UNCERTAINTY MODELING AND INTERPRETABILITY IN CONVOLUTIONAL NEURAL NETWORKS FOR POLYP SEGMENTATION

Kristoffer Wickstrøm*, Michael Kampffmeyer, Robert Jenssen

UiT The Arctic University of Norway
Dept. of Physics and Technology
UiT Machine Learning Group

ABSTRACT

Convolutional Neural Networks (CNNs) are propelling advances in a range of different computer vision tasks such as object detection and object segmentation. Their success has motivated research in applications of such models for medical image analysis. If CNN-based models are to be helpful in a medical context, they need to be precise, interpretable, and uncertainty in predictions must be well understood. In this paper, we develop and evaluate recent advances in uncertainty estimation and model interpretability in the context of semantic segmentation of polyps from colonoscopy images. We evaluate and enhance several architectures of Fully Convolutional Networks (FCNs) for semantic segmentation of colorectal polyps and provide a comparison between these models. Our highest performing model achieves a 76.06% mean IOU accuracy on the EndoScene dataset, a considerable improvement over the previous state-of-the-art.

Index Terms— Polyp Segmentation, Deep Learning, Fully Convolutional Networks, Uncertainty Modeling, CNN interpretability

1. INTRODUCTION

Colon cancer prevention is currently primarily done with the help of regular colonoscopy screenings. However, depending on the size and type, roughly $8 - 37\%$ of polyps are missed during the process [1]. Missed polyps can have fatal consequences, as they are potential precursors to colon cancer, which causes the third most cancer deaths globally [2]. Hence, increasing the detection rate of polyps is an important topic of research. Towards this end, automated detection procedures have been proposed [3, 4], which have the advantage of not being influenced by factors such as the fatigue of medical personnel towards the end of long operations. However, they present additional novel challenges. For instance,

to enable effective use of such methods, medical staff should be able to understand why the model believes that a specific region contains a polyp (interpretability) and there should be some notion of uncertainty in predictions.

The last couple of years have seen several works on automatic procedures based on Deep Convolutional Neural Networks (CNN) for health domain tasks such as interstitial lung disease classification [5], cell detection [6], estimation of cardiothoracic ratio [7], and polyp detection [8]. CNNs have improved the state of the art in a number of computer vision tasks such as image classification [9], object detection [10] and semantic segmentation [11, 12]. Recently, CNNs have shown promising performance for the task of polyp segmentation [2]. However, despite the promising results obtained on polyp segmentation, model interpretability and uncertainty quantification have been lacking, and recent advances in deep learning have not been incorporated [2].

In this paper, we enhance and evaluate two recent CNN architectures for pixel-to-pixel based segmentation of colorectal polyps, referred to as Fully Convolutional Networks (FCNs). Furthermore, we incorporate and develop recent advances in the field of deep learning for semantic segmentation of colorectal polyps in order to model uncertainty and increase model interpretability leading to novel uncertainty maps [13, 14] in a polyp segmentation context as well as the visualization of descriptive regions in the input image using Guided Backpropagation [15]. To the author's knowledge, these techniques have not been previously explored in the field of semantic segmentation of colorectal polyps.

2. ENHANCED FULLY CONVOLUTIONAL NETWORKS FOR POLYP SEGMENTATION

We choose two architectures for the task of polyp segmentation, namely the Fully Convolutional Network 8 (FCN-8) [11] and the more recent SegNet [16]. Previous use of FCNs for polyp segmentation have shown promising results, and we hypothesized that the inclusion of recent advances in deep learning would improve these results further. SegNet has been shown to achieve comparative results to FCNs in some appli-

cation domains but is a less memory intensive approach with fewer parameters to optimize.

The FCN is a CNN architecture particularly well suited to tackle per pixel prediction problems like semantic segmentation. FCNs employ an encoder-decoder architecture and are capable of end-to-end learning. The encoder extracts useful features from an image and maps it to a low-resolution representation. The decoder is tasked with mapping the low-resolution representation back into the same resolution as the input image. Upsampling in FCNs is performed using either bi-linear interpolation or by transposed (or fractional strided) convolutions, where the convolution filters are learned as part of the optimization procedure. Learned upsample filters add additional parameters to the network architecture, but tend to provide better overall results [11]. Upsampling can further be improved by including skip connections, which combine coarse level semantic information with higher resolution segmentation from previous network layers. Due to the lack of fully connected layers, inference can be performed on images of arbitrary size.

The SegNet architecture builds on the general idea of FCNs but proposes a novel approach to upsampling. Instead of learning the upsampling, SegNet utilizes the pooling indices from the encoder to upsample activations in the decoder, thereby producing sparse feature maps. These sparse representations are then processed by additional convolutional layers to produce dense activations and predictions. The advantage of SegNet compared to FCN is the reduction in learnable parameters, only 29.5 vs. 134.5 million, as the upsampling filters in FCNs tend to be large.

For our experiments, we propose an Enhanced FCN-8 (EFCN-8) architecture which leverages recent development in the deep learning field. The architecture is depicted in Figure 1. The second architecture that we propose is an Enhanced variant of SegNet (ESegNet), depicted in Figure 2. To improve overall performance, we propose to include several recent advances in deep learning which were not present in the original architectures. For the FCN-8, we make use of Batch Normalization [17] after each layer. Batch Normalization is a procedure that normalizes the output of each layer, allowing for a larger learning rate that accelerate the training procedure. For ESegNet, we include Dropout [18] after the three central encoders and decoders inspired by [13]. Dropout is a regularization technique that randomly set units in a layer to zero and can be interpreted as an ensemble of several networks. Including Dropout serves two purposes. It regularizes the model which encourage better generalization capabilities and, as we will see in Section 3, enable estimation of uncertainty in the model's prediction. The encoder of both models corresponds to the network VGG16 [19], which allows us to initialize the encoder with pretrained weights from a VGG16 model that was previously trained on the ImageNet dataset, an approach referred to as transfer learning. Utilizing pretrained weights was incorporated in the original architectures,
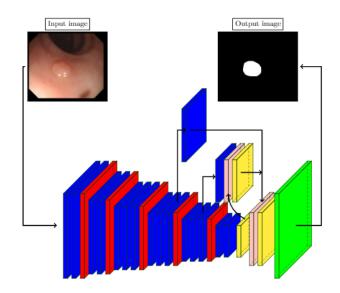

Fig. 1. An illustration of the Enhanced Fully Convolutional Network-8. Color codes description: Blue - Convolution (3x3), Batch Normalization and ReLU; Yellow - Upsampling; Pink - Summing; Red - Pooling (2x2); Green - Soft-max. Dropout was included according to [19].

but not included in recent work on segmentation of colorectal polyp using FCNs [2].

3. UNCERTAINTY AND INTERPRETABILITY IN FULLY CONVOLUTIONAL NETWORKS FOR POLYP SEGMENTATION

Despite their success on a number of different tasks, CNNs are not without their flaws. Most CNNs are unable to provide any notion of uncertainty in their prediction and determining what features in the input affect its prediction is challenging as a result of the complexity of the model. Such limitations have not stopped CNNs from being applied on many computer vision tasks, yet they become especially apparent when developing CNNs for medical applications. Most physicians will be reluctant to make a diagnosis based on a single segmentation map with no notion of uncertainty and indication as to what features were the basis of the prediction. This section will describe two recently proposed methods which address the limitations just outlined.

3.1. Uncertainty

Modeling uncertainty is crucial for designing trustworthy automatic procedures, yet CNNs have no natural way of providing such uncertainties to accompany its prediction. In contrast, Bayesian models provide a framework which naturally includes uncertainty by modeling posterior distribution for the quantities in question. Given a dataset $\mathbf{X} = \{\mathbf{x}_1, ..., \mathbf{x}_N\}$
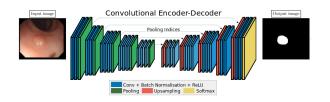
Fig. 2. Depiction of the standard SegNet architecture obtained from [16]. Our implementation include Dropout in the three central encoders and decoders for regularization and to enable uncertainty estimation.

with labels $\mathbf{Y} = \{\mathbf{y}_1, ..., \mathbf{y}_N\}$, the predictive distribution of a Bayesian neural network can be modeled as:

$$p(\mathbf{y}|\mathbf{x}, \mathbf{X}, \mathbf{Y}) = \int p(\mathbf{y}|\mathbf{x}, \mathbf{W})p(\mathbf{W}|\mathbf{x}, \mathbf{X}, \mathbf{Y})d\mathbf{W} \quad (1)$$

where $\mathbf{W}$ refers to the weights of the model, $p(\mathbf{y}|\mathbf{x}, \mathbf{W})$ is the softmax function applied to the output of the model, denoted by $f^{\mathbf{W}}(\mathbf{x})$, and $p(\mathbf{W}|\mathbf{x}, \mathbf{X}, \mathbf{Y})$ is the posterior over the weights which capture the set of plausible models parameters for the given data. Obtaining $p(\mathbf{y}|\mathbf{x}, \mathbf{W})$ only requires a forward pass of the network, but the inability to evaluate the posterior of the weights analytically makes Bayesian neural networks computationally unfeasible. To sidestep the problematic posterior of the weights, [20] proposed to incorporate Dropout [18] as a method for sampling sets of weights from the trained network to approximate the posterior of the weights. The predictive distribution in Equation 1 can then be approximated using Monte Carlo integration as follows:

$$p(\mathbf{y}|\mathbf{x}, \mathbf{X}, \mathbf{Y}) \approx \frac{1}{T}\sum_{t=1}^{T}\text{Softmax}(f^{\widehat{\mathbf{W}}_t}(\mathbf{x})) \quad (2)$$

where T is the number of sampled sets of weights and $\widehat{\mathbf{W}}_t$ is a set of sampled weights. In practice, the predictive distribution from Equation 2 can be estimated by running T forward passes of a model with Dropout applied to produce T predictions, which in turn can be used to estimate the uncertainty associated with the sample in question. The authors of [20] refer to this method of sampling from the posterior of the predictive distribution as Monte Carlo Dropout.

3.2. Interpretability

Another desirable property that CNNs lack is interpretability, i.e. being able to determine what features induces the network to produce a particular prediction. However, several recent works have proposed different methods to increase network interpretability [21, 22]. In this paper, we evaluate and develop the Guided Backpropagation [15] technique for FCNs on the task of semantic segmentation of colorectal polyp in order to assess which pixels in the input image the network

deems important for identifying polyps. We choose Guided Backpropagation as it is known to produce clear visualizations of salient input pixels and is more straightforward to employ compared to other methods.

The central idea of Guided Backpropagation is the interpretation of the gradients of the network with respect to an input image. In [23] they noted that, for a given image, the magnitude of the gradients indicate which pixels in the input image need to be changed the least to affect the prediction the most. By utilizing backpropagation, they obtained the gradients corresponding to each pixel in the input such that they could visualize what features the network considers essential. In [15] they argued that positive gradients with a large magnitude indicate pixels of high importance while negative gradients with a large magnitude indicate pixels which the networks want to suppress which, if included in the visualization of important pixels, can result in noisy images. To avoid this, Guided Backpropagation alters the backward pass of a neural network such that negative gradients are set to zero in each layer, thus allowing only positive gradients to flow backward through the network and highlighting pixels which the system finds important.

4. RESULTS

In this section, we present quantitative and qualitative results on semantic segmentation of colorectal polyps for both architectures, along with details regarding the training of the two models. We also present the results of using Monte Carlo Dropout to model the uncertainty associated with the predictions and the results of using Guided Backpropagation to visualize which pixels are considered important.

4.1. Training Approach

We evaluate our methods on the EndoScene [2] dataset for semantic segmentation of colorectal polyps, which consists of 912 RGB images obtained from colonoscopies from 36 patients. For each of the input images comes a corresponding annotated image provided by physicians, where pixels belonging to a polyp are marked in white and pixels belonging to the colon are marked in black. The first row of Figure 3 and 4 display examples from the dataset. We consider the two-class problem, where the task is to classify each pixel as polyp or part of the colon (background class). Following the approach of [2] we separate the dataset into training/validation/test set, where the training set consists of 20 patients and 547 images, the validation set consists of 8 patients and 183 images, and the test set consists of 8 patients and 182 images. All RGB input image are normalized to range $[0, 1]$.

For performance evaluation, we calculate Intersection over Union (IoU) and global accuracy (per-pixel accuracy) on the test set. For a given class c, prediction $\hat{y}_i$ and ground truth y_i, the IoU is defined as

$$\text{IoU}(c) = \frac{\sum_i (\hat{y}_i == c \land y_i == c)}{\sum_i (\hat{y}_i == c \lor y_i == c)} \quad (3)$$

where $\land$ is the logical *and* operation and $\lor$ is the logical *or* operation.

We initialize the decoder weights of both EFCN-8 and ESegNet using HeNormal initialization [24] and employ pre-trained weights for the encoders, as mentioned previously. All models were trained using ADAM [25] with a batch size of 10 and cross-entropy loss [11]. We use the validation set to apply early stopping and monitor polyp IoU score with a patience of 30. Class balancing was not applied as it gave no significant improvement and no weight decay was used.

Data augmentation was applied according to best practices to increase the number of training images artificially. We utilize a dynamic augmentation scheme that applies cropping, rotation, zoom, and shearing. During training, we crop images into 224x224 patches randomly chosen from the center or one of the corners, following the example of [9]. We apply random rotation between -90 and 90 degrees, random zoom from 0.8-1.2 and random shearing from 0-0.4.

4.2. Quantitative and Qualitative Results

In this section, we present the results for both architectures. Table 1 presents the quantitative results for our EFCN-8 and ESegNet on the test set along with previous results on this dataset obtained for a vanilla FCN-8 [2] and a previously state-of-the-art methods based on non-deep learning methods [4]. Row two of Figure 3 and 4 displays predictions from both models based on samples from the test set.

Model	# P(M)	IoU B	IoU P	IoU M	Acc M
SDEM [4]	<< 1	0.739	0.221	0.480	0.756
FCN-8 [2]	134.5	**0.946**	0.509	0.728	**0.949**
ESegNet	29.5	0.933	0.522	0.728	0.935
EFCN-8	134.5	**0.946**	**0.587**	**0.767**	**0.949**

Table 1. Results for the two-class problem of the EndoScene dataset. Abbreviations are: # P(M) = number of parameters in millions, IoU (Background, Polyp and Mean) and Accuracy Mean.

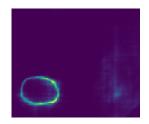## 4.3. Uncertainty and Interpretability Results

The third row of Figure 3 and 4 show the estimated standard deviation of each pixel in the predictions of both models. Using Monte Carlo Dropout, we obtain ten predictions which we used to estimate the standard deviation of each pixel. Row four of Figure 3 and 4 displays the results of using Guided Backpropagation to highlight what pixels in the input image both models consider important to their prediction.
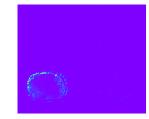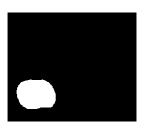

(a) input image

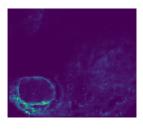(b) Ground truth

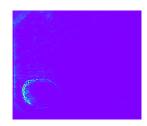
(c) EFCN-8 prediction

(d) ESegNet prediction

(e) EFCN-8 uncertainty

(f) ESegNet uncertainty

(g) EFCN-8 interpretability

(h) ESegNet interpretability

Fig. 3. Qualitative results on the Endoscene test set. The first row, from left to right, displays input image and its corresponding ground truth. The second row displays the prediction of both models. The third row displays the uncertainty associated with the prediction for both models and the fourth row displays which features are highlighted as important for both models.
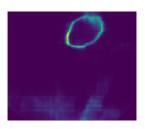
5. DISCUSSION

From Table 1 it is evident that the deep learning based models provide predictions of higher precision compared to previous methods based on traditional machine learning techniques. Also, the difference in performance between [2] FCN-8 and
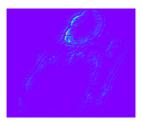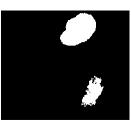
(a) input image

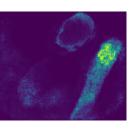(b) Ground truth

(c) EFCN-8 prediction

(d) ESegNet prediction

(e) EFCN-8 uncertainty

(f) ESegNet uncertainty

(g) EFCN-8 interpretability

(h) ESegNet interpretability

Fig. 4. Qualitative results on the Endoscene test set. The first row, from left to right, displays input image and its corresponding ground truth. The second row displays the prediction of both models. The third row displays the uncertainty associated with the prediction for both models and the fourth row displays which features are highlighted as important for both models.

our EFCN-8 indicated that the inclusion of recent techniques such as transfer learning and Batch Normalization is vital to increase the capabilities of deep models. Furthermore, ESegNet can achieve comparable results to the FCN-8 even though it has far fewer parameters. However, when recent techniques are included, we get a gap in performance between ESegNet

and the EFCN-8. This might imply that the increased complexity of the EFCN-8 is beneficial to performance.

For the qualitative results, Figure 3 displays an example where both models produce a correct prediction while Figure 4 shows an example where both models correctly segment the polyp present in the input image, but ESegNet also predicts a polyp where there is none. We have included the first example to illustrate that both models are capable of producing precise and correct predictions. The second example is included to highlight the difficulty when comparing models. For instance, given the predictions in Figure 4 and no ground truth, which prediction would you trust? We know the EFCN-8 achieved a higher performance on the test set, but that does not necessarily mean that it is correct in this particular case. Without further information, it would be difficult to choose between the two models without consulting a medical expert to asses the images.

However, if we could say something about the uncertainty of the two predictions we might choose the prediction with the lowest uncertainty associated with it. For the uncertainty estimates shown in row three of Figure 3, where both models successfully segment the polyp in the image, we see that the only pixels associated with high uncertainty are those around the border of the prediction. Such uncertainties are understandable as even physicians are unable to state precisely where the colon ends and the polyp starts. However, the uncertainty estimates which are shown in row three of Figure 4 tell a different story. Notice that both models exhibit similar uncertainty for the region where they both correctly segment a polyp, but ESegNet also has a large area of uncertainty for pixels associated with a ridge going along the colon. The falsely segmented polyp lies on this ridge of uncertainty, which might indicate that we should be careful about trusting the polyp prediction toward the bottom right of the image.

Another interesting question is, what pixels in the input image is influencing these predictions? In row four of Figure 3, we see that pixels associated with the edges of the polyp are highlighted, indicating that the models are leveraging edge information to identify the polyp. Also, notice that the EFCN-8 considers the entire top edge of the polyp while ESegNet only considers the left edge of the polyp, which might imply that the EFCN-8 has obtained a deeper understanding of the shape and form of polyp and can extract more useful information from the input image. In row four of Figure 4 we again see that both models are reacting to edges in the input image. But while the EFCN-8 can correctly identify pixels which belong to an actual polyp, ESegNet is also considering pixels which correspond to ridges of the colon.

Visualizing important pixels and modeling uncertainty is not just important to design automatic procedures which are trustworthy but also, as we have seen, allows for model analysis and model comparison. Of course, the results of such methods are still somewhat open to interpretation, and deep learning would benefit from a more theoretical framework for

analyzing models, yet including techniques such as Monte Carlo Dropout and Guided Backpropagation can lead towards a better understanding of CNNs.

6. CONCLUSION

In this paper, we improved and applied two established CNN architectures for pixel-wise segmentation, evaluated their performance on colorectal polyp segmentation and conclude that CNNs can achieve high performance in this context of medical image analysis. We also argued that modeling the uncertainty of the networks output and visualizing descriptive image regions in the input image can increase interpretability and make models based on deep learning more applicable for medical personnel.

The field of deep learning is continually improving and several recent architectures for semantic segmentation such as [26] show promising results. Increasing network interpretability is an active field of research where Relevance Propagation [22] is a particularly interesting approach for future experimentation. Post-processing procedures have also shown to improve performance of CNNs and including Conditional Random Fields could improve spatial coherence [12].

7. REFERENCES

[1] Jeroen C Van Rijn, Johannes B Reitsma, Jaap Stoker, Patrick M Bossuyt, Sander J Van Deventer, and Evelien Dekker, "Polyp miss rate determined by tandem colonoscopy: a systematic review," *The American journal of gastroenterology*, vol. 101, no. 2, pp. 343, 2006.

[2] David Vázquez, Jorge Bernal, F Javier Sánchez, Gloria Fernández-Esparrach, Antonio M López, Adriana Romero, Michal Drozdzal, and Aaron Courville, "A benchmark for endoluminal scene segmentation of colonoscopy images," *Journal of Healthcare Engineering*, vol. 2017, 2017.

[3] Jorge Bernal, F Javier Sánchez, Gloria Fernández-Esparrach, Debora Gil, Cristina Rodríguez, and Fernando Vilariño, "Wm-dova maps for accurate polyp highlighting in colonoscopy: Validation vs. saliency maps from physicians," *Computerized Medical Imaging and Graphics*, vol. 43, pp. 99–111, 2015.

[4] Jorge Bernal, Joan Manel Núñez, F Javier Sánchez, and Fernando Vilariño, "Polyp segmentation method in colonoscopy videos by means of msa-dova energy maps calculation," in *Workshop on Clinical Image-Based Procedures*. Springer, 2014, pp. 41–49.

[5] Q. Li, W. Cai, X. Wang, Y. Zhou, D. D. Feng, and M. Chen, "Medical image classification with convolutional neural network," in *2014 13th International Conference on Control Automation Robotics Vision (ICARCV)*, Dec 2014, pp. 844–848.

[6] Yao Xue and Nilanjan Ray, "Cell detection with deep convolutional neural network and compressed sensing," *arXiv preprint arXiv:1708.03307*, 2017.

[7] Nanqing Dong, Michael Kampffmeyer, Xiaodan Liang, Zeya Wang, Wei Dai, and Eric P Xing, "Unsupervised domain adaptation for automatic estimation of cardiothoracic ratio," in *International Conference on Medical Image Computing and Computer Assisted Intervention*. Springer, 2018.

[8] E. Ribeiro, A. Uhl, and M. Hfner, "Colonic polyp classification with convolutional neural networks," in *2016 IEEE 29th International Symposium on Computer-Based Medical Systems (CBMS)*, 2016, pp. 253–258.

[9] Alex Krizhevsky, Ilya Sutskever, and Geoffrey E Hinton, "Imagenet classification with deep convolutional neural networks," in *Advances in Neural Information Processing Systems*, 2012, pp. 1097–1105.

[10] Shaoqing Ren, Kaiming He, Ross Girshick, and Jian Sun, "Faster r-cnn: Towards real-time object detection with region proposal networks," *IEEE Transactions on Pattern Analysis and Machine Intelligence*, vol. 39, no. 6, pp. 1137–1149, 2017.

[11] Evan Shelhamer, Jonathan Long, and Trevor Darrell, "Fully convolutional networks for semantic segmentation," *IEEE Transactions on Pattern Analysis and Machine Intelligence*, vol. 39, no. 4, pp. 640–651, 2017.

[12] Liang-Chieh Chen, George Papandreou, Iasonas Kokkinos, Kevin Murphy, and Alan L Yuille, "Deeplab: Semantic image segmentation with deep convolutional nets, atrous convolution, and fully connected crfs," *IEEE Transactions on Pattern Analysis and Machine Intelligence*, vol. 40, no. 4, pp. 834–848, 2018.

[13] Alex Kendall, Vijay Badrinarayanan, and Roberto Cipolla, "Bayesian segnet: Model uncertainty in deep convolutional encoder-decoder architectures for scene understanding," *arXiv preprint arXiv:1511.02680*, 2015.

[14] Michael Kampffmeyer, Arnt-Borre Salberg, and Robert Jenssen, "Semantic segmentation of small objects and modeling of uncertainty in urban remote sensing images using deep convolutional neural networks," in *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition Workshops*, 2016, pp. 1–9.

[15] Jost Tobias Springenberg, Alexey Dosovitskiy, Thomas Brox, and Martin Riedmiller, "Striving for simplicity: The all convolutional net," *arXiv preprint arXiv:1412.6806*, 2014.

[16] Vijay Badrinarayanan, Alex Kendall, and Roberto Cipolla, "Segnet: A deep convolutional encoder-decoder architecture for image segmentation," *IEEE Transactions on Pattern Analysis and Machine Intelligence*, , no. 12, pp. 2481–2495, 2017.

[17] Sergey Ioffe and Christian Szegedy, "Batch normalization: Accelerating deep network training by reducing internal covariate shift," in *International Conference on Machine Learning*, 2015, pp. 448–456.

[18] Nitish Srivastava, Geoffrey Hinton, Alex Krizhevsky, Ilya Sutskever, and Ruslan Salakhutdinov, "Dropout: A simple way to prevent neural networks from overfitting," *Journal of Machine Learning Research*, vol. 15, pp. 1929–1958, 2014.

[19] Karen Simonyan and Andrew Zisserman, "Very deep convolutional networks for large-scale image recognition," *International Conference on Learning Representations*, 2015.

[20] Yarin Gal and Zoubin Ghahramani, "Dropout as a bayesian approximation: Representing model uncertainty in deep learning," in *International Conference on Machine Learning*, 2016, pp. 1050–1059.

[21] Matthew D. Zeiler and Rob Fergus, "Visualizing and understanding convolutional networks," *CoRR*, vol. abs/1311.2901, 2013.

[22] Sebastian Bach, Alexander Binder, Grégoire Montavon, Frederick Klauschen, Klaus-Robert Müller, and Wojciech Samek, "On pixel-wise explanations for non-linear classifier decisions by layer-wise relevance propagation," *PloS one*, vol. 10, no. 7, pp. e0130140, 2015.

[23] Karen Simonyan, Andrea Vedaldi, and Andrew Zisserman, "Deep inside convolutional networks: Visualising image classification models and saliency maps," *CoRR*, vol. abs/1312.6034, 2013.

[24] Kaiming He, Xiangyu Zhang, Shaoqing Ren, and Jian Sun, "Delving deep into rectifiers: Surpassing human-level performance on imagenet classification," *CoRR*, vol. abs/1502.01852, 2015.

[25] Diederik P. Kingma and Jimmy Ba, "Adam: A method for stochastic optimization," *CoRR*, vol. abs/1412.6980, 2014.

[26] Simon Jégou, Michal Drozdzal, David Vazquez, Adriana Romero, and Yoshua Bengio, "The one hundred layers tiramisu: Fully convolutional densenets for semantic segmentation," in *Computer Vision and Pattern Recognition Workshops (CVPRW), 2017 IEEE Conference on*. IEEE, 2017, pp. 1175–1183.