

## Review

# Deep learning and computer vision techniques for microcirculation analysis: A review

Maged Helmy,<sup>1,\*</sup> Trung Tuyen Truong,<sup>2</sup> Eric Jul,<sup>1,3,\*</sup> and Paulo Ferreira<sup>1</sup><sup>1</sup>Department of Informatics, University of Oslo, Oslo, Norway<sup>2</sup>Department of Mathematics, University of Oslo, Oslo, Norway<sup>3</sup>Department of Computer Science, University of Copenhagen, Copenhagen, Denmark\*Correspondence: [magedaa@uio.no](mailto:magedaa@uio.no) (M.H.), [ericbj@uio.no](mailto:ericbj@uio.no) (E.J.)<https://doi.org/10.1016/j.patter.2022.100641>

**THE BIGGER PICTURE** The analysis of microcirculation images has the potential to reveal early signs of life-threatening diseases. Quantifying the capillary distribution in microcirculation images can be used as a biological marker to assist patients. The quantification of these biological markers is labor-intensive, time-consuming, and subject to interobserver variability. Moreover, manual analysis has been reported to hinder the application of microvascular microscopy in a clinical environment. Several computer vision techniques with varying performances can be used to automate the analysis of these microcirculation images. Computer vision algorithms are faster than convolutional neural networks for capillary detection but have poorer accuracy. Convolutional neural networks are more accurate but slower and require many training data. Therefore, by creating a hybrid model combining both computer vision algorithms and convolutional neural networks, one can strike a balance between accuracy and speed.



**Concept:** Basic principles of a new data science output observed and reported

## SUMMARY

The analysis of microcirculation images has the potential to reveal early signs of life-threatening diseases such as sepsis. Quantifying the capillary density and the capillary distribution in microcirculation images can be used as a biological marker to assist critically ill patients. The quantification of these biological markers is labor intensive, time consuming, and subject to interobserver variability. Several computer vision techniques with varying performance can be used to automate the analysis of these microcirculation images in light of the stated challenges. In this paper, we present a survey of over 50 research papers and present the most relevant and promising computer vision algorithms to automate the analysis of microcirculation images. Furthermore, we present a survey of the methods currently used by other researchers to automate the analysis of microcirculation images. This survey is of high clinical relevance because it acts as a guidebook of techniques for other researchers to develop their microcirculation analysis systems and algorithms.

## INTRODUCTION

Arterioles, venules, and capillaries comprise most of the microvessels known as the microcirculation system.<sup>1</sup> These microvessels measure less than 20  $\mu\text{m}$ .<sup>1</sup> The cardiovascular system's main component, microcirculation, is in charge of delivering nutrients, hormones, and oxygen to the tissues.<sup>2,3</sup> The microcirculation system is essential for the body's ability to control blood pressure and sustain tissue function.<sup>4,5</sup> In addition, it aids in clearing out waste from the cell tissues.<sup>6</sup> Arterioles carry oxygenated blood to the capillaries, while venules carry away deoxygenated blood.<sup>1</sup> According to research conducted on microcirculation, a low number of healthy capillaries might indicate some underlying medical conditions.<sup>7–19</sup> In addition, the thin cell walls of the capillaries make it viable for car-

bon dioxide and oxygen to cross the capillary walls.<sup>20</sup> The delivery of oxygen helps the tissue cells' functional activity by providing the energy they need.<sup>21</sup> The parenchymal cells depend on the efficient functioning of the microcirculation to keep the cells alive and sustain their function.<sup>21</sup> Thus, it might be argued that the microcirculation system is an essential component of the cardiovascular system. The functional capillary density of the microcirculation system can be used to measure the oxygen transport factor to the cells, which is critical for cell survival;<sup>22</sup> thus, methods to measure the capillary density will be the main topic of this paper. Other vital microcirculation tasks include but are not limited to the management of solute exchange in the parenchymal space,<sup>1</sup> delivery of white blood cells to their target tissue,<sup>23</sup> delivery of all blood-borne hormones, and modulating hemostasis.<sup>1</sup>



One of the functions of capillaries is to deliver oxygen to the cells of all organs.<sup>24</sup> In the past decades, research has been conducted to elucidate whether monitoring microcirculation can be used to diagnose or assess the severity of various diseases.<sup>25,26</sup> One of the first study directions was to record blood movement in the nail fold capillaries to assess the severity of rheumatic diseases, such as systemic sclerosis, Raynaud's syndrome, and dermatomyositis.<sup>27</sup> These studies revealed that patients with rheumatic diseases had altered microcirculation, showing dilated and distorted capillary loops and areas with low capillary density compared with controls.<sup>27</sup> More recently, studies on the effect of nonrheumatic diseases on microcirculation have been conducted. Microcirculation has been monitored either on the skin surface<sup>28</sup> or sublingually<sup>2</sup> using video recordings of segments of blood flow within the capillaries. The impact of various diseases on microcirculation has been assessed, including sepsis<sup>29,30</sup> and, more recently, COVID.<sup>31,32</sup> These studies have reported various correlations between diseases and the density of capillaries, the velocity of blood flow, and the heterogeneity of perfusion.<sup>2,27,28,31,32</sup>

Recently, the microvascular community has been focused on standardizing the analysis of microcirculation images<sup>33</sup> by using automated methods<sup>34–36</sup> (see [state-of-the-art microcirculation image analysis techniques](#)). This is of exceptionally high importance as, currently, the gold standard for microvascular analysis is manual analysis by a trained researcher.<sup>34</sup> This process is time-consuming (approximately 2 min per microcirculation image) and prone to subjective bias. Moreover, the length of analysis and requirements for trained researchers prevent microvascular monitoring from being used in routine clinical applications.<sup>37</sup> In this article, we present the relevant methods in deep neural networks and traditional computer vision algorithms that can be used to achieve the automation of microcirculation image analysis.

Moreover, cardiogenic shock can deteriorate microcirculatory functions and cause disturbances. Therefore, speeding up the process of microcirculation analysis via accurate automation processes can assist in evaluating the efficiency of techniques used to treat a patient suffering from a cardiogenic shock. Other broader relevant diseases that affect microcirculation include but are not limited to shock reperfusion,<sup>38</sup> iatrogenic injury,<sup>38,39</sup> and pancreatitis.<sup>40</sup>

Traditional computer vision techniques require the user to find the optimal set of values to segment the capillaries in an image, while deep neural networks can automatically attempt to find those values based on the dataset provided.<sup>41</sup> Although deep neural networks can automate the manual segmentation process, traditional computer vision algorithms require less computational power and are faster than deep neural network methods.<sup>42</sup> Thus, we present both techniques in this review paper. That being said, we believe that a combination of both of these methods will be the future of microcirculation analysis in the clinical setting.<sup>36</sup>

This paper is intended to serve as a guidebook to inform researchers on what deep learning and computer vision techniques exist that can be used for the automated quantification of capillaries. The goal of the methods represented in this paper is to reduce the labor-intensive, time-consuming analysis from several minutes to seconds. Furthermore, these methods aim to reduce the subjectivity of the interobserver variability using a

standardized method. The [capillary dataset](#) section presents and discusses the capillary dataset. The [machine learning](#) section introduces machine learning and its four types. The [deep learning](#) section goes into the details of convolutional neural networks. The [deep learning object detection techniques](#) section introduces the two object detection techniques: regional proposal-based framework and unified-based framework. The [traditional computer vision object detection techniques](#) section introduces traditional computer vision object detection techniques. The [state-of-the-art microcirculation image analysis techniques](#) section presents the current methods and techniques published and used by researchers for capillary quantification. We then conclude our paper.

## CAPILLARY DATASET


Capillaroscopy is a method that noninvasively checks the dermal papillary capillaries using a microscopy system.<sup>43</sup> The image(s) obtained can then be evaluated for capillary density, dimension, and morphology.<sup>43</sup> [Figure 1](#) was obtained from Ruaro et al.<sup>44</sup> and best illustrates this. Ruaro et al.<sup>44</sup> describe systemic sclerosis as a disease that alters the microvascular structure, which can be seen using capillaroscopy. [Figure 1A](#) shows the typical pattern with no disease; [Figures 1B–1D](#) show how the capillaries react to different disease stages.<sup>44</sup>

Microcirculation data can prominently come from sublingual microcirculation,<sup>45</sup> the eyes,<sup>46</sup> or the dorsum region of the hand<sup>47</sup> or nail bed.<sup>48</sup>

Several microscopes are currently available that can achieve such images: Dino-Lite CapillaryScope,<sup>49</sup> Optilia Digital Capillaroscope,<sup>50</sup> Inspectis Digital Capillaroscope,<sup>51</sup> and Smart G-Scope.<sup>52</sup> Other equipment that can capture capillaries include dermatoscope, ophthalmoscope, and stereomicroscope; however, none was designed for capillary capture. Therefore, images from this equipment are of relatively lower quality<sup>53</sup> than images captured by microscopes. That said, contemporary microscopes designed for capillaroscopy still do not produce adequate image quality compared with those in a standard object recognition dataset. Some examples of standard object recognition datasets used to develop object recognition algorithms are ImageNet (14 million+ images and 21,841 categories)<sup>54</sup>, Common Objects in Context (COCO) (328,000+ images and 91 categories),<sup>54</sup> Places (10 million+ images and 434 categories),<sup>55</sup> and Open Images (9 million+ images and 6,000+ categories).<sup>56</sup> These datasets have several thousand images per category, and the objects to be identified are relatively not as pixelated as the capillary data. Thus, one of the main challenge is the dataset's quality. Since the capillaries are relatively smaller in size, measuring less than 20  $\mu\text{m}$  in diameter,<sup>24</sup> the shape of the capillary tends to be pixelated. Furthermore, in the literature, there is no clear definition of capillary shapes, which exponentially increases the challenge of not enough data being present for each shape/type of capillary, as shown in [Figure 2](#).

## MACHINE LEARNING

Machine learning (ML) is a sub-branch of AI that combines techniques from computer science and statistics.<sup>57</sup> ML aims to find



**Figure 1. Systemic sclerosis disease affects the microvascular structure**

In (A), we see a normal pattern, while in (B, C, and D), we notice changes in the capillary structure as the disease advances. This image is from Ruaro et al.<sup>44</sup>

patterns from data to make predictions about new data.<sup>58</sup> This process of finding a pattern from a set of data is known as training.<sup>59</sup> The product of this process is a model that is used to predict new data.<sup>59</sup> An ML framework can consist of seven processes,<sup>60–62</sup> which are data collection, data preparation, selecting and training the model, testing the selected model, evaluating the model using F1-score, and finally deploying the model to production.

ML techniques can be further divided into supervised learning, unsupervised learning, reinforcement learning, and transfer learning.<sup>63</sup>

Supervised learning is an ML technique where the algorithm learns from data labeled by a human expert.<sup>64</sup> In terms of microcirculation data, the model will attempt to highlight the capillaries with a bounding box based on the sample of labeled data provided to it. Supervised learning can be further divided into regression and classification. Regression predicts a value, while classification predicts a class. In the case of microcirculation analysis, we predict whether a region has capillaries or not, so using the classification algorithms from the supervised learning techniques is the most suitable method for capillary detection.

Unsupervised learning, as opposed to supervised learning, is giving a set of data with no labels to the algorithm.<sup>65</sup> The algorithm then attempts to find patterns between the data.<sup>66</sup> There are two types of unsupervised learning: clustering and association.<sup>67</sup> Clustering attempts to find subgroups in the data based on color, density, or other features. In contrast, association tries to discover relationships between patterns. For example, a system that tries to recommend to the user what to purchase next based on what other users have bought, that is, association. Another example is anomaly detection, where the system finds data that exhibits patterns outside the normal range of the data. We did not find any suitable unsupervised algorithms for capillary detection. Moreover, we have not yet found papers on capillary detection that use unsupervised algorithms for capillary detection.

Reinforcement learning finds patterns based on a reward system.<sup>68</sup> When the system predicts the right output, the weights of those nodes in the neural network are increased, while an incorrect prediction diminishes the weights of other nodes.<sup>69</sup> Reinforcement learning is common in game development for games such as chess and Go.<sup>70,71</sup> We did not find any suitable unsupervised algorithms for capillary detection.

Transfer learning uses a model trained on previous data and tweaks the weights of the last  $n$  (defined by a machine learning expert) number of layers to train it on new data.<sup>72</sup> An alternative way is to add new layers in addition to the existing layers of the network. It is assumed that the earlier layers of a transfer learning model detects the generic features of the image such as the width, height, and edges and that the latest layers learn the fine details of the classification.<sup>73</sup> Therefore by freezing the earlier layers while training the later layers on a new set of data, one can adapt any model to different datasets without having to retrain the algorithm from scratch.<sup>74</sup>

For microcirculation capillary quantification, our review work reveals that the most relevant ML techniques for microcirculation analysis are supervised classification learning and transfer learning.

## DEEP LEARNING


Deep learning is a sub-branch of ML that has vastly advanced the state of the art in speech recognition and visual object detection.<sup>75</sup> In its simplest form, a deep neural network is composed of multiple layers of neurons.<sup>75</sup> This typically consists of the input layer, hidden layer(s), and the output layer.<sup>75</sup> Deep learning outshines the other ML techniques because it can construct a feature extractor without the need for domain expertise.<sup>76</sup> Rather, the neural network increases prediction accuracy by adjusting the neuron weights to optimize an appropriate cost function using back-propagation techniques.<sup>77</sup> Deep learning architectures combine multiple non-linear modules to transform input data into a higher abstract level. That suppresses irrelevant information while intensifying the useful parts to increase prediction accuracy.<sup>78</sup> In this section, we describe several deep learning architectures and then delve into the details of convolutional neural network (CNN) architectures.

### Types of deep neural networks

There are many types of deep neural networks. In this section, we focus on the three most relevant for microcirculation analysis: recurrent neural networks (RNNs),<sup>79</sup> generative adversarial networks (GANs),<sup>80</sup> and CNNs.<sup>81</sup>

RNNs are a type of neural network that has a hidden state to find patterns in sequential data and uses the output of previous layers as an input for the current layer.<sup>79</sup> The nodes of an RNN are connected by directed cycles, allowing it to keep the state of the previous nodes. While a traditional neural network takes in the input and gives an output, an RNN assumes a relationship exists among the sequences of input data.<sup>82</sup> Thus, RNN is mainly used in speech recognition, translation, and sequential data analysis.<sup>82</sup>

Long short-term memory networks (LSTM)<sup>83</sup> and gated recurrent units (GRU)<sup>84</sup> are RNN architectures that deal with gradient vanishing challenges.<sup>85</sup> The vanishing/exploding gradient



**Figure 2. Example of capillary dataset captured using a relatively high-end microscope**

We see that the shapes of capillaries are not the same and can be faded. We also see instances with black spots and other cases with white spots, which is a reflection due to the oil. These are considered artifacts and are undesirable. Such issues make it challenging to develop a highly accurate algorithm that can generalize in detecting capillaries.

problem is when the size of the gradients of each layer does not equal 1. This will prevent the neural weights from converging. The modifications introduced by LSTM and GRU deal with this challenge by reducing the amount of irrelevant information propagating through the architecture by adding appropriate gates into the vanilla RNN.

GAN creates new instances of the data.<sup>86</sup> They are made up of a generator (G) and a discriminator (D), whose tasks are opposite to each other. The generator obtains a random noise input  $z$ , and its output is a fake datum sample  $G(z)$ . The discriminator D obtains as input both  $G(z)$  and a real sample  $x$ , and its output is the probability of whether  $G(z)$  is real. G can reliably produce new instances that mimic those in the real dataset. One of its main uses is in data augmentation to create more artificial data resembling the real data as much as possible. Furthermore, they can be used to enhance the image by scaling up the image and increasing its resolution.<sup>87</sup>

CNNs are the most prominent architecture used for medical image analysis, specifically for microcirculation analysis.<sup>77,88–90</sup> They were originally developed in 1998 to recognize zip codes<sup>91</sup> and digits. Gradually, they have become the most relevant architecture in image classification.

## CNNs

In this subsection, we dive into the details of a CNN.

### A CNN pipeline

A typical CNN image classification pipeline can consist of three main stages, which are:

- Dataset: where a set of images is labeled with their corresponding classes
- Learning: a model that has learned from every class of the data
- Evaluation: evaluating the performance of the model on a set of data that it has not been previously trained on<sup>92</sup>

The data at the first step can be further split into a training set and a validation set, so one can evaluate the model before new data come in. In this way, a pipeline can end up with three sub-datasets: a training set, a validation set, and a test set. A good CNN should be invariant to the listed challenges below:<sup>93</sup>

- Variation of viewpoint: as regards the camera, a single instance of an entity may be directed in various ways
- Variation of scale: the scale of visual groups is often variable
- Deformation: a term used to describe the process of changing many interesting things, which are not solid bodies and may be deformed dramatically
- Occlusion: a term used to describe a situation where objects of interest can be obscured. Just a small portion of an item (a few pixels) can be observable at any given moment
- Illumination: different brightness levels on different parts of the image can significantly affect algorithm performance
- Clutter in the background: the objects of interest can blend in with their surroundings, making them difficult to spot
- Variation within a class: the categories of interest, such as chair, may be very general. These artifacts come in various shapes and sizes, each with their distinct appearance

Data preprocessing can be done in two ways: mean subtraction and normalization.<sup>94</sup> The most popular form of preprocessing is mean subtraction.<sup>94</sup> It entails subtracting the mean from all the data's features, with the geometric understanding of centering the cloud of data (pixel values) in all dimensions (1 in case of white and black, 3 in case of RGB). Normalization is the process of bringing the sizes of the scale closer together. This normalization can be accomplished in two forms. The first is to zero-center the value and divide it by its SD. The second is to normalize each dimension such that the minimum and maximum values are between negative 1 and positive 1. In the next section, we talk about the different parts that make up a CNN: the fully connected (FC) layer, the convolutional layer, and the pooling layer.<sup>81</sup>

### The anatomy of a FC layer

In this subsection, we describe the details of a FC layer.<sup>94–96</sup>

Weight initialization: each neural network node consists of parameters known as weights, which take numerical values and contribute to the total of the other weighted input signals. It is recommended to initialize a neural model's weights with a positive random value to help it converge into the results with the epochs. There are other ways to initiate the weights in a neural network using Xavier weight initialization and He weight initialization. At every iteration, the algorithm tries to find the most appropriate set of weights for each node. For example, in stochastic descent, the neural network weight is gradually adjusted by a cost function.

A cost function is used in supervised learning to measure the difference between the predicted result and the expected value. There are two types of loss functions: one is used for

classification when the classes are of a fixed size and set, while regression is used to predict a quantified value.

Regularization: there are four ways to regularize a neural network: L2 regularization, L1 regularization, max norm constraints, and dropout. The most popular form of regularization is L2 regularization. It can be applied by penalizing all parameters' squared magnitudes. The intuitive understanding of the L2 regularization is that it strongly penalizes high weight vectors while favoring diffuse weight vectors. In L1 regularization, the values become almost invariant to noisy values since they use the most critical inputs and almost all of their irrelevant information is eliminated. In general, if a CNN architecture would be used for feature selection, then L1 regularization handles it best over other regularization methods.

Maxout constraints on the maximum average, which is another method to regularize the values by imposing an absolute upper limit on the magnitude of each neuron's weight vector, is enforced using projected gradient descent. In fact, this entails updating the parameters as usual and then clamping the weight vector to enforce the constraint. Dropout is an easy, efficient tool for removing a random percentage of neurons that the developer of the neuron can specify.

Activation function: this part in a CNN takes in a number and performs a mathematical operation on it. Many activation functions can be performed in a CNN; however, the common ones are sigmoid, tanh, ReLU, leaky ReLU, and maxout.

Hyper-parameter optimization is finding the right set of values for all of the above properties in a neural network. Mainly, finding the initial learning rate, the decay constant of the learning rate, and the regularization values can strengthen or weaken the neural network. During the forward pass, the score is calculated by applying the operations of all the blocks on the input value. On the backward pass, we compute an updated value of the weights that minimizes the loss function, which increases the overall accuracy of prediction.

### **The convolutional layer and the pooling layer**

In this section, we describe the details of the convolutional and pooling layer.<sup>93,94,97</sup>

Similar to the anatomy of a node in a FC layer explained in the previous section, a convolutional block is initialized with weights and has a cost function, activation function, and regularization. The difference is that convolutional net architectures presume that inputs are pictures, taking the spatial variance between the inputs into consideration. These processes and parameters decrease the network's overhead, reducing the number of parameters the program will execute as well as increasing the speed. If we were to pass a full HD image to a neural network ( $1920 \times 1080 \times 3$ ), there would be approximately 6.2 million weights to initialize—one weight for each pixel. Instead, we pass that image to a series of convolutional blocks first to reduce the number of weights from 6.2 million to possibly several hundred thousand weights, without compromising accuracy. This reduction of pixels is achieved by transforming the image into a more representative form using a CNN architecture.

A simple CNN architecture consists of four parts:

- input, where data is loaded as a matrix in its raw form (for example, a full HD image will be a matrix of  $1920 \times 1080 \times 3$ )

- a convolutional layer, which takes in a value and applies the dot product operation followed by some kind of non-linearity (activation function)
- the pooling layer, which applies downsampling to the matrix
- the FC layer, which computes a prediction associated with each class. With each iteration, the values for the convolutional and FC layers change as their weights are adjusted, while the activation function and pooling layer values stay constant throughout the whole process. The details of each part are described below

A convolutional layer consists of a set of filters, forming a matrix that is typically  $5 \times 5 \times 3$ . However, these values are strictly experimental and depend on the image used. A convolutional layer can have any number of filters specified by the machine learning expert. The filter values are randomly generated and are updated in the back-propagation to reduce the loss value, which increases the probability of correct classification. Each filter slides over the image with a stride value production of a smaller image known as the activation map. The number of strides dictates the number of weights to be initialized later: as the number of strides increases, the number of weights to be initialized at the FC decreases. For example, a stride of two means the filter will jump two pixels in the image before applying the dot product and skipping some pixels. There is a trade-off. Each filter can detect a different image property; for example, a filter can detect horizontal or vertical edges and types of colors. The number of filters in this layer is referred to as the filter depth. Each filter produces an activation map, and these are then stacked on top of each other. Padding is another concept in the convolutional layer that involves adding zeros around the borders of the input image to preserve the sizes of the input and output shapes. Combining the filter depth, stride, and padding, the output volume of the convolutional layer can be calculated. Output volume:  $(W - K + 2P/S) + 1$ , where  $W$  is the input height/length,  $K$  is the filter size,  $P$  is the padding, and  $S$  is the stride. For example, a single black-and-white image of dimension 200 with a stride of 1, padding of 0, and filter size of 5, with 32 filters, will turn the image from a  $200 \times 200 \times 1$  to a  $196 \times 196 \times 32$ . Each layer represents the activation map produced by the individual filter. This number is very high compared with the input value, and thus a pooling layer is applied to take down the number of parameters to initiate weights for. A pooling layer is typically inserted between consecutive convolutional layers before passing the final value into an FC layer. The pooling layer is applied individually on the activation map; the depth does not change, but the activation map's width and height are reduced. The pooling layer scales the image and takes in two parameters: the window size and stride. The larger these values are, the smaller the output image will be. There are several types of pooling, but the most common are average and max pooling. Average pooling takes the mean of a window of pixels, while max pooling takes the maximum value within the window. So a pooling layer with a  $2 \times 2$  window and a stride of 1 halves the image's dimensions from  $196 \times 196 \times 32$  to  $98 \times 98 \times 2$ . Applying several of these between consecutive convolutional layers inevitably reduces the dimensions of the activation map to the most relevant for prediction. The values from the last pooling layer are

flattened or converted from 2D to 1D and passed to the FC layer explained earlier. In the next section, we look into how the convolutional layers, activation function and pooling function are used as fundamental building blocks to predict the image classification.

### Types of CNNs

LeNet,<sup>91</sup> created in 1998, uses a five-level CNN: two convolutional layers with three FC layers. The convoluted layers were made up of a  $5 \times 5$  filter with a stride of 2 with sigmoid function, followed by an average pooling layer of  $2 \times 2$  with a stride of 1. The FC layer contained 120, 84, and 10 neurons, respectively, using softmax as an activation. This CNN's input data was a grayscale  $32 \times 32$ , which is relatively small with today's standard.

AlexNet,<sup>98</sup> released in 2012, outperformed LeNet. It used an eight-layer-deep CNN: five convolutional layers, two hidden layers, and one FC output layer. There are several significant differences between AlexNet and LeNet. AlexNet uses ReLU for the activation function. Moreover, AlexNet uses dropout instead of weight decay for regularization. AlexNet uses more neurons and different filter sizes for each convolutional net. For the FC layer, AlexNet uses 4,096, 4,096, and 1,000 neurons, respectively, compared with the 120, 84, and 10 of LeNet. AlexNet uses  $11 \times 11$ ,  $5 \times 5$ , and  $3 \times 3$  for the convolutional layer, while LeNet uses two  $5 \times 5$  filters.

Visual Geometry Group (VGGNet),<sup>99</sup> developed in 2014 at Oxford University, consisted of these basic CNN building blocks: a convolutional layer, activation function ReLU, and a maximum pooling layer. It used  $3 \times 3$  filters with a padding of 1 and a  $2 \times 2$  pooling with a stride of 2. Moreover, the paper authors experimented with several different architectures and concluded that deeper and narrower layers get better results than fewer and wider convolutional layers.

The above three architectures have a common pattern of using the convolutional layer followed by pooling with minimal tweaks. The next three architectures came later and use a slightly different design pattern.

GoogLeNet,<sup>100</sup> published in 2015, outperformed the previous three architectures, achieving close to human-level performance with its new inception module. An inception block uses four different blocks in the input images and then concatenates their output. The first three blocks apply a convolutional layer of different window sizes, while the fourth applies max-pooling and then the convolutional layer. The number of blocks in an inception module can be tested on different sets using hyperparameter tuning. GoogLeNet outperformed the others because it aims to extract the most spatial information possible from each layer. Instead of finding information from the previous block's output, inception aims to explore each image with different filter sizes. It is like taking the same photograph with different lens magnifications.

Residual neural network (ResNet)<sup>101</sup> was published in 2016 and designed to address the increasing complexity of making deeper neural networks. As the number of blocks increased, the accuracy gain per block decreased to where making the network even deeper started to adversely increase complexity and computational power and reduce accuracy. This was achieved using the residual block, which utilizes a skip connec-

tion with heavy batch normalization. Like the VGGnet convolutional layer design, a ResNet consists of two consecutive  $3 \times 3$  convolutional layers with an activation function. In addition, there is a connection between this block's input and the output, known as the skip connection. Like GoogLeNet, which uses four modules within the inception block, ResNet uses four modules within the residual block. It also uses a global average pooling layer before the FC layer.

For microcirculation image analysis, it is not enough to use a CNN architecture; these architectures only detect whether an image has a capillary. They cannot pinpoint the capillaries' location. To do that, we need to extend the CNN architecture with an object detection architectures.

### DEEP LEARNING OBJECT DETECTION TECHNIQUES

Object detection techniques aim to estimate the location and label of an object in an image.<sup>102</sup> The object detection part extends the CNN architecture. Object detection techniques can generally be split into two distinct categories. The first category, which is a two-step method, aims to first locate the object in the image (object localization) and then estimate the category of the object (object classification). These architectures can be referred to as the region proposal-based framework. The second category is a one-step method, which aims to locate and categorize the objects in one go. These architectures are known as the unified framework.<sup>103</sup> Before these methods were developed, the field was dominated by different techniques known as the scale-invariant feature transform (SIFT) technique (from 1999 to 2012).<sup>98,104,105</sup> Object detection architectures are benchmarked by measuring the mean average precision (mAP) and efficiency (speed of detection per frame) on standardized datasets.<sup>54–56,106–108</sup> In this section, seven selected architectures from the region proposal-based framework are described along with five selected architectures from the unified framework.

#### Region proposal-based framework

R-CNN (rich feature hierarchies for accurate object detection and semantic segmentation)<sup>109</sup>: when this paper was released in 2014, the best-performing object detection architectures had plateaued from 2010 to 2012, with an accuracy of 35% for the most popular datasets.<sup>105</sup> This algorithm achieved 20% higher accuracy than its predecessor with the VOC 2012 dataset.<sup>109</sup> This method was termed regions with CNN features, or R-CNN. It is also one of the first methods to propose a two-step approach, and many subsequent methods have been based on this approach from 2014 to this day.<sup>93</sup> Region proposal-based frameworks are inspired by the combination of deep CNN (DCNN)<sup>98</sup> and region proposals.<sup>110</sup> R-CNN takes in an image, applies a segmentation mask to it, and extracts the top 2,000 promising bounding boxes on that segmentation. The bounding boxes are of different scales, increasing the probability of identifying different sizes or shapes. It then computes the features of these boxes using a CNN and classifies each region with a linear SVM. On the other hand, this method can be slightly adjusted if the training data is low to apply a supervised pre-training CNN on the ImageNet followed by fine-tuning the low training data. The R-CNN consists of two concurrent pipelines: the region detector and the object detector. For the region

detector, R-CNN takes an image and applies a non-deep learning model called selective search to extract approximately 2,000 regions of interest (ROI). These regions present the places in the image where an object is more likely to reside. The proposed region is then warped or cropped to fit a specific dimension before being passed into the object detector. The object detector applies CNN + max-pooling with an activation function to calculate the feature map. The feature map is then passed to an FC layer to create a 4096-dimensional vector. This vector is passed to a classification head that tries to figure out the class, and the regression network that tries to refine the box coordinates. The classification head is optimized using cross-entropy loss, while the regression head is refined using L2 loss. The model is trained by optimizing the model first on the classification loss, then the regression loss. This can take up to several days, with large storage space, since all the features computed from the proposed regions require many gigabytes. This paper methodology consists of three modules.

- Generating the region proposals: selective search method is used to suggest the regions
- Extracting the features using CNN: a 4096-dimensional vector is extracted from each region generated by the previous step using the method applied by the DCNN.<sup>111</sup> The features are then computed by subtracting the mean from a  $227 \times 227$  image through a five convolutional layer and a two-FC layer. The output region is warped equally with  $p = 16$ ; however, the paper suggests that alternative values can be used
- Extracting the class using SVM: each region is scored using an SVM, and a greedy non-maximum suppression for each class is applied independently in the proposed region. If the two regions have an intersection-over-union higher than a threshold, one region will be rejected

If there is a lack of training sets, the paper suggests the addition of these two modules:

- Supervised pre-training: it starts by training the CNN on one of the large datasets as an image classification problem using the Caffe CNN library
- Domain-specific fine-tuning: the wrapped regions created in step two of the above method are used to fine-tune the CNN parameters using stochastic gradient descent.

R-CNN applies the results on the PASCAL VOC 2007 dataset and achieves a 53.7% mAP. These results are a big jump from the previous algorithms proposed for this dataset at that time (year 2010) where the highest achieved mAP was 35.1%. R-CNN is a big step toward building a high-quality object detection architecture after the SIFT era and in the DCNN era. This is noticeable in the jump in accuracy introduced by R-CNN. However, there are some drawbacks to using R-CNN. First, R-CNN has multi-stage, multi-step modules that need to be optimized individually to achieve good results, which increases the chances of introducing inaccuracies and makes training time notably longer. Second, R-CNN uses a FC layer that requires a fixed input shape. Moreover, approximately 2,000 regions are extracted, which, one can argue, is way too much in sparse images and way too little in denser images. Such disadvantages have led

to the development of successors such as SPP-Net (spatial pyramid pooling), Fast R-CNN, Faster R-CNN, region-based fully convolutional network (R-FCN), feature pyramid network (FPN), and Mask R-CNN which are presented in the next paragraphs.

SPP-Net (spatial pyramid pooling in deep convolutional networks for visual recognition)<sup>112</sup>: this method introduces two changes to the existing R-CNN architecture. First, it aims to tackle the challenge of having a fixed-size window since important information can be lost or distorted, reducing accuracy. Second, SPP-Net computes the feature maps for the images instead of repeatedly computing them on each ROI region as the R-CNN. The challenge of having a fixed-size window is tackled by adding a spatial pyramid to the top of the last convolutional layer before the FC layer. Instead of cropping or warping the image, this method aggregates the information by pooling the features and feeding it to the FC layer. The spatial pyramid pooling is an extension of the Bag-of-Words model released in 2006.<sup>113</sup> The difference between the R-CNN method and the spatial pyramid pooling methods can be illustrated as follows: the R-CNN method takes in an image, applies crop/warp, and passes it to the convolutional layer and then the FC layer. The SPP-Net method takes an image, passes it directly to the convolutional layer, applies the spatial pyramid pooling, and then passes it to the FC layer. The SPP-Net takes in the feature maps from the last layer of the convolutional layer to create feature maps of fixed-length feature vectors regardless of the input image size. Images with different sizes can be pooled and aggregated into a spatial pyramid, which is then passed to the FC layer. When this paper was released, four existing object detection architectures were compared with their non-SPP counterparts (ZF-5, Convnet\*5, Overfeat-5, Overfeat-7), and the CNN with SPP-Net showed state-of-the-art classification results on Pascal VOC 2007 and ranked at number 2 on the ILSVRC 2014 competition.<sup>93</sup> The spatial pyramid pooling method is more efficient than its predecessor since it obtains a significant speedup. The speedup is due to the fact that the CNN layer generates a feature map by running one iteration on the image. Furthermore, it is more accurate since it can learn feature maps from any scale without losing information to cropping or warping. The multi-level pooling makes the input images more robust to deformation. The main drawbacks of SPP-Net are that it is still a multi-stage, multi-step pipeline (feature extraction, network fine-tuning, SVM training, bounding box regressor, feature caching), making it relatively slow. Furthermore, the authors of the paper mention that the accuracy of SPP-Net layers drops when using deeper CNN, since tuning the network will not update the layers before the pyramid layer, leading to reduced accuracy and a very difficult challenge in implementing back-propagation.

Fast R-CNN<sup>114</sup>: this paper addresses the problems arising from the SPP-Net and R-CNN architectures. Until this paper's 2015 publication date, object detection methods required generating several hundred regions known as "proposals" to create a feature map; then, the proposals generated estimated the localization of the object. These proposals reduced speed and accuracy while increasing complexity. Similar to R-CNN, this method uses selective search to find the regions and then passes the regions to the object detector. The method also consists of two SVM heads: one for classification to get the class category and

the second regression to calculate the bounding box coordinates. The difference is that instead of running the CNN several times on the ROI, it runs it only once by introducing ROI pooling. Second, it streamlines the process on the object detector side, where it jointly classifies and learns the object's location simultaneously by using multiclass loss. This method generally achieves a higher mAP by having a single stage for the training with a multi-task loss. The increased accuracy is obtained by updating all layers. The speed is achieved by not requiring the features to be cached and because Fast R-CNN learns the softmax classifier and bounding box regression together rather than in two separate processes. These improvements have led to a major decrease in storage space needed. Unlike the R-CNN, Fast R-CNN creates a feature map from the entire image. Furthermore, Fast R-CNN method improves efficiency compared with the SPP-Net: 3× in training and 10× in testing. The authors report that “Fast R-CNN trains 9× faster than R-CNN on the VGG-16 and 213× faster at test-time, with a higher mAP on the PASCAL VOC 2007, 2010 and 2012 dataset ...”<sup>114</sup> These speed improvements result from a single process that updates all layers without requiring feature caching. Moreover, to reduce the time spent on the FC layers, Fast R-CNN uses a truncated singular value decomposition (SVD) to accelerate the testing procedures.<sup>115</sup> This method has significantly increased the speed and efficiency of object detection, firstly, by streamlining the whole process and, secondly, by applying SVD on the testing set. Thus, Fast R-CNN is more of a speed improvement than an accuracy improvement. On the same dataset that took 84 h to train, Fast R-CNN performed it in 9 h. A major drawback is that Fast R-CNN still relies on external region proposals that make the whole process relatively slow. It uses the selective search method to find the ROI. Furthermore, later research has concluded that convolutional layers are sufficient to localize objects; therefore, adding an FC layer slows down the process unnecessarily.

Faster R-CNN: toward real-time object detection with region proposal networks<sup>116</sup>: optimizations introduced by SPP-Net and Fast R-CNN have exposed the fact that using external region proposal methods slows down the process. Previous networks mainly relied on selective search<sup>110</sup> and Edge box<sup>117</sup> to create region proposals. This paper introduces a region proposal network (RPN), which aims to replace the selective search and Edge box. The RPN introduces an almost cost-free proposal computation. For RPN to compete with methods such as selective search, it has to predict the ROI of multiple scales and ratios from an image much faster. Thus, RPN introduces a novel concept of creating anchors on the feature maps. The RPN layer takes in the feature map and generates rectangular object bounds using CNN, which are the new ROI. Faster R-CNN can be trained end-to-end like Fast R-CNN. The RPN output tells the Fast R-CNN where to look. The Faster R-CNN architecture is complex because it has several interconnected parts. The RPN first initializes anchors of different ratios and scales on the feature maps created by the convolutional layer. The paper's author uses nine types of anchors when the ROI is decided on. The anchors are off three scales and three ratios. These anchors are mapped and fed into the two FC layers, where one layer is responsible for the category classification and the other for the box regression. RPN shares the convolutional feature with the

Fast R-CNN, enabling the same efficient computation as mentioned in the methodology of the previous paper. On the VGG-16 model,<sup>61</sup> the Faster R-CNN efficiently performs all steps on 5 fps with an accuracy exceeding all recorded results on the VOC 2007 dataset with 73.2% mAP, and on VOC 2012 with 70.4% mAP. Although this method is several times faster than Fast R-CNN, it still relies on applying several hundred ROI per image to detect the region of interest. This leads to computations not being shared after the ROI layer, reducing this method's overall efficiency.

R-FCN: object detection via region-based fully convolutional networks.<sup>118</sup> In Faster R-CNN, each region proposal had to be cropped and resized to be fed into the Fast R-CNN network. The R-FCN attempts to speed up the network by converting this process into fully convolutional. It aims to swap the costly per-region subnetworks with a fully convolutional one, thus allowing the computation to be shared across the whole image. Furthermore, the R-FCN differs from the Faster R-CNN in the ROI pooling layer. The R-FCN proposes a method to use convolutional layers to create an ROI subnetwork. It uses the RPN introduced in the previous method to extract features and pass them on to the R-FCN. The R-FCN then aggregates the output of the last convolutional layer and generates the scores for each ROI. Instead of cropping the regions from the feature map, the R-FCN inputs the feature map into the regression and classification heads, creating an ROI map on the feature map. R-FCN uses ResNet-101 as the backbone of its architecture.<sup>101</sup> ResNet-101 has 100 convolutional layers with a 1,000-FC layer. The average pooling layer and the convolutional layers are removed, and the convolutional layer is used to compute the feature maps. A layer applied to the last convolutional block generates the score maps. A sibling convolutional layer is also applied to calculate the bounding box regression. On the PASCAL VOC 2007, it achieves an 83.6% mAP with the 101-layer ResNet. It suggests the same accuracy as the Faster R-CNN but achieves 20 times the speed of its Faster R-CNN counterpart. Thus, R-FCN introduces two advantages over its predecessors: first, CNN is faster than FC layers. Second, the network becomes scale invariant since there is no FC to restrict the input image size.

FPN (feature pyramid networks for object detection)<sup>119</sup>: this method was designed to address an issue with Faster R-CNN. Faster R-CNN was generally made to address the scale-invariance problem introduced by Fast R-CNN. Faster R-CNN takes an input image and resizes it accordingly, meaning that the network has to run on the image several times with different box sizes, making it slow. The FPN deals with these different scales while maintaining the speed. The FPN is an extension of Faster R-CNN in the same manner that R-FCN is an extension of Faster R-CNN. Having a robust scale invariance is important for object detection since the network should be able to recognize an object at any distance from the camera. Faster R-CNN aimed to tackle this issue by creating anchor boxes. This proved time-consuming since the anchor boxes had to be applied to each ROI. The FPN, however, creates multiple feature maps that aim to represent the image at different scales. Hence, the feature map in RPN is replaced by the FPN, removing the necessity of having multi-scale anchor boxes. The regression and classification are applied across these multiple feature maps. The



FPN takes in an input image and outputs multiple feature maps representing smaller height and width but deeper channels known as the bottom-up pathway. The feature maps generated by the FPN goes through a  $1 \times 1$  convolutional layer with a depth of 256. The lateral connection is then applied, which adds the feature elements to the upsampled version of the feature map. Faster R-CNN runs on each scale, and predictions for each scale are generated. FPN comprises two paths: the bottom-up that uses ResNet and the top-down. In the bottom-up approach, CNN is applied to extract features. On the top-down pathway, the FPN constructs a higher resolution layer, but the object locations are no longer accurate because of the down- and upsampling. Therefore, FPN adds a lateral connection between the constructed layers to increase the probability of predicting locations. This method runs at 5 fps, as benchmarked by the previous methodology with a state-of-the-art result on the COCO 2016 dataset. Images have objects with different scales, making it challenging to detect them. When using several anchor boxes to detect objects with different scales, the ratio seems to be memory- and time-consuming. FPN seems to push the accuracy boundaries by introducing a pyramid of feature maps to detect objects of different sizes and scales in an image. It is important to highlight that FPN is a feature detector and not an object detector. Therefore, FPN has to be used with an object detector in its ROI.

Mask R-CNN<sup>120</sup>: this method extends Faster R-CNN by adding another layer to predict the object mask in parallel with the existing bounding box layer. This is a framework that enables instance segmentation on a state-of-the-art level. The mask branch added to the Faster R-CNN is a small FCN applied to each ROI, which predicts on a pixel-to-pixel basis. In brief, the Faster R-CNN has two stages: the RPN and the Fast R-CNN combined. The Mask R-CNN adopts the same notion as an identical first stage of RPN, and in the second stage, it outputs a mask for each ROI in parallel to the predicting class and box. The branch added to the second layer is an FCN on top of a CNN feature map. The ROI poolings lead to misalignment; therefore, the RoIAlign layer is proposed to preserve the pixel-level alignments. The main method Mask R-CNN introduces is the RoIAlign, which preserves the pixel-spatial correspondences and replaces the quantization from the ROI pooling with bilinear interpolation. The state-of-the-art results are achieved by ResNeXt101-FPN in the COCO dataset. The additional mask branch added introduces minor computational additions. Mask R-CNN is a very promising instance segmentation method that is very flexible and efficient for instance-level segmentation. However, as with the original Faster R-CNN, this architecture struggles with smaller-sized objects, mainly because of the feature maps' coarseness.

Other image classifications and object detections include but are not limited to NOC, Bayes, MR-CNN and S-CNN, HyperNet, ION, MSGR, StuffNet, OHEM, SDP+CRC, SubCNN, GBD-Net, PLANET, NIN, GoogLeNet, VGGNet, ResNet, DenseNet, RetinaNet, ResNet, Corner Net, Inception, Hourglass, Dilated Residual Networks, Xception, VGG, DetNet, Inception, dual path networks (DPN), FishNet, ResNeXt, and GLoRe.<sup>93,103,121</sup>

For microcirculation analysis, we conclude that deep convolutional neural networks have lifted much of the burden for feature engineering, which was the main focus in the pre-D-CNN era,

and changed the focus to designing more accurate and efficient network architecture. Despite the great successes, all methods suffer from the intense labor of creating the bounding boxes. All "newer" methods need exponentially more RAM and GPU in exchange for increased accuracy.

Furthermore, detecting small-size objects and localizing these objects remains a challenge. Using the stated architectures still requires an experienced machine learning engineer to select the appropriate parameters of the algorithms to learn the patterns of the small-sized objects. Several solutions have been suggested by the literature, including multi-task learning (Stuffnet),<sup>122</sup> multi-scale representation (IONet),<sup>123</sup> and context modeling (HyperNet).<sup>124</sup> On the other hand, methods have been proposed to deal with large data imbalances between the objects and the background, such as the online mining algorithms (OHEM).<sup>125</sup> For microcirculation analysis, we believe that a region proposal-based framework achieves better microcirculation data accuracy overall.

#### Unified-based framework

You Only Look Once (YOLO)<sup>126</sup>: YOLO is a unified-based framework for object detection suggested by Redmon et al.<sup>126</sup> The most significant difference between this architecture and the methods in the region proposal-based framework is the ability to track objects in real time. As mentioned earlier, Fast R-CNN proposes 2,000 regions to be predicted, while YOLO takes that down to 100 regions. On a Titan X GPU, YOLO can classify up to 45 frames per second compared with Fast R-CNN at 0.5 frames per second. YOLO takes a  $224 \times 224$  image as an input and divides the image into several grids. It then classifies each object within that grid by giving it two scores: what class it belongs to and confidence percentage. The classification is done by a 24-convolutional layer with a 2-FC layer. According to the tests, YOLO was ineffective at localization, and had low accuracy with comparison to R-CNN. Despite the high speed of YOLO, the low accuracy makes it an unsuitable choice for microcirculation analysis.

YOLOv2<sup>127</sup>: YOLOv2 addresses the precision issues brought by YOLOv1. It first replaces the CNN classifier with DarkNet19 instead of GoogLeNet. DarkNet19 is a simpler classifier utilizing 19 convolutional layers followed by 5 max-pooling layers, allowing for faster performance on the same dataset. It also removes the FC layer for prediction and uses the anchor boxes method instead, increasing the recall accuracy by 7%. Batch normalization is added between each convolutional layer, increasing the mAP by 2%. Furthermore, it increases the image input from  $224 \times 224$  to  $448 \times 448$ , which increases the mAP by an additional 4%. In Faster R-CNN, the size of the anchor boxes is defined beforehand, and YOLOv2 utilizes k-means clustering on the training set to find the right aspect ratio of the anchor boxes to use, increasing its accuracy by a further 5%.

YOLOv3<sup>128</sup>: this is an improved version over the YOLOv2 that increases overall accuracy with multi-scale labeling of small objects. YOLOv3 uses three separate feature maps to predict the ROIs. It also uses DarkNet53 with independent logistic classifiers, allowing it to detect multi-overlapping objects in the image. With these changes, YOLOv3 is suited to detect smaller objects within the grid, but it performs worse with medium to larger objects.

Single Shot MultiBox Detector (SSD)<sup>129</sup>: this improved the detection precision of a one-stage detector by implementing multi-reference and multi-resolution detection techniques. SSD detected objects of different sizes and scales across the network instead of just applying detection on the last layer. SSD maintains the speed of YOLO but has higher accuracy on the same standardized sets used to benchmark YOLO. SSD uses VGG16 as its backbone for image classification.

RetinaNet<sup>130</sup>: this introduces focal loss, which increases the prediction accuracy on small and medium objects compared with the previously mentioned detectors. In an image, the object of interest is relatively smaller than the background image. Therefore, the number of background images creates a class imbalance. The focal loss function aims to increase the weight of the minority class while reducing the weight associated with the majority class. RetinaNet archives comparable accuracies with the region proposal-based framework at the expense of speed.

CornerNet<sup>131</sup> challenges the use of anchor boxes by stating that they create the data imbalance issue in the first place. It also states that anchor boxes create unnecessary parameters that have to be tuned, which slows down the training and prediction time. Instead, CornerNet uses key points in a bounding box with a single convolutional neural network. It achieves the highest accuracy when compared with the standard benchmark dataset; however, it is slower than YOLO.

When examining the architectures in the unified framework, we generally notice a trade-off between speed and accuracy. With the above methods, as accuracy increased, the time for detection also increased. In microcirculation analysis, having an accurate method is more important than a fast method. Moreover, the difference in time analysis between the unified framework and the region proposed framework in microcirculation image analysis can boil down a few seconds. Therefore, we recommend the use of a region proposed framework for microcirculation analysis.

### Upscaling images using deep neural networks

From our review, the microscope videos have very low resolution. Upscaling the image might help the researcher annotate the data better. The upscaling process involves improving an image's details by increasing the dimensions and interpolating those extra pixels using a mathematical method. These mathematical methods include an enhanced deep super-resolution network (EDSR),<sup>132</sup> an efficient sub-pixel convolutional neural network (ESPCN),<sup>133</sup> a fast super-resolution convolutional neural network (FSRCNN),<sup>134</sup> and a Laplacian pyramid super-resolution network (LapSRN).<sup>135</sup> EDSR employs an architecture similar to ResNet without the batch normalization layer and the ReLU activation layer after the residual block. This architecture can be used to create a scale factor of 2. ESPCN extracts the feature maps and applies the upscaling at the end of the network. Like ESPCN, FSRCNN applies upscaling at the end of the network with a smaller filter size. LapSRN is based on the Laplacian pyramids concept, upscaling gradually through the network.

### TRADITIONAL COMPUTER VISION OBJECT DETECTION TECHNIQUES

In this part of the review, we present computer vision object detection techniques that can be used for microcirculation anal-

ysis. These presented methods do not use neural networks for classification. Such techniques are also known as feature descriptors; they were gaining momentum from the early 1990s until the rise of deep learning in 2012.<sup>103</sup> Although feature descriptors have fallen out of favor compared with deep learning with the benchmark datasets, they are still very relevant for microcirculation analysis. Their computational power and simplicity make these algorithms easier to implement on low-powered or battery-powered devices in hospitals.

Computer vision techniques aim to locate the image of interest from the background by distinguishing between edges, colors, textures, corners, and other image properties. Such traditional computer vision techniques need the values coded beforehand, which were found via trial-and-error methods and domain expertise. Below are three computer vision technique detection methods that can be used for microcirculation analysis.

The template matching-based object detection<sup>136</sup> methods consist of two steps. The first step is the template generation step, in which a template is generated by an expert based on the training set; the second step involves matching new data with that template-based image. A similar measure is then applied to detect similarities between these images. Statistical methods, such as the sum of absolute differences or Euclidean distances, can quantify the similarities between the template and test data. The template matching detection stage can be further categorized into methods: rigid template matching (RTM) and deformable template matching (DTM). Further modifications for the stated template methods include the SIFT, the speeded-up robust features, and the binary robust independent elementary features.

The main disadvantage of RTM is that it is sensitive to slight changes in viewpoint, shadows, and other challenges, as was stated earlier, while DTM needs a lot of geometrical engineering in the template beforehand. Moreover, these templates require two independent parameters to be tuned, the template to be generated from the training set, and the most suitable method for measuring similarities to be selected. This makes this approach time-consuming for the case of microcirculation analysis.

Another set of methods involves knowledge-based object detection.<sup>137</sup> These can be further divided into geometric knowledge and context knowledge. *A priori* knowledge of the shape is encoded into the geometric knowledge methods. However, this is extremely difficult with capillaries since the shapes are irregular. Context knowledge encodes the spatial relationship between the object and the background around how the neighboring pixels interact. Again, due to the different shades of skin and blood, this method is not preferred for microcirculation analysis.

Object-based image analysis (OBIA)<sup>138</sup> is the most promising for microcirculation analysis and comprises two parts: the image segmentation part and the object classification part. OBIA aims to group similar pixels together based on statistical methods. In the case of microcirculation analysis, we would like to highlight the pixels of capillaries and separate them from the background. Promising methods under the OBIA techniques are background subtraction methods, geometric transformation methods, and image thresholding techniques. These methods can be used in combination with each other or independently to detect capillaries in an image.

### Background subtraction methods

Background subtraction is a step in image preprocessing, where the goal is to remove the background and keep the object of interest.<sup>139</sup> The three methods stated next can take in an image of microcirculation and attempt to calculate an approximation of the location of the capillaries.

- Mixture of Gaussian method: this method uses Gaussian mixture-based background/foreground segmentation.<sup>140</sup> It takes in a pixel with a K Gaussian distribution and attempts to model the background. This method is based on using the L-recent window version after the sufficient statistics equation is calculated
- Improved mixture of Gaussian method: this method is also based on a Gaussian mixture-based background/foreground segmentation but uses recursive equations to update the parameters.<sup>141</sup> The previous method selects the background based on K Gaussian distribution, while this method uses an adaptive density estimation<sup>142</sup>
- Statistical background image estimation: this method uses Bayesian segmentation with Kalman filters and Gale-Shapley<sup>143</sup> matching to approximate the background image

Marcomini and Cunha<sup>144</sup> compare the performances of all three above-mentioned methods using accuracy rate, precision rate, and processing time, and conclude that the improved mixture of Gaussian methods had the best performance in their experimental dataset. This has also been shown to be the best method among the three for background selection in the CapillaryNet paper.<sup>36</sup>

### Image thresholding techniques

In its simplest form, thresholding involves changing a pixel value if it is above or below a certain value.<sup>145</sup> This value or threshold can be determined by several methods, and the change of value can also be calculated by different methods. In microcirculation, this can be beneficial for determining the set of values that represent the capillary and the other that represents the background. Listed below are five thresholding techniques that can be used for microcirculation analysis.

- Binary threshold: this method takes in two values—the threshold value and the value to be given if the value is higher than the threshold value. The values under the threshold value will be set to zero
- Truncating threshold: similar to the above method, it takes in two values. However, anything lower than the threshold value remains the same, while anything higher gets the assigned value
- Zero threshold: anything lower than the threshold value becomes zero, while anything higher stays the same

The above methods are fairly simple, and these values are determined by the user. However, they are not the optimal method if different parts of the same image have different illumination. The object of interest may have a higher or lower value depending on the light; therefore, the next two methods were designed to deal with this issue.

- Adaptive thresholding<sup>146</sup>: thresholding is applied locally on some pixels rather than globally on the whole image.

Thresholding can be calculated in two ways: mean of an area or weighted sum where the weights are decided by a Gaussian window. The size of the window is decided by the user. This way, every window-sized part of the image gets a threshold applied to it based on a calculation

- OTSU binarization<sup>147</sup>: this method is optimal for images with two peaks in their histogram. It finds a value between the two peaks in a histogram where the variance is minimal for both classes and applies thresholding based on that.

### Edges and lines

There are several methods for detecting edges and outlines of the capillaries. Below, we list those most relevant to microcirculation detection.

- Contours<sup>148</sup>: this involves drawing a line joining the pixels with the same color or intensity. In our case, this can help highlight the outline of a capillary. This method has the highest accuracy when thresholding is applied beforehand, so more pixels have similar values. In the below method, we use a marching square algorithm, which linearly interpolates the pixel values to find the algorithm output<sup>148</sup>
- The Canny edge detector<sup>149</sup> can detect and quantize the capillary area. This is a multi-stage detector that uses a Gaussian derivative to compute the gradients. The Gaussian attempts to reduce the noise in the image, and the curves are detected by selecting the maximum pixel value
- Skeletonization<sup>150</sup> is a method used to find the central pixels within the border image to get the object topology. This method iterates over the image several times, starting from the border of the object and moving toward the center until it terminates

### Histogram equalization

An image can be enhanced using histogram equalization methods. Histogram equalization<sup>151</sup> can be done using three methods: standard equalization, contrast stretching, or adaptive equalization.

- In a standard equalization, the most frequent value is spread out to roughly have a linear cumulative distribution graph
- In contrast stretching, the image pixels are rescaled to include all values between the 2nd and 98th percentile
- With adaptive equalization, changes in pixels occur locally based on a window size rather than the whole image

### Image denoising

Image can be enhanced by reducing the noise. This is called image denoising.<sup>152</sup> There are several ways to denoise an image; total variation filters, bilateral, wavelet denoising filters, and non-local means denoising algorithm.

- The total variation filter uses the L1 norm gradient to remove noise from the image

- The bilateral filter averages the pixels based on the weighted average of the window used by the user
- The wavelet denoising filter represents the image as a wave and analyzes the wavelet coefficients. The wavelet coefficients that are under a certain threshold are represented as noise and are removed
- A non-local means denoising algorithm estimates the value of a pixel based on similar patches from other similar areas in the image. This method can be applied either by spatial Gaussian weighting or uniform spatial weighting

### STATE-OF-THE-ART MICROCIRCULATION IMAGE ANALYSIS TECHNIQUES

In this section, we present the methods used by other researchers to develop their microcirculation analysis systems. The following methods utilize computer vision techniques to segment and, in some cases, quantify the capillaries. A summary table is shown in [Table 1](#).

Dobbe et al.<sup>153</sup> used a frame averaging method to remove the cell gaps within the capillary and applied an algorithm to detect capillaries. They also removed capillaries that were out of focus since they considered them to add noise to the frame averaging method. However, this can significantly reduce the capillary density values.

The study of Hilty and co-workers<sup>34,35</sup> is similar to that of Dobbe et al.,<sup>153</sup> but with some minor changes. Hilty et al.<sup>34</sup> use an algorithm to detect capillaries that are 20–30  $\mu\text{m}$  wide. This type of detection can sometimes lead to the detection of artifacts, such as hair or stains of similar sizes. Furthermore, the mean of the images across the whole video is not always the best representation value since different parts of the video might have different lighting or capillaries that can be out of optimal focus. Moreover, videos with slight motion will have to be completely disregarded since the central line is calculated across all frames instead of per frame.

Bezemer et al.<sup>155</sup> used 2D cross-correlation to fill in gaps in the images. However, this method also has some problems because it does not consider the dynamic changes in blood flow, which can reduce the prediction accuracy.

Tam et al.<sup>156</sup> have a method to detect capillaries that requires the user to select points on the image. The algorithm then decides if there is a capillary present. This method cannot be used in a clinical environment because it would take too long to analyze a microscopy video.

Geyman et al.<sup>157</sup> take a more manual approach to find the number of capillaries. They used software to click away the major blood vessels and then applied hardcoded calculations to detect the total number of capillaries based on the number of pixels in the region of interest. This is a manual method that is particularly subject to observer differences among datasets.

Demir et al.<sup>158</sup> used a method called contrast limited adaptive histogram equalization (CLAHE) to detect capillaries. This method makes it easier to see capillaries by equalizing the contrast in an image. CLAHE is usually combined with a median filter, which is a tool that removes outliers from data. They also used an adjustable threshold to make the detection more accurate. However,

this method is not perfect and sometimes requires manual adjustments depending on the lighting and skin thickness.

Cheng et al.<sup>159</sup> created a system that allows users to increase the contrast and smoothen images of capillaries manually, in order to make them easier to see. The system also generates macros, which are sets of instructions that can be applied to future images to save time. However, the macros generated by the system may not work well if the brightness or thickness of the skin in the images changes.

To find and quantify the capillaries in an image, the Tama et al.<sup>167</sup> study uses binarization, skeleton extraction, and segmentation. First, binarization is applied to the green channel since it is assumed to have the highest contrast between the capillaries and the background. Next, the top-hat transform method is used to reduce uneven illumination, followed by Wiener filtering to remove noisy pixels. Then, the Gaussian smoothing method is used to smoothen the image. Finally, the OTSU thresholding method is applied to segment the capillaries from the background. This method relies on the user finding a reference frame from the video with the highest contrast.

The work described next use ML techniques to segment and, in some cases, quantify the capillaries.

Prentašić et al.<sup>164</sup> trained a neural network to segment the microvasculature structure. The segmentation took approximately 2 min per single image, with an accuracy of 83%. The time taken and the high-end hardware used to analyze a single image make it unsuitable for clinical use since the users would like the results instantly.

Dai et al.<sup>160</sup> employed a custom neural network for segmentation, comparable with Prentašić et al.<sup>164</sup> Dai et al.,<sup>160</sup> on the other hand, utilized five CNN blocks instead of three. For picture improvement, they employed gamma correction and CLAHE. They reported an accuracy of 60.94%, which is insufficient for application.

Nivedha et al.<sup>165</sup> classified the capillaries using the image's green channel and a support vector machine. This approach required a manual step in which the user cropped the region of interest to enhance histogram equalization. They compared several denoising filters, including Gaussian, Wiener, median, and adaptive median. They concluded that the Gaussian filter is best suited to their data. Furthermore, they examined other segmentation methods, such as OTSU, k-means, and the watershed, and determined that the OTSU approach was best suited to their data. The segmented pictures were then fed into an SVM, which produced an accuracy of 83.33%.

Javia et al.<sup>166</sup> modify the ResNet18<sup>101</sup> to quantify capillaries and use the first 10 layers of the architecture. The main limitation of the ResNet architecture is that images have to be resized to 224  $\times$  224; however, most capillary images are less than 100  $\times$  100. This means images have to be scaled up, which makes this method inefficient and uses more resources than needed. They reported an accuracy of 89.45% on their data; however, ResNet18<sup>101</sup> has 11 million trainable parameters and, with such scaling up, training time can be up to several hours, and prediction time can be up to several minutes. This can make it slow and inefficient within a clinical setting. The training and test times were not reported in this paper.

To construct their neural network, Ye et al.<sup>168</sup> used transfer learning and the Single Shot Detector v.2. Because it is accurate

**Table 1. Summary of work in literature**

Name of Technique	Type of technique	Outcome	Summary of method
Dobbe et al. <sup>153</sup>	traditional computer vision techniques	no accuracy reported	before using a traditional computer vision algorithm to locate capillaries, Dobbe et al. utilize a frame-averaging approach to eliminate the plasma and white blood cell gaps within the capillary
Hilty and co-workers <sup>34,35</sup>	traditional computer vision techniques	no accuracy reported	Hilty and co-workers identify capillaries by producing a mean picture over all frames and then sending the resulting image through two pipelines: the first categorizing vessels with diameters of 20–30 μm as capillaries, and the second classifying vessels with diameters of up to 400 μm as venules. The capillaries are then equalized using an adaptive histogram after being run through a modified curvature-based area recognition technique <sup>154</sup>
Bezemer et al. <sup>155</sup>	traditional computer vision techniques	no accuracy reported	Bezemer et al. fill the blood flow gaps caused by plasma and white blood cells with 2D cross-correlation. This is a superior strategy since it reduces the number of frames that must be ignored
Tam et al. <sup>156</sup>	traditional computer vision techniques	no accuracy reported	Tam et al. detect capillaries through a semi-automated method that requires the user to select points on the image
Geyman et al. <sup>157</sup>	traditional computer vision techniques	no accuracy reported	a manual approach using a software to remove the major blood vessels and then using pre calculations to detect the total number of capillaries using the pixels in the region of interest
Demir et al. <sup>158</sup>	traditional computer vision techniques	no accuracy reported	uses contrast limited adaptive histogram equalization combined with a median filter
Cheng et al. <sup>159</sup>	traditional computer vision techniques	no accuracy reported	generates a macro by combining different types of traditional computer vision techniques used to detect capillaries
Dai et al. <sup>160</sup>	deep neural networks	accuracy of 60.94% reported	uses a shallow convolutional neural network
Hariyani et al. <sup>161</sup>	deep neural networks	accuracy of 64% reported	uses a U-net architecture combined with a dual attention module <sup>162,163</sup>
Prentašić et al. <sup>164</sup>	deep neural networks	accuracy of 83% reported	uses a shallow convolutional neural network
Nivedha et al. <sup>165</sup>	deep neural networks	accuracy of 83.3% reported	uses a non-linear support vector machine
Javia et al. <sup>166</sup>	deep neural networks	accuracy of 89.45% reported	uses a ResNet architecture
CapillaryNet <sup>36</sup>	mixture of traditional computer vision techniques and deep neural networks	accuracy of 93% reported	uses a combination of traditional computer vision techniques, including image background subtraction, image enhancement, and shallow convolutional neural networks

**Table 2. Name and type of technique recommended for each microcirculation data parameter**

Name of technique	Type of technique	Goal in microcirculation images
Unified-based framework	deep neural networks	capillary detection and quantification
Enhanced deep super-resolution network and Laplacian pyramid super-resolution network	deep neural networks	image upscaling
Template matching-based object detection	traditional computer vision techniques	capillary detection and quantification
Knowledge-based object detection	traditional computer vision techniques	capillary detection
Background subtraction methods	traditional computer vision techniques	supports in capillary quantification
Image thresholding techniques	traditional computer vision techniques	supports in capillary quantification
Edges and lines	traditional computer vision techniques	supports in capillary quantification
Histogram equalization	traditional computer vision techniques	image enhancement
Image denoising	traditional computer vision techniques	remove noise from pixelated capillary images
Image denoising	traditional computer vision techniques	remove noise from pixelated capillary images

and not overly sophisticated, this network is ideal for finding capillaries. The authors also calculated flow velocity using a spatio-temporal diagram analysis. This procedure is time-consuming, but it is precise.

Hariyani et al.<sup>161</sup> used a U-net architecture combined with a dual attention module to try to improve accuracy for detecting capillaries in images. However, the accuracy was only 64%, which is not high enough to be used in a clinical setting.

Semiautomatic analysis is required for more accurate approaches, but more automatic methods are less precise and hence inappropriate for clinical use. Furthermore, none of the previously supercited publications used parallel frameworks to determine capillary density. CapillaryNet is totally automated and can categorize microcirculation movies in *sim* 0.9 s with 93% accuracy,<sup>36</sup> whereas CapillaryX offers parallel frameworks for calculating capillary density.<sup>169</sup>

A summary of the most recommended techniques is shown in Table 2.

## Conclusions

In this paper, we present the most promising deep learning and computer vision techniques that can automate microcirculation analysis, specifically the quantification of capillaries. Automating the quantification of capillary density might reveal important biomarkers to clinical personnel that might assist in helping critically ill patients with life-threatening diseases. With the automation algorithms, the analysis time can be reduced from minutes to several seconds and decrease interobserver variability. We start by introducing the importance of analyzing microcirculation videos. We then present the two prominent ways of automating the analysis of microcirculation videos: traditional computer vision techniques and deep learning techniques. We discuss the types of deep neural networks, then dive into details about the convolutional neural networks. Convolutional neural networks are the preferred method for analyzing images since they have the highest accuracy in image classification competitions. We present why convolutional neural networks are good at what they do and what challenges they can overcome. We then present the anatomy of a convolutional neural network by discussing the FC layer, the convolutional layer, and the pooling

layer. Moreover, we present different types of convolutional neural networks that combine these three modules differently. Since convolutional neural networks can only classify images and cannot localize the regions of the capillaries, we present deep learning object detection techniques. The deep learning object detection techniques consist of two main frameworks: the unified-based framework and the regional proposed framework. We present seven different algorithms on the regional proposed framework and six different algorithms on the unified-based framework. We then discuss traditional computer vision object detection techniques, specifically, non-ML-based object detection methods, such as background subtraction methods, image thresholding techniques, edges and lines, and image enhancement techniques. Through the sections in this article, we have recommended the algorithms that can be used to develop an automated capillary detector and quantifier. Our contribution with this article is to assist researchers and developers with where to start looking if they are to develop an automated algorithm for capillary detection and quantification.

To finalize, combining deep neural networks with traditional computer vision algorithms is the recommended approach to automating capillary detection and quantification. The traditional computer vision step is used for segmentation and area estimation, while the deep neural network will be used to classify if a capillary exists within that area. Using purely deep neural networks for the whole phase can be slow (due to the millions of parameters needed for a deep neural network) and computationally and financially expensive (due to the GPU and advanced computers needed). Using pure traditional computer vision algorithms will reduce the overall accuracy of detecting capillaries since artifacts (i.e., dirt, hair, and other objects that are on the surface of the skin) can be mistakenly quantified.

## ACKNOWLEDGMENTS

We would like to thank ODI Medical AS, a MedTech company specializing in the analysis of microcirculation systems, for the time they allowed us to speak with their medical doctors and surgeons who spent several years in the microcirculation analysis field. Specifically, we would like to thank Prof. Dr. Knut Kvernebo, a distinguished cardiac surgeon, and Anastasiya Dykyy, a biologist and medical equipment expert.

The Research Council of Norway provided the necessary funds for this project: Industrial Ph.D. project no. 305716 and ODI Medical AS.

#### DECLARATION OF INTERESTS

Maged Helmy is 50% funded by ODI Medical AS and 50% funded by the Research Council of Norway.

#### REFERENCES

- Guven, G., Hilty, M.P., and Ince, C. (2020). Microcirculation: physiology, pathophysiology, and clinical application. *Blood Purif.* *49*, 143–150.
- De Backer, D., Hollenberg, S., Boerma, C., Goedhart, P., Büchele, G., Ospina-Tascon, G., Dobbe, I., and Ince, C. (2007). How to evaluate the microcirculation: report of a round table conference. *Crit. Care* *11*, R101.
- Shore, A.C. (2000). Capillaroscopy and the measurement of capillary pressure. *Br. J. Clin. Pharmacol.* *50*, 501–513.
- Bateman, R.M., Sharpe, M.D., and Ellis, C.G. (2003). Bench-to-bedside review: microvascular dysfunction in sepsis—hemodynamics, oxygen transport, and nitric oxide. *Crit. Care* *7*, 359–373.
- Zafrani, L., and Ince, C. (2015). Microcirculation in acute and chronic kidney diseases. *Am. J. Kidney Dis.* *66*, 1083–1094.
- Ovadia-Blechman, Z., Gritzman, A., Shuvi, M., Gavish, B., Aharonson, V., and Rabin, N. (2018). The response of peripheral microcirculation to gravity-induced changes. *Clin. Biomech.* *57*, 19–25.
- Parker, M.J.S., and McGill, N.W. (2019). The established and evolving role of nailfold capillaroscopy in Connective-Tissue disease. In *Connective Tissue Disease - Current State of the Art [Working Title]*, A. Takeda, ed. (IntechOpen), pp. 1–13.
- Nama, V., Onwude, J., Manyonda, I.T., and Antonios, T.F. (2011). Is capillary rarefaction an independent risk marker for cardiovascular disease in south asians? *J. Hum. Hypertens.* *25*, 465–466.
- Houben, A.J., Martens, R.J., and Stehouwer, C.D. (2017). Assessing microvascular function in humans from a chronic disease perspective. *J. Am. Soc. Nephrol.* *28*, 3461–3472.
- de Graaff, J.C., Ubbink, D.T., van der Spruit, J.A., Lagarde, S.M., and Jacobs, M.J. (2003). Influence of peripheral arterial disease on capillary pressure in the foot. *J. Vasc. Surg.* *38*, 1067–1074.
- Fagrell, B., and Intaglietta, M. (1997). Microcirculation: its significance in clinical and molecular medicine. *En. J. Intern. Med.* *241*, 349–362.
- Houtman, P.M., Kallenberg, C.G., Wouda, A.A., and The, T.H. (1985). Decreased nailfold capillary density in raynaud's phenomenon: a reflection of immunologically mediated local and systemic vascular disease? *Ann. Rheum. Dis.* *44*, 603–609.
- Schmeling, H., Stephens, S., Goia, C., Manlihot, C., Schneider, R., Luthra, S., Stringer, E., and Feldman, B.M. (2011). Nailfold capillary density is importantly associated over time with muscle and skin disease activity in juvenile dermatomyositis. *Rheumatology* *50*, 885–893.
- Duscha, B.D., Kraus, W.E., Keteyian, S.J., Sullivan, M.J., Green, H.J., Schachat, F.H., Phippen, A.M., Brawner, C.A., Blank, J.M., and Annex, B.H. (1999). Capillary density of skeletal muscle: a contributing mechanism for exercise intolerance in class II-III chronic heart failure independent of other peripheral alterations. *J. Am. Coll. Cardiol.* *33*, 1956–1963.
- Robbins, J.L., Jones, W.S., Duscha, B.D., Allen, J.D., Kraus, W.E., Reegensteiner, J.G., Hiatt, W.R., and Annex, B.H. (2011). Relationship between leg muscle capillary density and peak hyperemic blood flow with endurance capacity in peripheral artery disease. *J. Appl. Physiol.* *111*, 81–86.
- Moeini, M., Lu, X., Avti, P.K., Damseh, R., Bélanger, S., Picard, F., Boas, D., Kakkar, A., and Lesage, F. (2018). Compromised microvascular oxygen delivery increases brain tissue vulnerability with age. *Sci. Rep.* *8*, 8219.
- López, A., Grignola, J.C., Angulo, M., Alvez, I., Nin, N., Lacuesta, G., Baz, M., Cardinal, P., Prestes, I., Bouchacourt, J.P., et al. (2015). Effects of early hemodynamic resuscitation on left ventricular performance and microcirculatory function during endotoxic shock. *Intensive Care Med* *Exp* *3*, 49.
- De Backer, D., Creteur, J., Preiser, J.-C., Dubois, M.-J., and Vincent, J.-L. (2002). Microvascular blood flow is altered in patients with sepsis. *Am. J. Respir. Crit. Care Med.* *166*, 98–104.
- Wester, T., Awan, Z.A., Kvernebo, T.S., Salerud, G., and Kvernebo, K. (2014). Skin microvascular morphology and hemodynamics during treatment with veno-arterial extra-corporeal membrane oxygenation. *Clin. Hemorheol. Microcirc.* *56*, 119–131.
- Ellis, C.G., Jagger, J., and Sharpe, M. (2005). The microcirculation as a functional system. *Crit. Care* *9*, S3–S8.
- Pittman, R.N. (2005). Oxygen transport and exchange in the microcirculation. *Microcirculation* *12*, 59–70.
- den Uil, C.A., Klijjn, E., Lagrand, W.K., Brugs, J.J., Ince, C., Spronk, P.E., and Simoons, M.L. (2008). The microcirculation in health and critical disease. *Prog. Cardiovasc. Dis.* *51*, 161–170.
- Popel, A.S., and Johnson, P.C. (2005). Microcirculation and hemorheology. *Annu. Rev. Fluid Mech.* *37*, 43–69.
- De Backer, D., Hollenberg, S., Boerma, C., Goedhart, P., Büchele, G., Ospina-Tascon, G., Dobbe, I., and Ince, C. (2007). How to evaluate the microcirculation: report of a round table conference. *Crit. Care* *11*, R101.
- Cassoobhoy, A. (2020). What is the definition of capillaries?. visited on 10/27/2020. <https://www.webmd.com/heart-disease/heart-failure/qa/what-is-the-definition-of-capillaries>.
- Shore, A.C. (2000). Capillaroscopy and the measurement of capillary pressure. *Br. J. Clin. Pharmacol.* *50*, 501–513.
- Maricq, H.R., Spencer-Green, G., and LeRoy, E.C. (1976). Skin capillary abnormalities as indicators of organ involvement in scleroderma (systemic sclerosis), raynaud's syndrome and dermatomyositis. *Am. J. Med.* *61*, 862–870.
- Wester, T., Awan, Z.A., Kvernebo, T.S., Salerud, G., and Kvernebo, K. (2014). Skin microvascular morphology and hemodynamics during treatment with venoarterial extra-corporeal membrane oxygenation. *Clin. Hemorheol. Microcirc.* *56*, 119–131.
- De Backer, D., Creteur, J., Preiser, J.-C., Dubois, M.-J., and Vincent, J.-L. (2002). Microvascular blood flow is altered in patients with sepsis. *Am. J. Respir. Crit. Care Med.* *166*, 98–104.
- Top, A.P.C., Ince, C., de Meij, N., van Dijk, M., and Tibboel, D. (2011). Persistent low microcirculatory vessel density in nonsurvivors of sepsis in pediatric intensive care. *Crit. Care Med.* *39*, 8–13.
- Natalello, G., De Luca, G., Gigante, L., Campochiaro, C., De Lorenzis, E., Verardi, L., Paglionico, A., Petricca, L., Martone, A.M., Calvisi, S., et al. (2021). Nailfold capillaroscopy findings in patients with coronavirus disease 2019: broadening the spectrum of covid-19 microvascular involvement. *Microvasc. Res.* *133*, 104071.
- Kanoore Edul, V.S., Caminos Eguillor, J.F., Ferrara, G., Estenssoro, E., Siles, D.S.P., Cesio, C.E., and Dubin, A. (2021). Microcirculation alterations in severe covid-19 pneumonia. *J. Crit. Care* *61*, 73–75.
- Ince, C., Boerma, E.C., Cecconi, M., De Backer, D., Shapiro, N.I., Duran-teau, J., Pinsky, M.R., Artigas, A., Teboul, J.-L., Reiss, I.K.M., et al. (2018). Second consensus on the assessment of sublingual microcirculation in critically ill patients: results from a task force of the european society of intensive care medicine. *Intensive Care Med.* *44*, 281–299.
- Hilty, M.P., Guerci, P., Ince, Y., Toraman, F., and Ince, C. (2019). Microtools enables automated quantification of capillary density and red blood cell velocity in handheld vital microscopy. *Commun. Biol.* *2*, 1–15.
- Hilty, M.P., Akin, S., Boerma, C., Donati, A., Erdem, Ö., Giaccaglia, P., Guerci, P., Milstein, D.M., Montomoli, J., Toraman, F., et al. (2020). Automated algorithm analysis of sublingual microcirculation in an international multicenter database identifies alterations associated with disease and mechanism of resuscitation. *Crit. Care Med.* *48*, e864–e875.
- Helmy Abdou, M.A., Truong, T.T., Dykky, A., Ferreira, P., and Jul, E. (2022). Capillarynet: an automated system to quantify skin capillary

- density and red blood cell velocity from handheld vital microscopy. *Artif. Intell. Med.* *127*, 102287.
37. Martini, R. (2020). The compelling arguments for the need of microvascular investigation in covid-19 critical patients. *Clin. Hemorheol. Microcirc.* *75*, 27–34.
  38. Ocak, I., Kara, A., and Ince, C. (2016). Monitoring microcirculation. *Best Pract. Res. Clin. Anaesthesiol.* *30*, 407–418.
  39. Ashruf, J.F., Bruining, H.A., and Ince, C. (2013). New insights into the pathophysiology of cardiogenic shock: the role of the microcirculation. *Curr. Opin. Crit. Care* *19*, 381–386.
  40. Cuthbertson, C.M., and Christophi, C. (2006). Disturbances of the microcirculation in acute pancreatitis. *Br. J. Surg.* *93*, 518–530.
  41. O'Mahony, N., Campbell, S., Carvalho, A., Harapanahalli, S., Hernandez, G.V., Krpalkova, L., et al. (2019). Deep learning vs. traditional computer vision. In *Science and Information Conference*, K. Arai, ed. (Springer), pp. 128–144.
  42. Vouliodimos, A., Doulamis, N., Doulamis, A., and Protopapadakis, E. (2018). Deep Learning for Computer Vision: A Brief Review. *Comput. Intell. Neurosci.* *2018*, 7068349.
  43. Cutolo, M., Pizzorni, C., Secchi, M.E., and Sulli, A. (2008). Capillaroscopy. *Best Pract. Res. Clin. Rheumatol.* *22*, 1093–1108.
  44. Ruaro, B., Smith, V., Sulli, A., Decuman, S., Pizzorni, C., and Cutolo, M. (2015). Methods for the morphological and functional evaluation of microvascular damage in systemic sclerosis. *Korean J. Intern. Med.* *30*, 1–5.
  45. Dilken, O., Ergin, B., and Ince, C. (2020). Assessment of sublingual microcirculation in critically ill patients: consensus and debate. *Ann. Transl. Med.* *8*, 793.
  46. Kvernebo, A.K., Miyamoto, T., Sporastoyl, A.H., Wikslund, L.K., Måsoy, S.E., Drolsum, L., Moe, M.C., Salerud, G., Fukamachi, K., and Kvernebo, K. (2020). Quantification of ocular surface microcirculation by computer assisted video microscopy and diffuse reflectance spectroscopy. *Exp. Eye Res.* *201*, 108312.
  47. Kraemer, R., Kabbani, M., Sorg, H., Herold, C., Branski, L., Vogt, P.M., and Knobloch, K. (2012). Diabetes and peripheral arterial occlusive disease impair the cutaneous tissue oxygenation in dorsal hand microcirculation of elderly adults: implications for hand rejuvenation. *Dermatol. Surg.* *38*, 1136–1142.
  48. Hasegawa, K., Pereira, B.P., and Pho, R.W. (2001). The microvasculature of the nail bed, nail matrix, and nail fold of a normal human fingertip. *J. Hand Surg. Am.* *26*, 283–290.
  49. S. User. [Online]. <https://www.dino-lite.eu/index.php/en/products/medical/capillarscope>.
  50. Video capillaroscopy. [Online]. <https://www.optiliamedical.eu/products/2/capillarscope/11/Optilia%20Digital%20Capillaroscopy%20System,%2020Extensive%20kit/>.
  51. (2020). Digital capillaroscopy. [Online]. <https://www.inspect-is.com/capillaroscopypro/>.
  52. (2021). Smart g-scope™ europe - smart g-scope™ europe. [Online]. <https://g-scope.eu/>.
  53. Anders, H.J., Sigl, T., and Schattenkirchner, M. (2001). Differentiation between primary and secondary raynaud's phenomenon: a prospective study comparing nailfold capillaroscopy using an ophthalmoscope or stereomicroscope. *Ann. Rheum. Dis.* *60*, 407–409.
  54. Russakovsky, O., Deng, J., Su, H., Krause, J., Satheesh, S., Ma, S., Huang, Z., Karpathy, A., Khosla, A., Bernstein, M., et al. (2015). Imagenet large scale visual recognition challenge. *Int. J. Comput. Vis.* *115*, 211–252.
  55. Zhou, B., Lapedriza, A., Khosla, A., Oliva, A., and Torralba, A. (2018). Places: a 10 million image database for scene recognition. *IEEE Trans. Pattern Anal. Mach. Intell.* *40*, 1452–1464.
  56. Kuznetsova, A., Rom, H., Alldrin, N., Uijlings, J., Krasin, I., Pont-Tuset, J., Kamali, S., Popov, S., Mallocci, M., Kolesnikov, A., et al. (2020). The open images dataset v4. *Int. J. Comput. Vis.* *128*, 1956–1981.
  57. Jordan, M.I., and Mitchell, T.M. (2015). Machine learning: trends, perspectives, and prospects. *Science* *349*, 255–260.
  58. Zhang, X.-D. (2020). Machine learning. In *A Matrix Algebra Approach to Artificial Intelligence*, C.L. Chang, ed. (Springer), pp. 223–440.
  59. Burkov, A. (2019). The hundred-page machine learning book. Andriy Burkov Canada 7.
  60. Mayo, M. (2018). Frameworks for Approaching the Machine Learning Process (KDNuggets).
  61. Willemink, M.J., Koszek, W.A., Hardell, C., Wu, J., Fleischmann, D., Harvey, H., Folio, L.R., Summers, R.M., Rubin, D.L., and Lungren, M.P. (2020). Preparing medical imaging data for machine learning. *Radiology* *295*, 4–15.
  62. Ward, L., Agrawal, A., Choudhary, A., and Wolverton, C. (2016). A general-purpose machine learning framework for predicting properties of inorganic materials. *npj Comput. Mater.* *2*, 1–7.
  63. Kohli, M., Prevedello, L.M., Filice, R.W., and Geis, J.R. (2017). Implementing machine learning in radiology practice and research. *AJR Am. J. Roentgenol.* *208*, 754–760.
  64. Caruana, R., and Niculescu-Mizil, A. (2006). An empirical comparison of supervised learning algorithms. In *Proceedings of the 23rd International Conference on Machine Learning (MIT)*, pp. 161–168.
  65. Dayan, P., Sahani, M., and Deback, G. (1999). Unsupervised learning. *The MIT encyclopedia of the cognitive sciences*, 857–859.
  66. Friedman, J., Hastie, T., and Tibshirani, R. (2001). *The Elements of Statistical Learning, 1* (Springer Series in Statistics New York).
  67. Bousquet, O., von Luxburg, U., and Rätsch, G. (2011). *Advanced Lectures on Machine Learning: ML Summer Schools 2003, Canberra, Australia, February 2-14, 2003, Tübingen, Germany, August 4-16, 2003, Revised Lectures, 3176* (Springer).
  68. Sutton, R.S., and Barto, A.G. (2018). *Reinforcement Learning: An Introduction* (MIT Press).
  69. Szepesvári, C. (2010). Algorithms for reinforcement learning. *Synth. Lect. Artif. Intell. Mach. Learn.* *4*, 1–103.
  70. Wiering, M.A., and Van Otterlo, M. (2012). Reinforcement learning. *Adapt. Learn. Optim.* *12*.
  71. Silver, D., Hubert, T., Schrittwieser, J., Antonoglou, I., Lai, M., Guez, A., Lanctot, M., Sifre, L., Kumaran, D., Graepel, T., et al. (2018). A general reinforcement learning algorithm that masters chess, shogi, and go through self-play. *Science* *362*, 1140–1144.
  72. Pan, S.J., and Yang, Q. (2010). A survey on transfer learning. *IEEE Trans. Knowl. Data Eng.* *22*, 1345–1359.
  73. Weiss, K., Khoshgoftaar, T.M., and Wang, D. (2016). A survey of transfer learning. *J. Big Data* *3*, 9–40.
  74. Torrey, L., and Shavlik, J. (2009). Transfer learning. *handbook of research on machine learning applications*. IGI Global *3*, 17–35.
  75. LeCun, Y., Bengio, Y., and Hinton, G. (2015). Deep learning. *Nature* *521*, 436–444.
  76. Goodfellow, I., Bengio, Y., and Courville, A. (2016). *Deep Learning* (MIT press).
  77. Shen, D., Wu, G., and Suk, H.-I. (2017). Deep learning in medical image analysis. *Annu. Rev. Biomed. Eng.* *19*, 221–248.
  78. Deng, L., and Yu, D. (2014). Deep learning: methods and applications. *FNT. in Signal Processing* *7*, 197–387.
  79. Medsker, L.R., and Jain, L. (2001). Recurrent neural networks. *Design and Applications* *5*, 64–67.



80. Creswell, A., White, T., Dumoulin, V., Arulkumaran, K., Sengupta, B., and Bharath, A.A. (2018). Generative adversarial networks: an overview. *IEEE Signal Process. Mag.* 35, 53–65.
81. O’Shea, K., and Nash, R. (2015). An introduction to convolutional neural networks. Preprint at arXiv. <https://doi.org/10.48550/arXiv.1511.08458>.
82. Mandic, D., and Chambers, J. (2001). *Recurrent Neural Networks for Prediction: Learning Algorithms, Architectures and Stability* (Wiley).
83. Cheng, J., Dong, L., and Lapata, M. (2016). Long shortterm memory-networks for machine reading. Preprint at arXiv. <https://doi.org/10.48550/arXiv.1601.06733>.
84. Dey, R., and Salem, F.M. (2017). Gate-variants of gated recurrent unit (gru) neural networks. In 2017 IEEE 60th international midwest symposium on circuits and systems (MWSCAS) (IEEE), pp. 1597–1600.
85. Chung, J., Gulcehre, C., Cho, K., and Bengio, Y. (2014). Empirical evaluation of gated recurrent neural networks on sequence modeling. Preprint at arXiv. <https://doi.org/10.48550/arXiv.1412.3555>.
86. Pan, Z., Yu, W., Yi, X., Khan, A., Yuan, F., and Zheng, Y. (2019). Recent progress on generative adversarial networks (gans): a survey. *IEEE Access* 7, 36322–36333.
87. Saxena, D., and Cao, J. (2022). Generative adversarial networks (gans) challenges, solutions, and future directions. *ACM Comput. Surv.* 54, 1–42.
88. Anwar, S.M., Majid, M., Qayyum, A., Awais, M., Alnowami, M., and Khan, M.K. (2018). Medical image analysis using convolutional neural networks: a review. *J. Med. Syst.* 42, 1–13.
89. Litjens, G., Kooi, T., Bejnordi, B.E., Setio, A.A.A., Ciampi, F., Ghafoorian, M., Van Der Laak, J.A., Van Ginneken, B., and Sánchez, C.I. (2017). A survey on deep learning in medical image analysis. *Med. Image Anal.* 42, 60–88.
90. Yamashita, R., Nishio, M., Do, R.K.G., and Togashi, K. (2018). Convolutional neural networks: an overview and application in radiology. *Insights Imaging* 9, 611–629.
91. LeCun, Y., Bottou, L., Bengio, Y., and Haffner, P. (1998). Gradient-based learning applied to document recognition. *Proc. IEEE* 86, 2278–2324.
92. El-Amir, H., and Hamdy, M. (2019). *Deep Learning Pipeline: Building a Deep Learning Model with TensorFlow* (Apress).
93. Zhao, Z.-Q., Zheng, P., Xu, S.-t., and Wu, X. (2019). Object detection with deep learning: a review. *IEEE Trans. Neural Netw. Learn. Syst.* 30, 3212–3232.
94. Karpathy, A. (2018). Stanford university cs231n: convolutional neural networks for visual recognition. <http://cs231n.stanford.edu/syllabus.html>.
95. Yegnanarayana, B. (2009). *Artificial Neural Networks* (PHI Learning Pvt. Ltd.).
96. Yiqiao, Y. (2018). *Deep Learning Notes* (Columbia University).
97. Kim, P. (2017). Convolutional neural network. In *MATLAB deep learning* (Springer), pp. 121–147.
98. Krizhevsky, A., Sutskever, I., and Hinton, G.E. (2012). Imagenet classification with deep convolutional neural networks. *Adv. Neural Inf. Process. Syst.* 25, 1097–1105.
99. Simonyan, K., and Zisserman, A. (2014). Very deep convolutional networks for large-scale image recognition. Preprint at arXiv. <https://doi.org/10.48550/arXiv.1409.1556>.
100. Szegedy, C., Liu, W., Jia, Y., Sermanet, P., Reed, S., Anguelov, D., Erhan, D., Vanhoucke, V., and Rabinovich, A. (2015). Going deeper with convolutions. In *Proceedings of the IEEE conference on computer vision and pattern recognition* (IEEE), pp. 1–9.
101. He, K., Zhang, X., Ren, S., and Sun, J. (2016). Deep residual learning for image recognition. In *Proceedings of the IEEE conference on computer vision and pattern recognition* (IEEE), pp. 770–778.
102. Amit, Y., Felzenszwalb, P., and Girshick, R. (2020). *Object detection. Computer Vision: A Reference Guide*, 1–9.
103. Liu, L., Ouyang, W., Wang, X., Fieguth, P., Chen, J., Liu, X., and Pietikäinen, M. (2020). Deep learning for generic object detection: a survey. *Int. J. Comput. Vis.* 128, 261–318.
104. Lowe, D.G. (1999). Object recognition from local scaleinvariant features. In *Proceedings of the Seventh IEEE international Conference on Computer Vision*, 2 (IEEE), pp. 1150–1157.
105. Alom, M.Z., Taha, T.M., Yakopcic, C., Westberg, S., Sidike, P., Nasrin, M.S., Van Esesn, B.C., Awwal, A.A.S., and Asari, V.K. (2018). The history began from alexnet: a comprehensive survey on deep learning approaches. Preprint at arXiv. <https://doi.org/10.48550/arXiv.1803.01164>.
106. Everingham, M., Van Gool, L., Williams, C.K.I., Winn, J., and Zisserman, A. (2010). The pascal visual object classes (voc) challenge. *Int. J. Comput. Vis.* 88, 303–338.
107. Lin, T.-Y., Maire, M., Belongie, S., Hays, J., Perona, P., Ramanan, D., Doll’ar, P., and Zitnick, C.L. (2014). Microsoft coco: common objects in context. In *European Conference on Computer Vision* (Springer), pp. 740–755.
108. Hoiem, D., Chodpathumwan, Y., and Dai, Q. (2012). Diagnosing error in object detectors. In *European Conference on Computer Vision* (Springer), pp. 340–353.
109. Girshick, R., Donahue, J., Darrell, T., and Malik, J. (2014). Rich feature hierarchies for accurate object detection and semantic segmentation. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition* (IEEE), pp. 580–587.
110. Uijlings, J.R.R., van de Sande, K.E.A., Gevers, T., and Smeulders, A.W.M. (2013). Selective search for object recognition. *Int. J. Comput. Vis.* 104, 154–171.
111. Athiwaratkun, B., and Kang, K. (2015). *Feature Representation in Convolutional Neural Networks*. preprint. arXiv. <https://doi.org/10.48550/arXiv.1507.02313>.
112. He, K., Zhang, X., Ren, S., and Sun, J. (2015). Spatial pyramid pooling in deep convolutional networks for visual recognition. *IEEE Trans. Pattern Anal. Mach. Intell.* 37, 1904–1916.
113. Lazebnik, S., Schmid, C., and Ponce, J. (2006). Beyond bags of features: spatial pyramid matching for recognizing natural scene categories. In *2006 IEEE Computer Society Conference on Computer Vision and Pattern Recognition (CVPR’06) IEEE*, 2 (IEEE), pp. 2169–2178.
114. Girshick, R. (2015). Fast r-cnn. In *Proceedings of the IEEE International Conference on Computer Vision* (IEEE), pp. 1440–1448.
115. Xue, J., Li, J., and Gong, Y. (2013). Restructuring of deep neural network acoustic models with singular value decomposition. *Interspeech (ISCA)*, pp. 2365–2369.
116. Ren, S., He, K., Girshick, R., and Sun, J. (2015). Faster rcnn: towards real-time object detection with region proposal networks. *Adv. Neural Inf. Process. Syst.* 28, 91–99.
117. Zitnick, C.L., and Dollár, P. (2014). Edge boxes: locating object proposals from edges. In *European Conference on Computer Vision* (Springer), pp. 391–405.
118. Dai, J., Li, Y., He, K., and Sun, J. (2016). R-fcn: object detection via region-based fully convolutional networks. In *Advances in Neural Information Processing Systems*, D. Lee, M. Sugiyama, U. Luxburg, I. Guyon, and R. Garnett, eds. (MIT Press), pp. 379–387.
119. Lin, T.-Y., Doll’ar, P., Girshick, R., He, K., Hariharan, B., and Belongie, S. (2017). Feature pyramid networks for object detection. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition* (IEEE), pp. 2117–2125.
120. He, K., Gkioxari, G., Doll’ar, P., and Girshick, R. (2017). Mask r-cnn. In *Proceedings of the IEEE International Conference on Computer Vision* (IEEE), pp. 2961–2969.
121. Zhang, X., Yang, Y.-H., Han, Z., Wang, H., and Gao, C. (2013). Object class detection: a survey. *ACM Comput. Surv.* 46, 1–53.
122. Brahmabhatt, S., Christensen, H.I., and Hays, J. (2017). Stuffnet: using ‘stuff’ to improve object detection. In *2017 IEEE Winter Conference on Applications of Computer Vision (WACV)* (IEEE), pp. 934–943.

123. Bell, S., Zitnick, C.L., Bala, K., and Girshick, R. (2016). Inside-outside net: detecting objects in context with skip pooling and recurrent neural networks. In Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition (IEEE), pp. 2874–2883.
124. Kong, T., Yao, A., Chen, Y., and Sun, F. (2016). Hypernet: towards accurate region proposal generation and joint object detection. In Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition (IEEE), pp. 845–853.
125. Shrivastava, A., Gupta, A., and Girshick, R. (2016). Training region-based object detectors with online hard example mining. In Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition (IEEE), pp. 761–769.
126. Redmon, J., Divvala, S., Girshick, R., and Farhadi, A. (2016). You only look once: unified, real-time object detection. In Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition (IEEE), pp. 779–788.
127. Redmon, J., and Farhadi, A. (2017). Yolo9000: better, faster, stronger. In Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition (IEEE), pp. 7263–7271.
128. Redmon, J., and Farhadi, A. (2018). Yolov3: an incremental improvement. Preprint at arXiv. <https://doi.org/10.48550/arXiv.1804.02767>.
129. Hu, D. (2019). An introductory survey on attention mechanisms in nlp problems. In Proceedings of SAI Intelligent Systems Conference (Springer), pp. 432–448.
130. Lin, T.-Y., Goyal, P., Girshick, R., He, K., and Dollár, P. (2017). Focal loss for dense object detection. In Proceedings of the IEEE international conference on computer vision (IEEE), pp. 2980–2988.
131. Law, H., and Deng, J. (2018). Cornernet: detecting objects as paired keypoints. In Proceedings of the European conference on computer vision (ECCV) (The Computer Vision Foundation), pp. 734–750.
132. Lim, B., Son, S., Kim, H., Nah, S., and Mu Lee, K. (2017). Enhanced deep residual networks for single image superresolution. In Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition Workshops (IEEE), pp. 136–144.
133. Shi, W., Caballero, J., Huszar, F., Totz, J., Aitken, A.P., Bishop, R., Rueckert, D., and Wang, Z. (2016). Real-time single image and video super-resolution using an efficient sub-pixel convolutional neural network. In Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition (IEEE), pp. 1874–1883.
134. Dong, C., Loy, C.C., and Tang, X. (2016). Accelerating the super-resolution convolutional neural network. In European Conference on Computer Vision (Springer), pp. 391–407.
135. Lai, W.-S., Huang, J.-B., Ahuja, N., and Yang, M.-H. (2019). Fast and accurate image super-resolution with deep laplacian pyramid networks. *IEEE Trans. Pattern Anal. Mach. Intell.* *41*, 2599–2613.
136. Dufour, R.M., Miller, E.L., and Galatsanos, N.P. (2002). Template matching based object recognition with unknown geometric parameters. *IEEE Trans. Image Process.* *11*, 1385–1396.
137. Greig, D.W., and Denny, M. (2003). Knowledge-based methods for small-object detection in sar images. In SAR Image Analysis, Modeling, and Techniques V, International Society for Optics and Photonics, 4883, F. Posa, ed. (The International Society for Optics and Photonics), pp. 121–130.
138. Hossain, M.D., and Chen, D. (2019). Segmentation for objectbased image analysis (obia): a review of algorithms and challenges from remote sensing perspective. *ISPRS J. Photogrammetry Remote Sens.* *150*, 115–134.
139. Jodoin, P.M., Jodoin, P.-M., Emile, B., Laurent, H., and Rosenberger, C. (2010). Comparative study of background subtraction algorithms. *J. Electron. Imaging* *19*, 033003.
140. KaewTraKulPong, P., and Bowden, R. (2002). An improved adaptive background mixture model for real-time tracking with shadow detection. In Video-Based Surveillance Systems, P. Remagnino, G.A. Jones, N. Paragios, and C.S. Regazzoni, eds. (Springer), pp. 135–144.
141. Zivkovic, Z. (2004). Improved adaptive Gaussian mixture model for background subtraction. In Proceedings of the 17th International Conference on Pattern Recognition, 2004. ICPR 2004, 2 (IEEE), pp. 28–31.
142. Zivkovic, Z., and Van Der Heijden, F. (2006). Efficient adaptive density estimation per image pixel for the task of background subtraction. *Pattern Recogn. Lett.* *27*, 773–780.
143. Li, L., Huang, W., Gu, I.Y.-H., and Tian, Q. (2004). Statistical modeling of complex backgrounds for foreground object detection. *IEEE Trans. Image Process.* *13*, 1459–1472.
144. Marcomini, L., and Cunha, A.L. (2018). A comparison between background modelling methods for vehicle segmentation in highway traffic videos. Preprint at arXiv. <https://doi.org/10.48550/arXiv.1810.02835>.
145. Sankur, B., and Sankur, B. (2004). Survey over image thresholding techniques and quantitative performance evaluation. *J. Electron. Imaging* *13*, 146–165.
146. Bradley, D., and Roth, G. (2007). Adaptive thresholding using the integral image. *J. Graph. Tool.* *12*, 13–21.
147. Yousefi, J. (2011). Image Binarization Using Otsu Thresholding Algorithm (University of Guelph).
148. Bradski, G., and Kaehler, A. (2008). Learning OpenCV: Computer Vision with the OpenCV Library (O'Reilly Media, Inc.).
149. McIlhagga, W. (2011). The canny edge detector revisited. *Int. J. Comput. Vis.* *91*, 251–261.
150. Sherstyuk, A. (1999). Kernel functions in convolution surfaces: a comparative analysis. *Vis. Comput.* *15*, 171–182.
151. Kaur, M., Kaur, J., and Kaur, J. (2011). Survey of contrast enhancement techniques based on histogram equalization. *Int. J. Adv. Comput. Sci. Appl.* *2*.
152. Motwani, M.C., Gadiya, M.C., Motwani, R.C., and Harris, F.C. (2004). Survey of image denoising techniques. *Proceedings of GSPX* *27*, 27–30.
153. Dobbe, J.G.G., Streekstra, G.J., Atasever, B., Van Zijderveld, R., and Ince, C. (2008). Measurement of functional microcirculatory geometry and velocity distributions using automated image analysis. *Med. Biol. Eng. Comput.* *46*, 659–670.
154. Deng, H., Zhang, W., Mortensen, E., Dieterich, T., and Shapiro, L. (2007). Principal curvature-based region detector for object recognition. In 2007 IEEE Conference on Computer Vision and Pattern Recognition (IEEE), pp. 1–8.
155. Bezemer, R., Dobbe, J.G., Bartels, S.A., Boerma, E.C., Christiaan Boerma, E., Elbers, P.W.G., Heger, M., and Ince, C. (2011). Rapid automatic assessment of microvascular density in sidestream dark field images. *Med. Biol. Eng. Comput.* *49*, 1269–1278.
156. Tam, J., Martin, J.A., and Roorda, A. (2010). Noninvasive visualization and analysis of parafoveal capillaries in humans. *Invest. Ophthalmol. Vis. Sci.* *51*, 1691–1698.
157. Geyman, L.S., Garg, R.A., Suwan, Y., Trivedi, V., Krawitz, B.D., Mo, S., Pinhas, A., Tantraworasin, A., Chui, T.Y.P., Ritch, R., and Rosen, R.B. (2017). Peripapillary perfused capillary density in primary open-angle glaucoma across disease stage: an optical coherence tomography angiography study. *Br. J. Ophthalmol.* *101*, 1261–1268.
158. Demir, S.U., Hakimzadeh, R., Hargraves, R.H., Ward, K.R., Myer, E.V., and Najarian, K. (2012). An automated method for analysis of microcirculation videos for accurate assessment of tissue perfusion. *BMC Med. Imaging* *12*, 37–13.
159. Cheng, C., Lee, C.W., and Daskalakis, C. (2015). A reproducible computerized method for quantitation of capillary density using nailfold capillaroscopy. *J. Vis. Exp.* *104*, e53088.
160. Dai, G., He, W., Xu, L., Pazo, E.E., Lin, T., Liu, S., and Zhang, C. (2020). Exploring the effect of hypertension on retinal microvasculature using deep learning on east asian population. *PLoS One* *15*, e0230111.
161. Hariyani, Y.S., Eom, H., and Park, C. (2020). Da-capnet: dual attention deep learning based on u-net for nailfold capillary segmentation. *IEEE Access* *8*, 10543–10553.

162. Hu, J., Shen, L., and Sun, G. (2018). Squeeze-and-excitation networks. In Proceedings of the IEEE conference on computer vision and pattern recognition (IEEE), pp. 7132–7141.
163. Woo, S., Park, J., Lee, J.-Y., and Kweon, I.S. (2018). Cbam: convolutional block attention module. In Proceedings of the European conference on computer vision (ECCV) (The Computer Vision Foundation), pp. 3–19.
164. Prentašić, P., Heisler, M., Mammo, Z., Lee, S., Merkur, A., Navajas, E., Beg, M.F., Šarunic, M., and Lončarić, S. (2016). Segmentation of the foveal microvasculature using deep learning networks. *J. Biomed. Opt.* *21*, 075008.
165. Nivedha, R., Brinda, M., Suma, K., and Rao, B. (2016). Classification of nailfold capillary images in patients with hypertension using non-linear svm. In 2016 International Conference on Circuits, Controls, Communications and Computing (I4C) (IEEE), pp. 1–5.
166. Javia, P., Rana, A., Shapiro, N., and Shah, P. (2018). Machine learning algorithms for classification of microcirculation images from septic and non-septic patients. In 2018 17th IEEE International Conference on Machine Learning and Applications (ICMLA) (IEEE), pp. 607–611.
167. Tama, A., Mengko, T.R., and Zakaria, H. (2015). Nailfold capillaroscopy image processing for morphological parameters measurement. In 2015 4th International Conference on Instrumentation, Communications, Information Technology, and Biomedical Engineering (ICICI-BME) (IEEE), pp. 175–179.
168. Ye, F., Yin, S., Li, M., Li, Y., and Zhong, J. (2020). In vivo full-field measurement of microcirculatory blood flow velocity based on intelligent object identification. *J. Biomed. Opt.* *25*, 016003.
169. Abdou, M.A.H., Ferreira, P., Jul, E., and Truong, T.T. (2022). Capillaryx: a software design pattern for analyzing medical images in real-time using deep learning. Preprint at arXiv. <https://doi.org/10.48550/arXiv.2204.08462>.

#### About the author

**Maged Helmy** is a software engineer with 7+ years of experience in designing, developing, and shipping machine-learning enriched software. He is cloud certified by Microsoft and machine learning certified by Google. He completed his master's degree in systems engineering with a focus on machine learning and is a PhD candidate in software engineering with a focus on deep learning at the University of Oslo.