

# Towards an Efficient Deep Learning Model for COVID-19 Patterns Detection in X-ray Images

Eduardo J. S. Luz, Pedro Lopes Silva, Rodrigo Silva and Gladston J. P. Moreira

**Abstract**—Confronting the pandemic of COVID-19 caused by the new coronavirus, the SARS-CoV-2, is nowadays one of the most prominent challenges of the human species. A key factor in slowing down the virus propagation is rapid diagnosis and isolation of infected patients. Nevertheless, the standard method for COVID-19 identification, the RT-PCR, is time-consuming and in short supply due to the pandemic.

Researchers around the world have been trying to find alternative screening methods. In this context, deep learning applied to chest X-rays of patients has been showing a lot of promise for the identification of COVID-19. Despite their success, the computational cost of these methods remains high, which imposes difficulties in their accessibility and availability. Thus, in this work, we address the hypothesis that better performance in terms of overall accuracy and COVID-19 sensitivity can be achieved with much more compact models. In order to test this hypothesis, we propose a modification of the EfficientNet family of models. By doing this we were able to produce a high-quality model with an overall accuracy of 91.4%, COVID-19, sensitivity of 90% and positive prediction of 100% while having about 30 times fewer parameters than the baseline model, 28 and 5 times fewer parameters than the popular VGG16 and ResNet50 architectures, respectively.

**Index Terms**—COVID-19, Pneumonia, Deep Learning, Radiography Chest Images, EfficientNet

## I. INTRODUCTION

**I**N December 2019, many Chinese citizens in the province of Wuhan were affected by severe pneumonia. In January 2020, the cause was found to be a new virus of the coronavirus family (SARS-CoV-2) [1]. The virus quickly spread to other countries, and in a short time, it became a pandemic, changing the lives of many people around the globe. In February 2020, The World Health Organization (WHO) named the disease caused by SARS-COV-2 as COVID-19 and researchers from different fields have turned their efforts to fight it.

One of the problems regarding COVID-19 identification is its diagnosis. Although methods based on real-time reverse-transcriptase polymerase chain reaction (RT-PCR) are the gold

standard for diagnosing this type of virus [2], they are time-consuming and in short supply due to the pandemic. Besides, it has been reported in [3] that they tend to present low sensitivity ( $\sim 73\%$  for nasal swabs and  $\sim 61\%$  for throat swabs). In general, to ensure reliable results, multiple tests need to be done with samples spaced by days.

Chest radiography imaging (e.g., X-ray or computed tomography (CT) imaging) has also been used as an alternative screening method for COVID-19 [2], [4], [5]. In this case, a chest radiography examination is conducted by radiologists to look for visual indicators associated with COVID-19 viral infection. Notably, Ng *et al.* [5] found consistent patterns in the computed tomography (CT) and X-ray images of COVID-19 patients. However, they state that humans are not able to perceive these patterns on images of the chest for the vast majority of asymptomatic patients.

Deep Learning is a subset of machine learning in artificial intelligence (AI) concerned with algorithms inspired by the structure and function of the brain called artificial neural networks. Since deep learning techniques, in particular convolutional neural networks (CNNs), have been beating humans in various tasks of computer vision [6]–[8], it becomes a natural candidate for the analysis of chest radiography images.

Deep learning has already been explored for the detection and classification of pneumonia and other diseases on radiography. In [9], a 10-layer CNN was proposed to identify seven patterns observed in different interstitial lung diseases. In their study, 120 CT scans were divided into 14696 image patches, and the goal was to classify each patch into each pattern. The authors reported an 85% accuracy of the method. In [8], the authors proposed a 121-layer convolutional neural network trained on the ChestX-ray14 dataset [10] which contains over 100,000 frontal view X-ray images for 14 diseases. The authors reported a superior performance of the method when compared with four practicing academic radiologists. Recently, in [11], a convolutional neural network (CNN) with a branch for predicting a segmentation mask was used to classify chest X-ray from the RSNA dataset [12] into pneumonia-negative and pneumonia-positive. In addition, in the pneumonia positive cases, the method produces a bounding box around the lung opacities.

Addressing the COVID-19, in [13], a comparison among seven different well-known deep learning neural networks architectures was presented. In the experiments, they use a small data set with only 50 images in which 25 samples are from healthy patients and 25 from COVID-19 positive patients. The models were pre-trained with the ImageNet dataset [14], which is a generic image dataset with over 14 million im-

“This work was supported by PROPP/UFOP and Brazilian funding agencies CAPES, CNPq and FAPEMIG.”

E. Luz, P. Silva, R. Silva and G. Moreira are with the Department of Computer Science, Federal University of Ouro Preto, Campus Morro do Cruzeiro s/n, Ouro Preto-MG, Brazil, 35400-000 (e-mails: {eduluz, pedro.lopes1, rodrigo.silva, gladston}@ufop.edu.br).

This article has supplementary downloadable material available at <http://www.decom.ufop.br/csilab>, provided by the authors.

©2020 IEEE. Personal use of this material is permitted. Permission from IEEE must be obtained for all other uses, in any current or future media, including reprinting/republishing this material for advertising or promotional purposes, creating new collective works, for resale or redistribution to servers or lists, or reuse of any copyrighted component of this work in other works.

ages of all sorts, and only the classifier is trained with the radiography. In their experiments, the VGG19 [15] and the DenseNET201 [16] were the best performing architectures. In [17], a new architecture of CNN, called COVID-net, is created to classify X-ray images into normal, pneumonia, and COVID-19. Differently from the previous work, they use a much larger dataset consisting of 16,756 chest radiography images across 13,645 patient cases. The authors report an accuracy of 92.4% overall and sensitivity of 80% for COVID-19. In [18], the ResNet50 [19] is fine tuned for the problem of classifying X-ray images into normal, COVID-19, bacterial-pneumonia and viral pneumonia. The authors report better results when compared with the COVID-net, 96.23% accuracy overall, and 100% sensitivity for COVID-19. Nevertheless, it is important to highlight that the problem in [18] has an extra class and that the datasets are different. In [18] the dataset consists of 68 COVID-19 radiographs from 45 COVID-19 patients, 1203 healthy patients, 931 patients with a bacterial pneumonia and 660 patients with nonCOVID-19 viral pneumonia.

It is important to highlight that, at this point, there is no other peer-reviewed work that deals with COVID-19 screening through radiography images. Unfortunately, all the related work found so far, that is [13], [17], [18], are preprints available at <https://arxiv.org/>. From these, the only work which is fully reproducible, and because of that, the only one reliable at this point is [17] from which the authors have made all the code for model and protocols for their dataset construction available at <https://github.com/lindawangg/COVID-Net>.

Thus, this work aims to investigate deep learning models that are capable of finding patterns in CT / X-ray images of the chest, even if the patterns are imperceptible to the human eye, and to advance on a fundamental issue: computational cost.

The main model available thus far, the COVID-net [17], has over 110 million parameters amounting for over 2GB in memory. We believe that a mobile application that integrates deep learning models for the task of recognizing patterns in x-rays or CT must be easily accessible and readily available to the medical staff. And for that, the models must have a low footprint and low latency, that is, the models must require little memory and perform inference quickly to allow use on embedded devices and large scale, enabling integration with cell phones and medical equipment.

Following the experimentation protocol proposed in [17], our results suggest that it may be feasible to embed our proposed neural network model in a mobile device and make fast inferences. Despite its low computational cost, the proposed model achieves high accuracy (92.8 %) and detect infection caused by COVID-19 on chest X-rays with a Sensitivity of 90 % and Positivity Prediction of 100 %. The developments of this work may allow the future construction of an application for use by the medical team, through a camera on a regular cell phone.

The remainder of this paper consists of five sections. Section II defines the problem tackled in this paper. The methodology and the dataset are described in Section III. In Section IV, the results of a comprehensive set of computational experiments are presented. In Section V, propositions for future research

in the area are addressed. Finally, the article is concluded in Section VI.

## II. PROBLEM SETTINGS

The problem that will be dealt with by the proposed approach can be defined as: given a chest X-ray, determine if it belongs to a healthy patient, a patient with COVID-19, or a patient with other forms of pneumonia.

Fig. 1 shows typical chest X-ray samples in our dataset. As can be seen, the model does not make restrictions in regards to the view in which the X-ray was taken. Thus, given an image like these, the model will output one of the following three possible labels:

- *normal* - for healthy patients.
- *COVID-19* - for patients with COVID-19.
- *pneumonia* - for patients with non-COVID-19 pneumonias.

Following the rationale in [17], choosing these three possible predictions can help clinicians in deciding who should be prioritized for PCR testing for COVID-19 case confirmation. It might also help in treatment selection since COVID-19, and non-COVID-19 infections require different treatment plans.

## III. METHODOLOGY

In this section, we present the methodology for COVID-19 detection by means of a X-ray image. We detail the main datasets and and briefly describe the COVID-Net [17], our baseline method. Also, we describe the employed deep learning techniques as well as the learning methodology and evaluation.

### A. Datasets

1) *RSNA Pneumonia Detection Challenge dataset*: The RSNA Pneumonia Detection Challenge [12] is a competition that aims to locate lung opacities on chest radiographs. Pneumonia is associated with opacity in the lung, and some conditions such as pulmonary edema, bleeding, volume loss, lung cancer can also lead to opacity in lung radiography. Find patterns associated with pneumonia is a hard task. In that sense, the Radiological Society of North America (RSNA) has promoted the challenge, providing a rich dataset. Although The RSNA challenge is a segmentation challenge, here we are using the dataset for a classification problem. The dataset offers images for two classes: Normal and Pneumonia (non-normal). We are using a total of 16680 images of this dataset, of which 8066 are from Normal class and 8614 from the Pneumonia class.

2) *COVID-19 image data collection*: The “COVID-19 Image Data Collection” [20] is a collection of images, acquired from websites and research papers. The dataset was created by researchers from the University of Montreal with the help of the international research community to assure that it will be continuously updated. Nowadays, the dataset includes more than 120 X-ray images of patients who were affected by COVID-19 and other diseases, such as MERS, SARS, and ARDS. The dataset is public and also includes CT scans.




Fig. 1: Radiograph example of images from COVID-19 image data collection [20]. (a) X-ray of a 54-year-old male, infected with COVID-19 [20]. (b) X-ray of a 70-year-old female, infected with COVID-19 [20].

According to the authors, the dataset can be used to assess the advancement of COVID-19 in infected individuals, and also allow the identification of patterns related to COVID-19 helping in differentiating it from other types of pneumonia. Besides, chest X-ray images can be used as an initial screening for the COVID-19 diagnostic processes. So far, most of the images are from male individuals (approx. 60% of individuals are male), and the age group that concentrates most cases is 50 to 80 years old.

3) *COVID<sub>x</sub> dataset*: In [17], a new dataset is proposed by merging two other public datasets: “RSNA Pneumonia Detection Challenge dataset” and “COVID-19 Image Data Collection”. The new dataset, called *COVID<sub>x</sub>*, was designed for a classification problem and contemplates three classes: Normal, Pneumonia, and COVID-19. Most instances of the Normal and Pneumonia classes come from the “RSNA Pneumonia Detection Challenge dataset”, and all instances of the COVID-19 class come from the “COVID-19 Image Data Collection”. The dataset has a total of 16756 images from 13645 individuals and is split into two partitions, one for training purposes and one for testing (model evaluation). The distribution of images between the partitions is shown in Table I, and the source code to reproduce the dataset is publicly available<sup>1</sup>.

TABLE I: *COVID<sub>x</sub>* Images distribution among classes and partitions. The dataset is proposed in [17].

Type	Normal	Pneumonia	COVID-19	Total
Train	7966	8514	66	16546
Test	100	100	10	210

## B. COVID-Net

The COVID-Net architecture is based on the generative synthesis technique [21]. The generative synthesis consists of a *generative-inquisitor* pair, formulated to learn how to generate neural network architectures under constraints. For the particular case of the COVID-Net, a human specialist defined

two restrictions: (i) test accuracy  $\geq 80\%$  and (ii) network computational complexity  $\leq 2.5$  billion multiply-accumulate (MAC) operations. The human also has to determine the basic network block. For the COVID-Net, the main building block is the residual projection-expansion-projection-extension (PEPX), formed by the operations illustrated by Figure 2 inspired on residual blocks of [22].




Fig. 2: PEPX: Basic block for COVID-Net construction.

The COVID-Net was pre-trained on the ImageNet and then fine-tuned with the *COVID<sub>x</sub>* dataset, using transfer learning. Weights are updated by the Adam Optimizer, with a schedule rule decreasing the learning rate by a factor of 10 in the event of stagnation during training (‘patience policy’). For fine-tuning, learning rate started with  $2^{-5}$ , number of epochs = 10, batch size = 8, factor = 0.7 and patience = 5 were used. Translation, rotation, horizontal flip, and intensity shifts were also used to augment training data. Finally, a batch re-balancing strategy was used to promote a better distribution of data during the creation of mini-batches at training time.

## C. EfficientNet

The EfficientNet [23] is in fact a family of models defined over the baseline network described in Table II.

Its main component is known as the Mobile Inverted Bottleneck Conv (MBconv) Block introduced in [24] and depicted in Figure 3.

The rationale behind the EfficientNet family is to start from high quality yet compact baseline model and uniformly scale each of its dimensions systematically with a fixed set of scaling coefficients. Formally, an EfficientNet is defined by three dimensions: (i) depth; (ii) width; and (iii) resolutions as illustrated in Figure 4.

Starting from the baseline model in Table V each dimension is scaled by the parameter  $\phi$  according to Eq. 1. As stated in [23], Eq. 1 provides a nice balance between performance and computational cost.

<sup>1</sup><https://github.com/lindawangg/COVID-Net>

TABLE II: EfficientNet baseline network : B0 architecture.

Stage	Operator	Resolution	#channels	#layers
1	Conv3x3	224x224	32	1
2	MBCConv1,k3x3	112x112	16	1
3	MBCConv6,k3x3	112x112	24	2
4	MBCConv6,k5x5	56x56	40	2
5	MBCConv6,k3x3	28x28	80	3
6	MBCConv6,k5x5	14x14	112	3
7	MBCConv6,k5x5	14x14	192	4
8	MB Conv6,k3x3	7x7	320	1
9	Conv1x1/Pooling/FC	7x7	1280	1




Fig. 3: MBCov Block [24]. DWConv stands for depthwise conv, k3x3/k5x5 defines the kernel size, BN is batch norm,  $HxWxF$  means tensor shape (height, width, depth), and  $\tilde{A}1/2/3/4$  is the multiplier for number of repeated layers.




Fig. 4: Efficient net compound scaling on three parameters. (Adapted from [23])

$$\begin{aligned}
 \text{depth} &= \alpha^\phi \\
 \text{width} &= \beta^\phi \\
 \text{resolution} &= \gamma^\phi \\
 \text{s.t. } \alpha \cdot \beta^2 \cdot \gamma^2 &\approx 2 \\
 \alpha \geq 1, \beta \geq 1, \gamma &\geq 1
 \end{aligned}
 \tag{1}$$

Where  $\alpha, \beta, \gamma$  are constants obtained by a grid search. The coefficient  $\phi$  controls the available resources. The equation 1 determines the increase or decrease of model FLOPS when depth, width or resolution are modified.

Notably, in [23], a model from EfficientNet family was able to beat the powerful GPipe Network [16] on the ImageNet dataset [25] running with 8.4x fewer parameters and 6.1x faster.

#### D. Hierarchical Classification

In classification problems, it is common to have some sort of relationship among classes. Very often, on real problems, the classes (the category of an instance) are organized hierarchically, like a tree structure. According to Silla Jr. and Freitas [26], one can have three types of classification: flat classification, which ignores the hierarchy of the tree; local classification, in which there is a set of classifiers for each level of the tree (one classifier per node or level); and finally, global classification, in which one single classifier is built with the ability to classify any node in the tree, besides the leaves.

The most popular type of classification in the literature is the flat one. However, here we propose the use of local classification, which we call hierarchical classification. Thus, the target classes are located in the leaves of the tree, and in the intermediate nodes, we have classifiers. In this work, we need two classifiers, one at the root node, dedicated to discriminate between the Normal and Pneumonia classes, and another in the following level dedicated to discriminate between pneumonia types. The problem addressed here can be mapped as the topology depicted in Figure 5 in which there are two levels of classification. To make the class inference for a new instance, first, the instance is presented to the first classifier (in the root node). If it is predicted as “Normal”, the inference ends there. If the instance is considered “Pneumonia”, it is then presented to the second classifier, which will discern whether it is a Pneumonia caused by “COVID-19” or “Not”.

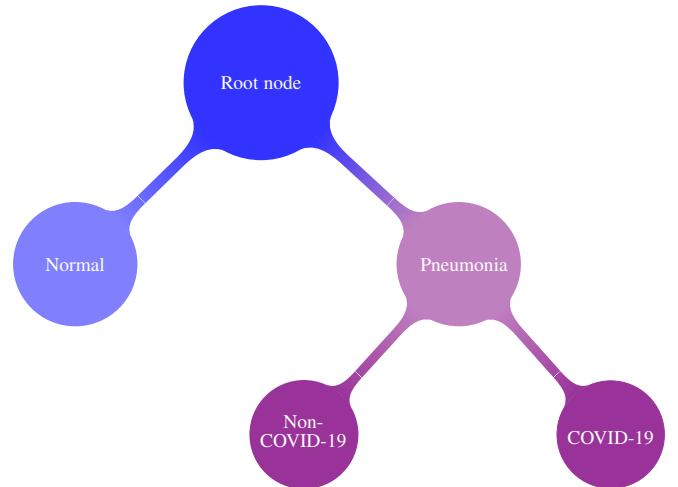


Fig. 5: Natural topology of the classes: Normal, Pneumonia, COVID-19. The Figure illustrates the Local-Per-Node hierarchical approach, in which there is a classifier on each parent node.

## E. Training

Deep learning models are complex, and therefore require a large number of instances to avoid overfitting, i.e., when the learned network performs well on the training set but underperform on the test set. Unfortunately, for most problems in real-world situations, data is not abundant. In fact, there are few scenarios in which there is an abundance of training data, such as the ImageNet [25], in which there are more than 14 million images of 21,841 classes/categories. To overcome this issue, researchers rely on two techniques: data augmentation and transfer learning. We also detail here the proposed models, based on EfficientNet.

1) *Image Pre-processing and Data Augmentation*: Several pre-processing techniques may be used for image cleaning, noise removal, outlier removal, etc. The only pre-processing applied in this work is a simple normalization of the image pixels intensity to the range [0 to 1]. In this manner, we rely on the filters of the convolutional network itself to perform possible data cleaning.

Data augmentation consists of expanding the training set with transformations of the images in the dataset [27] provided that the semantic information is not lost. In this work, we applied three transformations to the images: rotation, horizontal flip, and scaling, as such transformations would not hinder, for example, a physician to interpret the radiography.

2) *Proposed models*: The EfficientNet family has models of high performance and low computational cost. Since this research aims to find efficient models capable of being embedded in conventional smartphones, the EfficientNet family is a natural choice.

We explore the EfficientNets by adding more operator blocks on top of it. More specifically, we add four new blocks, as detailed in Table III.

Since the original EfficientNets were built to work on a different classification problem we add new fully connected layers (FC) responsible for last steps of the classification process. We also use batch normalization (BN), dropout, and swish activation functions for the reasons below.

The batch normalization constrains the output of the last layer in a range, forcing zero mean and standard deviation one. That acts as regularization, increasing the stability of the neural network, and accelerating the training [28].

The Dropout [29] is perhaps the most powerful method of regularization. The practical effect of dropout operation is to emulate a bagged ensemble of multiple neural networks by inhibiting a few neurons, at random, for each mini-batch during training. The number of inhibited neuronal units is defined by the dropout parameter, which ranges between 0 to 100 percent.

The most popular activation function is the Rectified Linear Unit (ReLU), which can be formally defined as  $f(x) = \max(0, x)$ . However, in the added block we have opted for the *swish activation function* [30] defined as:

$$f(x) = x \cdot (1 + \exp^{-x})^{-1} \quad (2)$$

Differently from the ReLU the swish activation is a smooth curve which is a desirable property since we will be using

a gradient descent algorithm to minimize the loss. Besides, again, differently from the ReLU, the swish activation does not zero out small negative values which may still be relevant for capturing patterns underlying the data [30].

TABLE III: Proposed models architecture, considering the EfficientNet B0 as base model. (NC = Number of Classes).

Stage	Operator	Resolution	#channels	#layers
1-9	EfficientNet B0	224x224	32	1
10	BN/Dropout	7x7	1280	1
11	FC/BN/Swich/Dropout	1	512	1
12	FC/BN/Swich	1	128	1
13	FC/Softmax	1	NC	1

3) *Transfer learning*: Instead of training a model from scratch, one can take advantage of the weights from a pre-trained network and accelerate or enhance learning. As discussed in [31], the initial layers of a model can be seen as feature descriptors for image representation, and the latter ones are related to instance categories. Thus, in many applications, several layers can be re-used. The task of transfer learning is then to define how and what layers of a pre-trained model to use. This technique has proved to be effective in several computer vision tasks, even when transferring weights from completely different domains [27], [32].

The steps for transfer of learning are:

- 1) Copying the weights from a pre-trained model to a new model;
- 2) Modifying the architecture of the new model to adapt it to the new problem, possibly including new layers;
- 3) Initialize the new layers;
- 4) Define which layers will pass through a new the learning process; and
- 5) training (updating the weights according to the loss function) with a suitable optimization algorithm.

We apply transfer learning to EfficientNets pre-trained on the ImageNet dataset [25]. It is clear that the ImageNet domain is much broader than the chest X-rays that will be presented to the models in this work. Thus, the imported network weights are taken just as an initial solution and are all fine-tuned by the optimizer over the new training phase. The rationale is that the imported models already have a lot of knowledge about all sorts of objects. By permitting all the weights to get fine-tuned we allow the model to specialize to the problem in hands. In the training phase, the weights are updated with the Adam Optimizer and a schedule rule decreasing the learning rate by a factor of 10 in the event of stagnation ('patience=2'). The learning rate started with  $10^{-4}$ , and the number of epochs fixed at 10.

## F. Model evaluation and metrics

The final evaluation is carried with the COVIDx dataset, and since the COVIDx comprises a combination of two other public datasets, we follow the script<sup>2</sup> provided in [17] to load

<sup>2</sup><https://github.com/lindawangg/COVID-Net>

the training and test sets. The data is then distributed according to the Table I.

With the aim to assess the generalization power of the proposed method, we design a second test, by adding 37 new COVID-19 images cases on the test set. Those images were collected from Italian patients [33], made available by the *Italian Society of Medical and Interventional Radiology*<sup>3</sup>.

In this work, three metrics are used to evaluate models: accuracy, COVID-19 sensitivity ( $Se_C$ ), and COVID-19 positive prediction ( $+P_C$ ). All three metrics are presented in Equation 3.

$$accuracy = \frac{TP_N + TP_P + TP_C}{\#samples}$$

$$Se_C = \frac{TP_C}{TP_C + FN_C} \quad (3)$$

$$+P_C = \frac{TP_C}{TP_C + FP_C}$$

wherein  $TP_N$  are the normal samples correctly classified,  $TP_P$  are the non-COVID-19 samples correctly classified,  $TP_C$  are the COVID-19 samples correctly classified,  $FN_C$  are the COVID-19 samples classified as normal or non-COVID-19,  $FP_C$  are the non-COVID-19 and normal samples classified as COVID-19. The number of multiply-accumulate (MAC) operations are used to measure the computational cost.

#### IV. EXPERIMENTS AND DISCUSSION

In this section, we present the experimental setup and results for both flat and hierarchical approaches. The execution environment of the computational experiments was conducted on an Intel(R) Core(TM) i7-5820K CPU @ 3.30GHz, 64Gb Ram, a Titan X with 12Gb, and the TensorFlow/Keras framework for Python.

##### A. Dataset setup

Three different training set configurations were analyzed: *i*) (Raw Dataset) - the raw dataset without any pre-processing; *ii*) (Raw Dataset + Data Augmentation) - the raw dataset with a data augmentation of 1,000 new images on COVID-19 samples and a limitation of 4,000 images for the two remaining classes; and *iii*) (Balanced Dataset) - the dataset with a 1,000 images per class achieved by data augmentation on COVID-19 samples and reducing on the other two classes. Work with an unbalanced dataset will bias the prediction model towards the classes with more samples, conducting in a worse model for COVID-19 classification.

In this work, we have to evaluate two scenarios: flat and hierarchical. Regardless of the scenarios, the three training sets remain the same (Raw, Raw + Data Augmentation, and Balanced). However, for the hierarchical case, there is an extra process to split the sets into two parts: the first part, the instances of Pneumonia and COVID-19 classes are joined and receive the same label (Pneumonia). In the second part, the

instances related to the Normal class are removed, leaving in the set only instances related to Pneumonia and COVID-19. Thus, two classifiers are built for the hierarchical case, and each one works with a different set of data (see Section III-D for more details).

##### B. Experimental details and results

We evaluate four families of convolutional neural networks: EfficientNet, MobileNet, VGG and ResNet. Their features are summarized in Table IV. Among the presented models, we highlight the low footprint of MobileNet and EfficientNet.

TABLE IV: Base models footprint details. (Mb = Megabytes)

Model	Input shape	#Params	Memory usage (Mb)
EfficientNet B0	224, 224, 3	5,330,564	21
EfficientNet B1	240, 240, 3	7,856,232	31
EfficientNet B2	260, 260, 3	9,177,562	36
EfficientNet B3	300, 300, 3	12,320,528	48
EfficientNet B4	380, 380, 3	19,466,816	76
EfficientNet B5	456, 456, 3	30,562,520	118
MobileNet	224, 224, 3	4,253,864	17
MobileNet V2	224, 224, 3	3,538,984	14
ResNet 50	224, 224, 3	25,636,712	99
VGG-16	224, 224, 3	138,357,544	528
VGG-19	224, 224, 3	143,667,240	549

Regarding the base models (B0-B5 models of EfficientNet family), the one with the lowest computational cost is the EfficientNet-B0. Thus, we started assess the impact of the three different training sets and the two forms of classification (flat and hierarchical) from this model. The results are shown in the Tables V and VI. As one can see, both positive prediction and sensitivity for the COVID-19 class are conflicting with accuracy. That can be explained by the unbalanced nature of data. Since there are more pneumonia, and normal x-ray samples than COVID-19, the neural network learning process tends to improve the classification of the majoritarian classes, since they have more weight for the loss calculation. This may also justify the results obtained by balancing the data.

TABLE V: EfficientNet B0 results over the three proposed training sets.

Classification	Raw Dataset	Raw Dataset + Data Augmentation	Balanced Dataset
Accuracy	92.3%	91.4%	91.4%
COVID-19 Sensitivity	50.0%	80.0%	90.0%
COVID-19 Positive Prediction	100.0%	66.7%	100.0%

As described in Section III-D, the hierarchical approach is also evaluated here. First, it classifies samples into pneumonia + COVID-19 and normal, and later, pneumonia and COVID-19. Table VI presents the same metrics used for the flat sce-

<sup>3</sup><https://www.sirm.org/>



nario. The hierarchical approach presented higher accuracy but no improvement in terms of positive prediction and sensitivity.

TABLE VI: Hierarchical results over the three proposed training sets.

Classification	Raw Dataset	Raw Dataset + Data Augmentation	Balanced Dataset
Accuracy	89.4%	92.8%	92.8%
COVID-19 Sensitivity	50.0%	80.0%	80.0%
COVID-19 Positive Prediction	83.0%	80.0%	80.0%




Fig. 6: Loss during training-time, EfficientNet-B0 and balanced data. Epochs vs Loss.

Analyzing both Tables V and VI, it is possible to see that the best configuration for the EfficientNet-B0 is the flat approach with a balanced training set. This scenario will be used to evaluate the remaining network architectures. The training loss for this scenario is presented in Figure 6.

The results of all evaluated architectures are summarized in Table VII. It can be seen that all the networks have comparable performances in terms of accuracy. However, the more expensive the model was, the more difficulty it had in classifying the COVID-19 class.

The cost of a model is related to the number of parameters. The higher the number of parameters, the higher the amount of data the model needs to adjust them. Thus, the lack of a bigger dataset may explain the difficulties faced by the more sophisticated models.

Table VIII presents a comparison of the proposed approach and the one proposed by Wang *et al.* [17] (COVID-net).

Even though the accuracy is comparable, the proposed approach presents an improvement on positive prediction without losing sensitivity. Besides, a significant reduction both in terms of memory (our model is >30 times smaller) and latency is observed. It is worth highlighting that Wang *et al.* [17] apply data augmentation to the dataset but it is not clear in their manuscript how many new images are created.

TABLE VII: Results over different network architectures as base model. Best scenario for COVID-19: all experiments with a balanced training set and flat classification.

Base Model	Accuracy	COVID-19 Sensitivity	COVID-19 Positive Prediction
EfficientNet B0	91.4%	90.0%	100.0%
EfficientNet B1	91.4%	60.0%	100.0%
EfficientNet B2	91.4%	90.0%	75.0%
EfficientNet B3	92.3%	90.0%	100.0%
EfficientNet B4	90.0%	50.0%	100.0%
EfficientNet B5	91.9%	50.0%	83.3%
MobileNet	92.3%	80.0%	88.9%
MobileNet V2	91.9%	70.0%	87.5%
ResNet-50	89.5%	30.0%	100.0%
VGG-16	80.5%	20.0%	22.2%
VGG-19	79.0%	90.0%	29.0%

The COVID-Net [17] is a very complex network, which demands a memory of 2.1GB and performs over 3.5 billion MAC operations implying three main drawbacks: computation-cost, time-consumption, and infrastructure costs. A 3.59 billion MAC operations model takes much more time and computations than a 4.7 million MAC model - in the order of almost 100 times -, and the same GPU necessary to run one COVID-Net model can run more than 50 models of the proposed approach (based on the EfficientNet B0) keeping a comparable (or even better) positive prediction and sensitivity for all classes. All this complexity can hinder the use of the model in the future, for instance, on mobile phones.

### C. Discussion

The presence of the COVID-19 infection can be observed through some opacity (white spots) on chest radiography imaging. In Figure 7, we present an X-ray image of a non-healthy person and its activation map. The main activation spots of the proposed approach have a considerable overlay with opacity points, which could indicate the presence of Pneumonia or COVID-19.




Fig. 7: The image (b) is the activation map using the EfficientNet B0 of image (a).

In regards to the proposed methodology, the hierarchical model presented a better overall result compared to the flat

TABLE VIII: Comparison of the proposed approach against SOTA.

Method	COVID-19 Sensitivity	COVID-19 Positive Prediction	#Params (M)	MACs	Memory required
Proposed approach	90.0%	<b>100.0%</b>	<b>4.8</b>	<b>4.5 millions</b>	<b>56Mb</b>
COVID-net [17]	90.0%	90.0%	126.6	3.5 billions	2.1Gb




Fig. 8: Activation map of EfficientNet B0 model. The first line are COVID-19 x-ray, and the second, other kind of pneumonia.

one. However, it has also shown a noticeable reduction in terms of sensitivity and positive prediction for the COVID-19 class. In Figure 9, the confusion matrices of flat and hierarchical approaches are presented. It is possible to observe that the hierarchical model classifies the normal class better. One hypothesis is that both Pneumonia and COVID-19 classes are similar (both kinds of pneumonia) and share key features. Thus, the lack of normal images reduces the diversity of the training set, interfering with model training. Besides, the computational cost is twice higher than flat classification since two models are required.




Fig. 9: Confusion matrix of flat and hierarchical approaches respectively with balanced training set. Class zero is the normal images, 1, pneumonia non-COVID-19, and, 2, COVID-19.

The reduced number of images on the test set, considering the protocol of [17], can lead to a metric overestimation. To examine the robustness of the best configuration on the

proposed approach, 37 new COVID-19 cases were added to the test set. The images were collected and made available by the Italian Society of Medical and Interventional Radiology [33]. With the addition of the new cases, a small decrease in sensitivity is noticeable, on the other hand, the COVID-19 positive prediction remained unchanged, as shown in Table IX. This analysis indicates the robustness of the proposed approach. In Figure 10, presents the confusion matrix for the new test.

TABLE IX: Results on the second test set, increased by 37 COVID-19 cases, from Italian patients [33].

Method	Accuracy	COVID-19 Sensitivity	COVID-19 Positive Prediction
Flat Approach	89.9%	82.9%	100.0




Fig. 10: Confusion matrix for the flat approach, trained with balanced training set, evaluated with the addition of 37 new images of COVID-19 cases. Class zero is the normal images, 1, pneumonia non-COVID-19, and, 2, COVID-19.

## V. FINDINGS AND FUTURE DIRECTION

We summarize our findings as follows.

- An efficient and low computational approach was proposed to classify infections caused by COVID-19 on chest X-ray images. Even with only a few images of the COVID-19 class, insightful results with a sensitivity of 90% and a positive prediction of 100% were obtained, with the evaluation protocol proposed in [17].
- The evaluation protocol proposed in [17] is based on the public dataset “COVID-19 Image Data Collection” [20], which is being expanded by the scientific community. With more images from the COVID-19 class, it will be possible to improve the training. However, the test partition tends to become more challenging. Our code is available at <https://github.com/ufopcsilab/EfficientNet-C19> to facilitate reproducibility and future comparisons.
- Even with the addition of new images from another set of data [33], the proposed model maintained its performance during the test, indicating robustness.



- The Internet of Medical Things (IOMT) [34] is now a hot topic on industry. However, the internet can be a major limitation for medical equipment, especially in poor countries. Our proposal is to move towards a model that can be fully embedded in conventional cell phones (edge computing), eliminating the use of the internet or cloud services. In that sense, the model achieved in this work requires only 55Mb of memory and has a viable inference time for a conventional cell phone processor.

## VI. CONCLUSION

In this paper, we exploit an efficient convolutional network architecture for detecting infection caused by COVID-19 through chest radiography images. Experiments were conducted to evaluate the neural network performance on two different datasets, the COVIDx and COVID-19 Database/Italian Society of Medical and Interventional Radiology, using two approaches: flat classification and hierarchical classification. Although the datasets are still incipient and, therefore, limited in the number of COVID-19 related images, it was possible to conduct efficient training on the deep neural networks, due to transfer learning and data augmentation techniques.

Concerning evaluation, the proposed approach brought improvements compared to baseline work, with a computational efficiency of more than 30 times higher (Acc of 91.4%, COVID-19 Sensitivity of 90% and Positivity Prediction of 100%).

We believe that the current proposal is a promising candidate for embedding in medical equipment or even physicians' mobile phones and help in screening the diagnosis of COVID-19, once more mature COVID-19 image datasets are made available.

## ACKNOWLEDGMENT

The authors thank UFOP and funding Brazilian agencies CAPES, FAPEMIG and CNPq. We gratefully acknowledge the support of NVIDIA Corporation with the donation of two Titan X Pascal GPU used for this research.

## REFERENCES

- [1] "WHO. coronavirus disease (covid-19) outbreak situation," <https://www.who.int/emergencies/diseases/novel-coronavirus-2019>, accessed: 2020-04-07.
- [2] P. Huang, T. Liu, L. Huang, H. Liu, M. Lei, W. Xu, X. Hu, J. Chen, and B. Liu, "Use of chest ct in combination with negative rt-pcr assay for the 2019 novel coronavirus but high clinical suspicion," *Radiology*, vol. 295, no. 1, pp. 22–23, 2020.
- [3] Y. Yang, M. Yang, C. Shen, F. Wang, J. Yuan, J. Li, M. Zhang, Z. Wang, L. Xing, J. Wei, L. Peng, G. Wong, H. Zheng, M. Liao, K. Feng, J. Li, Q. Yang, J. Zhao, Z. Zhang, L. Liu, and Y. Liu, "Evaluating the accuracy of different respiratory specimens in the laboratory diagnosis and monitoring the viral shedding of 2019-ncov infections," *medRxiv*, 2020. [Online]. Available: <https://www.medrxiv.org/content/early/2020/02/17/2020.02.11.20021493>
- [4] T. Ai, Z. Yang, H. Hou, C. Zhan, C. Chen, W. Lv, Q. Tao, Z. Sun, and L. Xia, "Correlation of chest ct and rt-pcr testing in coronavirus disease 2019 (covid-19) in china: a report of 1014 cases," *Radiology*, p. 200642, 2020.
- [5] M.-Y. Ng, E. Y. Lee, J. Yang, F. Yang, X. Li, H. Wang, M. M.-s. Lui, C. S.-Y. Lo, B. Leung, P.-L. Khong *et al.*, "Imaging profile of the covid-19 infection: radiologic findings and literature review," *Radiology: Cardiothoracic Imaging*, vol. 2, no. 1, p. e200034, 2020.
- [6] Y. LeCun, Y. Bengio, and G. Hinton, "Deep learning," *nature*, vol. 521, no. 7553, pp. 436–444, 2015.
- [7] H. Touvron, A. Vedaldi, M. Douze, and H. Jégou, "Fixing the train-test resolution discrepancy: Fixefficientnet," *arXiv preprint arXiv:2003.08237*, 2020.
- [8] P. Rajpurkar, J. Irvin, K. Zhu, B. Yang, H. Mehta, T. Duan, D. Ding, A. Bagul, C. Langlotz, K. Shpanskaya *et al.*, "Chexnet: Radiologist-level pneumonia detection on chest x-rays with deep learning," *arXiv preprint arXiv:1711.05225*, 2017.
- [9] M. Anthimopoulos, S. Christodoulidis, L. Ebner, A. Christe, and S. Mougiakakou, "Lung pattern classification for interstitial lung diseases using a deep convolutional neural network," *IEEE transactions on medical imaging*, vol. 35, no. 5, pp. 1207–1216, 2016.
- [10] X. Wang, Y. Peng, L. Lu, Z. Lu, M. Bagheri, and R. Summers, "Hospital-scale chest x-ray database and benchmarks on weakly-supervised classification and localization of common thorax diseases," in *IEEE CVPR*, 2017.
- [11] A. K. Jaiswal, P. Tiwari, S. Kumar, D. Gupta, A. Khanna, and J. J. Rodrigues, "Identifying pneumonia in chest x-rays: A deep learning approach," *Measurement*, vol. 145, pp. 511–518, 2019.
- [12] "Radiological Society of North America. RSNA pneumonia detection challenge." <https://www.kaggle.com/rsna-pneumonia-detection-challenge/data>, accessed: 2020-04-01.
- [13] E. E.-D. Hemdan, M. A. Shouman, and M. E. Karar, "Covid-net: A framework of deep learning classifiers to diagnose covid-19 in x-ray images," *arXiv preprint arXiv:2003.11055*, 2020.
- [14] J. Deng, W. Dong, R. Socher, L.-J. Li, K. Li, and L. Fei-Fei, "ImageNet: A Large-Scale Hierarchical Image Database," in *CVPR09*, 2009.
- [15] K. Simonyan and A. Zisserman, "Very deep convolutional networks for large-scale image recognition," *arXiv preprint arXiv:1409.1556*, 2014.
- [16] Y. Huang, Y. Cheng, A. Bapna, O. Firat, D. Chen, M. Chen, H. Lee, J. Ngiam, Q. V. Le, Y. Wu *et al.*, "Gpipe: Efficient training of giant neural networks using pipeline parallelism," in *Advances in Neural Information Processing Systems*, 2019, pp. 103–112.
- [17] L. Wang and A. Wong, "Covid-net: A tailored deep convolutional neural network design for detection of covid-19 cases from chest radiography images," *arXiv preprint arXiv:2003.09871*, 2020.
- [18] M. Farooq and A. Hafeez, "Covid-resnet: A deep learning framework for screening of covid19 from radiographs," *arXiv preprint arXiv:2003.14395*, 2020.
- [19] C. Szegedy, S. Ioffe, V. Vanhoucke, and A. A. Alemi, "Inception-v4, inception-resnet and the impact of residual connections on learning," in *Thirty-first AAAI conference on artificial intelligence*, 2017.
- [20] J. P. Cohen, P. Morrison, and L. Dao, "Covid-19 image data collection," *arXiv preprint arXiv:2003.11597*, 2020.
- [21] A. Wong, M. J. Shafiee, B. Chwyl, and F. Li, "Ferminets: Learning generative machines to generate efficient neural networks via generative synthesis," *arXiv preprint arXiv:1809.05989*, 2018.
- [22] K. He, X. Zhang, S. Ren, and J. Sun, "Deep residual learning for image recognition," in *Proceedings of the IEEE conference on computer vision and pattern recognition*, 2016, pp. 770–778.
- [23] M. Tan and Q. V. Le, "Efficientnet: Rethinking model scaling for convolutional neural networks," *arXiv preprint arXiv:1905.11946*, 2019.
- [24] M. Sandler, A. Howard, M. Zhu, A. Zhmoginov, and L.-C. Chen, "Mobilenetv2: Inverted residuals and linear bottlenecks," in *Proceedings of the IEEE conference on computer vision and pattern recognition*, 2018, pp. 4510–4520.
- [25] O. Russakovsky, J. Deng, H. Su, J. Krause, S. Satheesh, S. Ma, Z. Huang, A. Karpathy, A. Khosla, M. Bernstein *et al.*, "Imagenet large scale visual recognition challenge," *International journal of computer vision*, vol. 115, no. 3, pp. 211–252, 2015.
- [26] C. N. Silla and A. A. Freitas, "A survey of hierarchical classification across different application domains," *Data Mining and Knowledge Discovery*, vol. 22, no. 1-2, pp. 31–72, 2011.
- [27] I. Goodfellow, Y. Bengio, and A. Courville, *Deep learning*. MIT press, 2016.
- [28] S. Ioffe and C. Szegedy, "Batch normalization: Accelerating deep network training by reducing internal covariate shift," *arXiv preprint arXiv:1502.03167*, 2015.
- [29] N. Srivastava, G. Hinton, A. Krizhevsky, I. Sutskever, and R. Salakhutdinov, "Dropout: a simple way to prevent neural networks from overfitting," *The journal of machine learning research*, vol. 15, no. 1, pp. 1929–1958, 2014.
- [30] P. Ramachandran, B. Zoph, and Q. V. Le, "Searching for activation functions," *arXiv preprint arXiv:1710.05941*, 2017.

- [31] M. Oquab, L. Bottou, I. Laptev, and J. Sivic, "Learning and transferring mid-level image representations using convolutional neural networks," in *Proceedings of the IEEE conference on computer vision and pattern recognition*, 2014, pp. 1717–1724.
- [32] E. Luz, G. Moreira, L. A. Z. Junior, and D. Menotti, "Deep periocular representation aiming video surveillance," *Pattern Recognition Letters*, vol. 114, pp. 2–12, 2018.
- [33] I. S. of Medical and I. Radiology, "COVID-19. database," <https://www.sirm.org/en/category/articles/covid-19-database/>, accessed: 2020-04-12.
- [34] G. J. Joyia, R. M. Liaqat, A. Farooq, and S. Rehman, "Internet of medical things (iomt): applications, benefits and future challenges in healthcare domain," *Journal of Communications*, vol. 12, no. 4, pp. 240–7, 2017.