Deep Learning-based COVID-19 diagnostics of low-quality CT images

Daniel Ferber, Felipe Vieira, João Dalben, Mariana Ferraz,
Nicholas Sato, Gabriel Oliveira, Rafael Padilha and Zanoni Dias

November 25, 2021

Institute of Computing, University of Campinas, Campinas, SP, Brazil

Agenda

Introduction

Introduction

- COVID-19 has impacted healthcare systems, economic activity and citizen lives around the world
- Machine learning and deep learning classifiers could help the fight against COVID-19
- Blood cell counts, chest X-rays, and computed tomography (CT) scans can be used for training such classifiers
- CT scans present a higher sensitivity than X-rays for this problem and have been successfully used in computer-aided diagnosis
- The quality of a CT sample depends on many aspects, such as the conditions of the sensor, how the subject poses in the patient bed and which clothes they are wearing
- In this work, we evaluate several preprocessing techniques to reduce the impact of low-quality CT scans into deep learning-based diagnostic tools for COVID-19

Dataset

COVIDx-CT Dataset

Table 1: Image and patient distributions per set of the COVIDx-CT dataset [1].

	Images			**Patients**		
	Train	Val	Test	Train	Val	Test
Normal	27,201	9,107	9,450	144	47	52
CP	22,061	7,400	7,395	420	190	125
NCP	12,520	4,529	4,346	300	95	116
Total	61,782	21,036	21,191	864	332	293

Examples of Images

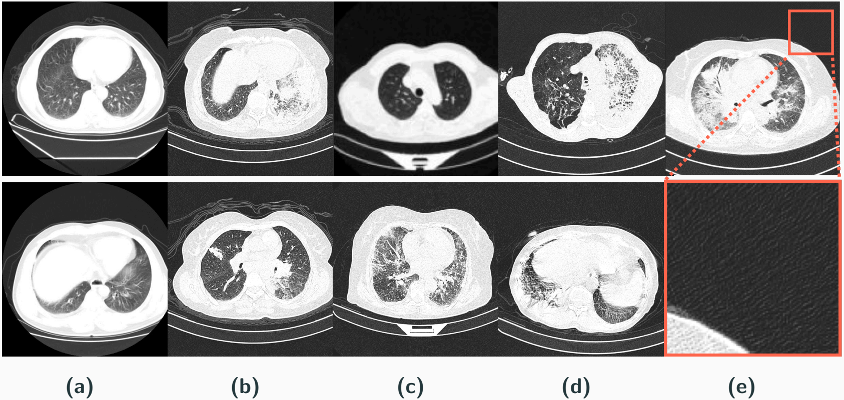

Figure 1: Examples of artifacts present in CT images of the COVIDx-CT dataset [1]: **(a)** rounded borders, **(b)** traces of clothes around chest area, **(c)** structure of the patient bed in the bottom part of the images, **(d)** reflection and bright artifacts outside chest area, and **(e)** background noise.

Methodology

Pipeline

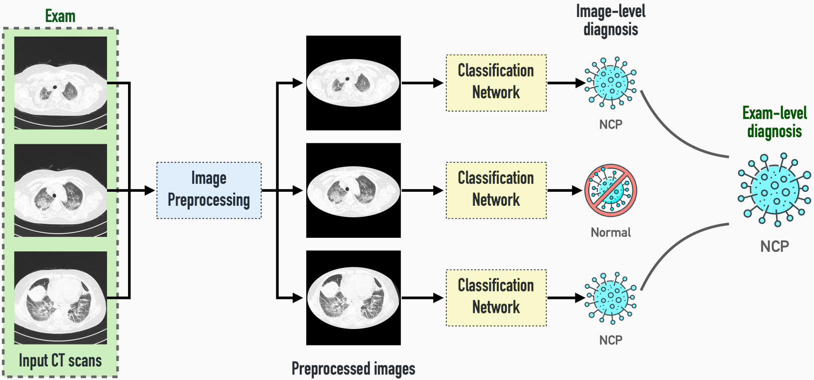

Figure 2: Overview of our pipeline.

Data Processing

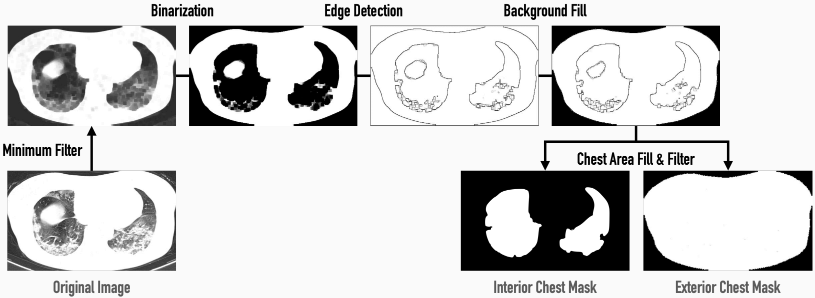

Figure 3: Preprocessing steps performed to extract chest masks.

Data Processing

- Rectangular
- Rectangular-Centered
- Elliptical
- Elliptical-Centered
- Exterior Chest
- Interior Chest

Data Processing

Figure 4: Original image and examples of generated masked images.

CNN architectures

- Baseline architecture
  - Three convolutional blocks, followed by a block of fully-connected layers
  - Each convolutional block has two convolutional layers, followed by a max pooling with 2$\times$2 kernel and dropout operation with rate 0.2
  - The number of convolutional filters starts at 32 for the first block and is multiplied by 2 in each subsequent block
  - The fully-connected block has three fully-connected layers with 256, 128 and 3 neurons
- ResNet50 architecture [2]

Experimental Evaluation

Pipeline

1. Data processing with the baseline architecture
2. ResNet50 with different number of unfrozen layers
3. Image with different resolutions
4. Classification of the test set
5. Interpretation of the model decisions

Baseline - Data Preprocessing

Table 2: Balanced accuracy and COVID-19 false negatives with respect to the validation set, considering baseline models trained with each preprocessing strategy. Best result is highlighted in **bold** and second best is <u>underlined</u>.

Preprocessing	**$\uparrow$ Acc (%)**	**$\downarrow$ COVID-19 FN (%)**
No Preprocessing	<u>94.51</u>	2.39
Rectangular	90.93	3.49
Rectangular-Centered	94.28	2.63
Elliptical	85.67	6.64
Elliptical-Centered	92.42	2.61
Exterior Chest	93.39	<u>2.04</u>
Interior Chest	**94.53**	**1.31**

ResNet50 - Unfrozen Layers

Table 3: Balanced accuracy and COVID-19 false negatives with respect to the validation set, considering the ResNet50 architecture trained with different amount of unfrozen layers for the *Exterior Chest* and *Interior Chest* preprocessing strategies. Best result for each metric is highlighted in **bold**.

Unfrozen	**Exterior Chest**		**Interior Chest**	
Layers	↑ Acc (%)	↓ FN (%)	↑ Acc (%)	↓ FN (%)
0	93.83	2.58	93.66	2.55
12	94.57	2.57	94.37	2.54
24	95.20	1.85	94.85	2.49
32	96.00	**0.87**	**95.39**	**2.30**
40	**96.21**	1.72	95.13	2.53

ResNet50 - Image Resolution

Table 4: Evaluation of ResNet50 models, trained with 128×128 and 224×224 resolution images preprocessed by the *Exterior Chest* strategy, with 32 unfrozen layers. Best result is highlighted in **bold**.

Resolution	$\uparrow$ Acc (%)	$\downarrow$ FN (%)
128×128	96.00	0.87
224×224	**97.65**	**0.78**

Classification of the Test Set

- 97.84% of balanced accuracy on image-level
- 99.50% of balanced accuracy on exam-level

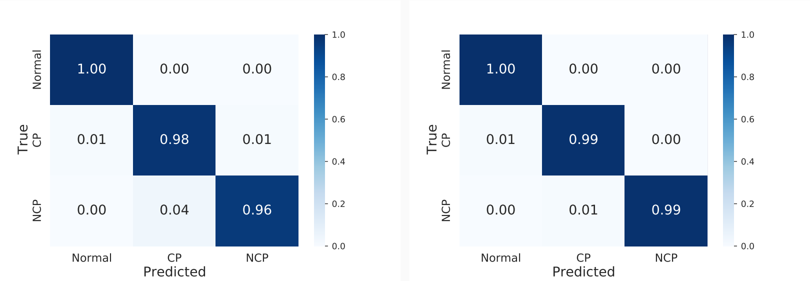

(a) Image-level **(b)** Exam-level

Figure 5: Confusion matrix for image- and exam-levels on the test set.

Interpreting Model Decisions

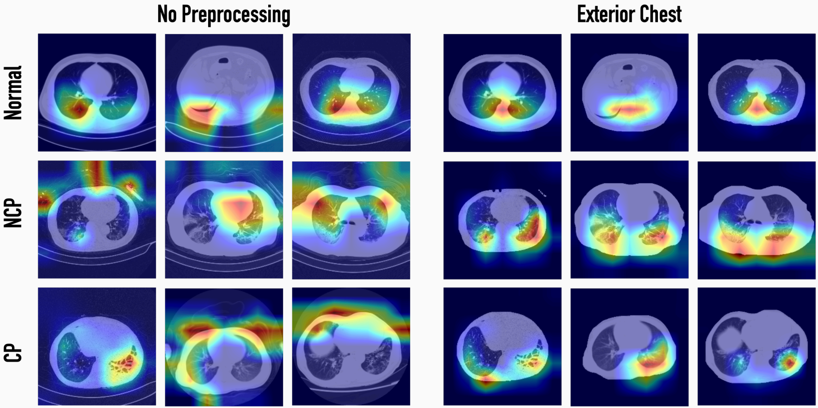

Figure 6: Class-activation maps [3] for *Normal*, *NCP* and *CP* images from models without any preprocessing and optimized with *Exterior Chest* masking strategy.

Conclusions and Future Work

Conclusions

- The importance of removing noisy information from the images
- The complexity of the models and high resolution images improved the results

Future Work

- Ensemble techniques
- Different CNN architectures
- Complementarity of preprocessing strategies

References

[1] Hayden Gunraj, Linda Wang, and Alexander Wong.
COVIDNet-CT: A tailored deep convolutional neural network design for detection of COVID-19 cases from chest CT images.
7, 2020.

[2] Kaiming He, Xiangyu Zhang, Shaoqing Ren, and Jian Sun.
Deep residual learning for image recognition.
pages 770–778, 2016.

[3] Ramprasaath R Selvaraju, Michael Cogswell, Abhishek Das, Ramakrishna Vedantam, Devi Parikh, and Dhruv Batra.
Grad-cam: Visual explanations from deep networks via gradient-based localization.
pages 618–626, 2017.