

RESEARCH ARTICLE

Variable generalization performance of a deep learning model to detect pneumonia in chest radiographs: A cross-sectional study

John R. Zech¹✉, Marcus A. Badgeley²✉, Manway Liu², Anthony B. Costa³, Joseph J. Titano⁴, Eric Karl Oermann^{3*}

1 Department of Medicine, California Pacific Medical Center, San Francisco, California, United States of America, **2** Verily Life Sciences, South San Francisco, California, United States of America, **3** Department of Neurological Surgery, Icahn School of Medicine, New York, New York, United States of America, **4** Department of Radiology, Icahn School of Medicine, New York, New York, United States of America

✉ These authors contributed equally to this work.

* eric.oermann@mountsinai.org



OPEN ACCESS

Citation: Zech JR, Badgeley MA, Liu M, Costa AB, Titano JJ, Oermann EK (2018) Variable generalization performance of a deep learning model to detect pneumonia in chest radiographs: A cross-sectional study. *PLoS Med* 15(11): e1002683. <https://doi.org/10.1371/journal.pmed.1002683>

Academic Editor: Aziz Sheikh, Edinburgh University, UNITED KINGDOM

Received: April 10, 2018

Accepted: September 28, 2018

Published: November 6, 2018

Copyright: © 2018 Zech et al. This is an open access article distributed under the terms of the [Creative Commons Attribution License](https://creativecommons.org/licenses/by/4.0/), which permits unrestricted use, distribution, and reproduction in any medium, provided the original author and source are credited.

Data Availability Statement: The code pipeline used to train and test the model across multiple institutions is available at <https://github.com/jrzech/cxr-generalize>. The NIH ChestX-ray14 dataset was curated and made publicly available by the National Institutes of Health (NIH) Clinical Center (<https://nihcc.app.box.com/v/ChestXray-NIHCC>). The Open-I dataset of chest radiographs from the Indiana University Hospital network was curated and made publicly available by the National Library of Medicine, NIH (<https://openi.nlm.nih.gov/faq>).

Abstract

Background

There is interest in using convolutional neural networks (CNNs) to analyze medical imaging to provide computer-aided diagnosis (CAD). Recent work has suggested that image classification CNNs may not generalize to new data as well as previously believed. We assessed how well CNNs generalized across three hospital systems for a simulated pneumonia screening task.

Methods and findings

A cross-sectional design with multiple model training cohorts was used to evaluate model generalizability to external sites using split-sample validation. A total of 158,323 chest radiographs were drawn from three institutions: National Institutes of Health Clinical Center (NIH; 112,120 from 30,805 patients), Mount Sinai Hospital (MSH; 42,396 from 12,904 patients), and Indiana University Network for Patient Care (IU; 3,807 from 3,683 patients). These patient populations had an age mean (SD) of 46.9 years (16.6), 63.2 years (16.5), and 49.6 years (17) with a female percentage of 43.5%, 44.8%, and 57.3%, respectively. We assessed individual models using the area under the receiver operating characteristic curve (AUC) for radiographic findings consistent with pneumonia and compared performance on different test sets with DeLong's test. The prevalence of pneumonia was high enough at MSH (34.2%) relative to NIH and IU (1.2% and 1.0%) that merely sorting by hospital system achieved an AUC of 0.861 (95% CI 0.855–0.866) on the joint MSH–NIH dataset. Models trained on data from either NIH or MSH had equivalent performance on IU (*P* values 0.580 and 0.273, respectively) and inferior performance on data from each other relative to an internal test set (i.e., new data from within the hospital system used for training data; *P* values both <0.001). The highest internal performance was achieved by combining training and test data from MSH and NIH (AUC 0.931, 95% CI 0.927–0.936), but this model

php). Retrospective data used in this study from Mount Sinai Health System cannot be released under the terms of our Institutional Review Board approval to protect patient confidentiality. Researchers interested in accessing Mount Sinai data through the Imaging Research Warehouse Initiative may contact Zahi Fayad, PhD at zahi.fayad@mssm.edu.

Funding: The Department of Radiology at the Icahn School of Medicine at Mount Sinai (<http://icahn.mssm.edu/about/departments/radiology>) supported this project financially via internal department funding (author JJT). No other authors received specific funding for this work. The funder had no role in study design, data collection and analysis, decision to publish, or preparation of the manuscript.

Competing interests: I have read the journal's policy and the authors of this manuscript have the following competing interests: MAB and ML are currently employees at Verily Life Sciences, which played no role in the research and has no commercial interest in it. EKO and ABC receive funding from Intel for unrelated work.

Abbreviations: AUC, area under the receiver operating characteristic curve; CAD, computer-aided diagnosis; CNN, convolutional neural network; ICU, intensive care unit; ILSVRC, ImageNet Large Scale Visual Recognition Competition; IU, Indiana University Network for Patient Care; MNIST, Modified National Institute of Standards and Technology; MSH, Mount Sinai Hospital; NIH, National Institutes of Health Clinical Center; NLP, natural language processing; NPV, negative predictive value; PACS, picture archiving and communication system; PPV, positive predictive value.

demonstrated significantly lower external performance at IU (AUC 0.815, 95% CI 0.745–0.885, $P = 0.001$). To test the effect of pooling data from sites with disparate pneumonia prevalence, we used stratified subsampling to generate MSH–NIH cohorts that only differed in disease prevalence between training data sites. When both training data sites had the same pneumonia prevalence, the model performed consistently on external IU data ($P = 0.88$). When a 10-fold difference in pneumonia rate was introduced between sites, internal test performance improved compared to the balanced model (10× MSH risk $P < 0.001$; 10× NIH $P = 0.002$), but this outperformance failed to generalize to IU (MSH 10× $P < 0.001$; NIH 10× $P = 0.027$). CNNs were able to directly detect hospital system of a radiograph for 99.95% NIH (22,050/22,062) and 99.98% MSH (8,386/8,388) radiographs. The primary limitation of our approach and the available public data is that we cannot fully assess what other factors might be contributing to hospital system–specific biases.

Conclusion

Pneumonia-screening CNNs achieved better internal than external performance in 3 out of 5 natural comparisons. When models were trained on pooled data from sites with different pneumonia prevalence, they performed better on new pooled data from these sites but not on external data. CNNs robustly identified hospital system and department within a hospital, which can have large differences in disease burden and may confound predictions.

Author summary

Why was this study done?

- Early results in using convolutional neural networks (CNNs) on X-rays to diagnose disease have been promising, but it has not yet been shown that models trained on X-rays from one hospital or one group of hospitals will work equally well at different hospitals.
- Before these tools are used for computer-aided diagnosis in real-world clinical settings, we must verify their ability to generalize across a variety of hospital systems.

What did the researchers do and find?

- A cross-sectional design was used to train and evaluate pneumonia screening CNNs on 158,323 chest X-rays from the National Institutes of Health Clinical Center (NIH; $n = 112,120$ from 30,805 patients), Mount Sinai Hospital (42,396 from 12,904 patients), and Indiana University Network for Patient Care ($n = 3,807$ from 3,683 patients).
- In 3 out of 5 natural comparisons, performance on chest X-rays from outside hospitals was significantly lower than on held-out X-rays from the original hospital system.
- CNNs were able to detect where a radiograph was acquired (hospital system, hospital department) with extremely high accuracy and calibrate predictions accordingly.

What do these findings mean?

- The performance of CNNs in diagnosing diseases on X-rays may reflect not only their ability to identify disease-specific imaging findings on X-rays but also their ability to exploit confounding information.
- Estimates of CNN performance based on test data from hospital systems used for model training may overstate their likely real-world performance.

Introduction

There is significant interest in using convolutional neural networks (CNNs) to analyze radiology, pathology, or clinical imaging for the purposes of computer-aided diagnosis (CAD) [1–5]. These studies are generally performed utilizing CNN techniques that were pioneered on well-characterized computer vision datasets, including the ImageNet Large Scale Visual Recognition Competition (ILSVRC) and the Modified National Institute of Standards and Technology (MNIST) database of hand-drawn digits [6,7]. Training CNNs to classify images from these datasets is typically done by splitting the dataset into three subsets: train (data directly used to learn parameters for models), tune (data used to choose hyperparameter settings, also commonly referred to as “validation”), and test (data used exclusively for performance evaluation of models learned using train and tune data). CNNs are trained to completion with the first two, and the final set is used to estimate the model’s expected performance on new, previously unseen data.

An underlying premise of the test set implying future generalizability to new data is that the test set is reflective of the data that will be encountered elsewhere. Recent work in computer vision has demonstrated that the true generalization performance of even classic CIFAR-10 photograph classification CNNs to new data may be lower than previously believed [8]. In the biomedical imaging context, we can contrast “internal” model performance on new, previously unseen data gathered from the same hospital system(s) used for model training with “external” model performance on new, previously unseen data from different hospital systems [9,10]. External test data may be different in important ways from internal test data, and this may affect model performance, particularly if confounding variables exist in internal data that do not exist in external data [11]. In a large-scale deep learning study of retinal funduscopy, Ting and colleagues noted variation in the performance of CNNs trained to identify ocular disease across external hospital systems, with the areas under the receiver operating characteristic curve (AUCs) ranging from 0.889 to 0.983 and image-level concordance with human experts ranging from 65.8% to 91.2% on external datasets [4]. Despite the rapid push to develop deep learning systems on radiological data for academic and commercial purposes, to date, no study has assessed whether radiological CNNs actually generalize to external data. If external test performance of a system is inferior to internal test performance, clinicians may erroneously believe systems to be more accurate than they truly are in the deployed context, creating the potential for patient harm.

The primary aim of this study was to obtain data from three separate hospital systems and to assess how well deep learning models trained at one hospital system generalized to other external hospital systems. For the purposes of this assessment, we chose the diagnosis of pneumonia on chest X-ray for both its clinical significance as well as common occurrence and significant interest [2]. We reproduced the CheXNet model of Rajpurkar and colleagues, whose

internal performance on National Institutes of Health Clinical Center (NIH) data has previously been reported [2]. We extended upon this work to evaluate the model's internal performance when trained on data from a different hospital system and to demonstrate how this model generalized to external hospital systems not used for model training. By training and testing models on different partitions of data across three distinct institutions, we sought to establish whether a truly generalizable model could be learned, as well as which factors affecting external validity could be identified to aid clinicians when assessing models for potential clinical deployment.

Methods

Datasets

This study was approved by the Mount Sinai Health System Institutional Review Board; the requirement for patient consent was waived for this retrospective study that was deemed to carry minimal risk. Three datasets were obtained from different hospital groups: NIH (112,120 radiographs from 1992 to 2015), Indiana University Network for Patient Care (IU; 7,470 radiographs, date range not available), and Mount Sinai Hospital (MSH; 48,915 radiographs from 2009 to 2016) [1,12]. This study did not have a prospective analysis plan, and all analyses performed are subsequently described.

CNNs

Deep learning encompasses any algorithm that uses multiple layers of feed-forward neural networks to model phenomena [13]. Classification CNNs are a type of supervised deep learning model that take an image as input and predict the probability of predicted class membership as output. A typical use of CNNs is classifying photographs according to the animals or objects they contain: a chihuahua, a stove, a speedboat, etc. [6]. Many different CNN architectures have been proposed, including ResNet-50 and DenseNet-121 used in this paper, and improving the performance of these models is an active area of research [14,15]. In practice, CNNs are frequently pre-trained on large computer vision databases, such as ImageNet, rather than being randomly initialized and trained de novo. After pretraining, the CNNs are then fine-tuned on the dataset of interest. This process of pretraining followed by fine-tuning reduces training time, promotes model convergence, and can regularize the model to reduce overfitting. A difficulty of using these models is that there are few formal guarantees as to their generalization performance [16]. In this paper, we use CNNs both to preprocess and to predict pneumonia in radiographs.

Preprocessing: Frontal view filtering

NIH data contained only frontal chest radiographs, whereas IU and MSH data contained both frontal and lateral chest radiographs and were found to contain inconsistent frontal and lateral labels on manual review. A total of 402 IU and 490 MSH radiographs were manually labeled as frontal/lateral and randomly divided into groups (IU: 200 train, 100 tune, 102 test; MSH: 200 train, 100 tune, 190 test) and used to train ResNet-50 CNNs to identify frontal radiographs [14]. A total of 187/190 MSH and 102/102 IU test radiographs were accurately classified. The datasets were then filtered to frontal radiographs using these CNNs, leaving a total of 158,323 radiographs (112,120 NIH, 42,396 MSH, and 3,807 IU) available for analysis (S1 Fig).

Preprocessing: Generating labels for pathology

IU radiographs were manually labeled by curators after review of the accompanying text radiology reports [12]. NIH radiographs were labeled automatically using a proprietary natural

language processing (NLP) system based on expanding sentences as parse trees and using handcrafted rules based on the MESH vocabulary to identify statements indicating positive pathology [1].

MSH radiographs did not initially include labels, so a subset of radiographic reports was manually labeled to train an NLP algorithm that could infer labels for the full dataset. A total of 405 radiographic reports were manually labeled for cardiomegaly, emphysema, effusion, hernia, nodule, atelectasis, pneumonia, edema, and consolidation. To evaluate the NLP algorithm's performance, these were split into train and test groups (283 and 122, respectively). A previously described NLP concept extraction model based on 1- and 2-gram bag-of-words features with Lasso logistic regression was trained to identify reports mentioning pneumonia [17]. AUC, sensitivity, and specificity at a 50% classification threshold are reported in [S1 Table](#). The NLP model was then refit with all 405 manually labeled reports and used to process all unlabeled reports. As reports positive for hernia occurred too infrequently to use this NLP algorithm, reports were automatically labeled as positive for hernia if the word “hernia” appeared in the report.

Preprocessing: Separation of patients across train, tune, and test groups

As NIH and MSH data contained patient identifiers, all NIH and MSH patients were separated into fixed train (70%), tune (10%), and test (20%) groups ([S2 Fig](#)). IU data did not contain patient identifiers. In the case of pneumonia detection, 100% of IU data was reserved for use as an external test set. IU data were used for training only to detect hospital system and in this case were separated into fixed train (70%), tune (10%), and test (20%) groups using an identifier corresponding to accession number (e.g., which radiographs were obtained at the same time on the same patient). Test data were not available to CNNs during model training, and all results reported in this study are calculated exclusively on test data.

Preprocessing: Identifying MSH portable scans from inpatient wards and the emergency department

Of 42,396 MSH radiographs, 39,574 contained a label indicating whether they were portable radiographs; 31,838 were labeled as portable. We identified a subset of 31,076 MSH portable radiographs that documented the department of acquisition, with 28,841 from inpatient wards and 2,235 from the emergency department.

Model training

PyTorch 0.2.0 and torchvision were used for model training [18]. All images were resized to 224×224 . CNNs used for experiments were trained with DenseNet-121 architecture with an additional dense layer ($n = 15$) attached to the original bottleneck layer and sigmoid activation (for binary classification) or a linear layer with output dimension equal to that of the classification label followed by softmax activation (for $n > 2$ multiclass prediction) [15]. This additional dense layer was added to facilitate extraction of bottleneck features in a reduced dimension. A DenseNet architecture with weights pretrained to ImageNet was chosen to facilitate comparison with recent work on pneumonia detection in radiographs by Rajpurkar and colleagues and for its state-of-the-art results on standard computer vision datasets [2]. All models were trained using a cross-entropy loss function with parameter update by stochastic gradient descent with momentum, initial learning rate 0.01, momentum 0.9, and weight decay 0.0001. Learning rate was decayed by a factor of 10 after each epoch with no improvement in validation loss, and training was stopped after three epochs with no improvement in validation loss.

Internal and external performance testing

To assess how individual models trained using single datasets would generalize compared to a model trained simultaneously on multiple datasets, we trained CNNs to predict nine overlapping diagnoses (cardiomegaly, emphysema, effusion, hernia, nodule, atelectasis, pneumonia, edema, and consolidation) using three different train set combinations: NIH, MSH, and a joint NIH–MSH train set. We were interested only in the prediction of pneumonia and included other diagnoses to improve overall model training and performance. For each model, we calculated AUC, accuracy, sensitivity, and specificity for four different test sets: joint NIH–MSH, NIH only, MSH only, and IU. We report differences in test AUC for all possible internal–external comparisons. We consider the joint MSH–NIH test set the internal comparison set for the jointly trained model. We additionally report differences in test AUC between a jointly trained MSH–NIH model and individual MSH–NIH test sets. The classification threshold was set to ensure 95% sensitivity on each test set to simulate model use for a theoretical screening task. After external review of this analysis, a trivial model that ranked cases based only on the average pneumonia prevalence in each hospital system’s training data and completely ignored radiographic findings was evaluated on the MSH–NIH test set to evaluate how hospital systems alone can predict pneumonia in the joint dataset. We include calibration plots across all reported comparisons as supporting information.

Hospital system and department prediction

After training models for pneumonia and evaluating their performance across sites, additional analysis was planned to better understand a CNN’s ability to detect site and department and how that could affect pneumonia prediction. We trained a CNN to predict hospital systems from radiographs to assess whether location information was directly detectable from the radiograph alone. Radiographs from all three hospital systems were utilized to learn a model that could identify the hospital system from which a given radiograph was drawn. To develop this concept more granularly, for MSH radiographs, we further identified from which hospital unit individual radiographs were obtained (inpatient wards, emergency department). In all cases, we report the classification accuracy on a held-out test set.

Sample activation maps

We created 7×7 sample activation maps, following Zhou and colleagues, to attempt to understand which locations in chest radiographs provided strong evidence for hospital system [19]. For this experiment, we specifically identify radiographs from the NIH. For a sample of NIH test radiographs ($n = 100$), we averaged the softmax probability for each subregion calculated as

$$P(\text{hospital} = \text{NIH} | \text{radiograph}_{i,j}) = \frac{e^{Y_{i,j} \text{ NIH}}}{e^{Y_{i,j} \text{ NIH}} + e^{Y_{i,j} \text{ MSH}} + e^{Y_{i,j} \text{ IU}}}, \tag{1}$$

where i,j corresponds to the subregion at the i th row and j th column of the final convolutional layer ($7 \times 7 = 49$ subregions), where each

$$Y_{i,j} \text{ Hospital System} = \sum (B_k \text{ Hospital System} * X_{k,i,j}) + B_0 \text{ Hospital System}, \tag{2}$$

where the sum is performed over the K final convolutional layers, and $X_{k,i,j}$ represents the activation at the i th row and j th column of the k th final convolutional layer. To characterize how many different subregions were typically involved in NIH hospital system classification, we report the mean, minimum, and maximum number of subregions that predicted NIH

decisively (probability $\geq 95\%$). To illustrate the contribution of particularly influential features (e.g., laterality labels) to classification, we present several examples of heatmaps generated by calculating $Y_{i,j} \text{ NIH} - Y_{i,j} \text{ MSH} - Y_{i,j} \text{ IU}$ for all i,j subregions in an image and subtracting the mean. This additional calculation was necessary to distinguish their positive contribution in the context of many subregions contributing positively to classification probability.

Engineered relative risk experiment

We wished to assess the hypothesis that a CNN could learn to exploit large differences in pathology prevalence between two hospital systems in training data by calibrating its predictions to the baseline prevalence at each hospital system rather than exclusively discriminating based on direct pathology findings. This would lead to strong performance on a test dataset consisting of imbalanced data from both hospital systems but would fail to generalize to data from an external hospital system. To test this hypothesis, we simulated experimental cohorts that differed only in relative disease prevalence and performed internal and external evaluations as described above. Five cohorts of 20,000 patients consisting of 10,000 MSH and 10,000 NIH patients were specifically sampled to artificially set different levels of pneumonia prevalence in each population while maintaining a constant overall prevalence: NIH Severe (NIH 9.9%, MSH 0.1%), NIH Mild (NIH 9%, MSH 1%), Balanced (NIH 5%, MSH 5%), MSH Mild (MSH 9%, NIH 1%), and MSH Severe (MSH 9.9%, NIH 0.1%). The sampling routine also ensured that males and females had an equal prevalence of pneumonia. We refer to these as “engineered prevalence cohorts.” Train, tune, and test splits consistent with prior modeling were maintained for these experiments. CNNs were trained on each cohort in the fashion previously described, and test AUCs on internal joint MSH–NIH and external IU data were compared.

Statistical methods

To assess AUC differences between classification models, we used either the paired or unpaired version of DeLong’s test for ROC curves as appropriate [20]. Comparisons between proportions were performed utilizing χ^2 tests, and proportion CIs were calculated using the Clopper–Pearson interval. All P values were assessed at an alpha of 0.05. Statistical analysis was performed using R version 3.4 with the pROC package and scikit-learn 0.18.1 [21,22].

Results

Datasets

The average age of patients in the MSH cohort was 63.2 years (SD 16.5 years), compared to 49.6 years (SD 17 years) in the IU cohort and 46.9 years (SD 16.6 years) in the NIH cohort (Table 1).

Positive cases of pneumonia were remarkably more prevalent in MSH data (34.2%) than in either NIH (1.2%, $\chi^2 P < 0.001$) or IU (1.0%, $P < 0.001$) data.

Internal and external performance testing

Overall, the internal performance of pneumonia detection CNNs significantly exceeded external performance in 3 out of 5 natural comparisons (Fig 1, Table 2). CNNs trained to detect pneumonia at NIH had internal test AUC 0.750 (95% CI 0.721–0.778), significantly worse external test AUC 0.695 at MSH (95% CI 0.683–0.706, $P < 0.001$), and comparable external test AUC 0.725 at IU (95% CI 0.644–0.807, $P = 0.580$). CNNs trained to detect pneumonia at MSH had internal test AUC 0.802 (95% CI 0.793–0.812), significantly worse external test AUC

Table 1. Baseline characteristics of datasets by site.

Characteristic	IU	MSH	NIH
Patient demographics			
No. patient radiographs	3,807	42,396	112,120
No. patients	3,683	12,904	30,805
Age, mean (SD), years	49.6 (17.0)	63.2 (16.5)	46.9 (16.6)
No. females (%)	643 (57.3%)	18,993 (44.8%)	48,780 (43.5%)
Image diagnosis frequencies			
Pneumonia, No. (%)	39 (1.0%)	14,515 (34.2%)	1,353 (1.2%)
Emphysema, No. (%)	62 (1.6%)	1,308 (3.1%)	2,516 (2.2%)
Effusion, No. (%)	142 (3.7%)	19,536 (46.1%)	13,307 (11.9%)
Consolidation, No. (%)	26 (0.7%)	25,318 (59.7%)	4,667 (4.2%)
Nodule, No. (%)	104 (2.7%)	569 (1.3%)	6,323 (5.6%)
Atelectasis, No. (%)	307 (8.1%)	16,713 (39.4%)	11,535 (10.3%)
Edema, No. (%)	45 (1.2%)	7,144 (16.9%)	2,303 (2.1%)
Cardiomegaly, No. (%)	328 (8.6%)	14,285 (33.7%)	2,772 (2.5%)
Hernia, No. (%)	46 (1.2%)	228 (0.5%)	227 (0.2%)

*Sex data available for 1,122 / 3,807 IU, 42,383 / 42,396 MSH; age data available for 112,077 / 112,120 NIH.

Abbreviations: IU, Indiana University Network for Patient Care; MSH, Mount Sinai Hospital; NIH, National Institutes of Health Clinical Center; No., number.

<https://doi.org/10.1371/journal.pmed.1002683.t001>

0.717 at NIH (95% CI 0.687–0.746, $P < 0.001$), and comparable external test AUC 0.756 at IU (95% CI 0.674–0.838, $P = 0.273$). A jointly trained MSH–NIH model had internal test AUC 0.931 (95% CI 0.927–0.936), significantly greater than external test AUC 0.815 at Indiana (95% CI 0.745–0.885, $P = 0.001$). The jointly trained model had stronger internal performance compared to either constituent site individually (MSH AUC 0.805, 95% CI 0.796–0.814, $P < 0.001$; NIH AUC 0.733, 95% CI 0.703–0.762, $P < 0.001$) (Table 2). A trivial model that ranked cases based only on the average pneumonia prevalence in each hospital system achieved AUC 0.861 (95% CI 0.855–0.866) on the joint MSH–NIH test set. Calibration plots comparing model performance across hospital systems demonstrated a wide range of calibration slopes (minimum 0.047 for train MSH and predict NIH, maximum 10.4 for train NIH and predict MSH; S3–S14 Figs).

Hospital system and department prediction

A CNN trained to identify hospital systems accurately identified 22,050 / 22,062 (99.95%, 95% CI 0.9991–0.9997) of NIH, 8,386 / 8,388 (99.98%, 95% CI 0.9991–1.0000) of MSH, and 737 / 771 (95.59%, 95% CI 0.9389–0.9693) of IU test radiographs, and hospital system was a larger source of variation in image data principal components than pneumonia (S15 Fig). To identify radiographs originating at a specific hospital system, such as NIH, CNNs used features from many different image regions (Fig 2A); the majority of image subregions were individually able to predict the hospital system with $\geq 95\%$ certainty (35.7 / 49, 72.9%, minimum 21, maximum 49, $N = 100$ NIH radiographs). Laterality labels were particularly influential (Fig 2B and 2C).

A CNN trained to identify individual departments within MSH accurately identified 5,805 / 5,805 (100%, 95% CI 0.9993–1.0000) of inpatient radiographs and 449 / 449 (100%, 95% CI 0.9918–1.0000) of emergency department radiographs. Patients who received portable

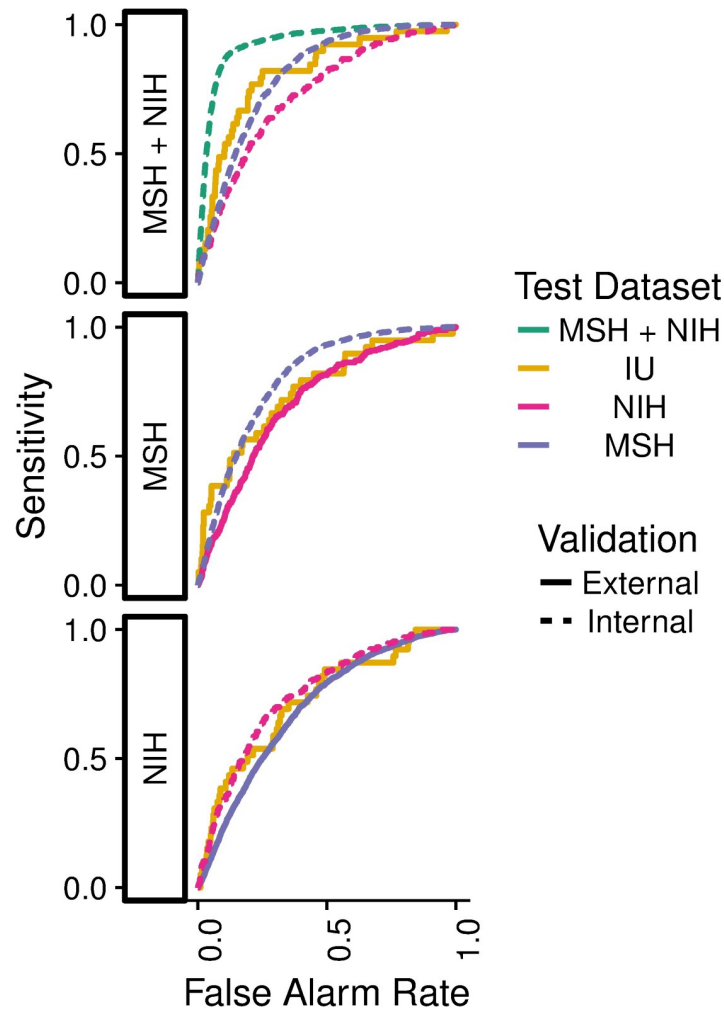


Fig 1. Pneumonia models evaluated on internal and external test sets. A model trained using both MSH and NIH data (MSH + NIH) had higher performance on the combined MSH + NIH test set than on either subset individually or on fully external IU data. IU, Indiana University Network for Patient Care; MSH, Mount Sinai Hospital; NIH, National Institutes of Health Clinical Center.

<https://doi.org/10.1371/journal.pmed.1002683.g001>

radiographs on an inpatient floor had a higher prevalence of pneumonia than those in the emergency department (41.1% versus 32.8%, respectively, $P < 0.001$).

Engineered relative risk experiment

Artificially increasing the difference in the prevalence of pneumonia between MSH and NIH led to CNNs that performed increasingly well on internal testing but not external testing (S2 Table). CNNs trained on engineered prevalence cohorts of NIH and MSH data showed stronger internal AUC on a joint NIH–MSH test set when the prevalence of pneumonia was imbalanced between the two hospital systems in the training dataset with MSH Severe AUC 0.899 (95% CI 0.885–0.914, $P < 0.001$), MSH Mild AUC 0.860 (95% CI 0.839–0.882, $P < 0.001$), NIH Mild AUC 0.807 (95% CI 0.778–0.836, $P = 0.002$), and NIH Severe AUC 0.849 (95% CI 0.826–0.871, $P < 0.001$) than when it was balanced with AUC 0.739 (95% CI 0.707–0.772) (Fig 3A and 3B).

Table 2. Internal and external pneumonia screening performance for all train, tune, and test hospital system combinations. Parentheses show 95% CIs.

Train/ Tune Site	Comparison Type*	Test Site (Images)	AUC	Accuracy	Sensitivity	Specificity	PPV	NPV
NIH	Internal	NIH (N = 22,062)	0.750 (0.721–0.778)	0.255 (0.250–0.261)	0.951 (0.917–0.973)	0.247 (0.241–0.253)	0.015 (0.013–0.017)	0.998 (0.996–0.999)
	External	MSH (N = 8,388)	0.695 (0.683–0.706)	0.476 (0.465–0.486)	0.950 (0.942–0.958)	0.212 (0.201–0.223)	0.401 (0.390–0.413)	0.884 (0.866–0.901)
	External	IU (N = 3,807)	0.725 (0.644–0.807)	0.190 (0.178–0.203)	0.974 (0.865–0.999)	0.182 (0.170–0.195)	0.012 (0.009–0.017)	0.999 (0.992–1.000)
	Superset *	MSH + NIH (N = 30,450)	0.773 (0.766–0.780)	0.462 (0.456–0.467)	0.950 (0.942–0.957)	0.403 (0.397–0.409)	0.160 (0.155–0.166)	0.985 (0.983–0.987)
	Superset *	MSH + NIH + IU (N = 34,257)	0.787 (0.780–0.793)	0.470 (0.464–0.475)	0.950 (0.942–0.957)	0.418 (0.413–0.424)	0.148 (0.144–0.153)	0.987 (0.985–0.989)
MSH	Internal	MSH (N = 8,388)	0.802 (0.793–0.812)	0.617 (0.607–0.628)	0.950 (0.942–0.958)	0.432 (0.419–0.446)	0.482 (0.469–0.495)	0.94 (0.930–0.949)
	External	NIH (N = 22,062)	0.717 (0.687–0.746)	0.184 (0.179–0.190)	0.951 (0.917–0.973)	0.175 (0.170–0.18)	0.014 (0.012–0.016)	0.997 (0.994–0.998)
	External	IU (N = 3,807)	0.756 (0.674–0.838)	0.099 (0.089–0.109)	0.974 (0.865–0.999)	0.090 (0.081–0.099)	0.011 (0.008–0.015)	0.997 (0.984–1.000)
	Superset *	MSH + NIH (N = 30,450)	0.862 (0.856–0.868)	0.562 (0.557–0.568)	0.950 (0.942–0.957)	0.516 (0.510–0.522)	0.19 (0.184–0.197)	0.989 (0.987–0.990)
	Superset *	MSH + NIH + IU (N = 34,257)	0.871 (0.865–0.877)	0.577 (0.572–0.582)	0.950 (0.942–0.957)	0.537 (0.532–0.543)	0.180 (0.174–0.185)	0.990 (0.989–0.992)
MSH + NIH	Internal	MSH + NIH (N = 30,450)	0.931 (0.927–0.936)	0.732 (0.727–0.737)	0.950 (0.942–0.957)	0.706 (0.700–0.711)	0.279 (0.271–0.288)	0.992 (0.990–0.993)
	Subset †	NIH (N = 22,062)	0.733 (0.703–0.762)	0.243 (0.237–0.249)	0.951 (0.917–0.973)	0.234 (0.229–0.240)	0.015 (0.013–0.017)	0.997 (0.996–0.999)
	Subset †	MSH (N = 8,388)	0.805 (0.796–0.814)	0.630 (0.619–0.640)	0.950 (0.942–0.958)	0.451 (0.438–0.465)	0.491 (0.478–0.504)	0.942 (0.933–0.951)
	External	IU (N = 3,807)	0.815 (0.745–0.885)	0.238 (0.224–0.252)	0.974 (0.865–0.999)	0.230 (0.217–0.244)	0.013 (0.009–0.018)	0.999 (0.994–1.000)
	Superset *	MSH + NIH + IU (N = 34,257)	0.934 (0.929–0.938)	0.732 (0.727–0.737)	0.95 (0.942–0.957)	0.709 (0.703–0.714)	0.258 (0.250–0.266)	0.993 (0.991–0.994)

*Superset = a test dataset containing data from the same distribution (hospital system) as the training data as well as external data.

† Subset = a test dataset containing data from fewer distributions (hospital systems) than the training data.

Abbreviations: AUC, area under the receiver operating characteristic curve; IU, Indiana University Network for Patient Care; MSH, Mount Sinai Hospital; NIH, National Institutes of Health Clinical Center; NPV, negative predictive value; PPV, positive predictive value.

<https://doi.org/10.1371/journal.pmed.1002683.t002>

Internal MSH–NIH performance of all models trained on imbalanced cohorts was significantly better than their corresponding external performance on IU (external MSH Severe AUC 0.641, 95% CI 0.552–0.730, $P < 0.001$; MSH Mild AUC 0.650, 95% CI 0.548–0.752, $P < 0.001$; NIH Mild AUC 0.703, 95% CI 0.616–0.790, $P = 0.027$; NIH Severe AUC 0.683, 95% CI 0.591–0.775, $P < 0.001$). Internal MSH–NIH performance did not significantly exceed external IU performance for Balanced (0.739, 95% CI 0.707–0.772 versus 0.732, 95% CI 0.645–0.819, $P = 0.880$).

Discussion

We have demonstrated that pneumonia-screening CNNs trained on data from individual or multiple hospital systems did not consistently generalize to external sites, nor did they make predictions exclusively based on underlying pathology. Given the significant interest in using deep learning to analyze radiological imaging, our findings should give pause to those considering rapid deployment of such systems without first assessing their performance in a variety

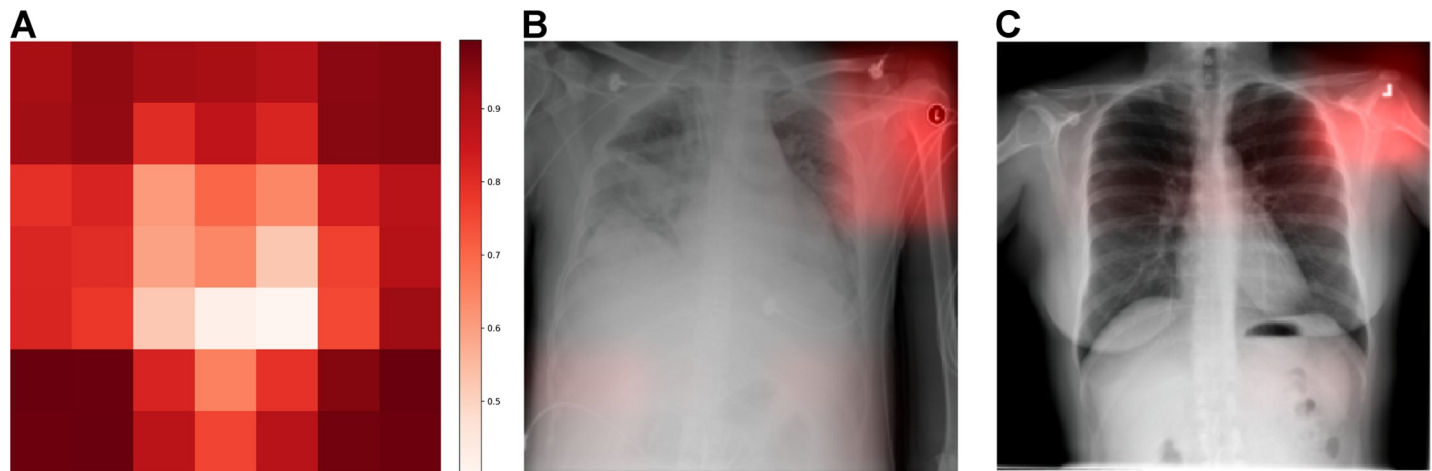


Fig 2. CNN to predict hospital system detects both general and specific image features. (A) We obtained activation heatmaps from our trained model and averaged over a sample of images to reveal which subregions tended to contribute to a hospital system classification decision. Many different subregions strongly predicted the correct hospital system, with especially strong contributions from image corners. (B-C) On individual images, which have been normalized to highlight only the most influential regions and not all those that contributed to a positive classification, we note that the CNN has learned to detect a metal token that radiology technicians place on the patient in the corner of the image field of view at the time they capture the image. When these strong features are correlated with disease prevalence, models can leverage them to indirectly predict disease. CNN, convolutional neural network.

<https://doi.org/10.1371/journal.pmed.1002683.g002>

of real-world clinical settings. To our knowledge, no prior studies have assessed whether radiological CNNs generalized to external datasets. We note that the issue of not generalizing externally is distinct from typical train/test performance degradation, in which overfitting to training data leads to lower performance on testing data: in our experiments, all results are reported on held-out test data exclusively in both internal and external comparisons. Performance of the jointly trained MSH–NIH model on the joint test set (AUC 0.931) was higher than performance on either individual dataset (AUC 0.805 and 0.733, respectively), likely because the model was able to calibrate to different prevalences across hospital systems in the joint test set but not individual test sets. A simple calibration-based non-CNN model that used hospital system pneumonia prevalence only to make predictions and ignored image features achieved AUC 0.861 because of the large difference in pneumonia prevalence between the MSH and NIH test sets. Calibration plots confirmed that a model trained on NIH data was poorly calibrated to MSH and vice versa.

By engineering cohorts of varying prevalence, we demonstrated that the more predictive a hospital system was of pneumonia, the more it was exploited to make predictions, which led to poor generalization on external datasets. We noted that metallic tokens indicating laterality often appeared in radiographs in a site-specific way, which made site identification trivial. However, CNNs did not require this indicator: most image subregions contained features indicative of a radiograph’s origin. These results suggest that CNNs could rely on subtle differences in acquisition protocol, image processing, or distribution pipeline (e.g., image compression) and overlook pathology. Radiological imaging of the same modality from different hospital systems can have markedly different CNN image feature distributions, as we highlighted in S3 Fig. Many of these features encode hospital processes, and many (e.g., site, scanner) may be associated with the prevalence of disease. This can lead to strong internal performance that is not realized on data from new sites. Even in the absence of recognized confounders, we would caution, following Recht and colleagues, that “current accuracy numbers are brittle and susceptible to even minute natural variations in the data distribution” [8].

A difficulty of using deep learning models in medicine is that they use a massive number of parameters, making it difficult to identify the specific variables driving predictions and

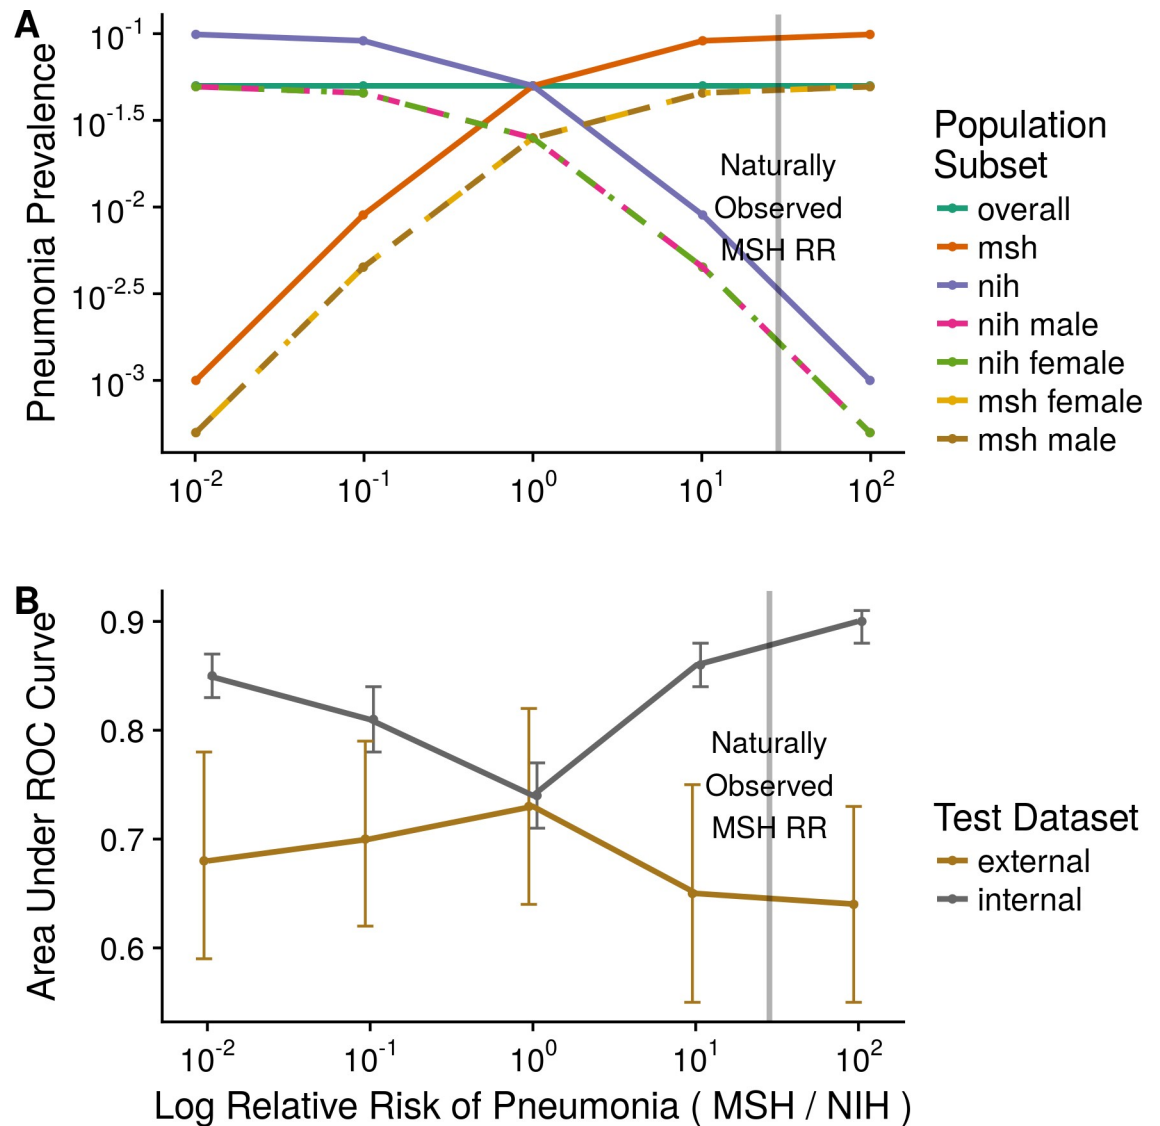


Fig 3. Assessing how prevalence differences in aggregated datasets encouraged confounder exploitation. (A) Five cohorts of 20,000 patients were systematically subsampled to differ only in relative pneumonia risk based on the clinical training data sites. Model performance was assessed on test data from the internal hospital systems (MSH, NIH) and from an external hospital system (IU). (B) Although models perform better in internal testing in the presence of extreme prevalence differences, this benefit is not seen when applied to data from new hospital systems. The natural relative risk of disease at MSH, indicated by a vertical line, is quite imbalanced. IU, Indiana University Network for Patient Care; MSH, Mount Sinai Hospital; NIH, National Institutes of Health Clinical Center; ROC, receiver operating characteristic; RR, relative risk.

<https://doi.org/10.1371/journal.pmed.1002683.g003>

complicating the traditional understanding of model overfitting. Best-practice guidelines within the biomedical literature put a heavy emphasis on identifying all features used and transparently reporting model coefficients to promote the development of reproducible, clinically relevant models [23,24]. This level of transparency is fundamentally incompatible with deep learning. The DenseNet-121 used in our analysis had 6,963,081 parameters, and the methods available to interrogate which features were driving model predictions remain inexact (e.g., heatmaps). With this large number of parameters, we note that the phenomenon of deep learning models failing to generalize externally is distinct from the traditional understanding of overfitting in the statistical literature. If allowed to train indefinitely, a CNN would overfit

in a traditional sense and would fail to generalize both internally and externally. We followed standard practices in training deep learning models to ensure internal generalization between the training data and held-out test data. Specifically, we used a tune set to assess internal generalization performance after every step of optimization, and we stopped training early at the point when the tune performance stopped improving. We note that in this process, each of the train, tune, and test sets were drawn from the same distribution. The ability of a radiological CNN to perform well in external testing on “plausibly related” populations, which may differ significantly in their underlying data distributions, will be critical for the real-world utility and adaptation of such models [24].

Even the development of customized deep learning models that are trained, tuned, and tested with the intent of deploying at a single site are not necessarily a solution that can control for potential confounding variables. At a finer level, we found that CNNs could separate portable radiographs from the inpatient wards and emergency department in MSH data with 100% accuracy and that these patient groups had significantly different prevalences of pneumonia. It was determined after the fact that devices from different manufacturers had been used in the inpatient units (Konica Minolta) and emergency department (Fujifilm), and the latter were stored in the picture archiving and communication system (PACS) in an inverted color scheme (i.e., air appears white) along with distinctive text indicating laterality and use of a portable scanner. While these identifying features were prominent to the model, they only became apparent to us after manual image review. If certain scanners within a hospital are used to evaluate patients with different baseline disease prevalences (e.g., intensive care unit [ICU] versus outpatient), these may confound deep learning models trained on radiological data. Fully external testing—ideally on a collection of data gathered from a varied collection of hospitals—can reveal and account for such sampling biases that may limit the generalizability of a model.

The development of CNN architectures specifically adapted to accommodate radiological imaging is an important step towards building stronger models. Entire high-resolution radiological images are often aggressively downsampled (e.g., to 224×224 pixels) to facilitate “transfer learning,” i.e., fine-tuning preexisting CNN architectures that have been pretrained to ImageNet [15]. While practically convenient, these models are not optimal for the radiological context, as they ignore essential domain information about the problem (e.g., pneumonia is present in the lungs), and the low-resolution images they require eliminate valuable radiographic findings. Both factors can lead to an increased reliance on confounding factors in making predictions. CNN architectures designed specifically to accommodate radiological imaging have demonstrated promising early results, and more work in this area is needed [25–27].

While our analysis found degradation of model performance on external test sets, we note that it is possible for external test set performance to be either better or worse than internal. Many different aspects of dataset construction (e.g., inclusion criteria, labeling procedure) and the underlying clinical data (pathology prevalence and severity, confounding protocolized variables) can affect performance. For example, a model trained on noisily labeled data that included all available imaging might reasonably be expected to have lower internal test performance than if tested externally on a similar dataset manually selected and labeled by a physician as clear examples of pathological and normal cases.

In addition to site-specific confounding variables that threaten generalizability, there are other factors related to medical management that may exist everywhere but undermine the clinical applicability of a model. As has been noted, chest tubes that treat pneumothorax frequently appear in studies positive for pneumothorax in NIH data; a CNN for pneumothorax may learn to detect obvious chest tubes rather than a subtler pneumothorax itself [28]. If such a model were deployed in an emergency department, it might inaccurately negatively diagnose patients presenting with pneumothorax because they lacked a chest tube. A CNN for

pneumonia could potentially learn that radiographs obtained with portable scanners were more likely to contain pneumonia and assign such radiographs a higher probability of disease. Models exploiting such confounding factors that reflect common medical practice may even generalize well but ultimately have limited clinical utility. If CNN-based systems are to be used for medical diagnosis, they must be tailored to carefully considered clinical questions, prospectively tested at a variety of sites in real-world-use scenarios, and carefully assessed to determine how they impact diagnostic accuracy.

There are several limitations to this study. Most notably, without more granular details on the underlying patient populations, we are unable to fully assess what factors might be contributing to the hospital system–specific biasing of the models. The extremely high incidence of pneumonia in the MSH dataset is also a point of concern; however, we attribute this to differences in the underlying patient populations and variability in classification thresholds for pathology. First, a majority of MSH radiographs were portable inpatient scans, ordered for patients too unstable to travel to the radiology department for a standard radiograph. In contrast, all IU radiographs were outpatient. While the inpatient/outpatient mix from NIH is not reported, we believe it likely contains a substantial outpatient percentage given that the incidence of pneumonia is similar to IU. Second, our NLP approach for MSH assigned positive ground truth labels more liberally than NIH or IU, marking a study as positive for pathology when a radiologist explicitly commented on it as a possibility in a report, indicating that the radiographic appearance was consistent with the finding. Different radiologists may have different thresholds at which they explicitly include a possible diagnosis in their reports. Researchers working in this area will continually have to make decisions about their classification threshold for labeling a study positive or negative. We believe that either of these two factors can drive large differences in prevalences of pathology across datasets, and this variation can confound diagnostic CNNs.

An additional limitation was that radiologic diagnoses are made in the context of a patient's history and clinical presentation, something not incorporated into our approach. Positive findings on chest radiograph are necessary but not sufficient for the diagnosis of pneumonia, which is only made when the patient also exhibits a “constellation of suggestive clinical features” [29]. Finally, the relatively small size and low number of pneumonia cases in IU data led to wide CIs in IU test AUC and may have limited our ability to detect external performance degradation in some cases. Nevertheless, many key comparisons achieved statistical significance with even this smaller external dataset.

Conclusion

Pneumonia-screening CNNs achieved better internal than external performance in 3 out of 5 natural comparisons. When models were trained on pooled data from sites with different pneumonia prevalence, they performed better on new pooled data from these sites but not on external data. CNNs robustly identified hospital system and department within a hospital, which can have large differences in disease burden and may confound predictions.

Supporting information

S1 Table. Performance of NLP algorithm on 30% test data. NLP, natural language processing. (XLSX)

S2 Table. Internal and external pneumonia screening performance with MSH–NIH training cohorts with varying engineered prevalences of pneumonia by hospital system. MSH,

Mount Sinai Hospital; NIH, National Institutes of Health Clinical Center.
(XLSX)

S1 Fig. Preprocessing diagram.
(TIF)

S2 Fig. STARD cohort splitting diagram.
(TIF)

S3 Fig. Calibration plot (train MSH, predict IU). IU, Indiana University Network for Patient Care; MSH, Mount Sinai Hospital.
(TIF)

S4 Fig. Calibration plot (train MSH, predict MSH). MSH, Mount Sinai Hospital.
(TIF)

S5 Fig. Calibration plot (train MSH, predict MSH–NIH). MSH, Mount Sinai Hospital; NIH, National Institutes of Health Clinical Center.
(TIF)

S6 Fig. Calibration plot (train MSH, predict NIH). MSH, Mount Sinai Hospital; NIH, National Institutes of Health Clinical Center.
(TIF)

S7 Fig. Calibration plot (train MSH–NIH, predict IU). IU, Indiana University Network for Patient Care; MSH, Mount Sinai Hospital; NIH, National Institutes of Health Clinical Center.
(TIF)

S8 Fig. Calibration plot (train MSH–NIH, predict MSH). MSH, Mount Sinai Hospital; NIH, National Institutes of Health Clinical Center.
(TIF)

S9 Fig. Calibration plot (train MSH–NIH, predict MSH–NIH). MSH, Mount Sinai Hospital; NIH, National Institutes of Health Clinical Center.
(TIF)

S10 Fig. Calibration plot (train MSH–NIH, predict NIH). MSH, Mount Sinai Hospital; NIH, National Institutes of Health Clinical Center.
(TIF)

S11 Fig. Calibration plot (train NIH, predict IU). IU, Indiana University Network for Patient Care; NIH, National Institutes of Health Clinical Center.
(TIF)

S12 Fig. Calibration plot (train NIH, predict MSH). MSH, Mount Sinai Hospital; NIH, National Institutes of Health Clinical Center.
(TIF)

S13 Fig. Calibration plot (train NIH, predict MSH–NIH). MSH, Mount Sinai Hospital; NIH, National Institutes of Health Clinical Center.
(TIF)

S14 Fig. Calibration plot (train NIH, predict NIH). NIH, National Institutes of Health Clinical Center.
(TIF)

S15 Fig. Distribution of images' embedded principal components. Principal component analysis was performed on each image's bottleneck features and colored to reveal the distribution of hospital systems and pneumonia. MSH images are largely separable from radiographs acquired at IU and NIH. IU, Indiana University Network for Patient Care; MSH, Mount Sinai Hospital; NIH, National Institutes of Health Clinical Center. (TIF)

Author Contributions

Conceptualization: John R. Zech, Marcus A. Badgeley, Eric Karl Oermann.

Data curation: John R. Zech, Marcus A. Badgeley, Anthony B. Costa, Joseph J. Titano, Eric Karl Oermann.

Formal analysis: John R. Zech, Marcus A. Badgeley, Manway Liu.

Investigation: John R. Zech, Marcus A. Badgeley.

Methodology: John R. Zech, Marcus A. Badgeley, Eric Karl Oermann.

Project administration: Anthony B. Costa, Joseph J. Titano.

Resources: Anthony B. Costa.

Software: John R. Zech, Marcus A. Badgeley.

Supervision: Eric Karl Oermann.

Validation: John R. Zech.

Visualization: John R. Zech, Marcus A. Badgeley.

Writing – original draft: John R. Zech, Marcus A. Badgeley, Eric Karl Oermann.

Writing – review & editing: John R. Zech, Marcus A. Badgeley, Manway Liu, Anthony B. Costa, Joseph J. Titano, Eric Karl Oermann.

References

1. Wang X, Peng Y, Lu L, Lu Z, Bagheri M, Summers RM. ChestX-ray8: Hospital-scale Chest X-ray Database and Benchmarks on Weakly-Supervised Classification and Localization of Common Thorax Diseases; 2017 [cited 1 July 2018]. Preprint. Available from: <https://arxiv.org/abs/1705.02315>.
2. Rajpurkar P, Irvin J, Zhu K, Yang B, Mehta H, Duan T, et al. CheXNet: Radiologist-Level Pneumonia Detection on Chest X-Rays with Deep Learning; 2017 [cited 1 July 2018]. Preprint. Available from: <https://arxiv.org/abs/1711.05225>.
3. Gulshan V, Peng L, Coram M, Stumpe MC, Wu D, Narayanaswamy A, et al. Development and Validation of a Deep Learning Algorithm for Detection of Diabetic Retinopathy in Retinal Fundus Photographs. *JAMA*. 2016; 316: 2402–2410. <https://doi.org/10.1001/jama.2016.17216> PMID: 27898976
4. Ting DSW, Cheung CY-L, Lim G, Tan GSW, Quang ND, Gan A, 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*. 2017; 318: 2211–2223. <https://doi.org/10.1001/jama.2017.18152> PMID: 29234807
5. Kermany DS, Goldbaum M, Cai W, Valentim CCS, Liang H, Baxter SL, et al. Identifying Medical Diagnoses and Treatable Diseases by Image-Based Deep Learning. *Cell*. Elsevier; 2018; 172: 1122–1131. <https://doi.org/10.1016/j.cell.2018.02.010> PMID: 29474911
6. Russakovsky O, Deng J, Su H, Krause J, Satheesh S, Ma S, et al. ImageNet Large Scale Visual Recognition Challenge; 2014 [cited 1 July 2018]. Preprint. Available from: <https://arxiv.org/abs/1409.0575>.
7. Lecun Y, Bottou L, Bengio Y, Haffner P. Gradient-based learning applied to document recognition. In *Intelligent signal processing*. IEEE Press. 2001. p. 306–351.

8. Recht B, Roelofs R, Schmidt L, Shankar V. Do CIFAR-10 Classifiers Generalize to CIFAR-100?; 2018 [cited 1 July 2018]. Preprint. Available from: <https://arxiv.org/abs/1806.00451>.
9. Rothwell PM. External validity of randomised controlled trials: “To whom do the results of this trial apply?” *Lancet*. 2005; 365: 82–93. [https://doi.org/10.1016/S0140-6736\(04\)17670-8](https://doi.org/10.1016/S0140-6736(04)17670-8) PMID: 15639683
10. Pandis N, Chung B, Scherer RW, Elbourne D, Altman DG. CONSORT 2010 statement: extension checklist for reporting within person randomised trials. *BMJ*. 2017; 357: j2835. <https://doi.org/10.1136/bmj.j2835> PMID: 28667088
11. Cabitza F, Rasoini R, Gensini GF. Unintended Consequences of Machine Learning in Medicine. *JAMA*. 2017; 318: 517–518. <https://doi.org/10.1001/jama.2017.7797> PMID: 28727867
12. Demner-Fushman D, Kohli MD, Rosenman MB, Shooshan SE, Rodriguez L, Antani S, et al. Preparing a collection of radiology examinations for distribution and retrieval. *J Am Med Inform Assoc*. 2016; 23: 304–310. <https://doi.org/10.1093/jamia/ocv080> PMID: 26133894
13. LeCun Y, Bengio Y, Hinton G. Deep learning. *Nature*. 2015; 521: 436–444. <https://doi.org/10.1038/nature14539> PMID: 26017442
14. He K, Zhang X, Ren S, Sun J. Deep Residual Learning for Image Recognition; 2015 [cited 1 July 2018]. Preprint. Available from: <https://arxiv.org/abs/1512.03385>.
15. Huang G, Liu Z, Weinberger KQ, van der Maaten L. Densely Connected Convolutional Networks; 2016 [cited 1 July 2018]. Preprint. Available from: arXiv: <https://arxiv.org/abs/1608.06993>.
16. Zhang C, Bengio S, Hardt M, Recht B, Vinyals O. Understanding deep learning requires rethinking generalization; 2016 [cited 1 July 2018]. Preprint. Available from: <https://arxiv.org/abs/1611.03530>.
17. Zech J, Pain M, Titano J, Badgeley M, Schefflein J, Su A, et al. Natural Language-based Machine Learning Models for the Annotation of Clinical Radiology Reports. *Radiology*. 2018; 287(2): 570–580. <https://doi.org/10.1148/radiol.2018171093> PMID: 29381109
18. Paszke A, Gross S, Chintala S, Chanan G, Yang E, DeVito Z, et al. Automatic differentiation in PyTorch; 2017 [cited 1 July 2018]. Preprint. Available from: <https://openreview.net/forum?id=BJJsrMfCZ>.
19. Zhou B, Khosla A, Lapedriza A, Oliva A, Torralba A. Learning Deep Features for Discriminative Localization; 2015. Preprint. Available from: <https://arxiv.org/abs/1512.04150>.
20. DeLong ER, DeLong DM, Clarke-Pearson DL. Comparing the areas under two or more correlated receiver operating characteristic curves: a nonparametric approach. *Biometrics*. 1988; 44: 837–845. PMID: 3203132
21. Robin X, Turck N, Hainard A, Tiberti N, Lisacek F, Sanchez J-C, et al. pROC: an open-source package for R and S+ to analyze and compare ROC curves. *BMC Bioinformatics*. 2011; 12: 77. <https://doi.org/10.1186/1471-2105-12-77> PMID: 21414208
22. Pedregosa F, Varoquaux G, Gramfort A, Michel V, Thirion B, Grisel O, et al. Scikit-learn: Machine Learning in Python. *J Mach Learn Res*. 2011; 12: 2825–2830.
23. Collins GS; Reitsma JB, Altman DG, Moons KGM. Transparent Reporting of a multivariable prediction model for Individual Prognosis Or Diagnosis (TRIPOD): The TRIPOD Statement. *Ann Intern Med* 2015; 162:55–63. <https://doi.org/10.7326/M14-0697> PMID: 25560714
24. Steyerberg EW and Vergouwe Y. Towards better clinical prediction models: seven steps for development and an ABCD for validation. *European Heart Journal* 2014; 35:1925–1931. <https://doi.org/10.1093/eurheartj/ehu207> PMID: 24898551
25. Gale W, Oakden-Rayner L, Carneiro G, Bradley AP, Palmer LJ. Detecting hip fractures with radiologist-level performance using deep neural networks; 2017 [cited 1 July 2018]. Preprint. Available from: <https://arxiv.org/abs/1711.06504>.
26. Mutasa S, Chang PD, Ruzal-Shapiro C, Ayyala R. MABAL: a Novel Deep-Learning Architecture for Machine-Assisted Bone Age Labeling. *J Digit Imaging*. 2018; 31(4): 513–519. <https://doi.org/10.1007/s10278-018-0053-3> PMID: 29404850
27. Geras KJ, Wolfson S, Gene Kim S, Moy L, Cho K. High-Resolution Breast Cancer Screening with Multi-View Deep Convolutional Neural Networks; 2017 [cited 1 July 2018]. Preprint. Available from: <https://arxiv.org/abs/1703.07047>.
28. Oakden-Rayner L. Exploring the ChestXray14 dataset: problems. 18 Dec 2017 [cited 26 Jan 2018]. In: Luke Oakden-Rayner [Internet]. Available from: <https://lukeoakdenrayner.wordpress.com/2017/12/18/the-chestxray14-dataset-problems/>.
29. Mandell LA, Wunderink RG, Anzueto A, Bartlett JG, Campbell GD, Dean NC, et al. Infectious Diseases Society of America/American Thoracic Society consensus guidelines on the management of community-acquired pneumonia in adults. *Clin Infect Dis*. 2007; 44 Suppl 2: S27–72.