
MAVIDH SCORE: A COVID-19 SEVERITY SCORING USING INTERPRETABLE CHEST X-RAY PATHOLOGY FEATURES

A PREPRINT

Douglas P. S. Gomes
Anwaar Ulhaq
Manoranjan Paul
Tanmoy Debnath

Machine Vision and Digital Health (MAVIDH) group,
School of Computing and Mathematics
Charles Sturt University
NSW, Australia

Michael J. Horry
Subrata Chakraborty

Centre for Advanced Modelling and Geospatial
Information Systems (CAMGIS),
School of Information, Systems, and Modeling
Faculty of Engineering and IT
University of Technology
Sydney, NSW

D.M. Motiur Rahaman
National Wine and Grape Industry Centre
Charles Sturt University
NSW, Australia

Manash Saha
Wollongong Hospital
Wollongong, NSW

ABSTRACT

The application of computer vision for COVID-19 diagnosis is complex and challenging, given the risks associated with patient misclassifications. Arguably, the primary value of medical imaging for COVID-19 lies rather on patient prognosis. Radiological images can guide physicians assessing the severity of the disease, and a series of images from the same patient at different stages can help to gauge disease progression. Hence, a simple method based on lung-pathology features for scoring disease severity from Chest X-rays is proposed here. As the primary contribution, this method shows to be correlated to patient severity in different stages of disease progression comparatively well to other existing methods. An original approach for data selection is also proposed, allowing the simple model to learn the severity-related features. It is hypothesized that the resulting competitive performance presented here is related to the method being feature-based rather than reliant on lung involvement or compromise as others in the literature. The fact that it is simpler and interpretable than other end-to-end, more complex models, also sets aside this work. As the data set is small, bias-inducing artifacts that could lead to overfitting are minimized through an image normalization and lung segmentation step at the learning phase. A second contribution comes from the validation of the results, conceptualized as the scoring of patients groups from different stages of the disease. Besides performing such validation on an independent data set, the results were also compared with other proposed scoring methods in the literature. The expressive results show that although imaging alone is not sufficient for assessing severity as a whole, there is a strong correlation with the scoring system, termed as MAVIDH score, with patient outcome. The method has the potential to help physicians rating and following disease progression in COVID-19 patients, especially when making decisions on resource-scarce environments.

Keywords computer vision · COVID-19 · progression · severity score

1 Introduction

COVID-19 remains an imminent threat to society with many countries experiencing a severe second wave of the virus Giacomo et al. [2020]. One widely documented characteristic of COVID-19 infection is the wide range of symptoms experienced by infected persons ranging from entirely asymptomatic through admission to the general ward for a

range of symptoms including fever, cough, fatigue, headache and diarrhoea Huang et al. [2020] to severe pneumonia requiring admission to ICU with mechanical ventilation Guan et al. [2020]. Reported case-fatality rates vary from 1% to greater than 7%, usually due to respiratory failure Vincent and Taccone [2020]. Fatal cases tend to progress rapidly, particularly in the case of elderly patients where the average survival time after admission can be as low as five days Wang et al. [2020]. Given this wide range of patient symptoms and potentially rapid progress to death in severe cases, it is imperative that a patient's condition be objectively tracked for severity so that scarce critical care resources may be efficiently deployed to improve patient outcomes.

Whilst care in modern ICUs will result in death rates towards the lower end of the range, life-sustaining therapies will, in practice, be limited by lack of personnel, beds or materials and equipment. There is an emerging body of evidence showing a close association between greater mortality and overwhelmed healthcare infrastructures Ho and Neo [2020]. This limitation of resources leads clinicians to make prognostic decisions based on criterion such as old age, fragility and comorbidity that can lead to the death of patients with poor prognosis in favour of patients with better progression outlook Vincent and Taccone [2020]. Understanding patient prognosis is therefore critical to patient outcomes at both an individual and group level.

Pathological multivariate scoring has been shown to be predictive of COVID-19 patient admission to ICU and death Allenbach et al. [2020], Fan et al. [2020]. This approach requires the collection of pathology data points such as CRP levels, lymphocyte counts, platelet counts, interleukin, and procalcitonin levels. Although effective, this approach may not be practical in a triage situation, since the collection of necessary pathological data points is both resource-intensive and time-consuming. Having a more accessible indicator for the severity progression would therefore be desirable and potentially useful.

Chest medical imaging has proven to be useful in managing more serious COVID-19 infections since progressive respiratory failure caused by massive alveolar damage is the main source of COVID-19 mortality Xu et al. [2020]. In particular, Chest X-rays (CXR) and Computed Tomography (CT) imaging are useful tools in the management of moderate to severe COVID-19 cases since these methods help clinicians to establish a baseline pulmonary status and identify underlying pulmonary conditions that may contribute to the patients' risk, as well as assessing COVID-19 progression Rubin et al. [2020]. The CXR imaging mode has the advantage of being less expensive per scan in comparison to CT Flores et al. [2017] and available as portable apparatus that is easier to disinfect than CT equipment (which is typically fixed in a dedicated radiology room) Wong et al. [2020]. Although not as accurate as CT or Ultrasound Horry et al. [2020], the sensitivity of CXR imaging increases over the course of COVID-19 infection with serial CXR imaging, especially after day 6 of symptom onset Stephanie et al. [2020].

One approach to COVID-19 severity and progression scoring is the automated or semi-automated analysis of medical images taken over consecutive periods. Several such scoring techniques appear in the literature for both the CXR and CT imaging modes Wasilewski et al. [2020], Yang et al. [2020]. These techniques typically divide the medical image into geographical regions with each region manually assessed by radiological staff according to specific criteria. Semi-qualitative manual scoring of CXR has been shown to have prognostic value for COVID-19 progression in low, moderate/high and highly severe cases corresponding to scores 1, 3 and 4 Baratella et al. [2020], Borghesi and Maroldi [2020]. Notably, there is no indication that these methods are reliable for moderate cases outside of this score range.

Since the manual interpretation of medical images is a highly specialized and resource-intensive skill, many researchers have investigated the utility of deep learning systems in determining a COVID-19 severity score. Such systems have proven to be successful in quantifying COVID-19 lung compromise at a point in time Blain et al. [2020], Cohen et al. [2020a], sometimes with good correlation between results from the CXR and CT imaging modes Amer et al. [2020]. Automation of the semi-qualitative approach has been proposed with results that are promising but leave much room for improvement, in part due to the difficulty associated with establishing the ground truth to such methods due to human radiological interpretation Signoroni et al. [2020]. The use of deep learning can have limited applications in medical use cases when used as support for decisions that affect a patient's clinical outcome because of their non-explainable, black-box nature. One popular technique for interpretation on convolutional networks is the saliency map, which provides a heatmap overlay of network attention calculated on a gradient basis Simonyan and Zisserman [2014]. Nevertheless, there are still serious concerns as to whether saliency mapping techniques accurately reflect trained model parameters Adebayo et al. [2018] since they do not have shown to be robust under rigorous examination in the context of medial imaging Arun et al. [2020].

Given the issues of assessing disease progression and explainability, this work presents an interpretable, fully-automated, CXR severity scoring for COVID-19 named MAVIDH (Machine Vision and Digital Health research group) score. The method uses machine learning techniques to extract semantic features from the CXR and score COVID-19 patients based on a particular selection of ICU-admission data. The method is validated against an independent data set and through a comparison with other existing works. Such a comparison is made in a proposed framework comprised of images from patients in different stages of disease progression that allows judging the score based on the different

expected severity from particular groups. The analysis shows that there is a relationship between CXR features and patient stratification. It is envisaged that this research may lead to the development of analytic tooling that will help clinicians stratify COVID-19 patients to achieve better outcomes while being valuable in assisting clinics in managing resource requirements relating to ICUs.

2 Related work

The amount of papers proposing computer vision techniques and other similar efforts to assist the control of the COVID-19 pandemic has established a trend in research Ulhaq et al. [2020], Manna et al. [2020]. Due to the popularity and accessibility of deep learning algorithms for classification tasks, one can find several works proposing algorithms for COVID-19 diagnosis based on radiological images. Nevertheless, there is a constant concern by physicians and experts that such methods have little practical use due to the potential bias, uncertainty, and consequential risks of relying on such algorithms for a diagnosis Wynants et al. [2020], Bachtiger et al. [2020]. There are efforts to mitigate such problems like segmenting and aligning the lung images, which helps Tartaglione et al. [2020], Rabinovich et al. [2007] but does not solve the diagnostic concerns. Some therefore argue that the most plausible fair use of such technologies is in gauging disease progressions and severity Cohen et al. [2020b].

Regarding disease severity assessment, there have been a number of efforts to produce multivariate COVID-19 risk stratification scores Wynants et al. [2020]. Not all of these methods are computer-vision based, and some rely on potentially onerous data to estimate the disease severity. The access to medical imaging such as CXR and CT, represents an alternate method of COVID-19 patient risk stratification, which is to use medical images to quantify lung abnormalities. X-rays, in particular, are flexible and inexpensive technology that could help physicians assess disease progression and severity. A particularly popular approach for using CXR in disease severity is the Brixia score Borghesi and Maroldi [2020], which is a method that uses a heuristic scoring system to quantify a severity score for COVID-19 pneumonia. It can be seen as a semi-qualitative assessment of lung disease by ranking pulmonary involvement over upper, middle and lower zones per lung on an 18 point severity scale. The Brixia score is an example of what will be referred to here as a lung-involvement score since it tries to estimate to what extent the lungs are compromised by pathologies like lesions and opacities.

Works using Brixia-like scores in both automatic Amer et al. [2020], Signoroni et al. [2020], Cohen et al. [2020a] and non-automatic Allenbach et al. [2020], Baratella et al. [2020] approaches can be found in the literature. The intricate end-to-end methods such as Amer et al. [2020] names their score as ‘Pneumonia Ratio’, for example; others like Cohen et al. [2020b], Blain et al. [2020] base its severity score on features named Opacity and Geographic extension. These methods seem to have some predictive value, but most rely on the extension of lung involvement and scores given by human experts, which they try to regress. In this paper, however, a hypothesis tested is that a method that does not rely on lung involvement and expert labelled data but rather on features from known lung pathologies and data selection from patients in different stages of the disease can work well at tracking disease progression. This testing is performed by comparing the results of such a method with other existing works. However, the comparison is only possible with other methods that fit the same definition as this one: automatic, computer-vision based, using CXRs to produce the severity score, and that also had code available to be implemented in the same comparable framework. To those requirements, only two other works were found fitting such description,

Cohen et al. Cohen et al. [2020a] employed three blinded radiologists to stage disease severity using the extent of lung involvement and degree of opacity to establish ground truth. Then, a deep convolutional neural network, trained on a number of public datasets covering 18 common radiological findings, was used as a feature extraction layer for COVID-19 CXR images. The features were then sequentially connected to a linear regression layer to predict the extent of lung involvement and degree of opacity to labels from experts as ground truth. The results showed relative fair measures for the regression with a mean absolute error of 1.14 for the geographic extent score and only 0.78 for the lung opacity score. Still, no validation on patients from different stages of disease progression was performed.

Taking a different approach to Cohen, Signoroni et al. Signoroni et al. [2020] used a large clinical dataset of 5000 CXR images to train a deep learning-based implementation of the semi-quantitative Brixia score. The results showed that it achieved an accuracy equivalent, or better than human radiologists with arguably greater consistency. After segmenting and aligning lung fields from source CXR images, a preprocessing pipeline was implemented to equalize and denoise the dataset before using these images to train variations of a ResNet-18 based CNN (BS-Net). The network showed comparable performance against an independent COVID-19 dataset in portability studies indicating potential usefulness in other clinical settings.

The work proposed here has comparably slightly less complexity than Signoroni et al. [2020], which is fully end-to-end, but it is a bit more intricate than Cohen et al. [2020a], which does not use lung segmentation as a preprocessing step. The few degrees of lesser complexity allows the method presented here to be more interpretable than more complex

alternatives as it uses a specialized CNN to extract features but, as in Cohen et al. [2020a], uses such features in simpler but potentially more robust and explainable learning methods.

3 Data and methods

3.1 Data sets

The main aspects of this work are the two data sets of X-rays images from different sources and the original methodology for severity scoring based on low-complexity, semi-interpretable, machine learning. The critical data set for assessing the hypotheses presented here was recently proposed and contained a set of multiple-instances images of patients at different stages of the disease Winther et al. [2020]. By having samples of images at different moments of disease progression from the same patient, this data set can be seen as a source of potential insights regarding the role of X-rays in assessing disease severity. Differently from most published in the literature up to date, this data set has rich metadata containing the distance (in days) from when the image was taken to the hospital and ICU admission. It is the most detailed data in this respect, to the best of our knowledge. Given the authors' affiliation to Hannover Medical School in Germany, this data set will be referred here as 'Hanno'. At no stage, this data set was used to learn the methods presented here; the Hanno data set was used only for validating the proposed hypotheses and comparing results to other methods. The other data set, used for learning the features relevant to the disease severity, is the popular set of X-ray images by Cohen et al. Cohen et al. [2020b], referred here by the first author's name. The Cohen data set has been used in many works in the literature focused on disease diagnosis, but much less in progression and severity scoring. The data set has limited metadata on the ICU admission in regards to the moment in time of when the image was taken but has some potentially relevant categorical classes concerning the stage and patient severity.

In short, the Hanno data set comes with 234 images, but the quality is much varied. Cleansing is performed on images where the lung field is not clearly visible, and a significant part is excluded. In total, 154 images resulted from such a selection. From these, 54 are images from patients that were not admitted in ICU, while the remaining 100 images are from patients taken in different stages of disease progression. This set of images is used in the same conditions for all the comparison to other methods and validation. Although it usually only contains a few images from the same patient, the rich metadata with the offset in days from key moments of disease progression can be used to group patients in defined periods of progression so an original validation approach can be adopted; more on the grouping method in the following subsections.

The Cohen data set is a somewhat popular set of X-Ray images, first announced in April 2020. Although it was first described in March 2020, it has more than 200 citations, 2000 favourites in Github, and many implementations. A significant part of citing works use this data set to address the diagnosis problem. Nevertheless, the diagnosis of COVID through imaging is a delicate problem, and its solution should present high robustness to risk in order to be useful. A number of authors have written on the challenges of such a use Arun et al. [2020], DeGrave et al. [2020], Wynants et al. [2020] with some pointing that probably the most promising use of X-rays would be at assessing disease severity and progression in a prognostic approach Cohen et al. [2020b], Manna et al. [2020]. With this in mind, this methodology processes a different use of the data set, where classes of patients with different ICU-related outcomes are adopted. By using this approach, it is hoped that the learned model will learn a reasonable probability estimator of patient outcome that could be used as a score. In all, after selecting the images with reasonable quality where the lungs and its features were visible and fitted our particular criteria for the patient stage, 100 images were selected. These were previously validated by the authors of this paper in a work describing how a similar approach could point to potentially relevant semantic features correlated to ICU outcome Gomes et al. [2020].

Regarding model selection, it is worth further commenting on the major factors to the methodological decisions proposed here, which is to avoid overfitting and prioritize interpretability. The main reasons for that are two-fold: (1) the data sets used are limited, and one should therefore minimize the risk of overfitting, and (2) differently from methods using end-to-end deep learning for X-ray analyses, interpretability can be much valuable in subjects like healthcare. The over-cautious care for overfitting leads to the choice of low-complexity methods (low Vapnik–Chervonenkis dimension), which are also usually more interpretable. The feature extraction role is performed by a deep, pre-trained specialized CNN, but the regression and classification results are given by a simple logistic regression model. Such a specialized CNN was trained to detect common lung pathologies with a combination of large data sets of tens of thousands of X-ray images. Its outputs are features learned to have semantic meaning, i.e., each of them can be seen as probability scores of different lung pathologies. It is thus more interpretable in the sense that the features are real pathology scores, and one can see how much weight is given to each feature in the subsequent severity scoring. This construction relates deeply with one of the hypotheses discussed here, which is that features representing specific pathologies in the lung may correlate better with the severity of symptoms than methods that estimate lung compromise, thus resulting in a




Figure 1: Training and validation curves from the lung segmentation network — Skip U-Net.

more realistic scoring. Since the feature extraction method is performed by such a deep, specialized network, one can say that it does fit into the learning paradigm known in the literature as transfer learning Shermin et al. [2019].

3.2 Feature extraction pipeline

Before extracting the features with such specialized CNN, the selected images are fed through a preprocessing pipeline, which has the goal of normalizing and reducing potential artifacts that could lead to biases in learning. Such a pipeline comprises sequential blocks with defined tasks: histogram equalization, lung segmentation, and cropping of the lung area. All these are done automatically with processes that share the same parameters for all images. Histogram equalization is the main normalization, and it is applied at the beginning and end of the pipeline. At the input, the normalization is standard, and it equalizes the pixel distribution so they are not concentrated in a small range, which could induce bias. Before the output, an image normalization technique called Contrast Limited Adaptive Histogram Equalization (CLAHE) is applied, which could be seen as a local histogram normalization. This technique is used to highlight the small local pathology features, so their contrast could be improved.

The lung segmentation step is subsequently applied to the images in the preprocessing pipeline. As data sets are limited, and there are many bias-inducing features in the images like bones, medical devices, and letters, segmenting the lungs can be seen as increasing the signal-to-noise ratio for the learning algorithm. This practice is corroborated in previous studies Rabinovich et al. [2007] in the literature where one can also find comprehensive reviews on lung area segmentation Candemir and Antani [2019]. Earlier works on lung segmentation presented dice similarity scores up to 0.989 Hwang and Park [2017]; a much simpler approach using the U-Net architecture was also proposed, achieving 0.974 Ronneberger et al. [2015]. The latter was trained on the JSRT dataset Shiraishi et al. [2000] consisting of 385 CXR images with gold-standard masks. In this work, we present a deep learning-based lung segmentation that surpasses such scores.

The lung segmentation network learned here is a variation of the U-Net architecture using skip connections with a ResNet backbone He et al. [2016]. This adoption is in contrast to the U-Net trained on simple convolutional stacks with max pooling and VGG architecture Simonyan and Zisserman [2014]. The authors hypothesized that the skip connections allow the features in earlier layers to be reused, thereby increasing the performance of the segmentation. Such a network design achieved a maximum validation dice similarity coefficient of 0.988 at epoch 93, as illustrated in Fig. 1, which depicts its learning curves. The performance boost was probably also achieved given the use of additional data sets adding up to 1185 CXRs samples Jaeger et al. [2014].

The pipeline then ends with lung-area cropping, which is a combination of closing of the masks and cropping any pixels outside their area. The closing of the masks helps to get rid of any artifacts escaping the lung segmentation network, and it is performed by a combination of morphological closing, contour and flood filling. The cropping has a desirable effect of normalizing the images for lung size; the lung area sometimes occupies only part of the image, either by the size the extent that the image covers of a patient or by factors like the size of the person or age. Examples of images in different stages of the preprocessing pipeline are illustrated in Fig. 2.




Figure 2: Chest X-Ray images in different stages of the pre-processing pipeline. a) Original images with different contrast and histogram. b) Segmented lungs without the normalization techniques. c) Images with lungs segmented and normalization applied.

3.3 Learning and validating the progression score

The methodology framework can be divided into two parts: learning and validation. The fact that each one of these parts is performed with data sets from different sources helps to attest the method’s ability to generalise. However, a meaningful result can only be achieved if overfitting can be limited, especially given the number of samples in the Cohen data set. Regarding the use of such a data set in the learning stage, one of the original ideas presented here is that, different from other methods that use it in a classification approach with ‘COVID’ and ‘Non-COVID’ classes, this methodology filters the metadata in particular classes: patients that were admitted to ICU but had the image taken before admission (‘future icu’) and patients that recovered without intensive care (‘not icu’). This filtering is particularly interesting because the first class is populated by samples of patients that, despite not being present in ICU at the time the image was taken, were eventually later admitted. Their images could potentially present insightful information of predecessor features to disease severity.

At the learning stage, the methodology using the semantic features for progression scoring is performed by first extracting the semantic features (specialised CNN), standardising them, and fitting a logistic regression model to classify the data in one of two mentioned classes (‘future icu’ and ‘not icu’). If the accuracy of the classifier showed it to have relevant invariance to the classes, one can hypothesise its estimated class probability (output from the logistic regression) could be correlated to the patient’s severity state. That would not be necessarily surprising since such the classifier was learned to attempt to predict if a particular patient, given their X-ray, will eventually be admitted to ICU or not. However, it is also not surprising that one would be sceptical of the validity of such a progression score. This potential concern is the reason for also proposing an original validation approach and comparison to other methods in the literature.

The original validation approach proposed here comprises two parts: (1) a heuristic method for grouping scores from patients’ images taken in similar stages of symptom progression, and (2) a comparison of the group scoring to other methods in the literature by implementing their code in the Hanno data set. It is worth noting that the grouping in the first step is only possible because, differently from Cohen (used for training), the Hanno data set contains rich metadata on the offset (number of days) from specific moments in the disease progression. The proposed heuristics here is to group X-ray images from 4 distinct periods or classes of patients:

- Group 1: Images from patients not admitted to ICU.
- Group 2: Images from patients in the vicinity (between 1 day before or after) of ICU admission.




Figure 3: Graphical abstract of the methodology. Images taken from the Cohen data set are particularly selected, processed and used for learning a severity scoring model. The images from the Hanno data set are processed and scored by the learned method and by other two from the literature. The assessment of the results is performed via a grouping method also proposed here.

- Group 3: Images from patients currently in ICU (between 1 day past ICU admission and 1 day before ICU release).
- Group 4: Images taken in the vicinity (in between 1 day before or after) of ICU release.


By the definition of these groups, one can hypothesise that despite the actual value of the score, it should respect a particular trend between groups to have at least minimal significance. For example, Group 1 and 4 should have the lowest scores on average, with patients that did not go to ICU (Group 1), being even lower than patients in the vicinity of ICU release (Group 4). Moreover, one should expect that Group 2, populated by images of patients in the vicinity of ICU admission, to have the high scores on average, as well for patients in Group 3 (currently in ICU).

Although following the expected trend is not definite evidence of the score validity, it indeed corroborates the hypothesis that it is any indicator correlated to disease severity. The contribution is reinforced if such a scoring method presents yet competitive results to other methods in the related literature. This point should be highlighted by noting that the data sets for training and validation come from different distributions, and any resulting validation should be more dependable than trivial train-test splits on the same data set. It is also important to note that the comparison to other methods was only possible because the authors made their code and neural network weights available. This initiative is praiseworthy and should be incentivised since it can increase the progress rate in such an urgent field. To better clarify how the pieces of the methodology work together, Fig. 3 presents the synthesised methodology in diagram form.


4 Results

4.1 Observational findings


The close inspection of the data sets given by the experience from developing this methodology generated insights on the potential role of X-rays to score disease severity. Although it is a somewhat trivial notion that X-rays will not translate all the information to a definite score the disease progression, some discrete insightful examples can illustrate



(a) Example of a case where the images loosely correlate to the development of the disease.



(b) Example of a case where the quality of the images is counter-intuitive to the actual disease progression.



(c) Example of a case where the images taken in very distinct moments do not show evident visual differences.

Figure 4: Noticeable cases of images taken from patients in distinct moments of their disease progression.

some of the limitations; the authors have not encountered such observations in other works, especially regarding the X-ray diagnosis-related literature. These honest illustrations are not necessarily unexpected since other comorbidities and patients' conditions will significantly affect the severity of their symptoms. Some notable cases are depicted in Fig. 4 as examples that correspond (Fig. 4a) or not (Fig. 4b-4c) to the expected stages of the disease.

Given such findings, the question then becomes to what extent X-rays can track the severity in the disease progression. This question, as the primary motivation of this paper, is important because although other methods in the literature rate the severity of lung damage with features like opacity and geographic extension, they do not attest if such scores actually correlate to the progression of the patients' symptoms or need for intensive care. The following results from grouping the patients in similar stages and comparing their score is set to address this issue.

4.2 Scoring

As a preamble to the results, it is worth noting a crucial aspect that makes this approach original, which is the data selection for learning. The scoring method itself is simple and direct, but as in every other problem, data is still paramount; model complexity is inconsequential when the labelled data does not present the information one needs. As it will be soon demonstrated, having a simpler but interpretable method also has considerable advantages. The main difference in the methodology adopted in data selection is the fact that images were not classified in 'COVID' vs. 'non-COVID', for example. They also did not have the lesions previously visually rated by humans in an attempt to model the intensity of score of lung damage. The hypothesis here is that some features may be more important than others in predicting severity than the overall area of involvement and that a specialised neural network could capture such information. The samples were therefore classified in two classes: images taken from patients previously to their




Figure 5: Confusion matrix from the training data set, Cohen (a), and from the validation data set, Hanno (b).

ICU admission, and patients that were not admitted to ICU but were symptomatic enough to have their images taken. Both classes being patients with COVID. A classifier trained on these two binary classes will approximate the task of rating the probability that a patient will end up needing intensive care or not.

4.3 Separability

Before using the selected classes as a scoring model, one should first attest if such particular classes are separable. Given the limited data set size, one should take extreme care to constrain overfitting and thus choose a simple classifier not simply to shatter the data set. Since a simpler classifier means a simpler set of admissible functions relying on fewer features, they have the advantage of being more robust with a stronger chance of generalising. As previously mentioned for scoring, attesting the separability between classes will also be done by using the Cohen data for fitting the classifier and the external (Hanno) data set for validation.

The classifier chosen, given the model selection criteria mentioned above, was a shallow and limited decision tree. The only parameters changed from the scikit-learn library Pedregosa et al. [2011] defaults were the maximum depth of the tree to 3, and the minimum leaves in a node to 10. This setting aggressively limits the ability of the tree to branch into nodes based only on a few observations, consequently limiting overfitting. In fact, the following separability figures resulted from a decision tree with only 5 nodes (5 from the 18 features). The in-depth work of using shallow decision trees for analysing the semantic features relevant for classification was done in a previous work by the authors Gomes et al. [2020] and will thus not be discussed in detail here.

Fitting the shallow decision tree led to a classification accuracy of 82% on the training data (Cohen). For cross-validation in a leave-two-out scenario, the resulting accuracy was 73%. Note that although being from the same data set, the accuracy resulted from cross-validation does not include data used in training. It should be noted that validating on the Hanno data set it is not done in the strict sense. This data set does not have metadata on if the patient was eventually admitted in ICU in the future as does Cohen, but it has labels on which images are from patients in ICU or not. Such validation is then more related to the ability to detect severity than predicting the outcome. Having noted that, the resulting accuracy on the Hanno data set was 77%. Both confusion matrices, from validating on the Cohen and Hanno data sets are illustrated in Fig. 5. The accuracies and confusion matrices show that the data are not entirely separable, but it also indicates that there is a correlation between these 5 features ('Effusion', 'Consolidation', 'Pneumonia', 'Fracture', 'Pleural Thickening') and symptom severity.

4.4 MAVIDH score assessment

As stated in the methodology, the assessment of the MAVIDH score model was done by creating a method for grouping the patients in specific moments of symptom progression. As the score only approximately correlates with the patient severity, the grouping method helps to assess the correlation to the expected trend by comparing different statistical metrics from each group. In total, there were 57 images from patients not admitted in ICU, 37 in the vicinity of being admitted, 112 images of patients in ICU, and 31 in the vicinity of being released.

Trained with the particular proposed labels, the binary regression model outputs a probability given by the logistic function, which is used here as the MAVIDH severity score. Since every image has a respective score, the box and whisker plot was chosen to illustrate the different statistical metrics from each of the groups in Fig. 6. The box plot




Figure 6: Group score box plot of the method proposed here (MAVIDH score).

presents some worth mentioning points regarding the agreement to the expected trend, like the fact that the upper quartile of patients that were not in ICU (0.22) was smaller than all the other groups (0.345, 0.370, and 0.324). The fact that patients in the vicinity of ICU admission have more spread and higher median score (0.496) is also somewhat expected as these patients should be close to or already presenting severe symptoms. In the group of patients in ICU, the interquartile range consolidates at a higher level and tighter range (0.203) as compared to others, showing that severity is more consistent between images. The patients in the vicinity of being released have a slightly smaller median than other ICU-related groups and much less distance between the upper quartile and maximum value than other groups.

4.5 Comparison with other methods

The method presented by Cohen et al. [2020a] has an interesting similarity with the one presented here, which is that both use the same feature extraction network. However, as Signoroni et al. [2020], it tries to regress the score based on expert labelled data. In the case of Cohen et al. [2020a], the task is generate scores based on opacity and geographic extension labels rather than tracking patient severity. It should be noted that the following comparisons and commentary do not concern the method’s quality or ability to assess these features; they are rather related to their correlation to patient severity in the groups designed here.

Both the opacity and geographic extension were calculated for all images in the Hanno data set using the author’s code. The opacity feature range between 0 to 6, while the geographic extension from 0 to 8. As seen in the Fig. 7, they do not necessarily track the progression in an expected trend. Differently from the method proposed here, the range of score for ‘non-ICU’ patients are much less concentrated at a lower level and span most of the range of other groups. Moreover, the group that has the highest and consolidated score is the one of patients in the vicinity of release, which is contrary to the expected trend.

The second method to be compared presents a strong case for severity scoring. As the method being proposed here, it integrates lung segmentation but also performs alignment and rotational invariance. It is not, however, trained with severity labels but with data labelled by humans on lung comprise. The scores presented in Fig. 8 were calculated by using the code and network weights kindly shared by the authors. Resulted from trying to regress scores on sections of the lung in the Brixia-score framework Borghesi and Maroldi [2020], the output score ranges from 0 to 18, as it rates six different parts in scores from 0 to 3. As illustrated by the resulting box-plot scores, it is somewhat consistent at tracking the expected progression trend. Some differences from the score results from the work here can be noticed. For one, the fact that for patients not admitted to ICU, the scores are not all concentrated at a lower level. The upper quartile is not smaller than all the other groups’ lower quartile, and the consolidation of score happens in the group in the vicinity of admission rather than in the group confirmed to be in ICU. However, the scores do present a significant trend, which in the same direction as this work, attest the existence of some significance of X-ray severity scoring to help assess patient symptom severity. The authors of this scoring method should be praised since their results show some significance at tracking patient severity progression, even though that was not exactly the learning task. Nevertheless, it must be noted that this method is entirely end-to-end with interpretability reliant only on gradient maps.

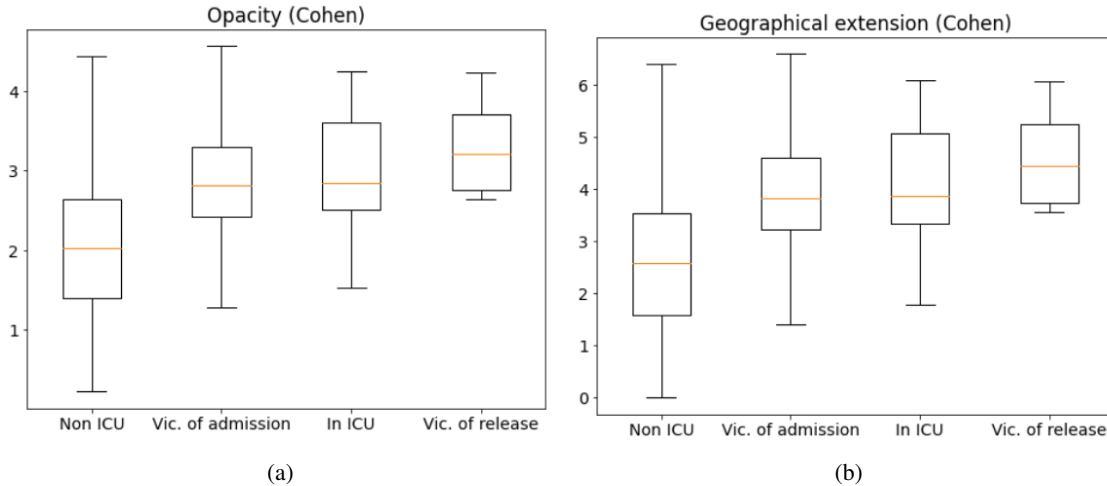


Figure 7: Group score box plot of the a) Opacity and b) Geographical extension features from Cohen et al. [2020a].





Figure 8: Group score box plot from the method presented in Signoroni et al. [2020].

4.6 Examples of scoring in the validation data set


As a result of applying the scoring method presented here, Figs. 9-11 illustrate the scores from a particular patient through time with their segmented X-rays. The horizontal dimension of the images represent the days from the hospital admission, and a dashed line is plotted through the scores to show tendencies. Most patients had only a few images; the following illustrations are some exceptions.

The first, illustrated in Fig. 9, was one of the few patients that had the first X-ray taken multiple days before the hospital admission. The scores and X-rays show that signs of severity were present much before the ICU admission, and it improved after intensive care. The second, shows a score getting progressively worse for a patient that was in ICU for many days. As shown in Fig. 10, the last X-ray image from this patient was taken at yet 23 days from ICU release, meaning that the patient was still in severe conditions and needed more time to recover. Lastly, Fig 11 shows an example of a relatively ambiguous result where similar images appear to have distant scores, attesting for the limitations of the method.

As the grouping analysis images have shown, despite the method chosen, there is probably much space for improvement in a potential COVID-19 disease severity scoring. The goal here was to show evidence that there is a probable relevant correlation in the information present in X-ray images to the severity of COVID patients rather than claiming an optimal score. It must be noted that such results were achieved with very limited data sets, and the fact that it was validated through different data sets should add to its expressivity. The authors hope to leverage better and bigger data sets to continue the investigation on the important features for disease severity that could help physicians to analyse medical imaging while assessing the patient's progression.




(a)




(b)

Figure 9: Patient no. 4 scoring a), and segmented X-rays b).



(a)



(b)

Figure 10: Patient no. 9 scoring a), and segmented X-rays b).





Figure 11: Patient no. 13 scoring a), and segmented X-rays b).

5 Conclusions

The present work has focused on presenting a feature-based, semi-interpretable, COVID-19 disease severity scoring and comparing its significance to other methods in the literature. The scoring method comprises a feature extraction pipeline with image normalisation, lung segmentation, and feature extraction by a specialised network trained to extract semantic features related to lung pathologies. The features were used in a logistic regression model to output the estimated probability of a patient going severe. One notable contribution was the data selection for such learning, which is filtered and labelled in a specific way so the model can learn the severity-related information. Such comparison is performed through a grouping methodology applied on a data set (Hanno) with metadata regarding the hospital and ICU admission offset. The analysis showed that the MAVIDH score proposed here had advantages at tracking the expected trend of the conceptualised groups. This is not to say that other methods are deficient since they were not trained for this task but rather to regress overall rate lung compromise. The results here attest the existence of a correlation between the developed score with patient severity through the disease progression. Although such a score does not perfectly correlate to severity, seen by the variance in the groups, the authors believe that it is a notable result given the limited data and the fact that presents comparable results to other more complex methods in the literature. It comprises a sensible, simple, and reasonably robust severity assessment support for a much urgent problem. The authors hope to improve this solution further with the availability of more high-quality and detailed labelled data in the near future.

References

- Cacciapaglia Giacomo, Cot Corentin, and Francesco Sannino. Second wave covid-19 pandemics in europe: a temporal playbook. *Scientific Reports (Nature Publisher Group)*, 10(1), 2020.
- Chaolin Huang, Yeming Wang, Xingwang Li, Lili Ren, Jianping Zhao, Yi Hu, Li Zhang, Guohui Fan, Jiuyang Xu, Xiaoying Gu, et al. Clinical features of patients infected with 2019 novel coronavirus in wuhan, china. *The lancet*, 395(10223):497–506, 2020.

- Wei-jie Guan, Zheng-yi Ni, Yu Hu, Wen-hua Liang, Chun-quan Ou, Jian-xing He, Lei Liu, Hong Shan, Chun-liang Lei, David SC Hui, et al. Clinical characteristics of coronavirus disease 2019 in china. *New England journal of medicine*, 382(18):1708–1720, 2020.
- Jean-Louis Vincent and Fabio S Taccone. Understanding pathways to death in patients with covid-19. *The Lancet Respiratory Medicine*, 8(5):430–432, 2020.
- Lang Wang, Wenbo He, Xiaomei Yu, Dalong Hu, Mingwei Bao, Huafen Liu, Jiali Zhou, and Hong Jiang. Coronavirus disease 2019 in elderly patients: Characteristics and prognostic factors based on 4-week follow-up. *Journal of Infection*, 2020.
- Esther Peiyang Ho and Han-Yee Neo. Covid 19: prioritise autonomy, beneficence and conversations before score-based triage. *Age and Ageing*, 2020.
- Yves Allenbach, David Saadoun, Georgina Maalouf, Matheus Vieira, Alexandra Hellio, Jacques Boddart, Helene Gros, Joe Elie Salem, Matthieu Resche-Rigon, Lucie Biard, et al. Multivariable prediction model of intensive care unit transfer and death: a french prospective cohort study of covid-19 patients. *medRxiv*, 2020.
- Guohui Fan, Chao Tu, Fei Zhou, Zhibo Liu, Yeming Wang, Bin Song, Xiaoying Gu, Yimin Wang, Yuan Wei, Hui Li, et al. Comparison of severity scores for covid-19 patients with pneumonia: a retrospective study. *European Respiratory Journal*, 56(3), 2020.
- Zhe Xu, Lei Shi, Yijin Wang, Jiyuan Zhang, Lei Huang, Chao Zhang, Shuhong Liu, Peng Zhao, Hongxia Liu, Li Zhu, et al. Pathological findings of covid-19 associated with acute respiratory distress syndrome. *The Lancet respiratory medicine*, 8(4):420–422, 2020.
- Geoffrey D Rubin, Christopher J Ryerson, Linda B Haramati, Nicola Sverzellati, Jeffrey P Kanne, Suhail Raoof, Neil W Schluger, Annalisa Volpi, Jae-Joon Yim, Ian BK Martin, et al. The role of chest imaging in patient management during the covid-19 pandemic: a multinational consensus statement from the fleischner society. *Chest*, 2020.
- John Paul Flores, Alejandro Moreno-Koehler, Matthew Finkelman, Jaime Caro, and Gary Strauss. Ma03.05 cost effectiveness analysis of ct vs chest x-ray (cxr) vs no screening for lung cancer (lc) in the plco and nlst randomized population trials (rpts). *Journal of Thoracic Oncology*, 12(1):S354–S355, 2017.
- Ho Yuen Frank Wong, Hiu Yin Sonia Lam, Ambrose Ho-Tung Fong, Siu Ting Leung, Thomas Wing-Yan Chin, Christine Shing Yen Lo, Macy Mei-Sze Lui, Jonan Chun Yin Lee, Keith Wan-Hang Chiu, Tom Chung, et al. Frequency and distribution of chest radiographic findings in covid-19 positive patients. *Radiology*, page 201160, 2020.
- M. J. Horry, S. Chakraborty, M. Paul, A. Ulhaq, B. Pradhan, M. Saha, and N. Shukla. Covid-19 detection through transfer learning using multimodal imaging data. *IEEE Access*, 8:149808–149824, 2020. doi:10.1109/ACCESS.2020.3016780.
- Stephanie Stephanie, Thomas Shum, Heather Cleveland, Suryanarayana R Challa, Allison Herring, Francine L Jacobson, Hiroto Hatabu, Suzanne C Byrne, Kumar Shashi, Tetsuro Araki, et al. Determinants of chest x-ray sensitivity for covid-19: A multi-institutional study in the united states. *Radiology: Cardiothoracic Imaging*, 2(5):e200337, 2020.
- Piotr G Wasilewski, Bartosz Mruk, Samuel Mazur, Gabriela Póltorak-Szymczak, Katarzyna Sklinda, and Jerzy Walecki. Covid-19 severity scoring systems in radiological imaging—a review. *Polish Journal of Radiology*, 85:e361, 2020.
- Ran Yang, Xiang Li, Huan Liu, Yanling Zhen, Xianxiang Zhang, Qiuxia Xiong, Yong Luo, Cailiang Gao, and Wenbing Zeng. Chest ct severity score: an imaging tool for assessing severe covid-19. *Radiology: Cardiothoracic Imaging*, 2(2):e200047, 2020.
- Elisa Baratella, Paola Crivelli, Cristina Marrocchio, Alessandro Marco Bozzato, Andrea De Vito, Giordano Madeddu, Laura Saderi, Marco Confalonieri, Luigi Tenaglia, and Maria Assunta Cova. Severity of lung involvement on chest x-rays in sars-coronavirus-2 infected patients as a possible tool to predict clinical progression: an observational retrospective analysis of the relationship between radiological, clinical, and laboratory data. *Jornal Brasileiro de Pneumologia*, 46(5), 2020.
- Andrea Borghesi and Roberto Maroldi. Covid-19 outbreak in italy: experimental chest x-ray scoring system for quantifying and monitoring disease progression. *La radiologia medica*, page 1, 2020.
- Maxime Blain, Michael T Kassin, Nicole Varble, Xiaosong Wang, Ziyue Xu, Daguang Xu, Gianpaolo Carrafello, Valentina Vespro, Elvira Stellato, Anna Maria Ierardi, et al. Determination of disease severity in covid-19 patients using deep learning in chest x-ray images. *Diagn Interv Radiol*, 2020.
- Joseph Paul Cohen, Lan Dao, Paul Morrison, Karsten Roth, Yoshua Bengio, Beiyi Shen, Almas Abbasi, Mahsa Hoshmand-Kochi, Marzyeh Ghassemi, Haifang Li, et al. Predicting covid-19 pneumonia severity on chest x-ray with deep learning. *arXiv preprint arXiv:2005.11856*, 2020a.

- Rula Amer, Maayan Frid-Adar, Ophir Gozes, Jannette Nassar, and Hayit Greenspan. Covid-19 in cxr: from detection and severity scoring to patient disease monitoring. *arXiv preprint arXiv:2008.02150*, 2020.
- Alberto Signoroni, Mattia Savardi, Sergio Benini, Nicola Adami, Riccardo Leonardi, Paolo Gibellini, Filippo Vaccher, Marco Ravanelli, Andrea Borghesi, Roberto Maroldi, et al. End-to-end learning for semiquantitative rating of covid-19 severity on chest x-rays. *arXiv preprint arXiv:2006.04603*, 2020.
- Karen Simonyan and Andrew Zisserman. Very deep convolutional networks for large-scale image recognition. *arXiv preprint arXiv:1409.1556*, 2014.
- Julius Adebayo, Justin Gilmer, Michael Muelly, Ian Goodfellow, Moritz Hardt, and Been Kim. Sanity checks for saliency maps. In *Advances in Neural Information Processing Systems*, pages 9505–9515, 2018.
- Nishanth Arun, Nathan Gaw, Praveer Singh, Ken Chang, Mehak Aggarwal, Bryan Chen, Katharina Hoebel, Sharut Gupta, Jay Patel, Mishka Gidwani, et al. Assessing the (un) trustworthiness of saliency maps for localizing abnormalities in medical imaging. *arXiv preprint arXiv:2008.02766*, 2020.
- Anwaar Ulhaq, Jannis Born, Asim Khan, Douglas Pinto Sampaio Gomes, Subrata Chakraborty, and Manoranjan Paul. Covid-19 control by computer vision approaches: A survey. *IEEE Access*, 8:179437–179456, 2020.
- Sayan Manna, Jill Wruble, Samuel Z Maron, Danielle Toussie, Nicholas Voutsinas, Mark Finkelstein, Mario A Cedillo, Jamie Diamond, Corey Eber, Adam Jacobi, et al. Covid-19: A multimodality review of radiologic techniques, clinical utility, and imaging features. *Radiology: Cardiothoracic Imaging*, 2(3):e200210, 2020.
- Laure Wynants, Ben Van Calster, Gary S Collins, Richard D Riley, Georg Heinze, Ewoud Schuit, Marc MJ Bonten, Darren L Dahly, Johanna AA Damen, Thomas PA Debray, et al. Prediction models for diagnosis and prognosis of covid-19: systematic review and critical appraisal. *bmj*, 369, 2020.
- Patrik Bachtiger, Nicholas S Peters, and Simon LF Walsh. Machine learning for covid-19—asking the right questions. *The Lancet Digital Health*, 2(8):e391–e392, 2020.
- Enzo Tartaglione, Carlo Alberto Barbano, Claudio Berzovini, Marco Calandri, and Marco Grangetto. Unveiling covid-19 from chest x-ray with deep learning: a hurdles race with small data. *arXiv preprint arXiv:2004.05405*, 2020.
- Andrew Rabinovich, Andrea Vedaldi, and Serge J Belongie. *Does image segmentation improve object categorization?* Citeseer, 2007.
- Joseph Paul Cohen, Paul Morrison, Lan Dao, Karsten Roth, Tim Q Duong, and Marzyeh Ghassemi. Covid-19 image data collection: Prospective predictions are the future. *arXiv preprint arXiv:2006.11988*, 2020b.
- Hinrich B. Winther, Hans Laser, Svetlana Gerbel, Sabine K. Maschke, Jan B. Hinrichs, Jens Vogel-Claussen, Frank K. Wacker, Marius M. Höper, and Bernhard C. Meyer. Covid-19 image repository, May 2020. URL https://figshare.com/articles/dataset/COVID-19_Image_Repository/12275009/1.
- Alex J DeGrave, Joseph D Janizek, and Su-In Lee. Ai for radiographic covid-19 detection selects shortcuts over signal. *medRxiv*, 2020.
- Douglas PS Gomes, Anwaar Ulhaq, Manoranjan Paul, Michael J Horry, Subrata Chakraborty, Manas Saha, Tanmoy Debnath, and DM Rahaman. Potential features of icu admission in x-ray images of covid-19 patients. *arXiv preprint arXiv:2009.12597*, 2020.
- Tasfia Shermin, Shyh Wei Teng, Manzur Murshed, Guojun Lu, Ferdous Sohel, and Manoranjan Paul. Enhanced transfer learning with imagenet trained classification layer. In *Pacific-Rim Symposium on Image and Video Technology*, pages 142–155. Springer, 2019.
- Sema Candemir and Sameer Antani. A review on lung boundary detection in chest x-rays. *International journal of computer assisted radiology and surgery*, 14(4):563–576, 2019.
- Sangheum Hwang and Sunggyun Park. Accurate lung segmentation via network-wise training of convolutional networks. In *Deep learning in medical image analysis and multimodal learning for clinical decision support*, pages 92–99. Springer International Publishing, 2017.
- Olaf Ronneberger, Philipp Fischer, and Thomas Brox. U-net: Convolutional networks for biomedical image segmentation. In *International Conference on Medical image computing and computer-assisted intervention*, pages 234–241. Springer International Publishing, 2015.
- Junji Shiraishi, Shigehiko Katsuragawa, Junpei Ikezoe, Tsuneo Matsumoto, Takeshi Kobayashi, Ken-ichi Komatsu, Mitate Matsui, Hiroshi Fujita, Yoshie Kodera, and Kunio Doi. Development of a digital image database for chest radiographs with and without a lung nodule: receiver operating characteristic analysis of radiologists’ detection of pulmonary nodules. *American Journal of Roentgenology*, 174(1):71–74, 2000.

- Kaiming He, Xiangyu Zhang, Shaoqing Ren, and Jian Sun. Deep residual learning for image recognition. In *Proceedings of the IEEE conference on computer vision and pattern recognition*, pages 770–778, 2016.
- Stefan Jaeger, Sema Candemir, Sameer Antani, Yi-Xiáng J Wáng, Pu-Xuan Lu, and George Thoma. Two public chest x-ray datasets for computer-aided screening of pulmonary diseases. *Quantitative imaging in medicine and surgery*, 4(6):475, 2014.
- F. Pedregosa, G. Varoquaux, A. Gramfort, V. Michel, B. Thirion, O. Grisel, M. Blondel, P. Prettenhofer, R. Weiss, V. Dubourg, J. Vanderplas, A. Passos, D. Cournapeau, M. Brucher, M. Perrot, and E. Duchesnay. Scikit-learn: Machine learning in Python. *Journal of Machine Learning Research*, 12:2825–2830, 2011.