Anatomy-Aware Hierarchical Multiple Instance Learning for Interpretable COPD Diagnosis and Phenotype Analysis

Raha Ahmadi[1✉], Don Sin[2], Stephen Lam[3], Rachel Eddy[2], and Roger Tam[1]

[1] The University of British Columbia, Vancouver, BC, Canada
raha2001@student.ubc.ca, roger.tam@ubc.ca
[2] Centre for Heart Lung Innovation, The University of British Columbia, Vancouver, BC, Canada
{Don.Sin, Rachel.Eddy}@hli.ubc.ca
[3] BC Cancer, Vancouver, BC, Canada
slam2@bccancer.bc.ca

Abstract. Chronic obstructive pulmonary disease (COPD) is a leading cause of mortality worldwide. COPD causes irreversible lung damage, making early detection essential for slowing disease progression. However, current clinical assessments primarily rely on pulmonary function tests, which require respiratory expertise and cannot fully capture underlying structural abnormalities. Chest computed tomography (CT) provides rich anatomical information, yet manual interpretation remains challenging due to COPD's subtle, diffuse, and heterogeneous disease patterns. Multiple instance learning (MIL) has demonstrated promising performance for weakly supervised COPD classification, but existing methods largely focus on prediction accuracy and offer limited anatomical interpretability, which reduces trustworthiness and phenotype-level insight. In this work, we propose AnatoMIL, an anatomy-aware hierarchical MIL framework for joint COPD detection, severity estimation, and phenotype analysis from 2D CT slices. Our method incorporates lung lobe information to guide hierarchical feature aggregation and employs a dual-head strategy for binary classification and ordinal severity modeling. Furthermore, we analyze the relationship between severity stages and lobe-level contribution profiles to characterize spatial disease patterns. The proposed model is trained and validated on a lung cancer screening dataset with 2,357 scans and externally evaluated on a dataset with 1,753 scans from an independent observational COPD study, demonstrating robust performance and strong generalization. This framework enables interpretable and robust COPD detection and phenotype analysis from CT imaging and has the potential to improve the ability to diagnose and understand the disease patterns of COPD. Code is available at https://github.com/rahaahmadi/AnatoMIL.

Keywords: Chronic Obstructive Pulmonary Disease · Hierarchical Multiple Instance Learning · Disease Phenotyping

1 Introduction

Chronic obstructive pulmonary disease (COPD) is a progressive respiratory disorder characterized by persistent airflow limitation and structural lung damage [1]. As a leading cause of mortality worldwide, early detection is critical, particularly before substantial loss of lung function, to slow disease progression and improve patient outcomes [5, 13]. Although pulmonary function testing (PFT) remains the clinical gold standard for diagnosis, it requires specialized spirometry equipment, trained personnel, and patient cooperation [11]. In addition, PFT may underdiagnose early-stage disease before measurable airflow obstruction develops [6], and does not characterize regional structural abnormalities.

Chest computed tomography (CT) provides detailed anatomical information and is widely available in routine clinical practice. CT can reveal subtle parenchymal changes, regional emphysema, and airway abnormalities that are not captured by spirometry [4, 18, 26]. However, COPD presents as diffuse, heterogeneous, and often subtle structural alterations [12], making consistent visual interpretation challenging even for experienced radiologists and limiting the scalability of manual assessment.

Deep learning methods have shown strong potential for automated COPD diagnosis from CT imaging [21, 23]. Both 2D and 3D convolutional neural networks (CNNs) have been explored for detection and staging tasks [2, 7, 8, 17]. While 3D CNNs directly process volumetric CT data and capture contextual information across slices, CT volumes are large and often require substantial computational resources and GPU memory. In contrast, 2D CNNs operate on individual slices, offering a computationally efficient alternative. Moreover, since clinical CT interpretation is typically performed through sequential inspection of 2D slices, slice-based mechanisms provide a natural and potentially more interpretable modeling framework. However, slice-based modeling introduces an inherent limitation. Clinical labels, such as COPD status or GOLD stage, are defined at the patient level, whereas structural abnormalities are heterogeneously distributed throughout the lung. Not every slice exhibits visible disease patterns, and both the spatial extent and severity vary across patients. Modeling slices independently may therefore dilute disease-specific signals.

Multiple instance learning (MIL) addresses this challenge by representing each CT scan as a bag of instances and learning to aggregate slice-level features into a patient-level prediction without requiring explicit slice annotations. Several MIL approaches have been proposed for COPD classification. Attention-based models emphasize informative slices or patches when forming patient-level representations [15, 22, 24], while graph-based frameworks model relationships among instances through self-attention mechanisms to capture structural dependencies [3]. Hierarchical aggregation strategies further refine bag-level representations through intermediate pooling steps [25].

Despite this progress, existing MIL approaches for COPD have two key shortcomings that limit their clinical utility. First, most methods treat instances as an unstructured set and aggregate them using flat pooling mechanisms that ignore the anatomical organization of the lungs. The lungs consist of five distinct

lobes, which are known to be differentially affected by disease. By overlooking this structure, current approaches produce attention weights and predictions that lack anatomical grounding, making it difficult for clinicians to relate model outputs to familiar anatomical landmarks. Second, limited interpretability is a major barrier to clinical adoption. Clinicians want to know not only whether COPD is present, but also which lung regions are driving the model's decisions. Motivated by this need, we propose an anatomy-aware hierarchical MIL framework for joint COPD detection, ordinal severity estimation, and phenotype analysis. Our method organizes CT slices into the five lung lobes using automated segmentation and aggregates features hierarchically, yielding interpretable lobe-level contribution weights that indicate which regions drive the prediction. A dual-head strategy performs binary classification alongside ordinal severity modeling, predicting severity scores and learning decision thresholds aligned with established staging to support early diagnosis. In addition, lobe-level attention profiles are analyzed to characterize spatial disease patterns and region-specific involvement across disease stages.

Our experiments demonstrate that anatomy-aware aggregation achieves robust diagnostic performance while providing regional insights. Our contributions are threefold. We propose AnatoMIL, an anatomy-aware hierarchical MIL framework that **(i) jointly addresses COPD detection and anatomically grounded interpretability** through lobe-level aggregation, **(ii) implements ordinal severity modeling** for stage-aware prediction, and **(iii) utilizes attention profiles** to compute a spatial phenotype for each subject that can be used to contextualize their disease state.

2 Method

We formulate COPD diagnosis from CT scans as a multiple instance learning (MIL) problem. Each volume $X = \{x_k\}_{k=1}^{K}$ is treated as a bag of K 2D slices and $Y = (y, s)$ contains the patient-level labels, consisting of a binary diagnosis label y and an ordinal GOLD severity stage s. To incorporate anatomical structure into the MIL framework, we assign each slice x_k a lobe-fraction vector $f_k \in R^L$, where $f_{k,\ell}$ denotes the fraction of lung tissue in slice k belonging to lobe ℓ, with $\sum_{\ell=1}^{L} f_{k,\ell} \leq 1$. These fractions encode the anatomical composition of each slice and guide hierarchical feature aggregation. As shown in Fig. 1, our framework comprises three components: 1) an anatomy-aware hierarchical MIL module that performs lobe-guided slice aggregation, 2) a dual-head prediction module that jointly models binary COPD diagnosis and ordinal severity, and 3) a phenotype analysis module that analyzes lobe-level contribution profiles to characterize spatial disease patterns.

2.1 Anatomy-Aware Hierarchical MIL

To preserve structural information while maintaining computational efficiency, we model each CT scan as a bag of 2D slices. Rather than aggregating slices

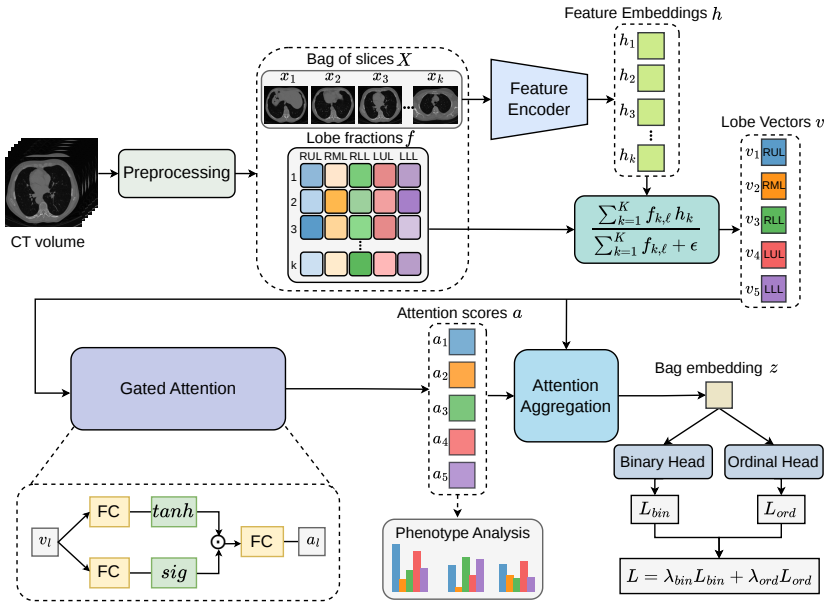

Fig. 1. Overview of the proposed anatomy-aware hierarchical MIL framework. The model includes 1) lobe-guided hierarchical aggregation, 2) dual-head prediction for binary COPD detection and ordinal severity modeling, and 3) phenotype analysis based on lobe contribution profiles. The five lobes are Left and Right Upper (LUL, RUL), Left and Right Lower (LLL, RLL), and Right Middle (RML).

directly using flat pooling, we introduce an intermediate anatomical hierarchy based on lung lobes. This design reflects the clinical observation that COPD involvement is often regionally heterogeneous.

Feature extraction. Each slice x_k is processed by a convolutional backbone to obtain a feature representation $h_k = \phi_\theta(x_k) \in R^D$. These slice-level embeddings encode structural characteristics such as emphysema patterns and airway abnormalities. Since supervision is provided only at the patient level, slice representations are learned indirectly through downstream aggregation and prediction objectives.

Slice-to-lobe aggregation. Conventional MIL methods treat slices as an unordered set and aggregate them directly. In contrast, we explicitly incorporate lung anatomy by grouping slice features according to lobe membership. Using the lobe-fraction vectors $f_{k,\ell}$ as soft assignments, slice features are aggregated into L lobe-level representations:

$$v_\ell = \frac{\sum_{k=1}^{K} f_{k,\ell}\, h_k}{\sum_{k=1}^{K} f_{k,\ell} + \epsilon}, \quad \ell = 1, \ldots, L. \tag{1}$$

This operation produces region-level descriptors that summarize structural characteristics within each anatomical lobe. By introducing this intermediate aggregation step, the model captures regional disease patterns before forming a global patient representation.

Lobe-level attention and bag embedding. After obtaining lobe-level representations $V = [v_1, \ldots, v_L] \in R^{L \times D}$, we model their relative importance using a gated attention mechanism:

$$a_\ell = \frac{\exp\, w^\top \left(\tanh(V_\ell U^\top) \odot \sigma(V_\ell W^\top)\right)}{\sum_{j=1}^{L} \exp\, w^\top \left(\tanh(V_j U^\top) \odot \sigma(V_j W^\top)\right)}, \tag{2}$$

where $w \in R^d$, $U, W \in R^{d \times D}$, and $\odot$ denotes element-wise multiplication. The patient-level embedding is computed as

$$z = \sum_{\ell=1}^{L} a_\ell v_\ell \in R^D. \tag{3}$$

The attention vector $a = (a_1, \ldots, a_L)$ serves as a lobe contribution profile, providing an anatomically grounded descriptor of regional disease involvement. Because attention weights are defined over clinically meaningful anatomical units, they offer structured interpretability and are retained for downstream phenotype analysis.

2.2 Dual-Head Prediction: Detection and Staging

Binary detection head. The patient embedding z is mapped to a scalar logit through a linear layer and optimized with binary cross-entropy loss $\mathcal{L}_{\text{bin}}$, yielding the probability of COPD presence.

Ordinal severity head. To model the ordered structure of GOLD staging, we adopt an all-threshold ordinal formulation. For a patient at stage s, the binary target vector $t \in 0, 1^{S-1}$ is defined element-wise as $t_j = \mathbf{1}[s > j]$. The model outputs $S-1$ logits $g = W_{\text{ord}} z + b_{\text{ord}}$, where W_{ord} is the learnable weight matrix and b_{ord} is the corresponding bias vector. Each logit is supervised independently using binary cross-entropy. A continuous severity score is computed as

$$\xi = \frac{1}{S-1} \sum_{j=1}^{S-1} \sigma(g_j), \tag{4}$$

and the discrete stage prediction is obtained by counting the number of sigmoid outputs exceeding 0.5. The ordinal loss $\mathcal{L}_{\text{ord}}$ is defined as the mean binary cross-entropy across all $S-1$ thresholds.

Training objective. The full objective is:

$$\mathcal{L} = \lambda_{\text{bin}} \mathcal{L}_{\text{bin}} + \lambda_{\text{ord}} \mathcal{L}_{\text{ord}}. \tag{5}$$

The overall loss combines the binary detection loss and the ordinal severity loss as a weighted sum. The ordinal branch serves as an auxiliary task that

Table 1. Quantitative results (%) of different MIL methods for binary COPD classification and severity staging on the PanCan (5-fold cross-validation) and ECLIPSE (external validation) cohorts (mean ± standard deviation).

Method	Binary		Severity Stage			
	AUC ↑	F1 ↑	Acc ↑	AUC (≥2) ↑	AUC (≥3) ↑	QWK ↑
PanCan (5-fold CV)						
ABMIL [10]	$84.67_{\pm 1.91}$	$78.51_{\pm 2.48}$	$60.93_{\pm 1.55}$	$83.63_{\pm 2.36}$	$91.02_{\pm 3.05}$	$58.82_{\pm 3.12}$
HierarchicalMIL [25]	$85.44_{\pm 1.83}$	$78.53_{\pm 1.73}$	$61.14_{\pm 1.23}$	$84.71_{\pm 2.13}$	$91.28_{\pm 2.39}$	$60.29_{\pm 4.77}$
TransMIL [14]	$86.43_{\pm 1.71}$	$78.83_{\pm 3.18}$	$63.60_{\pm 2.65}$	$85.42_{\pm 2.65}$	$90.92_{\pm 3.01}$	$62.43_{\pm 3.68}$
AnatoMIL	$\mathbf{86.88}_{\pm 2.25}$	$\mathbf{79.05}_{\pm 2.32}$	$\mathbf{63.65}_{\pm 3.83}$	$\mathbf{86.15}_{\pm 2.75}$	$\mathbf{92.57}_{\pm 2.93}$	$\mathbf{62.82}_{\pm 4.00}$
ECLIPSE (External, 5 CV models trained on PanCan)						
ABMIL [10]	$95.39_{\pm 0.24}$	$94.15_{\pm 1.03}$	$48.57_{\pm 2.47}$	$97.27_{\pm 0.20}$	$87.50_{\pm 0.52}$	$73.27_{\pm 4.08}$
HierarchicalMIL	$95.96_{\pm 0.34}$	$94.76_{\pm 0.71}$	$49.13_{\pm 1.85}$	$97.34_{\pm 0.17}$	$87.13_{\pm 0.51}$	$73.43_{\pm 1.72}$
TransMIL [14]	$96.18_{\pm 0.32}$	$95.11_{\pm 0.36}$	$51.44_{\pm 2.05}$	$97.37_{\pm 0.25}$	$87.44_{\pm 0.82}$	$74.68_{\pm 2.44}$
AnatoMIL	$\mathbf{96.30}_{\pm 0.22}$	$\mathbf{95.20}_{\pm 0.19}$	$\mathbf{51.66}_{\pm 2.61}$	$\mathbf{97.38}_{\pm 0.14}$	$\mathbf{87.60}_{\pm 0.42}$	$\mathbf{74.72}_{\pm 3.46}$

encourages the embedding space to reflect disease progression, which can improve the robustness of binary diagnosis.

2.3 Phenotype Analysis

After training, we collect the lobe-level attention vectors from all patients in the dataset. Let $a^{(n)}{}_{n=1}^{N}$ denote the lobe contribution profiles of all patients. Each profile $a^{(n)} \in R^L$ encodes the relative contribution of the L anatomical lobes to the patient-level prediction, forming a compact descriptor of spatial disease distribution.

3 Experiments

Datasets. Experiments are conducted on the PanCan lung cancer screening cohort [16] (n=2,357) for model development and internal validation, targeting binary COPD classification and severity staging. For GOLD staging, the distribution is imbalanced, with Stage 0 to Stage 4 proportions of 49.7%, 17.8%, 26.9%, 4.8%, and 0.7%, respectively. Due to the very small proportion of Stage 4 cases, Stages 3 and 4 were merged during dual-head training. External validation was subsequently performed on the independent ECLIPSE cohort [19] (n=1,753), which was specifically collected for COPD evaluation. In ECLIPSE, the GOLD stage distribution is 21.9%, 3.4%, 23.3%, 33.2%, and 18.3% for Stage 0 to Stage 4, respectively.

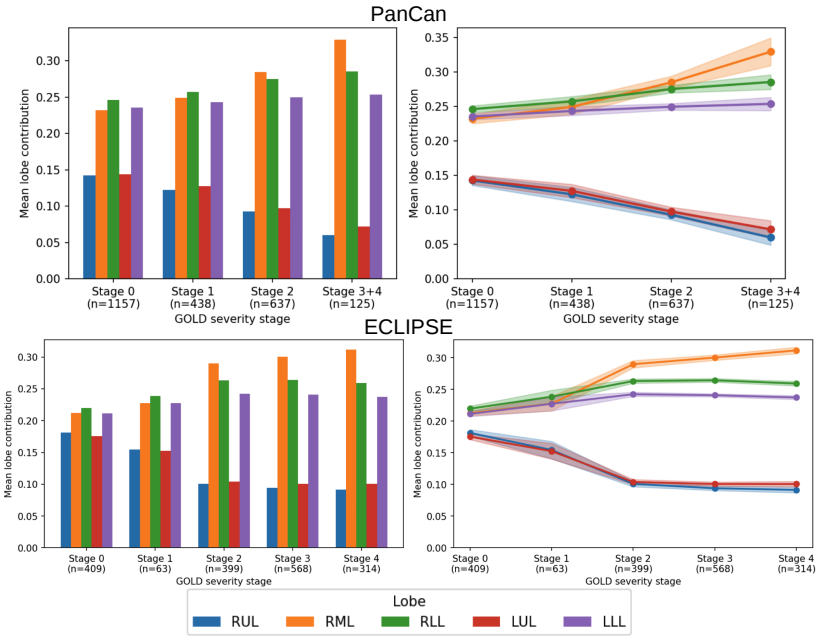

Fig. 2. Lobe-level contribution profiles across GOLD severity stages for the PanCan (top) and ECLIPSE (bottom) cohorts. Contributions shift from upper lobes (RUL, LUL) toward the middle and lower lobes (particularly RML) with increasing severity.

Implementation Details. CT volumes were resampled to 1 mm isotropic spacing. Lungs were segmented using a pretrained U-Net [9] and cropped to the lung region. For each scan, 24 axial slices were uniformly sampled from the lung extent. A 10 mm minimum-intensity projection was computed per slice and converted into a 3-channel image using three HU windows (1000/−500, 1600/−600, 500/−950) to capture emphysematous, dense, and overall lung patterns. Lobe segmentation is performed on the full 3D CT volume using a pretrained TotalSegmentator model [20], and since the MIL framework operates on 2D axial slices, soft lobe-fraction vectors corresponding to the Right Upper Lobe (RUL), Right Middle Lobe (RML), Right Lower Lobe (RLL), Left Upper Lobe (LUL), and Left Lower Lobe (LLL) are computed at the slice level for hierarchical aggregation. We use a ResNet-18 pretrained on ImageNet (224×224 input) optimized with AdamW (learning rate 1e−4, batch size 32) for up to 50 epochs with early stopping and a cosine learning-rate schedule. The binary and ordinal losses are weighted as $\lambda_{\text{bin}} = 1$ and $\lambda_{\text{ord}} = 0.5$. Five-fold cross-validation is performed, stratified by the binary label (or GOLD stage for the dual-head model).

Main Results. We report AUC and F1 for binary classification. For ordinal severity staging, we report accuracy, AUC, and Quadratic Weighted Kappa (QWK), where QWK measures agreement between predicted and true ordinal labels, assigning larger penalties to larger stage discrepancies. As shown in Table 1, our method achieves competitive performance relative to state-of-the-art MIL baselines across both internal and external cohorts. All methods were evaluated under a standardized preprocessing and feature extraction pipeline to ensure a fair comparison. The proposed anatomy-aware MIL performs on par with a transformer-based MIL approach, while slightly surpassing the others for most metrics. More importantly, the main contribution of our framework is incorporating explicit lung anatomy into the learning process. This allows structured lobe-level interpretability and spatial phenotype analysis, while maintaining competitive predictive performance.

Phenotype Analysis. We analyzed lobe-level attention profiles to characterize spatial patterns across GOLD stages (Fig. 2). Early-stage profiles show relatively balanced contributions across lobes, whereas advanced stages exhibit a redistribution toward the middle and lower lobes (RML, RLL, LLL), with the Right Middle Lobe showing the most pronounced increase. Specifically, RUL and LUL progressively decrease with increasing severity, RML shows a clear monotonic increase, and RLL and LLL exhibit a moderate upward trend at the milder stages, but are flatter above GOLD 2. Both datasets show similar trends, despite their large difference in severity stage distributions. Based on current knowledge of COPD, these regional differences likely reflect varying influences of the two main pathological processes of COPD, which are emphysema and small airway disease. Examining the lobe-level contribution profile for a given patient can potentially help the clinician determine whether the patient's pathology is consistent with the phenotypic profile of other similarly staged patients, which would enhance interpretability and trust. However, because attention weights are normalized, the contributions are relative and do not represent absolute regional changes across disease stages. Future work integrating lobe-level contributions with quantitative emphysema or airway metrics may clarify whether shifts in attention correspond to measurable structural progression. Overall, the consistency observed across cohorts supports the regional interpretability of the proposed method.

4 Conclusion

In this work, we presented AnatoMIL, an anatomy-aware hierarchical MIL framework for interpretable COPD detection and severity modeling from CT. By incorporating explicit lung lobe structure into the aggregation process, the proposed method preserves anatomical organization while maintaining competitive predictive performance across internal and external cohorts. Beyond classification accuracy, our framework provides structured lobe-level contribution profiles that enable spatial phenotype analysis and offer clinically meaningful insight into regional disease involvement. The consistent patterns observed across co-

horts suggest that embedding anatomical priors into weakly supervised learning can enhance interpretability without sacrificing robustness. Future work will focus on integrating regional contribution profiles with quantitative emphysema and airway metrics to further investigate the relationship between learned attention patterns and structural disease progression.

Disclosure of Interests. The authors have no competing interests to declare that are relevant to the content of this article.

References

1. Agustí, A., Celli, B.R., Criner, G.J., Halpin, D., Anzueto, A., Barnes, P., Bourbeau, J., Han, M.K., Martinez, F.J., De Oca, M.M., et al.: Global initiative for chronic obstructive lung disease 2023 report: GOLD executive summary. Journal of the Pan African Thoracic Society **4**(2), 58–80 (2022)
2. Ahmed, J., Vesal, S., Durlak, F., Kaergel, R., Ravikumar, N., Rémy-Jardin, M., Maier, A.: COPD classification in CT images using a 3D convolutional neural network. In: Bildverarbeitung für die Medizin 2020: Algorithmen–Systeme–Anwendungen. Proceedings des Workshops vom 15. bis 17. März 2020 in Berlin. pp. 39–45. Springer (2020)
3. Chen, L., Feng, Q., Yin, X., Min, X., Shi, L., Yang, D., Chen, Y.W., Zhang, D., Zhu, W.: A graph convolutional multiple instance learning on a hypersphere manifold approach for diagnosing chronic obstructive pulmonary disease in CT images. IEEE Journal of Biomedical and Health Informatics **26**(12), 6058–6069 (2022)
4. Elbehairy, A.F., Marshall, H., Naish, J.H., Wild, J.M., Parraga, G., Horsley, A., Vestbo, J.: Advances in COPD imaging using CT and MRI: linkage with lung physiology and clinical outcomes. European Respiratory Journal **63**(5) (2024)
5. Gershon, A.S., Thiruchelvam, D., Chapman, K.R., Aaron, S.D., Stanbrook, M.B., Bourbeau, J., Tan, W., To, T., Network, C.R.R., et al.: Health services burden of undiagnosed and overdiagnosed COPD. Chest **153**(6), 1336–1346 (2018)
6. Habteslassie, D., Khorramnia, S., Muruganandan, S., Romeo, N., See, K., Hannan, L.M.: Missed diagnosis or misdiagnosis: How often do hospitalised patients with a diagnosis of chronic obstructive pulmonary disease (COPD) have spirometry that supports the diagnosis? Internal medicine journal **53**(4), 510–516 (2023)
7. Hatt, C., Galban, C., Labaki, W., Kazerooni, E., Lynch, D., Han, M.: Convolutional neural network based COPD and emphysema classifications are predictive of lung cancer diagnosis. In: International Workshop on Reconstruction and Analysis of Moving Body Organs. pp. 302–309. Springer (2018)
8. Ho, T.T., Kim, T., Kim, W.J., Lee, C.H., Chae, K.J., Bak, S.H., Kwon, S.O., Jin, G.Y., Park, E.K., Choi, S.: A 3D-CNN model with CT-based parametric response mapping for classifying COPD subjects. Scientific Reports **11**(1), 34 (2021)
9. Hofmanninger, J., Prayer, F., Pan, J., Röhrich, S., Prosch, H., Langs, G.: Automatic lung segmentation in routine imaging is primarily a data diversity problem, not a methodology problem. European radiology experimental **4**(1), 50 (2020)
10. Ilse, M., Tomczak, J., Welling, M.: Attention-based deep multiple instance learning. In: International conference on machine learning. pp. 2127–2136. PMLR (2018)
11. Jia, J., Marges, E.R., De Vries-Bouwstra, J.K., Ninaber, M.K., Kroft, L.J., Schouffoer, A.A., Staring, M., Stoel, B.C.: Automatic pulmonary function estimation from

chest CT scans using deep regression neural networks: The relation between structure and function in systemic sclerosis. IEEE Access **11**, 135272–135282 (2023)
12. Lynch, D.A., Austin, J.H., Hogg, J.C., Grenier, P.A., Kauczor, H.U., Bankier, A.A., Barr, R.G., Colby, T.V., Galvin, J.R., Gevenois, P.A., et al.: CT-definable subtypes of chronic obstructive pulmonary disease: A statement of the fleischner society. Radiology **277**(1), 192–205 (2015)
13. Riley, C.M., Sciurba, F.C.: Diagnosis and outpatient management of chronic obstructive pulmonary disease: A review. Jama **321**(8), 786–797 (2019)
14. Shao, Z., Bian, H., Chen, Y., Wang, Y., Zhang, J., Ji, X., et al.: TransMIL: Transformer based correlated multiple instance learning for whole slide image classification. Advances in neural information processing systems **34**, 2136–2147 (2021)
15. Sun, J., Liao, X., Yan, Y., Zhang, X., Sun, J., Tan, W., Liu, B., Wu, J., Guo, Q., Gao, S., et al.: Detection and staging of chronic obstructive pulmonary disease using a computed tomography-based weakly supervised deep learning approach. European Radiology **32**(8), 5319–5329 (2022)
16. Tammemagi, M.C., Schmidt, H., Martel, S., McWilliams, A., Goffin, J.R., Johnston, M.R., Nicholas, G., Tremblay, A., Bhatia, R., Liu, G., et al.: Participant selection for lung cancer screening by risk modelling (the pan-canadian early detection of lung cancer [PanCan] study): a single-arm, prospective study. The lancet oncology **18**(11), 1523–1531 (2017)
17. Tang, L.Y., Coxson, H.O., Lam, S., Leipsic, J., Tam, R.C., Sin, D.D.: Towards large-scale case-finding: training and validation of residual networks for detection of chronic obstructive pulmonary disease using low-dose CT. The Lancet Digital Health **2**(5), e259–e267 (2020)
18. Thiboutot, J., Yuan, W., Park, H.C., Lerner, A.D., Mitzner, W., Yarmus, L.B., Li, X., Brown, R.H.: Current advances in COPD imaging. Academic radiology **26**(3), 335–343 (2019)
19. Vestbo, J., Anderson, W., Coxson, H.O., Crim, C., Dawber, F., Edwards, L., Hagan, G., Knobil, K., Lomas, D.A., MacNee, W., et al.: Evaluation of COPD longitudinally to identify predictive surrogate end-points (ECLIPSE). European Respiratory Journal **31**(4), 869–873 (2008)
20. Wasserthal, J., Breit, H.C., Meyer, M.T., Pradella, M., Hinck, D., Sauter, A.W., Heye, T., Boll, D.T., Cyriac, J., Yang, S., et al.: Totalsegmentator: Robust segmentation of 104 anatomic structures in CT images. Radiology: Artificial Intelligence **5**(5), e230024 (2023)
21. Wu, Q., Guo, H., Li, R., Han, J.: Deep learning and machine learning in CT-based COPD diagnosis: Systematic review and meta-analysis. International journal of medical informatics **196**, 105812 (2025)
22. Wu, Y., Qi, S., Feng, J., Chang, R., Pang, H., Hou, J., Li, M., Wang, Y., Xia, S., Qian, W.: Attention-guided multiple instance learning for COPD identification: To combine the intensity and morphology. Biocybernetics and Biomedical Engineering **43**(3), 568–585 (2023)
23. Wu, Y., Xia, S., Liang, Z., Chen, R., Qi, S.: Artificial intelligence in COPD CT images: identification, staging, and quantitation. Respiratory Research **25**(1), 319 (2024)
24. Xue, M., Jia, S., Chen, L., Huang, H., Yu, L., Zhu, W.: CT-based COPD identification using multiple instance learning with two-stage attention. Computer Methods and Programs in Biomedicine **230**, 107356 (2023)
25. Zhang, H., Zhao, M., Liu, M., Luo, J., Guan, Y., Zhang, J., Xia, Y., Zhang, D., Zhou, X., Fan, L., et al.: Hierarchical multiple instance learning for COPD grading

with relatively specific similarity. In: International Conference on Medical Image Computing and Computer-Assisted Intervention. pp. 536–545. Springer (2024)
26. Zhang, L., Jiang, B., Wisselink, H.J., Vliegenthart, R., Xie, X.: COPD identification and grading based on deep learning of lung parenchyma and bronchial wall in chest CT images. The British Journal of Radiology **95**(1133), 20210637 (2022)