
Fast Contactless Detection of Post-COVID-19 from Eye Microvascular Images Using Deep Convolutional Neural Networks

[Aristotle G. Koutsiaris](#)*, [Konstantina Riri](#), [Stylianos Boutlas](#), Thomas N. Panagiotou, Maria Kotoula, Zoe Daniil, [Aristeidis H. Zibis](#), [Evangelia E. Tsironi](#)

Posted Date: 14 April 2026

doi: 10.20944/preprints202604.0923.v1

Keywords: COVID-19; virus; fast detection; contactless; eye; microvascular images; video capillaroscopy; AI; neural networks; deep CNN



Preprints.org is a free multidisciplinary platform providing preprint service that is dedicated to making early versions of research outputs permanently available and citable. Preprints posted at Preprints.org appear in Web of Science, Crossref, Google Scholar, Scilit, Europe PMC.

Copyright: This open access article is published under a [Creative Commons CC BY 4.0 license](#), which permit the free download, distribution, and reuse, provided that the author and preprint are cited in any reuse.

Disclaimer/Publisher's Note: The statements, opinions, and data contained in all publications are solely those of the individual author(s) and contributor(s) and not of MDPI and/or the editor(s). MDPI and/or the editor(s) disclaim responsibility for any injury to people or property resulting from any ideas, methods, instructions, or products referred to in the content.

Article

Fast Contactless Detection of Post-COVID-19 from Eye Microvascular Images Using Deep Convolutional Neural Networks

Aristotle G. Koutsiaris ^{1,*}, Konstantina Riri ², Stylianos Boutlas ³, Thomas N. Panagiotou ², Maria Kotoula ², Zoe Daniil ³, Aristidis H. Zibis ⁴ and Evangelia E. Tsironi ²

¹ Medical Informatics and Biomedical Imaging (MIBI) Laboratory, Faculty of Medicine, School of Health Sciences, University of Thessaly, Biopolis Campus, Larissa, Greece

² Department of Ophthalmology, University Hospital of Larissa, Faculty of Medicine, School of Health Sciences, University of Thessaly, Larissa, Greece

³ Department of Respiratory Medicine, University Hospital of Larissa, Faculty of Medicine, School of Health Sciences, University of Thessaly, Larissa, Greece

⁴ Department of Anatomy, Faculty of Medicine, School of Health Sciences, University of Thessaly, Biopolis Campus, Larissa, Greece

* Correspondence: ariskout@otenet.gr

Abstract

Background/Objectives: Artificial intelligence (AI) and its subfield deep learning (DL) have rapidly evolved into a central tool in modern medicine. The purpose of this work was to examine if DL neural networks can discriminate efficiently the microvessel network of post-COVID-19 patients from healthy individuals from. **Methods:** A non-contact, digital slit-lamp video capillaroscopy system was used to record high magnification images from the bulbar conjunctival microcirculation of 12 COVID-19 survivors (named "COVID-19 Group") and 12 healthy volunteers (named "Control Group"). Four pretrained convolutional neural networks (CNNs) were fine-tuned by transfer learning and their performance was assessed by standard binary classification evaluation criteria. **Results:** A scene-centric CNN named GoogLeNet-Places365 excelled on all evaluation criteria with an average testing accuracy, sensitivity, specificity and AUC (area under the curve) of 92%, 92%, 91%, and 0.971, respectively. **Conclusions:** Post-COVID effects on the eye microcirculation can be detected by deep CNNs, and there is now evidence for the first time, that AI could provide a risk-free, painless, contactless, fast, and accurate detection method of viral effects that does not depend on the optical clarity of the eye.

Keywords: COVID-19; virus; fast detection; contactless; eye; microvascular images; video capillaroscopy; AI; neural networks; deep CNN

1. Introduction

Artificial intelligence (AI) has rapidly evolved into a central tool in modern medicine, transforming imaging, diagnosis, pathology, surgery, treatment planning, patient monitoring, treatment personalization, and healthcare efficiency. From a broader healthcare perspective, AI systems are expected to demonstrate improved resource allocation reducing costs [1].

Machine learning (ML) can be considered a subfield of artificial intelligence (AI) and deep learning (DL) a subfield of ML that uses artificial neural networks with multiple layers. In diagnostic medical imaging, DL with its exceptional ability to process complex visual data, reduces diagnostic errors [2]. Meta-analysis has demonstrated that AI can match clinicians in medical imaging accuracy in the fields of ophthalmology, respiratory imaging and breast imaging [3]. Convolutional Neural

Networks (CNNs), a subtype of artificial neural networks, can classify skin cancers with a level of accuracy comparable to board-certified dermatologists [4].

The seminal paper that introduced the modern concept of a CNN and its application to image recognition was published in 1998 [5]. The next big step was AlexNet, a significantly deeper and larger network than any previous CNN, which won the 2012 ImageNet Large Scale Visual Recognition Challenge (ILSVRC-2012) by an impressive margin, improving the error rate from 26.2% (the previous best) down to 15.3% [6]. Two years later, researchers at Google presented GoogLeNet (later renamed Inception-v1), a CNN with a core innovation the “inception module”, which executed multiple different convolutions and pooling operations in parallel [7]. GoogLeNet submission to ILSVRC-2014 used 12 times less parameters than AlexNet while being significantly more accurate [7]. One year later, at ILSVRC-2015, He et al.[8] introduced residual networks that are easier to optimize and gain accuracy from increased depth. Afterward, the combination of the inception architecture with the residual connections was studied and the hybrid versions “Inception-ResNet-v1” and “Inception-ResNet-v2” were presented [9]. In 2019, Tan & Le [10] proposed a new scaling method that balances network depth, width, and resolution and introduced a new family of models called “EfficientNets” which were significantly more efficient than previous CNNs.

For a human level performance in visual recognition tasks, apart from the CNN architecture, access to a sufficiently large training dataset, is of crucial importance. The rise of multi-million-item sets has improved dramatically the AI performance in classification of both objects (object recognition) and scenes (scene recognition). The scene or place in which an object appears may be equally or more important than the object itself, depending on the application. Scenes such as a field, park, runway, and a lobby are composed of parts and like objects have specific attributes [11].

In 2009, Deng et al. [12] introduced “ImageNet”, a large-scale database of images, consisting of 12 subtrees containing 3.2 million images, to be used as a training resource and a benchmark dataset for object recognition. In 2014, Zhou et al. [13] introduced the “Places” scene-centric database, a collection of over 7 million scene photographs. A few years later, they published [11] an improved version of “Places” with 10 million scene photographs, labeled with 434 scene semantic categories, comprising about 98 percent of scene types a human can encounter in the world.

The emergence of large datasets containing millions of items has allowed machine learning algorithms to achieve classification performance on visual object and scene recognition tasks that approaches human-level accuracy. The “transfer learning” process in deep learning, refines (fine-tunes) feature extraction of a pretrained network to be more specific to the current application, and has the advantage of increased flexibility with less effort. All CNNs used in this work were fine-tuned networks which had been pretrained on subsets of the aforementioned databases (ImageNet or Places).

The purpose of this work was to examine if fine-tuned CNNs can discriminate efficiently healthy individuals from post-COVID-19 patients, using conjunctival images acquired by high magnification conjunctival video capillaroscopy. The results demonstrated, for the first time, that a virus affects the microvascular topology in ways that only the sophisticated nature of a deep CNN can discriminate. On surplus, there is now evidence that AI could provide a painless, risk-free, contactless, fast, and accurate post-COVID-19 detection method that does not depend on the optical clarity of the eye.

2. Materials and Methods

2.1. Study Design

This machine learning, case-control study was based on images acquired at the University Hospital of Larissa from April 2021 to February 2022 with a protocol approved by the University Hospital of Larissa Scientific Committee (11155/2021). The research was performed according to the Declaration of Helsinki and informed consent was obtained from all participants in the study.

2.2. Subjects

At the temporal side of the bulbar conjunctiva of the right eye, microvascular images were recorded from 12 normal men volunteers (named "Control Group") and 12 COVID-19 men survivors shortly after exiting the hospital (named "COVID-19 Group"). Image recordings were performed in a temperature controlled environment (21 to 23 °C) after waiting 30 minutes for adaptation.

Inclusion criteria for the Control Group were no systemic or ocular disease, no alcohol or smoking habit. Control subjects were not under any medication and had never been tested positive for SARS-CoV-2.

Inclusion criteria for the COVID-19 Group were hospitalization for COVID-19, oxygen therapy, Acute Respiratory Distress Syndrome (ARDS) with HI (Horowitz Index) less than 300 mmHg and a 28-day maximum time interval between hospital exit and image recording (duration from discharge). Exclusion criterion for the COVID-19 Group was rheumatic disease.

Clinical data for both groups, including body height and body weight, diastolic and systolic arterial pressure (DP and SP), and cardiac beats per minute (BPM), were measured on site before and after image recording. Body surface area (BSA, Du Bois formula) and body mass index (BMI) were calculated from the body height and weight. Mean arterial pressure (MP) was calculated from DP and SP ($MP = DP + (SP - DP) / 3$). Subjects with diastolic blood pressure greater than 95 mmHg were excluded from the study (both groups) as hypertensive.

2.3. Experimental Set Up

The image acquisition experimental setup consisted of a slit lamp (Nikon FS-3V) connected with a high-speed CCD camera (12 bit, PCO Computer Optics GmbH, Germany) and a PC (Pentium 4, 3 GHz) [14]. The conjunctival video capillaroscopy system produced digital images of 320 × 240 pixels at a frame rate of 96 frames per second (fps) with an enhanced maximum magnification of 242x and a digital resolution of $1.257 \pm 0.004 \mu\text{m}/\text{pixel}$.

Image processing and deep learning in-silico AI experiments were performed on a desktop with an Intel CPU (i9-9900K) running at 3.6 (1.9 – 4.7) GHz, with 48 GB of RAM running at 2.13 GHz, and a 250 GB SSD. The image processing and deep learning AI software was MATLAB R2024a.

2.4. Deep Learning Training by Transfer Learning

A variety of pretrained convolutional neural networks (CNNs) were initially tested and 4 of them were selected for fine-tuning by transfer learning on our imaging data. Three CNNs, namely GoogLeNet, InceptionResNet-v2, and EfficientNet-b0, were pretrained on the ImageNet-1K dataset, a subset comprising roughly 1.2 million images from the full ImageNet database [12], enabling classification across 1,000 distinct object categories. One CNN, namely GoogLeNet-Places365 was pretrained on the Places365 dataset [11], which includes approximately 1.8 million images from the Places database, allowing it to classify images into 365 different scene categories. The output layers of the pretrained networks were modified to match the new binary classification task of 2 output classes (CONTROLS and COVID-19).

2.4.1. Architectural Hyperparameters and Size

Some of the architectural hyperparameters of the 4 pretrained deep neural networks (CNNs) used in this work, such as number of layers (NOL), depth (D), input image size (IIS), and input image channels (IIC) are shown in Table 1. GoogLeNet and GoogLeNet-Places365 are effectively the same network pretrained on different datasets, so, architectural hyperparameters are exactly the same.

Table 1. Architectural hyperparameters of different pretrained CNNs.

Architectural Hyperparameter	GoogLeNet	InceptionResNet-v2	EfficientNet-b0
NOL	144	824	290
D	22	164	82
IIS (pixels)	224 x 224	299 x 299	224 x 224
IIC	3	3	3

NOL = number of layers, D = depth, IIS = Input Image Size (pixels), IIC = Input Image Channels.

Architectural hyperparameters are quantifiable characteristics of the network's design, and correlate directly with how many learnable parameters (model parameters) the network will have. The amount of model parameters (size) of the pretrained deep neural networks (CNNs) used in this work, are shown in Table 2. It is clear that EfficientNet-b0 is lighter in comparison to the other networks, requiring only 20 MB of disk drive memory.

Table 2. Comparing the size of different pretrained CNNs.

Size	GoogLeNet	InceptionResNet-v2	EfficientNet-b0
Model Parameters (millions)	144	824	290
PM (MB)	22	164	82

PM = parameter memory in megabytes (MB).

2.4.2. Importing and Splitting Image Data

30 to 52 images were selected for each subject from image recordings of various microvessel fields of the same side (temporal side) of the bulbar conjunctiva. Fields were selected with as many microvessels as possible because the emphasis was in the microvessel topology. Images were examined to rule out repetitions of the same microvascular field but large translocations of the same field, if appeared, were included in order to incorporate as much information as possible.

Images were organized by class into two separated folders: one for the Control Group (CONTROLS) and one for the COVID-19 Group (COVID-19). There was no image pre-processing. A MATLAB image datastore was created for labelling images. From each group datastore, ninety percent (90%) was used for training and ten percent (10%) for testing, after randomly shuffling the datastore images. An augmented image datastore was created to convert the size and channels of the training and testing images to the input image size and channels of the corresponding pretrained CNN network (see Table 1).

2.4.3. Evaluation Criteria

Two evaluation criteria were used to assess the speed of each convolutional network. Training time (TT) was defined as the amount of time required to train the pretrained convolutional neural network and is measured in minutes. Decision time (DT) was defined as the amount of time required for the trained network to decide the category of a single input image, and is measured in seconds.

The following evaluation criteria were applied for assessing classification performance. Sensitivity (Se) or true positive rate or recall is defined in percent:

$$Se = \frac{TP}{TP + FN} 100 (\%) \quad (1)$$

where TP and FN is the number of true positive and false negative outcomes, respectively.

Precision (Pr) or positive predictive value is defined in percent:

$$Pr = \frac{TP}{TP + FP} 100 (\%) \quad (2)$$

where FP is the number of false positive outcomes.

F1 score or Dice similarity coefficient is defined as the harmonic mean of Se and Pr in percent:

$$F1 = \frac{2TP}{2TP + FP + FN} 100 (\%) \quad (3)$$

Specificity (Sp) or true negative rate is defined in percent:

$$Sp = \frac{TN}{TN + FP} 100 (\%) \quad (4)$$

where TN is the number of true negative outcomes.

Accuracy (Ac) is defined in percent:

$$Ac = \frac{TP + TN}{TP + TN + FP + FN} 100 (\%) \quad (5)$$

2.4.4. Fine-Tuning with Training Hyperparameters

Training hyperparameters used for fine-tuning in each in silico experiment were learning rate (LR), mini batch size (MBS), and maximum number of epochs (MEP). Initial LR, MBS, and MEP values were 0.01, 128, and 20, respectively. Initial LR and MBS were gradually reduced, and initial MEPs were gradually increased in random combinations. The performance of each combination for each pretrained network was assessed by the evaluation criteria presented in the previous Section 2.4.3. When a good combination of training hyperparameters was spotted, at least one experiment with validation data was performed to identify the onset of overfitting (optimum MEP). In these cases, the image data split was eighty-one percent (81%) for training, nine percent (9%) for validation, and ten percent (10%) for testing.

2.5. Statistical Analysis

Microsoft Excel 2016 (Professional edition) and SOFA (version 1.4, Paton-Simpson & Associates Ltd) were used for statistical analyses. To account for all variable distributions, normal or not, differences between groups were examined using the Mann-Whitney U test. The significance level was set at 0.05. After identifying the optimum maximum number of epochs (MEP), 10 successive in silico experiments were performed with the randomized image data split reported in Section 2.4.2 (90% for training and 10% for testing). In this manner, the average performance of each fine-tuned CNN was evaluated across 10 data splits using the criteria outlined in Section 2.4.3.

The overall discriminatory power of the fine-tuned CNN with the best average performance was evaluated by the area under the curve (AUC) of the receiver operating characteristic (ROC) curve. An AUC value of 1.0 represents perfect discriminatory power, while a value of 0.5 represents random guessing. For each data split, the ROC curve was generated by plotting the true positive rate (sensitivity) against the false positive rate (1 - specificity) across the full range of classification thresholds. This was repeated for all 10 train-test data splits to account for variability due to randomized data splits, and the mean AUC with a 95% confidence interval (CI) was calculated using the t-distribution.

For evaluating the performance at the Youden-optimal point, the Youden index (J) was estimated for every point on the ROC curve:

$$J = Se + Sp - 100 (\%) \quad (6)$$

For each split independently, the optimal classification threshold was determined by maximizing J, which is equivalent to finding the ROC point furthest from the diagonal line of chance. At this per-split Youden optimal threshold, the model was evaluated by the criteria of Section 2.4.3. Averages and standard deviations of these criteria were estimated across repeated experiments.

3. Results

3.1. Group Characteristics

A summary of the demographic and clinical statistical data for the 2 groups is presented in Table 3. There was no statistically significant difference in age, body mass index (BMI), body surface area

(BSA), cardiac beats per minute (BPM), arterial mean pressure (MP), and number of images per subject (NI) between the 2 groups. In each group, the changes in BPM and MP before and after the recording procedure were not significant.

Table 3. Demographic and clinical data of the Control Group and COVID-19 Group.

Variable	Control Group	COVID-19 Group	p value
N	12	12	-
Age (years)	50 ± 11	56 ± 7	0.07
BMI (kgr/m ²)	27 ± 2	29 ± 4	0.08
BSA (m ²)	2.0 ± 0.1	2.1 ± 0.1	0.10
BPM	69 ± 6	63 ± 10	0.17
MP (mmHg)	95 ± 8	93 ± 9	0.71
NI	41 ± 7	42 ± 6	0.71
TNI	486	504	-

N = number of subjects, BMI = Body Mass Index, BSA = Body Surface Area (Du Bois), BPM = cardiac Beats Per Minute, MP = arterial Mean Pressure, NI = Number of Images per subject, and TNI = Total Number of Images. Data are expressed as Mean ± Standard deviation. There was no significant difference in age, BMI, BSA, BPM, MP, and NI between groups.

A summary of the clinical and pharmacological profile of the COVID-19 Group during and after hospitalization is presented in Table 4. All patients (COVID-19 Group) had pneumonia with different lung distribution, and needed oxygen therapy with High Flow Nasal Cannula (HFNC). During hospitalization, the average Horowitz Index (HI) of the COVID-19 Group was 134 mmHg (Table 4), which implies moderate ARDS. During hospitalization, 11 patients (91%) received remdesivir and 6 patients (50%) received tocilizumab due to rapid clinical deterioration and high levels of C-reactive protein and ferritin (Table 4). In addition, 11 patients (91%) received LMWH (Low Molecular Weight Heparin) in antithrombotic prophylaxis dosage. Two of them received a full dose of enoxaparin since they developed atrial fibrillation and a high CHADS2 (Congestive heart failure, Hypertension, Age ≥75 years, Diabetes mellitus, Stroke [double weight]) score. One of them was also on antiplatelet therapy with ASA (Acetylsalicylic Acid) because of coronary artery disease (Table 4).

After hospital discharge, all patients had extensive radiological findings with deteriorated pulmonary function due to previous ARDS. Six patients (50%) needed oxygen therapy at home for at least one month (Table 2). All patients were on antithrombotic prophylaxis until fully mobilized (LMWH is recommended up to 40 days in severely and critically ill patients after discharge). Those 2 patients with atrial fibrillation who were receiving a therapeutic dose of LMWH continued their antithrombotic therapy after discharge with a full dose of rivaroxaban (Table 4).

Table 4. Clinical and pharmacological data of the COVID-19 Group (12 subjects) during and after hospitalization.

Variable	COVID-19 GROUP during hospitalization	COVID-19 GROUP after hospitalization
HI (mmHg)	134 ± 58	-
Oxygen therapy, n (%)	12 (100%)	6 (50%)
ICU, n (%)	2 (17%)	-
Antiviral therapy (remdesivir), n (%)	11 (91%)	-
Anti-IL6R (tocilizumab), n (%)	6 (50%)	-
AP, n (%)	11 * (91%)	12 * (100%)
ASA, n (%)	1 (9%)	-
Duration from discharge (days)	-	8 ± 7

Data are expressed as Mean ± Standard deviation or as number of subjects (n) and the corresponding percentage (%). Four subjects in the COVID-19 Group (33%) were diabetic and three subjects were smokers (25%). HI =

Horowitz Index, ICU = Intensive Care Unit, IL6R = InterLeukine-6 Receptor, AP = Antithrombotic Prophylaxis (* 2 of those patients developed atrial fibrillation and received full therapeutic dose of enoxaparin during hospitalization and rivaroxaban after hospitalization), ASA = AcetylSalicylic Acid. Duration from discharge: time interval between hospital exit and image recording.

3.2. Fine-Tuning and Performance Comparison

Different training combinations of the learning rate (LR), mini batch size (MBS), and maximum number of epochs (MEP) were tested for each of the 4 CNNs: GoogLeNet, GoogLeNet-Places365, InceptionResNet-v2, and EfficientNet-b0. To take into account the effect of the data split, many different in silico experiments were made with randomized train-test splits for each training hyperparameter combination. The training hyperparameter combinations LR, MBS, and MEP that gave the best average performance on testing data are shown in Table 5 for each of the 4 pretrained CNNs.

Table 5. Training hyperparameter combinations giving the best average performance.

Training Hyperparameter	GoogLeNet	GoogLeNet-Places365	InceptionResNet-v2	EfficientNet-b0
LR	$5 \cdot 10^{-4}$	$1 \cdot 10^{-3}$	$1 \cdot 10^{-2}$	$1 \cdot 10^{-3}$
MBS	128	32	128	32
MEP	75	22	30	15

LR = Learning Rate, MBS = Mini Batch Size, MEP = Maximum number of Epochs.

In Table 6, the best average performance of each fine-tuned pretrained CNN on testing data is presented, after using the combinations shown in Table 5. The best testing accuracy of 92% ($\pm 3\%$) was achieved by the GoogLeNet-Places365 network, followed by GoogLeNet and InceptionResNet-v2 with 88% ($\pm 5\%$). It is also noted, that GoogLeNet-Places365 was the network with the shortest training time (TT) (average value of 23.5 minutes) and the fastest decision time (DT) (average value of 0.56 seconds).

Table 6. Comparing performance among 4 different fine-tuned pretrained CNNs. .

Evaluation Criterion	GoogLeNet	GoogLeNet-Places365	InceptionResNet-v2	EfficientNet-b0
TT (min)	63.4 ± 0.7	23.5 ± 0.5	201.2 ± 0.9	38.2 ± 0.4
DT (s)	0.57 ± 0.01	0.56 ± 0.01	4.34 ± 0.19	1.35 ± 0.07
Se (%)	89 ± 6	92 ± 5	85 ± 7	90 ± 7
Pr (%)	88 ± 5	91 ± 3	89 ± 4	85 ± 3
F1 (%)	89 ± 5	92 ± 3	87 ± 5	88 ± 5
Sp (%)	88 ± 5	91 ± 4	90 ± 3	84 ± 2
Ac (%)	88 ± 5	92 ± 3	88 ± 5	87 ± 4

TT = Training Time (minutes), DT = Decision Time (seconds), Se = Sensitivity, Pr = Precision, F1 = F1 score, Sp = Specificity, Ac = Accuracy. Data are expressed as Mean \pm Standard deviation of ten (10) in silico experiments for each column.

3.3. Discrimination Performance

The network with the best testing accuracy in Table 6 (GoogLeNet-Places365) demonstrated stable discriminative power across 10 repeated train-test splits (in silico experiments). The mean AUC was 0.971, with a 95% confidence interval (CI) of [0.954 – 0.988], indicating robust separation between COVID-19 and non-COVID-19 cases across repeated experiments and limited sensitivity to data partitioning (Figure 1).

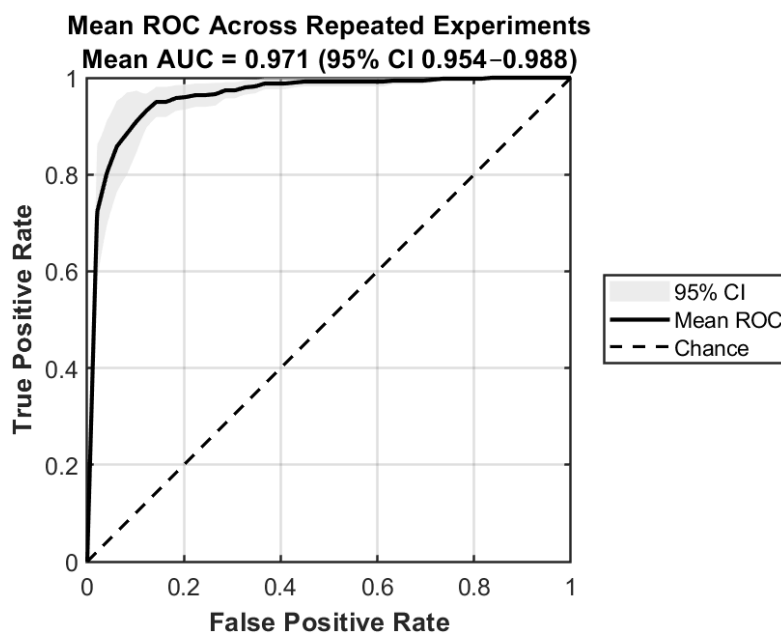


Figure 1. The receiver operating characteristic (ROC) curve diagram for the GoogLeNet-Places365. The average ROC and its 95% confidence interval across 10 train-test splits (experiments) are shown in a solid black line and a light grey background, respectively. The average area under the curve (AUC) was 0.971, with a 95% confidence interval (CI) of [0.954 – 0.988].

3.4. Performance at the Youden-Optimal Threshold

Using the Youden-optimal threshold per-split, the GoogleNetPlaces365 achieved a mean sensitivity of 92% (95% CI: 88% – 95%) and a mean specificity of 95% (95% CI: 93 – 98) across experiments (Figure 2). The corresponding mean \pm sd (standard deviation) precision, F1-score, and accuracy were $95 \pm 3.5\%$, $93 \pm 3.4\%$, and $93 \pm 3.3\%$, respectively.

Performance at the Youden-Optimal Operating Point

Mean Sensitivity = 0.918 (95% CI 0.883–0.953)

Mean Specificity = 0.951 (95% CI 0.925–0.977)

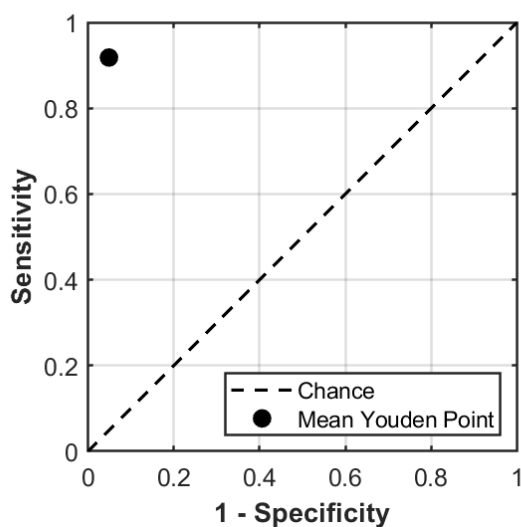


Figure 2. Performance of the GoogLeNet-Places365 convolutional neural network at the Youden-Optimal point. The Mean Youden point across 10 experiments is shown in a solid black dot.

4. Discussion

AI will revolutionize the way we practice medicine, transforming traditional approaches in many medical specialties and healthcare applications. In 2011, more than 1000 AI medical papers were uploaded to PubMed. Since then, the number has increased spectacularly, reaching approximately 33,000 papers in 2025.

Regarding COVID-19 (COroNaVirus Disease 2019), an acute respiratory disease caused by the SARS-CoV-2 virus (Severe Acute Respiratory Syndrome CoronaVirus 2), it emerged in late 2019 and rapidly spread on a worldwide scale, causing the 2019-2023 global pandemic. The pandemic had an overwhelming impact worldwide, leading to vast disruptions in healthcare systems, education, and social life. It ignited an extraordinary level of scientific international collaboration, enabling the rapid deployment of novel mRNA vaccines and antiviral treatments. Today, the virus continues to circulate, but it is typically managed through vaccination and public health surveillance.

Despite the effective management of the disease, some people cannot fully recover, presenting long-term symptoms which are usually named “Long-COVID-19” or “post-acute COVID-19 syndrome” without a universally accepted explanation. However, hemodynamic and thrombotic effects have been observed in the blood and microvessels of COVID-19 patients during [16,17] and after [15] the disease. Long-lasting microvascular thrombosis is functionally equivalent to microvascular loss reported by many quantitative case–control studies from multiple tissues [18–21]. Based on these findings, a pathophysiological framework for the explanation of Long COVID was grounded in persisting microvascular blood supply reduction [22], and 76% of the total incidence rate of Long COVID symptoms was directly attributed to this mechanism [23].

A key aspect of the smallest microvasculature that has not been examined yet in post COVID patients, is the microvessel network “spatial distribution” or “topological structure” (topology). The fundamental emerging question is whether COVID-19 affects the microvessel network topology. Based on this question, the central target in this work was to consider the microvessel topology as a “scene” and examine whether the microvessel topological structure of post-COVID patients can be discriminated from that of healthy individuals using the superior capabilities of artificial intelligence in image classification.

For that purpose, a variety of pretrained deep convolutional neural networks were assessed and the best average testing accuracy across 10 random data splits of 92% ($\pm 3\%$) was achieved by the GoogLeNet-Places365 network. GoogLeNet-Places365, a scene-centric network, excelled on all other performance criteria of Table 6, and performed much better than object-centric CNNs with a 4% difference in average testing accuracy from the second-best performance. Larger differences were observed in other evaluation criteria (Table 6), such as a 7% difference in Sensitivity (92% for GoogLeNet-Places365 versus 85% for InceptionResNet-v2) and Specificity (91% for GoogLeNet-Places365 versus 84% for EfficientNet-b0).

Overall, the GoogLeNet-Places365 exhibited both excellent discriminative ability and consistent performance at clinically relevant conditions with an average ROC AUC of 0.971 across 10 train-test splits (Figure 1). At the Youden-Optimal operating point, the high values and limited dispersion of sensitivity and specificity under varying data partitions indicates stable operating characteristics (Figure 2).

Regarding input image size (IIS), it is generally accepted that the bigger the better for testing accuracy (Tan and Le, 2019). Inception ResNet-v2 has a much larger IIS (299x299) than GoogLeNet and EfficientNet-b0 (224x224) (Table 1). In addition, Inception ResNet-v2 is significantly deeper (Table 1) and larger in size (Table 2) than the other networks. However, these architectural advantages of Inception ResNet-v2 did not result in better classification performance over the other networks (Table 6). GoogLeNet-Places365 performed significantly better at all classification evaluation criteria. Therefore, for this particular case, it is concluded that the critical factor was not depth and size, but the pretraining on a dataset of scenes and backgrounds (Places365) instead of objects.

Comparing the decision times (DT) of the 4 fine-tuned pretrained networks in Table 6, we observe that all average DTs are less than a few seconds, and the most accurate network (GeegLeNet-Places365) requires less than 1 second to make the final prediction using only one image. These decision times are significantly shorter (more than a thousand times) than the approximately 15 minutes needed by current antigen (rapid) COVID-19 detection tests (detection of viral proteins). They are also 4 orders of magnitude shorter than the hours required by the current gold standard of molecular PCR tests (detection of the virus's genetic material). The future prospect of a much faster COVID-19 test is reinforced by the noncontact nature of this AI imaging method in contrast to the requirement of liquid human samples (nose, throat, saliva) by the current tests. On surplus, since the images come from the bulbar conjunctiva and not the retina, the outcome does not depend on the optical clarity of the eye, which may be reduced significantly in cases of cataract or vitreous hemorrhage.

A limitation of the study was the small sample size and the subsequent medium to small image dataset.

5. Conclusions

In conclusion, it was shown that post-COVID effects on the eye microcirculation can be detected by deep CNNs with a relatively high average accuracy and excellent discriminative power. The critical factor for achieving higher accuracy was not depth and size, but the pretraining on a dataset of scenes and backgrounds. This AI-based post-COVID detection method could revolutionize the way we detect viral infections since it is risk-free, painless, contactless, fast, accurate, and does not depend on the optical clarity of the eye.

Author Contributions: Conceptualization, A.G.K.; methodology, A.G.K.; software, A.G.K.; formal analysis, A.G.K.; investigation, A.G.K., K.R., S.B., T.N.P.; resources, A.G.K., S.B., M.K., and E.E.T.; data curation, A.G.K.; writing—original draft preparation, A.G.K., K.R., and S.B.; writing—review and editing, A.G.K., M.K., Z.D., A.H.Z., and E.E.T.; visualization, A.G.K.; supervision, A.G.K., Z.D., A.H.Z., and E.E.T.; project administration, A.G.K., and E.E.T. All authors have read and agreed to the published version of the manuscript.

Funding: This research received no external funding.

Institutional Review Board Statement: The study was conducted in accordance with the Declaration of Helsinki and approved by the University Hospital of Larissa Scientific Committee (11155/13-4-2021, date 2021-04-13).

Informed Consent Statement: Informed consent was obtained from all subjects involved in the study.

Data Availability Statement: The original contributions presented in this study are included in the article. Further inquiries can be directed to the corresponding author.

Acknowledgments: We thank OpenAI and Google for their free versions of GPT Scholar and Gemini respectively, which helped us in improving coding and clarifying various AI related subjects.

Acknowledgments: The research is conducted in the operating framework of the University of Thessaly Innovation, Technology Transfer Unit and Entrepreneurship Center "One Planet Thessaly", under the "University of Thessaly Grants for Scientific Publication Support" action and is funded by the Special Account of Research Grants of the University of Thessaly.

Conflicts of Interest: The authors declare no conflicts of interest.

Abbreviations

The following abbreviations are used in this manuscript:

AI	Artificial intelligence
AP	Antithrombotic prophylaxis
ARDS	Acute Respiratory Distress Syndrome
ASA	Acetylsalicylic acid
AUC	Area under the curve
BMI	Body mass index
BPM	Beats per minute
BSA	Body surface area
CCD	Charge coupled device
CHADS	Congestive heart failure, hypertension, age, diabetes mellitus, stroke
CI	Confidence interval
CNN	Convolutional Neural Network
COVID-19	Corona virus disease of 2019
CPU	Central processing unit
DP	Diastolic pressure
DT	Decision time
HFNC	High flow nasal cannula
HI	Horowitz index
ICU	Intensive care unit
IIS	Input image size
IIC	Input image channels
IL6R	InterLeukine-6 Receptor,
ILSVRC	Imagenet large scale visual recognition challenge
LMWH	Low molecular weight heparin
LR	Learning rate
MBS	Mini batch size
MEP	Maximum number of epochs
ML	Machine learning
MP	Mean pressure
NI	Number of images per subject
NOL	Number of layers
PC	Personal computer
PCR	Polymerase chain reaction
PM	Parameter memory
RAM	Random access memory
ROC	Receiver operating characteristic
SD	Standard deviation
SSD	Solid state drive
SP	Systolic pressure
TNI	Total number of images
TT	Training time
WHO	World Health Organization

References

1. von Gerich, H.; Helenius, M.; Hörhammer, I.; others Economic Aspects of Machine Learning in Healthcare: A Systematic Review. *Int. J. Med. Inf.* **2025**, *190*, 105532, doi:10.1016/j.ijmedinf.2025.105532.
2. Parra-Hernandez, R.M.; Dominguez-Morales, M.J.; others Artificial Intelligence in Histopathology: A Systematic Review. *Artif. Intell. Rev.* **2023**, *56*, 9055–9090, doi:10.1007/s10462-023-10555-2.
3. Aggarwal, R.; Sounderajah, V.; Martin, G.; Ting, D.S.W.; Karthikesalingam, A.; King, D.; Ashrafian, H.; Darzi, A. Diagnostic Accuracy of Deep Learning in Medical Imaging: A Systematic Review and Meta-Analysis. *Npj Digit. Med.* **2021**, *4*, 65, doi:10.1038/s41746-021-00438-z.

4. Esteva, A.; Kuprel, B.; Novoa, R.A.; Ko, J.; Swetter, S.M.; Blau, H.M.; Thrun, S. Dermatologist-Level Classification of Skin Cancer with Deep Neural Networks. *Nature* **2017**, *542*, 115–118, doi:10.1038/nature21056.
5. Lecun, Y.; Bottou, L.; Bengio, Y.; Haffner, P. Gradient-Based Learning Applied to Document Recognition. *Proc. IEEE* **1998**, *86*, 2278–2324, doi:10.1109/5.726791.
6. Krizhevsky, A.; Sutskever, I.; Hinton, G.E. ImageNet Classification with Deep Convolutional Neural Networks. *Commun. ACM* **2017**, *60*, 84–90, doi:10.1145/3065386.
7. Christian Szegedy; Wei Liu; Yangqing Jia; Pierre Sermanet; Scott Reed; Dragomir Anguelov; Dumitru Erhan; Vincent Vanhoucke; Andrew Rabinovich Going Deeper with Convolutions. In Proceedings of the 2014 IEEE Conference on Computer Vision and Pattern Recognition; IEEE: Columbus, OH, USA, June 2014; pp. 1–9.
8. He, K.; Zhang, X.; Ren, S.; Sun, J. Deep Residual Learning for Image Recognition. In Proceedings of the 2016 IEEE Conference on Computer Vision and Pattern Recognition (CVPR); IEEE: Las Vegas, NV, USA, June 2016; pp. 770–778.
9. Szegedy, C.; Ioffe, S.; Vanhoucke, V.; Alemi, A. Inception-v4, Inception-ResNet and the Impact of Residual Connections on Learning. *Proc. AAAI Conf. Artif. Intell.* **2017**, *31*.
10. Tan, M.; Le, Q.V. EfficientNet: Rethinking Model Scaling for Convolutional Neural Networks. **2019**, doi:10.48550/ARXIV.1905.11946.
11. Zhou, B.; Lapedriza, A.; Khosla, A.; Oliva, A.; Torralba, A. Places: A 10 Million Image Database for Scene Recognition. *IEEE Trans. Pattern Anal. Mach. Intell.* **2018**, *40*, 1452–1464, doi:10.1109/TPAMI.2017.2723009.
12. Deng, J.; Dong, W.; Socher, R.; Li, L.-J.; Kai Li; Li Fei-Fei ImageNet: A Large-Scale Hierarchical Image Database. In Proceedings of the 2009 IEEE Conference on Computer Vision and Pattern Recognition; IEEE: Miami, FL, June 2009; pp. 248–255.
13. Zhou, B.; Lapedriza, A.; Xiao, J.; Torralba, A.; Oliva, A. Learning Deep Features for Scene Recognition Using Places Database. *Adv. Neural Inf. Process. Syst. NIPS* **2014**, *27*.
14. Moka, S.; Koutsiaris, A.G.; Garas, A.; Messinis, I.; Tachmitzi, S.V.; Giannoukas, A.; Tsironi, E.E. Blood Flow Velocity Comparison in the Eye Capillaries and Postcapillary Venules between Normal Pregnant and Non-Pregnant Women. *Microvasc. Res.* **2020**, *127*, 103926, doi:10.1016/j.mvr.2019.103926.
15. Koutsiaris, A.G.; Riri, K.; Boutlas, S.; Panagiotou, T.N.; Kotoula, M.; Daniil, Z.; Tsironi, E.E. COVID-19 Hemodynamic and Thrombotic Effect on the Eye Microcirculation after Hospitalization: A Quantitative Case-Control Study. *Clin. Hemorheol. Microcirc.* **2022**, *82*, 379–390, doi:10.3233/CH-221554.
16. Popazu, C.; Romila, A.; Petrea, M.; Grosu, R.M.; Lescai, A.-M.; Vlad, A.L.; Oprea, V.D.; Baltă, A.A. Ștefania Overview of Inflammatory and Coagulation Markers in Elderly Patients with COVID-19: Retrospective Analysis of Laboratory Results. *Life* **2025**, *15*, 370, doi:10.3390/life15030370.
17. Zlojutro, B.; Jandric, M.; Momcicevic, D.; Dragic, S.; Kovacevic, T.; Djajic, V.; Stojilkovic, M.P.; Skrbic, R.; Djuric, D.M.; Kovacevic, P. Dynamic Changes in Coagulation, Hematological and Biochemical Parameters as Predictors of Mortality in Critically Ill COVID-19 Patients: A Prospective Observational Study. *Clin. Hemorheol. Microcirc.* **2023**, *83*, 137–148, doi:10.3233/CH-221583.
18. Cutolo, M.; Sulli, A.; Smith, V.; Gotelli, E. Emerging Nailfold Capillaroscopic Patterns in COVID-19: From Acute Patients to Survivors. *Reumatismo* **2023**, *74*, doi:10.4081/reumatismo.2022.1555.
19. Kazantzis, D.; Machairoudia, G.; Theodossiadis, G.; Theodossiadis, P.; Chatziralli, I. Retinal Microvascular Changes in Patients Recovered from COVID-19 Compared to Healthy Controls: A Meta-Analysis. *Photodiagnosis Photodyn. Ther.* **2023**, *42*, 103556, doi:10.1016/j.pdpdt.2023.103556.
20. Osiaevi, I.; Schulze, A.; Evers, G.; Harmening, K.; Vink, H.; Kümpers, P.; Mohr, M.; Rovas, A. Persistent Capillary Rarefaction in Long COVID Syndrome. *Angiogenesis* **2023**, *26*, 53–61, doi:10.1007/s10456-022-09850-9.
21. Sulli, A.; Gotelli, E.; Bica, P.F.; Schiavetti, I.; Pizzorni, C.; Aloè, T.; Grosso, M.; Barisione, E.; Paolino, S.; Smith, V.; et al. Detailed Videocapillaroscopic Microvascular Changes Detectable in Adult COVID-19 Survivors. *Microvasc. Res.* **2022**, *142*, 104361, doi:10.1016/j.mvr.2022.104361.

22. Koutsiaris, A.G. A Blood Supply Pathophysiological Microcirculatory Mechanism for Long COVID. *Life* **2024**, *14*, 1076, doi:10.3390/life14091076.
23. Koutsiaris, A.G.; Karakousis, K. Long COVID Mechanisms, Microvascular Effects, and Evaluation Based on Incidence. *Life* **2025**, *15*, 887, doi:10.3390/life15060887.

Disclaimer/Publisher's Note: The statements, opinions and data contained in all publications are solely those of the individual author(s) and contributor(s) and not of MDPI and/or the editor(s). MDPI and/or the editor(s) disclaim responsibility for any injury to people or property resulting from any ideas, methods, instructions or products referred to in the content.