

RESEARCH ARTICLE

An Accurate and Explainable Deep Learning System Improves Interobserver Agreement in the Interpretation of Chest Radiograph

HIEU H. PHAM^{1,2}, (Member, IEEE), HA Q. NGUYEN^{3,4}, HIEU T. NGUYEN^{3,4}, LINH T. LE⁵, AND LAM KHANH⁶

¹College of Engineering and Computer Science, VinUniversity, Hanoi 10000, Vietnam

²VinUni-Illinois Smart Health Center, VinUniversity, Hanoi 10000, Vietnam

³Smart Health Center, VinBigData JSC, Hanoi, Vietnam

⁴Vingroup Big Data Institute, Hanoi, Vietnam

⁵Department of Radiology, Hanoi Medical University Hospital, Hanoi 100000, Vietnam

⁶Department of Radiology, Hospital 108, Hanoi 122001, Vietnam

Corresponding author: Hieu H. Pham (hieu.ph@vinuni.edu.vn)

This work was supported by the Vingroup Big Data Institute (VinBigData).

ABSTRACT Interpretation of chest radiographs (CXR) is a difficult but essential task for detecting thoracic abnormalities. Recent artificial intelligence (AI) algorithms have achieved radiologist-level performance on various medical classification tasks. However, only a few studies addressed the localization of abnormal findings from CXR scans, which is essential in explaining the image-level classification to radiologists. Additionally, the actual impact of AI algorithms on the diagnostic performance of radiologists in clinical practice remains relatively unclear. To bridge these gaps, we developed an explainable deep learning system called VinDr-CXR that can classify a CXR scan into multiple thoracic diseases and, at the same time, localize most types of critical findings on the image. VinDr-CXR was trained on 51,485 CXR scans with radiologist-provided bounding box annotations. It demonstrated a comparable performance to experienced radiologists in classifying 6 common thoracic diseases on a retrospective validation set of 3,000 CXR scans, with a mean area under the receiver operating characteristic curve (AUROC) of 0.967 (95% confidence interval [CI]: 0.958–0.975). The VinDr-CXR was also externally validated in independent patient cohorts and showed its robustness. For the localization task with 14 types of lesions, our free-response receiver operating characteristic (FROC) analysis showed that the VinDr-CXR achieved a sensitivity of 80.2% at the rate of 1.0 false-positive lesion identified per scan. A prospective study was also conducted to measure the clinical impact of the VinDr-CXR in assisting six experienced radiologists. The results indicated that the proposed system, when used as a diagnosis supporting tool, significantly improved the agreement between radiologists themselves with an increase of 1.5% in mean Fleiss' Kappa. We also observed that, after the radiologists consulted VinDr-CXR's suggestions, the agreement between each of them and the system was remarkably increased by 3.3% in mean Cohen's Kappa. Altogether, our results highlight the potentials of the proposed deep learning system as an effective assistant to radiologists in clinical practice. Part of the dataset used for developing the VinDr-CXR system has been made publicly available at <https://physionet.org/content/vindr-cxr/1.0.0/>.

INDEX TERMS Chest X-ray interpretation, deep learning, image classification, object detection.

I. INTRODUCTION

Common chest pathologies affect several hundred million people worldwide and kill several million cases every

The associate editor coordinating the review of this manuscript and approving it for publication was Cristian A. Linte.

year [1], [2]. They are the leading cause of death and impose an immense worldwide health burden. Diagnosis of thoracic diseases is a crucial clinical task for physicians. Currently, chest X-ray (CXR) radiography is the primary imaging modality used for screening, triaging, and diagnosing varieties of lung conditions [3] such as pneumothorax,

pneumonia, tuberculosis (TB), pleural effusion, atelectasis, emphysema, and cancers, etc. However, the CXR interpretation is a complicated task, which requires an in-depth understanding of radiologic signs in thoracic imaging [4], [5], [6], [7]. A previous study [8] reported that 22% of all errors in diagnostic radiology were made in the CXR interpretation. A recent work [9] showed that 19%–26% of lung cancers visible on CXR images were missed at the first reading. Furthermore, interpreting CXR scans usually is highly dependent on the observer and has a poor interagreement between physicians [10]. The interobserver agreement was considered poor to moderate depending on the type of findings [11]; this rate can be lower in local hospitals, leading to unfavorable consequences.

Advanced machine learning algorithms have recently shown their significant potential in medical image analysis [12], [13], [14], [15], [16], [17], [18], [19], [20], [21], [22], [23], [24], [25], [26], [27], [28] due to the availability of large-scale datasets [16], [29], [30], [31], [32], [33], [34], [35], [36]. In particular, previous studies [31], [37], [38], [39], [40], [41], [42], [43], [44], [45] have indicated that a deep learning (DL) system trained on a large-scale, annotated medical imaging dataset can reach a level of performance comparable to practicing radiologists in detecting common thoracic diseases [31], [37], [38], [39], [40], [41], analyzing retinal images [42], [43], or diagnosing skin cancers [44], [45].

However, the actual impact of DL systems in clinical practice remains unclear, and large-scale clinical evaluations of these such systems are limited. Hence, despite many promising results and increasing performances that have been published, very few DL algorithms have reached clinical implementation. We observe that multiple factors slow or impede artificial intelligence (AI) transfer into clinical practice. First, the development of an accurate and robust DL system requires a large number of annotated CXR scans from diverse sources. The creation of large-scale, high-quality datasets of annotated images is costly and challenging. Meanwhile, public datasets are limited, and their labels are unreliable since they are produced by automated rule-based labelers [29], [31], [46]. Previous evidences [39], [47], [48] showed that training DL systems on small datasets and low-quality annotations raises concerns about the robustness of those systems in real clinical contexts. Second, few clinical evaluations of DL-assisted diagnostic algorithms have been performed, and most of them are not even prospective and at high risk of bias [49]. Third, although DL systems can outperform physicians in specific clinical tasks [38], [50], the lack of explanatory power [51], [52] became a key obstacle to convince medical experts, who must understand how and why DL models have made a prediction. We believe that an accurate assessment of CXRs requires both detection of abnormal findings and a correct decision at the disease level. Hence, the provision of accurate and interpretable visualizations of lung abnormalities is a crucial step towards the clinical translation of DL systems.

To address these gaps, we introduce in this study a fully automated DL system, namely VinDr-CXR, for chest radiograph interpretation. The VinDr-CXR is designed to simultaneously classify six common lung diseases and localize 14 important findings from CXR scans. The development and evaluation of VinDr-CXR are based on large-scale medical imaging analysis and state-of-the-art DL algorithms. Specifically, we use a patient dataset from multiple hospitals in Vietnam, containing 51,485 manually annotated CXR studies, to train the DL system. To evaluate the performance of the proposed system, we compare the model's performance with that of human experts in a benchmark study using the consensus annotations provided by 5 experienced radiologists on a retrospective dataset as the reference standard. To validate the generalization capability of VinDr-CXR, we compute its performance on various external datasets. The results confirm that our framework is accurate and robust across multiple populations and settings. To demonstrate the clinical value of VinDr-CXR as an assistant to radiologists, we conduct a prospective study at two hospitals in Vietnam and investigate the inter-rater agreement on CXR interpretation with and without VinDr-CXR's assistance. We further calculate the change in the agreement rate between VinDr-CXR itself and each radiologist before and after he/she consults the system's suggestions. Experiments indicate that the proposed DL system provides meaningful supports for radiologists in detecting thoracic diseases in a real-world clinical setting.

II. MATERIALS AND METHODS

Our study was approved by the Institutional Review Boards (IRB) of the Hanoi Medical University Hospital (HMHU) and Hospital 108 (H108). In addition, the requirement for obtaining informed patient consent was waived due to the observational nature of this study. Under these approvals, raw CXR images in the Digital Imaging and Communications in Medicine (DICOM) format were collected retrospectively. Protected health information (PHI) has been de-identified to comply with the regulations of the U.S. HIPAA [53], European GDPR [54], and the local privacy laws [55]. In this section, we describe the methodology for developing and validating the VinDr-CXR system. First, we provide an overview of the proposed approach. Next, we describe details of datasets used in this study. Then, the development of DL algorithms is introduced. Last, we describe the experimental design for our reader study and statistical analysis methods used for model evaluation.

A. OVERVIEW OF APPROACH

We present in this section VinDr-CXR, a DL-based framework for the per-radiograph classification of common lung diseases and the detection of abnormal findings on CXRs. This framework includes two major components. First, an image-level classification network accepts a CXR scan as input and predicts whether it could be normal or abnormal. Second, a lesion-level detection network receives an

abnormal CXR scan as input from the classifier and provides the location of abnormal findings via bounding box predictions. An overview of the proposed approach is illustrated in Figure 1A. The core of the VinDr-CXR system is based on state-of-the-art DL networks for image classification and object detection tasks, named EfficientNet [56] and EfficientDet [57], respectively. The classification network is trained with image-level labels to distinguish six common lung diseases, including *Pneumonia*, *Tuberculosis*, *Lung Tumor*, *Pleural Effusion*, *Other Diseases*, and *No Finding*. Meanwhile, the detection network is trained with lesion-level annotations to localize 14 critical findings from the CXR images, i.e., *Cardiomegaly*, *Opacity*, *Consolidation*, *Atelectasis*, *Pneumothorax*, *Pleural Effusion*, *Aortic Enlargement*, *Interstitial Lung Disease (ILD)*, *Infiltration*, *Nodule/Mass*, *Pulmonary Fibrosis*, *Pleural Thickening*, *Calcification*, and *Other Lesions*. To develop DL algorithms, a total of 51,485 anteroposterior (AP) and posteroanterior (PA) CXR scans have been retrospectively collected from the HMUH and H108, which then manually annotated by expert radiologists. To evaluate the diagnostic accuracy of the VinDr-CXR, we compare its prediction with ground truth on an internal validation set of 3,000 studies, which was separate from the training set. Additionally, external datasets including CheXpert ($N = 200$) CheXphoto ($N = 200$) are used to verify the robustness of the VinDr-CXR for cross-site validation. Finally, we investigate the actual impact of the VinDr-CXR on clinical practice through a large reader study ($N = 400$). The inter-rater agreement among radiologists as well as the rate of agreement between VinDr-CXR and radiologists are then be assessed. Figure 1B shows an overview of the development and evaluation of the VinDr-CXR framework.

1) DATASETS FOR VinDr-CXR DEVELOPMENT AND VALIDATION

a: VinDr-CXR DATASET

To develop and validate the VinDr-CXR, we retrospectively collected a total of more than 100,000 anterior–posterior (AP) and posterior–anterior (PA) CXR scans of adult patients (aged > 10 years). The imaging data were in the DICOM format and performed at two major hospitals (i.e., the HMUH and H108) between January 1, 2018 and December 31, 2020. In addition, CXR studies were acquired from a wide diversity of scanners and manufacturers such as Phillips, GE, Fujifilm, Siemens, Toshiba, Canon, and Samsung. Out-of-distribution samples (e.g., images with poor quality or invalid) were excluded via a manual inspection process. A group of 17 expert radiologists has labeled a portion of the raw dataset, and each has at least ten years of experience. Overall, the dataset contains 54,485 CXR studies (mean age 43.77 years; 47.79% female patients) that meet the study criteria. A development set consists of 51,485 studies used to optimize DL algorithms and an internal validation set of 3,000 studies to evaluate models' performance. All CXRs were performed in independent patients for the validation

set and did not overlap with the development set to avoid bias. Each CXR scan in the development set was labeled by three independent radiologists, while a panel of five board-certified radiologists labeled each case in the internal validation set, and their consensus established the reference standard. Ground truth was established at the image-level for the classification task and pixel-level for the localization task. Our dataset was labeled for the presence of 28 labels. We defined for the first time two different types of labels in CXRs: (1) global labels that are image-level labels representing diseases or impressions and (2) local labels (lesion-level annotations) that are critical findings or lesions that occur in CXRs. The participating radiologists provide precisely the location of lesions or abnormalities via bounding box annotations for local labels. The data collection and annotation process is summarized in Figure 1B. The cohort demographic information and statistics of development and internal validation data sets are summarized in Table 1. Full details of the dataset collection and labeling process can be found in our previous study [33].

b: EXTERNAL DATASETS

We used two public CXR datasets to assess the accuracy and efficiency of the VinDr-CXR system across populations. Our external validation tests used data from patients in the CheXpert [31] ($N = 200$) and CheXphoto [58] ($N = 200$) datasets. These datasets shared several disease labels with the VinDr-CXR dataset, such as *Pleural Effusion*, *Pneumonia*, and *No Finding*. The CheXpert [31] is a large public dataset for the CXR interpretation performed between 2002 and 2017 at Stanford Hospital, USA. The CheXpert validation set contains 200 studies, for which the ground-truth label of each study is obtained by taking the majority vote of three board-certified radiologists. The CheXphoto [58] is a recently published CXR dataset for the automated interpretation of photos of CXR through cell phone photography. The CheXphoto validation set comprises natural photos of all 200 studies in the CheXpert validation set. It can be used as a resource for testing the robustness of DL algorithms on smartphone photos of CXRs. Additional details of the external datasets are provided in Irvin et al. [31] and Phillips et al. [58].

2) MODEL DEVELOPMENT AND TRAINING

To develop the VinDr-CXR system, we used a total of 51,485 annotated CXR scans from the development set. The system takes a CXR as input, and outputs are both disease classification and lesion localization. The whole VinDr-CXR architecture consists of two subnetworks, including a classification network and a detection network (see Figure 1A). We trained the EfficientNet-B6 [56] model on the CXR images with a size of 1024×1024 pixels to classify common chest diseases. The EfficientNet [56] was well-known as a state-of-the-art DL architecture for image recognition tasks. It can achieve a high level of accuracy while requiring less computational cost for model training. We used mean binary cross-entropy loss to optimize the network in a

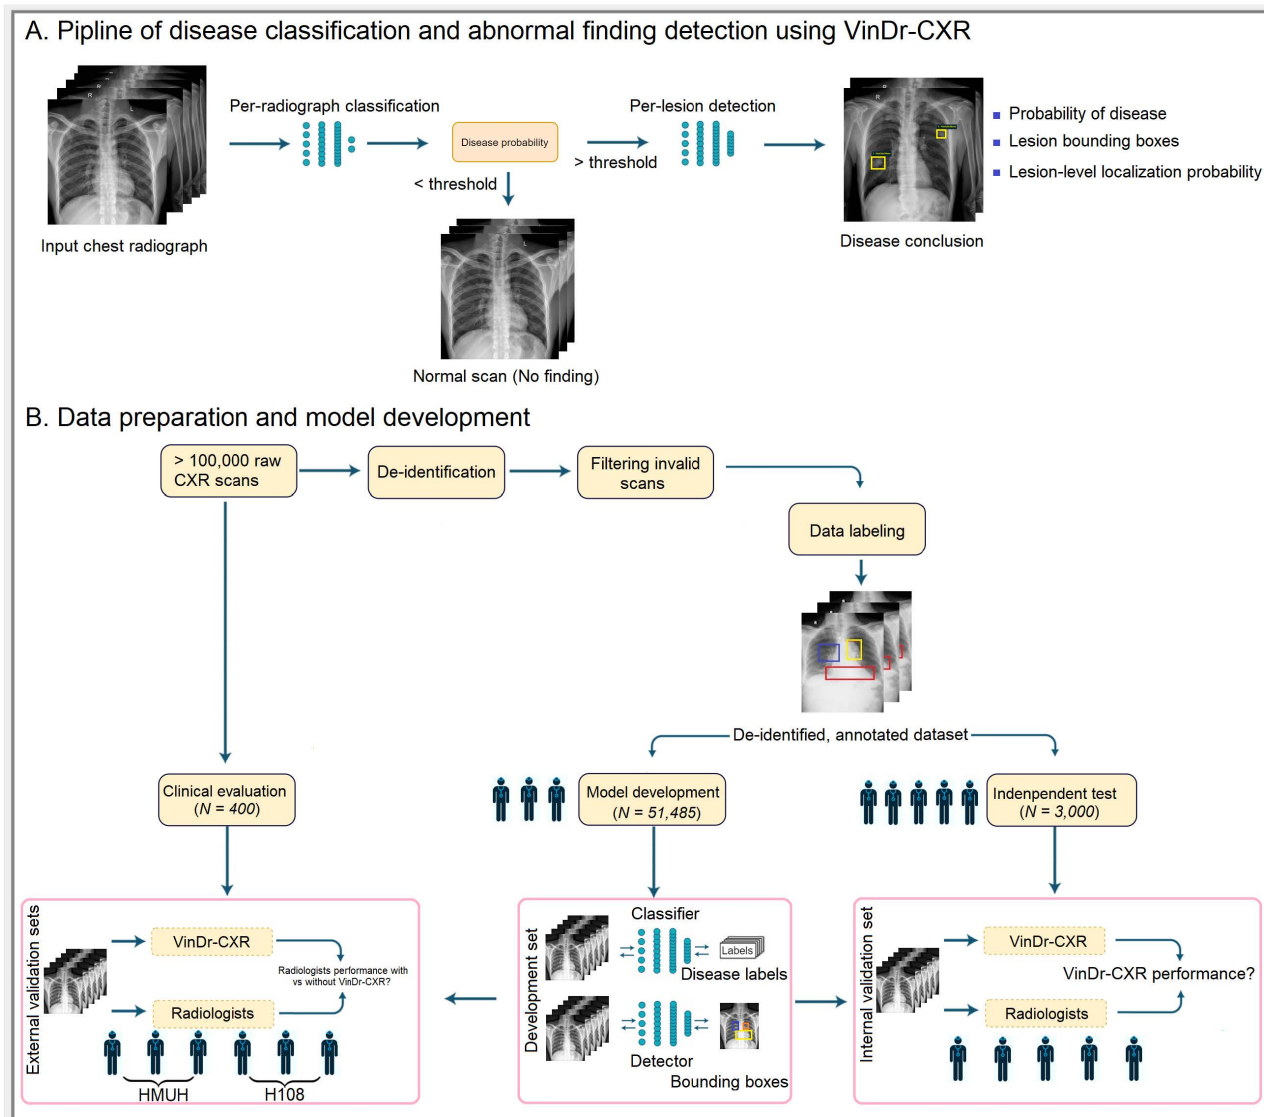


FIGURE 1. Overview of the proposed approach. (A) The pipeline of the VinDr-CXR framework for lung disease classification and localization using CXR images. A classifier takes as input one CXR scan and predicts its probability of abnormality. A detector takes an abnormal scan as input and provides bounding boxes with probabilities of abnormal findings at the lesion-level. The final prediction of the whole framework comes up with a complete description consisting of a disease conclusion, and abnormal findings. **(B)** Summary of the VinDr-CXR development and evaluation. The raw images in the DICOM standard were collected retrospectively from the picture archiving and communication system (PACS) of the HMUH and H108 hospitals. The raw images were then de-identified and all out-of-distribution samples were removed. An in-house web-based labeling tool called VinDr Lab was designed to annotate imaging data. In the labeling process, each CXR scan from the development set was labeled by a group of three experienced radiologists for the presence of 22 preliminary findings and six disease impressions. Meanwhile, CXR scans from the internal validation set were annotated by a consensus of five experienced radiologists. Finally, 51,485 CXR images were used to develop DL algorithms, and 3,000 studies were used for validation. External evaluations were also performed on CheXpert [31] (N = 200) and CheXphoto [58] (N = 200) datasets to evaluate the robustness of the VinDr-CXR across multiple hospitals. A reader study was designed and performed at the HMUH and H108 (N = 400) to investigate the impact of the VinDr-CXR in clinical practice.

supervised manner using image-level annotations. To localize abnormal findings, we deployed EfficientDet-D6 [57], a recent advance of DL-based detector for the object detection tasks. The per-lesion annotations provided by radiologists were used to optimize the EfficientDet-D6 [57] network. To reduce the impact of class imbalance, we adopted the focal loss [59] to optimize the detection network’s weights. Several data augmentation strategies have been applied to minimize the risk of over-fitting in both networks. Both the classification and detection networks were implemented

using Python 3.7 (<https://www.python.org/>), PyTorch 1.6 (<https://pytorch.org/>), and trained on an NVIDIA V100 32GB GPU. A detailed description of the model development and training is provided in the Appendices.

3) READER STUDY

To validate the effectiveness of the proposed DL approach, we conducted a reader study to assess the actual impact of the VinDr-CXR on the agreement of participating radiologists. We describe the reader study as follows.

TABLE 1. Characteristics of the datasets used for VinDr-CXR development and internal validation.

	Characteristics	Development set	Internal validation set
Collection statistics	Years	2018 to 2020	2018 to 2020
	Number of studies	51,485	3,000
	Number of images	51,941	3,000
	Number of abnormal studies	26,887	949
	Number of normal studies	24,598	2,051
	Number of annotators per scan	3	5
	Image size (pixel×pixel, mean)	2,729 × 2,395	2,748 × 2,394
	Age (years, mean)*	51.28	31.80
	Male (%)*	52.77	55.90
	Female (%)*	47.23	44.10
	Data size (GiB)	509.5	31.3
Local labels	1. Aortic Enlargement	9,183 (17.68%)	221 (07.37%)
	2. Atelectasis	1,346 (2.59%)	96 (3.20%)
	3. Cardiomegaly	8,359 (16.09%)	310 (10.33%)
	4. Calcification	3,154 (6.07%)	232 (7.73%)
	5. Clavicle Fracture	300 (0.58%)	2 (0.07%)
	6. Consolidation	2,125 (4.09%)	126 (4.20%)
	7. Edema	61 (0.12%)	0 (0%)
	8. Emphysema	1,007 (1.94%)	4 (0.13%)
	9. Enlarged PA	451 (0.87%)	9 (0.30%)
	10. Interstitial Lung Disease (ILD)	5,946 (11.45%)	316 (10.53%)
	11. Infiltration	3,254 (6.26%)	79 (2.63%)
	12. Lung Cavity	175 (0.34%)	10 (0.33%)
	13. Lung Cyst	129 (0.25%)	3 (0.10%)
	14. Lung Opacity	4,680 (9.01%)	95 (3.17%)
	15. Mediastinal Shift	706 (1.36%)	20 (0.67%)
	16. Nodule/Mass	5,521 (10.63%)	286 (9.53%)
	17. Pulmonary Fibrosis	6,919 (13.32%)	358 (11.93%)
	18. Pneumothorax	321 (0.62%)	25 (0.83%)
	19. Pleural Thickening	7,899 (15.21%)	240 (8.00%)
	20. Pleural Effusion	4,554 (8.76%)	136 (4.53%)
	21. Rib Fracture	1,241 (2.39%)	17 (0.57%)
	22. Other Lesions	3,851 (7.41%)	112 (3.73%)
Global labels	23. Lung Tumor	1,650 (3.18%)	80 (2.67%)
	24. Pneumonia	3,827 (7.37%)	246 (8.20%)
	25. Tuberculosis	2,130 (4.10%)	164 (5.47%)
	26. Other Diseases	18,848 (36.29%)	657 (21.90%)
	27. COPD	388 (0.75%)	2 (0.07%)
	28. No Finding	16,461 (31.7%)	2,051 (68.37%)

(*) The calculations were performed based on the number of CXR scans for which sex and age were known. For global labels, the number of positive examples was reported based on the majority vote of binarized labels provided by radiologists. For local labels, the percentage rate of number bounding boxes on the number of studies was reported. Several abnormalities (Clavicle Fracture, Edema, Emphysema, Enlarged PA, Lung Cavity, Lung Cyst, Mediastinal Shift, Rib Fracture, and COPD) were not considered during model training because the number of positive examples in the internal validation set was very limited. The “No Finding” label was intended to capture the absence of all findings and pathologies.

a: DATA COLLECTION

For clinical evaluation, 400 CXR examinations were collected retrospectively from the HMUH and H108 under IRB approvals. These examinations were acquired between March 2021 and June 2021 after the training process of the VinDr-CXR has been completed. Among 400 CXR studies, half ($N = 200$) was obtained from the HMUH, and the rest ($N = 200$) was from the H108. Data sampling was conducted based on the actual distributions at the hospitals. This ensures that the imaging data are representative of the real-world conditions in which the DL algorithms will be deployed. In addition, the CXR scans will be used to evaluate the agreement among participating readers. Hence, we did not establish a reference standard for the collected data.

b: READER SELECTION

We recruited a group of six board-certified radiologists from the radiology departments of the HMUH and H108 to participate in our observer performance test. All participating

radiologists were trained in CXR interpretation and had an average of 15.5 years of clinical experience interpreting thoracic diseases (range 10–22 years). In addition, the readers read an average of 25,000 CXR scans each year (range 15,000–40,000). Table 2 shows the characteristics of radiologists who participated in our reader study.

c: READER STUDY DESIGN

The reader study was conducted in two sessions. In the first session, participating readers read the CXR scans independently without the VinDr-CXR assistance. During the second session, the readers re-evaluated all CXR scans with the assistance of the VinDr-CXR. Specifically, the radiologists were provided the VinDr-CXR predictions in the form of bounding boxes, which locate abnormalities (see Figure 10; Appendices). They considered the model’s prediction and modified the diagnostics. During this process, the readers

TABLE 2. Characteristics of the participating radiologists. Mean annual diagnostic volumes were estimated based on the number of CXR scans interpreting.

Reader	Year's experience	Annual diagnostic volume (studies)
Radiologist 1 (H108)	10	40,000
Radiologist 2 (H108)	15	14,000
Radiologist 3 (H108)	11	45,000
Radiologist 1 (HMHU)	21	20,000
Radiologist 2 (HMHU)	22	15,000
Radiologist 3 (HMHU)	14	15,000
Average	15.5	25,000

were blinded to the relevant clinical information such as the original reports and previous medical histories of the patients or other patient records. To maximize human performance, the readers can perform the task on our browser-based viewer with zoom in or out, panning, and many other support tools. The reader study was set up to ensure that all radiologists can view and interpret the CXR studies in an environment similar to their routine workflow in clinical practice. Changes in radiologists' agreement were then assessed to investigate the impact of the VinDr-CXR assistance.

B. STATISTICAL ANALYSIS

Diagnostic performance metrics, including area under the receiver operating characteristic curve (AUC), sensitivity, specificity, *F1*-score, false-positive rate (FPR), and false-negative rate (FNR), were used to assess the accuracy of the VinDr-CXR for the classification task. For each indicator, 95% confidence interval (CI) was estimated with bootstrapping (10,000 replications). To evaluate the VinDr-CXR's ability to detect and localize lung lesions, we used the FROC (the sensitivities of models under different false positive rates as 0.25, 0.5, 1, 2, and 4). Cohen's Kappa statistics [60] and percentage agreement rate were used to evaluate the level of agreement between the VinDr-CXR system and participating radiologists, as well as to assess the agreement between pairs of radiologists. To assess inter-rater agreement among a group of radiologists, the Fleiss' Kappa [61] score was used. The Kappa values were interpreted as following guidelines [62]: (< 0.00): poor; (0.00–0.20): none to slight; (0.21–0.40): fair; (0.41–0.60): moderate; (0.61–0.80): substantial; and (0.81–1.00): almost perfect agreement. All statistical analyses were performed using Python (version 3.9.2 – <https://www.python.org/>) and scikit-learn (version 0.24.2 – <https://scikit-learn.org/>).

III. EXPERIMENTAL RESULTS

This section summarizes our main findings in this study. We first report the performance of the VinDr-CXR on the per-radiograph classification of common lung disease tasks on the internal and external test cohorts. We then provide experimental results for the lesion-level localization task. Finally, we show quantitative assessments on the impact of the VinDr-CXR in clinical practice.

A. EVALUATION OF VinDr-CXR PERFORMANCE

The following subsections detail the quantitative results of the VinDr-CXR system for the classification of diseases and

detection of critical lesions on internal and external validation datasets.

1) VinDr-CXR PROVIDES ACCURATE PER-RADIOGRAPH CLASSIFICATION OF COMMON LUNG DISEASE

The performance of the VinDr-CXR for the classification of common lung diseases was assessed on the internal validation set of 3,000 CXR studies, in which there were 948 patients with abnormal findings or diseases and 2052 patients without any pathologies (Table 1). The system achieved a mean AUROC of 0.967 (95% CI: 0.958, 0.975) over six global disease labels: 0.989 (0.983, 0.994) for Pleural Effusion, 0.978 (0.965, 0.988) for Lung Tumor, 0.969 (0.959, 0.978) for Pneumonia, 0.975 (0.964, 0.983) for Tuberculosis, and 0.920 (0.909, 0.931) for Other Diseases. The sensitivity, specificity, and *F1*-score of the VinDr-CXR were 0.933 (0.898, 0.964), 0.900 (0.887, 0.911), and 0.631 (0.589, 0.672), respectively. The system showed a FPR of 0.101 (0.089, 0.114) and a FNR of 0.067 (0.057, 0.102) over all target diseases. The overall accuracies of the system in differentiating between normal and abnormal CXRs were 0.972 (0.966, 0.978) in AUROC and 0.939 (0.931, 0.947) in *F1*-score. Our experimental results on CXR data from the internal validation set showed high sensitivity and specificity in classifying six disease labels. Detailed performances for individual diseases over all evaluation metrics on the internal validation cohort are reported in Table 3. Figure 2 shows six ROC curves of the system on the internal validation set for six global diseases.

2) VinDr-CXR SHOWS ROBUST CLASSIFICATION PERFORMANCE ON EXTERNAL TEST COHORTS

To investigate the consistency of the VinDr-CXR performance across multiple populations, we performed external validation tests using two independent datasets, including CheXpert [31] and CheXphoto [58]. The AUROC score, sensitivity, specificity, *F1*-score, FPR, and FNR of the VinDr-CXR with no additional training cost on 200 studies [normal: 27 cases, abnormal: 173 cases] of the CheXpert [31] validation set were 0.892 (0.839, 0.939), 0.866 (0.792, 0.932), 0.614 (0.556, 0.672), 0.453 (0.3650, 0.536), 0.386 (0.328, 0.444), and 0.134 (0.067, 0.208), respectively. On the CheXphoto [58] validation set [$N = 200$, normal: 27 cases, abnormal: 173 cases], the VinDr-CXR achieved an AUC of 0.888 (0.832, 0.936), a sensitivity of 0.871 (0.800, 0.936), a specificity of 0.614 (0.557, 0.670), a *F1*-score of 0.460 (0.372, 0.542), a FPR of 0.386 (0.330, 0.444), and a FNR of 0.129 (0.065, 0.200). The performance of the VinDr-CXR over all disease labels is reported in Table 4. Although a slight drop has been observed, the performance of the VinDr-CXR on external test sets remained at a high level, showing its robustness across different patient cohorts. These experimental results show evidence that training a DL system with a large-scale, high-quality dataset could reach a high diagnostic accuracy across populations without additional training cost.

TABLE 3. Per-radiograph classification performance of the VinDr-CXR on the internal validation set ($N = 3,000$). AUC = Area under the receiver operating characteristic curve, FPR = False-positive rate or false alarm rate, FNR = False-negative rate.

Label	AUC	Sensitivity	Specificity	F1-score	FPR	FNR
Pleural Effusion	0.989 (0.983, 0.994)	0.955 (0.912, 0.991)	0.934 (0.925, 0.943)	0.524 (0.463, 0.582)	0.066 (0.057, 0.075)	0.045 (0.009, 0.088)
Lung Tumor	0.978 (0.965, 0.988)	0.937 (0.877, 0.986)	0.937 (0.928, 0.946)	0.448 (0.381, 0.512)	0.063 (0.054, 0.072)	0.063 (0.014, 0.123)
Pneumonia	0.969 (0.959, 0.978)	0.959 (0.933, 0.982)	0.877 (0.864, 0.889)	0.576 (0.535, 0.616)	0.123 (0.111, 0.136)	0.041 (0.018, 0.067)
Tuberculosis	0.975 (0.964, 0.983)	0.903 (0.854, 0.945)	0.936 (0.927, 0.945)	0.602 (0.550, 0.651)	0.064 (0.055, 0.073)	0.097 (0.055, 0.146)
Other Diseases	0.920 (0.909, 0.931)	0.925 (0.904, 0.945)	0.796 (0.780, 0.812)	0.698 (0.674, 0.722)	0.204 (0.188, 0.220)	0.075 (0.055, 0.096)
No Finding	0.972 (0.966, 0.978)	0.920 (0.908, 0.932)	0.913 (0.895, 0.930)	0.939 (0.931, 0.947)	0.087 (0.070, 0.105)	0.080 (0.068, 0.092)
Mean	0.967 (0.958, 0.975)	0.933 (0.898, 0.964)	0.900 (0.887, 0.911)	0.631 (0.589, 0.672)	0.101 (0.089, 0.114)	0.067 (0.057, 0.102)

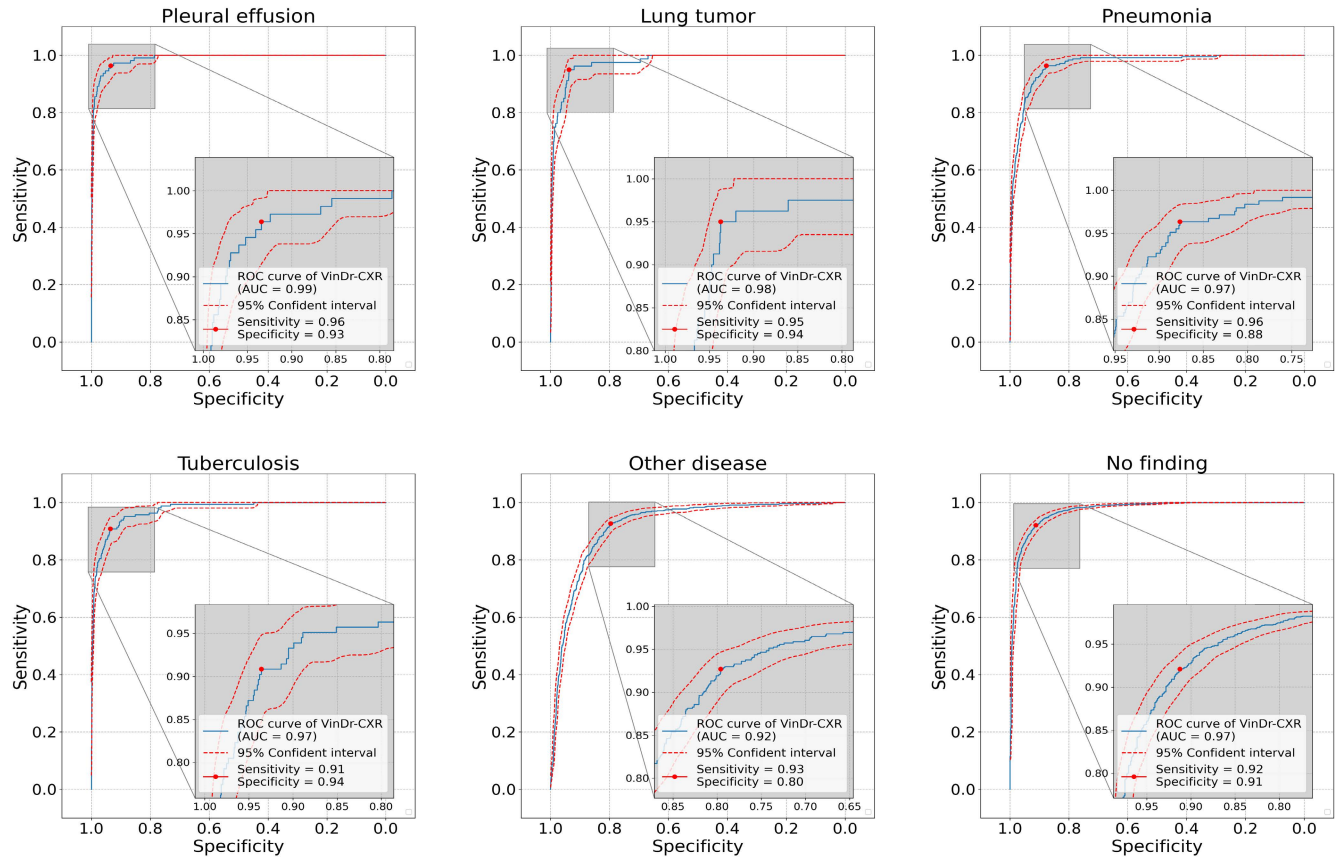


FIGURE 2. Receiver operating curves (ROC) of the VinDr-CXR system on the internal validation cohort. The solid blue lines show the ROC curves of the system and the dashed red lines show the 95 percentile intervals of the curves based on 10,000 bootstrap samples. We determined the optimal threshold for the VinDr-CXR system by maximizing Youden's index [63] for each disease label. We observe that the DL system achieved consistently high classification performances across all the target diseases (range AUROC: 0.920 – 0.989).

TABLE 4. Per-radiograph classification performance of the VinDr-CXR on Pleural Effusion, Pneumonia and No Finding from the external validation sets CheXpert [31] ($N = 200$) and CheXphoto [58] ($N = 200$) datasets.

Dataset&Label	AUC	Sensitivity	Specificity	F1-score	FPR	FNR
CheXpert [31]						
Pleural Effusion	0.895 (0.850, 0.934)	0.940 (0.877, 0.987)	0.727 (0.657, 0.795)	0.715 (0.637, 0.788)	0.273 (0.205, 0.343)	0.060 (0.013, 0.123)
Pneumonia	0.891 (0.824, 0.949)	1.000 (1.000, 1.000)	0.248 (0.193, 0.306)	0.085 (0.033, 0.143)	0.752 (0.694, 0.807)	0.000 (0.000, 0.000)
No Finding	0.891 (0.843, 0.933)	0.659 (0.500, 0.810)	0.868 (0.818, 0.914)	0.559 (0.424, 0.678)	0.132 (0.086, 0.182)	0.341 (0.190, 0.500)
Mean	0.892 (0.839, 0.939)	0.866 (0.792, 0.932)	0.614 (0.556, 0.672)	0.453 (0.365, 0.536)	0.386 (0.328, 0.444)	0.134 (0.067, 0.208)
CheXphoto [58]						
Pleural Effusion	0.889 (0.843, 0.930)	0.955 (0.899, 1.000)	0.739 (0.670, 0.804)	0.731 (0.654, 0.802)	0.261 (0.196, 0.331)	0.045 (0.000, 0.101)
Pneumonia	0.887 (0.816, 0.947)	1.000 (1.000, 1.000)	0.230 (0.176, 0.288)	0.083 (0.032, 0.140)	0.770 (0.712, 0.824)	0.000 (0.000, 0.000)
No Finding	0.887 (0.836, 0.930)	0.658 (0.500, 0.808)	0.873 (0.824, 0.919)	0.566 (0.430, 0.684)	0.127 (0.081, 0.176)	0.342 (0.194, 0.500)
Mean	0.888 (0.832, 0.936)	0.871 (0.800, 0.936)	0.614 (0.557, 0.670)	0.460 (0.372, 0.542)	0.386 (0.330, 0.444)	0.129 (0.065, 0.200)

Abbreviations: AUC = Area under the receiver operating characteristic curve, FPR = False-positive rate or false alarm rate, FNR = False-negative rate or miss detection rate.

3) VinDr-CXR PROVIDES ACCURATE LESION-LEVEL LOCALIZATION

For the per-lesion localization task, the VinDr-CXR's ability to detect and localize abnormal findings was evaluated on

3,000 CXR scans from the internal validation set using FROC analysis [64]. In this experiment, a detection is considered a true positive if the detected bounding box overlaps with the corresponding ground-truth bounding box more than 40%




FIGURE 3. Some accurate and erroneous predictions of the VinDr-CXR on several representative CXR images from the internal validation set. (A) The VinDr-CXR correctly identified four lesions, including Pleural Thickening, Calcification, Aortic Enlargement, and Cardiomegaly on the scan. (B) The system correctly identified a normal scan or true negative case. (C) The VinDr-CXR correctly identified a Nodule/Mass. However, it missed another Nodule/Mass on the left lung. (D) The VinDr-CXR correctly identified Nodule/Mass, but it seemed wrong to detect Calcification.

using the intersection over union (IoU) metric. Otherwise, it is considered to be a false positive. As shown in Table 5, the proposed system achieved a sensitivity of 80.2% (81.4, 84.9) at 1.0 false-positive marks per image. The FROC (average recall rate at the false positives as 0.25, 0.5, 1.00, 2.00, and 4.00) of the VinDr-CXR system was 78.36% (76.46, 80.16). FROC curves, which show the sensitivity of the system as a function of the number of false positives marks per image, of some representative findings are shown in Figure 4.

4) ANOMALY DETECTION VISUALIZATION

The per-lesion detection performance of the VinDr-CXR system can be inspected visually through Figure 3. We investigated the characteristics and errors of the VinDr-CXR detector by visualizing several representative cases containing both

correctly detected lesions and lesions that the DL system missed. In this experiment, we used the VinDr-CXR at a sensitivity of 0.802 and 1.0 false-positive marks per image to generate predictions. We found that the system was able to correctly identified almost all critical lesions. Meanwhile, most false-positive detections were small and non-dangerous lung lesions such as calcifications.

B. ASSESSMENT OF VinDr-CXR PERFORMANCE IN CLINICAL PRACTICE

1) INTER-AGREEMENT AMONG RADIOLOGISTS WITH AND WITHOUT VinDr-CXR ASSISTANCE

We conducted a multi-reader study at the HMUH and H108 to investigate the impact of the proposed DL system on

TABLE 5. Abnormality detection performance of the VinDr-CXR on the internal validation dataset.

Finding	Sensitivity@0.25	Sensitivity@0.5	Sensitivity@1.0	Sensitivity@2.0	Sensitivity@4.0
Cardiomegaly	0.965 (0.943, 0.984)	0.965 (0.943, 0.984)	0.968 (0.947, 0.986)	0.968 (0.947, 0.986)	0.968 (0.947, 0.986)
Opacity	0.617 (0.517, 0.718)	0.756 (0.660, 0.848)	0.842 (0.764, 0.915)	0.895 (0.827, 0.956)	0.906 (0.848, 0.959)
Consolidation	0.841 (0.764, 0.912)	0.898 (0.834, 0.954)	0.937 (0.884, 0.980)	0.937 (0.884, 0.980)	0.937 (0.884, 0.980)
Atelectasis	0.642 (0.538, 0.747)	0.698 (0.596, 0.796)	0.772 (0.678, 0.859)	0.772 (0.678, 0.859)	0.772 (0.678, 0.859)
Pneumothorax	0.639 (0.467, 0.800)	0.680 (0.522, 0.828)	0.680 (0.522, 0.828)	0.680 (0.522, 0.828)	0.680 (0.522, 0.828)
Pleural Effusion	0.898 (0.840, 0.950)	0.927 (0.877, 0.970)	0.934 (0.891, 0.972)	0.942 (0.902, 0.977)	0.942 (0.902, 0.977)
Aortic Enlargement	0.838 (0.789, 0.885)	0.882 (0.839, 0.923)	0.905 (0.864, 0.941)	0.905 (0.864, 0.942)	0.909 (0.870, 0.945)
ILD	0.664 (0.606, 0.723)	0.782 (0.730, 0.833)	0.858 (0.816, 0.899)	0.923 (0.890, 0.954)	0.937 (0.905, 0.964)
Infiltration	0.801 (0.714, 0.884)	0.861 (0.787, 0.930)	0.911 (0.843, 0.969)	0.911 (0.843, 0.969)	0.911 (0.843, 0.969)
Nodule/Mass	0.579 (0.510, 0.647)	0.663 (0.601, 0.727)	0.735 (0.670, 0.796)	0.769 (0.710, 0.827)	0.773 (0.714, 0.830)
Pulmonary Fibrosis	0.568 (0.514, 0.623)	0.627 (0.576, 0.678)	0.707 (0.657, 0.757)	0.775 (0.729, 0.819)	0.796 (0.751, 0.839)
Pleural Thickening	0.494 (0.425, 0.564)	0.608 (0.540, 0.675)	0.714 (0.651, 0.776)	0.797 (0.739, 0.852)	0.850 (0.797, 0.900)
Calcification	0.598 (0.527, 0.670)	0.685 (0.613, 0.755)	0.774 (0.713, 0.833)	0.802 (0.744, 0.858)	0.802 (0.744, 0.858)
Other Lesions	0.265 (0.197, 0.338)	0.311 (0.238, 0.388)	0.372 (0.299, 0.448)	0.484 (0.408, 0.562)	0.551 (0.474, 0.627)
No Finding	0.921 (0.909, 0.933)	0.921 (0.909, 0.933)	0.921 (0.909, 0.933)	0.921 (0.909, 0.933)	0.921 (0.909, 0.933)
Mean	0.689 (0.668, 0.710)	0.751 (0.731, 0.770)	0.802 (0.784, 0.819)	0.832 (0.814, 0.849)	0.844 (0.826, 0.860)

Model performance was evaluated using the FROC score at different the number of false-positive predictions per image (0.25, 0.5, 1.0, 2.0 and 4.0). Data in parentheses are 95% confidence intervals.

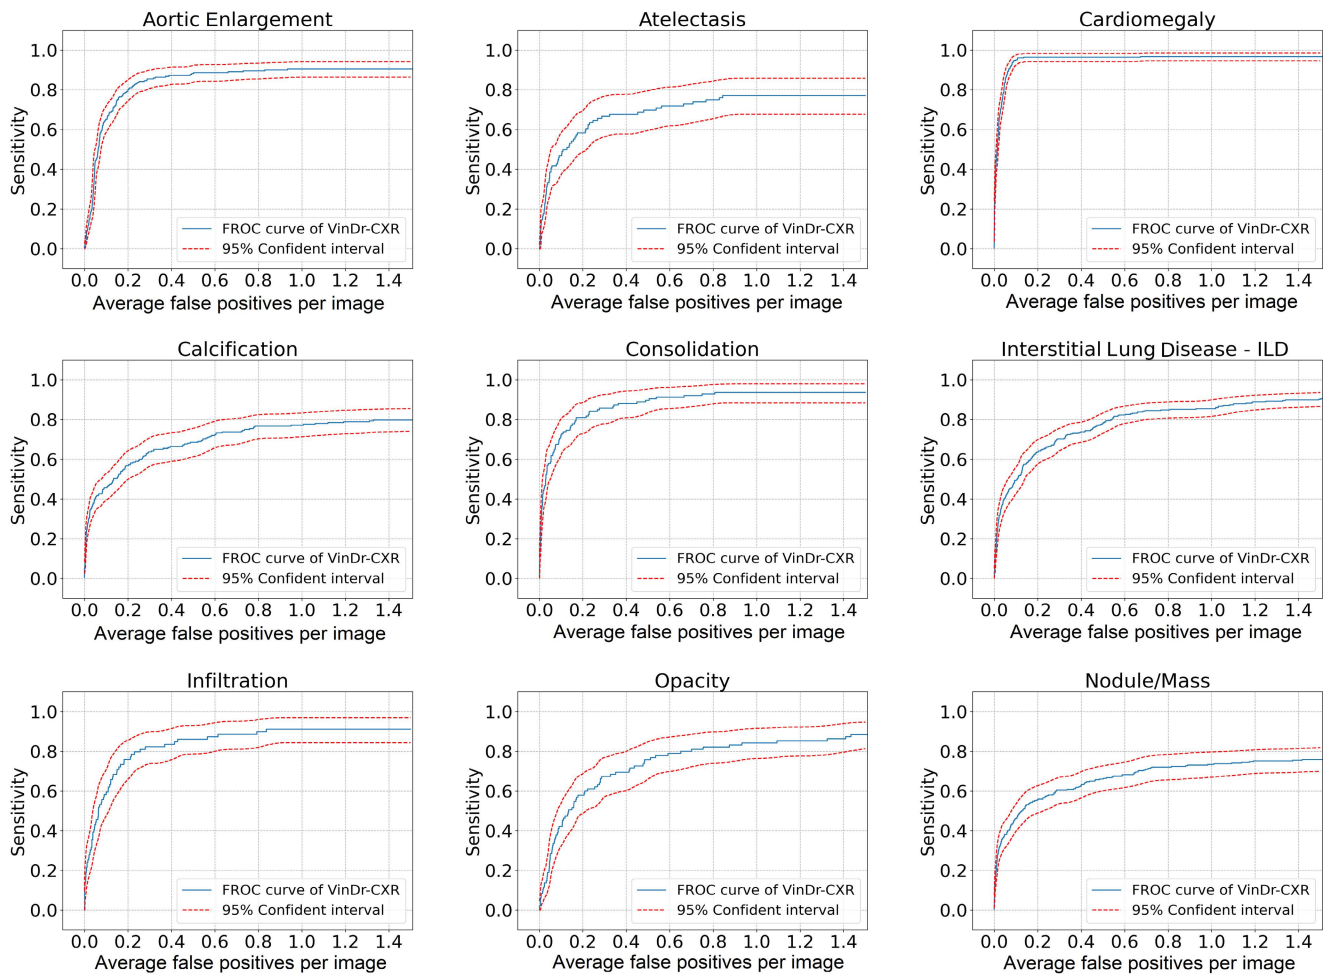


FIGURE 4. Free-response receiver operating characteristic (FROC) of the VinDr-CXR detector for several findings on the internal validation set. Sensitivity is calculated on a per lesion basis. The blue curve shows the results using all validation images, while the red dashed lines show the 95 percentile intervals from 10,000 bootstrap samples.

radiologists’ performance in the regular clinical workflow. At each hospital, three experienced radiologists assessed a total of 200 CXR images without VinDr-CXR assistance at the first read. The images were read with the assistance of VinDr-CXR in the second read. The agreement between

the VinDr-CXR and radiologists, as well as inter-observer agreement among radiologists were assessed using the percentage agreement rate and Kappa statistics. The details of the proposed reader study design are provided in the Methods section. In the first read, concordance among three

radiologists on data from the H108 showed percentage agreement rates between 89.9%–90.7% and Cohen's Kappa values between 0.501 (0.349, 0.654)–0.540 (0.392, 0.688). Agreement among three H108's radiologists was moderate with a Fleiss' Kappa of 0.529 (0.453, 0.605). In the second read, with the support of VinDr-CXR system, agreement among three raters showed a slightly higher percentage agreement rate. Specifically, percentage agreement rates ranged between 90.2%–90.7% and Cohen's Kappa values between 0.519 (0.351, 0.678)–0.556 (0.412, 0.700) that indicated moderate inter-individual agreement. Agreement among three H108's radiologists was moderate with a Fleiss' Kappa of 0.545 (0.465, 0.625), corresponding to an 3.0% improvement in Fleiss' Kappa compared to the first read.

We observed the same results on the data from the H108. Details of clinical evaluation results at the H108 are provided in Table 10 and Table 11 (Appendices). In the first read, the percentage agreement rates between each pair of radiologists ranged from 90.2%–90.8%; Cohen's Kappa values were between 0.367 (0.192, 0.543)–0.483 (0.304, 0.662), and a Fleiss' Kappa of 0.404 (0.329–0.480). In the second read, percentage agreement rates were between 90.2%–90.8%, and Cohen's Kappa values were between 0.367 (0.192, 0.543)–0.483 (0.304, 0.662). The agreement between three H108's radiologists was a Fleiss' Kappa of 0.418 (0.342, 0.494), corresponding to an improvement of 3.4% compared to the first read.

2) DIAGNOSTIC AGREEMENT BETWEEN VinDr-CXR AND RADIOLOGISTS WITH AND WITHOUT ASSISTANCE

In this experiment, the rate of agreement between the VinDr-CXR and radiologists in detecting abnormal lung findings from CXRs were assessed. Table 8 and Table 9 show the agreement rate between the VinDr-CXR and H108's radiologists without and with assistance, respectively. In the first read, the percentage agreement rates were 88.8%–90.2%, and the Cohen's Kappa values ranged from 0.462 (0.283, 0.640)–0.506 (0.327, 0.685). With the assistance of VinDr-CXR, the rates of agreement ranged from 90.5%–91.1% in percentage agreement and from 0.524 (0.348, 0.699)–0.546 (0.370, 0.717) in Cohen's Kappa values. At the H108, the percentage agreement rates were 88.8%–90.2%, and the Cohen's Kappa values ranged from 0.462 (0.283, 0.640)–0.506 (0.327, 0.685) without the VinDr-CXR assistance. Meanwhile, the rates of agreement ranged from 90.5%–91.1% in percentage agreement and from 0.524 (0.348, 0.699)–0.546 (0.370, 0.717) in Kappa values with the VinDr-CXR assistance (Detailed in Table 12 and Table 13 in the Appendices). After consulting the VinDr-CXR output, significant improvements in Kappa scores have been observed across two hospitals. These results indicated that the VinDr-CXR assistance resulted in a significant increase in the agreement between the DL system and radiologists.

3) IMPACT OF VinDr-CXR ON RADIOLOGIST DIAGNOSTIC AGREEMENT

We show evidence that the VinDr-CXR helped improve the degree of agreement among radiologists for the task of detection of lung lesions (see Figure 5A). Specifically, Fleiss' Kappa values improved by 1.4% and 1.6% for H108 and H108 readers, respectively. The inter-rater reliability has improved by 1.0%–2.9%, except for Reader 1 and Reader 2 from the H108 ($\Delta = -0.2\%$). Additionally, we found that the rate of VinDr-CXR agreement with the participating radiologists was slightly higher than the rate of agreement among radiologists. With the support of VinDr-CXR, the agreement between the proposed DL system and radiologists significantly improved. As shown in Figure 5B, increments of agreement degree ranged 1.8%–6.5% after consulting the VinDr-CXR predictions. Figure 9 in the Appendices shows several representative images from our reader study, which indicates the change in diagnostic decision after viewing the VinDr-CXR recommendation.

C. VISUAL INTERPRETABILITY OF VinDr-CXR

Explainability is an absolute necessity for the broad deployment of AI models in clinical practice. In a recent publication, the United States Food and Drug Administration (FDA) mentioned that explainability is required [65] for any AI-based computer-aided diagnosis (CAD) system. In summary, an interpretable AI/DL system can show the links between the features extracted by the system and its predictions [66]. Particularly, those links can be understood by a human expert. Explainable DL systems help human experts understand the underlying reasoning of DL systems and identify individual cases in which an AI model potentially gives incorrect predictions. In this study, the proposed VinDr-CXR is not only able to provide disease conclusions (global labels), but a helpful explanation involves abnormal findings (local labels) with their corresponding exact locations. Beyond the classification output, the VinDr-CXR can provide localization information that locates abnormalities accurately on CXR scans. The bounding boxes provided by the DL system may be an essential consideration supporting classification outputs. To illustrate the interpretability of the system, we utilized the trained classification network to compute and visualize the saliency maps for several examples from the internal validation set. To this end, we extracted feature maps produced by the last convolutional layer of the VinDr-CXR classifier model. We then used principal component analysis (PCA) [67] to reduce the channels of the feature map into a single channel. Then, this single-channel map is converted to a saliency map for visualization. Figure 6 shows the original CXRs and lesion bounding boxes annotated by our expert radiologists for Lung Tumor and Tuberculosis (TB) patients. The corresponding saliency maps obtained from the VinDr-CXR system are also provided. We observed the following insights that help better

TABLE 6. Inter-rater agreement among radiologists at the H108 without the VinDr-CXR assistance.

Findings	Rater 1 vs. Rater 2		Rater 1 vs. Rater 3		Rater 2 vs. Rater 3		Fleiss' κ
	Agree	Cohen's κ	Agree	Cohen's κ	Agree	Cohen's κ	
Lung Tumor	0.920	0.562 (0.375, 0.748)	0.955	0.643 (0.426, 0.859)	0.905	0.490 (0.296, 0.683)	0.552 (0.472, 0.632)
Pneumonia	0.895	0.698 (0.580, 0.817)	0.875	0.543 (0.387, 0.700)	0.880	0.626 (0.496, 0.756)	0.626 (0.546, 0.706)
Tuberculosis	0.945	0.674 (0.494, 0.853)	0.925	0.596 (0.413, 0.778)	0.950	0.763 (0.623, 0.904)	0.681 (0.601, 0.761)
Other Diseases	0.830	0.651 (0.544, 0.758)	0.810	0.595 (0.483, 0.707)	0.820	0.617 (0.507, 0.727)	0.619 (0.539, 0.699)
No Finding	0.945	0.886 (0.821, 0.952)	0.940	0.878 (0.812, 0.945)	0.945	0.889 (0.825, 0.952)	0.884 (0.804, 0.964)
Aortic Enlargement	0.915	0.662 (0.514, 0.811)	0.895	0.546 (0.377, 0.716)	0.930	0.669 (0.509, 0.830)	0.625 (0.545, 0.705)
Atelectasis	0.925	-0.009 (-0.027, 0.008)	0.980	-0.008 (-0.019, 0.004)	0.935	0.216 (-0.046, 0.478)	0.084 (0.004, 0.164)
Calcification	0.905	0.654 (0.512, 0.795)	0.870	0.255 (0.058, 0.452)	0.835	0.288 (0.122, 0.453)	0.419 (0.339, 0.499)
Cardiomegaly	0.915	0.574 (0.393, 0.756)	0.965	0.810 (0.673, 0.946)	0.950	0.734 (0.578, 0.891)	0.703 (0.623, 0.783)
Consolidation	0.940	0.308 (0.026, 0.591)	0.940	0.506 (0.264, 0.749)	0.940	0.308 (0.026, 0.591)	0.388 (0.308, 0.468)
ILD	0.845	0.552 (0.412, 0.692)	0.820	0.348 (0.189, 0.507)	0.775	0.196 (0.044, 0.349)	0.372 (0.292, 0.452)
Infiltration	0.815	0.288 (0.139, 0.436)	0.860	0.239 (0.052, 0.426)	0.865	0.562 (0.419, 0.706)	0.376 (0.296, 0.456)
Lung Opacity	0.830	0.167 (-0.014, 0.347)	0.870	0.054 (-0.067, 0.175)	0.890	-0.019 (-0.042, 0.005)	0.072 (-0.008, 0.152)
Nodule/Mass	0.895	0.512 (0.331, 0.694)	0.910	0.521 (0.328, 0.715)	0.915	0.591 (0.417, 0.766)	0.541 (0.461, 0.621)
Pleural Effusion	0.930	0.793 (0.690, 0.896)	0.905	0.682 (0.549, 0.815)	0.935	0.800 (0.697, 0.903)	0.760 (0.680, 0.840)
Pleural Thickening	0.835	0.284 (0.116, 0.451)	0.890	0.255 (0.037, 0.473)	0.865	0.458 (0.291, 0.624)	0.337 (0.257, 0.417)
Pneumothorax	1.000	1.000 (1.000, 1.000)	1.000	1.000 (1.000, 1.000)	1.000	1.000 (1.000, 1.000)	1.000 (1.000, 1.000)
Pulmonary Fibrosis	0.870	0.505 (0.340, 0.669)	0.890	0.514 (0.337, 0.691)	0.900	0.599 (0.441, 0.756)	0.539 (0.458, 0.619)
Other Lesions	0.920	0.386 (0.145, 0.627)	0.935	0.546 (0.327, 0.765)	0.925	0.476 (0.250, 0.702)	0.471 (0.391, 0.551)
Mean	0.899	0.534 (0.389, 0.679)	0.907	0.501 (0.349, 0.654)	0.903	0.540 (0.392, 0.688)	0.529 (0.453, 0.605)

TABLE 7. Inter-rater agreement among the H108's radiologists with the VinDr-CXR assistance.

Findings	Rater 1 vs. Rater 2		Rater 1 vs. Rater 3		Rater 2 vs. Rater 3		Fleiss' κ
	Agree	Cohen's κ	Agree	Cohen's κ	Agree	Cohen's κ	
Lung Tumor	0.915	0.525 (0.333, 0.717)	0.960	0.671 (0.458, 0.884)	0.895	0.436 (0.239, 0.632)	0.524 (0.444, 0.604)
Pneumonia	0.920	0.778 (0.675, 0.881)	0.905	0.690 (0.561, 0.818)	0.885	0.655 (0.529, 0.781)	0.709 (0.629, 0.789)
Tuberculosis	0.950	0.710 (0.541, 0.879)	0.930	0.646 (0.477, 0.814)	0.950	0.772 (0.637, 0.907)	0.711 (0.631, 0.791)
Other Diseases	0.815	0.621 (0.511, 0.731)	0.810	0.594 (0.480, 0.707)	0.825	0.633 (0.526, 0.740)	0.614 (0.534, 0.694)
No finding	0.950	0.897 (0.834, 0.959)	0.945	0.889 (0.825, 0.952)	0.945	0.889 (0.825, 0.952)	0.891 (0.811, 0.971)
Aortic Enlargement	0.920	0.687 (0.543, 0.830)	0.895	0.559 (0.393, 0.726)	0.935	0.708 (0.559, 0.857)	0.650 (0.570, 0.730)
Atelectasis	0.935	0.125 (-0.098, 0.348)	0.975	-0.008 (-0.021, 0.005)	0.940	0.312 (0.033, 0.590)	0.185 (0.105, 0.265)
Calcification	0.930	0.764 (0.646, 0.881)	0.835	0.192 (0.018, 0.365)	0.825	0.268 (0.105, 0.430)	0.438 (0.358, 0.518)
Cardiomegaly	0.925	0.639 (0.471, 0.807)	0.970	0.833 (0.703, 0.963)	0.955	0.775 (0.635, 0.916)	0.745 (0.665, 0.825)
Consolidation	0.930	0.197 (-0.056, 0.451)	0.935	0.484 (0.243, 0.724)	0.935	0.211 (-0.055, 0.477)	0.320 (0.240, 0.400)
ILD	0.845	0.552 (0.412, 0.692)	0.820	0.373 (0.214, 0.533)	0.795	0.296 (0.137, 0.454)	0.411 (0.331, 0.491)
Infiltration	0.810	0.276 (0.127, 0.425)	0.835	0.229 (0.059, 0.399)	0.885	0.646 (0.514, 0.777)	0.402 (0.322, 0.482)
Lung Opacity	0.830	0.167 (-0.014, 0.347)	0.885	0.262 (0.063, 0.462)	0.865	-0.055 (-0.086, -0.023)	0.131 (0.051, 0.211)
Nodule/Mass	0.895	0.557 (0.389, 0.726)	0.905	0.573 (0.401, 0.745)	0.920	0.647 (0.487, 0.806)	0.592 (0.512, 0.672)
Pleural Effusion	0.940	0.825 (0.730, 0.920)	0.925	0.759 (0.643, 0.875)	0.945	0.834 (0.739, 0.928)	0.807 (0.727, 0.887)
Pleural Thickening	0.840	0.318 (0.151, 0.485)	0.885	0.197 (-0.015, 0.410)	0.855	0.420 (0.253, 0.587)	0.321 (0.241, 0.401)
Pneumothorax	0.995	0.798 (0.410, 1.000)	0.995	0.798 (0.410, 1.000)	1.000	1.000 (1.000, 1.000)	0.855 (0.775, 0.935)
Pulmonary Fibrosis	0.865	0.518 (0.360, 0.675)	0.875	0.530 (0.368, 0.691)	0.900	0.661 (0.525, 0.798)	0.571 (0.491, 0.651)
Other Lesions	0.920	0.386 (0.145, 0.627)	0.940	0.593 (0.385, 0.801)	0.920	0.457 (0.233, 0.681)	0.482 (0.402, 0.562)
Mean	0.902	0.544 (0.374, 0.704)	0.907	0.519 (0.351, 0.678)	0.904	0.556 (0.412, 0.700)	0.545 (.465, .625)

TABLE 8. Agreement between the VinDr-CXR and H108's radiologists without assistance.

Findings	Rater 1 vs. AI		Rater 2 vs. AI		Rater 3 vs. AI	
	Agree	Cohen's κ	Agree	Cohen's κ	Agree	Cohen's κ
Lung Tumor	0.945	0.493 (0.235, 0.751)	0.885	0.329 (0.129, 0.530)	0.940	0.469 (0.215, 0.723)
Pneumonia	0.900	0.688 (0.561, 0.815)	0.905	0.735 (0.623, 0.846)	0.875	0.566 (0.418, 0.715)
Tuberculosis	0.965	0.769 (0.604, 0.934)	0.950	0.723 (0.560, 0.885)	0.940	0.695 (0.534, 0.856)
Other Diseases	0.725	0.399 (0.277, 0.521)	0.705	0.355 (0.231, 0.479)	0.795	0.487 (0.353, 0.620)
No Finding	0.905	0.810 (0.730, 0.890)	0.900	0.800 (0.718, 0.882)	0.935	0.870 (0.802, 0.938)
Aortic Enlargement	0.880	0.359 (0.174, 0.544)	0.895	0.361 (0.159, 0.563)	0.915	0.378 (0.150, 0.605)
Atelectasis	0.910	0.091 (-0.076, 0.259)	0.895	0.308 (0.087, 0.529)	0.900	0.067 (-0.093, 0.226)
Calcification	0.875	0.326 (0.127, 0.525)	0.860	0.418 (0.252, 0.585)	0.875	0.008 (-0.136, 0.152)
Cardiomegaly	0.920	0.607 (0.432, 0.782)	0.905	0.542 (0.360, 0.724)	0.935	0.662 (0.491, 0.832)
Consolidation	0.920	0.458 (0.235, 0.681)	0.910	0.219 (-0.006, 0.444)	0.930	0.526 (0.309, 0.743)
ILD	0.830	0.488 (0.340, 0.636)	0.845	0.537 (0.394, 0.680)	0.840	0.385 (0.219, 0.550)
Infiltration	0.830	0.226 (0.062, 0.391)	0.905	0.712 (0.592, 0.833)	0.870	0.530 (0.372, 0.687)
Lung Opacity	0.860	0.338 (0.147, 0.528)	0.860	0.255 (0.058, 0.453)	0.880	-0.019 (-0.043, 0.005)
Nodule/Mass	0.860	0.383 (0.200, 0.566)	0.875	0.489 (0.316, 0.661)	0.870	0.409 (0.225, 0.593)
Pleural Effusion	0.975	0.921 (0.853, 0.989)	0.945	0.839 (0.747, 0.931)	0.920	0.735 (0.613, 0.857)
Pleural Thickening	0.905	0.337 (0.109, 0.565)	0.850	0.390 (0.220, 0.559)	0.915	0.494 (0.286, 0.701)
Pneumothorax	0.985	0.565 (0.125, 1.000)	0.985	0.565 (0.125, 1.000)	0.985	0.565 (0.125, 1.000)
Pulmonary Fibrosis	0.900	0.558 (0.386, 0.730)	0.850	0.398 (0.224, 0.572)	0.870	0.384 (0.194, 0.575)
Other Lesions	0.945	0.563 (0.332, 0.794)	0.955	0.643 (0.426, 0.859)	0.940	0.568 (0.350, 0.787)
Mean	0.897	0.494 (0.308, 0.679)	0.888	0.506 (0.327, 0.685)	0.902	0.462 (0.283, 0.640)

understand the decision-making process of the VinDr-CXR system. First, the visualization of normal cases was irregular

with a symmetric high colormap, and there was no increased signal over all parts of the lung (Figure 6A). Second, across

TABLE 9. Agreement between the VinDr-CXR and H108’s radiologists with assistance.

Findings	Rater 1 vs. AI		Rater 2 vs. AI		Rater 3 vs. AI	
	Agree	Cohen’s κ	Agree	Cohen’s κ	Agree	Cohen’s κ
Lung Tumor	0.950	0.519 (0.259, 0.780)	0.885	0.329 (0.129, 0.530)	0.950	0.558 (0.313, 0.802)
Pneumonia	0.925	0.776 (0.668, 0.884)	0.905	0.735 (0.623, 0.846)	0.910	0.703 (0.575, 0.830)
Tuberculosis	0.970	0.807 (0.657, 0.957)	0.950	0.723 (0.560, 0.885)	0.940	0.707 (0.553, 0.862)
Other Diseases	0.735	0.414 (0.291, 0.537)	0.700	0.356 (0.234, 0.477)	0.815	0.543 (0.415, 0.670)
No finding	0.910	0.820 (0.742, 0.898)	0.900	0.800 (0.718, 0.882)	0.935	0.870 (0.802, 0.938)
Aortic Enlargement	0.880	0.359 (0.174, 0.544)	0.900	0.408 (0.209, 0.606)	0.915	0.417 (0.199, 0.635)
Atelectasis	0.910	0.091 (-0.076, 0.259)	0.905	0.374 (0.151, 0.597)	0.905	0.146 (-0.057, 0.349)
Calcification	0.905	0.550 (0.376, 0.723)	0.865	0.449 (0.286, 0.612)	0.870	0.002 (-0.138, 0.142)
Cardiomegaly	0.925	0.625 (0.451, 0.799)	0.920	0.634 (0.470, 0.799)	0.935	0.662 (0.491, 0.832)
Consolidation	0.925	0.506 (0.289, 0.722)	0.905	0.146 (-0.057, 0.349)	0.940	0.594 (0.386, 0.801)
ILD	0.840	0.518 (0.372, 0.664)	0.845	0.537 (0.394, 0.680)	0.850	0.449 (0.286, 0.611)
Infiltration	0.845	0.304 (0.134, 0.475)	0.905	0.712 (0.592, 0.833)	0.890	0.627 (0.485, 0.769)
Lung Opacity	0.860	0.338 (0.147, 0.528)	0.860	0.255 (0.058, 0.453)	0.905	0.366 (0.125, 1.000)
Nodule/Mass	0.875	0.489 (0.316, 0.661)	0.880	0.516 (0.347, 0.685)	0.880	0.487 (0.311, 0.663)
Pleural Effusion	0.985	0.954 (0.901, 1.000)	0.945	0.839 (0.747, 0.931)	0.930	0.773 (0.659, 0.886)
Pleural Thickening	0.915	0.407 (0.178, 0.636)	0.855	0.420 (0.253, 0.587)	0.910	0.451 (0.237, 0.664)
Pneumothorax	0.990	0.745 (0.405, 1.000)	0.985	0.565 (0.125, 1.000)	0.985	0.662 (0.491, 0.832)
Pulmonary Fibrosis	0.905	0.587 (0.420, 0.754)	0.870	0.508 (0.347, 0.670)	0.900	0.599 (0.441, 0.756)
Other Lesions	0.945	0.563 (0.332, 0.794)	0.955	0.643 (0.426, 0.859)	0.935	0.546 (0.328, 0.764)
Mean	0.905	0.546 (0.370, 0.717)	0.891	0.524 (0.348, 0.699)	0.911	0.527 (0.342, 0.711)

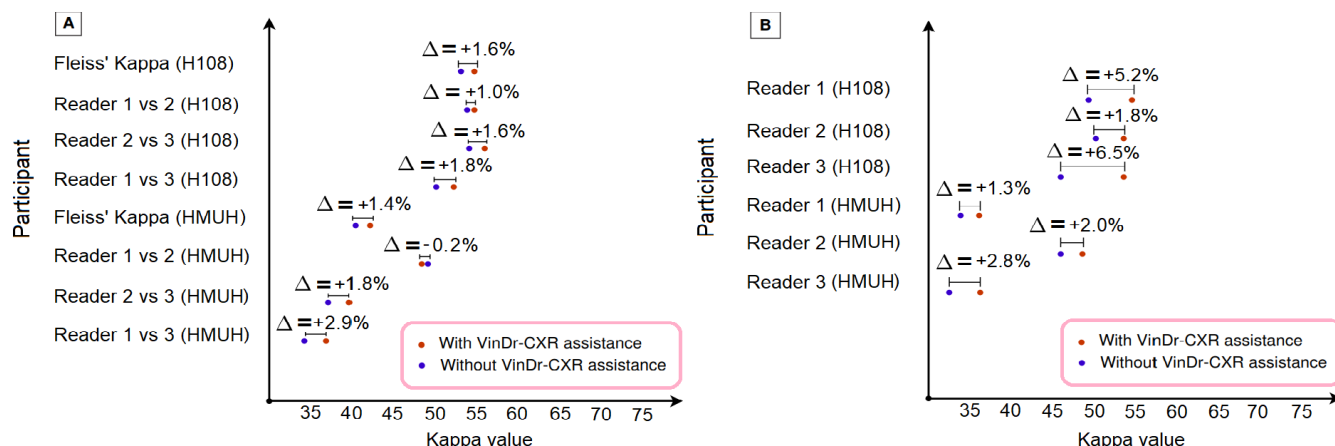


FIGURE 5. Impact of the VinDr-CXR on radiologists’ performance. (A) Change in inter-radiologist agreement before and after consulting the VinDr-CXR predictions. The VinDr-CXR assistance significantly improved agreement between radiologists with an increase of 1.5% in mean Fleiss’ Kappa. (B) Change in individual radiologists’ agreement with the system before and after consulting the VinDr-CXR predictions. The VinDr-CXR assistance resulted in a significant increase in the agreement between the AI system and radiologists with an improvement of 3.3% in mean Cohen’s Kappa. In particular, all differences were statistically significant.

all abnormal scans, the saliency maps highlighted parts of the CXRs that contain abnormal patterns such as Nodule/Mass, Calcification, and Opacity, which are clinically correlated with disease conclusions, including Lung Tumor and TB. In other words, the attention regions in the visualization maps were consistent with the annotated abnormal findings provided by our radiologists, as well as predictions by the VinDr-CXR detector model (Figure 6B–F).

IV. DISCUSSION

The VinDr-CXR may find helpful in several clinical scenarios. (1) The system was able to discriminate correctly between normal and abnormal CXRs. It, therefore, could be used as a tool to automate screening of common lung diseases for primary diagnosis (e.g., Tuberculosis and Malignancies) at scale. (2) Non-specialist clinicians could also potentially use the VinDr-CXR at healthcare centers to

provide teleradiology primary CXR reading or to support rapidly triage cases. (3) The VinDr can be used as a second-reader or assistance tool for clinicians. In this setting, the system can provide diagnosis assistance at different levels including image-level diagnosis and lesion-level diagnosis. This makes the proposed DL algorithms more transparent and more explainable, allowing clinicians to understand and explore the results to improve diagnosis outcomes.

On the technical aspect, this study strongly supports the following statements. First, a DL network can learn effectively and accurately predict common thoracic diseases and findings if trained on large-scale, multi-institutional, and expert-annotated imaging datasets. Second, although the current study was conducted in Vietnam, the proposed VinDr-CXR showed its robustness on external datasets acquired from hospitals in the United States. This result showed evidence that the model generalization of a DL

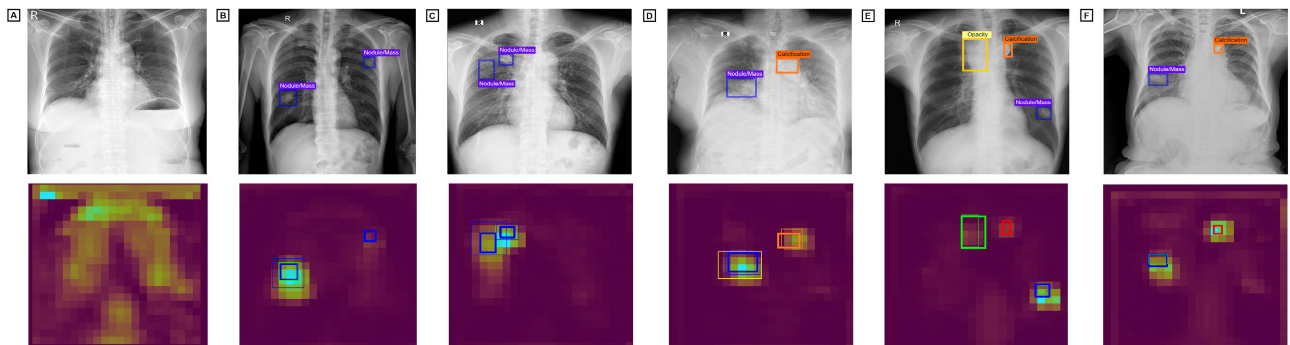


FIGURE 6. Visual explainability of the VinDr-CXR system. Visualization of CXRs and saliency maps overlaid on the original CXR images for several patients chosen from the internal validation set. The CXRs and affected areas were provided by our expert radiologist (top row). The corresponding saliency maps for CXRs were found by the VinDr-CXR classification model (bottom row). The blue, red, and green boxes cover the regions of Nodule/Mass, Calcification, and Opacity, respectively. Boxes with thick and thin lines denote ground-truth boxes and detected boxes, respectively. This figure indicates that the saliency maps highlighted parts of the CXRs that contain abnormal patterns (local labels), which are clinically correlated with disease conclusions (global labels). Best viewed in color.

system could consistent across geographical settings. Several studies [68], [69] showed that DL models for the CXR interpretation failed to generalize to image sources from new institutions and hospitals. In contrast, our finding is similar to Cohen *et al.* [70] who presented evidence that generalization over difference distributions is not due to a shift in the images but instead a shift in the labels. Third, our external validation experiment showed evidence that a DL system trained on digital CXR could generalize well on CXR captured by smartphone cameras without additional training cost. This opens the opportunity to integrate the DL system like VinDr-CXR into the vast spectrum of clinical workflows across the world, including developing regions that is still using films.

In terms of novelty, the VinDr-CXR system shows several contributions. First, while most previous studies [31], [37], [38], [39], [40], [41] used machine-generated annotations that contain many CXR images with uncertainty labels to train DL systems, our model was trained on radiologist-generated annotations for both development and validation data sets. Next, the proposed DL system evaluated both the detection and classification tasks, while most previous studies have only evaluated image-level classification performance without specifying the location of abnormal findings. Last, this study conducted a large-scale, clinical evaluation to investigate the actual impact of a DL system on the variability in radiologist performance in the interpretation of CXRs. We showed that the system significantly improved agreement among radiologists. To the best of our knowledge, we are the first to show that a DL system trained on a large-scale, annotated dataset can offer clinical value by helping to improve the rate of agreement among physicians. Furthermore, we also observed that VinDr-CXR assistance resulted in a significant increase in the agreement between the DL system and radiologists. Note that most literature refers to comparisons between human performing and AI models on the CXR interpretation, usually in diagnostic accuracy on the same clinical validation dataset. We suggest that these comparisons do not offer valuable insights into these systems' impact on clinical practice.

This study is not without limitations. First, the development and evaluation datasets only contain frontal CXR scans. Meanwhile, several clinical findings require lateral views. The next version of the dataset may consider adding the lateral views to train the VinDr-CXR system. Second, in clinical practice, physicians diagnose diseases based on both the patient's clinical history and visual information from CXRs. The VinDr-CXR, however, used only image information for providing diagnosis results without taking clinical and laboratory information into account. Third, the current study's most significant limitation is that we did not directly measure the actual impact of the VinDr-CXR on the sensitivity and specificity of participating radiologists due to the lack of gold reference ground truth. We showed in this study that the DL system was able to reduce clinical disagreements among radiologists. However, there is no clear evidence that the DLS helps improve the sensitivity or specificity of the radiologist in CXR interpretation.

V. CONCLUSION

We reported in this paper the development and validation of an AI-based system called VinDr-CXR for identifying 14 abnormal findings and classifying six common lung diseases from CXRs. To achieve this goal, we collected and annotated a large-scale CXR dataset of 53,485 studies across two major hospitals. Using radiologist-generated annotations as the reference standard to train the VinDr-CXR, we showed that it could achieve a performance level on par with a group of experienced radiologists in classifying common thoracic diseases. The system also achieved high diagnostic accuracy for most abnormal findings. Extensive validation experiments confirmed that the VinDr-CXR system generalizes well to datasets acquired from different patient cohorts. Most importantly, our clinical evaluations determined the value of the VinDr-CXR in clinical practice when it improved the agreement among physicians. The proposed system could be directly applied in different clinical settings, e.g., supporting physicians in triaging cases or using as a second reader. Although many DL-based models for predicting lung




FIGURE 7. (A) EfficientNet architecture: scaling method that uniformly scales all three dimensions with a fixed ratio. (B) EfficientDet architecture employs EfficientNet as the backbone network, BiFPN as the feature network, and shared class/box prediction network. These figures were reproduced from the original papers.

TABLE 10. Inter-rater agreement among radiologists at the HMUH without the VinDr-CXR assistance.

Findings	Rater 1 vs. Rater 2		Rater 1 vs. Rater 3		Rater 2 vs. Rater 3		Fleiss' κ
	Agree	Cohen's κ	Agree	Cohen's κ	Agree	Cohen's κ	
Lung Tumor	0.955	0.617 (0.388, 0.846)	0.945	0.591 (0.372, 0.810)	0.960	0.646 (0.419, 0.873)	0.616 (0.536, 0.696)
Pneumonia	0.865	0.536 (0.388, 0.684)	0.815	0.312 (0.157, 0.468)	0.870	0.281 (0.080, 0.482)	0.383 (0.303, 0.463)
Tuberculosis	0.925	0.675 (0.522, 0.827)	0.915	0.385 (0.164, 0.605)	0.880	0.298 (0.111, 0.486)	0.474 (0.394, 0.554)
Other Diseases	0.775	0.470 (0.344, 0.596)	0.660	0.234 (0.105, 0.363)	0.735	0.440 (0.312, 0.567)	0.373 (0.293, 0.453)
No Finding	0.880	0.758 (0.668, 0.848)	0.865	0.730 (0.636, 0.825)	0.905	0.811 (0.731, 0.890)	0.766 (0.686, 0.846)
Aortic Enlargement	0.955	0.549 (0.289, 0.810)	0.915	-0.032 (-0.056, -0.007)	0.940	-0.027(-0.047, -0.008)	0.207 (0.127, 0.287)
Atelectasis	0.980	0.490 (0.060, 0.920)	0.975	-0.008 (-0.021, 0.005)	0.975	-0.008 (-0.021, 0.005)	0.210 (0.130, 0.290)
Calcification	0.925	0.610 (0.432, 0.789)	0.920	0.556 (0.363, 0.749)	0.945	0.735 (0.586, 0.884)	0.638 (0.558, 0.718)
Cardiomegaly	0.890	0.277 (0.076, 0.478)	0.925	0.318 (0.067, 0.568)	0.885	0.404 (0.207, 0.601)	0.333 (0.253, 0.413)
Consolidation	0.935	0.216 (-0.046, 0.478)	0.955	0.164 (-0.144, 0.471)	0.930	0.330 (0.066, 0.593)	0.249 (0.169, 0.329)
ILD	0.920	0.488 (0.275, 0.702)	0.815	0.131 (-0.028, 0.290)	0.845	0.352 (0.175, 0.529)	0.303 (0.223, 0.383)
Infiltration	0.885	0.000 (0.000, 0.000)	0.980	0.000 (0.000, 0.000)	0.885	0.118 (-0.056, 0.292)	0.030 (-0.050, 0.110)
Lung Opacity	0.785	0.229 (0.072, 0.386)	0.785	0.103 (-0.020, 0.225)	0.870	0.020 (-0.114, 0.154)	0.116 (0.036, 0.196)
Nodule/Mass	0.955	0.688 (0.498, 0.877)	0.915	0.475 (0.267, 0.682)	0.900	0.490 (0.298, 0.682)	0.540 (0.460, 0.620)
Pleural Effusion	0.925	0.509 (0.298, 0.721)	0.945	0.674 (0.494, 0.853)	0.930	0.425 (0.178, 0.673)	0.547 (0.467, 0.627)
Pleural Thickening	0.940	0.377 (0.106, 0.648)	0.975	0.275 (-0.164, 0.714)	0.925	0.102 (-0.100, 0.304)	0.245 (0.165, 0.325)
Pneumothorax	1.000	1.000 (1.000, 1.000)	1.000	1.000 (1.000, 1.000)	1.000	1.000 (1.000, 1.000)	1.000 (1.000, 1.000)
Pulmonary Fibrosis	0.905	0.545 (0.367, 0.724)	0.895	0.343 (0.125, 0.560)	0.870	0.367 (0.177, 0.557)	0.423 (0.343, 0.503)
Other Lesions	0.890	0.175 (-0.024, 0.375)	0.920	0.165 (-0.073, 0.402)	0.890	0.332 (0.120, 0.544)	0.232 (0.152, 0.312)
Mean	0.910	0.485 (0.301, 0.669)	0.901	0.338 (0.171, 0.504)	0.902	0.374 (0.217, 0.532)	0.404 (0.329, 0.480)

diseases have improved diagnostic accuracy, in some cases surpassing radiologists' performance, there is little evidence showing that deployment of these models has improved patient outcomes. Therefore, further research is needed to validate the model prospectively and determine its utility in clinical settings. For example, the diagnostic and clinical effects of the VinDr-CXR needs to be assessed in large-scale test cohorts to determine the change in sensitivity and specificity of radiologists for the CXR interpretation in routine clinical practice.

APPENDICES

A. DEVELOPMENT OF DEEP LEARNING ALGORITHMS

1) NETWORK ARCHITECTURES

Deep learning (DL) [71], a subfield of machine learning, is a computational model that composes multiple processing layers and uses data-driven rules to learn representations of data with multiple levels of abstraction. DL networks showed their breakthrough successes in a wide variety of diagnostic tasks in medical imaging analysis [13], [37], [45], [72]. In this study, we applied two well-known DL networks for classifying common thoracic diseases and detecting abnormal findings in CXR images. We deployed EfficientNet-B6 [56] for the task of disease classification and EfficientDet-D6 [57] for the task of lesion detection. These network architectures

were well-known as the most commonly used and most successful DL networks for image classification and object detection. By balancing network depth, width, and resolution, the EfficientNet [56] can lead to much better accuracy and efficiency than other state-of-the-art CNN models. Meanwhile, the EfficientDet-D6 [57] used a weighted bi-directional feature pyramid network (BiFPN), allowing easy and fast multiscale feature fusion and combined with the feature learning capacity of EfficientNet-B6 [56]. This network architecture design, which requires less computational resources for training, achieved much better efficiency than prior state-of-the-art detectors. Therefore, we found that EfficientNet-B6 [56] and EfficientDet-D6 [57] well-suited for medical applications, including the CXR analysis. Figure 7 illustrates the key ideas behind these two network architectures.

2) IMPLEMENTATION DETAILS

We followed the original implementation of EfficientNet-B6 [56] and EfficientDet-D6 [57] with some minor modifications. Specifically, the last fully-connected layer of EfficientNet-B6 [56] has changed to output a vector of six dimensions corresponding to six classes. Mean binary cross-entropy loss was used to train these six classes simultaneously instead of a single multi-class cross-entropy loss in

TABLE 11. Inter-rater agreement among radiologists at the HMUH with the VinDr-CXR assistance.

Findings	Rater 1 vs. Rater 2		Rater 1 vs. Rater 3		Rater 2 vs. Rater 3		Fleiss' κ
	Agree	Cohen's κ	Agree	Cohen's κ	Agree	Cohen's κ	
Lung Tumor	0.955	0.617 (0.388, 0.846)	0.950	0.640 (0.433, 0.846)	0.965	0.702 (0.494, 0.910)	0.652 (0.572, 0.732)
Pneumonia	0.865	0.536 (0.388, 0.684)	0.820	0.338 (0.181, 0.494)	0.875	0.323 (0.121, 0.524)	0.403 (0.323, 0.483)
Tuberculosis	0.925	0.685 (0.537, 0.833)	0.915	0.385 (0.164, 0.605)	0.870	0.279 (0.100, 0.459)	0.471 (0.391, 0.551)
Other Diseases	0.775	0.470 (0.344, 0.596)	0.665	0.241 (0.111, 0.371)	0.740	0.449 (0.322, 0.576)	0.379 (0.299, 0.459)
No Finding	0.880	0.758 (0.668, 0.848)	0.870	0.740 (0.647, 0.833)	0.900	0.801 (0.719, 0.882)	0.766 (0.686, 0.846)
Aortic Enlargement	0.945	0.565 (0.337, 0.793)	0.920	0.243 (0.001, 0.486)	0.935	0.103 (-0.134, 0.341)	0.340 (0.260, 0.420)
Atelectasis	0.975	0.432 (0.021, 0.842)	0.970	-0.008 (-0.022, 0.005)	0.975	-0.008 (-0.021, 0.005)	0.186 (0.106, 0.266)
Calcification	0.930	0.656 (0.490, 0.823)	0.940	0.694 (0.531, 0.856)	0.940	0.716 (0.564, 0.868)	0.689 (0.609, 0.769)
Cardiomegaly	0.885	0.267 (0.071, 0.463)	0.930	0.335 (0.076, 0.593)	0.885	0.405 (0.210, 0.601)	0.333 (0.253, 0.413)
Consolidation	0.935	0.216 (-0.046, 0.478)	0.955	0.164 (-0.144, 0.471)	0.930	0.330 (0.066, 0.593)	0.249 (0.169, 0.329)
ILD	0.930	0.575 (0.376, 0.773)	0.815	0.161 (-0.002, 0.324)	0.855	0.412 (0.238, 0.585)	0.361 (0.281, 0.441)
Infiltration	0.880	0.000 (0.000, 0.000)	0.980	0.000 (0.000, 0.000)	0.880	0.112 (-0.055, 0.280)	0.026 (-0.054, 0.106)
Lung Opacity	0.775	0.212 (0.057, 0.367)	0.780	0.099 (-0.020, 0.219)	0.865	0.017 (-0.113, 0.146)	0.106 (0.026, 0.186)
Nodule/Mass	0.950	0.662 (0.469, 0.855)	0.915	0.500 (0.299, 0.701)	0.905	0.525 (0.338, 0.712)	0.554 (0.474, 0.634)
Pleural Effusion	0.925	0.509 (0.298, 0.721)	0.945	0.674 (0.494, 0.853)	0.930	0.425 (0.178, 0.673)	0.547 (0.467, 0.627)
Pleural Thickening	0.930	0.338 (0.084, 0.593)	0.975	0.275 (-0.164, 0.714)	0.915	0.089 (-0.092, 0.270)	0.219 (0.139, 0.299)
Pneumothorax	1.000	1.000 (1.000, 1.000)	1.000	1.000 (1.000, 1.000)	1.000	1.000 (1.000, 1.000)	1.000 (1.000, 1.000)
Pulmonary fibrosis	0.890	0.484 (0.301, 0.666)	0.895	0.343 (0.125, 0.560)	0.875	0.403 (0.216, 0.590)	0.413 (0.333, 0.493)
Other Lesions	0.895	0.189 (-0.013, 0.392)	0.915	0.158 (-0.066, 0.383)	0.890	0.362 (0.153, 0.571)	0.248 (0.168, 0.328)
Mean	0.908	0.483 (0.304, 0.662)	0.903	0.367 (0.192, 0.543)	0.902	0.392 (0.227, 0.557)	0.418 (0.342, 0.494)

TABLE 12. Agreement between the VinDr-CXR system and HMUH's radiologists without assistance.

Findings	Rater 1 vs. AI		Rater 2 vs. AI		Rater 3 vs. AI	
	Agree	Cohen's κ	Agree	Cohen's κ	Agree	Cohen's κ
Lung Tumor	0.940	0.427 (0.164, 0.690)	0.975	0.693 (0.439, 0.947)	0.945	0.451 (0.182, 0.719)
Pneumonia	0.845	0.513 (0.366, 0.661)	0.890	0.559 (0.396, 0.721)	0.870	0.417 (0.239, 0.594)
Tuberculosis	0.910	0.628 (0.472, 0.784)	0.945	0.793 (0.675, 0.910)	0.865	0.271 (0.096, 0.446)
Other Diseases	0.765	0.390 (0.248, 0.532)	0.840	0.640 (0.528, 0.752)	0.685	0.314 (0.181, 0.447)
No finding	0.860	0.721 (0.625, 0.816)	0.870	0.742 (0.652, 0.833)	0.845	0.689 (0.589, 0.790)
Aortic Enlargement	0.945	0.616 (0.409, 0.823)	0.930	0.430 (0.188, 0.672)	0.910	0.154 (-0.058, 0.366)
Atelectasis	0.970	0.235 (-0.167, 0.636)	0.980	0.490 (0.060, 0.920)	0.985	0.395 (-0.146, 0.936)
Calcification	0.900	0.237 (0.011, 0.462)	0.895	0.354 (0.146, 0.562)	0.920	0.463 (0.246, 0.680)
Cardiomegaly	0.875	0.203 (0.017, 0.389)	0.925	0.674 (0.520, 0.828)	0.870	0.340 (0.146, 0.535)
Consolidation	0.960	0.318 (-0.026, 0.662)	0.925	0.310 (0.052, 0.568)	0.935	0.201 (-0.073, 0.476)
ILD	0.850	0.178 (-0.006, 0.361)	0.880	0.433 (0.245, 0.621)	0.815	0.286 (0.113, 0.459)
Infiltration	0.885	0.000 (0.000, 0.000)	0.880	0.410 (0.217, 0.604)	0.885	0.118 (-0.056, 0.292)
Lung Opacity	0.785	0.179 (0.030, 0.329)	0.880	0.269 (0.060, 0.479)	0.900	0.045 (-0.124, 0.215)
Nodule/Mass	0.880	0.378 (0.190, 0.566)	0.875	0.442 (0.261, 0.623)	0.855	0.390 (0.211, 0.570)
Pleural Effusion	0.945	0.640 (0.447, 0.833)	0.940	0.423 (0.154, 0.691)	0.970	0.754 (0.565, 0.943)
Pleural Thickening	0.945	-0.028 (-0.045, -0.011)	0.895	-0.045 (-0.072, -0.018)	0.970	0.239 (-0.158, 0.635)
Pneumothorax	0.990	0.496 (-0.104, 1.000)	0.990	0.496 (-0.104, 1.000)	0.990	0.496 (-0.104, 1.000)
Pulmonary Fibrosis	0.880	0.302 (0.093, 0.511)	0.865	0.375 (0.188, 0.562)	0.865	0.196 (-0.003, 0.396)
Other Lesions	0.935	-0.033 (-0.052, -0.015)	0.895	0.235 (0.023, 0.447)	0.905	0.051 (-0.126, 0.228)
Mean	0.898	0.337 (0.141, 0.528)	0.909	0.459 (0.244, 0.670)	0.894	0.330 (0.090, 0.565)

TABLE 13. Agreement between the VinDr-CXR system and HMUH's radiologists with assistance.

Findings	Rater 1 vs. AI		Rater 2 vs. AI		Rater 3 vs. AI	
	Agree	Cohen's κ	Agree	Cohen's κ	Agree	Cohen's κ
Lung Tumor	0.940	0.427 (0.164, 0.690)	0.975	0.693 (0.439, 0.947)	0.950	0.523 (0.267, 0.778)
Pneumonia	0.845	0.513 (0.366, 0.661)	0.890	0.559 (0.396, 0.721)	0.875	0.447 (0.271, 0.623)
Tuberculosis	0.910	0.628 (0.472, 0.784)	0.945	0.798 (0.683, 0.913)	0.865	0.271 (0.096, 0.446)
Other Diseases	0.765	0.390 (0.248, 0.532)	0.840	0.640 (0.528, 0.752)	0.690	0.322 (0.188, 0.455)
No Finding	0.860	0.721 (0.625, 0.816)	0.870	0.742 (0.652, 0.833)	0.850	0.700 (0.601, 0.799)
Aortic Enlargement	0.955	0.718 (0.544, 0.893)	0.940	0.542 (0.315, 0.769)	0.915	0.231 (-0.003, 0.464)
Atelectasis	0.965	0.205 (-0.162, 0.571)	0.980	0.490 (0.060, 0.920)	0.985	0.395 (-0.146, 0.936)
Calcification	0.905	0.343 (0.119, 0.566)	0.895	0.354 (0.146, 0.562)	0.925	0.511 (0.302, 0.721)
Cardiomegaly	0.875	0.203 (0.017, 0.389)	0.940	0.743 (0.605, 0.881)	0.875	0.354 (0.158, 0.550)
Consolidation	0.960	0.318 (-0.026, 0.662)	0.925	0.310 (0.052, 0.568)	0.935	0.201 (-0.073, 0.476)
ILD	0.865	0.275 (0.080, 0.469)	0.885	0.466 (0.281, 0.651)	0.820	0.315 (0.143, 0.487)
Infiltration	0.885	0.000 (0.000, 0.000)	0.885	0.446 (0.255, 0.636)	0.885	0.118 (-0.056, 0.292)
Lung Opacity	0.780	0.174 (0.027, 0.321)	0.885	0.318 (0.108, 0.527)	0.900	0.045 (-0.124, 0.215)
Nodule/Mass	0.875	0.364 (0.176, 0.551)	0.875	0.442 (0.261, 0.623)	0.870	0.461 (0.286, 0.636)
Pleural Effusion	0.945	0.640 (0.447, 0.833)	0.940	0.423 (0.154, 0.691)	0.970	0.754 (0.565, 0.943)
Pleural Thickening	0.945	-0.028 (-0.045, -0.011)	0.895	0.045 (-0.116, 0.205)	0.970	0.239 (-0.158, 0.635)
Pneumothorax	0.990	0.496 (-0.104, 1.000)	0.990	0.496 (-0.104, 1.000)	0.990	0.496 (-0.104, 1.000)
Pulmonary Fibrosis	0.880	0.302 (0.093, 0.511)	0.860	0.364 (0.178, 0.549)	0.865	0.196 (-0.003, 0.396)
Other Lesions	0.940	-0.030 (-0.048, -0.012)	0.895	0.235 (0.023, 0.447)	0.915	0.223 (-0.017, 0.463)
Mean	0.899	0.350 (0.158, 0.538)	0.911	0.479 (0.259, 0.694)	0.897	0.358 (0.115, 0.595)

the original works. For the detector, we only changed the classification head to return scores for 14 lesion types while the loss functions and other parts of the architecture were preserved.

3) DATA AUGMENTATION

Several pre-processing steps were performed on CXRs before passing to the networks. Images in the DICOM format were converted to 8-bit PNG, then padded and




FIGURE 8. Examples of CXR scans from internal and external datasets. (A) An original CXR scan in the DICOM format from the VinDr-CXR dataset. (B) An original CXR scan in Portable Graphics Format (.PNG) format from CheXpert [31] dataset. (C) A CXR image produced by a Nokia phone from CheXphoto [58] dataset. (D) A CXR image produced by iPhone from CheXphoto [58] dataset.




FIGURE 9. Representative cases from our reader study. The first column shows CXRs and lesions marked by participating radiologists for the first assessment. The middle column shows lesions identified by the VinDr-CXR system. The last column shows the final decision of the radiologists after consulting the VinDr-CXR's result. In many cases, radiologists have changed their previous decisions by adding or removing lesions.

resized to 1024×1024 pixels. Subsequently, 1-channel images were transformed to 3-channel ones by repeating the channel three times. In the last step, images with the pixels intensity in range $[0, 255]$ were normalized by subtracting $(123.675, 116.28, 103.53)$ then dividing by $(58.395, 57.12, 57.375)$ in a channel-wise manner.

4) TRAINING PROCEDURES

Regarding the training procedure, the classifier's weights were initialized with weights trained on the ImageNet

dataset [73], a large-scale natural image dataset for the classification task. Pre-processed training images were randomly transformed using resize-cropping, shift-scale-rotating, horizontal flipping, brightness-contrast adjustments, then grouped into batches of 16. The classifier was trained for multi-label binary classification tasks (presence or absence of each disease) by optimizing the mean binary cross-entropy loss of all diseases. A variant of the stochastic gradient descent algorithm, Adam optimizer [74], was used with a base learning rate of $2 \times 10e-4$. The learning rate was then linearly




FIGURE 10. VinDr-CXR interface. The clinical findings detected by the VinDr-CXR system are listed on the interface, and bounding box predictions are displayed on the image.

increased in the first epoch then gradually annealed to 0 at the end of the 50-*th* epoch, following the cosine function. For the detector, a similar training method was employed. Detector's weights were initialized with parameters trained on the COCO dataset [75], a large dataset with common objects marked by bounding boxes. The optimization objective incorporates both regression loss for box categories by weighted summation. A batch size of 8 was used due to larger memory consumption compared to the classifier and the total training schedule was 60 epochs.

B. CLINICAL EVALUATION RESULTS AT THE HMUH

We provide in this section experimental results of the clinical evaluation at the HMUH. Table 10 shows the inter-rater agreement among radiologists at the HMUH without the VinDr-CXR assistance. Table 11 shows the inter-rater agreement among radiologists at the HMUH with the VinDr-CXR assistance. Table 12 and Table 13 report the agreement between the VinDr-CXR system and HMUH's radiologists without and with assistance, respectively.

C. EXAMPLES OF CXR SCANS FROM INTERNAL AND EXTERNAL DATASETS

Figure 8 shows several examples of CXR scans from internal (VinDr-CXR) and external datasets (CheXpert [31] and CheXphoto [58]).

D. REPRESENTATIVE CASES FROM THE READER STUDY

Several representative cases from our reader study are provided in Figure 9. In many cases, radiologists have changed their previous decisions by adding or removing lesions.

E. VinDr-CXR INTERFACE

Figure 10 shows VinDr-CXR interface that allows displayed detected findings on the image via bounding box predictions.

F. DATA AVAILABILITY

To facilitate a wide range of research topics in computer vision and medical imaging, we made the VinDr-CXR dataset (18,000 studies) publicly available through PhysioNet at <https://physionet.org/content/vindr-cxr/1.0.0/>. The image and annotation quality of the dataset can be visually check via our project webpage at <https://vindr.ai/datasets/cxr>. The CheXpert dataset is publicly available at <https://stanfordmlgroup.github.io/competitions/chexpert/>. The CheXphoto dataset is publicly available at <https://stanfordmlgroup.github.io/competitions/chexphoto/>.

G. CODE AVAILABILITY

Implementation of our work is based on the following open source repositories: Pytorch: <https://pytorch.org>; OpenCV <https://opencv.org/>; Pydicom: <https://pydicom.github.io/>. The source code used to train the VinDr-CXR system is a part of a commercial software product and not available to the public. The commercial version of VinDr-CXR can be freely tried through an online demonstration at <https://vindr.ai/>. All performance metrics were calculated with the support of scikit-learn <https://scikit-learn.org/>. Our labeling framework called VinDr Lab was made as open-source software and available for downloading at <https://vindr.ai/vindr-lab>.

ACKNOWLEDGMENT

The authors would like to thank the Hanoi Medical University Hospital (HMUH) and the Hospital 108 (H108) for providing them access to their image databases and for agreeing to make the VinDr-CXR dataset publicly available, and also would like to thank all of their radiologists, physicians, and technicians, who participated in this project.

REFERENCES

- [1] World Health Organization (WHO). (2019). *Global Tuberculosis Report*. Accessed: Mar. 5, 2021. [Online]. Available: <https://apps.who.int/iris/bitstream/handle/10665/329368/9789241565714-eng.pdf?ua=1>

- [2] World Health Organization. *Cancer Statistics 2019*. Accessed: Mar. 4, 2021. [Online]. Available: <https://www.who.int/news-room/fact-sheets/detail/cancer>
- [3] J. Come and M. Kumaran, *Chest X-ray Made Easy E-Book*. Amsterdam, The Netherlands: Elsevier, 2015.
- [4] L. Delrue, R. Gosselin, B. Ilsen, A. Van Landeghem, J. de Mey, and P. Duyck, "Difficulties in the interpretation of chest radiography," in *Comparative Interpretation of CT and Standard Radiography of the Chest*. Berlin, Germany: Springer, 2011, pp. 27–49.
- [5] R. Fitzgerald, "Error in radiology," *Clin. Radiol.*, vol. 56, no. 12, pp. 938–946, Dec. 2001.
- [6] D. J. Manning, S. C. Ethell, and T. Donovan, "Detection or decision errors? Missed lung cancer from the posteroanterior chest radiograph," *Brit. J. Radiol.*, vol. 77, no. 915, pp. 231–235, Mar. 2004.
- [7] D. P. Carmody, C. F. Nodine, and H. L. Kundel, "An analysis of perceptual and cognitive factors in radiographic interpretation," *Perception*, vol. 9, no. 3, pp. 339–344, Jun. 1980.
- [8] J. J. Donald and S. A. Barnard, "Common patterns in 558 diagnostic radiology errors," *J. Med. Imag. Radiat. Oncol.*, vol. 56, no. 2, pp. 173–178, Apr. 2012.
- [9] G. Gavelli and E. Giampalma, "Sensitivity and specificity of chest X-ray screening for lung cancer," *Cancer*, vol. 89, no. S11, pp. 2453–2456, Dec. 2000.
- [10] J.-M. Peng et al., "Does training improve diagnostic accuracy and inter-rater agreement in applying the Berlin radiographic definition of acute respiratory distress syndrome? A multicenter prospective study," *Crit. Care*, vol. 21, no. 1, pp. 1–8, Dec. 2017.
- [11] D. C. Moncada, Z. V. Rueda, A. Macías, T. Suárez, H. Ortega, and L. A. Vélez, "Reading and interpretation of chest X-ray in adults with community-acquired pneumonia," *Brazilian J. Infectious Diseases*, vol. 15, no. 6, pp. 540–546, Dec. 2011.
- [12] H. S. Alghamdi, G. Amoudi, S. Elhag, K. Saeedi, and J. Nasser, "Deep learning approaches for detecting COVID-19 from chest X-ray images: A survey," *IEEE Access*, vol. 9, pp. 20235–20254, 2021.
- [13] G. Litjens, T. Kooi, B. E. Bejnordi, A. A. A. Setio, F. Ciompi, M. Ghafoorian, J. A. Van Der Laak, B. Van Ginneken, and C. I. Sánchez, "A survey on deep learning in medical image analysis," *Med. Image Anal.*, vol. 42, pp. 60–88, Dec. 2017.
- [14] H. H. Pham, T. T. Le, D. Q. Tran, D. T. Ngo, and H. Q. Nguyen, "Interpreting chest X-rays via CNNs that exploit hierarchical disease dependencies and uncertainty labels," *Neurocomputing*, vol. 437, pp. 186–194, May 2021.
- [15] I. Allaouzi and M. B. Ahmed, "A novel approach for multi-label chest X-ray classification of common thorax diseases," *IEEE Access*, vol. 7, pp. 64279–64288, 2019.
- [16] H. T. Nguyen, H. H. Pham, N. T. Nguyen, H. Q. Nguyen, T. Q. Huynh, M. Dao, and V. Vu, "VinDr-SpineXR: A deep learning framework for spinal lesions detection and classification from radiographs," in *Proc. Int. Conf. Med. Image Comput. Comput.-Assist. Intervent.* Cham, Switzerland: Springer, 2021, pp. 291–301.
- [17] J. D. Arias-Londoño, J. A. GómeZ-García, L. Moro-Velázquez, and J. I. Godino-Llorente, "Artificial intelligence applied to chest X-ray images for the automatic detection of COVID-19. A thoughtful evaluation approach," *IEEE Access*, vol. 8, pp. 226811–226827, 2020.
- [18] T. T. Tran, H. H. Pham, T. V. Nguyen, T. T. Le, H. T. Nguyen, and H. Q. Nguyen, "Learning to automatically diagnose multiple diseases in pediatric chest radiographs using deep convolutional neural networks," in *Proc. IEEE/CVF Int. Conf. Comput. Vis. Workshops (ICCVW)*, Oct. 2021, pp. 3314–3323.
- [19] H.-J. Kwon, D.-H. Shin, and K. Chung, "PGGAN-based anomaly classification on chest X-ray using weighted multi-scale similarity," *IEEE Access*, vol. 9, pp. 113315–113325, 2021.
- [20] S. B. Tran, H. T. X. Nguyen, H. H. Pham, and H. Q. Nguyen, "Transparency strategy-based data augmentation for BI-RADS classification of mammograms," 2022, *arXiv:2203.10609*.
- [21] H. C. Nguyen, T. T. Le, H. H. Pham, and H. Q. Nguyen, "VinDr-RibCXR: A benchmark dataset for automatic segmentation and labeling of individual ribs on chest X-rays," 2021, *arXiv:2107.01327*.
- [22] M. Ahishali, A. Degerli, M. Yamac, S. Kiranyaz, M. E. H. Chowdhury, K. Hameed, T. Hamid, R. Mazhar, and M. Gabbouj, "Advance warning methodologies for COVID-19 using chest X-ray images," *IEEE Access*, vol. 9, pp. 41052–41065, 2021.
- [23] K. H. Le, T. V. Tran, H. H. Pham, H. T. Nguyen, T. T. Le, and H. Q. Nguyen, "Learning from multiple expert annotators for enhancing anomaly detection in medical image analysis," 2022, *arXiv:2203.10611*.
- [24] T. Rahman, A. Khandakar, M. A. Kadir, K. R. Islam, K. F. Islam, R. Mazhar, T. Hamid, M. T. Islam, S. Kashem, Z. B. Mahbub, M. A. Ayari, and M. E. H. Chowdhury, "Reliable tuberculosis detection using chest X-ray with deep learning, segmentation and visualization," *IEEE Access*, vol. 8, pp. 191586–191601, 2020.
- [25] B. T. Dao, T. V. Nguyen, H. H. Pham, and H. Q. Nguyen, "Phase recognition in contrast-enhanced CT scans based on deep learning and random sampling," *Med. Phys.*, vol. 2022, pp. 1–26, Mar. 2022.
- [26] N. H. Nguyen, H. Q. Nguyen, N. T. Nguyen, T. V. Nguyen, H. H. Pham, and T. N.-M. Nguyen, "Deployment and validation of an AI system for detecting abnormal chest radiographs in clinical settings," *Frontiers Digit. Health*, vol. 4, p. 130, Jul. 2022.
- [27] S. Sakib, T. Tazrin, M. M. Fouda, Z. M. Fadlullah, and M. Guizani, "DL-CRC: Deep learning-based chest radiograph classification for COVID-19 detection: A novel approach," *IEEE Access*, vol. 8, pp. 171575–171589, 2020.
- [28] H. T. X. Nguyen, S. B. Tran, D. B. Nguyen, H. H. Pham, and H. Q. Nguyen, "A novel multi-view deep learning approach for BI-RADS and density assessment of mammograms," 2021, *arXiv:2112.04490*.
- [29] X. Wang, Y. Peng, L. Lu, Z. Lu, M. Bagheri, and R. M. Summers, "ChestX-ray8: Hospital-scale chest X-ray database and benchmarks on weakly-supervised classification and localization of common thorax diseases," in *Proc. IEEE Conf. Comput. Vis. Pattern Recognit. (CVPR)*, Jul. 2017, pp. 2097–2106.
- [30] A. Bustos, A. Pertusa, J.-M. Salinas, and M. de la Iglesia-Vayá, "PadChest: A large chest X-ray image dataset with multi-label annotated reports," *Med. Image Anal.*, vol. 66, Dec. 2020, Art. no. 101797.
- [31] J. Irvin, P. Rajpurkar, M. Ko, Y. Yu, S. Ciurea-Ilcus, C. Chute, H. Marklund, B. Haghgoo, R. Ball, K. Shpanskaya, J. Seekins, D. A. Mong, S. S. Halabi, J. K. Sandberg, R. Jones, D. B. Larson, C. P. Langlotz, B. N. Patel, M. P. Lungren, and A. Y. Ng, "CheXpert: A large chest radiograph dataset with uncertainty labels and expert comparison," in *Proc. AAAI Conf. Artif. Intell.*, 2019, vol. 33, no. 1, pp. 590–597.
- [32] A. E. W. Johnson, T. J. Pollard, S. J. Berkowitz, N. R. Greenbaum, M. P. Lungren, C.-Y. Deng, R. G. Mark, and S. Horng, "MIMIC-CXR, a de-identified publicly available database of chest radiographs with free-text reports," *Sci. Data*, vol. 6, no. 1, pp. 1–8, Dec. 2019.
- [33] H. Q. Nguyen et al., "VinDr-CXR: An open dataset of chest X-rays with radiologist's annotations," *Sci. Data*, vol. 9, no. 1, pp. 1–7, Dec. 2022.
- [34] N. H. Nguyen, H. H. Pham, T. T. Tran, T. N. M. Nguyen, and H. Q. Nguyen, "VinDr-PCXR: An open, large-scale chest radiograph dataset for interpretation of common thoracic diseases in children," 2022, *arXiv:2203.10612*.
- [35] H. T. Nguyen, H. Q. Nguyen, H. H. Pham, K. Lam, L. T. Le, M. Dao, and V. Vu, "VinDr-Mammo: A large-scale benchmark dataset for computer-aided diagnosis in full-field digital mammography," *MedRxiv*, vol. 2022, pp. 1–11, Jan. 2022.
- [36] H. H. Pham, D. V. Do, and H. Q. Nguyen, "DICOM imaging router: An open deep learning framework for classification of body parts from DICOM X-ray scans," *MedRxiv*, vol. 2021, pp. 1–4, Jan. 2021.
- [37] P. Rajpurkar, J. Irvin, K. Zhu, B. Yang, H. Mehta, T. Duan, D. Ding, A. Bagul, C. Langlotz, K. Shpanskaya, M. P. Lungren, and A. Y. Ng, "CheXNet: Radiologist-level pneumonia detection on chest X-rays with deep learning," 2017, *arXiv:1711.05225*.
- [38] P. Rajpurkar et al., "Deep learning for chest radiograph diagnosis: A retrospective comparison of the CheXNeXt algorithm to practicing radiologists," *PLOS Med.*, vol. 15, no. 11, Nov. 2018, Art. no. e1002686.
- [39] A. Majkowska, S. Mittal, D. F. Steiner, J. J. Reicher, S. M. McKinney, G. E. Duggan, K. Eswaran, P.-H. C. Chen, Y. Liu, S. R. Kalidindi, A. Ding, G. S. Corrado, D. Tse, and S. Shetty, "Chest radiograph interpretation with deep learning models: Assessment with radiologist-adjudicated reference standards and population-adjusted evaluation," *Radiology*, vol. 294, no. 2, pp. 421–431, Feb. 2020.
- [40] P. Rajpurkar, A. Joshi, A. Pareek, P. Chen, A. Kiani, J. Irvin, A. Y. Ng, and M. P. Lungren, "CheXpedition: Investigating generalization challenges for translation of chest X-ray algorithms to the clinical setting," 2020, *arXiv:2002.11379*.
- [41] Y.-X. Tang, Y.-B. Tang, Y. Peng, K. Yan, M. Bagheri, B. A. Redd, C. B. Brandon, Z. Lu, M. Han, J. Xiao, and R. M. Summers, "Automated abnormality classification of chest radiographs using deep convolutional neural networks," *npj Digit. Med.*, vol. 3, no. 1, pp. 1–8, Dec. 2020.

- [42] D. S. W. Ting and et al., "Development and validation of a deep learning system for diabetic retinopathy and related eye diseases using retinal images from multiethnic populations with diabetes," *JAMA*, vol. 318, no. 22, pp. 2211–2223, 2017.
- [43] G. T. Zago, R. V. Andreão, B. Dorizzi, and E. O. T. Salles, "Retinal image quality assessment using deep learning," *Comput. Biol. Med.*, vol. 103, pp. 64–70, Dec. 2018.
- [44] Y. Liu, A. Jain, C. Eng, D. H. Way, K. Lee, P. Bui, K. Kanada, G. de Oliveira Marinho, J. Gallegos, S. Gabriele, and V. Gupta, "A deep learning system for differential diagnosis of skin diseases," *Nature Med.*, vol. 26, no. 6, pp. 900–908, 2020.
- [45] A. Esteva, B. Kuprel, R. A. Novoa, J. Ko, S. M. Swetter, H. M. Blau, and S. Thrun, "Dermatologist-level classification of skin cancer with deep neural networks," *Nature*, vol. 542, no. 7639, pp. 115–118, 2017.
- [46] A. Smit, S. Jain, P. Rajpurkar, A. Pareek, A. Y. Ng, and M. P. Lungren, "CheXbert: Combining automatic labelers and expert annotations for accurate radiology report labeling using BERT," 2020, *arXiv:2004.09167*.
- [47] L. Oakden-Rayner. (2017). *Exploring the ChestXray14 Dataset: Problems*. Accessed: May 4, 2021. [Online]. Available: <https://lukeoakdenrayner.wordpress.com/2017/12/18/the-chestxray14-dataset-problems/>
- [48] L. Oakden-Rayner, "Exploring large-scale public medical image datasets," *Academic Radiol.*, vol. 27, no. 1, pp. 106–112, 2020. [Online]. Available: <http://www.sciencedirect.com/science/article/pii/S107663321930488X>
- [49] M. Nagendran, Y. Chen, C. A. Lovejoy, A. C. Gordon, M. Komorowski, H. Harvey, E. J. Topol, J. P. A. Ioannidis, G. S. Collins, and M. Maruthappu, "Artificial intelligence versus clinicians: Systematic review of design, reporting standards, and claims of deep learning studies," *BMJ*, vol. 368, p. m689, Mar. 2020.
- [50] S. Trajanovski, D. Mavroicidis, C. L. Swisher, B. G. Gebre, B. S. Veeling, R. Wiemker, T. Klinder, A. Tahmasebi, S. M. Regis, C. Wald, B. J. McKee, S. Flacke, H. Macmahon, and H. Pien, "Towards radiologist-level cancer risk assessment in CT lung screening using deep learning," *Computerized Med. Imag. Graph.*, vol. 90, Jun. 2021, Art. no. 101883.
- [51] J. Amann, A. Blasimme, E. Vayena, D. Frey, and V. I. Madai, "Explainability for artificial intelligence in healthcare: A multidisciplinary perspective," *BMC Med. Informat. Decis. Making*, vol. 20, no. 1, pp. 1–9, Dec. 2020.
- [52] A. Holzinger, G. Langs, H. Denk, K. Zatloukal, and H. Müller, "Causability and explainability of artificial intelligence in medicine," *WIREs Data Mining Knowl. Discovery*, vol. 9, no. 4, p. e1312, Jul. 2019.
- [53] US Department of Health and Human Services. (2003). *Summary of the HIPAA Privacy Rule*. [Online]. Available: <https://www.hhs.gov/hipaa/for-professionals/privacy/laws-regulations/index.html>
- [54] European Parliament and Council of European Union. (2016). *Regulation (EU) 2016/679 (General Data Protection Regulation)*. Accessed: Apr. 11, 2021. [Online]. Available: <https://gdpr-info.eu/>
- [55] Vietnamese National Assembly. (2009). *Regulation 40/2009/QH12 (Law on Medical Examination and Treatment)*. Accessed: May 11, 2021. [Online]. Available: <http://vbpl.vn/hanoi/Pages/vbpqen-toanvan.aspx?ItemID=10482>
- [56] M. Tan and Q. Le, "EfficientNet: Rethinking model scaling for convolutional neural networks," in *Proc. Int. Conf. Mach. Learn. (ICML)*, 2019, pp. 6105–6114.
- [57] M. Tan, R. Pang, and Q. V. Le, "EfficientDet: Scalable and efficient object detection," in *Proc. IEEE/CVF Conf. Comput. Vis. Pattern Recognit. (CVPR)*, Jun. 2020, pp. 10781–10790.
- [58] N. A. Phillips, P. Rajpurkar, M. Sabini, R. Krishnan, S. Zhou, A. Pareek, N. M. Phu, C. Wang, M. Jain, N. D. Du, S. Q. H. Truong, A. Y. Ng, and M. P. Lungren, "CheXphoto: 10,000+ photos and transformations of chest X-rays for benchmarking deep learning robustness," 2020, *arXiv:2007.06199*.
- [59] T.-Y. Lin, P. Goyal, R. Girshick, K. He, and P. Dollar, "Focal loss for dense object detection," in *Proc. IEEE Int. Conf. Comput. Vis. (ICCV)*, Oct. 2017, pp. 2980–2988.
- [60] J. Cohen, "A coefficient of agreement for nominal scales," *Educ. Psychol. Meas.*, vol. 20, no. 1, pp. 37–46, Apr. 1960.
- [61] J. L. Fleiss, "Measuring nominal scale agreement among many raters," *Psychol. Bull.*, vol. 76, no. 5, p. 378, 1971.
- [62] J. R. Landis and G. G. Koch, "The measurement of observer agreement for categorical data," *Biometrics*, vol. 33, pp. 159–174, Mar. 1977.
- [63] W. J. Youden, "Index for rating diagnostic tests," *Cancer*, vol. 3, no. 1, pp. 32–35, 1950.
- [64] P. Bunch, J. Hamilton, G. Sanderson, and A. Simmons, "A free-response approach to the measurement and characterization of radiographic-observer performance," *J. Appl. Photograph. Eng.*, vol. 4, no. 4, pp. 166–171, 1978.
- [65] FDA. *Proposed Regulatory Framework for Modifications to Artificial Intelligence/Machine Learning (AI/ML)-Based Software as a Medical Device (SaMD)*. Accessed: May 5, 2021. [Online]. Available: <https://www.fda.gov/files/medical%20devices/published/US-FDA-Artificial-Intelligence-and-Machine-Learning-Discussion-Paper.pdf>
- [66] Y. Xie, M. Chen, D. Kao, G. Gao, and X. Chen, "CheXplain: Enabling physicians to explore and understand data-driven, AI-enabled medical imaging analysis," in *Proc. CHI Conf. Hum. Factors Comput. Syst.*, Apr. 2020, pp. 1–13.
- [67] S. Karamizadeh, S. M. Abdullah, A. A. Manaf, M. Zamani, and A. Hooman, "An overview of principal component analysis," *J. Signal Inf. Process.*, vol. 4, no. 3, p. 173, 2013.
- [68] J. R. Zech, M. A. Badgeley, M. Liu, A. B. Costa, J. J. Titano, and E. K. Oerchman, "Variable generalization performance of a deep learning model to detect pneumonia in chest radiographs: A cross-sectional study," *PLOS Med.*, vol. 15, no. 11, Nov. 2018, Art. no. e1002683.
- [69] C. J. Kelly, A. Karthikesalingam, M. Suleyman, G. Corrado, and D. King, "Key challenges for delivering clinical impact with artificial intelligence," *BMC Med.*, vol. 17, no. 1, p. 195, Dec. 2019.
- [70] J. P. Cohen, M. Hashir, R. Brooks, and H. Bertrand, "On the limits of cross-domain generalization in automated X-ray prediction," 2020, *arXiv:2002.02497*.
- [71] Y. LeCun, Y. Bengio, and G. Hinton, "Deep learning," *Nature*, vol. 521, no. 7553, pp. 436–444, Sep. 2015.
- [72] N. Wu et al., "Deep neural networks improve radiologists' performance in breast cancer screening," *IEEE Trans. Med. Imag.*, vol. 39, no. 4, pp. 1184–1194, Apr. 2020.
- [73] A. Krizhevsky, I. Sutskever, and G. E. Hinton, "ImageNet classification with deep convolutional neural networks," in *Proc. Adv. Neural Inf. Process. Syst. (NIPS)*, vol. 25, Dec. 2012, pp. 1097–1105.
- [74] D. P. Kingma and J. Ba, "Adam: A method for stochastic optimization," 2014, *arXiv:1412.6980*.
- [75] T.-Y. Lin, M. Maire, S. Belongie, J. Hays, P. Perona, D. Ramanan, P. Dollár, and C. L. Zitnick, "Microsoft COCO: Common objects in context," in *Proc. Eur. Conf. Comput. Vis.* Cham, Switzerland: Springer, 2014, pp. 740–755.



HIEU H. PHAM (Member, IEEE) received the Engineering degree in industrial informatics from the Hanoi University of Science and Technology (HUST), Vietnam, in 2016, and the Ph.D. degree in computer science from the Toulouse Computer Science Research Institute (IRIT), University of Toulouse, France, in 2019. He is currently a Teaching Fellow at the College of Engineering and Computer Science (CECS), VinUniversity, and serves as an Associate Director at the VinUni-

Illinois Smart Health Center. His research interests include computer vision, machine learning, medical image analysis, and their applications in smart healthcare. He is the author and coauthor of 30 scientific articles appeared in about 20 conferences and journals, such as *Computer Vision and Image Understanding*, *Neurocomputing*, International Conference on Medical Image Computing and Computer-Assisted Intervention (MICCAI), Medical Imaging with Deep Learning (MIDL), IEEE International Conference on Image Processing (ICIP), and IEEE International Conference on Computer Vision (ICCV). He is also currently serving as Reviewer for *MICCAI*, *ICCV*, *CVPR*, *IET Computer Vision Journal (IET-CVI)*, *IEEE JOURNAL OF BIOMEDICAL AND HEALTH INFORMATICS (JBHI)*, and *Scientific Reports (Nature)*. Before joining VinUniversity, he worked at the Vingroup Big Data Institute (VinBigData), as a Research Scientist and the Head of the Fundamental Research Team. With this position, he led several research projects on medical AI, including collecting various types of medical data, managing and annotating data, and developing new AI solutions for medical analysis.



HA Q. NGUYEN was born in Haiphong, Vietnam, in 1983. He received the B.S. degree in mathematics from the Hanoi National University of Education, Hanoi, Vietnam, the S.M. degree in electrical engineering and computer science from the Massachusetts Institute of Technology, Cambridge, MA, USA, and the Ph.D. degree in electrical and computer engineering from the University of Illinois at Urbana–Champaign, Champaign, IL, USA, in 2005, 2009, and 2014, respectively.

From 2009 to 2011, he was a Lecturer in electrical engineering with International University, Vietnam National University, Ho Chi Minh City, Vietnam, from 2014 to 2017, he was a Postdoctoral Research Associate with the Biomedical Imaging Group, Ecole Polytechnique Federale de Lausanne, Lausanne, Switzerland, and from 2017 to 2018, he was a Signal Processing Engineer with the Viettel Research and Development Institute, Hanoi. He is currently the Head of the Medical Imaging Department, Vingroup Big Data Institute, Hanoi. His research interests include medical image analysis, machine learning, computational imaging, and data compression. He was a fellow of Vietnam Education Foundation, Cohort 2007. He was a recipient of the Best Student Paper Award (second prize) of the IEEE International Conference on Acoustics, Speech and Signal Processing, in 2014, for his paper (with P.A. Chou and Y. Chen) on the compression of human body sequences using graph wavelet filter banks.



HIEU T. NGUYEN received the B.S. degree in computer science from the Department of Computer Science, Hanoi University of Science and Technology (HUST), Hanoi, Vietnam. He is currently pursuing the Ph.D. degree in CS with Northeastern University, USA. He is also a Research Intern at the Medical Imaging Department, Vingroup Big Data Institute (VinBDI), Hanoi. His research interests include medical image analysis and deep learning techniques.



LINH T. LE is currently the Director of the Department of Radiology, Hanoi Medical University Hospital, Hanoi, Vietnam. He is also a Faculty Member at Hanoi Medical University Hospital. He is an Experienced Radiologist with 20 years working on medical imaging analysis (X-ray analysis, CT, MRI interpretation). He is currently working on chest X-ray imaging and many other imaging modalities.



LAM KHANH received the Ph.D. degree in medical imaging from the National Institute for Basic Biology, Japan. He is currently an Associate Professor, the Major General, and the Deputy Director of the 108 Military Central Hospital. His research interests include MEG and neuro MRI imaging. He is currently working on brain magnetic resonance imaging, interventional imaging, application of multi-sequence CT in diagnosis, and treatment.

...