



Since January 2020 Elsevier has created a COVID-19 resource centre with free information in English and Mandarin on the novel coronavirus COVID-19. The COVID-19 resource centre is hosted on Elsevier Connect, the company's public news and information website.

Elsevier hereby grants permission to make all its COVID-19-related research that is available on the COVID-19 resource centre - including this research content - immediately available in PubMed Central and other publicly funded repositories, such as the WHO COVID database with rights for unrestricted research re-use and analyses in any form or by any means with acknowledgement of the original source. These permissions are granted for free by Elsevier for as long as the COVID-19 resource centre remains active.



BS-Net: Learning COVID-19 pneumonia severity on a large chest X-ray dataset



Alberto Signoroni^{a,1,*}, Mattia Savardi^{a,1}, Sergio Benini^a, Nicola Adami^a, Riccardo Leonardi^a, Paolo Gibellini^b, Filippo Vaccher^b, Marco Ravanelli^b, Andrea Borghesi^b, Roberto Maroldi^b, Davide Farina^b

^a Department of Information Engineering, University of Brescia, Brescia, Italy

^b Department of Medical and Surgical Specialties, Radiological Sciences, and Public Health, University of Brescia, Brescia, Italy

ARTICLE INFO

Article history:

Received 6 August 2020

Revised 4 February 2021

Accepted 17 March 2021

Available online 31 March 2021

Keywords:

COVID-19 severity assessment

Chest X-rays

Semi-quantitative rating

End-to-end learning

Convolutional neural networks

ABSTRACT

In this work we design an end-to-end deep learning architecture for predicting, on Chest X-rays images (CXR), a multi-regional score conveying the degree of lung compromise in COVID-19 patients. Such semi-quantitative scoring system, namely *Brixia score*, is applied in serial monitoring of such patients, showing significant prognostic value, in one of the hospitals that experienced one of the highest pandemic peaks in Italy. To solve such a challenging visual task, we adopt a weakly supervised learning strategy structured to handle different tasks (segmentation, spatial alignment, and score estimation) trained with a “from-the-part-to-the-whole” procedure involving different datasets. In particular, we exploit a clinical dataset of almost 5,000 CXR annotated images collected in the same hospital. Our BS-Net demonstrates self-attentive behavior and a high degree of accuracy in all processing stages. Through inter-rater agreement tests and a gold standard comparison, we show that our solution outperforms single human annotators in rating accuracy and consistency, thus supporting the possibility of using this tool in contexts of computer-assisted monitoring. Highly resolved (super-pixel level) explainability maps are also generated, with an original technique, to visually help the understanding of the network activity on the lung areas. We also consider other scores proposed in literature and provide a comparison with a recently proposed non-specific approach. We eventually test the performance robustness of our model on an assorted public COVID-19 dataset, for which we also provide *Brixia score* annotations, observing good direct generalization and fine-tuning capabilities that highlight the portability of BS-Net in other clinical settings. The CXR dataset along with the source code and the trained model are publicly released for research purposes.

© 2021 Elsevier B.V. All rights reserved.

1. Introduction

Worldwide, the saturation of healthcare facilities, due to the high contagiousness of SARS-CoV-2 virus and the significant rate of respiratory complications (WHO, 2020), is indeed one among the most critical aspects of the ongoing COVID-19 pandemic. Under these conditions, it is extremely important to adopt all types of measures to improve the accuracy in monitoring the evolution of the disease and the level of coordination and communication

between different clinicians for the streamlining of healthcare procedures, from facility- to single patient-level.

In this context, thoracic imaging, specifically chest X-ray (CXR) and computed tomography (CT), is playing an essential role in the management of patients, especially those evidencing risk factors (from triage phases) or moderate to severe COVID-19 signs of pulmonary disease (Rubin et al., 2020). In particular, CXR is a widespread, relatively cheap, fast, and accessible diagnostic modality, which may be easily brought to the patient's bed, even in the emergency departments. Therefore, the use of CXR may be preferred to CT not only in limited resources environments, but where this becomes fundamental to avoid handling of more compromised patients, or to prevent infection spread due to patient movements from and towards radiology facilities (Manna et al., 2020). More-

* Corresponding author.

E-mail address: alberto.signoroni@unibs.it (A. Signoroni).

¹ Shared first authorship.




Fig. 1. Brixia score: (a) zone definition and (b–d) examples of annotations. Lungs are first divided into six zones on frontal chest X-rays. Line A is drawn at the level of the inferior wall of the aortic arch. Line B is drawn at the level of the inferior wall of the right inferior pulmonary vein. A and D upper zones; B and E middle zones; C and F lower zones. A score ranging from 0 (green) to 3 (black) is then assigned to each sector, based on the observed lung abnormalities. (For interpretation of the references to color in this figure legend, the reader is referred to the web version of this article.)

over, a serious X-ray dose concern arises in this context: due to the typical rapid worsening of the disease and the need for a prompt assessment of the therapeutic effects, serial image acquisitions for the same patient are often needed on a daily basis. For these reasons, notwithstanding a lower sensitivity compared to CT, CXR has been set up as a first diagnostic imaging option for COVID-19 severity assessment and disease monitoring in many healthcare facilities.

Due to the projective nature of the CXR image and the wide range of possible disease manifestations, a precise visual assessment of the entity and severity of the pulmonary involvement is particularly challenging. To this purpose scoring systems have been recently adopted to map radiologist judgments to numerical scales, leading to a more objective reporting and improved communication among specialists. In particular, from the beginning of the pandemic phase in Italy, the Radiology Unit 2 of ASST Spedali Civili di Brescia introduced a multi-valued scoring system, namely Brixia score, that was immediately implemented in the clinical routine (Borghesi and Maroldi, 2020). With this system, the lungs are divided into six regions, and the referring radiologist assigns to each region an integer rating from 0 to 3, based on the local assessed severity of lung compromise (see Fig. 1 and Section 3.1 for details). Other scores have been recently proposed as well, which are less resolute either in terms of intensity scale granularity (Toussie et al., 2020), or localization capacity (Wong et al., 2020; Cohen et al., 2020a) (see Section 3.2). Severity scores offer a mean to introduce the use of representation learning techniques to automatically assess disease severity using artificial intelligence (AI) approaches starting from CXR analysis. However, while in general computer-aided interpretation of radiological images based on Deep Learning (DL) can be an important asset for a more effective handling of the pandemic in many directions (Shi et al., 2021), the early availability of public datasets of CXR images from COVID-19 subjects catalyzed the research almost exclusively on assisted COVID-19 diagnosis (i.e., dichotomous or differential diagnosis versus other bacterial/viral forms of pneumonia). This happened despite the fact that normal CXRs do not exclude the possibility of COVID-19, and abnormal CXR is not enough specific for a reliable diagnosis (Rubin et al., 2020). Moreover, much research applying AI to CXR in the context of COVID-19 diagnosis considered small or private datasets, or lacked rigorous experimental methods, potentially leading to overfitting and performance overestimation (Castiglioni et al., 2020; Maguolo and Nanni, 2020; Tartaglione et al., 2020). So far, no much work has been done in other directions, such as COVID-19 severity assessment from CXR, despite this has been highlighted as one of the highest reasonable research efforts to be pursued in the field of AI-driven COVID-19 radiology (Laghi, 2020; Cohen et al., 2020c).

1.1. Aims and contributions

Main goal In this work, we aim at introducing the first learning-based solution specifically designed to obtain an effective and reliable assessment of the severity of COVID-19 disease by means of an automated CXR interpretation. This system is able to produce a robust multi-regional self-attentive scoring estimation on clinical data directly coming from different X-ray modalities (computed radiography, CR, and digital radiography by Direct X-ray detection, DX), acquisition directions (anteroposterior, AP, and posteroanterior, PA) and patient conditions (e.g., standing, supine, with or w/o the presence of life support systems).

Large CXR database For this purpose, we collected and operated on a large dataset of almost 5000 CXRs, which can be assumed representative of all possible manifestations and degrees of severity of the COVID-19 pneumonia, since these images come from the whole flow of hospitalized patients in one of the biggest healthcare facilities in northern-Italy during the first pandemic peak. This CXR dataset is fully annotated with scores directly coming from the patients' medical records, as provided by the reporting radiologist on duty among the about 30 specialists forming the radiology staff. Therefore, we had the possibility to work on a complete and faithful picture of an intense and emergency-driven clinical activity.

End-to-end severity assessment architecture We developed an original multi-block Deep Learning-based architecture designed to learn and execute different tasks at once such as: data normalization, lung segmentation, feature alignment, and the final multi-valued score estimation. Although the presence of specific task dedicated blocks, the overall logic is *end-to-end* in the sense that the gradient flows from the head (brixia score estimation), to the bottom (input CXR) without interruption. This is reached through a semantically hierarchical organization of the training, of which we give a high-level representation in Fig. 2. To squeeze all the information from the data, we organize the training according to a *from-the-part-to-the-whole* strategy. This consists of leveraging multiple CXR datasets (described in Section 4) for early training stages, then arriving at a global (involving all network portions) severity assessment training, based on the large collected dataset. Our end-to-end network architecture, called BS-Net and described in detail in Section 5, is thus characterized by the joint work of different sections, represented as a group of four hexagons in Fig. 2, comprising (1) a shared multi-task feature extraction backbone, (2) a state-of-art lung segmentation branch, (3) an original registration mechanism that acts as a “multi-resolution feature alignment” block operating on the encoding backbone, and (4) a multi-regional classification part for the final six-valued score estimation. All these blocks act together in the final training thanks to a loss specifically crated for this task. This loss guarantees also perfor-

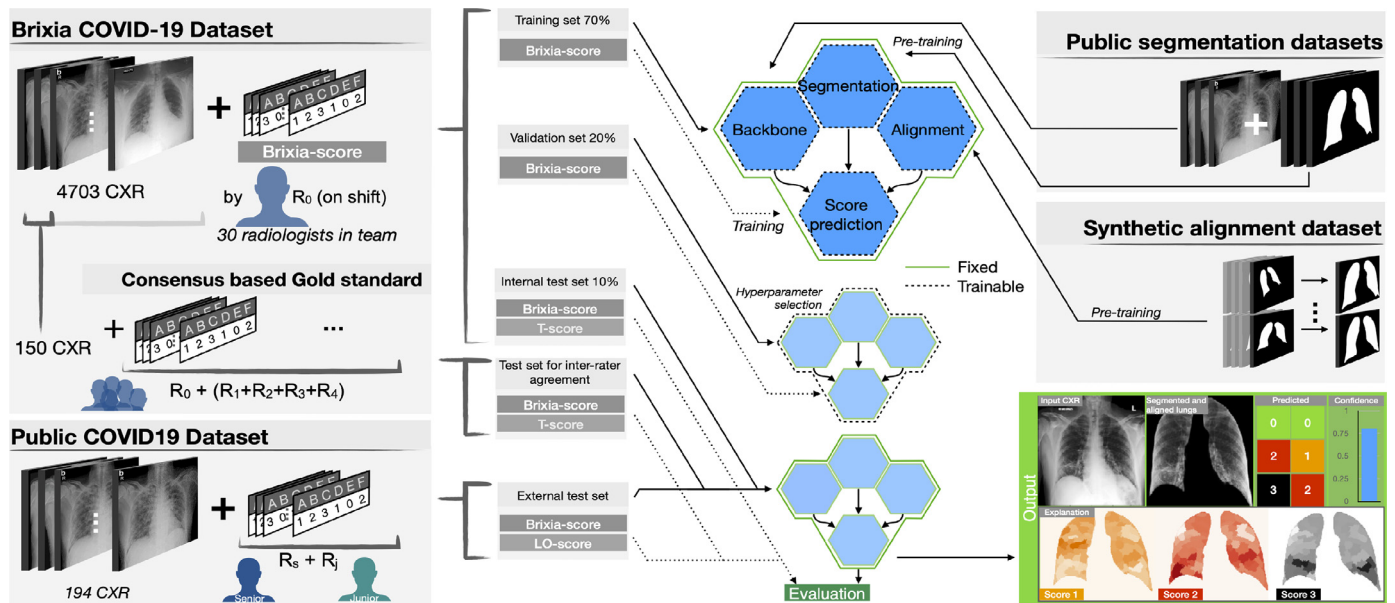


Fig. 2. Overview of the proposed method: representation of the two COVID-19 datasets (on the left) with associated *Brixia score* annotations, and of the other two datasets (on the right) used for the pre-training. Datasets splitting and usage is indicated (in the middle) for training/validation/test phases. The outputs of the proposed system are illustrated as well (bottom right).

mance robustness, comprising a differentiable version of the target discrete metric.

Weakly supervised learning framework The learning phase operates in a weakly-supervised fashion. This is due to the fact that difficulties and pitfalls in the visual interpretation of the disease signs on CXRs (spanning from subtle findings to heavy lung impairment), and the lack of detailed localization information, produces unavoidable inter-rater variability among radiologists in assigning scores (Zhou, 2017). Far from constituting in itself a “weakness”, these approaches have demonstrated to be highly valuable methods to leverage available knowledge in medical domains (Bontempi et al., 2020; Tajbakhsh et al., 2020; Wang et al., 2017; Xu et al., 2014; Karimi et al., 2019).

Explainability maps In the perspective of a responsible and transparent exploitation of the proposed solution, there is the need to establish a communication channel between the specialist and the AI system. Given the spatial distribution of the severity assessment, we need highly resolved explainability maps, also able to compensate for some limitations of conventional approaches based on Grad-CAM (Selvaraju et al., 2017). To this end we propose an original technique (described in Section 5.4) able to create highly structured explanation maps at a super-pixel level.

Experimental goals Experiments presented in Section 6 are carried out as follows. (1) In terms of system performance assessment, we select the best configurations and justify architectural choices also considering different possible alternatives from the literature. (2) To cope with the inter-rater variability and to demonstrate above radiologists’ performance, we involve a team of specialists to establish a consensus-based gold-standard as a reference for both single radiologist and our model ratings. (3) We consider other scores proposed in the literature and provide a comparison with a recently proposed non-specific approach for demonstrating both system versatility and performance. (4) We address the problem of portability of the proposed solution on CXRs coming from different worldwide contexts by providing new annotations and assessing our system behavior on a public CXR dataset of reference for COVID-19.

Data, code, and model distribution The whole dataset along with the source code and the trained model are available from <http://github.com/BrixIA>.

2. Related work

Since the very beginning of the pandemic, the attention and resources of researchers in digital technologies (Ting et al., 2020), AI and data science (Latif et al., 2020) have been captured by COVID-19. A review of artificial intelligence techniques in imaging data acquisition, segmentation, and diagnosis for COVID-19 has been made in Shi et al. (2021), where authors structured previous work according to different tasks, e.g., contactless imaging workflow, image segmentation, disease detection, radiomic feature extraction, etc.

The use of convolutional neural networks (CNN) to analyze CXRs for presumptive early diagnosis and better patient handling based on signs of pneumonia, has been proposed by Oh et al. (2020) at the very beginning of the outbreak, when a systematic collection of a large CXR dataset for deep neural network training was still problematic. As the epidemic was spreading, the increasing availability of CXR datasets from patients diagnosed with COVID-19 has polarized almost all the research efforts on diagnosis-oriented image interpretation studies. As it is hard to mention them all (there are well over 100 studies, most of them in a pre-print format), the reader can refer to an early reviewing effort Shi et al. (2021), while here we mention the ones most related to our work and derive most relevant emerging issues. Most presented methods exploit available public COVID-19 CXR datasets (Kalkreuth and Kaufmann, 2020). Constantly updated collections of COVID-19 CXR images are curated by the authors of (Linda Wang and Wong, 2020)² and (Cohen et al., 2020b).³ Prior to COVID-19, large CXR datasets have been released and also used in the context of open challenges for the analysis of different types of pneumonia Wang et al. (2017)⁴ or pulmonary and cardiac (cardiomegaly) diseases Irvin et al. (2019).⁵ These datasets are usually exploited as well in works related to COVID-19 diagnosis, as in Minaee et al. (2020), to pre-train existing networks and fine-tune them on a reduced set of COVID-19 CXRs, or to complete case cat-

² <https://github.com/lindawang/COVID-Net/blob/master/docs/COVIDx.md>.

³ <https://github.com/ieee8023/covid-chestxray-dataset>.

⁴ <https://www.kaggle.com/c/rsna-pneumonia-detection-challenge/data>.

⁵ <https://stanfordmlgroup.github.io/competitions/chexpert/>.

egories when the focus is on differential diagnosis to distinguish from other types of pneumonia (Oh et al., 2020; Li, Li and Zhu, 2020c; Linda Wang and Wong, 2020). However, beyond the fact that CXR modality should be more appropriate for patient monitoring and severity assessment than for primary COVID-19 detection, other issues severely affect the previous studies which employed CXR for COVID-19 diagnosis purposes (Burlacu et al., 2020). In particular, many of these data-driven works are interesting in principle, but almost all would require and benefit from extensions and validation on a higher number of CXRs from COVID-19 subjects. Working with datasets of a few hundred images, when analyzed with deep architectures, often results in severe overfitting and can encounter issues generated by unbalanced classes when larger datasets are used to represent other kinds of pneumonia (Pereira et al., 2020). Moreover, most of the studies issued in this emergency period are often based on sub-optimal experimental designs, and the numerous still unknowns factors about COVID-19 severely undermine the external validity and generalizability of the performance of diagnostic tests (Sardanelli and Di Leo, 2020). In (Maguolo and Nanni, 2020) it is pointed out that many CXR based systems seem to learn to recognize more the characteristics of the dataset rather than those related to COVID-19. This effect can be overcome, or at least mitigated, by working on more homogeneous and larger datasets, as in Castiglioni et al. (2020), or by preprocessing the data as in Pereira et al. (2020). In addition, automated lung segmentation can play an essential role in the design of a robust respiratory disease interpretation system, either diagnosis-oriented (Tartaglione et al., 2020) or, as we see in this work, for pneumonia severity assessment. Lung segmentation on CXR has been recently witnessing a convergence towards encoder-decoder architectures, such as in the U-Net framework (Ronneberger et al., 2015), and in the fully convolutional approach found in Long et al. (2015). All these approaches share skip connections as a common key element, originally found in ResNet (He et al., 2016) and DenseNet architectures (Huang et al., 2017). The idea behind skip connections is to combine coarse-to-fine activation maps from the encoder within the decoding flow. In doing so, these models are capable of efficiently using the extracted information from different abstraction layers. In Zhou et al. (2018), a nested U-Net is presented, bringing the idea of skip connection to its extreme. This approach is well suited to capture fine-grained details from the encoder network, and exploding them in the decoding branch. Last, in critical sectors like healthcare, the lack of understanding of complex machine-learned models is hugely problematic. Therefore, explainable AI approaches able to reveal to physicians where the model attention is directed are always desirable. However, while Grad-CAM (Selvaraju et al., 2017) and similar methods can work well on diagnostic tasks (see e.g., Oh et al., 2020; Reyes et al., 2020; Hryniewska et al., 2020; Karim et al., 2020; Rajaraman et al., 2020), this kind of approach is not enough informative to explain severity estimations, as it usually produces defocused heatmaps that hardly reveal fine details. Hence the need to find solutions to produce denser, more insightful, and more spatially detailed visual feedback to the clinician, keeping in mind perspectives of trustable deployment.

Despite some above evidenced issues, there is still an abundant ongoing effort toward AI-driven approaches for COVID-19 detection based on CT or CXR analysis, and this produced what has been recently dubbed as a deluge of articles in Summers (2021), where the need to move beyond opacity detection has been remarked. Cautions about a radiologic diagnosis of COVID-19 infection driven by deep learning have been also expressed by Laghi (2020), who states that a more interesting application of AI in COVID-19 infection is to provide for a more objective quantification of the disease, in order to allow the monitoring of the prognostic factors (i.e., lung compromising severity) for appropriate and timely

patient treatment, especially in critical conditions like the ones characterizing the management of health facilities overload. In the global race to contain and treat COVID-19, AI-based solutions have high potentials to expand the role of chest imaging beyond diagnosis, to facilitate risk stratification, disease progression monitoring, and trial of novel therapeutic targets (Kundu et al., 2020). However, the numerical disproportion of research works aiming at AI-driven image-based binary COVID-19 diagnosis, as well as the diffused availability of ready-to-use (with due fine-tuning or transfer learning) DL networks, should not bias observers' conviction (as it seems instead to happen in Summers, 2021) toward the idea that there is no need for purposefully technology design efforts, especially for different and clinically relevant tasks, as the one tackled here. Our work shows how a dedicated technical solution, which is up to the visual difficulty of a structured severity score assessment, can lead to a significant performance and robustness boost with respect to off-the-shelf methods. Anyway, so far only a few works present AI-driven solutions for COVID-19 disease monitoring and pneumonia severity assessment based on CXR, although this modality, for the aforementioned reasons, is part of the routine practice in many institutions, like the one from which this study originates. In (Cohen et al., 2020a), different kinds of features coming from a neural network pre-trained on non-COVID-19 CXR datasets are considered in their predictive value on the estimation of COVID-19 severity scores. In (Li et al., 2020b) a good correlation is reported between a lung based severity score judgement and machine prediction by using a transfer learning approach from a large non-COVID-19 dataset to a small COVID-19 one. However, authors acknowledge several limitations in the capability to handle image variability that can arise due to patient condition and image acquisition settings. Improved generalizability has been obtained in a subsequent work from same authors (Li et al., 2020a). In Zhu et al. (2020), transfer learning vs. conventional CNN learning has been compared on a data sample derived from Cohen et al., 2020b. In Blain et al. (2021) both interstitial and alveolar opacity are classified with a modular deep learning approach (segmentation stage followed by a fine-tuned classification network). Despite the limited data sample of 65 CXRs, the study shows correlations of severity estimation with age and comorbidity factors, intensity of care and radiologists' interpretations. In Amer et al., 2020 a geographic extent severity score (under the form of pneumonia area over lung area ratio) was estimated and correlated with experts' judgments on 94 CXRs, using pneumonia localization and lung segmentation networks. Geographic extent and opacity severity scores were predicted in Wong et al. (2020a) with a modified COVID-19 detection architecture (Linda Wang and Wong, 2020) and stratified Monte Carlo cross-validation, with measures of correlation with respect to expert annotations on 396 CXRs. These early studies, while establishing the feasibility of a COVID-19 severity estimation on CXRs, concurrently confirm the need of a dedicated design of methods for this challenging visual task, the urgency to operate on large annotated datasets coming from real clinical settings, and the need for expressive explainability solutions.

3. Scoring systems for severity assessment

CXR severity scoring based on the subdivision of lungs in different regions (Borghesi and Maroldi, 2020; Toussie et al., 2020; Wong et al., 2020) evidenced significant prognostic value when applied in serial monitoring of COVID-19 patients (Borghesi et al., 2020a,b; Maroldi et al., 2020). Since radiologists are asked to map a global or region-based qualitative judgment on a quantitative scale, this diagnostic image interpretation task can be defined as *semi-quantitative*, i.e. characterized by a certain degree of subjectivity. A detailed description of three semi-quantitative scoring systems we found in use follows.

3.1. Brixia score

The multi-region and multi-valued Brixia score was designed and implemented in routine reporting by the Radiology Unit 2 of ASST Spedali Civili di Brescia (Borghesi and Maroldi, 2020), and later validated for risk stratification on a large population in Borghesi et al. (2020). According to it, lungs in antero-posterior (AP) or postero-anterior (PA) views, are subdivided into six zones, three for each lung, as shown in Fig. 1(a). For each zone, a score 0 (no lung abnormalities), 1 (interstitial infiltrates), 2 (interstitial and alveolar infiltrates, interstitial dominant), or 3 (interstitial and alveolar infiltrates, alveolar dominant) is assigned, based on the characteristics and extent of lung abnormalities. The six scores can be aggregated to obtain a Global Score in the range [0,18]. Examples of scores assigned to different cases are showcased in Fig. 1(b–d). As in daily practice CXR exams are inevitably reported by different radiologists, this combined codification of the site and type of lung lesions makes the comparison of CXR exams faster and significantly more consistent, and this allows a better handling of patients.

3.2. Toussie score

In (Toussie et al., 2020) the presence/absence of pulmonary COVID-19 alterations is mapped on a 1/0 score associated to each of six pulmonary regions according to a subdivision scheme substantially reproducing the one of the Brixia score. We here refer to this score as T score which globally ranges from 0 to 6. By ignoring slight differences in terms of anatomic landmarks that guide the radiologist to determine the longitudinal lung subdivisions, the T score can be directly estimated from the Brixia score by just mapping the set of values {1, 2, 3} of the latter to the value 1 of the former.

3.3. GE-LO score

In Cohen et al. (2020a) two different COVID-19 severity scores are considered which are derived and simplified versions of a composite scoring system proposed by Warren et al. (2018) in the context of lung oedema. Both scores are composed by couples of values, one for each lung: (1) a geographic extent score, here GE score, in the integer range [0,8] and (2) a lung opacity score, here LO score in the integer range [0,6]. The GE score, introduced in Wong et al. (2020) for COVID-19 severity assessment, assigns for each lung a value depending on the extent of involvement by consolidation or ground glass opacity (0 = no involvement; 1 ≤ 25%; 2 = 25 – 50%; 3 = 50 – 75%; 4 ≥ 75% involvement). While this area-based quantification has no clear correspondence to the judgment made with the Brixia score, a possible mapping can be estimated with the LO score (e.g., by a simple linear regression, as we will see in Section 6.5), which assigns for each lung a value depending on degree of opacity (0 = no opacity; 1 = ground glass opacity; 2 = consolidation; 3 = white-out). A global score derived by a modified version of the ones introduced in Warren et al. (2018) has been also used in Li et al. (2020b).

3.4. AI-based prediction of severity scores

At first sight, an automatic assessment of a semi-quantitative prognostic score may seem easier than other tasks, such as differential diagnosis, or purely quantitative severity evaluations. Nevertheless, when dealing with semi-quantitative scores, major critical aspects arise.

First, the difficulty of establishing a ground truth information, since subjective differences in scoring are expressed by different

Table 1
Segmentation datasets.

	Training-set	Test-set	Split
Montgomery County	88	50	first 50
Shenzhen Hospital	516	50	first 50
JSRT database	124	123	original
Total	728	223	

radiologists while assessing and qualifying the presence of, sometimes subtle, abnormalities. This differs from more quantitative tasks that can be associated to measurable targets, as in the case of DL-based quantitative (volumetric) measure of opacities and consolidations on CT scans for lung involvement assessment in COVID-19 pneumonia (Huang et al., 2020; Lessmann et al., 2021; Gozes et al., 2020). The application of quantitative methods to AI-driven assessment of COVID-19 severity on CXR, however, is not advisable due to the projective nature of these images, with inherent ambiguities in relating opacity area measures to corresponding volumes.

A semi-quantitative scoring system can instead leverage the sensitivity of CXR as well as the ability of radiologists to detect COVID-19 pneumonia and communicate in an effective way its severity according to an agreed severity scale.

Second, the exact localization of the lung zones and of severity-related findings (even within each of the selected lung zones) remains implicit and related to the visual attention of the specialist in following anatomical landmarks (without any explicit localization information indicated with the score, nor any lung segmentation provided). This results in the difficulty to define reference spatial information usable as ground truth and implies significantly incomplete annotations with respect to the task.

Eventually, the same visual task related to global or localized COVID-19 severity assessment is challenging in itself, since CXR findings may be extremely variable (from no or subtle signs to extensive changes that modify anatomical patterns and borders), and the quality of the information conveyed by the images may be impaired from the presence of medical devices, or from sub-optimal patient positioning.

All these factors, if not handled, can impact in an unpredictable way on the reliability of an AI-based interpretation. This concomitant presence of quantitative and qualitative aspects, on a difficult visual analysis task, makes the six-valued Brixia score estimation on CXR particularly challenging.

4. Dataset

Training and validation of the proposed multi-network architecture (Fig. 2) take advantage of the use of multiple datasets, which are described in the following.

4.1. Segmentation datasets

For the segmentation module we exploit and merge three different datasets: Montgomery County (Jaeger et al., 2014), Shenzhen Hospital (Stirenko et al., 2018), and JSRT databases (Shiraishi et al., 2000) with the lung mask annotated by van Ginneken et al. (2006), for a total of about 1000 images. When indicated we adopt the original training/test set splitting (as for the JSRT database); otherwise, we consider the first 50 images as test set, and the remaining as training set (see Table 1).

4.2. Alignment dataset

CXRs acquired in a real clinical setting lack of standardized levels of magnification and alignment of the lungs. Moreover, possible patient positions are different (standing, sitting, prone, supine)

Table 2

Alignment dataset: synthetic transformations. The parameters refer to the implementation in Albumentation (Buslaev et al., 2020). In the last column is expressed the probability of application of each transformation.

	Parameters (up to)	Probability
Rotation	25 degree	0.8
Scale	10%	0.8
Shift	10%	0.8
Elastic transformation	alpha = 60, sigma = 12	0.2
Grid distortion	step = 5, limit = 0.3	0.2
Optical distortion	distort = 0.2, shift = 0.05	0.2




Fig. 3. Brixia score distribution with sex stratification on the Brixia COVID-19 dataset (left), and on the dataset of Cohen et al., 2020b (right).

and, according to subject conditions, it is not always feasible to produce images with an ideal shooting of the chest. To avoid the inclusion of anatomical parts not belonging to the lungs in the AI pipeline, which would increase the task complexity and introduce unwanted biases, we integrate in the network an alignment block. This exploits the same images used for the segmentation stage to create a synthetic dataset formed by artificially transformed images (see Table 2), including random rotations, shifts, and zooms, which are used in first phases of the training, in an on-line augmentation fashion, using the framework provided in Buslaev et al. (2020).

4.3. Brixia COVID-19 dataset

We collected a large dataset of CXR images corresponding to the entire amount of images taken for both triage and patient monitoring in sub-intensive and intensive care units during one month (between March 4th and April 4th 2020) of pandemic peak at the ASST Spedali Civili di Brescia, and contains all the variability originating from a real clinical scenario. It includes 4707 CXR images of COVID-19 subjects, acquired with both CR and DX modalities, in AP or PA projection, and retrieved from the facility RIS-PACS system. All data are directly imported from DICOM files, consisting in 12-bit gray-scale images, and mapped to float32 between 0 and 1. To mitigate the grayscale variability in the dataset, we normalize the appearance of the CXR by sequentially applying an adaptive histogram equalization (CLAHE, clip:0.01), a median filtering to cope with noise (kernel size: 3), and a clipping outside the 2nd and 98th percentile.

All image reports include the Brixia score as a string of six digits indicating the scores assigned to each region. The Global Score is simply the sum of the six regional scores, and its distribution can be appreciated in Fig. 3(left). Each image has been annotated with the six-valued score by the radiologist on shift (here referred to as R_0), belonging to a team of about 50 radiologists operating in different radiology units of the hospital with a very wide range of years of experience and different specific expertise in imaging of the chest. All images were collected and anonymized, and their usage for this study had the approval of the local Eth-

Table 3

Brixia COVID-19 dataset details.

Parameter	Value
Modality	CR (62%) - DX (38%)
View position	AP (87%) - PA (13%)
Manufacturers	Carestream, Siemens
Image size	$(1056 \div 3050) \times (1186 \div 3376)$
No. of images	4703
Training set	3311 images
Validation set	945 images
Test set	447 images

ical Committee (#NP4121) that also granted an authorization to release the whole anonymized dataset for research purposes. The main characteristics of the dataset are summarised in Table 3.

4.3.1. Consensus-based gold standard

To assess the level of inter-rater agreement among human annotators, we asked other 4 radiologists to rate a subset of 150 images belonging to the test set of the Brixia COVID-19 dataset. While R_0 is the original clinical annotation in the patient report, we name R_1, R_2, R_3 , and R_4 the four radiologists that provided additional scores. Their expertise is variegated so as to represent the whole staff experience: we have one resident at the 2nd year of training, and three staff radiologists with 9, 15 and 22 years of experience (reported numerical ordering of R_i does not necessarily correspond to seniority order). A Gold Standard score is then built based on a majority criterion, by exploiting the availability of multiple ratings (by R_0, R_1, R_2, R_3 , and R_4), using seniority in case of equally voted scores. Building such Gold Standard is useful, on the one hand, to grasp the inbuilt level of error in the training set, and on the other hand, to gain a reference measure for human performance and inter-rater variability assessments.

4.4. Public COVID-19 dataset

To later demonstrate the robustness and the portability of the proposed solution, we exploit the public repository by Cohen et al., 2020b, which contains CXR images of patients which are positive or suspected of COVID-19.⁶ This dataset is an aggregation of CXR images collected in several centers worldwide, at various spatial resolutions, and other unknown image quality parameters, such as modality and window-level settings. In order to contribute to such public dataset, two expert radiologists, a board-certified staff member R_s and a trainee R_j , with 22 and 2 years of experience, respectively, produced the related Brixia score annotations for CXR in this collection, exploiting labelbox,⁷ an online solution for labelling. After discarding few problematic cases (e.g., images with a significant portion missing, too low resolution/quality, the impossibility of scoring for external reasons, etc.), the obtained dataset is composed of 192 CXRs, completely annotated according to the Brixia score system.⁸ Its Global Score distribution is shown in Fig. 3 (right).

5. End-to-end multi-network model

5.1. Proposed architecture

To predict the pneumonia severity of a given CXR, we propose a novel architecture where different blocks cooperate, in an end-to-end scheme, to segment, align, and predict the Brixia score.

⁶ We downloaded a copy on May 11th, 2020.

⁷ <https://labelbox.com/>.

⁸ Available from both <https://github.com/BrixIA/Brixia-score-COVID-19> and <https://github.com/ieee8023/covid-chestxray-dataset>.

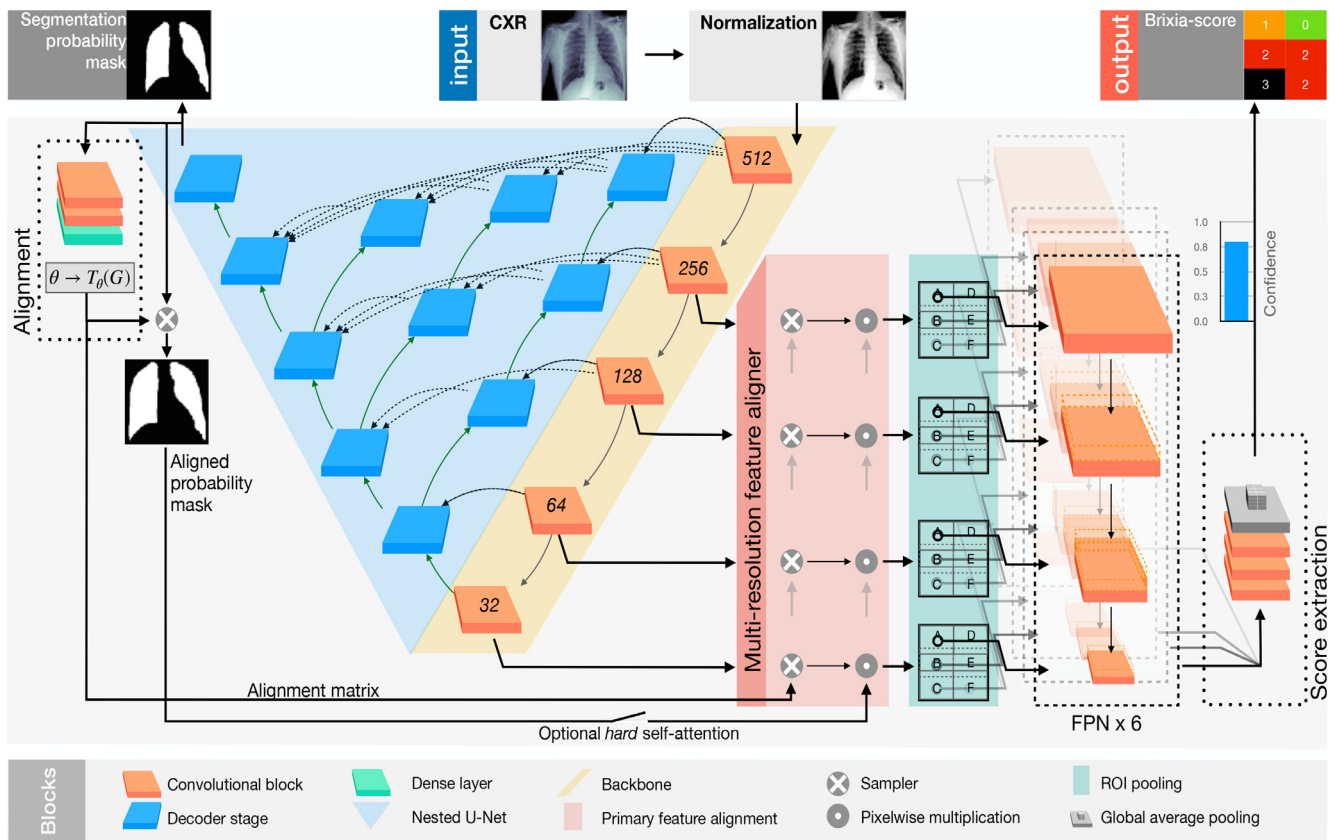


Fig. 4. Detailed scheme of the proposed architecture. In particular, in the top-middle the CXR to be analyzed is fed to the network. The produced outputs are: the segmentation mask of the lungs (top-left); the aligned mask (middle-left); the Brixia score (top-right).

Each block solves one of the specific tasks in which the severity score estimation can be subdivided, while the proposed global end-to-end loss provides the glue that concurs to the creation of a single end-to-end model. The global scheme is depicted in Fig. 4, while details on single parts follow.

Backbone The input image is initially processed by a cascade of convolutional blocks, referred to as Backbone (in yellow in Fig. 4). This cascade is used both as the encoder section of the segmentation network, and as the feature extractor for the Feature Pyramid Network of the classification branch. To identify the best solution at this stage, we tested different backbones among the state-of-the-art, i.e., ResNet (He et al., 2016), VGG (Simonyan and Zisserman, 2014), DenseNet (Huang et al., 2017), and Inception (Szegedy et al., 2017).

Segmentation Lung segmentation is performed by a nested version of U-net, also called U-Net++ (Zhou et al., 2018), a specialized architecture oriented to medical image segmentation (in blue in Fig. 4). It is composed of an encoder-decoder structure, where the encoder branch (i.e., the Backbone) exploits a nested interconnection to the decoder.

Alignment The segmentation probability map produced by the U-net++ decoder stage is used to estimate the alignment transformation. Alignment is achieved through a spatial transformer network (Jaderberg et al., 2015) able to estimate the spatial transform matrix in order to center, rotate, and correctly zoom the lungs. The spatial transformer network estimates a six-valued affine transform that is used to resample (by bilinear interpolation) and align all features at various scales produced by the multi-resolution backbone, before they are forwarded to the ROI (Region Of Interest) Pooling.

This alignment block is pre-trained on the synthetic alignment dataset in a weakly-supervised setting, using a Dice loss. The

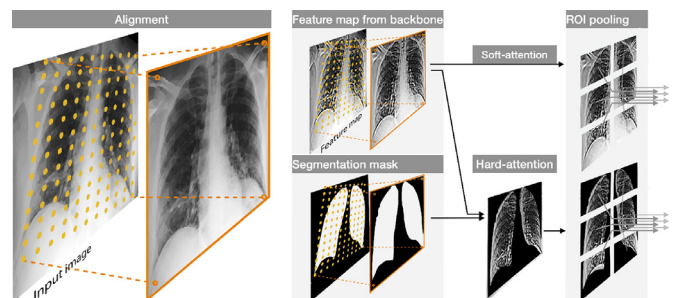


Fig. 5. Example of the alignment through the resampling grid produced by the transformation matrix, and its application to both the segmentation mask and the feature maps. On the right, the hard-attention mechanism and the ROI Pooling operation.

weakly-supervision is due to the fact that we do not provide transformation matrices as ground-truth labels, but original mask images before the synthetic transformation. Moreover, the labels are anyway noisy since images in the segmentation dataset (that compose the base for the used synthetic database) may be not perfectly aligned. More precisely we assume that the level of misalignment already present in dataset images is in general negligible with respect to the transforms we artificially apply. This allows a meaningful pre-training of the alignment block, as will be shown in the results, able to compensate the inaccurate patient positioning often present due to the various critical conditions of CXR acquisitions. In Fig. 4 we show the action of the alignment on the features at various scales generated by the backbone in a block we call “multi-resolution feature aligner”, while in Fig. 5 we give a representation of this block, where aligned features are pro-

duced by bi-linear resampling on the original maps. This resampling scheme is not only used to align the backbone features, but can be also helpful for realigning the segmentation map in an optional hard-attention configuration, as explained in the following.

Hard vs. Soft self-attention A hard self-attention mechanism can be applied by masking the aligned features with the segmentation mask (obtained as softmax probability map). This option has the advantage of switching off possible misleading regions outside the lungs, favouring the flow of relevant information only. Therefore the network can operate in two configurations: either with a hard self-attention scheme (HA), in which the realigned soft-max segmentation mask (from 0 to 1) is used as a (product) weighting mask in the multi-resolution feature aligner, or with a soft version (SA), where the segmentation mask is only used to estimate the alignment transform, but not to mask the aligned backbone features.

ROI pooling The aligned (and optionally hard masked) features are finally used to estimate the 3×2 matrix containing the Brixia score. To this purpose, ROI Pooling is performed on a fixed grid with the same dimensions. In particular, from the aligned features map produced by the multi-resolution feature aligner, the ROI Pooling extracts the 6 Brixia score regions (with a vertical overlap of 25%, and no horizontal overlap, since the left/right boundary between lungs is easily identified, while the vertical separation presents a larger variability). This pooling module introduces a-priori information regarding the location of the six regions (i.e., A, B, \dots, F), while leaving to the network the role to correctly rearrange the lungs feature maps by means of the alignment block. As output, this block returns 6 feature maps (one for each lung region) for each level in the backbone. The combination of alignment and ROI pooling produces (especially with hard attention) a self-attentive mechanism useful to propagate only the correct area of the lungs towards the final classification stage.

Scoring head The final scoring module exploits the idea of Feature Pyramid Networks (FPN) (Lin et al., 2017) for the combination of multi-scale feature maps. As depicted in Fig. 4, we combine feature maps that come from various levels of the network, therefore with different semantic information at various resolutions. The multi-resolution feature aligner produces input feature maps that are well focused on the specific area of interest. Eventually, the output of the FPN layer flows in a series of convolutional blocks to retrieve the output map ($3 \times 2 \times 4$, i.e., 3 rows, 2 columns, and 4 possible scores $[0, \dots, 3]$). The classification is performed by a final Global Average Pooling layer and a SoftMax activation.

5.2. Loss function and model training

The Loss function we use for training, is a sparse categorical cross entropy (SCCE) with a (differentiable) mean absolute error (MAE^d) contribution:

$$L = \alpha \cdot \text{SCCE} + (1 - \alpha) \cdot \text{MAE}^d \quad (1)$$

where α controls how much weight is given to SCCE and MAE^d, which are defined as follows:

$$\text{SCCE} = -\frac{1}{C} \sum_c y_c \log(\hat{y}_c) \quad (2)$$

$$\text{MAE}^d = \frac{1}{C} \left\| y - \sum_{c \in C} \frac{e^{\beta \hat{y}_c}}{\sum_{j \in C} e^{\beta \hat{y}_j}} c \right\| \quad (3)$$

where y is the reference Brixia score, \hat{y} is the predicted one, and $c \in [0, 1, \dots, 3]$ is the score class. To make the mean absolute error differentiable (MAE^d), β can be chosen to be an arbitrary large value.

The selection of such loss function is coherent with the choice to configure the Brixia score problem as a joint multi-class

classification and regression. Tackling our score estimation as a classification problem allows to associate to each score a confidence value: this can be useful to either produce a weighted average, or to introduce a quality parameter that the system, or the radiologist, can take into account. Moreover, the MAE^d component, while being meaningful for the scoring system, adds robustness to outliers and noise.

Due to the nature and complexity of the proposed architecture, the training of network weights takes place at several stages, according to a *from-the-part-to-the-whole* strategy, according to a task-driven policy. This does not contradict the end-to-end nature of the system. Indeed, not only pre-training each sub-network in a structured multi-level training is a possibility, but it is also the most advisable way to proceed, as also recognized in Glasmachers (2017). The different sections of the overall network are therefore first pre-trained on specific tasks: U-Net++ is trained using the lung segmentation masks in the segmentation datasets; the alignment block is trained using the synthetic alignment dataset (in a weakly-supervised setting); the classification portion (Scoring head) is trained using Brixia COVID-19 dataset, while blocking the weights of the other components (i.e., Backbone, Segmentation, Alignment). Then, a complete fine-tuning on Brixia COVID-19 dataset follows, making all weights (about 20 Million) trainable. The network hyper-parameters eventually undergo a selection that maximizes the MAE score on the validation set.

5.3. Implementation details

The network has an input size of 512×512 . The selected backbone is a ResNet-18 He et al. (2016), because it offers the best trade-off between the expressiveness of the extracted features and the memory footprint (as in the case of the input size). In the network, we use the rectified linear unit (ReLU) activation functions for the convolutional layer of the backbone, and the U-Net++, while the Swish activation function by Ramachandran et al. (2017) is used for the remaining blocks. We extensively make use of online augmentation throughout the learning phases. In particular, we apply all the geometric transformations described in Section 4.2, plus random brightness and contrast, as well. Moreover, we randomly flip images horizontally (and the score, accordingly) with a probability of 0.5. We exploit, for training purposes, two machines equipped with Titan® V GPUs. We train the model by jointly optimizing the sparse categorical cross-entropy function and MAE, with a selected $\alpha = 0.7$. Convergence is achieved after roughly 6 h of training (80 epochs), using Adam (Kingma and Ba, 2014) with an initial learning rate of 3×10^{-2} , halving it on flattening. The batch size is set to 8.

5.4. Super-pixel explainability maps

To evaluate whether the network is predicting the severity score on the basis of correctly identified lung areas, we need a method capable of generating explainability maps with a sufficiently high resolution. Unfortunately, with the chosen network architecture, the popular Grad-CAM approach (Selvaraju et al., 2017), and similar ones, generates poorly localized, and spatially blurred, visual explanations of activation regions, as it happens for example also in Karim et al. (2020); Oh et al. (2020). Moreover, we would face a concrete difficulty in our context as Grad-CAM would generate $6(\text{regions}) \times 4(\text{classes})$ maps that would be difficult to combine for an easy and fast inspection by the radiologist during the diagnosis. For these reasons, we designed a novel method for generating useful explainability maps, loosely inspired by the LIME (Ribeiro et al., 2016) initial phases. The creation of the explainability map starts with the input image division into N super-pixels i.e., regions that share similar intensity and pattern, extracted as

Table 4
Lung segmentation performance.

	Backbone	Dice coefficient	IoU
U-Net+	ResNet-18	0.971	0.945
U-Net	ResNet-18	0.969	0.941

in Vedaldi and Soatto (2008). Starting from the input image, we create N image replicas in which a single super-pixel i (from 1 to N) is masked to zero. We call p_0 the probability map that the model produces starting from the original image ($2 \times 3 \times 4$ values, i.e., one for each of the 4 severity classes in every lung sector that composes the Brixia score). We call instead p_i the probability map produced from the i th replica. We then accumulate for the differences between all the super-pixel masked predictions p_i and the original prediction p_0 . Given S the set of super-pixels, the output explanation map E is obtained as:

$$E = \sum_i^{|S|} S_i \cdot (p_i - p_0) \quad (4)$$

Intuitively, the obtained maps highlight the regions that most account for the score in exam. Examples can be appreciated in Fig. 2 (bottom-right) as well as in Fig. 10: the more intense the color, the more important the region contribution to the score decision.

6. Results

Through an articulated experimental validation we first consider how single components of the architecture operate (Sections 6.1 and 6.2) and we give a complete picture of the severity assessment performance referred to the whole collected dataset (Section 6.3). Then we deal with the inter-rater variability issue and demonstrate that the proposed solution overcomes the radiologists' performance on a consensus based Gold Standard reference (Section 6.4). We then widen the scope of our results in two directions: (1) we consider other proposed severity scores and give measures, observations, and comparisons that clearly support the need of dedicated solutions to the COVID-19 severity assessment (Section 6.5); (2) we consider the portability (both direct or mediated by a fine-tuning) of our model on public data collected in the most different locations and conditions, verifying a high degree of robustness and generalization of our solution (Section 6.7). Qualitative evidences of the role that the new explainability solution can have in a responsible and transparent use of the technology are then proposed (Section 6.8). We conclude with some ablation studies and technology variation experiments aiming essentially at showing how the complexity of our multi-network architecture is neither over- nor under-sized, but it is adequate to the needs and complexity of the target visual task (Section 6.9). All results presented in this section are discussed in the following Section 7.

6.1. Lung segmentation

The performance of the segmentation stage shows totally comparable results with respect to state-of-the-art methods (Candemir and Antani, 2019), both DL based (Frid-Adar et al., 2018; Oh et al., 2020; Zhou et al., 2018), and hybrid (Candemir et al., 2014). Table 4 reports the results for the U-Net++ (and U-Net comparison) in terms of Dice coefficient and Intersection over Union (IoU), aka Jaccard index, obtained on the test set of the segmentation datasets (Section 4.1). Training curves for the segmentation task (on both training set, and validation set), tracking the IoU, are shown in Fig. 8-a for BS-Net in HA configuration.

6.2. Alignment

After training on the synthetic dataset described in Section 4.2, we report the following alignment results: Dice coefficient = 0.873, IoU = 0.778. The Dice coefficient and IoU are calculated using the classical definition by considering the original masks (before synthetic misalignment) and the ones re-aligned by means of the affine transform estimated by the network. Training curves for the alignment task (on both training set, and validation set), tracking the IoU, are shown in Fig. 8b, always for HA configuration. Convergence behaviors are clearly visible, with a consolidating residual distance between training and validation curves.

Despite the difficulty of the task and the fact that simulated transforms are usually overemphasised with respect to misalignment found in real data, the measured alignment performance on the synthetic dataset produces a significant performance boost, as it clearly emerges from the ad-hoc ablation studies on rotations (see test in Section 6.9).

Residual errors are typically in the form of slight rotations and zooms, which are not critical and well tolerable in terms of overall self-attentive behavior. Moreover, during a visual check of the Brixia COVID-19 test dataset, no impairing misalignments (lung outside the normalized region) were observed.

6.3. Score prediction on Brixia COVID-19 dataset

To evaluate the overall performance of the network, we analyze the Brixia score predictions with respect to the score assigned by the radiologist(s) R_0 , i.e. the one(s) who originally annotated the CXR during the clinical practice. Discrepancies found on the 449 test images of Brixia COVID-19 dataset are evaluated in terms of Mean Error (MEr), Mean Absolute Error (MAE) with its standard deviation (SD), and Correlation Coefficient (CC).

Four networks are considered for comparison, three of which are different configurations of BS-Net: the hard attention (HA) one, the soft attention (SA) one, and one ensemble of the two previous configurations (ENS). In particular ENS configuration exploits both HA and SA paths to make the final prediction (by averaging their output probabilities): these realizations are the best model with respect to the validation set, and the obtained model after the last training iteration. The fourth compared network is ResNet-18 (the backbone of our framework) as an all-in-one solution (with no dedicated segmentation, nor alignment stage), since it is one among the most adopted architectures in studies involving CXR analyses, also until today in the COVID-19 context (Minaee et al., 2020b).

Table 5 lists all performance values referred to each of the six regions (A to F) of the Brixia score (range [0–3]), to the average on single regions, and to the Global Score (range [0–18]). In a consistency assessment perspective, Figs. 6 (top) and (bottom) show the confusion matrices for the four networks related to the score value assignments for single lung regions, and for their sum (Global Score), respectively. From Table 5 it clearly emerges that the ensemble decision strategy (ENS) succeeds in combining the strengths of the soft and hard attention policies, with the best average MAE on the six regions of 0.441. Conversely, the straightforward end-to-end approach offered by means of the all-in-one ResNet-18 is always the worst option compared to the three BS-Net configs.

The error distribution analysis on MAE, depicted in Fig. 7, shows the prevalence of lower error values on both single lung regions and Global Score estimations. The joint view gives another evidence that single-region errors unlikely sum up as constructive interference. Training curves for this prediction task (on both training set, and validation set), tracking the Mean Absolute Error (MAE) on the set of lung sectors are shown in Fig. 8c (HA con-

Table 5
Brixia score prediction performance parameters for the four considered models on the Brixia COVID-19 dataset (only blind test set results reported). Parameters are evaluated on each single lung region (A-F), averaged on all the lung regions and on the Global Score (P -value $\ll 0.00001$ everywhere).

		A	B	C	D	E	F	Avg. on regions	Global score
BS-Net-Ens	MEr	0.169	-0.038	-0.056	0.125	-0.045	-0.192	-0.006	-0.036
		0.107	-0.087	-0.171	0.082	-0.129	-0.343	-0.090	-0.541
		0.156	-0.147	-0.085	0.107	-0.016	-0.238	-0.037	-0.223
		0.356	-0.038	-0.056	0.125	-0.045	-0.192	0.100	0.601
BS-Net-Ens	MAE	0.459	0.448	0.412	0.374	0.459	0.494	0.441	1.728
		0.499	0.501	0.506	0.408	0.499	0.566	0.496	1.846
		0.481	0.477	0.481	0.370	0.488	0.532	0.471	1.826
		0.543	0.486	0.506	0.452	0.584	0.530	0.517	1.951
BS-Net-Ens	SD	0.604	0.524	0.540	0.541	0.574	0.609	0.565	1.429
		0.638	0.560	0.579	0.576	0.594	0.634	0.597	1.514
		0.613	0.575	0.583	0.552	0.594	0.616	0.589	1.505
		0.657	0.579	0.591	0.632	0.657	0.601	0.619	1.710
BS-Net-Ens	CC	0.675	0.779	0.731	0.682	0.737	0.672	0.713	0.862
		0.635	0.733	0.675	0.633	0.718	0.636	0.672	0.847
		0.665	0.742	0.679	0.662	0.722	0.645	0.686	0.845
		0.598	0.739	0.667	0.562	0.655	0.643	0.644	0.842



Fig. 6. Consistency/confusion matrices based on lung regions score values (top, 0–3), and on Global Score values (bottom, 0–18).

figuration). Convergence behaviors are clearly visible, with a consolidating residual distance between training and validation curves.

6.4. Performance assessment on consensus-based gold standard

Results in Table 5 are computed with respect to the score assigned by a single radiologist (i.e., R_0) among the ones in the whole staff. With the aim of providing a more reliable reference, we consider the Consensus-based Gold Standard (CbGS) dataset (Section 4.3.1). This allows to recompute the BS-Net (ENS) performance on a subset of 150 images, which are annotated with multiple ratings (by R_0, R_1, R_2, R_3 , and R_4) from which a consensus reference score is derived. Table 6 (top-center) clearly shows a significant improvement on MAE with respect to the comparison versus R_0 only (0.424 vs. 0.452).

Inter-rater agreement: human vs. machine performance In Fig. 9 we assess the inter-rater agreement by listing MAE and SD values referred to all possible pairs of raters, including R_0 , and also BS-Net (ENS), as a further “virtual rater”. Looking at how BS-Net behaves (orange boxes in Fig. 9), we observe the level of agreement reached between the network and any other rater is, almost always, higher than the inter-rater agreements between any pair of human raters. For example, by considering R_0 as a common reference, we have an equal performance only in the case of R_4 . Table 6 confirms how BS-Net (top) performs on average significantly better (MAE 0.424) with respect to the global indicators (bottom) coming from averaging all pairwise R_i vs. R_j ($i \neq j$) comparisons (MAE 0.528).

The inter-rater agreement values (between human raters) indicates a moderate level of agreement (both averaged two-raters




Fig. 7. Single and joint MAE distribution for lung regions and Global Score predictions obtained by BS-Net (ENS).

Table 6

Results on the Consensus-based Gold Standard dataset (150 images): (top) performance of BS-Net (ENS) computed on the Consensus-based Gold Standard; (center) performance of BS-Net (HS) vs. original rater R_0 ; (bottom) averaged performance of all R_i vs. CbGS

	MEr	MAE	SD	CC
BS-Net (ENS) vs. CbGS				
Avg. on reg.	-0.133	0.424	0.580	0.743
Global score	0.800	1.787	1.354	0.907
BS-Net (ENS) vs. R_0				
Avg. on reg.	-0.019	0.452	0.575	0.754
Global score	-0.113	1.847	1.553	0.834
Average on all pairs of radiologists				
Avg. on reg.	-0.131	0.528	0.614	0.736
Global score	-0.784	2.592	1.965	0.835

Cohen's kappa value and multi-rater Fleiss' kappa value, based on single cell scores, are around 0.4). On the one hand, this is a confirmation that the Brixia score, probably succeeds in reaching the tough compromise between maximizing the score expressiveness (spatial and rating granularity), while keeping under control the level of subjectivity (inter-rater variability). On the other hand, measured inter-rater variability levels constitute a clear limitation that bound the network learning abilities (weak supervision), concurrently allowing to assess whether the network surpasses single radiologists performance. The opportunity for the network to learn not only from a single radiologist but virtually from a fairly large community of specialists (being R_0 a varying radiologist) is, de facto, the margin that we have been able to exploit.




Fig. 8. Training curves related to BS-Net-HA. Segmentation (a); Alignment (b); Brixia score prediction - best single model (c).

	Avg. on regions					Global Score				
	R_1	R_2	R_3	R_4	BS-Net	R_1	R_2	R_3	R_4	BS-Net
R_0	0.55 (0.63)	0.47 (0.56)	0.58 (0.65)	0.45 (0.58)	0.45 (0.58)	2.58 (1.93)	2.31 (1.79)	2.86 (2.19)	2.05 (1.61)	1.85 (1.55)
R_1		0.37 (0.52)	0.56 (0.68)	0.60 (0.62)	0.52 (0.64)		1.53 (1.38)	2.84 (2.35)	2.93 (2.07)	2.28 (1.65)
R_2			0.58 (0.65)	0.54 (0.59)	0.45 (0.55)			3.07 (2.40)	2.85 (1.88)	2.08 (1.58)
R_3				0.58 (0.66)	0.56 (0.67)				2.89 (2.05)	2.60 (1.71)
R_4					0.42 (0.55)					1.77 (1.43)

Fig. 9. Pairwise inter-rater results in terms of MAE (and SD). In the most right column (orange), the inter-rater results with predictions by BS-Net-Ens.

Table 7

T score performance (faithful simulation of the score proposed in Toussie et al., 2020).

	MEr	MAE	SD	CC
BS-Net-ENS on the whole test set				
Avg. on reg.	0,009	0,147	0,349	0,51
Global score	0,056	0,742	0,921	0,728
BS-Net-ENS on the CbGS set				
Avg. on reg.	0,074	0,174	0,38	0,563
Global score	0,447	0,94	1,07	0,81

6.5. Performance assessment on Toussie score

In Table 7 we show the performance related to six-valued T score, akin to the score presented in Toussie et al. (2020), derived by thresholding the Brixia score as described in Section 3.2. We provide results on both the whole Brixia COVID-19 test set and on the Consensus-Based Gold Standard set. Interestingly, the correlation increases from 0.73 on the whole test set to 0.81 on the CbGS set, while MAE, despite maintaining acceptable values, increases as well (probably due to a residual from a non perfect mapping between the two scores).

6.6. Performance assessment on GE-LO score

In Table 8 we simulate the computation of the LO score by Brixia score mapping followed by linear regression (see Section 3.2). Thanks to annotations we provided for CXRs in the dataset by (Cohen et al., 2020c) we had the possibility to perform a direct comparison with the non-specific method de-

Table 8
Opacity score linear regression from global Brixia score.

	Correlation	R ²	MAE	MSE	coef	intercept
from						
Best (Cohen et al., 2020a)	0.80 ± 0.05	0.60 ± 0.17	1.14 ± 0.11	2.06 ± 0.34	–	–
BS-Net on Cohen set	0.84 ± 0.05	0.54 ± 0.17	0.67 ± 0.10	0.68 ± 0.18	0.31 ± 0.01	0.15 ± 0.08
BS-Net on our subset	0.85 ± 0.08	0.53 ± 0.24	0.67 ± 0.12	0.67 ± 0.22	0.31 ± 0.01	0.09 ± 0.09
R _s on our subset	0.90 ± 0.06	0.72 ± 0.13	0.55 ± 0.10	0.45 ± 0.15	0.28 ± 0.01	0.55 ± 0.07

Table 9

Portability tests on the public dataset (Cohen et al., 2020b). MAE and its SD are listed for both reporting radiologist R_j and BS-Net-HA. The network has been used in three training conditions: (1) as is, originally trained on the Brixia COVID-19 dataset (full), (2) fine-tuned on the public dataset (subset, fine-tuning), and (3) completely retrained, classification part only, on the public dataset (subset, from scratch).

	Test	A	B	C	D	E	F	Avg. on regions	Global score
Radiologist R _j	full	0.135	0.182	0.156	0.115	0.177	0.021	0.131	0.786
	subset	0.149	0.234	0.191	0.000	0.191	0.021	0.131	0.787
BS-Net-HA	full	0.177	−0.167	−0.177	0.167	−0.208	−0.651	−0.143	−0.859
	subset	0.125	−0.104	−0.188	0.167	−0.063	−0.583	−0.108	−0.646
* from scratch	subset	0.208	0.396	0.104	0.250	0.042	0.063	0.177	1.063
* fine-tuning	subset	0.146	0.167	−0.042	0.229	0.167	−0.208	0.076	0.458
Radiologist	full	0.396	0.401	0.438	0.333	0.396	0.469	0.405	1.974
	subset	0.404	0.447	0.489	0.340	0.362	0.447	0.415	1.851
BS-Net-HA	full	0.521	0.438	0.552	0.385	0.438	0.776	0.518	2.214
	subset	0.458	0.479	0.521	0.375	0.479	0.625	0.490	2.396
* from scratch	subset	0.375	0.646	0.604	0.458	0.542	0.479	0.517	2.188
* fine-tuning	subset	0.479	0.500	0.500	0.354	0.458	0.500	0.465	2.000

scribed in Cohen et al. (2020a) from which we report only the best results. We report the results produced by BS-Net-ENS on the intersection between the subset of the Cohen dataset considered in our work (Section 4.4) and the CXRs we found used in Cohen et al. (2020a) (a retrospective cohort of 94 PA CXR images). We also produce the results considering the whole subset for which we produced Brixia score annotations (obtaining virtually equivalent results). We also report LO score related results starting from Brixia score assigned by our experts. The performance boost produced by a prediction from a specifically designed solution is evident, and this is coherent also to the considerations and limitations acknowledged by the authors of Cohen et al. (2020a).

6.7. Public COVID-19 datasets: portability tests

The aggregate public CXR dataset (Cohen et al., 2020b), described in Section 4.4, has been judged as inherently well representative with respect to the various manifestations degrees of COVID-19. This dataset is quite heterogeneous and of a different nature with respect to our dataset acquired in clinical conditions (see Section 4.4). In particular the non-Dicom format, and the presence of low-resolution images and screenshots, make this dataset a challenging test bed to assess model portability and simulate a worst-case (or stress test) scenario. Exploiting the two independent Brixia score annotations of this dataset (one from a senior R_s, and another from a junior radiologist R_j), we performed a portability study, with the aim of deriving some useful guidance for extended use of our model on data generated in other facilities. In particular, we carried out three tests on BS-Net (HA configuration) measuring the performance on: (1) the whole set of 192 annotated images (full); (2) a reduced test set with fine-tuning, after random 75/25 splitting (subset, fine-tuning); (2) a reduced test set with partial retraining of the network, after random 75/25 splitting, with segmentation and alignment blocks trained on their specific datasets (subset, from scratch). In reporting results, we considered the senior radiologist R_s as the reference in order to have

the possibility to assess the second rater R_j performance the same way we assess the network performance. Table 9 lists all results from the above described tests.

Looking at results on the full dataset, we can derive that, even by directly applying the model trained on the Brixia COVID-19 dataset on a completely different dataset (and collected in a highly uncontrolled way), the network confirms the meaningfulness of the learning task and shows a fair robustness to work in different context even in an uncontrolled way. On the other hand, the skilled human observer confirms higher generalization capability, with a MAE of 0.405 on the full dataset. On the reduced subset, when retrained from scratch, the network is not able anymore to produce even the results obtained by the same model trained on the Brixia COVID-19 dataset: a clear evidence of the need to work with a large dataset and of the adequate capacity of our model. This is further confirmed by looking at the fine-tuning results, where the network reaches the best performance (MAE 0.465) by exploiting the already trained baseline with only a tiny amount of training data.

6.8. Explainability maps

In Fig. 10 we illustrate some explainability maps generated on Brixia COVID-19 data: three are exact predictions (top), while two are chosen among the most difficult cases (bottom). Along with maps, we also report the related lung segmentation and alignment images. Such maps, obtained as described in Section 5.4, clearly highlight the regions that triggered a specific score: they are drawn with the color of the class score (i.e., 1 = orange, 2 = red, 3 = black) in case they significantly contributed to the decision for that score, while a white colored region means that it gave no contribution to such score decision. The first row of Fig. 10 offers a clear overview of the agreement level between the network and the radiologist in shift (R₀). Conversely, the two cases in the bottom row of Fig. 10 evidences both under- and over-estimations in single sectors, despite producing a correct Global Score in the case on the left.

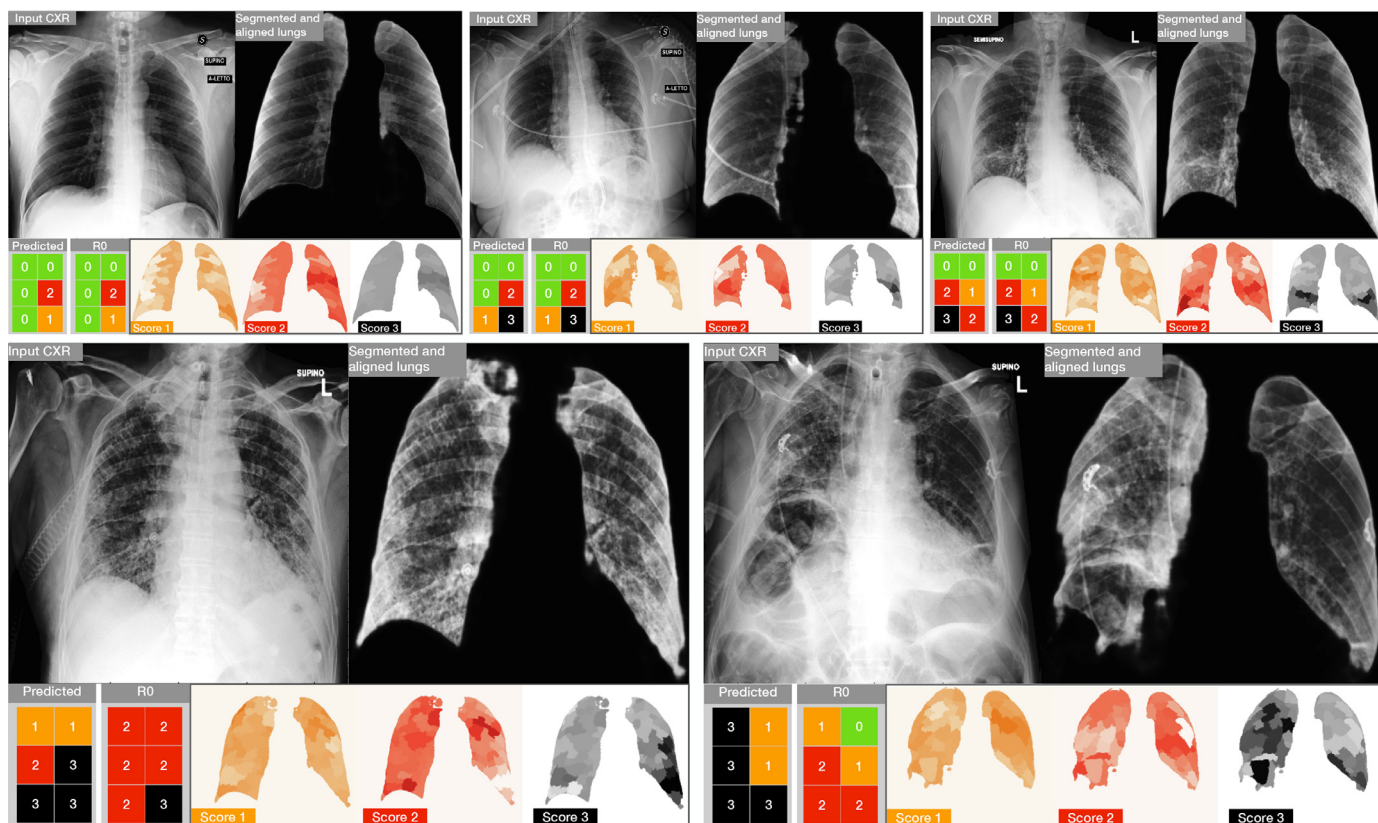


Fig. 10. Results and related explainability maps obtained on five examples from the Brixia COVID-19 test set. (top) Three examples of accurate predictions. (bottom) Two critical cases in which the prediction is poor with respect to the original clinical annotation R_0 . For each block, the most left image is the input CXR, followed by the aligned and masked lungs. In the second row we show the predicted Brixia score with respect to the original clinical annotation R_0 , and the explainability map. In such maps the relevance is colored so that white means that the region does not contribute to that prediction, while the class color (i.e., 1 = orange, 2 = red, 3 = black) means that the region had an important role in the prediction of the T score class.

6.9. Ablation and variation studies

We conducted various ablation and technology variation studies. First, we assess the actual contribution of the last training phase that involves all weights in the end-to-end configuration. Then, adopting BS-Net in the HA configuration, we conducted two sets of experiments regarding some modifications of feature extraction and data augmentation strategies, which are carried out on the training (3313 CXRs) and on the validation (945 CXRs) sets of the Brixia COVID-19 database. Finally, we carried out an experiment to evaluate the impact of the multi-resolution feature aligner.

End-to-end training With reference to the structured “from-the-part-to-the-whole” training strategy, the contribution of the last end-to-end training stage is significant, since it accounts for about 6–7% of the MAE performance, for both Soft and Hard Attention configurations.

Feature extraction This set of experiments investigates on the type and complexity of the feature map extraction leading to the Brixia score estimation (see Scoring head in Section 5.1). We then compare the adopted FPN-based solution with three different configurations. The first, simplified, configuration gets rid of the multi-scale approach, so that the score estimation exploits the features extracted by the backbone head. The second, more complex, configuration envisages an articulated multi-scale approach based on the EfficientDet (Tan et al., 2019), where BiFPN (Bidirectional FPN) blocks are introduced for an easy and fast multi-scale feature fusion. The third configuration adds one resolution level to the FPN to allow the flow of 1024×1024 images. Results on this benchmark are reported in Table 10. While the first simplified configuration produces an improved MAE on the training set,

Table 10
Performance on the training and validation set for different Feature Pyramid Networks (or lack of). (*) No complete fine-tuning due to memory limitation.

	Train MAE (SD)	Val MAE (SD)
FPN (adopted)	0.455 (0.578)	0.469 (0.583)
Backbone head	0.395 (0.542)	0.475 (0.592)
BiFPN	0.486 (0.593)	0.504 (0.598)
FPN 1024 (*)	0.498 (0.564)	0.498 (0.589)

Table 11
Performance on training and valid sets for different augmentation policies.

	Pre-proc.	Train MAE (SD)	Val MAE (SD)
No augmentation	y	0.306 (0.490)	0.528 (0.610)
Bright. & contrast	y	0.341 (0.518)	0.550 (0.628)
Geometric transf.	y	0.272 (0.460)	0.541 (0.623)
All together	n	0.437 (0.585)	0.571 (0.645)
All together	y	0.455 (0.578)	0.469 (0.583)

we observe a poorer generalization capability on the validation set. The complex configuration, instead, produces a worsening on both data sets. Therefore FPN confirms to be a good intermediate solution between the latter two configurations. Eventually, adding one resolution level produces again a performance worsening (this is a confident indication despite, for memory limitations, we did not succeed in performing a complete fine-tuning).

Data augmentation The investigated matter regards whether the pre-processing and augmentation policies we adopted are effective, or if there are some prevailing or redundant constituents. In Table 11 we present the results of performed experiments which




Fig. 11. MAE on regions (left) and MAE of the global score (right) versus synthetic rotation. The blue curve is from the network ‘without’ the alignment block, while the orange is ‘with’ the alignment block enabled. The shaded areas correspond to the 95% confidence interval. (For interpretation of the references to color in this figure legend, the reader is referred to the web version of this article.)

combines different augmentation policies: it is clearly evident that partial augmentation policies are not adequate, since they produce performance worsening on the validation set, and tend to create overfitting gaps between training and validation performance. Conversely, their joint use produces best results and not such a gap. Moreover, we can appreciate the impact of the pre-processing equalization, which is able to correctly handle the grayscale variability in the dataset.

Multi-resolution feature aligner In order to demonstrate the robustness of the proposed network to misalignment and the contribution of the adopted multi-resolution feature alignment solution, we augmented the test set by synthetically rotate (from 30 to 30 degrees) all its images and we measured the network performance *with* and *without* the compensation estimated by the alignment block.

Fig. 11 shows how the MAE (on single regions, as well as globally) varies according to the value of the induced rotation angle. Looking at this picture, two important observations arise. First, even without induced rotations (vertical blue line at 0 degrees), the use of the alignment block produces significantly better results, and this demonstrates the effectiveness of the compensation of the real image misalignments originally present in the dataset. Second, the alignment module is able to compensate and improve the performance in a wide range of induced rotations, while without alignment, the influence of rotation angle on performance degradation is evident. Moreover, the flattened error region, approximately corresponding to the range $[-20^\circ, 20^\circ]$, is in line with the ± 25 degree range used for the alignment block pre-training (Table 2), and compatible with the actual range of rotations compensated by the alignment block during testing (we measured they span over a min/max range of ± 15 degrees). Overall, the obtained performance improvement on MAE is about 20–25% with respect to the case without alignment stage. It is therefore clear from these results that the alignment network actually learns to compensate misalignment and its exploitation for multi-resolution feature alignment produces a significant performance boost.

7. Discussion

We have introduced an end-to-end image analysis system for the assessment of a semi-quantitative rating based on a highly difficult visual task. The estimated lung severity score is, by itself, the result of a compromise: on the one hand, the need for a clinically expressive granularity of the different stages of the disease; on the other hand, the built-in subjectivity in the interpretation of CXR images, stemming from the intrinsic limits of such imaging modality, and from the high variability of COVID-19 manifestations. As an additional complication, the available Brixia score, even if coming from expert personnel, has neither ground truth characteristics (as ratings are affected by inter-observer variability), nor it

is highly accurate in terms of spatial indication (since scores are related to generic rectangular regions).

The best performance in the prediction of the Brixia score is obtained, in all tests, by using an ensemble decision strategy (ENS). Reported mean absolute errors on the six regions are: MAE = 0.441, when compared against the clinical annotations by radiologist R_0 on the whole test set of 450 CXRs; and MAE = 0.424 when compared against a Consensus-based Gold Standard set of 150 images annotated by 5 radiologists. Naively speaking, being the scoring scale defined on integers, any MAE measured on single regions below 0.5 could be interpreted as acceptable. This might appear as a simplified reasoning, but interviewed radiologists, with hundreds of cases of experience of such semi-quantitative rating system, also indicated, from a clinical perspective, ± 0.5 as an acceptable error on each region of the Brixia score, and ± 2 as an acceptable error on the Global Score from 0 to 18. These indications are also backed up by the prognostic value and associated use of the score as a severity indicator that comes from the experimental evidence and clinical observations during the first period of its application (Borghesi et al., 2020a).

The above a-priori interpretation of “acceptable” error on a single region is clearly not sufficient. It is also for this reason that we built a CbGS. This is useful to have a reference measure of the inter-rater agreement between human annotators, acting as a boundary measure of human performance. This is relevant since, being a source of error in our weakly supervised approach, such inter-rater agreement also determines an implicit limit to the performance we can expect from the network. Tests on the CbGS confirm that, on average, the level of agreement reached between the network and any other human rater is slightly above the inter-rater agreements between any pair of human raters, thus statistically evidencing that BS-Net performance overcomes radiologists’ in accomplishing the task. This is a fundamental basis to think at clinically oriented studies in a perspective of a responsible deployment of the technology in (human-machine collaborative or computer-aided diagnosis) clinical settings. A MAE under 0.5 for both the network and radiologists is also an indirect evidence of the fact that the Brixia score rating system (on four severity values) is a good trade-off between the two opposite needs of having a fine granularity of the rating scale, and a good inter-rater agreement. Moreover, what comes out from the comparison with other scoring systems is that the Brixia score can be considered as a good super-set with respect to others. This is provided by either a good spatial granularity (over six regions), and a good sensibility (over four levels).

The fact that the ensemble decision strategy combines the strengths of the soft and hard attention policies deserves some further elaboration. In fact, if on one hand, removing the context is fine to avoid possible bias on the decision, on the other hand, the context info can help when the segmentation is too restrictive. For example, retrocardiac consolidations that could be visible are removed from the lung segmentation and therefore their assessment is not allowed in the hard attention approach. This residual complementarity between the two options is exploited by the ensemble decision and this can explain the significant improvement. This also clearly justifies the more structured approach we designed, which demonstrates to pay-off right from the soft-attention configuration. Another aspect that emerges looking at MEr is that our architecture does not over/under-estimate, while the other compared architecture (ResNet-18) tends to overestimate. A last qualitative evidence in favor of the use of the composite loss function of Eq. (1), comprising an additional component related to the MAE, is that the performance we obtain are unbiased and stable over several repetitions of the whole training process.

Although aware that collecting a multi-center database would allow further generalizability tests and possibly increase the sys-

tem robustness and portability, our clinical dataset is already highly representative of a large set of variability parameters: therefore, we do not expect to be varying too much between centers in terms of pneumonia severity range and distribution, adverse image acquisition conditions (ICU devices and patient state conditioned positioning), acquisition devices (fixed vs. portable X-ray equipment, CR vs. DX, different manufacturers), patient age. Moreover, despite the geography of lungs does not depend on human phenotype, our patient sample also reflects the multi-ethnic composition of North-Italy population. Portability tests performed on public datasets can be easily turned into clear guidelines for the use of the proposed model on data originated at different facilities and/or on related clinical contexts. In fact, from collected evidence, there are all the makings of a highly portable model: performance are robust to change of settings within the COVID scenario, while transfer learning on quality data is highly advisable for performance optimization in different domains. Moreover, the results of this study suggests the possibility of using our system for monitoring the severity of potentially any pathology that manifests itself with a similar 'basic syntax' in terms of opacification of the parenchyma. However, the *Brixia score* has been designed specifically for COVID severity assessment, and the only possible direct application of our model is the one concerning other viral pneumonia that lead to a ARDS (adult respiratory distress syndrome) clinical picture, for example that of the H1N1 virus or the one of other coronaviruses. For bacterial / mycobacterial, fungal, and other viral infections, or in case of pulmonary edema (where, unlike COVID pneumonia, pleural effusion dominates), other scoring might be preferable and domain adaptation would be needed.

For what concerns explainability maps, this information can be used to increase trust in the outcome, and point in possible confounding situations. With the proposed super-pixel-based approach we pave the way to new kinds of representations. In particular, we used a three-map representation to show the accurate localization power and the richness of generated information. In the following we account for both specific examples and more general considerations coming from our radiologists about the nature and the use of the proposed explainability solutions, as collected during CXR interpretation activities. Looking for example at Fig. 10, the case on the bottom-left part, despite producing a correct Global Score (considered a positive fact by the radiologists, given the specific case), evidences both under- and over-estimations in single sectors. Reviewed by a senior radiologist, this case evidenced some problematic areas even for the radiologists that are often called to express an average score about what they see in different areas. For this reason the super-pixel explainability maps constitute a useful machine-human communication tool. In particular the sector A interested by the presence of an assisted ventilation device can determine a common problem for both machine an expert, but the machine proves to not diverge toward unduly high ranking, while the expert ruled in favor to the machine in the other upper lung sector D, while confirms a slight overestimation in sector E. In general, we can appreciate the fact that external equipment do not harm the prediction, nor produces unwanted biases as in Fig. 10 (top-middle, bottom-left, bottom-right). The last case (bottom right) evidences a segmentation error on the right lung, since part of the colon (filled by air, pulling up the diaphragm and compressing the lung) enters in the segmentation mask, determining a wrong evaluation of the corresponding area. However, the radiologist positively considered the fact that the system sees a difference of two levels between the upper lung portions, which the expert considered a more coherent judgment at a second view. From these qualitative observations and explainability map support, we can conclude that, despite the automated AI-driven scoring is not meant to eliminate radiologists' evaluations, it can be used to aid and streamline the reporting workflow, and improving the timeli-

ness of first evaluation, by proposing a preliminary interpretation of findings.

Starting from mid December 2020 the proposed system is undergoing an experimental deployment in the two Diagnostic Radiology Units of ASST Spedali Civili of Brescia (Italy), where all radiologists can see the severity score estimation and explainability maps during the reporting of all CXRs from incoming (admission/ER triage) and hospitalized COVID-19 patients. In general, we had a valuable positive general feedback from radiologists that considered these maps as a possible complementary source of visual attention and of second level reasoning in a possible scenario of computer-assisted diagnosis. In the following we quote the opinion of a senior radiologist in the team, when interviewed about the proposed explainability approach: "the maps concentrate in a small space a huge amount of information; in the overburdened workflow of a radiology department (not only during a pandemic, but also in the routine practice) it is quite unlikely that an experienced physician interrogates such a map. The reporting time of a CXR is in the range of a few seconds, thus the assessment of the explainability maps in every single case would significantly delay the reporting process; on the other hand, the maps are far from being useless. First of all, the neural network and the radiologist examine the same object (a CXR) using completely different perspectives and approaches; the map may be conceived as a mediator that allows subjects that speak different languages to communicate. In addition, in selected and complex cases (particularly during a multidisciplinary discussion) the maps may support the decision offering a valuable and non-subjective second opinion. Finally, in an academic setting the maps may guide young radiologists in identifying the abnormalities that sum up to a final diagnosis and to a definition of severity." Thus, even if explainability maps can be sometimes complex to interpret, finding good representations able to effectively convey visual information and to act as a mediator between clinicians and the AI model is of recognized high relevance. Combined with the fact that explainability solutions have an important role in emerging regulatory framework for trustable AI, it is our intention to ground on interesting properties of our super-pixel-based maps to further explore, in a dedicated work, their optimization and adoption according to user experience in clinical practice.

The main limitations of this work are related to residual errors. Despite highly aligned with radiologists' performance and not evidencing statistical biases, a case-by-case deeper comparison of possible causes of error could guide further improvement of the image analysis architecture. In particular, since the segmentation task is sometimes ambiguous and hard to accomplish, better solutions should be designed to handle anomalies in the segmentation results.

8. Conclusions and future directions

Our study for the estimation of a semi-quantitative rating of lung severity is justified and driven by the strong clinical reasons related to the role of CXR images for the management of COVID-19 patient monitoring, especially in conditions of overloaded health-care facilities. Working with a very large dataset of almost 5000 images allowed to develop a solution which deals with all aspects of the problem, also thanks to other datasets which are exploited in a hierarchical end-to-end training process. The prospected solution is designed to work in a weakly supervised context and to be capable of manifesting self-attentive abilities on data directly coming from all CXR modalities/devices and in all clinical and patient care conditions. Having tested architectural variants, and targeted ablation studies, the network performance ultimately surpasses qualified radiologists in rating accuracy and consistency. We

also collected evidences of the robustness and portability of the model to other clinical settings.

Regarding future work, observed results constitute a strong basis in a perspective of more clinically oriented validations and in a perspective of trustable and responsible deployment. An interesting clinical scenario to evaluate is the following: whether, in case the BS-Net follows the same patient, it could exhibit greater self-coherent behavior if compared to the case where serial CXR acquisitions are reported by different radiologists (according to availability and working shifts). Another direction to consider is the possibility to exploit CT images from COVID-19 patients to derive semi-quantitative ground truth information directly from the quantitative volumetric assessments.

Source code and data availability

All annotated CXRs and the proposed system (both source code and pre-trained models) are made available for research purposes at the BrixIA project page <https://brixia.github.io>.

Declaration of Competing Interest

The authors declare that they have no known competing financial interests or personal relationships that could have appeared to influence the work reported in this paper.

CRediT authorship contribution statement

Alberto Signoroni: Conceptualization, Methodology, Validation, Formal analysis, Investigation, Writing - original draft, Writing - review & editing, Visualization, Supervision, Project administration, Funding acquisition. **Mattia Savardi:** Conceptualization, Methodology, Software, Validation, Investigation, Writing - original draft, Writing - review & editing, Visualization. **Sergio Benini:** Conceptualization, Writing - original draft, Writing - review & editing. **Nicola Adami:** Validation, Formal analysis, Writing - review & editing, Visualization. **Riccardo Leonardi:** Writing - review & editing. **Paolo Gibellini:** Investigation, Resources. **Filippo Vaccher:** Validation, Data curation. **Marco Ravanelli:** Data curation. **Andrea Borghesi:** Data curation, Resources. **Roberto Maroldi:** Resources, Data curation, Supervision, Project administration, Funding acquisition. **Davide Farina:** Conceptualization, Investigation, Resources, Data curation, Writing - original draft, Writing - review & editing, Supervision.

Acknowledgments

The authors would like to thank colleagues of the Information and IT System Operational Units of ASST Spedali Civili di Brescia, Francesco Scazzoli, Silvio Finardi, Roberto Marini, and Sabrina Vicari for providing the infrastructure. A very special thanks to Marco Renato Roberto Merli, and Andrea Carola of Philips S.p.A., and to Fabio Capra, Daniele Soldi, Gian Stefano Bosio, and Simone Gaggero of ELCO. S.r.l (Cairo Montenotte, Italy) for their outstanding technical support.

References

Blain, M., Kassim, M.T., Varble, N., Wang, X., Xu, Z., Xu, D., Carrafiello, G., Vespro, V., Stellato, E., Ierardi, A.M., et al., 2021. Determination of disease severity in COVID-19 patients using deep learning in chest X-ray images. *Diagn. Interv. Radiol.* 27 (1), 20–27.

Bontempi, D., Benini, S., Signoroni, A., Svanera, M., Muckli, L., 2020. CEREBRUM: a fast and fully-volumetric convolutional Encoder-decoder for weakly-supervised segmentation of BRain structures from out-of-the-scanner MRI. *Med. Image Anal.* 62, 101688.

Borghesi, A., Maroldi, R., 2020. COVID-19 outbreak in Italy: experimental chest X-ray scoring system for quantifying and monitoring disease progression. *Radiol. Med.* 125, 509–513.

Borghesi, A., Zigliani, A., Golemi, S., Carapella, N., Maculotti, P., Farina, D., Maroldi, R., 2020. Chest X-ray severity index as a predictor of in-hospital mortality in coronavirus disease 2019: a study of 302 patients from Italy. *Int. J. Infect. Dis.* 96, 291–293.

Amer, R., Frid-Adar, M., Gozes, O., Nassar, J., Greenspan, H., 2020. COVID-19 in CXR: from detection and severity scoring to patient disease monitoring. *arXiv:2008.02150* doi: [10.1109/JBHI.2021.3069169](https://doi.org/10.1109/JBHI.2021.3069169).

Burlacu, A., Crisan-Dabija, R., Popa, I. V., Artene, B., Birzu, V., Pricop, M., Plesoianu, C., Generali, D., 2020. Curbing the AI-induced enthusiasm in diagnosing COVID-19 on chest X-rays: the present and the near-future. *medRxiv*.

Borghesi, A., Zigliani, A., Masciullo, R., Golemi, S., Maculotti, P., Farina, D., Maroldi, R., 2020. Radiographic severity index in COVID-19 pneumonia: relationship to age and sex in 783 Italian patients. *Radiol. Med.* 125, 461–464.

Buslaev, A., Igloukov, V.I., Khvedchenya, E., Parinov, A., Druzhinin, M., Kalinin, A.A., 2020. Alumentations: fast and flexible image augmentations. *Information* 11 (2), 125.

Candemir, S., Antani, S., 2019. A review on lung boundary detection in chest X-rays. *Int. J. Comput. Assist. Radiol. Surg.* 14 (4), 563–576.

Candemir, S., Jaeger, S., Palaniappan, K., Musco, J.P., Singh, R.K., Xue, Z., Karargyris, A., Antani, S., Thoma, G., McDonald, C.J., 2014. Lung segmentation in chest radiographs using anatomical atlases with nonrigid registration. *IEEE Trans. Med. Imaging* 33 (2), 577–590.

Castiglioni, I., Ippolito, D., Interlenghi, M., Monti, C. B., Salvatore, C., Schiaffino, S., Polidori, A., Gandola, D., Messa, C., Sardanelli, F., 2020. Artificial intelligence applied on chest X-ray can aid in the diagnosis of COVID-19 infection: a first experience from lombardy, Italy. *medRxiv*.

Cohen, J. P., Dao, L., Morrison, P., Roth, K., Bengio, Y., Shen, B., Abbasi, A., Hoshmand-Kochi, M., Ghassemi, M., Li, H., Duong, T. Q., 2020a. Predicting COVID-19 pneumonia severity on chest X-ray with deep learning. *arXiv:2005.11856*.

Cohen, J.P., Morrison, P., Dao, L., Roth, K., Duong, T.Q., Ghassemi, M., 2020. COVID-19 image data collection: prospective predictions are the future. *J. Mach. Learn. Biomed. Imaging* 2, 1–38.

Toussie, D., Voutsinas, N., Finkelstein, M., Cedillo, M.A., Manna, S., Maron, S.Z., Jacobi, A., Chung, M., Bernheim, A., Eber, C., Concepcion, J., Fayad, Z., Gupta, Y.S., 2020. Clinical and chest radiography features determine patient outcomes in young and middle age adults with COVID-19. *Radiology* 297 (1), E197–E206.

van Ginneken, B., Stegmann, M.B., Loog, M., 2006. Segmentation of anatomical structures in chest radiographs using supervised methods: a comparative study on a public database. *Med. Image Anal.* 10 (1), 19–40.

Cohen, J. P., Morrison, P., Dao, L., 2020b. COVID-19 image data collection. *arXiv:2003.11597*, <https://github.com/ieee8023/covid-chestxray-dataset>.

Glasmachers, T., 2017. Limits of End-to-End Learning. *arXiv:1704.08305*.

Gozes, O., Frid-Adar, M., Greenspan, H., Browning, P. D., Zhang, H., Ji, W., Bernheim, A., Siegel, E., 2020. Rapid ai development cycle for the coronavirus (COVID-19) pandemic: initial results for automated detection and patient monitoring using deep learning CT image analysis. *arXiv:2003.05037*.

Frid-Adar, M., Ben-Cohen, A., Amer, R., Greenspan, H., 2018. Improving the segmentation of anatomical structures in chest radiographs using U-net with an imagenet pre-trained encoder. In: *Image Analysis for Moving Organ, Breast, and Thoracic Images*. Springer, pp. 159–168.

He, K., Zhang, X., Ren, S., Sun, J., 2016. Deep residual learning for image recognition. In: *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*, pp. 770–778.

Hryniewska, W., Bombiński, P., Szatkowski, P., Tomaszewska, P., Przelaskowski, A., Biecek, P., 2020. Do not repeat these mistakes – a critical appraisal of applications of explainable artificial intelligence for image based COVID-19 detection. *arXiv:2012.08333*.

Huang, L., Han, R., Ai, T., Yu, P., Kang, H., Tao, Q., Xia, L., 2020. Serial quantitative chest CT assessment of COVID-19: deep-learning approach. *Radiol.: Cardiothorac. Imag.* 2 (2), e200075.

Huang, G., Liu, Z., Van Der Maaten, L., Weinberger, K.Q., 2017. Densely connected convolutional networks. In: *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*, pp. 4700–4708.

Irvin, J., Rajpurkar, P., Ko, M., Yu, Y., Ciurea-Illcus, S., Chute, C., Marklund, H., Haghgoo, B., Ball, R., Shpanskaya, K., et al., 2019. Chexpert: a large chest radiograph dataset with uncertainty labels and expert comparison. In: *Proceedings of the AAAI Conf. on Artificial Intelligence*, 33, pp. 590–597.

Jaderberg, M., Simonyan, K., Zisserman, A., Kavukcuoglu, K., 2015. Spatial transformer networks. In: *Advances in Neural Information Processing Systems*, 28. Curran Associates, Inc., pp. 2017–2025.

Jaeger, S., Candemir, S., Antani, S., Wang, Y.-X.J., Lu, P.-X., Thoma, G., 2014. Two public chest X-ray datasets for computer-aided screening of pulmonary diseases. *Quant. Imaging Med. Surg.* 4 (6), 475.

Kalkreuth, R., Kaufmann, P., 2020. COVID-19: a survey on public medical imaging data resources. *arXiv:2004.04569*.

Karim, M. R., Dhmen, T., Rehbholz-Schuhmann, D., Decker, S., Cochez, M., Beyan, O., 2020. DeepCOVIDexplainer: explainable COVID-19 predictions based on chest X-ray images. *arXiv:2004.04582*.

Karimi, D., Dou, H., Warfield, S. K., Gholipour, A., 2019. Deep learning with noisy labels: exploring techniques and remedies in medical image analysis. *arXiv:1912.02911*.

Kingma, D. P., Ba, J., 2014. Adam: a method for stochastic optimization. *arXiv preprint arXiv:1412.6980*.

Kundu, S., Elhalawani, H., Gichoya, J.W., Kahn, C.E., 2020. How might ai and chest imaging help unravel COVID-19' mysteries? *Radiol.: Artif. Intell.* 2 (3), e200053.

- Laghi, A., 2020. Cautions about radiologic diagnosis of COVID-19 infection driven by artificial intelligence. *Lancet Digit. Health* 2 (5), e225.
- Latif, S., Usman, M., Manzoor, S., Iqbal, W., Qadir, J., Tyson, G., Castro, I., Razi, A., Boulos, M.N.K., Weller, A., Crowcroft, J., 2020. Leveraging data science to combat COVID-19: a comprehensive review. *IEEE Trans. Artif. Intell.* 1 (1), 85–103.
- Lessmann, N., Sanchez, C.I., Beenen, L., Boulogne, L.H., Brink, M., Calli, E., Charbonnier, J.-P., Dofferhoff, T., van Everdingen, W.M., Gerke, P.K., Geurts, B., Gieterma, H.A., Groeneveld, M., van Harten, L., Hendrix, N., Hendrix, W., Huisman, H.J., Igum, I., Jacobs, C., Kluge, R., Kok, M., Krdzalic, J., Lassen-Schmidt, B., van Leeuwen, K., Meakin, J., Overkamp, M., van Rees Vellinga, T., van Rikxoort, E.M., Samperna, R., Schaefer-Prokop, C., Schalekamp, S., Scholten, E.T., Sital, C., Stager, L., Teuwen, J., Vaidhya Venkadesh, K., de Vente, C., Vermaat, M., Xie, W., de Wilde, B., Prokop, M., van Ginneken, B., 2021. Automated assessment of CO-RADS and chest CT severity scores in patients with suspected COVID-19 using artificial intelligence. *Radiology* 298 (1), E18–E28. doi:10.1148/radiol.20202439. PMID: 32729810
- Li, M. D., Arun, N. T., Aggarwal, M., Gupta, S., Singh, P., Little, B. P., Mendoza, D. P., Corradi, G. C. A., Takahashi, M. S., Ferracioli, S. F., Succi, M. D., Lang, M., Bizzo, B. C., Dayan, I., Kitamura, F. C., Kalpathy-Cramer, J., 2020a. Improvement and multi-population generalizability of a deep learning-based chest radiograph severity score for COVID-19. medRxiv. 10.1101/2020.09.15.20195453.
- Li, M.D., Arun, N.T., Gidwani, M., Chang, K., Deng, F., Little, B.P., Mendoza, D.P., Lang, M., Lee, S.I., OShea, A., Parakh, A., Singh, P., Kalpathy-Cramer, J., 2020. Automated assessment and tracking of COVID-19 pulmonary disease severity on chest radiographs using convolutional siamese neural networks. *Radiology* 2 (4), e200079.
- Li, X., Li, C., Zhu, D., 2020c. COVID-MobileXpert: On-device COVID-19 screening using snapshots of chest X-ray. arXiv:2004.03042.
- Linda Wang, Z. Q. L., Wong, A., 2020. COVID-Net: A tailored deep convolutional neural network design for detection of COVID-19 cases from chest radiography images. arXiv:2003.09871.
- Lin, T.-Y., Dollar, P., Girshick, R., He, K., Hariharan, B., Belongie, S., 2017. Feature Pyramid Networks for Object Detection. In: *The IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*. IEEE, pp. 936–944.
- Long, J., Shelhamer, E., Darrell, T., 2015. Fully convolutional networks for semantic segmentation. In: *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*, pp. 3431–3440.
- Maguolo, G., Nanni, L., 2020. A Critic Evaluation of Methods for COVID-19 Automatic Detection from X-Ray Images. arXiv:2004.12823.
- Manna, S., Wruble, J., Maron, S., Toussie, D., Voutsinas, N., Finkelstein, M., Cedillo, M.A., Diamond, J., Eber, C., Jacobi, A., Chung, M., Bernheim, A., 2020. COVID-19: a multimodality review of radiologic techniques, clinical utility, and imaging features. *Radiol.: Cardiothorac. Imag.* 2 (3), e200210.
- Maroldi, R., Rondi, P., Agazzi, G.M., Ravanelli, M., Borghesi, A., Farina, D., 2020. Which role for chest X-ray score in predicting the outcome in COVID-19 pneumonia? *Eur. Radiol.* 1–7. doi:10.1007/s00330-020-07504-2.
- Minaee, S., Kafieh, R., Sonka, M., Yazdani, S., Jamalipour Soufi, G., 2020. Deep-COVID: predicting COVID-19 from chest X-ray images using deep transfer learning. *Med. Image Anal.* 65, 101794.
- Minaee, S., Kafieh, R., Sonka, M., Yazdani, S., Jamalipour Soufi, G., 2020. Deep-COVID: predicting COVID-19 from chest X-ray images using deep transfer learning. *Med. Image Anal.* 65, 101794.
- Oh, Y., Park, S., Ye, J.C., 2020. Deep learning COVID-19 features on CXR using limited training data sets. *IEEE Trans. Med. Imaging* 39 (8), 2688–2700.
- Rajaraman, S., Siegelman, J., Alderson, P. O., Folio, L. S., Folio, L. R., Antani, S. K., 2020. Iteratively pruned deep learning ensembles for COVID-19 detection in chest X-rays. arXiv:2004.08379.
- Ramachandran, P., Zoph, B., Le, Q. V., 2017. Searching for activation functions. arXiv preprint arXiv:1710.05941.
- Pereira, R.M., Bertolini, D., Teixeira, L.O., Silla, C.N., Costa, Y.M.G., 2020. COVID-19 identification in chest X-ray images on flat and hierarchical classification scenarios. *Comput. Methods Prog. Biomed.* 194, 105532.
- Reyes, M., Meier, R., Pereira, S., Silva, C.A., Dahlweid, F.-M., Tengge-Kobligk, H.v., Summers, R.M., Wiest, R., 2020. On the interpretability of artificial intelligence in radiology: challenges and opportunities. *Radiol.: Artif. Intell.* 2 (3), e190043.
- Ribeiro, M.T., Singh, S., Guestrin, C., 2016. why should i trust you?: explaining the predictions of any classifier. In: *Proceedings of the 22nd ACM SIGKDD International Conference on Knowledge Discovery and Data Mining*. Association for Computing Machinery, New York, NY, USA, p. 11351144.
- Ronneberger, O., Fischer, P., Brox, T., 2015. U-net: convolutional networks for biomedical image segmentation. In: *International Conference on Medical Image Computing and Computer-Assisted Intervention*. Springer, pp. 234–241.
- Rubin, G.D., Ryerson, C.J., Haramati, L.B., Sverzellati, N., Kanne, J.P., Raouf, S., Schluger, N.W., Volpi, A., Yim, J.-J., Martin, I.B.K., Anderson, D.J., Kong, C., Altes, T., Bush, A., Desai, S.R., Goldin, J., Goo, J.M., Humbert, M., Inoue, Y., Kauczor, H.-U., Luo, F., Mazzone, P.J., Prokop, M., Remy-Jardin, M., Richeldi, L., Schaefer-Prokop, C.M., Tomiyama, N., Wells, A.U., Leung, A.N., 2020. The role of chest imaging in patient management during the COVID-19 pandemic: a multinational consensus statement from the Fleischner society. *Radiology (Simultaneously Published in Chest)* 296 (1), 172–180.
- Sardanelli, F., Di Leo, G., 2020. Assessing the value of diagnostic tests in the new world of COVID-19 pandemic. *Radiology* 296 (3), E193–E194.
- Selvaraju, R.R., Cogswell, M., Das, A., Vedantam, R., Parikh, D., Batra, D., 2017. Grad-CAM: visual explanations from deep networks via gradient-based localization. In: *The IEEE International Conference on Computer Vision (ICCV)*. IEEE, pp. 618–626.
- Shi, F., Wang, J., Shi, J., Wu, Z., Wang, Q., Tang, Z., He, K., Shi, Y., Shen, D., 2021. Review of artificial intelligence techniques in imaging data acquisition, segmentation and diagnosis for COVID-19. *IEEE Rev. Biomed. Eng.* 14, 4–15.
- Shiraishi, J., Katsuragawa, S., Ikezoe, J., Matsumoto, T., Kobayashi, T., Komatsu, K.-i., Matsui, M., Fujita, H., Kodera, Y., Doi, K., 2000. 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. *Am. J. Roentgenol.* 174 (1), 71–74.
- Simonyan, K., Zisserman, A., 2014. Very deep convolutional networks for large-scale image recognition. arXiv preprint arXiv:1409.1556.
- Stirenko, S., Kochura, Y., Alienin, O., Rokovy, O., Gordienko, Y., Gang, P., Zeng, W., 2018. Chest X-ray analysis of tuberculosis by deep learning with segmentation and augmentation. In: *2018 IEEE 38th International Conference on Electronics and Nanotechnology (ELNANO)*. IEEE, pp. 422–428.
- Summers, R.M., 2021. Artificial intelligence of COVID-19 imaging: a hammer in search of a nail. *Radiology (Published Online)* 298 (3), E162–E164. PMID: 33350895
- Szegedy, C., Ioffe, S., Vanhoucke, V., Alemi, A.A., 2017. Inception-v4, inception-resnet and the impact of residual connections on learning. In: *Thirty-First AAAI Conference on Artificial Intelligence*. ACM, pp. 4278–4284.
- Tajbakhsh, N., Jeyaseelan, L., Li, Q., Chiang, J.N., Wu, Z., Ding, X., 2020. Embracing imperfect datasets: A review of deep learning solutions for medical image segmentation. *Med. Image Anal.* 63, 101693.
- Tan, M., Pang, R., Le, Q. V., 2019. EfficientDet: Scalable and Efficient Object Detection. arXiv:1911.09070.
- Tartaglione, E., Barbano, C. A., Berzovini, C., Calandri, M., Grangetto, M., 2020. Unveiling COVID-19 from chest X-ray with deep learning: a hurdles race with small data. arXiv:2004.05405.
- Ting, D.S.W., Carin, L., Dzau, V., Wong, T.Y., 2020. Digital technology and COVID-19. *Nat. Med.* 26 (4), 459–461.
- Vedaldi, A., Soatto, S., 2008. Quick shift and kernel methods for mode seeking. In: *Forsyth, D., Torr, P., Zisserman, A. (Eds.), European Conference of Computer Vision (ECCV)*. Springer Berlin Heidelberg, Berlin, Heidelberg, pp. 705–718.
- Wang, X., Peng, Y., Lu, L., Lu, Z., Bagheri, M., Summers, R.M., 2017. ChestX-ray8: Hospital-scale chest X-ray database and benchmarks on weakly-supervised classification and localization of common thorax diseases. In: *2017 IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*, pp. 3462–3471.
- Warren, M.A., Zhao, Z., Koyama, T., Bastarache, J.A., Shaver, C.M., Semler, M.W., Rice, T.W., Matthay, M.A., Calfee, C.S., Ware, L.B., 2018. Severity scoring of lung oedema on the chest radiograph is associated with clinical outcomes in ARDS. *Thorax* 73 (9), 840–846.
- WHO, 2020. Coronavirus disease (COVID-19) outbreak. World Health Organization. <https://www.who.int/emergencies/diseases/novel-coronavirus-2019>.
- Wong, A., Lin, Z. Q., Wang, L., Chung, A. G., Shen, B., Abbasi, A., Hoshmand-Kochi, M., Duong, T. Q., 2020a. COVIDNet-S: Towards computer-aided severity assessment via training and validation of deep neural networks for geographic extent and opacity extent scoring of chest X-rays for SARS-CoV-2 lung disease severity. arXiv:2005.12855.
- Wong, H.Y.F., Lam, H.Y.S., Fong, A.H.-T., Leung, S.T., Chin, T.W.-Y., Lo, C.S.Y., Lui, M.M.-S., Lee, J.C.Y., Chiu, K.W.-H., Chung, T., Lee, E.Y.P., Wan, E.Y.F., Hung, F.N.L., Lam, T.P.W., Kuo, M., Ng, M.-Y., 2020b. Frequency and distribution of chest radiographic findings in COVID-19 positive patients. *Radiology* 296 (2), E72–E78.
- Xu, Y., Zhu, J.-Y., Chang, E.I.-C., Lai, M., Tu, Z., 2014. Weakly supervised histopathology cancer image segmentation and classification. *Med. Image Anal.* 18 (3), 591–604.
- Zhou, Z., Rahman Siddiquee, M.M., Tajbakhsh, N., Liang, J., 2018. UNet++: a nested U-Net architecture for medical image segmentation. In: *Deep Learning in Medical Image Analysis and Multimodal Learning for Clinical Decision Support*. Springer International Publishing, Cham, pp. 3–11.
- Zhou, Z.-H., 2017. A brief introduction to weakly supervised learning. *Natl. Sci. Rev.* 5 (1), 44–53.
- Zhu, J., Shen, B., Abbasi, A., Hoshmand-Kochi, M., Li, H., Duong, T.Q., 2020. Deep transfer learning artificial intelligence accurately stages COVID-19 lung disease severity on portable chest radiographs. *PLoS One* 15 (7), 1–11.