signals from the encoding network in the form of long-range skip connections, by highlighting specific image features and regions that are integrated into the decoding network. Successive upsampling in the decoding network enables polyp localisation at progressively higher resolution, with the final output producing the segmentation map defining, if present, the precise shape and location of the polyp. Short-range skip connections and deep supervision create additional pathways for information transfer, diversifying the features extracted and providing shortcuts for the loss to propagate backwards to the deeper layers when updating parameters.

3.2. Attention Gates and the Focus Gate

The concept of attention mechanisms in neural networks is inspired by cognitive attention, where relevant stimuli in the visual field are identified and selectively processed. In the context of neural networks, distinctions are made between hard and soft attention, as well as global and local attention [35,36]. Hard attention calculates attention scores for each region of the image to select the regions to attend. This requires a stochastic sampling process, which is a non-differentiable calculation relying on reinforcement learning to update parameters [37]. In contrast, soft attention is deterministic and assigns regions of interest (ROIs) with higher weight, with the benefit that this process is differentiable and therefore trainable by standard backpropagation [38,39]. The distinction between global and local attention refers to whether the whole input or only a subset of the input is attended [40]. For training of neural networks, a combination of soft and local attention is often favoured [41].

Attention Gates (AGs) provide neural networks with the capacity to selectively attend to inputs. The use of AG first originated in the context of machine translation as part of Natural Language Processing (NLP) [40,42–44], but has also more recently shown success in Computer Vision, with particular interest in the Attention U-Net for medical image

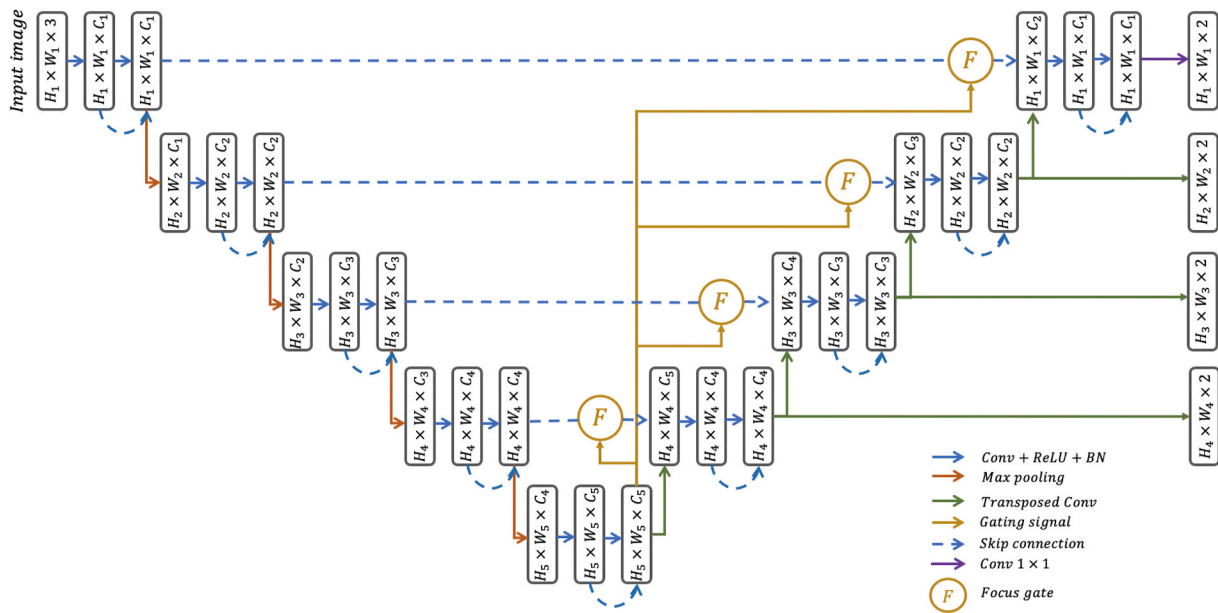


Fig. 1. Architecture of the proposed Focus U-Net. The gating signal originates from the deepest layer in the network and refines incoming skip connection input at each depth. Deep supervision is represented by the dark green arrows generating outputs at each depth.

segmentation [41,45].

The structure of the additive AG is illustrated in Fig. 2 [41]. This AG receives two inputs, the gating signal and associated skip connection generated at that level. The gating signal originates from the deepest

layer of the neural network, where feature representation is the greatest at the cost of significant down-sampling. In contrast, skip connections arise in more superficial layers, where feature representation is coarser, but image resolution is relatively spared. The AG uses contextual

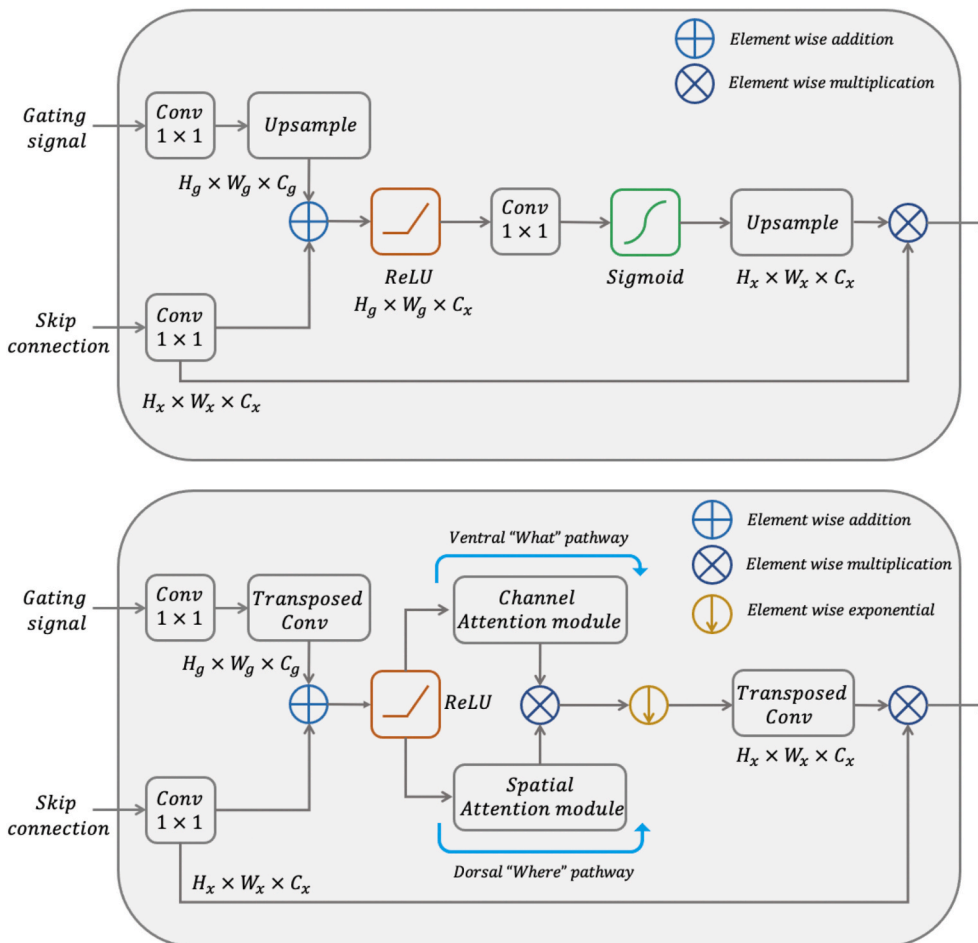


Fig. 2. Top: schematic of the additive AG. The gating signal and skip connection are first resized and then combined to form attention coefficients. Multiplication of the original skip connection with the attention coefficients provides spatial context highlighting ROIs. Bottom: schematic of the Focus Gate. The gating signal and skip connection are first resized and then combined prior to spatial and object-related feature extraction. The attention coefficients pass through an additional focal filter controlling the degree of background suppression. Finally, multiplying the original skip connection with the attention coefficients provides both spatial and feature context highlighting regions and features of interest.

information from the gating signal to prune the skip connection, highlighting ROIs and therefore reducing false positive predictions. To accomplish this, the initial stage involves simultaneous upsampling of the gating signal and downsampling of the skip connection to produce equivalent image dimensions enabling element-wise addition. Although computationally more expensive, additive attention has been shown to achieve higher accuracy than multiplicative attention [40].

The resulting matrix is passed through a ReLU activation, followed by global average pooling along the channel axis and final sigmoid activation, generating a matrix of attention weights, also known as the attention coefficients $i \in [0, 1]$.

The final step is an element-wise multiplication of the upsampled attention coefficients with the original skip connection input, providing spatial context to the skip connection prior to fusion with outputs from the decoder network.

Before describing the FG, illustrated in the bottom of Fig. 2, we first describe two of its main components, namely the channel attention module and the spatial attention module.

3.3. Channel attention module

The global average pooling operation in the additive AG extracts the spatial context to localise the ROIs. However, by pooling across the channel axis, information conveyed by the channels relating to objects features, such as edges and colour, is lost. On the contrary, by assigning weights along the channel axis, channel interdependencies may be

explicitly modelled, enabling networks to better recalibrate the features used for segmentation [46–49]. Squeeze-and-excitation (SE) blocks achieve this by initial feature aggregation using global average pooling along the spatial axis, known as the ‘squeeze’ operation, followed by two fully connected layers with ReLU and sigmoid activations producing the ‘excitation’ operation [46,50]. The two fully connected layers involve dimensionality reduction to control model complexity, with implications for computation and performance. Efficient Channel Attention (ECA) [51] avoids dimensionality reduction by modelling cross-channel interaction with an adaptive kernel size k , defined by:

$$k = \left\lfloor \frac{\log_2(C)}{\gamma} + \frac{b}{\gamma} \right\rfloor_{\text{odd}} \quad (1)$$

, where C is the channel dimension, while b and γ are set to 2 and 1, respectively.

A separate insight incorporated into the Convolutional Block Attention Module (CBAM) for channel attention involves using a global max pooling operation in addition to global average pooling, providing two complementary spatial contexts prior to feature recalibration [48].

The channel module used in the FG is illustrated in Fig. 3. We extend the ideas provided by ECA and CBAM by using initial global average and global max pooling to generate two separate spatial contexts, followed by feature recalibration using an adaptive convolutional kernel size avoiding dimensionality reduction. Finally, a sigmoid activation re-distributes the values between [0, 1], generating attention coefficients

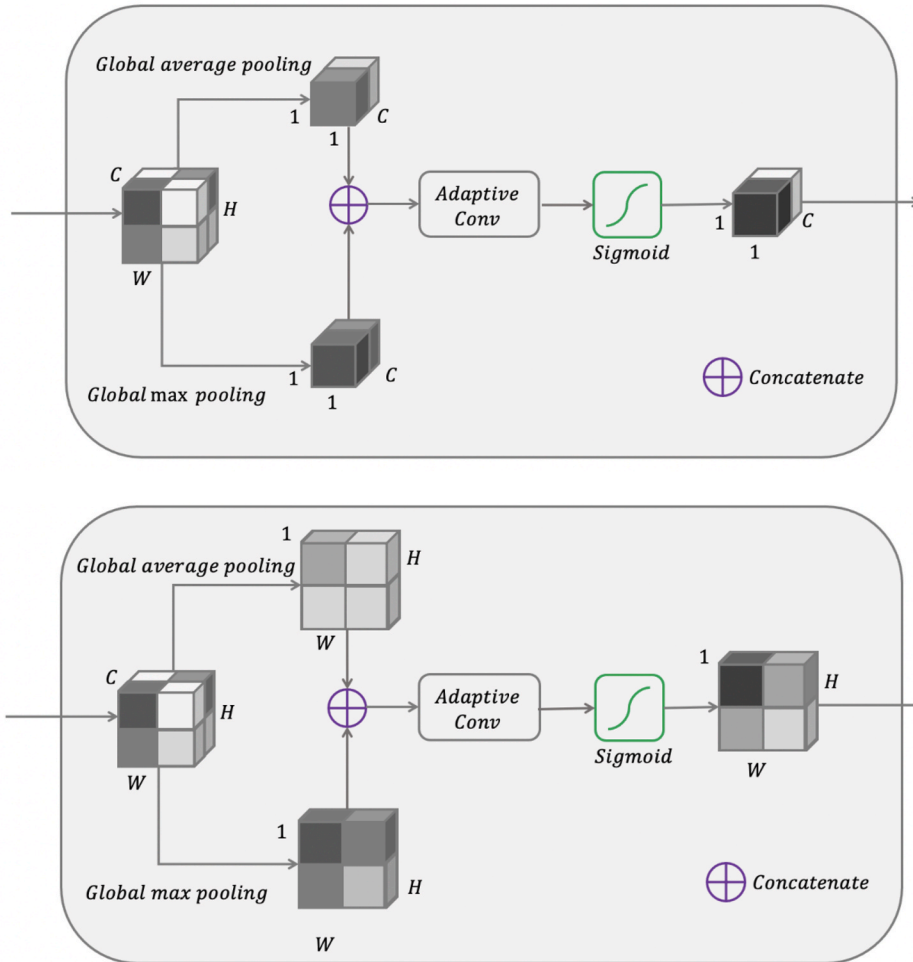


Fig. 3. Top: schematic of the channel attention module used in the Focus Gate. Global max pooling and global average pooling generate two complementary spatial contexts prior to channel weighting. Bottom: schematic of the spatial attention module used in the Focus Gate. Global max pooling and global average pooling generate two complementary channels contexts prior to spatial context generation.

along the channel axis.

3.4. Spatial attention module

Complementary to the channel attention module, spatial attention modules involve feature aggregation along the channel axis [47,48,52]. While dimensional reduction is not an issue for spatial attention modules, the replacement of fully connected layers with a convolutional layer requires an additional kernel size parameter. Larger kernel sizes provide a larger receptive field, with better performance but at the cost of computational efficiency [52]. The spatial attention module used in the FG is illustrated in Fig. 3. Again, we extend the ideas provided by ECA and CBAM by using initial global average and global max pooling along the channel axis, generating two separate channel contexts, followed by spatial recalibration with an adaptive convolutional kernel size. In contrast to ECA, the spatial dimension is inversely proportional to the channel dimension, and therefore we modify the original equation and determine kernel size k for the spatial attention module by:

$$k = \left\lfloor \frac{\log_2(C_{\max} + C_0 - C) + \frac{b}{\gamma}}{\gamma} \right\rfloor_{\text{odd}} \quad (2)$$

where C_{\max} is the maximum channel dimension of the network, C_0 is the channel dimension for the first layer, and C is the channel dimension for current layer. The parameters b and γ are set to 2 and 1, respectively [51]. This provides an efficient compromise by scaling the kernel size in proportion to the input dimension, with larger kernel sizes reserved for larger inputs.

3.5. Focus gate

Having introduced both spatial and channel attention modules, in this section we describe the structure of the FG (Fig. 2).

Similar to the attention gate, the gating signal is generated from the deepest layer of the network. The upsampling operation is replaced with a learnable kernel weight using a transposed convolution, but otherwise the skip connection and gating signal are resampled to matching

dimensions. Following element-wise addition and non-linear activation, spatial and channel attention coefficients are processed in parallel, analogous to processing of the dorsal “where” and ventral “what” pathways, respectively, of the two-streams hypothesis for visual processing [53]. The spatial and channel attention coefficients are combined with element-wise multiplication, and passed through a tunable filter involving element-wise exponential parameterised by the focal parameter prior to resampling.

The concept of a focal parameter originates from work on loss function optimisation, where the contributions of easy examples are downweighed enabling the learning of harder examples [54,55]. Here, we apply the focal parameter to the matrix of attention coefficients, enhancing the contrast between foreground and background objects by controlling the degree of background suppression. Following sigmoid activation, all attention coefficient values are redistributed $i \in [0, 1]$. This enables higher values of the focal parameter to significantly reduce the weights of irrelevant regions and features, while salient regions and important features are relatively spared. The effect of altering the focal parameter is illustrated in Fig. 4.

Careful tuning of the focal parameter is required, to suppress background regions while preserving attention for borders between foreground and background where attention coefficients take middle values.

3.6. Deep supervision

The vanishing and exploding gradients problems are well-recognised issues with training deep CNNs [56,57]. The Focus U-Net incorporates two separate, complementary mechanisms to overcome this. Firstly, short-range skip connections, in addition to the long-range skip connections characteristic of the U-Net, allow the error signal to propagate to earlier layers more directly. However, this comes at the cost of computational efficiency, and therefore for the Focus U-Net, we leverage filter factorisation introduced by the Inception network and incorporated into the MultiResUNet, providing an efficient implementation while maintaining performance gains [58,59].

In contrast, deep supervision encourages semantic discrimination of

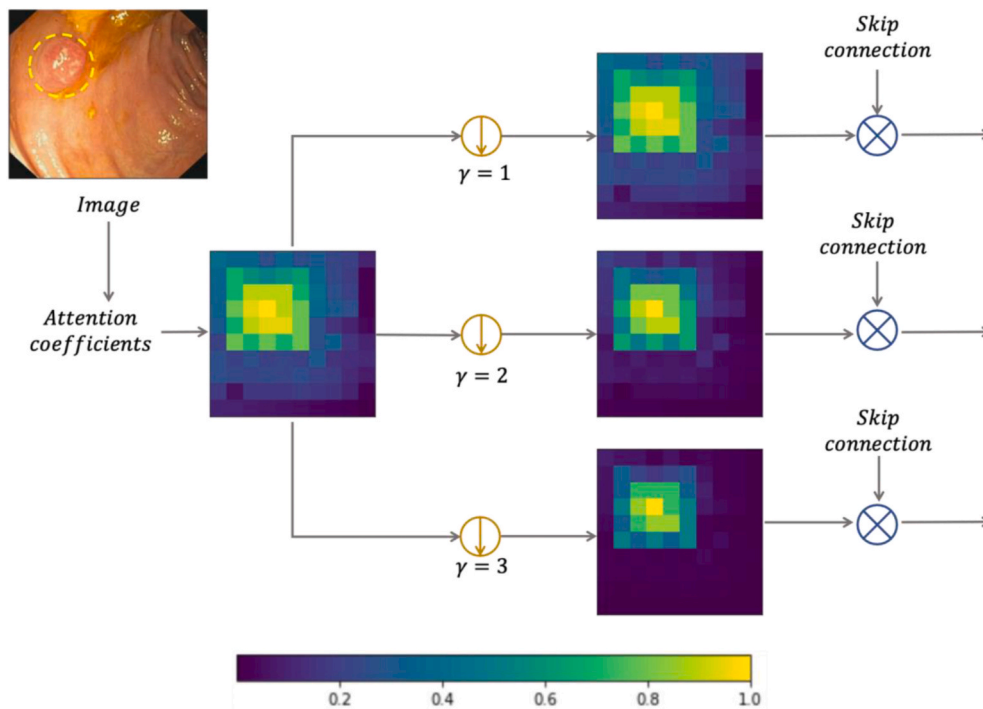


Fig. 4. The effect of modifying the focal parameter in the Focus Gate. The brighter colours in the heatmaps are associated with higher activations, here corresponding to the polyp location. Higher focal parameter values lead to increased background suppression.

intermediate feature maps at each level by assigning a loss to outputs at multiple layers [60,61]. Equal weighting of outputs produces sub-optimal results due to converging to solutions favouring improved performance of deeper layers at the cost of performance of the final layer. To overcome this, more complicated solutions have been developed, such as multi-scale training [55], or fine-tuning using a fully connected layer [41]. To preserve efficiency, here we assign weights w to different output layers according to Eq. (3):

$$w = 2^{-(\text{stride length} \times \text{stride width})}, \quad (3)$$

where the stride length and width refer to the final transposed convolution stride dimensions required to resample the feature map to the original image dimension.

Intuitively, higher weights are therefore assigned to the layers requiring a smaller degree of upsampling, with the greatest weight assigned to the final output, followed by an exponential decrease in weighting with increasing depth of the network.

3.7. Hybrid Focal loss

The training of neural networks is based on solving the optimisation problem defined by the loss function. For semantic segmentation tasks, a popular choice of loss function is the sum of the Dice similarity coefficient (DSC) loss and cross entropy (CE) loss:

$$\mathcal{L}_{\text{total}} = \text{DSC} + \mathcal{L}_{\text{CE}}, \quad (4)$$

with:

$$\text{DSC} = \frac{2\text{TP}}{2\text{TP} + \text{FP} + \text{FN}}, \quad (5)$$

$$\mathcal{L}_{\text{CE}}(\mathbf{y}, \hat{\mathbf{y}}) = -(\mathbf{y} \log(\hat{\mathbf{y}}) + (1 - \mathbf{y}) \log(1 - \hat{\mathbf{y}})), \quad (6)$$

where TP, FP and FN refer to true positives, false positives and false negatives respectively, and $\mathbf{y}, \hat{\mathbf{y}} \in \{0, 1\}^N$ where $\hat{\mathbf{y}}$ refers to the predicted value and \mathbf{y} refers to the ground truth label.

However, with class imbalanced tasks such as polyp segmentation, the resulting segmentation using the Dice loss often leads to high precision but low recall rate [62]. By weighting false negative predictions more heavily, the Tversky loss improves recall-precision balance:

$$\mathcal{L}_{\text{Tversky}} = 1 - \text{TI}, \quad (7)$$

where the Tversky index (TI) is defined as:

$$\text{TI} = \frac{\sum_{i=1}^N p_{0i} g_{0i}}{\sum_{i=1}^N p_{0i} g_{0i} + \alpha \sum_{i=1}^N p_{0i} g_{1i} + \beta \sum_{i=1}^N p_{1i} g_{0i}}, \quad (8)$$

p_{0i} is the probability of pixel i belonging to the foreground class and p_{1i} is the probability of pixel belonging to background class; g_{0i} is 1 for foreground and 0 for background and conversely g_{1i} takes values of 1 for background and 0 for foreground.

Complementary to weighting the positive and negative examples, applying focal parameters to both the Tversky and cross entropy loss enables the downweighting of background objects in favour of foreground object segmentation, and produces the Focal Tversky loss and Focal loss, respectively [54,55]:

$$\mathcal{L}_{\text{FTL}} = (1 - \text{TI})^\gamma, \quad (9)$$

$$\mathcal{L}_{\text{FL}} = -\alpha(1 - p_i)^\gamma \cdot \mathcal{L}_{\text{CE}}, \quad (10)$$

where α controls the class weights.

Finally, we define the Hybrid Focal loss (HFL) as the sum of the Focal Tversky loss and Focal loss:

$$\mathcal{L}_{\text{HFL}} = \mathcal{L}_{\text{FTL}} + \mathcal{L}_{\text{FL}}. \quad (11)$$

To mitigate suppression of the loss near convergence, we supervise the last layer without the focal parameters [55].

4. Materials and evaluation methods

4.1. Dataset descriptions

We assess the ability of the Focus U-Net to accurately segment polyps using five public datasets containing images of polyps taken during optical colonoscopy: CVC-ClinicDB [63], Kvasir-SEG [64], CVC-ColonDB [65], ETIS-Larib PolypDB [66] and EndoScene test set (CVC-T) [67]. CVC-ClinicDB, CVC-ColonDB, ETIS-Larib PolypDB and CVC-T were created by the Hospital Clinic, Universidad de Barcelona, Spain, while the Kvasir-SEG database was produced by the Vestre Viken Health Trust, Norway.

The CVC-ClinicDB database consists of 612 frames containing polyps with image resolution 288×368 pixels, generated from 23 video sequences from 13 different patients using standard optical colonoscopy interventions. The Kvasir-SEG database consists of 1000 polyp images collected and verified by experienced gastroenterologists. Images vary in size from 332×487 to 1920×1072 pixels. CVC-ColonDB consists of 300 images of resolution 500×574 pixels obtained from 15 video sequences with a random sample of 20 frames per sequence. ETIS-Larib PolypDB similarly consists of 300 images, with image resolution 1225×966 pixels. Lastly, CVC-T consists of 182 frames containing polyps from 8 patients derived from either the CVC-ClinicDB and CVC-ColonDB datasets, with image resolutions of 288×384 or 500×574 pixels.

4.2. Experimental setup and implementation details

For our experiments, we use the Medical Imaging Segmentation with Convolutional Neural Networks (MIScnn) open-source Python library [68]. For all datasets, images and associated ground truth masks are provided in the png file format. For the Kvasir-SEG dataset, we resize all images to 512×512 pixels following pre-processing methods used in previous models [30,31], but otherwise resize images to 288×384 pixels for all other datasets. Pixel values are normalised to $[0, 1]$ using the z-score. We perform full-image analysis with a batch size of 16, except for the Kvasir-SEG dataset where the large image sizes required the batch size to be reduced to 8. We use the Focus U-Net architecture as described previously with a final softmax activation layer.

For the ablation studies, we use the CVC-ClinicDB dataset, with five-fold cross validation using random assignment. We evaluate the baseline performance of the U-Net [23] and Attention U-Net [41], and sequentially assess the performance with subsequent additions of the Focus Gate, Hybrid Focal loss, short-range skip connections and deep supervision. Similar to hyperparameter selection of Focal loss functions [54, 55,69], we perform a grid-search, selecting values for the focal parameter $\gamma \in [1, 3]$. Model parameters are initialised with Xavier initialisation, and each model is trained for 100 epochs using Stochastic Gradient Descent with Nesterov momentum ($\mu = 0.99$). We set the initial learning rate at 0.01, and follow a polynomial learning decay rate schedule [70]:

$$\left(\frac{1 - \text{epoch}}{\text{epoch}_{\text{max}}} \right)^{0.9}. \quad (12)$$

For fairer comparison, we do not apply any data augmentation techniques at this stage.

In contrast, when attempting for state-of-the-art results on the CVC-ClinicDB dataset, we train our final model using five-fold cross validation for 500 epochs and use the following data augmentation techniques: scaling, rotation, elastic deformation, mirror and gamma transformations. For both evaluation on the Kvasir-SEG dataset and evaluation on the combination of all five public datasets, we follow the single train-test split used in Refs. [30,31], and train each model for 1000 epochs with the same data augmentation settings for result

reproducibility and comparisons.

For the Hybrid Focal loss, we follow the optimal hyperparameters reported in the original studies. We set $\alpha = 0.3$ and $\beta = 0.7$ for the Tversky index, $\alpha = 0.25$, $\gamma = 2$ for the Focal loss and $\alpha = 0.3$, $\beta = 0.7$ and $\gamma = 3/4$ for the Focal Tversky loss [54,55,62].

For all cases, the validation loss is evaluated at the end of each epoch, and the model with the lowest validation loss is selected as the final model. All experiments are programmed using Keras with TensorFlow backend and trained with NVIDIA P100 GPUs, with CUDA version 10.2 and cuDNN version 7.6.5. Source code is available at: <https://github.com/mlyg/Focus-U-Net>.

4.3. Evaluation metrics

To assess segmentation accuracy, we follow recommendations from Ref. [64], and use DSC and intersection over union (IoU) as the two main metrics. DSC is previously defined in Eq. (5), and IoU is defined as:

$$\text{IoU} = \frac{\text{TP}}{\text{TP} + \text{FP} + \text{FN}} \quad (13)$$

The IoU metric penalises single instances of poor pixel classification more heavily than DSC, providing similar but complementary perspectives on assessing segmentation accuracy. We further assess recall and precision:

$$\text{Recall} = \frac{\text{TP}}{\text{TP} + \text{FN}} \quad (14)$$

$$\text{Precision} = \frac{\text{TP}}{\text{TP} + \text{FP}} \quad (15)$$

In the context of polyp segmentation, recall, also known as sensitivity, measures the proportion of the pixels corresponding to the polyp that are correctly identified. In contrast, precision, also known as the positive predictive value, measures the proportion of pixels correctly labelled as representing the polyp. While both are accounted for in the DSC metric, measuring recall and precision provides additional information regarding the false positive and false negative rates.

Finally, we evaluate model efficiency by calculating the frames per second (FPS) using the mean inference time:

$$\text{FPS} = \frac{1}{\text{mean inference time}} \quad (16)$$

5. Experimental results

We first perform a series of ablation studies to evaluate individual components of the Focus U-Net, followed by separate evaluations on the CVC-ClinicDB and Kvasir-SEG datasets, and finally evaluate against a test set combining five public polyp datasets.

Table 1

Results from training on the CVC-ClinicDB for the U-Net, Attention U-Net and Focus U-Net. The best result is seen with the addition of the Focus Gate, short skip connections, deep supervision and use of the Hybrid Focal loss.

Model	Loss function	Focal parameter γ	Short skip connections	Deep supervision	mDSC	mIoU	Recall	Precision	FPS
U-Net	DSC + CE	–	×	×	0.828 ± 0.021	0.747 ± 0.022	0.817 ± 0.024	0.877 ± 0.021	27
Attention U-Net	DSC + CE	–	×	×	0.801 ± 0.019	0.705 ± 0.023	0.799 ± 0.012	0.844 ± 0.030	25
Focus U-Net	DSC + CE	1	×	g ×	0.838 ± 0.018	0.755 ± 0.018	0.833 ± 0.016	0.876 ± 0.028	25
Focus U-Net	DSC + CE	1.25	×	×	0.844 ± 0.011	0.762 ± 0.011	0.844 ± 0.018	0.876 ± 0.020	25
Focus U-Net	DSC + CE	1.5	×	×	0.842 ± 0.025	0.755 ± 0.027	0.845 ± 0.038	0.866 ± 0.016	25
Focus U-Net	DSC + CE	2	×	×	0.817 ± 0.025	0.728 ± 0.030	0.817 ± 0.023	0.859 ± 0.026	25
Focus U-Net	DSC + CE	3	×	×	0.825 ± 0.022	0.736 ± 0.022	0.820 ± 0.035	0.863 ± 0.022	25
Focus U-Net	DSC + CE	1.25	✓	×	0.867 ± 0.018	0.800 ± 0.011	0.852 ± 0.023	0.908 ± 0.017	25
Focus U-Net	HFL	1.25	✓	×	0.869 ± 0.013	0.797 ± 0.014	0.870 ± 0.017	0.892 ± 0.010	2
Focus U-Net	HFL	1.25	✓	✓	0.875 ± 0.016	0.801 ± 0.017	0.878 ± 0.013	0.889 ± 0.018	25

The results from the ablation study are shown in Table 1.

Performance gains are observed with successive addition of each component, and with all components present there is a significant improvement with a DSC score of 0.875 ± 0.016 compared to the U-Net (0.828 ± 0.021) and Attention U-Net (0.801 ± 0.019). The Focus U-Net achieves similar FPS performance to the Attention U-Net and comparable FPS performance to the U-Net, justifying the improvement in accuracy with minimal efficiency losses.

The results for the CVC-ClinicDB dataset are shown in Table 2.

The Focus U-Net achieves state-of-the-art results with a mDSC score of 0.941 and a mIoU score of 0.893, outperforming the ResUNet++ with Conditional Random Field (CRF) and DoubleU-Net. Focus U-Net also has the best recall-precision balance, while PolypSegNet achieves the highest precision at the cost of recall, and conversely the ResUNet++ with CRF achieves high recall at the cost of precision.

Next, we evaluate our model on the Kvasir-SEG dataset. The results are shown in Table 3. The Focus U-Net achieves state-of-the-art results with a mDSC score of 0.910 and mIoU score of 0.845. The highest mIoU is achieved by HardNet-MSEG [31].

Finally, Table 4 shows the results for the evaluation on five public polyp datasets. It is worth noting that, for a fair comparison, the evaluations metrics are shown only for the approaches that focused on segmentation performance over computational efficiency.

The Focus U-Net achieves the highest score across four of the five datasets with a mDSC of 0.938 and mIoU of 0.889 on the CVC-ClinicDB, mDSC of 0.878 and mIoU of 0.804 for CVC-ColonDB, mDSC of 0.832 and mIoU of 0.757 for ETIS-LaribPolypDB, mDSC of 0.920 and mIoU of 0.860 for CVC-T and mDSC of 0.910 and mIoU of 0.853 for the Kvasir-SEG dataset. Importantly, the combined score takes into account the relative number of images of each dataset, and the Focus U-Net achieves a mDSC of 0.878 and mIoU of 0.702, a 14% increase in mDSC over the

Table 2

Results for the CVC-ClinicDB dataset. Boldface numbers denote the highest values for each metrics.

Model	mDSC	mIoU	Recall	Precision
FCN-8S [21]	0.810	–	0.748	0.883
Multi-scale patch-based CNN [71]	0.813	–	0.786	0.809
FCN [72]	0.830	–	0.773	0.900
MultiResUNet [59]	–	0.821	–	–
cGAN [73]	0.872	0.795	–	–
U-Net [23]	0.878	0.788	0.787	0.933
Multiple encoder-decoder network [74]	0.889	0.894	–	–
PraNet [30]	0.898	0.840	–	–
PolypSegNet [28]	0.915	0.862	0.911	0.962
ResUNet++ with CRF [26]	0.920	0.890	0.939	0.846
Double U-Net [27]	0.924	0.861	0.846	0.959
FANet [34]	0.934	0.894	0.934	0.940
Focus U-Net	0.941	0.893	0.956	0.930

Table 3
Results for the Kvasir-SEG dataset. Boldface numbers denote the highest values for each metrics.

Model	mDSC	mIoU	Recall	Precision
U-Net [23]	0.715	0.433	0.631	0.922
Double U-Net [27]	0.813	0.733	0.840	0.861
FCN8 (VGG16 backbone) [20]	0.831	0.737	0.835	0.882
PSPNet (ResNet50 backbone) [75]	0.841	0.744	0.836	0.890
HRNet [76]	0.845	0.759	0.859	0.878
ResUNet++ with CRF [26]	0.851	0.833	0.876	0.823
DeepLabv3+ (ResNet101 backbone) [77]	0.864	0.786	0.859	0.906
U-Net (ResNet34 backbone) [23]	0.876	0.810	0.944	0.862
FANet [34]	0.880	0.81	0.906	0.901
PolypSegNet [28]	0.887	0.826	0.925	0.917
HarDNet-MSEG [31]	0.904	0.848	0.923	0.907
Focus U-Net	0.910	0.845	0.916	0.917

previous state-of-the-art (HarDNet-MSEG, mDSC = 0.768) and 15% increase in mIoU (HarDNet-MSEG, mIoU = 0.702). The greatest improvements are observed in the most challenging datasets, namely the

Table 4
Results for evaluation on five public polyp datasets. The combined score takes into account the number of images used for each dataset, enabling cross dataset comparisons.

Model	CVC-ClinicDB (61)		CVC-ColonDB (380)		ETIS-LaribPolypDB (196)		CVC-T (60)		Kvasir-SEG (100)		Combined Score	
	mDSC	mIoU	mDSC	mIoU	mDSC	mIoU	mDSC	mIoU	mDSC	mIoU	mDSC	mIoU
SFA [78]	0.700	0.607	0.469	0.347	0.297	0.217	0.467	0.329	0.723	0.611	0.476	0.367
U-Net++ [24]	0.794	0.729	0.483	0.410	0.401	0.344	0.707	0.624	0.821	0.743	0.546	0.476
ResUNet++ [25]	0.796	0.796	-	-	-	-	-	-	0.813	0.793	-	-
U-Net [23]	0.823	0.755	0.512	0.444	0.398	0.335	0.710	0.627	0.818	0.746	0.561	0.493
PraNet [30]	0.899	0.849	0.709	0.640	0.628	0.567	0.871	0.797	0.898	0.840	0.740	0.675
HarDNet-MSEG [31]	0.932	0.882	0.731	0.660	0.677	0.613	0.887	0.821	0.912	0.857	0.768	0.702
Focus U-Net	0.938	0.889	0.878	0.804	0.832	0.757	0.920	0.860	0.910	0.853	0.878	0.809

CVC-ColonDB and ETIS-LaribPolypDB datasets. For CVC-ColonDB, a 20% increase in mDSC is observed over the previous state-of-the-art (HarDNet-MSEG, mDSC = 0.731) and an 87% increase in mDSC compared to the previously top performing Selective Feature Aggregation (SFA) model. Even more significantly, for ETIS-LaribPolypDB, a 23% increase in mDSC is observed over the previous state-of-the-art (HarDNET-MSEG, mDSC = 0.677) and a 180% increase in mDSC compared to SFA.

Examples of polyp segmentations for the five datasets are shown in Fig. 5.

The accuracy of segmentations obtained from the intermediate layer highlights the ability for the deepest layers to localise the polyp effectively. The Focus U-Net generalises well with consistently accurate segmentations across all datasets. For the images corresponding to the poorest segmentation quality, these are either objectively challenging polyps to identify, or in many cases poor-quality images such as in the CVC-ColonDB example.

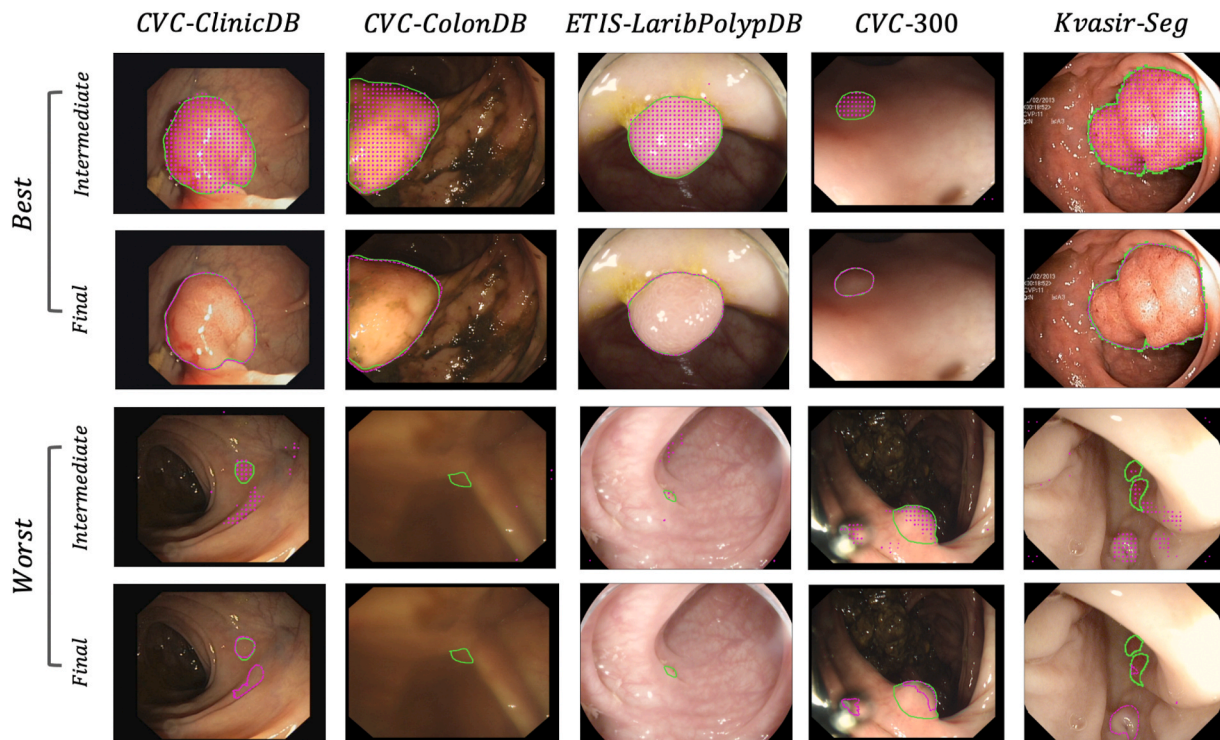


Fig. 5. Examples of the best and worst cases in terms of DSC from each of the five public datasets. The solid green line represents the ground truth mask. The dashed magenta line corresponds to the predictions yielded by the Focus U-Net. The intermediate prediction is derived from the deepest layer of the network, while the final prediction is used for evaluation.

6. Discussion and conclusion

In this paper, we introduce a novel dual attention-gated U-Net architecture, named the Focus U-Net, which uses a Focus Gate to encourage learning of salient regions combined with a focal parameter controlling suppression of irrelevant background regions. Moreover, with the additions of short-range skip connections and deep supervision, as well as optimisation based on the Hybrid Focal loss, the Focus U-Net outperforms the state-of-the-art results across five public polyp datasets. Importantly, the proposed architecture performs consistently well across all datasets, demonstrating an ability to generalise to unseen data from different datasets. Visualising the resulting polyp segmentations confirms the segmentation quality, with poorer segmentations associated with a combination of either poorer image quality or objectively more challenging polyps to identify.

The proposed Focus U-Net is the latest addition to lightweight, yet accurate, polyp segmentation models, achieving state-of-the-art results with a mDSC of 0.878 and mIoU score of 0.809 when evaluated on the combination of five public datasets, a 14% and 15% improvement over the previous state-of-the-art results from HarDNet-MSEG with a mDSC of 0.768 and mIoU of 0.702, respectively.

While these results are promising, it is important to determine whether such a model may be applied in clinical practice. Given that colonoscopy involves recordings of live video, a model with a fast inference time is required to process images in real-time. Accordingly, the Focus U-Net architecture is efficiently designed, with both efficient channel and spatial attention mechanisms, as well as a lightweight U-Net backbone, with comparable FPS performance to the standard U-Net [79]. With polyp miss rates as high as 26% reported for small adenomas [9], the primary advantage of AI-assisted colonoscopy is to aid clinicians in reducing polyp miss-rate detection. However, a secondary advantage with segmentation-based computer-aided detection is providing an accurate and operator-independent estimate of the polyp size; an important factor in guiding biopsy decisions that may be required during colonoscopy.

There are several limitations associated with our current study. Firstly, the datasets used to train our model consist of images all containing polyps, in contrast to in practice, where the majority of live video data would not contain a polyp. However, in terms of model training, it has been observed that training with images in the absence of polyps results in poorer generalisation [80]. In terms of model performance, we would expect a higher false positive rate. This is not as undesirable as the converse of a high false negative rate, because the purpose of the computer-aided system is to focus the operator to attend to highlighted regions that may contain missed polyps [81].

While colonoscopy remains the gold-standard for investigating suspected CRC, CT virtual colonography is a relatively newer method for bowel cancer screening that offers non-invasive visualisation of the colon [82]. The flexibility of our model does not restrict usage to polyps in visible light and is equally applicable for polyp detection using CT colonography. However, these newer modalities present additional challenges, such as fluid submersion obscuring polyps [83]. In fact, the scope for using the Focus U-Net architecture is not limited for colorectal polyps, and is applicable for any image segmentation problem where there is the issue of class imbalance and requirement for efficiency.

Declaration of competing interest

All authors in this paper do not have any conflict of interest.

Acknowledgements

This work was partially supported by The Mark Foundation for Cancer Research and Cancer Research UK Cambridge Centre [C9685/A25177], the CRUK National Cancer Imaging Translational Accelerator (NCITA) [C42780/A27066] and the Wellcome Trust Innovator Award,

United Kingdom [215733/Z/19/Z]. Additional support was also provided by the National Institute of Health Research (NIHR) Cambridge Biomedical Research Centre [BRC-1215-20014] and the Cambridge Mathematics of Information in Healthcare (CMIH) [funded by the EPSRC grant EP/T017961/1]. The views expressed are those of the authors and not necessarily those of the NHS, the NIHR, or the Department of Health and Social Care.

CBS in addition acknowledges support from the Leverhulme Trust project on 'Breaking the non-convexity barrier', the Philip Leverhulme Prize, the Royal Society Wolfson Fellowship, the EPSRC grants EP/S026045/1, EP/N014588/1, European Union Horizon 2020 research and innovation programmes under the Marie Skłodowska-Curie grant agreement No. 777826 NoMADS and No. 691070 CHiPS, the Cantab Capital Institute for the Mathematics of Information and the Alan Turing Institute.

This work was performed using resources provided by the Cambridge Service for Data Driven Discovery (CSD3) operated by the University of Cambridge Research Computing Service (www.csd3.cam.ac.uk), provided by Dell EMC and Intel using Tier-2 funding from the Engineering and Physical Sciences Research Council (capital grant EP/P020259/1), and DIRAC funding from the Science and Technology Facilities Council (www.dirac.ac.uk).

References

- [1] F. Bray, J. Ferlay, I. Soerjomataram, R.L. Siegel, L.A. Torre, A. Jemal, Global cancer statistics 2018: GLOBOCAN estimates of incidence and mortality worldwide for 36 cancers in 185 countries, *CA A Cancer J. Clin.* 68 (6) (2018) 394–424.
- [2] P. Rawla, T. Sunkara, A. Barsouk, Epidemiology of colorectal cancer: incidence, mortality, survival, and risk factors, *Przegląd Gastroenterol.* 14 (2) (2019) 89.
- [3] B. Vogelstein, E.R. Fearon, S.R. Hamilton, S.E. Kern, A.C. Preisinger, M. Leppert, A. M. Smits, J.L. Bos, Genetic alterations during colorectal-tumor development, *N. Engl. J. Med.* 319 (9) (1988) 525–532.
- [4] K.M. Kuntz, I. Lansdorp-Vogelaar, C.M. Rutter, A.B. Knudsen, M. Van Ballegooijen, J.E. Savarino, E.J. Feuer, A.G. Zauber, A systematic comparison of microsimulation models of colorectal cancer: the role of assumptions about adenoma progression, *Med. Decis. Making* 31 (4) (2011) 530–539.
- [5] H. Brenner, M. Hoffmeister, C. Stegmaier, G. Brenner, L. Altenhofen, U. Haug, Risk of progression of advanced adenomas to colorectal cancer by age and sex: estimates based on 840 149 screening colonoscopies, *Gut* 56 (11) (2007) 1585–1589.
- [6] E.H. Schreuders, A. Ruco, L. Rabeneck, R.E. Schoen, J.J. Sung, G.P. Young, E. J. Kuipers, Colorectal cancer screening: a global overview of existing programmes, *Gut* 64 (10) (2015) 1637–1649.
- [7] M. Arnold, M.S. Sierra, M. Laversanne, I. Soerjomataram, A. Jemal, F. Bray, Global patterns and trends in colorectal cancer incidence and mortality, *Gut* 66 (4) (2017) 683–691.
- [8] I.A. Issa, M. Noureddine, Colorectal cancer screening: an updated review of the available options, *World J. Gastroenterol.* 23 (28) (2017) 5086.
- [9] N.H. Kim, Y.S. Jung, W.S. Jeong, H.-J. Yang, S.-K. Park, K. Choi, D.I. Park, Miss rate of colorectal neoplastic polyps and risk factors for missed polyps in consecutive colonoscopies, *Int. Res.* 15 (3) (2017) 411.
- [10] J.C. Van Rijn, J.B. Reitsma, J. Stoker, P.M. Bossuyt, S.J. Van Deventer, E. Dekker, Polyp miss rate determined by tandem colonoscopy: a systematic review, *Am. J. Gastroenterol.* 101 (2) (2006) 343–350.
- [11] A. Leufkens, M. Van Oijen, F. Vleggaar, P. Siersema, Factors influencing the miss rate of polyps in a back-to-back colonoscopy study, *Endoscopy* 44 (2012) 470–475, 05.
- [12] C. Le Berre, W.J. Sandborn, S. Aridhi, M.-D. Devignes, L. Fournier, M. Smail-Tabbone, S. Danese, L. Peyrin-Biroulet, Application of artificial intelligence to gastroenterology and hepatology, *Gastroenterology* 158 (1) (2020) 76–94.
- [13] S. Thakkar, N.M. Carleton, B. Rao, A. Syed, Use of artificial intelligence-based analytics from live colonoscopies to optimize the quality of the colonoscopy examination in real time: proof of concept, *Gastroenterology* 158 (5) (2020) 1219–1221.
- [14] S. Krishnan, X. Yang, K. Chan, S. Kumar, P. Goh, Intestinal abnormality detection from endoscopic images, in: Proc. 20th Annual International Conference of the IEEE Engineering in Medicine and Biology Society (EMBC), vol. 2, IEEE, 1998, pp. 895–898.
- [15] S. Hwang, J. Oh, W. Tavanapong, J. Wong, P.C. De Groen, Polyp detection in colonoscopy video using elliptical shape feature, in: Proc. IEEE International Conference on Image Processing (ICIP), vol. 2, IEEE, 2007, pp. II–465.
- [16] J. Bernal, J. Sánchez, F. Vilarino, Towards automatic polyp detection with a polyp appearance model, *Pattern Recogn.* 45 (9) (2012) 3166–3182.
- [17] S.A. Karkanis, D.K. Iakovidis, D.E. Maroulis, D.A. Karras, M. Tzivras, Computer-aided tumor detection in endoscopic video using color wavelet features, *IEEE Trans. Inf. Technol. Biomed. e* 7 (3) (2003) 141–152.

- [18] M.T. Coimbra, J.S. Cunha, MPEG-7 visual descriptors—contributions for automated feature extraction in capsule endoscopy, *IEEE Trans. Circ. Syst. Video Technol.* 16 (5) (2006) 628–637.
- [19] L. Yu, H. Chen, Q. Dou, J. Qin, P.A. Heng, Integrating online and offline three-dimensional deep learning for automated polyp detection in colonoscopy videos, *IEEE J. Biomed. Health Inform.* 21 (1) (2016) 65–75.
- [20] J. Long, E. Shelhamer, T. Darrell, Fully convolutional networks for semantic segmentation, in: *Proc. IEEE/CVF Conference on Computer Vision and Pattern Recognition, CVPR*, 2015, pp. 3431–3440.
- [21] M. Akbari, M. Mohrekeh, E. Nasr-Esfahani, S.R. Soroushmehr, N. Karimi, S. Samavi, K. Najarian, Polyp segmentation in colonoscopy images using fully convolutional network, in: *Proc. 40th Annual International Conference of the IEEE Engineering in Medicine and Biology Society (EMBC)*, IEEE, 2018, pp. 69–72.
- [22] P. Brandao, E. Mazomenos, G. Ciuti, R. Calió, F. Bianchi, A. Mencias, P. Dario, A. Koulaouzidis, A. Arezzo, D. Stoyanov, Fully convolutional neural networks for polyp segmentation in colonoscopy, in: *Proc. Medical Imaging 2017: Computer-Aided Diagnosis*, vol. 10134, International Society for Optics and Photonics, 2017, p. 101340F.
- [23] O. Ronneberger, P. Fischer, T. Brox, U-Net, Convolutional networks for biomedical image segmentation, in: *Proc. International Conference on Medical Image Computing and Computer-Assisted Intervention (MICCAI)*, Springer, 2015, pp. 234–241.
- [24] Z. Zhou, M.M.R. Siddiquee, N. Tajbakhsh, J. Liang, UNet++: redesigning skip connections to exploit multiscale features in image segmentation, *IEEE Trans. Med. Imag.* 39 (6) (2019) 1856–1867.
- [25] D. Jha, P.H. Smedsrud, M.A. Riegler, D. Johansen, T. De Lange, P. Halvorsen, H. D. Johansen, ResUNet++: an advanced architecture for medical image segmentation, in: *Proc. IEEE International Symposium on Multimedia (ISM)*, IEEE, 2019, pp. 225–2255.
- [26] D. Jha, P.H. Smedsrud, D. Johansen, T. de Lange, H.D. Johansen, P. Halvorsen, M. A. Riegler, A comprehensive study on colorectal polyp segmentation with ResUNet ++, conditional random field and test-time augmentation, *IEEE J. Biomed. Health Inform.* 25 (6) (2021) 2029–2040.
- [27] D. Jha, M.A. Riegler, D. Johansen, P. Halvorsen, H.D. Johansen, DoubleU-Net, A deep convolutional neural network for medical image segmentation, in: *Proc. IEEE 33rd International Symposium on Computer-Based Medical Systems (CBMS)*, IEEE, 2020, pp. 558–564.
- [28] T. Mahmud, B. Paul, S.A. Fattah, PolypSegNet: a modified encoder-decoder architecture for automated polyp segmentation from colonoscopy images, *Comput. Biol. Med.* 128 (2021) 104119.
- [29] D. Jha, S. Ali, N.K. Tomar, H.D. Johansen, D. Johansen, J. Rittscher, M.A. Riegler, P. Halvorsen, Real-time polyp detection, localization and segmentation in colonoscopy using deep learning, *IEEE Access* 9 (2021) 40496–40510, <https://doi.org/10.1109/ACCESS.2021.3063716>.
- [30] D.-P. Fan, G.-P. Ji, T. Zhou, G. Chen, H. Fu, J. Shen, L. Shao, PraNet: parallel reverse attention network for polyp segmentation, in: *Proc. International Conference on Medical Image Computing and Computer-Assisted Intervention (MICCAI)*, Springer, 2020, pp. 263–273.
- [31] C.-H. Huang, H.-Y. Wu, Y.-L. Lin, HardNet-MSEG: A Simple Encoder-Decoder Polyp Segmentation Neural Network that Achieves over 0.9 Mean Dice and 86 FPS, *arXiv preprint arXiv:2101.07172*.
- [32] P. Chao, C.-Y. Kao, Y.-S. Ruan, C.-H. Huang, Y.-L. Lin, HardNet: a low memory traffic network, in: *Proc. IEEE International Conference on Computer Vision, ICCV*, 2019, pp. 3552–3561.
- [33] G.-P. Ji, Y.-C. Chou, D.-P. Fan, G. Chen, H. Fu, D. Jha, L. Shao, Progressively Normalized Self-Attention Network for Video Polyp Segmentation, *arXiv preprint arXiv:2105.08468*.
- [34] N. K. Tomar, D. Jha, M. A. Riegler, H. D. Johansen, D. Johansen, J. Rittscher, P. Halvorsen, S. Ali, FANet: A Feedback Attention Network for Improved Biomedical Image Segmentation, *arXiv preprint arXiv:2103.17235*.
- [35] C. Han, L. Rundo, K. Murao, T. Noguchi, Y. Shimahara, Z.Á. Milacski, S. Koshino, E. Sala, H. Nakayama, S. Satoh, MADGAN: unsupervised medical anomaly detection GAN using multiple adjacent brain MRI slice reconstruction, *BMC Bioinf.* 22 (2021) 31.
- [36] S. Chen, Y. Zou, P.X. Liu, IBA-U-Net, Attentive BConvLSTM U-Net with redesigned inception for medical image segmentation, *Comput. Biol. Med.* (2021) 104551.
- [37] G. Tagliarini, E. Page, A neural-network solution to the concentrator assignment problem, in: *Proc. Neural Information Processing Systems, NIPS*, 1987, pp. 775–782.
- [38] S. Jetley, N. A. Lord, N. Lee, P. H. Torr, Learn to Pay Attention, *arXiv preprint arXiv:1804.02391*.
- [39] F. Wang, M. Jiang, C. Qian, S. Yang, C. Li, H. Zhang, X. Wang, X. Tang, Residual attention network for image classification, in: *Proc. IEEE/CVF Conference Computer Vision and Pattern Recognition, CVPR*, 2017, pp. 3156–3164.
- [40] M.-T. Luong, H. Pham, C. D. Manning, Effective Approaches to Attention-Based Neural Machine Translation, *arXiv preprint arXiv:1508.04025*.
- [41] J. Schlemper, O. Oktay, M. Schaap, M. Heinrich, B. Kainz, B. Glocker, D. Rueckert, Attention gated networks: learning to leverage salient regions in medical images, *Med. Image Anal.* 53 (2019) 197–207.
- [42] D. Bahdanau, K. Cho, Y. Bengio, Neural Machine Translation by Jointly Learning to Align and Translate, *arXiv preprint arXiv:1409.0473*.
- [43] T. Shen, T. Zhou, G. Long, J. Jiang, S. Pan, C. Zhang, DISAN: directional self-attention network for rnn/cnn-free language understanding, in: *Proc. AAAI Conference on Artificial Intelligence*, vol. 32, 2018, pp. 5446–5455.
- [44] A. Vaswani, N. Shazeer, N. Parmar, J. Uszkoreit, L. Jones, A. N. Gomez, L. Kaiser, I. Polosukhin, Attention Is All You Need, *arXiv preprint arXiv:1706.03762*.
- [45] O. Oktay, J. Schlemper, L. L. Folgoc, M. Lee, M. Heinrich, K. Misawa, K. Mori, S. McDonagh, N. Y. Hammerla, B. Kainz, et al., Attention U-Net: Learning where to Look for the Pancreas, *arXiv preprint arXiv:1804.03999*.
- [46] J. Hu, L. Shen, G. Sun, Squeeze-and-excitation networks, in: *Proc. IEEE/CVF Conference on Computer Vision and Pattern Recognition, CVPR*, 2018, pp. 7132–7141.
- [47] J. Fu, J. Liu, H. Tian, Y. Li, Y. Bao, Z. Fang, H. Lu, Dual attention network for scene segmentation, in: *Proc. IEEE/CVF Conference on Computer Vision and Pattern Recognition, CVPR*, 2019, pp. 3146–3154.
- [48] S. Woo, J. Park, J.-Y. Lee, I.S. Kweon, CBAM: convolutional block attention module, in: *Proc. European Conference on Computer Vision, ECCV*, 2018, pp. 3–19.
- [49] Z. Gao, J. Xie, Q. Wang, P. Li, Global second-order pooling convolutional networks, in: *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, 2019, pp. 3024–3033.
- [50] L. Rundo, C. Han, Y. Nagano, J. Zhang, R. Hataya, C. Militello, A. Tangherloni, M. S. Nobile, C. Ferretti, D. Besozzi, et al., USE-Net: incorporating squeeze-and-excitation blocks into U-Net for prostate zonal segmentation of multi-institutional MRI datasets, *Neurocomputing* 365 (2019) 31–43.
- [51] Q. Wang, B. Wu, P. Zhu, P. Li, W. Zuo, Q. Hu, ECA-Net, Efficient channel attention for deep convolutional neural networks, in: *Proc. IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*, IEEE, 2020, pp. 11531–11539, <https://doi.org/10.1109/CVPR42600.2020.01155>.
- [52] A.G. Roy, N. Navab, C. Wachinger, Concurrent spatial and channel ‘squeeze & excitation’ in fully convolutional networks, in: *Proc. Int. Conference on Medical Image Computing and Computer-Assisted Intervention (MICCAI)*, Springer, 2018, pp. 421–429.
- [53] M.A. Goodale, A.D. Milner, Separate visual pathways for perception and action, *Trends Neurosci.* 15 (1) (1992) 20–25.
- [54] T.-Y. Lin, P. Goyal, R. Girshick, K. He, P. Dollár, Focal loss for dense object detection, in: *Proc. IEEE International Conference on Computer Vision, ICCV*, 2017, pp. 2980–2988.
- [55] N. Abraham, N.M. Khan, A novel focal Tversky loss function with improved attention U-Net for lesion segmentation, in: *Proc. IEEE 16th International Symposium on Biomedical Imaging (ISBI)*, IEEE, 2019, pp. 683–687.
- [56] X. Glorot, Y. Bengio, Understanding the difficulty of training deep feedforward neural networks, in: *Proc. 13th Int. Conference on Artificial Intelligence and Statistics (AISTATS)*, JMLR Workshop and Conference Proceedings, 2010, pp. 249–256.
- [57] R. Pascanu, T. Mikolov, Y. Bengio, On the difficulty of training recurrent neural networks, in: *Proc. Int. Conference on Machine Learning (ICML)*, PMLR, 2013, pp. 1310–1318.
- [58] C. Szegedy, V. Vanhoucke, S. Ioffe, J. Shlens, Z. Wojna, Rethinking the inception architecture for computer vision, in: *Proc. IEEE Conference on Computer Vision and Pattern Recognition, CVPR*, 2016, pp. 2818–2826, <https://doi.org/10.1109/CVPR.2016.308>.
- [59] N. Irbetaz, M.S. Rahman, MultiResUNet: rethinking the U-Net architecture for multimodal biomedical image segmentation, *Neural Network*. 121 (2020) 74–87.
- [60] C.-Y. Lee, S. Xie, P. Gallagher, Z. Zhang, Z. Tu, Deeply-supervised nets, in: *Proc. 18th Int. Conference on Artificial Intelligence and Statistics (AISTATS)*, PMLR, 2015, pp. 562–570.
- [61] Q. Dou, L. Yu, H. Chen, Y. Jin, X. Yang, J. Qin, P.-A. Heng, 3D deeply supervised network for automated segmentation of volumetric medical images, *Med. Image Anal.* 41 (2017) 40–54.
- [62] S.S.M. Salehi, D. Erdogmus, A. Gholipour, Tversky loss function for image segmentation using 3D fully convolutional deep networks, in: *Proc. Int. Workshop on Machine Learning in Medical Imaging (MLMI)*, Springer, 2017, pp. 379–387.
- [63] J. Bernal, F.J. Sánchez, G. Fernández-Esparrach, D. Gil, C. Rodríguez, F. Vilariño, WM-DOVA maps for accurate polyp highlighting in colonoscopy: validation vs. saliency maps from physicians, *Comput. Med. Imag. Graph.* 43 (2015) 99–111.
- [64] D. Jha, P.H. Smedsrud, M.A. Riegler, P. Halvorsen, T. de Lange, D. Johansen, H. D. Johansen, S.E.G. Kvasir, A segmented polyp dataset, in: *Proc. International Conference on Multimedia Modeling (MMM)*, Springer, 2020, pp. 451–462.
- [65] N. Tajbakhsh, S.R. Gurudu, J. Liang, Automated polyp detection in colonoscopy videos using shape and context information, *IEEE Trans. Med. Imag.* 35 (2) (2015) 630–644.
- [66] J. Silva, A. Histace, O. Romain, X. Dray, B. Granado, Toward embedded detection of polyps in WCE images for early diagnosis of colorectal cancer, *Int. J. Comput. Assist. Radiol. Surg.* 9 (2) (2014) 283–293.
- [67] D. Vázquez, J. Bernal, F.J. Sánchez, G. Fernández-Esparrach, A.M. López, A. Romero, M. Drozdal, A. Courville, A benchmark for endoluminal scene segmentation of colonoscopy images, *J. Healthc. Eng.* (2017).
- [68] D. Müller, F. Kramer, MIScnn: a framework for medical image segmentation with convolutional neural networks and deep learning, *BMC Med. Imag.* 21 (1) (2021) 1–11.
- [69] M. Yeung, E. Sala, C.-B. Schönlieb, L. Rundo, Unified Focal Loss: Generalising Dice and Cross Entropy-Based Losses to Handle Class Imbalanced Medical Image Segmentation, *arXiv preprint arXiv:2102.04525*.
- [70] F. Isensee, P.F. Jaeger, S.A. Kohl, J. Petersen, K.H. Maier-Hein, nnU-Net: a self-configuring method for deep learning-based biomedical image segmentation, *Nat. Methods* 18 (2) (2021) 203–211.
- [71] D. Banik, D. Bhattacharjee, M. Nasipuri, A multi-scale patch-based deep learning system for polyp segmentation, in: *Advanced Computing and Systems for Security*, Springer, 2020, pp. 109–119.
- [72] Q. Li, G. Yang, Z. Chen, B. Huang, L. Chen, D. Xu, X. Zhou, S. Zhong, H. Zhang, T. Wang, Colorectal polyp segmentation using a fully convolutional neural

- network, in: Proc. 10th Int. Congress on Image and Signal Processing, Biomedical Engineering and Informatics (CISP-BMED), IEEE, 2017, pp. 1–5.
- [73] J. Poorneshwaran, S.S. Kumar, K. Ram, J. Joseph, M. Sivaprakasam, Polyp segmentation using generative adversarial network, in: Proc. 41st Annual International Conference of the IEEE Engineering in Medicine and Biology Society (EMBC), IEEE, 2019, pp. 7201–7204.
- [74] Q. Nguyen, S.-W. Lee, Colorectal segmentation using multiple encoder-decoder network in colonoscopy images, in: Proc. IEEE First Int. Conference on Artificial Intelligence and Knowledge Engineering (AIKE), IEEE, 2018, pp. 208–211.
- [75] H. Zhao, J. Shi, X. Qi, X. Wang, J. Jia, Pyramid scene parsing network, in: Proc. IEEE/CVF Conference on Computer Vision and Pattern Recognition, CVPR, 2017, pp. 2881–2890.
- [76] J. Wang, K. Sun, T. Cheng, B. Jiang, C. Deng, Y. Zhao, D. Liu, Y. Mu, M. Tan, X. Wang, et al., Deep high-resolution representation learning for visual recognition, *IEEE Trans. Pattern Anal. Mach. Intell.*
- [77] L.-C. Chen, Y. Zhu, G. Papandreou, F. Schroff, H. Adam, Encoder-decoder with atrous separable convolution for semantic image segmentation, in: Proc. European Conference on Computer Vision, ECCV, 2018, pp. 801–818.
- [78] Y. Fang, C. Chen, Y. Yuan, K.-y. Tong, Selective feature aggregation network with area-boundary constraints for polyp segmentation, in: Proc. Int. Conference on Medical Image Computing and Computer-Assisted Intervention (MICCAI), Springer, 2019, pp. 302–310.
- [79] Z. Wang, Y. Zou, P.X. Liu, Hybrid dilation and attention residual U-Net for medical image segmentation, *Comput. Biol. Med.* 134 (2021) 104449.
- [80] P. Brandao, O. Zisimopoulos, E. Mazomenos, G. Ciuti, J. Bernal, M. Visentini-Scarzanella, A. Menciassi, P. Dario, A. Koulaouzidis, A. Arezzo, et al., Towards a computed-aided diagnosis system in colonoscopy: automatic polyp segmentation using convolution neural networks, *J. Med. Robot. Res.* 3 (2018) 1840002, 02.
- [81] Y. Mori, S.-e. Kudo, M. Misawa, K. Takeda, T. Kudo, H. Itoh, M. Oda, K. Mori, Artificial intelligence for colorectal polyp detection and characterization, *Curr. Treat. Options Gastroenterol.* 18 (2) (2020) 200–211, <https://doi.org/10.1007/s11938-020-00287-x>.
- [82] R.M. Summers, J. Yao, P.J. Pickhardt, M. Franaszek, I. Bitter, D. Brickman, V. Krishna, J.R. Choi, Computed tomographic virtual colonoscopy computer-aided polyp detection in a screening population, *Gastroenterology* 129 (6) (2005) 1832–1844.
- [83] C. Bräuer, P. Lefere, S. Gryspeerdt, H. Ringl, A. Al-Mukhtar, P. Apfaltrer, D. Berzaczy, B. Fügler, J. Furtner, C. Müller-Mang, et al., CT colonography: size reduction of submerged colorectal polyps due to electronic cleansing and CT-window settings, *Eur. Radiol.* 28 (11) (2018) 4766–4774, <https://doi.org/10.1007/s00330-018-5416-0>.