# Automatic Semi-quantitative Histological Assessment of Tissue Traits Using a Smart Web Application

Olympia Giannou, Dimitra E. Zazara, Anastasios D. Giannou, Petra Clara Arck, Georgios Pavlidis

# Automatic semi-quantitative histological assessment of tissue traits using a smart web application

Olympia Giannou[1,*], Dimitra E. Zazara[2,3,*], Anastasios D. Giannou[4,5], Petra Clara Arck[2], Georgios Pavlidis[1]

[1] Computer Engineering & Informatics Dept., Polytechnic School, University of Patras, Patras, Greece.
[2] Division for Experimental Feto-Maternal Medicine, Department of Obstetrics and Fetal Medicine, University Medical Center Hamburg-Eppendorf (UKE), Hamburg, Germany.
[3] University Children's Hospital, UKE, Hamburg, Germany.
[4] I. Department of Medicine, UKE, Hamburg, Germany.
[5] Department of General, Visceral and Thoracic Surgery, UKE, Hamburg, Germany.

[*] equal co-first authors

**Abstract.** A smart web application suitable for classifying goblet cell hyperplasia and level of mucus production in stained lung tissues from mice with experimentally induced allergic asthma. Multiple trainer-model approaches are investigated and proposed in this manuscript, based on machine learning techniques, which provide a technological evolution in the analysis of traits of biomedical imaging. Several schemes, which consist of pre-trained image classifiers on ImageNet, are analyzed and compared each other. Lung tissue images of mice with allergic asthma, depicting mucus-containing periodic acid-Schiff (PAS) positive bronchial cells, are fed as input datasets. The performance of each model is evaluated, based on a variety of metrics: accuracy, recall, precision, cross entropy, f1-score, confusion matrix. Such a web tool could contribute to biomedical research by providing an automated standardized way to determine phenotypic severity of histological traits based on a semi-quantitative scoring scale.

**Keywords:** lung, asthma, tissue, machine learning, classification, CNN, web application.

## 1 Introduction

Asthma is one of the most common chronic respiratory diseases worldwide [1]. The main hallmarks of asthma include chronic airway inflammation, increased mucus production and airway remodeling thereby resulting in narrowing of the airways and recurring episodes of respiratory symptoms such as wheezing, coughing and dyspnea. The causes, pathogenesis and treatment strategies of asthma comprise a key focus of biomedical research with animal models for experimental asthma induction being an essential tool for this purpose. Typical endpoints in experimental approaches focusing on the phenotypic traits of asthma are the severity of airway inflammation, mucus production and related goblet cell hyperplasia, as well as peribronchial fibrosis. The assessment of these parameters usually relies on the semi-quantitative evaluation of re-

spectively stained lung tissue sections by more than one independent and blinded observers, based on an empirically defined scoring scale. Thus, to avoid a potential bias and reduce the time required for such a task, an automated, relatively unbiased and easily reproducible semi-quantitative evaluation of such histological phenotypic features is required. In this study, based on images from stained lung tissues isolated from mice suffering from experimental asthma, Machine Learning (ML) and Convolutional Neural Networks (CNN) were employed to develop a smart application, which could automatically evaluate the phenotypic severity of the induced disease based on the provided histological image.

In this way, features of the tissue images are uncovered and automatically extracted based on previous available knowledge during training on different data (transfer learning approach [2]), patterns are recognized and classification of them is elaborated through ML models, so that human visual interpretation, commonly used until now, is avoided. This automated way of disease severity estimation decreases the possibility of human bias and offers a more accurate objective evaluation, independent on the disagreement of different researchers.

Our objective in this work is to present an automated diagnostic tool to efficiently classify the phenotypic severity of an experimentally-induced condition, such as asthma, on a scale from 0 to 4, based on stained tissue sections. The feature of interest examined in this study was mucus production and the abundance of mucus-containing bronchial epithelial cells in the mouse airways. A similar approach may also be followed in other similar histological staining and related semi-quantitative assessments, focusing on e.g. the severity of tissue inflammation, the extent of tissue infiltration etc. The proposed CNN model is selected following a comparison of its performance with this of several different pre-trained model architectures (Table 3) on our dataset of mouse lung tissue sections, using a variety of metrics (accuracy, recall, precision, cross entropy, f1-score, confusion matrix). The advantages and disadvantages of each architecture are also discussed and compared in Section 4. The structure of this manuscript is as follows: section 2 describes the related work; the materials and methods are presented in section 3; results and comparisons are given in section 4; conclusions are discussed in section 5.

## 2 Related work

Image classification can be efficiently achieved using the approach of transfer learning (TL) and pre-trained image classifier models [3], since the learned features are stored in weighted network structures, freezing a tensor graph [4]. Mainly, the final layer is trained on the new dataset, using less high computing resources and training time. Several studies have used TL on CNNs to achieve image trait classification following different model architectures, pre-trained on the ILSVRC, CIFAR or MNIST image dataset [5].

Tissue image classification using CNNs is generally employed in many biomedical applications. Tumor tissue classification was done with ELM, SVM and DCNN, reaching 97.91% accuracy obtained by DCNN-5 and 97.39% by DCNN-7 [6]. M4-CNN-MaxFeat-based RF achieved 97.1% accuracy in classifying whole lung tissue cancer [7]. Patterns such as emphysema, ground glass, fibrosis, micronodules and healthy lung

tissue were classified with an accuracy of 92.2%, 100%, 86.7%, 93.8% and 92.9% respectively [8]. Same classes were detected using PASA, kNN and SVM architectures [9], showing a poorer performance. Multiple classifiers were applied to lung tissue categorization in [10] while SVM classifier was able to correctly classify 88.3% of the instances and 96.4% when detecting healthy tissue versus all other pathological classes. Covid-19 automatic lung tissue images were classified and segmented in [11]. Covid-19 pneumonia was detected using lung ultrasound and severity was classified in scores 0 to 3 using CNN [12]. Histology image classification using ResNet in celiac disease, lung adenocarcinoma and renal cell carcinoma datasets had an accuracy of 87.06%, 94.51% and 90.16%, respectively [13]. Classification of lung adenocarcinoma, squamous cell carcinoma and benign tissue was achieved using CAD with an accuracy of approximately 99% [14]. Colorectal tissue image classification was achieved with a 93-98% throughput using VGG16 and CapsNet architectures [15].

Previous studies resulted in the classification of various diseases, such as pneumonia and cancer, of organs such as lung or liver, focusing on severity scoring in several cases. To our knowledge, our study is the first to present an automated general-purpose biomedical research tool, able to estimate the severity of histological phenotypic traits, based on the analysis of whole histopathology tissue images (WHTI). Additionally, the objective of this study is to present and compare the performance of several image classifiers, using all known available metrics for their evaluation.

# 3 Material and Methods

## 3.1 Dataset – Data Collection

The dataset contained images of periodic-acid-Schiff (PAS)-stained lung tissue sections from mice with experimentally-induced allergic asthma. For experimental asthma induction, mice were initially sensitized and subsequently challenged with ovalbumin, as described previously. Tissue collection and PAS-staining were performed as described previously [16]. The abundance of mucus-containing bronchial epithelial cells was evaluated in twenty airways in the lung parenchyma of each mouse, based on the severity scoring scale described in Table 1, following published protocols [17], and a mean score for mucus production for each mouse was calculated.

**Table 1.** Description of severity score scale.

| Severity score | Description |
|---|---|
| 0 | *< 0.5% PAS-positive cells* |
| 1 | *5-25%* |
| 2 | *25-50%* |
| 3 | *50-75%* |
| 4 | > 75% |

### 3.2 Methodology

**Input:** Our dataset in Section 3.1 was fed as input in a variety of CNN architectures for image pattern classification. Images were split into 3 subsets in all experiments: training (80%), validation (10%) and testing (10%). Rescaling, cropping and data augmentation techniques were applied to provide well-defined image patterns in the training model architectures. The performance was evaluated on two images at a time (batch size) and the learning rate was equivalent to 0.01. The dataset consisted of five scoring classes, namely 0-4 according to published protocols [17] and each class contained 40 images. The tissue images are fed to the classifier as whole tissue images as they are displayed in Figures 1 and 3, in which the main tissue area of interest has been isolated.

**Convolutional Base (convolutional layers, hidden layers, pooling, fully connected layers):** Deep learning architectures and pre-trained image classifier models [3] following TL were used to preserve the learned features and the final layer was trained on our dataset, building a scheme which could accurately predict the severity score of an input-tissue-image of mucus-containing PAS-positive bronchial cells. During the training procedure, memorization of patterns of each image was done and corresponding bottlenecks were calculated to help the classifier to detect the different classes before resulting in the predicted score through the final CNN layer. The following open-source image classifier models were used [3]: InceptionV1-4, Inception-ResNet-V2, VGGNet-16 and -19 (descendants of AlexNet), MobileNetV1-2, NASNet and PNASNet (mobile and large version), pre-trained on the ILSVRC, CIFAR or MNIST dataset [5]. All models were implemented in Python 3.8 and Tensorflow 2.1. The tissue images of our dataset were rescaled to fit each pre-trained model image size architecture (Table 3). In the pre-trained model architectures, used and compared in this study, the final regression layer was trained on our dataset, following the transfer learning method. Our experiments were elaborated in a standard-performance computer (32GB RAM, Intel i7@3.6GHz).

**Output:** The described scheme is able to classify input images and determine the severity score to which each image belongs. The score values fluctuate from 0 to 4 according to the Table 2. More specifically, the abundance of mucus-containing bronchial epithelial cells on each tissue image is detected, a score is assigned to it and reported to the scientist via the web application 'Automatic Tissue Severity Detection' (ATSD), designed for the purpose of this study. Our web application ATSD allows researchers to upload the tissue image of mice with allergic asthma, depicting mucus-containing periodic acid-Schiff (PAS) positive bronchial cells and to request the corresponding severity score in scale 0-4 (Section 3.1). Firstly, this image is sent as an HTTP request to the REST API server (Django framework [18]). The server feeds this request to the back-end CNN VGG16. VGG16 was identified as the best performing one on our dataset (Section 4) based on all required representative metrics (Section 3.3). This model predicts the score and sends this diagnosis result through the REST API server to the front-end graphical user interface of the web application ATSD (Fig.3) for the user (researchers). The REST API server and the CNN model are implemented in Python 3.8. The graphical interfaces of the web application in Fig. 3 are implemented in html.

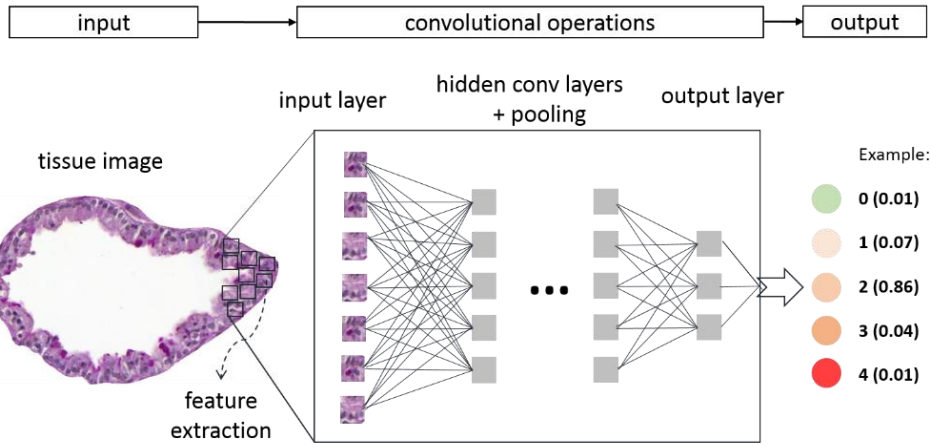


**Fig. 1.** The proposed overall concept to determine the severity score of our input lung tissue image dataset.

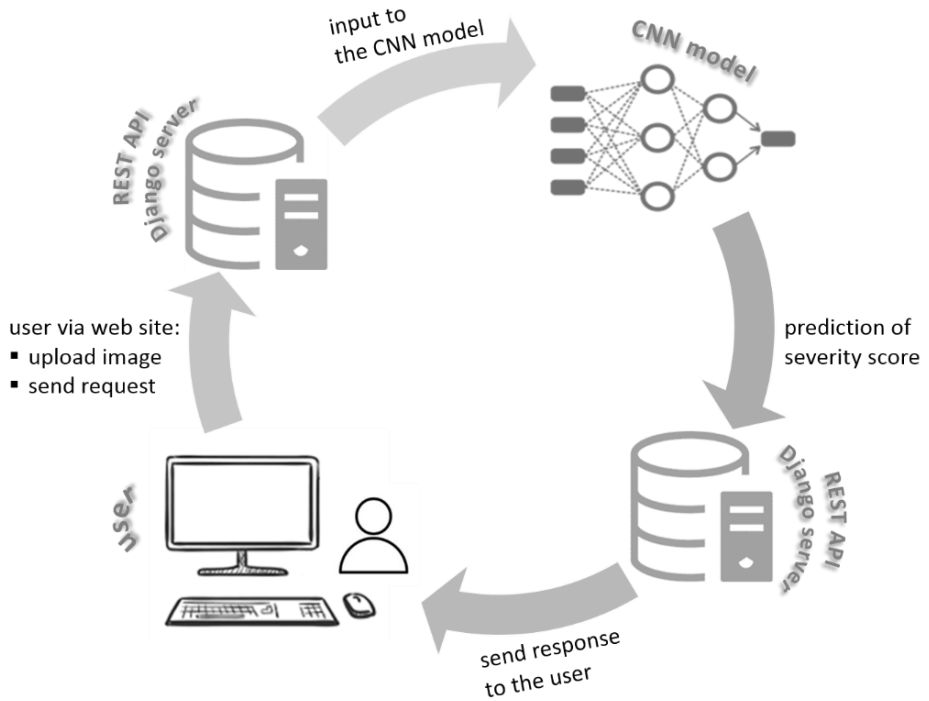


**Fig. 2.** Workflow of our web application ATSD to determine the severity score of histological phenotypic traits demonstrated on tissue images through a Django REST API.

Through our user-friendly web application ATSD, the user can upload a whole tissue image and then request prediction of its severity score. Finally, the input tissue image accompanied with the predicted score are displayed.

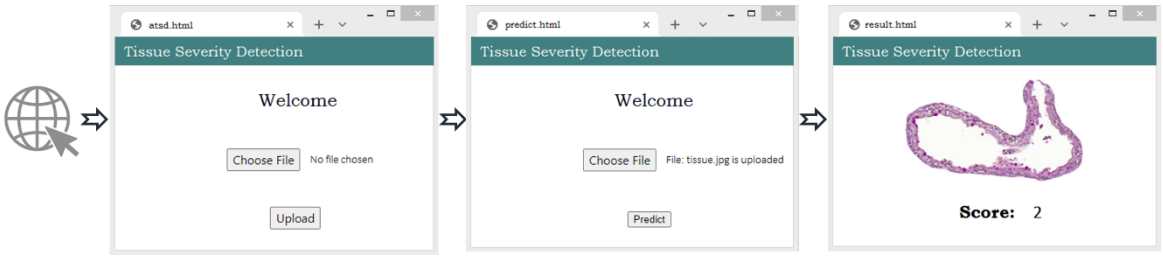


**Fig. 3.** Graphical User Interface (GUI) of our web application ATSD.

### 3.3 Metrics

The performance of all different architectures, reported in Section 3.2, was evaluated by calculating a variety of metrics: accuracy, f1-score, recall, cross-entropy (loss), precision and confusion matrixes [19], as shown in Table 2.

**Table 2.** Interpretation of metrics.

| Metrics | best value | worst value | Description |
|---|---|---|---|
| confusion matrix | *1 in diagonal axis* | *0 in diagonal axis* | *diagonal axis: if the predicted positive and predicted negative labels are true* |
| accuracy | *1* | *0* | *fraction of images correctly judged as true* |
| cross entropy | *0* | *> 0* | *difference/dissimilarities between two probability distributions* |
| precision | *1* | *0* | *ability to return only relevant objects* |
| recall | *1* | *0* | *ability to identify all relevant objects* |
| f1-score | 1 | 0 | harmonic mean of precision and recall |

## 4 Results and Comparisons

The performance of accurately determining the severity score of histological phenotypic traits demonstrated on lung tissue images has been evaluated for all reported architectures in Section 3.2.

The Inception network is complex, consisting of many techniques to increase performance and its evolutions and improvements lead to many different versions: V1 [20], V2 and V3 [21], V4 and ResNet [22]. V1 is used when salient parts (objects of interest) of an image vary in size. V2 and V3 show more accuracy and less computational complexity with reduced bottlenecks. V4 and Inception-ResNet simplify the previous versions and have different stems. The computational costs between V1 and V3, as well as between V2 and V4 are similar. The differences in the architecture between

ResNet-V1 and –V2 [23], [24] are: the addition of a second non-linearity after the addition between x and F(x) on V1 and the removal of the last non-linearity on V2. Additionally, ResNetV2 applies ReLU and batch normalization in a different point compared to ResNetV1. The number of layers used at each case are shown in Table 3. VGGNet, and specifically VGG16 and VGG19, [25] have an increased depth compared to the other classifier CNNs in order to achieve a better performance. However, the increase in layers fails to significantly increase the performance of VGGNet, as in the case of ResNet, since the weights are updated through back-propagation, which finally only slightly changes each weight. ResNet achieves a higher performance while the number of layers increases, avoiding the gradient exploding. MobilenetV1 (2 layers in total) [26] is less complex and has a smaller size than any other classifier. For this reason, it is more frequently introduced for mobile devices. MobileNetV2 (2 layers in total) [27] uses inverted residual structure and the non-linearities are removed in narrow layers. Therefore, in general, the performance of MobileNetV2 is slightly better than this of MobilenetV1. PNASNet-5 [28] has a slightly better performance on ILSVRC than NASNet-A [29] and both surpass MobileNetV1 and InceptionV1. In NASNet-A, the blocks are searched by a learning method and are not predefined, while, in PNASNet-5, the search is elaborated in a progressive order, starting firstly from simple cell structures. The mobile versions of PNASNet and NASNet use 224x224 pixel- instead of 331x331 pixel-images in their large versions.

**Table 3.** Performance of the applied image classifier architectures in determining the severity of airway mucus production in our lung tissue image dataset. (300 training steps)

| **Architecture [3]** | **Input Image Size (pixels)** | **Accuracy %** | **Cross Entropy** | **Precision** | **Recall** | **F1-score** |
|---|---|---|---|---|---|---|
| Inception V1 | *224x224* | *76.9* | *1.14* | *0.65* | *0.51* | *0.58* |
| Inception V2 | *299x299* | *76.2* | *1.15* | *0.62* | *0.50* | *0.56* |
| Inception V3 | *299x299* | *69.2* | *1.14* | *0.64* | *0.51* | *0.57* |
| Inception V4 | *299x299* | *77.0* | *1.13* | *0.65* | *0.50* | *0.57* |
| Inception-ResNet-v2 | *299x299* | *76.7* | *1.15* | *0.67* | *0.52* | *0.59* |
| ResNetV1 50 | *224x224* | *84.6* | *1.14* | *0.67* | *0.52* | *0.59* |
| ResNetV1 101 | *224x224* | *77.0* | *1.13* | *0.63* | *0.51* | *0.57* |
| ResNetV1 152 | *224x224* | *69.0* | *1.11* | *0.61* | *0.50* | *0.55* |
| ResNetV2 50 | *224x224* | *76.1* | *1.12* | *0.68* | *0.49* | *0.59* |
| ResNetV2 101 | *224x224* | 76.9 | 1.14 | 0.64 | 0.51 | 0.57 |
| ResNetV2 152 | *224x224* | 76.8 | 1.14 | 0.68 | 0.55 | 0.61 |
| **VGG 16** | **224x224** | **89.1** | **0.82** | **0.72** | **0.62** | **0.67** |
| VGG 19 | 224x224 | 69.3 | 1.14 | 0.61 | 0.50 | 0.55 |
| MobileNet_v1_1.0_224 | 224x224 | 59.9 | 3.7 | 0.39 | 0.41 | 0.40 |
| MobileNet_v2_1.0_224 | 224x224 | 60.5 | 2.4 | 0.44 | 0.47 | 0.45 |
| NASNet-A_Mobile_224 | 224x224 | 61.5 | 1.22 | 0.55 | 0.43 | 0.49 |
| NASNet-A_Large_331 | 331x331 | 67.8 | 1.17 | 0.57 | 0.45 | 0.51 |
| PNASNet-5_Mobile_224 | 224x224 | 74.2 | 1.31 | 0.61 | 0.45 | 0.53 |
| PNASNet-5_Large_331 | 331x331 | 76.6 | 1.37 | 0.63 | 0.47 | 0.55 |

Experiments demonstrated that VGG16 fits better to our dataset in terms of all metrics, as depicted in Table 3. The metric values among all model architectures are not profoundly different. However, VGG16 outperfomed.

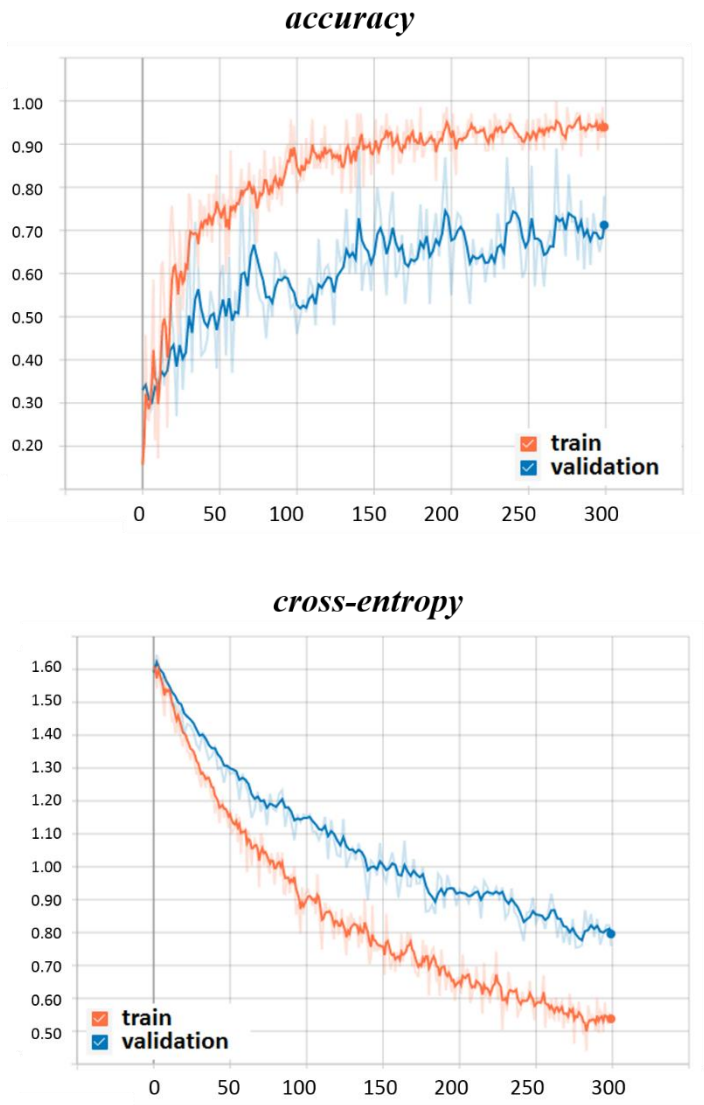


**Fig. 4.** Based on all metrics, VGG16 was identified as the best architecture in accurately determining the severity score of the examined tissue traits. Accuracy (top), cross-entropy (bottom) over the period in steps (300). Training period in orange, validation period in blue. The values are smoothened by the margin of 0.7, real values are shadowed in fainter color. X-axis: steps in steps; Y-axis: accuracy (up), cross-entropy (down).

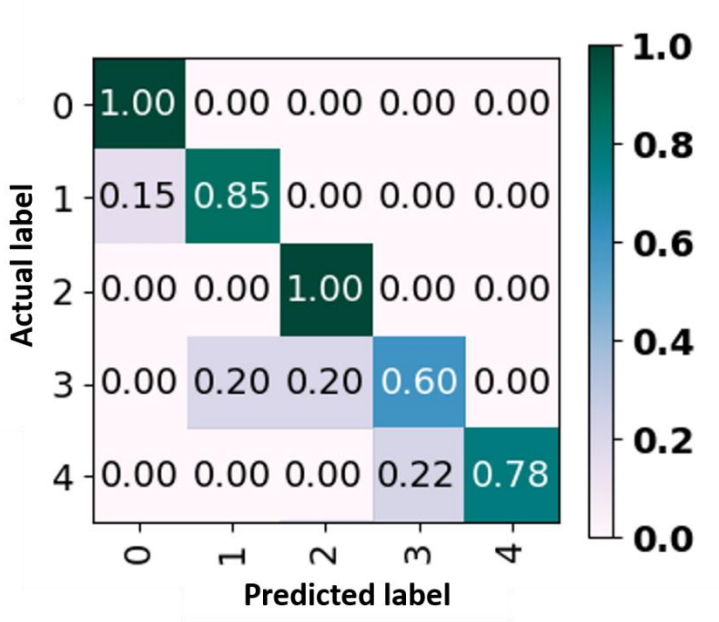


**Fig. 5.** VGG16 was identified as the best architecture in accurately determining the severity score of the examined tissue traits. Confusion matrix. The scale of probability in which the predicted severity score is closer to the actual one is fluctuated from 1 (dark green) to 0 (white-grey).

It is observed that it is more difficult for our classifier to correctly predict the severity score 3 (probability=0.6).

## 5 Conclusions

Using a machine learning approach, we were here able to develop an accurate scheme based on the architecture VGG16 that can be used as an efficient automatic research tool for the semi-quantitative assessment of the extent of histological phenotypic traits based on whole tissue images. A variety of published open-source pre-trained image classifiers were tested on our tissue images and the model VGG16 was identified as the best performing one based on all required representative metrics. This scheme achieves 89.1% accuracy with 0.82 loss and a good f1-score of 0.67, close to 1, which is the harmonic mean between precision and recall. Thus, the here presented web application can efficiently determine the abundance of mucus-containing bronchial cells in asthmatic mouse lungs and serves as an example for addressing similar histological assessments based on the same principles. Overall, such applications may be used as a general-purpose tool in biomedical research, allowing an automated, standardized, fast and highly accurate image analysis.

## 6 Future work

Future work aims to investigate the universality of the developed scheme in a variety of images and histological assays. Specifically, we seek to examine whether the severity score assigned by our web ATSD application each time agrees with the actual score according to standard histological features and the experience of the observer. Optimization of the process will also result in a faster and more efficient outcome.

## 7 Acknowledgement

The authors thank Thomas Andreas for excellent technical assistance. We are also grateful to the Microscopy Imaging Facility of the University Medical Center Hamburg-Eppendorf. This work was supported in part by the German Research Foundation to DEZ and PCA (AR232/25-2 and SO1413/1-2) and State Research Funding - FV45, Authority for Science, Research and Equality, Hanseatic City of Hamburg, Germany to PCA. DEZ is also supported by a clinician scientist scholarship awarded by the Medical Faculty of the University of Hamburg. ADG is supported by an Else Kröner Memorial fellowship, the Werner Otto and the "Hamburger Krebsgesellschaft" foundation.

## References

1. Reddel HK, Bateman ED, Becker A, Boulet L-P, Cruz AA, Drazen JM, et al. A summary of the new GINA strategy: a roadmap to asthma control. Eur Respir J 2015;46:622-39.
2. Hussain M, et al. "A study on cnn transfer learning for image classification," in UK Workshop, Computational Intelligence. *Springer*, (2018) pp. 191–202.
3. Classifiers: github.com/tensorflow/models/tree/master/re-search/slim#pre-trained-models
4. Convolutional Neural Networks: https://www.tensorflow.org/tutorials/
5. ILSVRC:https://image-net.org/challenges/LSVRC/2012/ , CIFAR: https://www.cs.toronto.edu/~kriz/cifar.html, MNIST: http://yann.lecun.com/exdb/mnist/
6. Yang, Xulei, Si Yong Yeo, Jia Mei Hong, Sum Thai Wong, Wai Teng Tang, Zhen Zhou Wu, Gary Lee et al. "A deep learning approach for tumor tissue image classification." IASTED Biomedical Engineering (2016).
7. Wang, Xi, Hao Chen, Caixia Gan, Huangjing Lin, Qi Dou et al. "Weakly supervised learning for whole slide lung cancer image classification." (2018).
8. Depeursinge, Adrien, Daniel Sage, Asmâa Hidki, Alexandra Platon, Pierre-Alexandre Poletti, Michael Unser, and Henning Muller. "Lung tissue classification using wavelet frames." In 2007 29th Annual International Conference of the IEEE Engineering in Medicine and Biology Society, pp. 6259-6262. IEEE, 2007.
9. Song, Yang, Weidong Cai, Yun Zhou, and David Dagan Feng. "Feature-based image patch approximation for lung tissue classification." IEEE transactions on medical imaging 32, no. 4 (2013): 797-808.
10. Depeursinge, Adrien, Jimison Iavindrasana, Asmâa Hidki, Gilles Cohen, Antoine Geissbuhler, Alexandra Platon, Pierre-Alexandre Poletti, and Henning Müller. "Comparative performance analysis of state-of-the-art classification algorithms applied to lung tissue categorization." Journal of digital imaging 23, no. 1 (2010): 18-30.
11. Zaffino, Paolo, Aldo Marzullo, et al. "An open-source covid-19 ct dataset with automatic lung tissue classification for radiomics." Bioengineering 8, no. 2 (2021): 26.

12. La Salvia, Marco, Gianmarco Secco, Emanuele Torti, Giordana Florimbi, Luca Guido, Paolo Lago, Francesco Salinaro, Stefano Perlini, and Francesco Leporati. "Deep learning and lung ultrasound for Covid-19 pneumonia detection and severity classification." *Computers in Biology and Medicine* 136 (2021): 104742.
13. DiPalma, Joseph, Arief A. Suriawinata, Laura J. Tafe, Lorenzo Torresani, and Saeed Hassanpour. "Resolution-based distillation for efficient histology image classification." *Artificial Intelligence in Medicine* 119 (2021): 102136.
14. Nishio, Mizuho, Mari Nishio, Naoe Jimbo, and Kazuaki Nakane. "Homology-based image processing for automatic classification of histopathological images of lung tissue." *Cancers* 13, no. 6 (2021): 1192.
15. Nguyen, Huu-Giao, Annika Blank, Heather E. Dawson, Alessandro Lugli, and Inti Zlobec. "Classification of colorectal tissue images from high throughput tissue microarrays by ensemble deep learning methods." *Scientific rep.* 11, no. 1 (2021): 1-11.
16. Zazara DE, Wegmann M, Giannou AD, Hierweger AM, Alawi M, Thiele K, Huber S, Pincus M, Muntau AC, Solano ME, Arck PC. A prenatally disrupted airway epithelium orchestrates the fetal origin of asthma in mice. J Allergy Clin Immunol. 2020 Jun;145(6):1641-1654.
17. Myou S, Leff AR, Myo S, Boetticher E, Tong J, Meliton AY, et al. Blockade of inflammation and airway hyperresponsiveness in immune-sensitized mice by dominant-negative phosphoinositide 3-kinase-TAT. J Exp Med. 2003;198(10):1573-82.
18. Django framework: https://www.djangoproject.com/
19. Metrics: https://en.wikipedia.org/wiki/Evaluation_of_binary_classifiers, https://en.wikipedia.org/wiki/Confusion_matrix , https://en.wikipedia.org/wiki/Precision_and_recall
20. Szegedy, Christian, Wei Liu, Yangqing Jia, Pierre Sermanet, Scott Reed, Dragomir Anguelov, et al. "Going deeper with convolutions." In Proceedings of the IEEE conference on computer vision and pattern recognition, pp. 1-9. 2015.
21. Szegedy, Christian, Vincent Vanhoucke, Sergey Ioffe, Jon Shlens, and Zbigniew Wojna. "Rethinking the inception architecture for computer vision." In Proceedings of the IEEE conference on computer vision and pattern recognition, pp. 2818-2826. 2016.
22. Szegedy, Christian, Sergey Ioffe, Vincent Vanhoucke, and Alexander A. Alemi. "Inception-v4, inception-resnet and the impact of residual connections on learning." In Thirty-first AAAI conference on artificial intelligence. 2017.
23. He, Kaiming, Xiangyu Zhang, Shaoqing Ren, and Jian Sun. "Deep residual learning for image recognition." In *Proceedings of the IEEE conference on computer vision and pattern recognition*, pp. 770-778. 2016.
24. He, Kaiming, Xiangyu Zhang, et al.. "Identity mappings in deep residual networks." In European conference on computer vision, pp. 630-645. Springer, Cham, 2016.
25. Simonyan, Karen, and Andrew Zisserman. "Very deep convolutional networks for large-scale image recognition." *arXiv preprint arXiv:1409.1556* (2014).
26. Howard, Andrew G., Menglong Zhu, Bo Chen, Dmitry Kalenichenko, Weijun Wang, Tobias Weyand, Marco Andreetto, and Hartwig Adam. "Mobilenets: Efficient convolutional neural networks for mobile vision applications." arXiv preprint arXiv:1704.04861 (2017).
27. Sandler, Mark, Andrew Howard, Menglong Zhu, Andrey Zhmoginov, and Liang-Chieh Chen. "Mobilenetv2: Inverted residuals and linear bottlenecks." In Proceedings of the IEEE conference on computer vision and pattern recognition, pp. 4510-4520. 2018.
28. Liu, Chenxi, Barret Zoph, Maxim Neumann, et al. "Progressive neural architecture search." In Proceedings of the European conference on computer vision (ECCV), pp. 19-34. 2018.
29. Zoph, Barret, Vijay Vasudevan, Jonathon Shlens, and Quoc V. Le. "Learning transferable architectures for scalable image recognition." In Proceedings of the IEEE conference on computer vision and pattern recognition, pp. 8697-8710. 2018.