Detection and Segmentation of Lesions in Chest CT Scans for The Prediction of COVID-19

Seminar Presentation at the University of Kent School of Computing

Alex Ter-Sarkisov,
Lecturer in Artificial Intelligence,
Department of Computer Science,
City, University of London

November 23, 2020

Main Findings

1. Prediction of COVID-19 vs Common Pneumonia (CP) vs Control/Negative on large datasets is a challenging problem.
2. Detected areas with lesions in chest CT scans can be used for predicting the class of the whole image (COVID-19 vs CP vs Control).
3. The use of the advanced methodology like instance detection and segmentation both improves the accuracy of the prediction and reduces the demand for the training data.
4. Instance model can be extended to make global (image class) predictions with very high accuracy: COVID-19 sensitivity of 93.88% and overall accuracy of 95.64%

Plan of the presentation

1. Chest CT scans vs X-rays,
2. Prediction of COVID-19 from chest CT scans,
3. Deep Learning for the detection + segmentation algorithm
4. COVID-CT-Mask-Net model
5. Results
6. Papers + OS Code

Chest CT scans vs X-Rays

CT vs X-rays

- Axial slices vs frontal,
- Slices are merged into a single prediction,
- X-ray: faster, easier to interpret,
- CT: more accurate

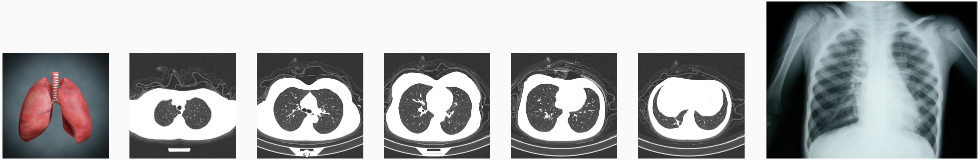

#23 #46 #69 #92 #115

Figure 1: Chest CT Scan slices vs single X-ray. Source: CNCB, OS

Chest CT Scans For COVID-19

Chest CT Scans For COVID-19

1. More accurate than the X-ray, less labor-intensive than rRT-PCR,
2. Automated prediction of COVID-19 from chest CT scans: a method of supporting doctors and radiologists
3. Pneumonia, incl COVID-19 have various manifestations in chest CT scans (lesions)

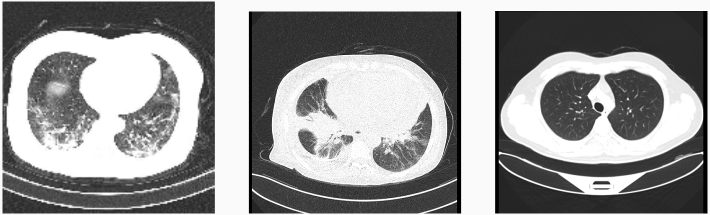

(a) (b) (c)

Figure 2: Chest CT Scan. (a): COVID-19 positive, (b): Common Pneumonia, (c): Normal. Source: CNCB

Chest CT Scans For COVID-19

1. The differences between COVID-19 and other types of pneumonia are observable (e.g. D.Zhao et al, 2020; X.Li et al, 2020;), but often not statistically significant.
2. Two types of lesions correlated with both conditions: Ground Glass Opacity (GGO) and Consolidation (C).
3. To distinguish between the two conditions, the *configuration* of each lesion type is important: location (uni- vs bilateral), distribution (peripheral, diffuse), range, number, attenuation (fade).

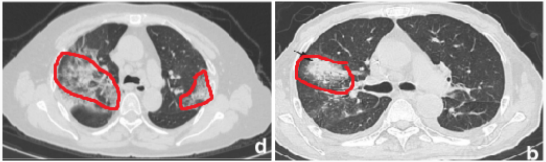

Figure 3: COVID-19 (left) vs CP. Source: X.Li et al, 2020

Lesion Detection using Deep Learning

Lesion Detection

Two types of DL solutions exist:

1. Straightforward (feature extractor + class logits, e.g. in Gunraj et al, 2020, Butt et al, 2020) using ResNet,ResNeXt, DenseNet or a tailored solution (e.g. COVNet, COVNet-CT).
2. Feature extractor + semantic segmentation (e.g. UNet): predicted masks are concatenated with the feature maps to predict the image class, e.g. JCS, Y-H Wu et al, 2020.

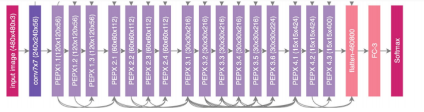

(a)

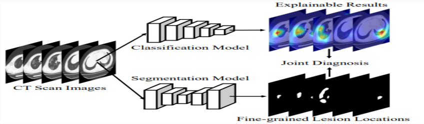

(b)

Figure 4: (a): COVIDNet, source: L.Wang et al, 2020 , (b): JCS Sourece: Y-H Wu et al, 2020

Lesion Detection

The main disadvantage of the existing methods is their dependence on large amounts of data for training, data augmentation tricks, etc. Many were trained on small amounts of data, and therefore their ability to generalize requires additional testing.

1. We can train DL algorithms that detect (bounding box) and segment (contours) lesions in chest CT scans. Using these algorithms has a number of advantages:
   1.1 They generalize well to the unseen (new) data,
   1.2 A number of OS solutions (e.g. Torchvision models pretrained on MS COCO 2017) is available that can be finetuned (transfer learning) to the problem at hand,
   1.3 Training is fast of GPUs
2. The main advantage is their architecture is that it reduces the size of the data required for the training: in addition to using batches of images, they *extract training batches from each image*.

Lesion Detection

1. Faster R-CNN (Ren et al, 2015) and Mask R-CNN (He et al, 2017) solve the problem of detecting (predicting the bounding box+class) and segmenting each object (predict a mask of the object) *independently*, i.e. the model understands objects at the *instance* level rather than image or pixel.
2. This is in contrast to both image classification like ResNet and semantic (pixel-level) segmentation like UNet. Instead of nameless feature maps and image-wide score maps Mask R-CNN can handle partial occlusion and disconnected objects.
3. As a result, Mask R-CNN does a very accurate instance segmentation and prediction.

Lesion Detection

To train Mask R-CNN on COVID-19 data we need the input images and the image masks, from which everything else is extracted:

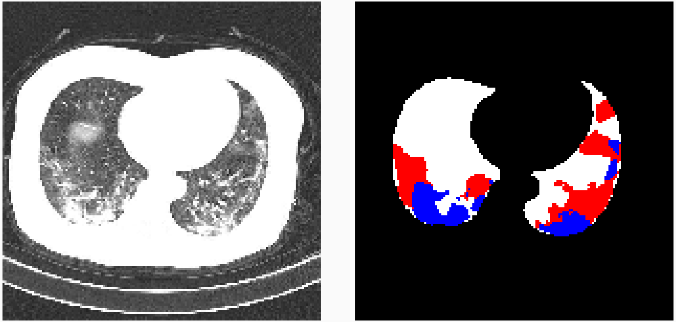

(a) (b)

Figure 5: (a): Input Image , (b): Mask with two types of lesions. Source: CNCB

Lesion Detection

The Model:

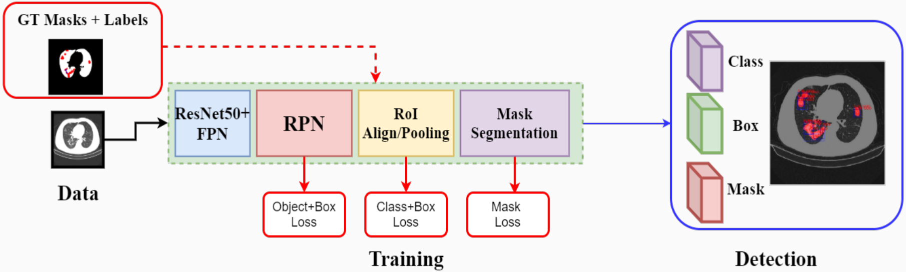

Figure 6: Mask R-CNN training

Lesion Detection

The output/prediction + *independent* score maps, RoI score$_\theta$=0.75:

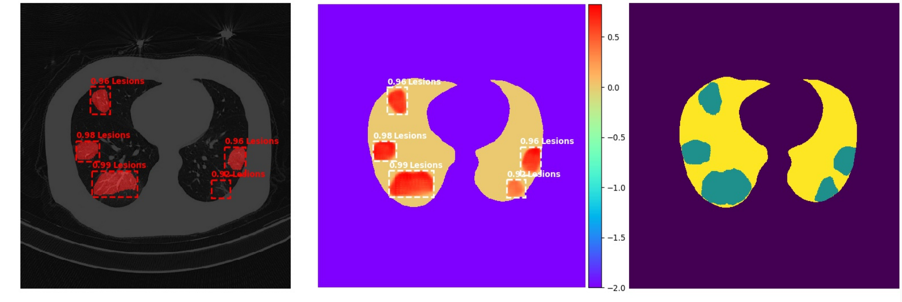

Figure 7: Mask R-CNN output. Left: input imag with the overlaid predictions (class, mask, box), center: normalized scoremaps overlaid with the mask, right: gt masks

COVID-CT-Mask-Net

COVID-19 Classifier

1. Mask R-CNN localizes objects, but this says nothing about the *global* (image) class, i.e. Mask R-CNN doesn't *understand* the image level
2. Specifically, there are many overlaps between COVID-19 and other types of pneumonia. It is impossible to classify a slice as either class from just one region. The accurate prediction requires consideration of all regions of interest (RoIs) detected by Mask R-CNN.
3. Solution: consider *all* positive RoIs regardless of their score and extract the image class from them/learn their distribution,
4. Solution: COVID-CT-Mask-Net extends Mask R-CNN to classify whole images

COVID-19 Classifier

The model (Mask R-CNN + Image Classification Module):

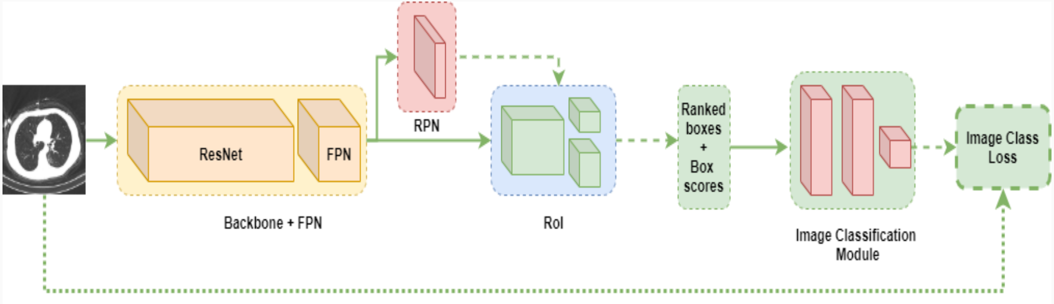

Figure 8: COVID-CT-Mask-Net. Normal arrows: tensors/features, broken arrows: batches/predictions, dotted arrow: image label

COVID-19 Classifier

The model: RoI Batch To Feature Vector

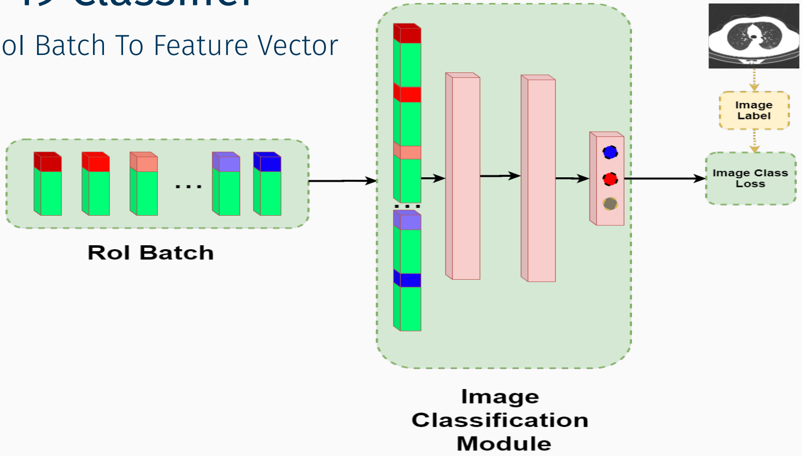

Figure 9: COVID-CT-Mask-Net. Conversion of the RoI batch to a feature vector.

COVID-19 Classifier

The model (RoI score$_\theta$=-0.01):

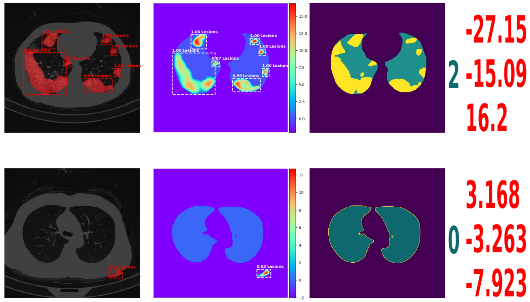

Figure 10: COVID-CT-Mask-Net. Output of the classification module.

COVID-19 Classifier

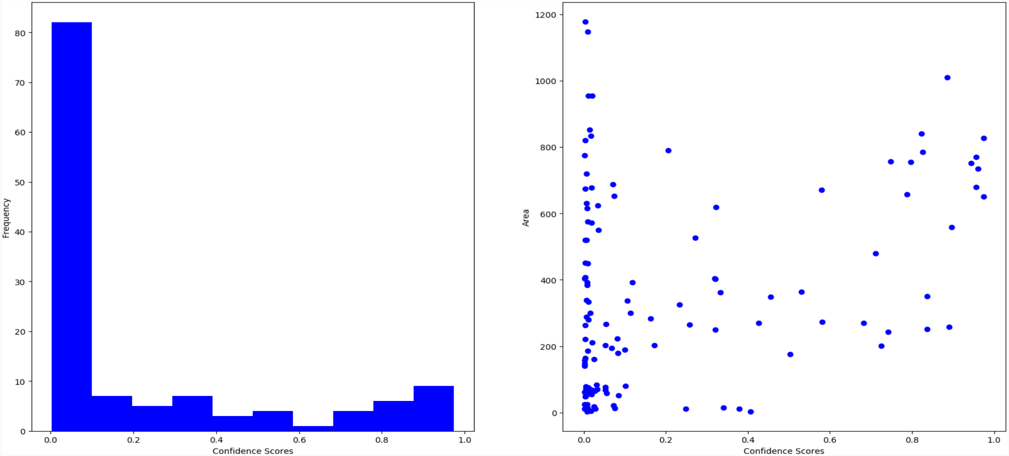

Figure 11: COVID-CT-Mask-Net. Histrogram + score vs size of the positive RoIs: COVID-19

COVID-19 Classifier

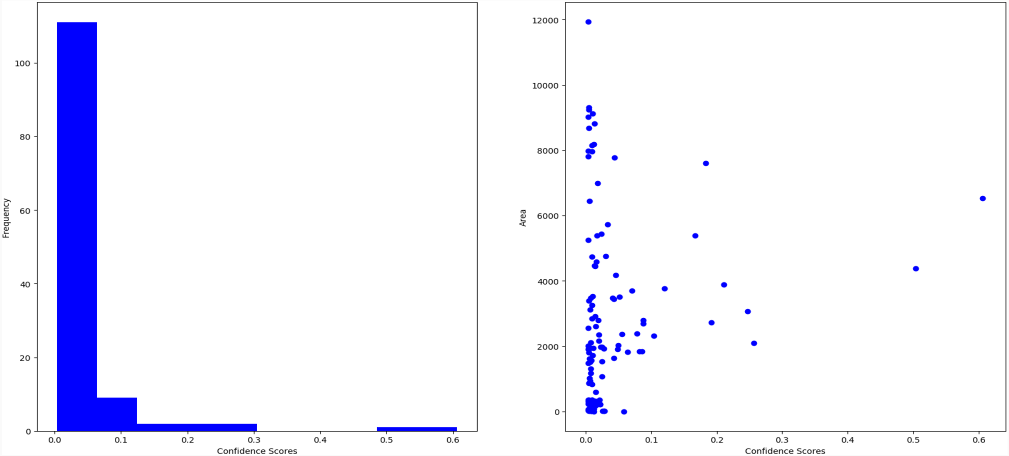

Figure 12: COVID-CT-Mask-Net. Histogram + score vs size of the positive RoIs: CP

COVID-19 Classifier

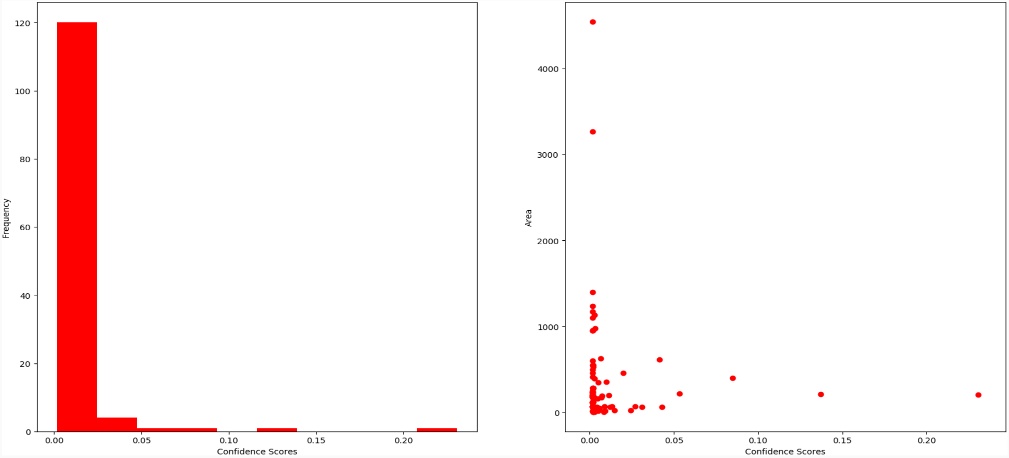

Figure 13: COVID-CT-Mask-Net. Histogram + score vs size of the positive RoIs: Control

Results

Results

Model	#Total parameters	#Trainable parameters	Training	Validation	Test	Ratio Test/Train
Mask R-CNN (ResNet50+FPN)	31.78M		600	150	-	-
Mask R-CNN (ResNet18+FPN)	14.52M		600	150	-	-
Mask R-CNN (ResNet18$_{T1}$+FPN)	6.12M				-	-
Mask R-CNN (ResNet18$_{T2}$+FPN)	4.02M				-	-
Mask R-CNN (ResNet34+FPN)	24.63M				-	-
Mask R-CNN (ResNet34$_{T1}$+FPN)	11.45M				-	-
Mask R-CNN (ResNet34$_{T2}$+FPN)	4.68M				-	-
COVID-CT-Mask-Net (ResNet50+FPN)	34.14M	2.36M	3K	20.6K	21.1K	7.06
COVID-CT-Mask-Net (ResNet18+FPN)	14.75M	0.6M	3K	20.6K	21.1K	7.06
COVID-CT-Mask-Net (ResNet18$_{T1}$+FPN)	6.35M					
COVID-CT-Mask-Net (ResNet18$_{T2}$+FPN)	4.25M					
COVID-CT-Mask-Net (ResNet34+FPN)	24.86M					
COVID-CT-Mask-Net (ResNet34$_{T_1}$+FPN)	11.74M					
COVID-CT-Mask-Net (ResNet34$_{T2}$+FPN)	4.92M					
COVIDNet-CT (best) [GWW20]	1.8M	1.8M	60K	20.6K	21.1K	0.353
COVNet [LQX$^+$20]	25.61M	25.61M	3K	370	438	0.129
ResNet18 [BGCB20]	11.69M	11.69M	528		90	0.17

Figure 14: Model sizes

Results

Model	AP@0.5	AP@0.75	mAP@[0.5:0.95:0.05]	Model size
Lightweight model (truncated ResNet34+FPN)	59.88%	45.06%	44.76%	11.45M
Lightweight model (truncated ResNet18+FPN)	49.95%	37.78%	39.32%	6.12M
Full model (merged masks)	61.92%	45.22%	44.68%	31.78M
Full model (GGO + C masks)	50.20%	41.98%	38.71%	31.78M

Figure 15: Accuracy of the Segmentation models (MS COCO criterion)

Results

Model	Control	CP	COVID-19	Overall accuracy
Lightweight model (truncated ResNet34+FPN)	92.89%	91.70%	91.76%	92.89%
Lightweight model (truncated ResNet18+FPN)	96.98%	91.63%	91.35%	93.95%
Full model (merged masks)	97.74%	96.69%	92.68%	96.33%
Full model (both masks)	96.91%	95.06%	93.88%	95.64%

Figure 16: Accuracy of the Classification models

Results

Summary of the advantages of COVID-CT-Mask-Net:

1. Requires a fraction of the training data to achieve a very high COVID-19 sensitivity and accuracy,
2. Trains a fraction of model parameters,
3. Generalizes well to the unseen data,
4. Understands COVID-19 and CP correlates explicitly,
5. Can be extended to do single-shot segmentation and detection.

Papers + Code + Data

Preprints on medRxiv


```
@article {Ter-Sarkisov2020.10.30.20223586,
	author = {Ter-Sarkisov, Aram},
	title = {Lightweight Model For The Prediction of COVID-19 Through The Detection And Segmentation
	of Lesions in Chest CT Scans},
	year = {2020},
	doi = {10.1101/2020.10.30.20223586},
	publisher = {Cold Spring Harbor Laboratory Press},
	journal = {medRxiv}
}

@article {Ter-Sarkisov2020.10.23.20218461,
	author = {Ter-Sarkisov, Aram},
	title = {Detection and Segmentation of Lesion Areas in Chest CT Scans For The Prediction of COVID-19},
	year = {2020},
	doi = {10.1101/2020.10.23.20218461},
	publisher = {Cold Spring Harbor Laboratory Press},
	journal = {medRxiv}
}

@article {Ter-Sarkisov2020.10.11.20211052,
	author = {Ter-Sarkisov, Aram},
	title = {COVID-CT-Mask-Net: Prediction of COVID-19 from CT Scans Using Regional Features},
	year = {2020},
	doi = {10.1101/2020.10.11.20211052},
	publisher = {Cold Spring Harbor Laboratory Press},
	journal = {medRxiv}
}
```

Github code

Models:`https://github.com/AlexTS1980/COVID-CT-Mask-Net`

Data: `http://ncov-ai.big.ac.cn`

Splits: `https://github.com/haydengunraj/COVIDNet-CT`