

Resolution scaling governs DINOv3 transfer performance in chest radiograph classification

Soroosh Tayebi Arasteh (1,2,3,4), Mina Shaigan (5), Christiane Kuhl (2), Jakob Nikolas Kather (6,7,8), Sven Nebelung* (1,2), Daniel Truhn* (1,2)

- (1) Lab for AI in Medicine, Department of Diagnostic and Interventional Radiology, University Hospital RWTH Aachen, Aachen, Germany.
- (2) Department of Diagnostic and Interventional Radiology, University Hospital RWTH Aachen, Aachen, Germany.
- (3) Department of Urology, Stanford University, Stanford, CA, USA.
- (4) Department of Radiology, Stanford University, Stanford, CA, USA.
- (5) Institute for Computational Genomics, Joint Research Center for Computational Biomedicine, University Hospital RWTH Aachen, Aachen, Germany.
- (6) Else Kroener Fresenius Center for Digital Health, Technical University Dresden, Dresden, Germany.
- (7) Department of Medicine I, University Hospital Dresden, Dresden, Germany.
- (8) National Center for Tumor Diseases (NCT), University Hospital Heidelberg, Heidelberg, Germany.

* Sven Nebelung and Daniel Truhn are shared senior authors.

Abstract

Self-supervised learning (SSL) has advanced visual representation learning, but its value in chest radiography, a high-volume imaging modality with fine-grained findings, remains unclear. Meta’s DINOv3 extends earlier SSL models through Gram-anchored self-distillation. Whether these design choices improve transfer learning for chest radiography has not been systematically tested. We benchmarked DINOv3 against DINOv2 and ImageNet initialization across seven datasets ($n > 814,000$). Two representative backbones were evaluated: ViT-B/16 and ConvNeXt-B. Images were analyzed at 224×224 , 512×512 , and 1024×1024 pixels. We additionally assessed frozen features from a 7B model. The primary outcome was mean AUROC across labels. At 224×224 , DINOv3 and DINOv2 achieved comparable performance on adult datasets. Increasing resolution to 512×512 yielded consistent improvements for DINOv3 over both DINOv2 and ImageNet. In contrast, results in pediatric cohort showed no differences across initializations. Across all settings, ConvNeXt-B outperformed ViT-B/16. Models using frozen DINOv3-7B features underperformed relative to fully finetuned 86-89M-parameter backbones, highlighting the importance of domain adaptation. Scaling to 1024×1024 did not further improve accuracy. Resolution-related gains were most evident for boundary-dependent and small focal abnormalities. In chest radiography, higher input resolution is critical for leveraging the benefits

of modern self-supervised models. 512×512 pixels represent a practical upper limit where DINOv3-initialized ConvNeXt-B networks provide the strongest performance, while larger inputs offer minimal return on cost. Clinically, these findings support use of finetuned, mid-sized backbones at 512×512 for chest radiograph interpretation, with the greatest gains expected in detecting subtle or boundary-centered lesions relevant to emergency and critical care settings.

Correspondence

Soroosh Tayebi Arasteh, Dr.-Ing., Dr. rer. medic.
Lab for AI in Medicine
Department of Diagnostic and Interventional Radiology
University Hospital RWTH Aachen
Pauwelsstr. 30
52070 Aachen, Germany
Email: soroosh.arasteh@rwth-aachen.de

Introduction

Chest radiography is the most widely performed imaging examination worldwide and a first-line tool for detecting pulmonary and cardiac abnormalities. Subtle or low-contrast findings, such as interstitial lung disease, reticular changes, or diffuse pulmonary opacification, can be difficult to recognize, motivating the use of automated analysis to assist interpretation and triage. Artificial intelligence (AI) has become an integral component of medical imaging^{1–3}, with chest radiographs serving as one of the most extensively studied modalities for evaluating new algorithms^{4–6}. Early advances relied on supervised deep learning, where models were pretrained on large annotated datasets such as ImageNet^{7,8} and then fine-tuned for radiographic tasks. Although this strategy improved performance compared with training from scratch, it remains constrained by the domain mismatch between natural and medical images and by its dependence on costly manual annotations. Constructing large, expertly labeled radiograph collections continues to be a major bottleneck, motivating the exploration of label-efficient alternatives.

Self-supervised learning (SSL) has emerged as a paradigm to address this challenge. By constructing pretraining objectives that do not depend on manual labels, SSL enables the use of massive unlabeled datasets to learn transferable visual representations^{9,10}. Methods such as MoCo¹¹, SimCLR¹², BYOL¹³, and SwAV¹⁴ have demonstrated strong performance on natural images, and their application to medical imaging has shown promising gains in classification and segmentation tasks. However, most medical studies to date have been limited in scale, typically using tens rather than hundreds of thousands of radiographs, and existing benchmarks leave open questions regarding the robustness and generalizability of SSL for clinical imaging¹⁵.


The introduction of transformer- based¹⁶ architectures has further accelerated progress. Vision transformers (ViTs)¹⁷ and modern convolutional backbones such as ConvNeXt¹⁸ have redefined the state-of-the-art in computer vision, and their transfer to radiology has underscored the value of flexible, high-capacity architectures for medical data. Within this landscape, the DINO¹⁹ family of SSL methods (self-distillation with no labels) has been particularly influential. DINOv2²⁰, pretrained on hundreds of millions of natural images, established itself as a strong general-purpose representation learner. In our prior work²¹, we showed that DINOv2 could not only match but in many cases surpass supervised ImageNet pretraining when transferred to chest radiograph classification. Building on this foundation, Meta recently released DINOv3²², which introduces Gram-anchored self-distillation and explicit high-resolution adaptation. These modifications are designed to preserve fine-grained visual information during long training schedules and improve scaling to larger input sizes, addressing precisely the resolution constraints that often limit medical imaging models. Yet, despite these architectural advances, it remains unknown whether DINOv3’s improvements translate to medical imaging tasks. Early studies across different modalities^{23,24} suggest that scaling laws from natural images do not always transfer to medical data²⁵. Our work addresses this open question through a large-scale, systematic evaluation focused on chest radiograph classification, a domain uniquely sensitive to resolution scaling.

Here, we present the first systematic evaluation of DINOv3 for chest radiograph classification across seven datasets comprising more than 814,000 anteroposterior (AP) or posteroanterior (PA) radiographs (see **Figure 1**). Our benchmark spans multiple axes of diversity: two backbone families (the transformer-based ViT-B/16 and the fully convolutional ConvNeXt-B), input resolutions from 224×224 to 1024×1024 pixels (extending beyond prior studies typically limited to $\leq 336 \times 336$ ^{21,26–28}), and a broad label space covering up to 21 distinct imaging findings. These include common abnormalities such as cardiomegaly, pleural effusion, pneumonia, and atelectasis, as well as less frequent but clinically important findings such as pulmonary fibrosis, emphysema, hernia, kyphosis, lung nodules or masses, infiltrates, and fractures. The datasets vary in size from fewer than 10,000 to more than 200,000 radiographs, span multiple continents, and include both adult and pediatric cohorts, providing a robust testbed for generalization. In addition, we evaluate frozen representations from the 7B-parameter DINOv3 teacher model—a vision-only, self-supervised backbone trained without paired text. To our knowledge, this is the first evaluation in radiology of a self-supervised vision encoder used as a frozen feature extractor for chest radiograph classification at the billion-parameter scale, distinct from prior work on vision–language models for report generation^{29,30}. Collectively, these contributions establish a comprehensive benchmark for transferring state-of-the-art SSL to chest radiographs. Our findings reveal a clear resolution–performance relationship: while DINOv2 remains slightly stronger at 224×224 , DINOv3 consistently outperforms both DINOv2 and ImageNet at 512×512 , particularly with ConvNeXt backbones. We further show that frozen features from the 7B-parameter DINOv3 underperform compared with full finetuning of much smaller 86–89M models, underscoring the importance of domain-specific adaptation. Finally, scaling beyond 512×512 yields no measurable advantage despite substantial computational cost, suggesting a practical upper bound for DINOv3 transfer in chest radiography. Together, these results highlight both the promise and current limitations of transferring billion-scale SSL vision models to healthcare and provide actionable guidance for integrating high-resolution SSL into medical imaging pipelines.


Results

We benchmarked ImageNet, DINOv2, and DINOv3 initializations under full finetuning across six publicly available datasets—Pedi-CXR³¹ ($n = 9,125$, 3 labels), VinDr-CXR³² ($n = 18,000$, 14 labels), ChestX-ray14³³ ($n = 112,120$, 14 labels), PadChest³⁴ ($n = 110,525$, 17 labels), CheXpert³⁵ ($n = 157,878$, 10 labels), and MIMIC-CXR⁶ ($n = 213,921$, 10 labels)—as well as one internal dataset, UKA-CXR^{21,36–40} ($n = 193,361$, 6 labels). These cohorts varied in size, label diversity, and population, ranging from a small pediatric cohort to large-scale multi-label adult datasets. Across all, we evaluated two backbone families (ViT-B/16 and ConvNeXt-B) and different input resolutions. **Table 1** summarizes dataset characteristics, **Table 2** reports overall performance, and **Table 3** lists exact p-values between all pairwise comparisons. Overall performance distributions across datasets are shown in **Supplementary Figure 1**, while accuracy, sensitivity, and specificity are detailed in **Supplementary Figure 2** and per-label metrics in **Supplementary Tables 1–7**.

a Workflow schematic



b Scale of the benchmark



c Example radiographs



Figure 1: Study overview. (a) Workflow schematic of the experimental design. Seven chest radiograph datasets were included: Pedi-CXR (training n = 7,728; test n = 1,397), VinDr-CXR (training n = 15,000; test n = 3,000), ChestX-ray14 (training n = 86,524; test n = 25,596), PadChest (training n = 88,480; test n = 22,045), CheXpert (training n = 128,356; test n = 29,320), MIMIC-CXR (training n = 170,153; test n = 43,768), and UKA-CXR (training n = 153,537; test n = 39,824). Models were trained with two backbone families (ViT-B/16, ConvNeXt-B), three initialization strategies (ImageNet, DINOv2, DINOv3), and frozen features from the DINOv3-7B teacher, evaluated at two input resolutions (224×224 and 512×512). **(b)** Scale of the benchmark, totaling 814,930 anteroposterior or posteroanterior chest radiographs across 21 diagnostic labels from adult and pediatric cohorts. **(c)** Example radiographs from the UKA-CXR dataset.

Table 1. Characteristics of the datasets utilized in this study. Summary of patient cohorts, image counts, demographics, and label sets for all seven datasets: Pedi-CXR, VinDr-CXR, ChestX-ray14, PadChest, CheXpert, MIMIC-CXR, and UKA-CXR. Reported values include the number of patients and radiographs, split into training and test sets, as well as patient age distributions (median, mean $\pm$ standard deviation (SD), and range) and sex ratios (female/male, given separately for training and test sets). The labels used for multi-label classification are listed as defined in each dataset. Dataset locations and the distribution of image projections (anteroposterior vs. posteroanterior) are also reported. Whenever available, the “no finding” label was preserved as a separate category to indicate a completely normal radiograph without any imaging abnormality, not merely the absence of the labels considered in this study. Patient-wise splits were used in all datasets to ensure no overlap between training and test cohorts. Only anteroposterior or posteroanterior images are considered in this study. N/A = not available. * The youngest patient was a one-month-old baby.

Dataset	Number of patients (n)	Number of radiographs (n) Total Training set Test set	Patient age (years) Median Mean $\pm$ SD Range	Patient sex (female/male [%]) Training set Test set	Labels used in this study	Location	Projections (%)
							Anteroposterior Posteroanterior
Pedi-CXR	N/A	9,125 7,728 1,397	2 4 $\pm$ 3 (0.1*, 10)	42.4/57.6 40.9/59.1	no finding, pneumonia, bronchitis/bronchiolitis	Vietnam	0.0 100.0
VinDr-CXR	N/A	18,000 15,000 3,000	42 54 $\pm$ 18 (2, 91)	47.8/52.2 44.1/55.9	cardiomegaly, pleural effusion, pneumonia, atelectasis, no finding, consolidation, pneumothorax, pleural thickening, lung opacity, pulmonary fibrosis, nodule/mass	Vietnam	0.0 100.0
ChestX-ray14	30,805	112,120 86,524 25,596	49 47 $\pm$ 17 (1, 96)	42.4/57.6 41.9/58.1	cardiomegaly, effusion, pneumonia, atelectasis, no finding, consolidation, pneumothorax, fibrosis, emphysema, hernia, pleural thickening, edema, nodule, mass	USA	40.0 60.0
PadChest	67,213	110,525 88,480 22,045	63 59 $\pm$ 20 (1, 105)	50.0/50.0 48.2/51.8	cardiomegaly, pleural effusion, pneumonia, atelectasis, no finding, consolidation, pneumothorax, emphysema, hernia, scoliosis, congestion, aortic elongation, kyphosis, COPD signs, pleural thickening, nodule mass, infiltrates	Spain	17.1 82.9
CheXpert	65,240	157,878 128,356 29,320	61 60 $\pm$ 18 (18, 91)	41.4/58.6 39.0/61.0	cardiomegaly, pleural effusion, pneumonia, atelectasis, no finding, consolidation, pneumothorax, lung opacity, lung lesion, fracture	USA	84.5 15.5
MIMIC-CXR	65,379	213,921 170,153 43,768	N/A N/A N/A	N/A N/A N/A	cardiomegaly, pleural effusion, pneumonia, atelectasis, no finding, consolidation, pneumothorax, lung opacity, lung lesion, fracture	USA	58.2 41.8
UKA-CXR	54,176	193,361 153,537 39,824	68 66 $\pm$ 15 (1, 111)	34.4/65.6 36.3/63.7	cardiomegaly, congestion, pleural effusion, pneumonic infiltrates, atelectasis, no finding	Germany	100.0 0.0

At 224×224 , DINOv2 significantly outperformed ImageNet in five of six adult datasets (all $p = 0.0060$), whereas DINOv3 showed significant improvement in all six ($p \leq 0.012$) and no gain over DINOv2 (0/6 significantly higher). Thus, DINOv3 did not exceed DINOv2 at low resolution. At 512×512 , DINOv3 achieved a clear and statistically significant advantage over both ImageNet (6/6 datasets, $p \leq 0.0060$) and DINOv2 (5/6 datasets, $p \leq 0.019$), confirming the intended benefit of its high-resolution Gram-anchored self-distillation, while DINOv2 surpassed ImageNet in 5/6 datasets ($p \leq 0.044$). The pediatric dataset (Pedi-CXR) differed from all adult cohorts, showing no significant improvement in any comparison at 224×224 or 512×512 ($p \geq 0.31$). Scaling to 1024×1024 did not yield systematic gains over 512×512 , confirming 512×512 as the optimal balance between accuracy and computational cost. ConvNeXt-B consistently outperformed ViT-B across resolutions, and the strongest overall performance was obtained with DINOv3 + ConvNeXt-B ($p \leq 0.013$ vs. ImageNet; while not significant for only VinDr-CXR dataset at 224×224). Finally, in adult datasets, frozen features from the 7B-parameter DINOv3 model underperformed relative to finetuned 86–89M models ($p \leq 0.0060$ in all cases except VinDr-CXR at 512 with ViT-B, $p = 0.23$; not significant) and were also inferior to ImageNet across datasets, underscoring the continued need for domain-specific finetuning in medical imaging.

Resolution scaling ($224 \rightarrow 512 \rightarrow 1024$)

Resolution scaling revealed a clear shift in the relative strength of DINOv2 and DINOv3 across datasets and backbones (**Figure 2, Tables 2 and 3**). At 224×224 , DINOv2 significantly outperformed ImageNet in five of six adult datasets (all $p = 0.0060$), confirming the benefit of self-supervised initialization. DINOv3 also showed significant improvement over ImageNet in all six adult datasets ($p \leq 0.012$) but did not exceed DINOv2 in any (0/6 significantly higher) and underperformed in three (3/6 significantly lower). For example, on CheXpert (ViT-B), DINOv2 achieved 80.29 ± 0.17 vs. 79.98 ± 0.17 for DINOv3, and on MIMIC-CXR, 80.86 ± 0.16 vs. 80.76 ± 0.15 .

At 512×512 , DINOv3 achieved a clear and statistically significant advantage over both ImageNet (6/6 datasets, $p \leq 0.0060$) and DINOv2 (5/6 datasets, $p \leq 0.019$), with absolute AUROC gains of 0.8–1.5 percentage points. Representative examples include CheXpert (81.90 vs. 81.39) and MIMIC-CXR (82.49 vs. 80.72). PadChest and ChestX-ray14 showed similar significant trends, while UKA-CXR exhibited smaller yet consistent gains. DINOv2 remained significantly superior to ImageNet in 5/6 datasets ($p \leq 0.044$). The pediatric dataset (Pedi-CXR) differed from all adult cohorts, showing no significant improvement across any initialization at 224 or 512 ($p \geq 0.31$). These findings confirm that DINOv3's advantages emerge primarily at higher resolution, consistent with its design for fine-grained feature preservation via Gram-anchored self-distillation. Resolution-dependent calibration curves are shown in **Supplementary Figure 3**.

Table 2. Overall performance across datasets and initialization strategies. Average area under the receiver operating characteristic curve (AUROC) derived from 1,000 bootstrap resamples for full finetuning of ViT-B/16 and ConvNeXt-B backbones across two resolutions (224 × 224 and 512 × 512 pixels). Results are reported for models initialized from ImageNet, DINOv2, and DINOv3. Frozen DINOv3-7B results are shown separately for comparison. Results are shown for all datasets: Pedi-CXR (training n = 7,728; test n = 1,397), VinDr-CXR (training n = 15,000; test n = 3,000), ChestX-ray14 (training n = 86,524; test n = 25,596), PadChest (training n = 88,480; test n = 22,045), CheXpert (training n = 128,356; test n = 29,320), MIMIC-CXR (training n = 170,153; test n = 43,768), and UKA-CXR (training n = 153,537; test n = 39,824).. Values are presented as mean ± standard deviation [95% confidence intervals (CIs)].

Full finetuning										
	ImageNet				DINOv2		DINOv3			
	ViT		ConvNeXt		ViT		ViT		ConvNeXt	
	224	512	224	512	224	512	224	512	224	512
Pedi-CXR	73.20 ± 1.17 [70.94, 75.54]	73.66 ± 1.21 [71.12, 75.97]	72.61 ± 1.24 [70.19, 75.01]	74.00 ± 1.13 [71.79, 76.22]	73.39 ± 1.17 [71.12, 75.59]	74.13 ± 1.18 [71.96, 76.43]	72.71 ± 1.23 [70.23, 75.06]	73.94 ± 1.17 [71.61, 76.33]	73.80 ± 1.14 [71.58, 75.95]	73.64 ± 1.21 [71.24, 75.89]
VinDr-CXR	88.31 ± 0.61 [87.05, 89.45]	86.42 ± 0.65 [85.06, 87.63]	88.01 ± 0.59 [86.83, 89.11]	89.68 ± 0.53 [88.64, 90.70]	89.16 ± 0.74 [87.47, 90.48]	89.09 ± 0.54 [88.03, 90.13]	90.15 ± 0.56 [89.02, 91.24]	90.26 ± 0.49 [89.30, 91.16]	88.04 ± 0.55 [86.96, 89.13]	90.49 ± 0.48 [89.54, 91.45]
ChestXray14	78.97 ± 0.20 [78.56, 79.38]	79.54 ± 0.23 [79.09, 80.01]	79.78 ± 0.21 [79.37, 80.18]	81.55 ± 0.20 [81.13, 81.92]	80.08 ± 0.24 [79.58, 80.54]	80.01 ± 0.22 [79.56, 80.44]	80.13 ± 0.23 [79.67, 80.56]	81.35 ± 0.24 [80.87, 81.77]	80.37 ± 0.22 [79.95, 80.80]	82.28 ± 0.24 [81.83, 82.73]
PadChest	87.04 ± 0.20 [86.65, 87.43]	87.56 ± 0.20 [87.16, 87.95]	87.45 ± 0.20 [87.03, 87.83]	88.45 ± 0.19 [88.05, 88.82]	88.00 ± 0.21 [87.59, 88.40]	88.48 ± 0.20 [88.06, 88.84]	87.39 ± 0.22 [86.95, 87.82]	88.94 ± 0.19 [88.57, 89.30]	88.00 ± 0.19 [87.61, 88.33]	89.33 ± 0.17 [89.01, 89.67]
CheXpert	79.68 ± 0.17 [79.39, 80.02]	80.65 ± 0.16 [80.31, 80.96]	79.52 ± 0.17 [79.19, 79.81]	81.56 ± 0.16 [81.25, 81.89]	80.29 ± 0.17 [79.95, 80.60]	81.39 ± 0.16 [81.07, 81.70]	79.98 ± 0.17 [79.66, 80.30]	81.90 ± 0.17 [81.56, 82.21]	80.48 ± 0.17 [80.15, 80.79]	82.50 ± 0.15 [82.19, 82.80]
MIMIC-CXR	79.91 ± 0.16 [79.61, 80.24]	81.38 ± 0.15 [81.08, 81.68]	80.42 ± 0.15 [80.13, 80.72]	82.15 ± 0.15 [81.86, 82.44]	80.86 ± 0.16 [80.56, 81.16]	80.72 ± 0.16 [80.41, 81.05]	80.76 ± 0.15 [80.45, 81.05]	82.49 ± 0.15 [82.21, 82.80]	81.21 ± 0.15 [80.92, 81.51]	82.66 ± 0.16 [82.35, 82.97]
UKA-CXR	87.71 ± 0.11 [87.50, 87.92]	87.96 ± 0.11 [87.75, 88.16]	87.84 ± 0.11 [87.62, 88.05]	88.37 ± 0.11 [88.16, 88.58]	88.08 ± 0.11 [87.87, 88.29]	87.99 ± 0.11 [87.77, 88.19]	87.95 ± 0.11 [87.75, 88.16]	88.45 ± 0.10 [88.24, 88.65]	88.18 ± 0.11 [87.97, 88.39]	88.52 ± 0.11 [88.31, 88.73]
Frozen DINOv3-7B										
	224 x 224					512 x 512				
Pedi-CXR	68.82 ± 1.28 [66.38, 71.27]					69.35 ± 1.30 [66.85, 71.84]				
VinDr-CXR	82.80 ± 0.78 [81.19, 84.30]					86.77 ± 0.63 [85.50, 87.95]				
ChestXray14	76.12 ± 0.22 [75.69, 76.52]					77.66 ± 0.26 [77.12, 78.17]				
PadChest	84.60 ± 0.21 [84.19, 84.98]					86.03 ± 0.20 [85.65, 86.42]				
CheXpert	76.83 ± 0.17 [76.49, 77.16]					78.88 ± 0.17 [78.54, 79.20]				
MIMIC-CXR	77.06 ± 0.16 [76.76, 77.37]					78.91 ± 0.17 [78.59, 79.23]				
UKA-CXR	83.83 ± 0.13 [83.57, 84.08]					85.43 ± 0.12 [85.20, 85.66]				

Table 3: Pairwise bootstrap p-values across datasets, initialization strategies, and input resolutions.

Two-sided p-values were obtained from paired bootstrap tests using 1,000 resampled AUROC pairs per model, ensuring identical resampling across initialization strategies for fair comparison. Results are reported separately for ViT-B/16 and ConvNeXt-B under full finetuning, as well as for frozen DINOv3-7B encoders with linear classifiers. Each comparison was performed at 224×224 , 512×512 , and—where available— 1024×1024 input resolutions. p-values were adjusted for multiple comparisons within coherent families of related tests (e.g., per-resolution comparisons across the six adult datasets) using the Benjamini–Hochberg false discovery rate (FDR) procedure, with FDR-adjusted $p < 0.05$ considered statistically significant. “N/A” indicates configurations that were not evaluated. Results are shown for all adult datasets: Pedi-CXR (training $n = 7,728$; test $n = 1,397$), VinDr-CXR (training $n = 15,000$; test $n = 3,000$), ChestX-ray14 (training $n = 86,524$; test $n = 25,596$), PadChest (training $n = 88,480$; test $n = 22,045$), CheXpert (training $n = 128,356$; test $n = 29,320$), MIMIC-CXR (training $n = 170,153$; test $n = 43,768$), and UKA-CXR (training $n = 153,537$; test $n = 39,824$). Dataset descriptions are provided in **Table 1**.

Backbone/ fine tuning	Resolution	Pedi-CXR	VinDr- CXR	ChestX ray14	PadChes t	CheXpert	MIMIC- CXR	UKA- CXR	Pedi- CXR
ViT full fine tuning	224 x 224	ImageNet vs DinoV3	0.31	0.0060	0.012	0.0060	0.0060	0.0060	0.0060
		ImageNet vs DinoV2	0.61	0.11	0.0060	0.0060	0.0060	0.0060	0.0060
		DinoV3 vs DinoV2	0.78	0.050	0.0030	0.40	0.0030	0.20	0.0080
	512 x 512	ImageNet vs DinoV3	0.37	0.0060	0.0060	0.0060	0.0060	0.0060	0.0060
		ImageNet vs DinoV2	0.32	0.0060	0.0060	0.0060	0.0060	0.044	0.26
		DinoV3 vs DinoV2	0.54	0.019	0.0060	0.0060	0.0060	0.0060	0.0060
	1024 x 1024	ImageNet vs DinoV3	0.0010	N/A	N/A	0.0020	N/A	0.0020	N/A
		ImageNet vs DinoV2	0.0010	N/A	N/A	0.0020	N/A	0.40	N/A
		DinoV3 vs DinoV2	0.0010	N/A	N/A	0.0020	N/A	0.0020	N/A
ConvNeXt full fine tuning	224 x 224	ImageNet vs DinoV3	0.080	0.45	0.0010	0.0010	0.0010	0.0010	0.0010
	512 x 512	ImageNet vs DinoV3	0.66	0.013	0.0010	0.0010	0.0010	0.0010	0.0010
	1024 x 1024	ImageNet vs DinoV3	0.76	N/A	N/A	0.0020	N/A	0.038	N/A
ViT frozen features	224 x 224	Frozen vs DinoV3	0.0010	0.0060	0.0060	0.0060	0.0060	0.0060	0.0060
		Frozen vs DinoV2	0.0010	0.0060	0.0060	0.0060	0.0060	0.0060	0.0060
		Frozen vs ImageNet	0.0010	0.0060	0.0060	0.0060	0.0060	0.0060	0.0060
	512 x 512	Frozen vs DinoV3	0.0010	0.0060	0.0060	0.0060	0.0060	0.0060	0.0060
		Frozen vs DinoV2	0.0010	0.0060	0.0060	0.0060	0.0060	0.0060	0.0060
		Frozen vs ImageNet	0.0010	0.23	0.0060	0.0060	0.0060	0.0060	0.0060
ConvNeXt frozen features	224 x 224	Frozen vs DinoV3	0.0010	0.0060	0.0060	0.0060	0.0060	0.0060	0.0060
		Frozen vs ImageNet	0.0010	0.0060	0.0060	0.0060	0.0060	0.0060	0.0060
	512 x 512	Frozen vs DinoV3	0.0010	0.0060	0.0060	0.0060	0.0060	0.0060	0.0060
		Frozen vs ImageNet	0.0010	0.0060	0.0060	0.0060	0.0060	0.0060	0.0060

To probe further scaling, we extended experiments to 1024×1024 on three representative datasets (MIMIC-CXR, ChestX-ray14, and Pedi-CXR). Within this setting, DINOv3 again significantly outperformed the baselines wherever evaluated (MIMIC-CXR, ChestX-ray14; $p = 0.0020$). However, relative to 512×512 the absolute AUROC changes were small and inconsistent, for example, MIMIC-CXR showed a modest rise for DINOv3–ConvNeXt-B ($82.66 \rightarrow 83.34$), while ChestX-ray14 changed little or decreased slightly for some backbones, and Pedi-CXR remained essentially unchanged. We did not perform formal cross-resolution hypothesis tests (512×512 vs. 1024×1024). Overall, scaling beyond 512×512 did not yield systematic additional gains across datasets or backbones, supporting 512×512 as the practical balance between performance and computational cost.

Backbone effects (ConvNeXt-B vs. ViT-B)

Backbone comparisons showed that ConvNeXt-B consistently outperformed ViT-B across datasets and resolutions (**Figure 3, Tables 2 and 3**). At 224×224 , ConvNeXt-B provided modest but significant gains on 5/6 adult datasets (all $p = 0.0010$). For example, on MIMIC-CXR, DINOv3 achieved 81.21 ± 0.15 vs. 80.42 ± 0.15 for ImageNet. Similar but smaller advantages were seen on UKA-CXR (88.18 ± 0.11 vs. 87.84 ± 0.11) and ChestX-ray14 (80.37 ± 0.22 vs. 79.78 ± 0.21). At 512×512 , the advantage of ConvNeXt-B widened further, reaching statistical significance in all adult datasets (all $p \leq 0.013$). On ChestX-ray14, DINOv3 achieved 82.28 ± 0.24 vs. 81.55 ± 0.20 for ImageNet. The pediatric dataset (Pedi-CXR) again showed no meaningful backbone-related separation ($p \geq 0.66$). Complete results are provided in **Supplementary Figure 4**.


Overall, ConvNeXt-B outperformed ViT-B when paired with any initialization, with DINOv3 consistently amplifying this advantage. These findings demonstrate that DINOv3’s benefits are not limited to transformer-based architectures but extend robustly to modern convolutional backbones.

Frozen DINOv3-7B versus finetuned 86–89M models

We next compared frozen representations from the 7B-parameter DINOv3 teacher, used only with a linear classification head, to fully finetuned ViT-B (86 M) and ConvNeXt-B (87 M) models (**Figure 4**). Despite its scale, the frozen model consistently underperformed across datasets and resolutions.

At 224×224 , frozen DINOv3 lagged significantly behind finetuned models in all adult datasets (all $p \leq 0.0060$). On VinDr-CXR, frozen DINOv3 reached 82.80 ± 0.78 vs. 90.15 ± 0.56 for finetuned DINOv3–ViT, and similar gaps were observed on ChestX-ray14 (76.12 ± 0.22 vs. 80.37 ± 0.22) and PadChest (84.60 ± 0.21 vs. 88.94 ± 0.19). Even Pedi-CXR followed the same pattern ($p = 0.0010$).

a ViT



b ConvNeXt




Figure 2: Resolution scaling from 224 x 224 to 1024 x 1024 pixels. Bar plots for average area under the receiver operating characteristic curve (AUROC) values across all labels, comparing ImageNet, DINOv2, and DINOv3 initializations at 224 × 224, 512 × 512, and 1024 × 1024 resolution. Results are shown for two representative datasets: ChestX-ray14 (training n = 86,524; test n = 25,596), and MIMIC-CXR (training n = 170,153; test n = 43,768). **(a)** ViT-B/16 backbone. **(b)** ConvNeXt-B backbone.

At 512 × 512, the performance gap persisted across both backbones. Frozen DINOv3 achieved 69.35 ± 1.30 on Pedi-CXR, 86.77 ± 0.63 on VinDr-CXR, and 78.91 ± 0.17 on MIMIC-CXR, compared with 73.94 ± 1.17 , 90.49 ± 0.48 , and 82.66 ± 0.16 for finetuned DINOv3–ConvNeXt-B, respectively. Differences remained statistically significant in all adult datasets ($p = 0.0060$) and the pediatric dataset ($p = 0.0010$) except for VinDr-CXR with ViT-B ($p = 0.23$). Across datasets, frozen DINOv3 features were also consistently inferior to ImageNet initialization (all $p \leq 0.0060$). Complete results are provided in **Supplementary Figure 5**.

In summary, frozen billion-parameter encoders remain markedly inferior to finetuned 86–89 M models, demonstrating that model scale alone cannot replace task-specific adaptation in medical imaging.




Figure 3: Backbone comparison across datasets. Mean AUROC values across all labels with standard deviations from 1,000 bootstrap resamples for ViT-B/16 and ConvNeXt-B backbones at 512×512 resolution. Results for ImageNet, DINOv2, and DINOv3 initializations are shown side by side within each backbone. Results are shown for all adult datasets: VinDr-CXR (training $n = 15,000$; test $n = 3,000$), ChestX-ray14 (training $n = 86,524$; test $n = 25,596$), PadChest (training $n = 88,480$; test $n = 22,045$), CheXpert (training $n = 128,356$; test $n = 29,320$), MIMIC-CXR (training $n = 170,153$; test $n = 43,768$), and UKA-CXR (training $n = 153,537$; test $n = 39,824$).




Figure 4: Classification using frozen DINOv3-7B features versus full finetuning of smaller models. Bootstrap distributions of AUROC values ($n = 1,000$ resamples) comparing classifiers trained on frozen DINOv3-7B features with lightweight heads ($\sim 2\text{M}$ parameters) against full finetuning of ViT-B/16 and ConvNeXt-B backbones ($\sim 86\text{--}87\text{M}$ parameters). Across datasets, full finetuning consistently outperforms frozen representations, despite the much smaller backbone size. Results are shown for all seven datasets: Pedi-CXR (training $n = 7,728$; test $n = 1,397$), VinDr-CXR (training $n = 15,000$; test $n = 3,000$), ChestX-ray14 (training $n = 86,524$; test $n = 25,596$), PadChest (training $n = 88,480$; test $n = 22,045$), CheXpert (training $n = 128,356$; test $n = 29,320$), MIMIC-CXR (training $n = 170,153$; test $n = 43,768$), and UKA-CXR (training $n = 153,537$; test $n = 39,824$).

Dataset- and label-level performance patterns


Absolute AUROC values varied systematically across datasets but were not solely determined by dataset size (**Supplementary Figure 6**). At 512×512 with ConvNeXt-B + DINOv3, VinDr-CXR ($n = 18\text{ k}$) achieved the highest mean AUROC (90.49 ± 0.48 [95% CI: 89.54, 91.45]), followed by PadChest (89.33 ± 0.17 [89.01, 89.67]), consistent with their expert-curated or diverse annotations. In contrast, CheXpert and MIMIC-CXR, both labeled via NLP pipelines, reached lower values in the low-82% range (82.50 ± 0.15 and 82.66 ± 0.16 , respectively), reflecting the impact of label noise. ChestX-ray14 showed similar moderate performance (82.28 ± 0.24 [81.83, 82.73]), while UKA-CXR (88.52 ± 0.11 [88.31, 88.73]) benefited from its narrower label space. The pediatric cohort (Pedi-CXR) remained lowest (73.64 ± 1.21 [71.24, 75.89]), likely due to its smaller sample size and greater variability in acquisition conditions.

Despite these absolute differences, the relative ordering of initialization strategies was highly stable across datasets. At 512×512 , DINOv3 significantly outperformed both ImageNet and DINOv2 in five of six adult datasets ($p \leq 0.019$), whereas Pedi-CXR again showed no significant improvement ($p \geq 0.31$). DINOv2 also remained consistently superior to ImageNet in five of six datasets ($p \leq 0.044$). Full bootstrap distributions of Δ AUROC values are provided in **Supplementary Figure 7**.

Label-wise analyses (**Figure 5; Supplementary Tables 1–7**) revealed that large-structure findings such as pleural effusion and cardiomegaly achieved the highest AUROCs (often $> 90\%$), while focal or boundary-centered findings, notably pulmonary nodules and pneumothorax, benefited most from high-resolution inputs and DINOv3 initialization. Diffuse or textural findings such as consolidation and atelectasis showed moderate but consistent resolution-linked gains (typically $+3$ – 5 AUROC percentage points). Representative receiver operating characteristic (ROC) curves for selected labels (cardiomegaly, pleural effusion, pneumonia) are provided in **Supplementary Figure 8**.

In summary, dataset-level trends reflect differences in annotation quality more than dataset size, while label-level results confirm that DINOv3's high-resolution representations preferentially enhance fine-detail recognition in chest radiographs.

a 224 x 224



b 512 x 512




Figure 5: Label-wise performance analysis. (a) Heatmap of differences in AUROC (ΔAUROC) between DINOv3 and DINOv2 across all labels and datasets at 224×224 resolution. Results are shown for all seven datasets: Pedi-CXR (training $n = 7,728$; test $n = 1,397$), VinDr-CXR (training $n = 15,000$; test $n = 3,000$), ChestX-ray14 (training $n = 86,524$; test $n = 25,596$), PadChest (training $n = 88,480$; test $n = 22,045$), CheXpert (training $n = 128,356$; test $n = 29,320$), MIMIC-CXR (training $n = 170,153$; test $n = 43,768$), and UKA-CXR (training $n = 153,537$; test $n = 39,824$). (b) Corresponding heatmap at 512×512 resolution.

Discussion

In this study, we provide the first systematic benchmark of DINOv3 for chest radiograph classification across more than 814,000 images from seven diverse datasets. Our results yield four central insights. First, DINOv3 does transfer to medical imaging, but its advantages over DINOv2 are modest at 224×224 pixels and become consistent only when input resolution is scaled to 512×512 pixels. Second, backbone choice matters: while ConvNeXt-B did not always exceed ViT-B at the lower resolution, it provided clear and consistent gains at 512×512 resolution, with the combination of ConvNeXt-B and DINOv3 emerging as the strongest overall configuration. Third, simply freezing features from billion-parameter DINOv3 models proved insufficient; targeted finetuning of smaller 86–89M parameter networks remained decisively stronger. Fourth, experiments at 1024×1024 resolution revealed no measurable improvement over 512×512 pixels, indicating that for DINOv3 the practical resolution ceiling in chest radiography is 512×512 pixels. Taken together, these findings show that the benefits of modern SSL methods extend to radiology, but they depend critically on resolution scaling, backbone choice, and domain-specific adaptation rather than parameter count alone. From a clinical perspective, this pattern suggests that SSL models such as DINOv3 are most advantageous for detecting small, well-defined lesions (e.g., pulmonary nodules) and subtle, low-contrast abnormalities (e.g., interstitial changes or early pulmonary opacification), where higher spatial fidelity directly improves diagnostic confidence.

Resolution dependence is particularly noteworthy. At the standard 224×224 -pixel input size, DINOv2 often retained a slight advantage on adult datasets, with DINOv3 showing comparable but not consistently superior results. At 512×512 pixels, however, DINOv3 surpassed both DINOv2 and ImageNet across nearly all adult cohorts, in some cases by close to or exceeding one AUROC point. Experiments at 1024×1024 resolution on three representative datasets (Pedi-CXR, ChestX-ray14, and MIMIC-CXR) did not yield further gains over 512×512 , suggesting that the benefits of DINOv3 may saturate at this resolution under current training conditions. This observation aligns with recent evidence that scaling laws established on natural images do not straightforwardly extend to medical imaging²⁵, and the pattern is consistent with the design of DINOv3, which incorporates high-resolution adaptations through Gram-anchored distillation²². Medical images, and chest radiographs in particular, are of exquisitely high resolution in the range of 2000 to 4000 pixels per image dimension, and, thus, contain fine features such as interstitial markings, vascular structures, and subtle reticular or ground-glass changes, that may not be fully represented at lower resolutions⁴¹. At the same time, prior convolutional network studies have reported that inputs between 256×256 and 448×448 were sufficient for diagnostic performance in chest radiography^{21,42,43}, highlighting that the resolution demands of SSL-based high-capacity backbones may differ from earlier supervised CNN settings. The pediatric cohort was a notable exception, where limited sample size and narrower label ontology constrained performance and minimized differences between initialization strategies^{31,44,45}. Overall, these results align with prior observations that certain diagnostic imaging tasks, particularly those involving fine-grained structures, benefit from higher spatial fidelity than is typically required in natural-image benchmarks.

From a clinical perspective, the observed AUROC improvements, typically in the 0.5–1.0 range, translate to greater reliability for subtle or low-contrast findings that are easily overlooked on standard-resolution models. In particular, boundary-centered findings, such as pneumothorax, and small focal lesions, such as pulmonary nodules, benefited most from 512×512 inputs with DINOv3. These findings suggest that high-resolution self-supervised features can enhance the detection of subtle pathologies, supporting potential applications in triage, emergency, and critical-care settings where timely recognition of subtle changes is essential.

The backbone comparison reinforced the importance of architecture in medical imaging. ConvNeXt-B, a modern convolutional design, consistently outperformed ViT-B across all datasets and resolutions, with the advantage becoming more pronounced when paired with DINOv3. This suggests that architectural advances in convolutional networks remain highly relevant for radiology, particularly in combination with state-of-the-art self-supervised pretraining. While transformers have attracted much attention in the field²⁸, our findings indicate that convolutional backbones, when integrated with SSL, currently provide a favorable balance of accuracy, robustness, and adaptability in chest radiograph analysis.

The underperformance of frozen DINOv3-7B features relative to finetuned smaller models provides an important cautionary note. Despite their scale, billion-parameter encoders trained exclusively on natural images did not yield superior diagnostic accuracy when used without adaptation. Across all datasets—including those with structured labels such as UKA-CXR—full finetuning of much smaller 86–89 M parameter backbones consistently achieved higher AUROCs. This shows that model size alone does not ensure transferability to clinical imaging tasks. While full finetuning of the 7B-parameter model may eventually enhance performance, domain-specific adaptation remains essential. In practical terms, compact models that are carefully optimized for medical data can deliver greater diagnostic value and efficiency than massive frozen networks.

Dataset characteristics also shaped performance. While AUROCs are not directly comparable across datasets due to differences in label sets and difficulty, clear trends emerged within each cohort. For example, VinDr-CXR achieved very high AUROCs under SSL initialization, consistent with its carefully curated expert labels, while PadChest benefited strongly from resolution scaling, likely reflecting its diverse label ontology. In contrast, very large but noisily labeled datasets such as MIMIC-CXR and CheXpert reached lower absolute AUROCs within their own domains, underscoring that scale alone does not guarantee stronger results. Similarly, ChestX-ray14 remained limited by its NLP-derived labeling system, and UKA-CXR's narrower label set constrained its performance relative to more heterogeneous cohorts. The pediatric dataset, Pedi-CXR, showed the lowest absolute AUROCs, likely due to both its small size and the increased difficulty of pediatric imaging. Importantly, however, the relative ordering of initialization strategies was stable across all cohorts: ImageNet baselines consistently trailed SSL, and DINOv3 at 512 resolution generally ranked best. This reproducible trend across populations, label sets, and geographic origins strengthens the generality of our conclusions.

Because the datasets span multiple continents and acquisition protocols, potential sources of bias must be considered. Differences in projection type, particularly the higher proportion of AP portable studies in UKA-CXR and CheXpert compared with the PA images in

VinDr-CXR may influence apparent model performance. Findings that depend strongly on boundary delineation, such as pneumothorax or cardiomegaly, can appear less distinct in AP views, emphasizing the value of higher-resolution inputs. Future work should examine whether such resolution-sensitive benefits persist across projection types and demographic subgroups to ensure equitable model performance.

Our study has several limitations. First, although we analyzed more than 814,000 radiographs, our focus was limited to AP or PA chest radiographs; extending evaluation to lateral views and other modalities such as computed tomography, mammography, or ultrasound will be essential to determine whether high-resolution self-supervision provides similar benefits across imaging domains. Second, while we included datasets from three continents, institutional biases in labeling and acquisition protocols remain. In particular, NLP-derived labels in large public datasets such as MIMIC-CXR, CheXpert, and ChestX-ray14 likely attenuate apparent performance gains relative to expertly curated datasets; future work should explore weak-to-strong or report-guided label refinement to mitigate this limitation. Third, computational efficiency was not systematically assessed, and higher input resolutions inevitably incur additional cost. Although 1024×1024 experiments on three representative datasets and two backbones revealed modest gains in MIMIC-CXR, they provided no consistent improvement beyond 512×512 , suggesting diminishing returns under current training conditions^{42,43}. Broader exploration across architectures and pretraining strategies will be needed to determine whether this reflects a model-specific limitation or a general ceiling for resolution scaling in chest radiography. Finally, our analysis focused on classification tasks; whether the resolution-linked advantages of DINOv3 extend to segmentation, localization, or report generation remains to be tested.

In conclusion, this work demonstrates that DINOv3 provides measurable improvements for chest radiograph classification when scaled to higher resolutions and paired with modern backbones, but also reveals that frozen billion-scale vision models alone do not obviate the need for finetuning. While 512×512 emerged as the most effective setting in our benchmark, preliminary 1024×1024 experiments suggested no further gains despite markedly higher cost. For clinical AI, the path forward may lie less in sheer model size or resolution scaling, and more in the careful alignment of pretraining innovations with the spatial and diagnostic demands of medical imaging, particularly for fine-detail findings such as interstitial markings, early pulmonary edema, and subtle reticular or ground-glass opacities. These findings provide actionable guidance for integrating next-generation SSL into radiology workflows and establish a foundation for future studies exploring how scaling strategies, architectural choices, and domain adaptation can be balanced in practice.

Materials and methods

Ethics statement

All methods were carried out in accordance with relevant guidelines and regulations. Ethical approval for this retrospective study was obtained from the Ethics Committee of the Medical Faculty of RWTH Aachen University (Reference No. EK 028/19). The requirement for individual informed consent was waived by the committee.

Patient datasets

This study included a total of $n=814,930$ AP or PA chest radiographs from seven international cohorts encompassing both adult and pediatric populations. Patients ranged in age from infancy to over 111 years. The datasets span diverse geographic regions (Asia, Europe, and North America), label generation strategies (manual annotation, rule-based natural language processing (NLP), and hybrid approaches), and clinical contexts (inpatient, outpatient, intensive care, and pediatrics). A detailed overview of dataset characteristics is provided in **Table 1**. Below, we describe each dataset.

Pedi-CXR dataset

The Pedi-CXR³¹ dataset is the largest publicly available pediatric chest radiograph dataset with diagnostic labels. It contains 9,125 posteroanterior images from children under the age of 10 years (median: 2 years), collected in Vietnam. All radiographs were manually annotated by three radiologists with at least 10 years of experience. For this study, we followed the dataset's original split into training ($n = 7,728$) and test ($n = 1,397$) sets. Labels include pneumonia and related pediatric conditions (see **Table 1**).

VinDr-CXR dataset

The VinDr-CXR³² dataset comprises 18,000 adult radiographs, curated from more than 100,000 studies performed at two Vietnamese hospitals. Images were acquired on equipment from multiple manufacturers. Labeling was performed by 17 radiologists, with each image independently annotated by three experts. The dataset authors provided a patient-wise split into $n=15,000$ training and $n=3,000$ test images, which we used directly. Labels cover common thoracic diseases such as cardiomegaly, effusion, and pneumonia (see **Table 1**).

ChestX-ray14 dataset

The ChestX-ray14³³ dataset, released by the National Institutes of Health, contains $n=112,120$ AP or PA radiographs from 30,805 patients. Fourteen thoracic pathologies were labeled using a two-stage NLP pipeline applied to corresponding radiology reports. Following prior work, we generated a patient-wise 80%/20% split, resulting in $n=86,524$ training and $n=25,596$ test images. Labels span major cardiopulmonary conditions (see **Table 1**).

PadChest dataset

The PadChest³⁴ dataset includes n=110,525 AP or PA radiographs from the Hospital Universitario de San Juan in Alicante, Spain. Labels were derived from radiology reports in Spanish: 27,593 studies were manually annotated by radiologists, and the remainder were automatically labeled using a text classifier trained on this subset. We performed a patient-wise 80%/20% split, stratified by manual vs. automatic labeling, yielding n=88,480 training and n=22,045 test images. Labels are diverse and include both common and less frequent findings (see **Table 1**).

CheXpert dataset

The CheXpert³⁵ dataset consists of n=157,878 AP or PA chest radiographs from 65,240 patients at Stanford Hospital in CA, USA. Labels for 14 common radiographic findings were extracted using a rule-based NLP system that categorized mentions as positive, negative, or uncertain. Following established practice, uncertain and negative mentions were grouped as “negative.” We used a patient-wise 80%/20% split, resulting in 128,356 training and 29,320 test images. Labels include cardiomegaly, effusion, pneumonia, and others (see **Table 1**).

MIMIC-CXR dataset

The MIMIC-CXR⁶ dataset contains n=213,921 AP or PA radiographs from Beth Israel Deaconess Medical Center in Boston, MA, USA, collected between 2011 and 2016. Images were de-identified and linked to associated reports. Labels were generated automatically using the same NLP system as CheXpert³⁵, ensuring consistency. We created a patient-wise 80%/20% split, yielding n=171,137 training and n=42,784 test images. Labels overlap with CheXpert, covering major cardiopulmonary findings (see **Table 1**).

UKA-CXR dataset

The UKA-CXR^{21,36–40} dataset is an internal cohort from University Hospital RWTH Aachen, in Aachen, Germany. It includes n=193,361 adult AP radiographs collected between 2009 and 2020 across 10 intensive care units, using 18 radiography systems. Images were labeled by radiologists within the clinical reporting workflow, using a structured template with categories such as pleural effusion, pneumonia, atelectasis, congestion, and cardiomegaly. For this study, we defined a patient-wise 80%/20% split into training and test sets. Labels reflect routine diagnostic categories from clinical reporting (see **Table 1**).

Label system and preprocessing

As in our prior works^{21,36–40,45,46}, all datasets were mapped into a unified binary multilabel classification framework, where each image was assigned a positive or negative label for every included condition. Only AP or PA views were used in all experiments. Pedi-CXR, VinDr-CXR, ChestX-ray14, and PadChest were provided in binary format by design and were used directly.

In CheXpert, and consequently in MIMIC-CXR, the original four categories (“positive,” “negative,” “uncertain,” and “not mentioned”) were reduced to binary by treating “negative,” “uncertain,” and “not mentioned” as negative, and considering only “positive” as positive. For the UKA-CXR dataset, which contained multiple severity levels, “normal” and “uncertain” were classified as negative, while all severity categories above normal (for example, “mild,” “moderate,” or “severe” for effusion, and “borderline,” “enlarged,” or “massively enlarged” for cardiomegaly) were classified as positive. Additionally, UKA-CXR contained separate left- and right-sided labels for several findings; in these cases, the presence of a finding on either side was counted as positive. In PadChest, where annotations were generated through a combination of manual labeling and NLP, only the subset of labels overlapping with the target label system was retained, and these were binarized accordingly. Finally, whenever available, the “no finding” label was preserved as a separate category to indicate a completely normal radiograph without any imaging abnormality, not merely the absence of the labels considered in this study.

Images were supplied in mixed formats depending on the dataset. ChestX-ray14, PadChest, CheXpert, and MIMIC-CXR were already available as PNG/JPG files, while Pedi-CXR, VinDr-CXR and UKA-CXR were provided in DICOM format and converted to PNG/JPG prior to analysis. For DICOM images, metadata were checked to ensure correct polarity; if pixel intensities were stored in inverted form, they were re-inverted to maintain consistent orientation. All radiographs were to 224×224 , 512×512 , or 1024×1024 pixels. To normalize intensity values, each image was shifted such that the minimum pixel value corresponded to zero, scaled by the maximum, and clipped to the valid range before conversion to 8-bit grayscale⁶. Contrast was then enhanced by applying histogram equalization implemented with the OpenCV library^{6,37}. These steps yielded a uniform preprocessing pipeline across datasets, with patient-wise splits ensuring strict separation of training and test cohorts.

Experimental design

To ensure consistency across experiments, we applied a unified pre-processing workflow to all datasets and fixed the training/test splits throughout the study. The held-out test sets comprised $n = 1,397$ (Pedi-CXR), $n = 3,000$ (VinDr-CXR), $n = 25,596$ (ChestX-ray14), $n = 22,045$ (PadChest), $n = 29,320$ (CheXpert), $n = 43,768$ (MIMIC-CXR), and $n = 38,672$ (UKA-CXR), with no patient overlap between training and test partitions. For each dataset, multilabel classification was performed using the set of available imaging findings that met minimum prevalence thresholds, resulting in 3–17 labels per dataset (**Table 1**).

We benchmarked three initialization strategies: supervised ImageNet-21K⁷, self-supervised DINOv2, and the recently introduced DINOv3. Two backbone families were evaluated: the Vision Transformer base model (ViT-B/16, ~86M parameters) and ConvNeXt-B (~89M parameters), each trained at two input resolutions, 224×224 and 512×512 pixels. In addition, we examined frozen representations from the large DINOv3-7B model. For this setting, we added a compact multilayer classification head, referred to as DinoNet. DinoNet consisted of a layer

normalization⁴⁷ applied to the 4096-dimensional backbone features, followed by a Gaussian error linear unit (GELU)⁴⁸ activation and a dropout layer with $p = 0.3$. The output was then passed through a linear projection to a 512-dimensional embedding, followed by another dropout layer with $p = 0.3$, a second layer normalization, and finally a linear mapping to the target label space. This head added approximately 2.1 million trainable parameters, while the 7B backbone itself remained frozen.

For full finetuning, all backbone parameters were optimized using AdamW⁴⁹ with a learning rate of 10^{-5} and no weight decay. For frozen-feature experiments with DINOv3-7B, only the DINO-Net classifier was optimized, using AdamW with a learning rate of 10^{-4} and weight decay of 5×10^{-5} . Across all experiments, data augmentation consisted of random horizontal flips and random rotations up to 7° . The loss function was a binary weighted cross-entropy, with class weights set inversely proportional to the frequency of each label in the training set⁵⁰. Batch sizes were adapted to fit available GPU memory.

Self-supervised pretraining objectives (DINOv3)

To contextualize the initialization strategies evaluated in this study, we briefly summarize the pretraining objectives underlying DINOv3²², which extend the DINOv2²⁰ framework. These objectives were not retrained on medical images here, but they underpin the pretrained weights used in our experiments.

Image- and patch-level consistency

Following the student–teacher paradigm, an image x is augmented into two views, x_s (student) and x_t (teacher). Encoders f_s, f_t produce normalized embeddings $z_s, z_t \in R^d$, projected onto K prototypes with SoftMax normalization. The image-level objective enforces invariance of global representations:

$$L_{img} = - \sum_{k=1}^K p_t^{(k)}(x_t) \log p_s^{(k)}(x_s), \quad (1)$$

where p_s, p_t are the probability vectors of student and teacher outputs.

To capture local information, the patch-level objective compares masked student patches with the corresponding teacher features:

$$L_{patch} = - \sum_{j=1}^P \sum_{k=1}^K p_t^{(j,k)} \log p_s^{(j,k)}, \quad (2)$$

with P the number of visible patches.

Feature regularization

To prevent representational collapse, a KoLeo entropy regularizer⁵¹ is included:

$$L_{KoLeo} = \frac{1}{B} \sum_{i=1}^B \log \left(\frac{1}{B} \sum_{j=1}^B \exp(-\|z_i - z_j\|^2) \right), \quad (3)$$

where B is the batch size. The combined pretraining loss is

$$L_{pre} = L_{img} + L_{patch} + L_{KoLeo}. \quad (4)$$

Gram anchoring refinement

Extended pretraining of large ViTs can degrade patch-level similarity²². DINOv3 introduces Gram anchoring to stabilize dense features. Given current student features $X_s \in R^{p \times d}$ and stored teacher features X_G , the Gram loss is:

$$L_{Gram} = \|X_s X_s^T - X_G X_G^T\|_F^2, \quad (5)$$

penalizing divergence of pairwise similarity structures. The final refinement stage optimizes:

$$L_{ref} = \alpha L_{img} + \beta L_{patch} + \gamma L_{KoLeo} + \delta L_{Gram}. \quad (6)$$

where $\alpha, \beta, \gamma, \delta \in R^+$ are hyperparameters balancing the contributions of global consistency, local invariance, entropy regularization, and Gram anchoring. These weights are tuned during DINOv3 pretraining and fixed in the publicly released checkpoints we adopt.

High-resolution adaptation

Unlike earlier SSL methods, DINOv3 explicitly supports larger input sizes. Mixed-resolution training with rotary positional embeddings ensures that features remain spatially stable when scaling to higher resolutions. In this study, we leveraged these pretrained weights directly, testing both transformer (ViT-B/16) and convolutional (ConvNeXt-B) backbones distilled from the DINOv3 teacher.

Evaluation

The primary evaluation metric was the area under the receiver operating characteristic curve (AUROC), which provides a threshold-independent measure of discrimination in multilabel classification. Accuracy, sensitivity, and specificity were reported as complementary metrics. For each dataset, we summarized results using the mean AUROC across all labels, while per-label AUROC, accuracy, sensitivity, and specificity are provided in the supplementary information.

Thresholds for sensitivity and specificity were chosen according to Youden’s criterion⁵², i.e., the cut-off maximizing the difference between true-positive and false-positive rates.

Statistical analysis was performed using Python 3.9 with NumPy 1.22, SciPy 1.10, scikit-learn 1.2. Bootstrapping with 1,000 redraws was used to estimate means, standard deviations, and 95% confidence intervals (CI)^{21,53}. A paired design ensured identical resampling across initialization strategies to enable fair within-dataset comparisons. Statistical significance between model pairs was assessed using paired bootstrap tests on AUROC differences^{21,36,46}. To control for multiple comparisons across datasets and configurations, p-values were adjusted within coherent families of related tests (e.g., per-resolution comparisons across the six adult datasets) using the Benjamini–Hochberg false discovery rate (FDR) procedure, with statistical significance defined as FDR-adjusted as $p < 0.05$ ⁵⁴.

Data availability

The datasets used in this study are either publicly accessible, available under controlled access, or internal. The ChestX-ray14 and PadChest datasets are publicly available at <https://www.kaggle.com/datasets/nih-chest-xrays/data> and <https://bimcv.cipf.es/bimcv-projects/padchest/>, respectively. The VinDr-CXR and MIMIC-CXR datasets are restricted-access resources hosted on PhysioNet and can be obtained by agreeing to the relevant data protection requirements at <https://physionet.org/content/vindr-cxr/1.0.0/> and <https://physionet.org/content/mimic-cxr-jpg/2.0.0/>. The Pedi-CXR dataset (VinDr-PCXR) is also available through PhysioNet at <https://physionet.org/content/vindr-pcxr/1.0.0/>. The CheXpert dataset may be requested from Stanford University at <https://stanfordmlgroup.github.io/competitions/chexpert/>. The UKA-CXR dataset contains patient data from the University Hospital Aachen, Germany; access may be granted upon reasonable request to the corresponding authors and within a written cooperation agreement. A subset of the UKA-CXR dataset is publicly available on Hugging Face via <https://huggingface.co/TLAIM>.

Code availability and reproducibility

All source code, configuration files, and instructions to reproduce the experiments are available at <https://github.com/tayebiarasteh/vit-med>. Training and evaluation were performed strictly in full 32-bit floating point (FP32) precision. Experiments were conducted between August 13, 2025, and September 22, 2025.

Implementation details: Python 3.9 with PyTorch 2.8 and torchvision 0.23. Core libraries: NumPy 1.22, SciPy 1.10, scikit-learn 1.2, pandas 1.4, timm 0.6, and OpenCV (cv2) 4.7. Hugging Face tooling: transformers 4.56, huggingface-hub 0.34, datasets 2.19, accelerate 1.10, tokenizers 0.21, and safetensors 0.4.

Pretrained initialization weights were obtained from official public repositories:

- ImageNet-21K:
 - ViT-B/16 from timm (model identifier *vit_base_patch16_224_in21k*).
 - ConvNeXt-B from Hugging Face (loaded using the safetensors format): <https://huggingface.co/facebook/convnext-base-224-22k>
- DINOv2:
 - ViT-B/16 from Hugging Face, configured with scaled-dot product attention: <https://huggingface.co/facebook/dinov2-base>
- DINOv3:
 - ViT-B/16 from Hugging Face, configured with scaled-dot product attention: <https://huggingface.co/facebook/dinov3-vitb16-pretrain-lvd1689m>
 - ViT-7B from Hugging Face, configured with scaled-dot product attention: <https://huggingface.co/facebook/dinov3-vit7b16-pretrain-lvd1689m>
 - ConvNeXt-B from Hugging Face (loaded using the safetensors format): <https://huggingface.co/facebook/dinov3-convnext-base-pretrain-lvd1689m>

Additional information

Funding

JNK is supported by the German Cancer Aid (DECADE, 70115166), the German Federal Ministry of Education and Research (PEARL, 01KD2104C; CAMINO, 01EO2101; SWAG, 01KD2215A; TRANSFORM LIVER, 031L0312A; TANGERINE, 01KT2302 through ERA-NET Transcan; Come2Data, 16DKZ2044A; DEEP-HCC, 031L0315A), the German Academic Exchange Service (SECAI, 57616814), the German Federal Joint Committee (TransplantKI, 01VSF21048) the European Union's Horizon Europe and innovation programme (ODELIA, 101057091; GENIAL, 101096312), the European Research Council (ERC; NADIR, 101114631), the National Institutes of Health (EPICO, R01 CA263318) and the National Institute for Health and Care Research (NIHR, NIHR203331) Leeds Biomedical Research Centre. The views expressed are those of the author(s) and not necessarily those of the NHS, the NIHR or the Department of Health and Social Care. This work was funded by the European Union. Views and opinions expressed are however those of the author(s) only and do not necessarily reflect those of the European Union. Neither the European Union nor the granting authority can be held responsible for them. SN was supported by grants from the Deutsche Forschungsgemeinschaft (DFG) (NE 2136/3-1, LI3893/6-1, TR 1700/7-1). DT was supported by grants from the DFG (NE 2136/3-1, LI3893/6-1, TR 1700/7-1) and is supported by the German Federal Ministry of Education (TRANSFORM LIVER, 031L0312A; SWAG, 01KD2215B) and the European Union's Horizon Europe and innovation programme (ODELIA [Open Consortium for Decentralized Medical Artificial Intelligence], 101057091).

Author contributions

The formal analysis was conducted by STA, SN, and DT. The original draft was written by STA and edited by STA, MS, SN, and DT. The code was developed by STA. The experiments were performed by STA. The illustrations were designed by MS. The statistical analyses were performed by STA, SN, and DT. STA, CK, JNK, SN, and DT provided clinical expertise. STA, MS, JNK, and DT provided technical expertise. The study was defined by STA, SN, and DT. All authors read the manuscript and agreed to the submission of this paper.

Competing interests

STA is an editorial board at *Communications Medicine* and at *European Radiology Experimental*, a trainee editorial board at *Radiology: Artificial Intelligence*, and is partially employed by Synagen GmbH, Germany. JNK declares consulting services for Bioptimus, France; Owkin, France; DoMore Diagnostics, Norway; Panakeia, UK; AstraZeneca, UK; Scailyte, Switzerland; Mindpeak, Germany; and MultiplexDx, Slovakia. Furthermore, he holds shares in StratifAI GmbH, Germany, and in Synagen GmbH, Germany, has received a research grant by GSK, and has received honoraria by AstraZeneca, Bayer, Eisai, Janssen, MSD, BMS, Roche, Pfizer and Fresenius. DT received honoraria for lectures by Bayer, GE, Roche, AstraZeneca, and Philips and holds shares in StratifAI GmbH, Germany, and in Synagen GmbH, Germany. The other authors do not have any competing interests to disclose.

References


1. Rajpurkar, P., Chen, E., Banerjee, O. & Topol, E. J. AI in health and medicine. *Nat Med* **28**, 31–38 (2022).
2. Tayebi Arasteh, S. *et al.* Large language models streamline automated machine learning for clinical studies. *Nat Commun* **15**, 1603 (2024).
3. Haug, C. J. & Drazen, J. M. Artificial Intelligence and Machine Learning in Clinical Medicine, 2023. *N Engl J Med* **388**, 1201–1208 (2023).
4. Tayebi Arasteh, S. *et al.* The Treasure Trove Hidden in Plain Sight: The Utility of GPT-4 in Chest Radiograph Evaluation. *Radiology* **313**, e233441 (2024).
5. Chen, Z. *et al.* A Vision-Language Foundation Model to Enhance Efficiency of Chest X-ray Interpretation. Preprint at <https://doi.org/10.48550/arXiv.2401.12208> (2024).
6. Johnson, A. E. W. *et al.* MIMIC-CXR, a de-identified publicly available database of chest radiographs with free-text reports. *Sci Data* **6**, 317 (2019).
7. Deng, J. *et al.* ImageNet: A large-scale hierarchical image database. in *2009 IEEE Conference on Computer Vision and Pattern Recognition* 248–255 (IEEE, Miami, FL, 2009). doi:10.1109/CVPR.2009.5206848.
8. Ke, A., Ellsworth, W., Banerjee, O., Ng, A. Y. & Rajpurkar, P. CheXtransfer: performance and parameter efficiency of ImageNet models for chest X-Ray interpretation. in *Proceedings of the Conference on Health, Inference, and Learning* 116–124 (ACM, Virtual Event USA, 2021). doi:10.1145/3450439.3451867.
9. Krishnan, R., Rajpurkar, P. & Topol, E. J. Self-supervised learning in medicine and healthcare. *Nat. Biomed. Eng* **6**, 1346–1352 (2022).
10. Hendrycks, D., Mazeika, M., Kadavath, S. & Song, D. Using self-supervised learning can improve model robustness and uncertainty. in *NIPS'19: Proceedings of the 33rd International Conference on Neural Information Processing Systems* vol. 1403 15663–15674 (2019).
11. He, K., Fan, H., Wu, Y., Xie, S. & Girshick, R. Momentum Contrast for Unsupervised Visual Representation Learning. in *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)* 9729–9738 (2020).
12. Chen, T., Kornblith, S., Norouzi, M. & Hinton, G. A Simple Framework for Contrastive Learning of Visual Representations. in *International Conference on Machine Learning* vol. 119 (Vienna, Austria, 2020).
13. Grill, J.-B. *et al.* Bootstrap your own latent-a new approach to self-supervised learning. *Advances in neural information processing systems* **33**, 21271–21284 (2020).
14. Caron, M. *et al.* Unsupervised Learning of Visual Features by Contrasting Cluster Assignments. in *Advances in neural information processing systems* 33 9912–9924 (2020).
15. Wen, Y., Chen, L., Deng, Y. & Zhou, C. Rethinking pre-training on medical imaging. *Journal of Visual Communication and Image Representation* **78**, 103145 (2021).
16. Vaswani, A. *et al.* Attention Is All You Need. in *NIPS'17: Proceedings of the 31st International Conference on Neural Information Processing Systems* 6000–6010 (2017).
17. Dosovitskiy, A. *et al.* An Image is Worth 16x16 Words: Transformers for Image Recognition at Scale. Preprint at <http://arxiv.org/abs/2010.11929> (2021).
18. Liu, Z. *et al.* A convnet for the 2020s. in *Proceedings of the IEEE/CVF conference on computer vision and pattern recognition* 11976–11986 (2022).
19. Caron, M. *et al.* Emerging Properties in Self-Supervised Vision Transformers. in *Proceedings of the IEEE/CVF International Conference on Computer Vision (ICCV)* 9650–9660 (2021).

20. Oquab, M. *et al.* DINOv2: Learning Robust Visual Features without Supervision. Preprint at <http://arxiv.org/abs/2304.07193> (2023).
21. Tayebi Arasteh, S., Misera, L., Kather, J. N., Truhn, D. & Nebelung, S. Enhancing diagnostic deep learning via self-supervised pretraining on large-scale, unlabeled non-medical images. *Eur Radiol Exp* **8**, 10 (2024).
22. Siméoni, O. *et al.* DINOv3. Preprint at <https://doi.org/10.48550/arXiv.2508.10104> (2025).
23. Yang, S., Wang, H., Xing, Z., Chen, S. & Zhu, L. SegDINO: An Efficient Design for Medical and Natural Image Segmentation with DINO-V3. Preprint at <https://doi.org/10.48550/arXiv.2509.00833> (2025).
24. Li, Y., Wu, Y., Lai, Y., Hu, M. & Yang, X. MedDINOv3: How to adapt vision foundation models for medical image segmentation? Preprint at <https://doi.org/10.48550/arXiv.2509.02379> (2025).
25. Liu, C. *et al.* Does DINOv3 Set a New Medical Vision Standard? Preprint at <https://doi.org/10.48550/arXiv.2509.06467> (2025).
26. Khader, F. *et al.* Multimodal Deep Learning for Integrating Chest Radiographs and Clinical Parameters: A Case for Transformers. *Radiology* **309**, e230806 (2023).
27. Wang, B., Li, Q. & You, Z. Self-supervised learning based transformer and convolution hybrid network for one-shot organ segmentation. *Neurocomputing* **527**, 1–12 (2023).
28. He, K. *et al.* Transformers in medical image analysis. *Intelligent Medicine* **3**, 59–78 (2023).
29. Tanno, R. *et al.* Collaboration between clinicians and vision–language models in radiology report generation. *Nat Med* **31**, 599–608 (2025).
30. Sloan, P., Clatworthy, P., Simpson, E. & Mirmehdi, M. Automated radiology report generation: A review of recent advances. *IEEE Reviews in Biomedical Engineering* **18**, 368–387 (2024).
31. Nguyen, N. H., Pham, H. H., Tran, T. T., Nguyen, T. N. M. & Nguyen, H. Q. *VinDr-PCXR: An Open, Large-Scale Chest Radiograph Dataset for Interpretation of Common Thoracic Diseases in Children*. <http://medrxiv.org/lookup/doi/10.1101/2022.03.04.22271937> (2022) doi:10.1101/2022.03.04.22271937.
32. Nguyen, H. Q. *et al.* VinDr-CXR: An open dataset of chest X-rays with radiologist’s annotations. *Sci Data* **9**, 429 (2022).
33. Wang, X. *et al.* ChestX-ray8: Hospital-scale Chest X-ray Database and Benchmarks on Weakly-Supervised Classification and Localization of Common Thorax Diseases. in *2017 IEEE Conference on Computer Vision and Pattern Recognition (CVPR)* 3462–3471 (2017). doi:10.1109/CVPR.2017.369.
34. Bustos, A., Pertusa, A., Salinas, J.-M. & de la Iglesia-Vayá, M. PadChest: A large chest x-ray image dataset with multi-label annotated reports. *Medical Image Analysis* **66**, 101797 (2020).
35. Irvin, J. *et al.* CheXpert: A Large Chest Radiograph Dataset with Uncertainty Labels and Expert Comparison. *AAAI* **33**, 590–597 (2019).
36. Khader, F. *et al.* Artificial Intelligence for Clinical Interpretation of Bedside Chest Radiographs. *Radiology* **307**, e220510 (2022).
37. Tayebi Arasteh, S. *et al.* Collaborative training of medical artificial intelligence models with non-uniform labels. *Sci Rep* **13**, 6046 (2023).
38. Tayebi Arasteh, S. *et al.* Preserving fairness and diagnostic accuracy in private large-scale AI models for medical imaging. *Commun Med* **4**, 46 (2024).
39. Tayebi Arasteh, S. *et al.* Securing Collaborative Medical AI by Using Differential Privacy: Domain Transfer for Classification of Chest Radiographs. *Radiology. Artificial Intelligence* **6**, e230212 (2024).
40. Tayebi Arasteh, S., Isfort, P., Kuhl, C., Nebelung, S. & Truhn, D. Automatic Evaluation of Chest Radiographs – The Data Source Matters, But How Much Exactly? in *RöFo-*


- Fortschritte auf dem Gebiet der Röntgenstrahlen und der bildgebenden Verfahren* vol. 195 ab99 (Georg Thieme Verlag, RheinMain CongressCenter (RMCC) in Wiesbaden, 2023).
41. Chiarenza, A. et al. Chest imaging using signs, symbols, and naturalistic images: a practical guide for radiologists and non-radiologists. *Insights Imaging* **10**, 114 (2019).
 42. Sabottke, C. F. & Spieler, B. M. The Effect of Image Resolution on Deep Learning in Radiography. *Radiology: Artificial Intelligence* **2**, e190015 (2020).
 43. Haque, M. I. U. et al. Effect of image resolution on automated classification of chest X-rays. *J Med Imaging (Bellingham)* **10**, 044503 (2023).
 44. Capitanio, M. A. Pitfalls in Pediatric Chest Radiography. *Radiology* **137**, 656–656 (1980).
 45. Lotfinia, M., Tayebiarasteh, A., Samiei, S., Joodaki, M. & Tayebi Arasteh, S. Boosting multi-demographic federated learning for chest radiograph analysis using general-purpose self-supervised representations. *European Journal of Radiology Artificial Intelligence* **3**, 100028 (2025).
 46. Tayebi Arasteh, S. et al. Enhancing domain generalization in the AI-based analysis of chest radiographs with federated learning. *Sci Rep* **13**, 22576 (2023).
 47. Ba, J. L., Kiros, J. R. & Hinton, G. E. Layer Normalization. Preprint at <https://doi.org/10.48550/arXiv.1607.06450> (2016).
 48. Hendrycks, D. & Gimpel, K. Gaussian Error Linear Units (GELUs). Preprint at <https://doi.org/10.48550/arXiv.1606.08415> (2023).
 49. Loshchilov, I. & Hutter, F. Decoupled Weight Decay Regularization. in *Proceedings of Proceedings of Seventh International Conference on Learning Representations (ICLR) 2019* (New Orleans, LA, USA, 2019).
 50. Rezaei-Dastjerdehei, M. R., Mijani, A. & Fatemizadeh, E. Addressing Imbalance in Multi-Label Classification Using Weighted Cross Entropy Loss Function. in *2020 27th National and 5th International Iranian Conference on Biomedical Engineering (ICBME)* 333–338 (IEEE, Tehran, Iran, 2020). doi:10.1109/ICBME51989.2020.9319440.
 51. Sablayrolles, A., Douze, M., Schmid, C. & Jégou, H. Spreading vectors for similarity search. in *Proceedings of Proceedings of Seventh International Conference on Learning Representations (ICLR) 2019* (arXiv, New Orleans, LA, USA, 2019). doi:10.48550/ARXIV.1806.03198.
 52. Unal, I. Defining an Optimal Cut-Point Value in ROC Analysis: An Alternative Approach. *Comput Math Methods Med* **2017**, 3762651 (2017).
 53. Konietzschke, F. & Pauly, M. Bootstrapping and permuting paired t-test type statistics. *Stat Comput* **24**, 283–296 (2014).
 54. Tayebi Arasteh, S. et al. RadioRAG: Online Retrieval–Augmented Generation for Radiology Question Answering. *Radiology: Artificial Intelligence* **7**, e240476 (2025).

Supplementary information


a 224 x 224




b 512 x 512



Supplementary Figure 1: Overall performance distributions across datasets. (a) Violin plots of bootstrap distributions (n = 1,000 resamples) for average AUROC values across all labels, comparing ImageNet, DINOv2, and DINOv3 initializations at 224 × 224 resolution with the ViT-B/16 backbone. At this resolution, DINOv2 often retained a slight edge, with DINOv3 performing comparably. (b) Corresponding bootstrap distributions at 512 × 512 resolution. Results are shown for all seven datasets: Pedi-CXR (training n = 7,728; test n = 1,397), VinDr-CXR (training n = 15,000; test n = 3,000), ChestX-ray14 (training n = 86,524; test n = 25,596), PadChest (training n = 88,480; test n = 22,045), CheXpert (training n = 128,356; test n = 29,320), MIMIC-CXR (training n = 170,153; test n = 43,768), and UKA-CXR (training n = 153,537; test n = 39,824).




Supplementary Figure 2: Accuracy, sensitivity, and specificity across initialization strategies. Each panel corresponds to one dataset, showing bar plots for ImageNet, DINOv2, and DINOv3 at 224×224 and 512×512 resolutions. Bars display mean values with standard deviations from 1,000 bootstrap resamples.




Supplementary Figure 3: Calibration analysis. Reliability diagrams comparing expected versus observed probabilities for ImageNet and DINOv3 initializations under full finetuning. Plots are stratified by dataset (rows) and resolution (224×224 , 512×512 ; columns), with separate panels for ViT-B/16 (a) and ConvNeXt-B (b). The diagonal line indicates perfect calibration. Results are averaged across all labels within each dataset, with bootstrapped confidence shading from 1,000 resamples.


a 224 x 224



b 512 x 512




Supplementary Figure 4: Backbone comparison across datasets. (a) Mean AUROC values across all labels with standard deviations from 1,000 bootstrap resamples for ViT-B/16 and ConvNeXt-B backbones at 224 × 224 resolution. Results for ImageNet, DINOv2, and DINOv3 initializations are shown side by side within each backbone. (b) Corresponding results at 512 × 512 resolution. Results are shown for all seven datasets: Pedi-CXR (training n = 7,728; test n = 1,397), VinDr-CXR (training n = 15,000; test n = 3,000), ChestX-ray14 (training n = 86,524; test n = 25,596), PadChest (training n = 88,480; test n = 22,045), CheXpert (training n = 128,356; test n = 29,320), MIMIC-CXR (training n = 170,153; test n = 43,768), and UKA-CXR (training n = 153,537; test n = 39,824).




Supplementary Figure 5: Classification using frozen DINOv3-7B features versus full finetuning of smaller models. Bootstrap distributions of AUROC values ($n = 1,000$ resamples) comparing classifiers trained on frozen DINOv3-7B features with lightweight heads ($\sim 2\text{M}$ parameters) against full finetuning of ViT-B/16 and ConvNeXt-B backbones ($\sim 86\text{--}87\text{M}$ parameters). Across datasets, full finetuning consistently outperforms frozen representations, despite the much smaller backbone size. Results are shown for all seven datasets: Pedi-CXR (training $n = 7,728$; test $n = 1,397$), VinDr-CXR (training $n = 15,000$; test $n = 3,000$), ChestX-ray14 (training $n = 86,524$; test $n = 25,596$), PadChest (training $n = 88,480$; test $n = 22,045$), CheXpert (training $n = 128,356$; test $n = 29,320$), MIMIC-CXR (training $n = 170,153$; test $n = 43,768$), and UKA-CXR (training $n = 153,537$; test $n = 39,824$).


a Relationship between dataset size and mean AUROC



b ROC curves for pediatric dataset




c AUROC distributions for different resolutions




Supplementary Figure 6: Dataset and sample size effects on performance. (a) Relationship between dataset size and mean AUROC, with each point representing an initialization strategy (ImageNet, DINOv2, DINOv3). Results are shown for all seven datasets: Pedi-CXR (training $n = 7,728$; test $n = 1,397$), VinDr-CXR (training $n = 15,000$; test $n = 3,000$), ChestX-ray14 (training $n = 86,524$; test $n = 25,596$), PadChest (training $n = 88,480$; test $n = 22,045$), CheXpert (training $n = 128,356$; test $n = 29,320$), MIMIC-CXR (training $n = 170,153$; test $n = 43,768$), and UKA-CXR (training $n = 153,537$; test $n = 39,824$). (b) Receiver operating characteristic (ROC) curves for the pediatric dataset (Pedi-CXR), showing uniformly lower absolute AUROC across all initialization strategies, reflecting limited sample size and label diversity despite similar relative trends. (c) Bootstrap distributions of average AUROC values for PadChest and CheXpert datasets at 224×224 and 512×512 resolutions, illustrating that the margin of improvement from DINOv3 widens at higher resolution in this diverse, multi-label dataset.

Distribution of Δ AUROC values (DINOv3 – DINOv2)



Supplementary Figure 7: Distribution of Δ AUROC values (DINOv3 – DINOv2) for the ViT-B backbone.

Kernel density plots show bootstrap-sampled differences in average AUROC ($n = 1,000$ resamples) across all labels for each dataset, comparing 224×224 and 512×512 resolutions. A vertical dashed line at zero indicates no difference between methods. Results are shown for all seven datasets: Pedi-CXR (training $n = 7,728$; test $n = 1,397$), VinDr-CXR (training $n = 15,000$; test $n = 3,000$), ChestX-ray14 (training $n = 86,524$; test $n = 25,596$), PadChest (training $n = 88,480$; test $n = 22,045$), CheXpert (training $n = 128,356$; test $n = 29,320$), MIMIC-CXR (training $n = 170,153$; test $n = 43,768$), and UKA-CXR (training $n = 153,537$; test $n = 39,824$).



Supplementary Figure 8: Representative receiver operating characteristic (ROC) curves for selected labels (cardiomegaly, pleural effusion, pneumonia) at 512×512 resolution. Results are shown for all adult datasets: VinDr-CXR (training $n = 15,000$; test $n = 3,000$), ChestX-ray14 (training $n = 86,524$; test $n = 25,596$), PadChest (training $n = 88,480$; test $n = 22,045$), CheXpert (training $n = 128,356$; test $n = 29,320$), MIMIC-CXR (training $n = 170,153$; test $n = 43,768$), and UKA-CXR (training $n = 153,537$; test $n = 39,824$).

Supplementary Tables

Supplementary Table 1: Per-label performance on the Pedi-CXR dataset under full finetuning. Reported metrics include AUROC, accuracy, sensitivity, and specificity, given as mean $\pm$ standard (in percent) deviation across 1,000 bootstrap resamples for all 3 labels. Results are provided for ViT-B/16 and ConvNeXt-B backbones initialized from ImageNet, DINOv2, and DINOv3 at two input resolutions (224×224 and 512×512).

	ImageNet				DINOv2		DINOv3			
	ViT		ConvNeXt		ViT		ViT		ConvNeXt	
Label	224	512	224	512	224	512	224	512	224	512
AUROC										
No finding	75.0 $\pm$ 1.4	74.0 $\pm$ 1.4	73.2 $\pm$ 1.4	74.5 $\pm$ 1.4	74.5 $\pm$ 1.4	74.5 $\pm$ 1.4	73.0 $\pm$ 1.4	74.2 $\pm$ 1.4	74.7 $\pm$ 1.4	73.5 $\pm$ 1.4
Pneumonia	81.4 $\pm$ 1.7	81.6 $\pm$ 1.7	79.4 $\pm$ 1.8	82.9 $\pm$ 1.7	81.8 $\pm$ 1.8	83.0 $\pm$ 1.6	81.3 $\pm$ 1.8	82.3 $\pm$ 1.7	80.1 $\pm$ 1.8	81.2 $\pm$ 1.7
Bronchitis/br onchiolitis	63.2 $\pm$ 1.7	65.4 $\pm$ 1.9	65.2 $\pm$ 1.9	64.7 $\pm$ 1.8	64.0 $\pm$ 1.8	64.8 $\pm$ 1.8	63.9 $\pm$ 1.9	65.3 $\pm$ 1.9	66.6 $\pm$ 1.7	66.2 $\pm$ 1.8
Accuracy										
No finding	70.9 $\pm$ 2.0	68.7 $\pm$ 2.6	69.5 $\pm$ 2.5	70.8 $\pm$ 3.3	71.0 $\pm$ 1.9	70.1 $\pm$ 2.1	69.8 $\pm$ 2.1	71.5 $\pm$ 2.4	68.7 $\pm$ 2.6	70.3 $\pm$ 1.4
Pneumonia	74.6 $\pm$ 5.0	76.0 $\pm$ 5.7	75.9 $\pm$ 5.3	77.9 $\pm$ 2.4	75.5 $\pm$ 6.1	73.7 $\pm$ 4.3	76.1 $\pm$ 4.5	69.9 $\pm$ 3.6	72.7 $\pm$ 3.4	69.4 $\pm$ 4.2
Bronchitis/br onchiolitis	53.0 $\pm$ 6.5	58.0 $\pm$ 4.2	64.8 $\pm$ 6.3	59.7 $\pm$ 7.6	55.9 $\pm$ 3.6	60.3 $\pm$ 9.1	58.9 $\pm$ 4.9	63.6 $\pm$ 3.6	54.9 $\pm$ 4.5	62.8 $\pm$ 9.0
Sensitivity										
No finding	75.3 $\pm$ 5.3	69.1 $\pm$ 8.1	72.3 $\pm$ 7.9	76.6 $\pm$ 10.2	75.8 $\pm$ 4.9	71.9 $\pm$ 5.6	74.9 $\pm$ 6.3	76.0 $\pm$ 7.0	68.9 $\pm$ 7.8	75.9 $\pm$ 2.9
Pneumonia	74.4 $\pm$ 6.5	73.5 $\pm$ 7.2	69.1 $\pm$ 6.4	75.7 $\pm$ 3.9	72.2 $\pm$ 8.1	79.2 $\pm$ 5.7	72.9 $\pm$ 6.0	83.8 $\pm$ 4.8	75.9 $\pm$ 5.0	81.2 $\pm$ 5.9
Bronchitis/br onchiolitis	73.0 $\pm$ 10.2	70.8 $\pm$ 6.6	57.1 $\pm$ 9.7	63.4 $\pm$ 11.7	70.5 $\pm$ 5.5	63.5 $\pm$ 14.0	64.6 $\pm$ 7.8	62.1 $\pm$ 5.9	77.3 $\pm$ 7.0	59.8 $\pm$ 14.5
Specificity										
No finding	62.8 $\pm$ 5.4	67.8 $\pm$ 8.3	64.2 $\pm$ 8.3	60.2 $\pm$ 10.0	62.0 $\pm$ 5.2	66.8 $\pm$ 5.4	60.4 $\pm$ 6.8	63.2 $\pm$ 7.2	68.3 $\pm$ 8.0	60.1 $\pm$ 3.4
Pneumonia	74.6 $\pm$ 6.4	76.4 $\pm$ 7.3	76.8 $\pm$ 6.7	78.2 $\pm$ 3.0	76.0 $\pm$ 7.9	72.9 $\pm$ 5.5	76.5 $\pm$ 5.8	68.0 $\pm$ 4.6	72.3 $\pm$ 4.4	67.8 $\pm$ 5.5
Bronchitis/br onchiolitis	48.4 $\pm$ 10.2	55.0 $\pm$ 6.4	66.7 $\pm$ 9.8	58.9 $\pm$ 12.0	52.5 $\pm$ 5.5	59.5 $\pm$ 14.4	57.6 $\pm$ 7.7	64.0 $\pm$ 5.6	49.6 $\pm$ 7.0	63.5 $\pm$ 14.3

Supplementary Table 2: Per-label performance on VinDr-CXR dataset under full finetuning. Reported metrics include AUROC, accuracy, sensitivity, and specificity, given as mean $\pm$ standard (in percent) deviation across 1,000 bootstrap resamples for all 11 labels. Results are provided for ViT-B/16 and ConvNeXt-B backbones initialized from ImageNet, DINOv2, and DINOv3 at two input resolutions (224×224 and 512×512).

	ImageNet				DINOv2		DINOv3			
	ViT		ConvNeXt		ViT		ViT		ConvNeXt	
Label	224	512	224	512	224	512	224	512	224	512
AUROC										
Cardiomegaly	94.7 $\pm$ 0.5	92.6 $\pm$ 0.7	92.4 $\pm$ 0.7	94.0 $\pm$ 0.6	93.9 $\pm$ 0.6	91.0 $\pm$ 0.7	94.2 $\pm$ 0.6	92.4 $\pm$ 0.7	92.3 $\pm$ 0.7	93.0 $\pm$ 0.6
Pleural effusion	95.3 $\pm$ 1.2	96.2 $\pm$ 1.1	95.5 $\pm$ 1.3	96.5 $\pm$ 1.1	96.7 $\pm$ 0.8	97.5 $\pm$ 0.8	97.9 $\pm$ 0.4	97.9 $\pm$ 0.7	96.9 $\pm$ 0.8	97.4 $\pm$ 0.9
Pneumonia	90.9 $\pm$ 0.9	89.7 $\pm$ 1.1	89.7 $\pm$ 1.1	92.7 $\pm$ 1.0	92.1 $\pm$ 1.0	91.7 $\pm$ 1.0	92.5 $\pm$ 1.0	92.0 $\pm$ 1.0	89.2 $\pm$ 1.2	91.2 $\pm$ 1.1
Atelectasis	84.5 $\pm$ 2.1	71.3 $\pm$ 3.2	81.9 $\pm$ 2.4	82.3 $\pm$ 2.7	86.8 $\pm$ 2.2	86.7 $\pm$ 1.9	85.5 $\pm$ 2.3	84.4 $\pm$ 2.2	84.6 $\pm$ 2.2	89.3 $\pm$ 1.8
No finding	90.7 $\pm$ 0.6	89.7 $\pm$ 0.7	89.5 $\pm$ 0.7	90.0 $\pm$ 0.7	90.9 $\pm$ 0.6	89.1 $\pm$ 0.6	91.9 $\pm$ 0.6	90.5 $\pm$ 0.6	89.2 $\pm$ 0.7	90.9 $\pm$ 0.6
Consolidation	88.5 $\pm$ 1.8	87.3 $\pm$ 2.2	88.8 $\pm$ 1.8	90.1 $\pm$ 2.0	89.8 $\pm$ 2.0	92.7 $\pm$ 1.2	91.5 $\pm$ 1.6	90.0 $\pm$ 2.0	88.8 $\pm$ 1.8	92.7 $\pm$ 1.3
Pneumothorax	90.1 $\pm$ 3.6	94.7 $\pm$ 1.4	93.4 $\pm$ 2.3	96.5 $\pm$ 1.0	86.0 $\pm$ 6.1	91.3 $\pm$ 2.9	90.3 $\pm$ 3.9	96.0 $\pm$ 1.1	93.8 $\pm$ 1.5	95.9 $\pm$ 1.2
Pleural thickening	85.6 $\pm$ 1.5	84.2 $\pm$ 1.6	86.3 $\pm$ 1.4	88.4 $\pm$ 1.3	86.4 $\pm$ 1.5	85.7 $\pm$ 1.5	88.9 $\pm$ 1.2	89.1 $\pm$ 1.2	87.9 $\pm$ 1.3	88.7 $\pm$ 1.3
Lung opacity	85.7 $\pm$ 1.5	82.6 $\pm$ 1.8	85.5 $\pm$ 1.5	86.2 $\pm$ 1.4	86.8 $\pm$ 1.6	85.8 $\pm$ 1.6	87.9 $\pm$ 1.2	87.1 $\pm$ 1.4	85.8 $\pm$ 1.7	85.2 $\pm$ 1.5
Pulmonary fibrosis	83.5 $\pm$ 1.4	82.4 $\pm$ 1.6	83.3 $\pm$ 1.5	85.3 $\pm$ 1.4	86.8 $\pm$ 1.3	82.6 $\pm$ 1.5	85.8 $\pm$ 1.4	86.4 $\pm$ 1.3	81.4 $\pm$ 1.4	84.5 $\pm$ 1.5
Nodule/mass	81.8 $\pm$ 1.6	79.9 $\pm$ 1.8	81.9 $\pm$ 1.7	84.3 $\pm$ 1.6	84.5 $\pm$ 1.7	86.0 $\pm$ 1.4	85.2 $\pm$ 1.5	87.0 $\pm$ 1.6	78.5 $\pm$ 1.8	86.5 $\pm$ 1.5
Accuracy										
Cardiomegaly	82.8 $\pm$ 2.4	85.6 $\pm$ 3.6	83.4 $\pm$ 2.9	87.5 $\pm$ 1.6	84.1 $\pm$ 2.5	80.3 $\pm$ 2.4	87.2 $\pm$ 1.9	85.8 $\pm$ 1.7	85.2 $\pm$ 2.1	82.5 $\pm$ 2.4
Pleural effusion	92.6 $\pm$ 1.4	90.9 $\pm$ 2.6	91.8 $\pm$ 3.2	94.9 $\pm$ 1.0	89.7 $\pm$ 2.2	92.9 $\pm$ 2.0	91.7 $\pm$ 3.0	93.4 $\pm$ 2.5	93.4 $\pm$ 0.6	95.2 $\pm$ 1.1
Pneumonia	83.1 $\pm$ 3.9	83.0 $\pm$ 2.7	83.9 $\pm$ 2.7	87.9 $\pm$ 2.6	88.3 $\pm$ 3.0	86.0 $\pm$ 2.7	84.7 $\pm$ 3.0	87.0 $\pm$ 2.9	85.6 $\pm$ 2.6	84.5 $\pm$ 2.7
Atelectasis	76.6 $\pm$ 8.8	76.1 $\pm$ 11.8	81.7 $\pm$ 7.5	80.3 $\pm$ 3.9	82.1 $\pm$ 2.5	82.5 $\pm$ 5.1	81.6 $\pm$ 5.7	78.8 $\pm$ 4.5	78.0 $\pm$ 3.4	85.4 $\pm$ 2.9
No finding	83.1 $\pm$ 0.9	83.4 $\pm$ 1.1	82.9 $\pm$ 1.0	86.2 $\pm$ 0.7	85.7 $\pm$ 1.0	81.7 $\pm$ 1.2	86.1 $\pm$ 0.9	84.2 $\pm$ 1.2	83.4 $\pm$ 1.2	84.4 $\pm$ 1.1
Consolidation	83.3 $\pm$ 2.7	83.6 $\pm$ 4.2	81.8 $\pm$ 3.8	90.0 $\pm$ 1.8	89.0 $\pm$ 2.4	86.4 $\pm$ 2.3	83.7 $\pm$ 3.7	83.8 $\pm$ 3.1	83.8 $\pm$ 2.7	85.1 $\pm$ 1.6
Pneumothorax	85.5 $\pm$ 8.1	85.9 $\pm$ 4.6	87.5 $\pm$ 5.5	89.3 $\pm$ 1.5	90.2 $\pm$ 5.1	86.6 $\pm$ 6.9	84.3 $\pm$ 8.5	87.8 $\pm$ 2.5	85.5 $\pm$ 4.5	86.8 $\pm$ 3.6
Pleural thickening	76.7 $\pm$ 3.2	79.2 $\pm$ 2.8	79.6 $\pm$ 4.3	82.8 $\pm$ 2.8	77.5 $\pm$ 4.7	80.3 $\pm$ 2.6	81.0 $\pm$ 3.1	84.4 $\pm$ 3.4	77.3 $\pm$ 4.0	83.8 $\pm$ 2.5
Lung opacity	72.7 $\pm$ 4.1	71.2 $\pm$ 3.5	72.0 $\pm$ 4.1	73.5 $\pm$ 4.2	71.9 $\pm$ 4.7	75.6 $\pm$ 3.4	73.6 $\pm$ 4.5	75.9 $\pm$ 4.1	74.0 $\pm$ 5.2	69.0 $\pm$ 5.5
Pulmonary fibrosis	76.4 $\pm$ 4.0	79.4 $\pm$ 2.0	76.4 $\pm$ 4.6	83.1 $\pm$ 4.5	78.2 $\pm$ 4.6	80.4 $\pm$ 3.6	83.5 $\pm$ 1.6	76.1 $\pm$ 5.5	74.3 $\pm$ 4.9	77.5 $\pm$ 2.5
Nodule/mass	78.5 $\pm$ 5.3	83.4 $\pm$ 3.5	76.4 $\pm$ 2.3	82.9 $\pm$ 2.1	78.3 $\pm$ 2.8	76.7 $\pm$ 2.3	78.6 $\pm$ 3.3	82.3 $\pm$ 2.3	74.5 $\pm$ 4.6	80.1 $\pm$ 3.3
Sensitivity										
Cardiomegaly	94.1 $\pm$ 2.7	83.6 $\pm$ 4.2	87.1 $\pm$ 3.4	86.9 $\pm$ 2.4	88.9 $\pm$ 3.1	87.3 $\pm$ 3.1	86.8 $\pm$ 2.8	83.3 $\pm$ 2.5	85.4 $\pm$ 2.7	87.3 $\pm$ 3.0
Pleural effusion	87.6 $\pm$ 3.1	91.3 $\pm$ 2.8	87.8 $\pm$ 3.7	92.8 $\pm$ 2.3	92.5 $\pm$ 2.8	94.6 $\pm$ 2.1	93.1 $\pm$ 2.7	93.1 $\pm$ 2.9	93.1 $\pm$ 2.3	91.1 $\pm$ 2.5
Pneumonia	83.5 $\pm$ 4.5	81.7 $\pm$ 3.6	81.4 $\pm$ 3.4	86.1 $\pm$ 3.2	82.0 $\pm$ 3.7	83.5 $\pm$ 3.5	87.0 $\pm$ 3.6	83.2 $\pm$ 3.6	79.1 $\pm$ 3.6	84.9 $\pm$ 3.4
Atelectasis	76.8 $\pm$ 9.2	56.4 $\pm$ 12.3	70.6 $\pm$ 7.8	74.0 $\pm$ 5.8	80.8 $\pm$ 4.4	77.4 $\pm$ 5.8	78.9 $\pm$ 6.1	78.9 $\pm$ 5.2	78.8 $\pm$ 4.9	81.6 $\pm$ 4.4
No finding	83.4 $\pm$ 1.8	85.2 $\pm$ 2.4	84.2 $\pm$ 2.2	90.7 $\pm$ 1.1	88.5 $\pm$ 2.1	84.1 $\pm$ 2.8	88.5 $\pm$ 2.0	86.6 $\pm$ 2.9	86.0 $\pm$ 2.7	85.2 $\pm$ 2.5
Consolidation	80.2 $\pm$ 4.6	78.9 $\pm$ 5.0	82.6 $\pm$ 4.8	82.7 $\pm$ 3.8	81.7 $\pm$ 4.3	86.0 $\pm$ 3.5	86.3 $\pm$ 4.3	86.0 $\pm$ 4.0	80.7 $\pm$ 4.3	90.3 $\pm$ 3.1
Pneumothorax	76.0 $\pm$ 9.2	86.9 $\pm$ 5.6	81.1 $\pm$ 7.9	91.0 $\pm$ 4.8	69.1 $\pm$ 10.7	78.1 $\pm$ 8.4	77.5 $\pm$ 9.6	90.0 $\pm$ 5.1	86.6 $\pm$ 5.5	89.2 $\pm$ 5.0
Pleural thickening	81.6 $\pm$ 3.9	76.9 $\pm$ 4.2	77.5 $\pm$ 5.1	81.0 $\pm$ 3.7	80.5 $\pm$ 5.0	77.3 $\pm$ 3.8	82.0 $\pm$ 4.1	79.5 $\pm$ 4.2	85.3 $\pm$ 4.5	81.4 $\pm$ 3.7

Lung opacity	87.9 ± 4.6	82.7 ± 4.8	85.9 ± 4.6	85.9 ± 4.6	89.6 ± 4.8	83.5 ± 4.6	90.5 ± 4.1	86.3 ± 4.3	83.5 ± 5.2	88.7 ± 5.4
Pulmonary fibrosis	76.8 ± 4.8	72.3 ± 3.5	75.3 ± 5.3	71.0 ± 5.6	81.1 ± 5.3	69.6 ± 4.7	75.2 ± 3.1	81.2 ± 6.3	73.8 ± 5.7	77.3 ± 3.8
Nodule/mass	71.6 ± 5.8	63.3 ± 4.8	77.2 ± 3.7	72.9 ± 3.7	79.6 ± 4.1	82.8 ± 3.5	77.0 ± 4.2	81.0 ± 3.5	71.3 ± 5.5	78.2 ± 4.3
Specificity										
Cardiomegaly	81.5 ± 2.9	85.8 ± 4.4	83.0 ± 3.5	87.5 ± 1.9	83.5 ± 3.1	79.5 ± 3.0	87.2 ± 2.4	86.1 ± 2.1	85.2 ± 2.6	82.0 ± 2.9
Pleural effusion	92.8 ± 1.5	90.9 ± 2.7	91.9 ± 3.5	94.9 ± 1.0	89.6 ± 2.3	92.8 ± 2.1	91.6 ± 3.2	93.4 ± 2.6	93.4 ± 0.6	95.3 ± 1.1
Pneumonia	83.1 ± 4.6	83.1 ± 3.2	84.1 ± 3.2	88.0 ± 3.0	88.8 ± 3.5	86.2 ± 3.2	84.5 ± 3.5	87.3 ± 3.3	86.2 ± 3.0	84.4 ± 3.2
Atelectasis	76.6 ± 9.3	76.7 ± 12.5	82.0 ± 7.9	80.5 ± 4.2	82.1 ± 2.6	82.6 ± 5.4	81.7 ± 6.0	78.8 ± 4.7	78.0 ± 3.6	85.5 ± 3.0
No finding	82.6 ± 2.0	79.4 ± 2.6	80.0 ± 2.3	76.4 ± 1.6	79.5 ± 2.4	76.6 ± 2.9	80.9 ± 2.2	78.9 ± 3.0	77.7 ± 2.8	82.6 ± 2.7
Consolidation	83.4 ± 2.9	83.8 ± 4.5	81.8 ± 4.0	90.3 ± 1.9	89.2 ± 2.5	86.4 ± 2.4	83.6 ± 3.9	83.7 ± 3.2	83.9 ± 2.8	85.0 ± 1.7
Pneumothorax	85.5 ± 8.2	85.9 ± 4.7	87.6 ± 5.6	89.3 ± 1.5	90.3 ± 5.2	86.7 ± 7.0	84.4 ± 8.6	87.8 ± 2.6	85.5 ± 4.5	86.8 ± 3.6
Pleural thickening	76.5 ± 3.6	79.3 ± 3.1	79.7 ± 4.8	82.9 ± 3.1	77.4 ± 5.2	80.5 ± 2.9	80.9 ± 3.5	84.7 ± 3.8	76.8 ± 4.4	83.9 ± 2.8
Lung opacity	72.2 ± 4.3	70.8 ± 3.7	71.6 ± 4.2	73.1 ± 4.4	71.4 ± 4.9	75.4 ± 3.6	73.1 ± 4.8	75.6 ± 4.3	73.7 ± 5.5	68.4 ± 5.8
Pulmonary fibrosis	76.4 ± 4.6	79.9 ± 2.3	76.5 ± 5.3	84.1 ± 5.2	78.0 ± 5.3	81.2 ± 4.2	84.2 ± 1.8	75.7 ± 6.4	74.3 ± 5.7	77.5 ± 2.9
Nodule/mass	78.9 ± 5.9	84.7 ± 3.9	76.4 ± 2.5	83.5 ± 2.3	78.2 ± 3.1	76.3 ± 2.6	78.7 ± 3.7	82.4 ± 2.5	74.7 ± 5.2	80.3 ± 3.7

Supplementary Table 3: Per-label performance on ChestX-ray14 dataset under full finetuning. Reported metrics include AUROC, accuracy, sensitivity, and specificity, given as mean $\pm$ standard (in percent) deviation across 1,000 bootstrap resamples for all 14 labels. Results are provided for ViT-B/16 and ConvNeXt-B backbones initialized from ImageNet, DINOv2, and DINOv3 at two input resolutions (224×224 and 512×512).

	ImageNet				DINOv2		DINOv3			
	ViT		ConvNeXt		ViT		ViT		ConvNeXt	
Label	224	512	224	512	224	512	224	512	224	512
AUROC										
Cardiomegaly	88.1 $\pm$ 0.5	88.0 $\pm$ 0.5	86.7 $\pm$ 0.6	87.7 $\pm$ 0.6	89.3 $\pm$ 0.5	88.0 $\pm$ 0.5	86.9 $\pm$ 0.6	89.0 $\pm$ 0.5	88.6 $\pm$ 0.5	87.5 $\pm$ 0.6
Effusion	80.5 $\pm$ 0.3	80.8 $\pm$ 0.3	80.5 $\pm$ 0.3	82.3 $\pm$ 0.3	81.2 $\pm$ 0.3	81.8 $\pm$ 0.3	81.6 $\pm$ 0.3	81.7 $\pm$ 0.3	82.0 $\pm$ 0.3	83.2 $\pm$ 0.3
Pneumonia	69.9 $\pm$ 1.1	70.7 $\pm$ 1.1	70.2 $\pm$ 1.1	71.5 $\pm$ 1.1	70.0 $\pm$ 1.1	69.7 $\pm$ 1.0	71.0 $\pm$ 1.1	72.7 $\pm$ 1.0	69.8 $\pm$ 1.1	72.7 $\pm$ 1.1
Atelectasis	74.0 $\pm$ 0.5	74.7 $\pm$ 0.4	75.6 $\pm$ 0.4	77.0 $\pm$ 0.4	74.8 $\pm$ 0.5	75.8 $\pm$ 0.4	75.8 $\pm$ 0.4	77.1 $\pm$ 0.4	75.7 $\pm$ 0.4	77.8 $\pm$ 0.4
No finding	72.0 $\pm$ 0.3	72.1 $\pm$ 0.3	71.7 $\pm$ 0.3	72.8 $\pm$ 0.3	72.9 $\pm$ 0.3	73.3 $\pm$ 0.3	73.0 $\pm$ 0.3	73.3 $\pm$ 0.3	72.7 $\pm$ 0.3	74.4 $\pm$ 0.3
Consolidation	72.7 $\pm$ 0.5	72.6 $\pm$ 0.6	73.3 $\pm$ 0.6	73.0 $\pm$ 0.5	74.6 $\pm$ 0.5	74.4 $\pm$ 0.6	72.8 $\pm$ 0.6	74.3 $\pm$ 0.5	74.4 $\pm$ 0.6	75.1 $\pm$ 0.5
Pneumothorax	84.3 $\pm$ 0.4	86.4 $\pm$ 0.4	86.2 $\pm$ 0.4	87.7 $\pm$ 0.3	86.6 $\pm$ 0.3	87.1 $\pm$ 0.3	86.0 $\pm$ 0.4	88.0 $\pm$ 0.3	86.0 $\pm$ 0.3	86.9 $\pm$ 0.3
Fibrosis	81.0 $\pm$ 0.9	80.9 $\pm$ 1.0	81.5 $\pm$ 1.0	81.5 $\pm$ 1.0	81.9 $\pm$ 0.9	80.8 $\pm$ 1.0	80.8 $\pm$ 1.0	83.3 $\pm$ 0.9	81.0 $\pm$ 1.0	83.1 $\pm$ 0.9
Emphysema	85.9 $\pm$ 0.6	90.1 $\pm$ 0.5	88.6 $\pm$ 0.5	92.2 $\pm$ 0.4	90.8 $\pm$ 0.5	92.2 $\pm$ 0.4	90.0 $\pm$ 0.5	93.0 $\pm$ 0.4	91.0 $\pm$ 0.5	93.0 $\pm$ 0.4
Hernia	90.4 $\pm$ 1.5	88.7 $\pm$ 2.0	92.4 $\pm$ 1.7	94.2 $\pm$ 1.6	87.5 $\pm$ 2.4	86.1 $\pm$ 2.0	90.3 $\pm$ 2.0	87.4 $\pm$ 2.4	91.3 $\pm$ 1.7	91.6 $\pm$ 2.3
Pleural thickening	74.2 $\pm$ 0.7	74.9 $\pm$ 0.7	75.9 $\pm$ 0.7	77.7 $\pm$ 0.7	75.1 $\pm$ 0.7	74.9 $\pm$ 0.7	75.5 $\pm$ 0.7	77.3 $\pm$ 0.7	75.0 $\pm$ 0.7	78.2 $\pm$ 0.7
Edema	83.6 $\pm$ 0.6	83.7 $\pm$ 0.6	84.3 $\pm$ 0.6	84.6 $\pm$ 0.6	84.5 $\pm$ 0.5	83.7 $\pm$ 0.6	84.9 $\pm$ 0.6	85.2 $\pm$ 0.5	83.4 $\pm$ 0.6	84.7 $\pm$ 0.6
Nodule	70.8 $\pm$ 0.7	72.3 $\pm$ 0.7	71.5 $\pm$ 0.6	77.7 $\pm$ 0.6	72.7 $\pm$ 0.7	73.7 $\pm$ 0.7	73.5 $\pm$ 0.6	75.8 $\pm$ 0.7	73.7 $\pm$ 0.6	79.4 $\pm$ 0.6
Mass	78.1 $\pm$ 0.6	77.7 $\pm$ 0.6	78.6 $\pm$ 0.6	81.9 $\pm$ 0.5	79.4 $\pm$ 0.6	78.7 $\pm$ 0.6	79.5 $\pm$ 0.6	80.6 $\pm$ 0.5	80.5 $\pm$ 0.6	84.2 $\pm$ 0.5
Accuracy										
Cardiomegaly	81.0 $\pm$ 2.2	79.1 $\pm$ 1.7	75.3 $\pm$ 2.9	79.7 $\pm$ 1.7	77.7 $\pm$ 2.2	78.2 $\pm$ 2.8	79.2 $\pm$ 2.0	79.8 $\pm$ 2.0	83.6 $\pm$ 1.4	80.8 $\pm$ 1.5
Effusion	71.5 $\pm$ 1.1	71.3 $\pm$ 1.3	69.9 $\pm$ 0.7	74.2 $\pm$ 1.1	72.3 $\pm$ 0.8	72.3 $\pm$ 0.9	72.4 $\pm$ 1.6	71.1 $\pm$ 1.4	70.3 $\pm$ 1.2	74.5 $\pm$ 1.1
Pneumonia	64.7 $\pm$ 1.2	63.1 $\pm$ 7.8	65.0 $\pm$ 7.5	62.6 $\pm$ 4.7	66.4 $\pm$ 4.2	60.8 $\pm$ 3.5	68.8 $\pm$ 5.3	64.0 $\pm$ 10.6	68.8 $\pm$ 3.7	60.9 $\pm$ 2.6
Atelectasis	67.6 $\pm$ 1.3	67.8 $\pm$ 1.9	69.4 $\pm$ 1.4	69.5 $\pm$ 1.8	68.5 $\pm$ 1.6	67.1 $\pm$ 3.3	67.3 $\pm$ 1.7	68.5 $\pm$ 0.8	67.1 $\pm$ 2.8	67.7 $\pm$ 1.9
No finding	70.0 $\pm$ 0.5	70.1 $\pm$ 0.7	70.1 $\pm$ 0.3	70.7 $\pm$ 0.5	70.7 $\pm$ 0.6	71.3 $\pm$ 0.4	71.1 $\pm$ 0.6	71.1 $\pm$ 0.5	70.7 $\pm$ 0.6	71.3 $\pm$ 0.7
Consolidation	60.9 $\pm$ 3.3	60.0 $\pm$ 2.3	61.7 $\pm$ 1.3	59.6 $\pm$ 2.2	61.3 $\pm$ 4.9	63.1 $\pm$ 2.4	57.2 $\pm$ 1.3	63.7 $\pm$ 3.1	62.0 $\pm$ 3.1	65.7 $\pm$ 2.3
Pneumothorax	74.7 $\pm$ 3.0	78.3 $\pm$ 1.6	76.5 $\pm$ 1.9	79.2 $\pm$ 0.7	78.2 $\pm$ 2.0	80.1 $\pm$ 1.7	78.4 $\pm$ 0.9	79.9 $\pm$ 1.8	77.1 $\pm$ 1.4	77.1 $\pm$ 2.3
Fibrosis	71.2 $\pm$ 3.3	71.9 $\pm$ 3.3	75.2 $\pm$ 4.0	74.1 $\pm$ 4.8	74.4 $\pm$ 3.0	74.5 $\pm$ 2.6	72.8 $\pm$ 1.2	74.8 $\pm$ 5.1	73.5 $\pm$ 4.6	75.0 $\pm$ 2.9
Emphysema	79.6 $\pm$ 2.0	84.9 $\pm$ 0.7	81.2 $\pm$ 0.7	85.6 $\pm$ 2.2	84.4 $\pm$ 2.8	84.8 $\pm$ 2.1	84.5 $\pm$ 2.3	86.9 $\pm$ 1.4	84.8 $\pm$ 1.6	89.0 $\pm$ 0.6
Hernia	78.0 $\pm$ 4.1	81.1 $\pm$ 4.8	91.5 $\pm$ 3.0	91.0 $\pm$ 4.5	89.4 $\pm$ 2.6	80.6 $\pm$ 7.8	91.5 $\pm$ 3.9	87.7 $\pm$ 5.1	89.6 $\pm$ 2.7	94.5 $\pm$ 2.2
Pleural thickening	67.1 $\pm$ 4.0	67.9 $\pm$ 5.3	63.9 $\pm$ 1.1	69.7 $\pm$ 3.0	69.3 $\pm$ 3.1	71.9 $\pm$ 2.5	67.7 $\pm$ 2.4	71.2 $\pm$ 2.8	71.9 $\pm$ 3.8	67.3 $\pm$ 2.4
Edema	73.3 $\pm$ 1.6	72.3 $\pm$ 1.9	74.1 $\pm$ 3.0	70.2 $\pm$ 1.8	72.7 $\pm$ 1.6	74.0 $\pm$ 2.7	72.3 $\pm$ 0.7	72.5 $\pm$ 1.7	75.3 $\pm$ 1.5	72.2 $\pm$ 2.3
Nodule	68.1 $\pm$ 3.2	70.7 $\pm$ 5.0	61.2 $\pm$ 3.1	72.9 $\pm$ 3.2	74.7 $\pm$ 2.4	72.0 $\pm$ 5.2	71.8 $\pm$ 5.0	71.5 $\pm$ 1.9	71.0 $\pm$ 2.8	75.9 $\pm$ 2.9
Mass	67.0 $\pm$ 3.8	74.0 $\pm$ 2.3	72.0 $\pm$ 3.8	77.9 $\pm$ 2.1	76.9 $\pm$ 4.1	73.8 $\pm$ 2.3	73.7 $\pm$ 0.9	76.8 $\pm$ 1.6	76.5 $\pm$ 2.4	79.6 $\pm$ 2.0
Sensitivity										
Cardiomegaly	78.8 $\pm$ 2.4	81.1 $\pm$ 2.1	81.8 $\pm$ 3.1	79.4 $\pm$ 2.2	86.1 $\pm$ 2.4	81.6 $\pm$ 3.0	78.5 $\pm$ 2.3	82.4 $\pm$ 2.3	78.5 $\pm$ 1.8	80.5 $\pm$ 1.8
Effusion	75.7 $\pm$ 1.7	76.6 $\pm$ 1.9	78.8 $\pm$ 1.1	76.2 $\pm$ 1.6	76.9 $\pm$ 1.3	77.2 $\pm$ 1.5	77.0 $\pm$ 2.4	78.6 $\pm$ 2.2	81.7 $\pm$ 1.8	77.0 $\pm$ 1.7
Pneumonia	67.3 $\pm$ 2.3	68.5 $\pm$ 8.1	65.8 $\pm$ 7.7	70.8 $\pm$ 5.3	64.2 $\pm$ 4.7	70.1 $\pm$ 3.9	64.3 $\pm$ 5.4	69.2 $\pm$ 10.8	62.3 $\pm$ 4.1	74.8 $\pm$ 3.2
Atelectasis	68.2 $\pm$ 1.8	69.4 $\pm$ 2.5	68.4 $\pm$ 1.9	70.9 $\pm$ 2.5	68.9 $\pm$ 2.1	71.9 $\pm$ 4.4	71.3 $\pm$ 2.3	72.6 $\pm$ 1.2	71.2 $\pm$ 3.6	74.2 $\pm$ 2.6
No finding	55.2 $\pm$ 2.0	54.9 $\pm$ 2.9	56.1 $\pm$ 0.9	56.2 $\pm$ 1.9	54.6 $\pm$ 2.2	54.2 $\pm$ 1.4	55.6 $\pm$ 2.1	55.2 $\pm$ 2.1	53.9 $\pm$ 2.3	57.6 $\pm$ 2.6

Consolidation	75.7 ± 3.8	75.7 ± 2.6	75.8 ± 1.7	78.0 ± 2.5	77.7 ± 5.6	76.0 ± 2.9	81.1 ± 1.6	75.1 ± 3.5	76.1 ± 3.6	73.2 ± 2.6
Pneumothorax	79.8 ± 3.6	80.5 ± 2.1	82.3 ± 2.4	82.6 ± 1.1	80.2 ± 2.6	80.2 ± 2.2	79.9 ± 1.3	82.8 ± 2.2	80.5 ± 1.7	82.5 ± 2.9
Fibrosis	78.3 ± 3.4	76.0 ± 3.7	73.8 ± 4.3	74.3 ± 4.7	77.1 ± 3.3	74.8 ± 3.1	77.1 ± 2.3	76.9 ± 5.2	74.9 ± 4.9	77.5 ± 3.2
Emphysema	77.9 ± 2.4	82.7 ± 1.2	81.8 ± 1.3	85.0 ± 2.3	82.1 ± 3.1	85.4 ± 2.3	80.8 ± 2.4	86.1 ± 1.6	83.1 ± 1.8	85.0 ± 1.2
Hernia	86.9 ± 4.7	82.5 ± 5.2	78.0 ± 4.8	83.3 ± 4.8	73.1 ± 5.2	75.8 ± 7.3	75.4 ± 5.0	72.7 ± 5.9	78.9 ± 4.6	80.0 ± 4.1
Pleural thickening	69.2 ± 4.3	69.5 ± 5.7	76.8 ± 1.5	73.4 ± 3.2	69.1 ± 3.4	66.4 ± 2.9	72.2 ± 2.7	71.5 ± 3.1	65.8 ± 4.0	77.7 ± 2.7
Edema	80.7 ± 2.0	82.3 ± 2.1	79.7 ± 3.1	84.6 ± 2.1	83.7 ± 1.9	79.9 ± 2.8	84.4 ± 1.3	84.5 ± 2.0	78.8 ± 2.0	85.4 ± 2.4
Nodule	62.0 ± 3.8	61.3 ± 5.7	70.8 ± 3.5	68.0 ± 3.7	58.3 ± 2.9	62.5 ± 6.0	61.3 ± 5.6	66.9 ± 2.4	63.4 ± 3.4	69.2 ± 3.4
Mass	75.5 ± 4.5	67.5 ± 2.7	70.8 ± 4.3	70.9 ± 2.6	66.2 ± 4.7	70.7 ± 2.8	71.0 ± 1.5	69.4 ± 2.0	68.1 ± 2.9	73.2 ± 2.4
Specificity										
Cardiomegaly	81.1 ± 2.4	79.0 ± 1.9	75.1 ± 3.1	79.7 ± 1.9	77.3 ± 2.4	78.1 ± 3.0	79.2 ± 2.2	79.7 ± 2.2	83.8 ± 1.5	80.8 ± 1.6
Effusion	70.6 ± 1.7	70.1 ± 1.9	67.9 ± 1.1	73.8 ± 1.6	71.2 ± 1.3	71.2 ± 1.4	71.4 ± 2.5	69.5 ± 2.2	67.8 ± 1.8	74.0 ± 1.7
Pneumonia	64.7 ± 1.2	63.0 ± 8.1	65.0 ± 7.9	62.5 ± 4.9	66.4 ± 4.4	60.6 ± 3.6	68.9 ± 5.5	63.9 ± 11.1	68.9 ± 3.9	60.6 ± 2.8
Atelectasis	67.5 ± 1.7	67.6 ± 2.5	69.5 ± 1.8	69.3 ± 2.4	68.5 ± 2.1	66.4 ± 4.4	66.8 ± 2.3	67.9 ± 1.0	66.5 ± 3.7	66.8 ± 2.5
No finding	79.3 ± 2.0	79.7 ± 2.8	79.0 ± 0.8	79.7 ± 1.9	80.7 ± 2.2	82.0 ± 1.4	80.8 ± 2.1	81.1 ± 2.0	81.2 ± 2.2	80.0 ± 2.6
Consolidation	59.8 ± 3.9	58.9 ± 2.7	60.6 ± 1.5	58.1 ± 2.5	60.0 ± 5.7	62.1 ± 2.8	55.4 ± 1.4	62.9 ± 3.6	61.0 ± 3.5	65.2 ± 2.6
Pneumothorax	74.1 ± 3.7	78.1 ± 2.0	75.8 ± 2.4	78.8 ± 0.9	78.0 ± 2.5	80.1 ± 2.2	78.2 ± 1.1	79.6 ± 2.3	76.7 ± 1.7	76.5 ± 2.9
Fibrosis	71.0 ± 3.4	71.9 ± 3.4	75.2 ± 4.2	74.1 ± 4.9	74.3 ± 3.1	74.5 ± 2.6	72.8 ± 1.3	74.8 ± 5.2	73.4 ± 4.8	74.9 ± 3.0
Emphysema	79.7 ± 2.2	85.0 ± 0.8	81.2 ± 0.8	85.6 ± 2.4	84.5 ± 3.1	84.8 ± 2.3	84.7 ± 2.5	86.9 ± 1.5	84.9 ± 1.7	89.1 ± 0.7
Hernia	77.9 ± 4.2	81.1 ± 4.8	91.5 ± 3.0	91.0 ± 4.5	89.5 ± 2.6	80.6 ± 7.9	91.6 ± 3.9	87.7 ± 5.1	89.6 ± 2.7	94.5 ± 2.2
Pleural thickening	67.0 ± 4.3	67.8 ± 5.8	63.3 ± 1.1	69.5 ± 3.2	69.3 ± 3.4	72.1 ± 2.8	67.5 ± 2.6	71.2 ± 3.1	72.2 ± 4.1	66.8 ± 2.6
Edema	73.0 ± 1.7	72.0 ± 2.1	73.9 ± 3.2	69.7 ± 1.9	72.3 ± 1.7	73.8 ± 2.9	71.9 ± 0.8	72.1 ± 1.8	75.1 ± 1.6	71.7 ± 2.5
Nodule	68.5 ± 3.6	71.3 ± 5.7	60.5 ± 3.5	73.2 ± 3.6	75.8 ± 2.7	72.7 ± 6.0	72.5 ± 5.7	71.9 ± 2.1	71.5 ± 3.2	76.3 ± 3.4
Mass	66.4 ± 4.4	74.5 ± 2.6	72.1 ± 4.3	78.4 ± 2.4	77.7 ± 4.7	74.0 ± 2.6	73.8 ± 1.1	77.4 ± 1.8	77.2 ± 2.8	80.0 ± 2.3

Supplementary Table 4: Per-label performance on PadChest dataset under full finetuning. Reported metrics include AUROC, accuracy, sensitivity, and specificity, given as mean $\pm$ standard (in percent) deviation across 1,000 bootstrap resamples for all 17 labels. Results are provided for ViT-B/16 and ConvNeXt-B backbones initialized from ImageNet, DINOv2, and DINOv3 at two input resolutions (224×224 and 512×512).

	ImageNet				DINOv2		DINOv3			
	ViT		ConvNeXt		ViT		ViT		ConvNeXt	
Label	224	512	224	512	224	512	224	512	224	512
AUROC										
Cardiomegaly	91.7 $\pm$ 0.3	92.3 $\pm$ 0.3	92.3 $\pm$ 0.2	92.3 $\pm$ 0.3	92.7 $\pm$ 0.2	92.6 $\pm$ 0.3	92.2 $\pm$ 0.3	92.5 $\pm$ 0.3	92.2 $\pm$ 0.3	92.6 $\pm$ 0.2
Pleural effusion	95.5 $\pm$ 0.3	95.6 $\pm$ 0.3	95.6 $\pm$ 0.3	95.6 $\pm$ 0.3	95.8 $\pm$ 0.2	95.8 $\pm$ 0.3	95.6 $\pm$ 0.3	96.0 $\pm$ 0.3	95.2 $\pm$ 0.3	96.2 $\pm$ 0.2
Pneumonia	84.6 $\pm$ 0.6	85.2 $\pm$ 0.7	84.9 $\pm$ 0.6	86.5 $\pm$ 0.6	85.8 $\pm$ 0.6	85.7 $\pm$ 0.6	85.3 $\pm$ 0.6	87.3 $\pm$ 0.6	86.3 $\pm$ 0.6	86.5 $\pm$ 0.6
Atelectasis	81.8 $\pm$ 0.6	85.4 $\pm$ 0.6	82.8 $\pm$ 0.6	85.8 $\pm$ 0.5	84.9 $\pm$ 0.6	87.5 $\pm$ 0.5	83.4 $\pm$ 0.6	87.3 $\pm$ 0.5	84.5 $\pm$ 0.6	87.8 $\pm$ 0.5
No finding	85.6 $\pm$ 0.2	86.6 $\pm$ 0.2	86.3 $\pm$ 0.2	87.0 $\pm$ 0.2	86.5 $\pm$ 0.2	87.4 $\pm$ 0.2	86.6 $\pm$ 0.2	87.3 $\pm$ 0.2	86.9 $\pm$ 0.2	87.2 $\pm$ 0.2
Consolidation	89.8 $\pm$ 0.7	90.3 $\pm$ 0.7	89.0 $\pm$ 0.8	90.1 $\pm$ 0.7	89.7 $\pm$ 0.7	90.4 $\pm$ 0.8	90.0 $\pm$ 0.7	91.0 $\pm$ 0.7	90.1 $\pm$ 0.8	90.6 $\pm$ 0.6
Pneumothorax	88.7 $\pm$ 1.7	88.1 $\pm$ 1.8	87.0 $\pm$ 2.1	87.8 $\pm$ 2.0	89.4 $\pm$ 2.0	88.4 $\pm$ 1.9	87.0 $\pm$ 2.3	89.8 $\pm$ 1.6	89.2 $\pm$ 1.7	90.7 $\pm$ 1.3
Emphysema	88.7 $\pm$ 1.0	87.9 $\pm$ 1.1	88.5 $\pm$ 1.1	87.6 $\pm$ 1.1	89.7 $\pm$ 1.0	89.5 $\pm$ 1.1	87.7 $\pm$ 1.2	89.4 $\pm$ 1.0	88.9 $\pm$ 1.0	90.5 $\pm$ 1.0
Hernia	93.1 $\pm$ 0.9	93.5 $\pm$ 0.9	94.0 $\pm$ 0.8	95.0 $\pm$ 0.8	94.8 $\pm$ 0.7	94.7 $\pm$ 0.8	93.6 $\pm$ 0.9	94.5 $\pm$ 0.8	94.8 $\pm$ 0.8	96.6 $\pm$ 0.6
Scoliosis	87.9 $\pm$ 0.5	89.0 $\pm$ 0.5	90.4 $\pm$ 0.5	90.4 $\pm$ 0.5	90.2 $\pm$ 0.5	89.9 $\pm$ 0.5	90.5 $\pm$ 0.5	91.1 $\pm$ 0.5	90.6 $\pm$ 0.4	92.2 $\pm$ 0.4
Congestion	90.1 $\pm$ 1.1	88.8 $\pm$ 1.3	90.1 $\pm$ 1.3	90.5 $\pm$ 1.1	89.6 $\pm$ 1.2	89.8 $\pm$ 1.2	89.1 $\pm$ 1.3	89.9 $\pm$ 1.2	89.2 $\pm$ 1.2	90.0 $\pm$ 1.2
Aortic elongation	90.9 $\pm$ 0.3	90.9 $\pm$ 0.3	91.0 $\pm$ 0.3	91.5 $\pm$ 0.3	91.5 $\pm$ 0.3	91.5 $\pm$ 0.3	90.5 $\pm$ 0.3	91.7 $\pm$ 0.3	90.8 $\pm$ 0.3	91.6 $\pm$ 0.3
Kyphosis	89.2 $\pm$ 0.6	88.7 $\pm$ 0.8	89.7 $\pm$ 0.7	89.5 $\pm$ 0.6	89.4 $\pm$ 0.7	88.7 $\pm$ 0.7	89.6 $\pm$ 0.7	89.9 $\pm$ 0.6	89.4 $\pm$ 0.7	89.9 $\pm$ 0.7
COPD signs	84.2 $\pm$ 0.4	84.2 $\pm$ 0.4	84.8 $\pm$ 0.3	85.3 $\pm$ 0.3	84.4 $\pm$ 0.4	84.9 $\pm$ 0.4	84.8 $\pm$ 0.3	85.0 $\pm$ 0.3	84.3 $\pm$ 0.4	85.0 $\pm$ 0.3
Pleural thickening	83.6 $\pm$ 0.8	84.8 $\pm$ 0.7	84.5 $\pm$ 0.8	87.8 $\pm$ 0.7	84.6 $\pm$ 0.8	87.2 $\pm$ 0.7	83.4 $\pm$ 0.8	87.1 $\pm$ 0.7	85.8 $\pm$ 0.8	89.0 $\pm$ 0.6
Nodule mass	73.3 $\pm$ 0.8	74.4 $\pm$ 0.9	73.7 $\pm$ 0.8	77.6 $\pm$ 0.8	74.6 $\pm$ 0.8	77.0 $\pm$ 0.8	74.4 $\pm$ 0.8	77.5 $\pm$ 0.8	75.2 $\pm$ 0.8	78.7 $\pm$ 0.8
Infiltrates	81.0 $\pm$ 0.7	82.7 $\pm$ 0.7	82.0 $\pm$ 0.7	83.4 $\pm$ 0.6	82.4 $\pm$ 0.7	83.1 $\pm$ 0.6	82.0 $\pm$ 0.7	84.7 $\pm$ 0.6	82.6 $\pm$ 0.6	83.5 $\pm$ 0.6
Accuracy										
Cardiomegaly	81.2 $\pm$ 1.2	83.8 $\pm$ 1.1	82.6 $\pm$ 1.3	82.8 $\pm$ 1.4	82.9 $\pm$ 0.8	83.5 $\pm$ 1.1	84.1 $\pm$ 1.5	83.8 $\pm$ 1.6	82.6 $\pm$ 1.1	83.0 $\pm$ 0.8
Pleural effusion	89.2 $\pm$ 0.9	89.4 $\pm$ 0.7	88.7 $\pm$ 0.8	90.3 $\pm$ 0.9	89.4 $\pm$ 0.7	88.8 $\pm$ 0.9	90.7 $\pm$ 1.2	90.2 $\pm$ 1.2	89.3 $\pm$ 0.7	90.9 $\pm$ 1.0
Pneumonia	78.1 $\pm$ 1.0	78.1 $\pm$ 2.1	75.8 $\pm$ 2.3	77.5 $\pm$ 1.9	78.4 $\pm$ 3.4	79.2 $\pm$ 1.9	75.1 $\pm$ 1.7	81.0 $\pm$ 2.5	80.0 $\pm$ 2.5	76.2 $\pm$ 1.2
Atelectasis	75.0 $\pm$ 1.1	77.0 $\pm$ 1.7	73.7 $\pm$ 2.9	77.4 $\pm$ 3.0	75.9 $\pm$ 2.0	80.2 $\pm$ 1.7	74.8 $\pm$ 2.8	78.4 $\pm$ 1.2	74.8 $\pm$ 2.1	80.9 $\pm$ 1.2
No finding	77.0 $\pm$ 0.5	77.1 $\pm$ 0.9	77.2 $\pm$ 0.6	78.2 $\pm$ 0.4	77.3 $\pm$ 0.7	78.0 $\pm$ 0.6	77.3 $\pm$ 0.7	78.0 $\pm$ 0.6	78.1 $\pm$ 0.4	77.5 $\pm$ 0.5
Consolidation	80.8 $\pm$ 1.9	81.8 $\pm$ 2.4	77.9 $\pm$ 1.8	78.5 $\pm$ 2.0	76.5 $\pm$ 2.9	80.3 $\pm$ 2.2	79.3 $\pm$ 2.8	80.1 $\pm$ 2.2	78.0 $\pm$ 2.2	79.2 $\pm$ 2.3
Pneumothorax	82.2 $\pm$ 2.5	81.4 $\pm$ 4.0	79.4 $\pm$ 4.1	81.1 $\pm$ 8.1	83.4 $\pm$ 3.8	81.5 $\pm$ 5.1	87.2 $\pm$ 3.3	82.5 $\pm$ 6.8	83.1 $\pm$ 5.7	80.6 $\pm$ 2.4
Emphysema	84.3 $\pm$ 3.0	84.2 $\pm$ 1.4	82.1 $\pm$ 4.8	82.5 $\pm$ 4.5	82.5 $\pm$ 5.1	84.6 $\pm$ 3.0	80.0 $\pm$ 2.6	84.2 $\pm$ 4.9	79.2 $\pm$ 5.3	82.5 $\pm$ 1.8
Hernia	92.1 $\pm$ 2.4	93.8 $\pm$ 1.8	92.7 $\pm$ 2.9	92.6 $\pm$ 1.4	92.4 $\pm$ 2.3	93.3 $\pm$ 2.4	92.3 $\pm$ 1.9	93.1 $\pm$ 2.6	92.9 $\pm$ 1.3	92.5 $\pm$ 1.8
Scoliosis	82.2 $\pm$ 2.9	82.0 $\pm$ 3.2	81.4 $\pm$ 2.9	81.6 $\pm$ 1.8	82.7 $\pm$ 1.1	82.0 $\pm$ 1.3	84.6 $\pm$ 2.3	83.2 $\pm$ 2.8	83.6 $\pm$ 1.3	83.4 $\pm$ 1.6
Congestion	83.6 $\pm$ 1.5	79.8 $\pm$ 2.5	85.7 $\pm$ 1.0	84.0 $\pm$ 2.5	82.2 $\pm$ 2.9	83.7 $\pm$ 4.5	81.2 $\pm$ 5.1	84.2 $\pm$ 3.1	78.0 $\pm$ 6.6	81.7 $\pm$ 5.3
Aortic elongation	80.6 $\pm$ 1.4	79.2 $\pm$ 1.3	80.4 $\pm$ 1.1	81.1 $\pm$ 1.5	80.1 $\pm$ 1.7	81.3 $\pm$ 1.1	77.8 $\pm$ 1.4	80.5 $\pm$ 1.6	78.9 $\pm$ 1.9	80.7 $\pm$ 1.8
Kyphosis	82.8 $\pm$ 2.8	80.5 $\pm$ 2.2	80.2 $\pm$ 4.8	78.7 $\pm$ 2.3	80.2 $\pm$ 3.1	80.5 $\pm$ 2.1	81.5 $\pm$ 2.7	81.0 $\pm$ 2.0	79.7 $\pm$ 2.0	78.8 $\pm$ 1.7
COPD signs	73.4 $\pm$ 2.2	73.7 $\pm$ 0.8	72.9 $\pm$ 1.9	74.7 $\pm$ 1.1	73.8 $\pm$ 1.6	73.9 $\pm$ 2.6	74.4 $\pm$ 1.6	73.8 $\pm$ 0.7	73.8 $\pm$ 2.0	73.7 $\pm$ 2.5
Pleural thickening	76.3 $\pm$ 2.9	79.1 $\pm$ 3.9	80.4 $\pm$ 2.4	82.0 $\pm$ 2.9	76.0 $\pm$ 5.2	83.0 $\pm$ 1.5	77.6 $\pm$ 3.6	82.0 $\pm$ 3.9	79.2 $\pm$ 1.0	81.0 $\pm$ 2.2
Nodule mass	66.2 $\pm$ 1.9	67.8 $\pm$ 5.1	67.9 $\pm$ 4.8	74.5 $\pm$ 3.1	68.4 $\pm$ 4.5	69.2 $\pm$ 2.4	67.5 $\pm$ 3.1	76.3 $\pm$ 3.7	70.6 $\pm$ 2.4	76.2 $\pm$ 3.8
Infiltrates	71.2 $\pm$ 1.7	72.5 $\pm$ 2.7	75.7 $\pm$ 0.3	74.0 $\pm$ 2.1	73.2 $\pm$ 2.1	75.2 $\pm$ 1.9	72.7 $\pm$ 1.9	75.2 $\pm$ 0.8	71.6 $\pm$ 2.0	71.9 $\pm$ 2.0

Sensitivity										
Cardiomegaly	87.8 ± 1.6	85.8 ± 1.4	87.3 ± 1.6	87.5 ± 1.6	88.7 ± 1.0	87.6 ± 1.3	85.5 ± 1.8	87.0 ± 1.9	87.6 ± 1.4	88.5 ± 1.1
Pleural effusion	91.2 ± 1.1	91.9 ± 1.0	92.5 ± 1.0	91.5 ± 1.1	92.5 ± 0.9	93.3 ± 1.1	90.3 ± 1.4	92.1 ± 1.4	91.9 ± 1.0	91.5 ± 1.1
Pneumonia	78.0 ± 1.6	79.5 ± 2.4	79.5 ± 2.6	82.4 ± 2.1	78.1 ± 3.7	78.6 ± 2.3	81.7 ± 2.0	80.0 ± 2.8	77.0 ± 2.8	83.7 ± 1.7
Atelectasis	74.4 ± 1.6	79.7 ± 2.1	77.2 ± 3.3	78.5 ± 3.4	79.6 ± 2.5	80.2 ± 2.0	76.7 ± 3.2	81.5 ± 1.7	80.1 ± 2.5	81.3 ± 1.6
No finding	80.1 ± 1.1	82.4 ± 2.3	81.7 ± 1.5	81.8 ± 0.9	81.8 ± 1.9	83.3 ± 1.4	82.5 ± 1.9	83.5 ± 1.6	81.5 ± 0.9	84.7 ± 1.3
Consolidation	86.2 ± 2.3	85.4 ± 2.9	88.5 ± 2.2	90.3 ± 2.3	89.1 ± 3.0	86.8 ± 2.5	87.3 ± 2.9	89.1 ± 2.7	89.3 ± 2.5	89.7 ± 2.3
Pneumothorax	80.4 ± 4.5	80.8 ± 4.9	80.7 ± 5.7	77.7 ± 7.7	83.2 ± 5.0	80.7 ± 5.8	73.1 ± 5.5	80.9 ± 6.4	77.7 ± 6.2	87.0 ± 3.8
Emphysema	78.0 ± 3.5	77.2 ± 3.0	79.6 ± 4.9	76.4 ± 4.7	81.3 ± 5.2	79.8 ± 3.6	80.5 ± 3.6	78.0 ± 5.1	81.8 ± 5.3	85.4 ± 2.5
Hernia	80.5 ± 2.9	82.5 ± 2.5	82.5 ± 3.1	87.1 ± 2.0	84.6 ± 2.8	84.0 ± 2.7	83.0 ± 2.4	84.4 ± 2.8	85.8 ± 2.1	90.3 ± 2.2
Scoliosis	77.5 ± 3.2	78.7 ± 3.4	82.9 ± 3.1	83.9 ± 2.1	82.9 ± 1.6	83.0 ± 1.8	80.4 ± 2.6	83.6 ± 3.0	81.3 ± 1.7	86.5 ± 1.9
Congestion	84.7 ± 2.9	85.9 ± 3.2	83.8 ± 3.0	83.3 ± 3.2	84.7 ± 3.5	82.1 ± 4.5	82.4 ± 4.9	82.3 ± 3.6	84.9 ± 6.6	83.2 ± 5.2
Aortic elongation	87.6 ± 1.6	89.3 ± 1.6	87.3 ± 1.4	88.1 ± 1.7	89.1 ± 2.0	88.3 ± 1.3	89.6 ± 1.7	88.8 ± 1.9	88.3 ± 2.2	88.0 ± 2.1
Kyphosis	81.0 ± 3.0	83.2 ± 2.7	84.5 ± 4.7	85.6 ± 2.5	84.5 ± 3.4	82.7 ± 2.6	83.6 ± 2.9	84.3 ± 2.4	84.8 ± 2.5	88.0 ± 2.1
COPD signs	80.2 ± 2.9	80.7 ± 1.1	82.3 ± 2.5	81.3 ± 1.6	80.9 ± 2.1	80.5 ± 3.4	80.3 ± 2.2	82.3 ± 1.1	79.9 ± 2.7	81.4 ± 3.2
Pleural thickening	75.0 ± 3.2	74.2 ± 4.1	72.9 ± 2.7	77.7 ± 3.1	76.6 ± 5.4	75.8 ± 2.2	72.6 ± 4.0	75.6 ± 4.1	77.8 ± 1.9	82.1 ± 2.5
Nodule mass	69.2 ± 2.5	68.8 ± 5.6	67.0 ± 5.3	67.5 ± 3.6	67.1 ± 5.2	71.7 ± 3.0	69.8 ± 3.5	64.3 ± 4.1	67.5 ± 2.9	67.1 ± 4.2
Infiltrates	79.0 ± 2.2	79.4 ± 3.0	77.0 ± 1.4	80.0 ± 2.4	79.4 ± 2.6	78.7 ± 2.2	77.9 ± 2.2	80.7 ± 1.5	81.3 ± 2.4	82.9 ± 2.3
Specificity										
Cardiomegaly	80.5 ± 1.4	83.6 ± 1.4	82.1 ± 1.5	82.4 ± 1.6	82.4 ± 0.9	83.1 ± 1.3	84.0 ± 1.8	83.4 ± 1.9	82.2 ± 1.3	82.5 ± 0.9
Pleural effusion	89.1 ± 1.0	89.2 ± 0.8	88.5 ± 0.9	90.2 ± 1.0	89.2 ± 0.8	88.5 ± 1.1	90.7 ± 1.4	90.1 ± 1.3	89.1 ± 0.8	90.8 ± 1.1
Pneumonia	78.1 ± 1.1	78.1 ± 2.3	75.6 ± 2.5	77.3 ± 2.1	78.4 ± 3.7	79.2 ± 2.1	74.8 ± 1.9	81.0 ± 2.7	80.1 ± 2.7	75.9 ± 1.3
Atelectasis	75.1 ± 1.2	76.8 ± 1.9	73.5 ± 3.3	77.3 ± 3.4	75.7 ± 2.3	80.3 ± 1.9	74.6 ± 3.1	78.2 ± 1.4	74.5 ± 2.4	80.9 ± 1.4
No finding	75.5 ± 1.1	74.6 ± 2.4	75.0 ± 1.5	76.3 ± 0.9	75.1 ± 1.9	75.4 ± 1.4	74.8 ± 1.9	75.4 ± 1.6	76.4 ± 0.9	74.1 ± 1.3
Consolidation	80.7 ± 2.0	81.7 ± 2.4	77.8 ± 1.9	78.3 ± 2.0	76.4 ± 3.0	80.2 ± 2.2	79.2 ± 2.9	80.0 ± 2.2	77.9 ± 2.2	79.0 ± 2.3
Pneumothorax	82.3 ± 2.5	81.4 ± 4.0	79.4 ± 4.1	81.1 ± 8.1	83.4 ± 3.8	81.5 ± 5.2	87.2 ± 3.3	82.5 ± 6.9	83.1 ± 5.8	80.5 ± 2.4
Emphysema	84.3 ± 3.0	84.2 ± 1.5	82.1 ± 4.9	82.6 ± 4.6	82.5 ± 5.2	84.7 ± 3.0	80.0 ± 2.7	84.2 ± 5.0	79.2 ± 5.4	82.4 ± 1.8
Hernia	92.3 ± 2.5	94.0 ± 1.8	92.9 ± 3.0	92.7 ± 1.4	92.5 ± 2.4	93.5 ± 2.5	92.5 ± 2.0	93.3 ± 2.7	93.0 ± 1.4	92.5 ± 1.9
Scoliosis	82.5 ± 3.2	82.2 ± 3.5	81.4 ± 3.1	81.5 ± 1.9	82.7 ± 1.3	82.0 ± 1.4	84.8 ± 2.5	83.1 ± 3.1	83.8 ± 1.5	83.2 ± 1.8
Congestion	83.6 ± 1.5	79.8 ± 2.5	85.8 ± 1.1	84.0 ± 2.5	82.2 ± 3.0	83.8 ± 4.6	81.2 ± 5.2	84.2 ± 3.1	78.0 ± 6.7	81.7 ± 5.4
Aortic elongation	80.1 ± 1.6	78.3 ± 1.6	79.8 ± 1.3	80.6 ± 1.7	79.4 ± 2.0	80.7 ± 1.2	76.9 ± 1.7	79.9 ± 1.9	78.2 ± 2.2	80.1 ± 2.0
Kyphosis	82.8 ± 2.9	80.5 ± 2.4	80.0 ± 5.0	78.6 ± 2.4	80.1 ± 3.3	80.5 ± 2.2	81.4 ± 2.8	81.0 ± 2.1	79.6 ± 2.0	78.5 ± 1.8
COPD signs	72.4 ± 2.9	72.6 ± 1.0	71.5 ± 2.5	73.7 ± 1.5	72.7 ± 2.1	72.9 ± 3.5	73.5 ± 2.2	72.5 ± 0.9	72.9 ± 2.7	72.5 ± 3.3
Pleural thickening	76.3 ± 3.1	79.2 ± 4.1	80.6 ± 2.5	82.1 ± 3.1	76.0 ± 5.5	83.2 ± 1.6	77.7 ± 3.9	82.2 ± 4.1	79.2 ± 1.1	81.0 ± 2.3
Nodule mass	66.0 ± 2.1	67.8 ± 5.6	68.0 ± 5.3	74.8 ± 3.4	68.5 ± 4.9	69.1 ± 2.7	67.4 ± 3.4	76.8 ± 4.1	70.8 ± 2.6	76.6 ± 4.2
Infiltrates	70.9 ± 1.9	72.2 ± 2.9	75.6 ± 0.3	73.8 ± 2.3	72.9 ± 2.3	75.0 ± 2.1	72.5 ± 2.0	75.0 ± 0.8	71.2 ± 2.2	71.5 ± 2.2

Supplementary Table 5: Per-label performance on the CheXpert dataset under full finetuning. Reported metrics include AUROC, accuracy, sensitivity, and specificity, given as mean $\pm$ standard (in percent) deviation across 1,000 bootstrap resamples for all 10 labels. Results are provided for ViT-B/16 and ConvNeXt-B backbones initialized from ImageNet, DINOv2, and DINOv3 at two input resolutions (224×224 and 512×512).

	ImageNet				DINOv2		DINOv3			
	ViT		ConvNeXt		ViT		ViT		ConvNeXt	
Label	224	512	224	512	224	512	224	512	224	512
AUROC										
Cardiomegaly	87.3 $\pm$ 0.3	88.1 $\pm$ 0.3	87.0 $\pm$ 0.3	87.7 $\pm$ 0.3	87.8 $\pm$ 0.3	88.4 $\pm$ 0.3	87.3 $\pm$ 0.3	88.3 $\pm$ 0.3	88.0 $\pm$ 0.3	88.7 $\pm$ 0.3
Pleural effusion	87.5 $\pm$ 0.2	88.3 $\pm$ 0.2	87.2 $\pm$ 0.2	88.4 $\pm$ 0.2	88.3 $\pm$ 0.2	88.7 $\pm$ 0.2	87.6 $\pm$ 0.2	88.5 $\pm$ 0.2	88.1 $\pm$ 0.2	88.9 $\pm$ 0.2
Pneumonia	77.8 $\pm$ 0.8	78.9 $\pm$ 0.8	76.4 $\pm$ 0.9	78.2 $\pm$ 0.9	77.3 $\pm$ 0.8	78.2 $\pm$ 0.8	77.8 $\pm$ 0.9	79.5 $\pm$ 0.8	77.6 $\pm$ 0.8	79.5 $\pm$ 0.8
Atelectasis	69.0 $\pm$ 0.4	69.8 $\pm$ 0.4	69.0 $\pm$ 0.4	70.5 $\pm$ 0.4	69.9 $\pm$ 0.4	71.4 $\pm$ 0.4	69.2 $\pm$ 0.4	70.7 $\pm$ 0.4	70.3 $\pm$ 0.4	71.5 $\pm$ 0.4
No finding	87.5 $\pm$ 0.3	88.0 $\pm$ 0.3	87.7 $\pm$ 0.3	88.1 $\pm$ 0.3	88.1 $\pm$ 0.3	88.5 $\pm$ 0.3	88.0 $\pm$ 0.3	88.5 $\pm$ 0.3	88.1 $\pm$ 0.3	88.7 $\pm$ 0.3
Consolidation	74.8 $\pm$ 0.6	75.0 $\pm$ 0.6	74.1 $\pm$ 0.6	75.1 $\pm$ 0.6	75.3 $\pm$ 0.6	75.6 $\pm$ 0.5	74.6 $\pm$ 0.6	75.8 $\pm$ 0.6	75.2 $\pm$ 0.6	75.7 $\pm$ 0.6
Pneumothorax	86.3 $\pm$ 0.3	88.5 $\pm$ 0.3	86.5 $\pm$ 0.3	90.6 $\pm$ 0.3	87.3 $\pm$ 0.3	90.7 $\pm$ 0.3	87.5 $\pm$ 0.3	90.9 $\pm$ 0.3	88.2 $\pm$ 0.3	91.6 $\pm$ 0.3
Lung opacity	73.5 $\pm$ 0.3	74.7 $\pm$ 0.3	73.5 $\pm$ 0.3	74.7 $\pm$ 0.3	74.2 $\pm$ 0.3	74.9 $\pm$ 0.3	74.0 $\pm$ 0.3	75.0 $\pm$ 0.3	74.3 $\pm$ 0.3	75.2 $\pm$ 0.3
Lung lesion	76.3 $\pm$ 0.7	77.5 $\pm$ 0.7	76.8 $\pm$ 0.7	80.9 $\pm$ 0.7	76.8 $\pm$ 0.7	79.4 $\pm$ 0.7	76.2 $\pm$ 0.7	80.3 $\pm$ 0.7	78.3 $\pm$ 0.7	82.2 $\pm$ 0.7
Fracture	76.8 $\pm$ 0.7	77.5 $\pm$ 0.7	77.0 $\pm$ 0.7	81.4 $\pm$ 0.6	77.9 $\pm$ 0.7	78.1 $\pm$ 0.7	77.6 $\pm$ 0.7	81.4 $\pm$ 0.6	76.6 $\pm$ 0.7	83.0 $\pm$ 0.6
Accuracy										
Cardiomegaly	80.8 $\pm$ 1.4	79.5 $\pm$ 0.5	79.5 $\pm$ 1.1	80.0 $\pm$ 1.1	80.3 $\pm$ 1.3	80.8 $\pm$ 1.2	80.4 $\pm$ 0.9	80.4 $\pm$ 0.7	81.0 $\pm$ 1.0	81.0 $\pm$ 1.0
Pleural effusion	78.9 $\pm$ 0.4	79.9 $\pm$ 0.4	78.8 $\pm$ 0.3	80.1 $\pm$ 0.5	79.9 $\pm$ 0.5	80.2 $\pm$ 0.3	79.6 $\pm$ 0.3	80.0 $\pm$ 0.3	79.6 $\pm$ 0.3	80.3 $\pm$ 0.3
Pneumonia	71.5 $\pm$ 3.3	72.7 $\pm$ 2.6	73.2 $\pm$ 3.2	68.8 $\pm$ 3.5	73.3 $\pm$ 5.3	73.5 $\pm$ 3.9	67.4 $\pm$ 4.5	73.3 $\pm$ 3.8	71.8 $\pm$ 4.0	76.3 $\pm$ 3.1
Atelectasis	59.0 $\pm$ 2.1	58.7 $\pm$ 1.9	58.3 $\pm$ 5.2	59.1 $\pm$ 3.0	58.8 $\pm$ 1.7	61.0 $\pm$ 2.6	59.1 $\pm$ 1.1	57.9 $\pm$ 3.4	58.6 $\pm$ 2.0	60.7 $\pm$ 2.4
No finding	78.9 $\pm$ 0.8	80.3 $\pm$ 1.1	80.4 $\pm$ 0.8	79.8 $\pm$ 1.2	80.3 $\pm$ 1.1	81.3 $\pm$ 0.4	79.9 $\pm$ 0.7	81.0 $\pm$ 0.9	81.4 $\pm$ 0.5	81.1 $\pm$ 0.7
Consolidation	59.7 $\pm$ 2.1	62.7 $\pm$ 3.0	65.2 $\pm$ 3.1	62.9 $\pm$ 3.4	60.7 $\pm$ 3.3	64.7 $\pm$ 3.4	61.7 $\pm$ 2.0	64.7 $\pm$ 3.5	60.6 $\pm$ 2.8	59.7 $\pm$ 3.6
Pneumothorax	78.9 $\pm$ 0.8	81.4 $\pm$ 1.5	79.3 $\pm$ 1.4	83.9 $\pm$ 1.1	82.2 $\pm$ 1.0	85.0 $\pm$ 0.7	79.5 $\pm$ 1.4	84.5 $\pm$ 0.9	79.9 $\pm$ 1.6	84.9 $\pm$ 0.7
Lung opacity	66.9 $\pm$ 0.3	67.5 $\pm$ 0.3	66.8 $\pm$ 0.5	67.4 $\pm$ 0.4	67.3 $\pm$ 0.3	67.3 $\pm$ 0.4	66.9 $\pm$ 0.4	67.9 $\pm$ 0.3	67.2 $\pm$ 0.4	68.0 $\pm$ 0.3
Lung lesion	68.3 $\pm$ 3.5	64.4 $\pm$ 3.0	66.8 $\pm$ 3.7	70.5 $\pm$ 5.0	63.8 $\pm$ 3.6	73.7 $\pm$ 5.2	66.4 $\pm$ 4.3	73.3 $\pm$ 4.6	69.1 $\pm$ 3.5	75.5 $\pm$ 3.7
Fracture	71.5 $\pm$ 3.5	71.7 $\pm$ 2.6	68.4 $\pm$ 4.4	75.2 $\pm$ 4.7	70.8 $\pm$ 2.2	72.2 $\pm$ 5.3	68.2 $\pm$ 3.4	77.2 $\pm$ 1.6	71.2 $\pm$ 3.7	73.8 $\pm$ 3.3
Sensitivity										
Cardiomegaly	78.6 $\pm$ 1.8	82.9 $\pm$ 0.9	80.2 $\pm$ 1.5	80.8 $\pm$ 1.5	81.1 $\pm$ 1.9	81.3 $\pm$ 1.7	79.1 $\pm$ 1.3	81.9 $\pm$ 1.1	80.7 $\pm$ 1.3	82.2 $\pm$ 1.5
Pleural effusion	84.0 $\pm$ 1.1	83.0 $\pm$ 1.3	84.0 $\pm$ 0.9	84.0 $\pm$ 1.8	83.9 $\pm$ 1.7	84.3 $\pm$ 0.9	82.4 $\pm$ 1.2	84.7 $\pm$ 1.1	84.0 $\pm$ 1.0	84.5 $\pm$ 0.8
Pneumonia	70.2 $\pm$ 3.7	70.9 $\pm$ 3.0	67.2 $\pm$ 3.5	73.5 $\pm$ 3.9	67.0 $\pm$ 5.5	68.6 $\pm$ 4.3	75.6 $\pm$ 4.9	71.0 $\pm$ 4.2	70.8 $\pm$ 4.3	68.6 $\pm$ 3.4
Atelectasis	70.8 $\pm$ 3.0	72.4 $\pm$ 2.8	71.1 $\pm$ 7.4	72.8 $\pm$ 4.2	72.5 $\pm$ 2.5	71.8 $\pm$ 3.6	71.3 $\pm$ 1.5	75.0 $\pm$ 4.9	73.8 $\pm$ 2.9	72.4 $\pm$ 3.4
No finding	83.5 $\pm$ 1.1	82.7 $\pm$ 1.4	82.1 $\pm$ 1.2	83.5 $\pm$ 1.5	82.4 $\pm$ 1.5	82.5 $\pm$ 0.8	83.0 $\pm$ 1.0	82.7 $\pm$ 1.3	81.9 $\pm$ 0.8	83.5 $\pm$ 1.1
Consolidation	78.8 $\pm$ 2.5	75.6 $\pm$ 3.3	71.3 $\pm$ 3.5	75.5 $\pm$ 3.9	77.6 $\pm$ 3.6	73.8 $\pm$ 3.8	76.7 $\pm$ 2.3	74.6 $\pm$ 3.9	78.3 $\pm$ 3.2	79.4 $\pm$ 4.0
Pneumothorax	79.2 $\pm$ 1.2	80.1 $\pm$ 1.9	78.9 $\pm$ 1.8	83.0 $\pm$ 1.4	76.7 $\pm$ 1.4	81.2 $\pm$ 1.0	80.5 $\pm$ 1.8	82.7 $\pm$ 1.3	81.4 $\pm$ 2.1	83.8 $\pm$ 1.0
Lung opacity	76.1 $\pm$ 1.4	77.7 $\pm$ 1.3	74.9 $\pm$ 2.9	76.9 $\pm$ 2.3	74.9 $\pm$ 1.7	78.7 $\pm$ 2.0	76.7 $\pm$ 2.2	75.9 $\pm$ 1.5	77.6 $\pm$ 2.7	76.4 $\pm$ 1.3
Lung lesion	70.3 $\pm$ 4.0	76.9 $\pm$ 3.3	73.5 $\pm$ 4.1	75.7 $\pm$ 5.4	75.8 $\pm$ 4.0	69.6 $\pm$ 5.6	71.8 $\pm$ 4.6	71.6 $\pm$ 4.9	72.6 $\pm$ 3.8	72.9 $\pm$ 4.1
Fracture	68.2 $\pm$ 3.9	69.9 $\pm$ 2.9	71.5 $\pm$ 4.7	72.4 $\pm$ 5.2	71.2 $\pm$ 2.7	69.3 $\pm$ 5.8	73.8 $\pm$ 3.7	71.0 $\pm$ 2.1	67.7 $\pm$ 4.0	76.3 $\pm$ 3.8
Specificity										
Cardiomegaly	81.2 $\pm$ 1.9	79.0 $\pm$ 0.7	79.4 $\pm$ 1.5	79.8 $\pm$ 1.5	80.2 $\pm$ 1.8	80.7 $\pm$ 1.6	80.6 $\pm$ 1.2	80.2 $\pm$ 0.9	81.0 $\pm$ 1.3	80.8 $\pm$ 1.4

Pleural effusion	75.7 ± 1.2	77.9 ± 1.3	75.5 ± 0.9	77.7 ± 1.8	77.3 ± 1.8	77.6 ± 0.9	77.8 ± 1.2	77.0 ± 1.1	76.8 ± 1.1	77.6 ± 0.8
Pneumonia	71.6 ± 3.5	72.8 ± 2.7	73.4 ± 3.4	68.7 ± 3.7	73.5 ± 5.6	73.6 ± 4.1	67.1 ± 4.8	73.4 ± 4.0	71.8 ± 4.2	76.5 ± 3.3
Atelectasis	56.8 ± 3.0	56.2 ± 2.8	56.0 ± 7.5	56.6 ± 4.3	56.3 ± 2.5	59.0 ± 3.7	56.8 ± 1.5	54.8 ± 4.9	55.8 ± 2.9	58.6 ± 3.4
No finding	78.2 ± 1.0	80.0 ± 1.4	80.2 ± 1.1	79.3 ± 1.5	80.0 ± 1.4	81.1 ± 0.6	79.5 ± 0.9	80.7 ± 1.2	81.4 ± 0.7	80.8 ± 0.9
Consolidation	58.4 ± 2.4	61.9 ± 3.4	64.8 ± 3.5	62.1 ± 3.9	59.6 ± 3.7	64.2 ± 3.9	60.8 ± 2.3	64.0 ± 3.9	59.5 ± 3.1	58.4 ± 4.1
Pneumothorax	78.9 ± 1.1	81.6 ± 1.9	79.4 ± 1.8	84.0 ± 1.4	82.8 ± 1.3	85.4 ± 0.8	79.4 ± 1.8	84.7 ± 1.1	79.8 ± 2.1	85.0 ± 0.8
Lung opacity	59.7 ± 1.4	59.5 ± 1.3	60.4 ± 2.9	59.9 ± 2.3	61.5 ± 1.7	58.5 ± 2.0	59.3 ± 2.2	61.6 ± 1.5	59.1 ± 2.7	61.3 ± 1.4
Lung lesion	68.3 ± 3.8	64.0 ± 3.3	66.5 ± 4.0	70.2 ± 5.4	63.4 ± 3.9	73.8 ± 5.6	66.2 ± 4.6	73.4 ± 5.0	69.0 ± 3.8	75.6 ± 4.0
Fracture	71.6 ± 3.8	71.8 ± 2.8	68.2 ± 4.8	75.4 ± 5.1	70.7 ± 2.4	72.3 ± 5.8	68.0 ± 3.7	77.5 ± 1.8	71.4 ± 4.1	73.7 ± 3.6

Supplementary Table 6: Per-label performance on the MIMIC-CXR dataset under full finetuning. Reported metrics include AUROC, accuracy, sensitivity, and specificity, given as mean $\pm$ standard (in percent) deviation across 1,000 bootstrap resamples for all 10 labels. Results are provided for ViT-B/16 and ConvNeXt-B backbones initialized from ImageNet, DINOv2, and DINOv3 at two input resolutions (224×224 and 512×512).

Label	ImageNet				DINOv2		DINOv3		DINOv3	
	ViT		ConvNeXt		ViT		ViT		ConvNeXt	
Label	224	512	224	512	224	512	224	512	224	512
AUROC										
Cardiomegaly	81.2 $\pm$ 0.2	81.6 $\pm$ 0.2	80.7 $\pm$ 0.2	81.3 $\pm$ 0.2	81.5 $\pm$ 0.2	81.3 $\pm$ 0.2	81.1 $\pm$ 0.2	81.5 $\pm$ 0.2	81.4 $\pm$ 0.2	81.7 $\pm$ 0.2
Pleural effusion	90.7 $\pm$ 0.1	91.3 $\pm$ 0.1	90.7 $\pm$ 0.1	91.2 $\pm$ 0.1	91.0 $\pm$ 0.1	91.1 $\pm$ 0.2	90.7 $\pm$ 0.1	91.4 $\pm$ 0.1	91.1 $\pm$ 0.1	91.6 $\pm$ 0.1
Pneumonia	74.0 $\pm$ 0.5	75.2 $\pm$ 0.5	74.5 $\pm$ 0.5	75.9 $\pm$ 0.5	74.7 $\pm$ 0.5	74.5 $\pm$ 0.5	74.7 $\pm$ 0.5	76.5 $\pm$ 0.5	75.3 $\pm$ 0.5	76.9 $\pm$ 0.5
Atelectasis	81.6 $\pm$ 0.2	82.6 $\pm$ 0.2	81.7 $\pm$ 0.2	82.6 $\pm$ 0.2	81.8 $\pm$ 0.2	82.3 $\pm$ 0.2	81.7 $\pm$ 0.2	82.7 $\pm$ 0.2	82.0 $\pm$ 0.2	83.1 $\pm$ 0.2
No finding	85.4 $\pm$ 0.2	86.0 $\pm$ 0.2	85.3 $\pm$ 0.2	86.0 $\pm$ 0.2	85.6 $\pm$ 0.2	85.7 $\pm$ 0.2	85.5 $\pm$ 0.2	86.2 $\pm$ 0.2	85.8 $\pm$ 0.2	86.3 $\pm$ 0.2
Consolidation	81.9 $\pm$ 0.4	83.0 $\pm$ 0.4	82.1 $\pm$ 0.4	83.2 $\pm$ 0.4	82.5 $\pm$ 0.4	82.3 $\pm$ 0.4	82.4 $\pm$ 0.4	83.6 $\pm$ 0.4	82.5 $\pm$ 0.4	83.9 $\pm$ 0.4
Pneumothorax	86.0 $\pm$ 0.5	88.7 $\pm$ 0.4	86.8 $\pm$ 0.4	89.7 $\pm$ 0.4	87.6 $\pm$ 0.4	87.7 $\pm$ 0.4	87.0 $\pm$ 0.4	90.1 $\pm$ 0.4	87.8 $\pm$ 0.4	90.5 $\pm$ 0.4
Lung opacity	76.4 $\pm$ 0.3	77.6 $\pm$ 0.3	76.9 $\pm$ 0.3	78.4 $\pm$ 0.3	77.1 $\pm$ 0.3	77.1 $\pm$ 0.3	77.1 $\pm$ 0.3	78.5 $\pm$ 0.3	77.5 $\pm$ 0.3	78.7 $\pm$ 0.2
Lung lesion	74.1 $\pm$ 0.8	76.3 $\pm$ 0.8	74.8 $\pm$ 0.8	78.8 $\pm$ 0.7	75.7 $\pm$ 0.8	76.4 $\pm$ 0.8	75.5 $\pm$ 0.8	78.8 $\pm$ 0.8	76.9 $\pm$ 0.8	80.7 $\pm$ 0.7
Fracture	67.8 $\pm$ 0.9	71.5 $\pm$ 0.8	70.7 $\pm$ 0.9	74.5 $\pm$ 0.9	71.1 $\pm$ 0.9	68.8 $\pm$ 0.9	72.0 $\pm$ 0.8	75.6 $\pm$ 0.8	71.8 $\pm$ 0.9	73.4 $\pm$ 0.9
Accuracy										
Cardiomegaly	69.8 $\pm$ 1.1	69.8 $\pm$ 0.6	70.2 $\pm$ 1.1	70.8 $\pm$ 0.7	70.1 $\pm$ 0.9	70.2 $\pm$ 0.9	69.9 $\pm$ 0.9	69.2 $\pm$ 1.0	69.5 $\pm$ 0.7	69.8 $\pm$ 1.3
Pleural effusion	81.5 $\pm$ 0.9	82.1 $\pm$ 0.4	81.6 $\pm$ 0.4	81.7 $\pm$ 1.1	81.4 $\pm$ 0.7	82.3 $\pm$ 0.5	81.3 $\pm$ 0.8	81.8 $\pm$ 0.6	82.0 $\pm$ 0.4	82.2 $\pm$ 0.4
Pneumonia	70.0 $\pm$ 3.1	73.5 $\pm$ 2.4	66.7 $\pm$ 2.4	69.7 $\pm$ 2.5	71.0 $\pm$ 2.9	72.0 $\pm$ 2.2	73.4 $\pm$ 2.1	71.3 $\pm$ 1.7	71.1 $\pm$ 3.1	70.9 $\pm$ 2.0
Atelectasis	70.9 $\pm$ 0.5	71.5 $\pm$ 1.0	71.3 $\pm$ 1.0	72.8 $\pm$ 1.2	71.1 $\pm$ 1.0	70.3 $\pm$ 0.6	71.5 $\pm$ 0.6	72.8 $\pm$ 0.9	71.0 $\pm$ 0.5	73.6 $\pm$ 0.7
No finding	78.8 $\pm$ 0.2	79.1 $\pm$ 0.2	78.5 $\pm$ 0.5	79.5 $\pm$ 0.3	78.5 $\pm$ 0.3	78.6 $\pm$ 0.4	78.7 $\pm$ 0.5	79.7 $\pm$ 0.3	79.0 $\pm$ 0.4	79.8 $\pm$ 0.4
Consolidation	65.6 $\pm$ 2.4	70.0 $\pm$ 1.8	68.4 $\pm$ 2.1	67.5 $\pm$ 3.0	70.6 $\pm$ 2.6	67.8 $\pm$ 2.8	68.8 $\pm$ 1.6	69.4 $\pm$ 1.5	67.2 $\pm$ 1.8	69.8 $\pm$ 2.0
Pneumothorax	78.2 $\pm$ 3.0	82.5 $\pm$ 2.2	78.6 $\pm$ 1.7	83.7 $\pm$ 2.8	79.1 $\pm$ 1.4	79.2 $\pm$ 2.0	80.8 $\pm$ 2.8	85.7 $\pm$ 2.6	81.5 $\pm$ 2.2	86.4 $\pm$ 0.9
Lung opacity	66.7 $\pm$ 1.7	67.1 $\pm$ 1.7	66.7 $\pm$ 1.0	67.2 $\pm$ 0.8	67.1 $\pm$ 1.5	66.2 $\pm$ 1.0	66.1 $\pm$ 1.6	68.1 $\pm$ 1.1	67.1 $\pm$ 1.7	67.8 $\pm$ 1.0
Lung lesion	76.5 $\pm$ 3.6	76.0 $\pm$ 5.4	73.9 $\pm$ 5.4	76.4 $\pm$ 4.2	73.6 $\pm$ 4.3	73.6 $\pm$ 4.5	77.2 $\pm$ 1.7	74.6 $\pm$ 3.1	76.2 $\pm$ 2.7	76.0 $\pm$ 4.1
Fracture	67.5 $\pm$ 3.9	68.7 $\pm$ 6.4	65.3 $\pm$ 6.8	70.2 $\pm$ 5.2	72.9 $\pm$ 7.5	67.7 $\pm$ 8.2	69.4 $\pm$ 4.6	74.2 $\pm$ 3.0	69.7 $\pm$ 4.3	73.9 $\pm$ 2.7
Sensitivity										
Cardiomegaly	81.0 $\pm$ 1.9	81.8 $\pm$ 1.0	79.0 $\pm$ 1.7	79.5 $\pm$ 1.2	81.1 $\pm$ 1.5	80.1 $\pm$ 1.5	80.5 $\pm$ 1.5	82.3 $\pm$ 1.7	81.4 $\pm$ 1.1	81.3 $\pm$ 2.1
Pleural effusion	85.9 $\pm$ 1.6	86.6 $\pm$ 0.7	86.0 $\pm$ 0.7	86.7 $\pm$ 2.0	87.0 $\pm$ 1.2	85.4 $\pm$ 0.9	86.2 $\pm$ 1.5	87.6 $\pm$ 1.1	86.6 $\pm$ 0.7	87.8 $\pm$ 0.7
Pneumonia	65.2 $\pm$ 3.6	62.8 $\pm$ 2.8	69.7 $\pm$ 2.9	69.0 $\pm$ 3.0	65.0 $\pm$ 3.3	63.4 $\pm$ 2.6	62.5 $\pm$ 2.5	68.5 $\pm$ 2.1	65.2 $\pm$ 3.6	69.5 $\pm$ 2.3
Atelectasis	80.9 $\pm$ 0.9	82.1 $\pm$ 1.7	80.3 $\pm$ 1.7	80.4 $\pm$ 1.9	80.9 $\pm$ 1.6	83.2 $\pm$ 1.0	80.4 $\pm$ 0.9	80.4 $\pm$ 1.4	81.6 $\pm$ 0.9	80.4 $\pm$ 1.1
No finding	77.1 $\pm$ 0.7	78.6 $\pm$ 0.7	77.7 $\pm$ 1.8	77.5 $\pm$ 0.9	78.9 $\pm$ 0.8	78.6 $\pm$ 1.3	77.5 $\pm$ 1.7	76.7 $\pm$ 1.0	77.7 $\pm$ 1.5	77.5 $\pm$ 1.2
Consolidation	85.2 $\pm$ 2.6	82.4 $\pm$ 1.9	81.9 $\pm$ 2.3	84.6 $\pm$ 3.2	80.1 $\pm$ 2.8	82.6 $\pm$ 3.1	82.4 $\pm$ 1.8	83.5 $\pm$ 1.7	83.7 $\pm$ 2.0	84.1 $\pm$ 2.2
Pneumothorax	77.0 $\pm$ 3.3	78.4 $\pm$ 2.5	77.7 $\pm$ 2.0	79.1 $\pm$ 3.1	78.5 $\pm$ 1.7	78.6 $\pm$ 2.3	75.2 $\pm$ 3.2	77.6 $\pm$ 2.9	76.7 $\pm$ 2.4	77.9 $\pm$ 1.2
Lung opacity	74.6 $\pm$ 2.7	75.8 $\pm$ 2.7	76.0 $\pm$ 1.6	77.9 $\pm$ 1.4	75.2 $\pm$ 2.3	76.3 $\pm$ 1.6	77.2 $\pm$ 2.6	76.9 $\pm$ 1.7	75.7 $\pm$ 2.6	77.8 $\pm$ 1.7
Lung lesion	58.9 $\pm$ 4.0	62.6 $\pm$ 5.5	62.5 $\pm$ 5.9	67.0 $\pm$ 4.3	64.8 $\pm$ 4.7	65.6 $\pm$ 4.8	61.1 $\pm$ 2.2	69.2 $\pm$ 3.5	64.7 $\pm$ 3.2	70.7 $\pm$ 4.4
Fracture	58.6 $\pm$ 4.3	62.7 $\pm$ 6.6	64.9 $\pm$ 7.1	65.5 $\pm$ 5.5	58.1 $\pm$ 7.8	58.9 $\pm$ 8.5	63.1 $\pm$ 4.9	64.6 $\pm$ 3.5	63.1 $\pm$ 4.5	62.2 $\pm$ 3.2
Specificity										
Cardiomegaly	66.9 $\pm$ 1.9	66.7 $\pm$ 1.0	67.9 $\pm$ 1.8	68.5 $\pm$ 1.1	67.2 $\pm$ 1.5	67.6 $\pm$ 1.6	67.2 $\pm$ 1.6	65.8 $\pm$ 1.7	66.4 $\pm$ 1.1	66.9 $\pm$ 2.1
Pleural effusion	80.2 $\pm$ 1.6	80.8 $\pm$ 0.7	80.3 $\pm$ 0.7	80.2 $\pm$ 2.0	79.7 $\pm$ 1.2	81.3 $\pm$ 0.8	79.8 $\pm$ 1.5	80.0 $\pm$ 1.1	80.6 $\pm$ 0.6	80.5 $\pm$ 0.7

Pneumonia	70.4 ± 3.6	74.3 ± 2.7	66.5 ± 2.8	69.7 ± 2.9	71.4 ± 3.3	72.6 ± 2.6	74.2 ± 2.4	71.5 ± 2.0	71.5 ± 3.6	71.0 ± 2.3
Atelectasis	68.4 ± 0.8	68.9 ± 1.7	69.0 ± 1.7	70.9 ± 2.0	68.7 ± 1.6	67.1 ± 0.9	69.3 ± 0.9	70.9 ± 1.4	68.3 ± 0.8	71.9 ± 1.1
No finding	79.9 ± 0.6	79.3 ± 0.6	78.9 ± 1.8	80.7 ± 0.8	78.2 ± 0.8	78.6 ± 1.3	79.4 ± 1.8	81.4 ± 1.0	79.8 ± 1.5	81.0 ± 1.2
Consolidation	64.8 ± 2.6	69.5 ± 1.9	67.9 ± 2.3	66.8 ± 3.3	70.2 ± 2.8	67.2 ± 3.1	68.3 ± 1.7	68.9 ± 1.6	66.5 ± 1.9	69.2 ± 2.2
Pneumothorax	78.3 ± 3.4	82.6 ± 2.4	78.7 ± 1.8	83.9 ± 3.0	79.1 ± 1.5	79.3 ± 2.2	81.0 ± 3.1	86.1 ± 2.8	81.7 ± 2.4	86.8 ± 0.9
Lung opacity	64.8 ± 2.7	65.0 ± 2.8	64.5 ± 1.6	64.7 ± 1.3	65.2 ± 2.3	63.9 ± 1.6	63.5 ± 2.6	66.0 ± 1.7	65.1 ± 2.7	65.4 ± 1.6
Lung lesion	77.0 ± 3.8	76.4 ± 5.7	74.2 ± 5.7	76.6 ± 4.4	73.8 ± 4.6	73.8 ± 4.8	77.7 ± 1.8	74.8 ± 3.3	76.5 ± 2.9	76.2 ± 4.3
Fracture	67.7 ± 4.1	68.8 ± 6.7	65.4 ± 7.1	70.3 ± 5.5	73.2 ± 7.8	67.9 ± 8.6	69.6 ± 4.8	74.4 ± 3.1	69.8 ± 4.5	74.2 ± 2.8

Supplementary Table 7: Per-label performance on the UKA-CXR dataset under full finetuning. Reported metrics include AUROC, accuracy, sensitivity, and specificity, given as mean $\pm$ standard (in percent) deviation across 1,000 bootstrap resamples for all 6 labels. Results are provided for ViT-B/16 and ConvNeXt-B backbones initialized from ImageNet, DINOv2, and DINOv3 at two input resolutions (224×224 and 512×512).

	ImageNet				DINOv2		DINOv3			
	ViT		ConvNeXt		ViT		ViT		ConvNeXt	
Label	224	512	224	512	224	512	224	512	224	512
AUROC										
Cardiomegaly	85.3 $\pm$ 0.2	85.6 $\pm$ 0.2	85.1 $\pm$ 0.2	85.8 $\pm$ 0.2	85.9 $\pm$ 0.2	85.5 $\pm$ 0.2	85.4 $\pm$ 0.2	85.9 $\pm$ 0.2	86.0 $\pm$ 0.2	86.3 $\pm$ 0.2
Congestion	84.3 $\pm$ 0.3	84.8 $\pm$ 0.3	85.1 $\pm$ 0.3	85.4 $\pm$ 0.3	84.9 $\pm$ 0.3	84.8 $\pm$ 0.3	84.8 $\pm$ 0.3	85.7 $\pm$ 0.3	84.8 $\pm$ 0.3	85.6 $\pm$ 0.3
Pleural effusion	91.5 $\pm$ 0.2	91.6 $\pm$ 0.2	91.3 $\pm$ 0.2	91.8 $\pm$ 0.2	91.5 $\pm$ 0.2	91.7 $\pm$ 0.2	91.6 $\pm$ 0.2	91.8 $\pm$ 0.2	91.8 $\pm$ 0.2	91.9 $\pm$ 0.2
Pneumonic infiltrates	92.2 $\pm$ 0.2	92.6 $\pm$ 0.2	92.5 $\pm$ 0.2	93.3 $\pm$ 0.2	92.5 $\pm$ 0.2	92.6 $\pm$ 0.2	92.7 $\pm$ 0.2	93.4 $\pm$ 0.2	92.7 $\pm$ 0.2	93.2 $\pm$ 0.2
Atelectasis	86.4 $\pm$ 0.2	86.2 $\pm$ 0.2	86.3 $\pm$ 0.2	86.6 $\pm$ 0.2	86.6 $\pm$ 0.2	86.6 $\pm$ 0.2	86.5 $\pm$ 0.2	86.8 $\pm$ 0.2	86.7 $\pm$ 0.2	86.8 $\pm$ 0.2
Healthy	86.6 $\pm$ 0.2	86.9 $\pm$ 0.2	86.7 $\pm$ 0.2	87.2 $\pm$ 0.2	87.0 $\pm$ 0.2	86.7 $\pm$ 0.2	86.8 $\pm$ 0.2	87.2 $\pm$ 0.2	87.0 $\pm$ 0.2	87.3 $\pm$ 0.2
Accuracy										
Cardiomegaly	76.7 $\pm$ 0.2	76.9 $\pm$ 0.2	76.2 $\pm$ 0.3	77.0 $\pm$ 0.2	77.1 $\pm$ 0.2	76.9 $\pm$ 0.2	76.5 $\pm$ 0.2	77.0 $\pm$ 0.3	77.3 $\pm$ 0.2	77.3 $\pm$ 0.2
Congestion	72.8 $\pm$ 2.0	73.2 $\pm$ 1.9	75.2 $\pm$ 1.2	74.3 $\pm$ 1.7	72.4 $\pm$ 3.2	73.5 $\pm$ 0.9	73.1 $\pm$ 1.1	75.0 $\pm$ 1.2	73.1 $\pm$ 2.3	74.5 $\pm$ 2.1
Pleural effusion	81.3 $\pm$ 1.0	80.7 $\pm$ 0.6	81.1 $\pm$ 1.7	82.5 $\pm$ 0.6	80.5 $\pm$ 1.7	81.1 $\pm$ 1.3	81.4 $\pm$ 0.7	81.2 $\pm$ 1.5	82.4 $\pm$ 0.7	82.6 $\pm$ 1.2
Pneumonic infiltrates	83.4 $\pm$ 0.8	84.3 $\pm$ 0.5	85.7 $\pm$ 1.1	86.0 $\pm$ 0.9	83.8 $\pm$ 0.6	83.9 $\pm$ 0.7	84.2 $\pm$ 1.7	84.4 $\pm$ 0.6	82.8 $\pm$ 1.1	83.7 $\pm$ 0.6
Atelectasis	75.9 $\pm$ 1.0	75.5 $\pm$ 1.1	75.2 $\pm$ 1.1	75.0 $\pm$ 0.8	74.8 $\pm$ 1.0	75.8 $\pm$ 0.9	75.2 $\pm$ 0.8	74.9 $\pm$ 0.4	74.5 $\pm$ 1.5	76.3 $\pm$ 1.4
Healthy	78.4 $\pm$ 0.4	78.1 $\pm$ 0.6	78.2 $\pm$ 0.4	78.7 $\pm$ 0.3	78.7 $\pm$ 0.2	78.0 $\pm$ 0.3	78.3 $\pm$ 0.8	78.4 $\pm$ 0.4	78.3 $\pm$ 0.4	78.6 $\pm$ 0.3
Sensitivity										
Cardiomegaly	77.2 $\pm$ 2.0	76.3 $\pm$ 1.2	78.5 $\pm$ 2.4	80.0 $\pm$ 1.6	79.2 $\pm$ 1.2	78.6 $\pm$ 0.8	77.4 $\pm$ 2.0	78.3 $\pm$ 2.1	77.9 $\pm$ 1.0	81.6 $\pm$ 1.4
Congestion	81.2 $\pm$ 2.4	82.1 $\pm$ 2.3	79.5 $\pm$ 1.5	81.5 $\pm$ 2.0	82.5 $\pm$ 3.8	81.8 $\pm$ 1.1	81.4 $\pm$ 1.3	81.7 $\pm$ 1.5	81.5 $\pm$ 2.8	81.7 $\pm$ 2.5
Pleural effusion	86.7 $\pm$ 1.3	88.1 $\pm$ 0.9	86.5 $\pm$ 2.2	86.4 $\pm$ 0.8	87.3 $\pm$ 2.2	87.9 $\pm$ 1.7	87.4 $\pm$ 1.0	87.9 $\pm$ 1.9	86.3 $\pm$ 0.9	85.9 $\pm$ 1.6
Pneumonic infiltrates	85.2 $\pm$ 1.2	85.2 $\pm$ 0.8	82.8 $\pm$ 1.6	84.6 $\pm$ 1.2	85.4 $\pm$ 0.9	85.3 $\pm$ 1.0	84.7 $\pm$ 2.3	86.9 $\pm$ 0.8	87.1 $\pm$ 1.5	87.9 $\pm$ 0.8
Atelectasis	81.6 $\pm$ 1.4	81.7 $\pm$ 1.6	82.7 $\pm$ 1.6	83.7 $\pm$ 1.1	84.0 $\pm$ 1.3	82.2 $\pm$ 1.2	82.8 $\pm$ 1.1	84.4 $\pm$ 0.7	84.4 $\pm$ 2.1	81.9 $\pm$ 1.8
Healthy	76.9 $\pm$ 1.3	78.6 $\pm$ 2.4	78.3 $\pm$ 1.4	78.7 $\pm$ 1.1	76.4 $\pm$ 0.7	78.3 $\pm$ 0.9	77.5 $\pm$ 3.1	79.1 $\pm$ 1.6	78.5 $\pm$ 1.5	78.6 $\pm$ 1.0
Specificity										
Cardiomegaly	76.3 $\pm$ 1.9	77.5 $\pm$ 1.2	74.2 $\pm$ 2.5	74.4 $\pm$ 1.6	75.3 $\pm$ 1.2	75.3 $\pm$ 0.8	75.8 $\pm$ 2.0	75.9 $\pm$ 2.2	76.7 $\pm$ 1.1	73.5 $\pm$ 1.4
Congestion	72.0 $\pm$ 2.4	72.4 $\pm$ 2.2	74.8 $\pm$ 1.4	73.7 $\pm$ 2.0	71.5 $\pm$ 3.8	72.8 $\pm$ 1.1	72.4 $\pm$ 1.3	74.4 $\pm$ 1.5	72.3 $\pm$ 2.8	73.9 $\pm$ 2.5
Pleural effusion	80.5 $\pm$ 1.3	79.7 $\pm$ 0.8	80.3 $\pm$ 2.2	81.9 $\pm$ 0.8	79.5 $\pm$ 2.3	80.1 $\pm$ 1.7	80.5 $\pm$ 1.0	80.3 $\pm$ 2.0	81.8 $\pm$ 0.9	82.1 $\pm$ 1.6
Pneumonic infiltrates	83.0 $\pm$ 1.1	84.1 $\pm$ 0.7	86.2 $\pm$ 1.6	86.3 $\pm$ 1.2	83.5 $\pm$ 0.9	83.6 $\pm$ 1.0	84.1 $\pm$ 2.4	83.9 $\pm$ 0.8	82.1 $\pm$ 1.5	82.9 $\pm$ 0.8
Atelectasis	75.0 $\pm$ 1.3	74.5 $\pm$ 1.6	74.0 $\pm$ 1.5	73.6 $\pm$ 1.1	73.3 $\pm$ 1.4	74.8 $\pm$ 1.2	74.0 $\pm$ 1.0	73.3 $\pm$ 0.6	72.8 $\pm$ 2.1	75.3 $\pm$ 1.9
Healthy	79.4 $\pm$ 1.3	77.8 $\pm$ 2.4	78.1 $\pm$ 1.4	78.7 $\pm$ 1.1	80.2 $\pm$ 0.6	77.8 $\pm$ 0.9	78.7 $\pm$ 3.1	78.0 $\pm$ 1.6	78.2 $\pm$ 1.5	78.7 $\pm$ 1.0