

RESEARCH ARTICLE

Lung cancer diagnosis using deep attention-based multiple instance learning and radiomics

Junhua Chen | Haiyan Zeng | Chong Zhang | Zhenwei Shi | Andre Dekker | Leonard Wee* | Inigo Bermejo*

Department of Radiation Oncology (MAASTRO), GROW School for Oncology and Developmental Biology, Maastricht University Medical Centre+, Maastricht, The Netherlands

Correspondence

Junhua Chen, Department of Radiation Oncology (MAASTRO), GROW School for Oncology and Developmental Biology, Maastricht University Medical Centre+, Maastricht, 6229 ET, The Netherlands. Email: j.chen@maastrichtuniversity.nl

*Leonard Wee and Inigo Bermejo are equally senior author.

Funding information

China Scholarship Council scholarship, Grant/Award Number: 201906540036; China Scholarship Council scholar, Grant/Award Number: 201909370087; STRaTegy, Grant/Award Number: STW 14930; BIONIC, Grant/Award Number: NWO 629.002.205; TRAIN, Grant/Award Number: NWO 629.002.212; CARRIER, Grant/Award Number: NWO 628.011.212; The Hanarth Foundation for LW

Abstract

Background: Early diagnosis of lung cancer is a key intervention for the treatment of lung cancer in which computer-aided diagnosis (CAD) can play a crucial role. Most published CAD methods perform lung cancer diagnosis by classifying each lung nodule in isolation. However, this does not reflect clinical practice, where clinicians diagnose a patient based on a set of images of nodules, instead of looking at one nodule at a time. Besides, the low interpretability of the output provided by these methods presents an important barrier for their adoption.

Method: In this article, we treat lung cancer diagnosis as a multiple instance learning (MIL) problem, which better reflects the diagnosis process in the clinical setting and provides higher interpretability of the output. We selected radiomics as the source of input features and deep attention-based MIL as the classification algorithm. The attention mechanism provides higher interpretability by estimating the importance of each instance in the set for the final diagnosis. To improve the model's performance in a small imbalanced dataset, we propose a new bag simulation method for MIL.

Results and conclusion: The results show that our method can achieve a mean accuracy of 0.807 with a standard error of the mean (SEM) of 0.069, a recall of 0.870 (SEM 0.061), a positive predictive value of 0.928 (SEM 0.078), a negative predictive value of 0.591 (SEM 0.155), and an area under the curve (AUC) of 0.842 (SEM 0.074), outperforming other MIL methods. Additional experiments show that the proposed oversampling strategy significantly improves the model's performance. In addition, experiments show that our method provides a good indication of the importance of each nodule in determining the diagnosis, which combined with the well-defined radiomic features, to make the results more interpretable and acceptable for doctors and patients.

KEYWORDS

attention mechanism, lung cancer diagnosis, multiple instance learning, radiomics

1 | INTRODUCTION

According to the statistics from the World Health Organization (WHO), lung cancer is the most frequently diagnosed malignant carcinoma and the leading cause of cancer death worldwide, accounting for an estimated

2.09 million deaths in 2018.^{1,2} Early diagnosis and treatment can reduce a lung cancer patient's mortality significantly. A plausible method for early lung cancer diagnosis is the routine use of low-dose computed tomography (CT) scans.³ To date, radiologists typically need to visually inspect CT scans slice by slice, which

This is an open access article under the terms of the [Creative Commons Attribution](https://creativecommons.org/licenses/by/4.0/) License, which permits use, distribution and reproduction in any medium, provided the original work is properly cited.

© 2022 The Authors. *Medical Physics* published by Wiley Periodicals LLC on behalf of American Association of Physicists in Medicine.

is costly and time-consuming as well as susceptible to human error.^{4,5} Computer-aided diagnosis (CAD) for rapid early lung nodules classification based on low-dose CT imaging has therefore attracted much attention from researchers during the last decades.^{6,7}

The development of CAD for lung nodules classification has reached new peaks in last decade mainly due to breakthroughs in deep learning neural networks⁸ and its application to a wide range of medical image analysis tasks. Several deep learning-based lung nodule classification methods have been proposed in recent years, with steadily improving state-of-the-art performance. Shen et al.⁹ developed a multiscale convolutional neural network (CNN) to extract features (referred to as “deep features”¹⁰ in the literature) and then applied a supervised random forest classifier to the deep features, reporting an accuracy of 86%. Xie et al.¹¹ combined handcrafted features with deep features to classify each nodule as either benign or malignant, achieving an area under the curve (AUC) of 0.96. Alakwaa et al.¹² combined the LUNA16¹³ dataset with a subset of the National Lung Screening Trial (NLST),¹⁴ and then used a pretrained U-Net to segment potential nodules from a CT scan automatically. The segmented nodules were passed to a 3D CNN to detect early-stage lung cancer, achieving an AUC of 0.83 in a randomly-split test cohort from the abovementioned data. Ardila et al.¹⁵ developed an end-to-end set of 3D CNN modules to compute the overall risk of lung malignancy based on autodetection of nodules, using the full-size publicly available NLST dataset. In a retrospective reader study, their model outperformed six experienced radiologists with absolute reductions of 11% and 5% in false positives and false negatives, respectively.

The need for transparency, interpretability, and explainability in such computer-aided diagnostic recommendations will grow to become increasingly prominent in the immediate future. A crucial piece of law, the General Data Protection Regulation (GDPR), governs the rights of European Union (EU) citizens as human data subjects and addresses processing by automated means for decision-making anywhere in the world if it concerns an EU individual. Specifically, the GDPR enshrines the right of an individual to receive “meaningful information about the logic involved” in an automated decision concerning them, and on that basis to either legally challenge the decision, or exercise conscientious objection to the use of an automated means for deriving the decision.¹⁶

Although definition of “meaningful” is open for debate, it is clearly helpful to be able to point at specific regions of interest (ROIs) that were strongly triggering for the diagnostic recommendation, along with related features of lung cancer and nonlung cancer cases. In this way, a human radiologist can review the information in depth, and either confirm or overrule the recommendations of an automated system. Irrespective of a right to an explanation, a computerized diagnostic support system

with high transparency and high interpretability would be immensely valuable in clinical practice.

For automated diagnosis of lung cancer, a deep learning-based system can be applied in two levels: at nodule level, to identify potential malignant nodule(s) for further biopsy and performing diagnosis at patient level. Generally speaking, nodule classification methods need a label for each nodule to be able to train a model.^{9,11} However, labeling each nodule is more time-consuming and expensive than having a label for each patient, which is usually already available in hospital records. In this study, we focus on deep learning methods for lung cancer diagnosis that can make use of the existing data to develop a lung cancer CAD system that classifies patients based on multiple suspected nodules in the entire CT series without the need to assign a label to each nodule (i.e., each instance), and at the same time provide high visibility of the triggering features of its recommendation. Multiple instance learning (MIL) with attention mechanism^{17,18,19} fits this need well. In MIL, the nodules are grouped into “bags of instances” (assuming multiple nodules in one CT examination of the chest area). The task is hence to determine the diagnosis for the subject as a whole. Only the subject-level diagnoses (i.e., the bag labels) are needed, but not individual labels of every nodule found in the subject.²⁰ This approach is thus more amenable to real-world data mining in lung cancer, because the subject-level diagnosis is much more widely available than annotations on each nodule.

Research on MIL problems has progressed along instance-level versus embedding-level solutions,²¹ with the latter seeming to perform better at subject-level classification.²² Widely used embedding approaches include MI-SVM,²³ mi-Graph,²⁴ miVLAD,²⁵ and MI-Net,²⁵ but the shortcoming of these is the lack of transparency of triggering instance(s). An attention-based deep MIL^{21,33} has been recently introduced, which allows a deep learning model to estimate the contribution of each instance to the predicted subject label, using the well-established attention mechanism.²⁶

The objective of this work was to develop a lung cancer classification model at the subject (patient) level from multiple examined nodules, without the need to have specific expert findings reported at the level of each individual nodule. An MIL method with an additional deep attention mechanism was used to help draw an expert clinician’s eye toward the individual nodules that were strongly triggering for the model’s diagnostic recommendation. We propose that this will be important by way of offering better interpretability and the possibility of human expert verification of the internal logic of the algorithm. A selection of commonly used hand-crafted radiomics features was used as a source of image features,^{27,28,29} and we also compared a number of alternative MIL methods.³⁰ We have reused an existing open access data collection for training and cross-validation.

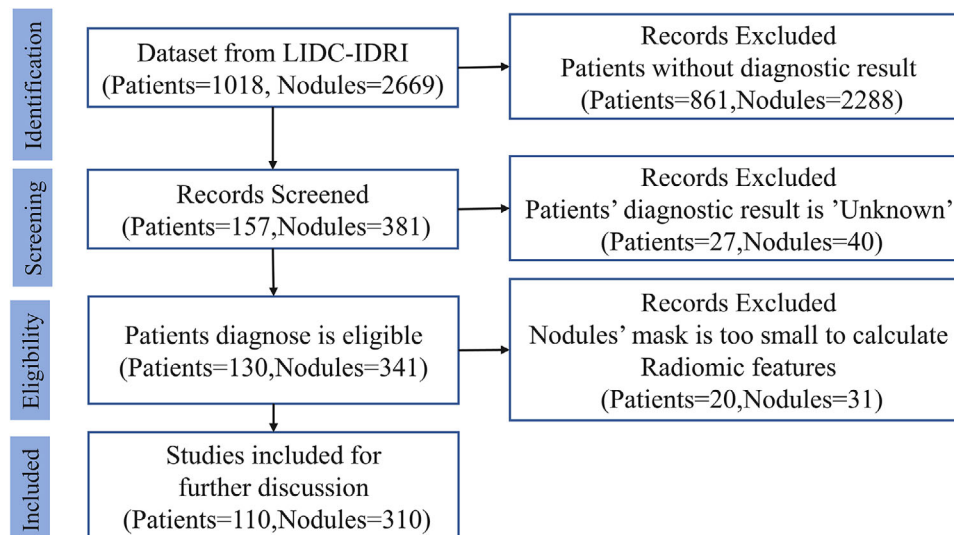


FIGURE 1 Sample selection flowchart describing the number of subjects and the number of nodules selected for this analysis

This article is organized as follows: the methods and classifier experiments are described in Section 2. Our results are given in Section 3. The significance of our findings and limitations of our current approach will be discussed in Section 4. An overall summary and conclusion are presented in Section 5. Finally, source code will open access for public at <https://gitlab.com/UM-CDS/combine-mil-and-radiomics-for-lung-screening>) and additional details of the system architecture will be given in the Supporting Information.

2 | METHODS

2.1 | Dataset

The primary data source of data is an open access collection from the Lung Image Database Consortium (LIDC-IDRI),³¹ accessed at The Cancer Imaging Archive (TCIA) during May 2020³² under a Creative Commons Attribution Non-Commercial 3.0 Unported (CC BY-NC) license. The details of subjects in LIDC-IDRI have been provided elsewhere,³¹ but briefly, the collection comprises 1018 clinical chest CT examinations from seven disjoint institutions. Radiologists working independently entered 7371 annotations, of which there were 2669 consensus nodules. We excluded subjects with unreported or unknown diagnosis, and excluded nodules below 3 mm in diameter according to current diagnosis protocols.^{34,35} This resulted in 110 unique subjects with a total of 310 nodules eligible for consideration. Binary masks for the nodules were provided in the data collection as an XML file. Numbers of subjects and nodules excluded, along with reason, are provided in Figure 1. From the summary of diagnostic findings in Table 1, we note that the majority of subjects and lung nodules in the dataset are positive for lung cancer; 75% and 77%,

respectively. Index of available patients for experiments in LIDC-IDRI can be found in Table 1 in the Supporting Information.

Table 1 summarizes the radiological findings available in the selected subset with definitive subject-level diagnosis and nodule-level classification.

2.2 | Image acquisition settings

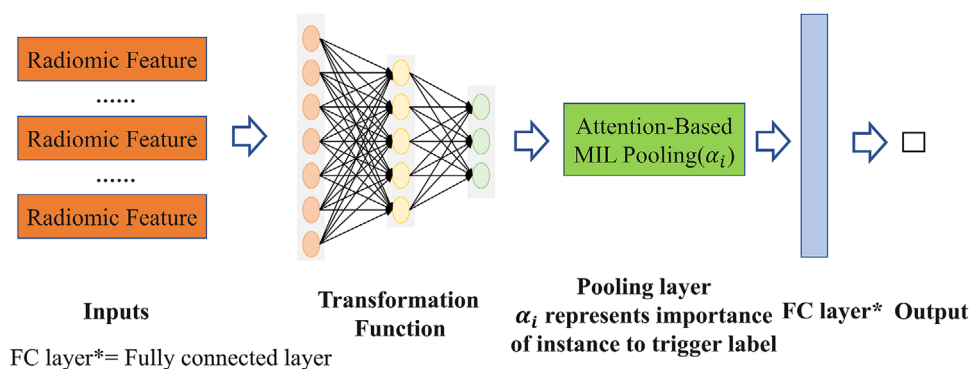
The LIDC-IDRI contains a heterogeneous set of CT of subjects from different institutions. We used axial CT images with dimension of 512×512 pixels. Radiation exposure of selected samples ranged from 3 to 534 mAs (median: 147.5 mAs), and reconstructed slice thicknesses ranging from 0.6 to 5.0 mm (median: 2.0 mm).

2.3 | Feature extraction

Radiomics features were extracted using an open-source Python library pyRadiomics (v2.2.0).³⁶ Images were resampled to 2 mm isotropic voxels prior to feature extraction. A total of 103 features were extracted. These consisted of 13 morphology (shape) features, 17 intensity-histogram (first-order) features, and 73 textural (Haralick) features. Binary masks for the GTV were generated from the XML file in the LIDC-IDRI collection, using an open-access library *pylidc*.³⁷ DICOM CT images were converted to 3D images by using SimpleITK (v1.2.4)³⁸ for pyRadiomics feature extraction. The mathematical definition of each feature has been given in the online documentation. Our pyRadiomics extraction settings (from the params.yaml file) have been included in Table S2 of the Supporting Information. All features included in this analysis have been listed in Table S3 of the Supporting Information.

TABLE 1 Number of patients and nodules according to ground truth diagnosis in the dataset

	Lung cancer	Not lung cancer	Total
Numbers of (% of total) patients	82 (75%)	28 (25%)	110
Numbers of (% of total) nodules	239 (77%)	71 (23%)	310

**FIGURE 2** Architecture of the attention-based deep MIL. Extracted radiomics features are used as the input to the transformation network, which is then pooled with attention. A fully connected final layer combines the attention-based pooling to give the output probability

2.4 | Classifier

We used an attention-based MIL for the lung cancer classifier component. This consists of two parts that can be trained end-to-end. First, the transformation network was implemented as three fully connected neural network layers with a dropout rate of 0.5. Additional details about this network are in Table S4 in the Supporting Information. To fix the dimension of the input layer of neurons, the 103 features per nodule were duplicated within the same subject until it was the same as the maximum number of nodules per subject, which we found to be 12 in this case. More specifically, each nodule in the same bag should be duplicated with the same probability. For example, if there are five nodules in a bag, three random nodules need to be duplicated once (i.e., appear twice in total), and two random nodules need to be duplicated twice (appear three times in total) in the final fixed feature bags. Therefore, the dimension of the input layer should be 103 and one bag consists of 12 vectors (103×12). Feature duplication was performed before model training and was also used in model testing.

Second, the attention-based pooling layer implemented the attention mechanism popularized by long short-term memory networks (LSTMs).³⁹ The attention mechanism is an important strategy that fits encoder input sequences into a fixed-length internal representation. The architecture of the classifier is illustrated schematically in Figure 2.

2.5 | Addressing class imbalance

Imbalance in the outcome frequency (i.e., lung cancer vs. not lung cancer) has been known to affect the classifier,

biasing this toward the dominant class. Several methods are available to address class imbalance⁴¹ in general, and we applied a novel sampling method to address class imbalance specifically for MIL. It is assumed that all nodules in noncancer subjects are, by clinical definition, noncancerous nodules. Synthetic noncancer patients were thus generated by randomly sampling a finite number of instances out of all the nodules in an aggregated pool of actual noncancer subjects. On the other hand, synthetic cancer patients could be generated by adding a random number of negative instances sampled from the instances pool (from both negative and positive bags) to the original positive bags. However, we did not simulate cancer patients in our experiments, because positive bags were majority in our dataset. This was only done for the training set; no class imbalance correction was applied in the testing set.

2.6 | Model development and validation

All work was executed on a Core i7 8565U CPU with 8GB of RAM. The optimizer for network training was stochastic gradient descent (SGD),⁴² with batch size 1 and the learning rate fixed at 0.0001. The neural network was trained for 500 epochs (taking 3–4 min) per experiment.

We performed experiments for the attention-based MIL in comparison with other MIL approaches—MI-SVM, mi-graph, miVLAD, MI-Net, and a naïve MIL algorithm that performs a simple aggregation of the predictions by replacing the attention-based MIL pooling with average MIL pooling.²⁰ The optimizer, batch size, learning rate, and training epochs were set same as attention-based MIL in MI-Net. The setting of hyperparameters in other methods was followed as

mentioned in original literatures.^{23,24,25} Same training and testing data were used in every running for all methods.

Model training was performed on all the available subjects, taking their respective diagnosis as the “bag label” and the nodules as the instances. We ran 20 repetitions of end-to-end training runs on the hand-crafted features with fivefold cross-validation in each run and there is no oversampling in testing dataset. For each repetition of fivefold cross-validation, samples were randomly sorted first and then split into fivefolds, so that each sample was used once for testing and four times for training in each repetition. We adjusted for the lower number of non-cancer diagnoses by generating synthetic noncancer patients as described above (Section 2.4). Specifically, we synthesized 60 additional noncancer subjects from the initial training dataset and added these to the actual 88 training subjects, resulting in a training set containing 148 subjects in total. No synthetic resampling was used for positive lung cancer subjects. We further conducted an additional sensitivity analysis to assess how oversampling to overcome class imbalance might have affected the model’s performance by using only the original data of 110 subjects.

The discriminative performance was assessed using the mean and standard error of the mean (SEM) of recall, accuracy, positive predictive value (PPV), and negative predictive value (NPV), respectively. For dichotomization of outcome, we used a probability threshold of 0.5 to separate lung cancer from nonlung cancer. The area under the receiver-operating characteristic curve (AUC) was computed for each model, and the definition of AUC can be found in Ref. 43. Let TP, TN, FP, and FN denote true positive, true negative, false positive, and false negative, respectively, and then we define recall, accuracy, PPV, and NPV as:

$$\text{recall} = \frac{TP}{TP + FN}, \quad (1)$$

$$\text{accuracy} = \frac{TP + TN}{TP + TN + FP + FN}, \quad (2)$$

$$\text{PPV} = \frac{TP}{TP + FP}, \quad (3)$$

$$\text{NPV} = \frac{TN}{TN + FN}. \quad (4)$$

All statistical analyses were done in Python (version 3.6.1).

3 | RESULTS

Figure 3 shows the violin plots comparing the results of attention-based MIL with (Figure 3a) and without (Figure 3b) synthetic minority oversampling. The estimated mean (with SEM in the parentheses) for recall, accuracy, PPV, and NPV and the AUC for the model including the class imbalance correction were: 0.870 (SEM 0.061), 0.807 (SEM 0.069), 0.928 (SEM 0.078), 0.591 (SEM 0.155), and 0.842 (SEM 0.071), respectively. Without the class imbalance correction, these values were 0.889 (SEM 0.061), 0.768 (SEM 0.059), 0.842 (SEM 0.071), 0.483 (SEM 0.209), and 0.696 (SEM 0.108), respectively. The main effect of the minority oversampling was to improve accuracy, PPV, NPV, and AUC. A representative (from a selected repetition) set of AUC curves for the different MIL methods with the same training and testing data can be found in Figure 4.

Table 2 summarizes the results of comparing different MIL approaches. Attention-based MIL without oversampling achieved the best recall, MI-Net achieved the best PPV, and attention-based MIL achieved the best accuracy, PPV, and AUC. Attention-based MIL was better than other methods in PPV and AUC significantly (Wilcoxon test, $p < 0.01$); however, attention-based MIL was worse than best result in recall and NPV (Wilcoxon test, $p = 0.02$ and $p < 0.01$, respectively). Moreover, attention-based MIL with oversampling is better than attention-based MIL without oversampling in all metrics significantly except recall (Wilcoxon test, $p < 0.01$ for accuracy, PPV, NPV, and AUC, $p = 0.02$ for recall).

The absence of AUCs for mi-graph and miVLAD is due to our reusing of the source code by the LAMDA lab, Nanjing University.⁴⁴ Their source code for mi-graph and miVLAD outputs only the classification label (not the probability), and therefore, the AUCs cannot be calculated.

To determine the level of oversampling, we ran sensitivity analyses. We gradually increased the number of included simulated noncancer subjects from 0 to 100 on steps of 20. We ran 20-repeat fivefold cross-validation for each experiment. The results of sensitivity analysis are shown in Figure 5.

As shown in Figure 5, including 60 simulation samples results in good performance for all metrics (especially for recall) with less computation compared with other settings with similar performance.

Given how important batch size is for CNNs training,⁴⁵ we ran a sensitivity analysis on this parameter. We ran 20-repeat fivefold cross-validation analyses with different batch sizes (1–4) for each experiment. The loss curves for model training with different batch

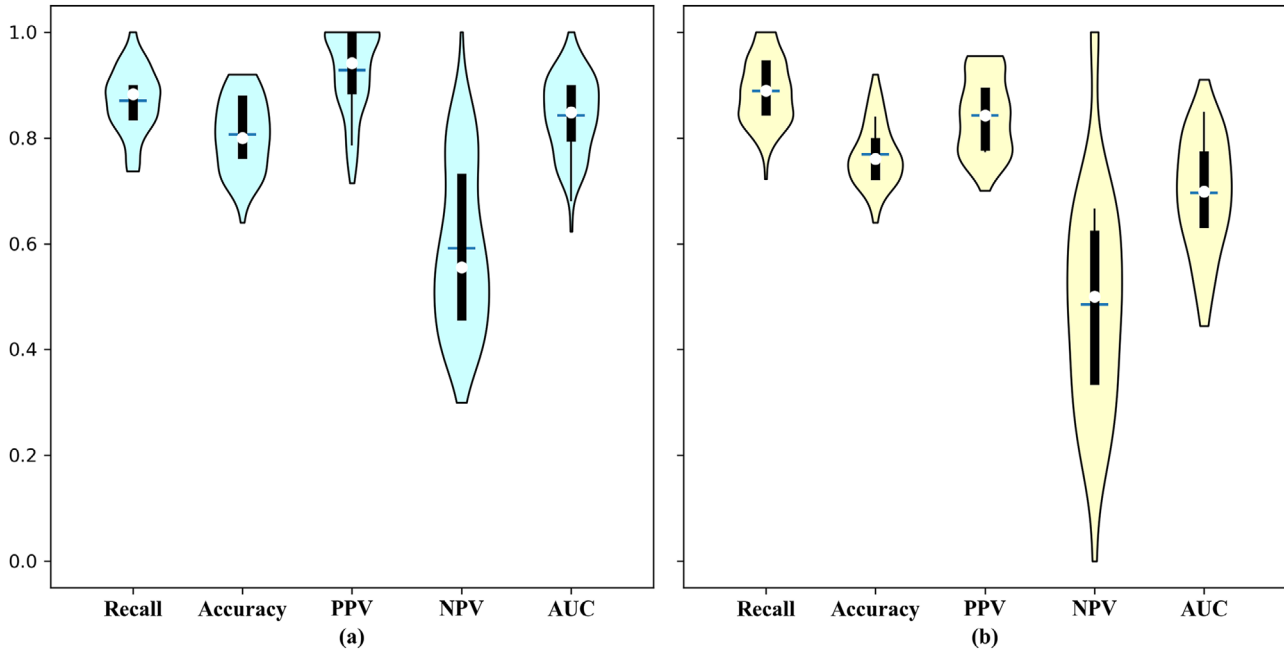


FIGURE 3 Violin plot of the experimental results (a) with oversampling and (b) without oversampling

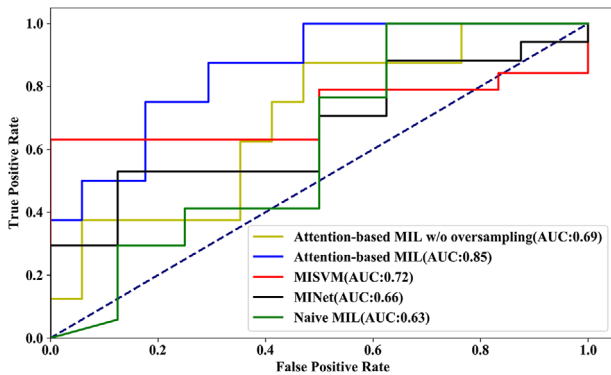


FIGURE 4 An example of AUC curves for different methods with same training and testing data. An AUC curves for attention-based MIL, attention-based MIL w/o oversampling, MI-SVM, MI-Net, and naïve MIL

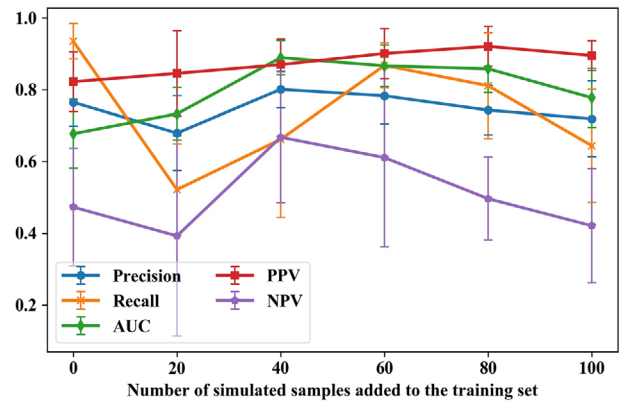


FIGURE 5 Results of sensitivity analysis for different levels of oversampling

sizes are shown in Figure 6a and the performance of models trained with different batch sizes is shown in Figure 6b.

As shown in Figure 6, the model trained with a batch size of 1 achieved the best performance according to all metrics except AUC (0.842 for batch size 1 vs. 0.849 for

TABLE 2 Results of the attention-based deep MIL approach with class imbalance correction, compared to other MIL methods (attention-based MIL w/o oversampling, MI-SVM, mi-graph, miVLAD, and MI-Net)

Methods	Attention-based MIL	Attention-based MILw/o oversampling	MI-SVM	mi-graph	miVLAD	MI-Net	Traditional MIL
Recall	0.870 ± 0.061	0.889 ± 0.061	0.756 ± 0.084	0.777 ± 0.048	0.871 ± 0.087	0.835 ± 0.109	0.850 ± 0.099
Accuracy	0.807 ± 0.069	0.768 ± 0.059	0.703 ± 0.080	0.749 ± 0.055	0.782 ± 0.063	0.727 ± 0.050	0.748 ± 0.065
PPV	0.928 ± 0.078	0.842 ± 0.071	0.560 ± 0.199	0.772 ± 0.042	0.835 ± 0.059	0.522 ± 0.265	0.835 ± 0.070
NPV	0.591 ± 0.155	0.483 ± 0.209	0.810 ± 0.080	0.713 ± 0.229	0.675 ± 0.160	0.838 ± 0.069	0.478 ± 0.233
AUC	0.842 ± 0.071	0.696 ± 0.108	0.625 ± 0.099	–	–	0.662 ± 0.093	0.681 ± 0.080

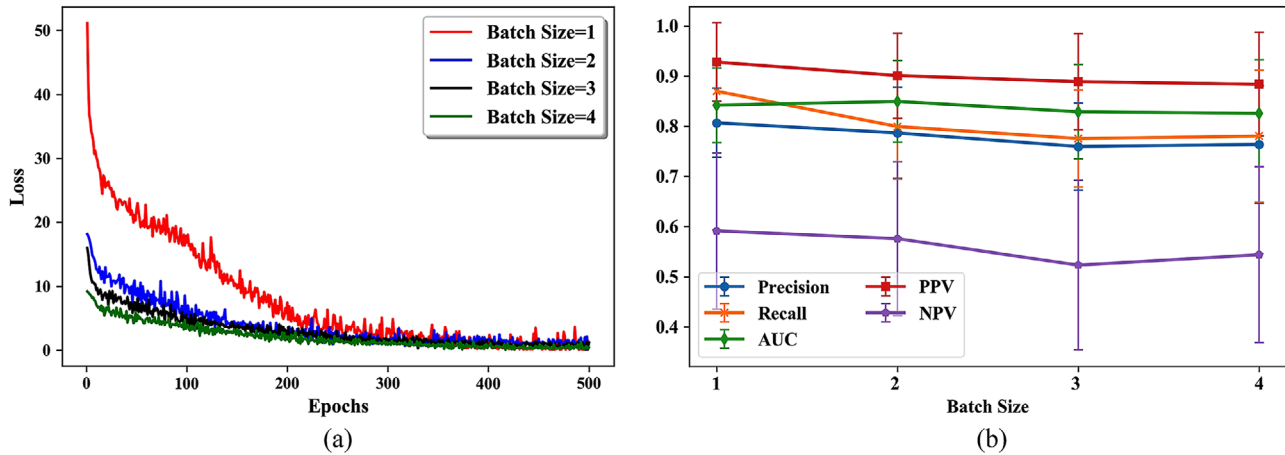


FIGURE 6 Results of sensitivity analysis for different batch sizes. (a) Loss curves for model training with different batch size and (b) performance of models trained with different batch sizes




FIGURE 7 An example of attention weights for two positive lung cancer subjects (LIDC-IDRI-1004 and 1011)

batch size 2) and the loss of all models converged at the end of the 500 epochs. Therefore, we set the batch size to 1 in this study.

Besides model performance, one of the most appealing aspects that we selected the attention-based MIL method for was to indicate the instances that might have been strongly influential on the classification. In this case, it would be the relative importance of each nodule when predicting the subject label as either lung cancer or not lung cancer. A couple of lung cancer examples are shown in Figure 7 for two subjects in the dataset, LIDC-IDRI-1004 and LIDC-IDRI-1011.

Alpha in Figure 7 means the strength of attention, value of alpha only meaningful in the same patient, and it is meaningless by comparing alphas across patients.

The order of nodules was arranged in random way within same patient.

The evaluation of the attention mechanism was performed by one of the coauthors—a radiologist with 3-year experience, who examined some sample patients' weights and agreed with the weighting. In these examples, it is clearly discernable from the weights (α_2 and α_3 larger than α_0 and α_1) that the two rightmost nodules pictured for subject LIDC-IDRI-1004 are much more strongly influential in the diagnostic evaluation compared to the two leftmost nodules. Similarly, for subject LIDC-IDRI-1011, three of the nodules are influential on the subject classification, but the nodule pictured rightmost is not influential at all (alpha < 0.01).

4 | DISCUSSION

Our objective was to propose a lung cancer classification at the subject level from multiple examined nodules, with an attention mechanism for improving the interpretability. The results show that our proposed classification achieves good performance compared to other MIL methods and that the unique characteristic of the deep attention-based MIL, namely, attention weights, potentially makes our method more interpretable for clinicians.

To see the effect of minority oversampling to overcome class imbalance, we tested the model with and without the oversampling. The results show that the oversampling improved the model's performance significantly in accuracy, PPV, NPV, and AUC by comparing Attention-based MIL without oversampling. However, there seem some decrease in recall.

We observed from Figure 3 that minority oversampling has a major effect on the AUC. In fact, the AUC sinks below 0.5 in some experiments without oversampling. This can be explained by the fact that the AUC is more sensitive to the classification performance of the model with the minority class than either accuracy or recall.⁴⁰

We proposed a new synthetic subject generation method that can be used to overcome class imbalance by oversampling the minority class. We did this by sampling from an aggregated pool of nodules from patients with the ground truth of "not lung cancer." To the best of our knowledge, such methods have not yet been proposed for MIL. This oversampling technique resulted in significant improvements on accuracy, PPV, NPV, and AUC. We believe that this strategy, which is based on the characteristics of MIL, can be used when training any MIL model from a class-imbalanced dataset.

The results show that our method could potentially be applied to automated lung cancer diagnosis, subject to further validation and studies in large datasets. However, we acknowledge that there are some limitations and weaknesses in the assumptions we had to make. First, due to the need of a mask that delineates the nodules to calculate radiomic features, our method would have to be dependent on lung nodule detection and segmentation methods such as the ones proposed by Huang et al.⁴⁶ and Anirudh et al.⁴⁷ This dependence on pre-existing or human expert segmentation is not new, and is problem that still affects many aspects of medical image analysis and supervised machine learning. Related to this is the potential for interobserver disagreement about the external outline of the nodule. This problem is well known and documented for large and locally advanced lung tumors, but with the small nodule volumes involved in this study, we have assumed that the interobserver problem does not strongly affect the extracted features. A further question we cannot address in this study is the problem of undetected nodules and very small nodules (diameter smaller than 3 mm) that were omitted; more-

over, images were resampled to 2 mm isotropic voxels prior to feature extraction, which is possibly also a reason why very small nodules are not appropriate. This work has assumed no false positives and no false negatives, so we cannot elucidate what happens with imperfect nodule detection.

The performance of our model appears sensitive to sampling effects; in other words, the performance of model fluctuates across repeated experiments, as shown in Figure 3. This is likely a direct consequence of the relatively small sample size of the dataset. Expanding the sample size by including small nodules is not immediately helpful because they do not add that many subjects and nodules to the sample, whereas hand-crafted features would not be stable when taken from very small volumes. The major root of the problem appears to be the lack of ground truth and annotated images. Related to this fact is that we currently did not find a suitable dataset for external, independent validation. Therefore, our results should be interpreted as preliminary indication of feasibility, and larger datasets need to be used to demonstrate wider generalizability of this work.

Due to the high fitting ability of neural networks and large epochs during training, the model returns 1 or 0 almost all of the time, which means that the overall model calibration was generally poor.^{48,49} Model calibration plot is shown in Figure S1 of the Supporting Information, and it appears that all MIL methods have poor calibration except MI-SVM.

In addition, we have not explored feature dimensionality reduction and applied feature redundancy analysis. This is in part due to the transformation network that does not require explicit feature selection steps prior to MIL pooling. The repeatability and reproducibility of handcrafted features are subjects of numerous investigations in radiomics and appear to be highly modality-specific. This work has not explored the stability of low-dose CT-derived image features, which tend to have quite a lot of noise present.^{50,51} This could affect the performance of our model in an external validation, and image harmonization or denoising strategies may be needed in future to support general extensibility.

Moreover, we were not able to test the performance of the models in an external dataset, which would have provided more reliable estimates of the models' potential performance in a different setting. On the other hand, the dataset used in this study (LIDC-IDRI) was collected over 10 years ago. With new emerging CT technologies and reconstruction methods, it is possible that different conclusions would be reached if the proposed method is applied to newer images currently being used in clinical practice. Further research on this aspect is required.

Finally, our oversampling strategy is sensitive to the quality of data's label at patient level. More specifically, if labels are incorrect (e.g., if one or two of the nodules has been misclassified by error and the subject is hence

a false negative), the noise will be amplified due to over-sampling.

For future work, an automated nodule detection and segmentation algorithm could be attached to this attention-based MIL classifier to fully complete the lung cancer diagnosis workflow. Second, methods for improving radiomic features' reliability in low-dose CT may be necessary for improving model's performance in unseen data. Third, large scale and comprehensive evaluation of the attention mechanism is needed in the future to assess its reliability and reproducibility. Fourth, a comparison between the proposed method and a traditional deep learning-based image classification algorithm would be of special interest. Finally, the proposed model needs to be externally validated to assess whether the model suffers from overfitting to the training data or whether it is widely generalizable to CT images from different scanners.

5 | CONCLUSION

We treated computer-aided diagnosis of lung cancer as an MIL problem, such that the classification as lung cancer or not is made at the subject level (i.e., the patient) without relying on classifications at the level of individual nodules (i.e., each of suspicious lung nodules). The addition of the attention mechanism was used to draw the clinician's eye toward features that were important for triggering the recommended diagnosis, with the aim of supporting interpretability and, importantly, verification by human experts of the algorithm's internal logic. We used radiomics as a source of interpretable image-derived features, and deep attention-based MIL was found to be a superior classifier compared to other MIL options with regard to accuracy, NPV, and AUC. A novel approach for minority oversampling, adapted for MIL problems, has been used to address the outcome class imbalance in the LIDC-IDRI dataset. We showed how an attention mechanism could be used as an indication of the importance of each nodule for triggering the diagnostic recommendation. Cross-validation was used to check for model performance, but more data are required to provide a robust test of wider generalizability.

ACKNOWLEDGMENTS

The authors would like to thank the authors of the open-access source code for mi-graph and miVLAD from LAMDA lab (<http://210.28.132.67/Data.ashx>). JC is supported by a China Scholarship Council scholarship (201906540036). HZ is supported by a China Scholarship Council scholar (201909370087). The coauthors acknowledge funding support from the following: STRaTegy (STW 14930), BIONIC (NWO 629.002.205), TRAIN (NWO 629.002.212), CARRIER (NWO 628.011.212), and The Hanarth Foundation for LW.

CONFLICT OF INTEREST STATEMENT

The authors have no conflicts to disclose.

REFERENCES

1. Cancer. [WWW document]. URL <https://www.who.int/news-room/fact-sheets/detail/cancer> [access on 20 June 2020], 2018.
2. Siegel RL, Miller KD, Jemal A. Cancer statistics, 2020. *CA Cancer J Clin.* 2020;70(1):7-30. <http://doi.org/10.3322/caac.21590>
3. National Lung Screening Trial Research Team. (2011). Reduced lung-cancer mortality with low-dose computed tomographic screening. *N Engl J Med.* 2011;365(5):395-409.
4. Kanazawa K, Kawata Y, Niki N, et al. Computer-aided diagnosis for pulmonary nodules based on helical CT images. *Comput Med Imaging Graph.* 1998;22(2):157-167.
5. Naqi SM, Sharif M, Jaffar A. Lung nodule detection and classification based on geometric fit in parametric form and deep learning. *Neural Comput Appl.* 2020;32(9):4629-4647.
6. Parveen SS, Kavitha C. A review on computer aided detection and diagnosis of lung cancer nodules. *Int J Comput Technol.* 2012;3(3a):393-400.
7. Yang Y, Feng X, Chi W, et al. Deep learning aided decision support for pulmonary nodules diagnosing: a review. *J Thorac Dis.* 2018;10(Suppl 7):S867.
8. Krizhevsky A, Sutskever I, Hinton GE. Imagenet classification with deep convolutional neural networks. In: *Advances in Neural Information Processing Systems.* 2012:1097-1105.
9. Shen W, Zhou M, Yang F, Yang C, Tian J. Multi-scale convolutional neural networks for lung nodule classification. In: *International Conference on Information Processing in Medical Imaging.* Springer; 2015:588-599.
10. Zhou B, Khosla A, Lapedriza A, Oliva A, Torralba A. Learning deep features for discriminative localization. In: *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition.* 2018:2921-2929.
11. Xie Y, Zhang J, Xia Y, Fulham M, Zhang Y. Fusing texture, shape and deep model-learned information at decision level for automated classification of lung nodules on chest CT. *Inf Fus.* 2018;42:102-110.
12. Alakwaa W, Nassef M, Badr A. Lung cancer detection and classification with 3D convolutional neural network (3D-CNN). *Lung Cancer.* 2017;8(8):409.
13. Setio AAA, Traverso A, De Bel T, et al. Validation, comparison, and combination of algorithms for automatic detection of pulmonary nodules in computed tomography images: the LUNA16 challenge. *Med Image Anal.* 2017;42:1-13.
14. Black WC, Gareen IF, Soneji SS, et al. Cost-effectiveness of CT screening in the national lung screening trial. *N Engl J Med.* 2014;371(19):1793-1802.
15. Ardila D, Kiraly AP, Bharadwaj S, et al. End-to-end lung cancer screening with three-dimensional deep learning on low-dose chest computed tomography. *Nat Med.* 2019;25(6):954-961.
16. Selbst A, Powles J. "Meaningful information" and the right to explanation. In: *Conference on Fairness, Accountability and Transparency.* PMLR; 2018:48-48.
17. Data protection in the EU. [WWW document]. URL https://ec.europa.eu/info/law/law-topic/data-protection/data-protection-eu_en [access on 5 July 2020], 2020.
18. Dietterich TG, Lathrop RH, Lozano-Pérez T. Solving the multiple instance problem with axis-parallel rectangles. *Artif Intell.* 1997;89(1-2):31-71.
19. Maron O, Lozano-Pérez T. A framework for multiple-instance learning. In: *Advances in Neural Information Processing Systems.* 1998:570-576.
20. Bhattacharjee K, Pant M, Zhang YD, Satapathy SC. Multiple instance learning with genetic pooling for medical data analysis. *Pattern Recognit Lett.* 2020;133:247-255.

21. Ilse M, Tomczak JM, Welling M. Attention-based deep multiple instance learning. In *35th International Conference on Machine Learning, ICML 2018*. International Machine Learning Society (IMLS); 2018:3376-3391.
22. Wang X, Yan Y, Tang P, Bai X, Liu W. Revisiting multiple instance neural networks. *Pattern Recognit*. 2018;74:15-24.
23. Andrews S, Tsochantaridis I, Hofmann T. Support vector machines for multiple-instance learning. In: *Advances in Neural Information Processing Systems*. 2003:577-584.
24. Zhou ZH, Sun YY, Li YF. Multi-instance learning by treating instances as non-iid samples. In: *Proceedings of the 26th Annual International Conference on Machine Learning*; 2009:1249-1256.
25. Wei XS, Wu J, Zhou ZH. Scalable algorithms for multi-instance learning. *IEEE Trans Neural Netw Learn Syst*. 2016;28(4):975-987.
26. Sutskever I, Vinyals O, Le QV. Sequence to sequence learning with neural networks. In: *Advances in Neural Information Processing Systems*. 2014:3104-3112.
27. Afshar P, Mohammadi A, Plataniotis KN, Oikonomou A, Benali H. From handcrafted to deep-learning-based cancer radiomics: challenges and opportunities. *IEEE Signal Process Mag*. 2019;36(4):132-160.
28. Aerts HJ, Velazquez ER, Leijenaar RT, et al. Decoding tumour phenotype by noninvasive imaging using a quantitative radiomics approach. *Nat Commun*. 2014;5(1):1-9.
29. Lambin P, Rios-Velazquez E, Leijenaar R, et al. Radiomics: extracting more information from medical images using advanced feature analysis. *Eur J Cancer*. 2012;48(4):441-446.
30. Susto GA, Schirru A, Pampuri S, McLoone S, Beghi A. Machine learning for predictive maintenance: a multiple classifier approach. *IEEE Trans Ind Inf*. 2014;11(3):812-820.
31. Armato III SG, McLennan G, Bidaut L, et al. The lung image database consortium (LIDC) and image database resource initiative (IDRI): a completed reference database of lung nodules on CT scans. *Med Phys*. 2011;38(2):915-931.
32. Clark K, Vendt B, Smith K, et al. The Cancer Imaging Archive (TCIA): maintaining and operating a public information repository. *J Digit Imaging*. 2013;26(6):1045-1057.
33. Han Z, Wei B, Hong Y, et al. Accurate screening of COVID-19 using attention based deep 3D multiple instance learning. *IEEE Trans Med Imaging*. 2020;39:2584-2594.
34. Han F, Wang H, Zhang G, et al. Texture feature analysis for computer-aided diagnosis on pulmonary nodules. *J Digit Imaging*. 2015;28(1):99-115.
35. Dhara AK, Mukhopadhyay S, Dutta A, Garg M, Khandelwal N. A combination of shape and texture features for classification of pulmonary nodules in lung CT images. *J Digit Imaging*. 2016;29(4):466-475.
36. Van Griethuysen JJ, Fedorov A, Parmar C, et al. Computational radiomics system to decode the radiographic phenotype. *Cancer Res*. 2017;77(21):e104-e107.
37. Hancock MC, Magnan JF. Lung nodule malignancy classification using only radiologist-quantified image features as inputs to statistical learning algorithms: probing the Lung Image Database Consortium dataset with two statistical learning methods. *J Med Imag*. 2016;3(4):044504.
38. Yaniv Z, Lowekamp BC, Johnson HJ, Beare R. SimpleITK image-analysis notebooks: a collaborative environment for education and reproducible research. *J Digit Imaging*. 2018;31(3):290-303.
39. Hochreiter S, Schmidhuber J. Long short-term memory. *Neural Comput*. 1997;9(8):1735-1780.
40. He H, Garcia EA. Learning from imbalanced data. *IEEE Trans Knowl Data Eng*. 2009;21(9):1263-1284.
41. Haixiang G, Yijing L, Shang J, Mingyun G, Yuanyue H, Bing G. Learning from class-imbalanced data: review of methods and applications. *Expert Syst Appl*. 2017;73:220-239.
42. Bottou L. Stochastic gradient descent tricks. In: Orr Gb, ed.: *Neural Networks: Tricks of the Trade*. Springer; 2012:421-436.
43. Huang J, Charles XL. Using AUC and accuracy in evaluating learning algorithms. *IEEE Trans Knowl Data Eng*. 2005;17(3):299-310.
44. Learning and Mining from Data. [WWW document]. URL <http://210.28.132.67/Data.ashx>. [accesses on 4 Jun 2021], 2021.
45. Yong H, Huang J, Meng D, Hua X, Zhang L. Momentum batch normalization for deep learning with small batch size. In: *European Conference on Computer Vision*. Springer; 2020:224-240.
46. Huang X, Shan J, Vaidya V. Lung nodule detection in CT using 3D convolutional neural networks. In: *2017 IEEE 14th International Symposium on Biomedical Imaging (ISBI 2017)*. IEEE; 2017:379-383.
47. Anirudh R, Thiagarajan JJ, Bremer T, Kim H. Lung nodule detection using 3D convolutional neural networks trained on weakly labeled data. In: *Medical Imaging 2016: Computer-Aided Diagnosis*. International Society for Optics and Photonics; 2016:978532.
48. Steyerberg EW, Vickers AJ, Cook NR, et al. Assessing the performance of prediction models: a framework for some traditional and novel measures. *Epidemiology*. 2010;21(1):128.
49. Steyerberg EW, Vergouwe Y. Towards better clinical prediction models: seven steps for development and an ABCD for validation. *Eur Heart J*. 2014;35(29):1925-1931.
50. Bagher-Ebadian H, Siddiqui F, Liu C, Movsas B, Chetty IJ. On the impact of smoothing and noise on robustness of CT and CBCT radiomics features for patients with head and neck cancers. *Med Phys*. 2017;44(5):1755-1770.
51. Moons KG, Altman DG, Reitsma JB, et al. Transparent Reporting of a multivariable prediction model for Individual Prognosis or Diagnosis (TRIPOD): explanation and elaboration. *Ann Intern Med*. 2015;162(1):W1-W73.

SUPPORTING INFORMATION

Additional supporting information may be found in the online version of the article at the publisher's website.

How to cite this article: Chen J, Zeng H, Zhang C, et al. Lung cancer diagnosis using deep attention-based multiple instance learning and radiomics. *Med Phys*. 2022;49:3134–3143. <https://doi.org/10.1002/mp.15539>