Deep Learning and Handcrafted Feature Based Approaches for Automatic Detection of Angiectasia*

Konstantin Pogorelov[1], Olga Ostroukhova[2], Andreas Petlund[3], Pål Halvorsen[1], Thomas de Lange[4], Håvard Nygaard Espeland[3], Tomas Kupka[3], Carsten Griwodz[1] and Michael Riegler[1]

***Abstract*—Angiectasia, formerly called angiodysplasia, is one of the most frequent vascular lesions and often the cause of gastrointestinal bleedings. Medical specialists assessing videos or images of examinations reach a detection performance of 16% for the detection of bleeding to 69% for the detection of angiectasia [1]. This shows that automatic detection to support medical experts can be useful. In this paper, we present several machine learning-based approaches for angiectasia detection in wireless video capsule endoscopy frames. In summary, the most promising results for pixel-wise localization and frame-wise detection are obtained by the proposed deep learning method using generative adversarial networks (GANs). Using this approach, we achieve a sensitivity of 88% and specificity of 99.9% for pixel-wise localization, and a sensitivity of 98% and a specificity of 100% for frame-wise detection. Thus, the results demonstrate the capability of using deep learning for automatic angiectasia detection in real clinical settings.**

***Index Terms*—Angiectasia, computer aided diagnosis, deep learning, machine learning, video capsular endoscopy**

I. Introduction

An obscure gastrointestinal (GI) bleeding is a common finding in the GI tract and caused by different diseases/conditions. The most challenging part is to detect the bleeding source in the small bowel either using video capsule endoscopes (VCEs) or via very invasive enteroscopy examinations. Superficial vascular lesions called angiectasia (see Figure 1 for an example) represent one of the most common source of bleeding in the small bowel and are therefore important to detect [2].

The most common procedure to detect angiectasia is to use VCEs. A VCE provides visualization of the GI tract by capturing images or recording a video by swallowing a pill-like disposable capsule equipped with one or more cameras. The camera pill contains a small processing device, a memory or wireless transmitter, and a battery. The VCE is swallowed by the patient, and it traverses and visualizes the GI tract for subsequent diagnosis and detection of GI diseases, such as angiectasia, by a doctor manually inspecting the video recordings. The latest generation of VCEs supports a maximum resolution of 520x520 pixels and is able to collect around 60,000 images per patient. Medical specialists assessing the images detect only around 69% of angiectasias (16% for the detection of bleeding to 69% for the detection of angiectasias) [1].

(a) Input frame (b) Ground truth mask (c) Segmentation mask

Fig. 1. Example of an angiectasia lesion marked with a green circle (a), a corresponding ground truth mask (b) and a segmentation mask generated using our GAN approach (c). Image taken from the GIANA dataset [3].

State-of-the-art software from the industry is able to reach an automatic detection sensitivity of 41% and a specificity of 67%. For a clinical scenario, this is clearly not reliable enough for automatic analysis. Both, sensitivity and specificity should be as close as possible to 100%, but at least larger than 85% for being used in a real clinical setting [1], [4]. Automatic detection of angiectasia is not very well researched, and there are only a few publications on the topic using saliency detection [5], [6]. However, no work has looked into machine learning using deep learning or handcrafted features. In this work, we therefore test different machine learning approaches to tackle automatic angiectasia detection in VCE videos. Using a publicly available and unbiased (equal number of negative and positive examples) dataset [3], we are testing algorithms (deep learning and handcrafted features-based) for frame-wise detection and pixel-wise localization. The best achieved results in this paper are a sensitivity of 88% and specificity of 99.9% for pixel-wise localization and a sensitivity of 98% and a specificity of 100% for frame-wise detection.

The rest of the paper is organized as follows: first, we give an overview of the related work in the field. This is followed by a description of our methods, which we next experimentally evaluate. Finally, we conclude the paper and give directions for future work.

*This work is founded by the FRINATEK project "EONS" #231687

[1]Konstantin Pogorelov, Carsten Griwodz, Pål Halvorsen and Michael Riegler are with Simula Research Laboratory and University of Oslo, Norway konstantin@simula.no

[2]Olga Ostroukhova is with Research Institute of Multiprocessor Computation Systems n.a. A.V. Kalyaev, Taganrog, Russia

[3]Andreas Petlund, Håvard Nygaard Espeland and Tomas Kupka are with ForzaSys AS, Norway

[4]Thomas de Lange is with Institute of Clinical Medicine, University of Oslo, Norway

II. Related work

To the best of our knowledge, there exists almost no related work about automatic detection of angiectasia in VCE images or videos. The only known work is from Deeba et al. [5], [6] who present a saliency-based approach. The method is two-staged and unsupervised. In the first step, a patch distinctness (PD) map and an Index of Hemoglobin (IHb) map are created. In the second step, the PD and IHb maps are combined to create the final saliency map. On this final map, a local maximum search is performed to find regions-of-interests containing the lesions. In [5], a sensitivity of 100%, a specificity of 82.5% and an accuracy of 90.1% are reported on a dataset containing 50 normal images and 50 images containing angiectasia. In [6], the dataset is extended to 3,602 images with 968 containing angiectasia and 2,634 normal ones. The images are taken from 9 different videos, whereas 5 are from wired endoscopy and 4 from VCE. For the VCE videos, a sensitivity of 94.44% and a specificity of 83.92% are reported. The localization score is 95.04% and measures the fraction of correctly detected regions compared to the regions containing angiectasia.

Furthermore, since bleeding angiectasia looks quite similar to regular GI bleeding, a very common condition, work addressing automatic bleeding detection should be considered. For bleeding detection, a lot of related work exist, and the main challenge is that the bleedings do not occur in specific patterns, shapes, textures or colors, which makes them hard to detect. Furthermore, bleeding is usually caused by other intestinal diseases like angiectasia or cancer, etc [7]. The main methods to detect bleeding are based on handcrafted color and textural features. In [8], chrominance-moments-based texture and uniform local binary patterns in combination with a multi-layer perception neural network classifier are used to localize the source of bleeding in the VCE video. Methods working on pixel-level are shown to be more accurate to distinguish between bleeding and non-bleeding. Yuan [9] utilizes color features on pixel level of VCE frames and thresholds the color space to segment bleeding from normal mucosa. In [10], the authors perform super-pixel based segmentation to reduce the computational complexity and at the same time achieve high accuracy. In general, pixel level methods have higher accuracy than frame based, but are computationally more costly. This is an important factor taking into account that usually more than 60,000 frames have to be processed for one patient.

Most recent work is focusing on deep learning for bleeding detection by utilizing convolutional neural networks (CNN). In [11], the authors present a bleeding detection approach that uses CNNs. The presented CNN consists of eight layers and is basically a simple variation of the Imagenet architecture [12]. They report a recall/sensitivity of 99.2%, a precision of 99.9% and F1 score of 99.55%. For the training, 8,200 images were used, and 1,800 images were used for the testing. Both the training and testing dataset were biased towards the negative class. Furthermore, no cross validation for the evaluation was performed. Therefore, it cannot be ruled out that the almost perfect performance is based on overfitting. This was followed by another study using the same approach [13], but with a different dataset and on pixel accuracy level for segmentation. As the main metric, region intersection over union (IU) was used. For active bleeding, an IU of 0.7750 and for inactive bleeding an IU of 0.7524 were achieved. Another recent work that does not use deep learning but classifier fusion is the work by Deeba et al. [14] that combines two optimized Support Vector Machine (SVM) classifiers to detect bleeding. The features used by the classifiers are based on the RGB and HSV color spaces. For parameter tuning and evaluation, a cross validation approach is used, and they report an average accuracy of 95%, sensitivity of 94% and specificity of 95.3% for a dataset of 8,872 VCE frames.

The presented related work contains only two papers about angiectasia and some related work in the field of bleeding detection. As one can observe, even if deep learning is in the rise, handcrafted features still achieve good performance if used in a clever way. In the context of angiectasia, one can see that the VCE datasets used in the related work are biased and too small. Therefore, the goal of this work is to compare and evaluate deep learning and hand crafted features based approaches on a large and unbiased dataset.

III. Pixel-wise segmentation approach

The segmentation approach presented in this paper is able to pixel-accurate mark the angiectasia in the given frame. Based on our previous experience [15], [16], we decided to use generative adversarial network (GAN) to perform the segmentation. GANs [17] are machine learning algorithms that are usually used in unsupervised learning and are implemented by using two neural networks competing with each other in a zero-sum game. We used a GAN model architecture initially developed for the retinal vessel segmentation in fundoscopic images called V-GAN as basis for our angiectasia segmentation approach. The V-GAN architecture [18] is designed for RGB images and provides a per-pixel image segmentation as output. To be able to use the V-GAN architecture in our angiectasia segmentation approach, we added an additional output layer to the generator network that implements an activation layer with a step function which is required to generate the binary segmentation output.

IV. Frame-wise detection approaches

Frame-wise detection approaches are designed to detect angiectasia on a frame level, i.e., if there is angiectasia in the frame or not. For frame-wise detection, we propose different methods where we conducted experiments using various configurations of our main methods. The main methods are global features (GF), deep features (DF) and a variation of our GAN approach. For the classification, we used the Random Tree (RT), Random Forrest (RF) and Logistic Model Tree (LMT) classifiers provided in the WEKA library [19].

Global features. For the GF method, we extracted handcrafted global features (describing the image on a global level, e.g., texture, color distribution, etc.) using the LIRE framework [20]. The features are Joint Composite Descriptor, Tamura, Color Layout, Edge Histogram, Auto Color Correlogram and Pyramid Histogram of Oriented Gradients. We performed early fusion by combining all extracted features resulting in a feature vector with the size of 1186.

Deep features. For the DF approach, we used different well known working deep learning architectures to extract either the features directly (FEA) or to classify the images and using the whole range of concepts and their probabilities as input for the classifiers (CON). The architectures that we used are ResNet50 [21], VGG19 [22], and InceptionV3 [23].

Data augmentation. For fair performance comparison of the GF and DF approaches with the GAN approach, we implemented the same data augmentation (AUG) scheme (rotation and flipping of frames) as used in the training process of the GAN. Rotation was performed with 20° steps for the original and the flipped frames, resulting in 35 new frames complementary to the original ones.

GAN. The GAN detection approach utilizes a simple threshold activation function, which takes the number of positively marked pixels in the frame as an input. In the cross-validation experiments, we evaluated the activation thresholds from one pixel to a quarter of the frame. The best results were achieved with a threshold value of 2 pixels, which has been used for the detection experiments.

V. Experiments

The data used for all the experiments is from the GIANA 2017 challenge [3], and it is publicly available for research purposes. The data consists of training (development) and test frame sets. The training set consists of 600 fully annotated frames from VCEs (300 with angiectasia and 300 without). The frames with angiectasia also have a pixel-wise ground truth (GT) mask depicting the exact lesion location inside each frame that allows both pixel-wise localization and frame-wise detection experiments. The test set consists of 600 unannotated frames. In order to perform validation and performance evaluation of the developed detection algorithm, we annotated the test set frame-wise with the help of an experienced researcher with medical pathology diagnosis background. The 600 frames from the development set are used for training and the 600 frames (300 with angiectasia and 300 normal) from the test set for verification. The advantages of the used dataset are (i) the number of images (compared to related work, this is the largest one for VCEs), (ii) the even split between positive and negative examples and (iii) that it is publicly available making it easy to compare different approaches. For evaluation of the experiments, we used the precision (PREC), recall/sensitivity (SENS), specificity (SPEC), accuracy (ACC), F1 score (F1), Matthew correlation coefficient (MCC) and processing speed in number of frames per second (FPS) metrics. A detailed description and reasoning for the used metrics can be found in [24]. The localization metrics are calculated pixel-wise using the provided GT masks. For the best working approach (GAN), we also report detailed results for the ten-fold cross-validation including 95% confidence intervals (CI). For the detection part, we use a ZeroR classifier as baseline which assigns the label from the majority class (most common label in the dataset) to all the instances.

TABLE I
Ten-fold cross-validation results of the pixel-wise angiectasia areas the GAN segmentation approach.

Fold	PREC	SENS	SPEC	ACC	F1	MCC
1	0.805	0.877	0.999	0.999	0.839	0.839
2	0.893	0.908	0.999	0.999	0.901	0.900
3	0.870	0.871	0.999	0.999	0.871	0.870
4	0.808	0.884	0.999	0.998	0.844	0.844
5	0.876	0.894	0.999	0.999	0.885	0.885
6	0.838	0.849	0.999	0.998	0.843	0.842
7	0.900	0.887	0.999	0.999	0.893	0.893
8	0.863	0.900	0.999	0.999	0.881	0.880
9	0.866	0.914	0.999	0.999	0.889	0.889
10	0.873	0.817	0.999	0.999	0.844	0.844
95% CI	**0.859 ±0.020**	**0.880 ±0.018**	**0.999 ±0.001**	**0.999 ±0.001**	**0.869 ±0.015**	**0.869 ±0.015**

TABLE II
Ten-fold cross-validation results of the angiectasia frame-wise detection using the GAN approach.

Fold	PREC	SENS	SPEC	ACC	F1	MCC
1	1.000	1.000	1.000	1.000	1.000	1.000
2	1.000	0.967	1.000	0.983	0.983	0.967
3	1.000	1.000	1.000	1.000	1.000	1.000
4	1.000	1.000	1.000	1.000	1.000	1.000
5	1.000	0.967	1.000	0.983	0.983	0.967
6	1.000	0.967	1.000	0.983	0.983	0.967
7	1.000	1.000	1.000	1.000	1.000	1.000
8	1.000	1.000	1.000	1.000	1.000	1.000
9	1.000	1.000	1.000	1.000	1.000	1.000
10	1.000	0.967	1.000	0.983	0.983	0.967
95% CI	**1.000 ±0**	**0.987 ±0.011**	**1.000** ±0	**0.993 ±0.005**	**0.993 ±0.005**	**0.987 ±0.011**

A. Results

Table I shows the results for the GAN localization algorithm (see Figure 1(b) and 1(c) for a comparison between the GT and the output of the GAN). On average, sensitivity and specificity are above the 85% margin recommended for a real clinical settings. This can be seen as very good results since we perform pixel-wise evaluation. The processing speed for the GAN approach is 1.5 FPS. The frame-wise detection performance of the GAN approach for the development set is presented in Table II. The detection outperforms significantly the 85% requirements. Both result sets are a strong indicators that our GAN approach performs well for the tasks of angiectasia localization and detection. Finally, in Table III, we report the frame-wise detection performance on the test set for all our runs. All tested approaches outperform the ZeroR baseline, but most of them do not even come close to the 85% margin for clinical use. The handcrafted features outperform the VGG19 and InceptionV3 approaches but not the RestNet50. From the classifiers LMT performs best most of the time, followed by RF. The best performing not-GAN approach is *AUG DF ResNet50 FEA + LMT*. The GAN approach achieves superior performance compared to all other detection methods for the frame-wise detection with a sensitivity of 98% and a specificity of 100%. The best

TABLE III
RESULTS FOR THE ANGIECTASIA FRAME-WISE DETECTION APPROACHES EVALUATED WITH THE ANNOTATED TEST SET.

Approach	PREC	SENS	SPEC	ACC	F1	MCC	FPS
GF+RT	0.570	0.568	0.568	0.568	0.566	0.138	130
GF+RF	0.628	0.623	0.623	0.623	0.620	0.252	105
GF+LMT	0.695	0.680	0.680	0.680	0.674	0.375	80
DF ResNet50 CON+RT	0.636	0.636	0.636	0.636	0.636	0.271	88
DF ResNet50 CON+RF	0.742	0.742	0.742	0.742	0.742	0.483	78
DF ResNet50 CON+LMT	0.734	0.732	0.732	0.732	0.731	0.465	53
DF ResNet50 FEA+RT	0.558	0.557	0.557	0.557	0.554	0.114	79
DF ResNet50 FEA+RF	0.721	0.720	0.720	0.720	0.720	0.441	70
DF ResNet50 FEA+LMT	0.748	0.738	0.738	0.738	0.736	0.486	46
DF VGG19 CON+RT	0.538	0.538	0.538	0.538	0.538	0.077	60
DF VGG19 CON+RF	0.594	0.593	0.593	0.593	0.592	0.187	49
DF VGG19 CON+LMT	0.545	0.545	0.545	0.545	0.544	0.090	32
DF VGG19 FEA+RT	0.515	0.515	0.515	0.515	0.515	0.030	54
DF VGG19 FEA+RF	0.548	0.548	0.548	0.548	0.548	0.097	47
DF VGG19 FEA+LMT	0.525	0.525	0.525	0.525	0.525	0.050	29
DF InceptionV3 CON+RT	0.537	0.537	0.537	0.537	0.537	0.073	66
DF InceptionV3 CON+RF	0.617	0.617	0.617	0.617	0.617	0.233	50
DF InceptionV3 CON+LMT	0.663	0.663	0.663	0.663	0.663	0.327	37
DF InceptionV3 FEA+RT	0.515	0.515	0.515	0.515	0.513	0.030	56
DF InceptionV3 FEA+RF	0.551	0.548	0.548	0.548	0.542	0.099	43
DF InceptionV3 FEA+LMT	0.533	0.533	0.533	0.533	0.533	0.067	30
AUG GF+RT	0.545	0.545	0.545	0.545	0.544	0.090	130
AUG GF+RF	0.650	0.643	0.643	0.643	0.639	0.293	105
AUG GF+LMT	0.627	0.625	0.625	0.625	0.624	0.252	80
AUG DF ResNet50 CON+RT	0.620	0.620	0.620	0.620	0.620	0.240	88
AUG DF ResNet50 CON+RF	0.787	0.787	0.787	0.787	0.787	0.574	78
AUG DF ResNet50 CON+LMT	0.765	0.763	0.763	0.763	0.763	0.529	53
AUG DF ResNet50 FEA+RT	0.553	0.553	0.553	0.553	0.553	0.107	79
AUG DF ResNet50 FEA+RF	0.727	0.723	0.723	0.723	0.722	0.450	70
AUG DF ResNet50 FEA+LMT	0.797	0.788	0.788	0.788	0.787	0.585	46
GAN	**1.000**	**0.980**	**1.000**	**0.990**	**0.990**	**0.980**	**1.5**
Baseline (ZeroR)	0.250	0.500	0.500	0.500	0.333	0.000	-

processing speed is reached by the GF approach using RT. In terms of fastest speed and best classification performance, *AUG DF ResNet50 CON + RF* performs best with a sensitivity of 78.7% , a specificity of 78.7% and a processing speed of 78 FPS. The processing speed of the GAN method for detection is the lowest with 1.5 FPS.

B. Conclusion

In this paper, we presented hand crafted and deep learning-based methods for automatic detection of angiectasia on a pixel- and frame-wise level. We compared several approaches (handcrafted and deep learning) and demonstrated, on a public available dataset, the capability of our proposed GAN approach to reach and exceed clinical requirements (sensitivity and specificity higher than 85%) for localization and detection performance. In summary, we achieved a sensitivity of 88% and a specificity of 99.9% for pixel-wise localization, and a sensitivity of 98% and a specificity of 100% for frame-wise detection. For future work, the improvement of the processing speed and verification with other pathologies for our best working approach is planed.

REFERENCES

[1] V. Baptista, N. Marya, A. Singh, A. Rupawala, B. Gondal, and D. Cave, "Continuing challenges in the diagnosis and management of obscure gastrointestinal bleeding," *World journal of gastrointestinal pathophysiology*, vol. 5, no. 4, p. 523, 2014.

[2] X. Bosch, E. Montori, M. Guerra-García, J. Costa-Rodríguez, M. H. Quintanilla, P. E. Tolosa-Chapasian, P. Moreno, N. Guasch, and A. López-Soto, "A comprehensive evaluation of the gastrointestinal tract in iron-deficiency anemia with predefined hemoglobin below 9mg/dl: A prospective cohort study," *Digestive and Liver Disease*, vol. 49, no. 4, pp. 417–426, 2017.

[3] J. Bernal and H. Aymeric, "Miccai endoscopic vision challenge angiodysplasia d & l," https://endovissub2017-giana.grand-challenge.org/home/, accessed: 2017-11-20.

[4] J. Regula, E. Wronska, and J. Pachlewski, "Vascular lesions of the gastrointestinal tract," *Best Practice & Research Clinical Gastroenterology*, vol. 22, no. 2, pp. 313–328, 2008.

[5] F. Deeba, S. K. Mohammed, F. M. Bui, and K. A. Wahid, "A saliency-based unsupervised method for angioectasia detection in capsule endoscopic images," *Proc. of the CMBES*, vol. 39, no. 1, 2016.

[6] ——, "A saliency-based unsupervised method for angiectasia detection in endoscopic video frames," *Journal of Medical and Biological Engineering*, pp. 1–11, 2017.

[7] M. Van Leerdam, E. Vreeburg, E. Rauws, A. Geraedts, J. Tijssen, J. Reitsma, and G. Tytgat, "Acute upper gi bleeding: did anything change?" *The American journal of gastroenterology*, vol. 98, no. 7, pp. 1494–1499, 2003.

[8] B. Li and M. Q.-H. Meng, "Computer-aided detection of bleeding regions for capsule endoscopy images," *IEEE Transactions on Biomedical Engineering*, vol. 56, no. 4, pp. 1032–1039, 2009.

[9] Y. Yuan, B. Li, and M. Q.-H. Meng, "Bleeding frame and region detection in the wireless capsule endoscopy video," *IEEE Journal of biomedical and health informatics*, vol. 20, no. 2, pp. 624–630, 2016.

[10] Y. Fu, W. Zhang, M. Mandal, and M. Q.-H. Meng, "Computer-aided bleeding detection in wce video," *IEEE Journal of biomedical and health informatics*, vol. 18, no. 2, pp. 636–642, 2014.

[11] X. Jia and M. Q.-H. Meng, "A deep convolutional neural network for bleeding detection in wireless capsule endoscopy images," in *Proc. of IEEE EMBC*, 2016, pp. 639–642.

[12] A. Krizhevsky, I. Sutskever, and G. E. Hinton, "Imagenet classification with deep convolutional neural networks," in *Advances in neural information processing systems*, 2012, pp. 1097–1105.

[13] X. Jia and M. Q.-H. Meng, "A study on automated segmentation of blood regions in wireless capsule endoscopy images using fully convolutional networks," in *Proc. of IEEE ISBI*, 2017, pp. 179–182.

[14] F. Deeba, M. Islam, F. M. Bui, and K. A. Wahid, "Performance assessment of a bleeding detection algorithm for endoscopic video based on classifier fusion method and exhaustive feature selection," *Biomedical Signal Processing and Control*, vol. 40, pp. 415–424, 2018.

[15] K. Pogorelov, M. Riegler, S. L. Eskeland, T. de Lange, D. Johansen, C. Griwodz, P. T. Schmidt, and P. Halvorsen, "Efficient disease detection in gastrointestinal videos–global features versus neural networks," *Multimedia Tools and Applications*, pp. 1–33, 2017.

[16] K. Ahmad, K. Pogorelov, M. Riegler, N. Conci, and H. Pal, "Cnn and gan based satellite and social media data fusion for disaster detection," in *Proc. of the MediaEval 2017 Workshop*, Dublin, Ireland.

[17] I. Goodfellow, J. Pouget-Abadie, M. Mirza, B. Xu, D. Warde-Farley, S. Ozair, A. Courville, and Y. Bengio, "Generative adversarial nets," in *Advances in neural information processing systems*, 2014.

[18] J. Son, S. J. Park, and K.-H. Jung, "Retinal vessel segmentation in fundoscopic images with generative adversarial networks," *arXiv preprint arXiv:1706.09318*, 2017.

[19] M. Hall, E. Frank, G. Holmes *et al.*, "The weka data mining software: an update," *ACM SIGKDD explorations newsletter*, vol. 11, no. 1, pp. 10–18, 2009.

[20] M. Lux, M. Riegler, P. Halvorsen, K. Pogorelov, and N. Anagnostopoulos, "Lire: open source visual information retrieval," in *Proc. of ACM MMSys*, 2016.

[21] K. He, X. Zhang, S. Ren, and J. Sun, "Deep residual learning for image recognition," in *Proc. of IEEE CVPR*, 2016, pp. 770–778.

[22] K. Simonyan and A. Zisserman, "Very deep convolutional networks for large-scale image recognition," *arXiv preprint arXiv:1409.1556*, 2014.

[23] C. Szegedy, V. Vanhoucke, S. Ioffe, J. Shlens, and Z. Wojna, "Rethinking the inception architecture for computer vision," in *Proc. of IEEE CVPR*, 2016, pp. 2818–2826.

[24] K. Pogorelov, K. R. Randel, C. Griwodz, S. L. Eskeland, T. de Lange, D. Johansen, C. Spampinato, D.-T. Dang-Nguyen, M. Lux, P. T. Schmidt, M. Riegler, and P. Halvorsen, "Kvasir: A multi-class image dataset for computer aided gastrointestinal disease detection," in *Proc. of ACM MMSYS*, 2017, pp. 164–169.