

Fibrosis-Net: A Tailored Deep Convolutional Neural Network Design For Prediction of Pulmonary Fibrosis Progression From Chest CT Images

Alexander Wong (✉ a28wong@uwaterloo.ca)

University of Waterloo

Jack Lu

DarwinAI Corp

Adam Dorfman

DarwinAI Corp

Paul McInnis

DarwinAI Corp

Mahmoud Famouri

DarwinAI Corp

Daniel Manary

DarwinAI Corp

James Ren Hou Lee

DarwinAI Corp

Michael Lynch

DarwinAI Corp

Research Article

Keywords: pulmonary fibrosis, convolutional, neural network, computed tomography

Posted Date: June 21st, 2021

DOI: <https://doi.org/10.21203/rs.3.rs-610010/v1>

License: © ⓘ This work is licensed under a Creative Commons Attribution 4.0 International License.

[Read Full License](#)

RESEARCH

Fibrosis-Net: A Tailored Deep Convolutional Neural Network Design for Prediction of Pulmonary Fibrosis Progression from Chest CT Images

Alexander Wong^{1,2,3*}, Jack Lu³, Adam Dorfman³, Paul McInnis³, Mahmoud Famouri³, Daniel Manary³, James Ren Hou Lee³ and Michael Lynch³

*Correspondence:

a28wong@uwaterloo.ca

¹ Vision and Image Processing Research Group, University of Waterloo, Waterloo, Canada

² Waterloo Artificial Intelligence Institute, University of Waterloo, Waterloo, Canada

³ DarwinAI Corp., Waterloo, Canada

Full list of author information is available at the end of the article

Abstract

Background:

Pulmonary fibrosis is a devastating chronic lung disease that causes irreparable lung tissue scarring and damage, resulting in progressive loss in lung capacity and has no known cure. A critical step in the treatment and management of pulmonary fibrosis is the assessment of lung function decline, with computed tomography (CT) imaging being a particularly effective method for determining the extent of lung damage caused by pulmonary fibrosis. Motivated by this, we introduce Fibrosis-Net, a deep convolutional neural network design tailored for the prediction of pulmonary fibrosis progression from chest CT images. More specifically, machine-driven design exploration was leveraged to determine a strong architectural design for CT lung analysis, upon which we build a customized network design tailored for predicting forced vital capacity (FVC) based on a patient's CT scan, initial spirometry measurement, and clinical metadata. Finally, we leverage an explainability-driven performance validation strategy to study the decision-making behaviour of Fibrosis-Net as to verify that predictions are based on relevant visual indicators in CT images.

Results:

Experiments using a patient cohort from the OSIC Pulmonary Fibrosis Progression Challenge showed that the proposed Fibrosis-Net is able to achieve a significantly higher modified Laplace Log Likelihood score than the winning solutions on the challenge. Furthermore, explainability-driven performance validation demonstrated that the proposed Fibrosis-Net exhibits correct decision-making behaviour by leveraging clinically-relevant visual indicators in CT images when making predictions on pulmonary fibrosis progress.

Conclusion:

Fibrosis-Net is able to achieve a significantly higher modified Laplace Log Likelihood score than the winning solutions on the OSIC Pulmonary Fibrosis Progression Challenge, and has been shown to exhibit correct decision-making behaviour when making predictions. Fibrosis-Net is available to the general public in an open-source and open access manner as part of the OpenMedAI initiative. While Fibrosis-Net is not yet a production-ready clinical assessment solution, we hope that its release will encourage researchers, clinicians, and citizen data scientists alike to leverage and build upon it.

Keywords: pulmonary fibrosis; convolutional; neural network; computed tomography

Introduction

Pulmonary fibrosis is a serious chronic lung disease in which permanent tissue scarring and damage occurs in the lungs. The increasing replacement of healthy lung tissues with fibrotic tissue results in progressive, irreversible reduction in lung function over time. There are currently no known cure and limited treatment options for pulmonary fibrosis, with treatment and management of the disease focused on the attenuation of lung function decline progression and improving quality of life. The rate of progression for a patient with pulmonary fibrosis is highly variable, ranging from little to no change over many years to rapid deterioration in a short period of time.

A very critical step in the treatment and management of pulmonary fibrosis is the assessment of lung function decline, as this guides clinicians to determining the best course of treatment and management ranging from oxygen therapy and pulmonary rehabilitation to pharmacological agents (pirfenidone [1] and nintedanib [2]) and lung transplantation [3]. Guidelines set out by the ATS/ERS/JRS/ALAT [4] describe a number of methods for diagnosis of pulmonary fibrosis. Invasive techniques, such as surgical lung biopsy, have associated risks to the health and lung function for the patient [5]. Transbronchial lung biopsy is a less invasive technique where small samples of the lung tissue are taken using video-assisted thoracoscopy or flexible bronchoscopy [6]. A number of methods have been utilized by clinicians for assessing lung function decline after diagnosis. For example, spirometry tests are frequently leveraged for measuring the FVC of the lung which is a key indicator of lung function [7, 8, 9, 10].

One of the most effective methods for assessing lung function decline and the extent of lung damage due to pulmonary fibrosis is computed tomography (CT) imaging. Several visual signs in CT scans have been identified and leveraged by radiologists to assess lung function decline from pulmonary fibrosis, with the most common visual indicator being honeycombing, which present as cystic spaces with irregularly thickened fibrotic tissue walls [11]. However, given the rate of progression of different patients can be highly variable, the ability to accurately predict the progress of pulmonary fibrosis remains a major challenge. This is further compounded by the fact that some common visual indicators such as honeycombing may not be present at certain stages of progression or even at all [12], while other atypical patterns mimicking other diseases may be present in the CT scans instead (e.g., predominance of ground-glass opacity, consolidation, nodules, and atypical distribution of lesions [13], as well as the presence of ground-glass attenuation [14]). As such, new methods for improving prediction accuracy when leveraging CT images as a tool for assessing lung function decline due to pulmonary fibrosis is highly desired.

Motivated by the potential of machine learning for computer-aided clinical decision support for pulmonary fibrosis, in this study we introduce Fibrosis-Net, a deep convolutional neural network design tailored specifically for the prediction of pulmonary fibrosis progression from chest CT images. More specifically, machine-driven design exploration was leveraged to determine a strong architectural design for CT lung analysis, upon which we build a customized network design tailored for predicting forced vital capacity (FVC) based on a patient's CT scan, initial spirometry

measurement, and clinical metadata. Furthermore, to explore the decision-making behaviour of Fibrosis-Net, we leverage an explainability-driven performance validation strategy to audit Fibrosis-Net to verify that predictions are based on relevant visual indicators in CT images. Fibrosis-Net is available to the general public in an open-source and open access manner as part of the OpenMedAI initiative, an open source initiative for medical artificial intelligence solutions that currently include the COVID-Net [15, 16, 17, 18, 19] initiative, Cancer-Net [20] initiative, and the TB-Net initiative [21]. While Fibrosis-Net is not yet a production-ready screening solution, we hope that its open source release will encourage researchers, clinicians, and citizen data scientists alike to leverage and build upon them.

The paper is organized as follows. Related work in the area of artificial intelligence for computer-aided clinical decision support for pulmonary fibrosis is discussed in Section . Section provides a detailed description of the data preparation and analysis process, the architecture design construction process, the proposed Fibrosis-Net network architecture design, and the explainability-driven performance validation process. Section presents both the quantitative performance validation results evaluating the efficacy of the proposed Fibrosis-Net using a patient cohort from the OSIC Pulmonary Fibrosis Progression Challenge [22] as well as visual validation results from the explainability-driven performance validation process used to study the decision-making behaviour of Fibrosis-Net. Finally, Conclusions are drawn and future work is discussed in Section and the broader impact of the proposed work is discussed.

Related Work

Motivated by the significant benefits that can be gained and the challenges involved from a clinical perspective, there has been a recent interest in leveraging artificial intelligence for computer-aided clinical decision support of pulmonary fibrosis based on CT images [23, 24, 25, 26]. For instance, Anthimopoulos et al. [27] use a deep convolutional neural network to analyse 2D patches from the CT image to classify reticulation, honeycombing, ground glass opacity (GGO), consolidation, and micronodules in lung tissue. Christodoulidis et al. [28] proposed a multi-source transfer learning approach with deep convolutional neural networks pre-trained with a selection of texture data sets, again with the goal of classifying 2D image patches of lung tissue in CT. Bermejo-Peláez et al. [29] describe a method using an ensemble of deep convolutional neural networks where the output of each network is summed up and weighted before being combined to form the overall output of the ensemble.

More recently, the significant potential and need for advancements in artificial intelligence-driven methods for computer-aided clinical decision support of pulmonary fibrosis was exemplified by the Kaggle Pulmonary Fibrosis Progression Challenge [22] launched by the Open Source Imaging Consortium (OSIC) to get the research community to accelerate advancement of machine learning for pulmonary fibrosis assessment. To the best of the authors' knowledge, the patient cohort curated by OSIC for the challenge is the largest publicly available cohort in literature. Amongst the many artificial intelligence solutions introduced as part of the challenge, the 1st place winning solution [22] proposed a weighted ensemble between a deep convolutional neural network with a state-of-the-art EfficientNet-B5 network architecture design [30] and a multiple quantile regressor to predict the

lung function decline of a patient based on a patient's CT scans, initial spirometry measurement, and clinical metadata.

To the best of the authors' knowledge, the notion of explainability-driven performance validation of the decision-making behaviour of artificial intelligence solutions for the prediction of pulmonary fibrosis progression have not been previously explored in literature and can be very valuable for driving greater clinical adoption of such solutions in a transparent and trusted manner. Furthermore, while explainability-driven performance validation strategies has been demonstrated to be very successful in past studies for the purpose of clinical classification [15, 17, 18, 21], the the best of the authors' knowledge this is the first study in literature to successfully leverage explainability-driven performance validation on a clinical regression problem.

Methods

Data preparation and analysis

To build the proposed Fibrosis-Net, we leverage the patient cohort from the OSIC Pulmonary Fibrosis Progression Challenge [22]. The data for this patient cohort consists of chest CT scans, forced vital capacity (FVC) measurements from frequent visits over the course of around 1-3 years, and associated clinical metadata (i.e., age, sex, smoking status, and patient's relative FVC measurement compared to the typical FVC measurement of a patient with similar characteristics). More specifically, in this study, the training set consists of 172 patient cases, while the test set consists of approximately 28 patient cases. To the best of the authors' knowledge, the patient cohort curated for this challenge is the largest publicly available cohort in literature.

Table 1 summarizes the demographic variables of the data from the patient cohort in the OSIC Pulmonary Fibrosis Progression Challenge [22] used in this study. A number of observations can be made based on the demographic distribution analysis. First, it can be observed that the patient cases in the cohort are distributed across the different age groups, with the mean age being 67.1 and the highest number of patients in the cohort are between the ages of 60-69. This distributional trend towards older adults is reflective of the fact that pulmonary fibrosis is typically diagnosed later in life due to the condition worsening as time goes by. Second, it can be observed that the majority of patient cases in the cohort are individuals who have smoked sometime in their life, with the majority of patients being ex-smokers. This distributional trend is reflective of the fact that a recognized risk factor for the development of pulmonary fibrosis is smoking. Furthermore, smoking can lead to more significant detrimental effects that can reduce the survival rate of patients with pulmonary fibrosis. Third, it can be observed that a majority of the patient cases in the cohort are male. This is consistent with clinical studies in literature showing that gender differences in the risk of pulmonary fibrosis [31] as well as at clinical presentation [32]. Finally, it can be observed that the patient cases span a wide range of relative FVC measurement values, which indicates different levels of lung function decline. This distribution diversity in relative FVC measurement values is desirable as it allows the deep neural network being trained to be exposed to a wider variety of pulmonary fibrosis progression scenarios for greater generalizability for different patient conditions.

A number of pre-processing steps were conducted to improve the consistency and quality of the CT images from the patient scans. More specifically, all CT imaging data was translated to Hounsfield units (HU), and windowing was performed with a window level of -650 HU and a window width of 1700 HU to better focus on clinically relevant lung features. Furthermore, synthetic padding and circular artifacts found in the CT imaging data for several patient cases within the patient cohort are mitigated to reduce the likelihood of erroneous visual features from being learned as predictive indicators. Finally, calibration value errors found in the data for several patient cases within the patient cohort are accounted for to further mitigate the likelihood of erroneous characteristics from being learned as predictive indicators. Example CT slices from the patient cohort are shown in Figure 1. It can be observed that the visual appearance of pulmonary fibrosis in different patient CT scans can be quite varied, and thus can be quite challenging to utilize for lung function decline prediction. The variable visual appearance in CT scans further motivate the exploration of deep learning strategies for tackling such a complex prediction task to facilitate for computer-assisted clinical decision support.

Machine-driven design exploration

The goal of the proposed Fibrosis-Net is to predict the forced vital capacity (FVC) of a patient (in ml) at a specific time-point in the future given a patient’s CT scan, initial spirometry measurement, and clinical metadata. In order to construct a highly customized deep convolutional neural network architecture design tailored specifically for high predictive performance, we take inspiration from [15] and a generative synthesis [33] approach was leveraged as the machine-driven design exploration strategy for determining a strong architectural design for CT lung analysis. In this approach, the problem of determining a tailored deep neural network architecture design is formulated as a constrained optimization problem based on a universal performance function \mathcal{U} (e.g., [34]) and a set of quantitative constraints. The aforementioned constrained optimization problem is solved in an iterative fashion, based on an initial network design prototype, the set of quantitative constraints, and data at hand, to produce network N .

More specifically, the backbone architecture design for CT lung analysis identified via machine-driven design exploration leveraged residual architecture design principles [35, 36] as an initial network design prototype, 2,116 patient cases acquired from around the world both with presence and absence of respiratory diseases for improve the quantity and diversity of CT scans, along with associated predictive performance constraints [17, 18]. It is upon this backbone architecture design that the proposed Fibrosis-Net network architecture design was built to be tailored specifically for predicting FVC based on the CT scan, initial spirometry measurement, and clinical metadata of a patient. The resulting Fibrosis-Net was implemented, trained, and evaluated using the TensorFlow deep learning library [37], with training conducted using the Adam optimization algorithm [38] with a mean absolute error (MAE) loss, learning rate of 1e-4, exponential decay of 0.99 every 100 steps, and a batch size of 8.

Network architecture

The proposed Fibrosis-Net architecture is shown in Figure 2. Given a stack of CT images from a patient’s CT scan, each CT image from the lower 55% subset of the CT scan (where pulmonary fibrosis typically presents itself in the lungs) is passed through a series of convolutional layers to create a condensed feature representation characterizing the CT image. This condensed feature representation of the CT image, along with clinical metadata, are then passed together into a dense layer to predict the linear rate of change in lung function. The predicted linear rates of change in lung function from all of the CT images in a patient’s CT scan are then passed into the FVC prediction layer, where the median predicted linear rate of change in lung function, the initial spirometry measurement of the patient, and clinical metadata (i.e., age, sex, and smoking status) are leveraged to predict the FVC at the desired time-point. As the final operations in the FVC prediction layer, a regressor fitted on the clinical metadata with Elastic Net regularization is also leveraged to produce a predicted FVC at the desired time-point, which is then combined with the convolutional-driven FVC prediction to obtain the final FVC prediction.

The proposed Fibrosis-Net possesses a highly efficient, heterogeneous design comprising largely of lightweight architectural components such as depthwise convolutions and pointwise convolutions. In particular, similar to [17, 18], both unstrided and strided projection-replication-projection-expansion design patterns (denoted as PRPE and PRPE-S for unstrided and strided patterns, respectively) are heavily utilized to strike a strong balance between representational capacity and architectural and computational efficiency. The efficiency of the proposed Fibrosis-Net makes it very well-suited for clinical scenarios where computational resources are limited, particularly when dealing with CT scans consisting of many CT images such as the pulmonary fibrosis progression prediction task in this study.

Furthermore, selective long-range connectivity is exhibited in the proposed Fibrosis-Net, with central connectivity hubs comprised of pointwise convolutions for enabling flow in information directly from earlier convolutional layers to the later convolutional layers. By keeping the quantity of long-range connections very sparse through centralized connectivity hubs, the selective long-range connectivity characteristics of Fibrosis-Net strikes a strong balance between improved representational capacity and architectural efficiency.

It can also be observed that Fibrosis-Net fuses learned visual features from a patient’s CT scan with clinical metadata at a later stage of the network architecture design. This enables the efficient utilization of important clinical knowledge captured within clinical metadata alongside important visual characteristics captured in a patient’s CT scan for more well-informed, comprehensive predictive capabilities. Finally, the utilization of a flexible FVC prediction layer in the proposed Fibrosis-Net architecture design takes into account a variable set of linear rates of change in lung function predicted by the dense layer depending on the quantity of CT images within a particular patient’s CT scan. As a result, this effectively allows for much greater flexibility in dealing with different real-world scenarios where the imaging protocol and imaging equipment parameters and configurations can vary greatly.

These architectural characteristics possessed by the proposed Fibrosis-Net illustrates the effectiveness of leveraging machine-driven design exploration for constructing customized deep neural network architecture designs that are specifically tailored for clinical decision support tasks.

Explainability-driven performance validation of Fibrosis-Net

Understanding the behaviour of a deep neural network when making predictions is very important when used in clinical decision support scenarios, given that such predictions will impact patient care and influence treatment and management planning. Inspired by this, we leverage an explainability-driven performance validation strategy to study the decision-making behaviour of Fibrosis-Net as to verify that predictions are based on relevant visual indicators in CT images. Similar to [15, 17, 18, 21], we leverage GSInquire [39] as the explainability method of choice for explainability-driven performance validation in this study to identify critical visual factors in CT images that Fibrosis-Net leverages to make FVC predictions. More specifically, GSInquire harnesses the same generative synthesis strategy [33] leveraged in the machine-driven design exploration process, as a previous study demonstrated the ability of GSInquire to provide explanations that better reflect the decision-making process of deep neural networks quantitatively when compared to other state-of-the-art explainability methods [39]. More importantly, GSInquire is, to the best of the authors' knowledge, is one of the only explainability methods in literature that can be leveraged for studying and validating clinical regression problems, and thus makes it highly desirable for this study.

In brief, GSInquire utilizes the generative synthesis [33] approach from the machine-driven design exploration process for identifying the backbone architecture design, where an inquisitor \mathcal{I} is leveraged to probe a network N with input x . The reactionary responses y from the probing process are leveraged by \mathcal{I} to produce quantitative interpretation z of the decision-making process of a deep neural network N given an input x in the same sub-space as x . The details pertaining to GSInquire for explaining the decision-making behaviour of deep neural networks on CXR images can be found in [15]. An interesting property of GSInquire that also makes it well-suited for explainability-driven performance validation is that it is capable of producing explanations identifying specific critical factors within an image that quantitatively impacts the decisions made by a deep neural network, thus making it more readily interpretable and more quantitative for validation purposes than the types of relative importance variations visualized by other methods.

The details regarding how GSInquire can be leveraged to produce interpretations of deep neural network decision-making behaviour for clinical prediction tasks can be found in [15]. Here, the interpretation z indicates the critical visual factors leveraged by Fibrosis-Net from CT images when making FVC predictions.

Results and discussion

To explore and evaluate the efficacy of the proposed Fibrosis-Net for the prediction of lung decline progression due to pulmonary fibrosis from chest CT images, we take a multi-pronged approach where we conduct: 1) an empirical quantitative performance evaluation of the deep neural network design to study its performance

compared to last state-of-the-art methods, as well as 2) a visual validation evaluation of the decision-making behaviour of the deep neural network design using an explainability-driven performance validation process. The quantitative and qualitative results are presented and discussed in detail below.

Quantitative performance validation results

We quantitatively evaluate the efficacy of the proposed Fibrosis-Net using the OSIC Pulmonary Fibrosis Progression Challenge test cohort [22]. As consistent with the evaluation procedure described in [22], the modified Laplace Log Likelihood score is used, and a comparative analysis is conducted against the three Kaggle winning methods in the OSIC Pulmonary Fibrosis Progression Challenge. The modified Laplace Log Likelihood score (denoted here as L) can be expressed as,

$$\sigma_{clipped} = \max(\sigma, 70 \text{ ml}), \quad (1)$$

$$\Delta = \min(|FVC_{true} - FVC_{predicted}|, 1000 \text{ ml}), \quad (2)$$

$$L = -\frac{\sqrt{2}\Delta}{\sigma_{clipped}} - \ln(\sqrt{2}\sigma_{clipped}). \quad (3)$$

The modified Laplace Log Likelihood score is negative in value and higher the score is, the better the performance of the method is for predicting pulmonary fibrosis progression.

The modified Laplace Log Likelihood scores for the proposed Fibrosis-Net and the winning methods are shown in Table 2. It can be observed that Fibrosis-Net achieves a significantly higher modified Laplace Log Likelihood score when compared to the winning solutions in the challenge. More specifically, Fibrosis-Net achieved a Laplace Log Likelihood score that exceeded the Kaggle 1st place winning solution by **0.0117**, which is significantly higher than the score gaps between the three winning solutions given the logarithmic scale of the modified Laplace Log Likelihood score. Based on these results, it can be seen that Fibrosis-Net can achieve state-of-the-art performance for lung decline progression and demonstrates the efficacy of machine-driven design exploration for constructing deep neural network designs tailored for clinical decision support tasks.

Explainability-driven performance validation results

As discussed earlier, we harnessed GSInquire [39] to conduct explainability-driven performance validation of Fibrosis-Net in order to study its behaviour when making predictions of lung function decline, as well as validate whether predictions are based on clinically-relevant imaging features rather than based on irrelevant features. Figure 3 illustrates example critical factors in CT images of pulmonary fibrosis patients as identified by GSInquire that are key to the decision-making behaviour of the proposed Fibrosis-Net. It can be observed that Fibrosis-Net is capable of leveraging clinically relevant visual indicators such as the presence and geographic extent of honeycombing in the lungs as presented in the CT images to make FVC predictions. As such, it can be clearly seen that the proposed Fibrosis-Net is driven by

correct, clinically relevant decision-making behaviour when making predictions of pulmonary fibrosis progression similar to those leveraged by clinicians [11]. These visual results also highlight the importance of harnessing explainability-driven performance validation when building and evaluating deep neural networks for clinical decision support tasks.

There are several important benefits to taking such an explainability-driven approach to performance validation, particularly for the proposed Fibrosis-Net where the purpose relates to clinical decision support. First of all, by leveraging explainability-driven performance validation, one can obtain greater transparency and understanding into the decision-making behaviour of a deep neural network to ensure that it is leveraging clinically-relevant imaging features to make decisions (i.e., “making the right decisions for the right reasons“). Second, one can gain much greater insight into potential gaps, biases, and errors in both the data as well as in the decision-making behaviour of a deep neural network (i.e., “making the right decisions for the wrong reasons“ based on irrelevant features such as synthetic padding, circular artifacts, etc.). Third, by providing greater transparency into the decision-making processing during prediction, one can provide a greater sense of trust for clinicians leveraging such deep neural networks for computer-assisted clinical decision support and drive greater clinical adoption of such artificial intelligence-driven technologies.

Based on both quantitative and qualitative results, it was demonstrated that Fibrosis-Net can not only make FVC predictions at a higher level of accuracy than state-of-the-art methods, but also do it in a more trustworthy, validated manner that leverages clinically relevant visual indicators within the CT images of a pulmonary fibrosis patient.

Conclusion and Discussion

In this study, we introduced Fibrosis-Net, a deep convolutional neural network design tailored for the prediction of pulmonary fibrosis progression from CT images. Designed with the help of a machine-driven design exploration strategy, Fibrosis-Net is available open source to the general public as part of the OpenMedAI initiative. Experimental results using a patient cohort from the OSIC Pulmonary Fibrosis Progression Challenge show that the proposed Fibrosis-Net can achieve state-of-the-art forced vital capacity prediction performance when compared to the winning solutions on the challenge. Furthermore, an explainability-driven performance validation of Fibrosis-Net showed that relevant visual indicators in the CT images were leveraged when producing predictions. Given the promise of the proposed Fibrosis-Net, we aim to explore this strategy for creating deep neural networks to perform other clinical decision support tasks for other pulmonary conditions such as chronic obstructive pulmonary disease prediction and pulmonary hypertension detection.

Given the devastating effects of pulmonary fibrosis on a individual’s health and well-being and the lack of a known cure, the research in lung function decline prediction presented in this study can have positive benefit to clinical scientists and researchers who are developing deep learning systems for supporting clinical workflows in a number of impactful ways. First, by illustrating the efficacy of machine-driven design for building highly tailored deep neural network architecture designs

for a prediction task beyond the types of clinical decision support tasks illustrate in past studies [15, 16, 17, 18], the hope is that other researchers and scientists may consider leveraging such an approach to accelerate and improve the design of deep learning solutions for different clinical scenarios. Second, by illustrating the efficacy of explainability-based performance validation on gaining a better understanding of the behaviour of Fibrosis-Net on making FVC predictions, the hope is that other researchers and scientists may consider leveraging explainability methods more frequently to improve transparency and trust. Third and finally, the proposed Fibrosis-Net is released in an open-access, open-source fashion, thus allowing for researchers, scientists, and clinicians to leverage this work for further investigation and build upon this work to accelerate the development of clinically viable systems.

It is important to note that Fibrosis-Net is by no means a production-ready clinical assessment solution, and is intended as a foundation for further research and development. Furthermore, it is important to note that the predictions made by Fibrosis-Net and other similar artificial intelligence clinical decision support tools should not be accepted blindly but rather be utilized to aid clinicians in the clinical decision support process. As such, much greater impact can be achieved with tools such as Fibrosis-Net are utilized in a human-in-the-loop manner.

Acknowledgements

We would like to thank Amazon Web Services, Open Source Imaging Consortium, and DarwinAI Corp. We would also like to thank Hayden Gunraj from the University of Waterloo.

Funding

Not applicable.

Availability of data and materials

This work is publically available at <https://github.com/darwinai/FibrosisNet>. Training and inference scripts are provided, along with a pre-trained model. The datasets generated and analysed during the current study are available at <https://www.kaggle.com/c/osic-pulmonary-fibrosis-progression>.

Ethics approval and consent to participate

The study has received ethics clearance from the University of Waterloo (42235).

Competing interests

A.W., J.L., A.D., P.M., M.F., D.M., J.R.H.L., and M.L. are affiliated with DarwinAI Corp.

Consent for publication

Not applicable.

Authors' contributions

A.W. conceived the experiments, J.L., A.D., P.M., M.F., D.M., J.R.H.L., and M.L. conducted the experiments, all authors analysed the results, and all authors reviewed the manuscript.

Additional Information

Fibrosis-Net is a part of the OpenMedAI open source initiative, available at <https://github.com/darwinai/FibrosisNet>.

Author details

¹ Vision and Image Processing Research Group, University of Waterloo, Waterloo, Canada. ² Waterloo Artificial Intelligence Institute, University of Waterloo, Waterloo, Canada. ³ DarwinAI Corp., Waterloo, Canada.

References

1. Taniguchi, H., *et al.*: Pirfenidone in idiopathic pulmonary fibrosis. *European Respiratory Journal* **35**, 821–829 (2010)
2. Richeldi, L., *et al.*: Efficacy and safety of nintedanib in idiopathic pulmonary fibrosis. *The New England Journal of Medicine* **370**, 2071–2082 (2014)
3. Kistler, K., *et al.*: Lung transplantation in idiopathic pulmonary fibrosis: a systematic review of the literature. *BMC Pulm Med.* **14**, 139 (2014)
4. Raghu, G., Remy-Jardin, M., Myers, J.L., Richeldi, L., Ryerson, C.J., Lederer, D.J., Behr, J., Cottin, V., Danoff, S.K., Morell, F., *et al.*: Diagnosis of idiopathic pulmonary fibrosis. an official ats/ers/jrs/alat clinical practice guideline. *American journal of respiratory and critical care medicine* **198**(5), 44–68 (2018)

5. Richeldi, L., Collard, H.R., Jones, M.G.: Idiopathic pulmonary fibrosis. *The Lancet* **389**(10082), 1941–1952 (2017)
6. Tomassetti, S., Wells, A.U., Costabel, U., Cavazza, A., Colby, T.V., Rossi, G., Sverzellati, N., Carloni, A., Carretta, E., Buccioli, M., *et al.*: Bronchoscopic lung cryobiopsy increases diagnostic confidence in the multidisciplinary diagnosis of idiopathic pulmonary fibrosis. *American journal of respiratory and critical care medicine* **193**(7), 745–752 (2016)
7. Watters, L., *et al.*: A clinical, radiographic, and physiologic scoring system for the longitudinal assessment of patients with idiopathic pulmonary fibrosis. *American Review of Respiratory Disease* **133**(1) (1985)
8. Russell, A., *et al.*: Daily home spirometry: An effective tool for detecting progression in idiopathic pulmonary fibrosis. *American Journal of Respiratory and Critical Care Medicine* **194**(8) (2016)
9. Wuyts, W., *et al.*: Daily home spirometry: A new milestone in the field of pulmonary fibrosis. *American Journal of Respiratory and Critical Care Medicine* **194**(8) (2016)
10. Du Bois, R.M., Weycker, D., Albera, C., Bradford, W.Z., Costabel, U., Kartashov, A., King Jr, T.E., Lancaster, L., Noble, P.W., Sahn, S.A., *et al.*: Forced vital capacity in patients with idiopathic pulmonary fibrosis: test properties and minimal clinically important difference. *American journal of respiratory and critical care medicine* **184**(12), 1382–1389 (2011)
11. Devaraj, A.: Imaging: how to recognise idiopathic pulmonary fibrosis. *European Respiratory Review* **23**, 215–219 (2014)
12. Gruden, J.F.: Ct in idiopathic pulmonary fibrosis: Diagnosis and beyond. *American Journal of Roentgenology* **206**, 495–507 (2016)
13. Souza, C.A., Müller, N.L., Flint, J., Wright, J.L., Chung, A.: Idiopathic pulmonary fibrosis: Spectrum of high-resolution ct findings. *American Journal of Roentgenology* **185**, 1531–1539 (2005)
14. Lynch, D.: Ground glass attenuation on ct in patients with idiopathic pulmonary fibrosis. *Chest Journal* **110**(2), 312–313 (1996)
15. Wang, L., Lin, Z.Q., Wong, A.: Covid-net: A tailored deep convolutional neural network design for detection of covid-19 cases from chest x-ray images. *Scientific Reports* (2020)
16. Wong, A., Lin, Z.Q., Wang, L., Chung, A.G., Shen, B., Abbasi, A., Hoshmand-Kochi, M., Duong, T.Q.: Covid-net 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. *Scientific Reports* (2021)
17. Gunraj, H., Wong, A.: Covidnet-ct: A tailored deep convolutional neural network design for detection of covid-19 cases from chest ct images. *Frontiers in Medicine* (2020)
18. Gunraj, H., Sabri, A., Koff, D., Wong, A.: Covid-net ct-2: Enhanced deep neural networks for detection of covid-19 from chest ct images through bigger, more diverse learning. *arXiv (2101.07433)* (2021)
19. Ebad, A., Xi, P., MacLean, A., Tremblay, S., Kohli, S., Wong, A.: COVID-US – An open-access benchmark dataset of ultrasound imaging data for AI-driven COVID-19 analytics (2021). 2103.10003
20. Lee, J.R.H., Pavlova, M., Famouri, M., Wong, A.: CancerNet-SCa: Tailored Deep Neural Network Designs for Detection of Skin Cancer from Dermoscopy Images (2020). 2011.10702
21. Wong, A., Lee, J.R.H., Rahmat-Khah, H., Sabri, A., Alaref, A.: TB-Net: A Tailored, Self-Attention Deep Convolutional Neural Network Design for Detection of Tuberculosis Cases from Chest X-ray Images (2021). 2104.03165
22. OSIC: Osic pulmonary fibrosis progression (2020)
23. Christe, A., *et al.*: Computer-aided diagnosis of pulmonary fibrosis using deep learning and ct images. *Investigative Radiology* **54**(10), 627–632 (2019)
24. Walsh, S., *et al.*: Deep learning for classifying fibrotic lung disease on high-resolution computed tomography: a case-cohort study. *The Lancet Respiratory Medicine* **6**(11) (2018)
25. Walsh, S., *et al.*: A deep learning algorithm for classifying fibrotic lung disease on high resolution computed tomography. *American Journal of Respiratory and Critical Care Medicine* **201**(A7645) (2020)
26. Levin, D.: Deep learning and the evaluation of pulmonary fibrosis. *Lancet Respir Med.* **6**(11), 803–805 (2018)
27. Anthimopoulos, M., Christodoulidis, S., Ebner, L., Christe, A., Mouggiakakou, S.: Lung pattern classification for interstitial lung diseases using a deep convolutional neural network. *IEEE transactions on medical imaging* **35**(5), 1207–1216 (2016)
28. Christodoulidis, S., Anthimopoulos, M., Ebner, L., Christe, A., Mouggiakakou, S.: Multisource transfer learning with convolutional neural networks for lung pattern analysis. *IEEE journal of biomedical and health informatics* **21**(1), 76–84 (2016)
29. Bermejo-Peláez, D., Ash, S.Y., Washko, G.R., Estépar, R.S.J., Ledesma-Carbayo, M.J.: Classification of interstitial lung abnormality patterns with an ensemble of deep convolutional neural networks. *Scientific reports* **10**(1), 1–15 (2020)
30. Tan, M., Le, Q.V.: EfficientNet: Rethinking Model Scaling for Convolutional Neural Networks (2020). 1905.11946
31. M, E., T, G., K, B., *et al.*: Effects of smoking, gender and occupational exposure on the risk of severe pulmonary fibrosis: a population-based casecontrol study. *BMJ Open* **4** (2014)
32. Kalafatis, D., Gao, J., Pesonen, I., *et al.*: Gender differences at presentation of idiopathic pulmonary fibrosis in sweden. *BMC Pulm Med* **19**(22) (2019)
33. Wong, A., Shafiee, M.J., Chwyl, B., Li, F.: FermiNets: Learning generative machines to generate efficient neural networks via generative synthesis (2018). *arXiv:1809.05989*
34. Wong, A.: Netscore: Towards universal metrics for large-scale performance analysis of deep neural networks for practical usage. *CoRR abs/1806.05512* (2018). 1806.05512
35. He, K., Zhang, X., Ren, S., Sun, J.: Deep residual learning for image recognition. In: 2016 IEEE Conference on Computer Vision and Pattern Recognition (CVPR), pp. 770–778 (2016)
36. He, K., Zhang, X., Ren, S., Sun, J.: Identity mappings in deep residual networks. In: Leibe, B., Matas, J., Sebe, N., Welling, M. (eds.) *Computer Vision - ECCV 2016*, pp. 630–645. Springer, Cham (2016)

Table 1 Summary of demographic variables of data used in this study.

Age	mean ± std	67.14 ± 7.01
	50-59	14.9%
	60-69	47.9%
	70-79	34.4%
	80-89	2.8%
Sex		
	Male	78.55%
	Female	21.45%
Smoking status		
	Currently smokes	5.4%
	Ex-smoker	66.3%
	Never smoked	28.3%
Relative FVC		
	≤50	3.4%
	50-69	37.0%
	70-89	36.8%
	90-109	15.2%
	110+	7.6%

Table 2 Comparison of Laplace Log Likelihood scores for the Kaggle winning methods and the proposed Fibrosis-Net on the test cohort from the OSIC Pulmonary Fibrosis Progression Challenge. Best results highlighted in **bold**.

Method	Laplace Log Likelihood
Kaggle 1st place [22]	-6.8305
Kaggle 2nd place [22]	-6.8311
Kaggle 3rd place [22]	-6.8336
Fibrosis-Net	-6.8188

37. Abadi, M., Agarwal, A., Barham, P., Brevdo, E., Chen, Z., Citro, C., Corrado, G.S., Davis, A., Dean, J., Devin, M., Ghemawat, S., Goodfellow, I., Harp, A., Irving, G., Isard, M., Jia, Y., Jozefowicz, R., Kaiser, L., Kudlur, M., Levenberg, J., Mané, D., Monga, R., Moore, S., Murray, D., Olah, C., Schuster, M., Shlens, J., Steiner, B., Sutskever, I., Talwar, K., Tucker, P., Vanhoucke, V., Vasudevan, V., Viégas, F., Vinyals, O., Warden, P., Wattenberg, M., Wicke, M., Yu, Y., Zheng, X.: TensorFlow: Large-Scale Machine Learning on Heterogeneous Systems. Software available from tensorflow.org (2015)

38. Kingma, D.P., Ba, J.: Adam: A Method for Stochastic Optimization (2017). 1412.6980

39. Lin, Z.Q., Shafiee, M.J., Bochkarev, S., Jules, M.S., Wang, X.Y., Wong, A.: Do Explanations Reflect Decisions? A Machine-centric Strategy to Quantify the Performance of Explainability Algorithms (2019). arXiv:1910.07387

Figure 1 Example CT slices from the patient cohort from the OSIC Pulmonary Fibrosis Progression Challenge [22].

Tables

Figure 2 The proposed Fibrosis-Net architecture design. Given a patient's CT scan, initial spirometry measurement, and clinical metadata, the proposed Fibrosis-Net predicts the forced vital capacity (FVC) of a patient at a specified time-point in the future. Fibrosis-Net exhibits an efficient network architecture design with light-weight components that strikes a strong balance between representational capacity and architecture and computational efficiency.

Figure 3 Example critical factors (highlighted as brighter regions) in example CT images of pulmonary fibrosis patients as identified by GSInquire [39]. It can be observed that Fibrosis-Net is capable of leveraging clinically relevant visual indicators such as the presence and geographic extent of honeycombing in the lungs as presented in the CT images to make FVC predictions.

Figures

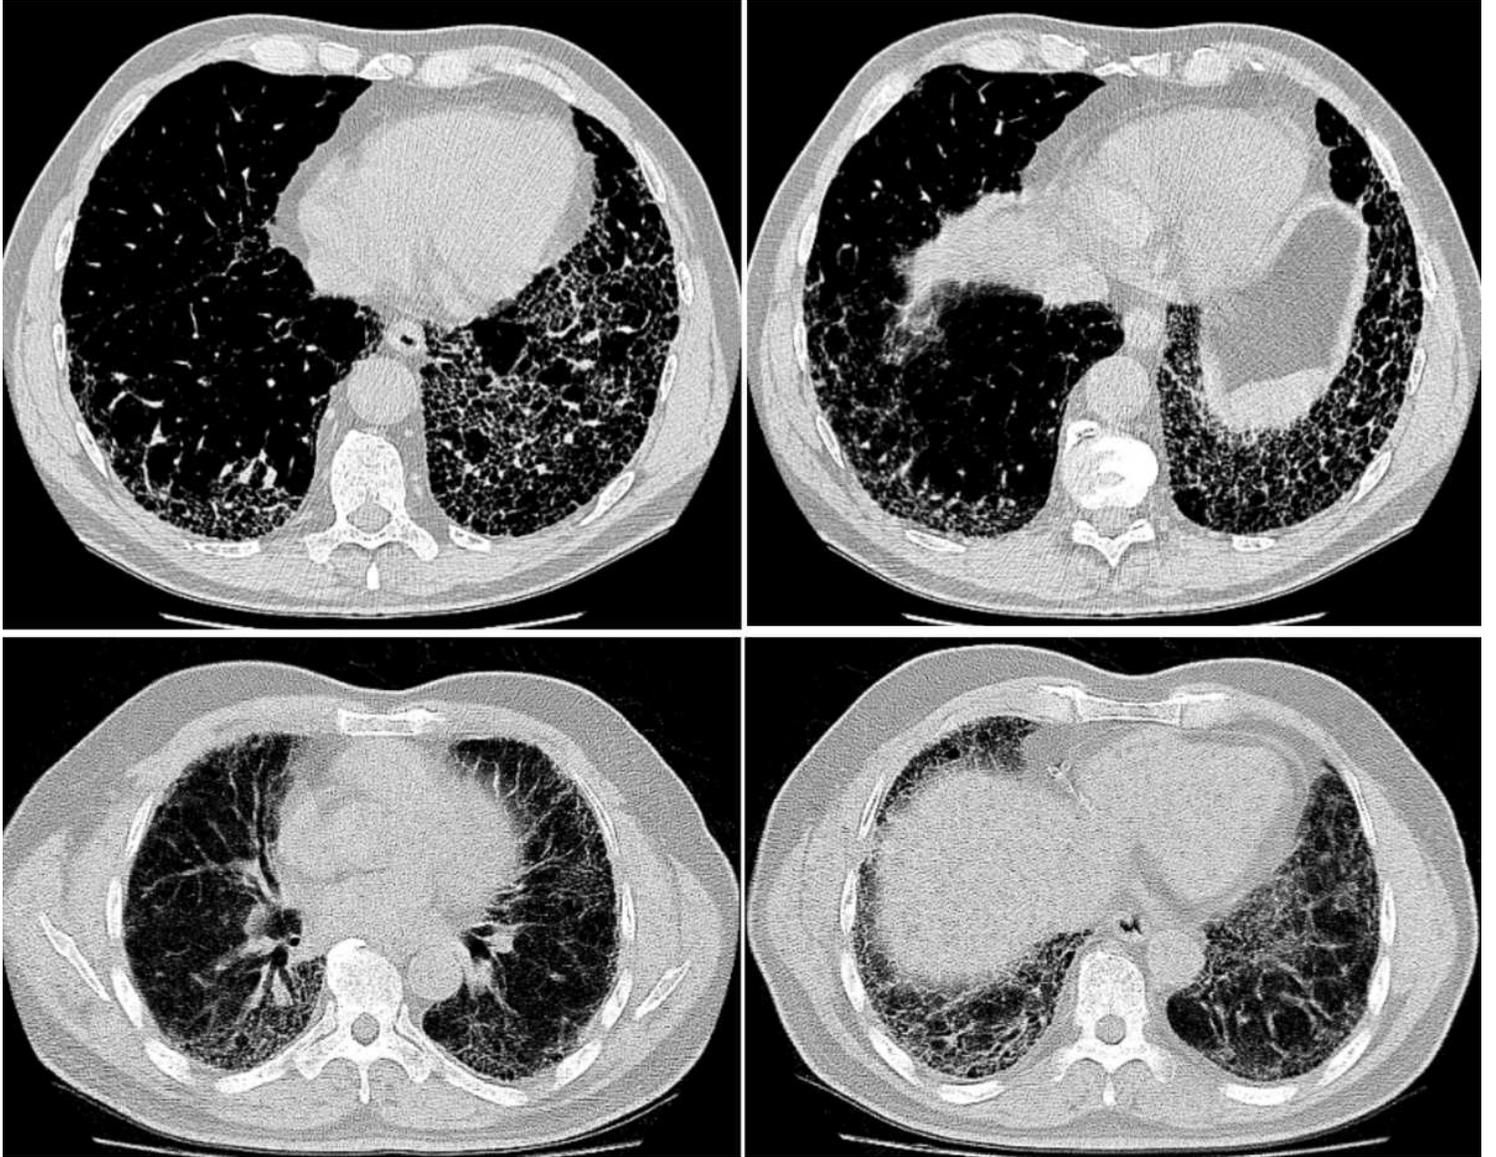


Figure 1

Example CT slices from the patient cohort from the OSIC Pulmonary Fibrosis Progression Challenge [22].

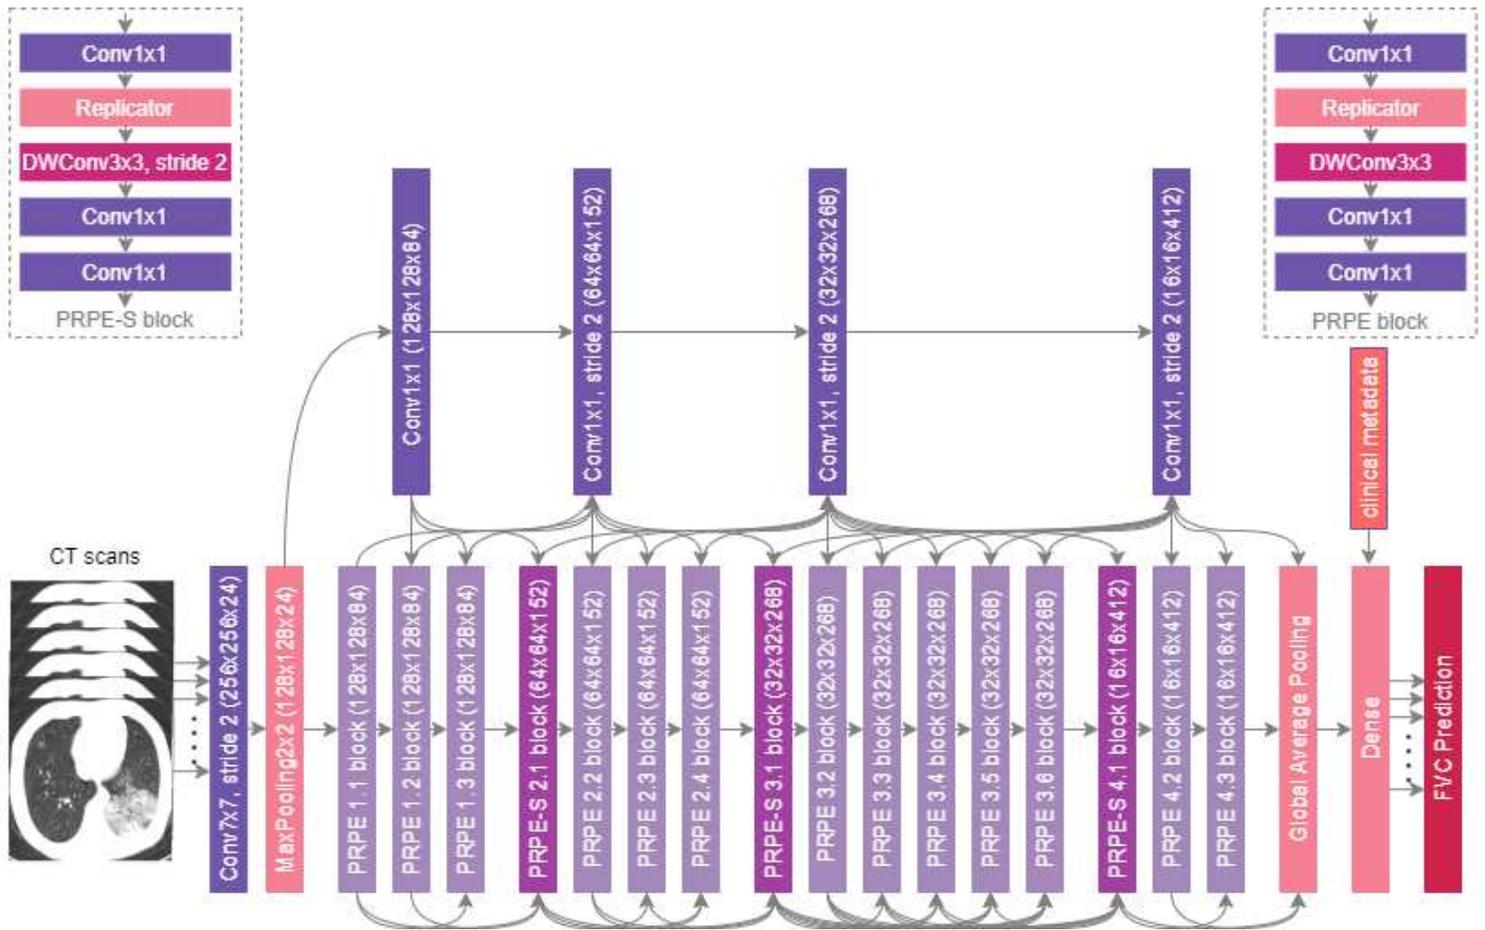


Figure 2

The proposed Fibrosis-Net architecture design. Given a patient's CT scan, initial spirometry measurement, and clinical metadata, the proposed Fibrosis-Net predicts the forced vital capacity (FVC) of a patient at a specified time-point in the future. Fibrosis-Net exhibits an efficient network architecture design with light-weight components that strikes a strong balance between representational capacity and architecture and computational efficiency.

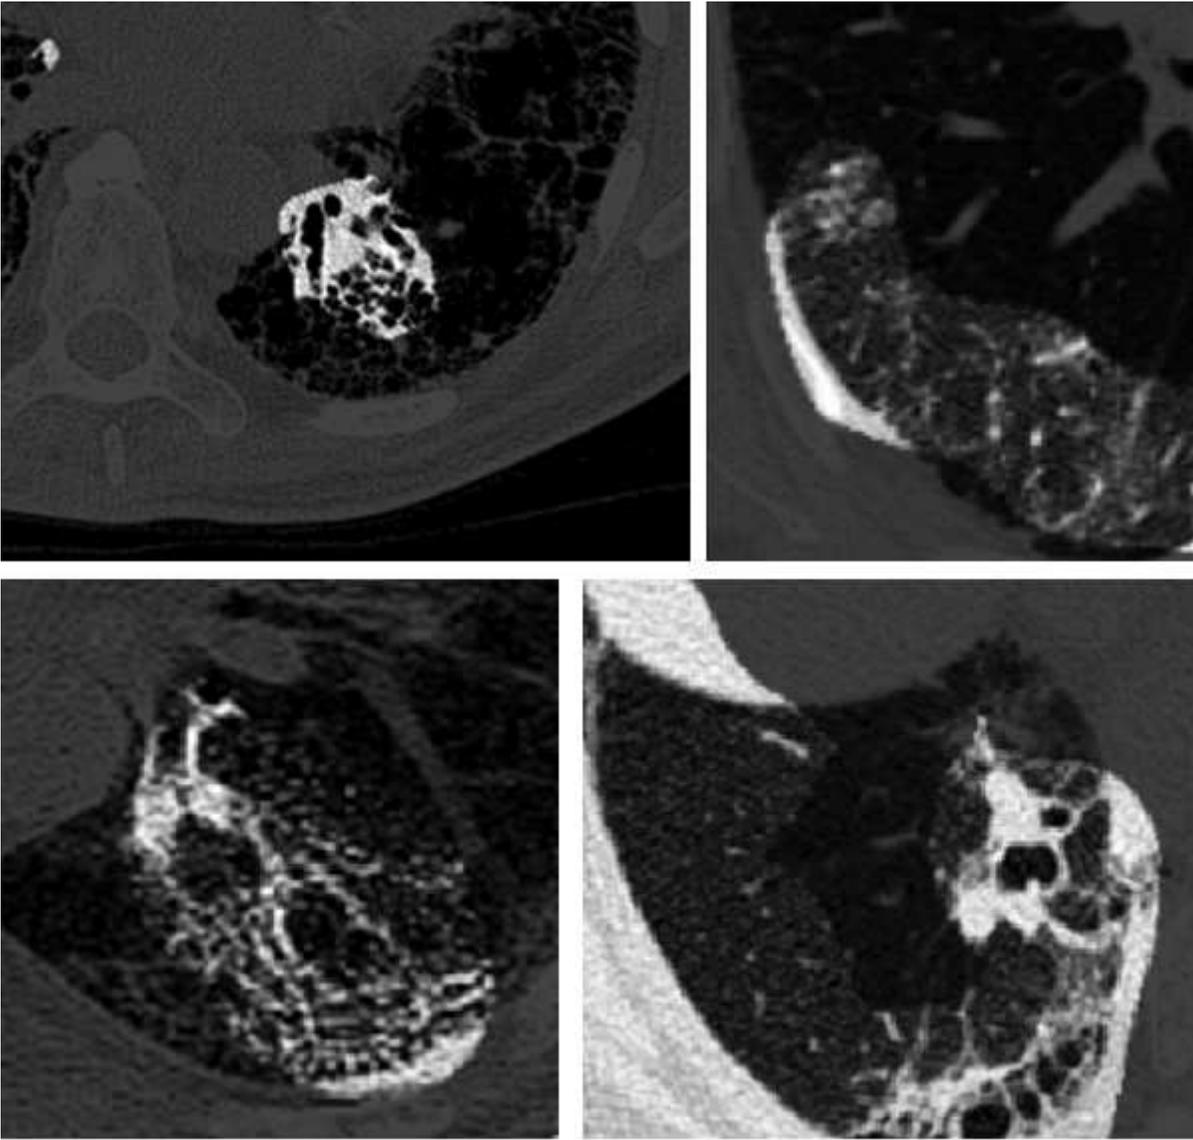


Figure 3

Example critical factors (highlighted as brighter regions) in example CT images of pulmonary fibrosis patients as identified by GSInquire [39]. It can be observed that Fibrosis-Net is capable of leveraging clinically relevant visual indicators such as the presence and geographic extent of honeycombing in the lungs as presented in the CT images to make FVC predictions.