

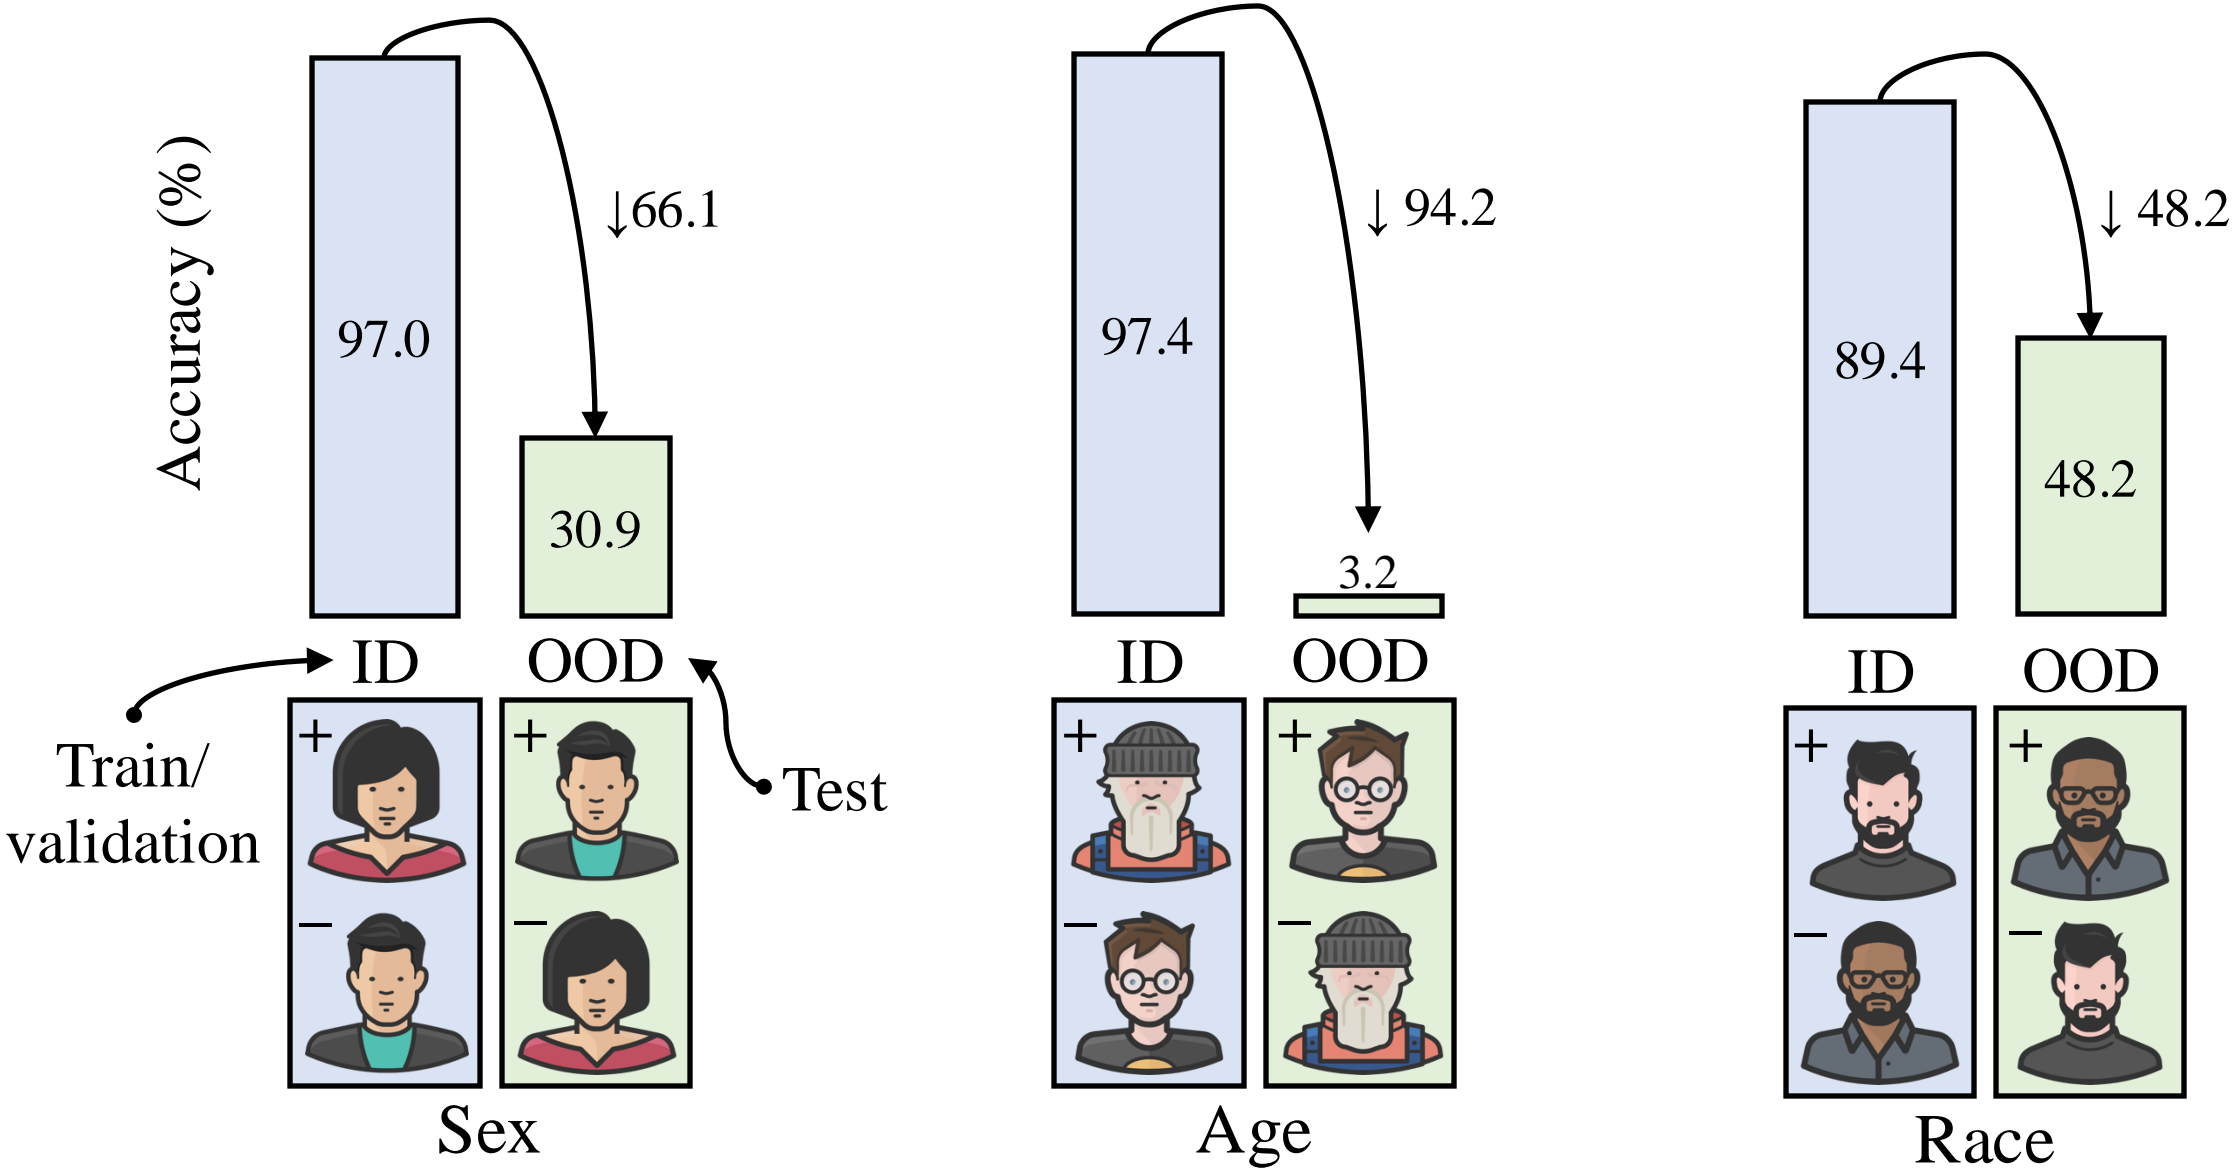


A Textbook Remedy for Domain Shifts: Knowledge Priors for Medical Image Analysis

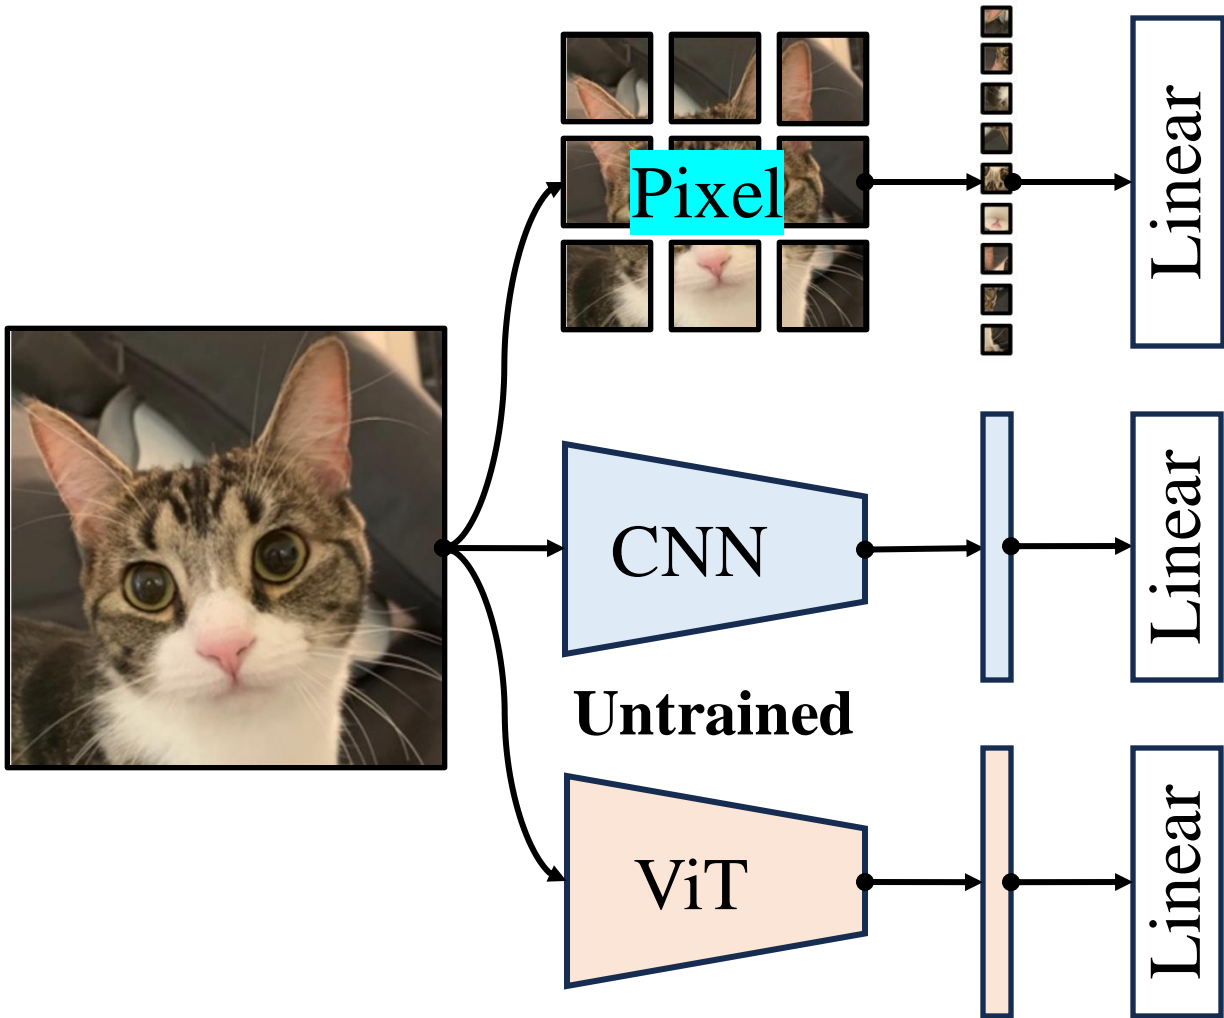
Yue Yang, Mona Gandhi, Yufei Wang, Yifan Wu, Michael S. Yao,
Chris Callison-Burch, James C. Gee, Mark Yatskar



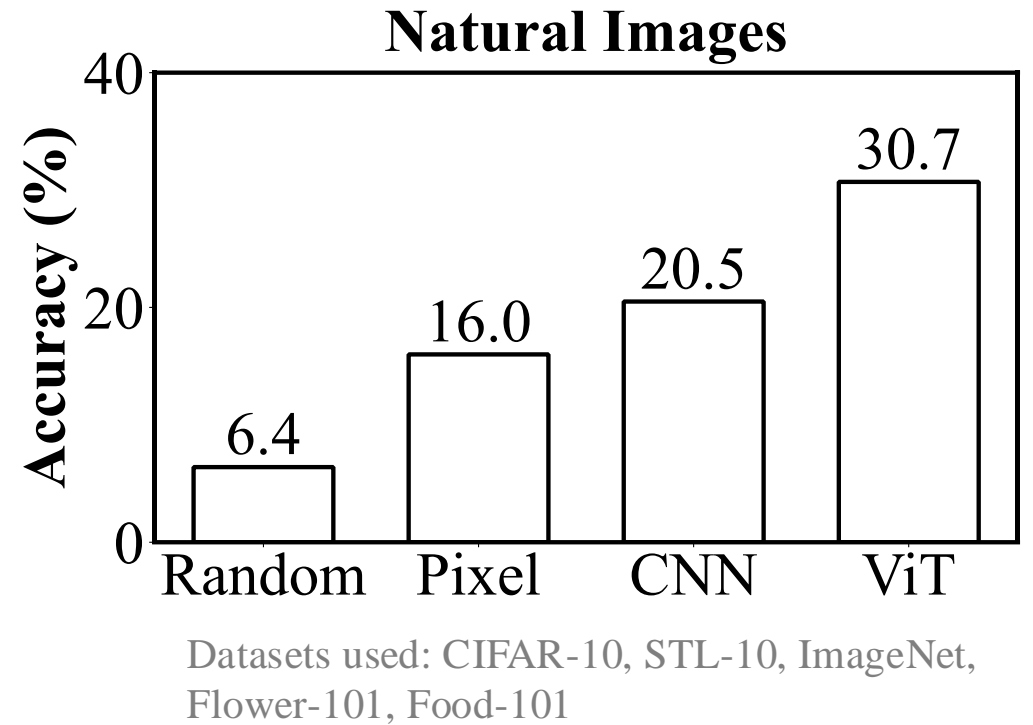
Black-box models generalize poorly on medical domain shifts.



Deep models have good priors for the **general domain**.



Vision backbones have a **deep image prior**, even when entirely untrained.



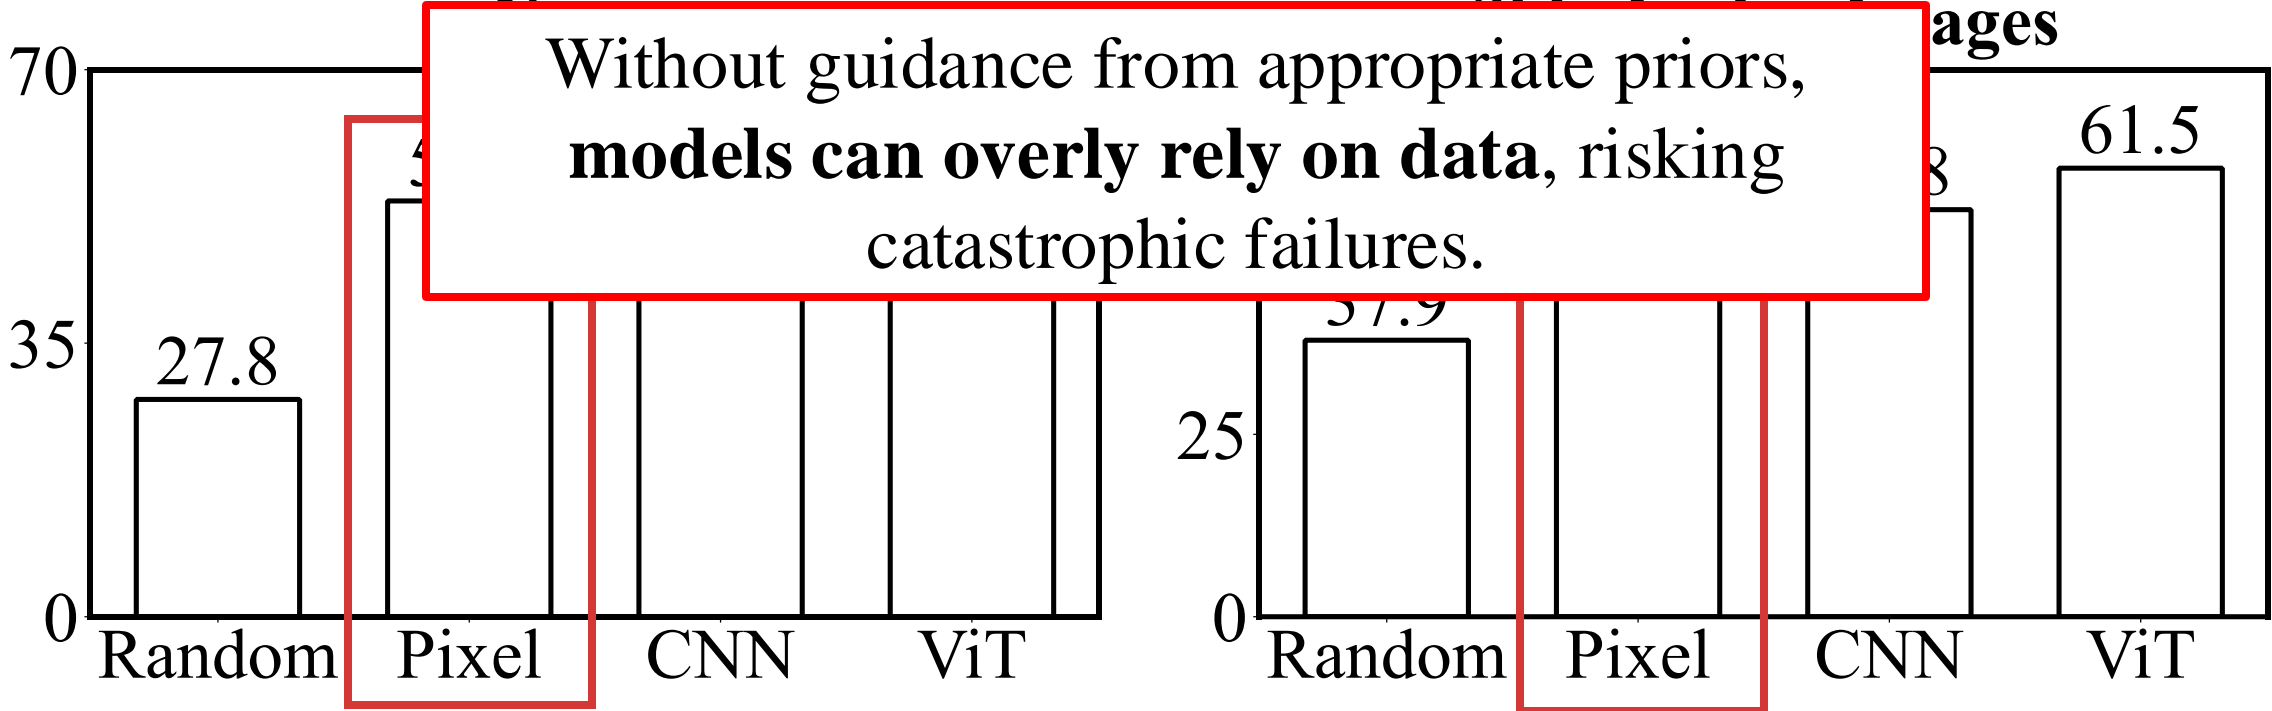
A. M. Saxe et al., On random weights and unsupervised feature learning. ICML 2011.

D. Ulyanov et al., Deep image prior. CVPR 2018.

Deep models **don't have** good priors for the **medical domain**.



Without guidance from appropriate priors, **models can overly rely on data**, risking catastrophic failures.



X-ray Datasets: Pneumonia, COVID-QU, NIH-CXR, Open-I, VinDr-CXR.

Skin Lesion Datasets: HAM10000, BCN20000, PAD-UFS-20, Melanoma, UWaterloo.

Inspired by Medical Education


Medical School



Textbooks

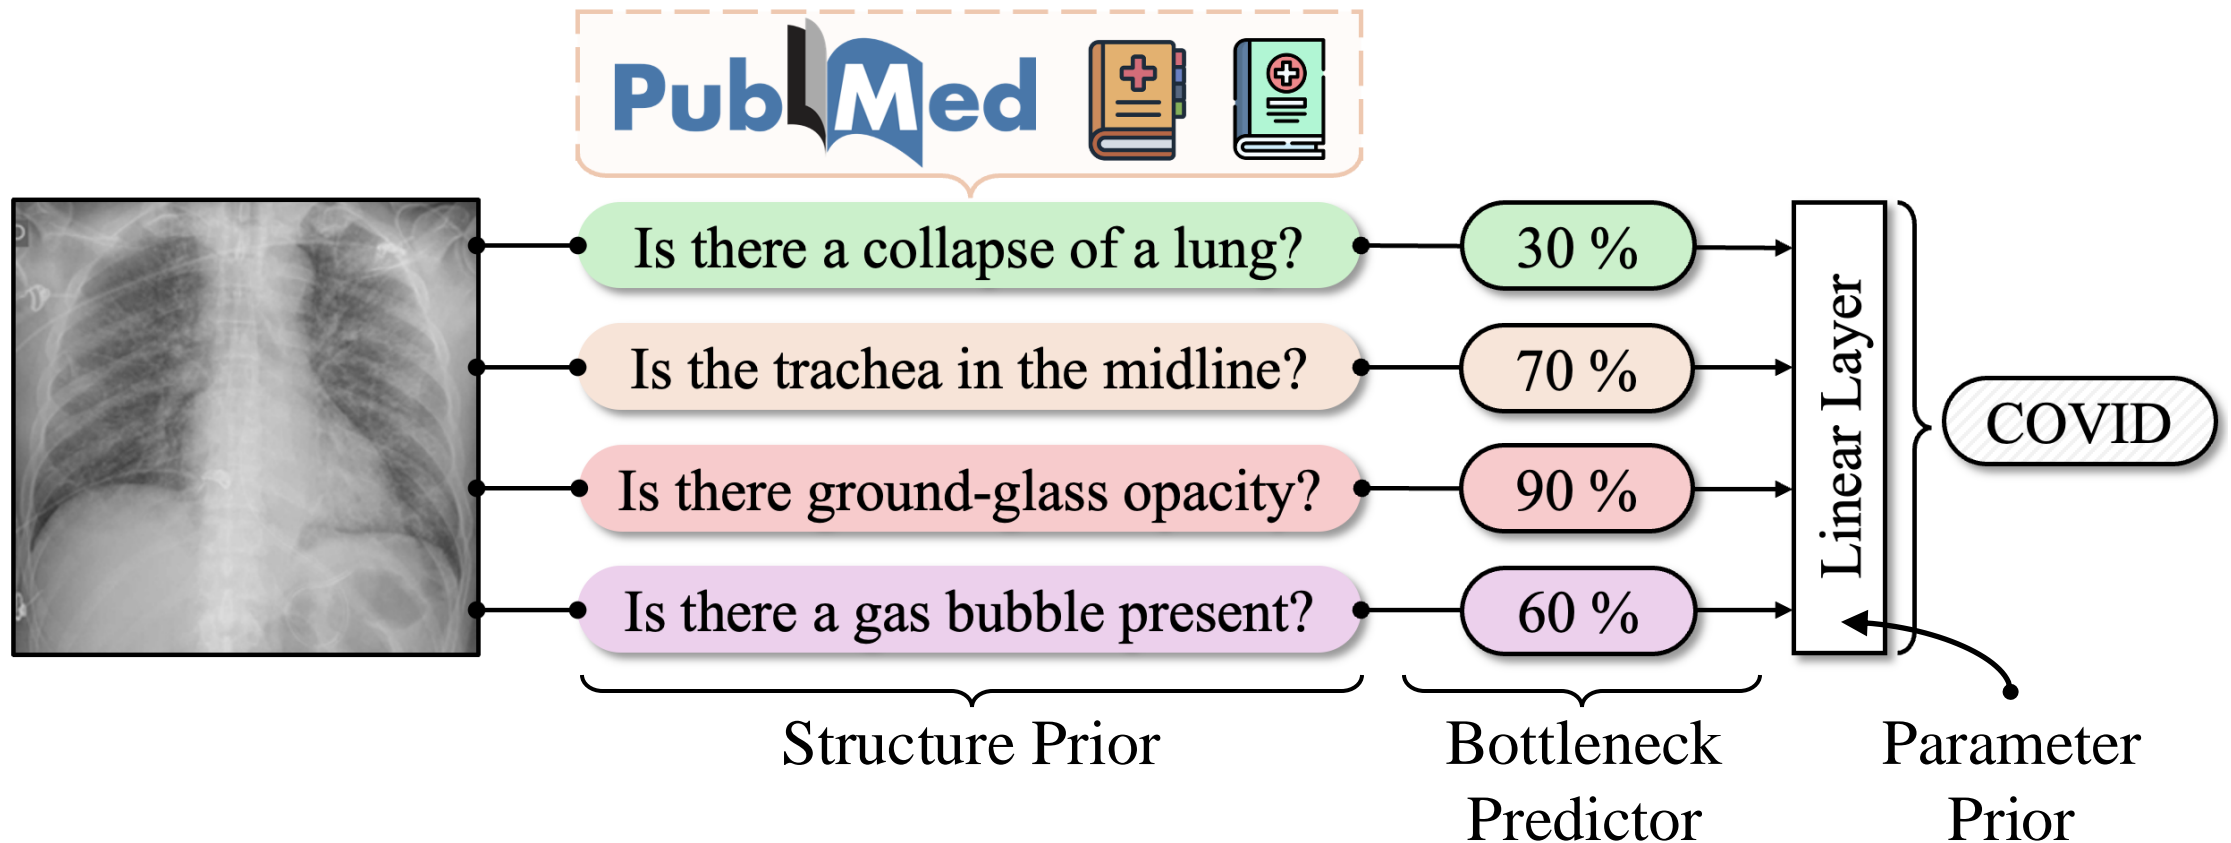


Residency in Hospital



Practice

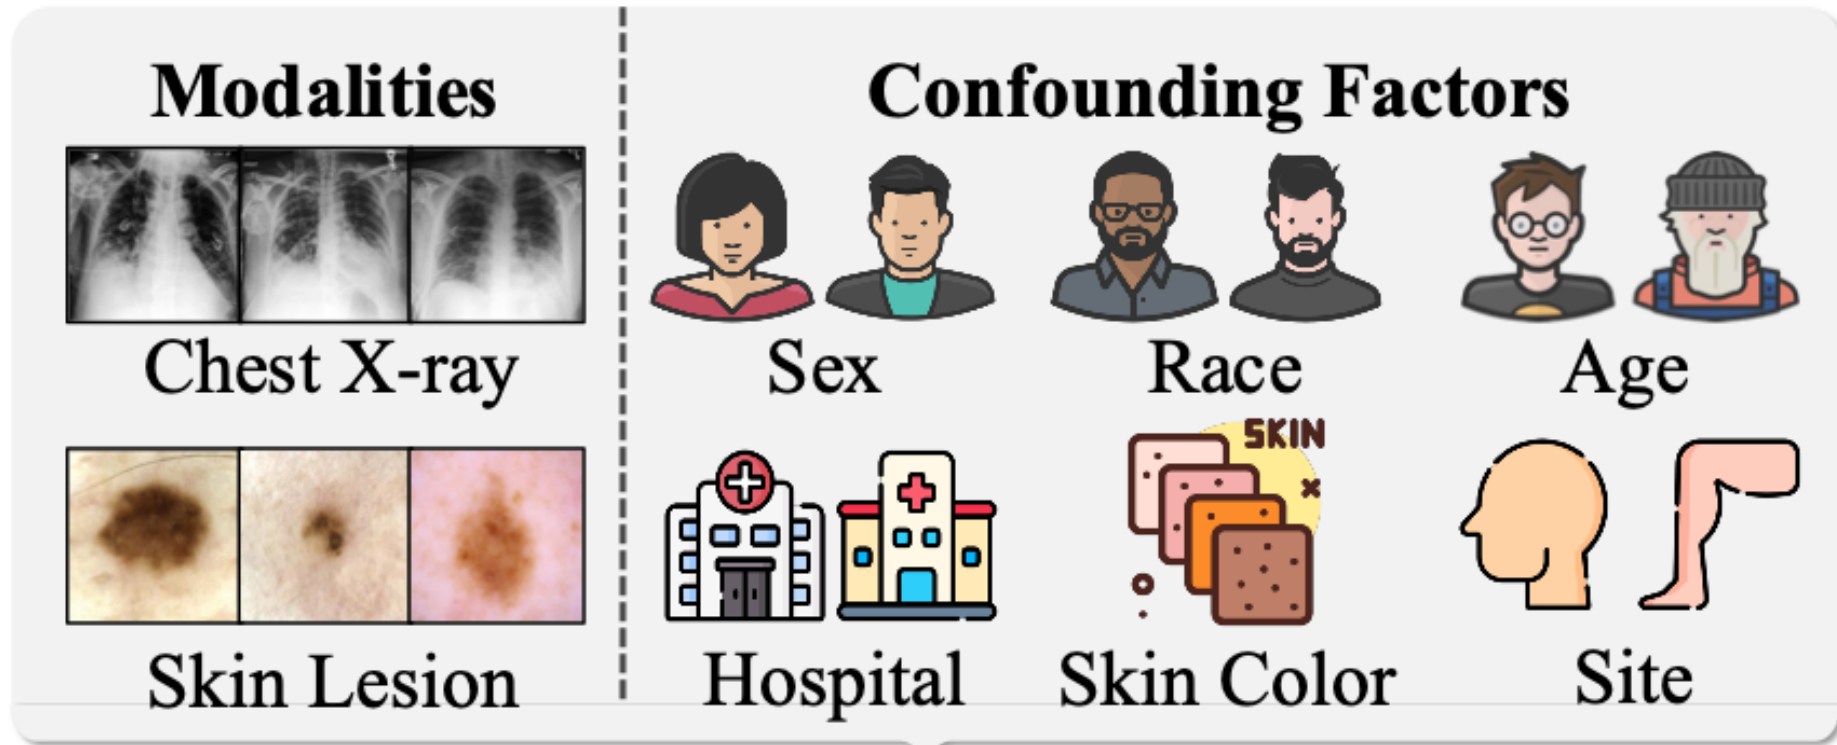
Ground Diagnose onto **Knowledge Bottlenecks**



[1] Koh et al. Concept Bottleneck Models. PMLR. 2020.

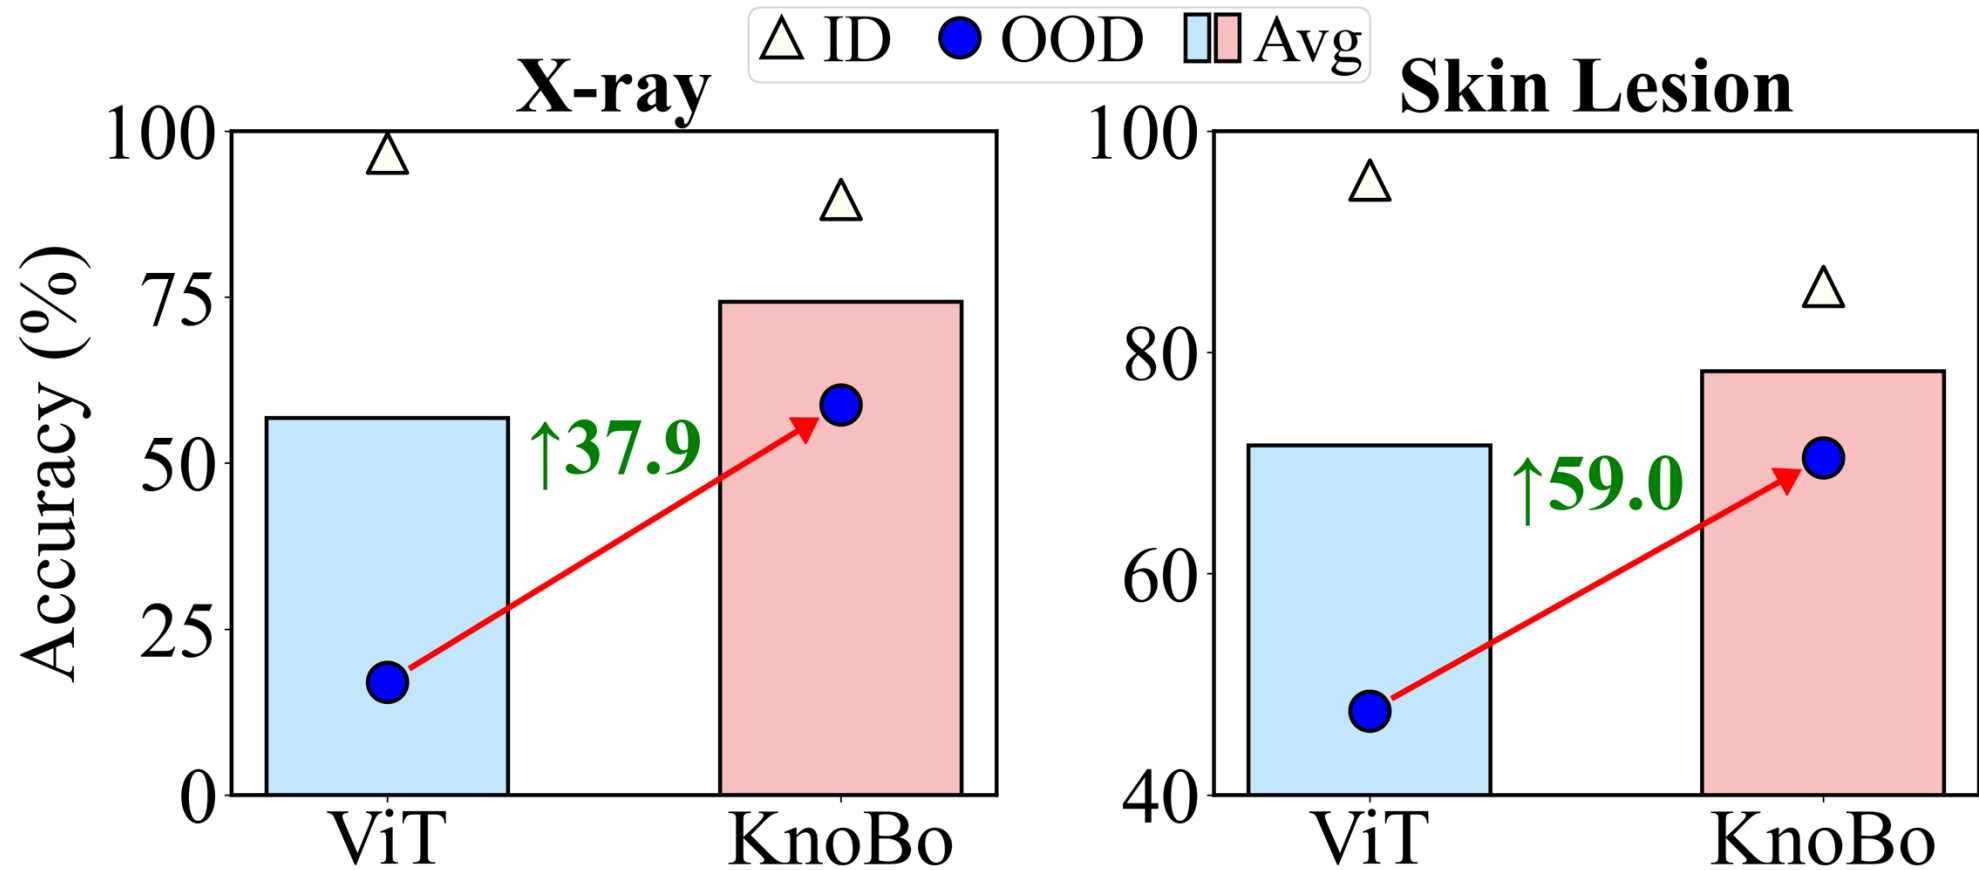
[2] Yang et al. Language Model Guided Concept Bottlenecks for Interpretable Image Classification. CVPR. 2023

Confounded Datasets


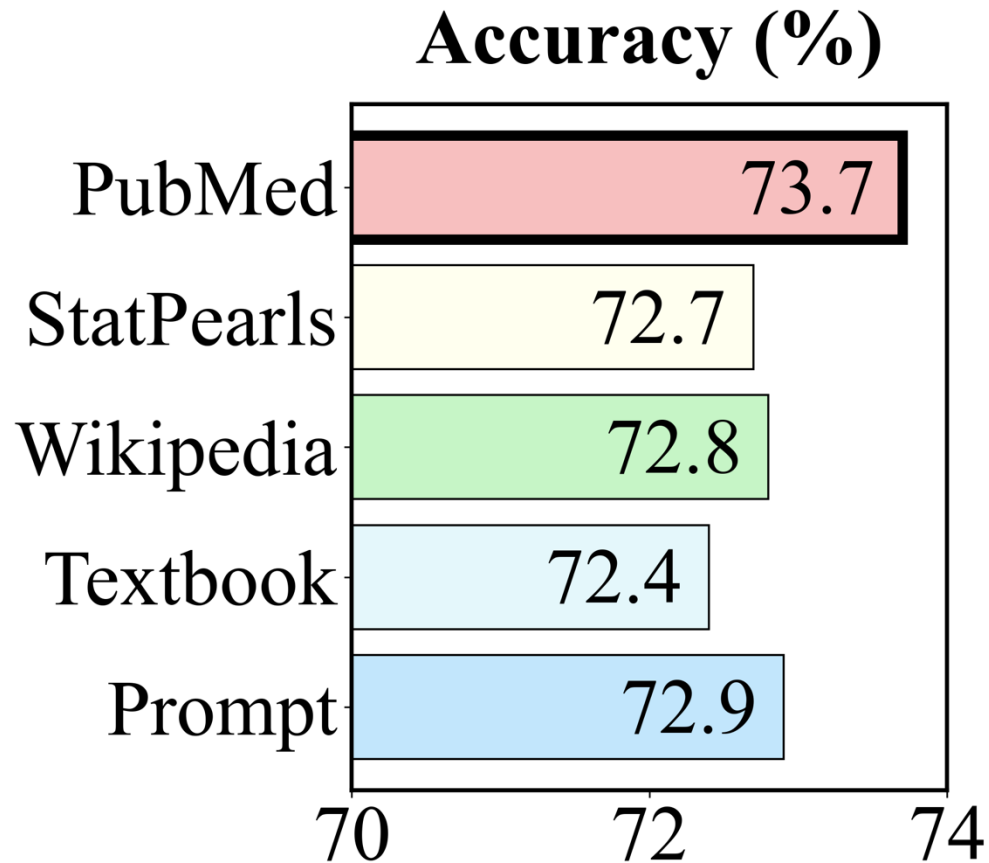


5 datasets for each modality

Results on Confounded Datasets



PubMed is a Promising Resource



Conclusion



Medical documents are reliable resources to extract medical knowledge.



Interpretable models with knowledge priors are **more robust in medical domains**.