Equity and Algorithm Development

Regina Barzilay

Disclosures

- SAB: Janssen, Amgen, Immunai, Fermenitch
- BOD: Dewpoint Therapeutics
- Research Funding: MLPDS, DTRA, Darpa, Quanta, Novo Nordisk, NSF, NCI

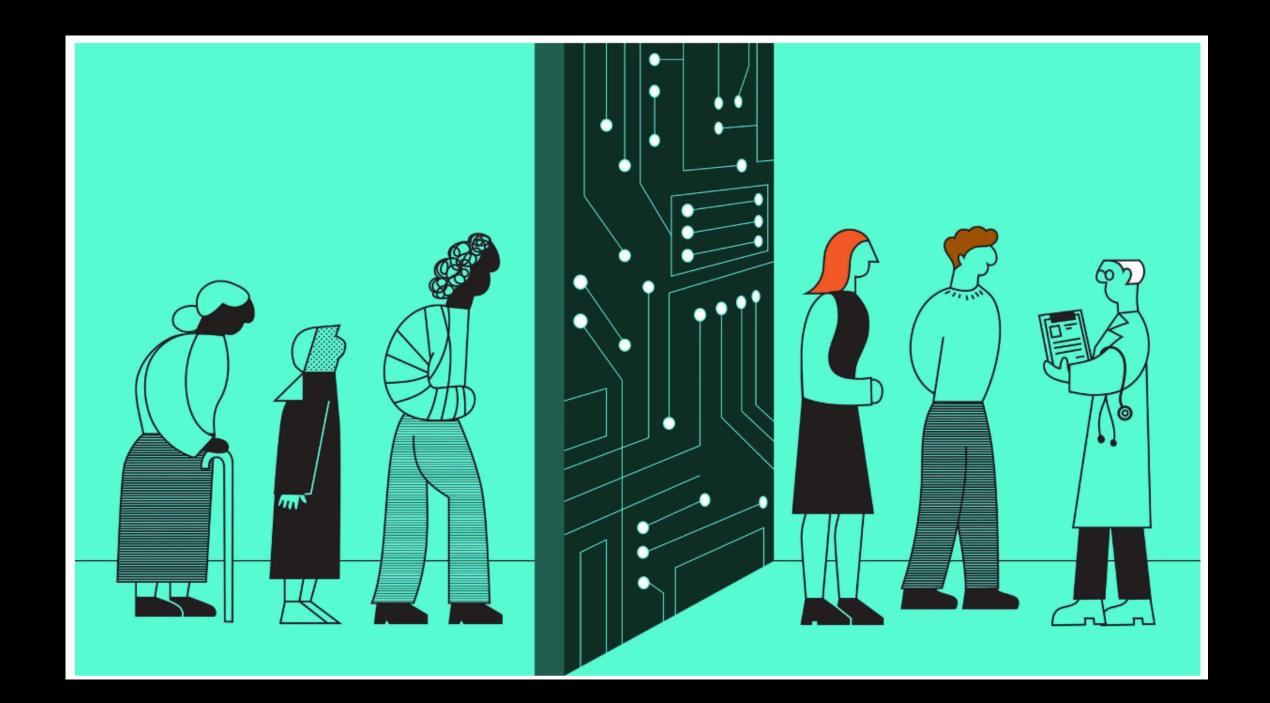
Handling Diversity

A.I. Could Worsen Health Disparities

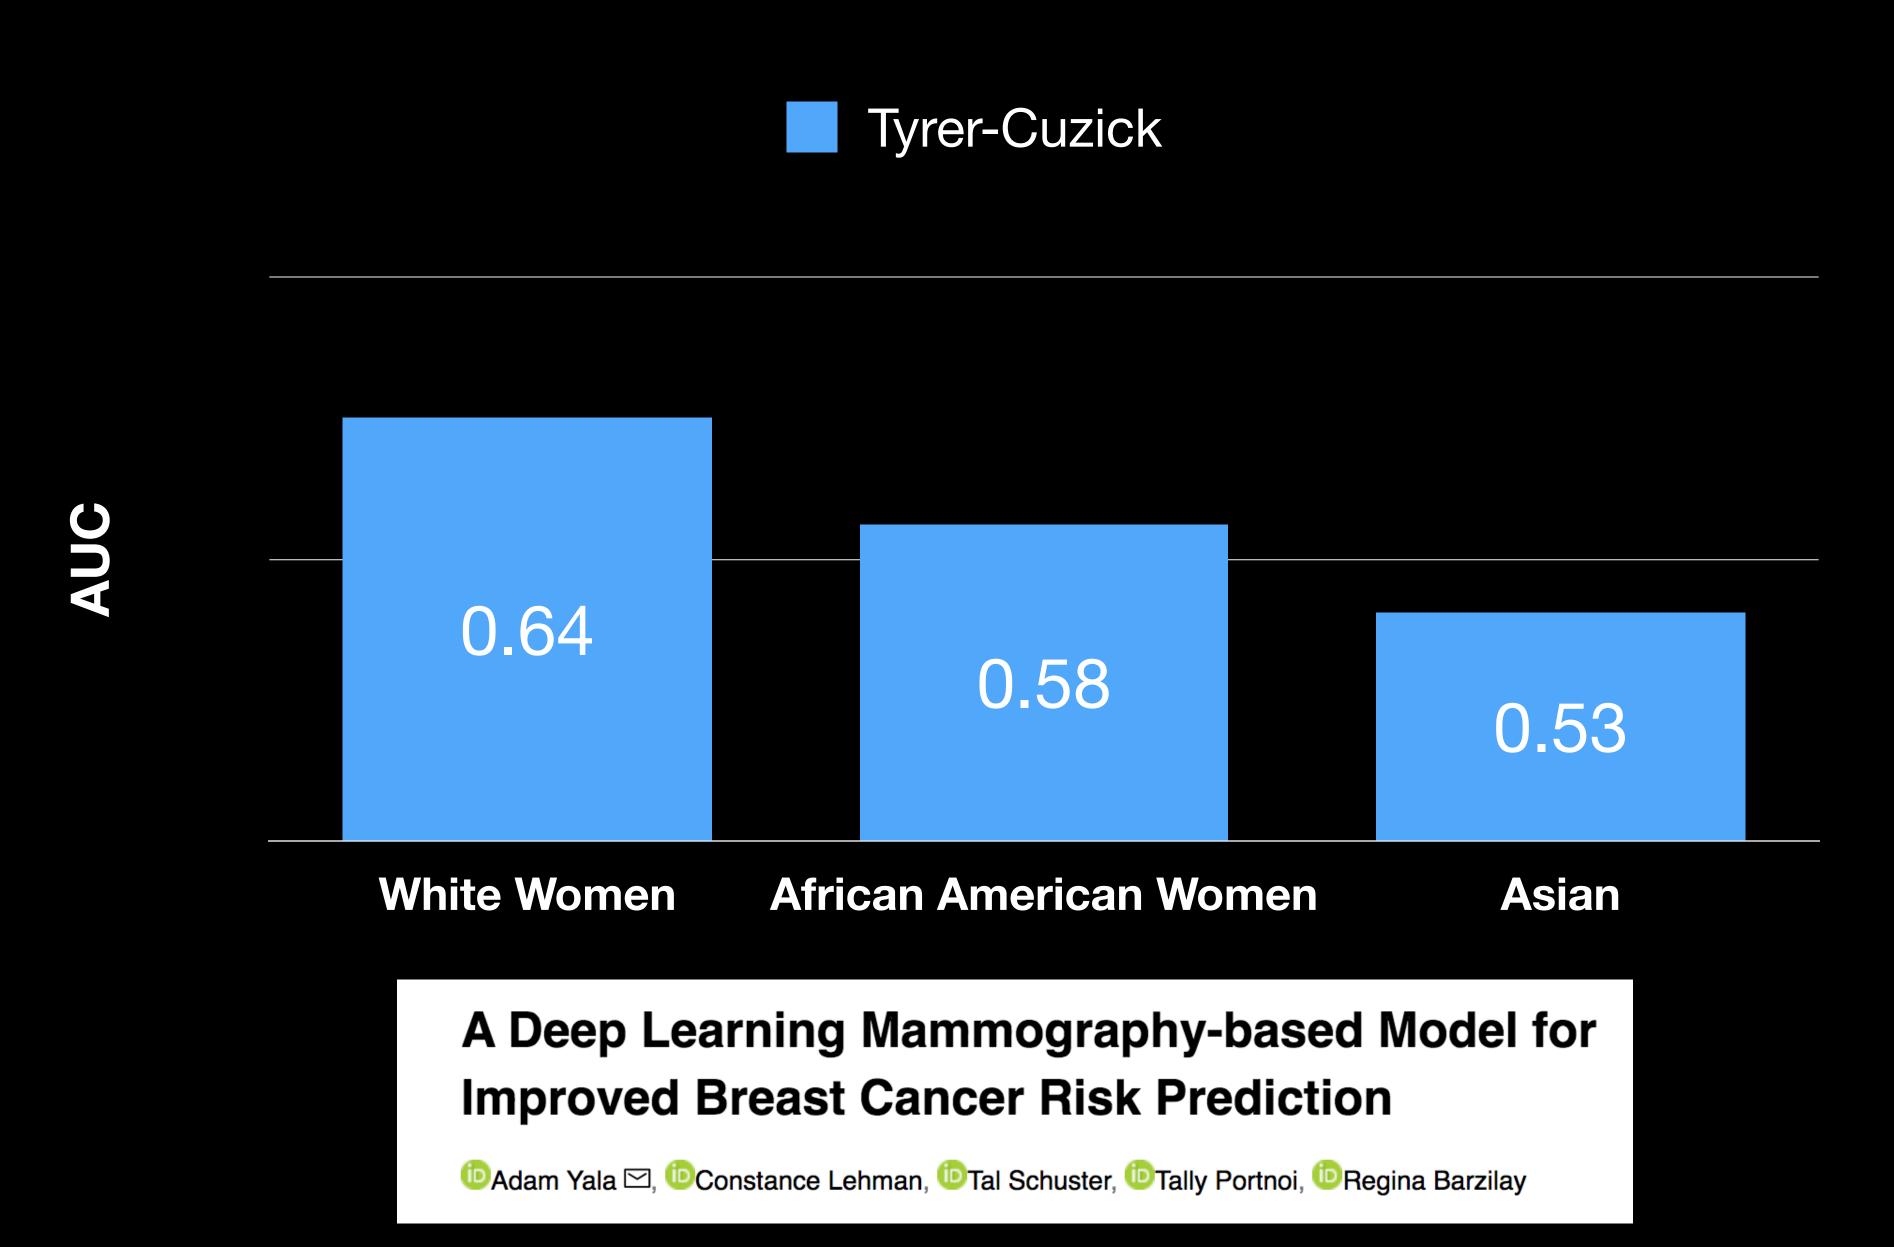
In a health system riddled with inequity, we risk making dangerous biases automated and invisible.

By Dhruv Khullar

Dr. Khullar is an assistant professor of health care policy and research.



Biased Models Predate AI



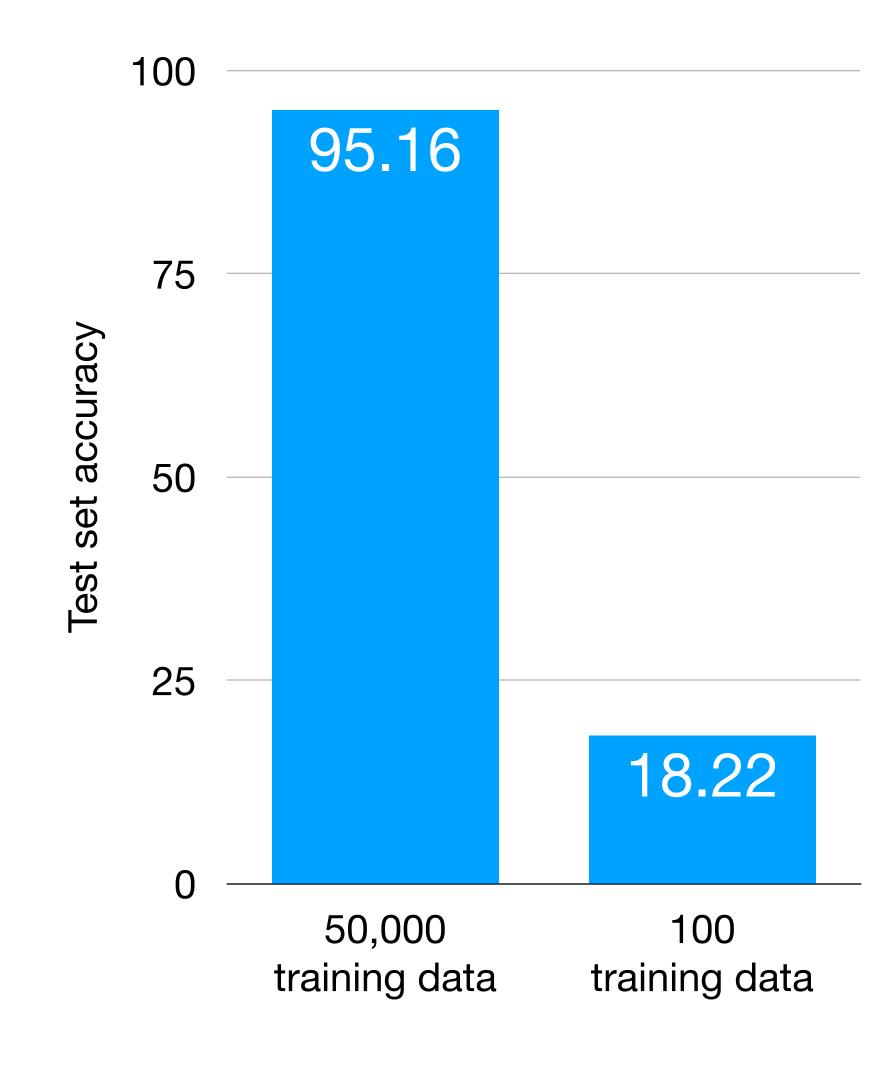
Problem 1: DNNs are Data Hungry



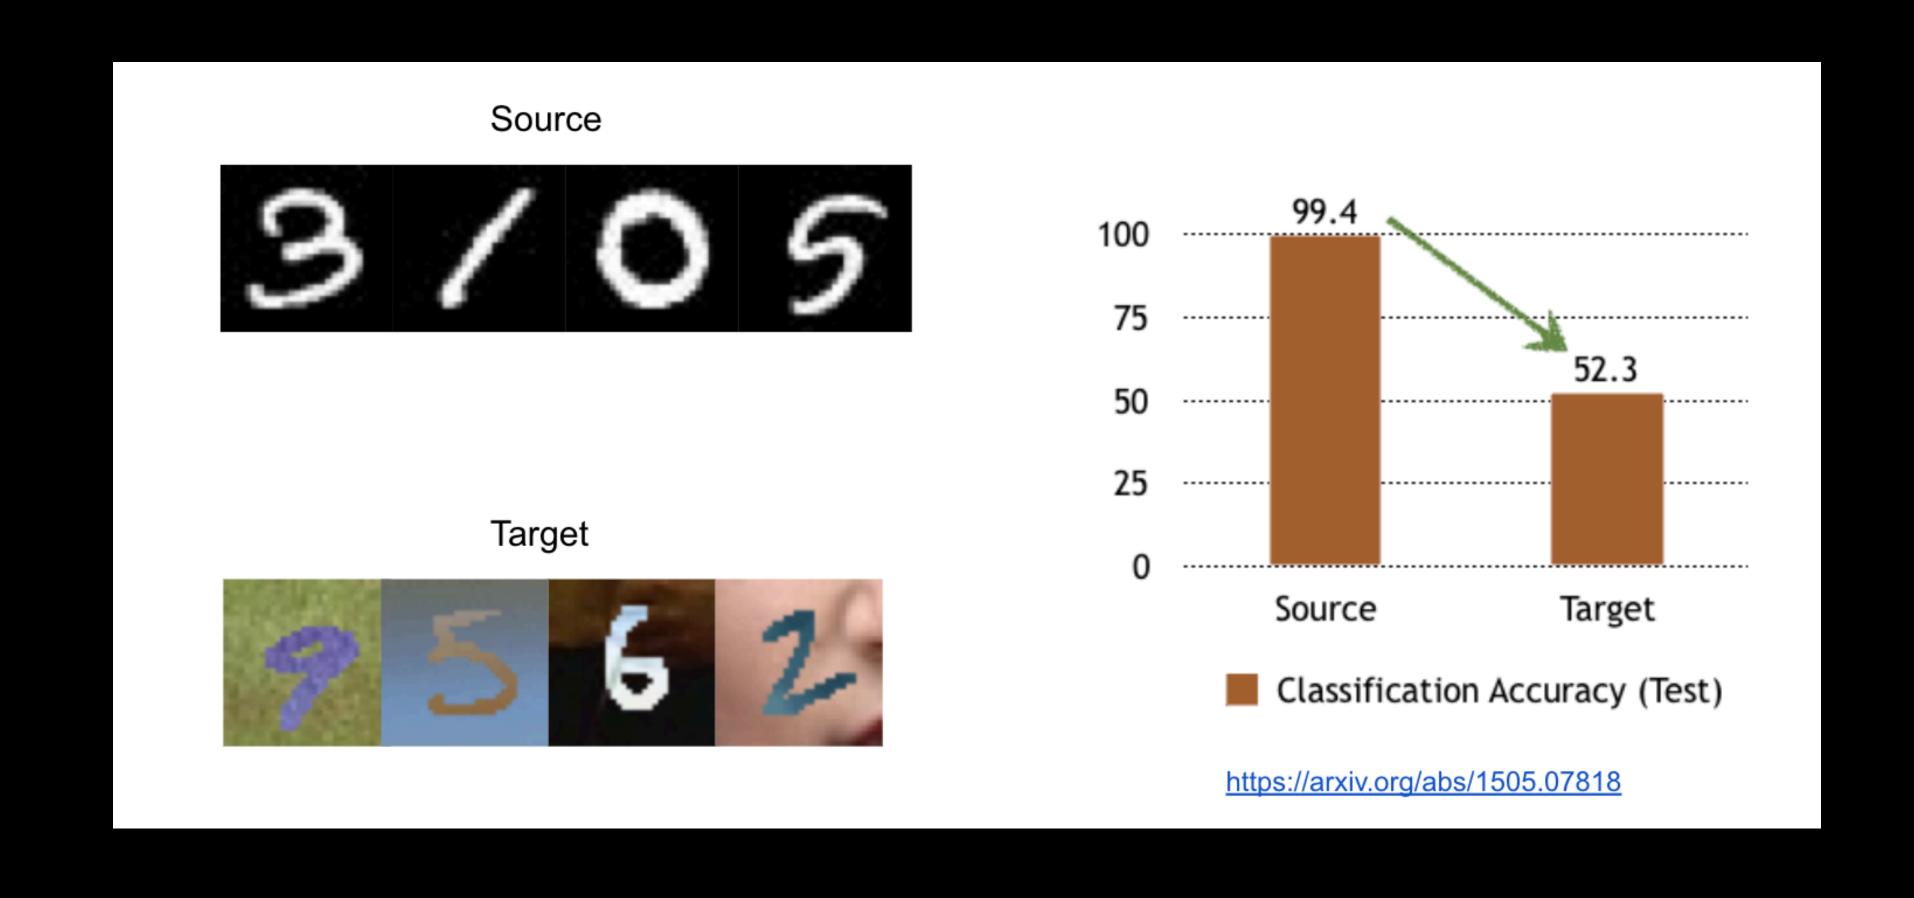
CIFAR-10 dataset

classes: 10

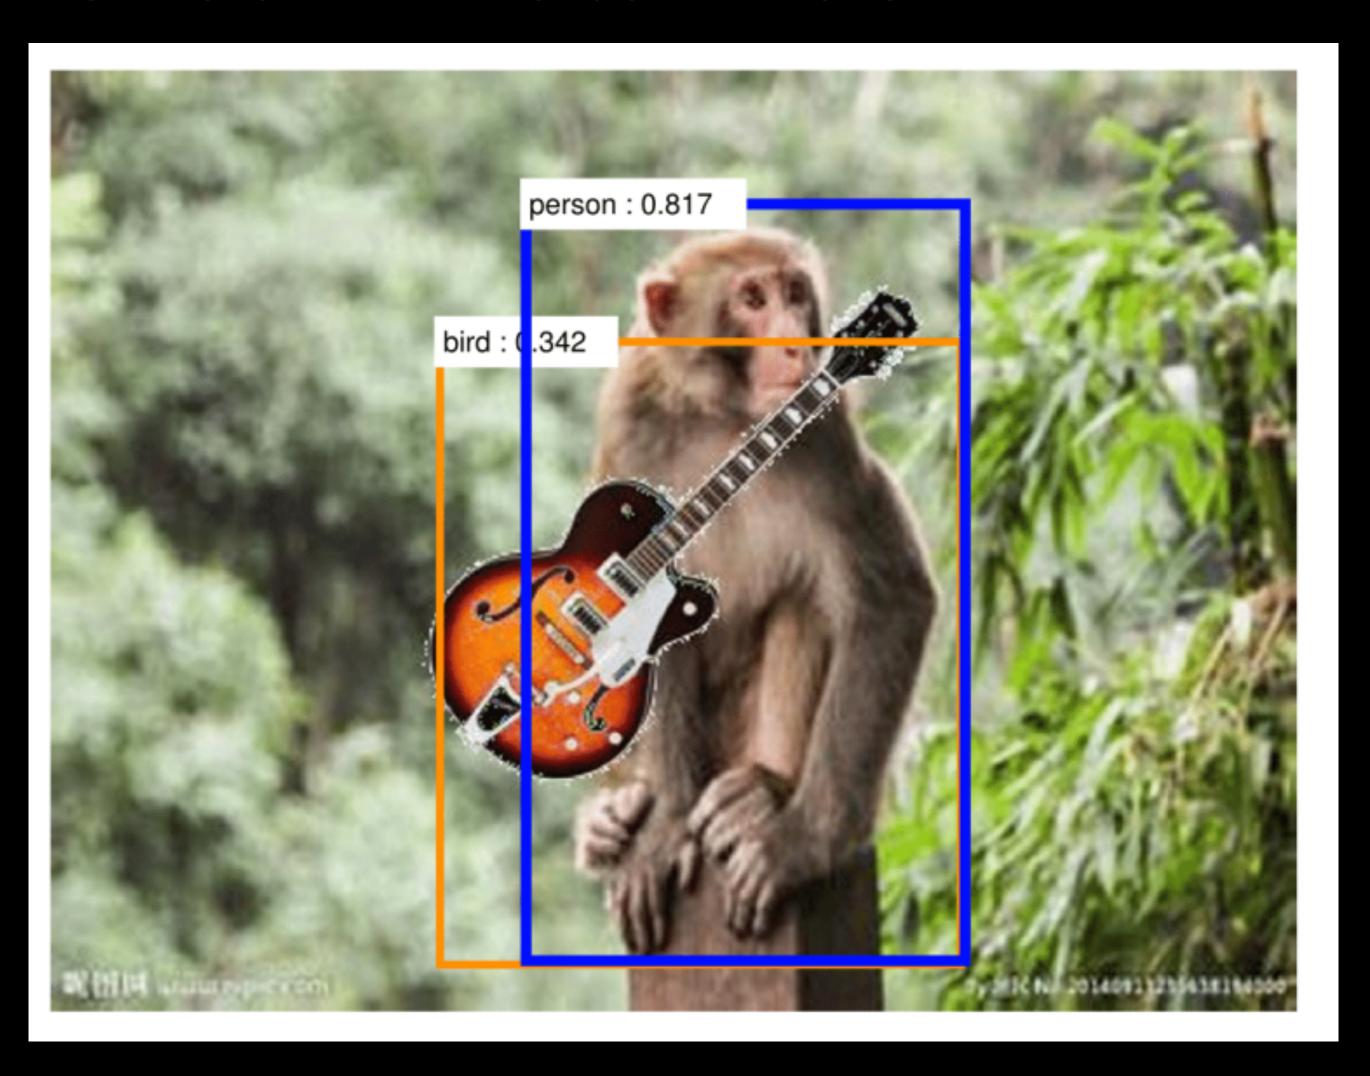
img per class: 5000



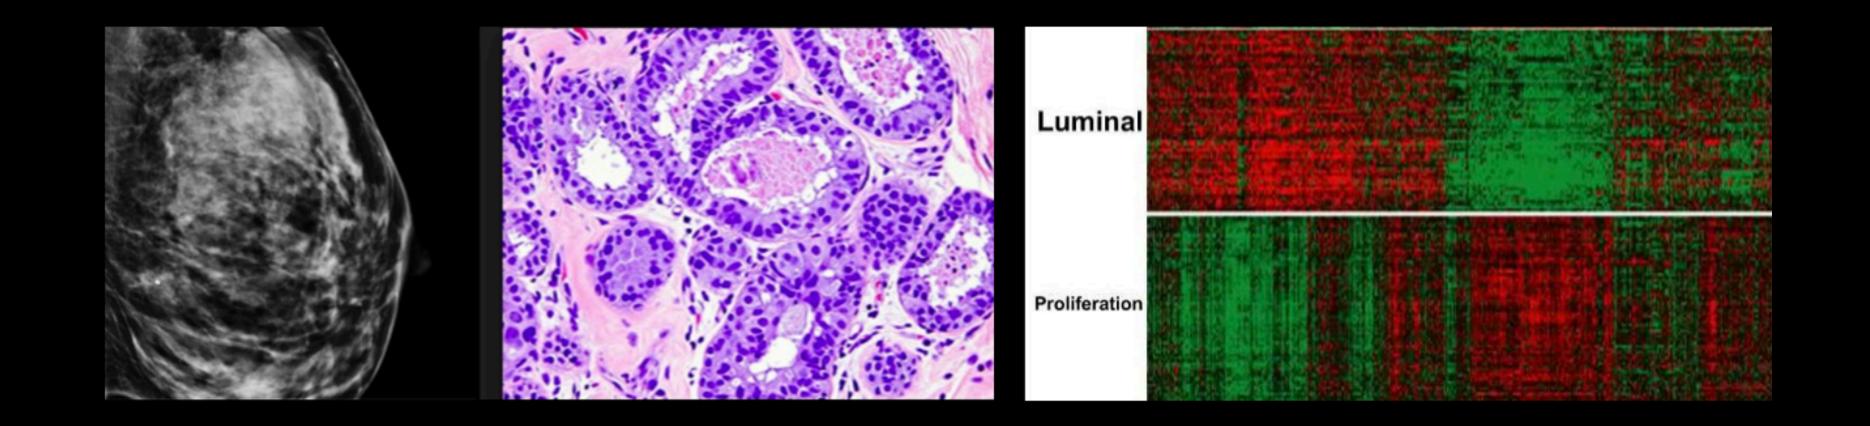
Problem 2: DNNs are not Robust Towards Distributional Shift



Problem 3: DNNs learn Data Bias



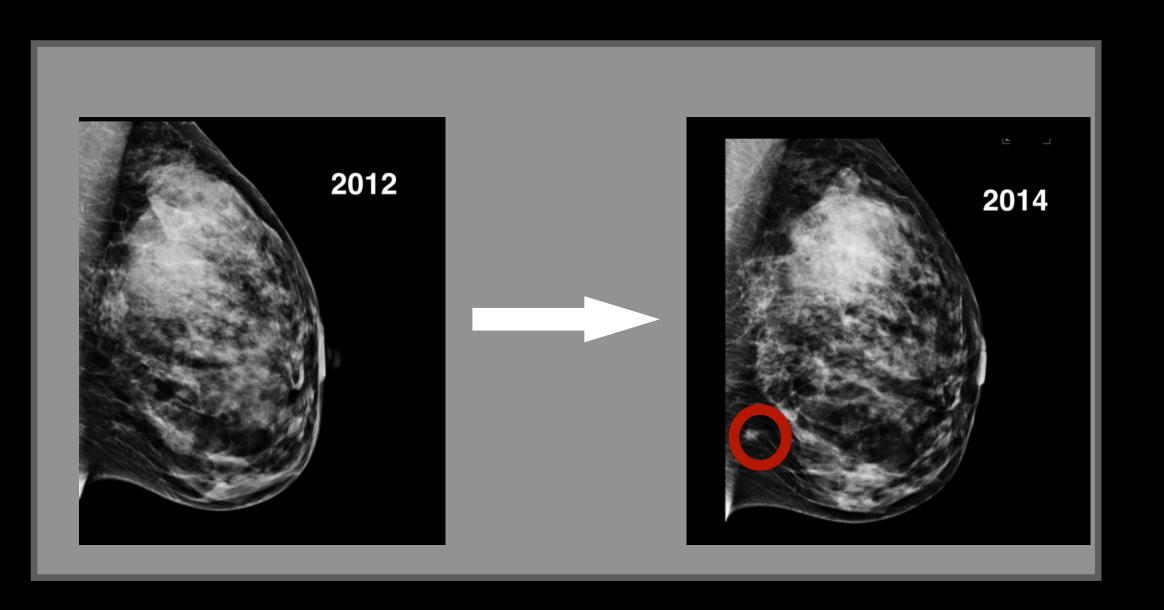
Most Diagnostics Tasks are Beyond Human Prediction Capacity

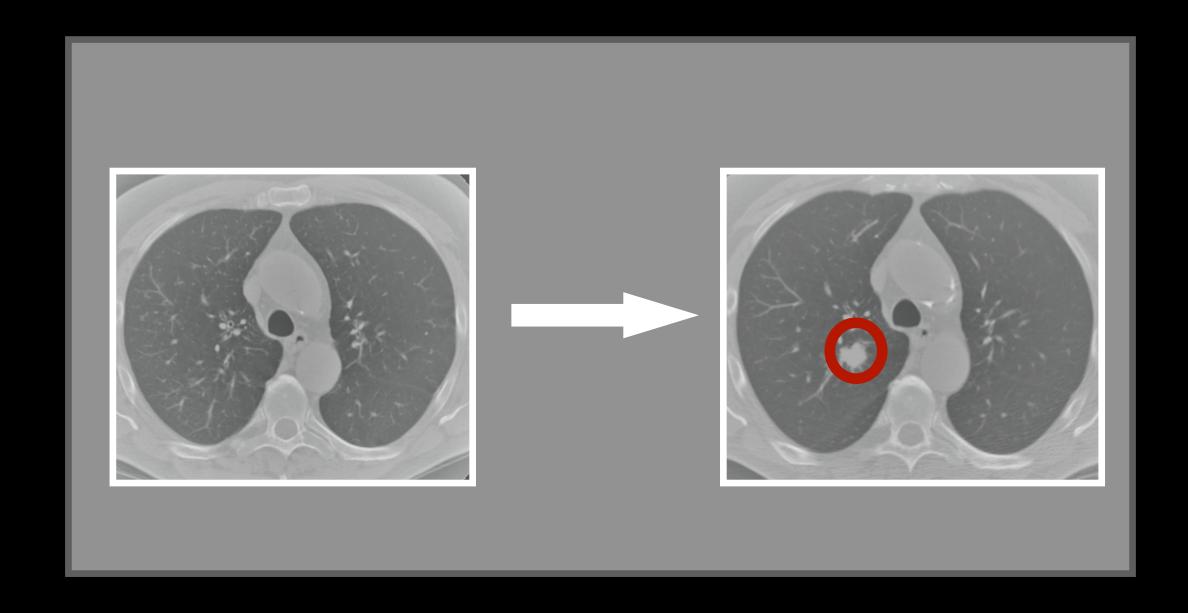




Predict Recurrences, Sensitivity to Treatment,
Population at Risk

Case Studies: Predicting Future Risk





SCIENCE TRANSLATIONAL MEDICINE | RESEARCH ARTICLE

CANCER

Toward robust mammography-based models for breast cancer risk

Adam Yala^{1,2}*, Peter G. Mikhael^{1,2}, Fredrik Strand^{3,4}, Gigin Lin⁵, Kevin Smith^{6,7}, Yung-Liang Wan⁵, Leslie Lamb⁸, Kevin Hughes⁹, Constance Lehman^{8†}, Regina Barzilay^{1,2†}

<u>Journal of Clinical Oncology</u> > <u>List of Issues</u> > <u>Newest Content</u> >

ORIGINAL REPORTS | Breast Cancer

Multi-Institutional Validation of a Mammography-Based Breast Cancer Risk Model

Optimizing risk-based breast cancer screening policies with reinforcement learning

Adam Yala , Peter G. Mikhael, Constance Lehman, Gigin Lin, Fredrik Strand, Yung-Liang Wan, Kevin Hughes, Siddharth Satuluru, Thomas Kim, Imon Banerjee, Judy Gichoya, Hari Trivedi & Regina Barzilay

Nature Medicine 28, 136–143 (2022) | Cite this article

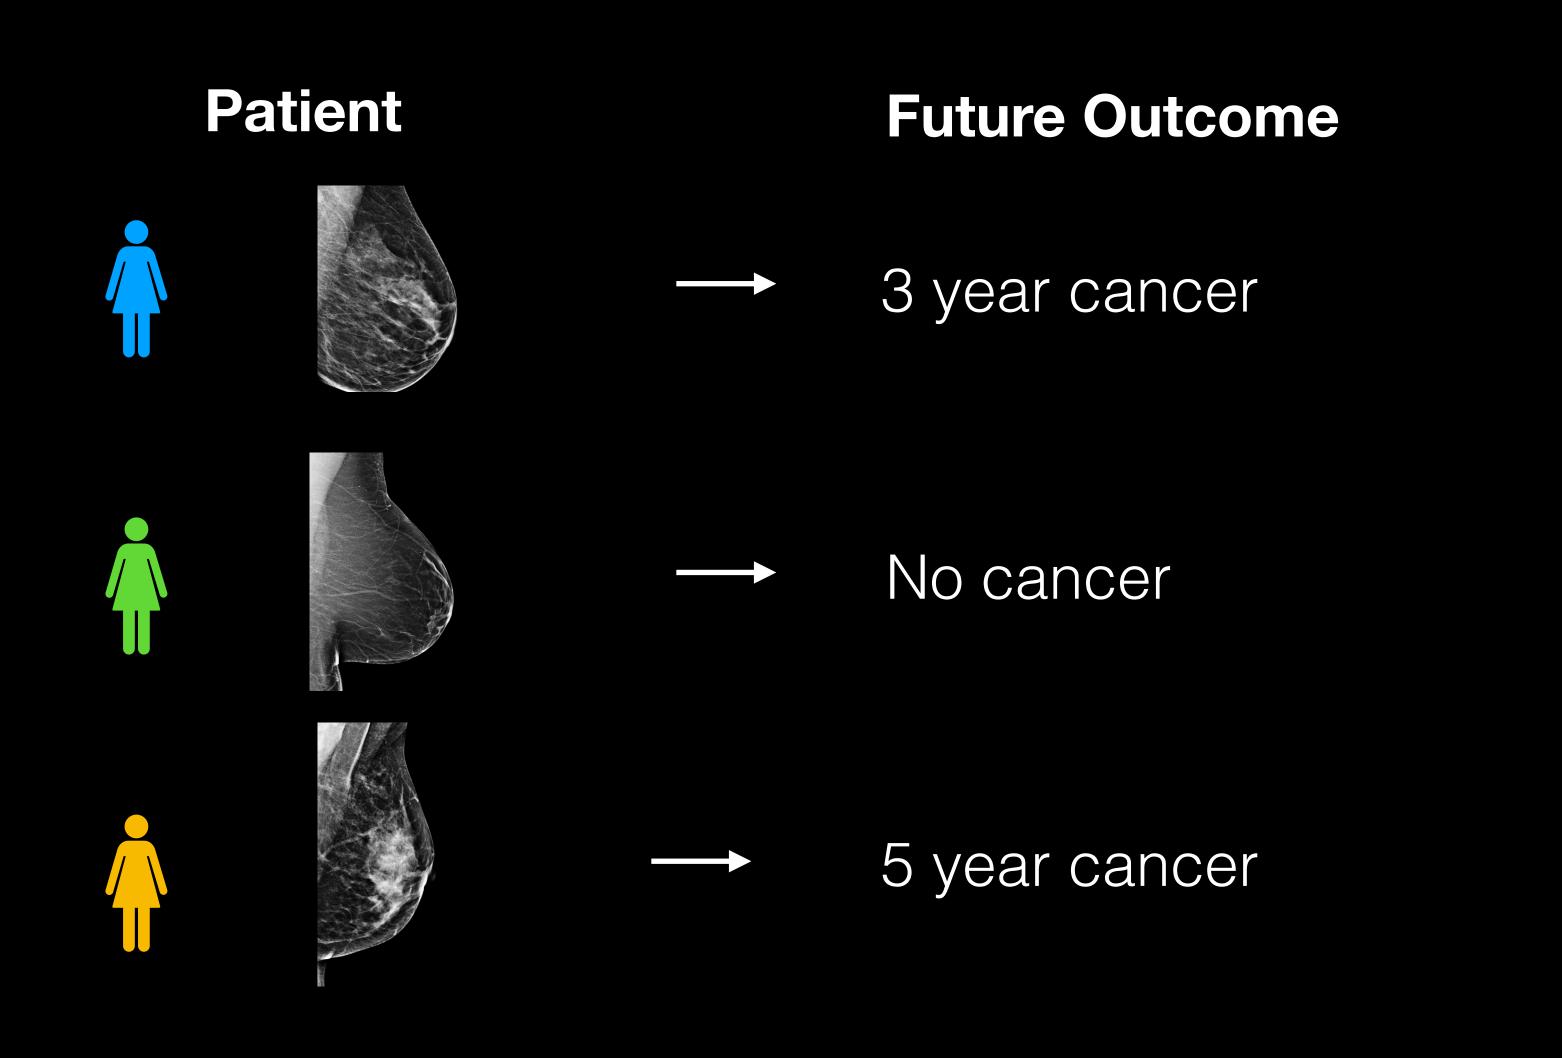


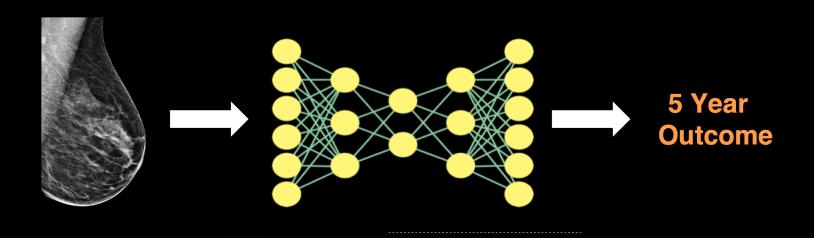
| Journal of Clinical Oncology > List of Issues > Newest Content >

ORIGINAL REPORTS | Thoracic Oncology

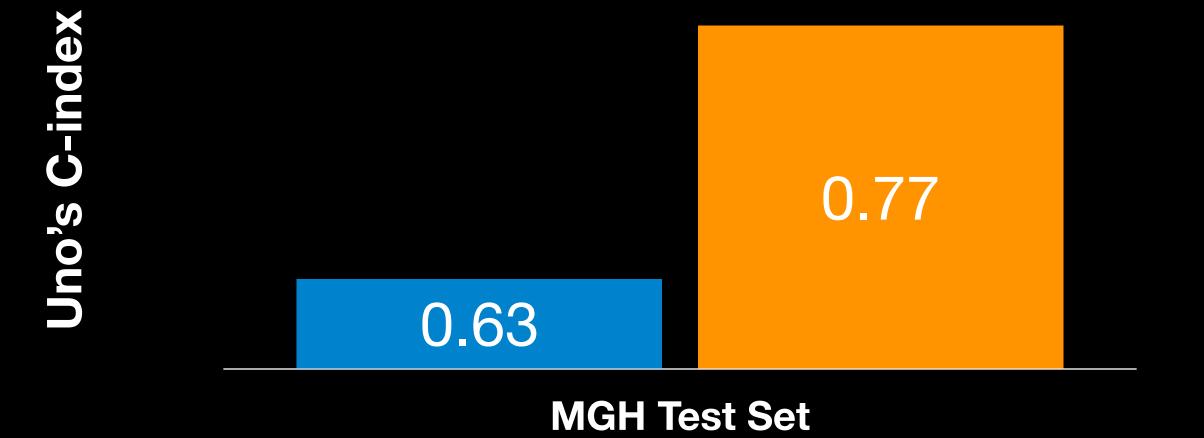
Sybil: A Validated Deep Learning Model to Predict Future Lung Cancer Risk From a Single Low-Dose Chest Computed Tomography

Learning to predict future cancer

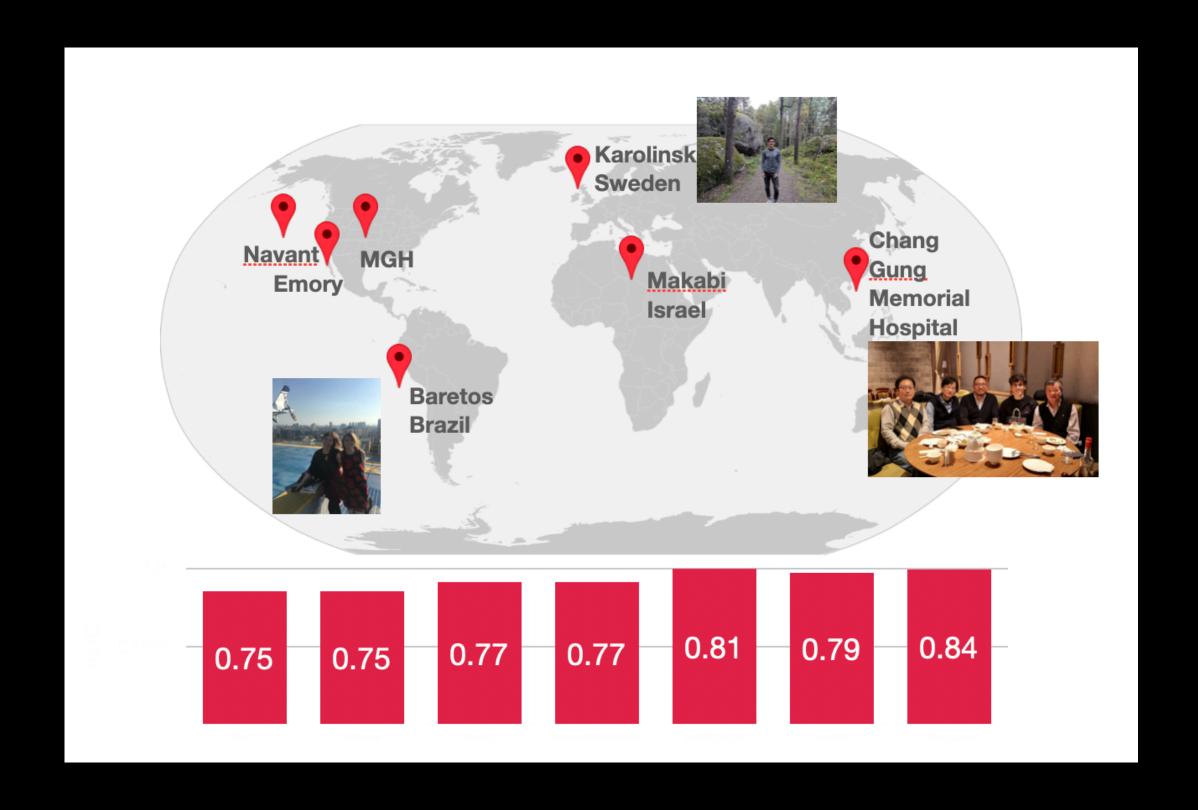


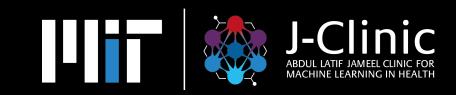


Tyrer-Cuzick (Prior State of the Art)
MIRAI (Ours - New Result)



2020: Multi-Institutional Validation





SCIENCE TRANSLATIONAL MEDICINE | RESEARCH ARTICLE

CANCER

Toward robust mammography-based models for breast cancer risk

Adam Yala^{1,2}*, Peter G. Mikhael^{1,2}, Fredrik Strand^{3,4}, Gigin Lin⁵, Kevin Smith^{6,7}, Yung-Liang Wan⁵, Leslie Lamb⁸, Kevin Hughes⁹, Constance Lehman^{8†}, Regina Barzilay^{1,2†} <u>Journal of Clinical Oncology</u> > <u>List of Issues</u> > <u>Newest Content</u> >

ORIGINAL REPORTS | Breast Cancer

Multi-Institutional Validation of a Mammography-Based Breast Cancer Risk Model Optimizing risk-based breast cancer screening policies with reinforcement learning

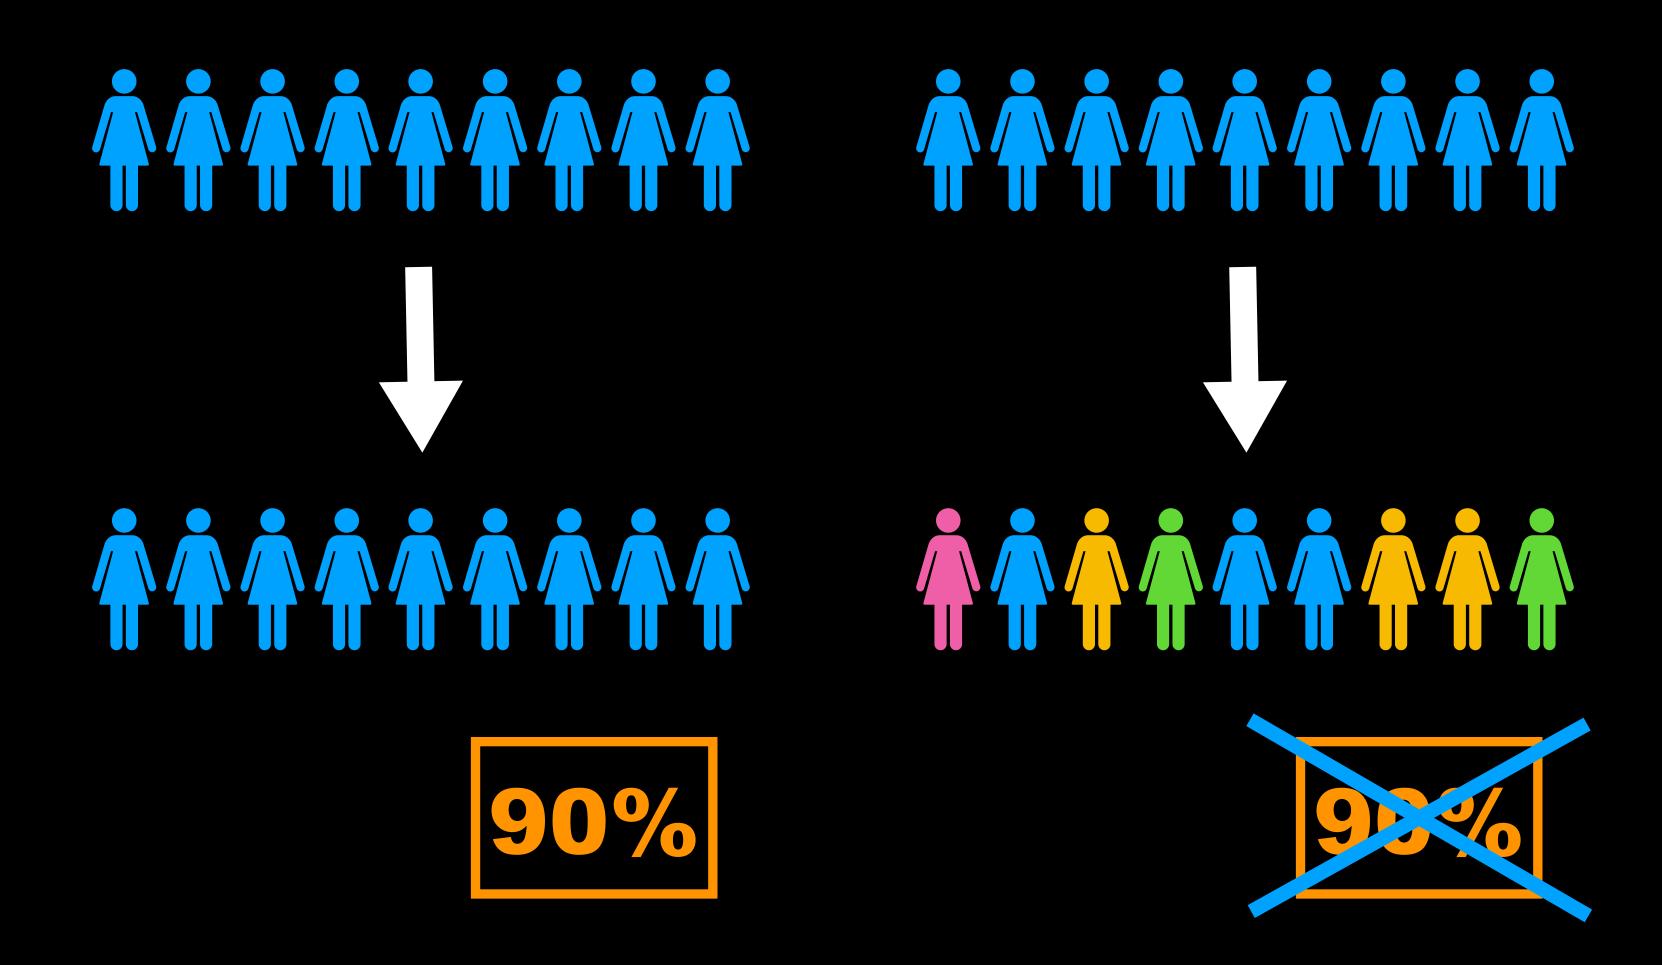
Adam Yala 🖂, Peter G. Mikhael, Constance Lehman, Gigin Lin, Fredrik Strand, Yung-Liang Wan, Kevin Hughes, Siddharth Satuluru, Thomas Kim, Imon Banerjee, Judy Gichoya, Hari Trivedi & Regina Barzilay

<u>Nature Medicine</u> **28**, 136–143 (2022) | <u>Cite this article</u>

2023: Hospital Network Deployment



Distributional Shift from Training



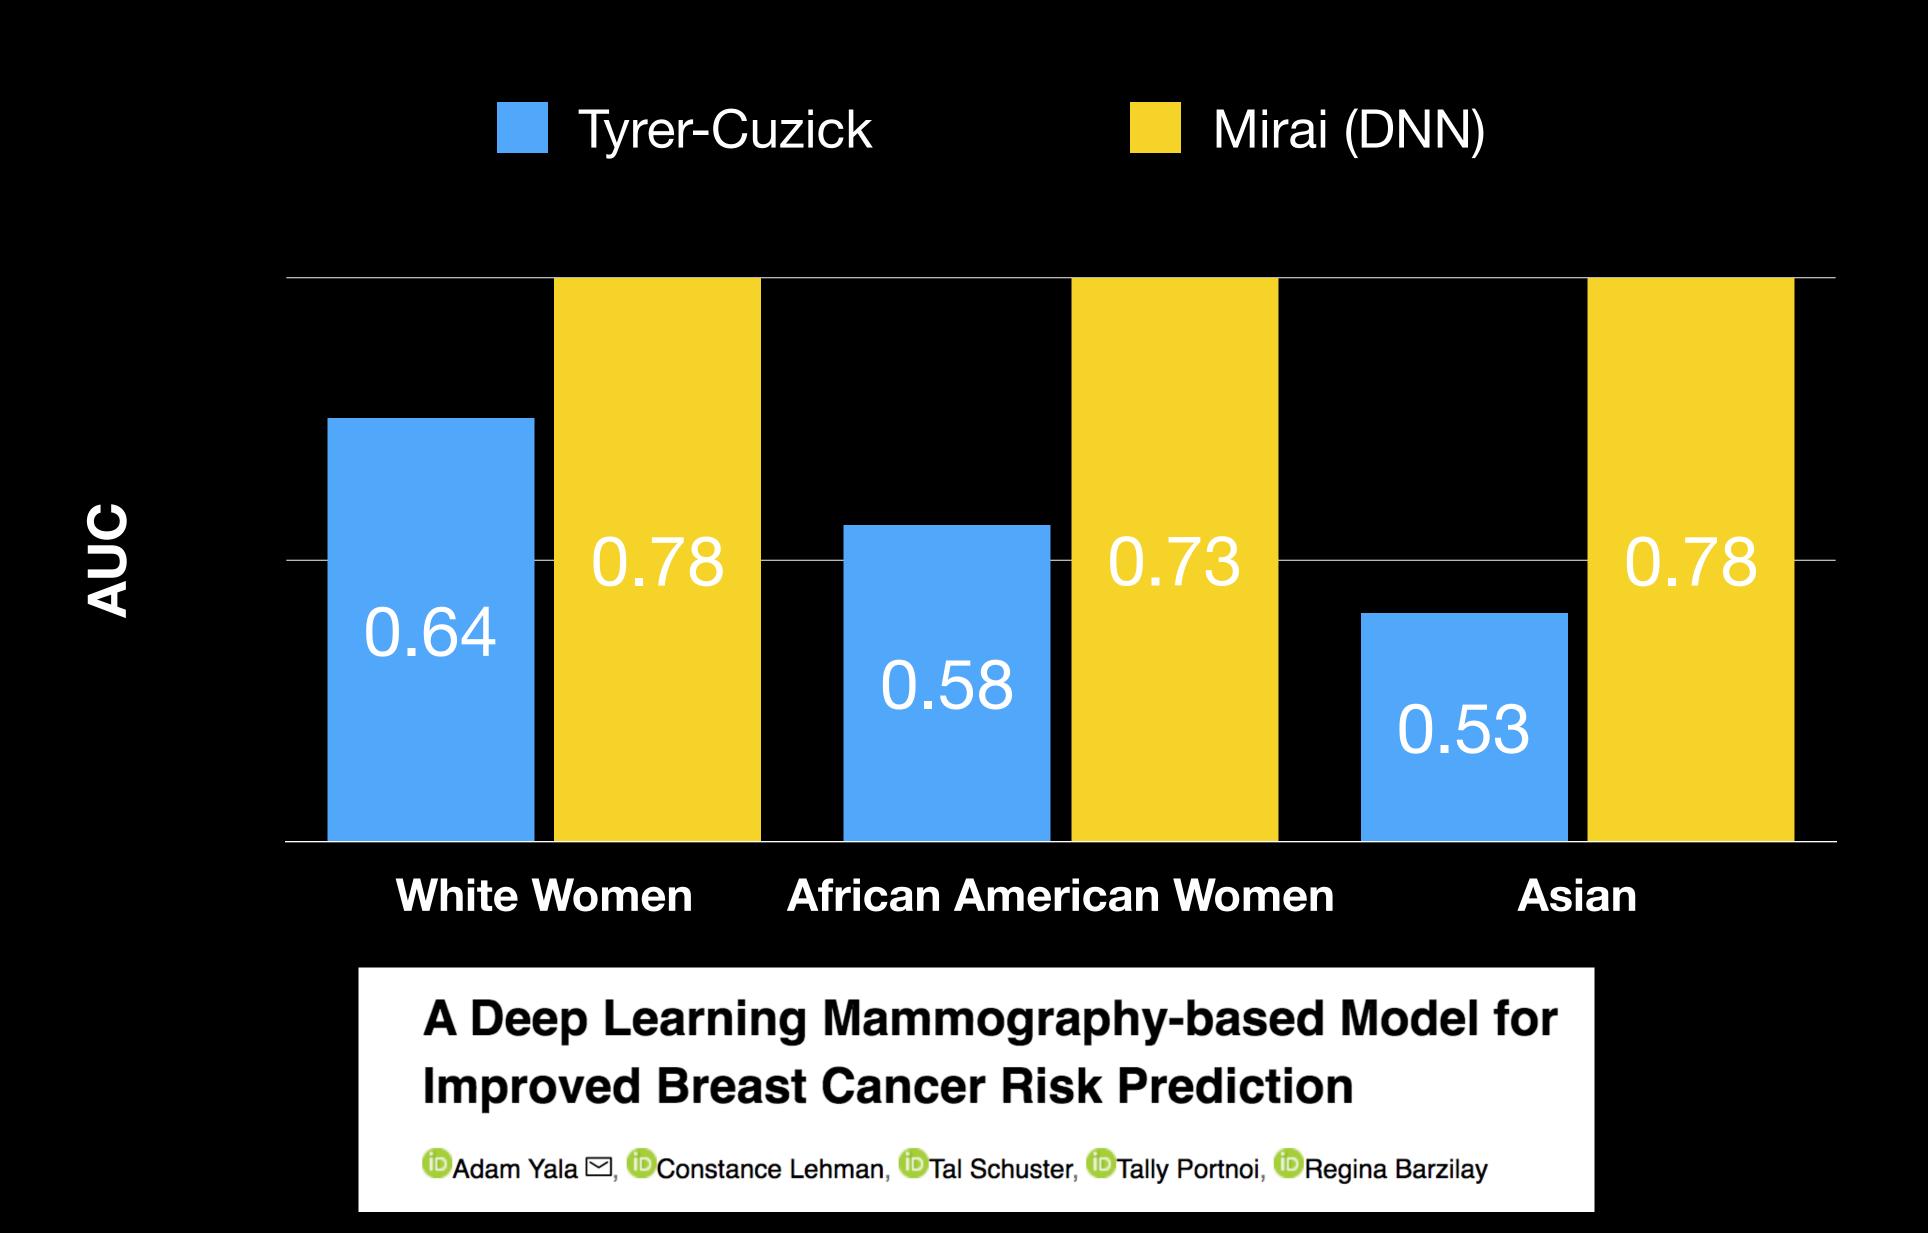
Distributional Shift from Training

Demographics

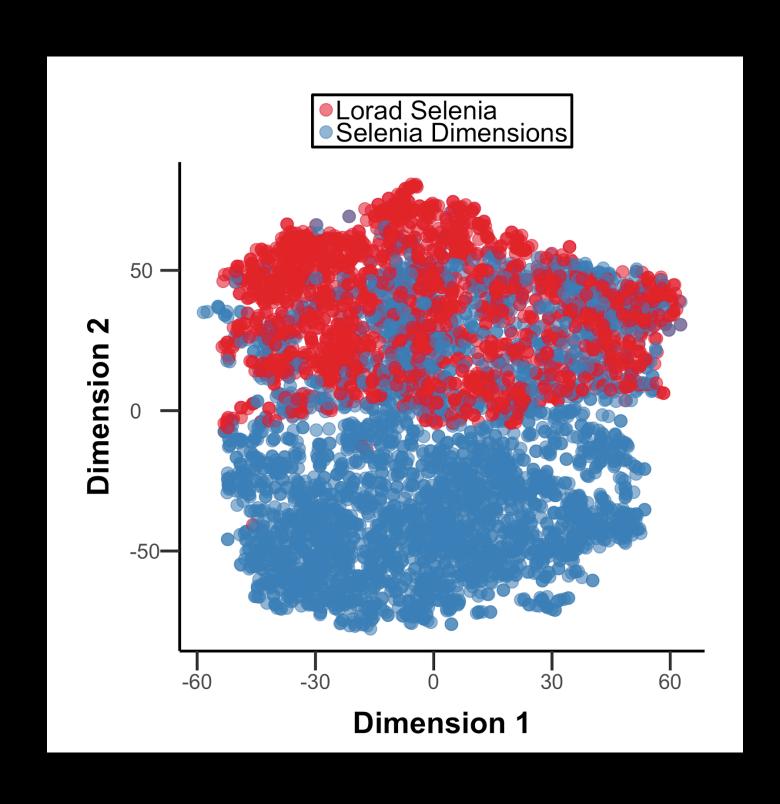
Clinical Setting

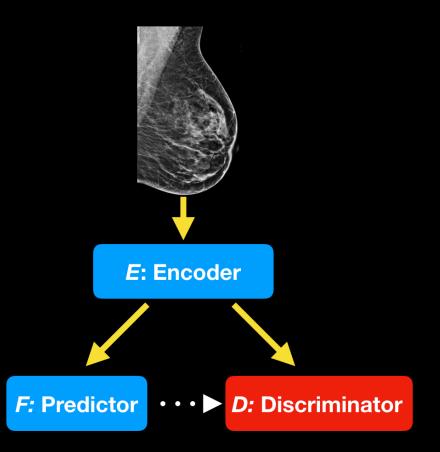
Device

Training from Diverse Data Helps



Algorithmic Solution 1: Anti-Bias Training Algorithms: Learning Invariant Representations



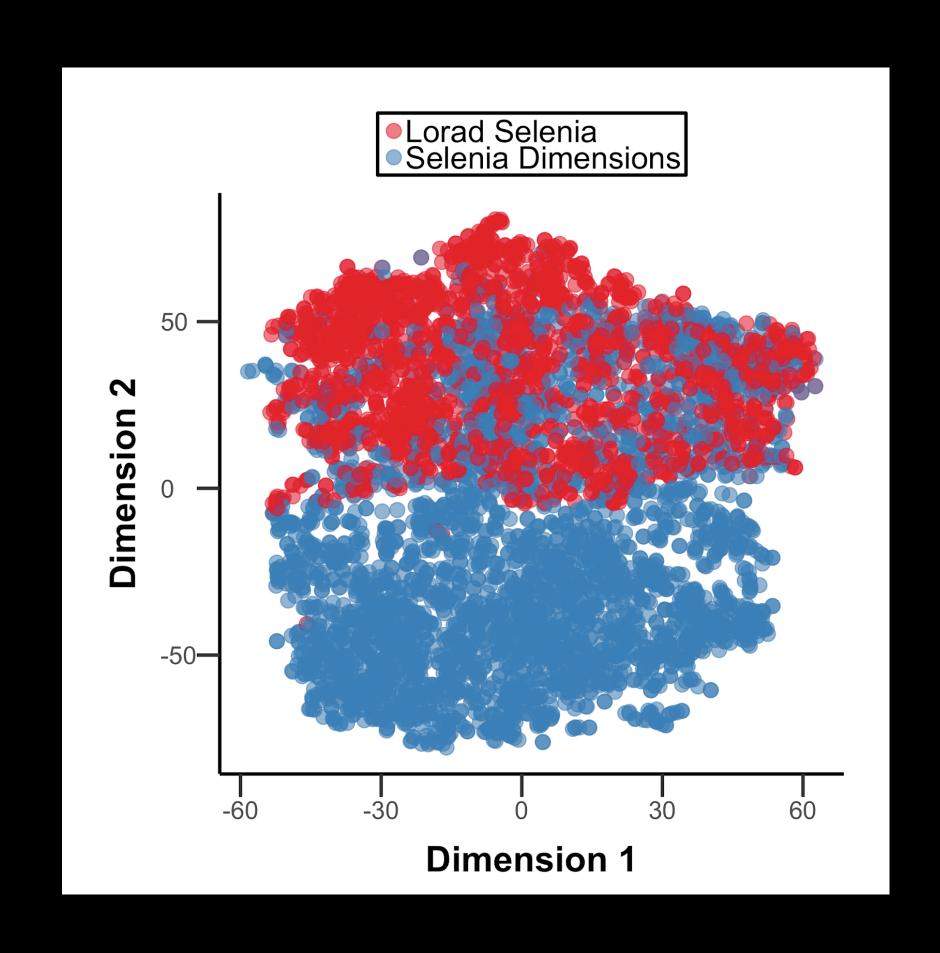


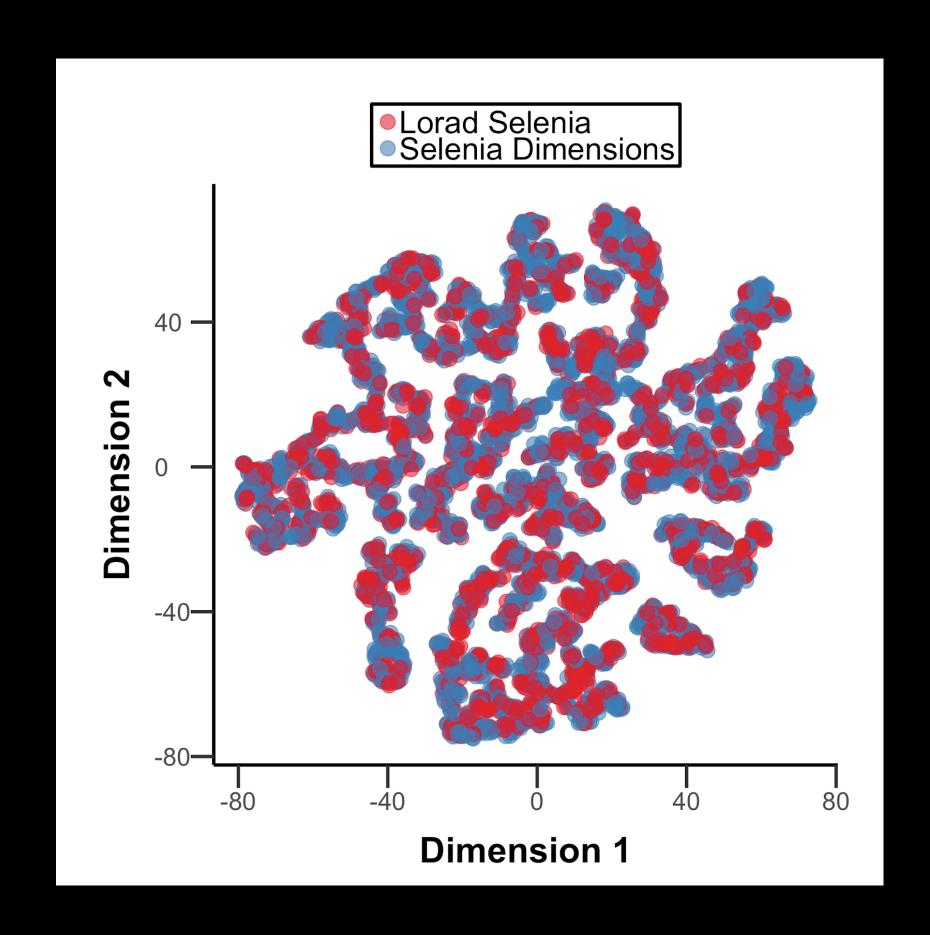
Predicting which device was used: d = D(y, h)

Objective:

Max accuracy **Predictor**Min accuracy **Discriminator**

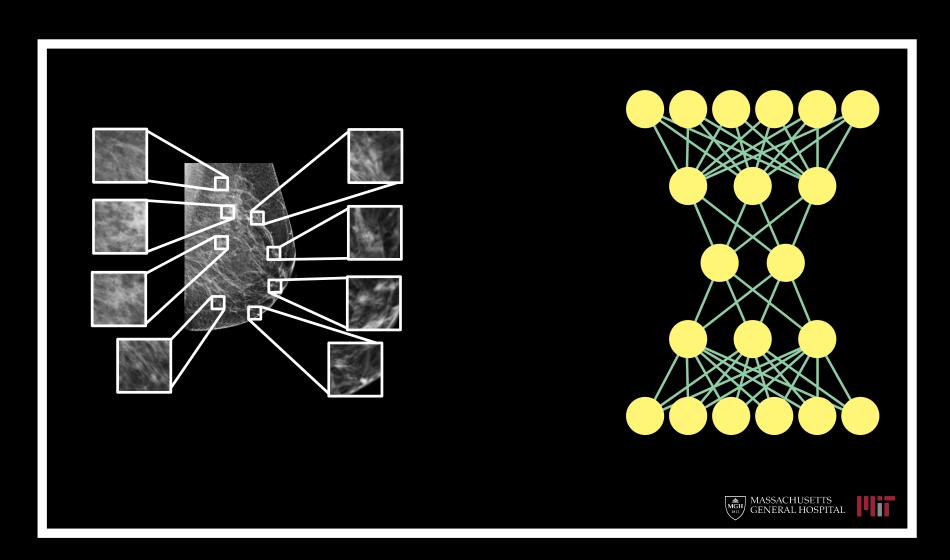
Anti-Bias Training Algorithms: Learning Invariant Representations





When Data Bias is Unknown

Case 1: Mystery of 99%



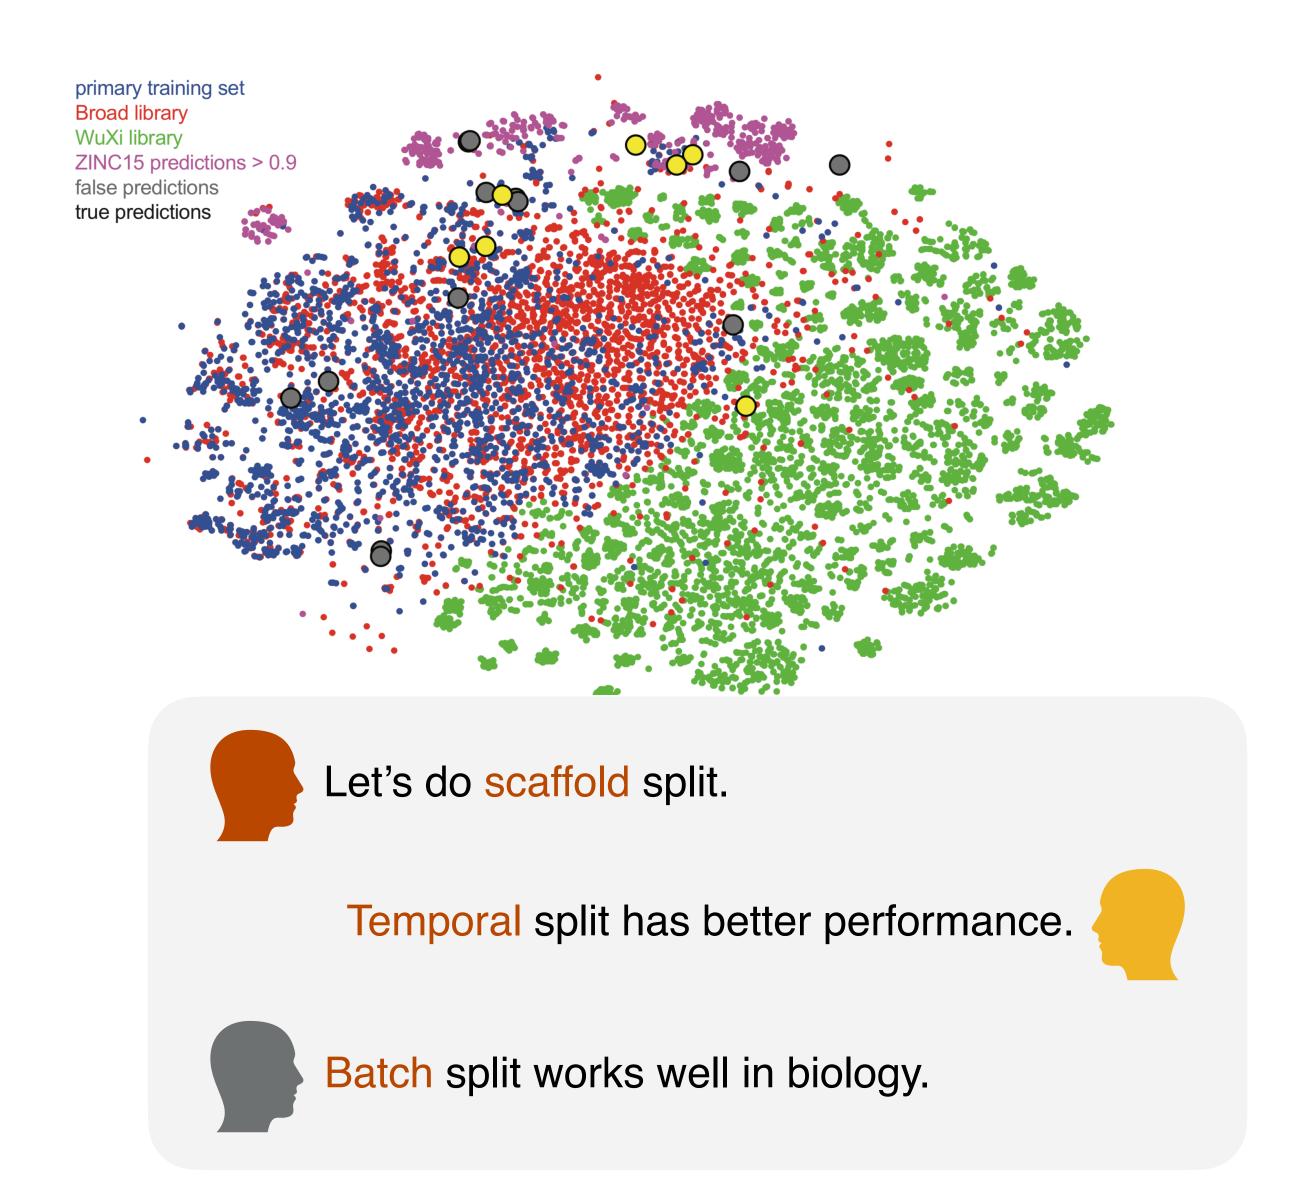
Case 2: Learning from Wrong Information

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

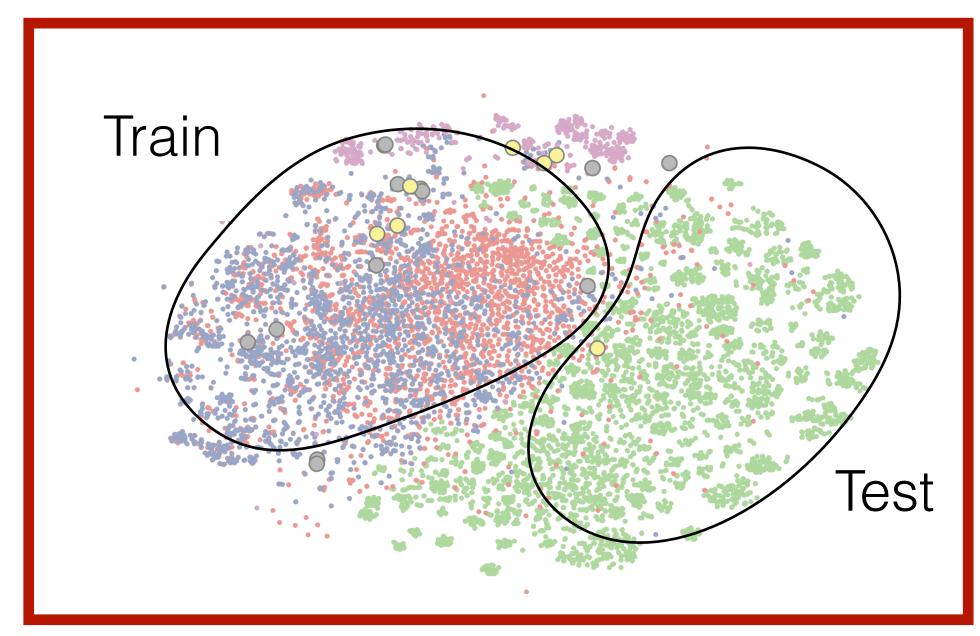
John R. Zech 🔤, Marcus A. Badgeley 🔤, Manway Liu, Anthony B. Costa, Joseph J. Titano, Eric Karl Oermann 🖾

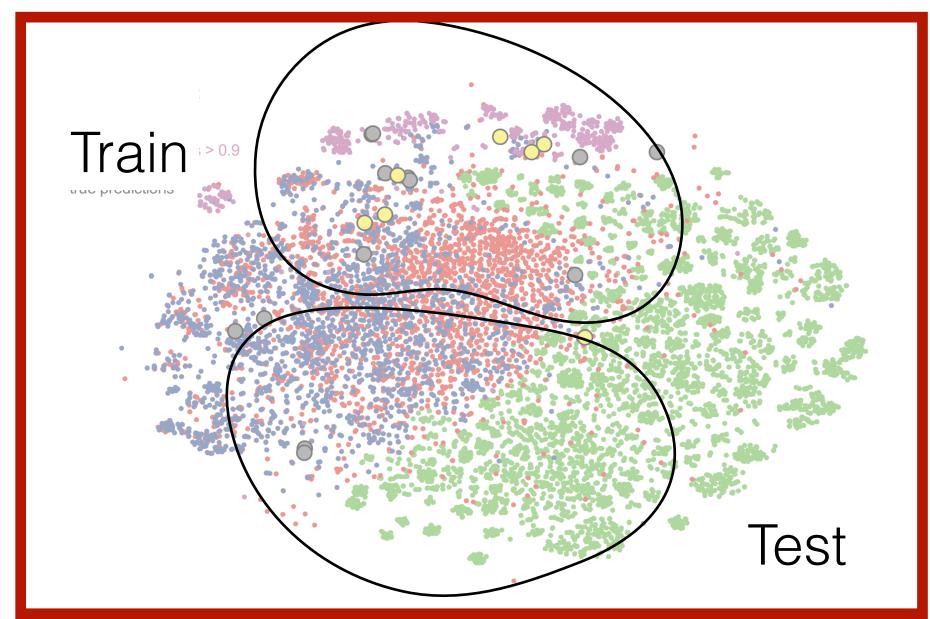
Published: November 6, 2018 • https://doi.org/10.1371/journal.pmed.1002683

Data Split: Art or Science?



Splitting Strategies Vary



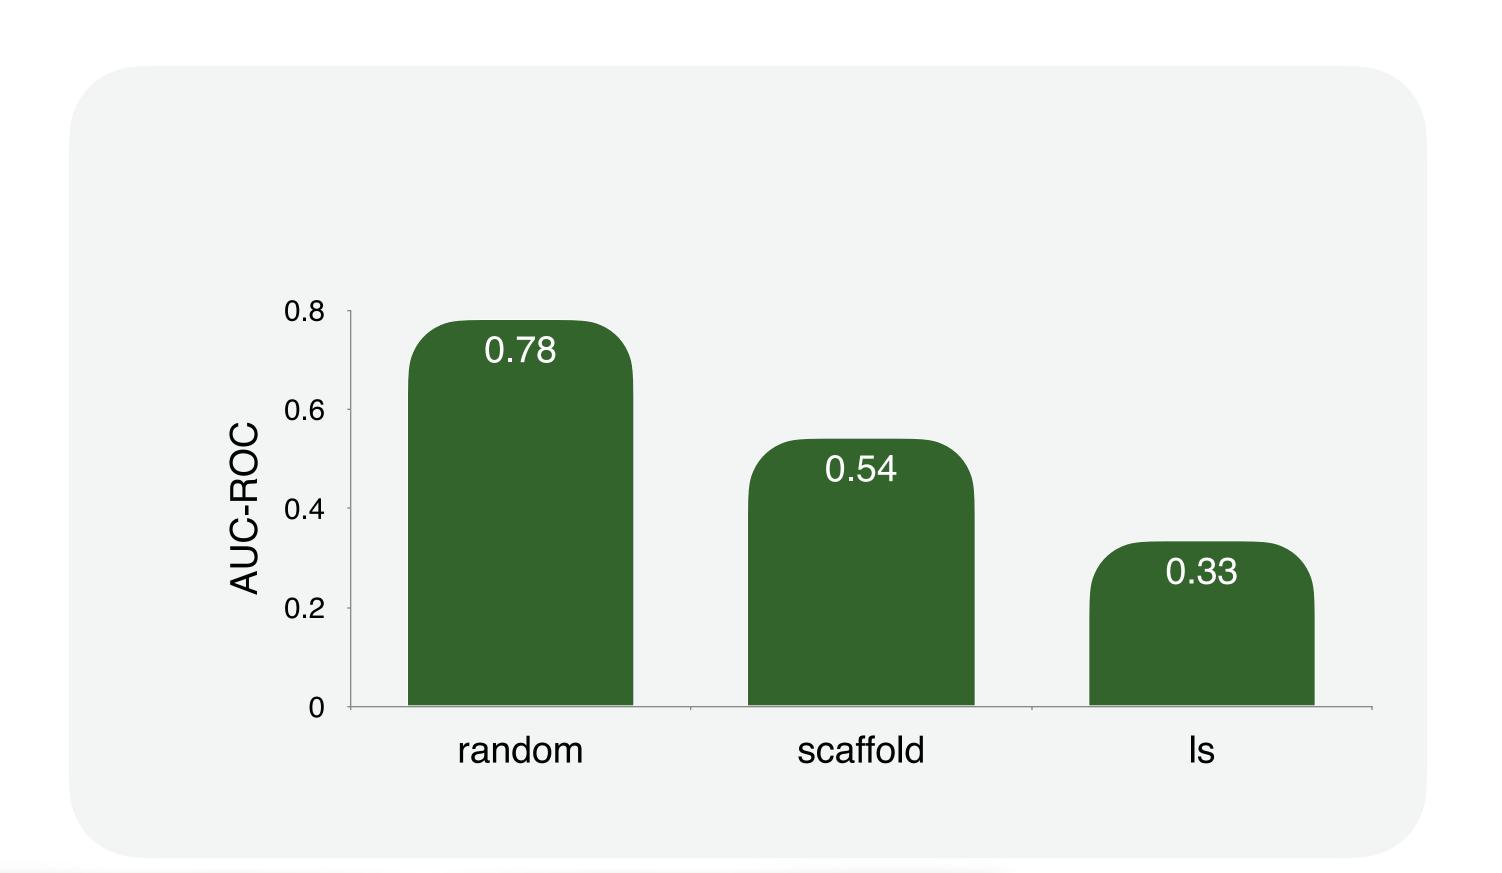


Scaffold Split

Random Split

- Standard in Chemistry: scaffold split
- Split is driven by graph similarity

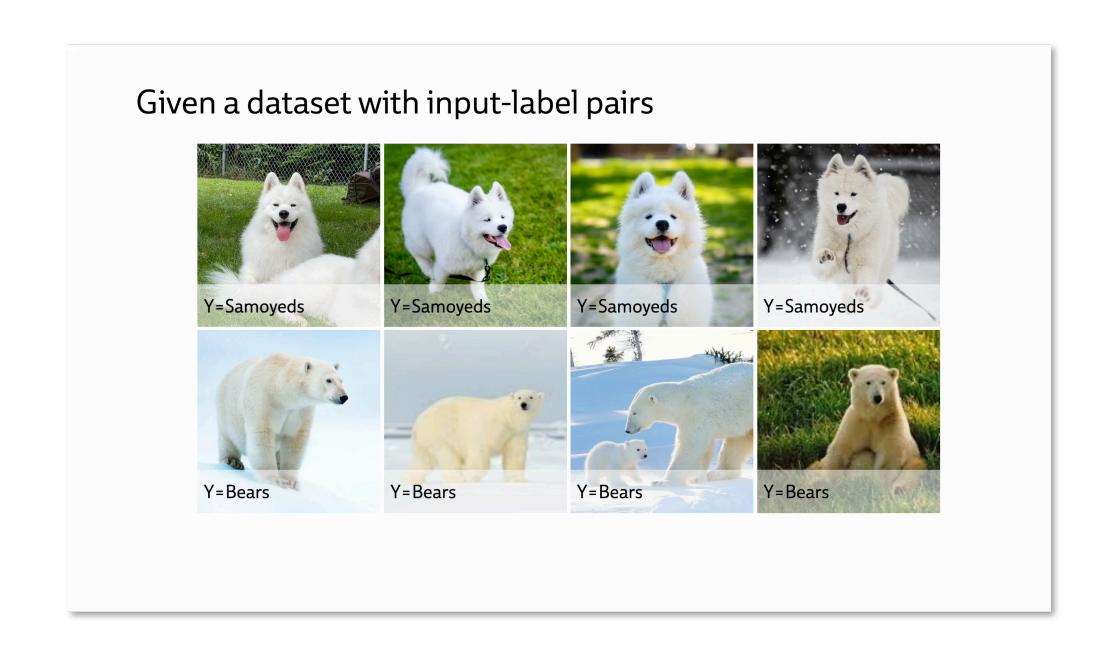
Finding Harder Splits



Learning to Split for Automatic Bias Detection

Yujia Bao, Regina Barzilay

Algortimic Soltion #2: Learning to Split for Automatic Bias Detection



Learn to split the dataset such that predictors trained on the training set cannot generalize to the testing set.

Bi-level optimization

Splitter

Generate a train-test split of the dataset.

Predictor

Minimize the loss on the training set.*

Estimate the generalization error on the testing set.

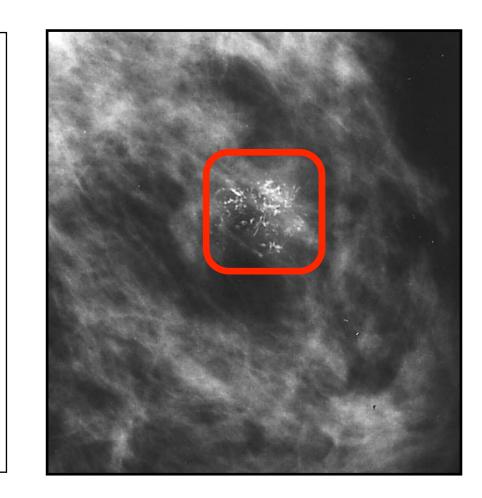
Maximize the generalization error subject to constraints

There is no evidence of extranodal extension.

BREAST (RIGHT) EXCISIONIAL RIOPSY:

ILIVASIVE DUCTAL CARCINOMA (SEE TABLE #1). DUCTAL CARCINOMA IN-SITU, GRADE 1. ATYPICAL DUCTAL HYPERPLASIA. LOBULAR NEOPLASIA (ATYPICAL LOBULAR HYPERPLASIA). TABLE OF PATHOLOGICAL FINDINGS #1 INVASIVE CARCINOMA

.



VIEWPOINT | VOLUME 3, ISSUE 11, E745-E750, NOVEMBER 2021

The false hope of current approaches to explainable artificial intelligence in health care

3. Patient-Centered Approach Incorporating Transparency to Users

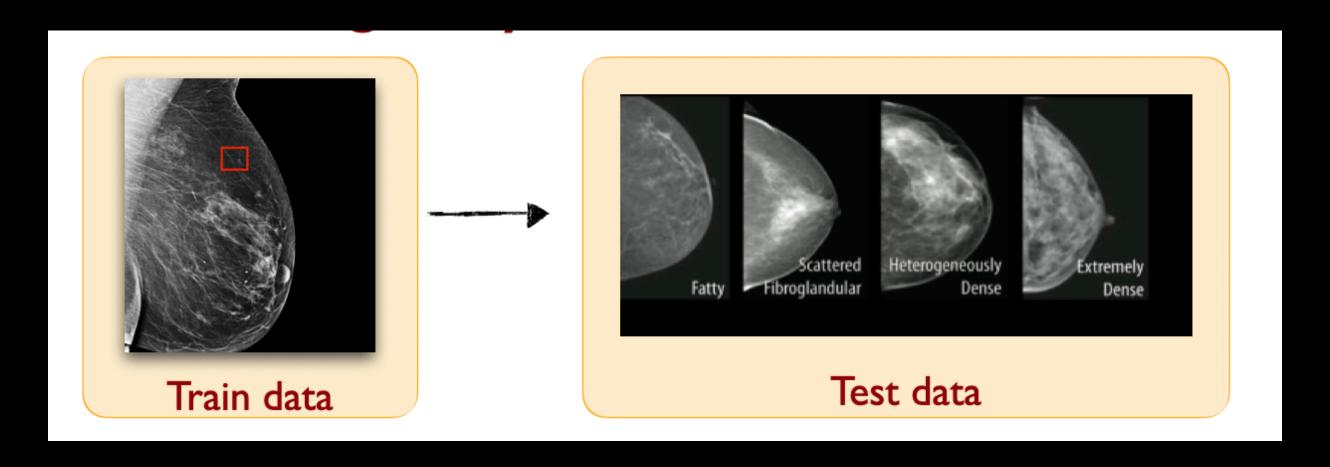
4

What we heard: Stakeholders called for further discussion with FDA on how AI/ML-based technologies interact with people, including their transparency to users and to patients more broadly.

What we'll do: Following up on the Agency's recent Patient Engagement Advisory Committee meeting on AI/ML-based devices, our next step will be to hold a public workshop on how device labeling supports transparency to users and enhances trust in AI/ML-based devices

The Agency acknowledges that AI/ML-based devices have unique considerations that necessitate a proactive patient-centered approach to their development and utilization that takes into account issues including usability, equity, trust, and accountability. One way that FDA is addressing these issues is through the promotion of the transparency of these devices to users, and to patients more broadly,

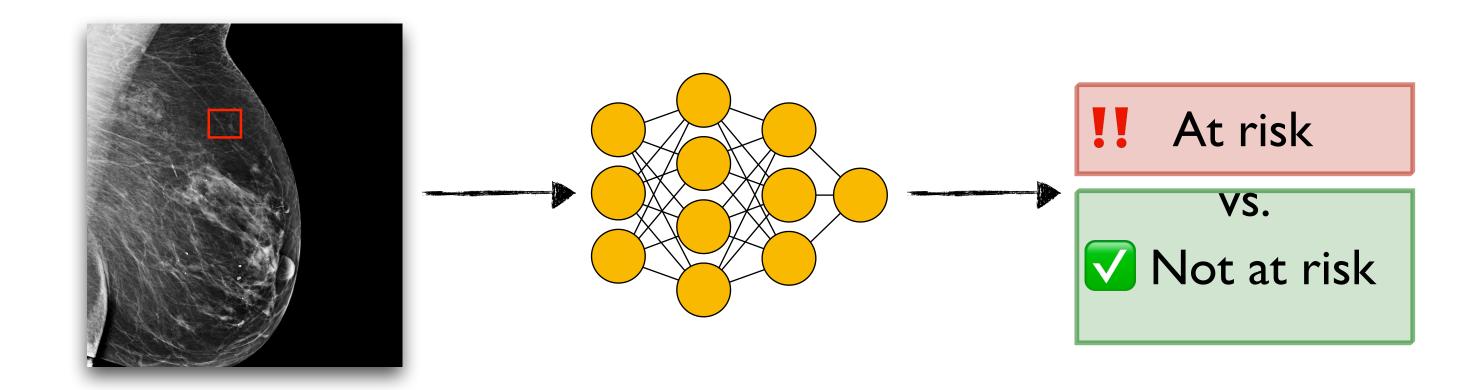
Selective Classification





Learn to abstain on uncallibrated examples

Algortimic Solution #3: Out-of-domain Calibration



- Suppose we train a model. When can we trust its predictions?
- We can answer these questions if predictions are well-calibrated.

Calibrated Selective Classification

Adam Fisch, Tommi S. Jaakkola, Regina Barzilay, January 2023 [openreview] [pdf] [bib] [code]

Algorithmic Solutions for Safe and Equitable AI development

Uncertainty Quantification

Anti-Bias Training

Distribution Shift Detection

MIT-MGB Al Cures Conference

April 24, 2023 | Samberg Center, MIT

Featured Speakers

Anne Klibanski

President & CEO, Mass General Brigham

Cynthia Barnhart

Paul Anderson

Acting Chief Academic Officer, MGB

Distinguished Professor & Faculty Lead, MIT

Polina Golland

David Sontag

Martha Jones

Emery Brown

Bruce Fischl

George Washko

Megha Gupta Obstetrics & Gynecology Specialist, BIDMC

JAMEEL CLINIC

Mass General Brigham

What is Al Cures?

patient populations. Hosted by MIT

Jameel Clinic, Al Cures showcases

edge research and their impact on patient care. Equally important is the process for collecting the data

algorithms, with a focus on privacy,

Follow @aihealthmit Follow @massgenbrigham

View the agenda at mitmgb.ai ightarrow

Register at mitmgb.ai

