Machine learning to personalize cancer care

Adam Yala, PhD

Assistant Professor

Computational
PRECISION HEALTH

Berkeley | UCSF

Personalized care as a computational problem

Patient

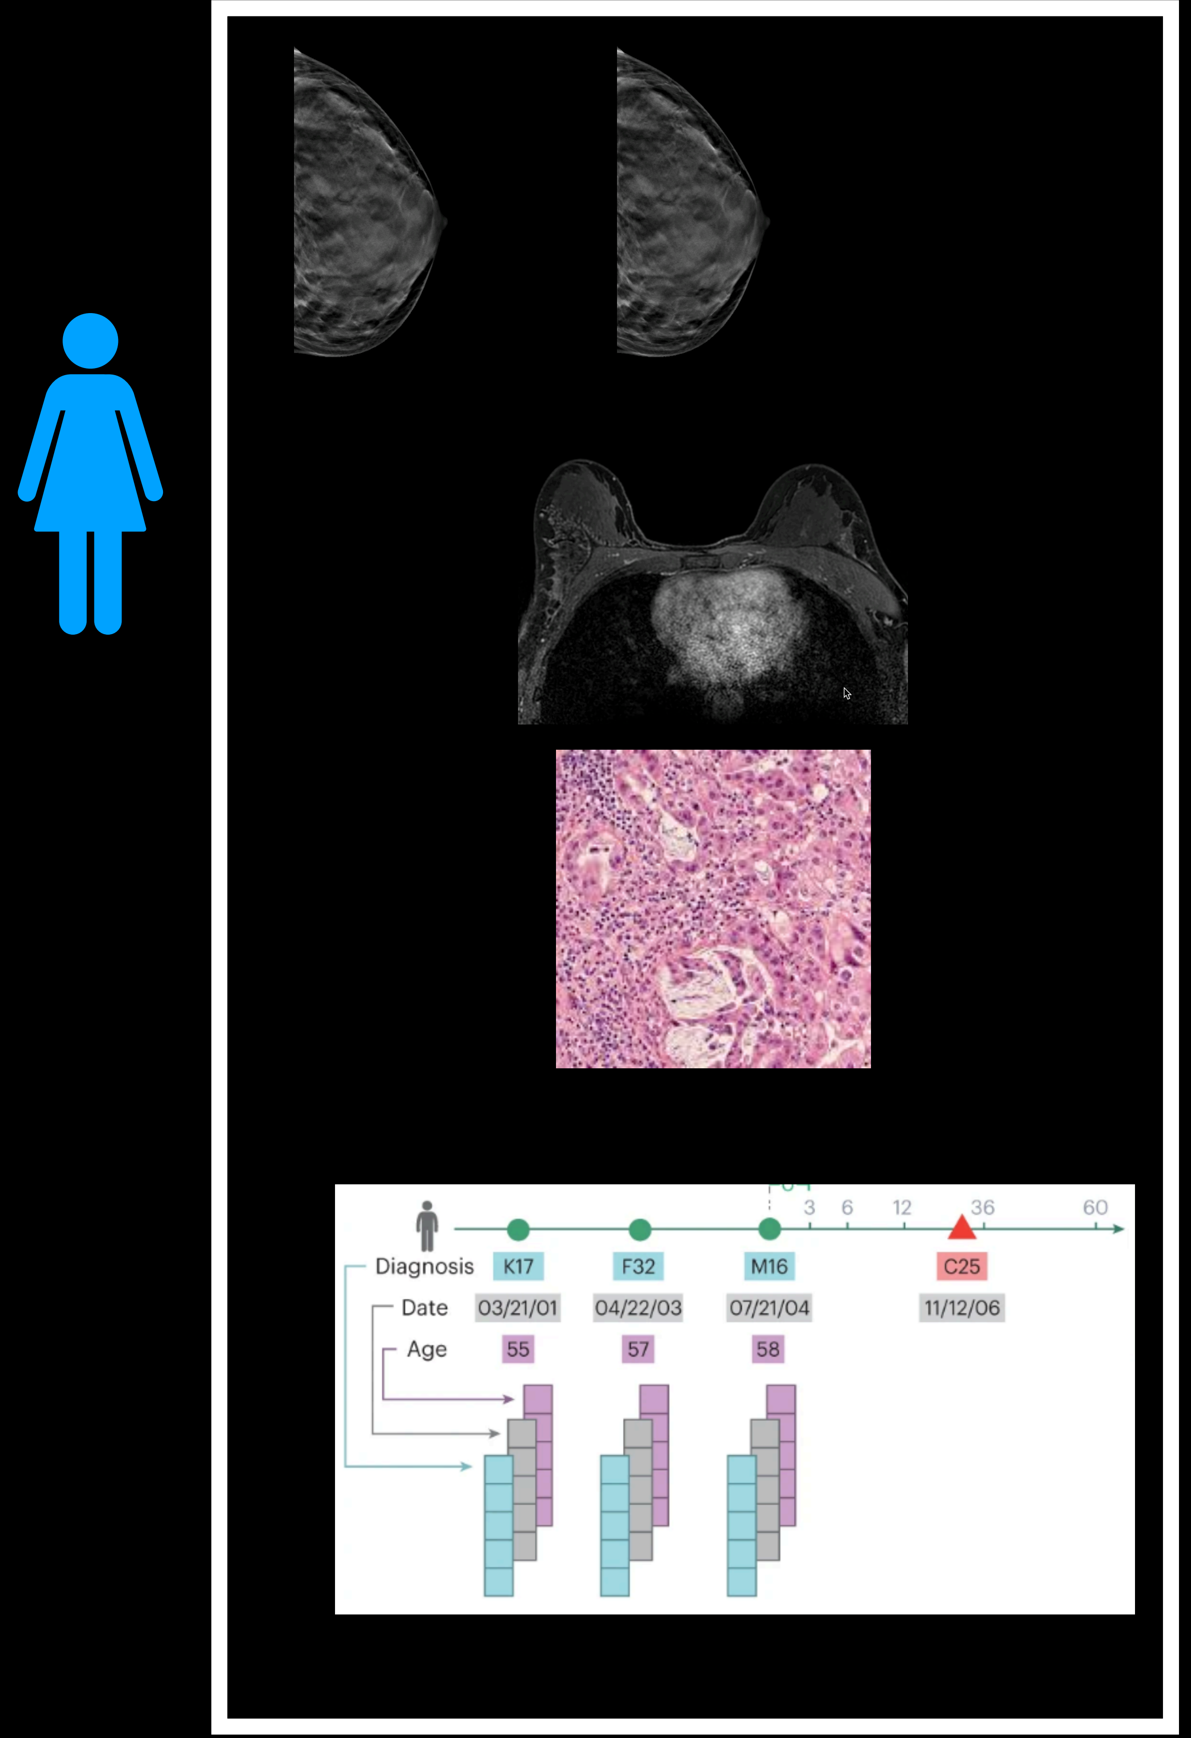

GB of data!

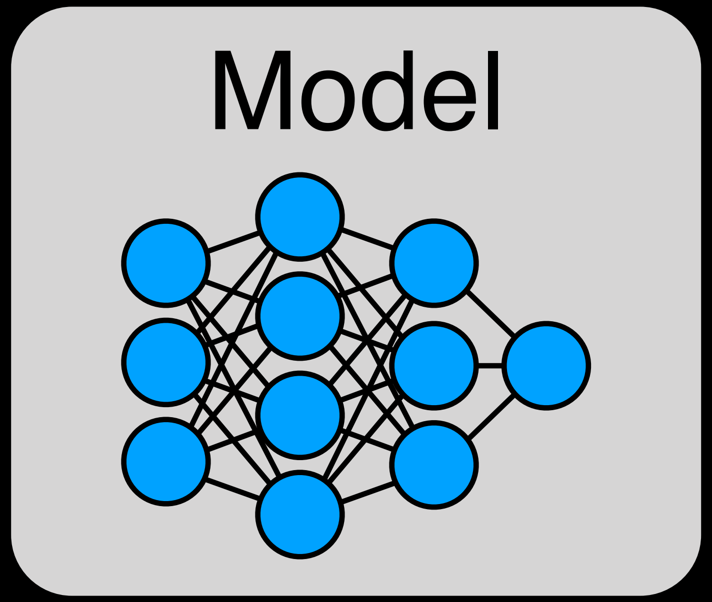

Intervention

Question: How do we use ***everything*** to recommend right intervention at right time?

Today: Towards AI-driven care

Prediction

Control

Translation

Today: Towards AI-driven care

Prediction

Motivating example: Screening today - one size fits all

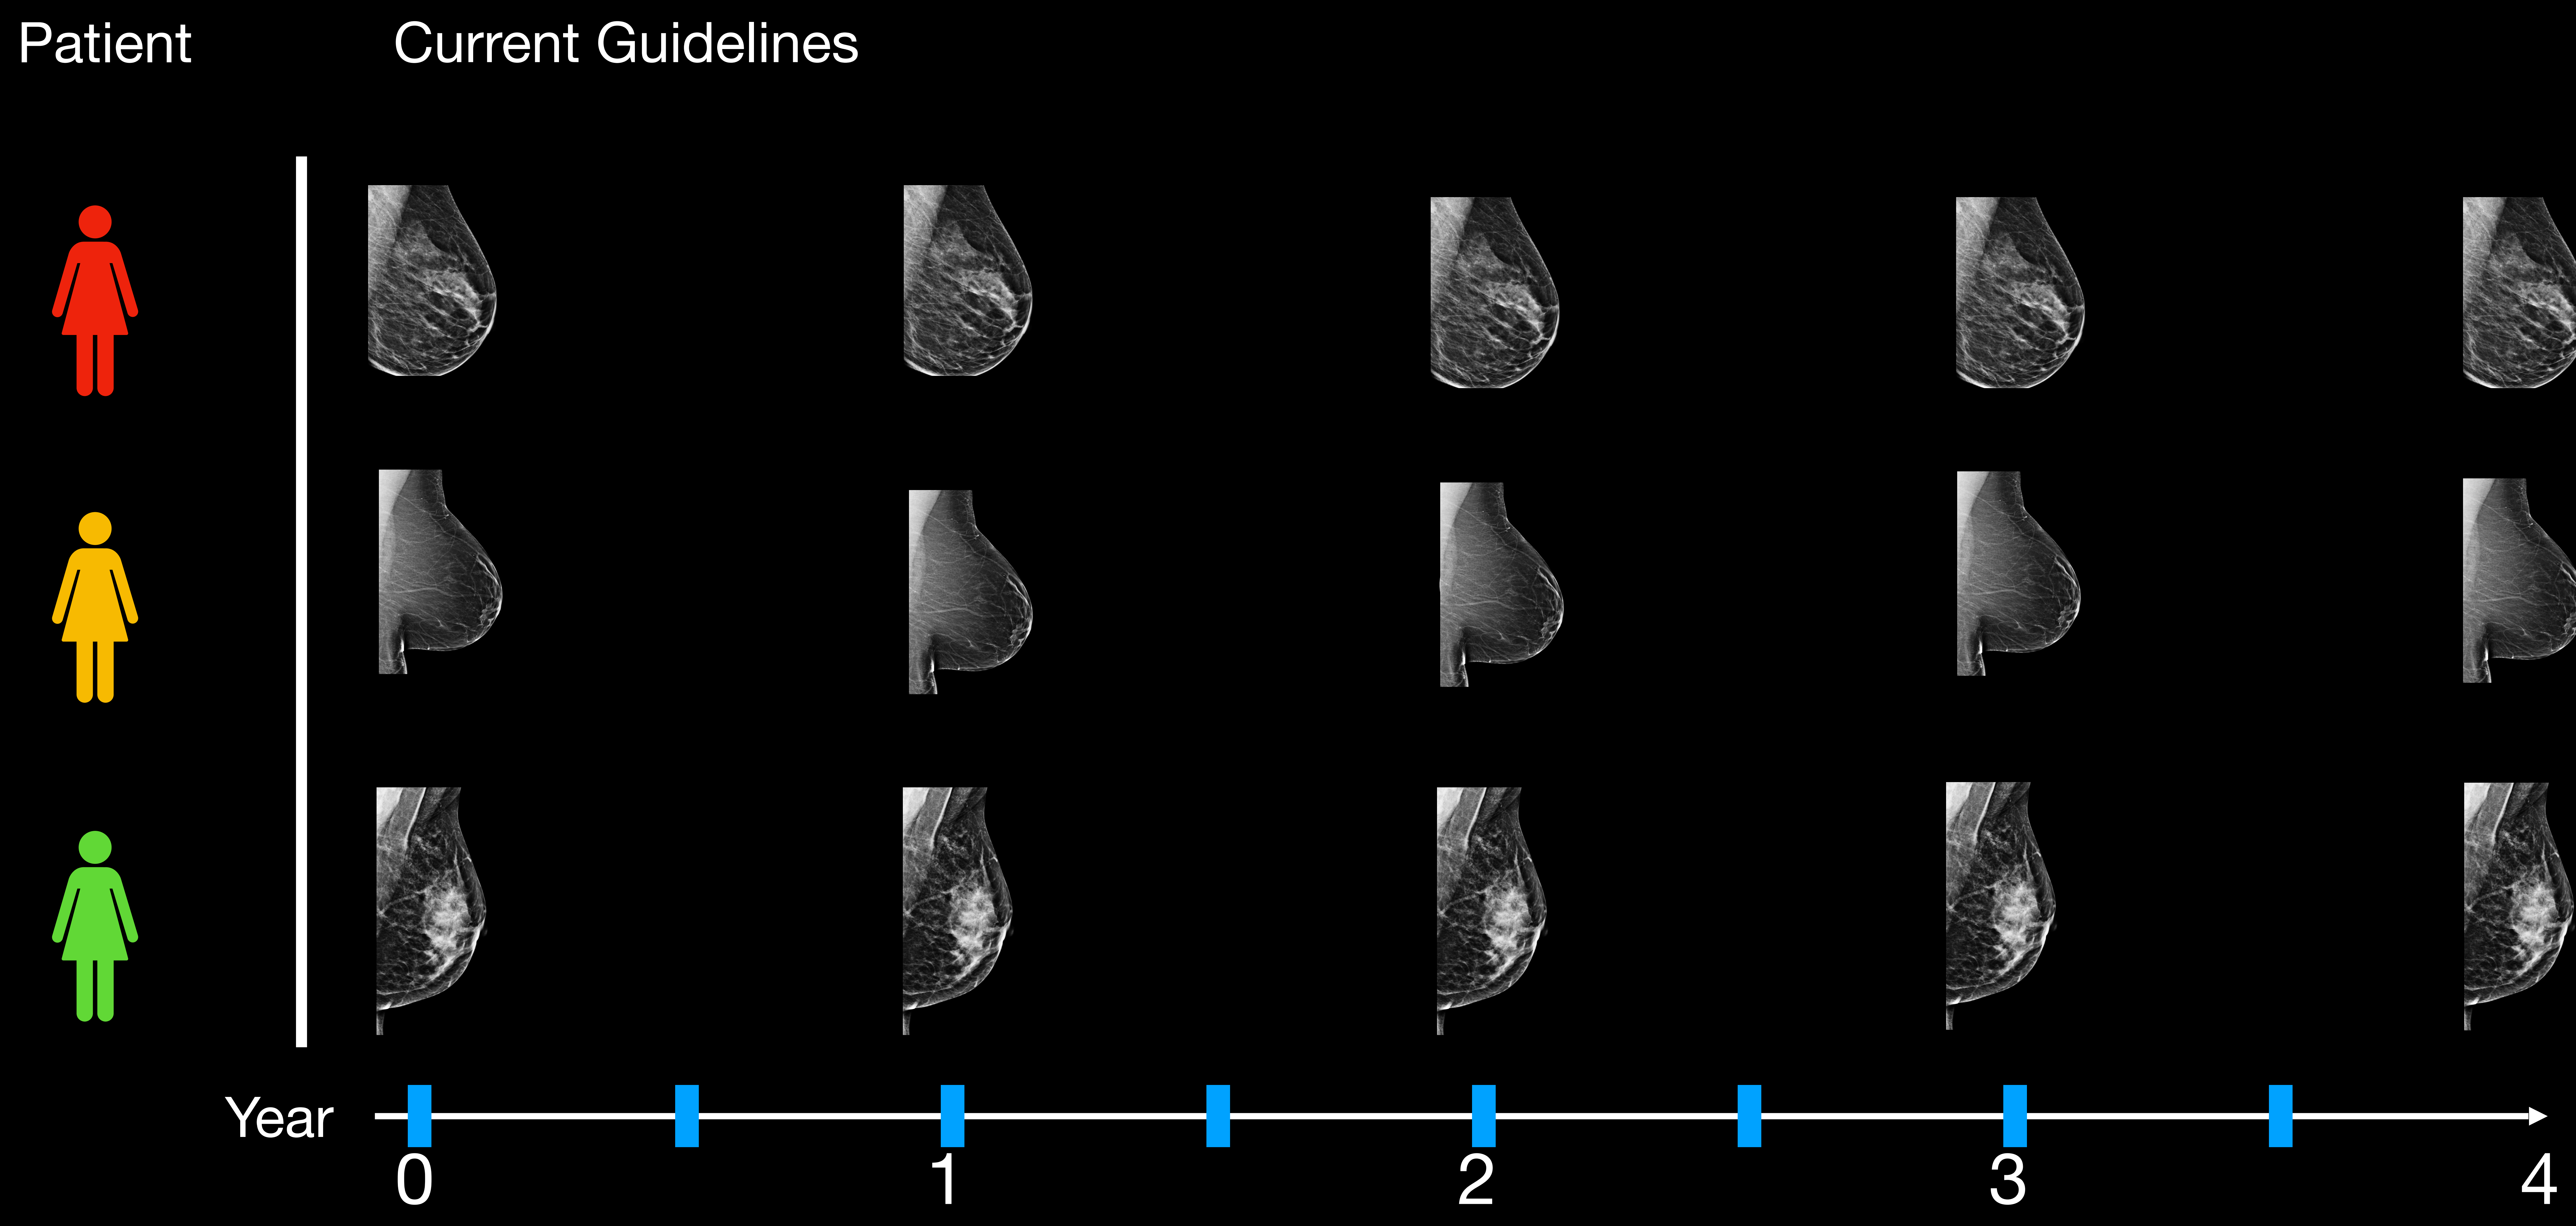

The harms of late diagnosis

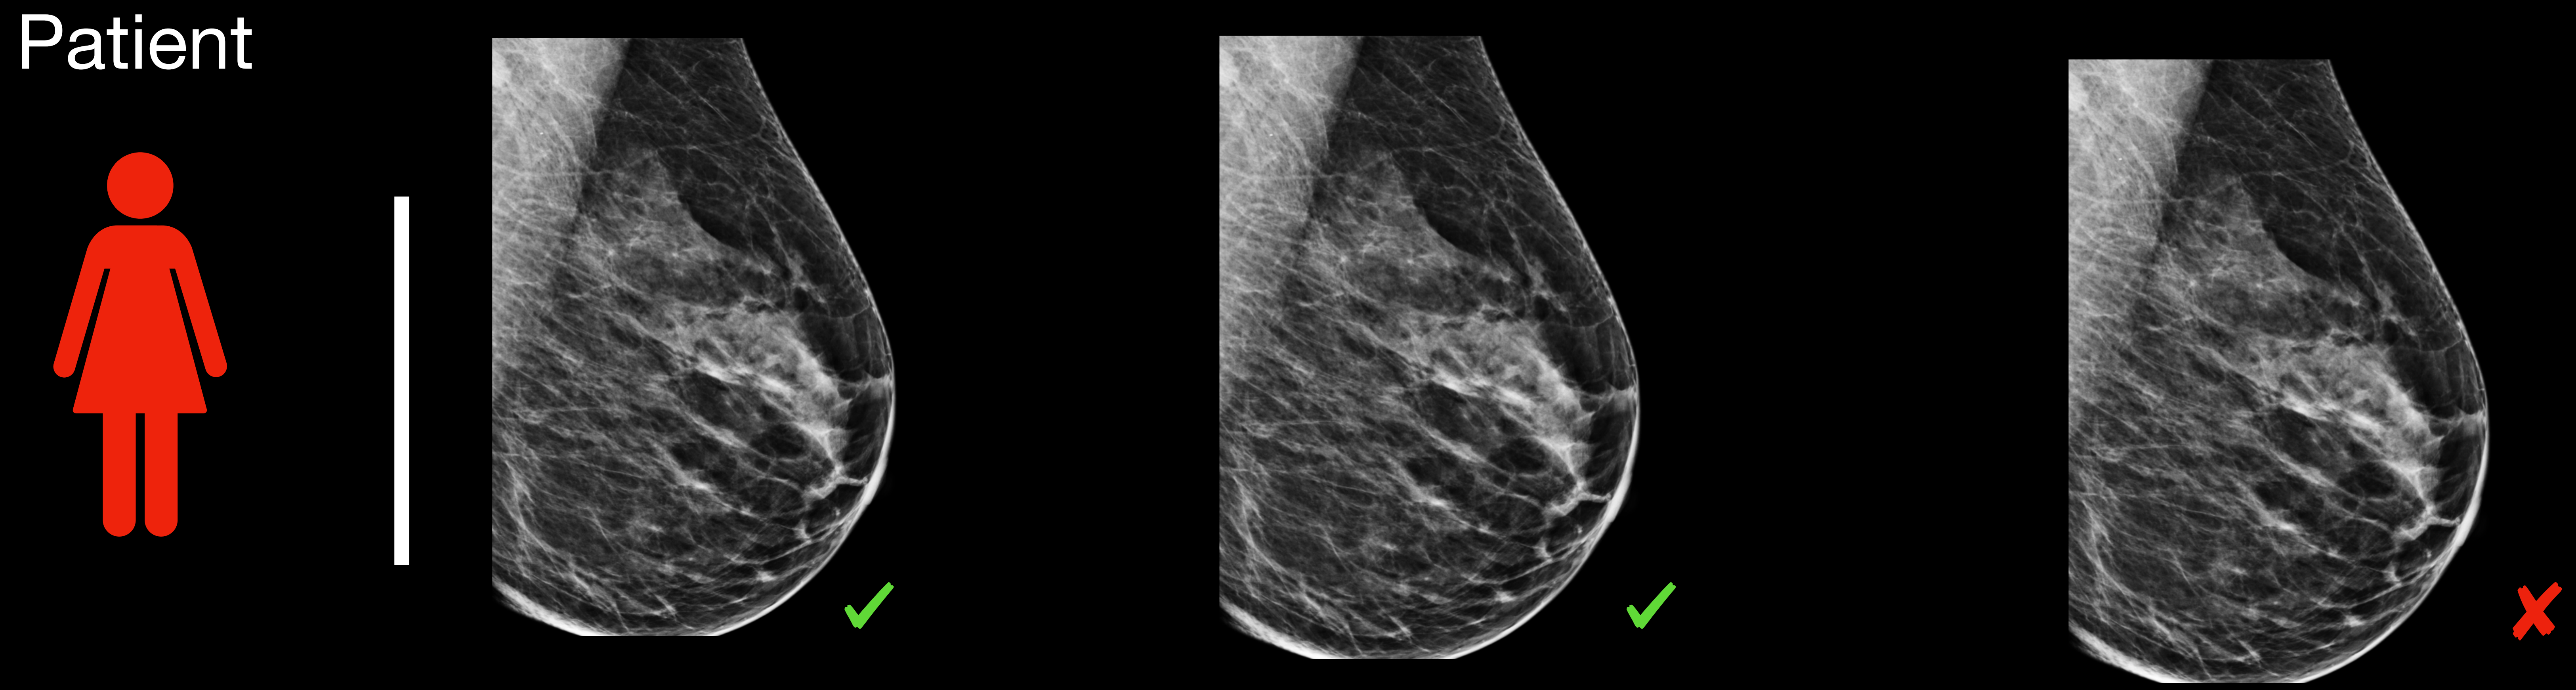

Morbid treatment options, poor chances of survival

We should have done more

The harms of over screening

Patient

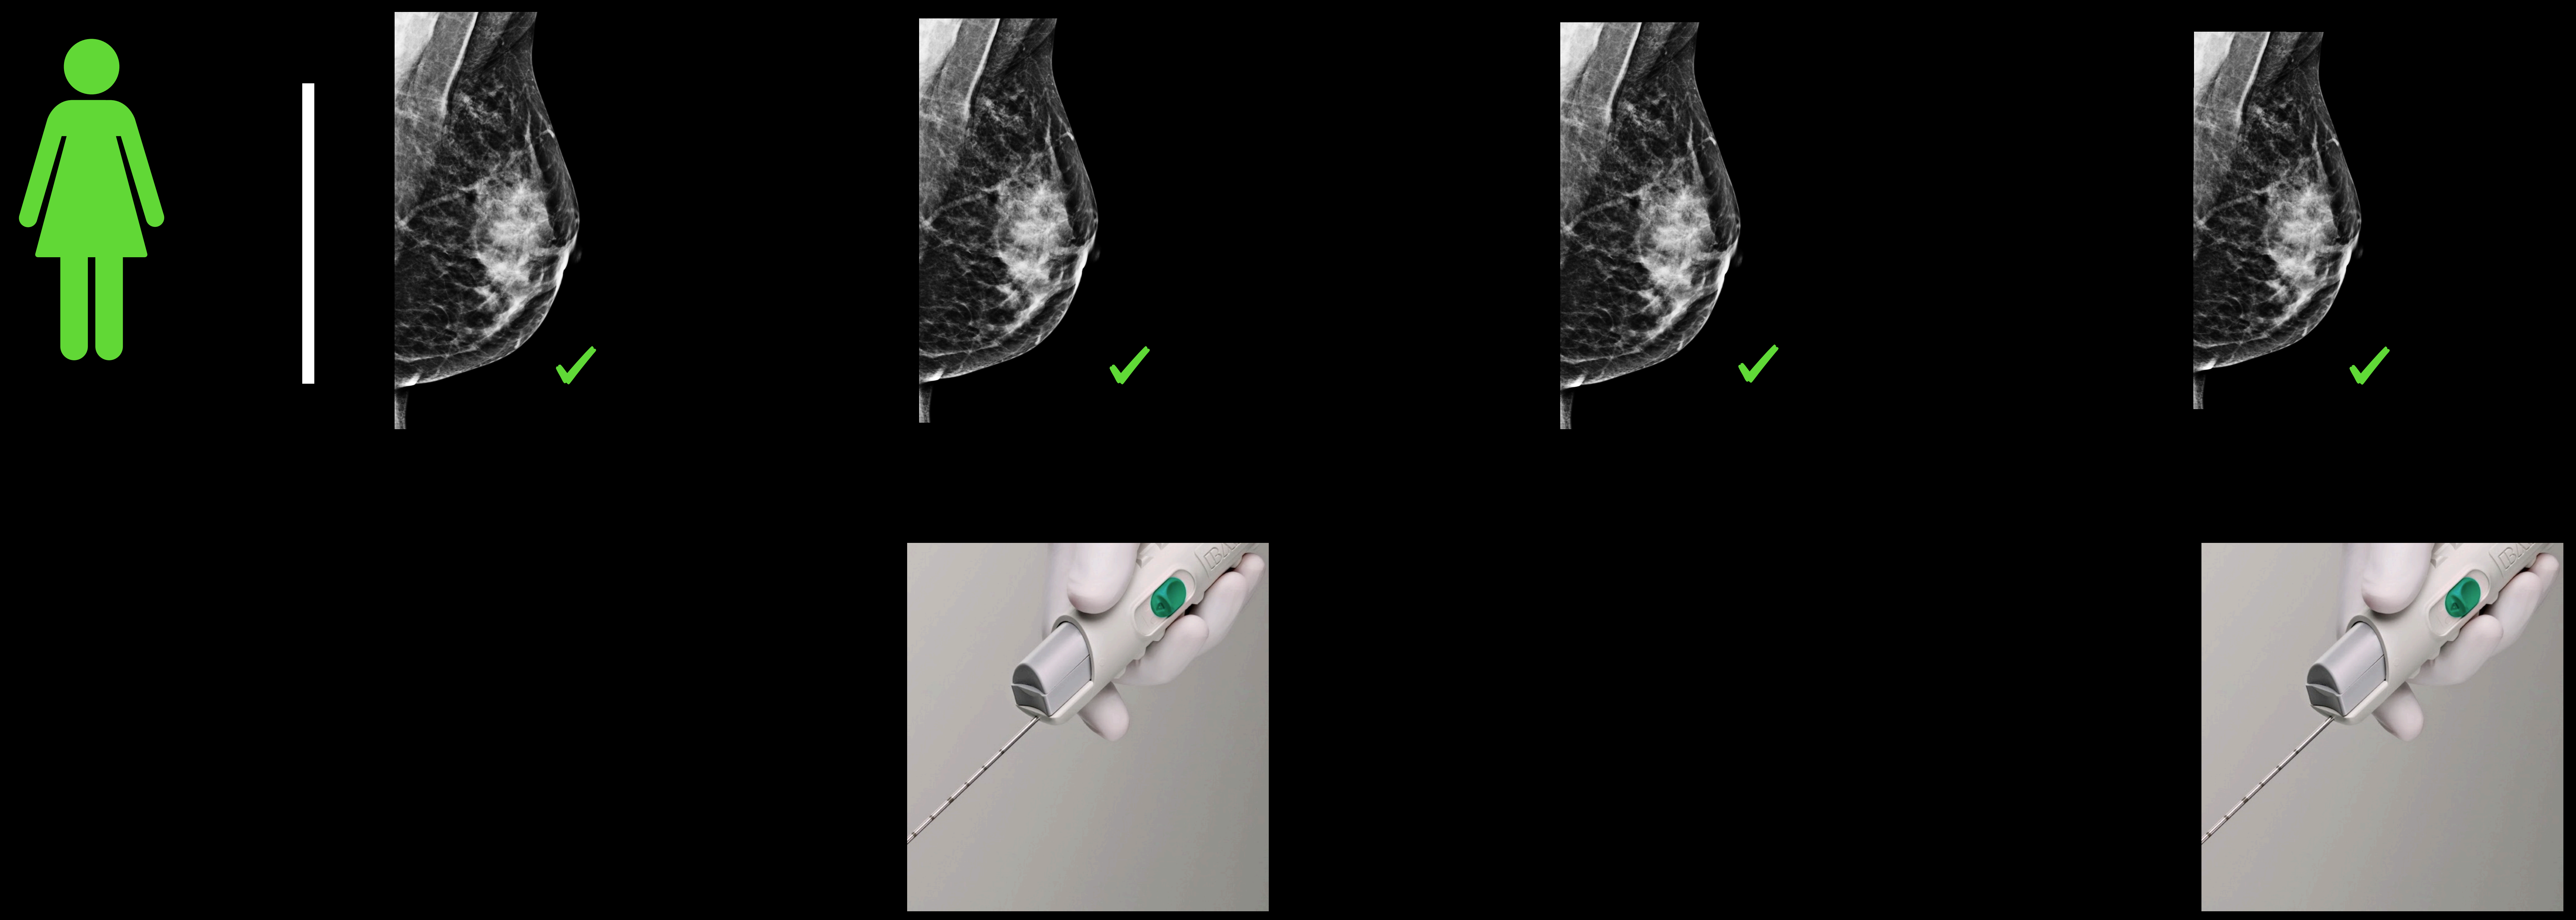

Unnecessary biopsies, terrible anxiety

We should have done less

How to catch cancer earlier

Predict Cancer Risk

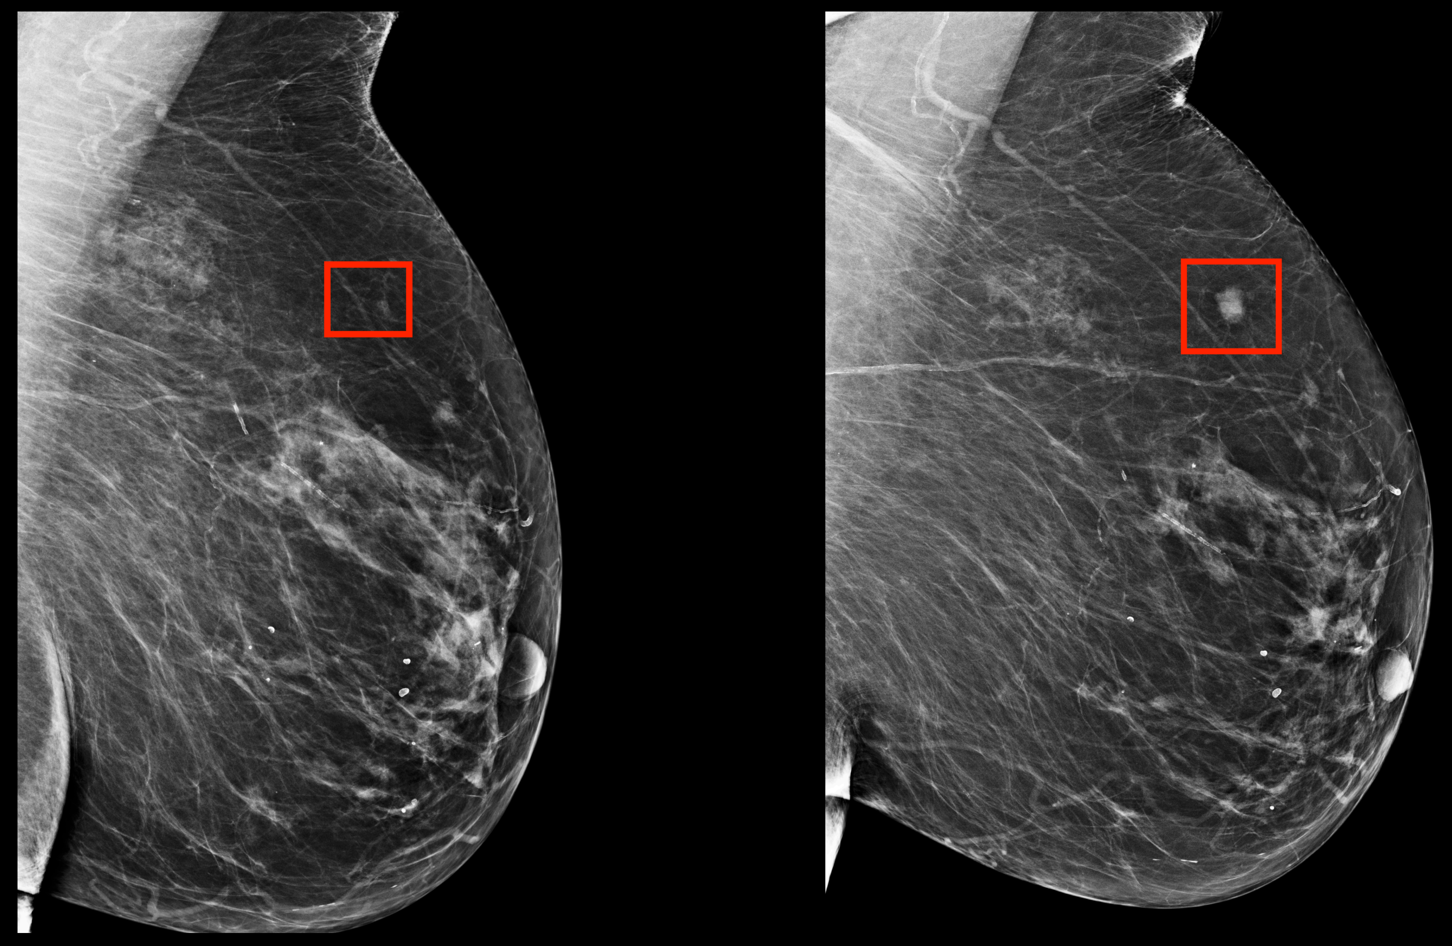

Year 0 Year 5

Create personalized screening policy

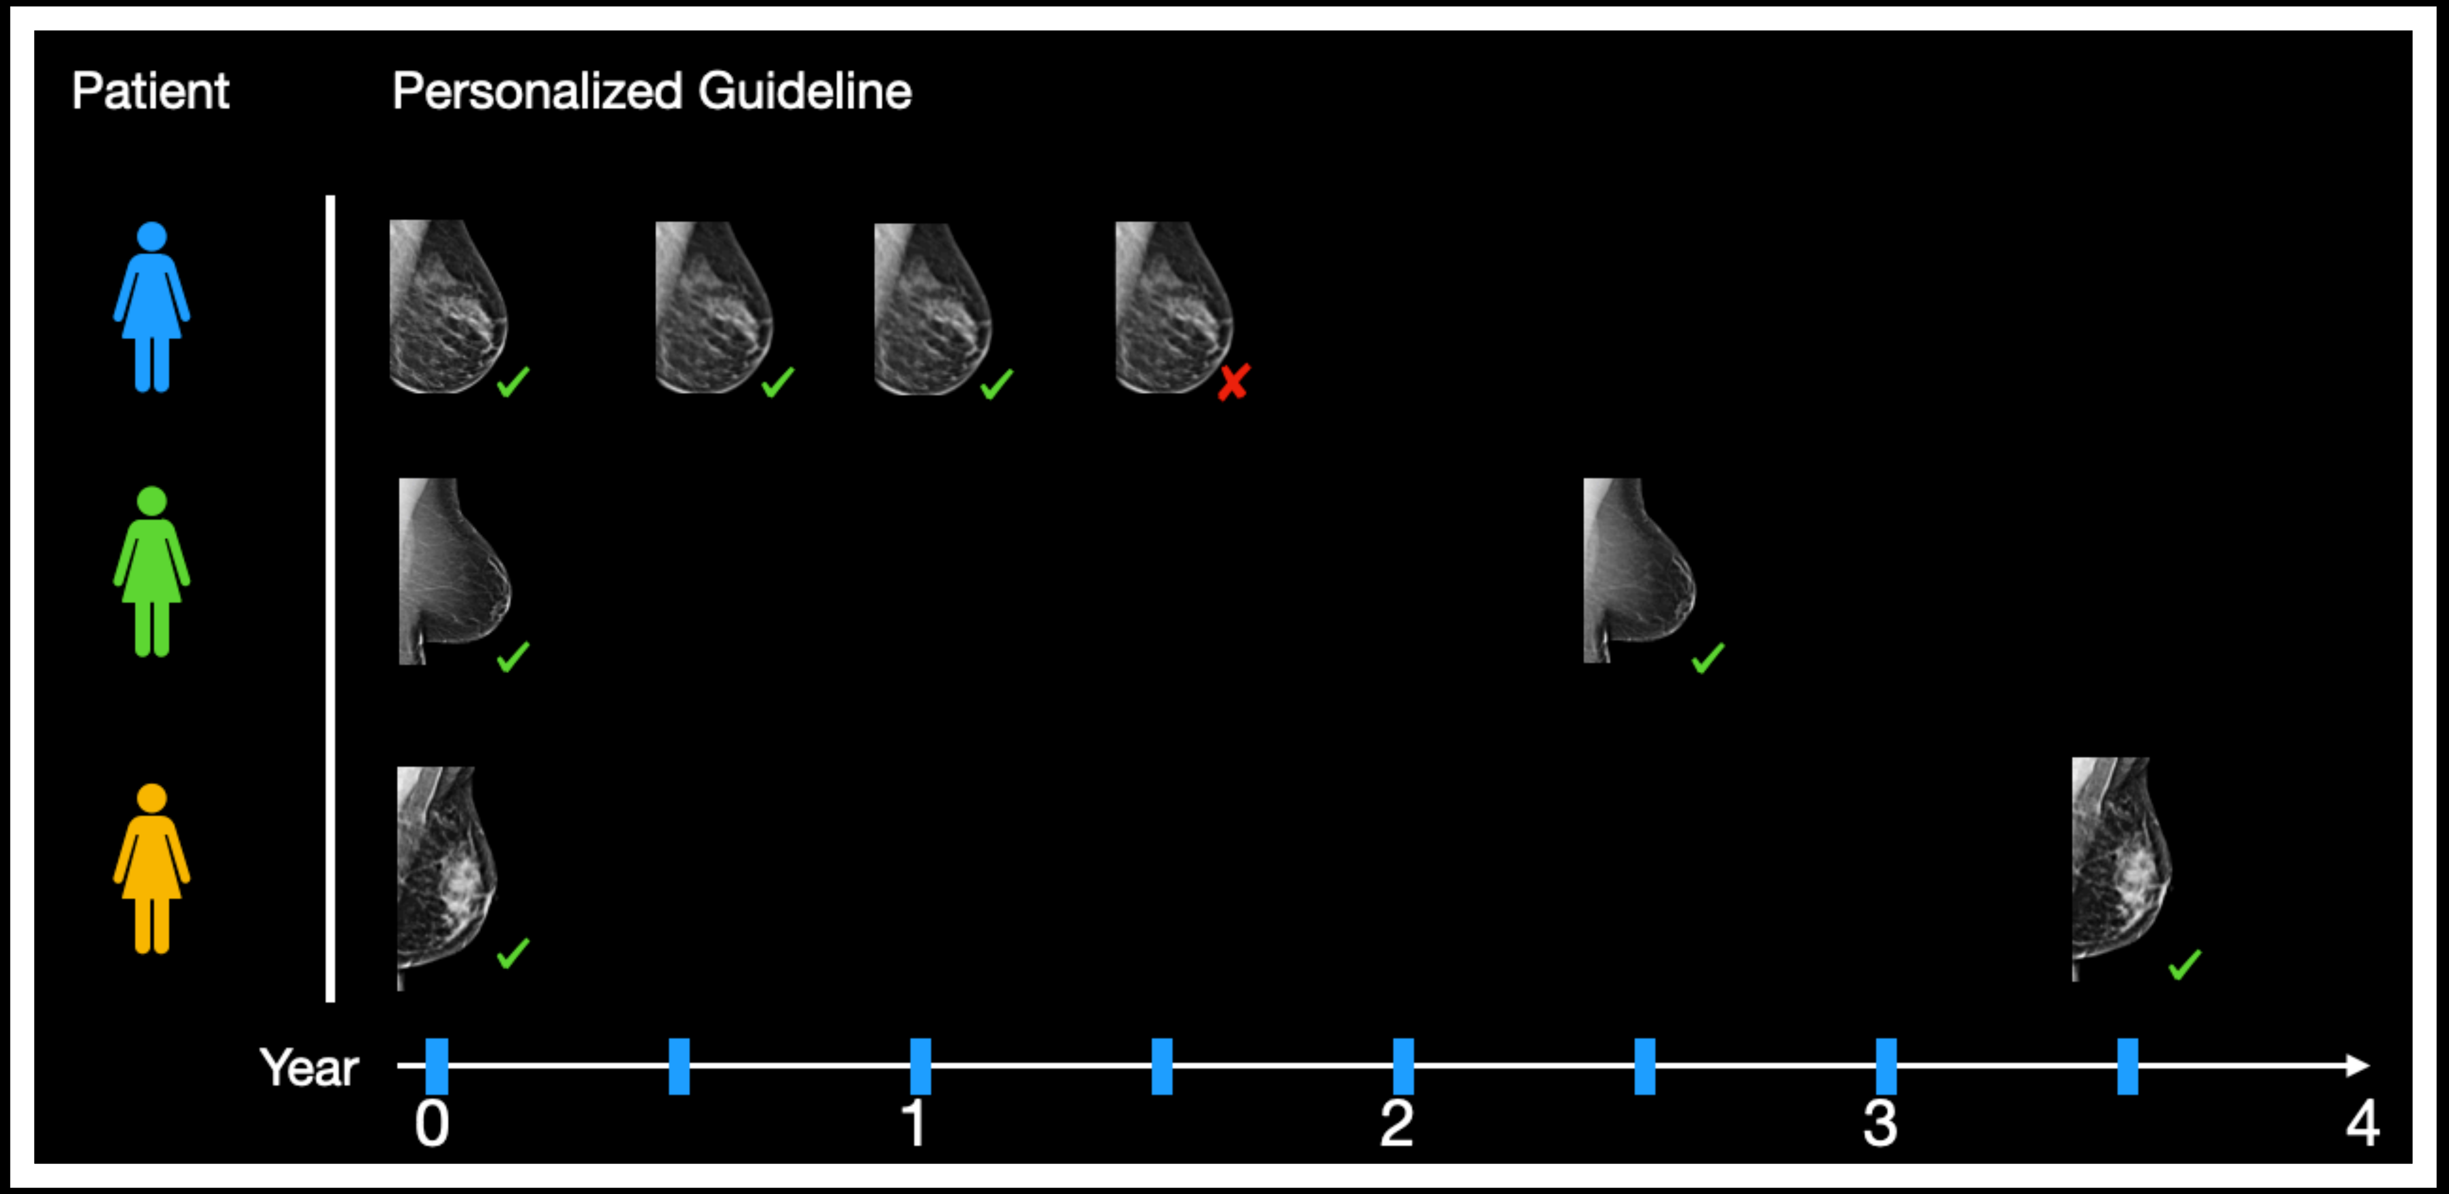

Obstacle: Current clinical tools are insufficient

Questionnaire based (<< 1KB of data per patient)

Age
Family History
Prior Breast Procedure
Breast Density

→ Model → **Risk**

Identify **<25%** of future cancers as "high risk"

>95% of "high risk" patients won't get cancer

Personalized screening as a computational problem

Patient

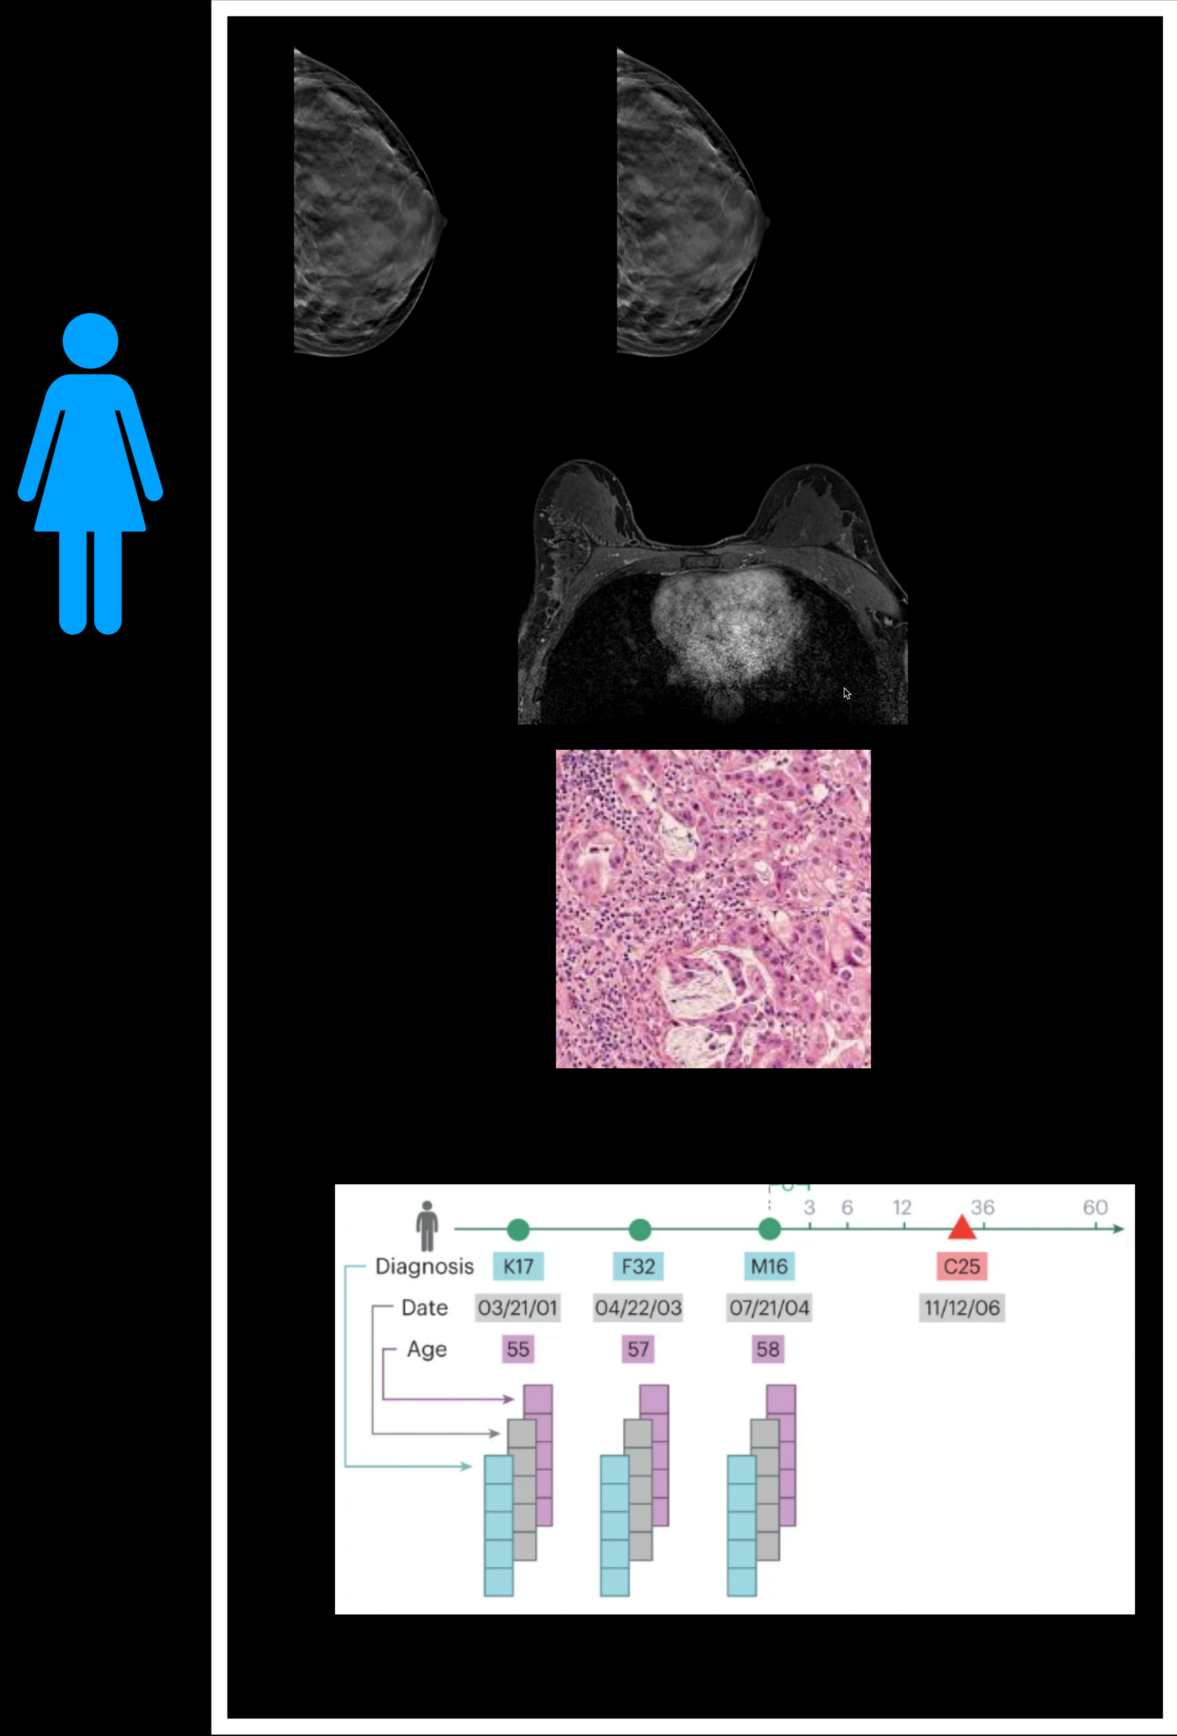

→

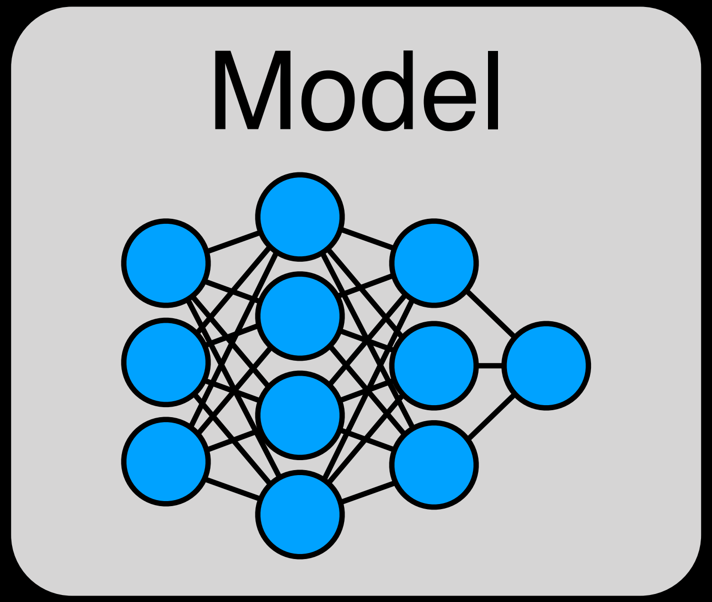

→ Screening plan

GB of data!

Where are we now? From bits to MB

Patient **Data**

Future Outcome

→ 3 year cancer

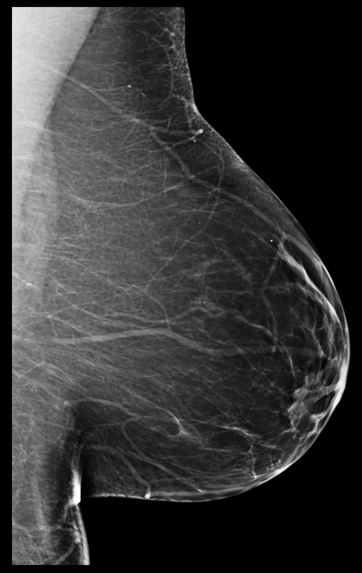

→ No cancer

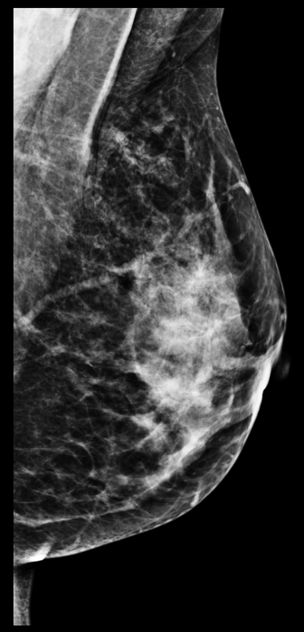

→ 5 year cancer

MB of data per patient

Journal of Clinical Oncology®
An American Society of Clinical Oncology Journal

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

Check for updates

Adam Yala, MEng[1,2]; Peter G. Mikhael, BS[1,2]; Fredrik Strand, MD, PhD[3,4]; Gigin Lin, MD, PhD[5]; Siddharth Satuluru, BS[6];

SCIENCE TRANSLATIONAL MEDICINE

Toward robust mammography-based models for breast cancer risk

Adam Yala[1,2]*, Peter G. Mikhael[1,2], Fredrik Strand[3,4], Gigin Lin[5], Kevin Smith[6,7], Yung-Liang Leslie Lamb[8], Kevin Hughes[9], Constance Lehman[8†], Regina Barzilay[1,2†]

Radiology

A Deep Learning Mammography-based Model for Improved Breast Cancer Risk Prediction

Adam Yala, MEng • Constance Lehman, MD, PhD • Tal Schuster, MS • Tally Portnoi, BS • Regina Barzilay, PhD

Journal of Clinical Oncology®
An American Society of Clinical Oncology Journal

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

Peter G. Mikhael, BSc[1,2]; Jeremy Wohlwend, ME[1,2]; Adam Yala, PhD[1,2]; Ludvig Karstens, MSc[1,2]; Justin Xiang, ME[1,2]; Angelo K. Takigami, MD[3,4]; Patrick P. Bourgouin, MD[3,4]; PuiYee Chan, PhD[5]; Sofiane Mrah, MSc[4]; Wael Amayri, BSc[4]; Yu-Hsiang Juan, MD[6,7]; Cheng-Ta Yang, MD[6,8]; Yung-Liang Wan, MD[6,7]; Gigin Lin, MD, PhD[6,7]; Lecia V. Sequist, MD, MPH[3,5];

Aside: How can we curate outcomes at scale?

Data Curation

Human-level information extraction from clinical reports with fine-tuned language models

Longchao Liu, Long Lian, Yiyan Hao, Aidan Pace, Elaine Kim, Nour Homsi, Yash Pershad, Liheng Lai, Thomas Gracie, Ashwin Kishtagari, Peter R Carroll, Alexander G Bick, Anobel Y Odisho, Maggie Chung, Adam Yala

doi: https://doi.org/10.1101/2024.11.18.24317466

Led by:  Joy Liu 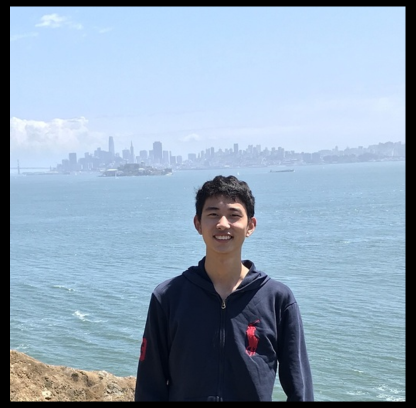 Tony Lian

Strata: Human-level performance for less than 5$ of compute

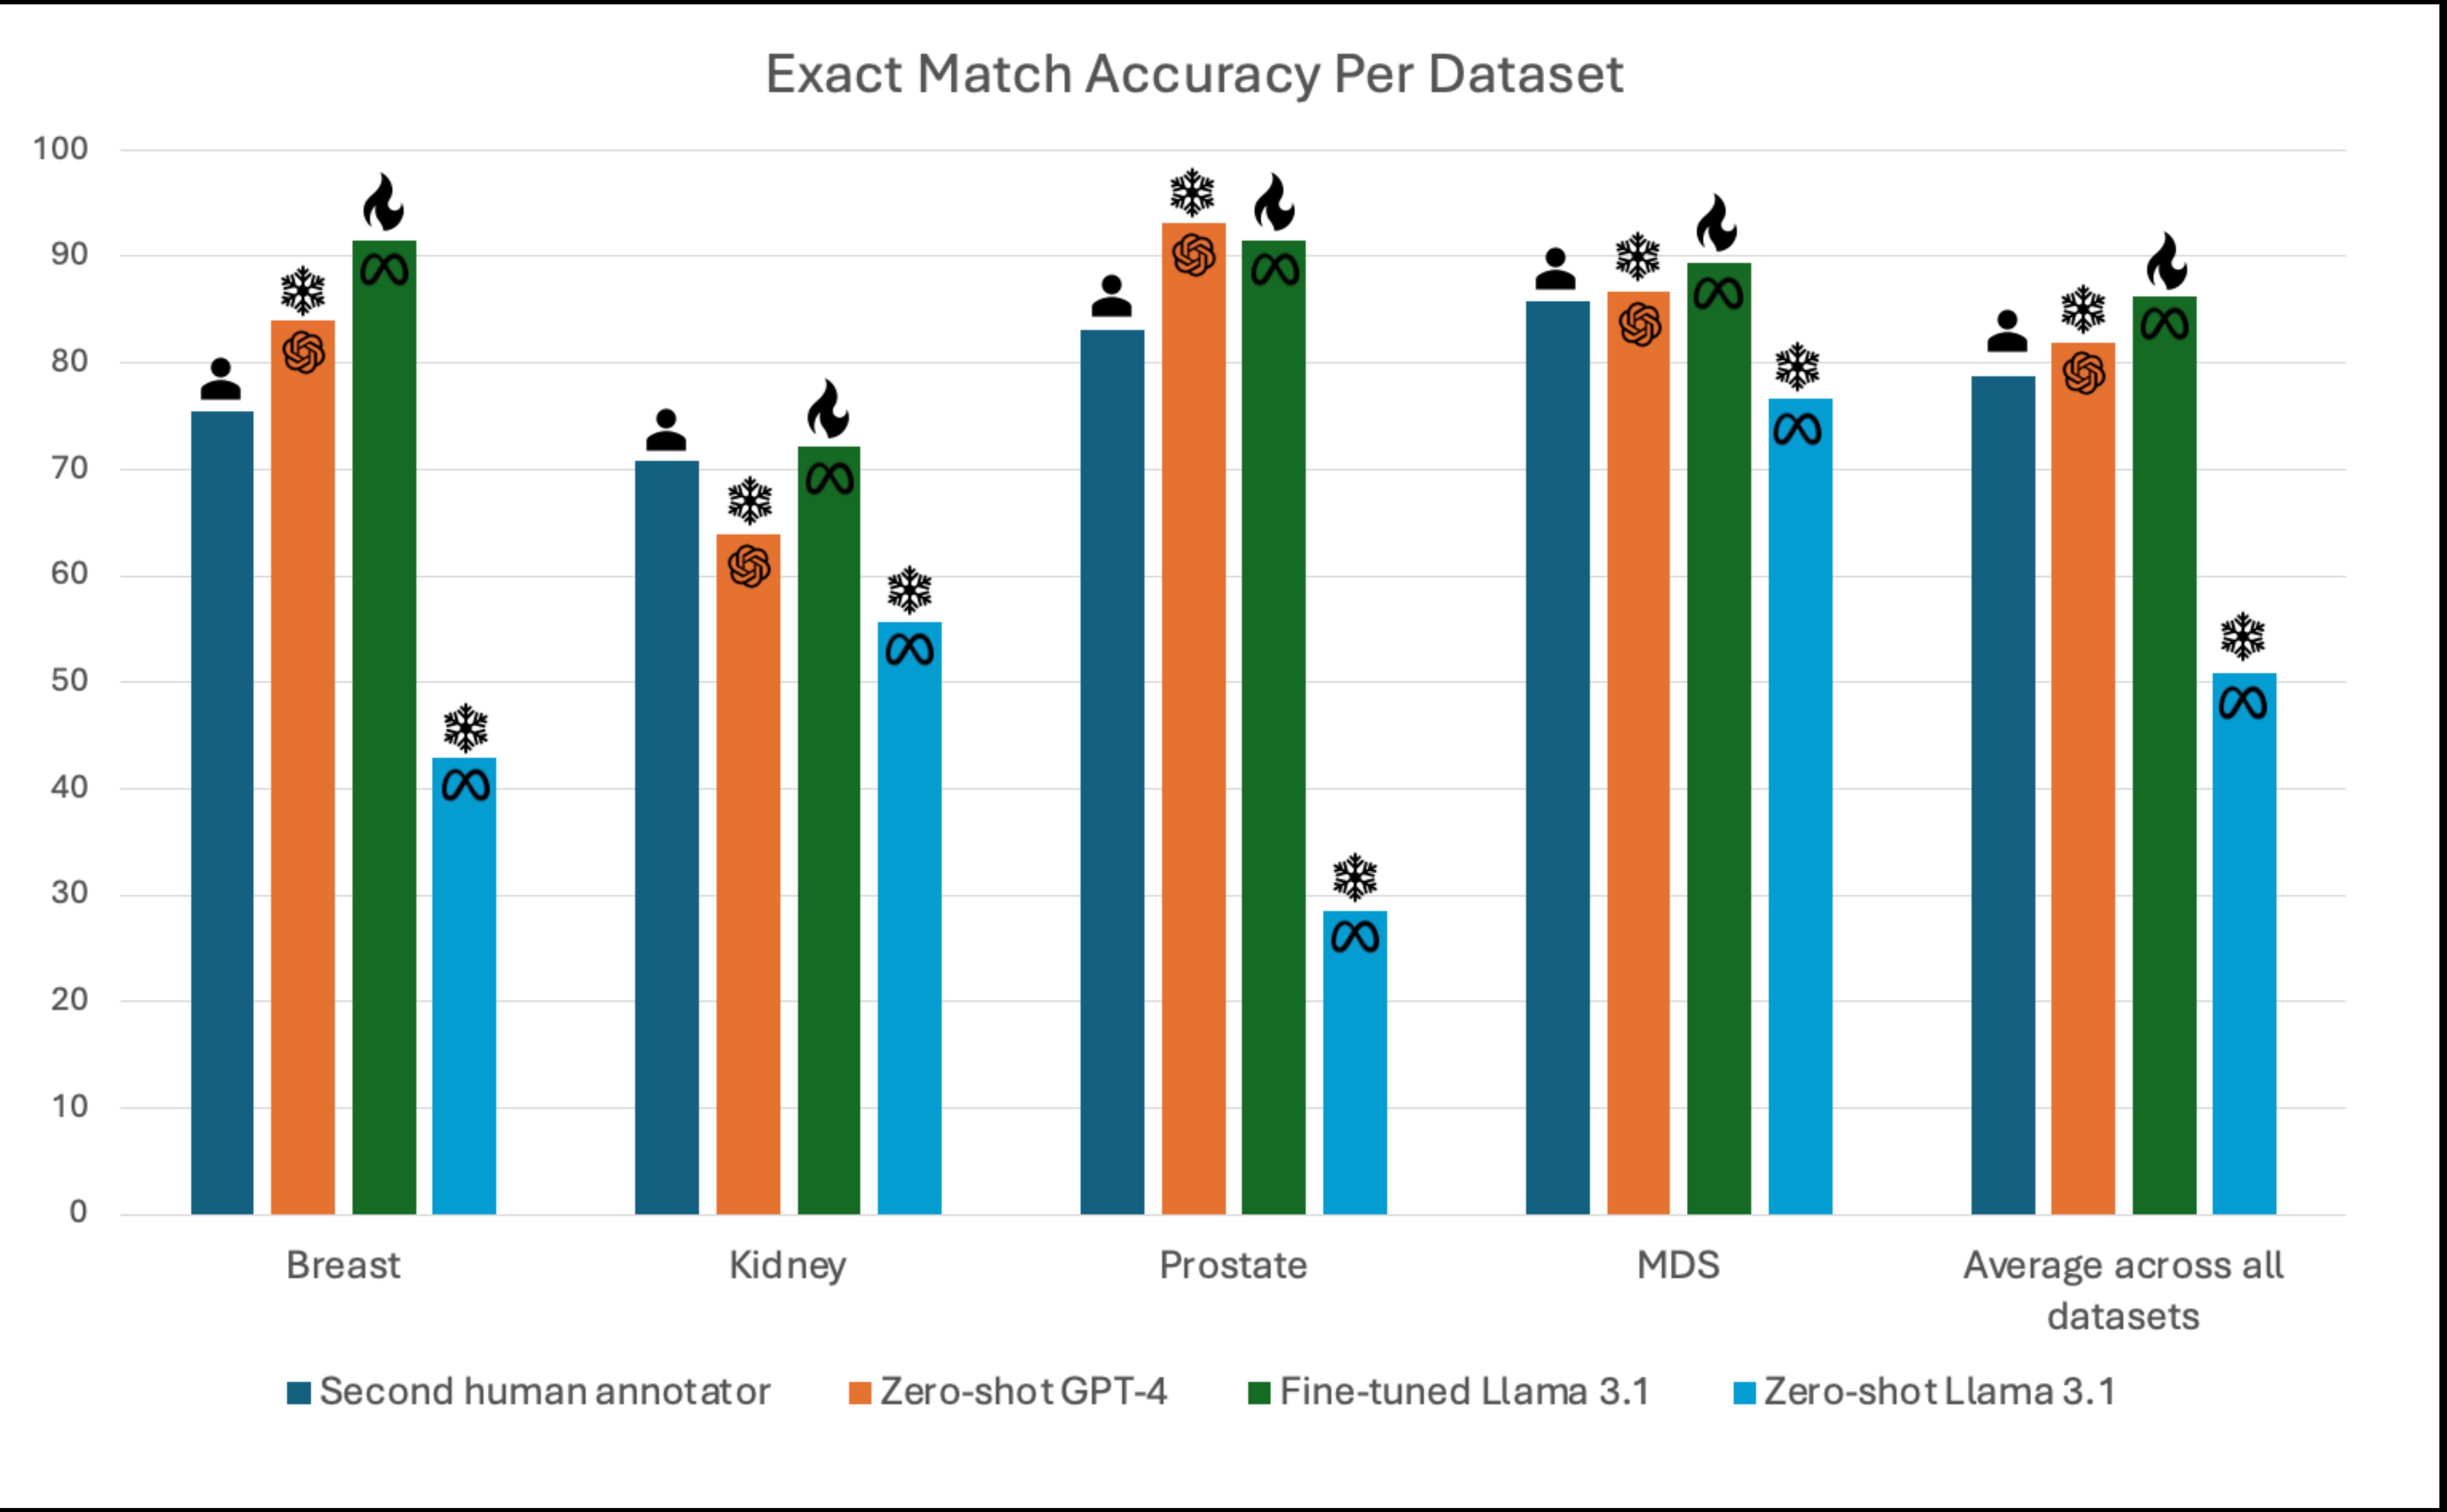

YalaLab / strata

Where are we now? From bits to MB

Patient **Data**

Future Outcome

→ 3 year cancer

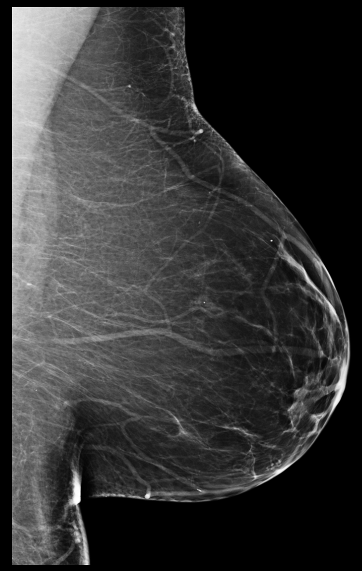

→ No cancer

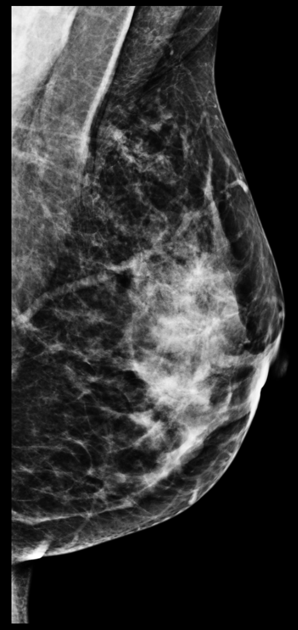

→ 5 year cancer

MB of data per patient

Journal of Clinical Oncology®
An American Society of Clinical Oncology Journal

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

Adam Yala, MEng[1,2]; Peter G. Mikhael, BS[1,2]; Fredrik Strand, MD, PhD[3,4]; Gigin Lin, MD, PhD[5]; Siddharth Satuluru, BS[6];

SCIENCE TRANSLATIONAL MEDICINE

Toward robust mammography-based models for breast cancer risk

Adam Yala[1,2]*, Peter G. Mikhael[1,2], Fredrik Strand[3,4], Gigin Lin[5], Kevin Smith[6,7], Yung-Liang Leslie Lamb[8], Kevin Hughes[9], Constance Lehman[8†], Regina Barzilay[1,2†]

Radiology

A Deep Learning Mammography-based Model for Improved Breast Cancer Risk Prediction

Adam Yala, MEng • Constance Lehman, MD, PhD • Tal Schuster, MS • Tally Portnoi, BS • Regina Barzilay, PhD

Journal of Clinical Oncology®
An American Society of Clinical Oncology Journal

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

Peter G. Mikhael, BSc[1,2]; Jeremy Wohlwend, ME[1,2]; Adam Yala, PhD[1,2]; Ludvig Karstens, MSc[1,2]; Justin Xiang, ME[1,2]; Angelo K. Takigami, MD[3,4]; Patrick P. Bourgouin, MD[3,4]; PuiYee Chan, PhD[5]; Sofiane Mrah, MSc[4]; Wael Amayri, BSc[4]; Yu-Hsiang Juan, MD[6,7]; Cheng-Ta Yang, MD[6,8]; Yung-Liang Wan, MD[6,7]; Gigin Lin, MD, PhD[6,7]; Lecia V. Sequist, MD, MPH[3,5];

Mirai: Image-based Risk model

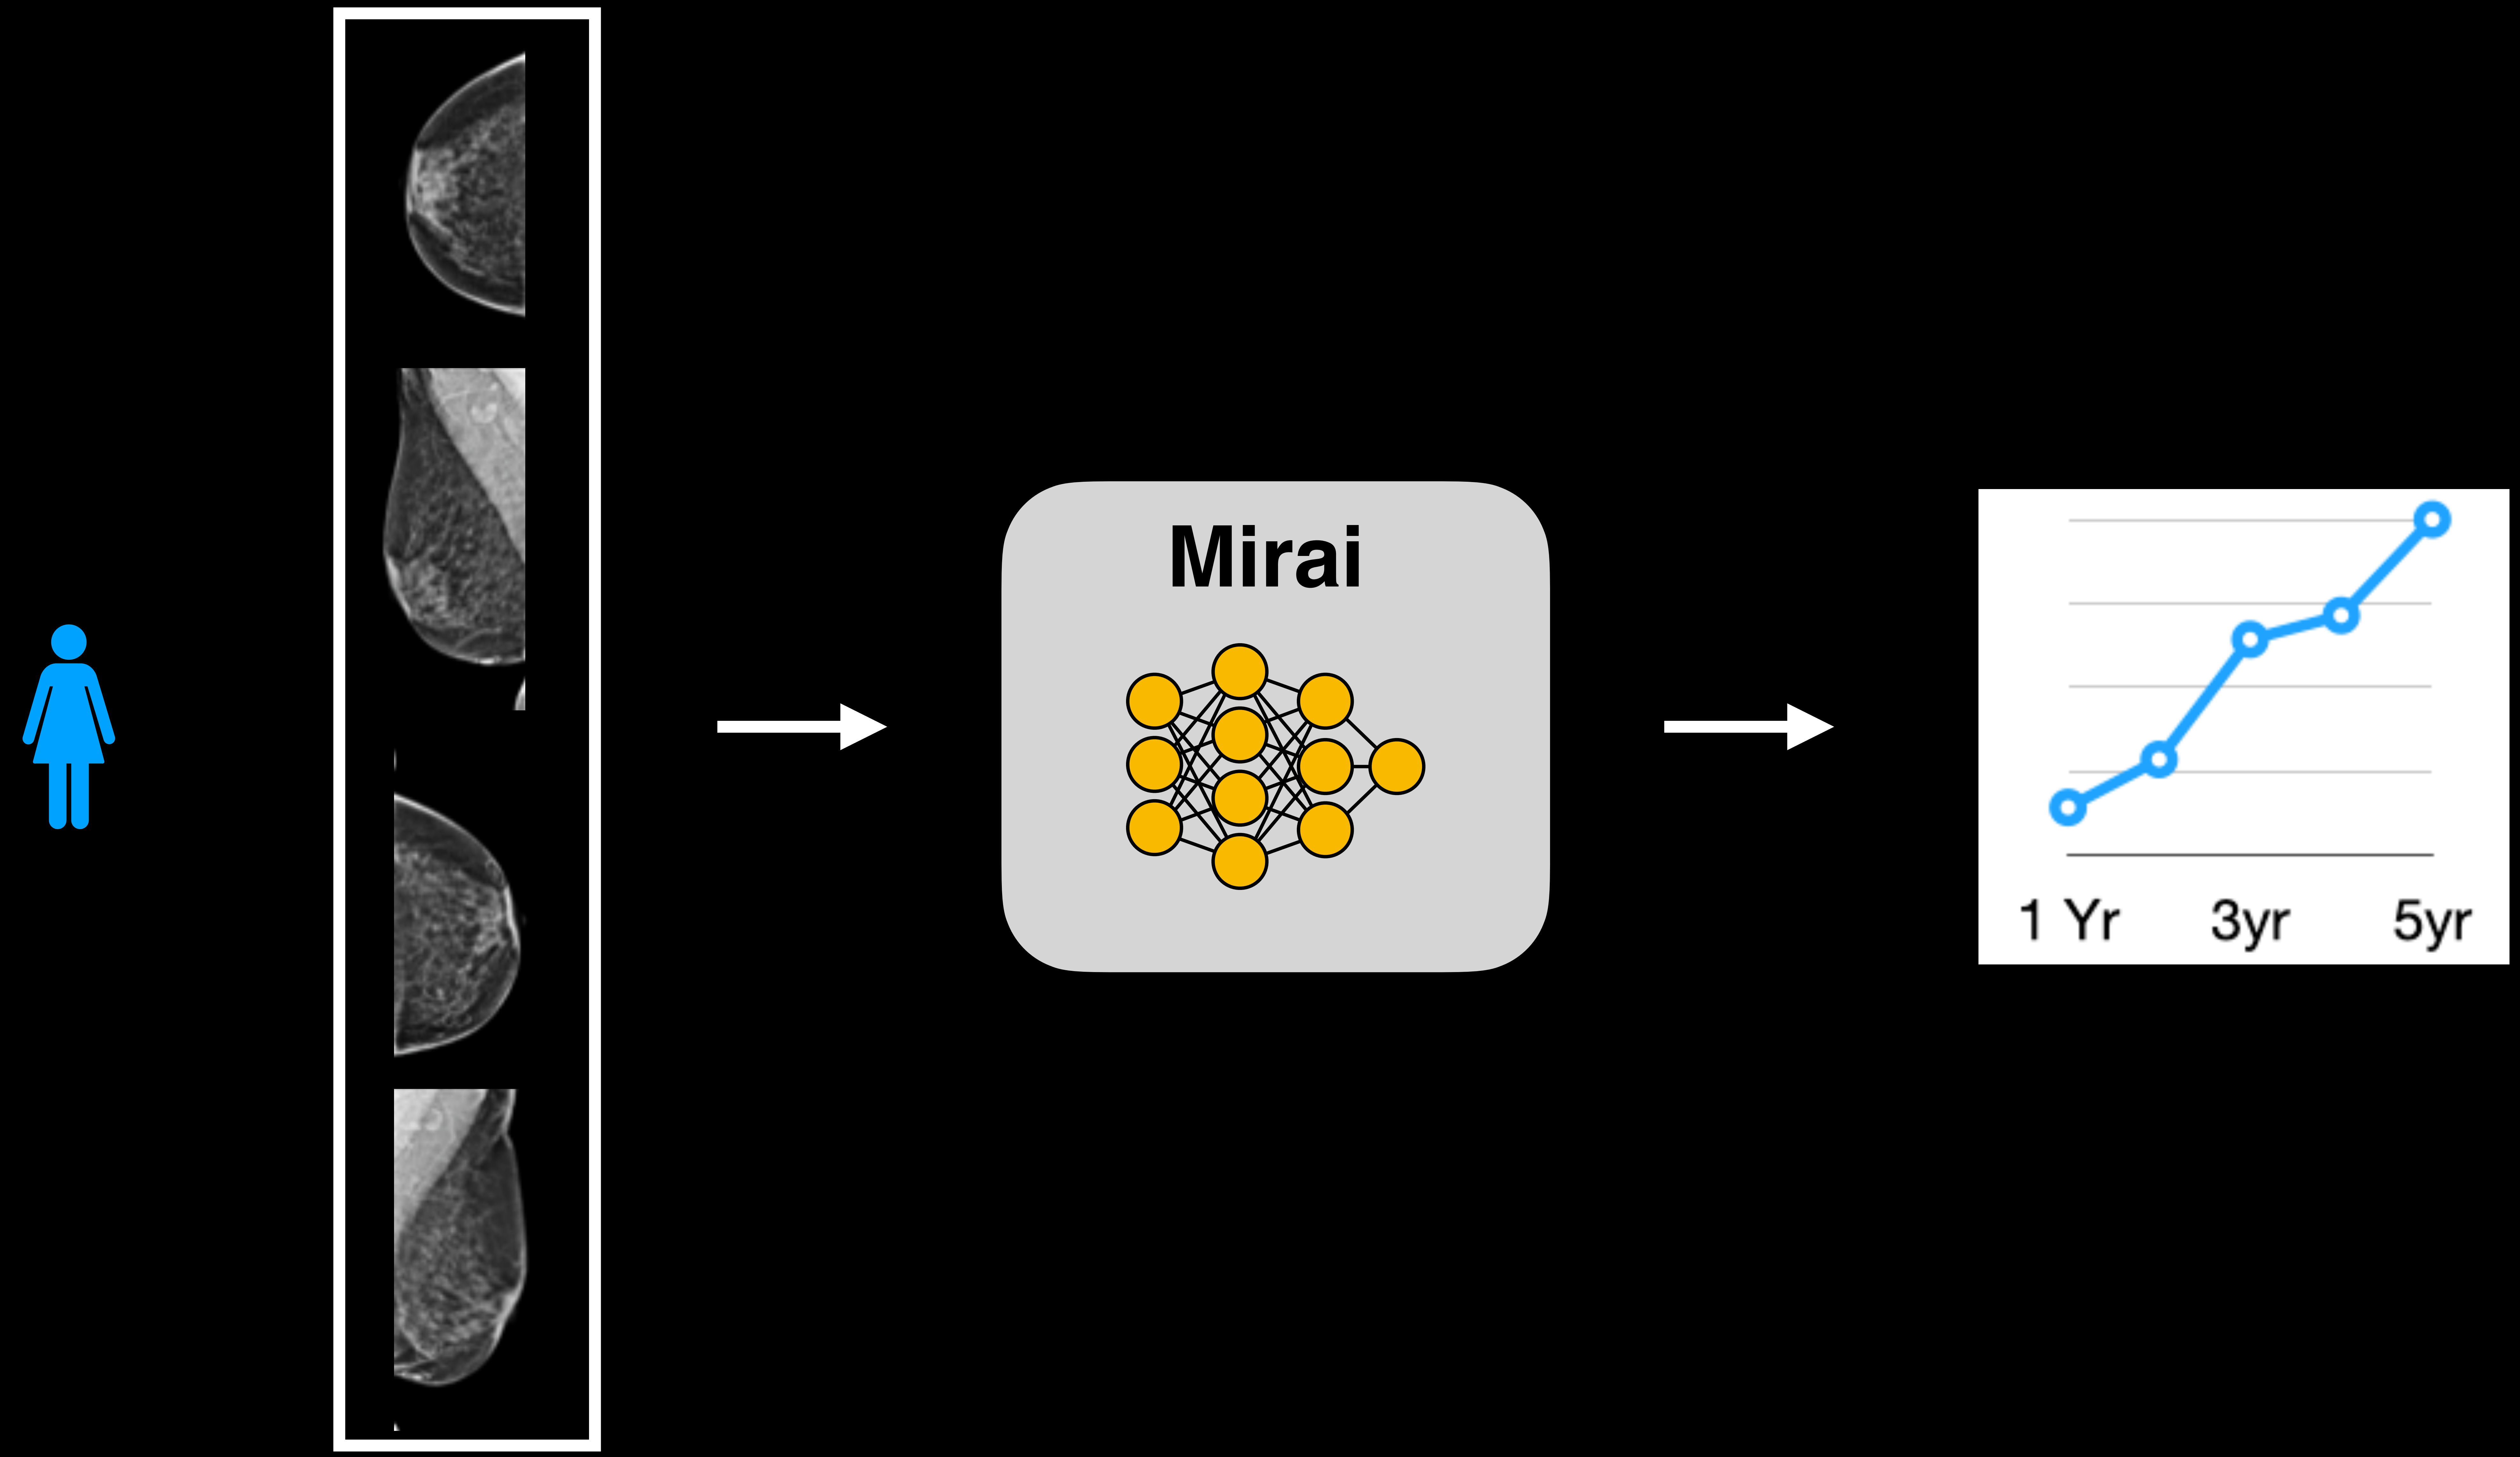

Maintains accuracy across diverse populations

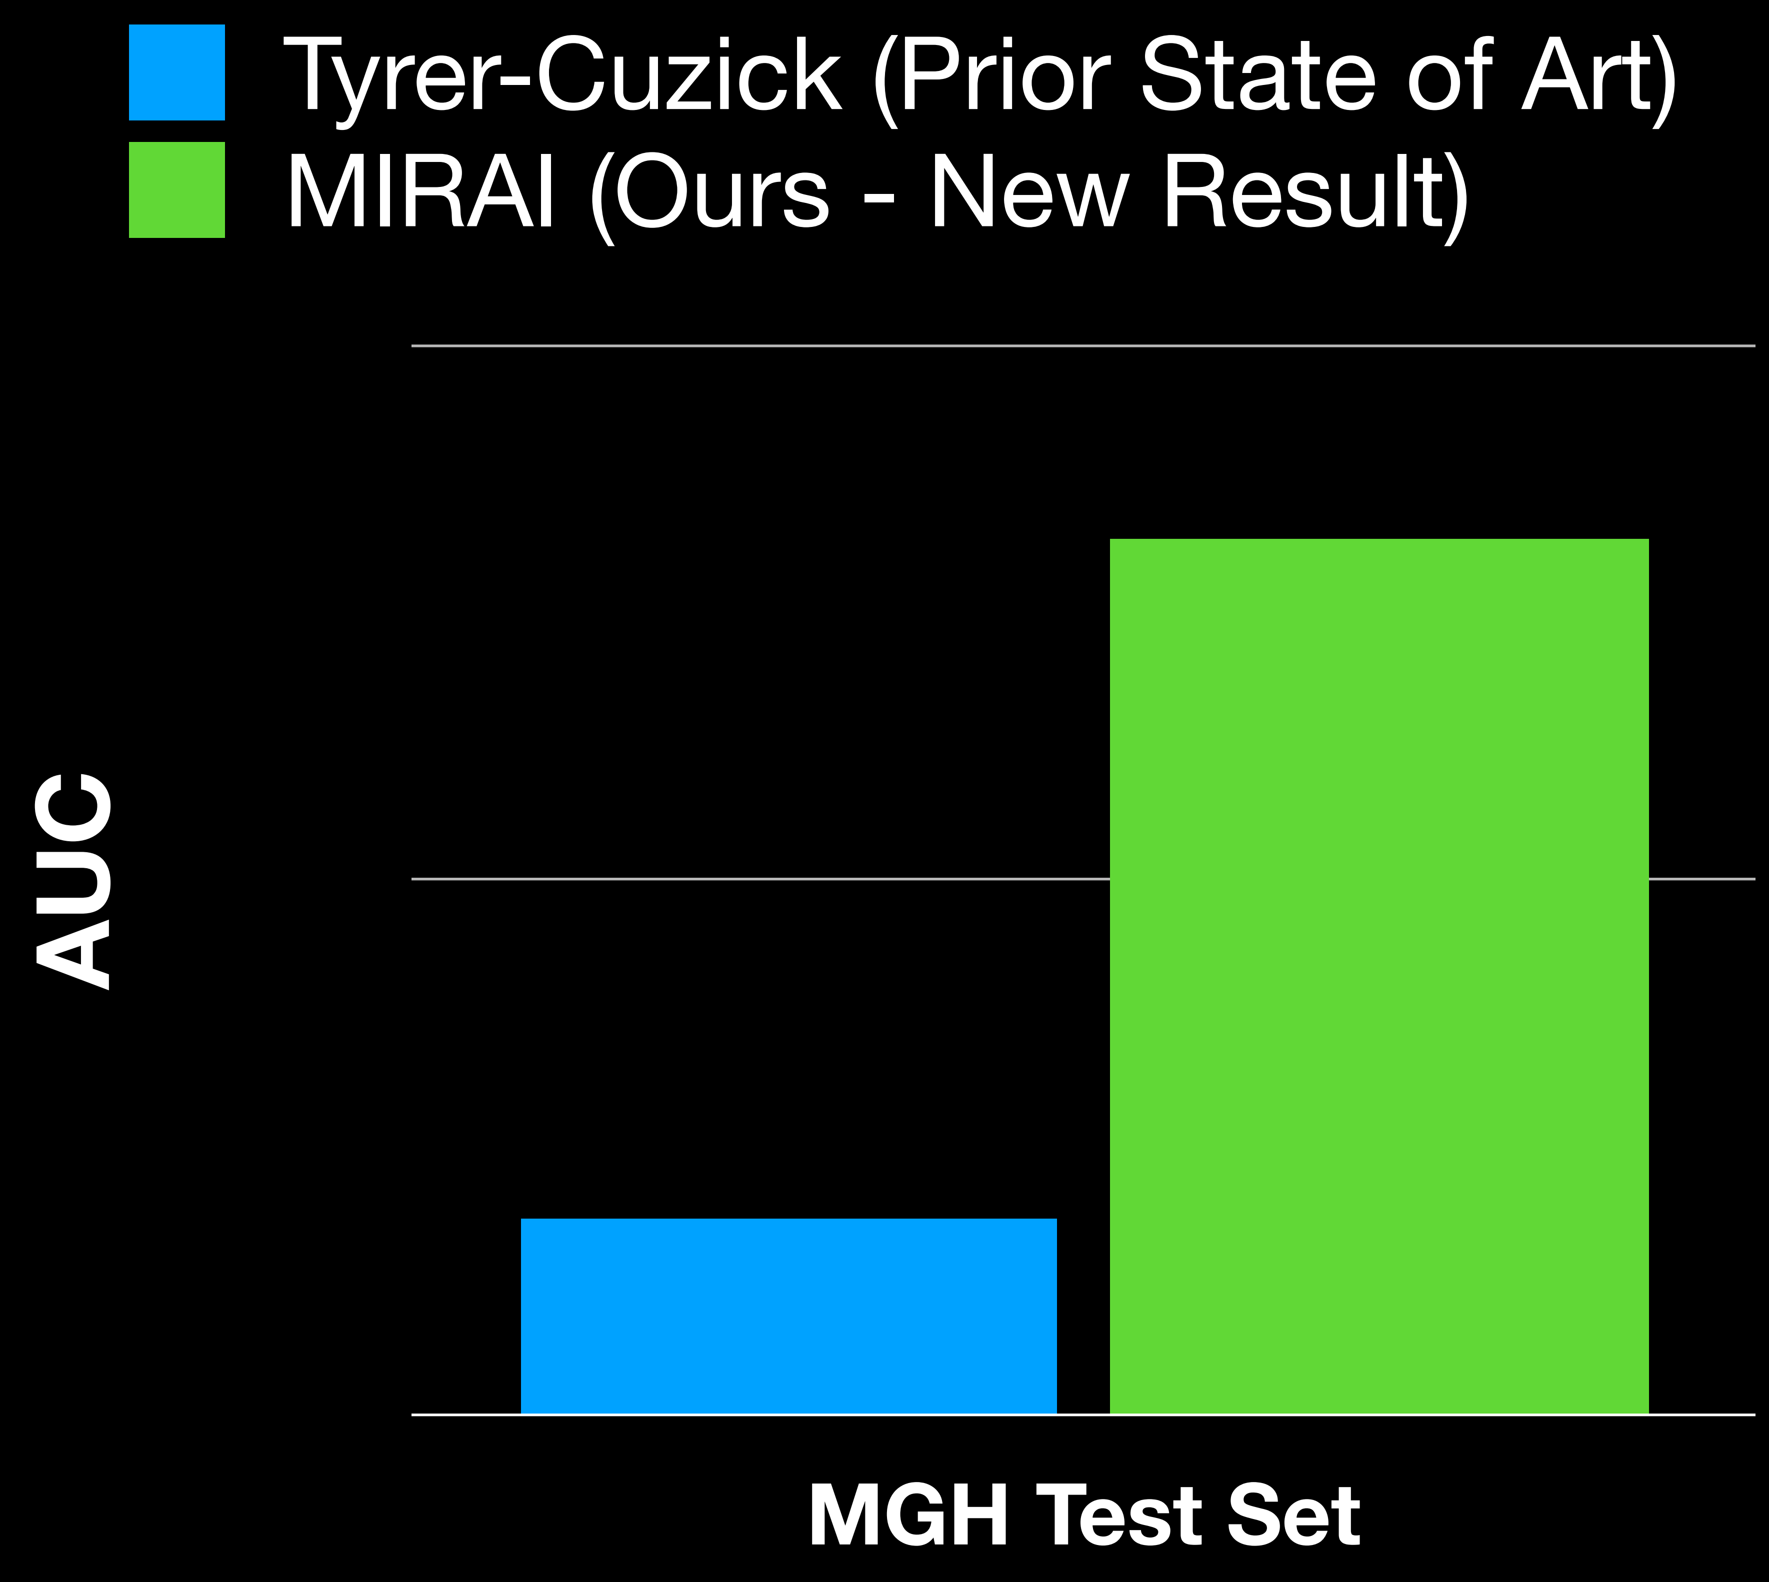

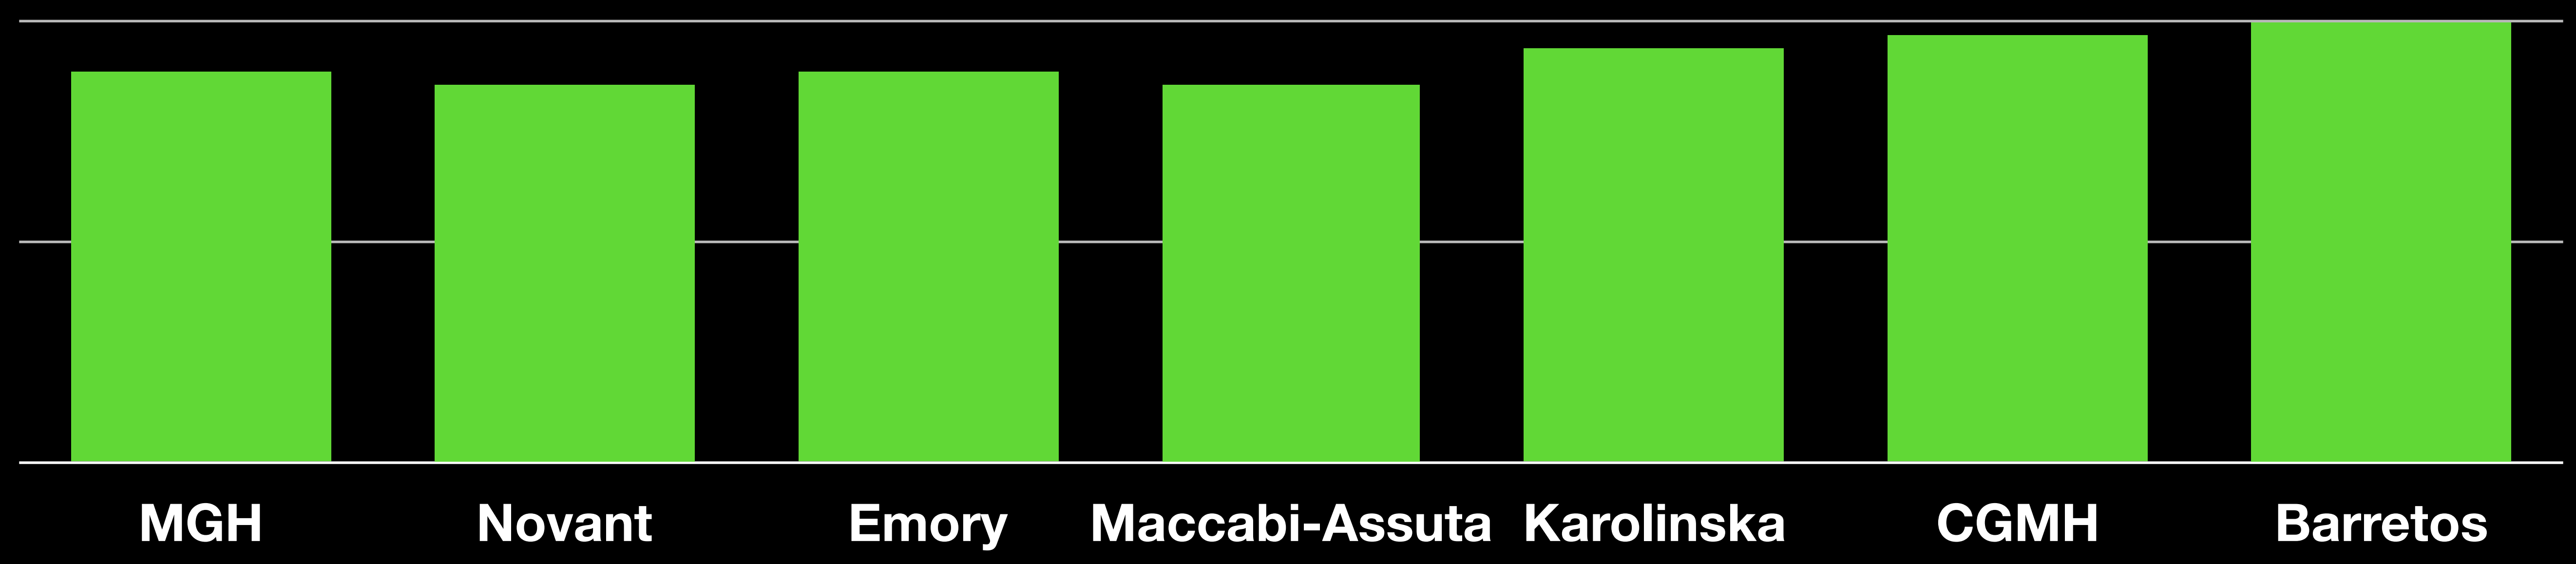

Jameel Clinic

MIRAI

1.5M+ mammograms

43 hospitals

14 countries

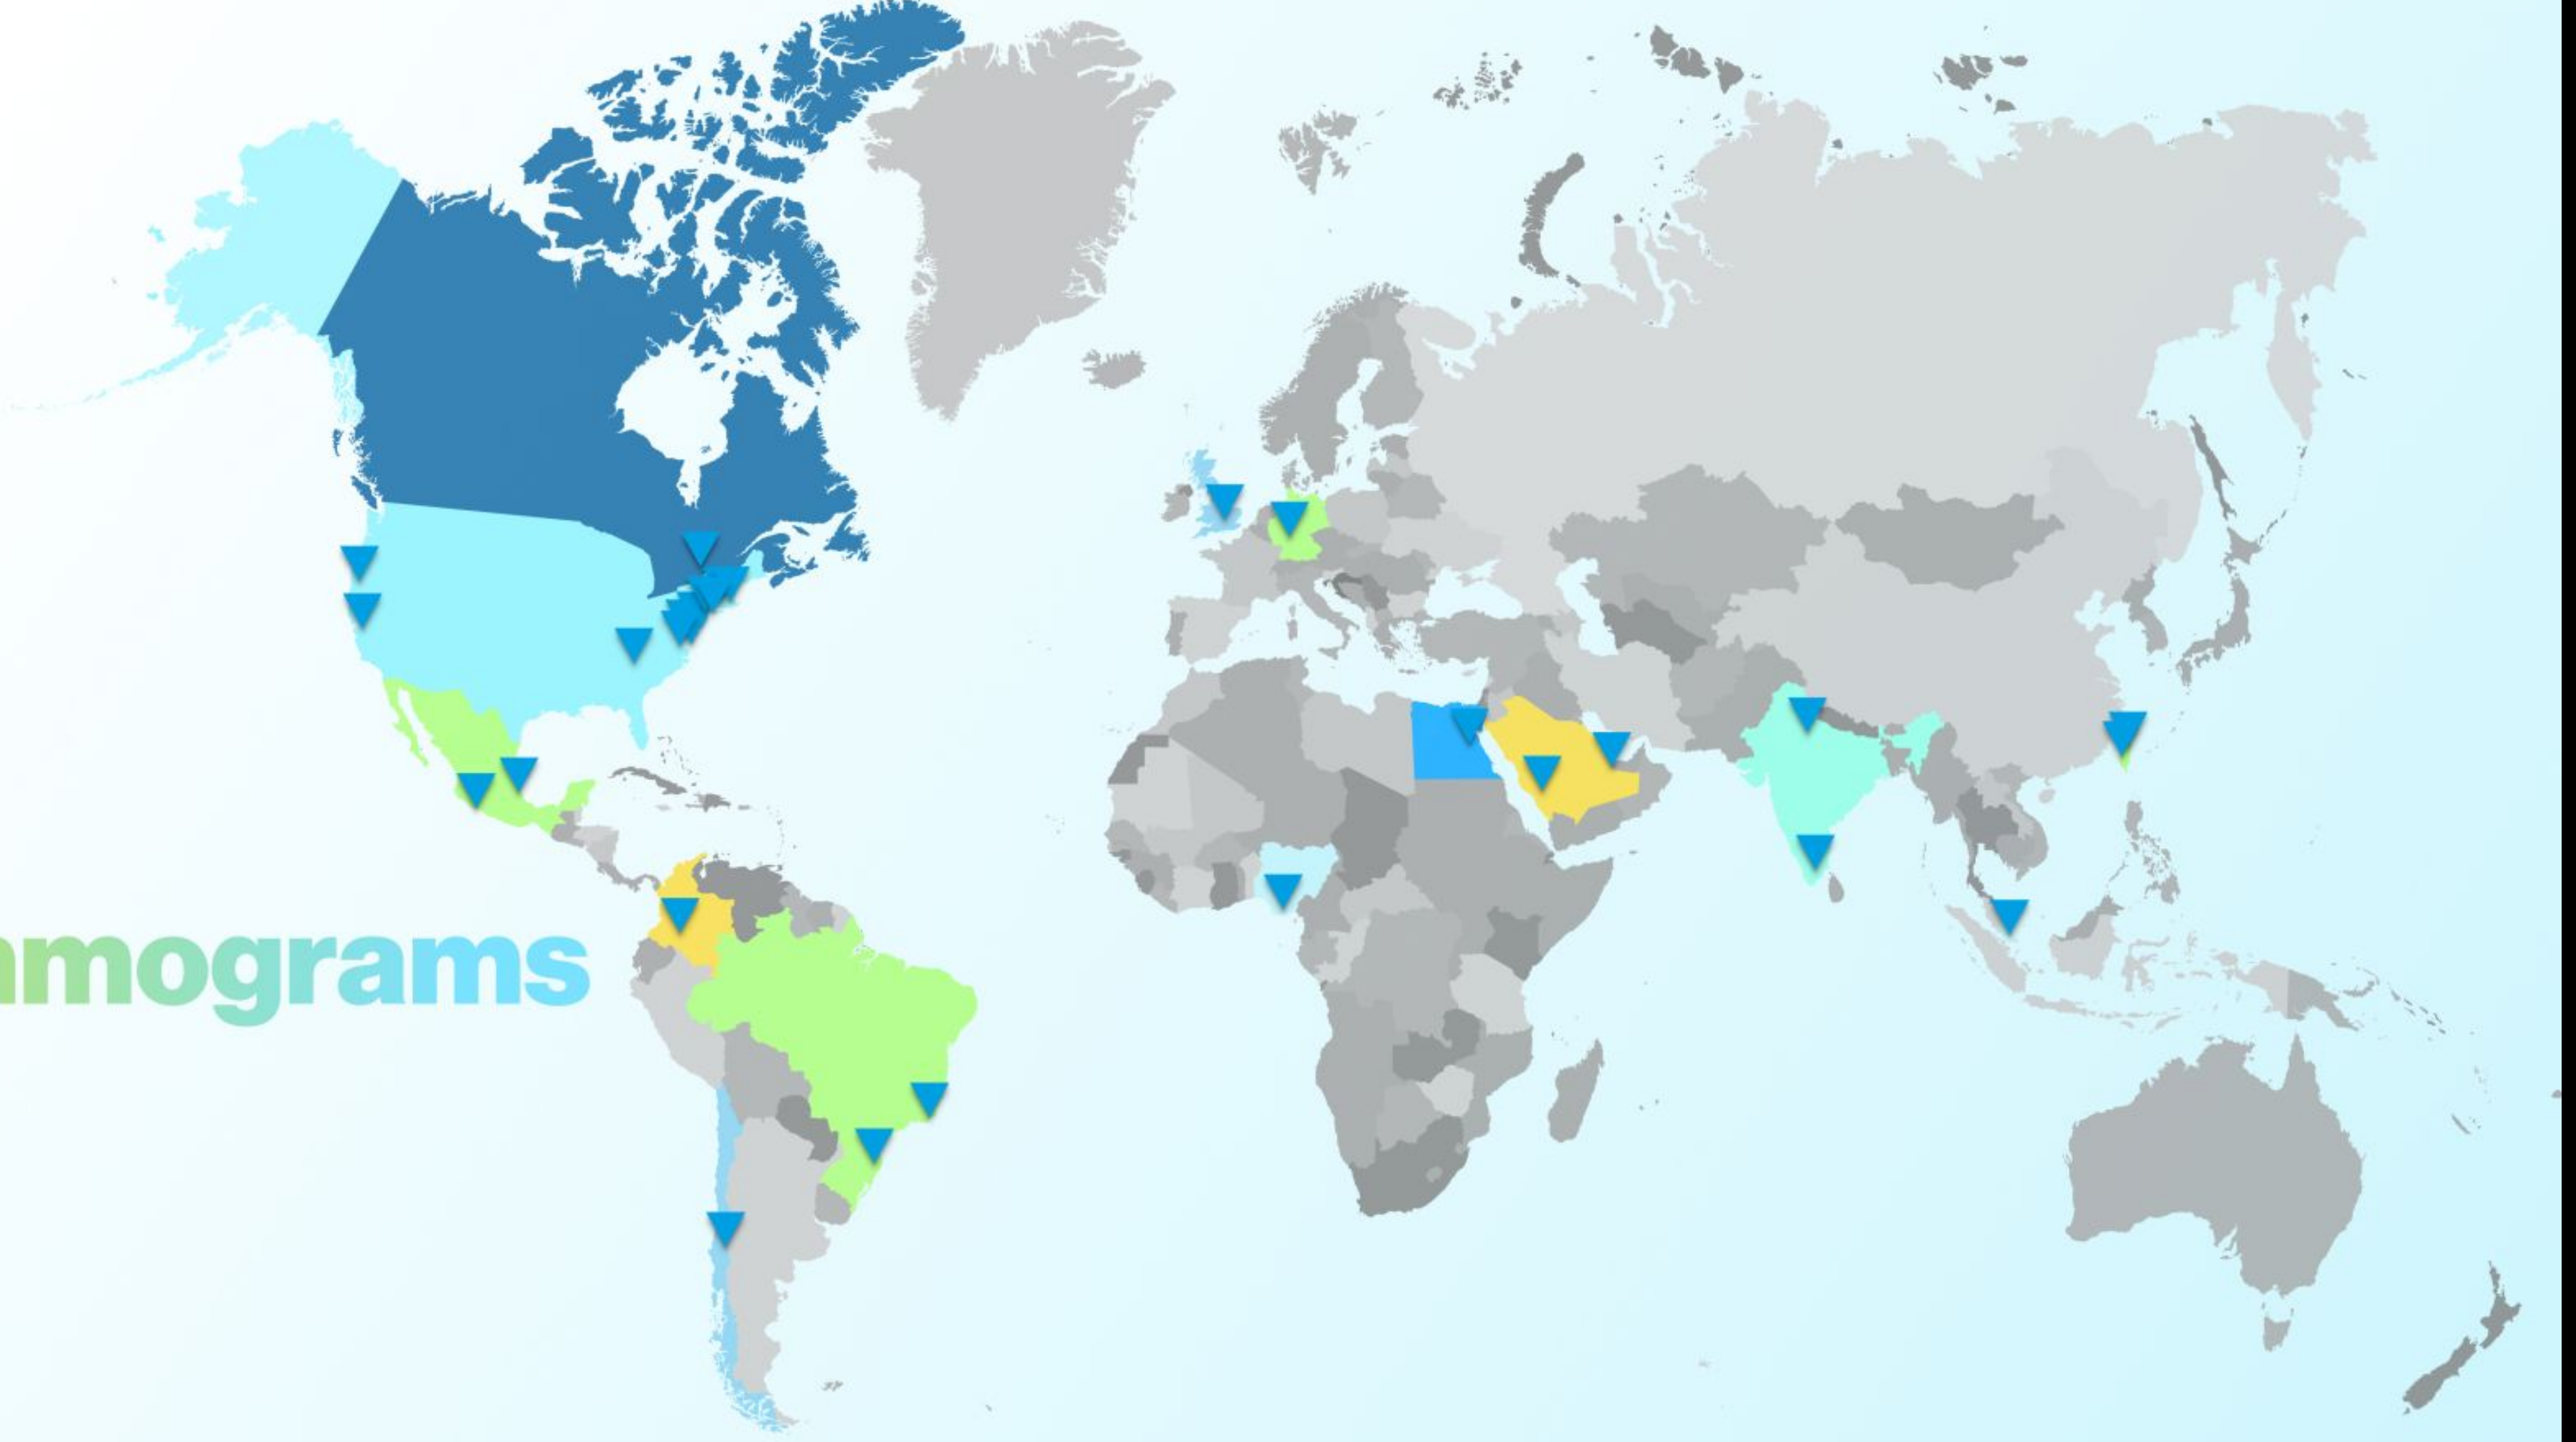

Mirai: Image-based Risk model

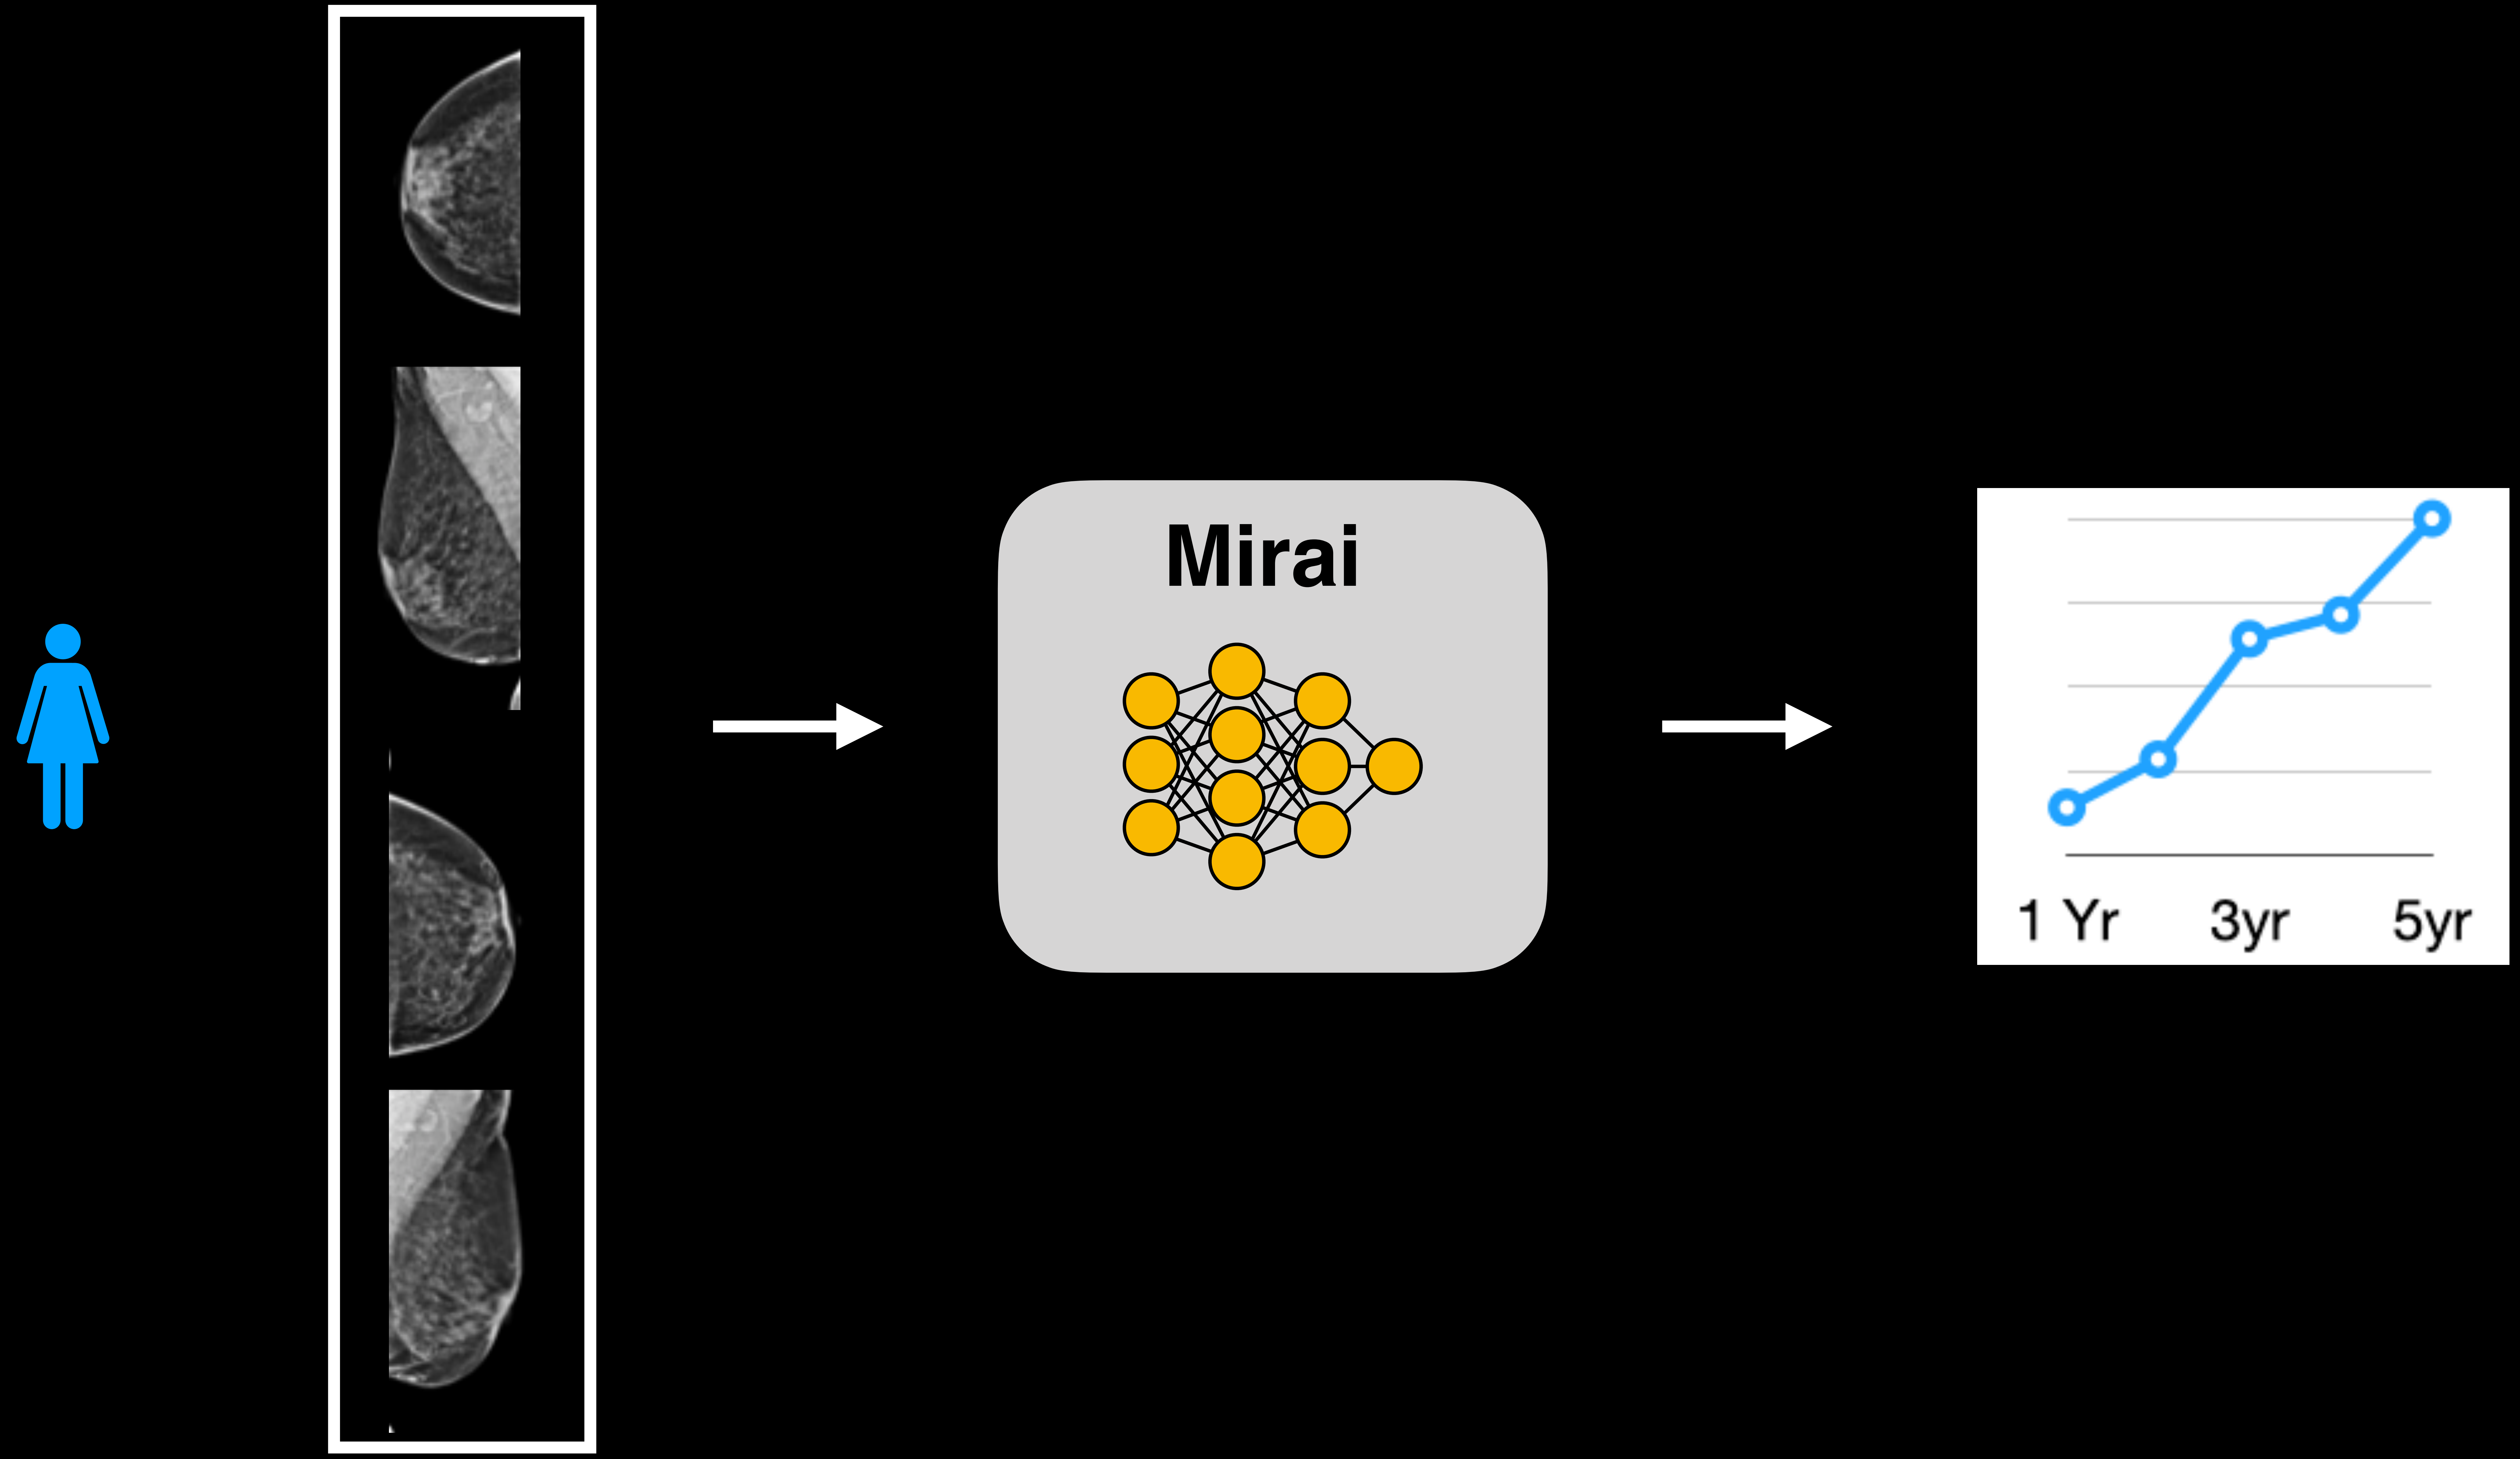

Problem 1: Device Invariance

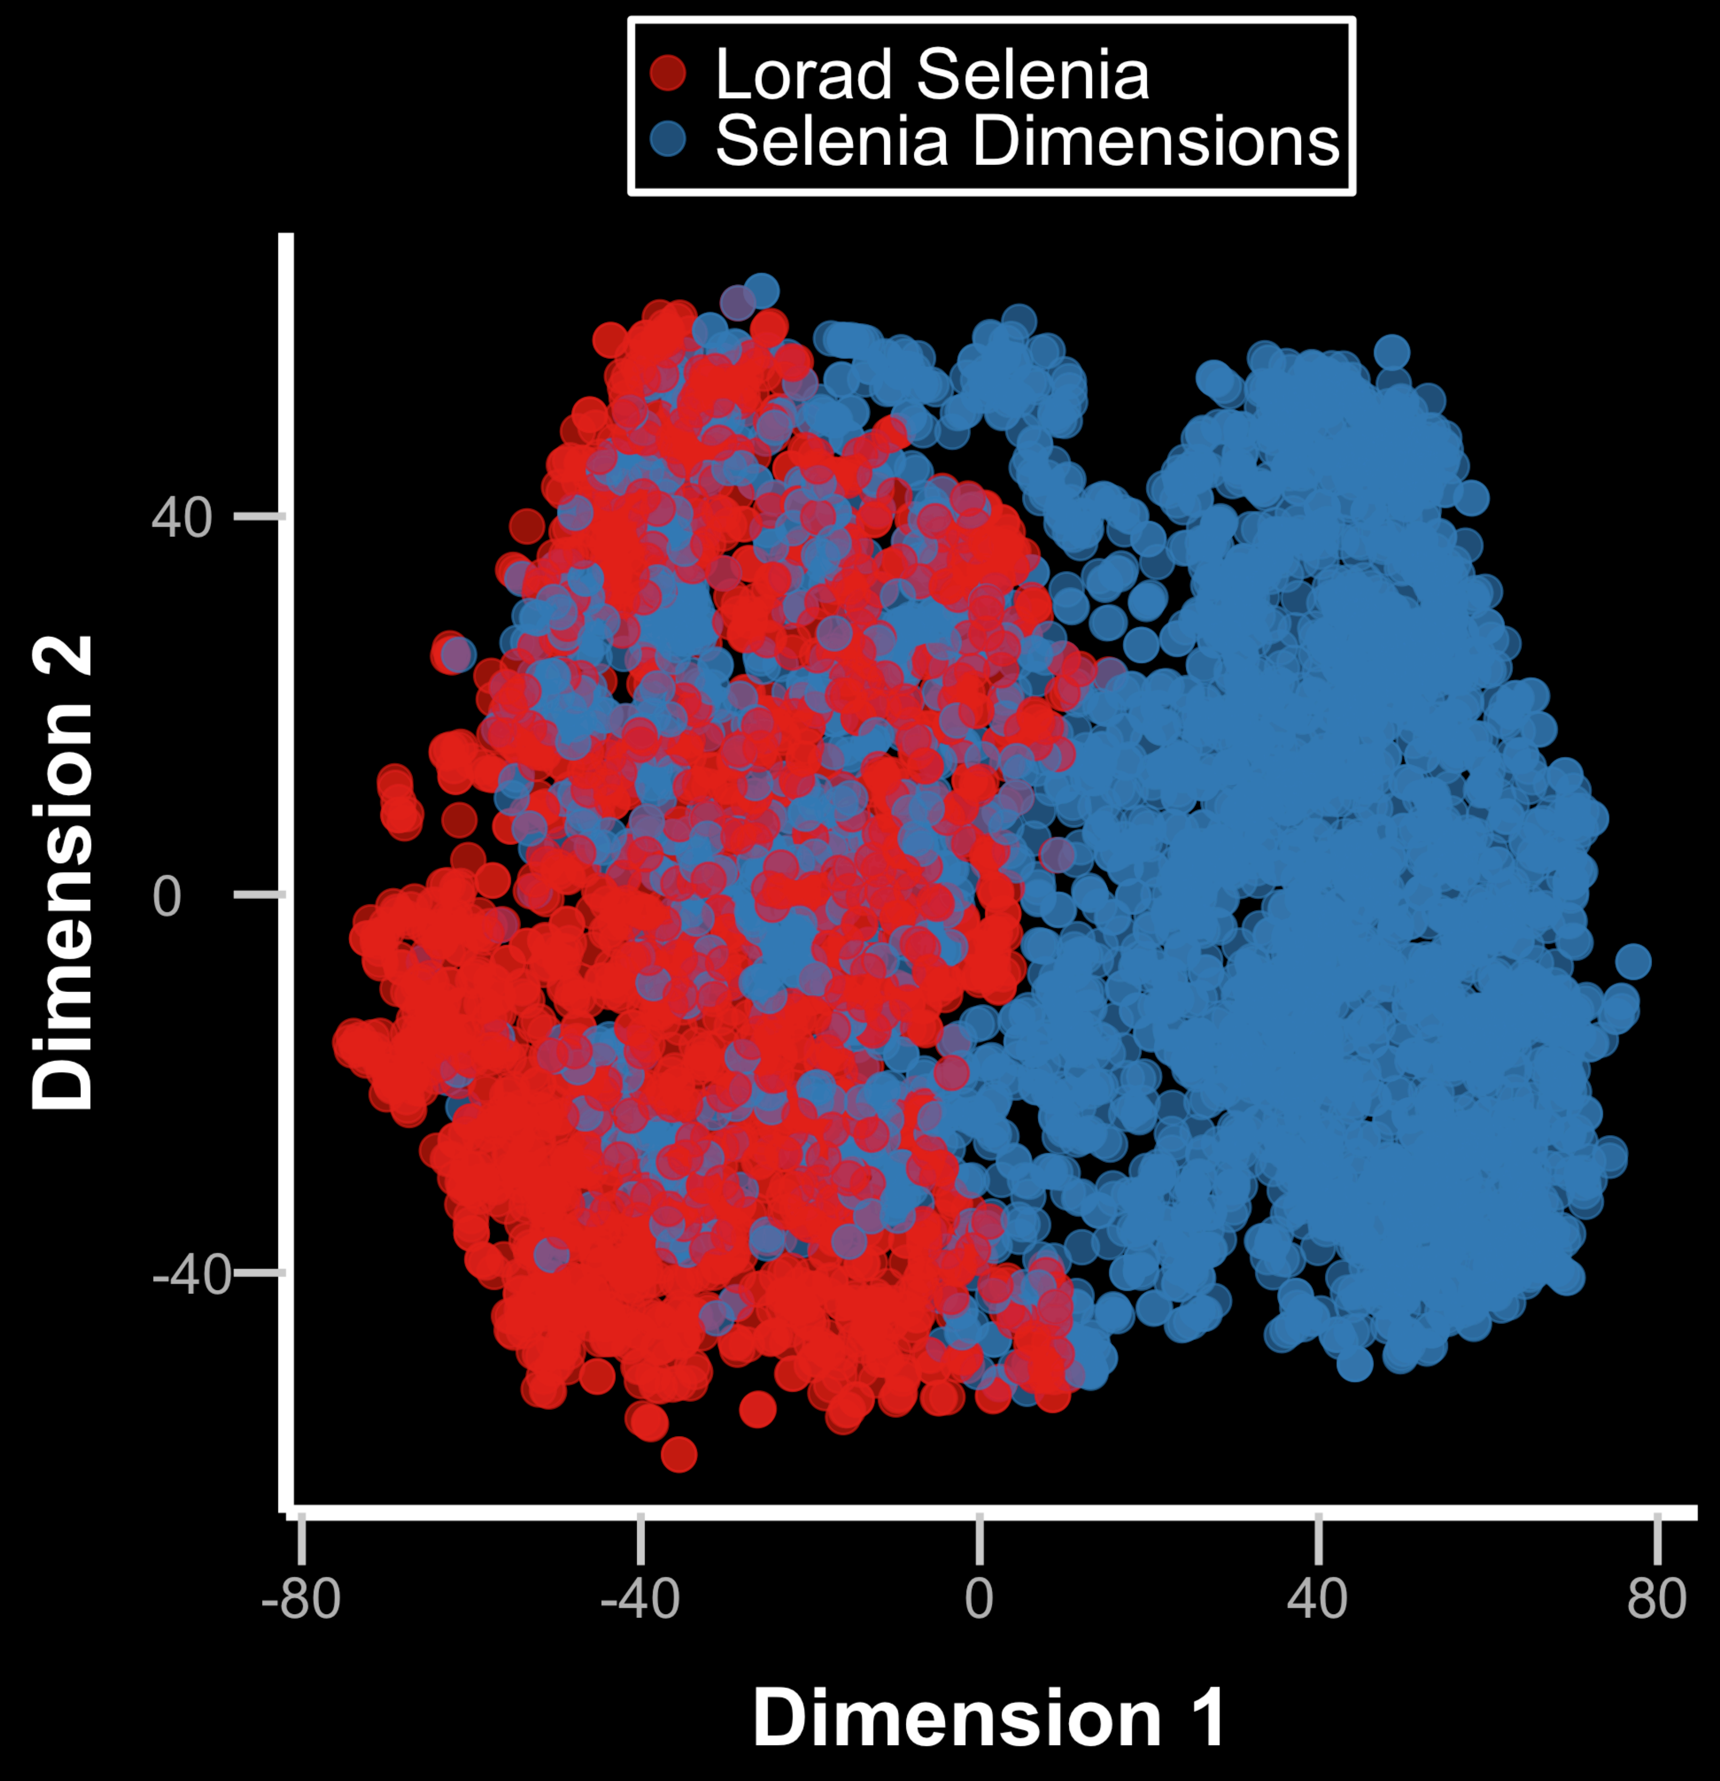

Problem 1: Device Invariance

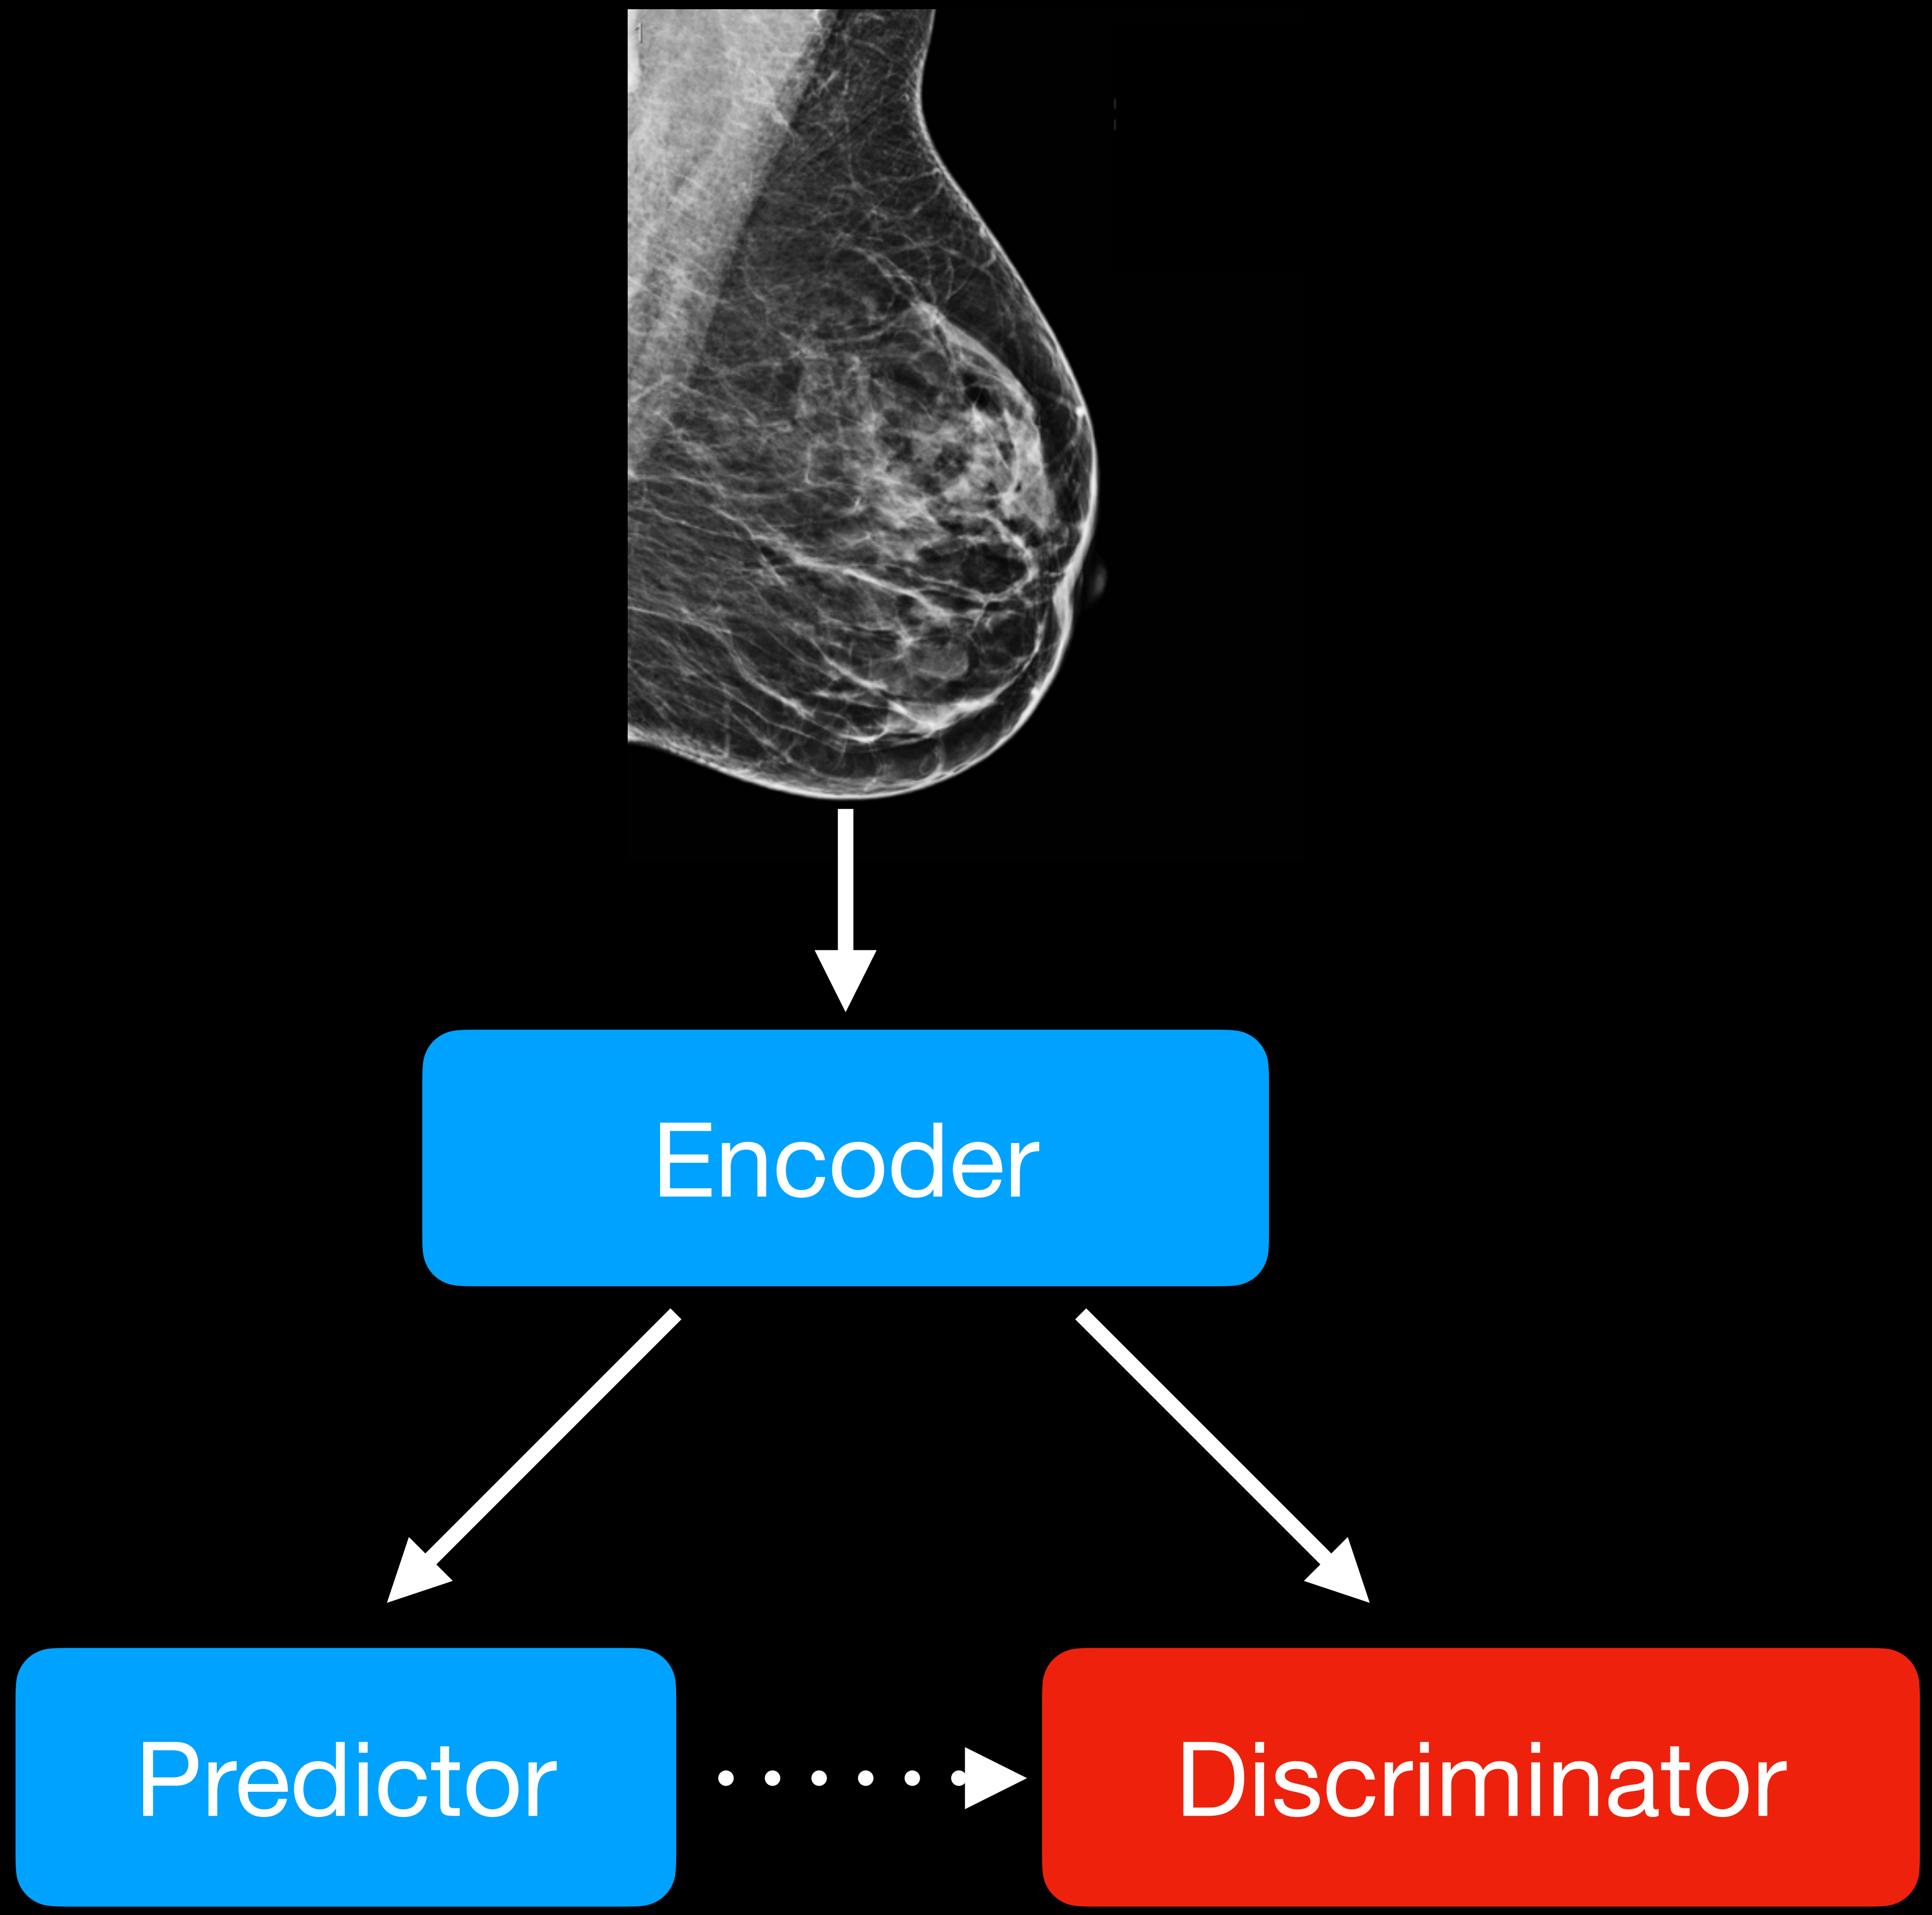

Objective:

Max accuracy **Predictor**

Min accuracy **Discriminator**

Problem 1: Device Invariance

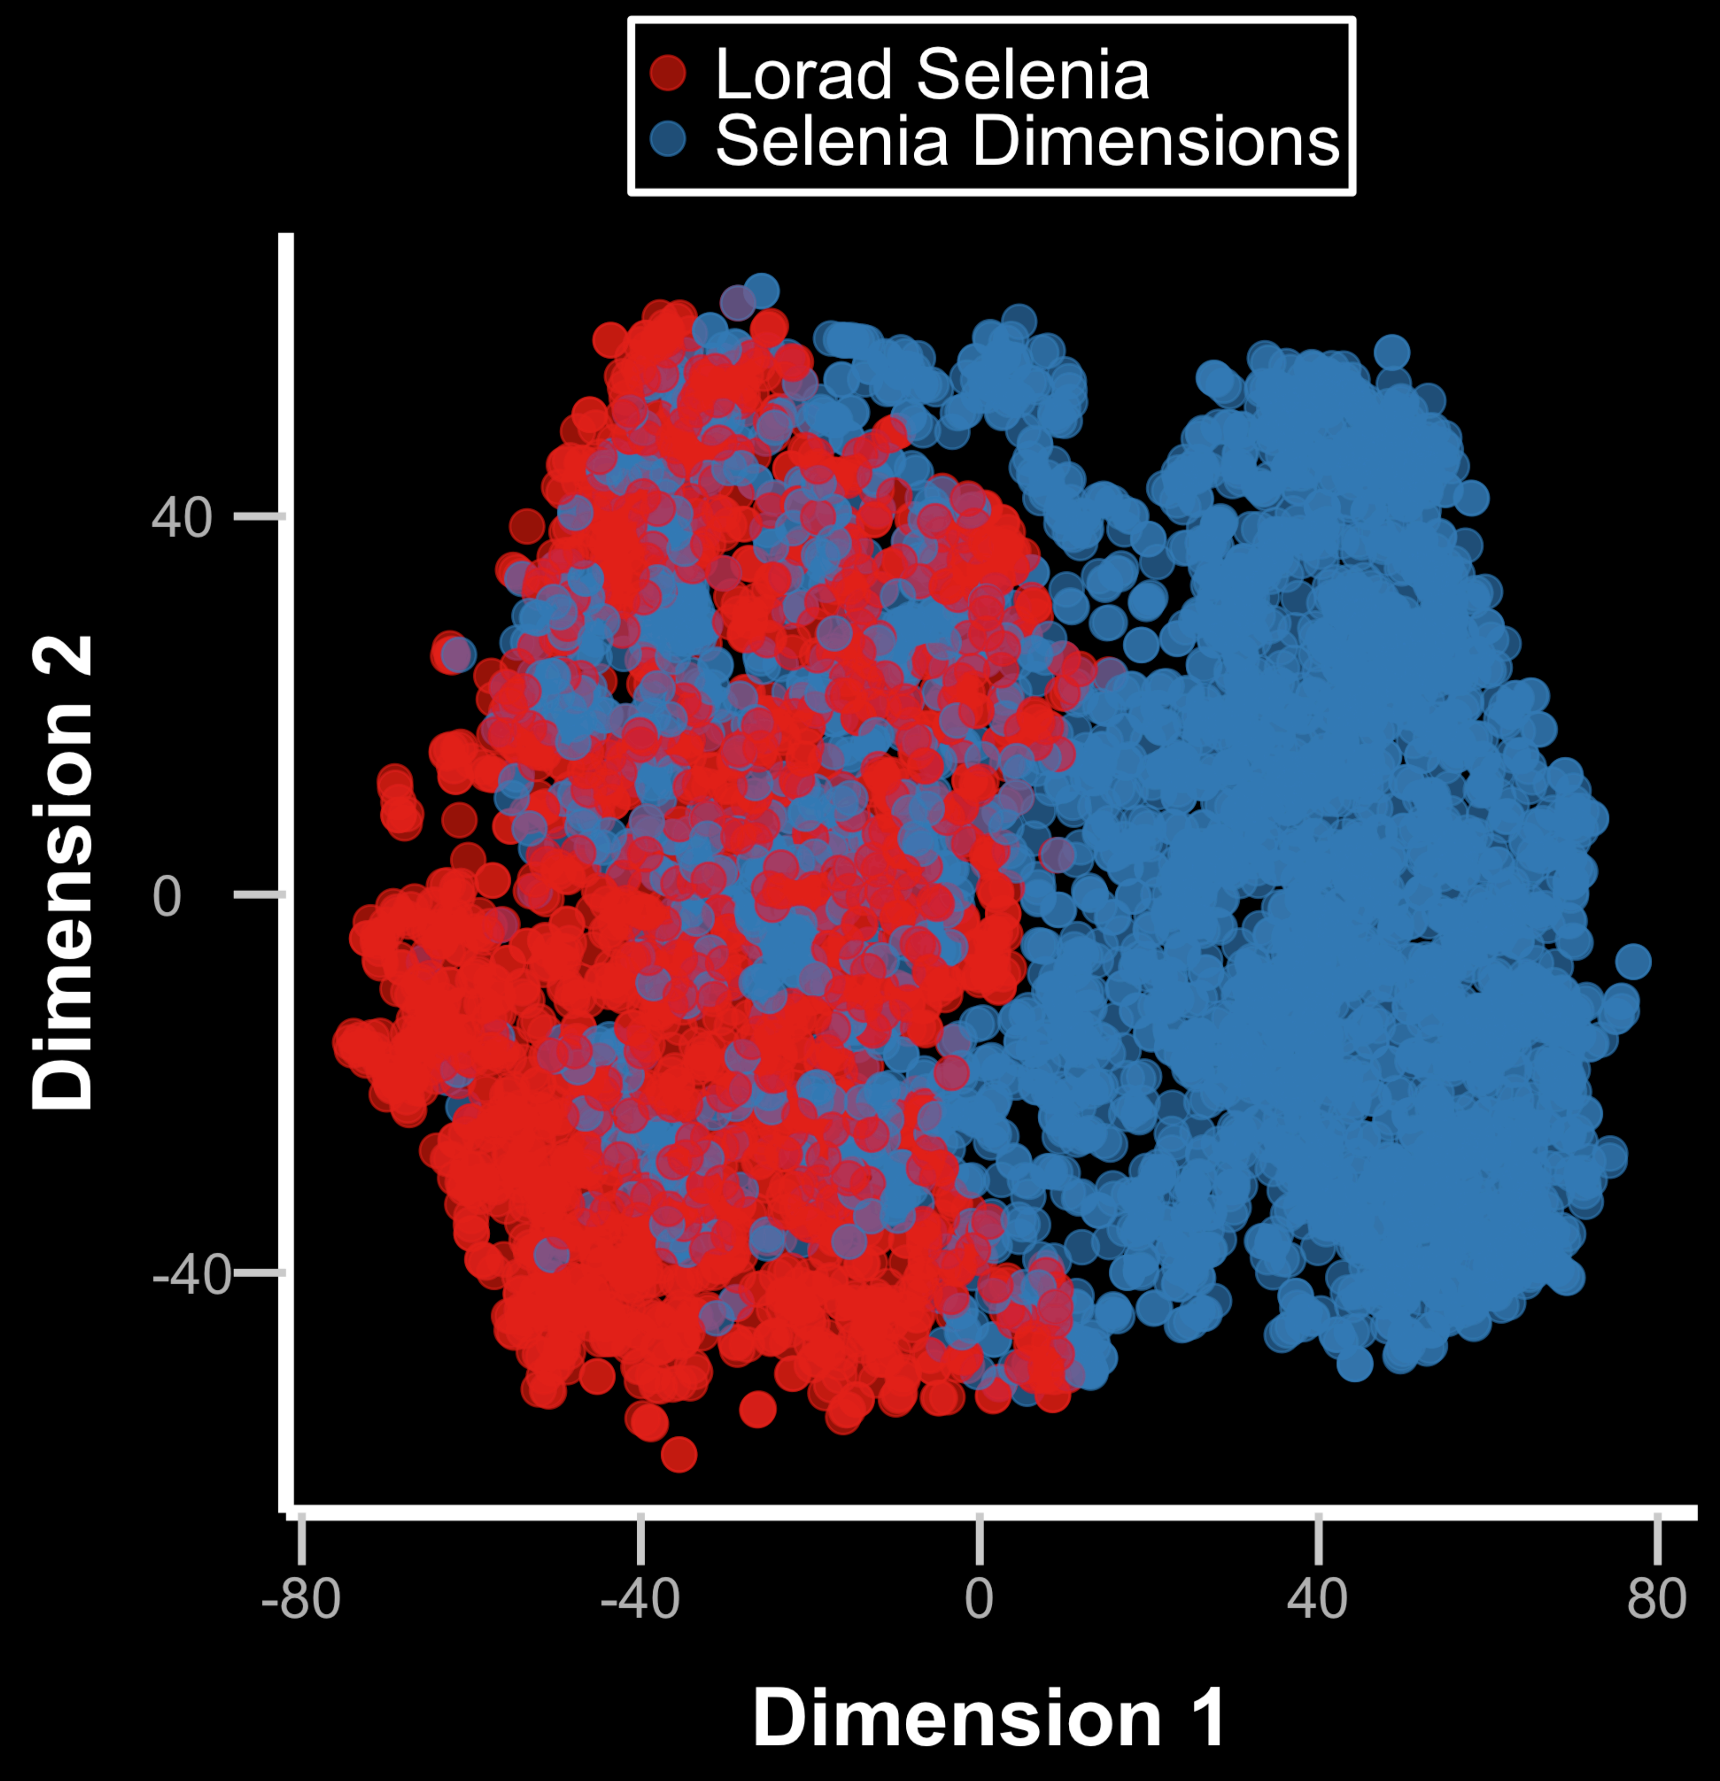

Without Adversary

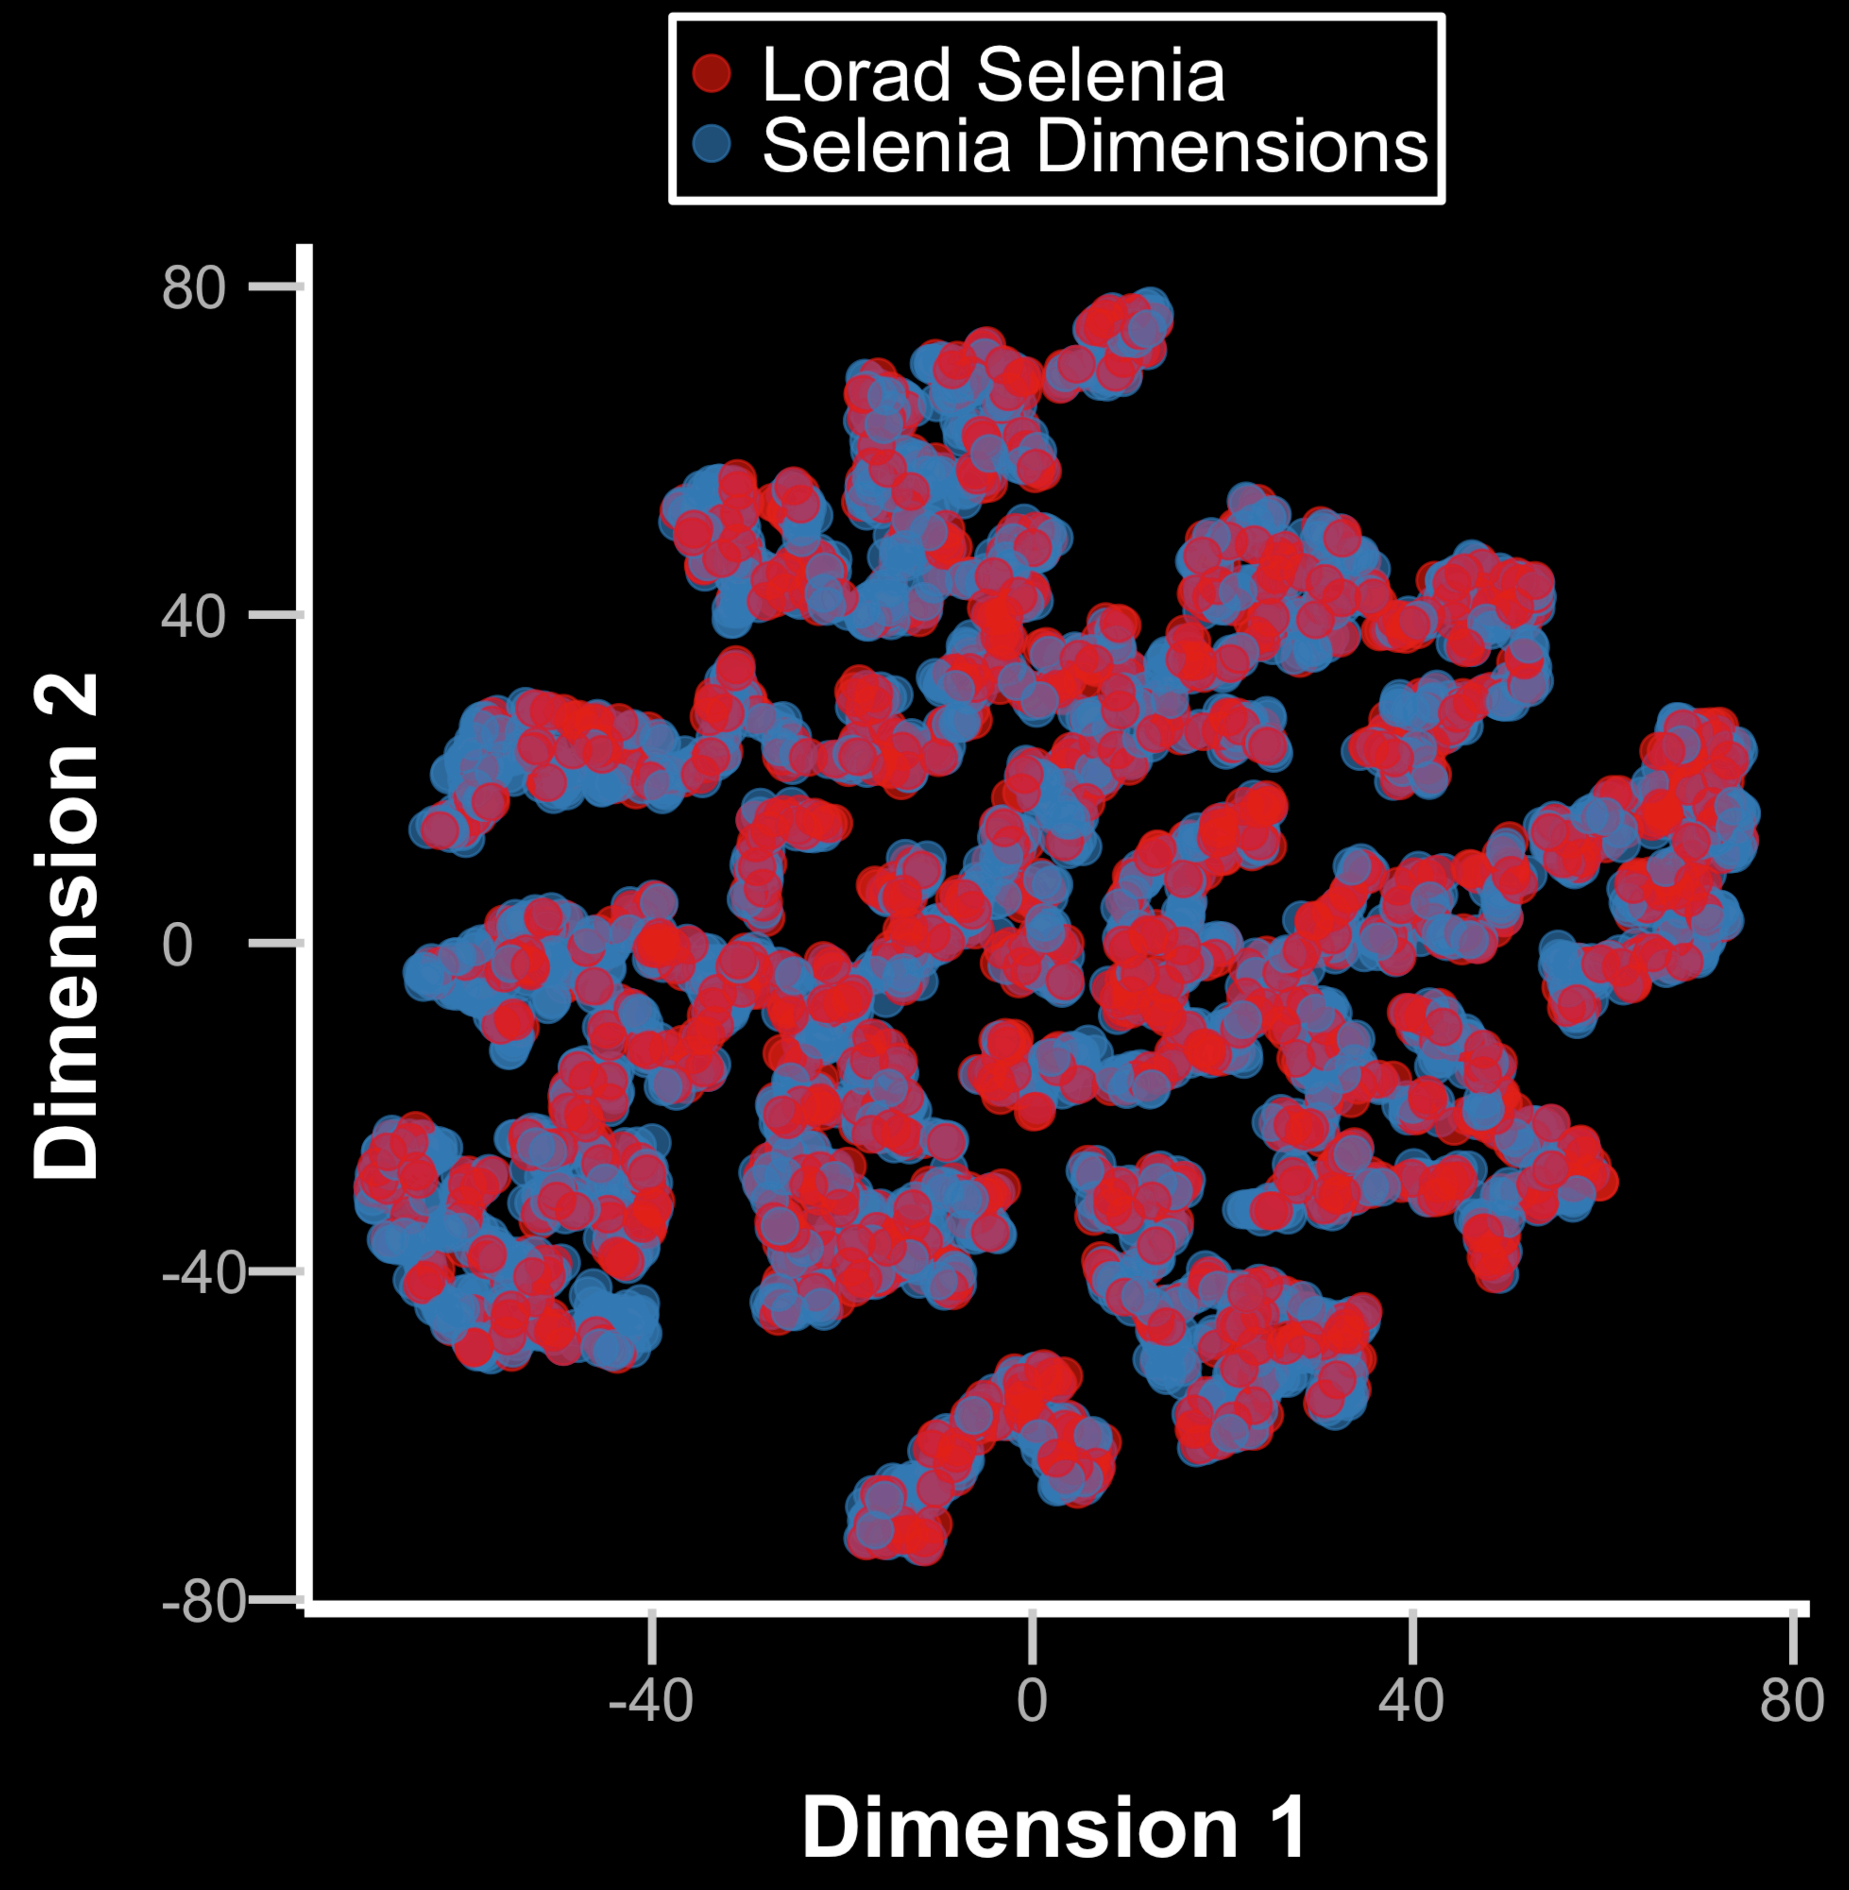

With Adversary

Problem 2: Missing risk factor data

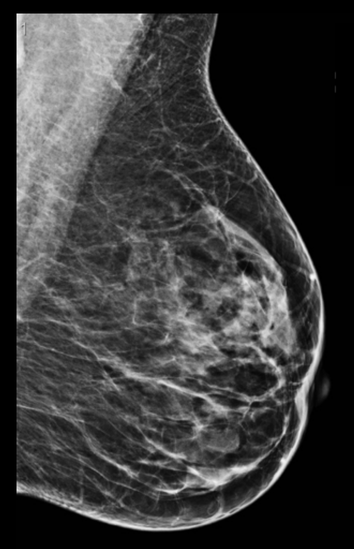

→ Encoder → Risk Factor Predictor →

Risk Factors
Age
Family History
Prior biopsy
Num children
⋮
Menopause status

Problem 2: Missing risk factor data

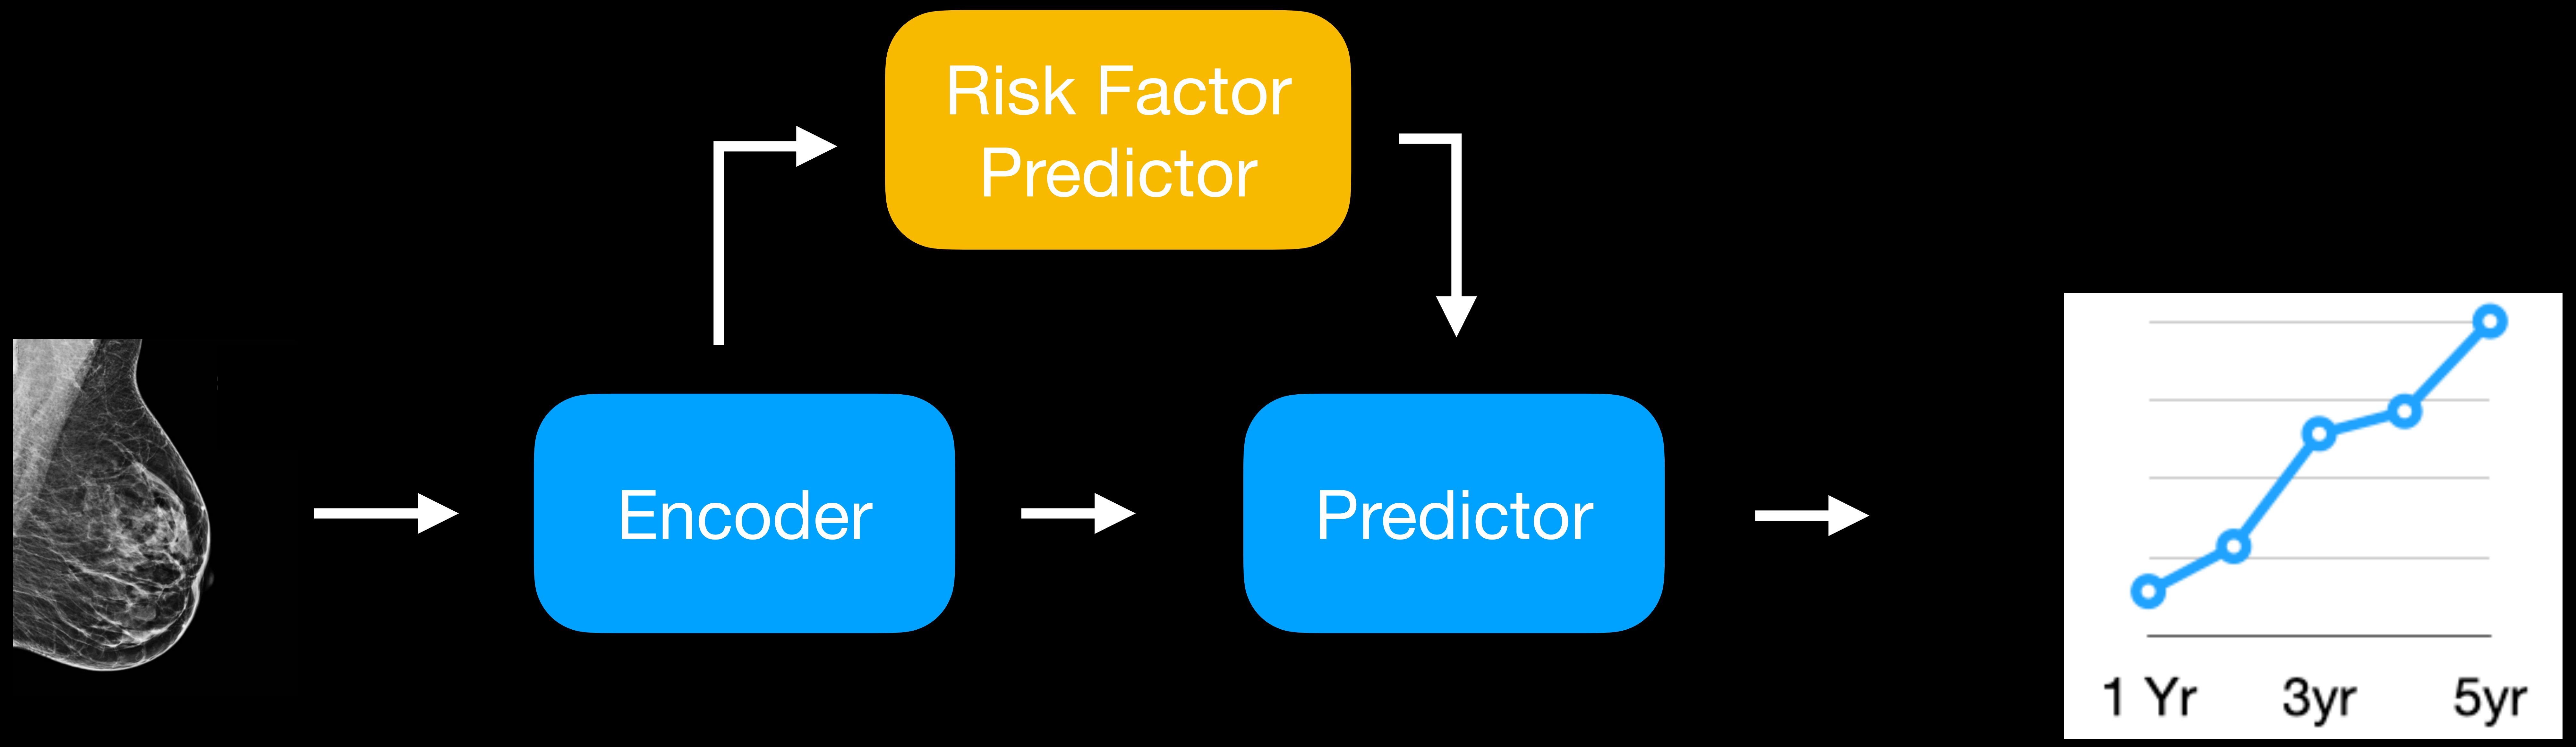

Problem 3: Modeling risk over time

→ Encoder → B, H_1, ⋮, H_5 →

Cumulative Risk

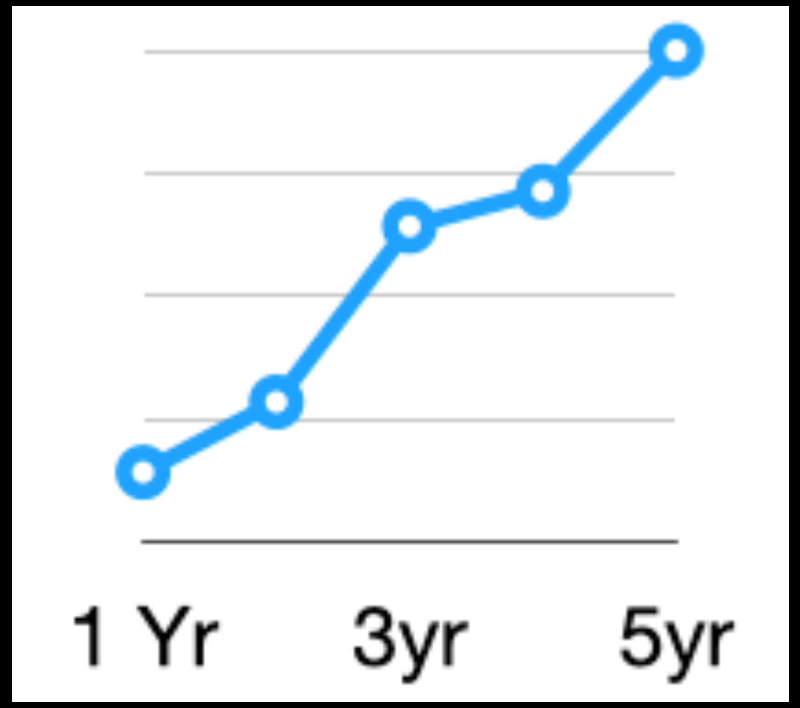

$$P(t_{cancer} = k \mid x) = B(E(x)) + \sum H_i(E(x))$$

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 Lin[5], Kevin Smith[6,7], Yung-Liang Wan[5], Leslie Lamb[8], Kevin Hughes[9], Constance Lehman[8†], Regina Barzilay[1,2†]

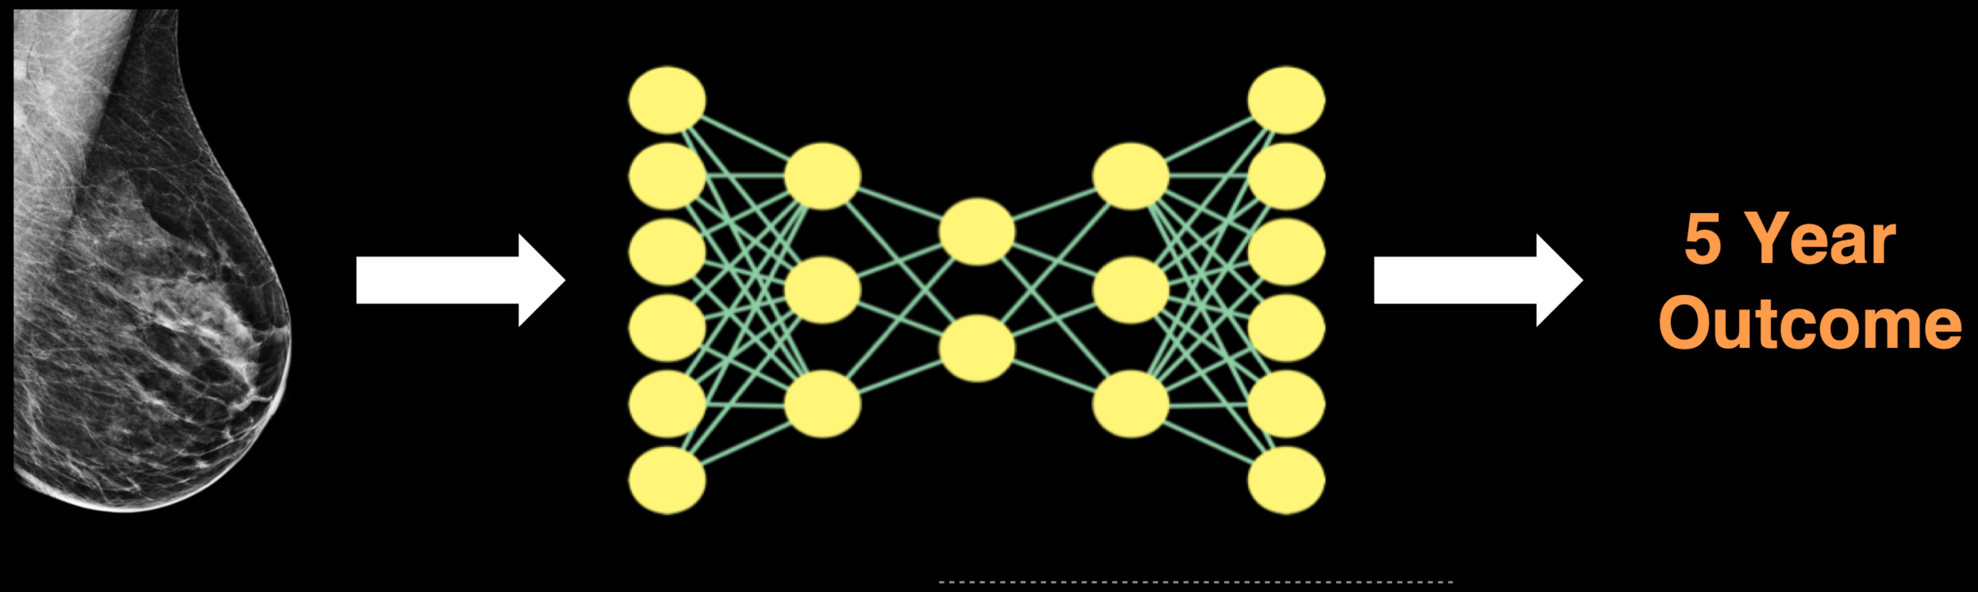

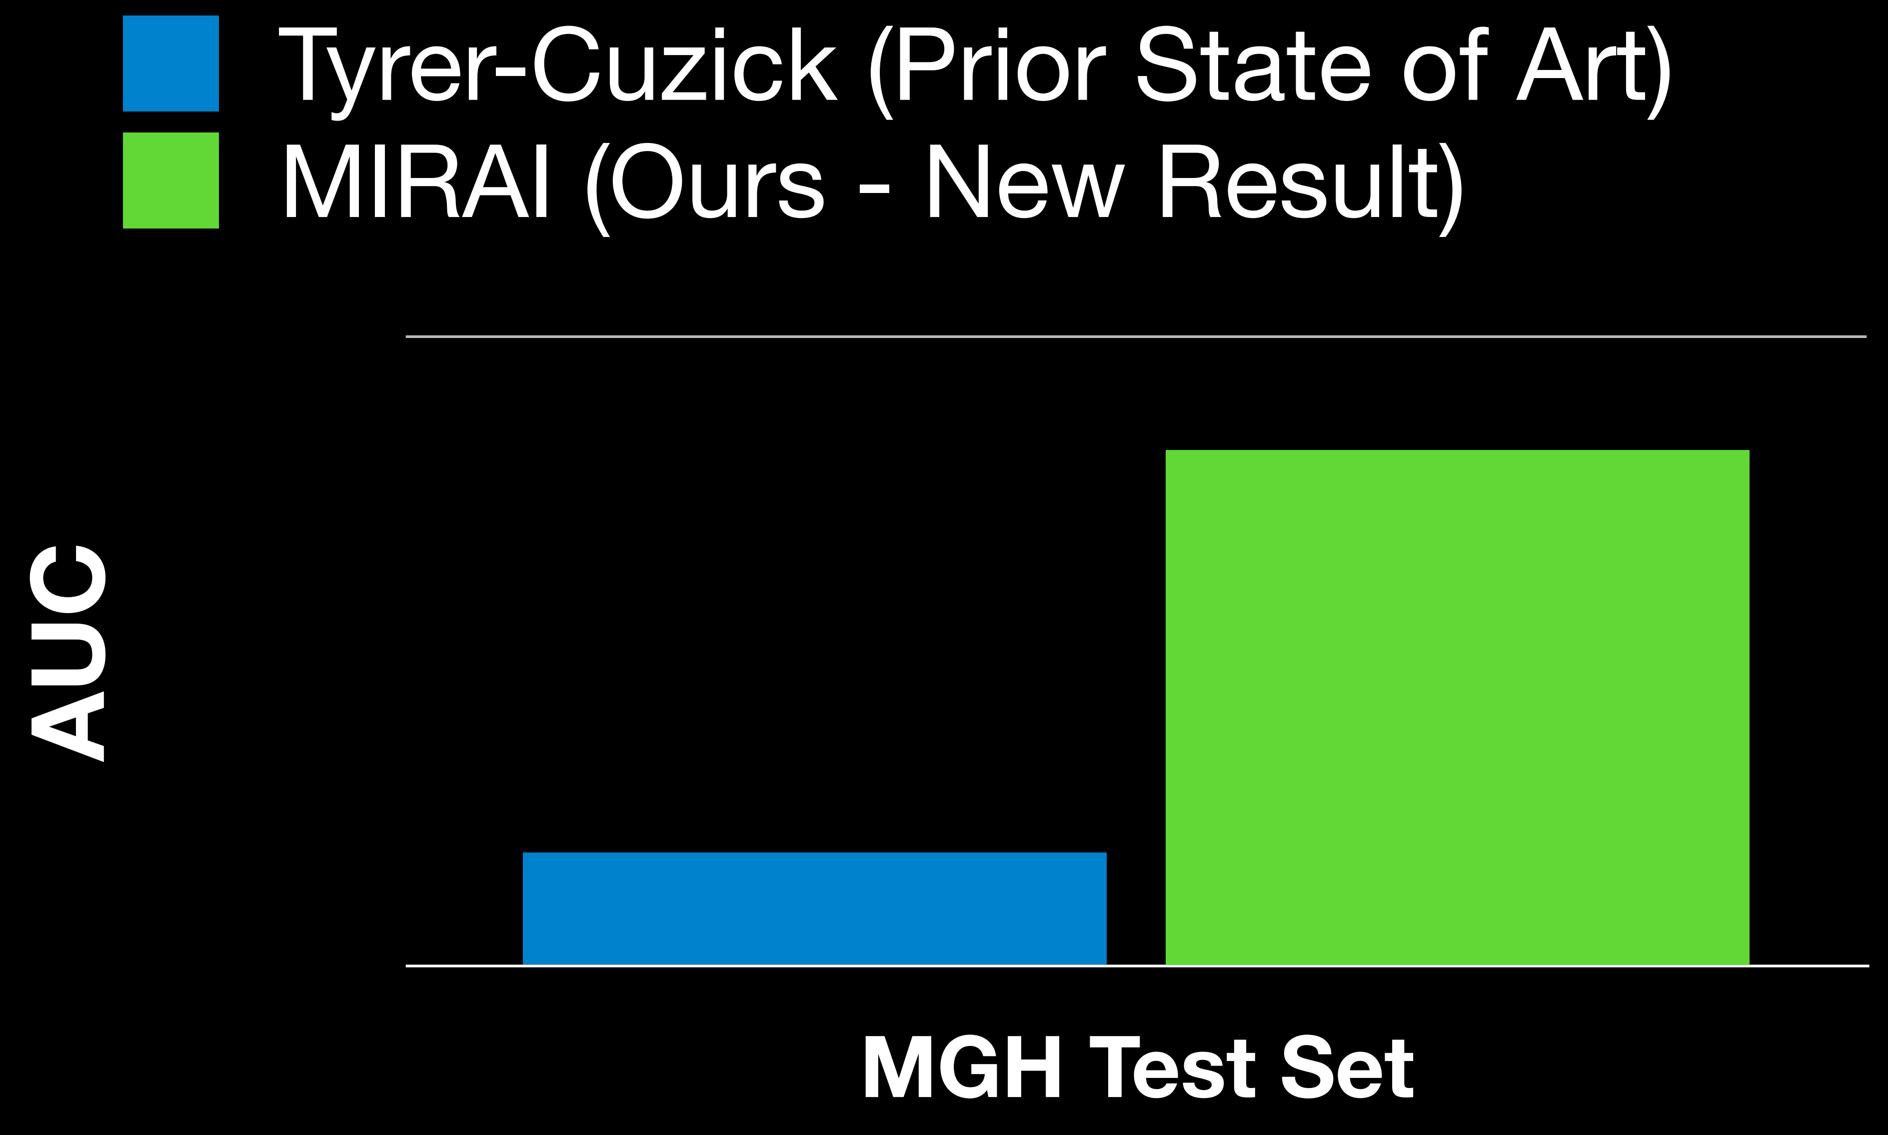

Journal of Clinical Oncology®
An American Society of Clinical Oncology Journal

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

Peter G. Mikhael, BSc[1,2]; Jeremy Wohlwend, ME[1,2]; Adam Yala, PhD[1,2]; Ludvig Karstens, MSc[1,2]; Justin Xiang, ME[1,2]; Angelo K. Takigami, MD[3,4]; Patrick P. Bourgouin, MD[3,4]; PuiYee Chan, PhD[5]; Sofiane Mrah, MSc[4]; Wael Amayri, BSc[4]; Yu-Hsiang Juan, MD[6,7]; Cheng-Ta Yang, MD[6,8]; Yung-Liang Wan, MD[6,7]; Gigin Lin, MD, PhD[6,7]; Lecia V. Sequist, MD, MPH[3,5]; Florian J. Fintelmann, MD[3,4]; and Regina Barzilay, PhD[1,2]

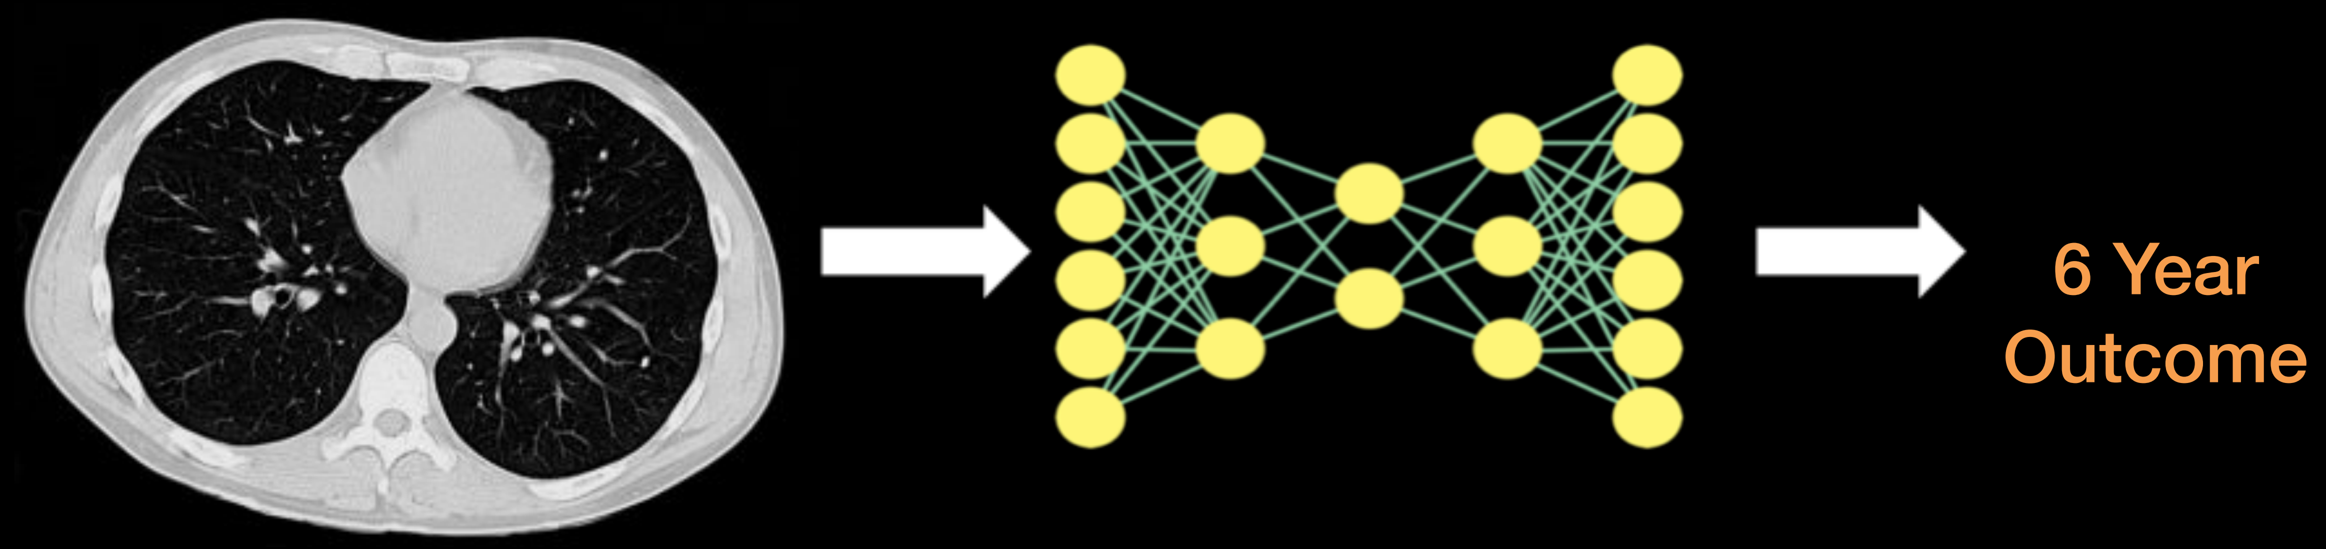

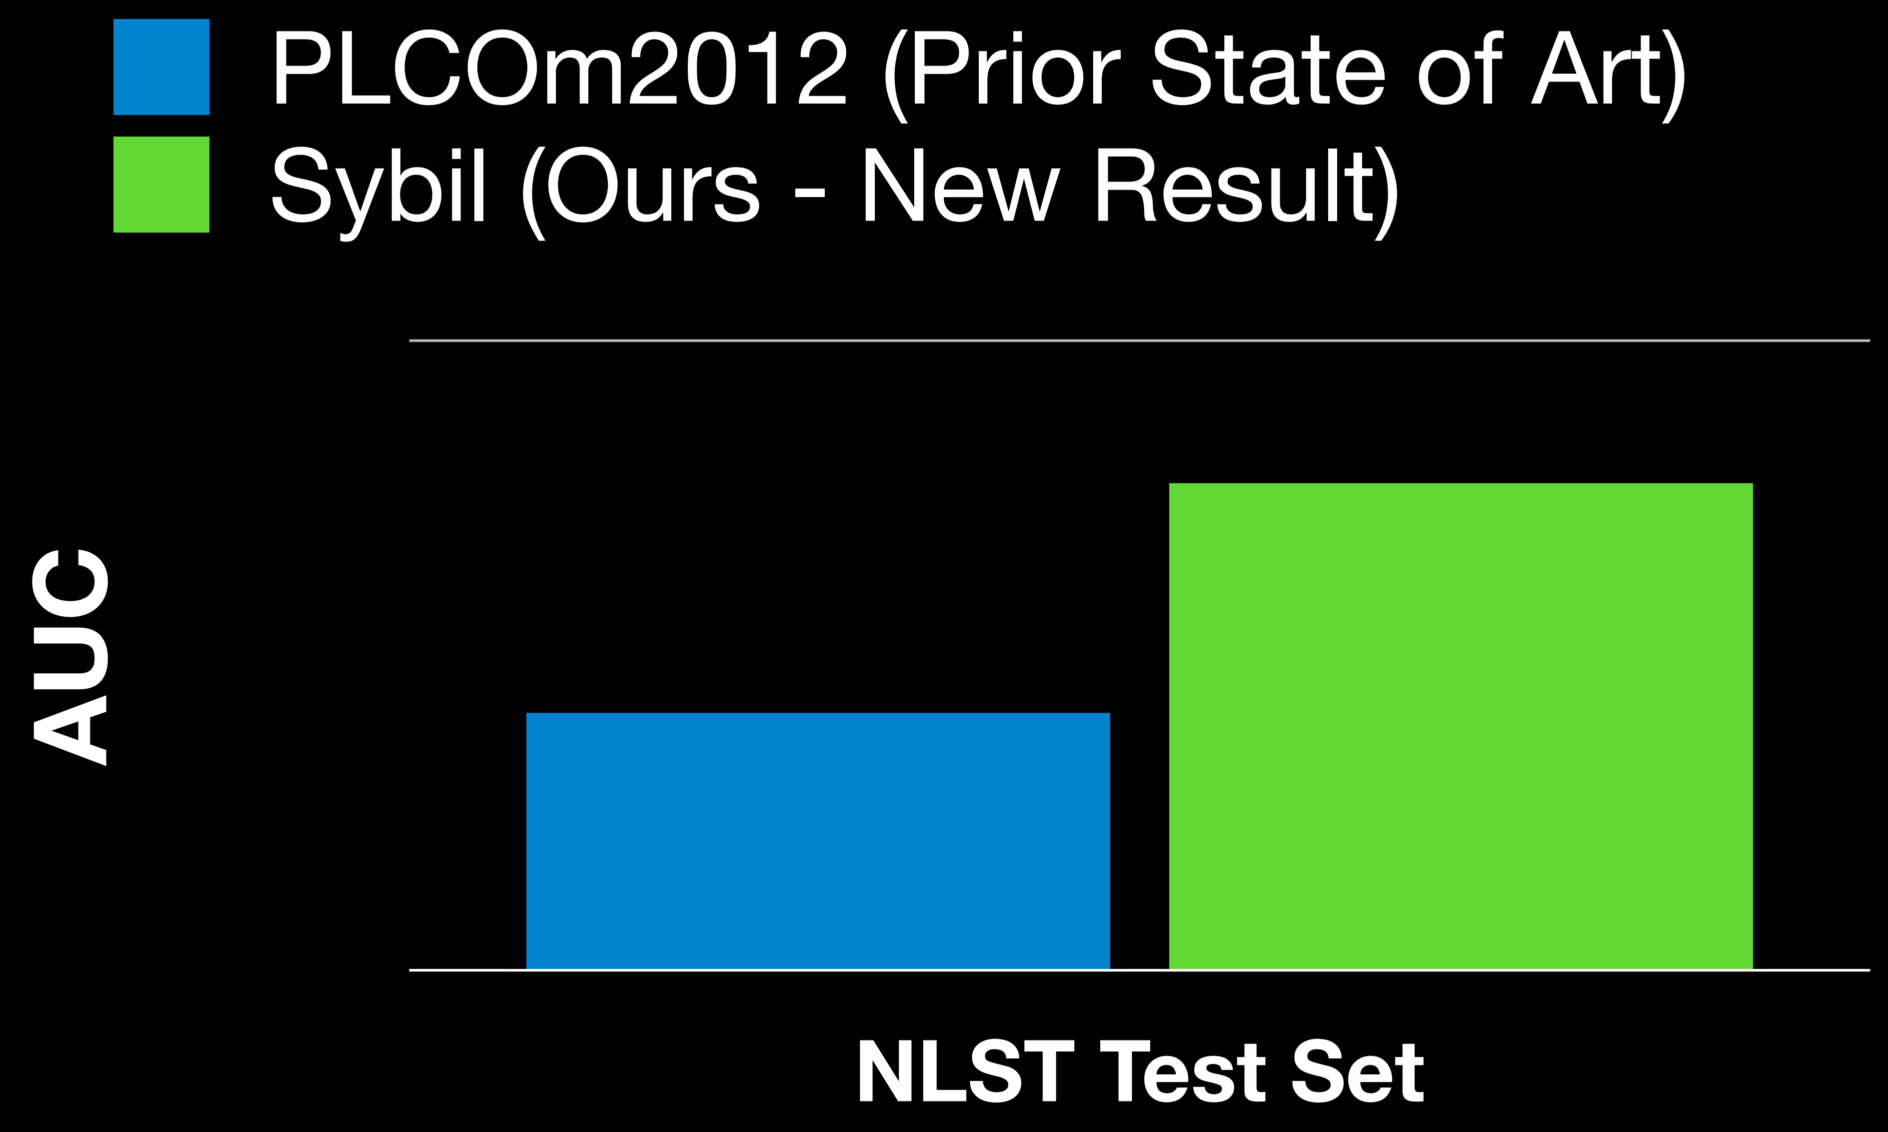

The harms of late diagnosis

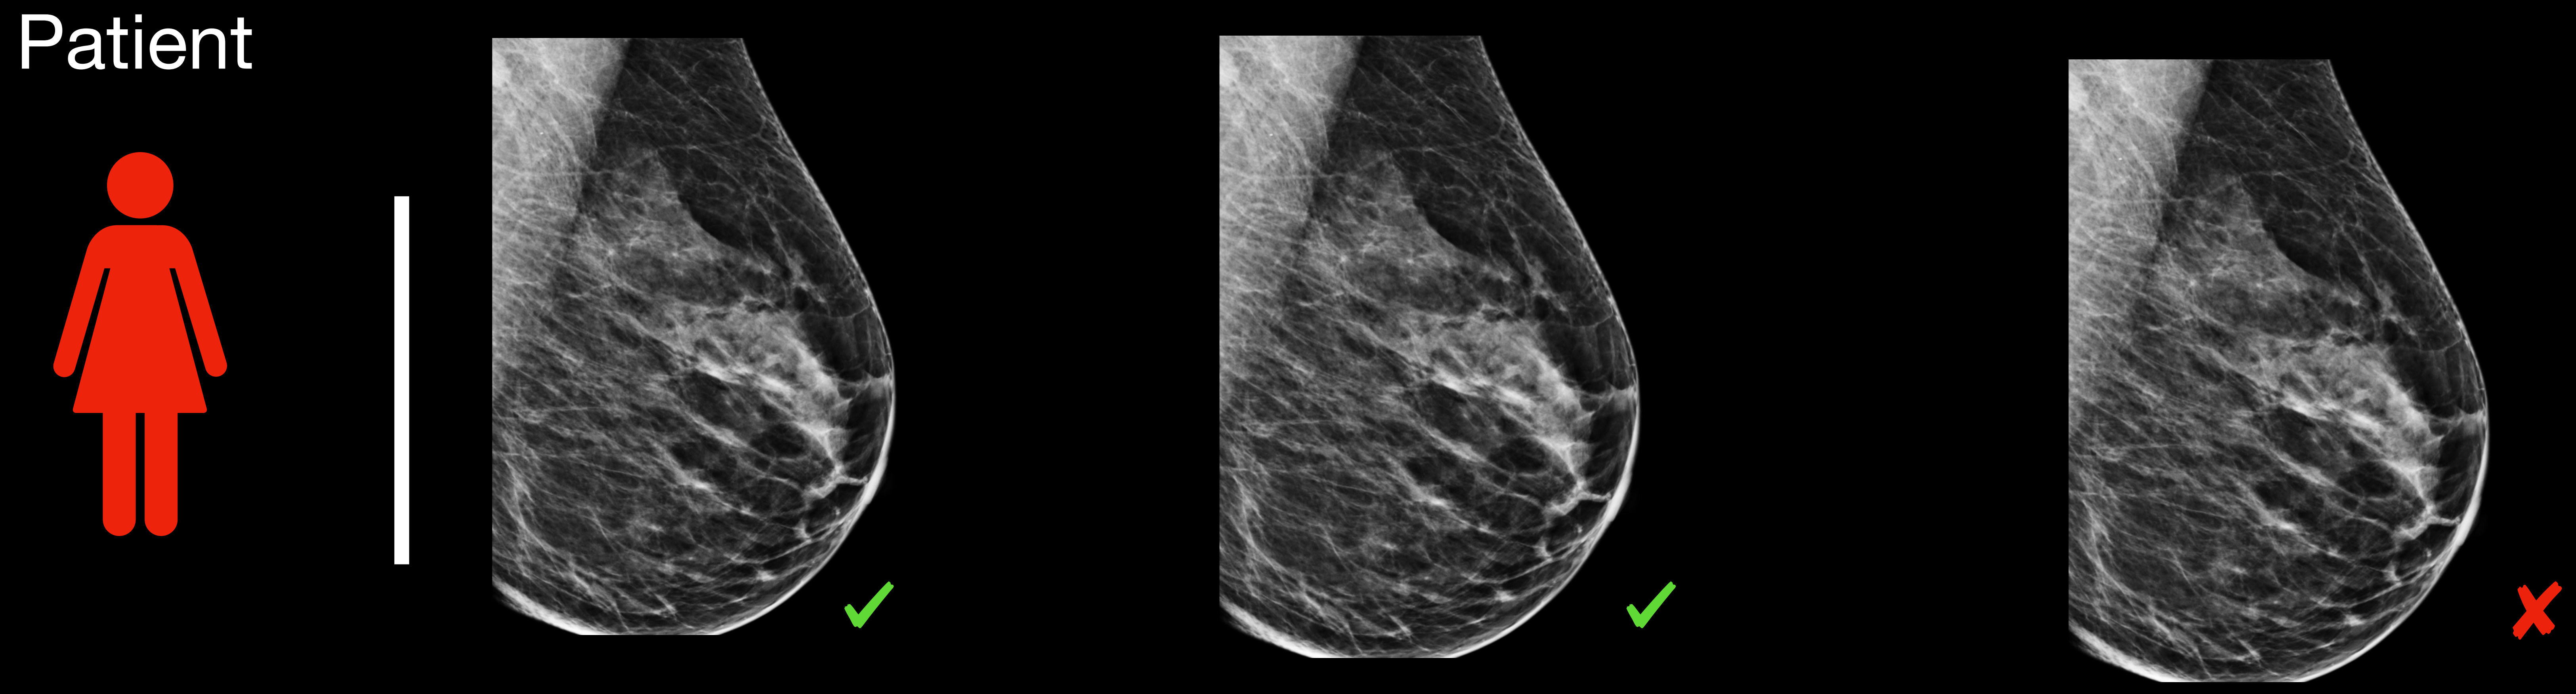

Morbid treatment options, poor chances of survival

We should have done more

Ongoing Prospective Trials: Mirai-MRI

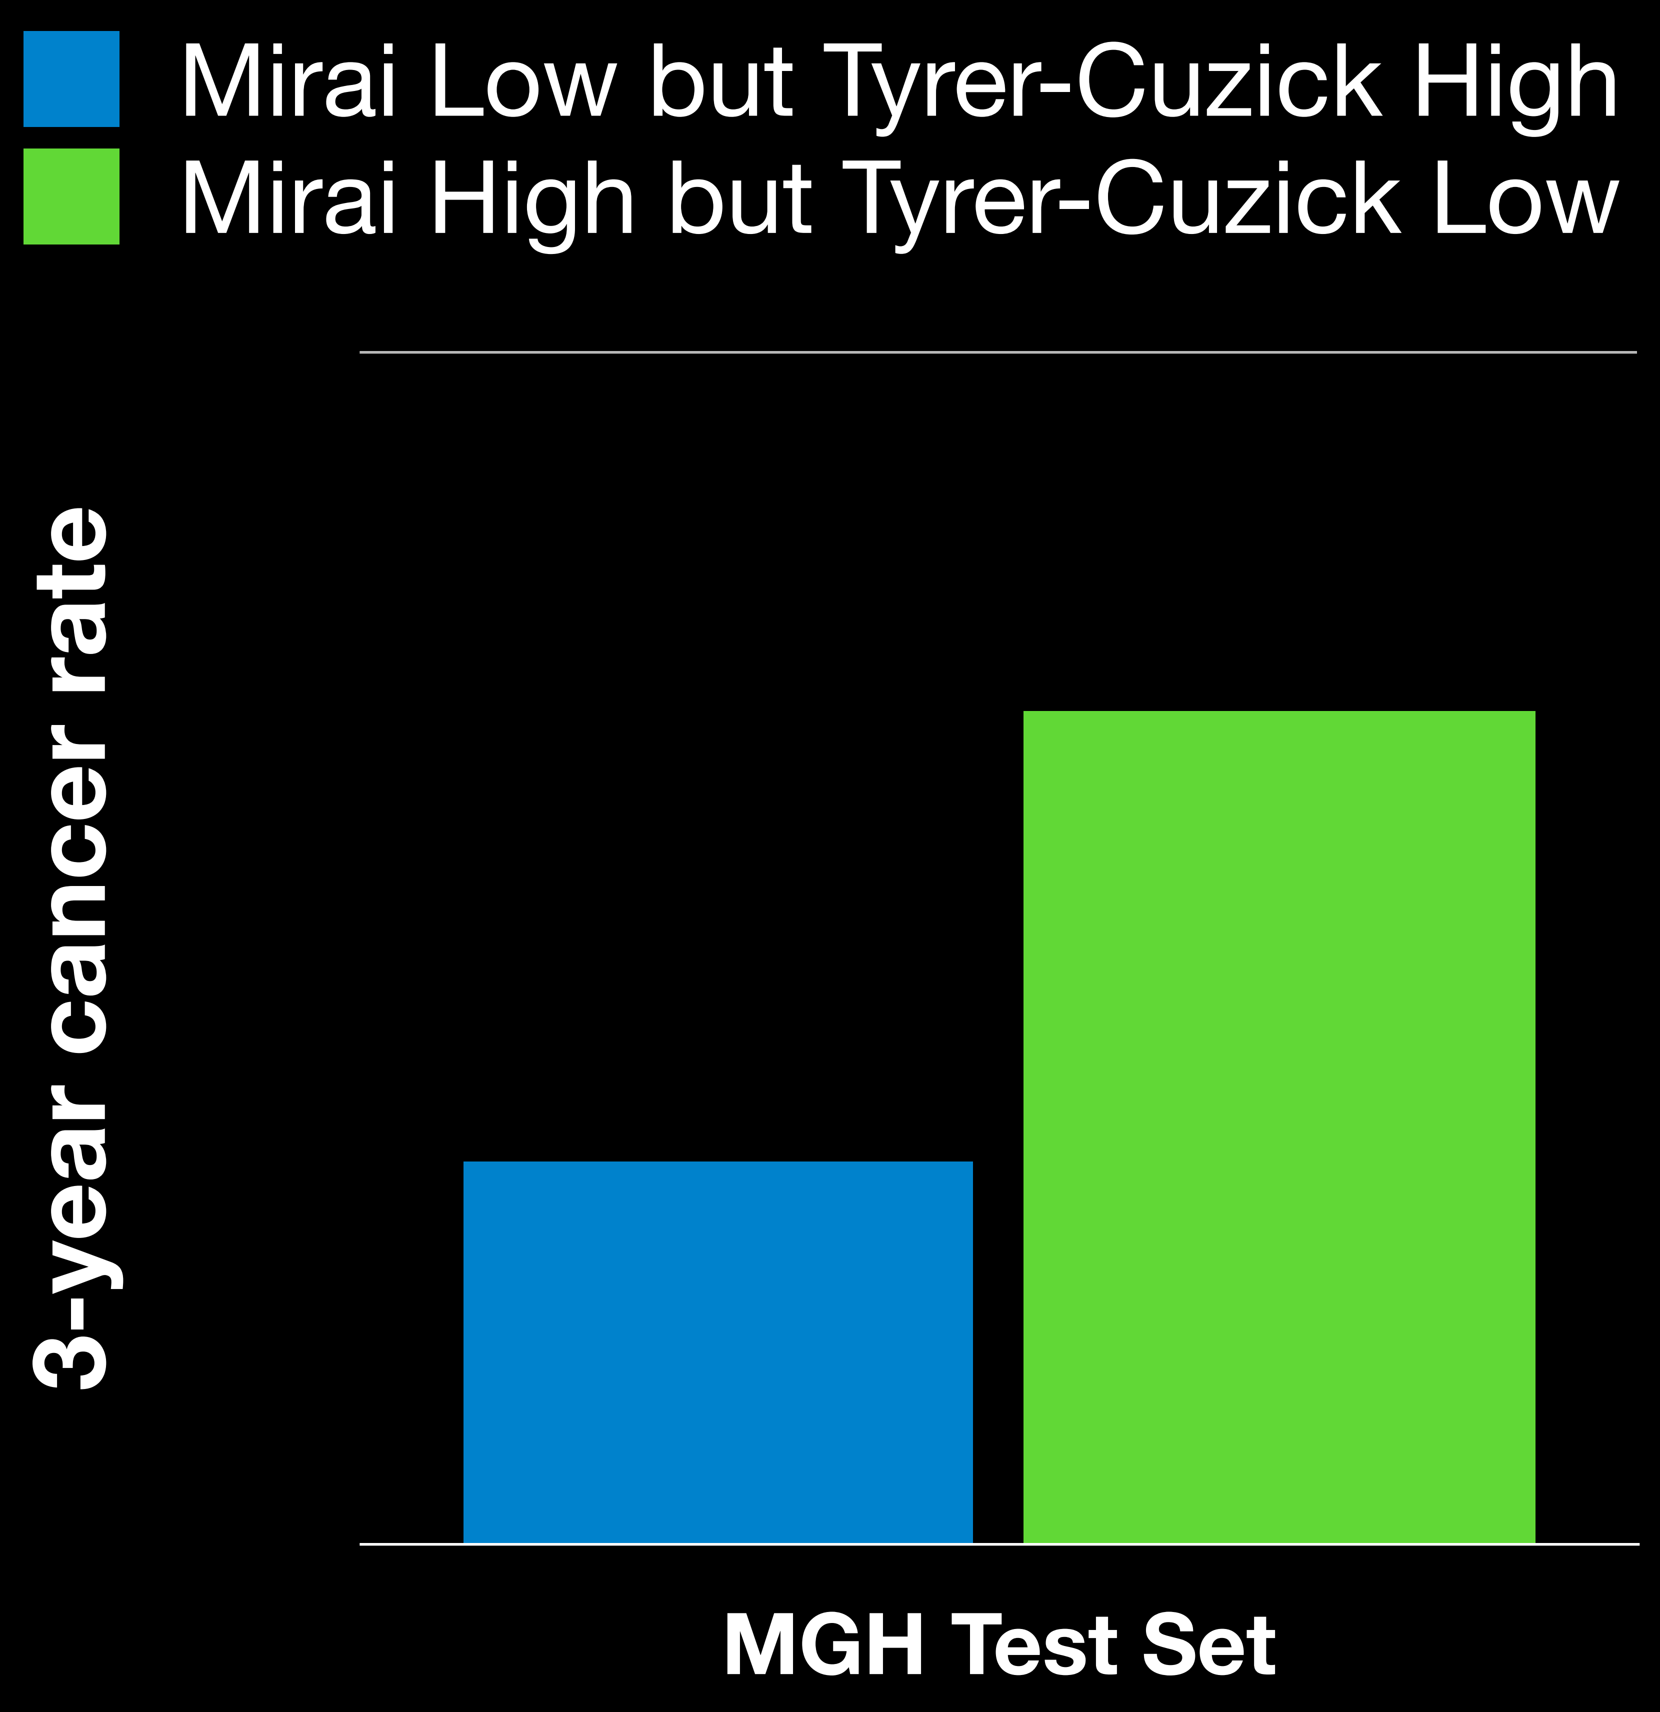

Retrospective analysis

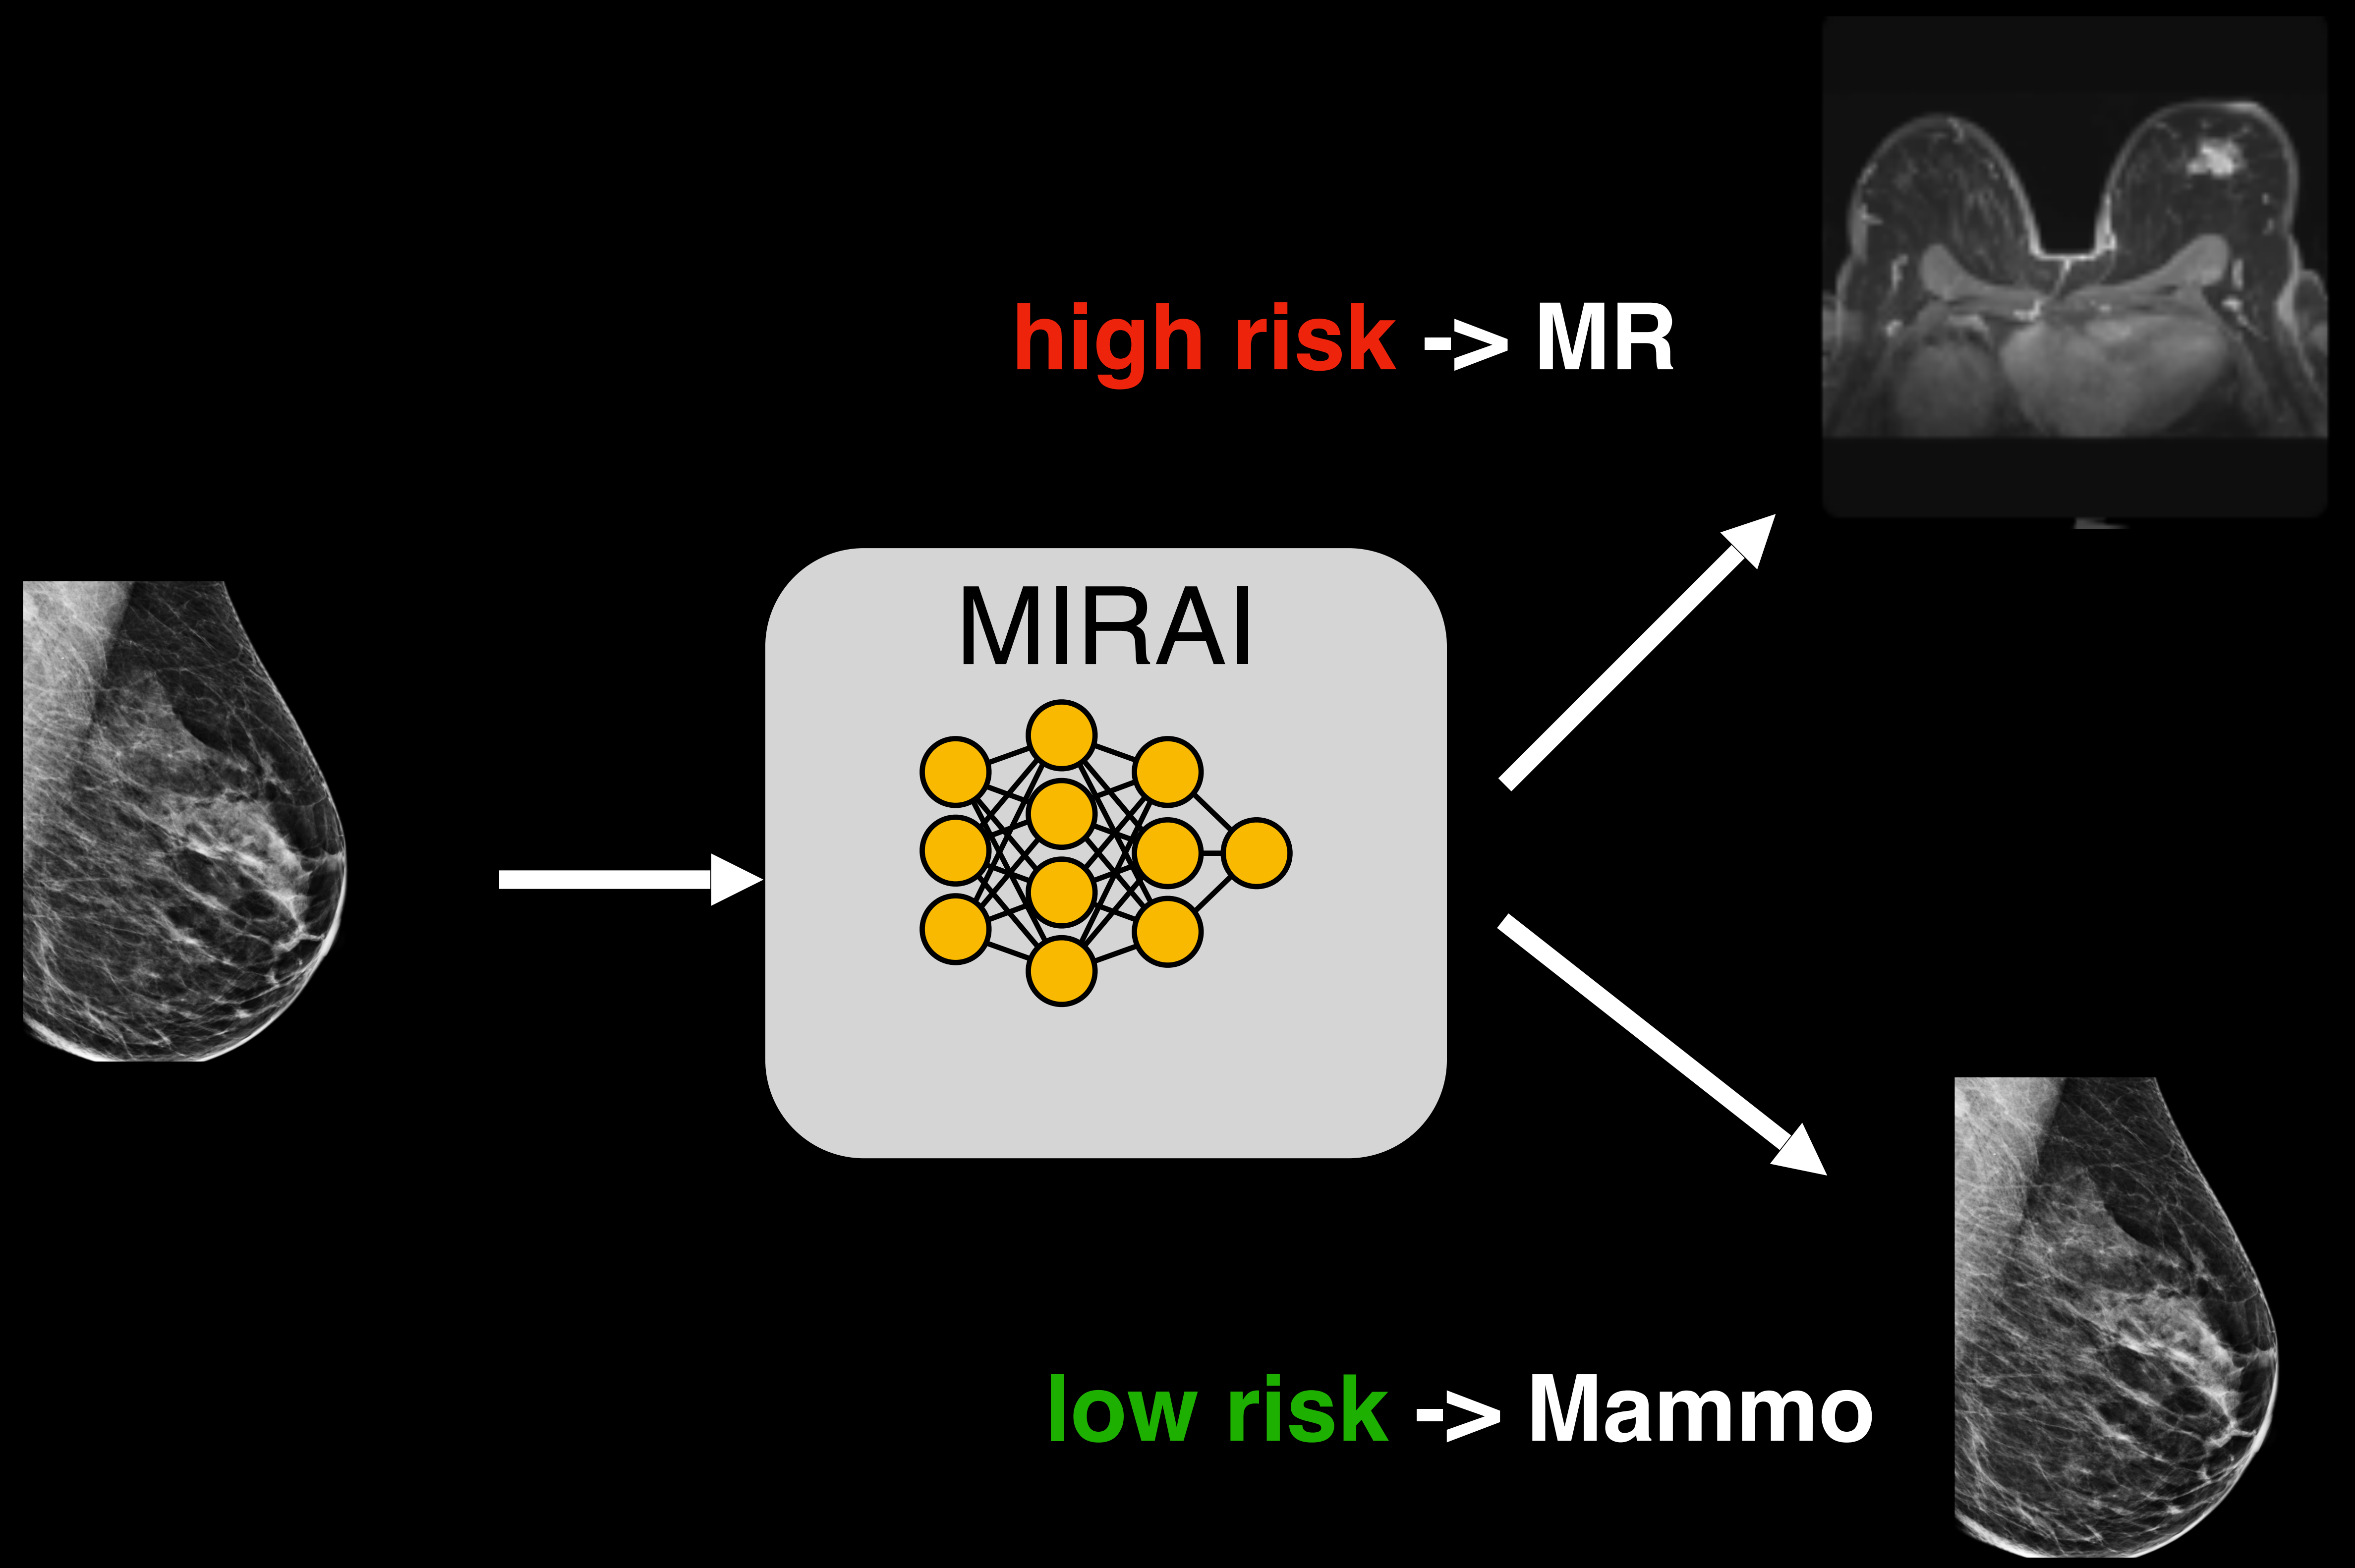

Mirai-based Supplemental Imaging

NCT 05968157

The harms of over screening

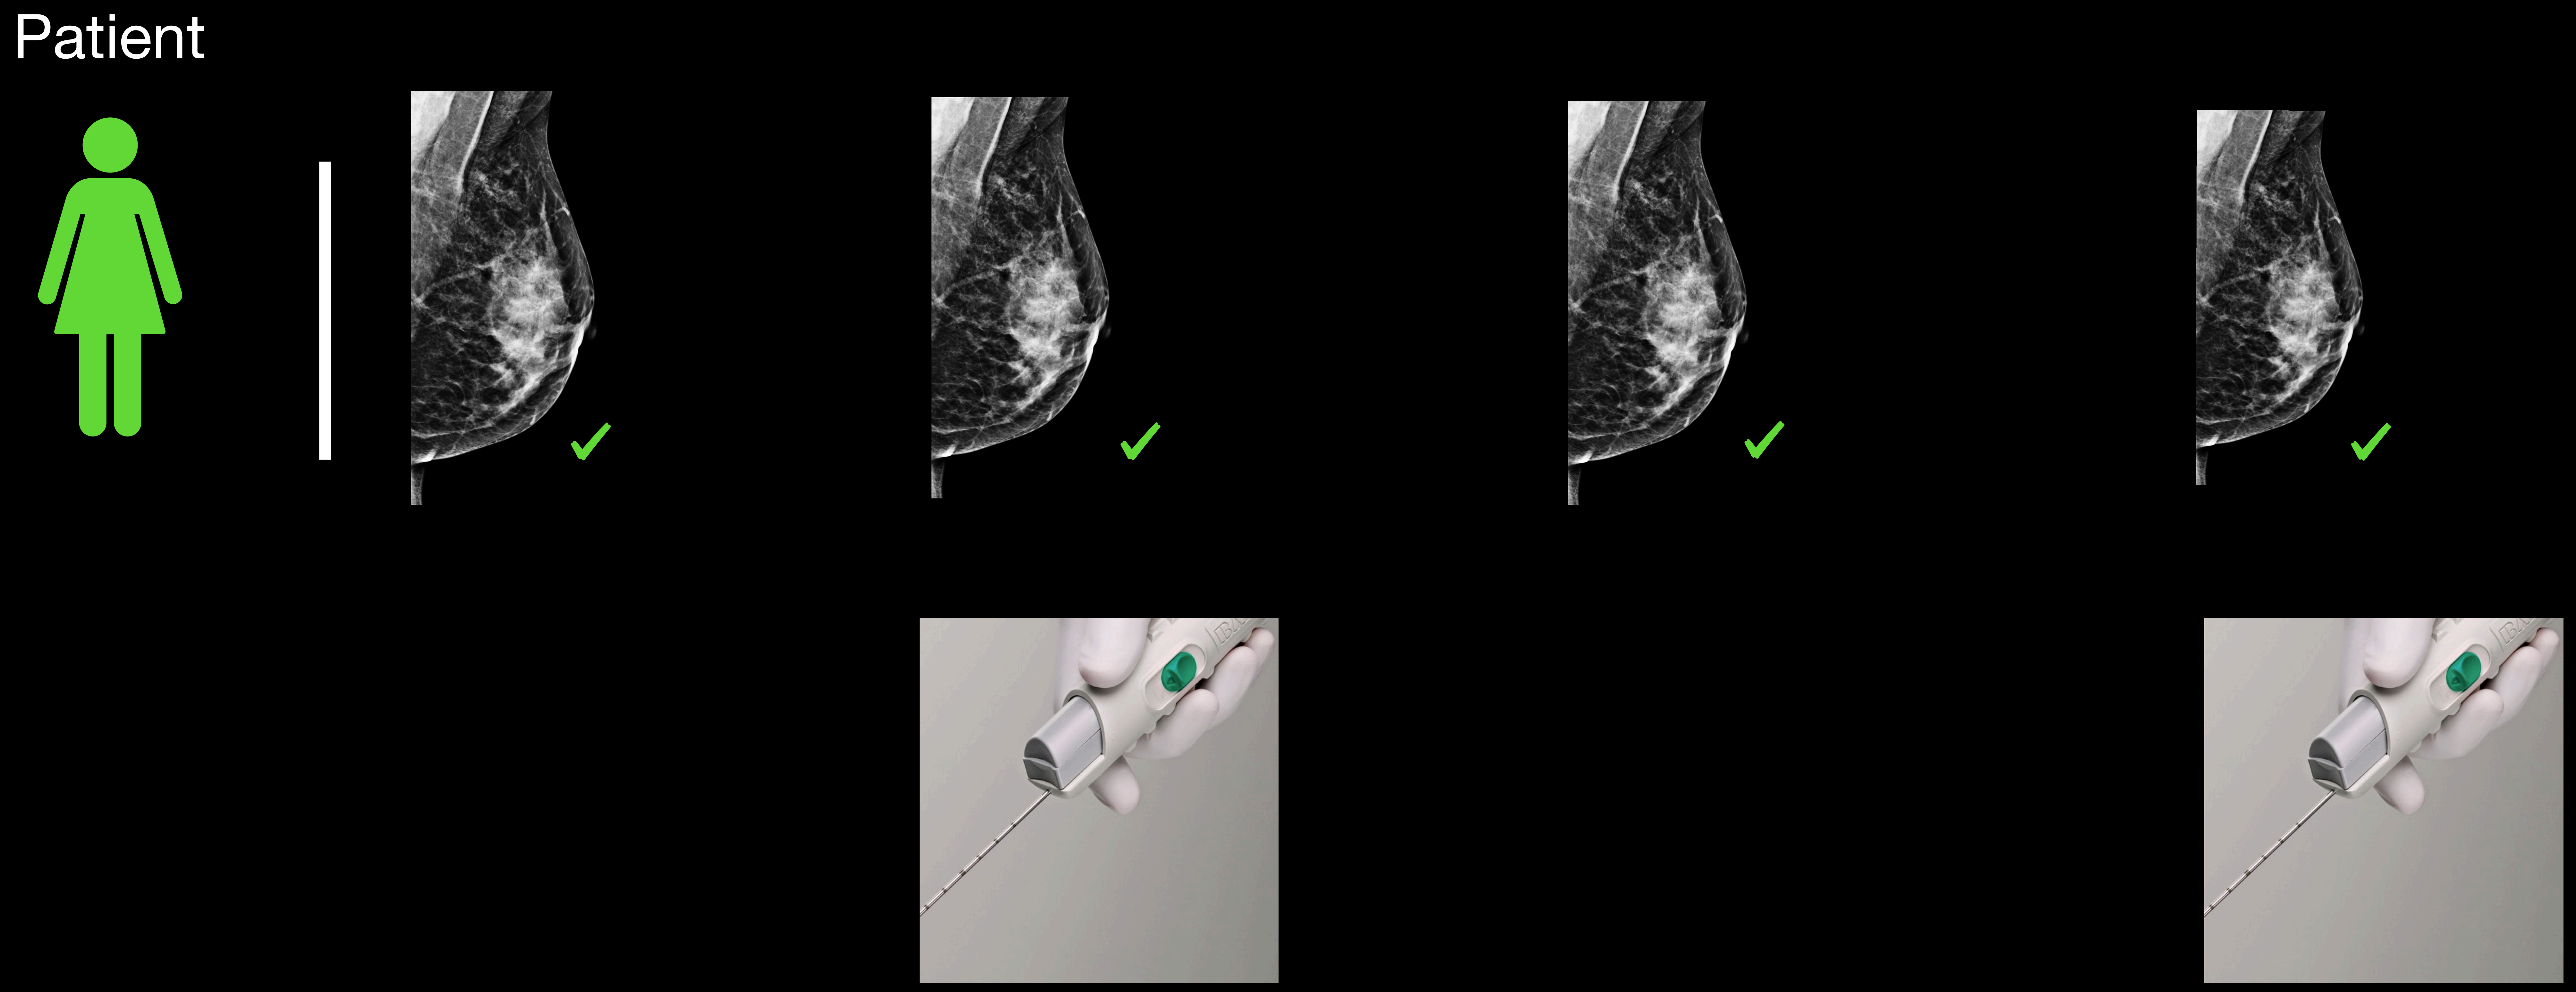

Unnecessary biopsies, terrible anxiety

We should have done less

Ongoing Prospective Trials: Mirai-SDA

SDA Workflow:

- **Realtime** AI-based cancer risk assessment
- Invite high risk patient for **same-day** diagnostic exam
- Prelim results: Reducing time to diagnosis from **38 days to 58 minutes**

Screen

<5 seconds

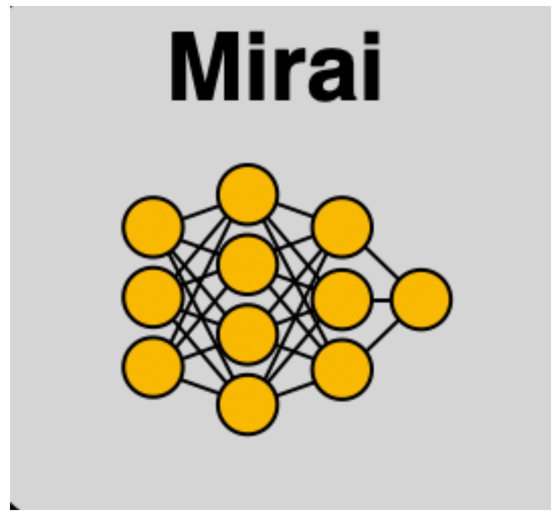

AI Risk

Top 10% Risk

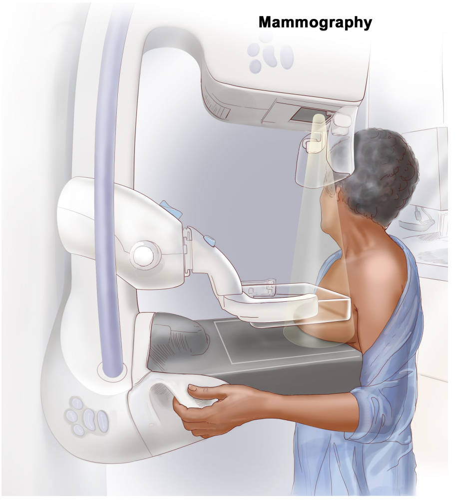

Same Day Diagnostic

Today: Towards AI-driven care

Control

Both Mirai-MRI and Mirai-SDA are heuristic control algorithms.

Opportunities to design guidelines as learned algorithms!


nature medicine
ARTICLES
https://doi.org/10.1038/s41591-021-01599-w

Check for updates

Optimizing risk-based breast cancer screening policies with reinforcement learning

Adam Yala[1,2], Peter G. Mikhael[1,2], Constance Lehman[3], Gigin Lin[4,5], Fredrik Strand[6,7], Yung-Liang Wan[4,5], Kevin Hughes[8], Siddharth Satuluru[9], Thomas Kim[10], Imon Banerjee[11], Judy Gichoya[9], Hari Trivedi[9] and Regina Barzilay[1,2]

Next AI Leap : Modeling full patient context

MB of data

GB of data

Tomosynthesis

Pathology

2D Mammogram

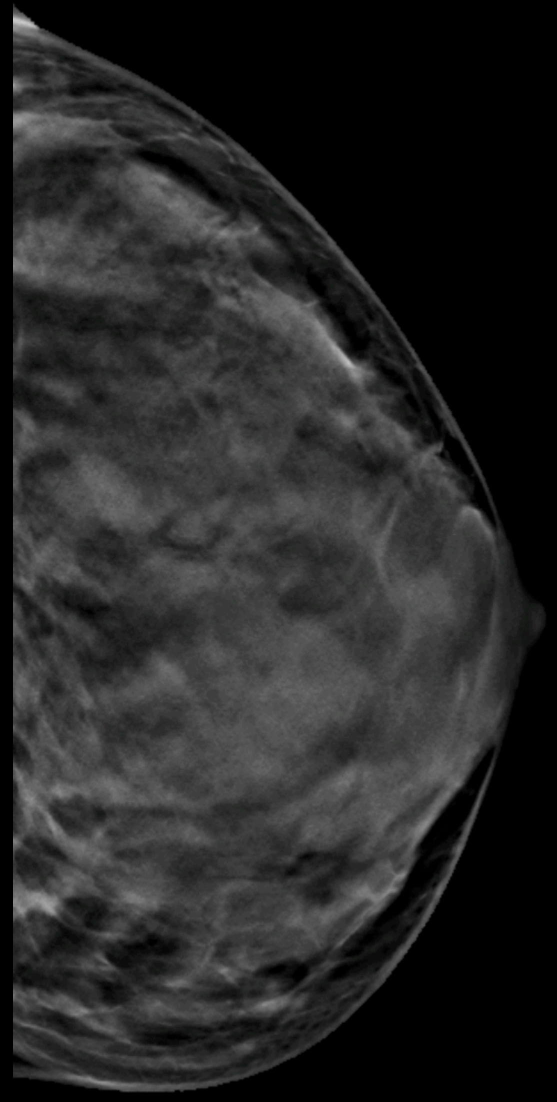

MRI

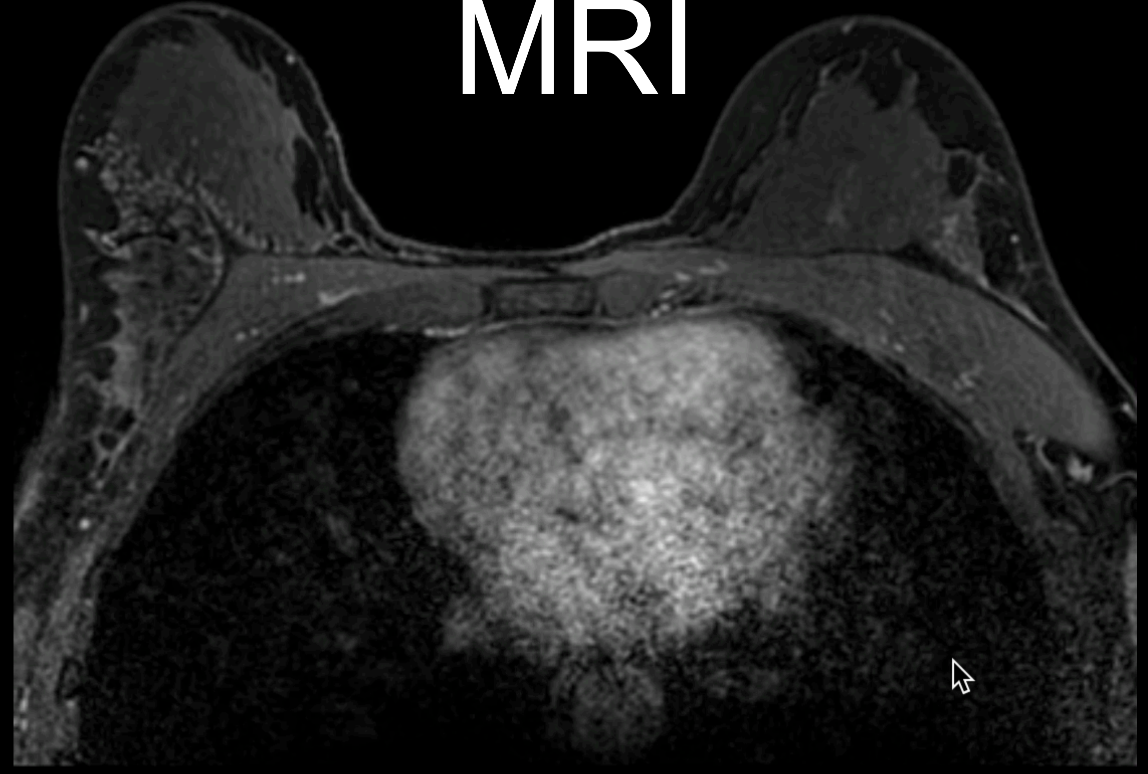

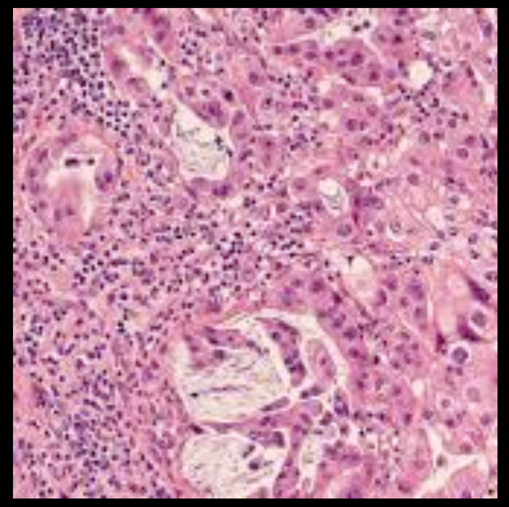

Full care record: notes, codes and labs

Modeling: How can we model *giant* inputs?

Prediction

***Atlas*: Multi-Scale Attention Improves Long Context Image Modeling**

Kumar Krishna Agrawal[*1†] **Long Lian**[*1] **Longchao Liu**[1] **Natalia Harguindeguy**[1 2] **Boyi Li**[1]
Alexander Bick[3] **Maggie Chung**[2] **Trevor Darrell**[1] **Adam Yala**[1 2]

Led by:

Krishna Agrawal

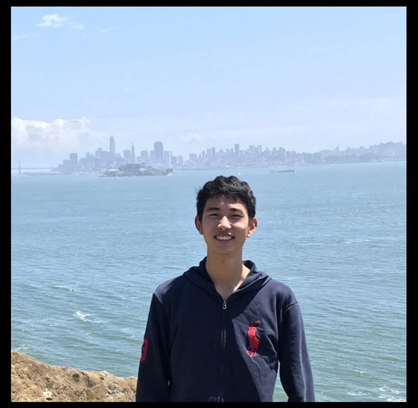

Tony Lian

Prelim: Images as sequences of tokens

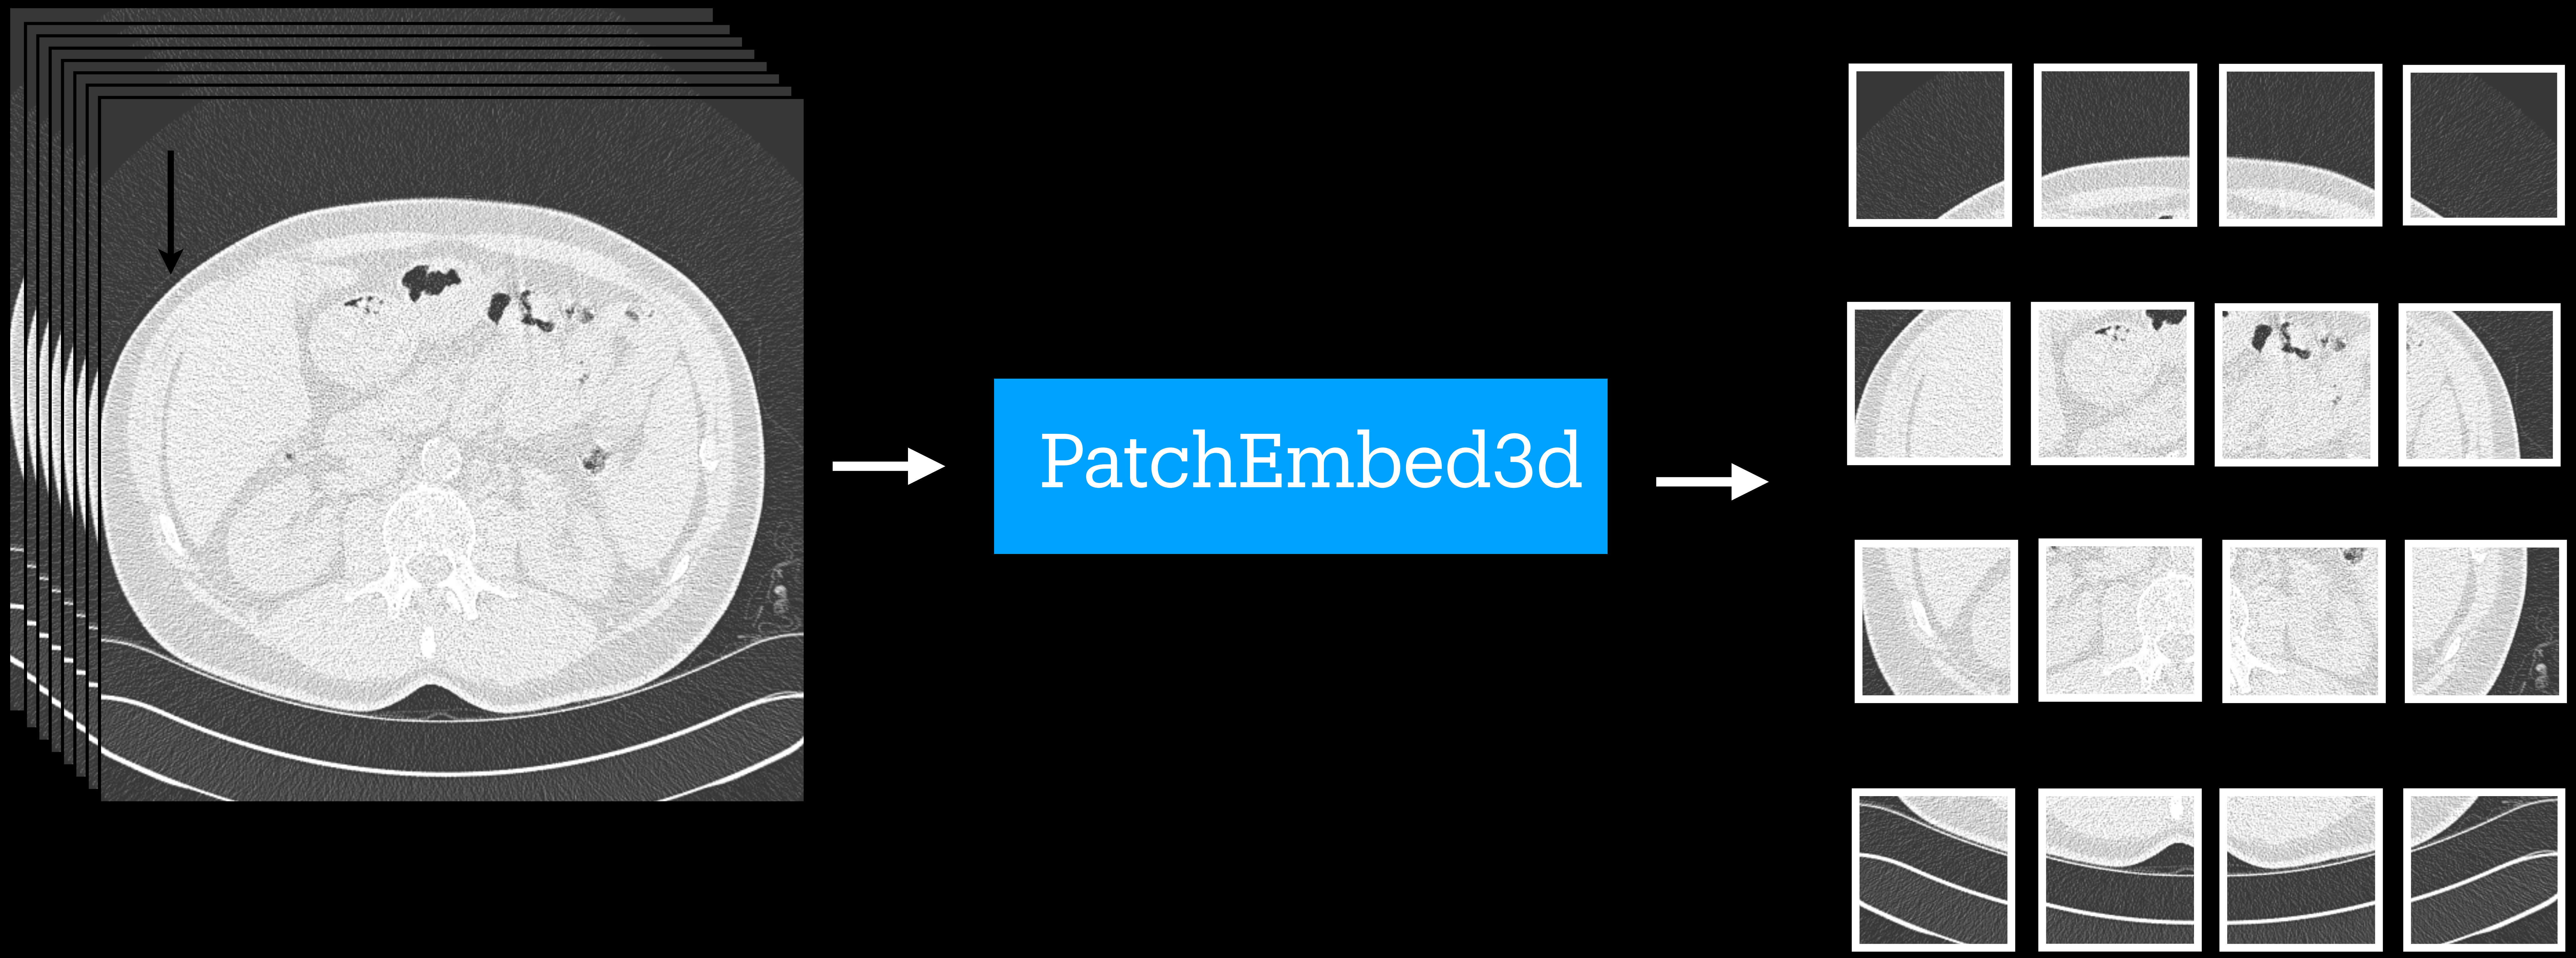

input resolution	200x512x512
patch-size	2x4x4
sequence-length	1.6M tokens

Prelim: Transformers and Self-attention

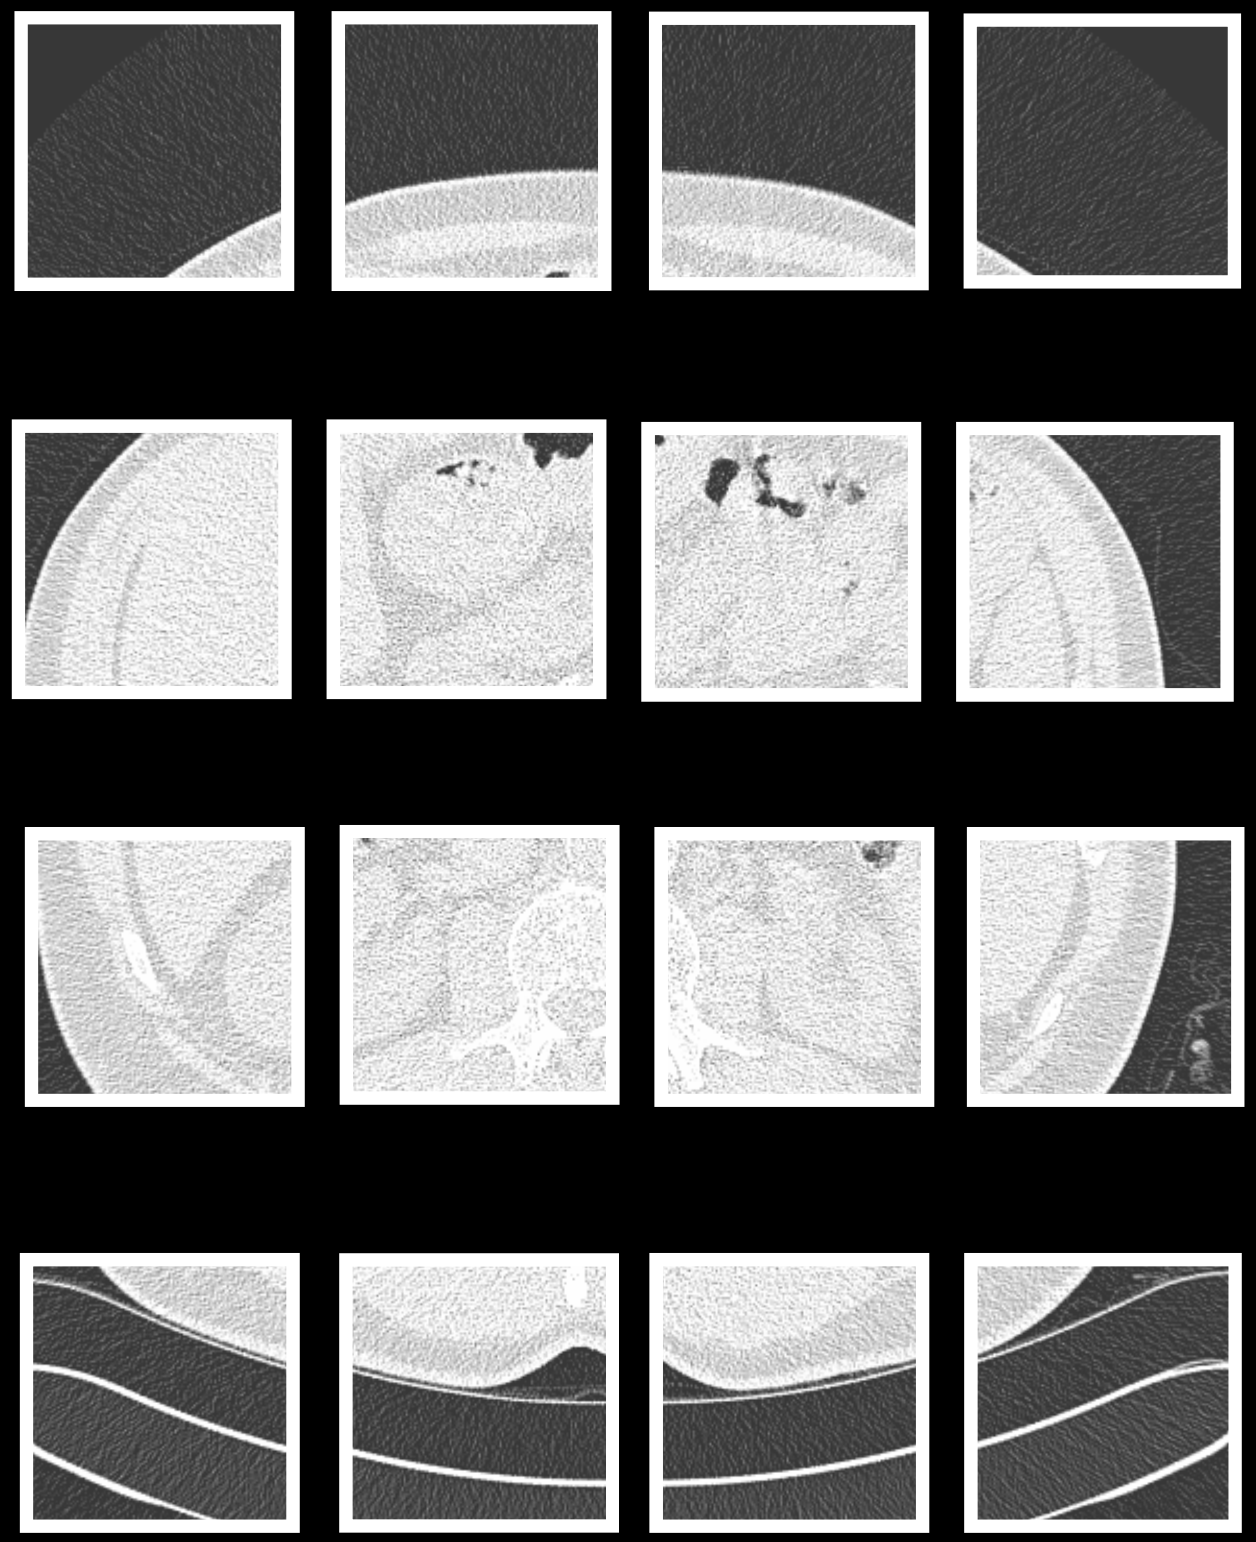

→

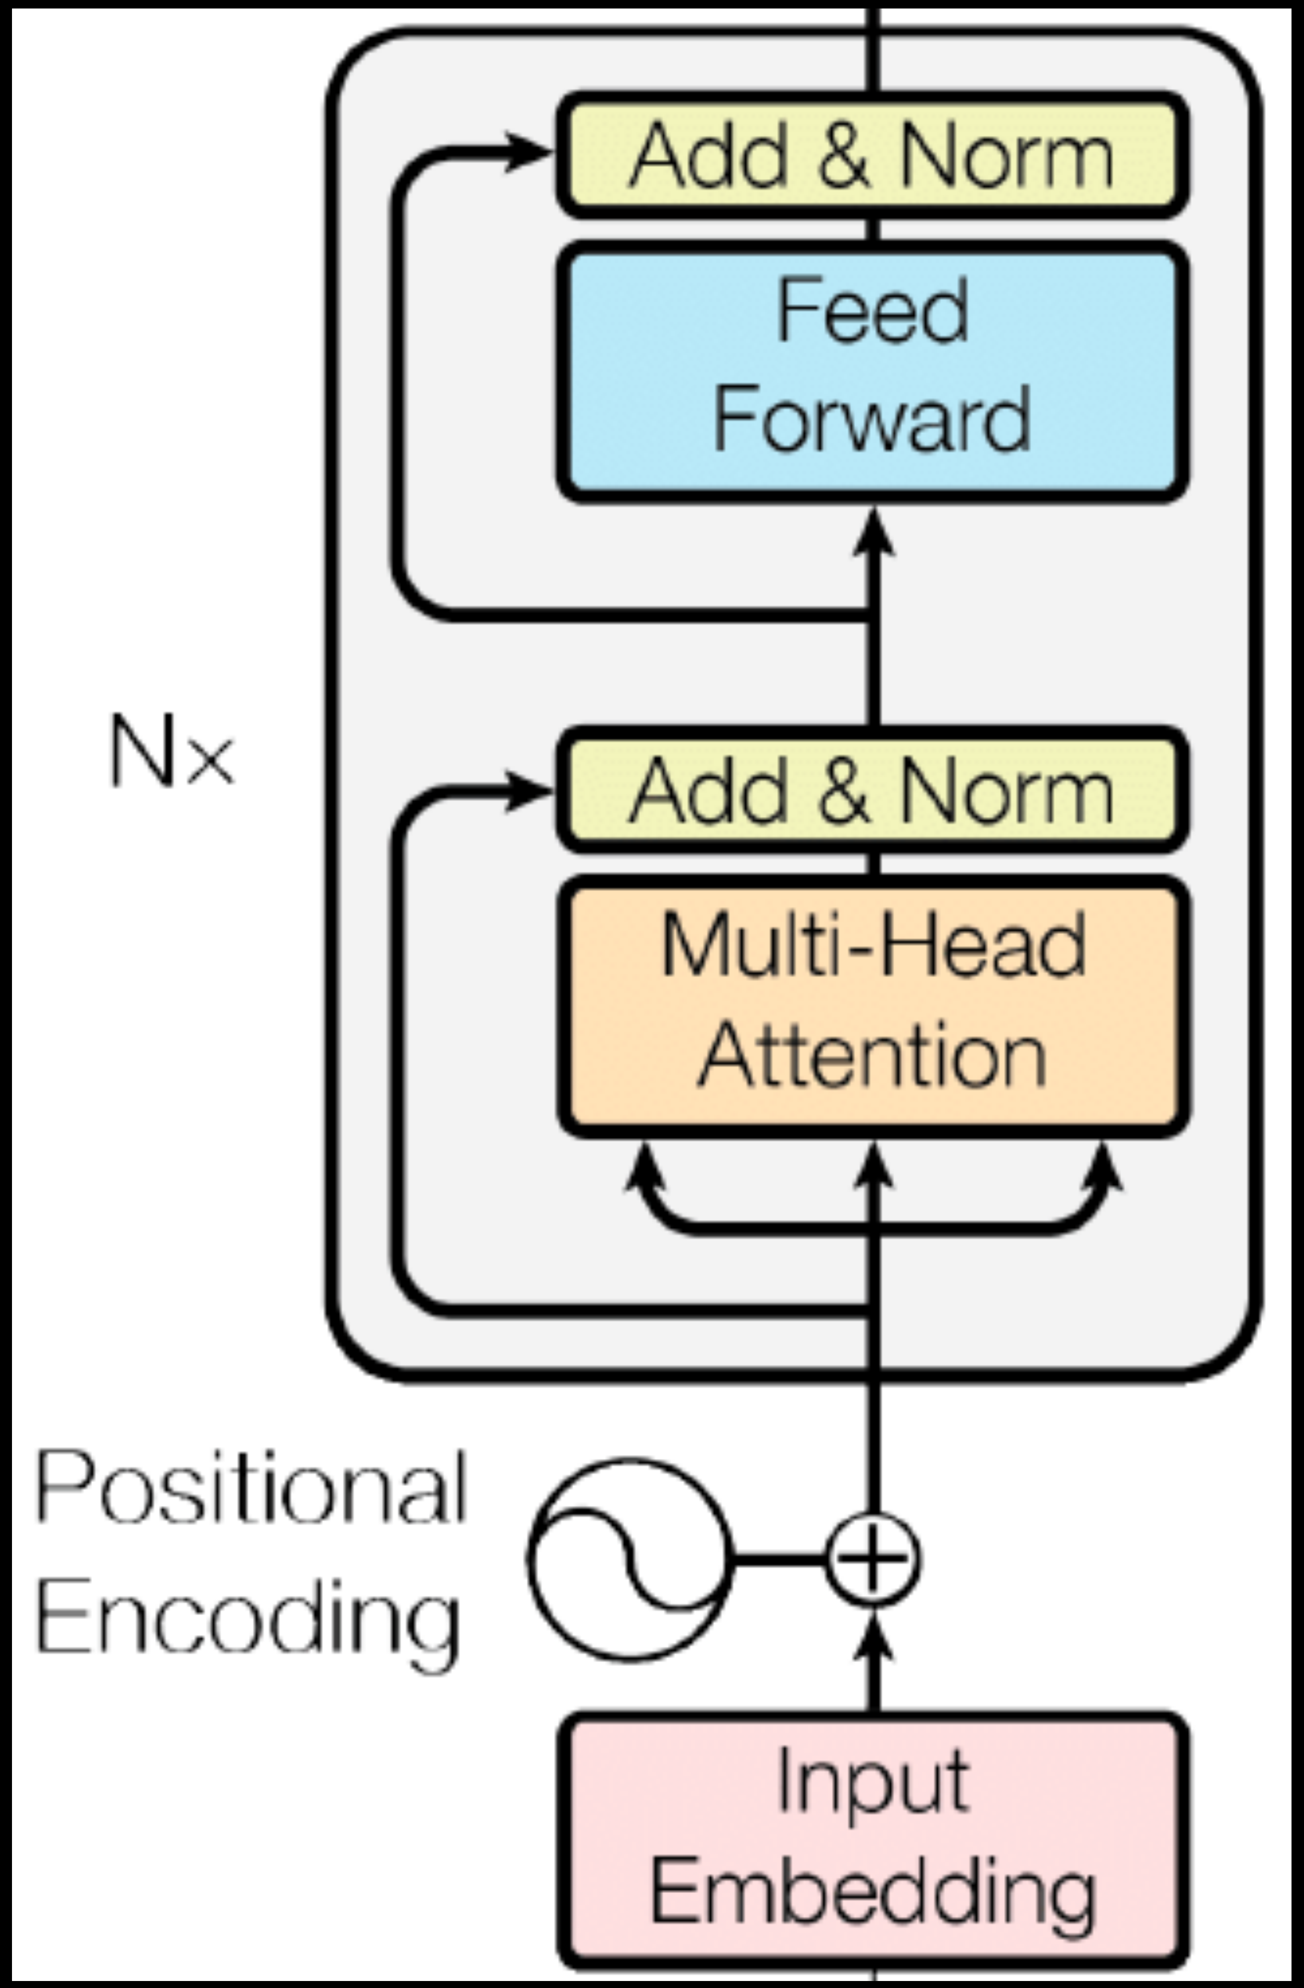

Compute cost: O(n^2)

Maximum path length: O(1)

CT N	1,600,000 tokens
ImageNet N	256 tokens

Prelim: Transformers and Self-attention

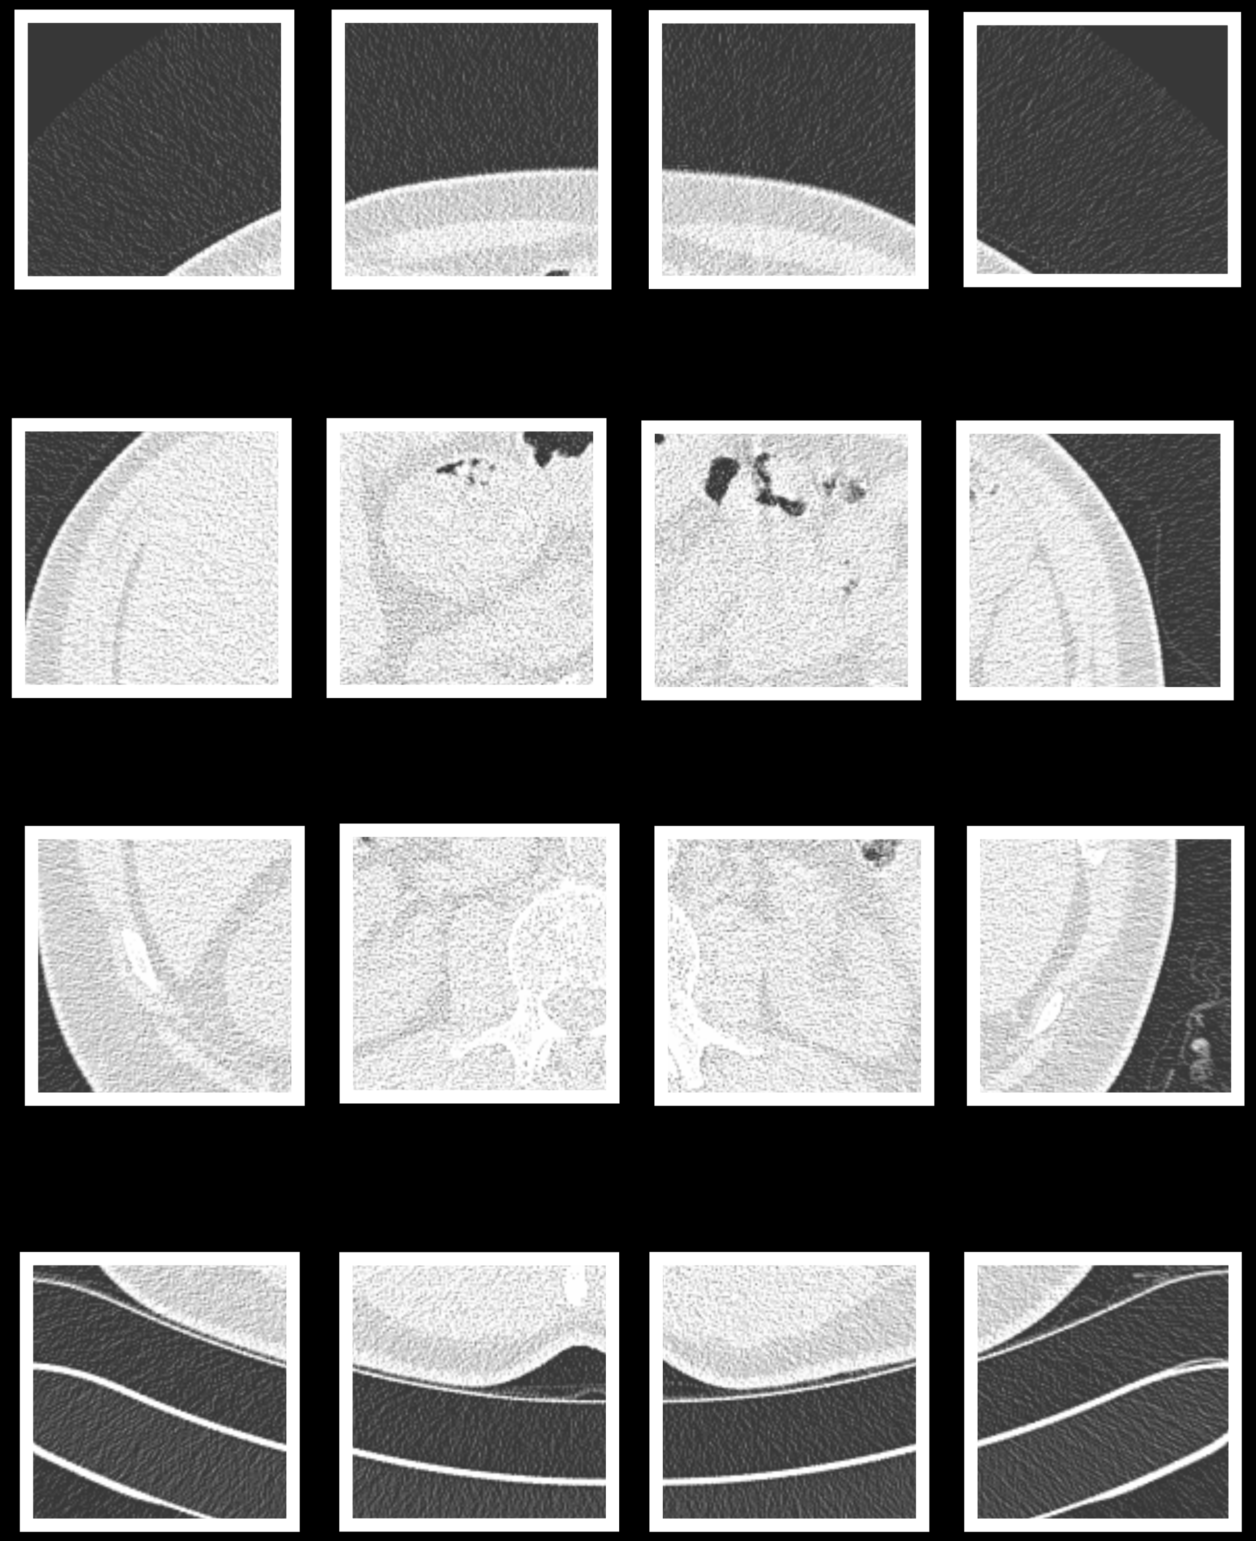

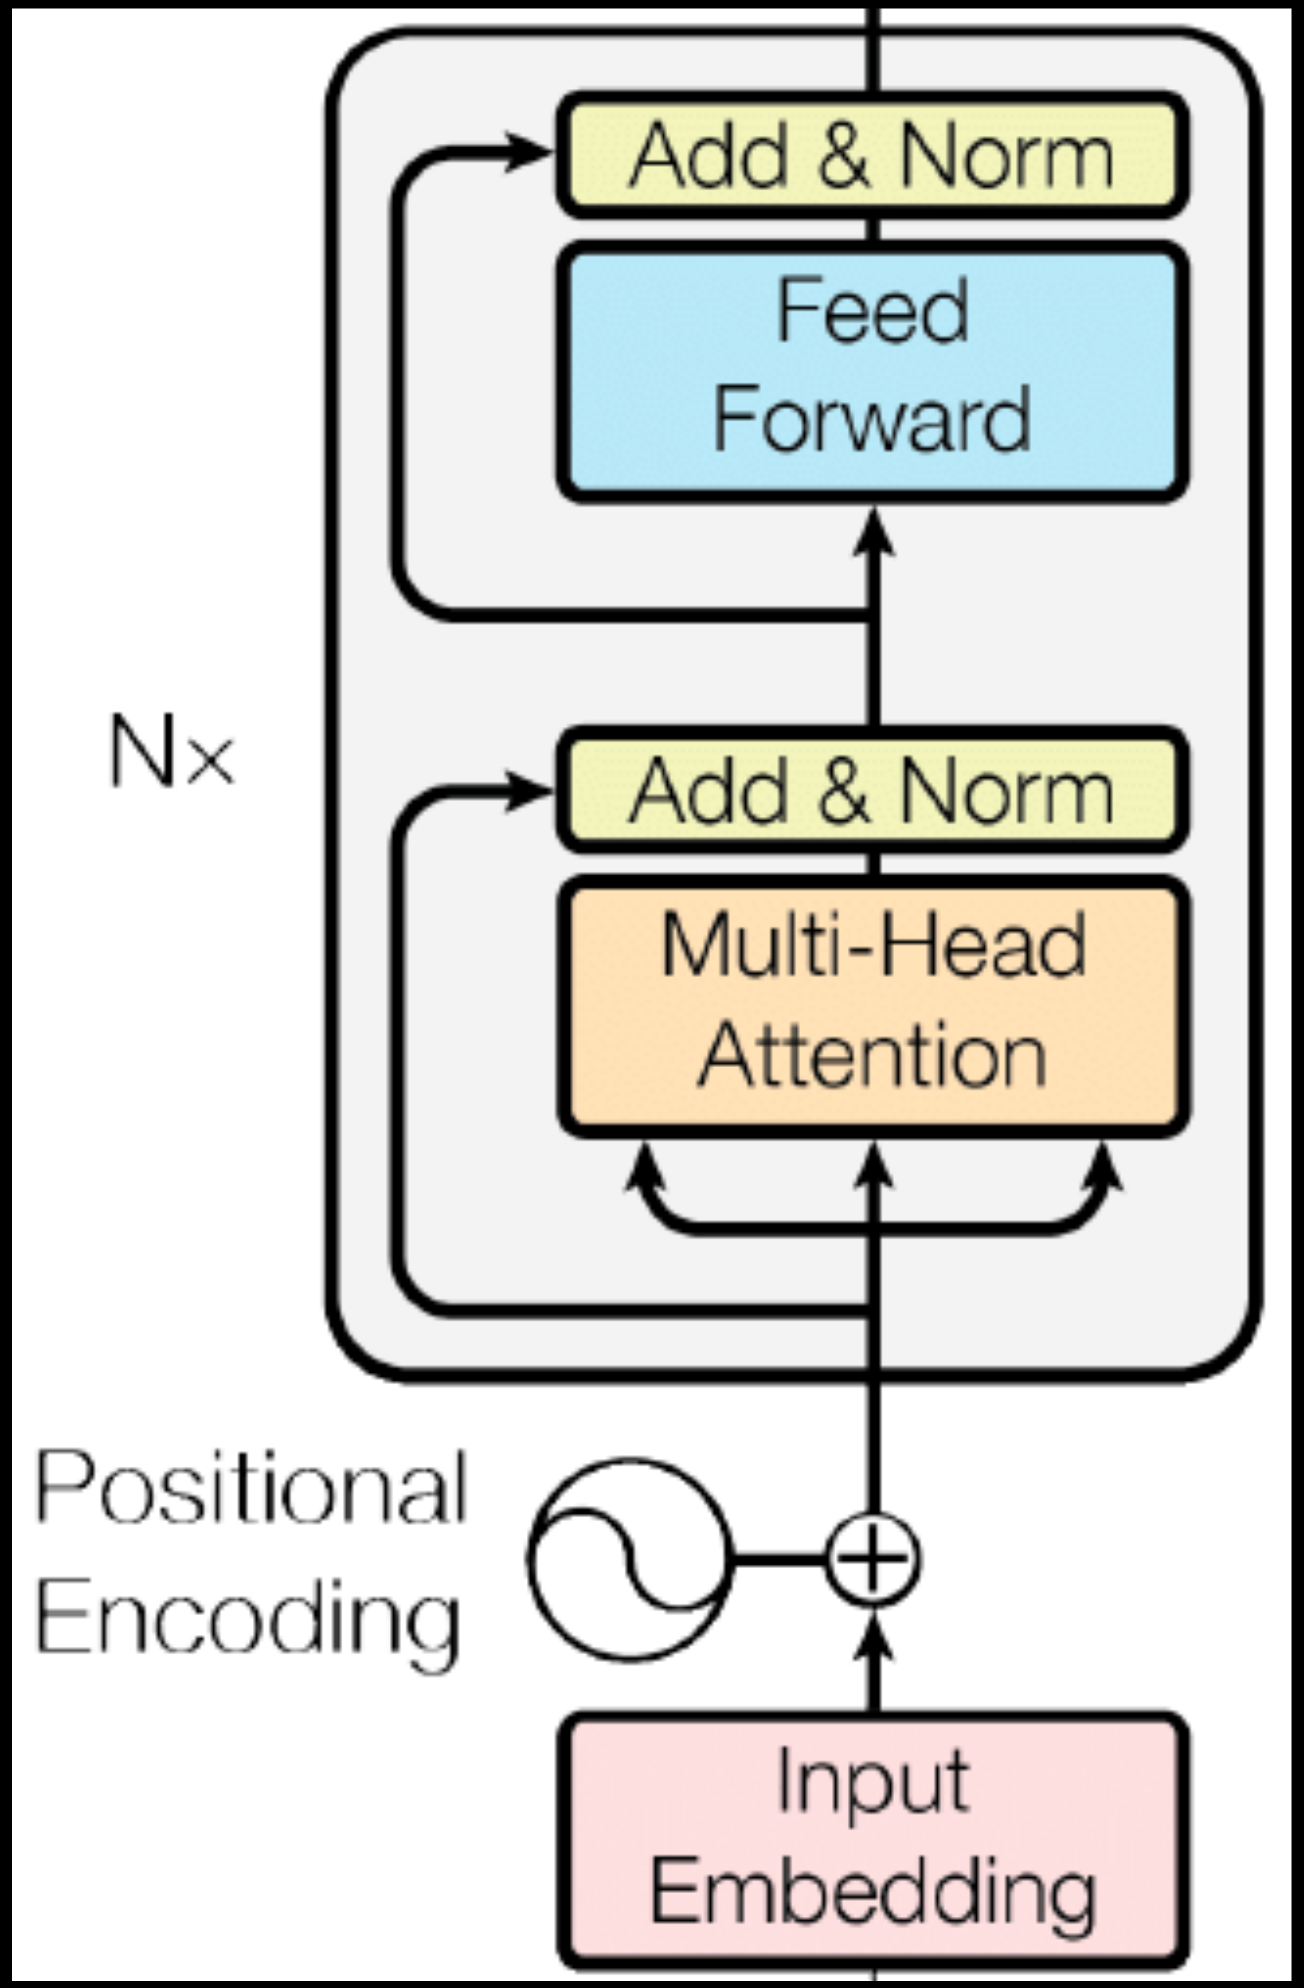

Compute cost: O(n^2)

Maximum path length: O(1)

Problem: Intractable at our scale!

62,000 x bigger than ImageNet
3.9 billion times more compute expensive

CT N 1,600,000 tokens
ImageNet N 256 tokens

Atlas : Overview

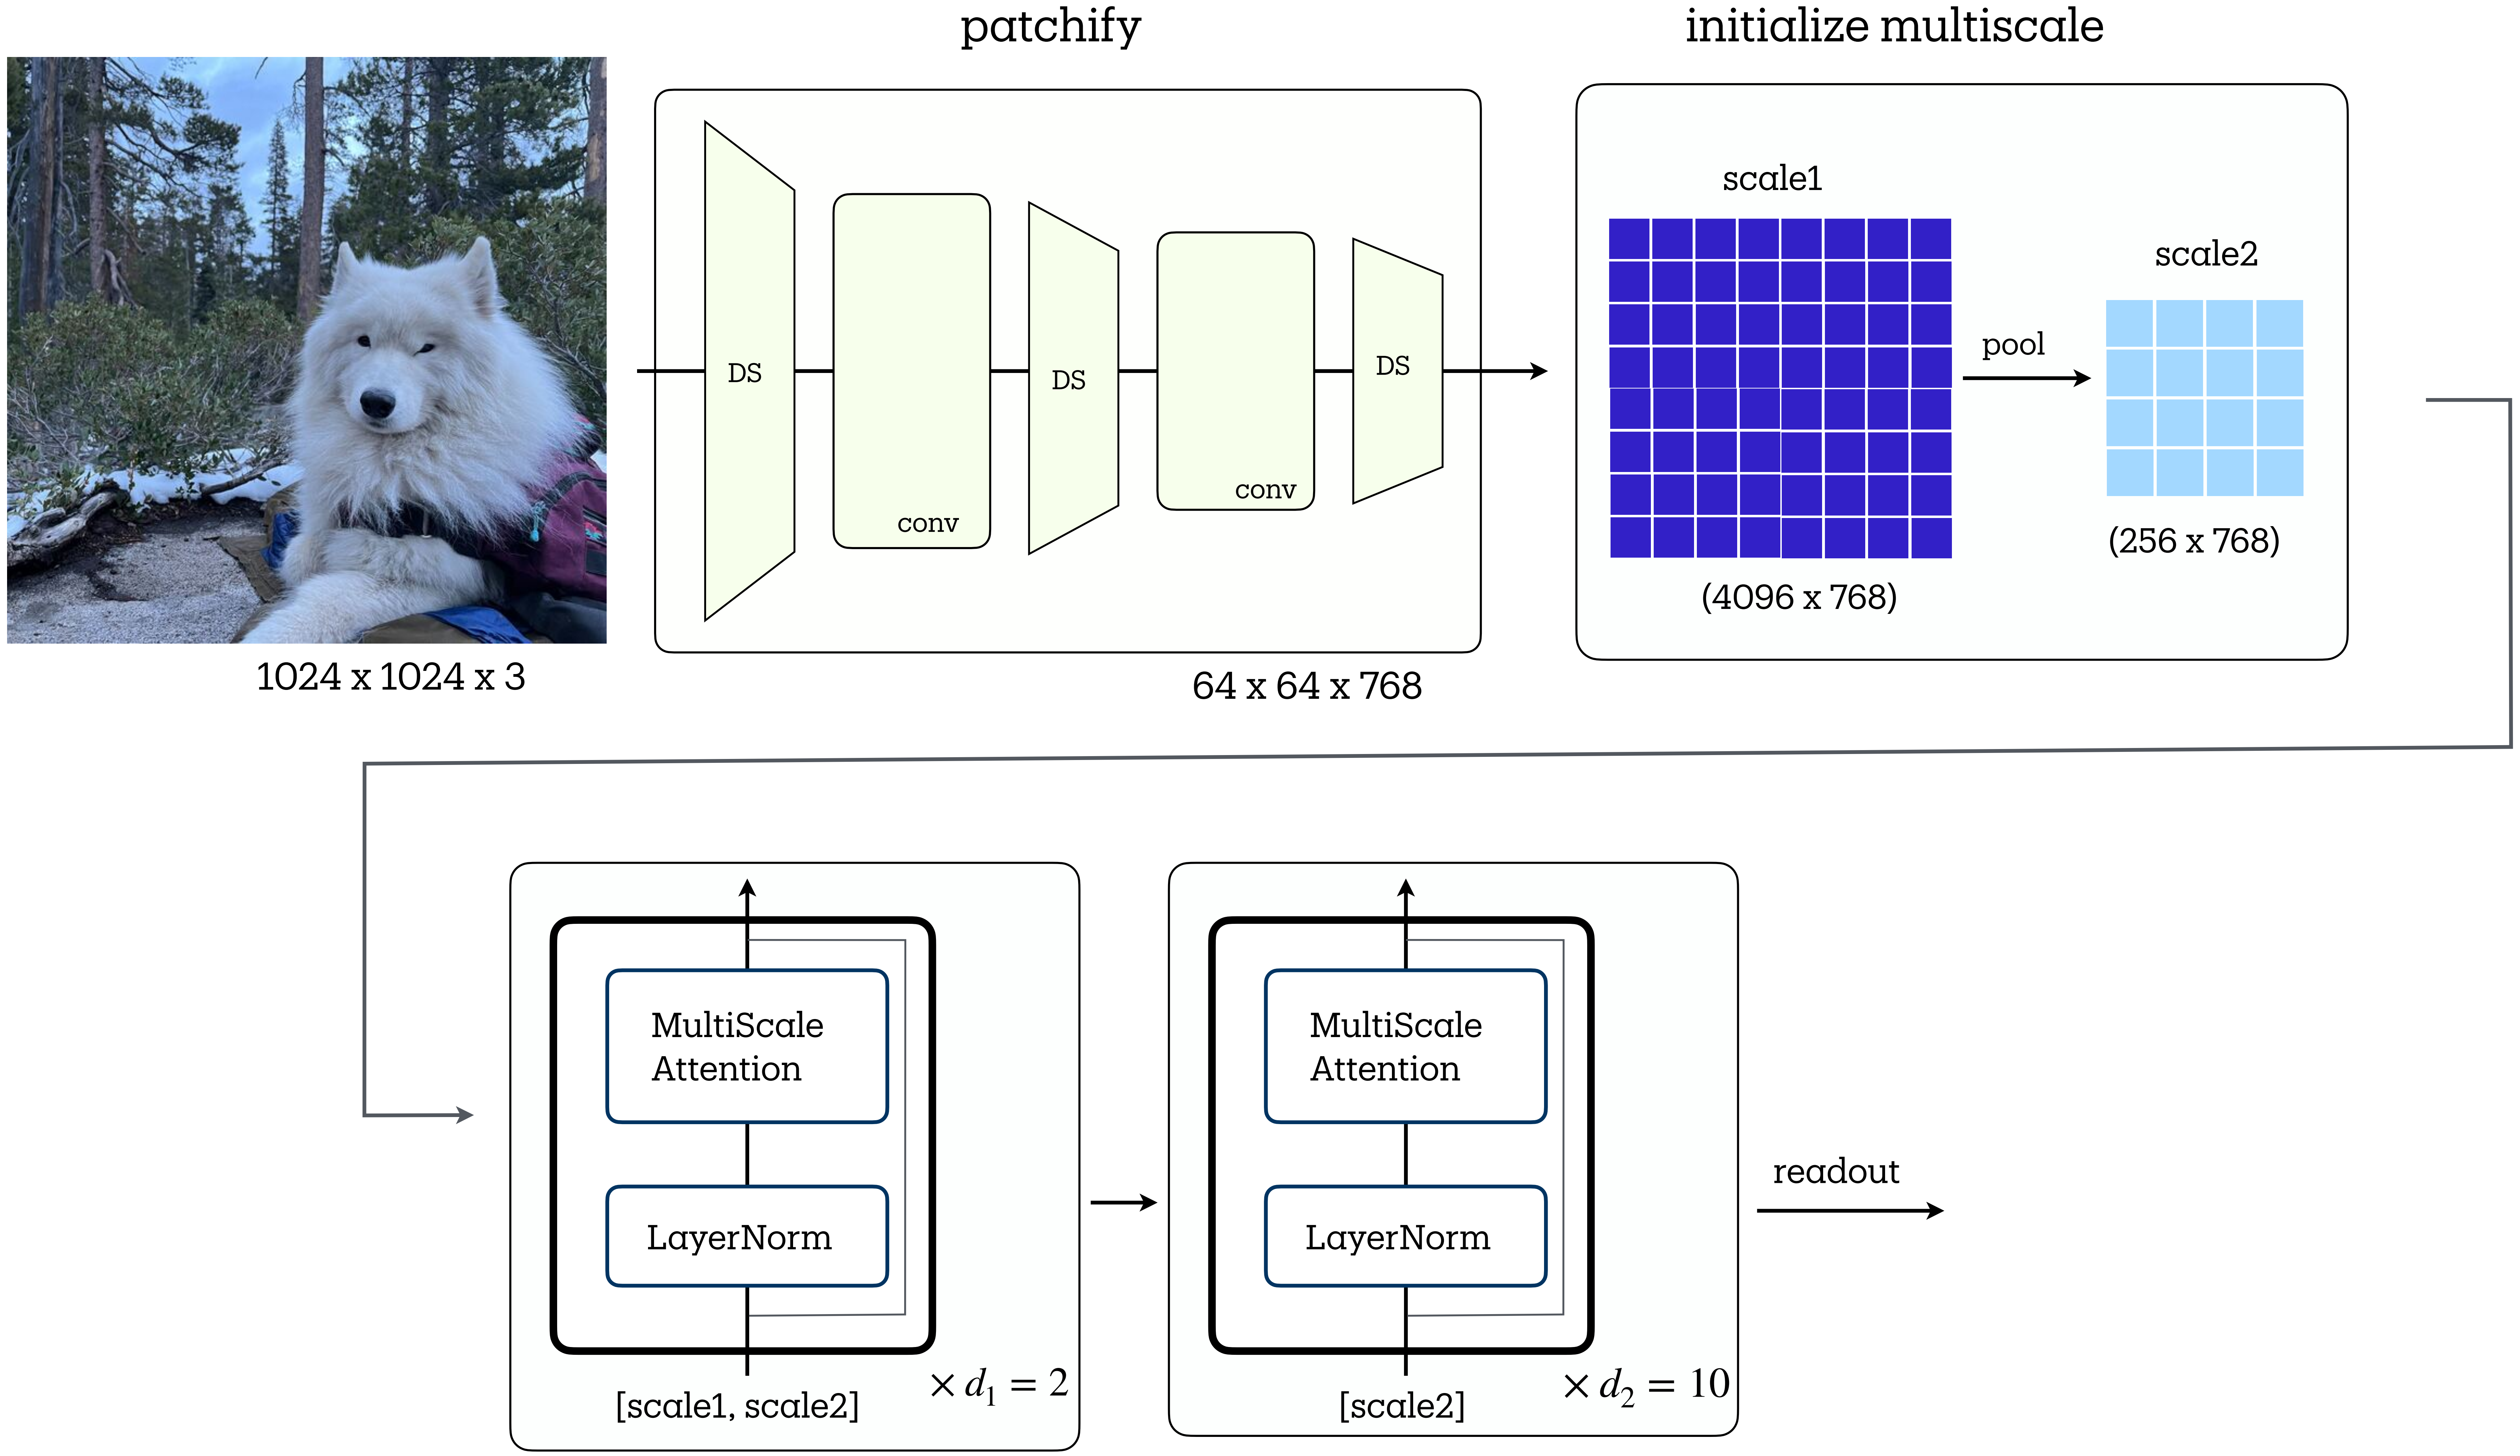

Computational Efficiency

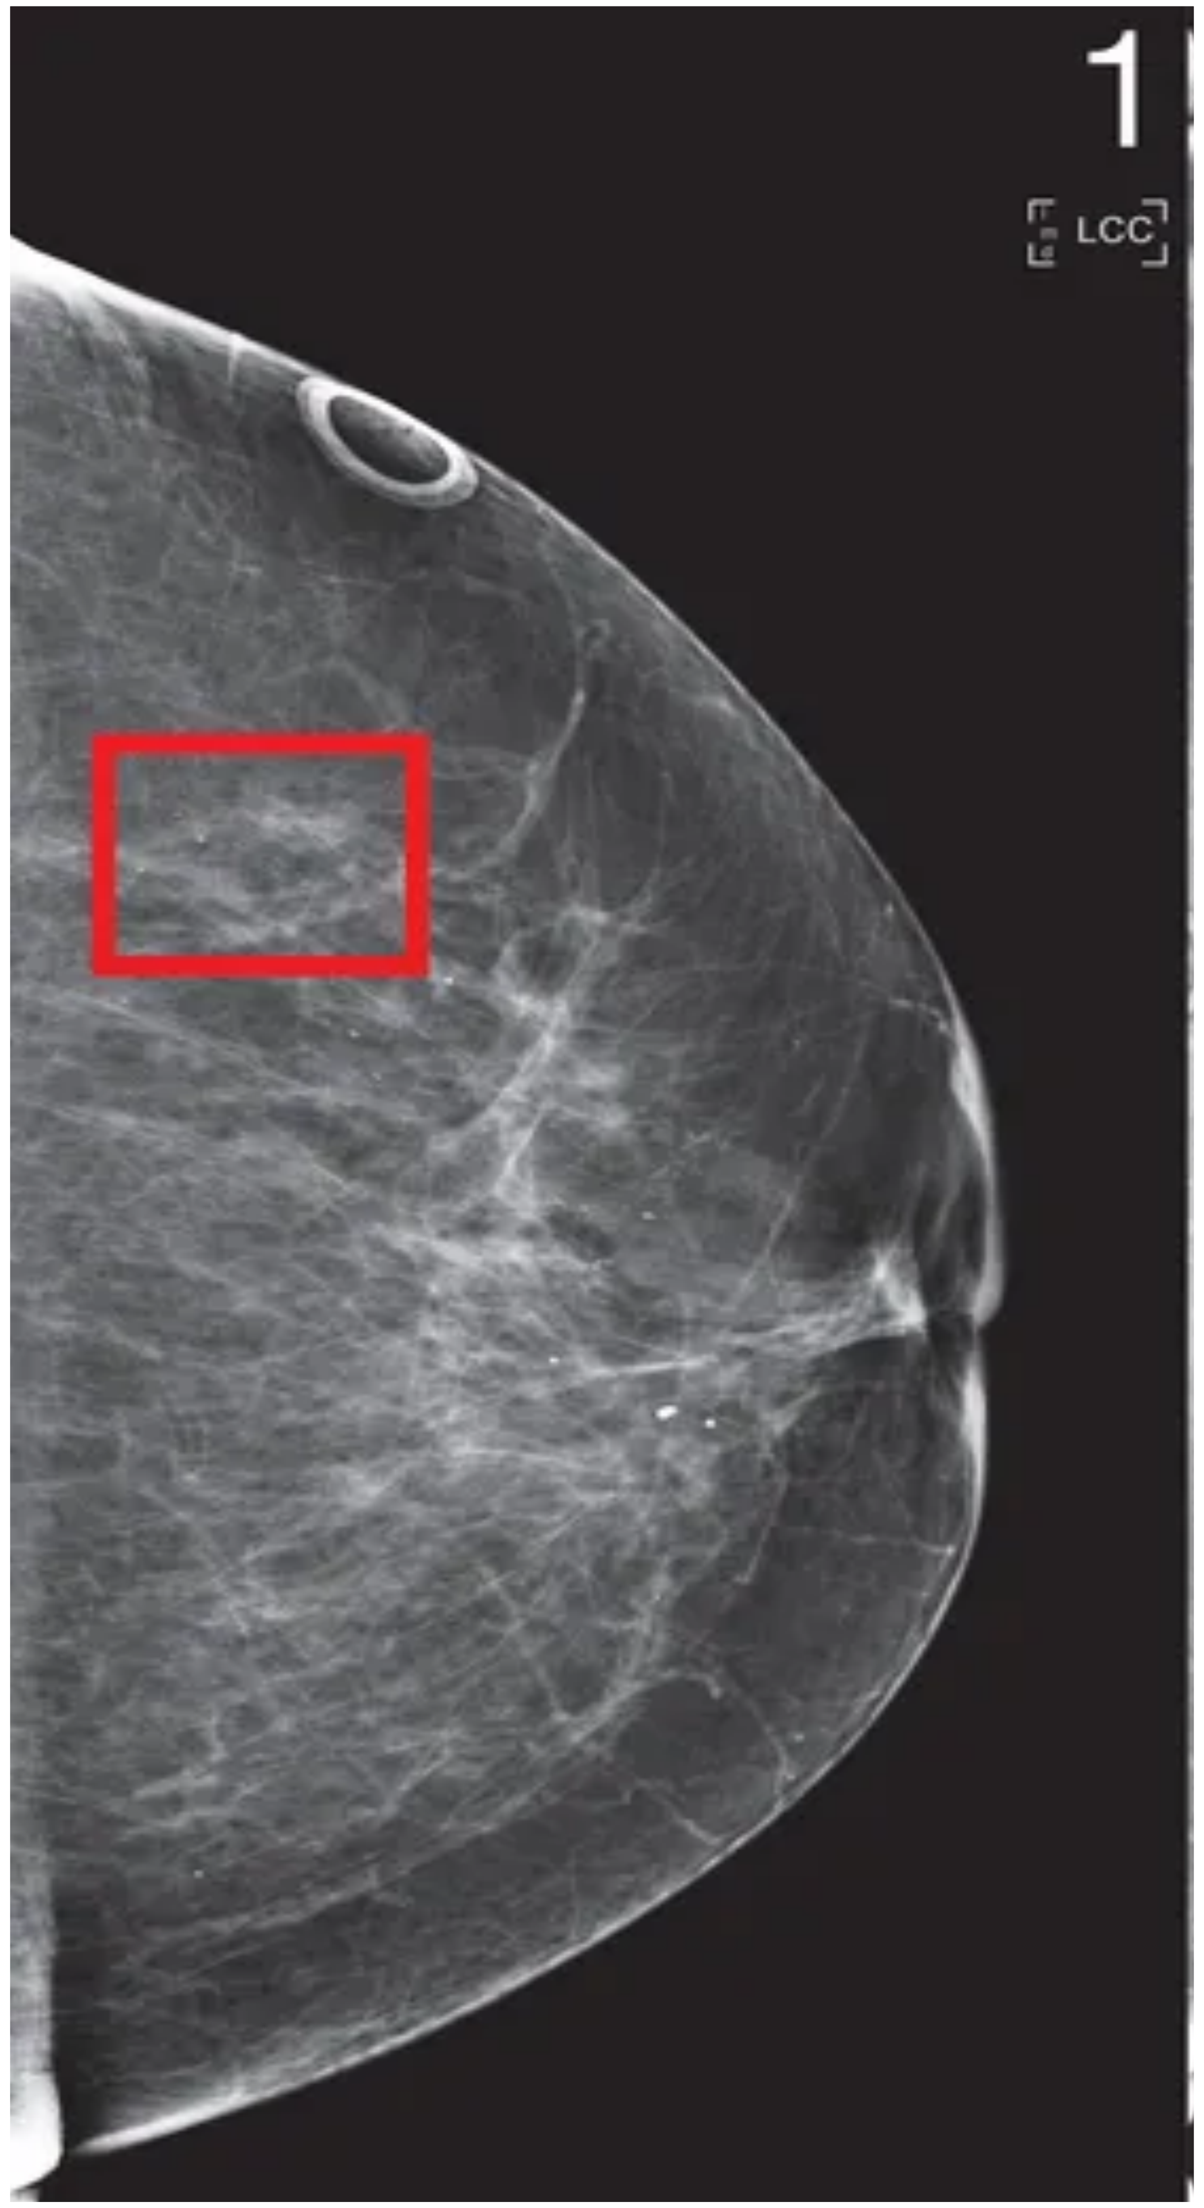

$2560 \times 3328 \sim 8.5M$ pixels

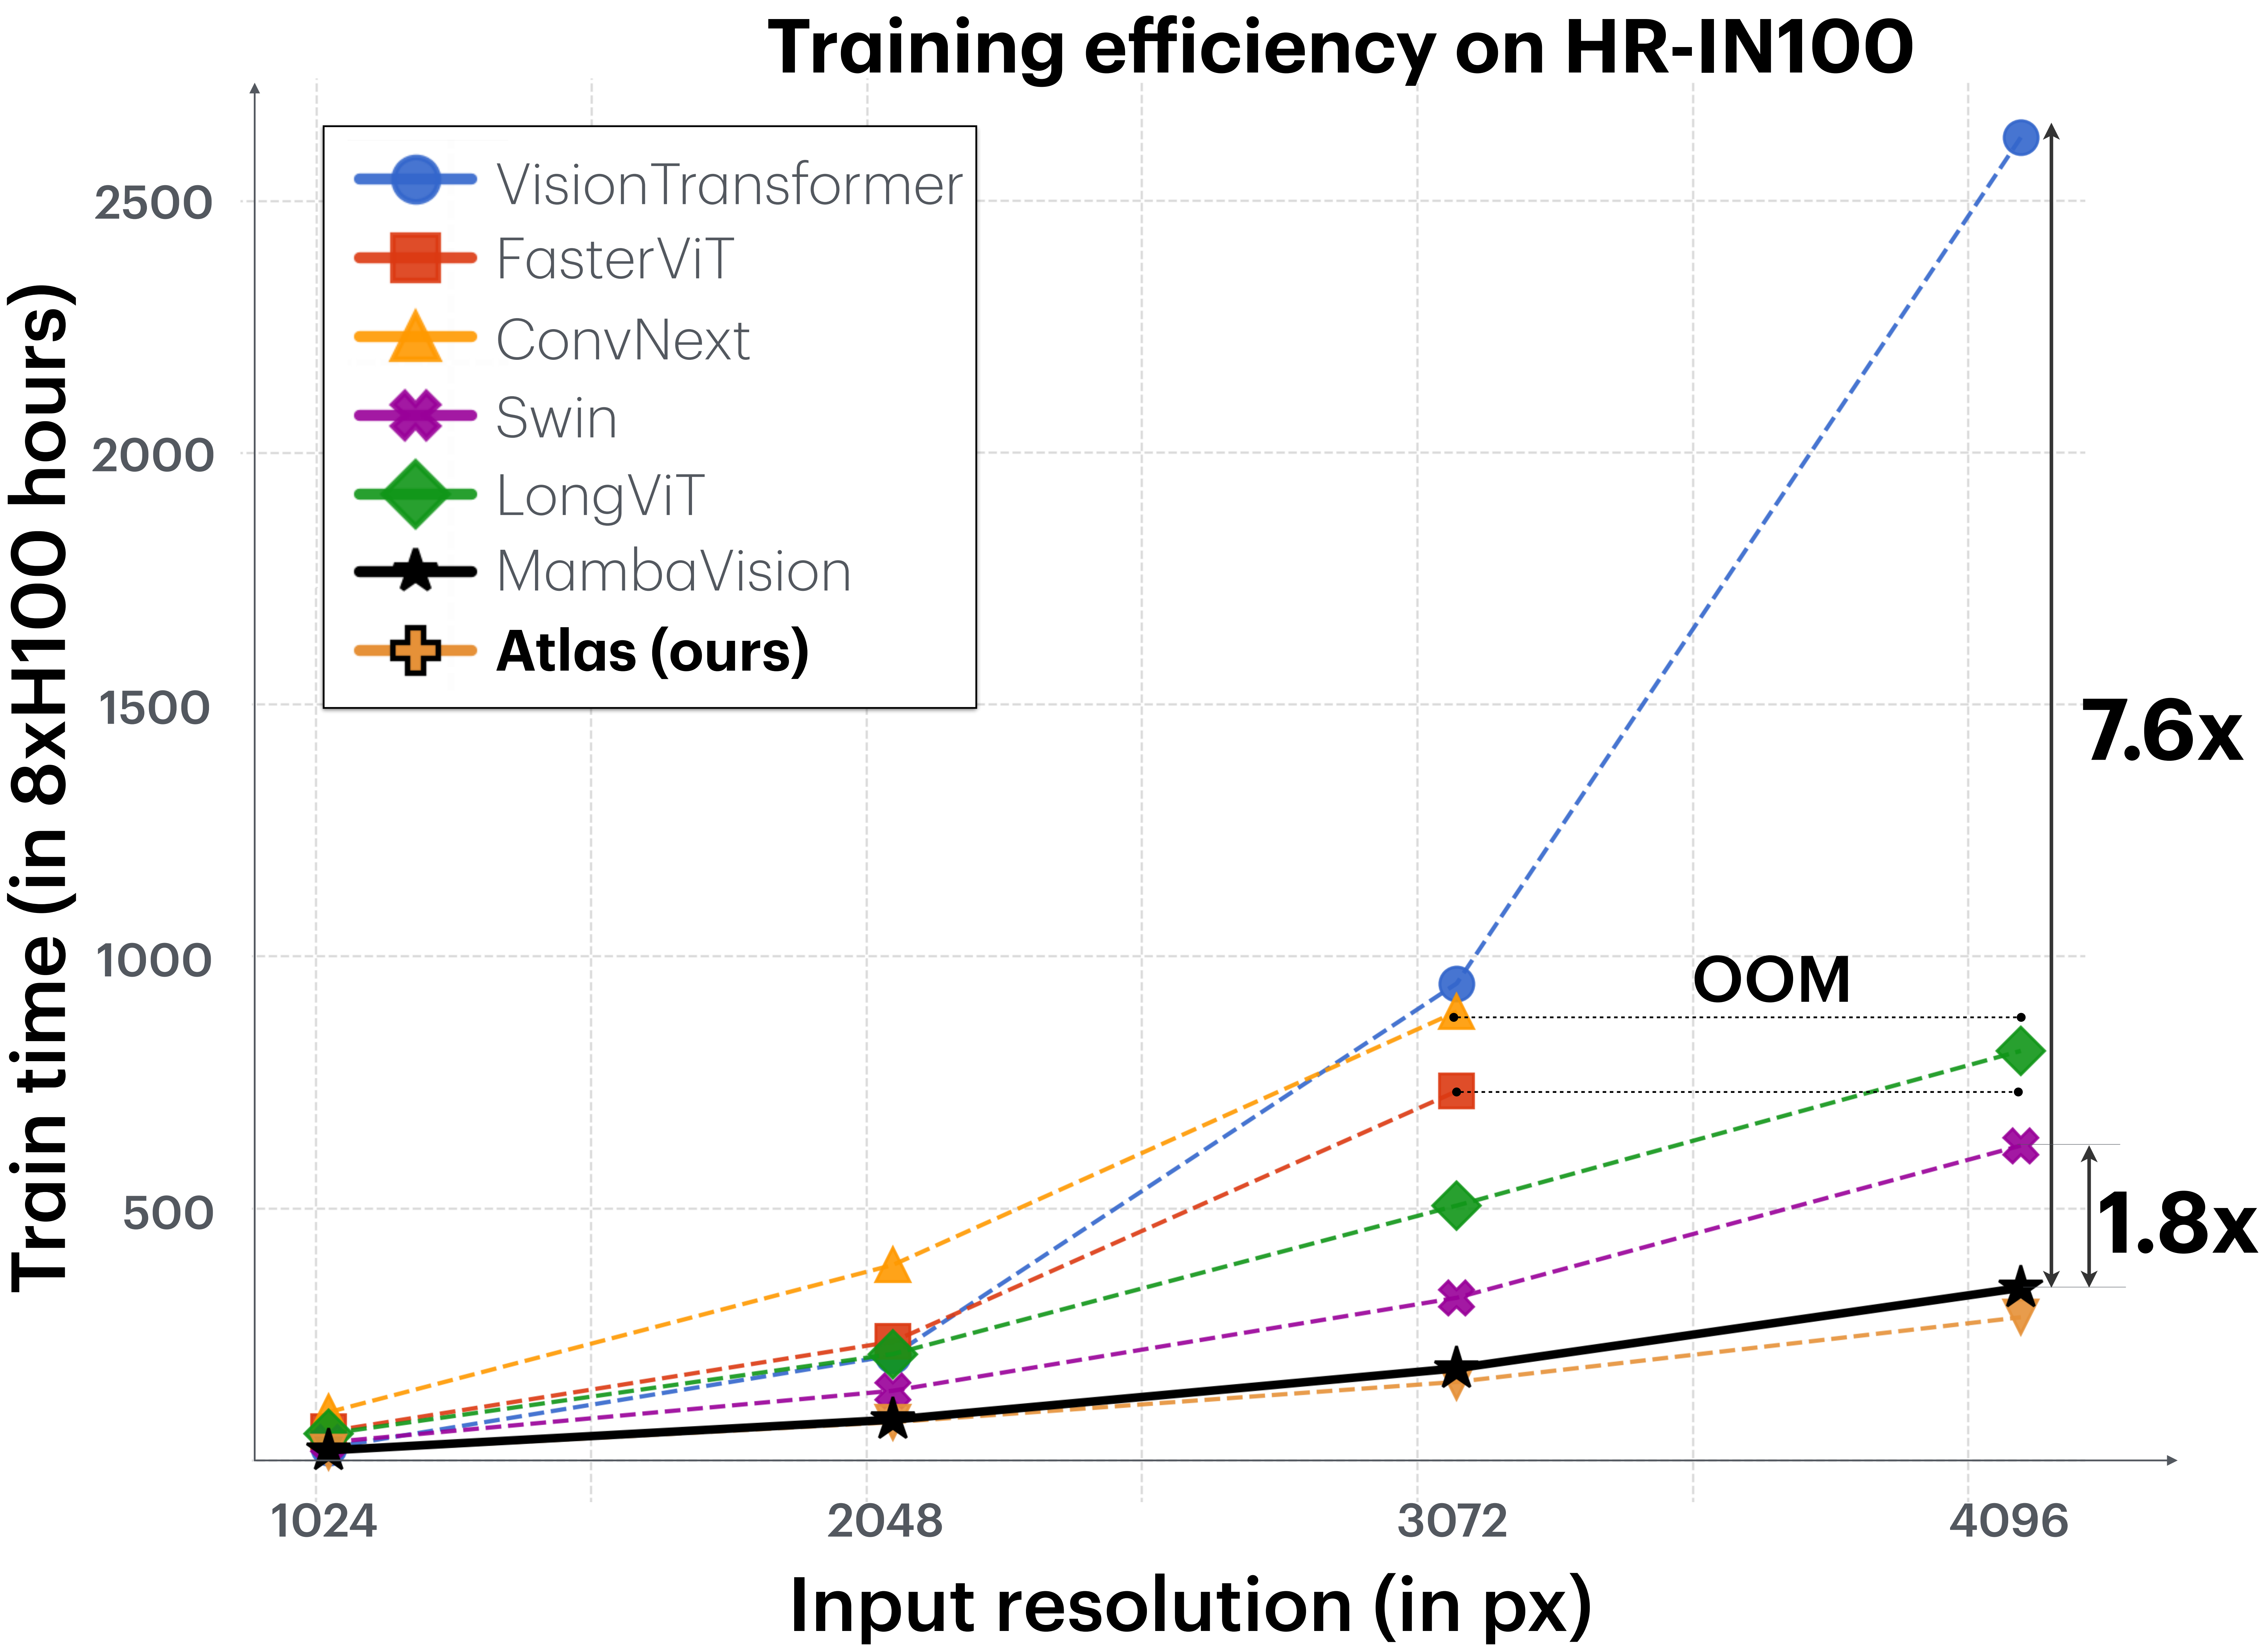

Atlas

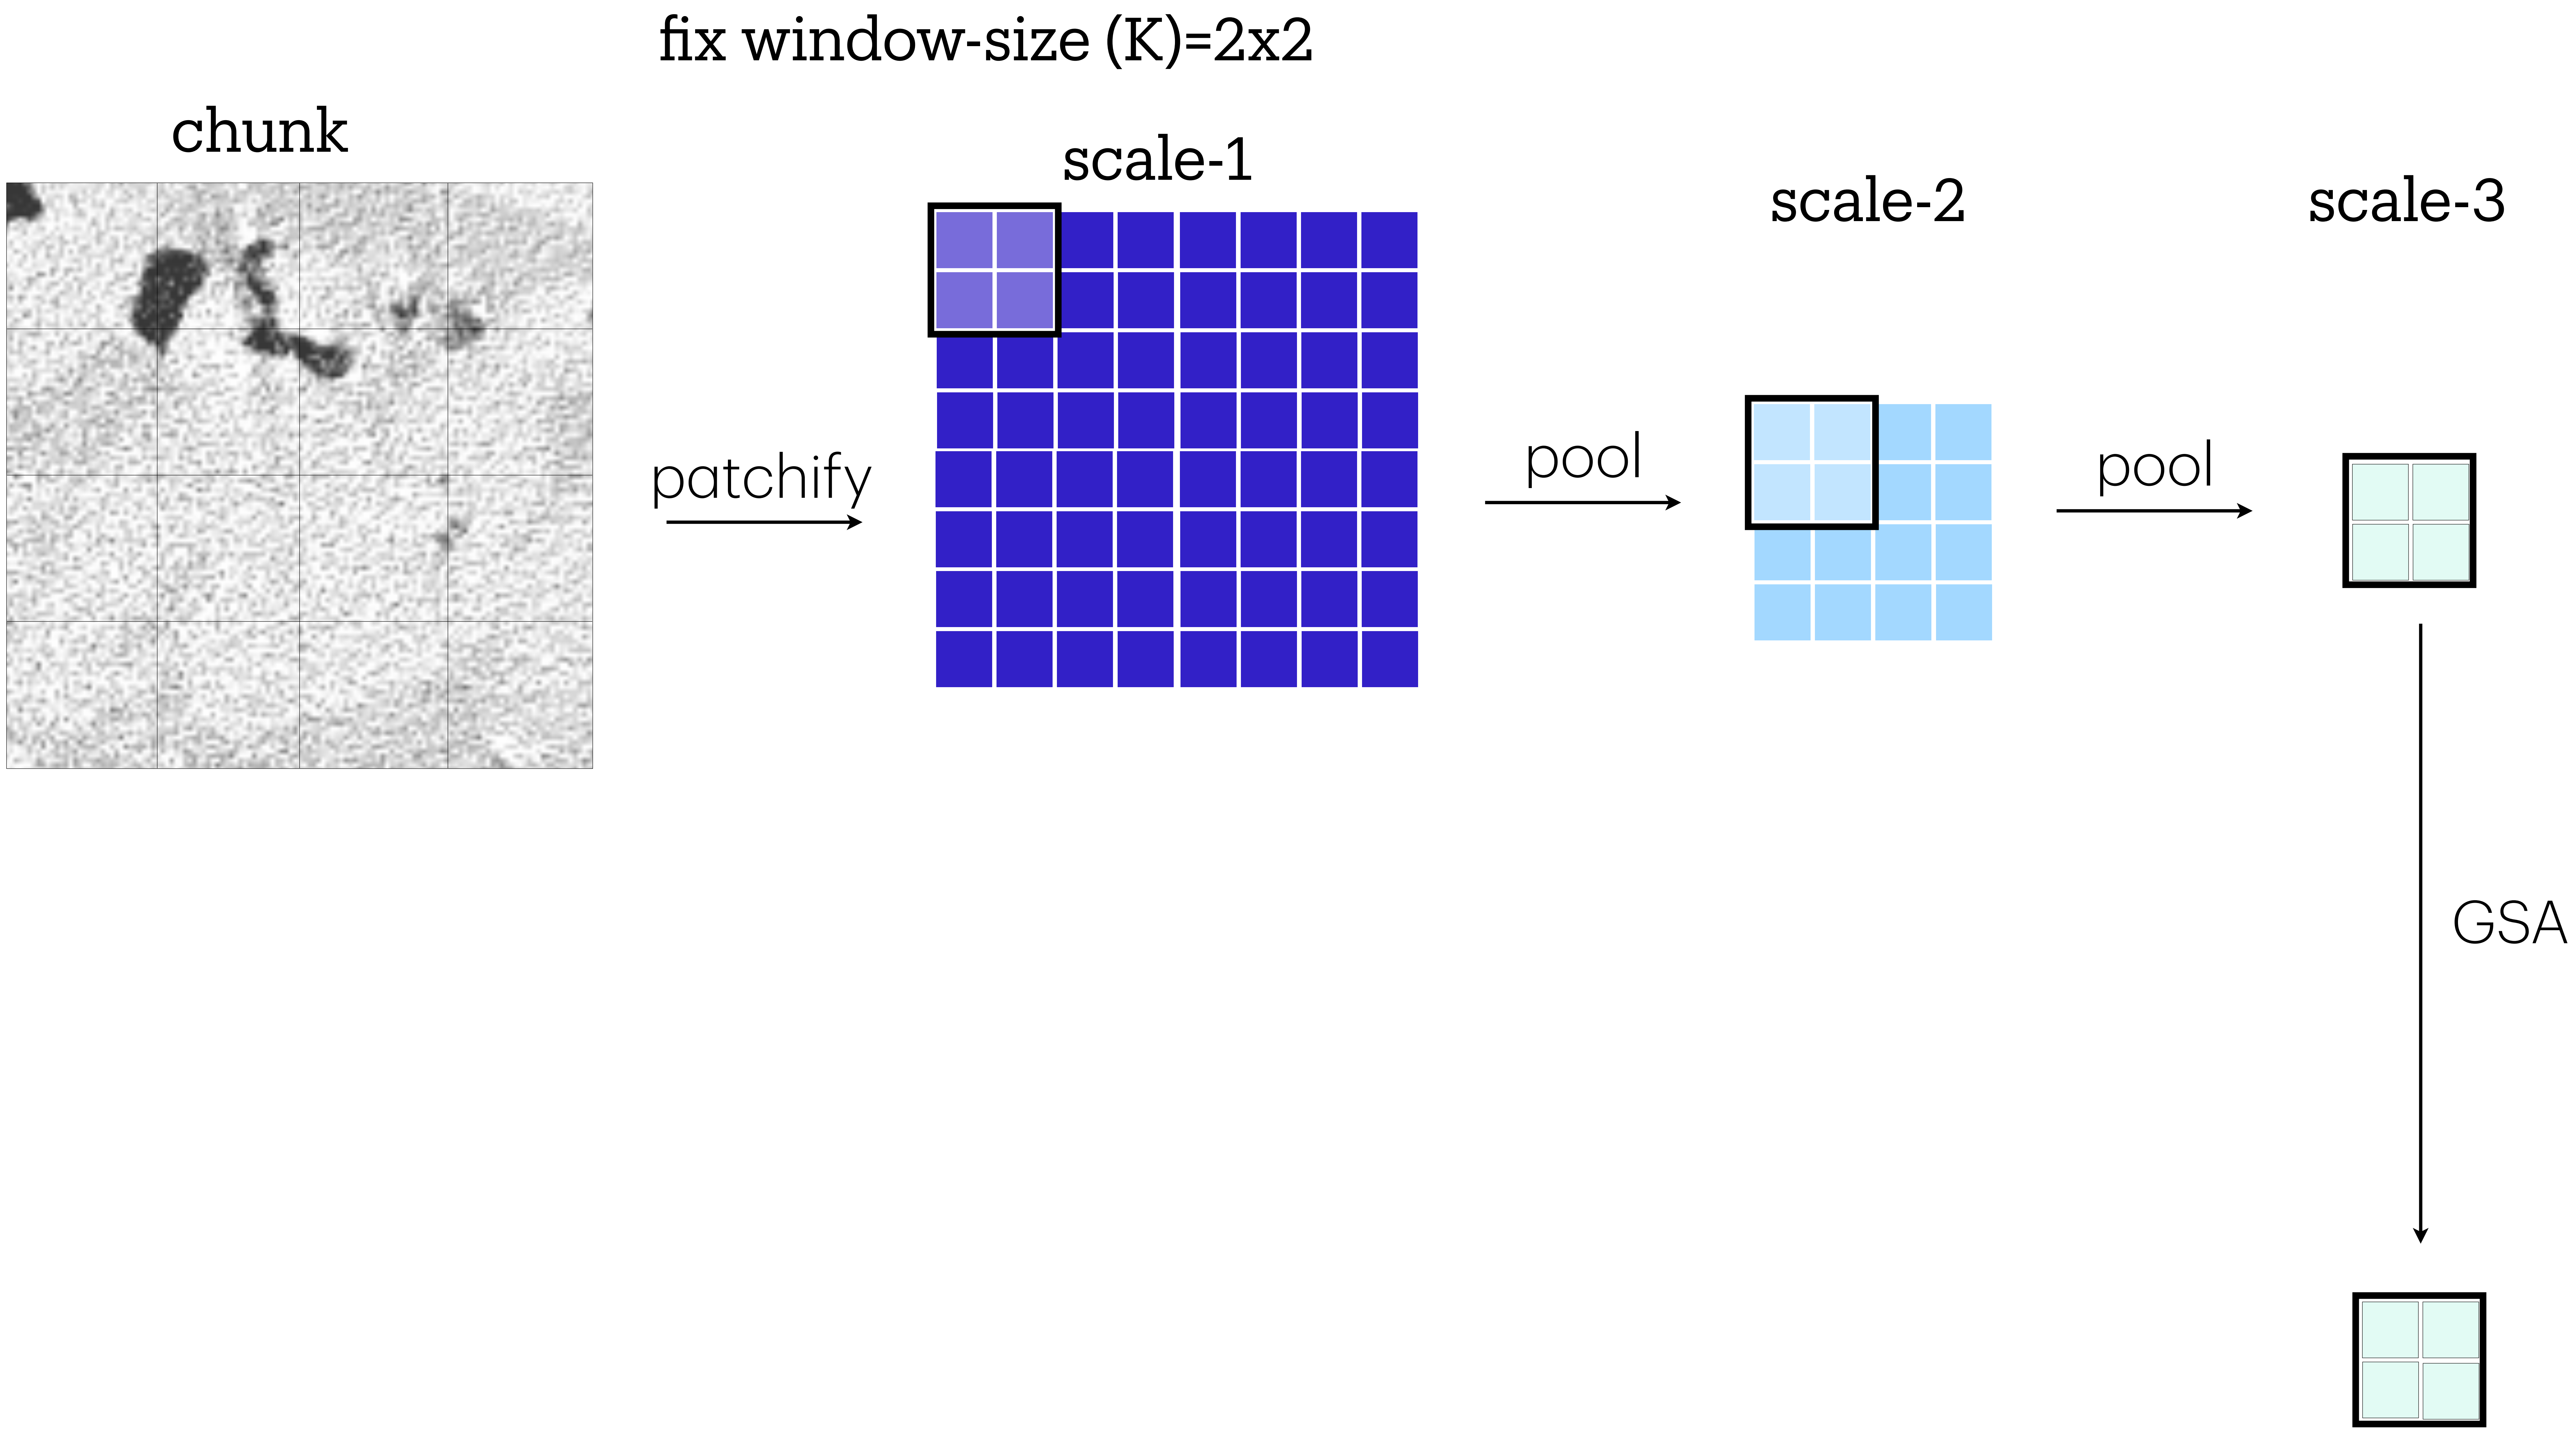

fix window-size (K)=2x2
chunk
patchify
scale-1
pool
scale-2
pool
scale-3
GSA

Atlas with Multi-Scale Attention

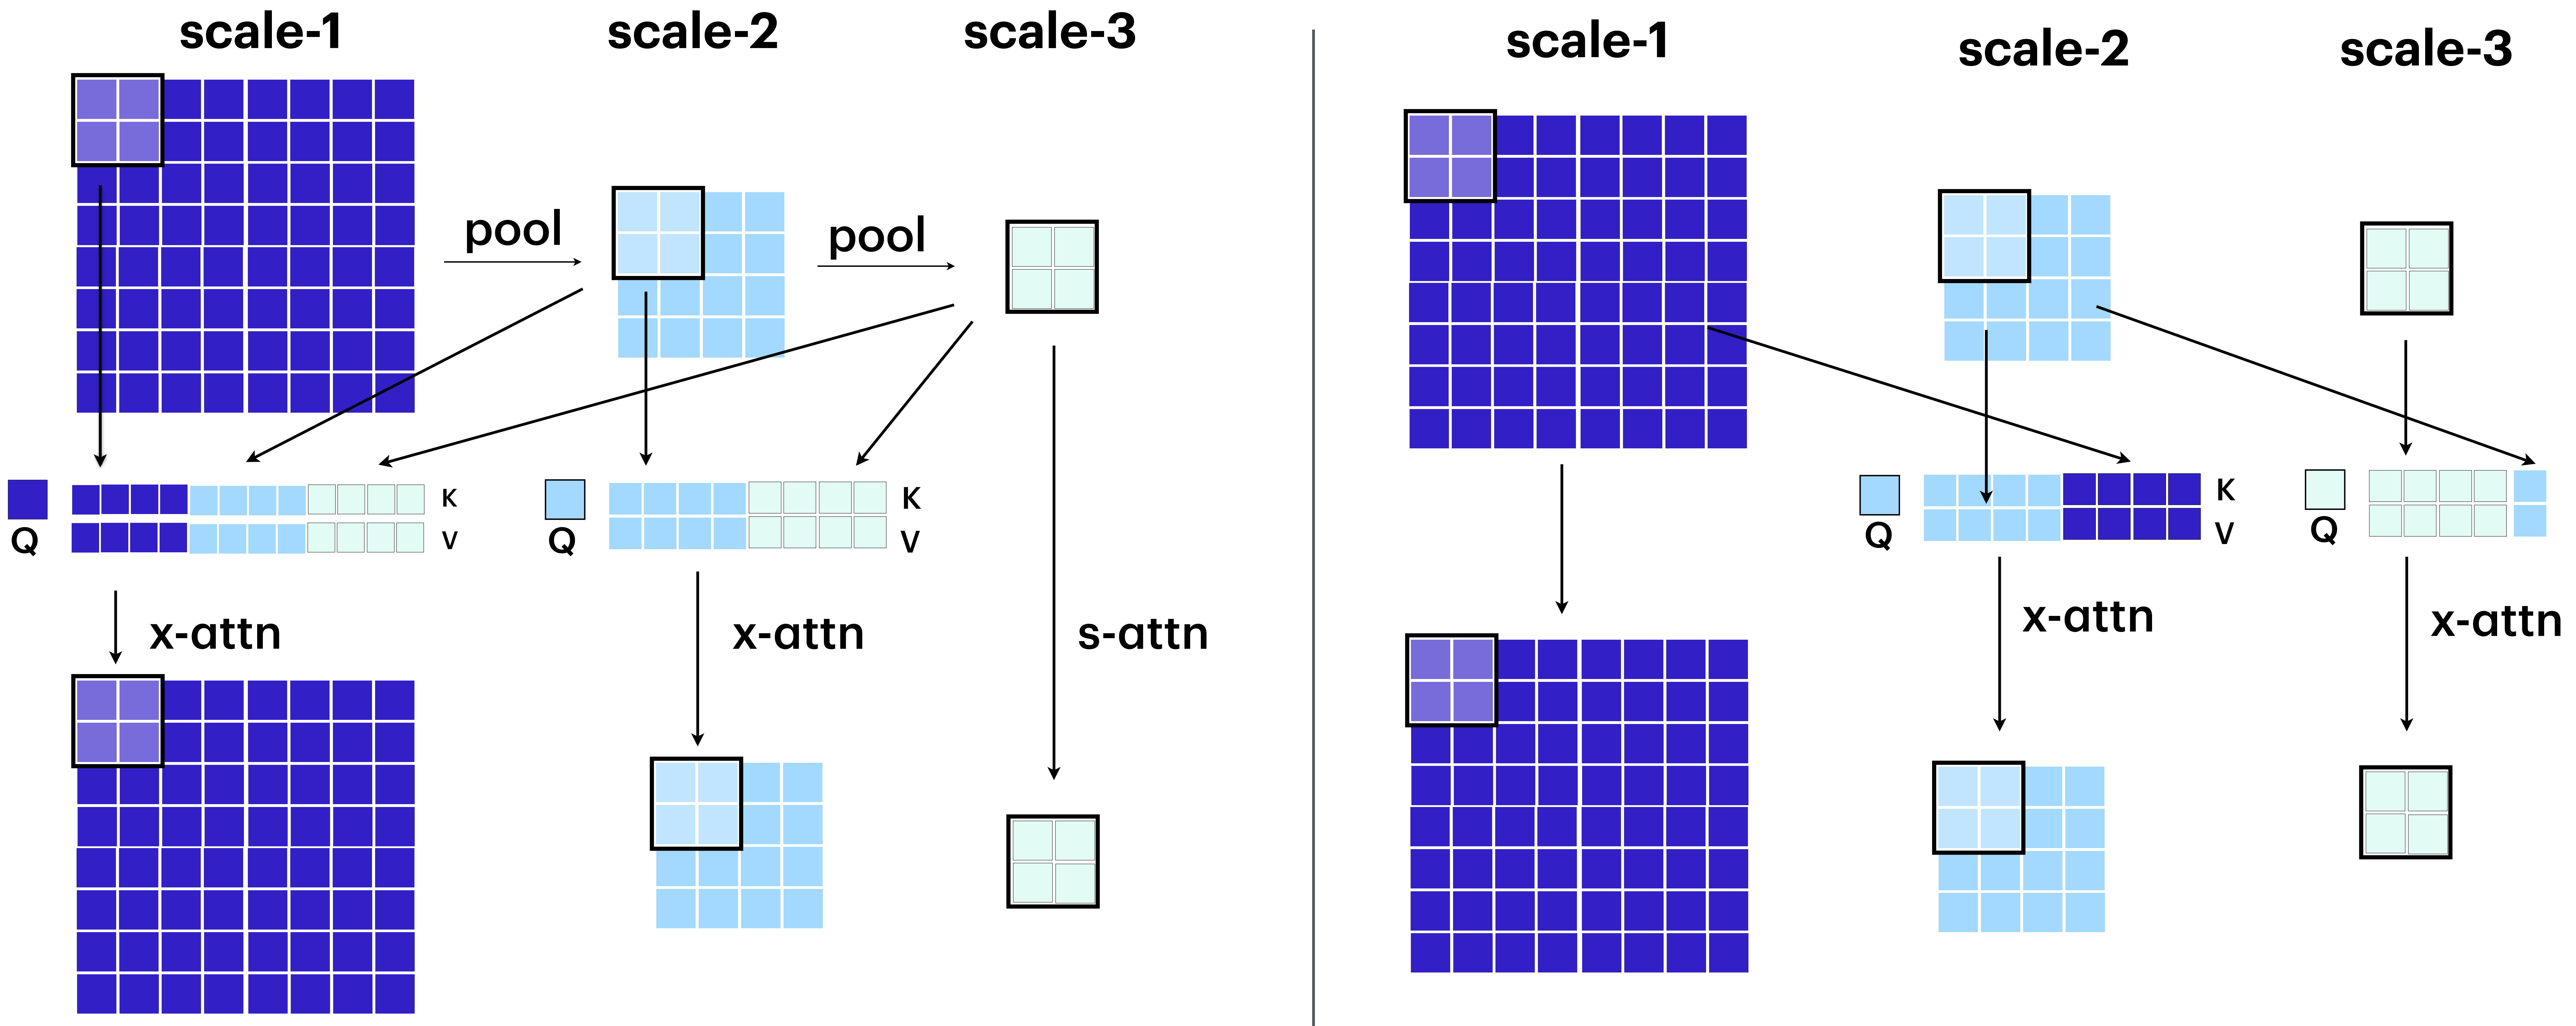

top-down **multi-scale communication pathways**

bottom-up **multi-scale communication pathways**

Atlas Architecture

key idea: drop early scales progressively

configuration : dX-dY

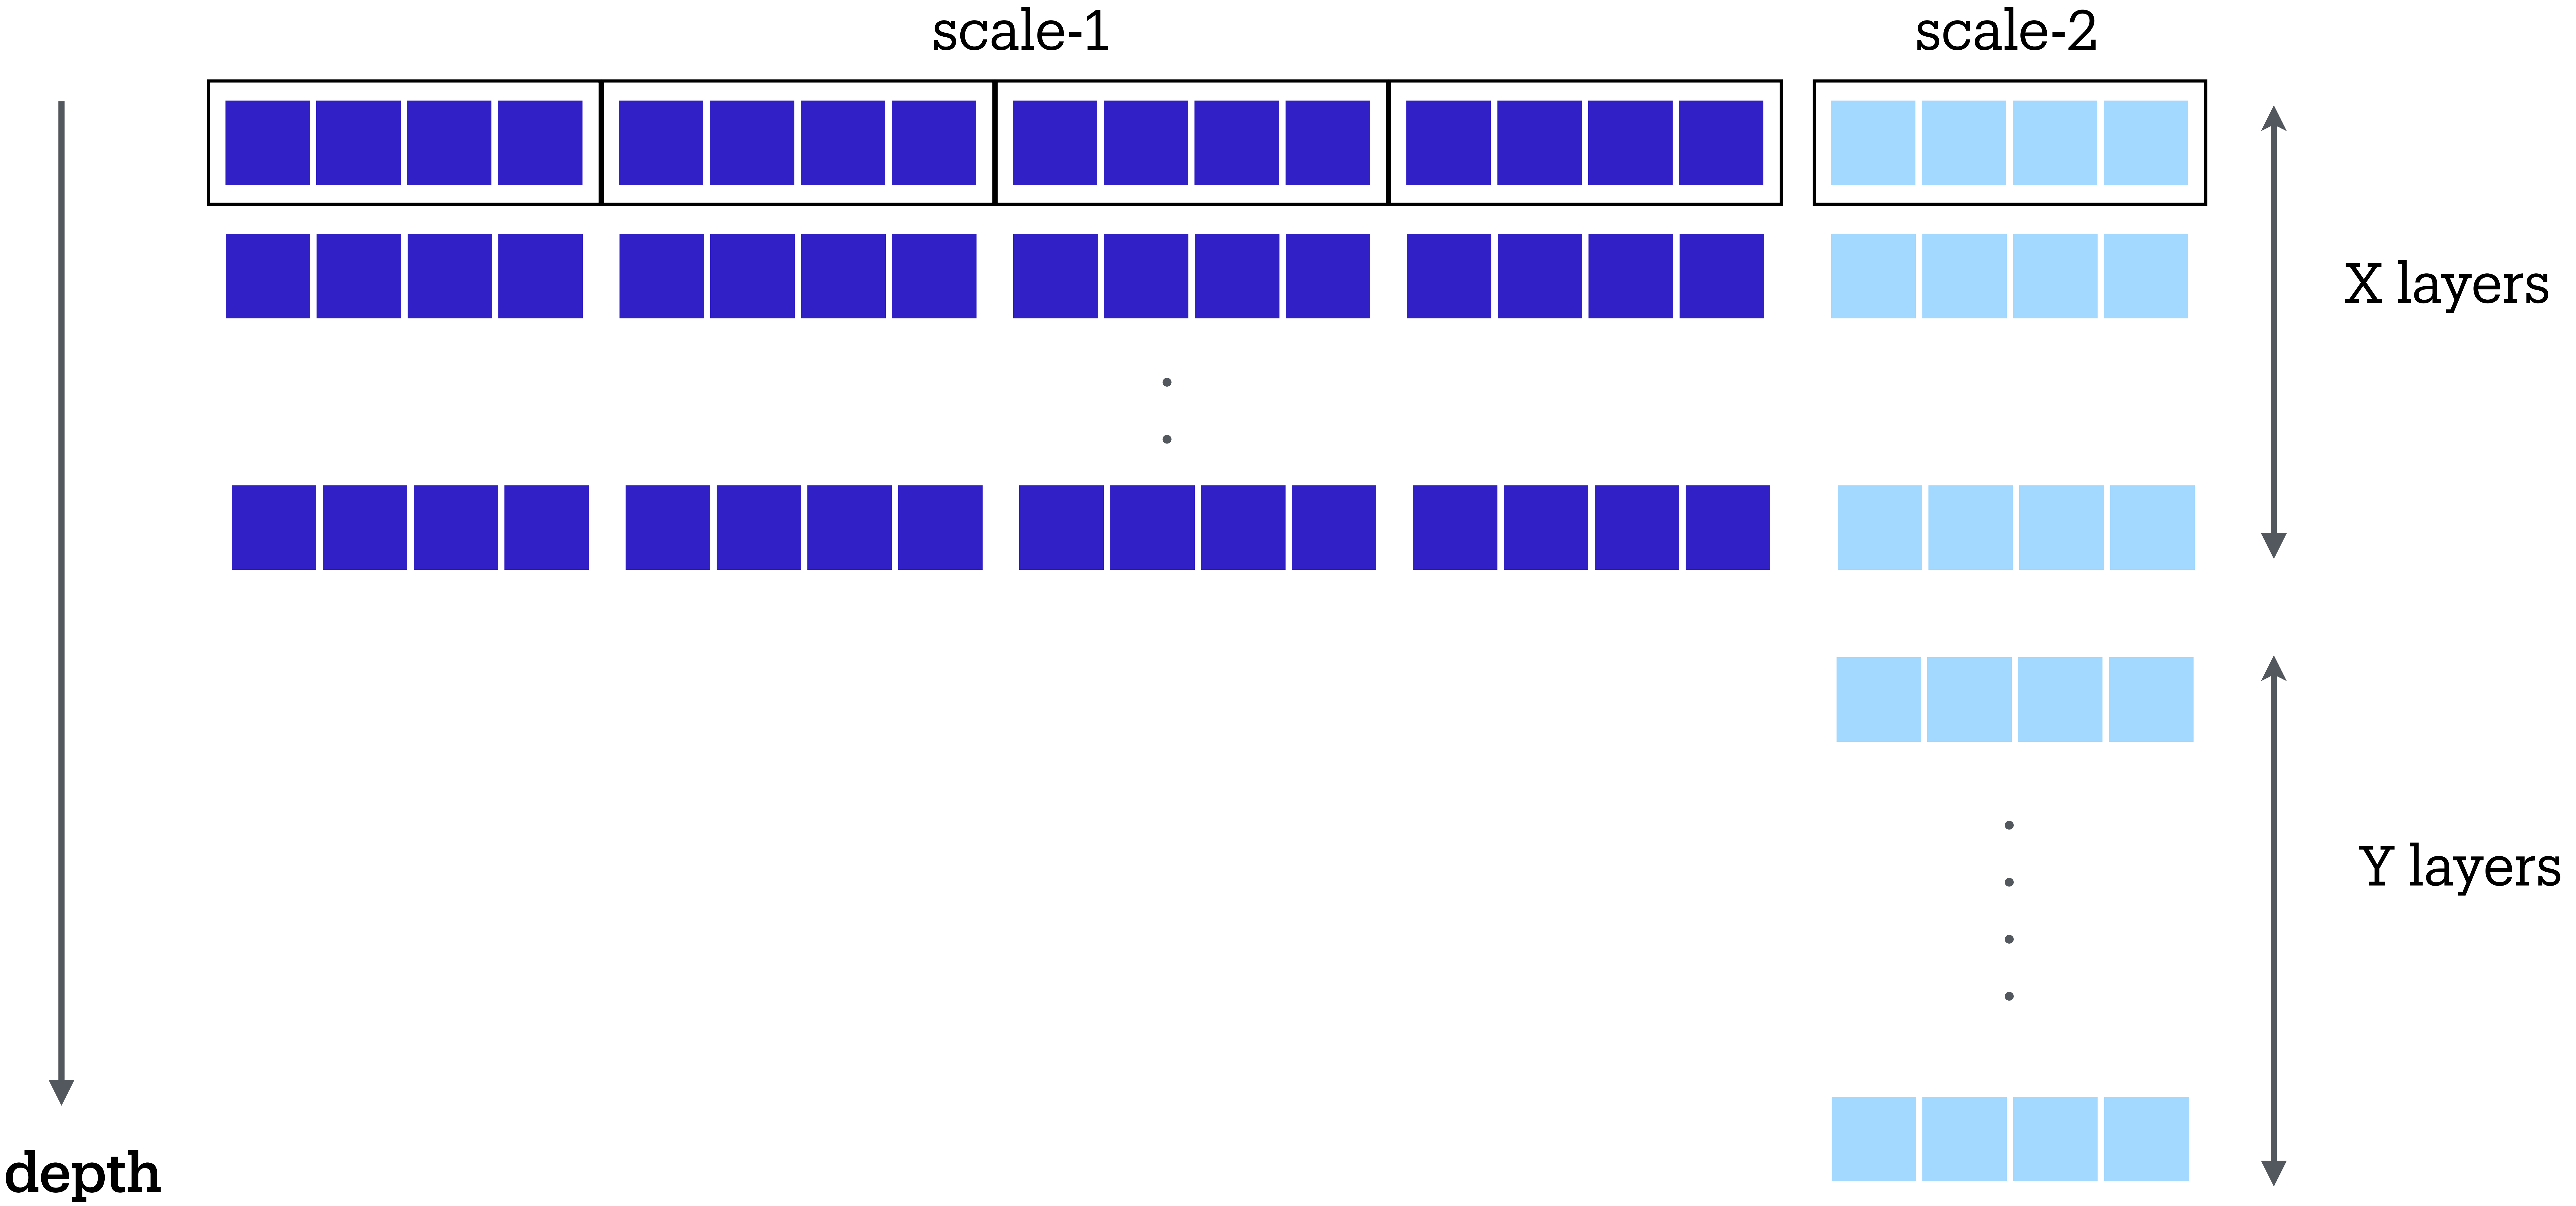

Atlas : Overview

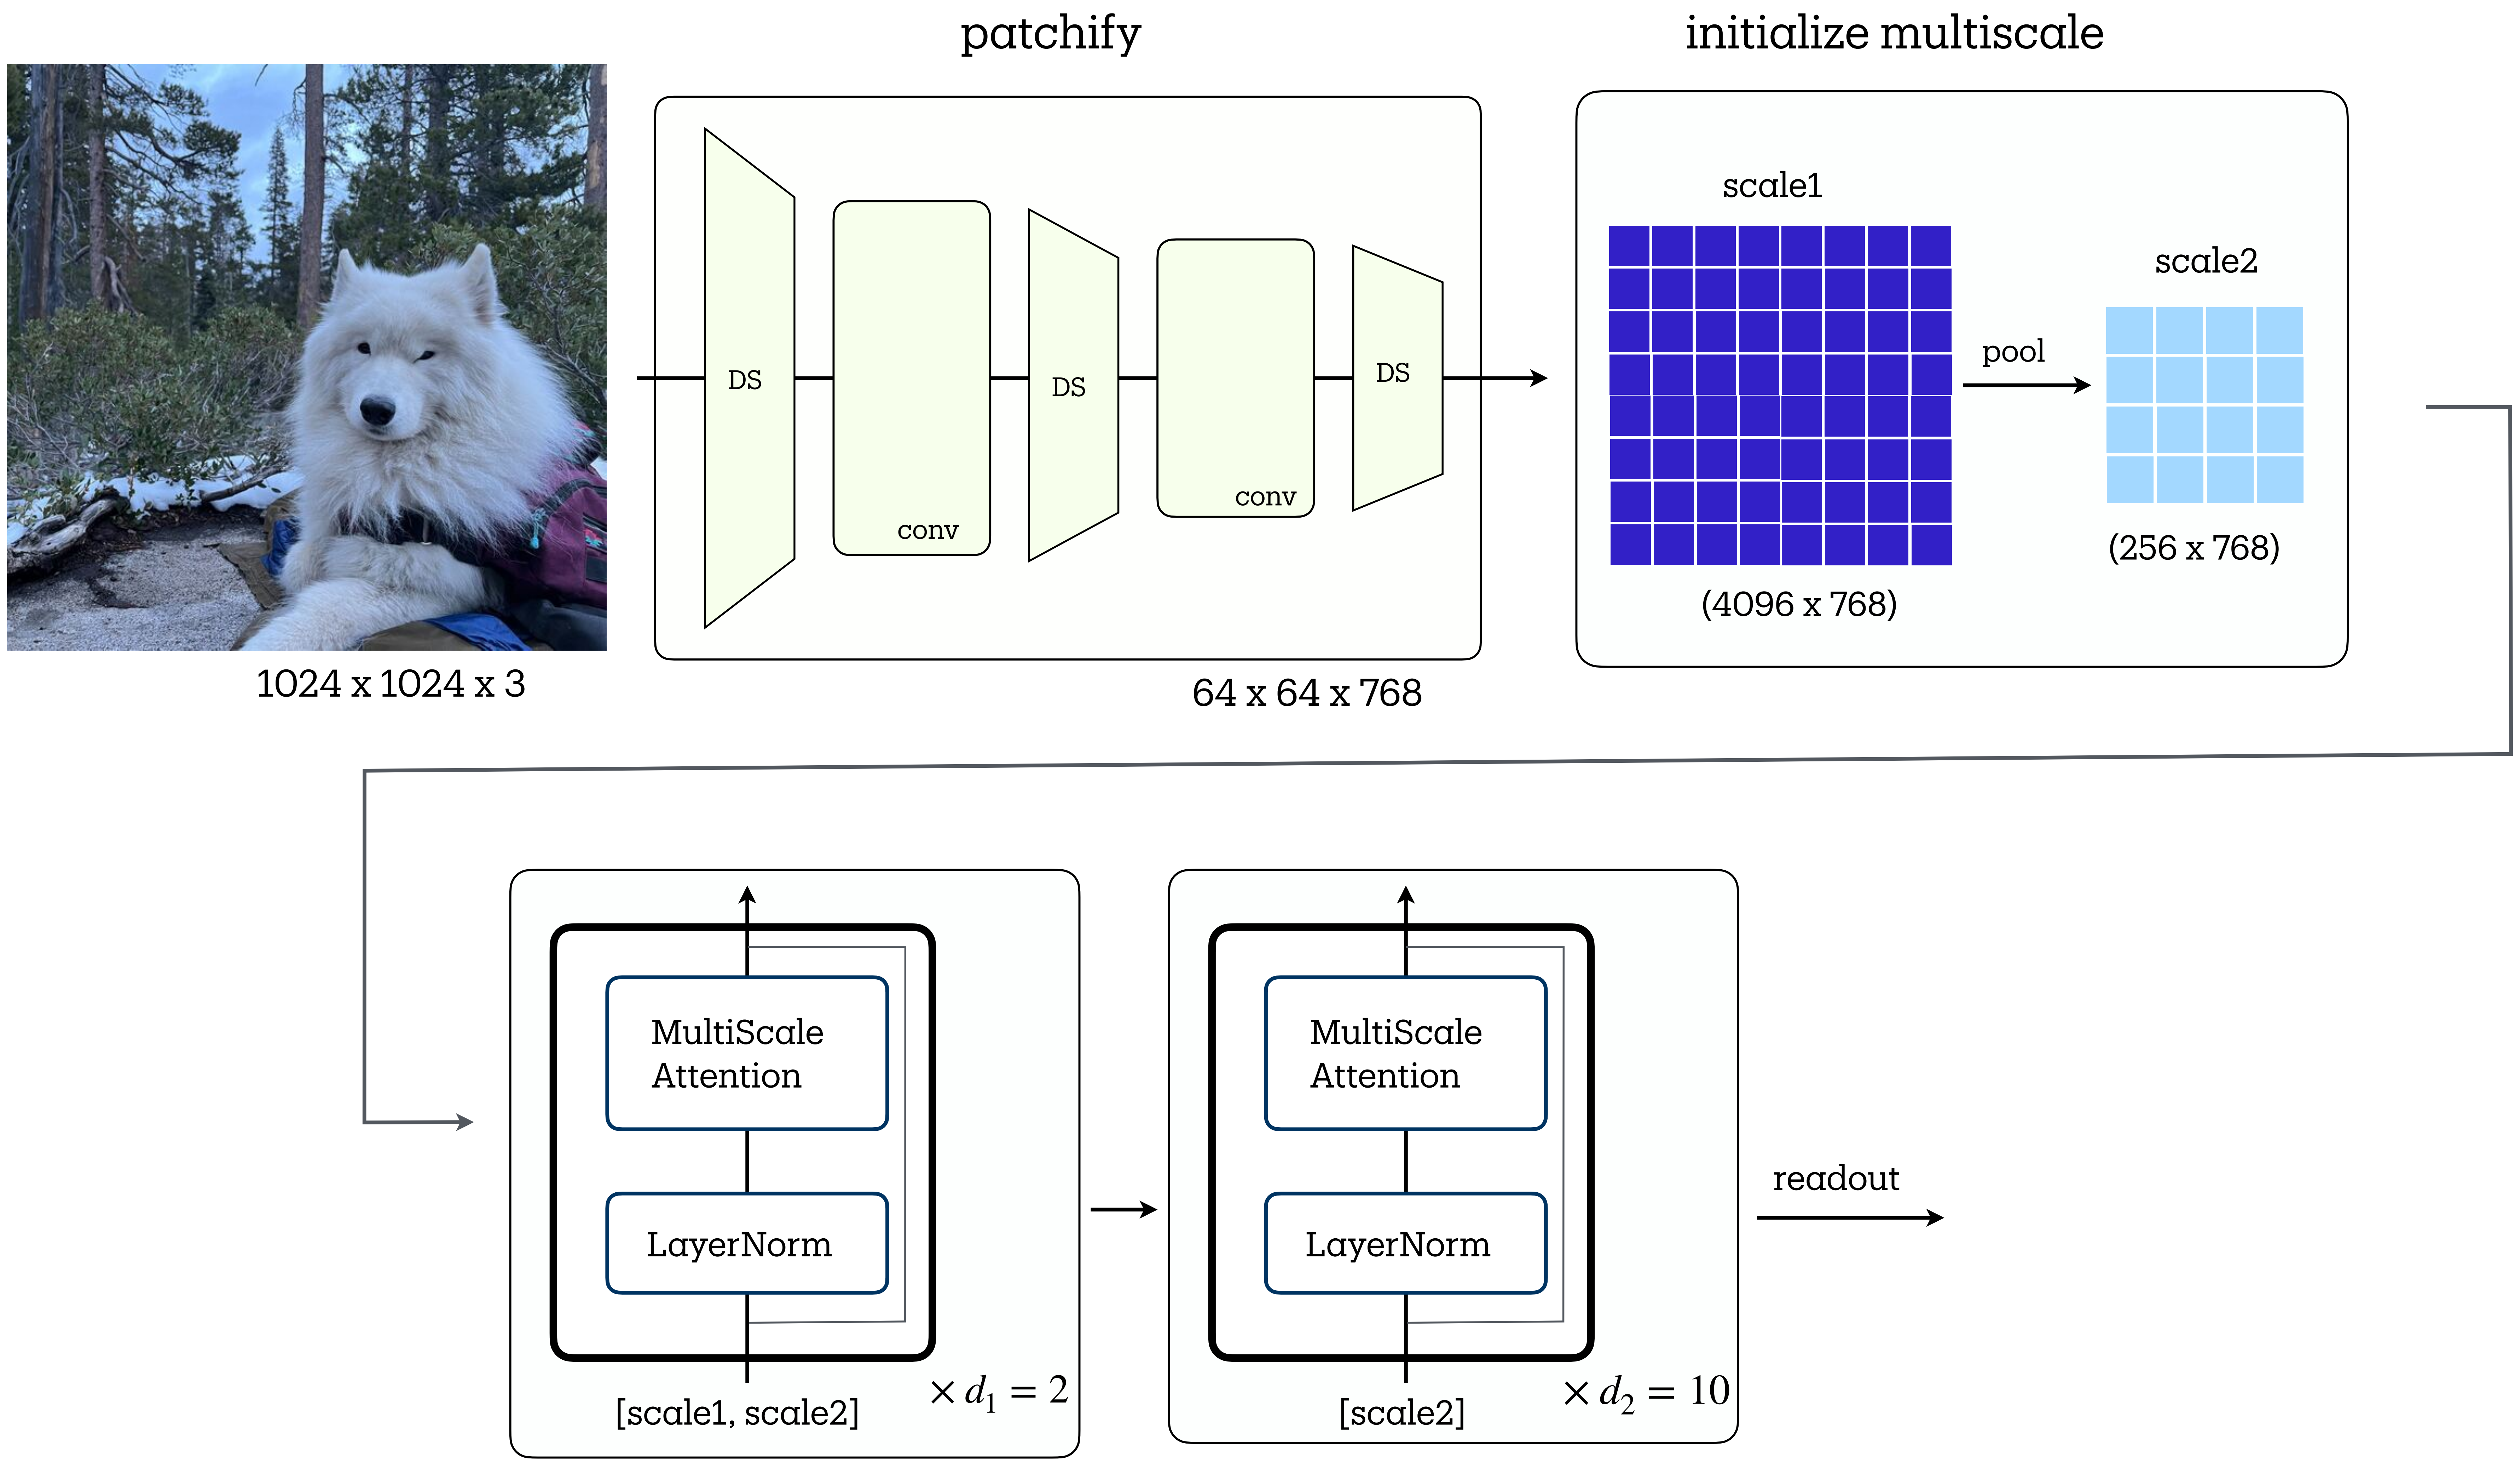

Performance Comparison

(a)

Architecture		Runtime (hr) ↓	Relative speedup ↓	Top-1 Acc. (%) ↑
Transformer	ViT-B	26.77	1.15x	90.66
	Swin-B	37.25	1.6x	90.89
	FasterViT-4	68.31	2.9×	83.66
	LongViT-B	52.23	2.2×	86.08
Convolutional	ConvNext-B	100.11	4.3×	91.92
Mamba	MambaVision-B	22.69	0.98×	84.86
Multi-Scale	Atlas-B	**23.12**	**1.00×**	**91.04**

Comparison of vision backbones on 1024x1024 image resolution on the HR-IN100 benchmark. Each model is evaluated on runtime (in hours), relative speed compared to Atlas, and Top-1 accuracy (in %). All models are base scale and were trained for 320 epochs until convergence on single 8 × H100 GPU node.

(b)

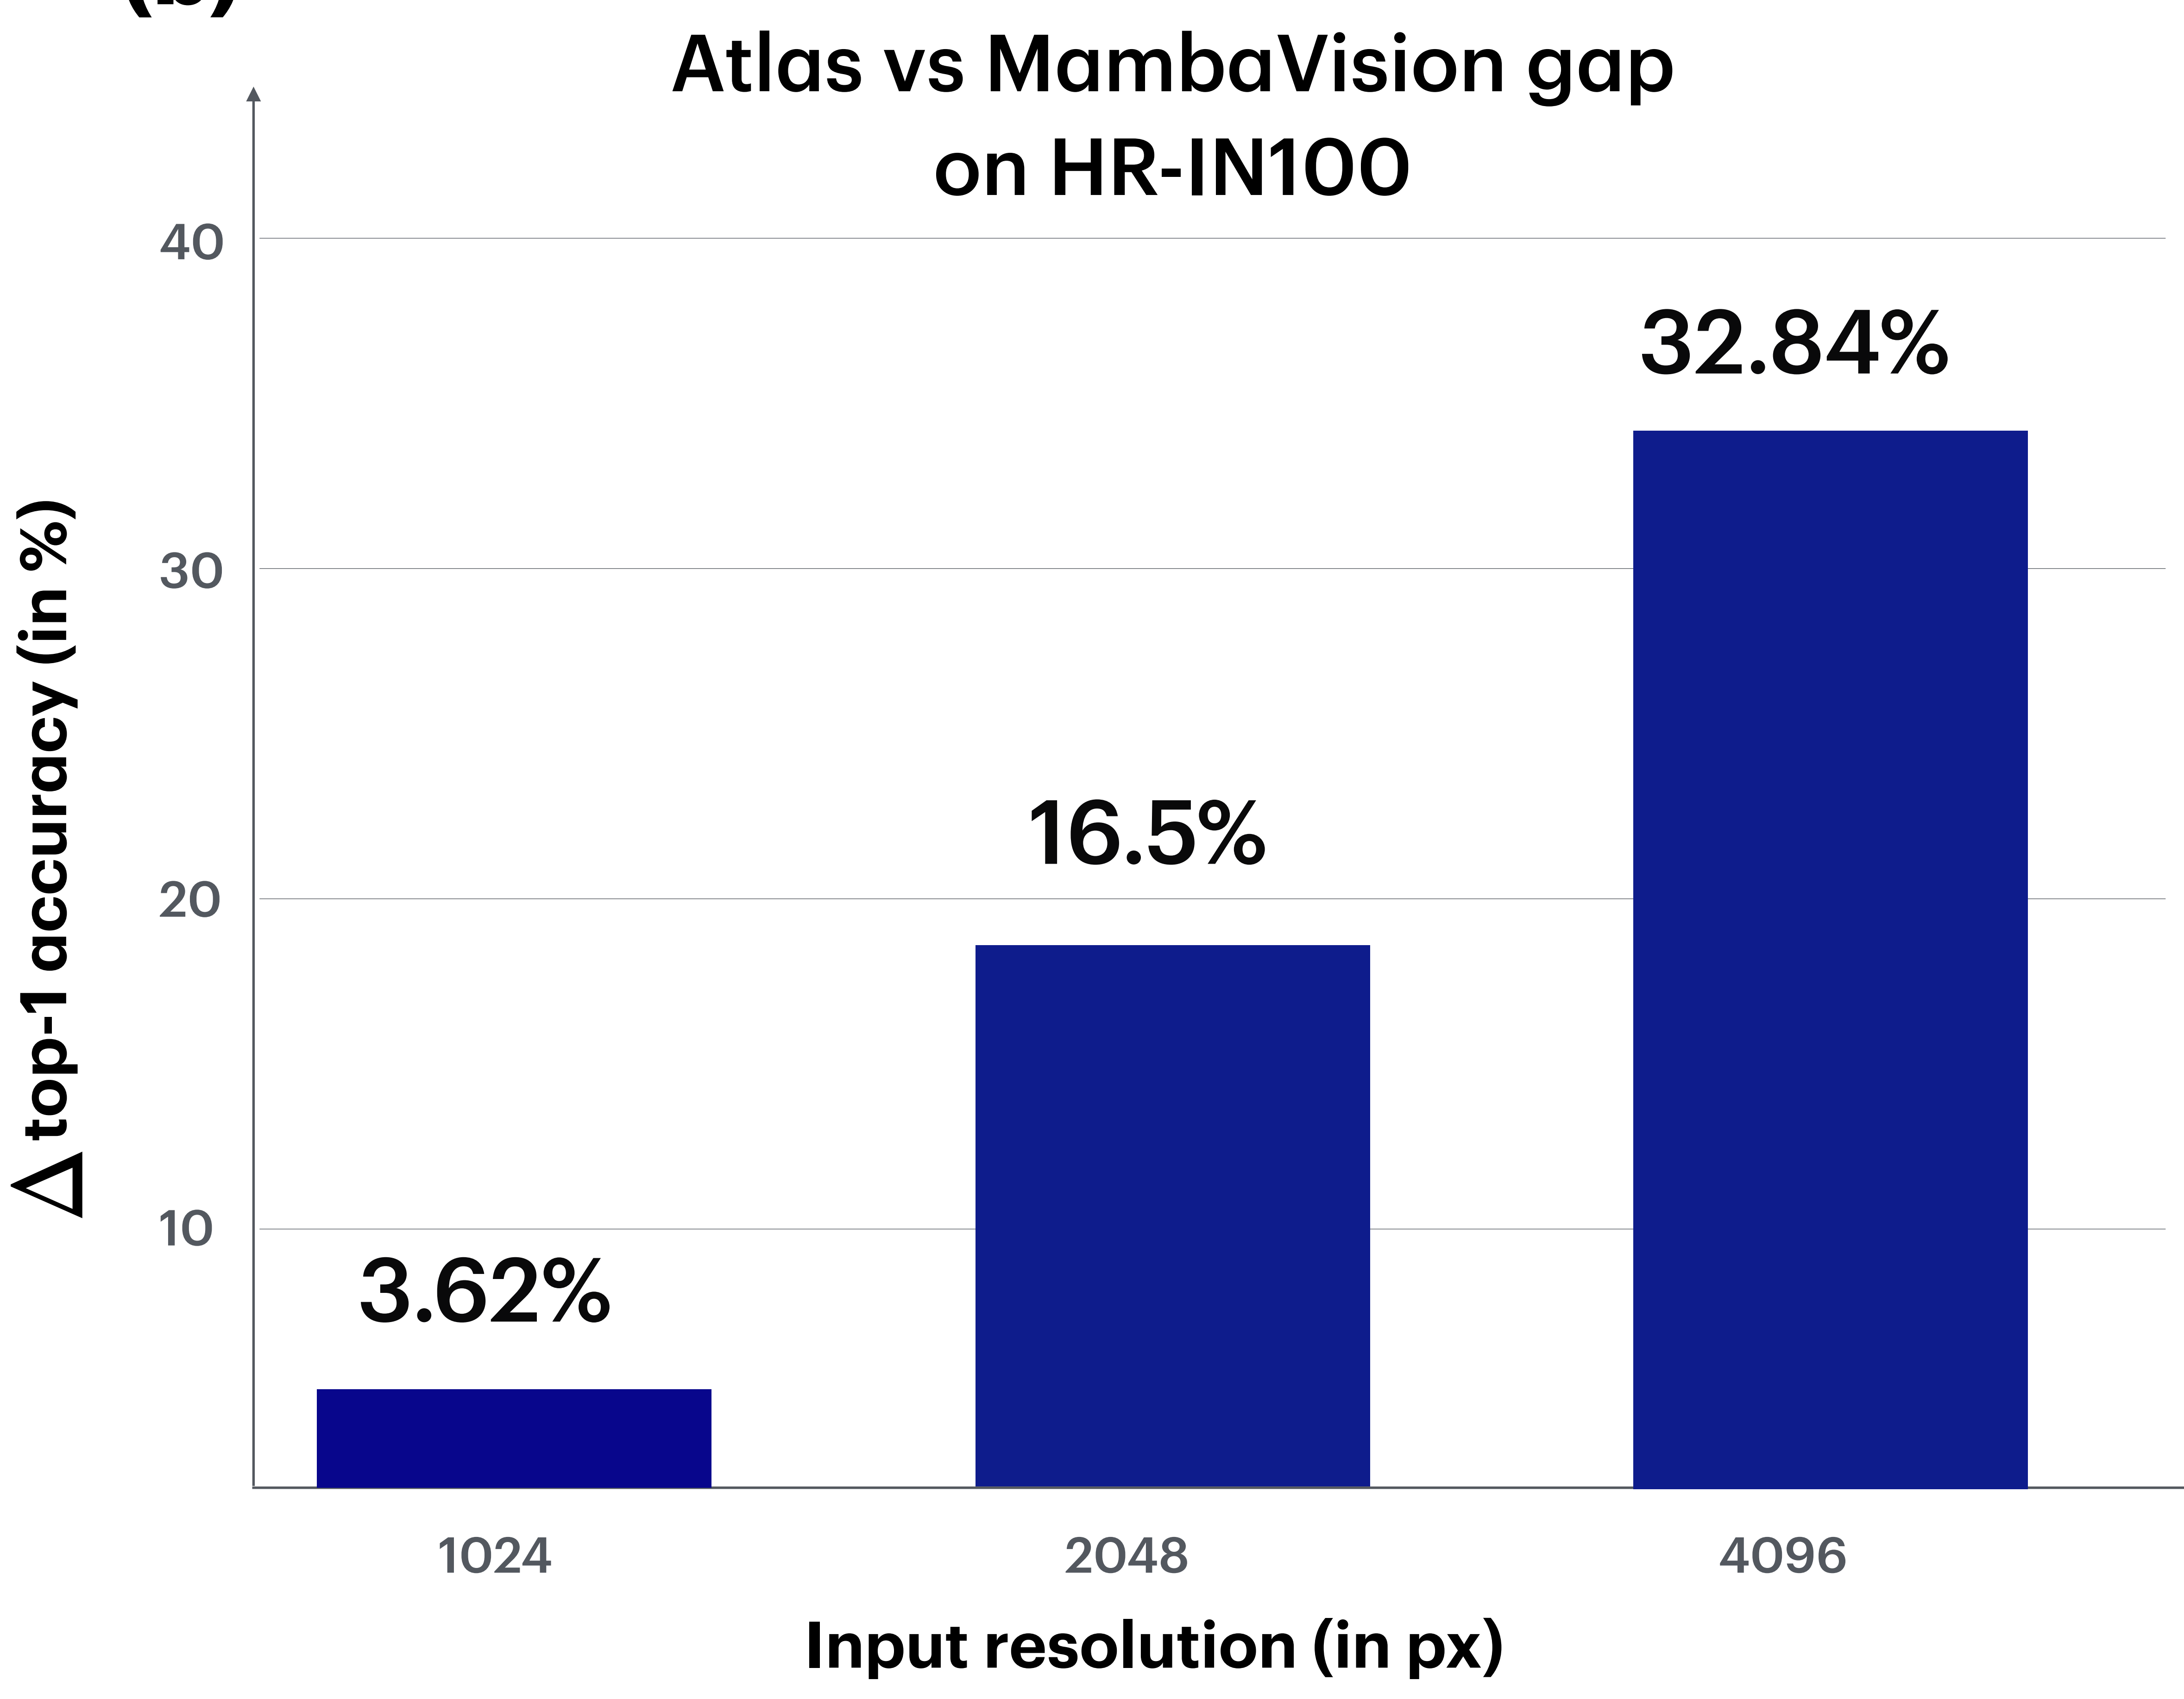

Pillar: Unified pre-training across 5 modalites

** New unpublished model*

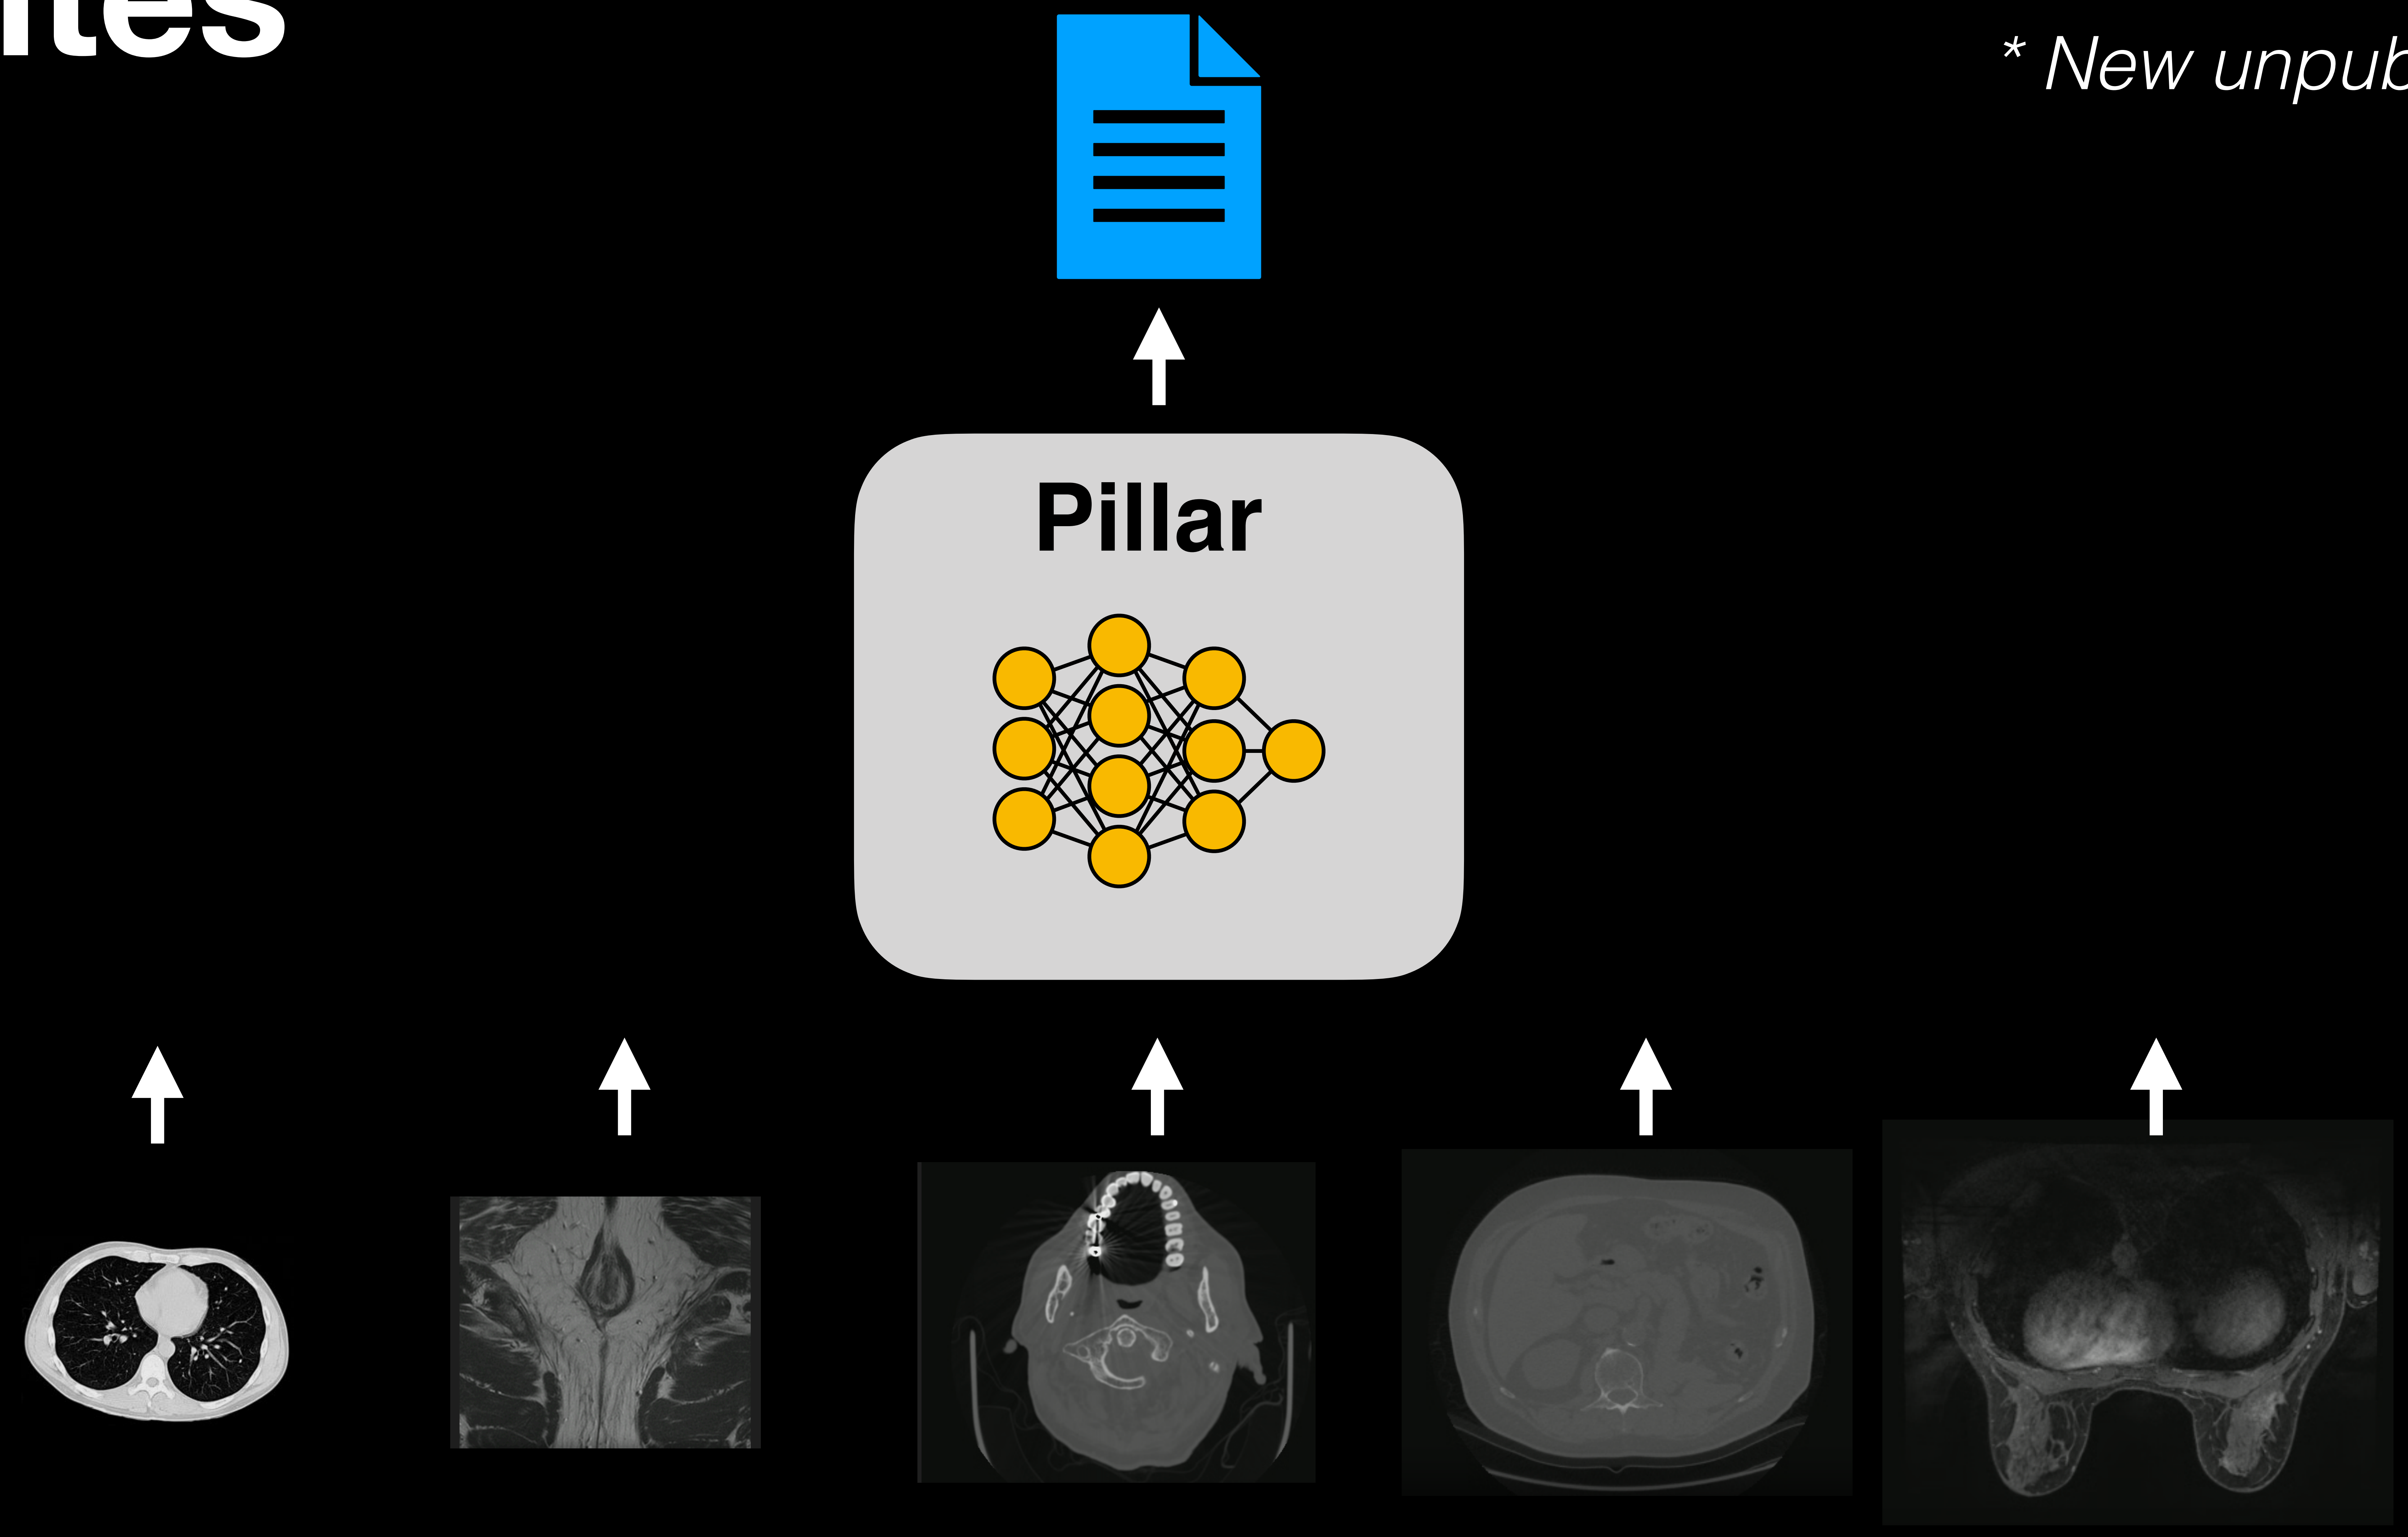

Leaps in performance 80+ abdomen CT findings

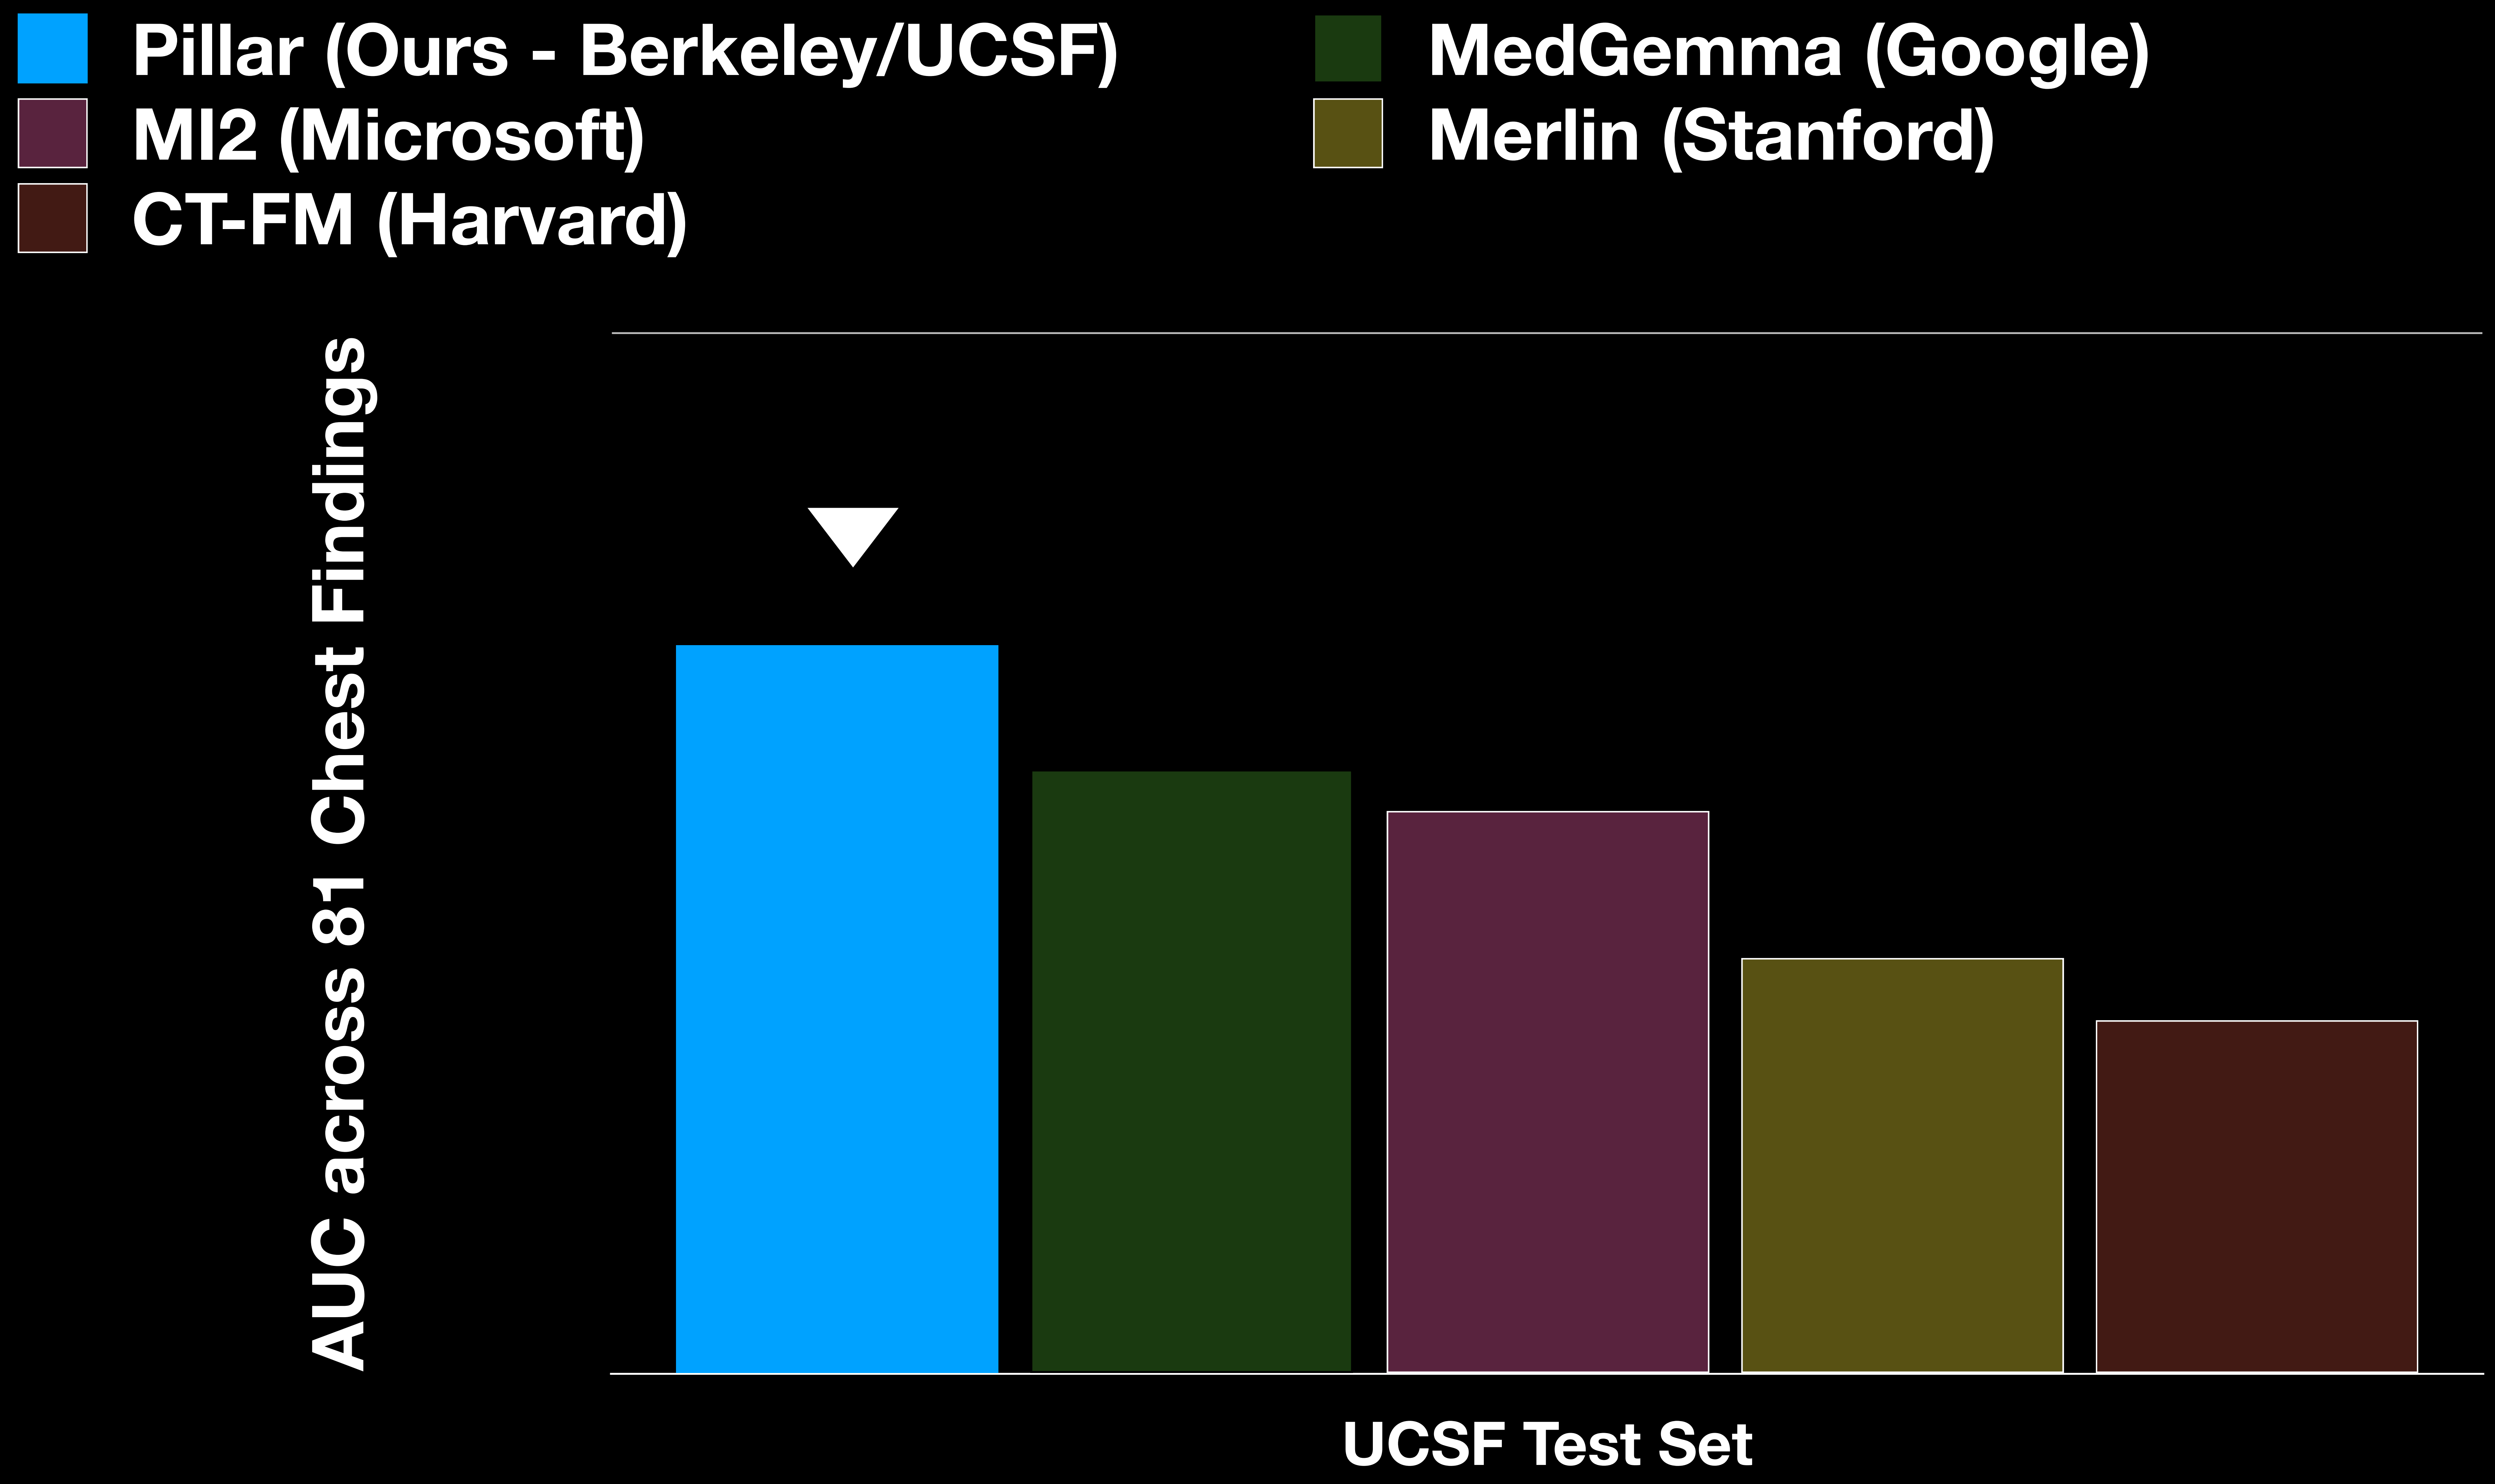

Today: Towards AI-driven care

Prediction

Control

Translation

Today: Towards AI-driven care

Control

How to catch cancer earlier

Predict Cancer Risk

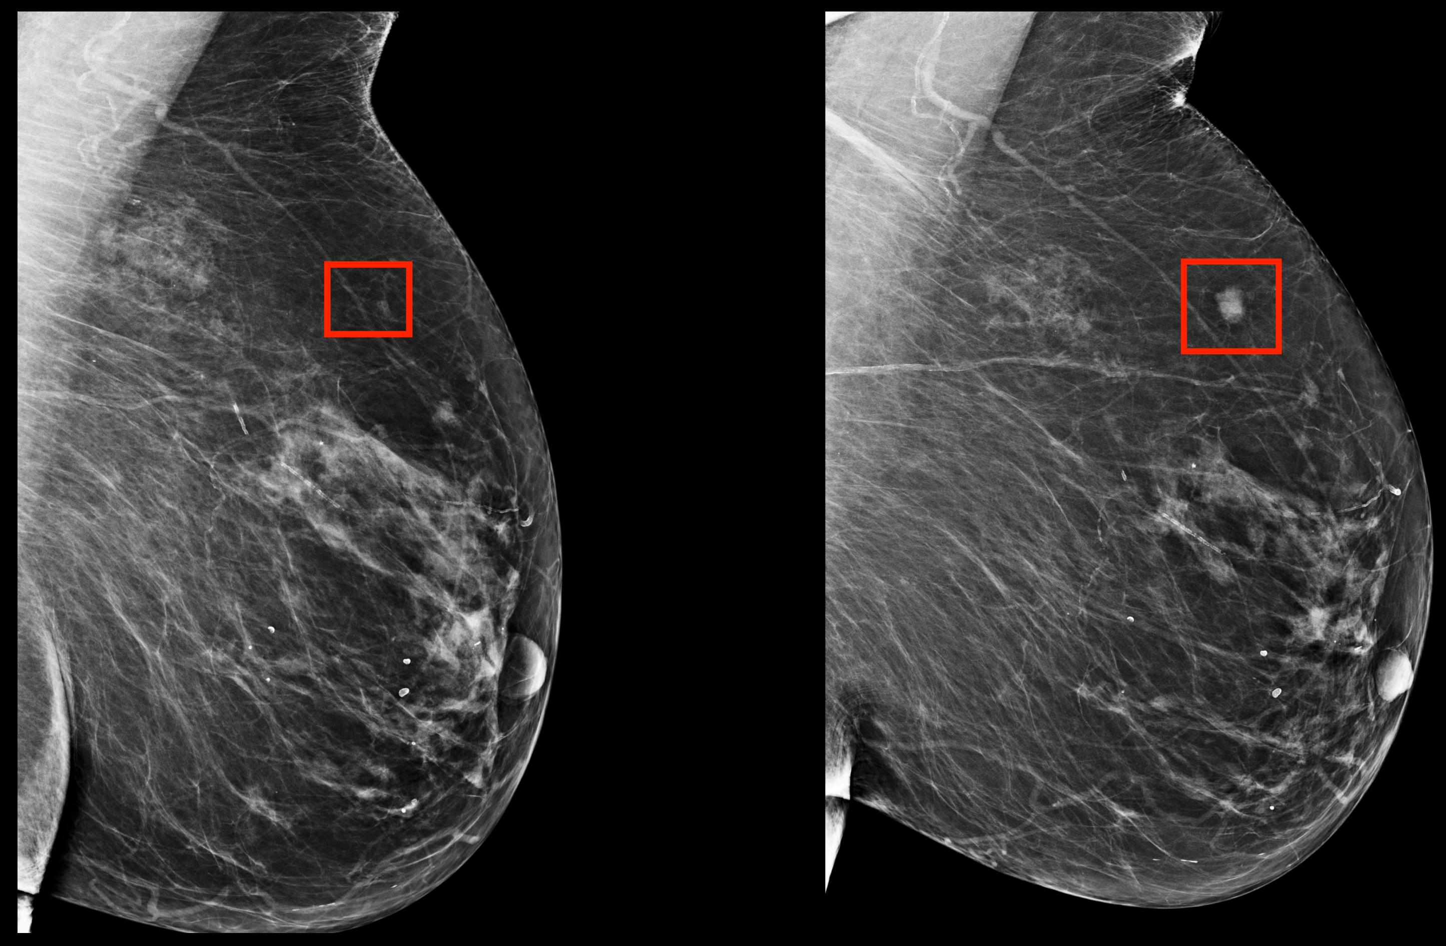

Year 0 Year 5

Create personalized screening policy

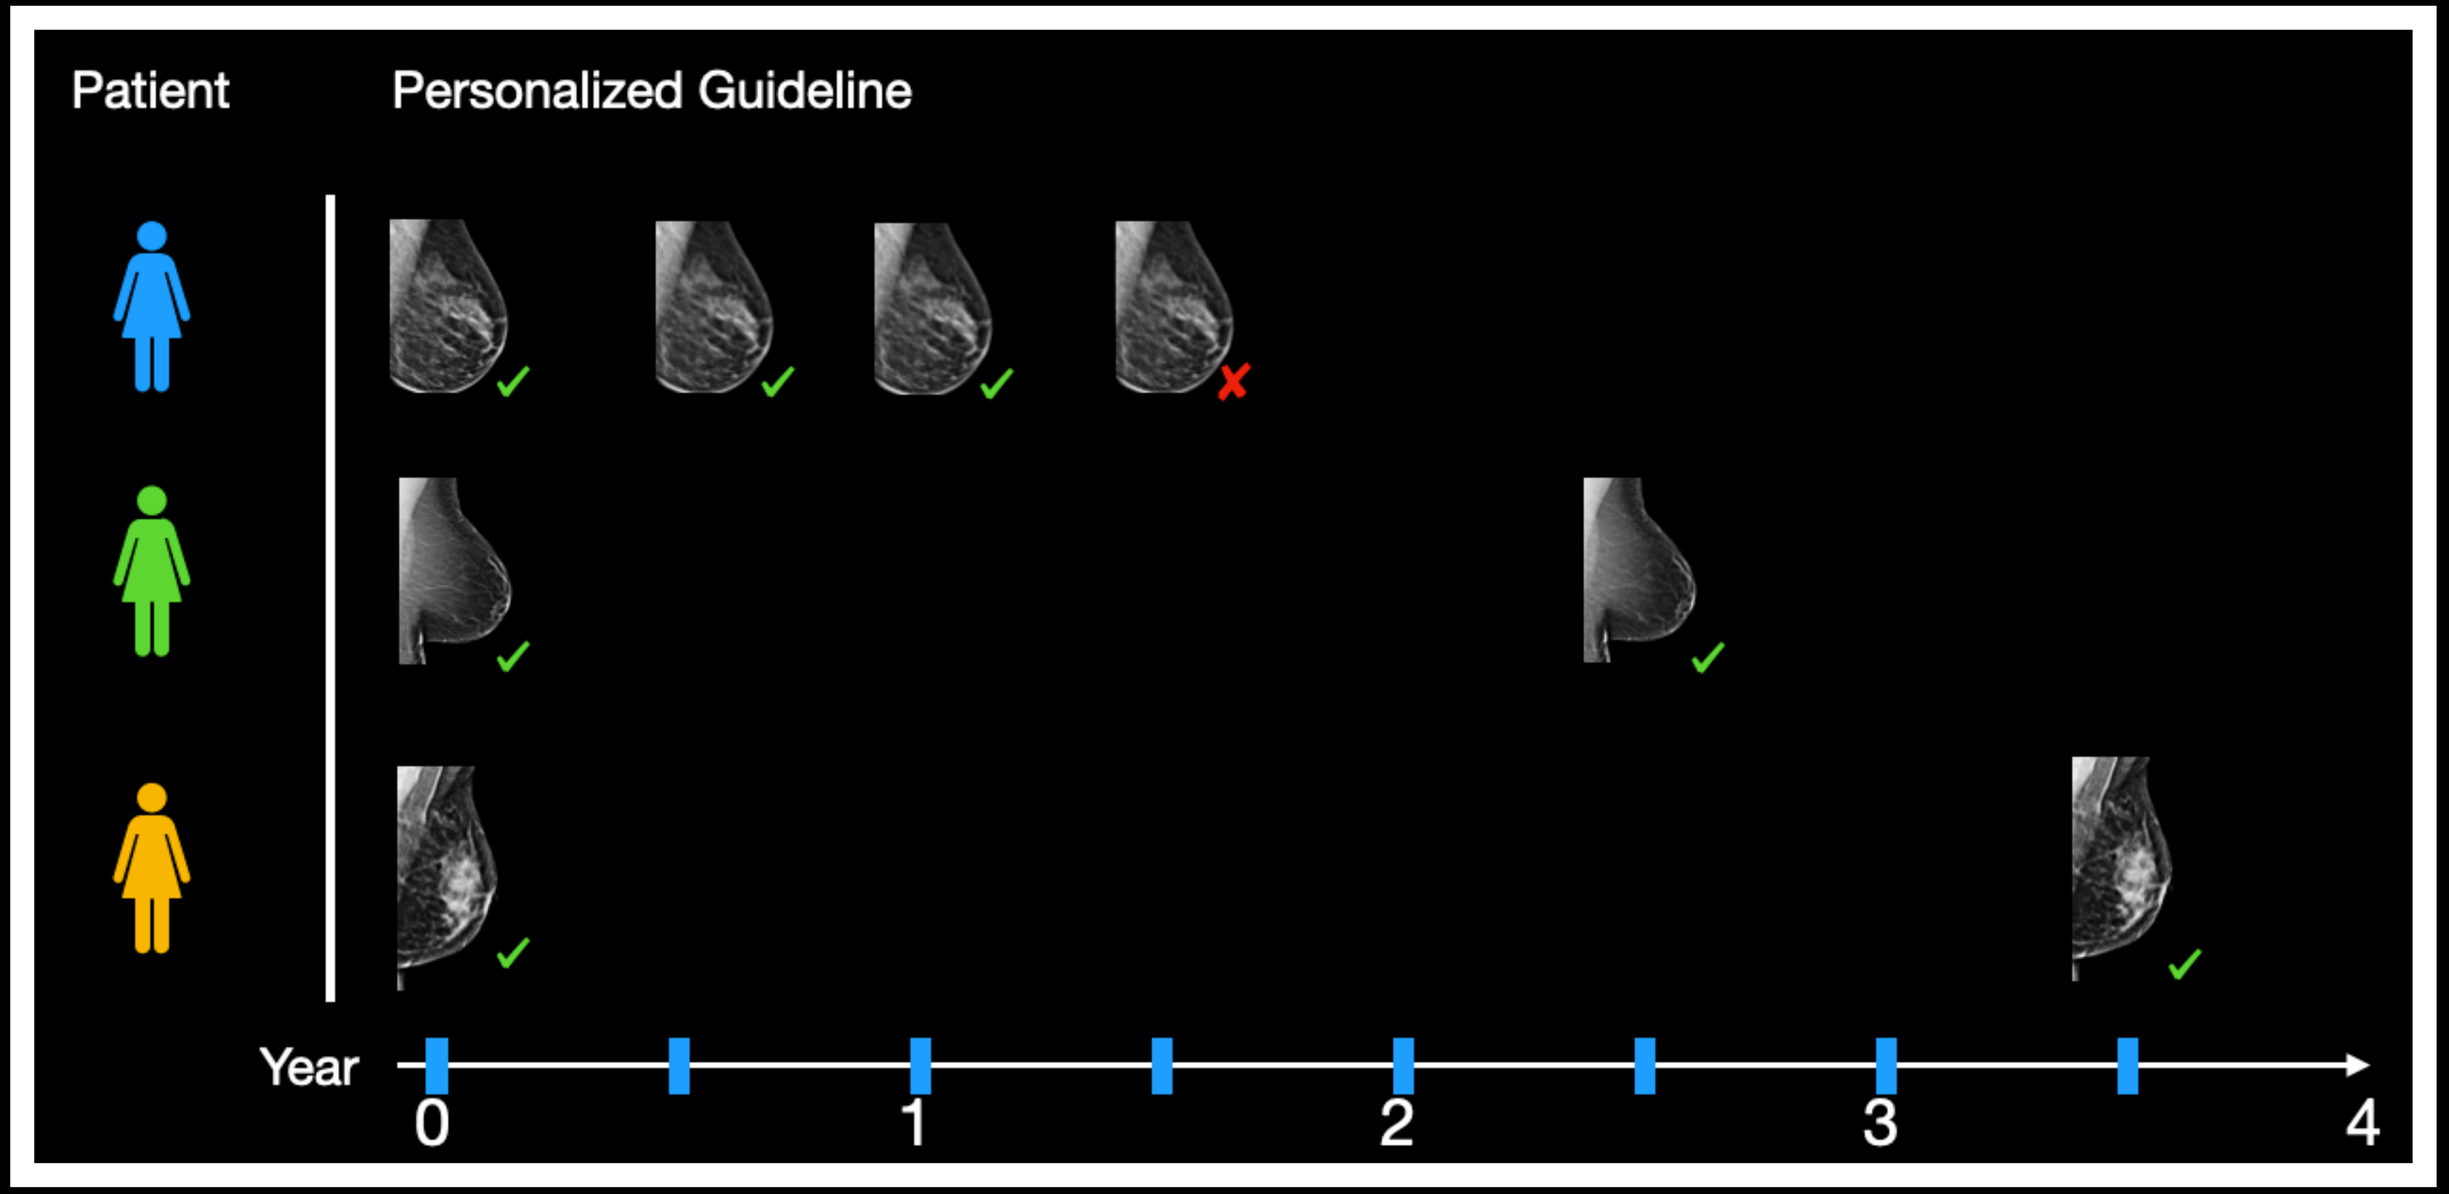

How to catch cancer earlier

nature medicine

ARTICLES

https://doi.org/10.1038/s41591-021-01599-w

Check for updates

Optimizing risk-based breast cancer screening policies with reinforcement learning

Adam Yala[1,2], Peter G. Mikhael[1,2], Constance Lehman[3], Gigin Lin[4,5], Fredrik Strand[6,7], Yung-Liang Wan[4,5], Kevin Hughes[8], Siddharth Satuluru[9], Thomas Kim[10], Imon Banerjee[11], Judy Gichoya[9], Hari Trivedi[9] and Regina Barzilay[1,2]

Create personalized screening policy

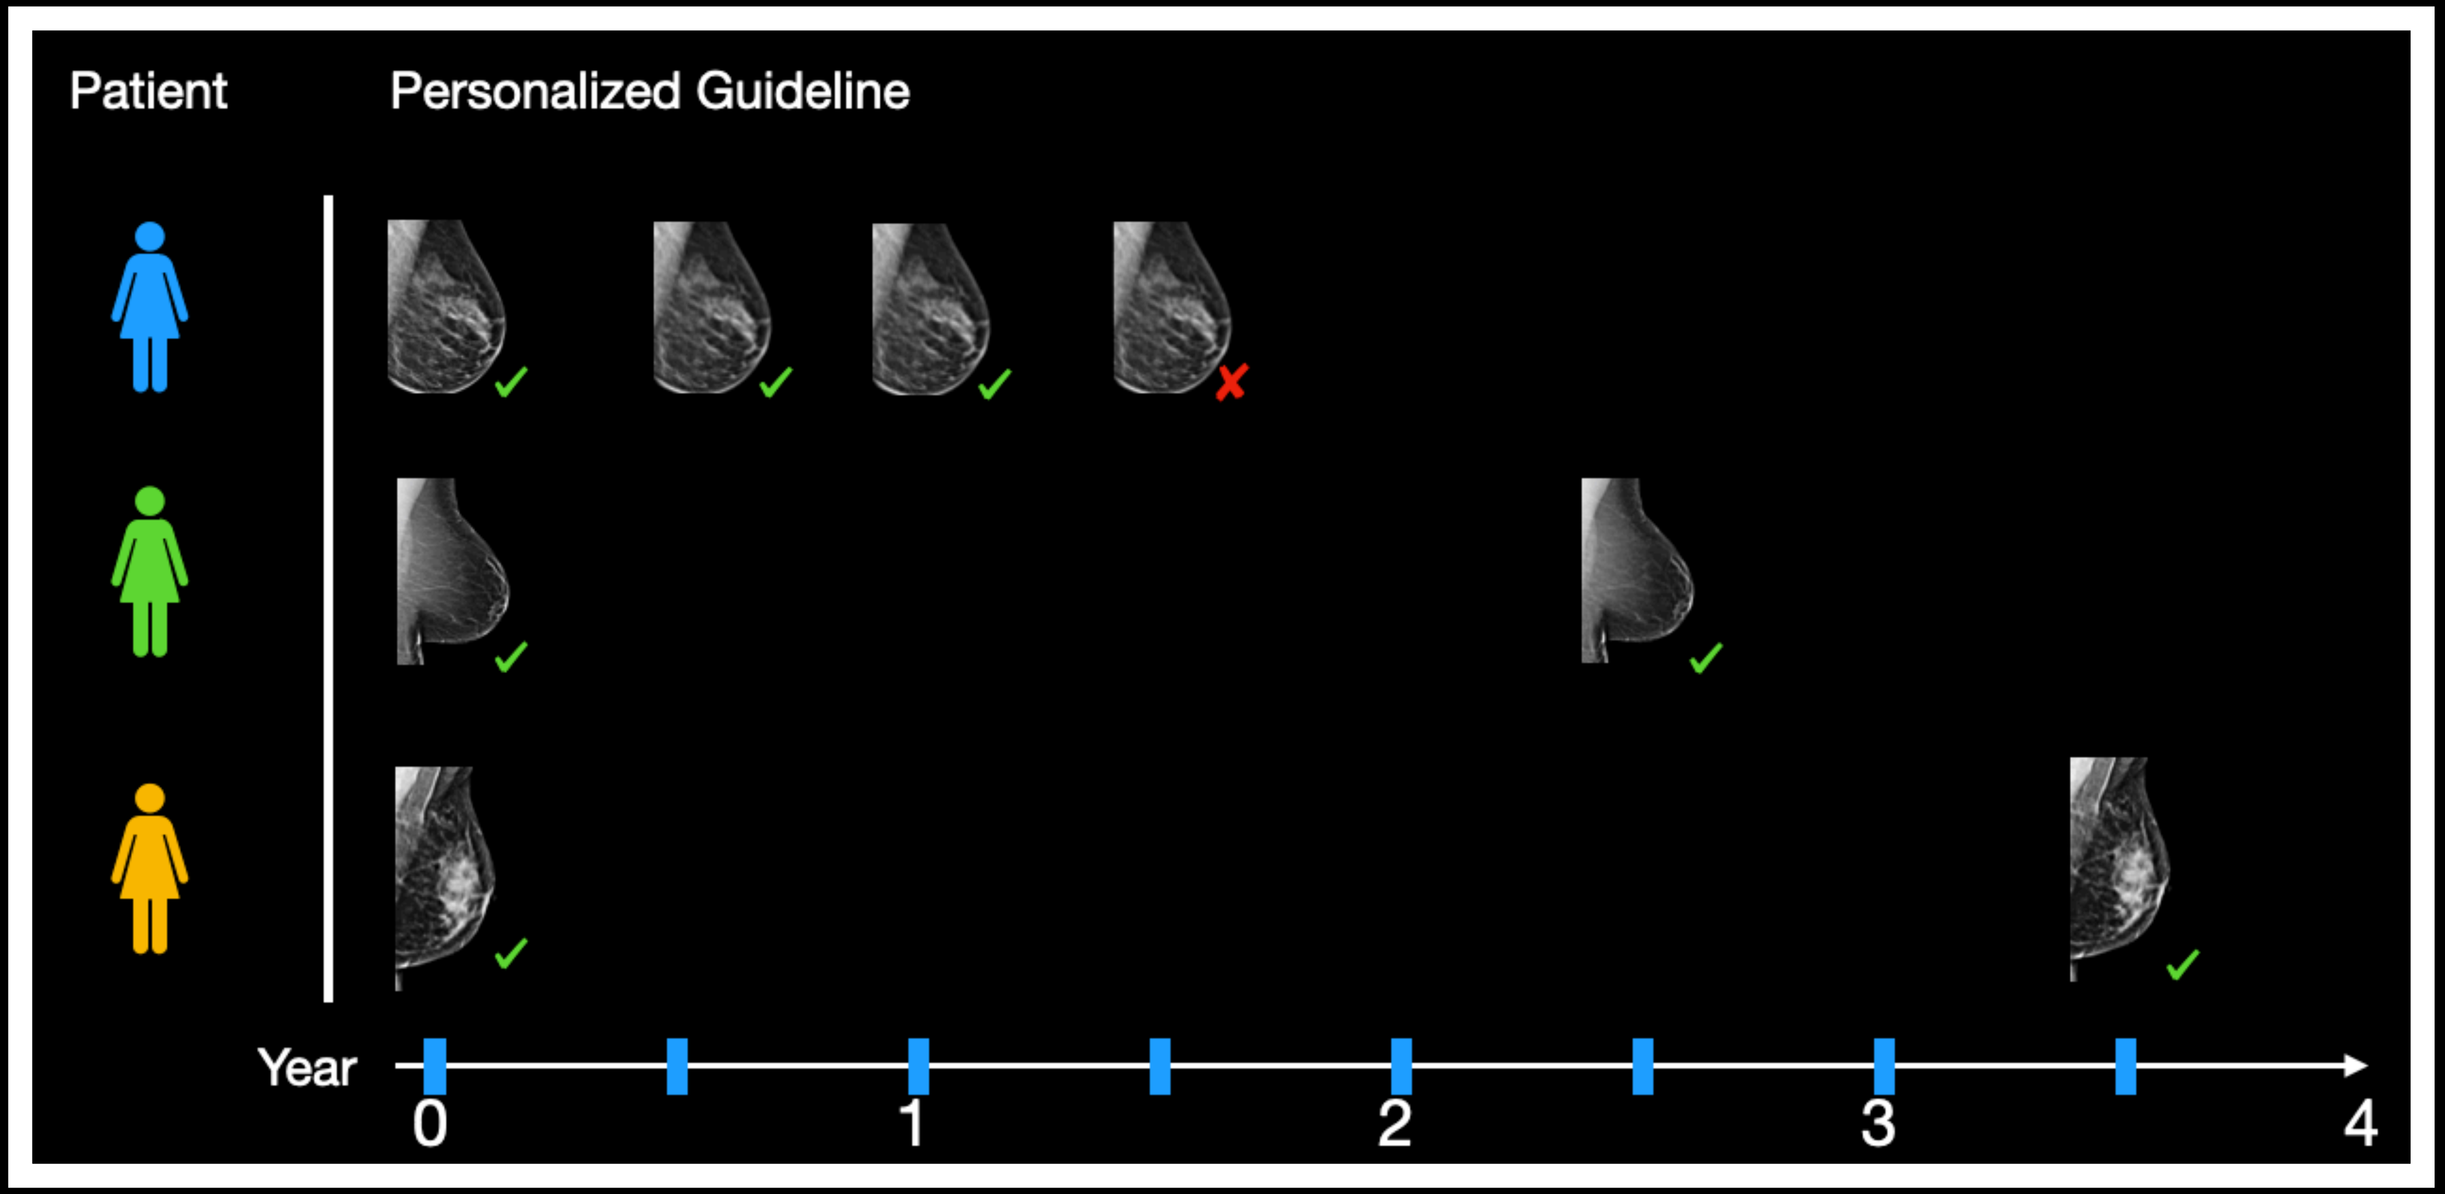

Screening today

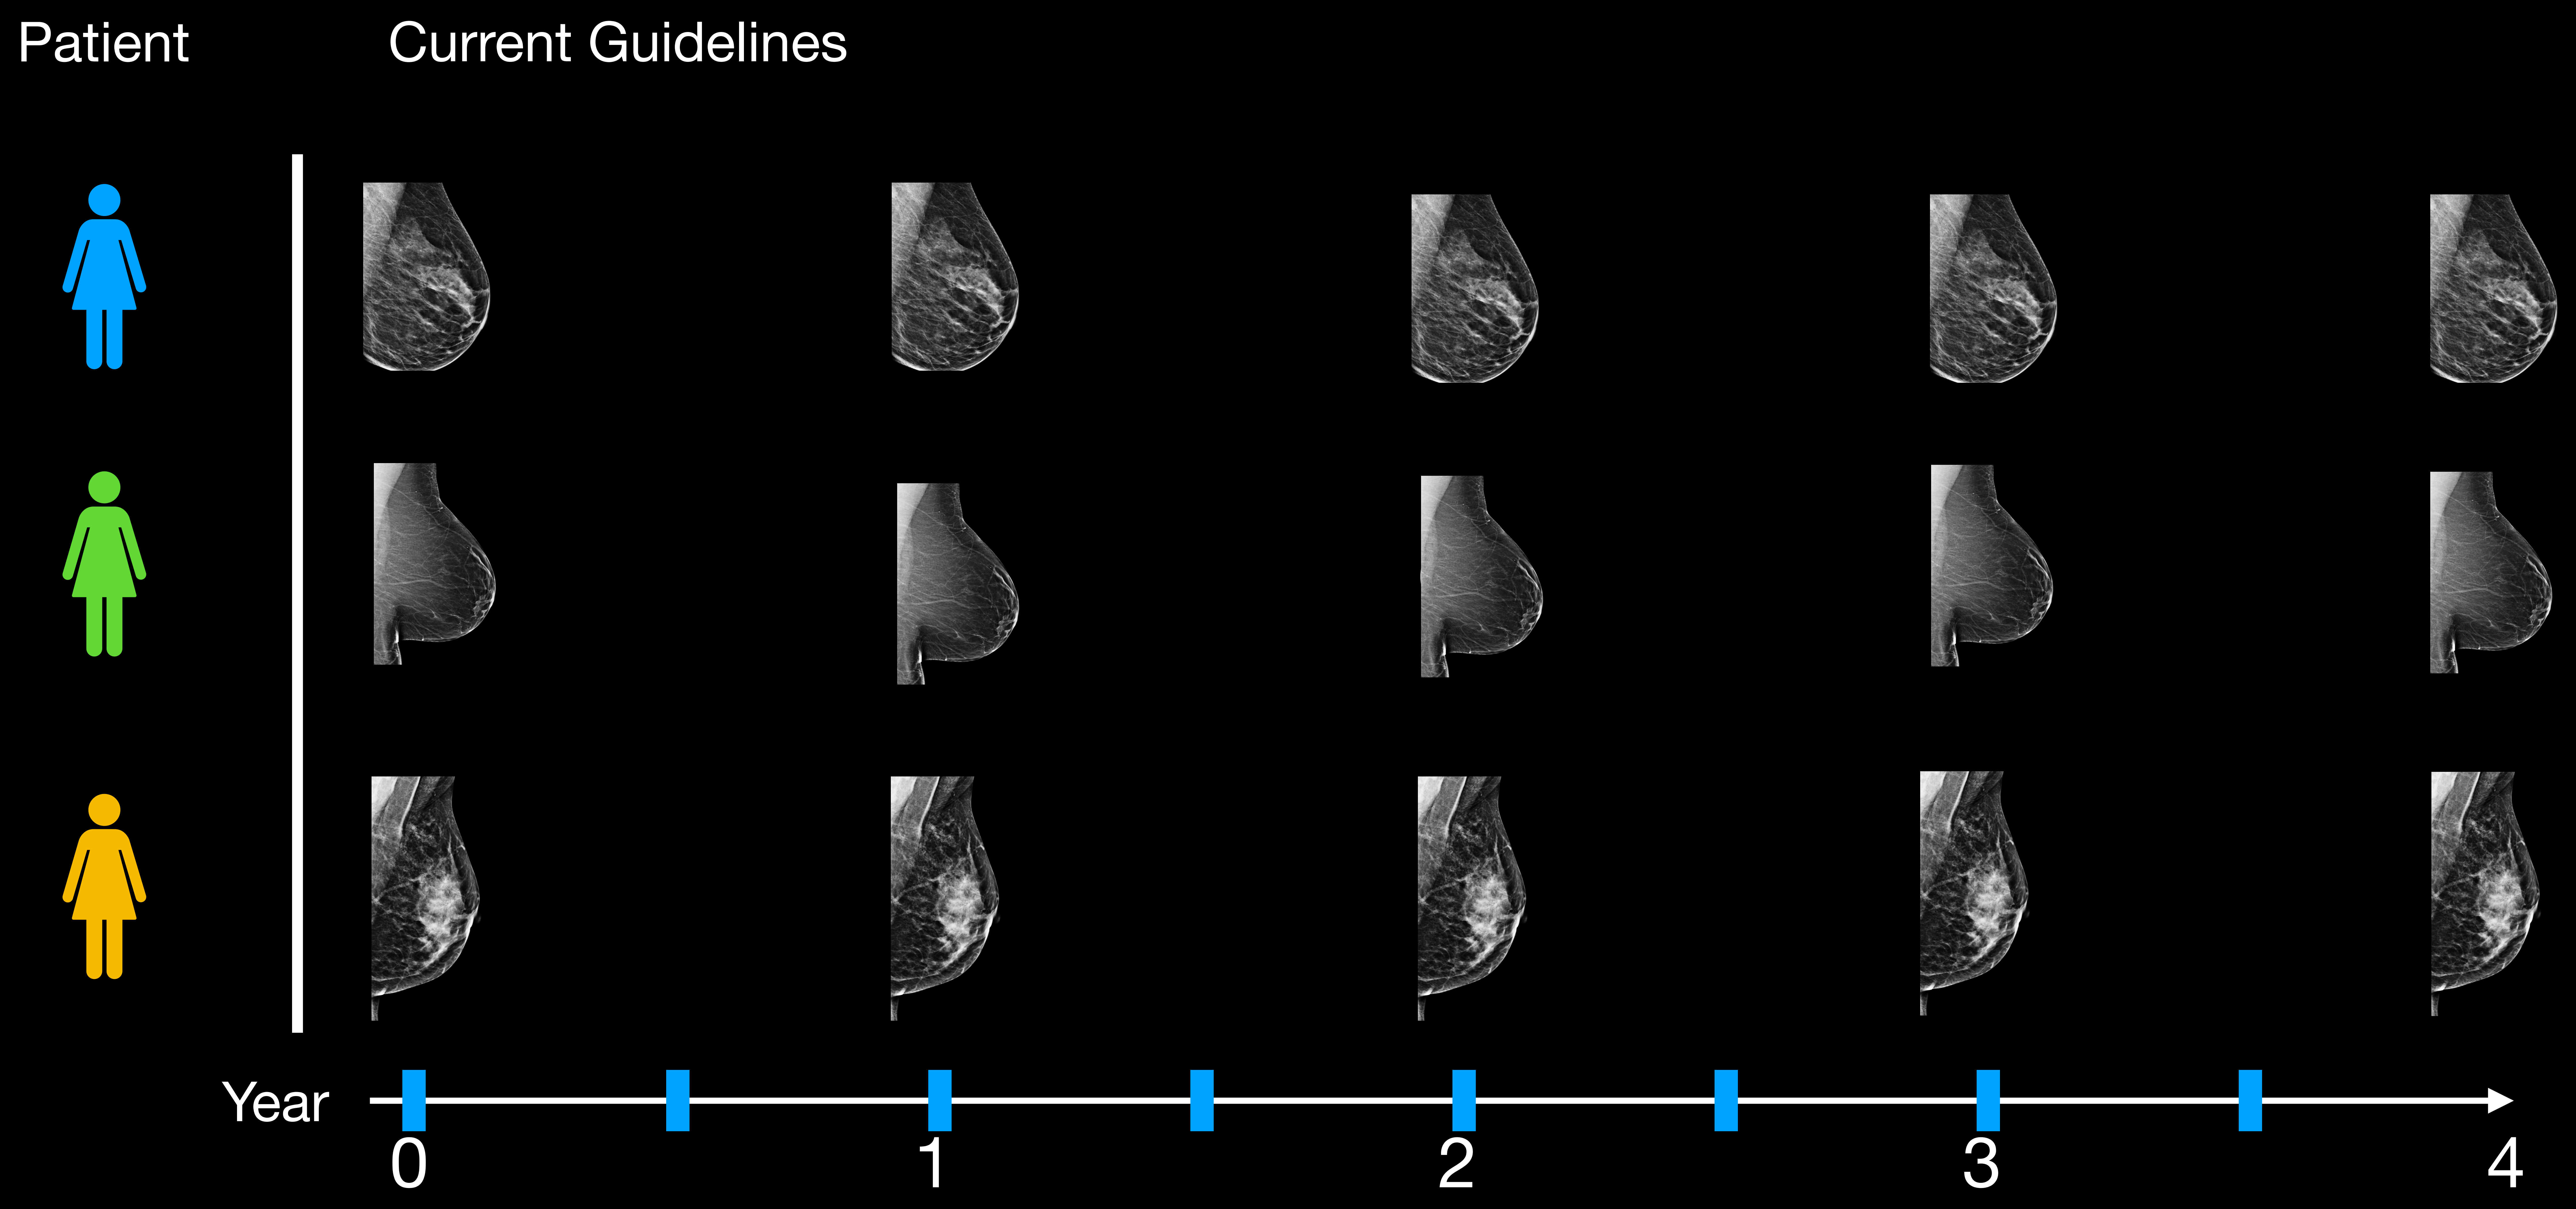

Same screening, different outcomes

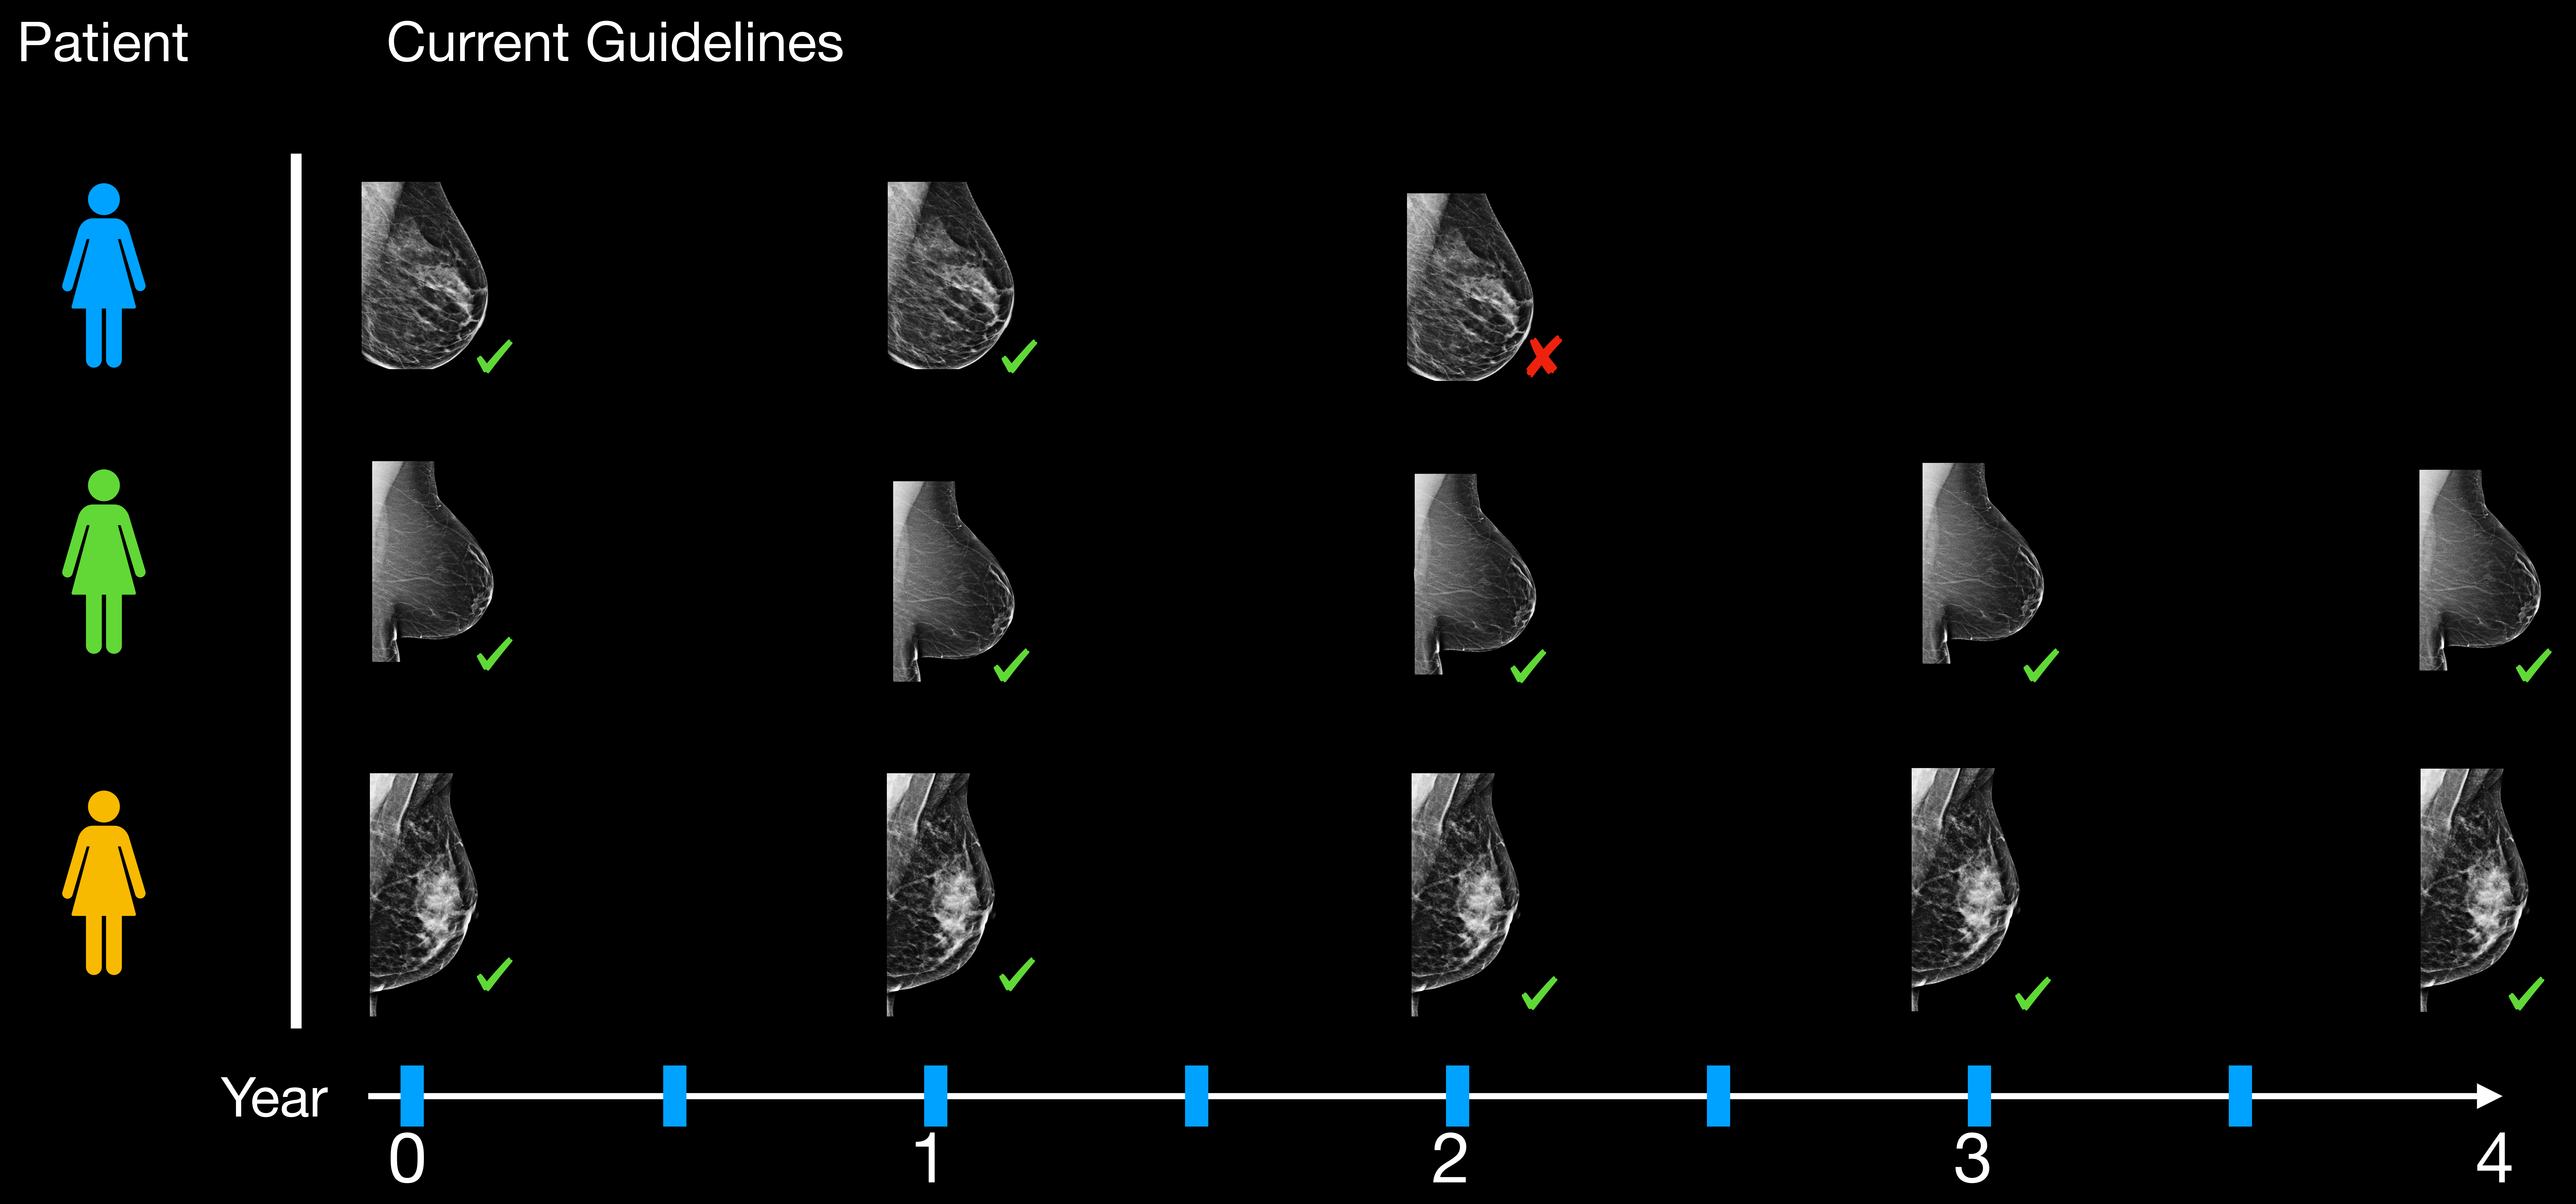

Challenges in current screening

Late Detection

Over Screening

Health Disparities

Tailor screening regime to patient need

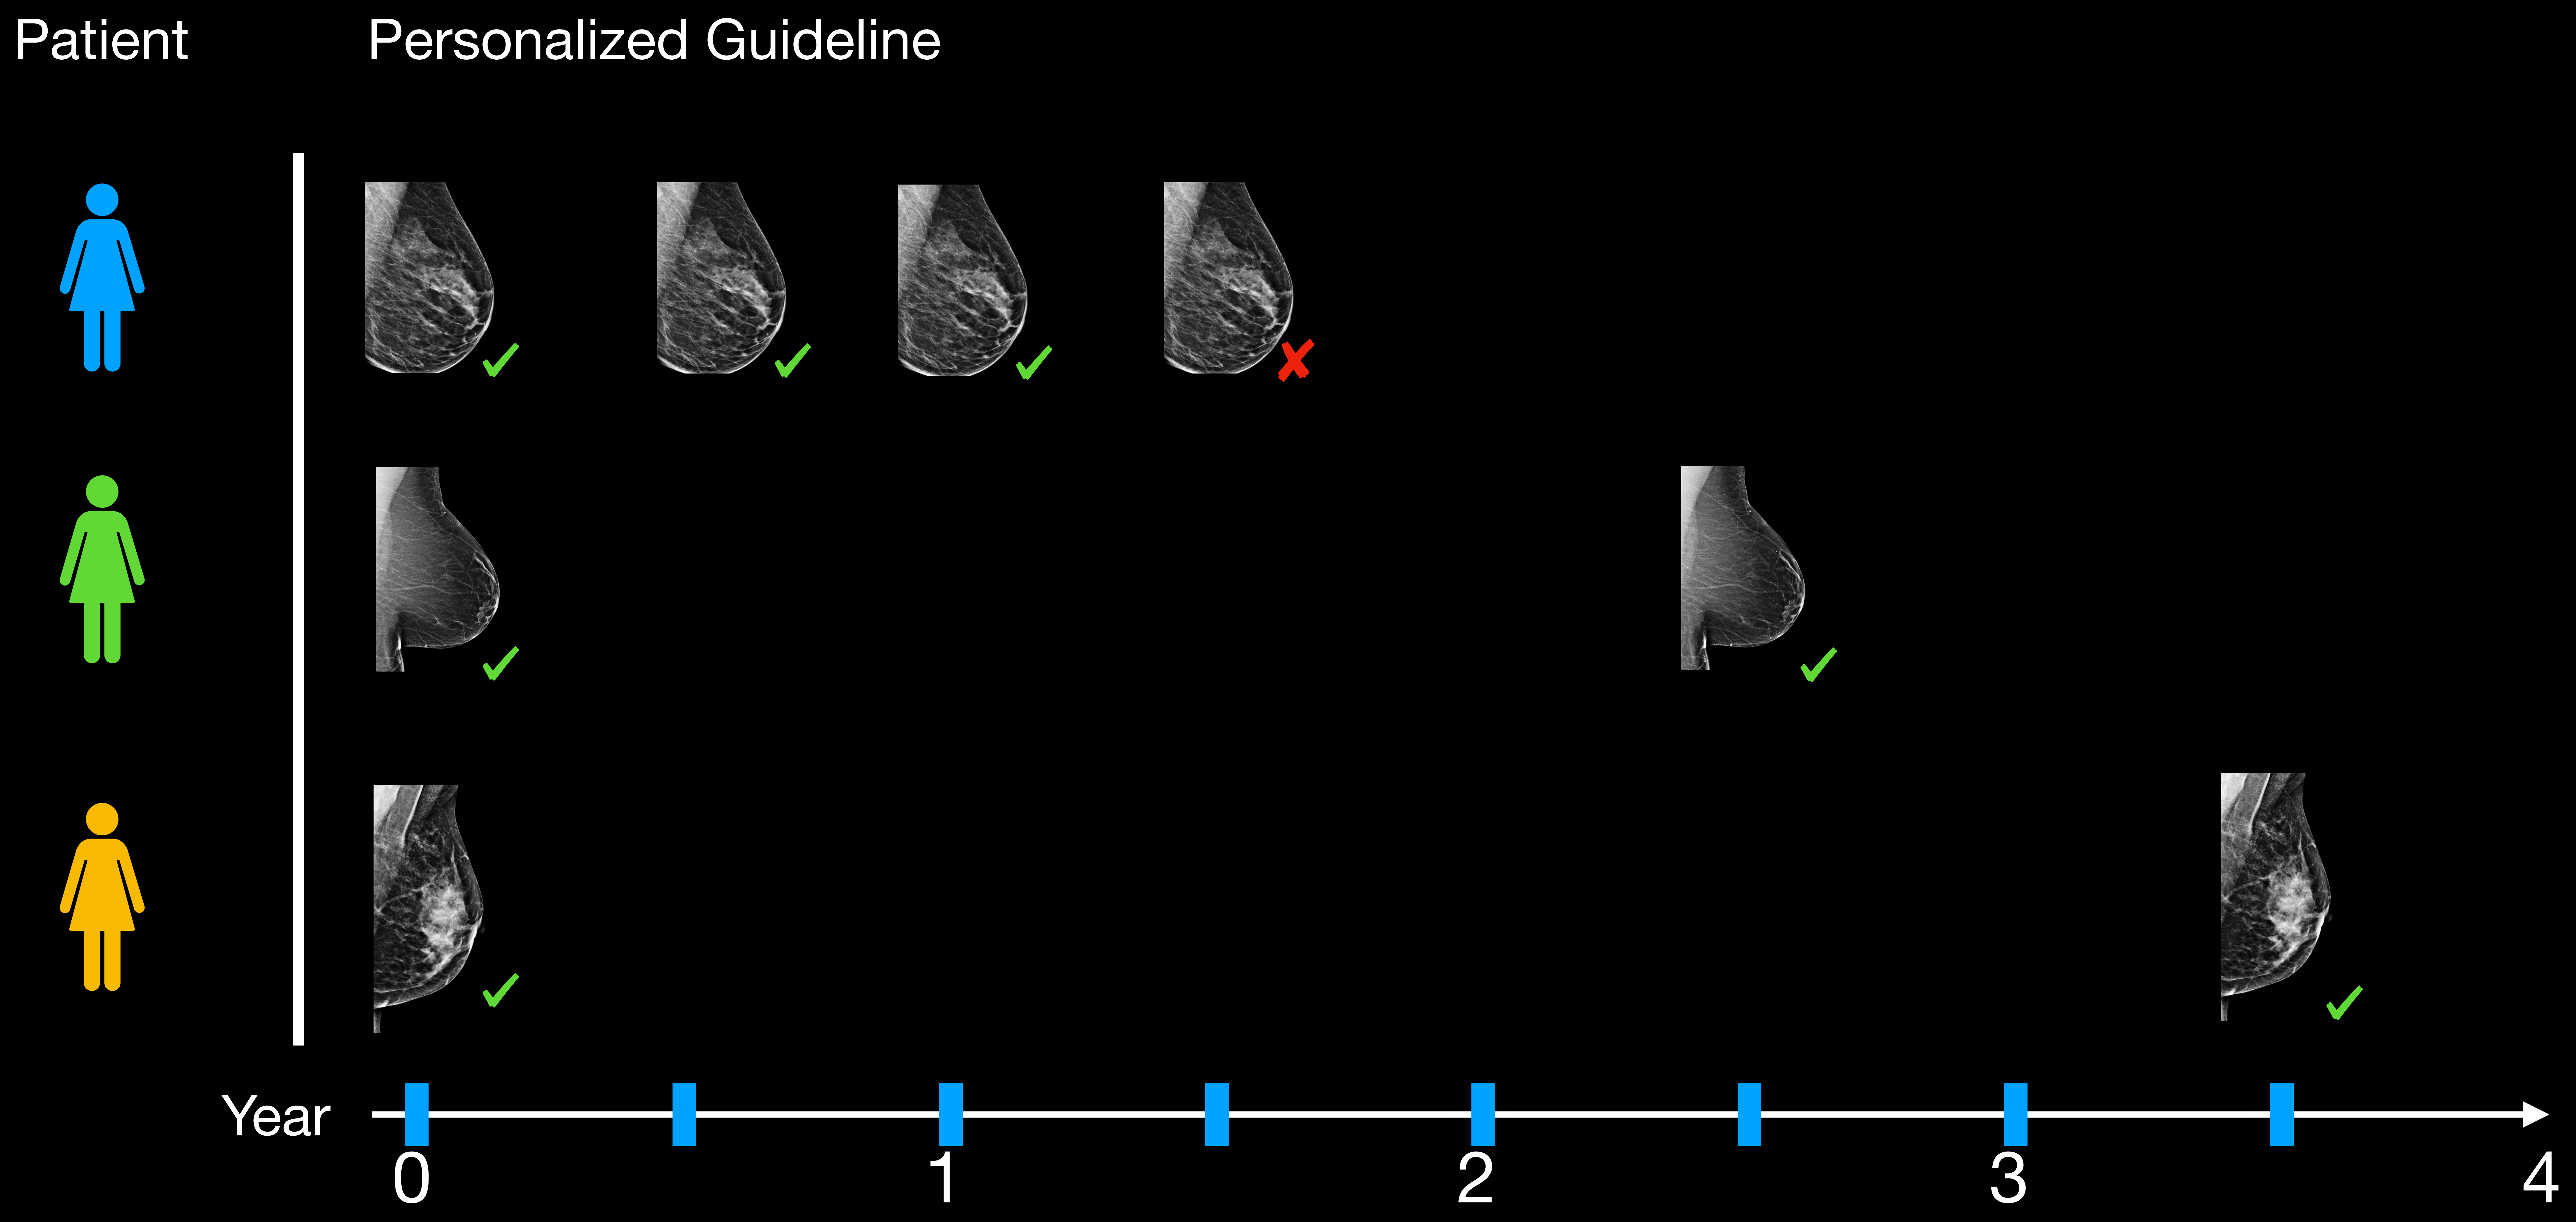

Current processes for policy design

NCCN National Comprehensive Cancer Network®

U.S. Preventive Services TASK FORCE

ACR® AMERICAN COLLEGE OF RADIOLOGY

...

Expert panel meetings by physician organizations

Most meet every five years

Multiple conflicting one-size fits all guidelines

No explicit validation across populations. Health disparities grow

Modeling for clinical guidelines

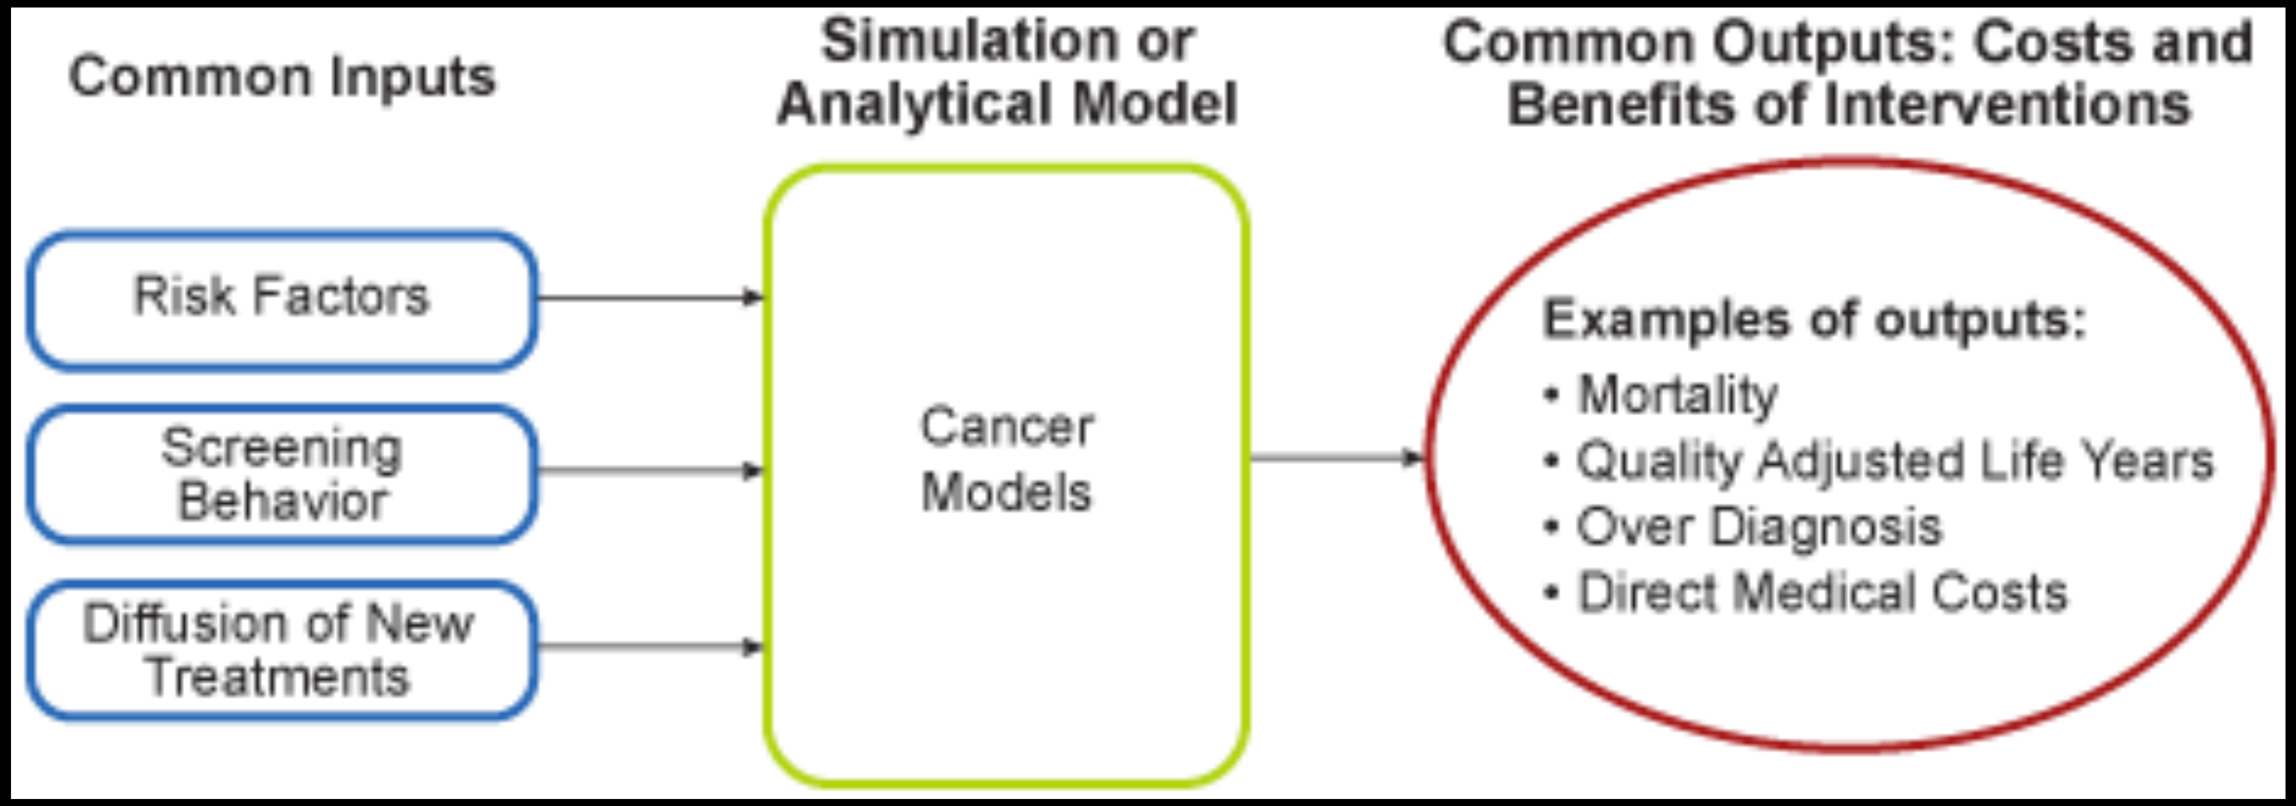

Assume full probabilistic models of disease

Simulate hypothetical patients under screening strategies

Suggest “correct” trade-off

Cannot incorporate new risk models or evaluated on real patients

Tailor screening regime to patient need

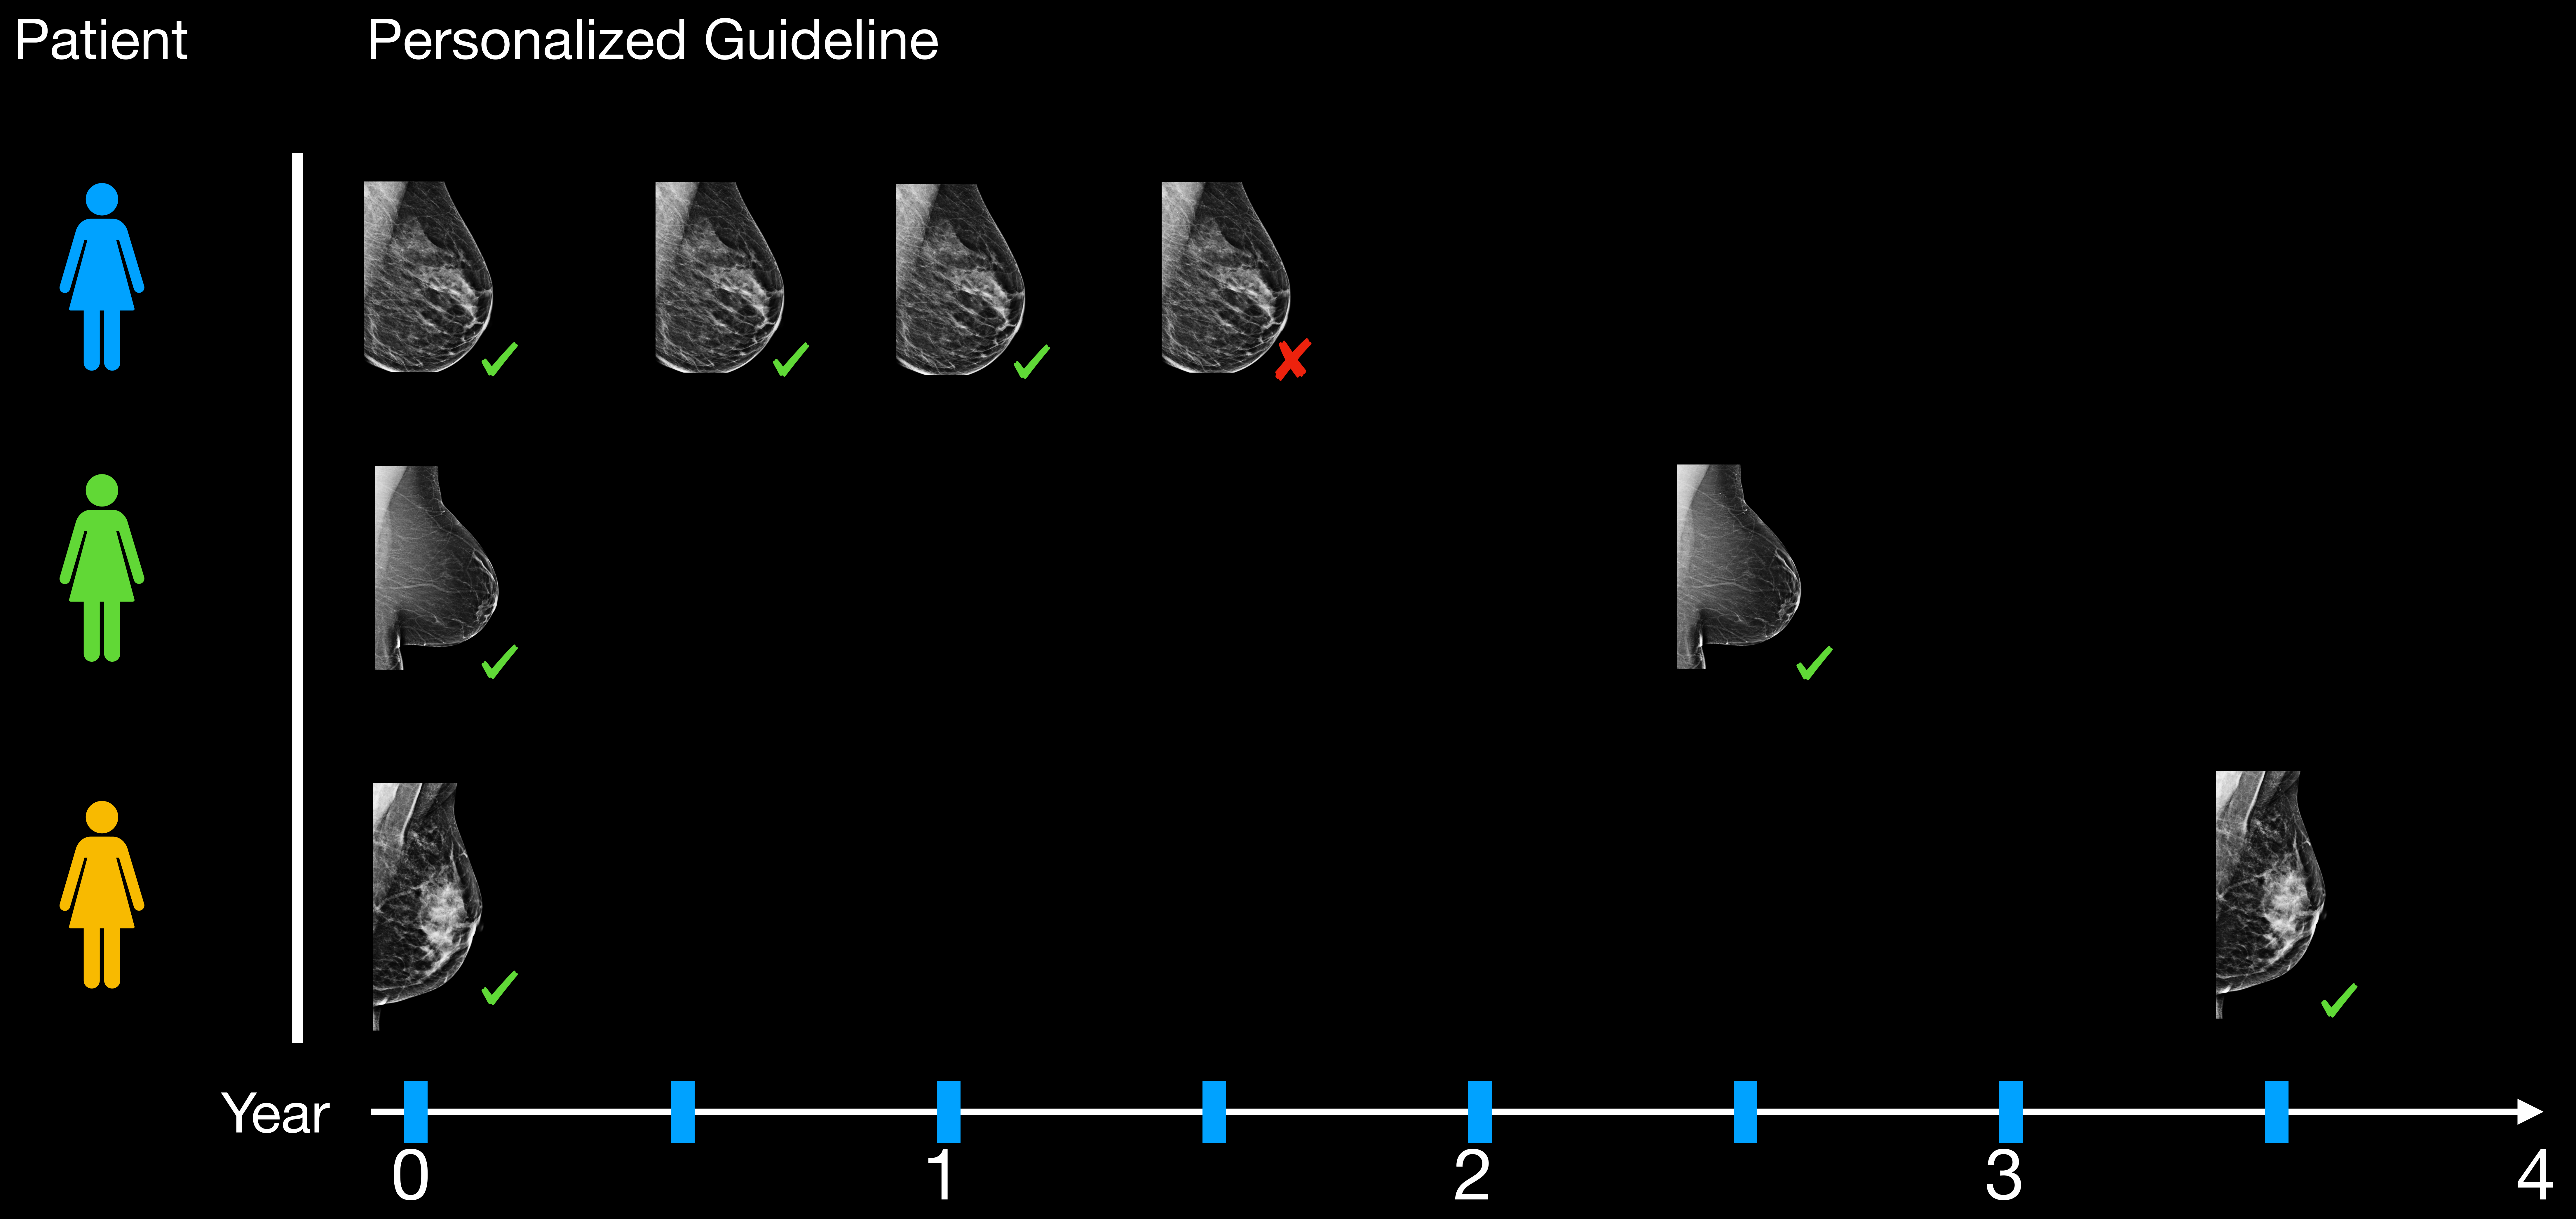

Policy Design as a Learning Task

Patient **Risk**

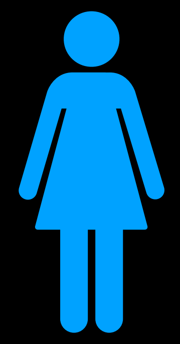 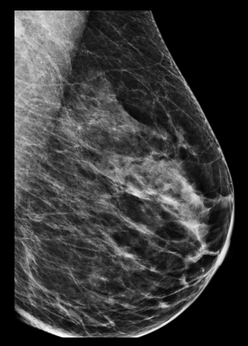 → 0.80 → Tempo 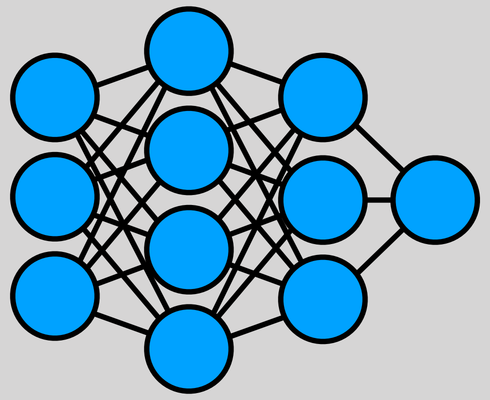 → **0.5 year followup**

Reward = λ_1 Early Detection Benefit - λ_2Screening Cost

Desiderata: *Testable, Adaptive, Flexible*

Simulate patient trajectories

Evaluate individual impact

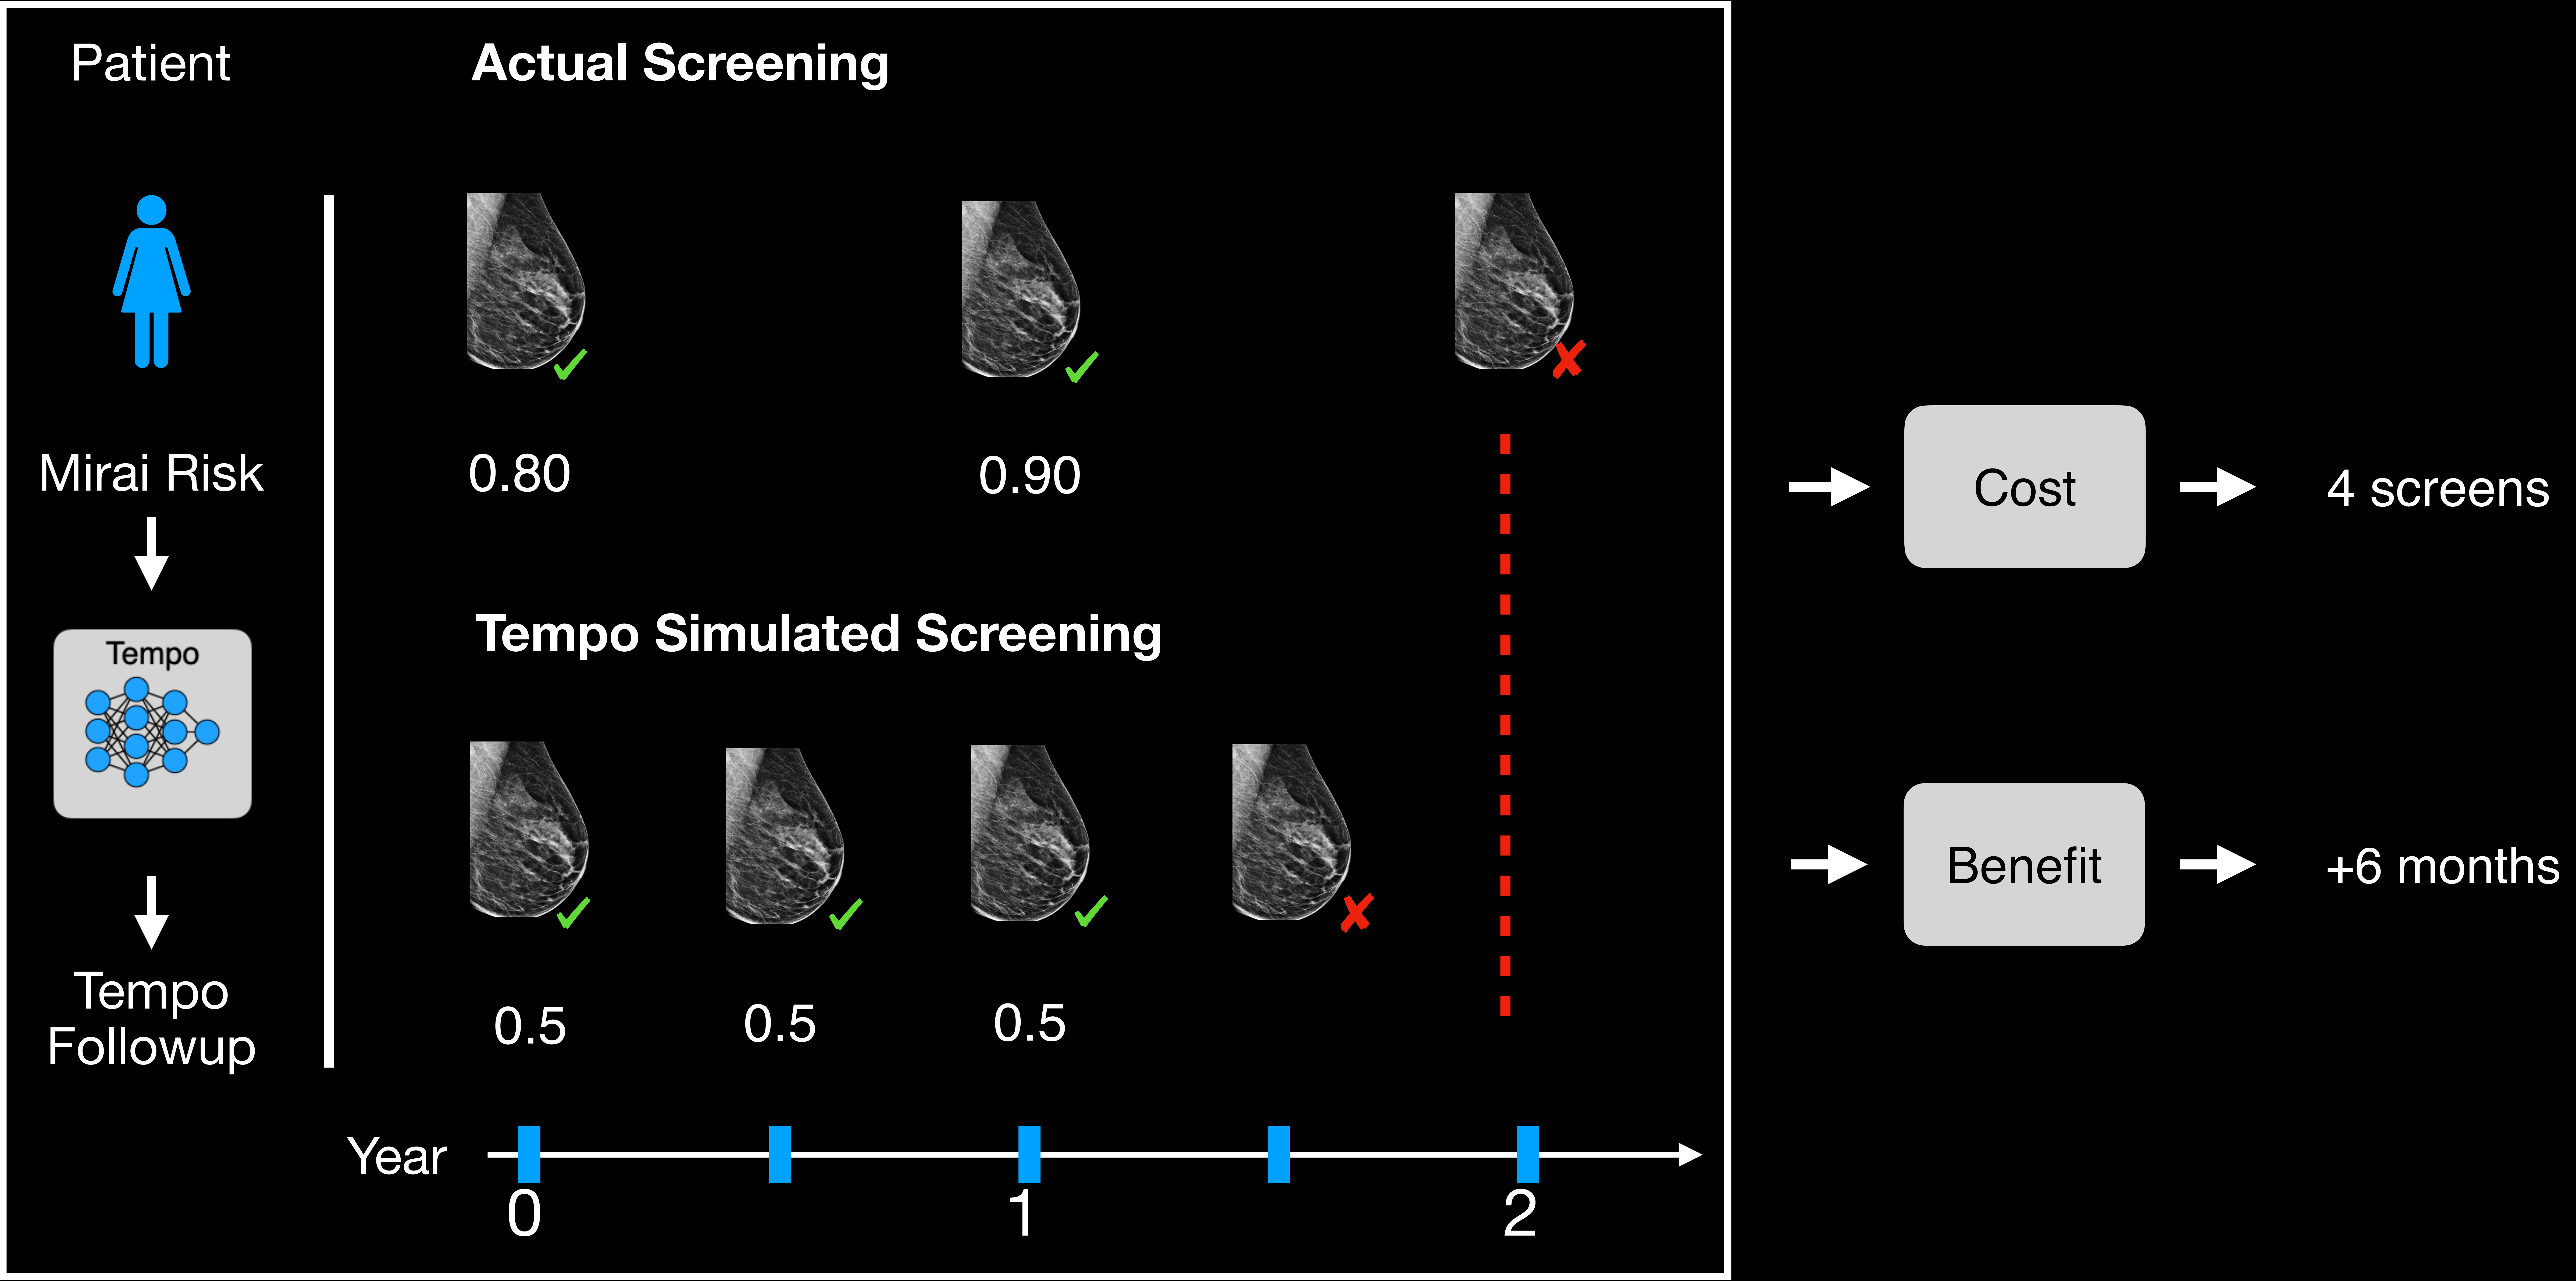

Evaluate individual impact

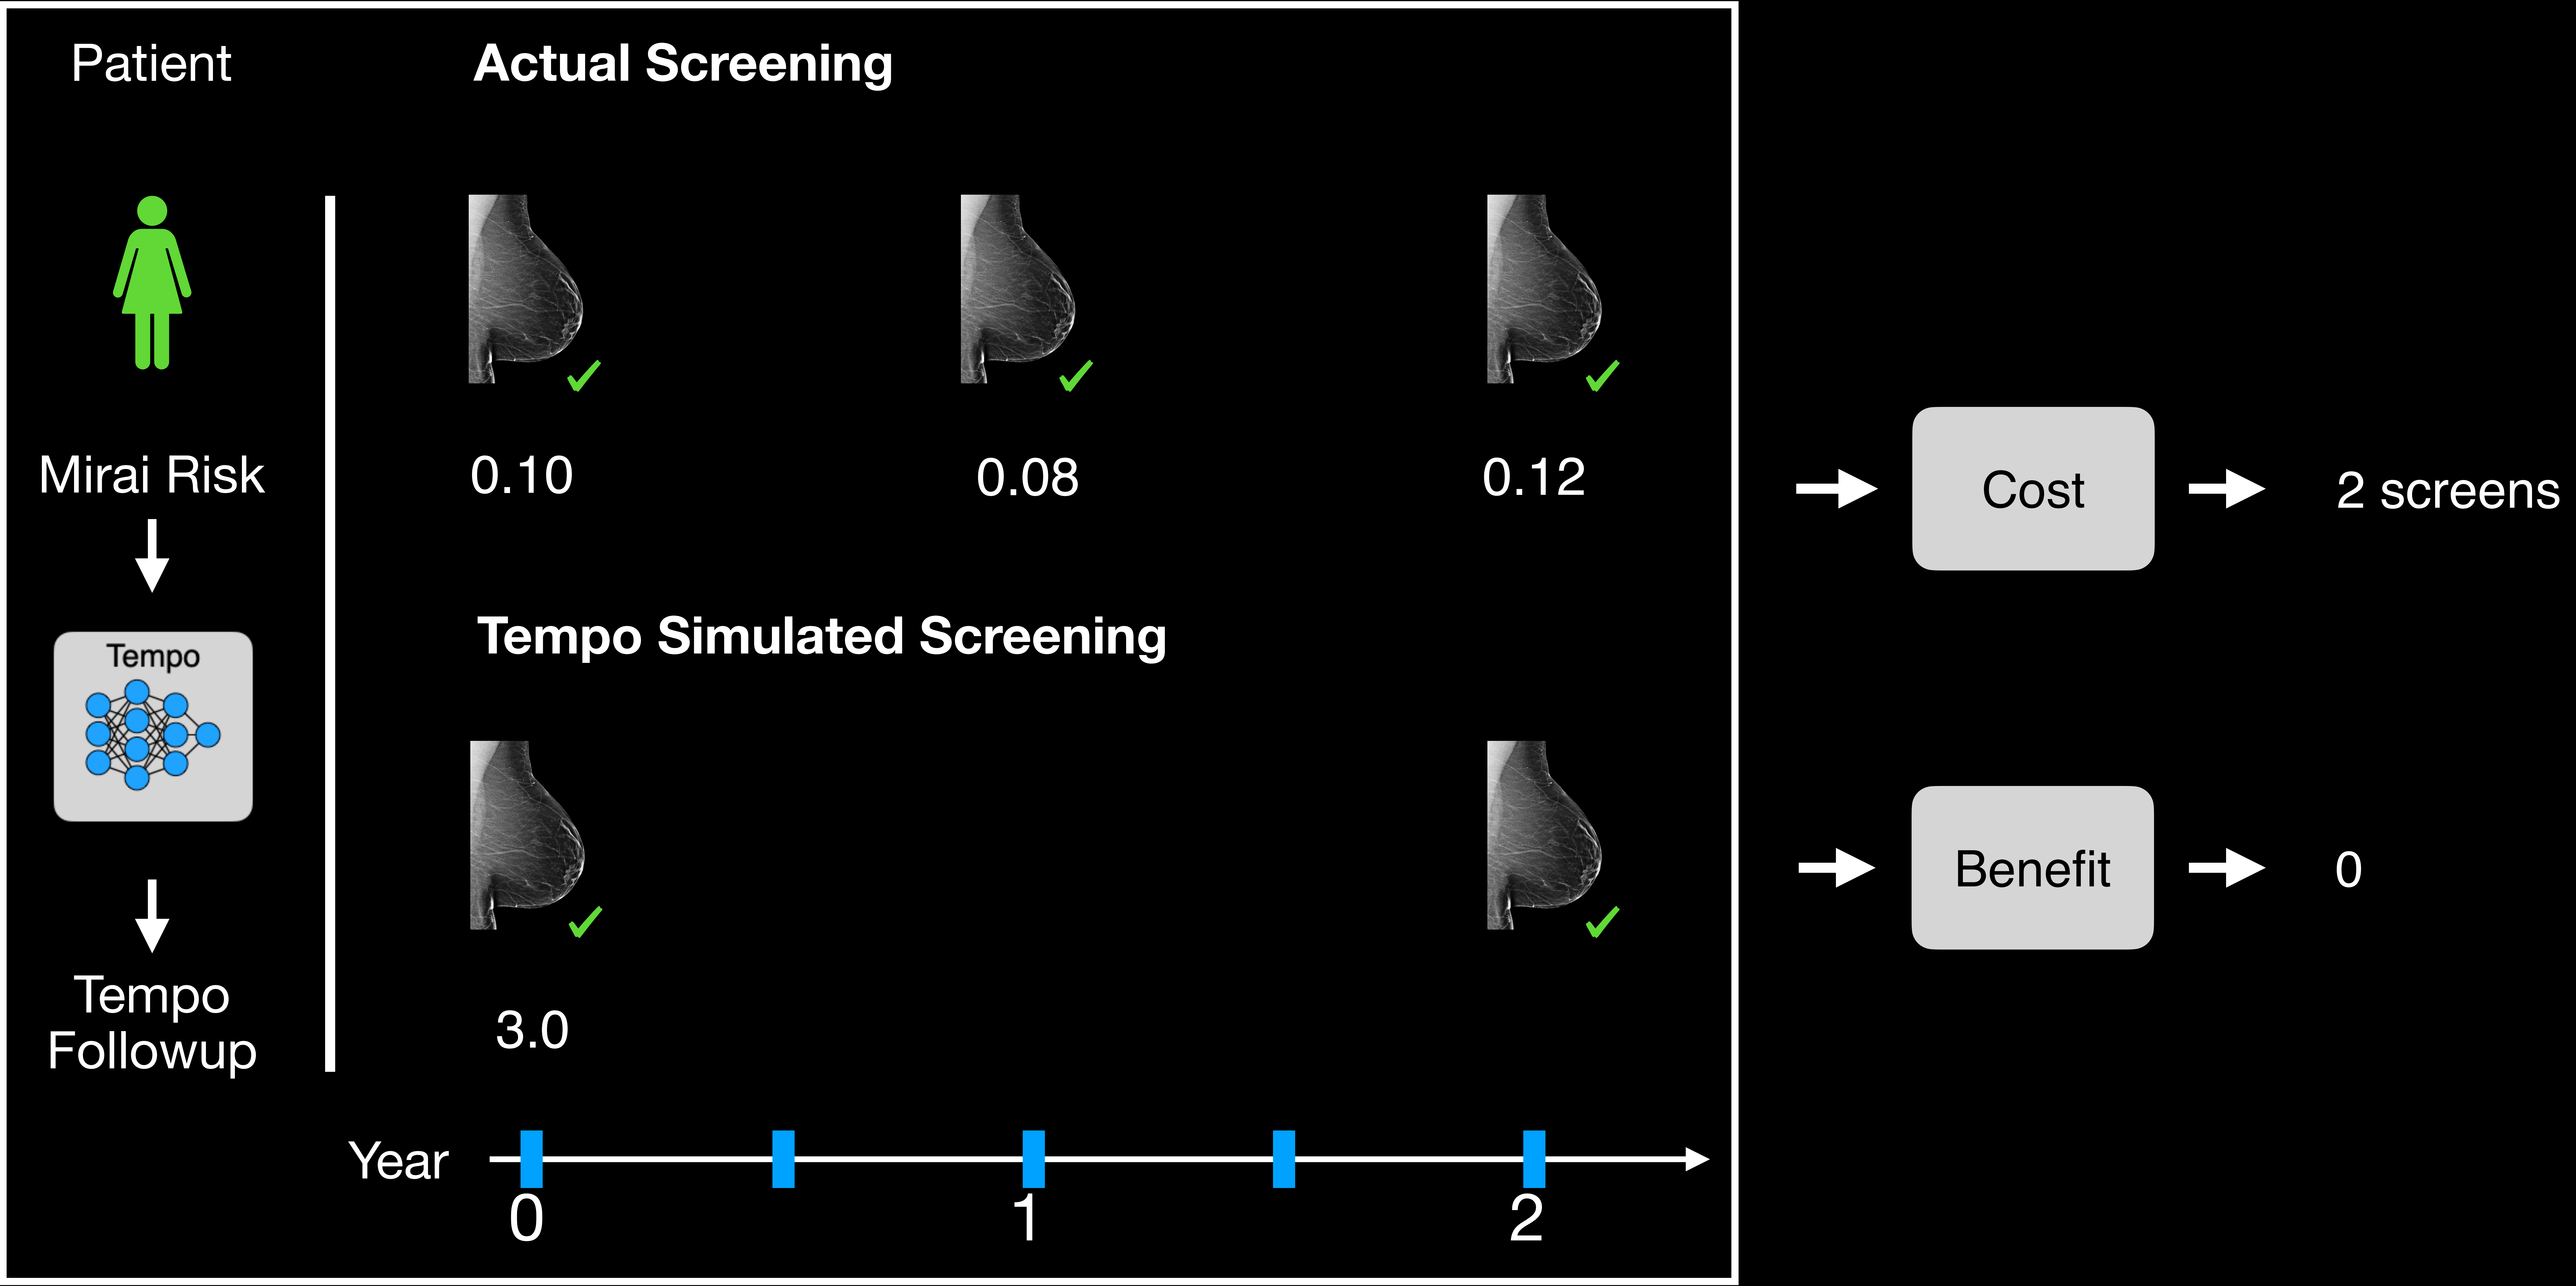

Optimized over population screening

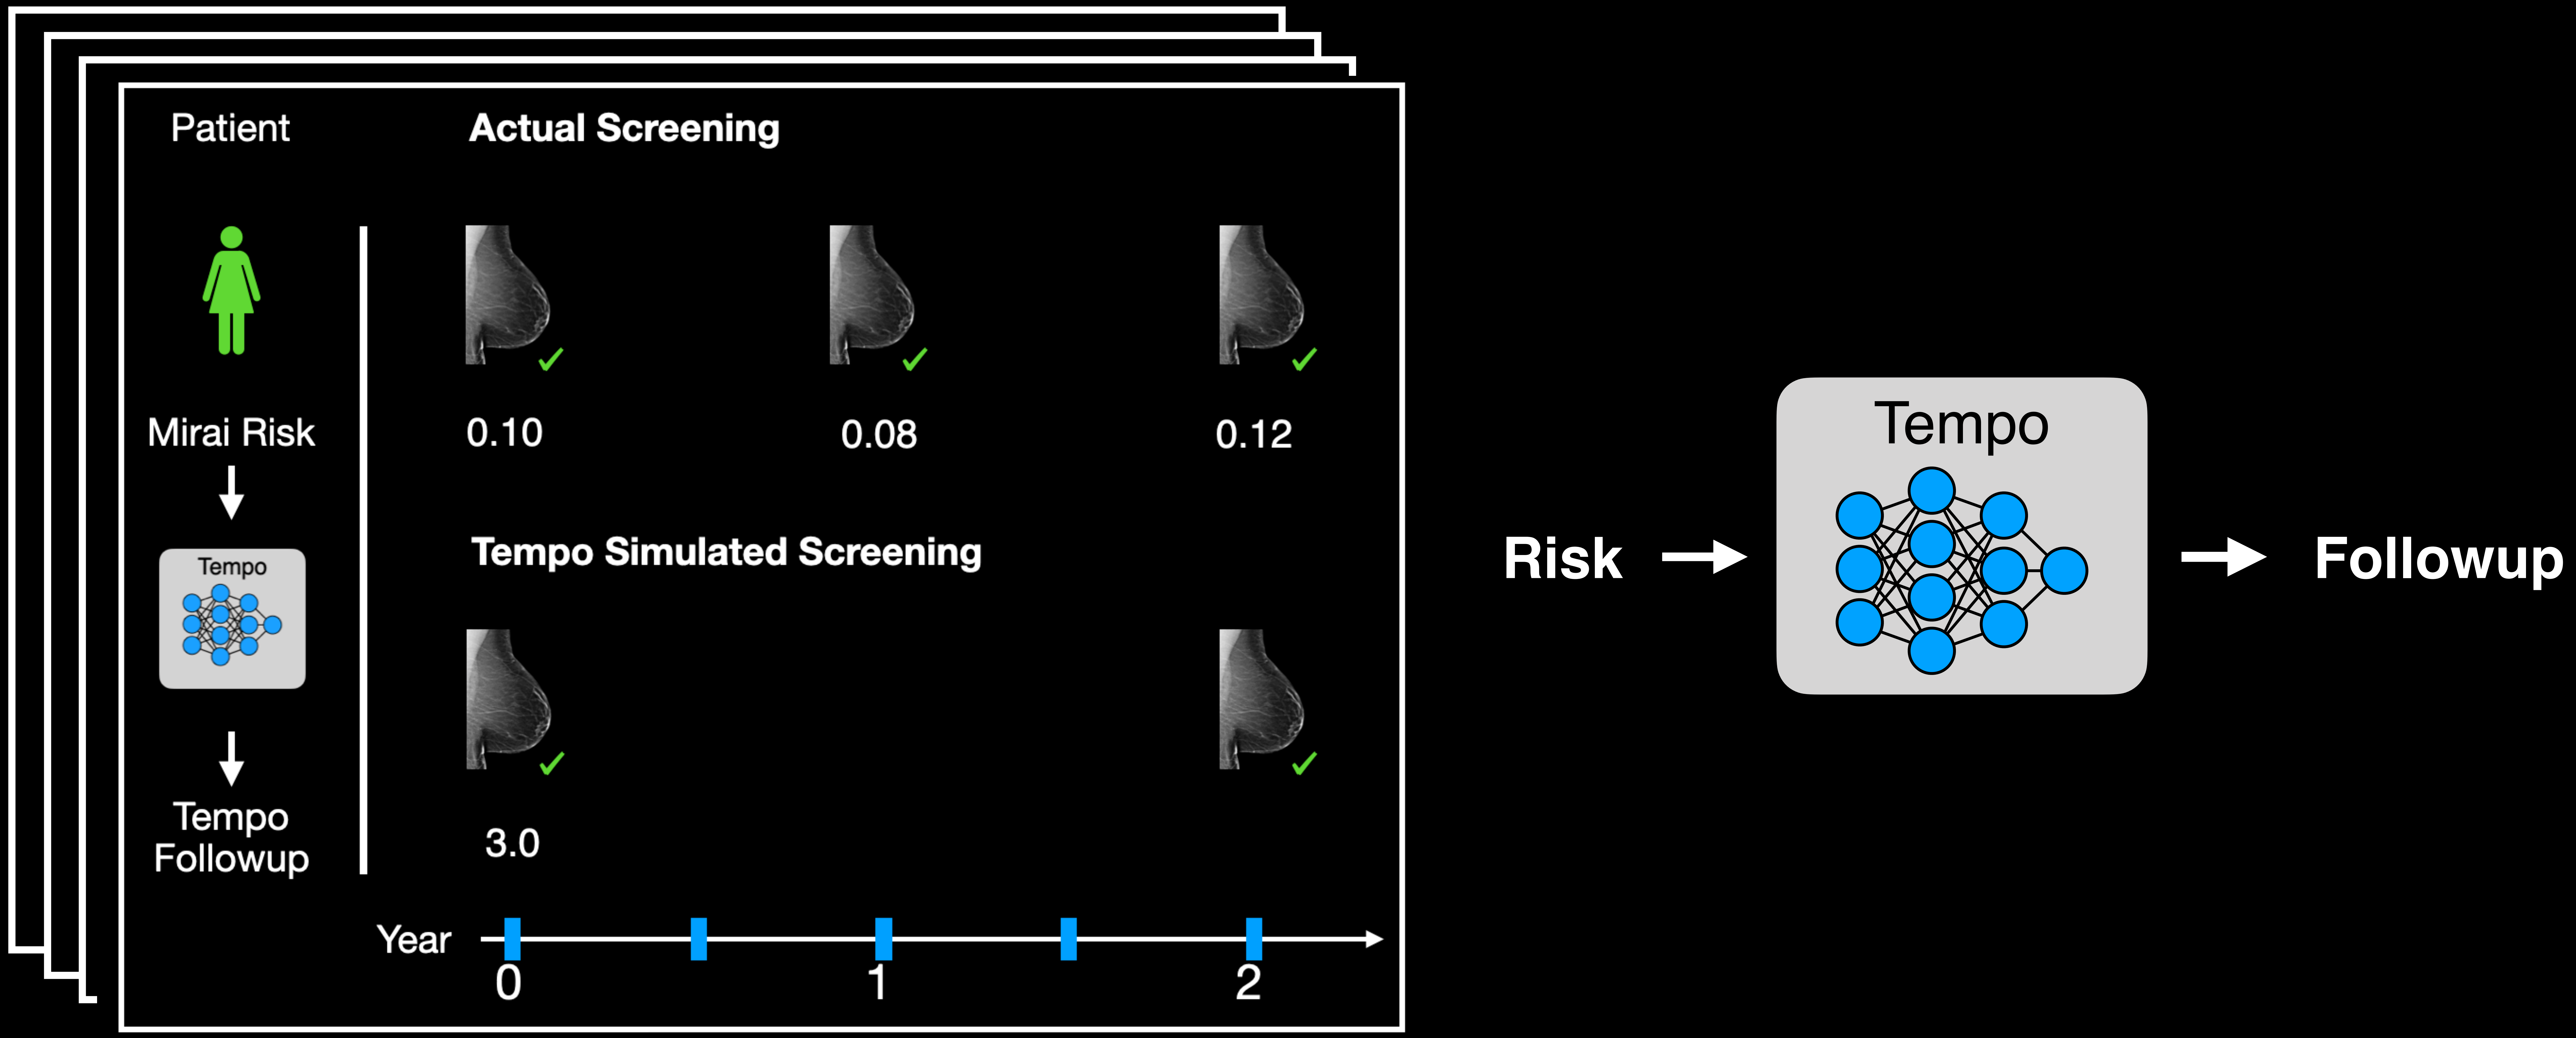

Estimating missing risk assessments

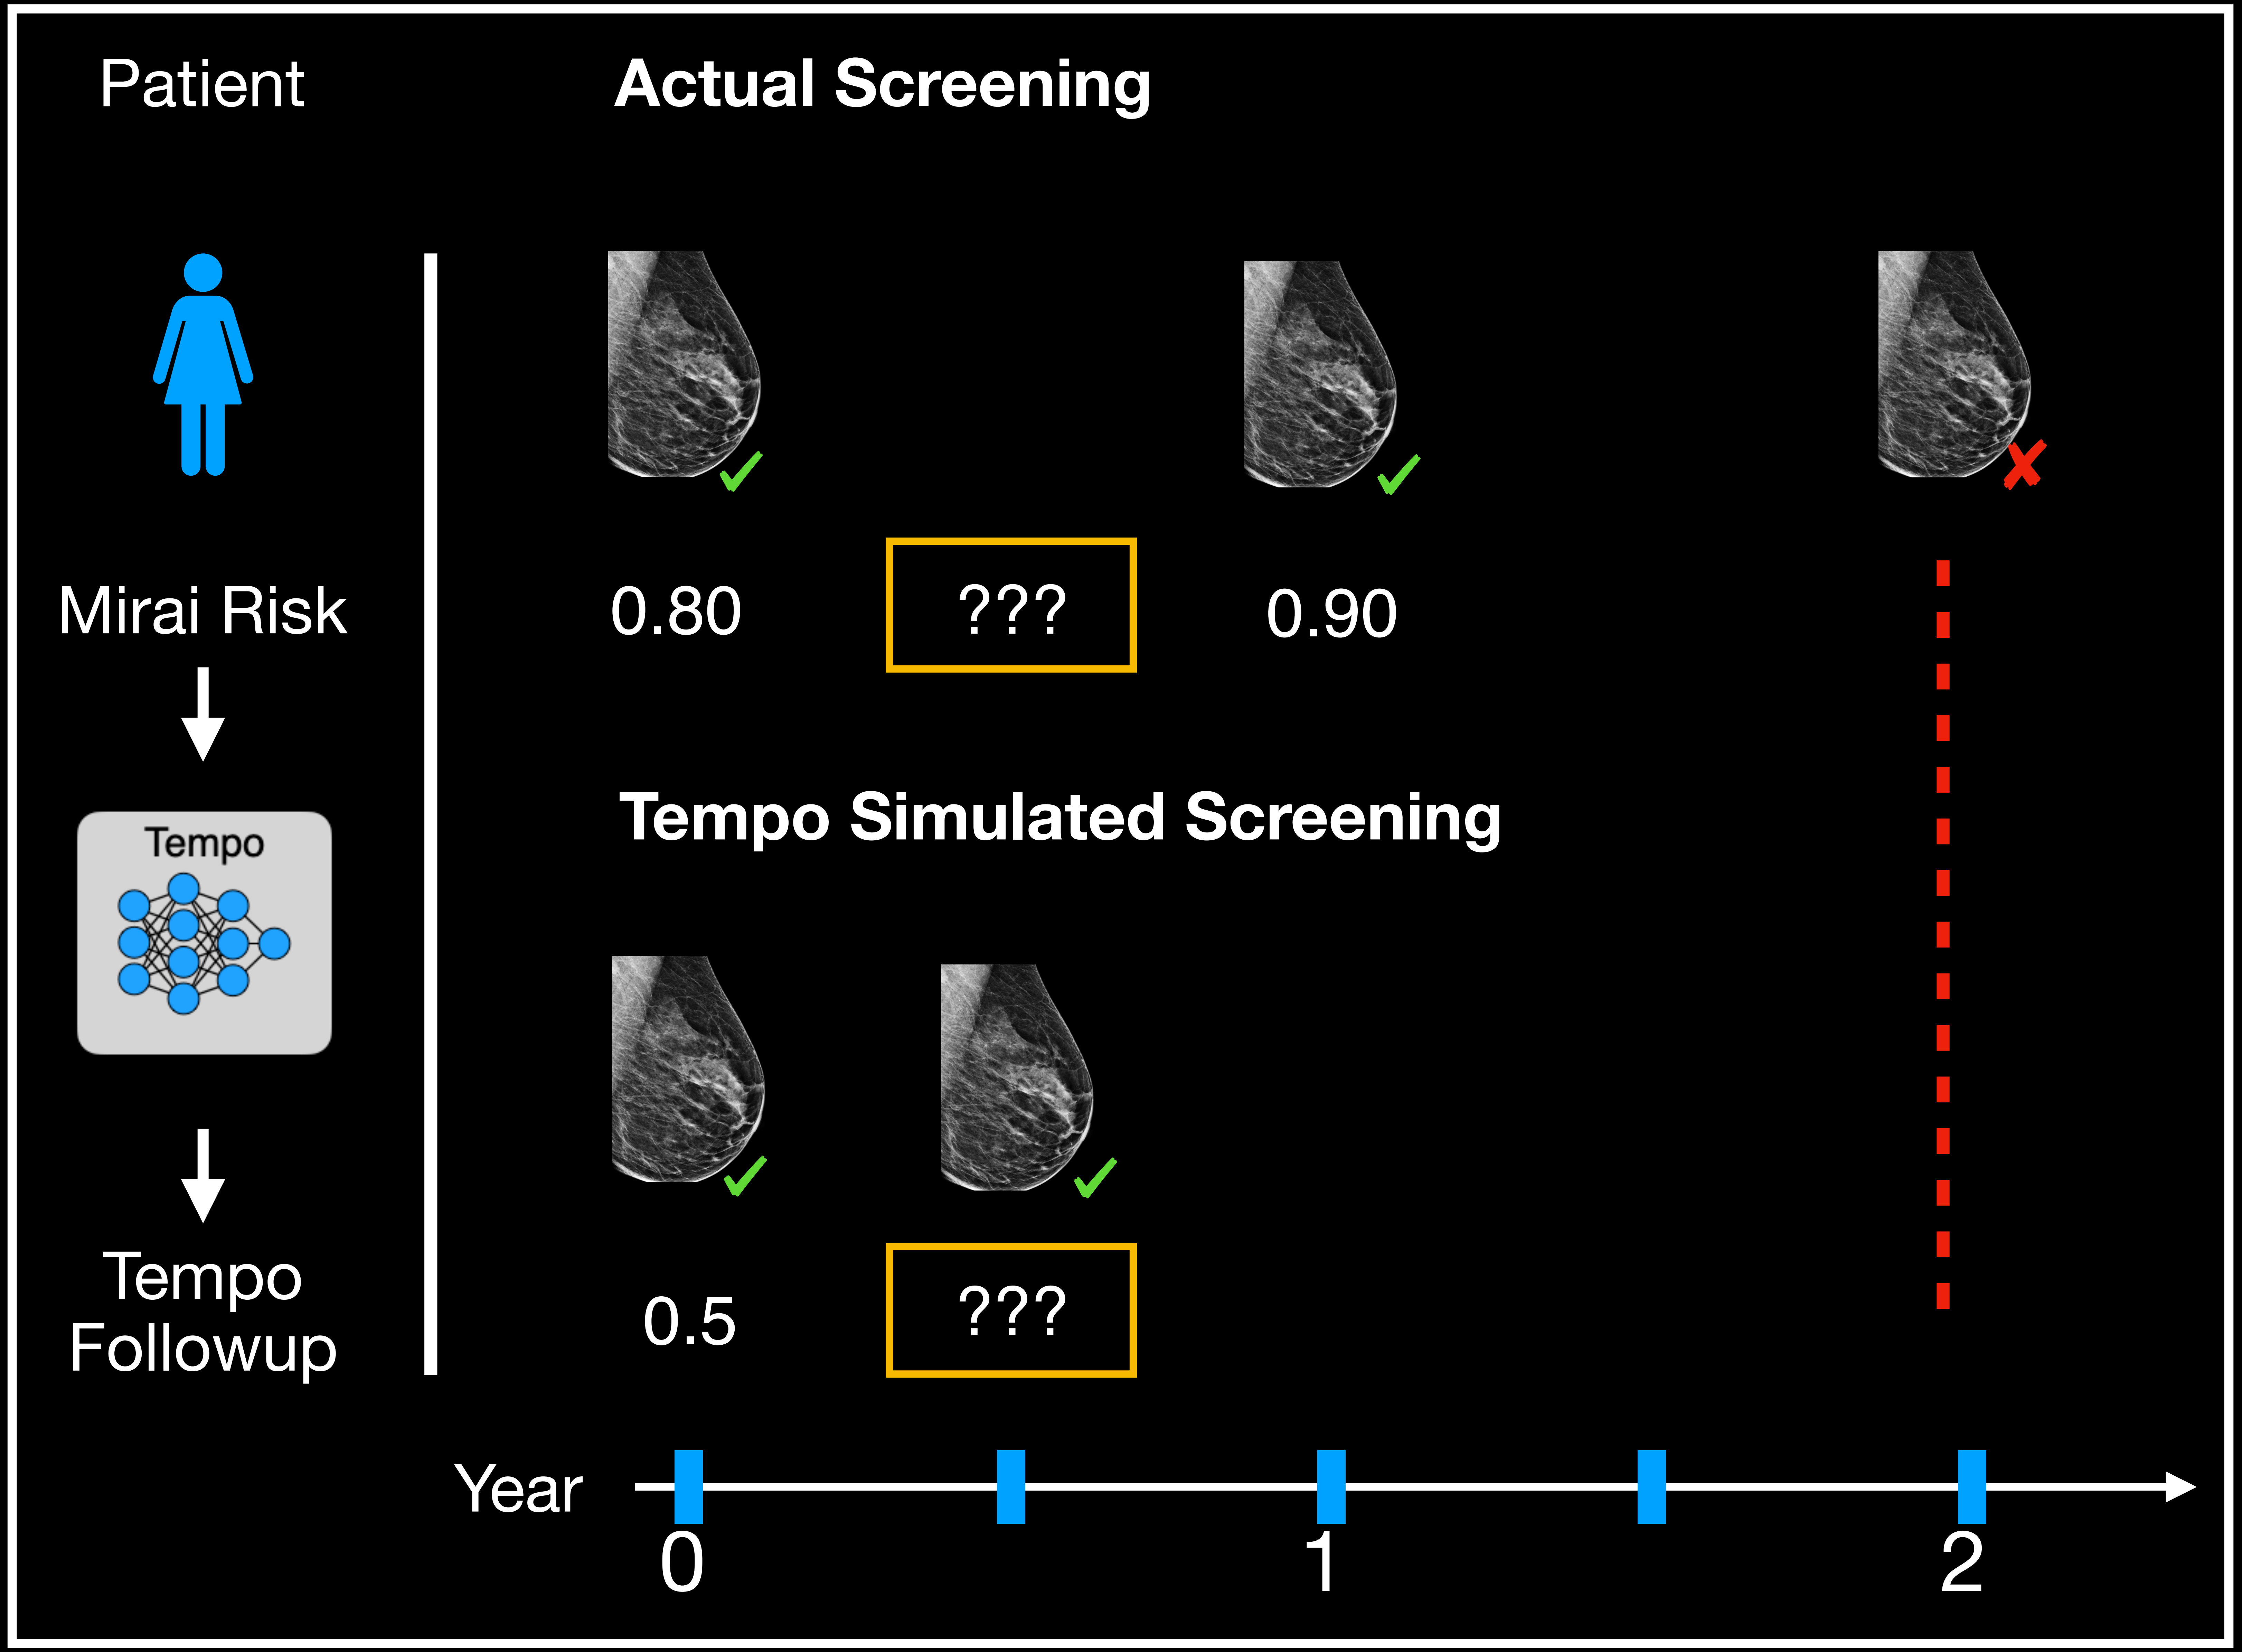

Learn $P(r_t \mid r_{t-1}, r_{t-2}, \ldots, r_0)$

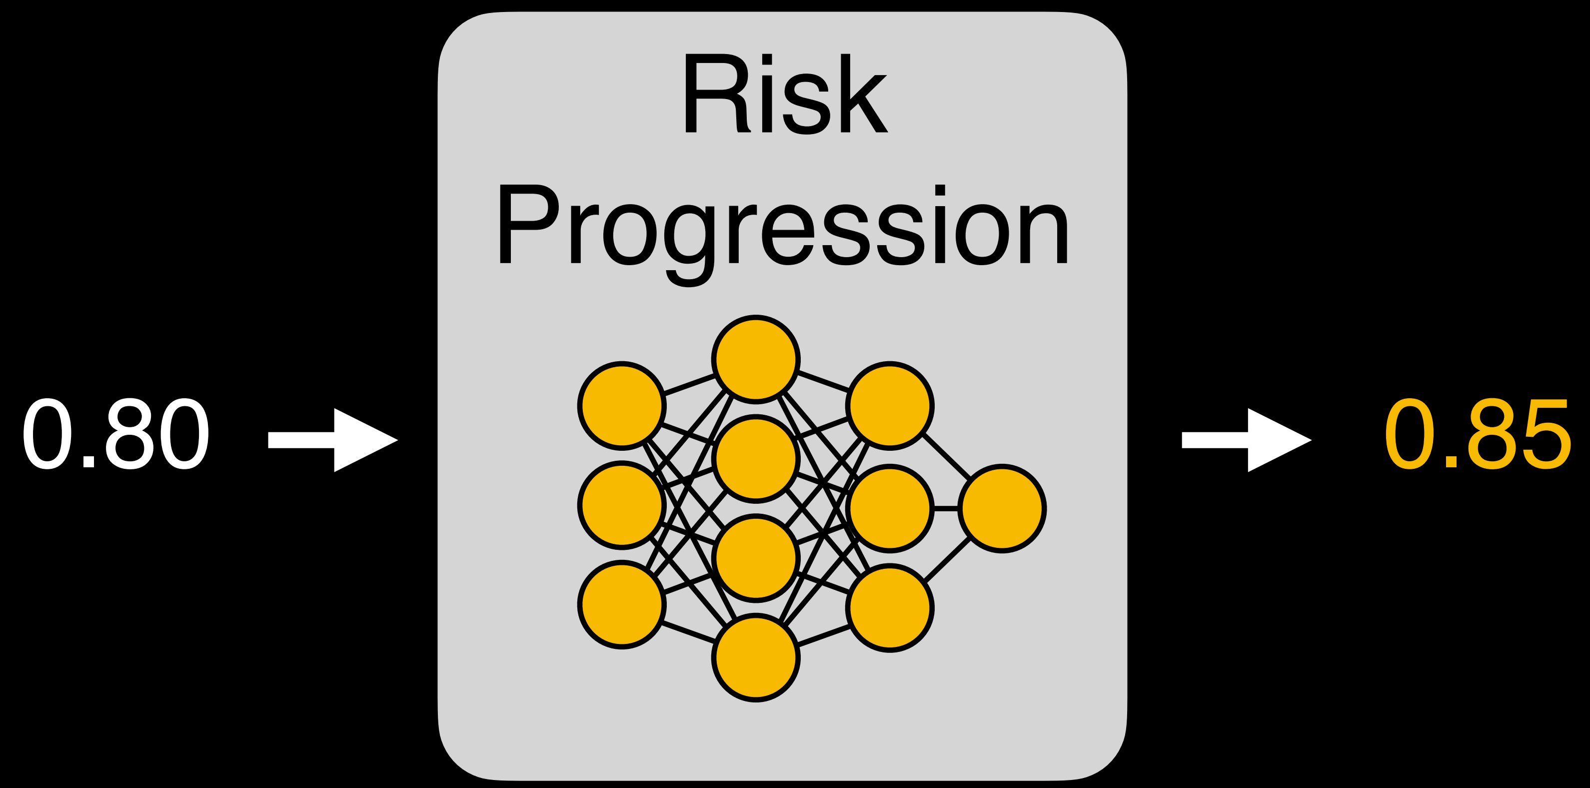

Estimating missing risk assessments

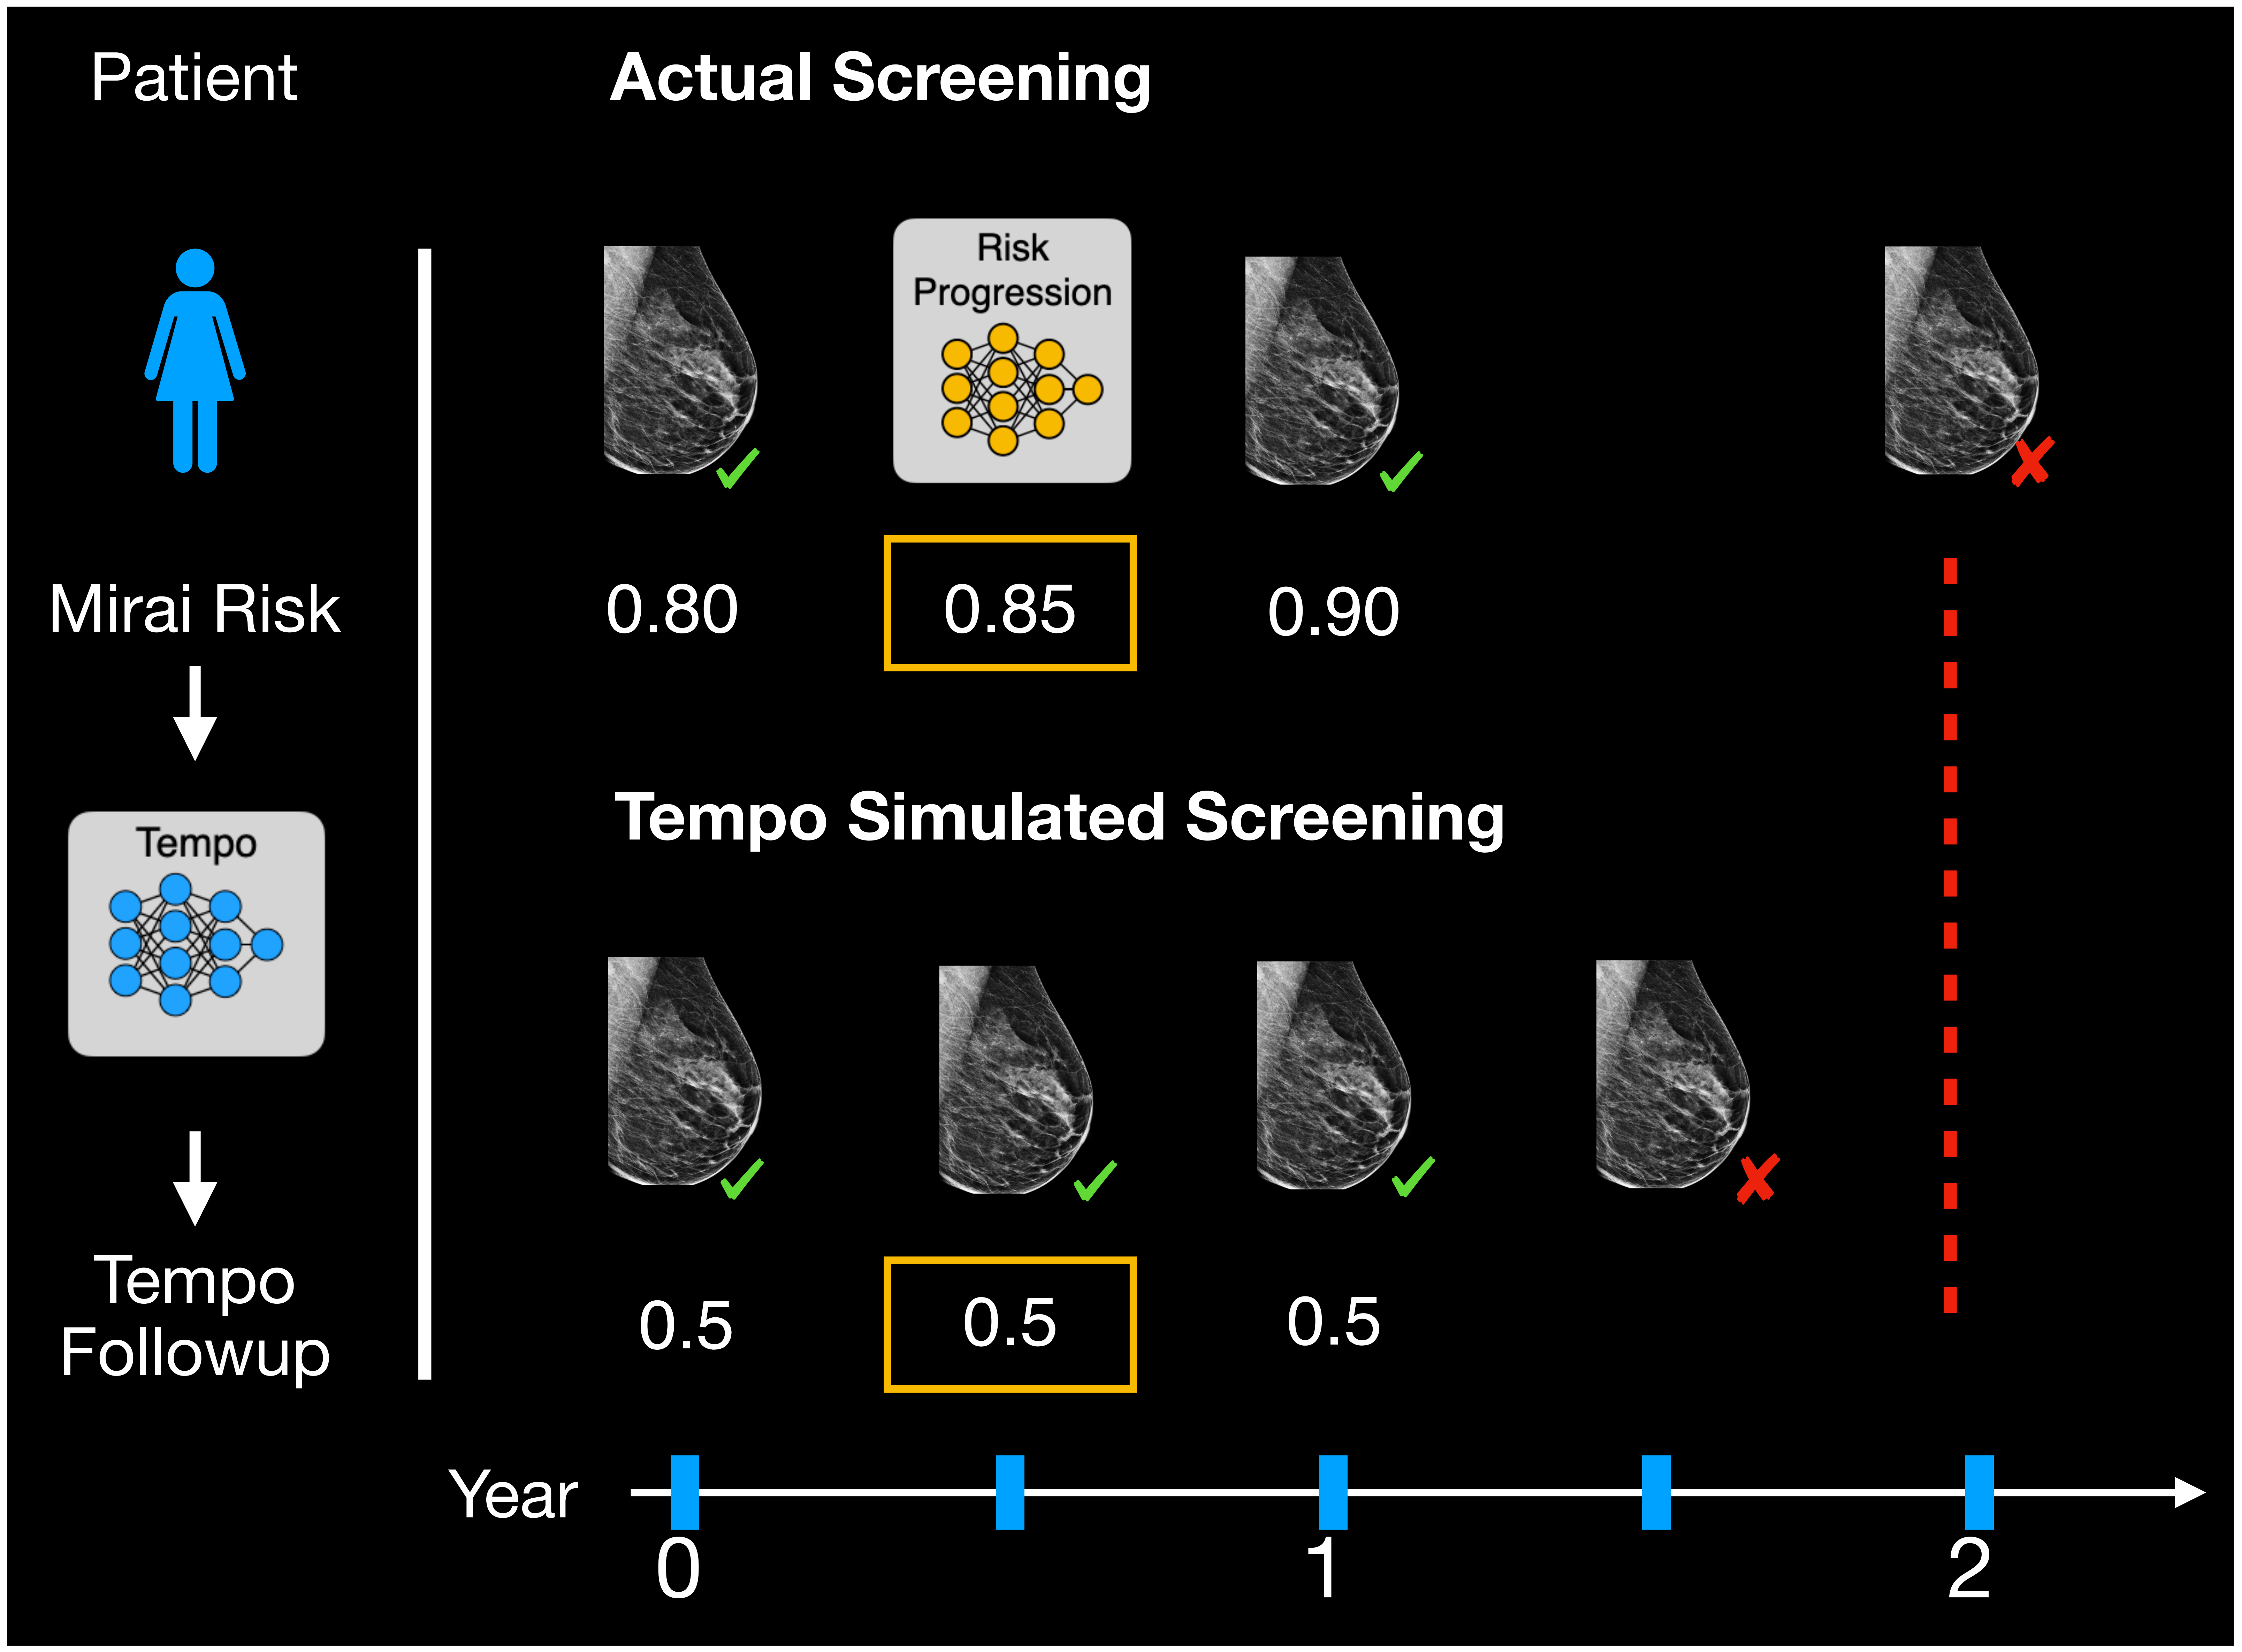

Learn $P(r_t \,|\, r_{t-1}, r_{t-2}, \ldots, r_0)$

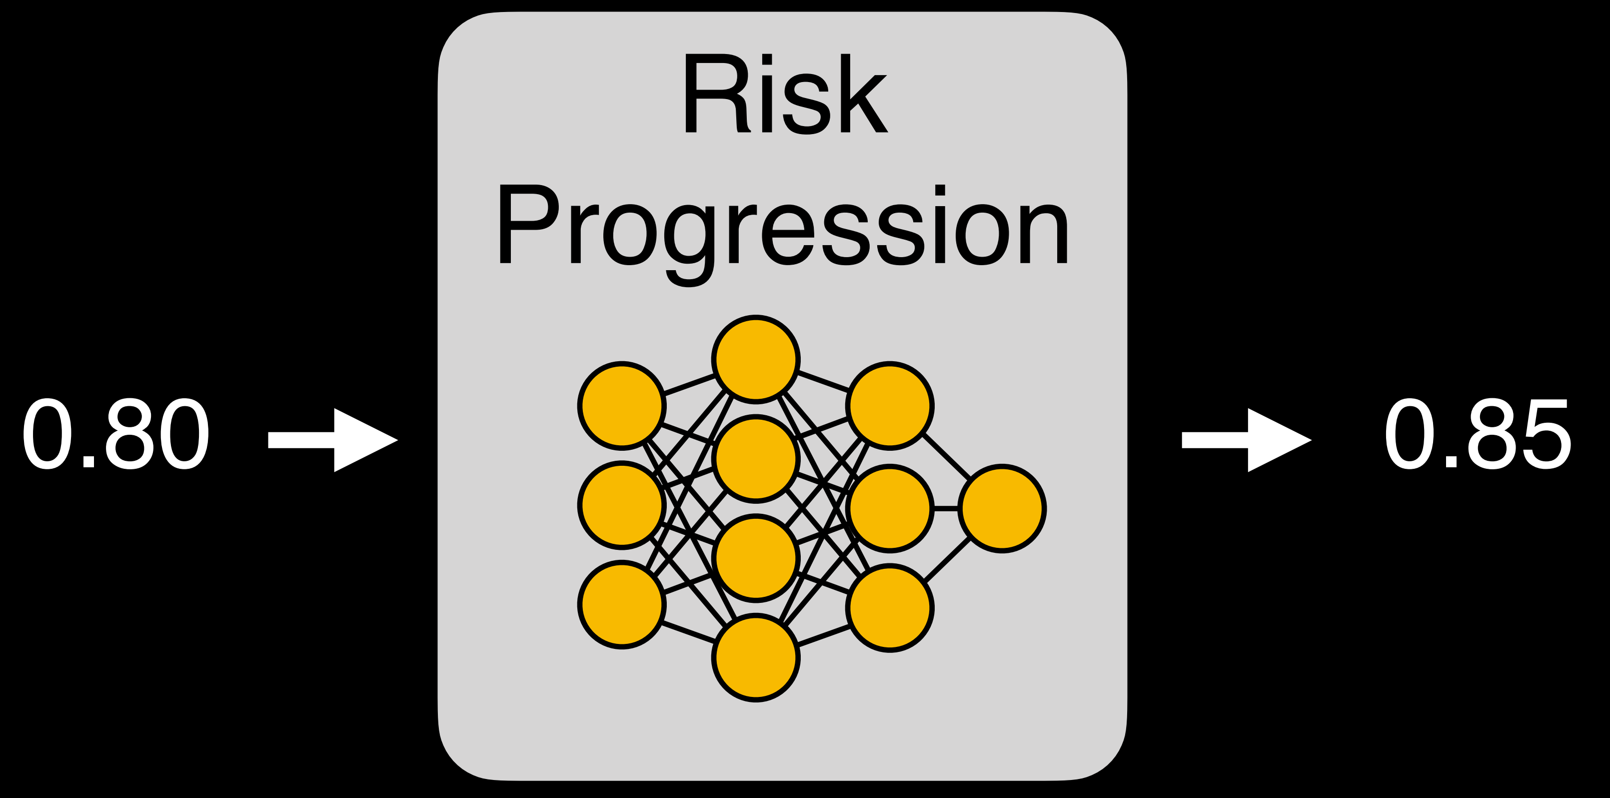

Baseline: Policy Design as Imitation

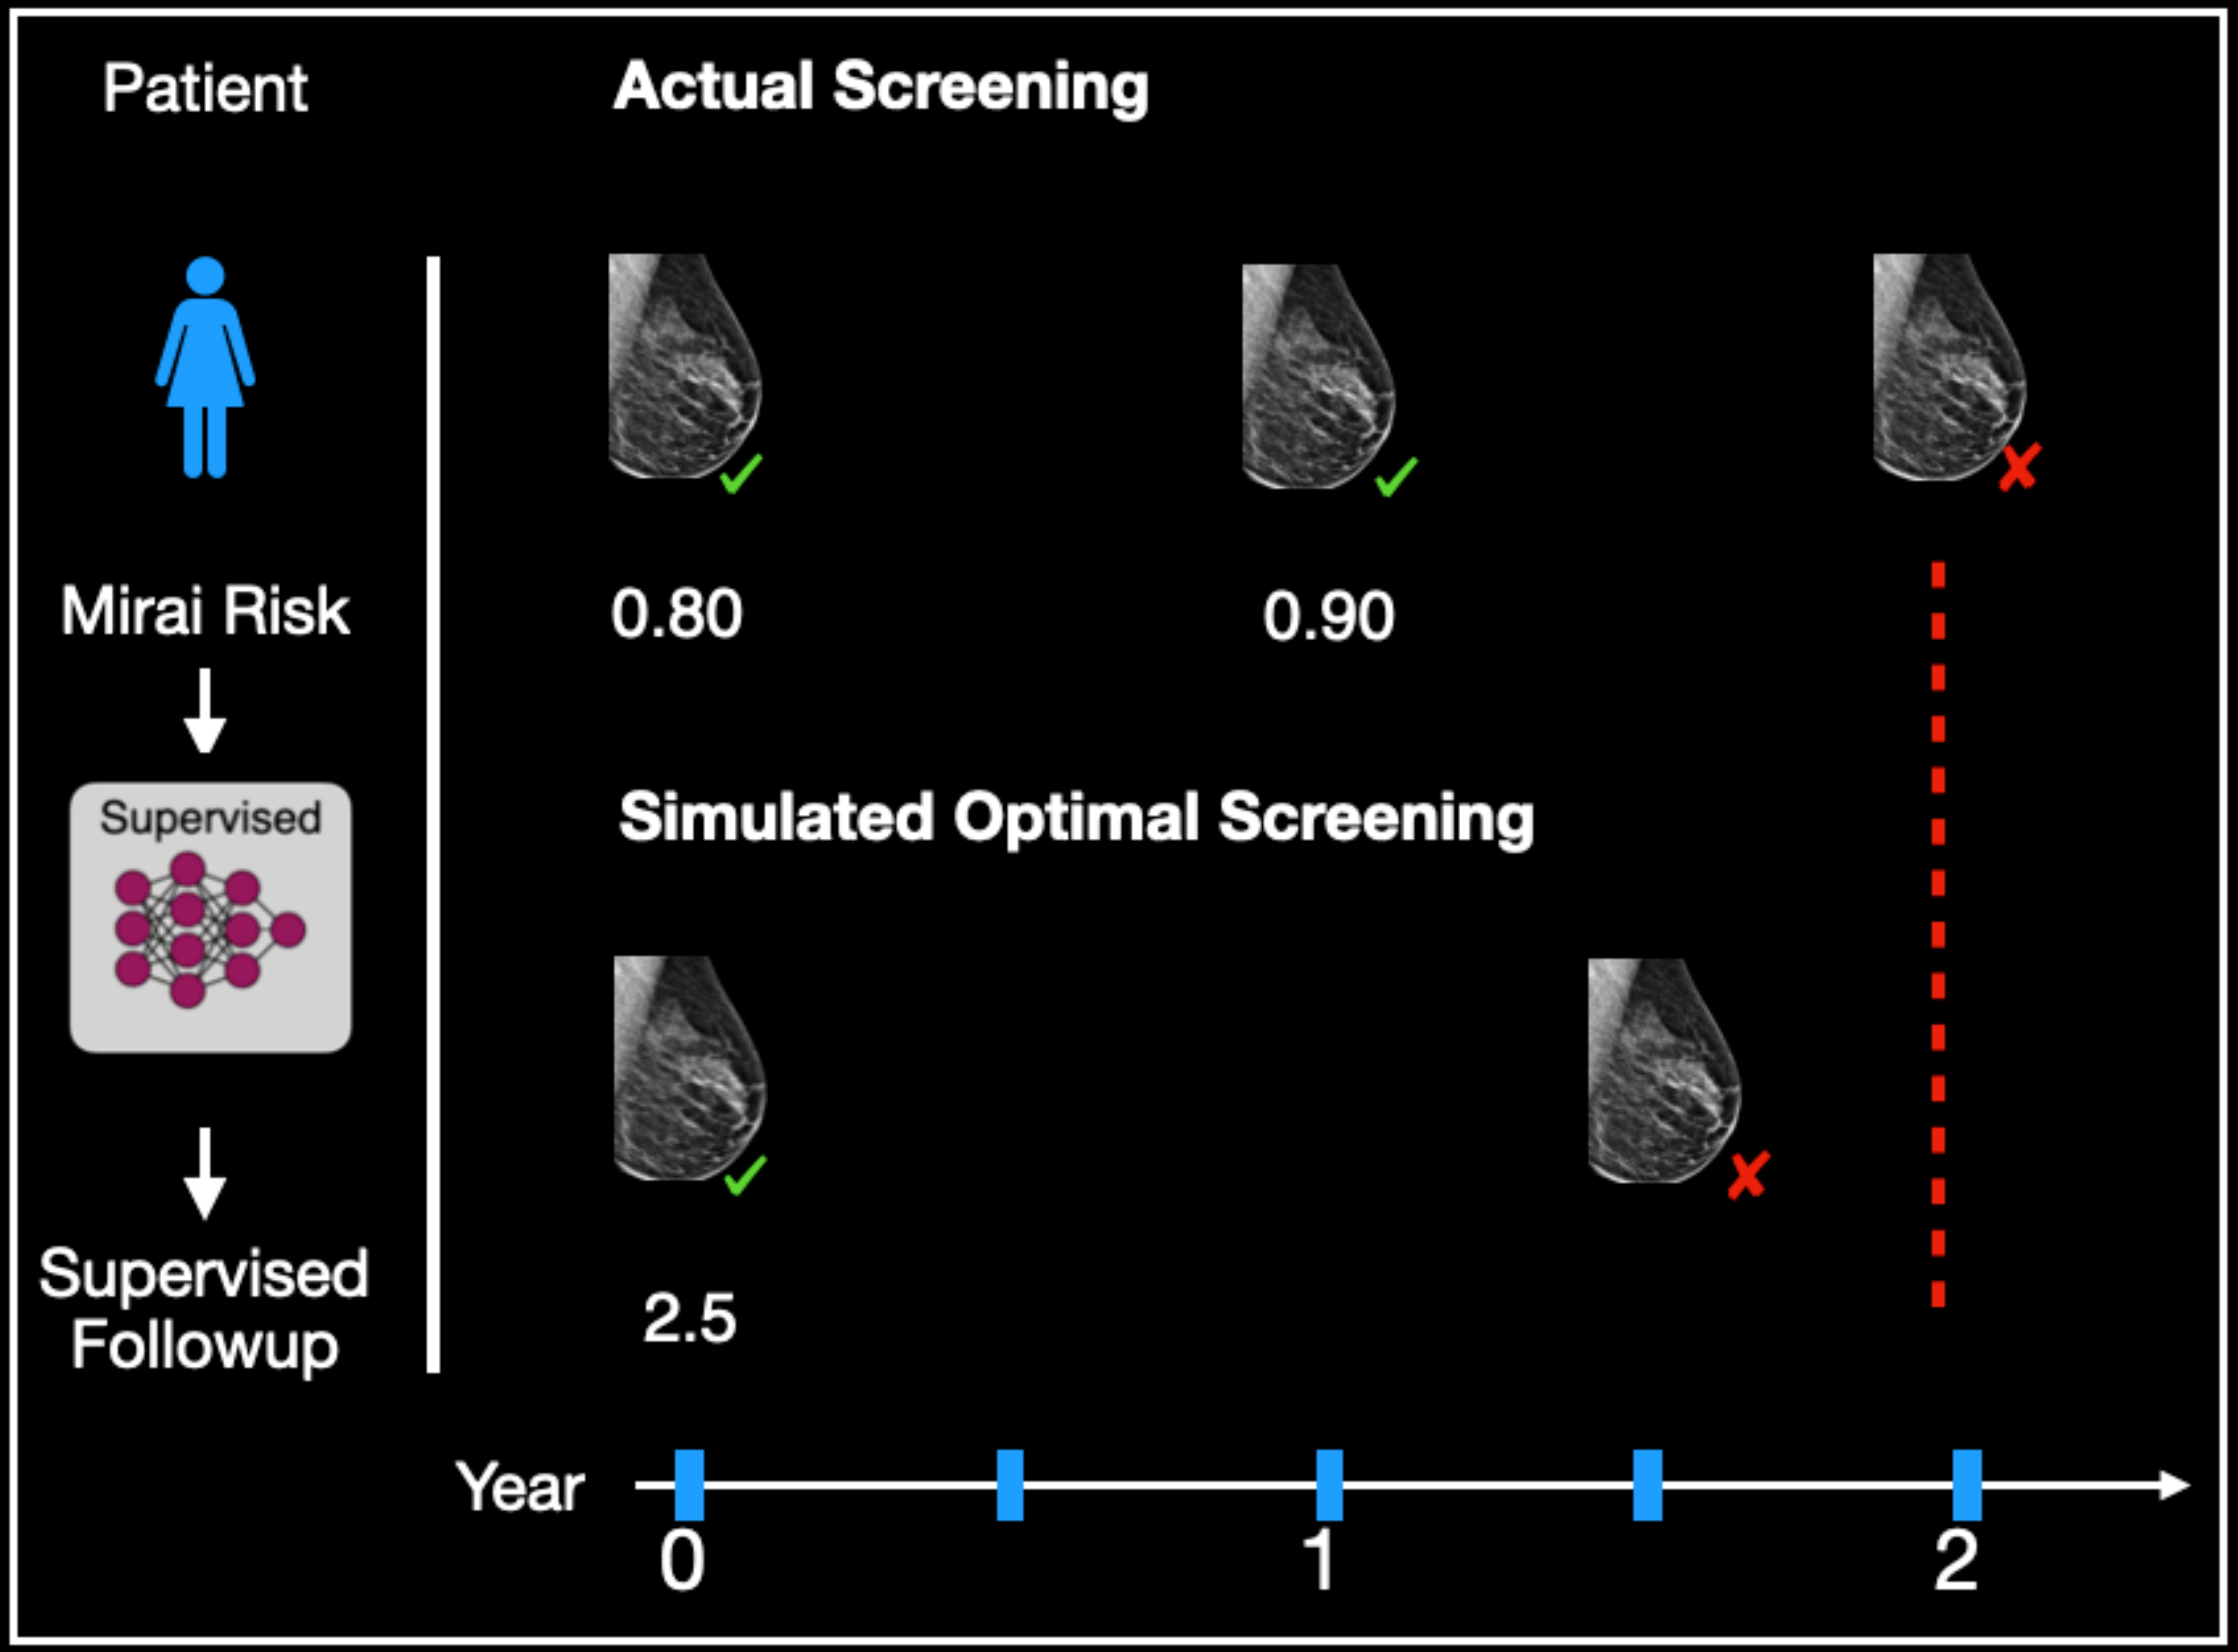

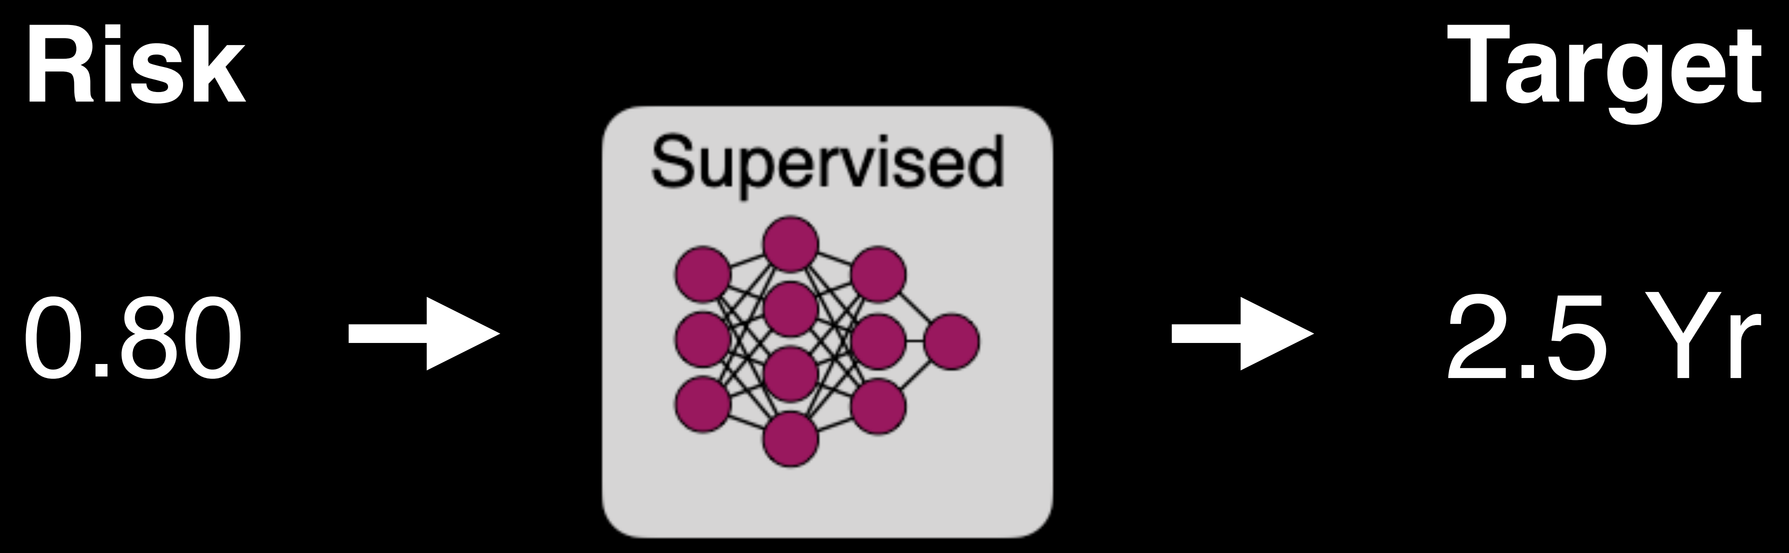

Reward = - || Prediction - Target ||

Tempo: Policy Design as Reinforcement Learning

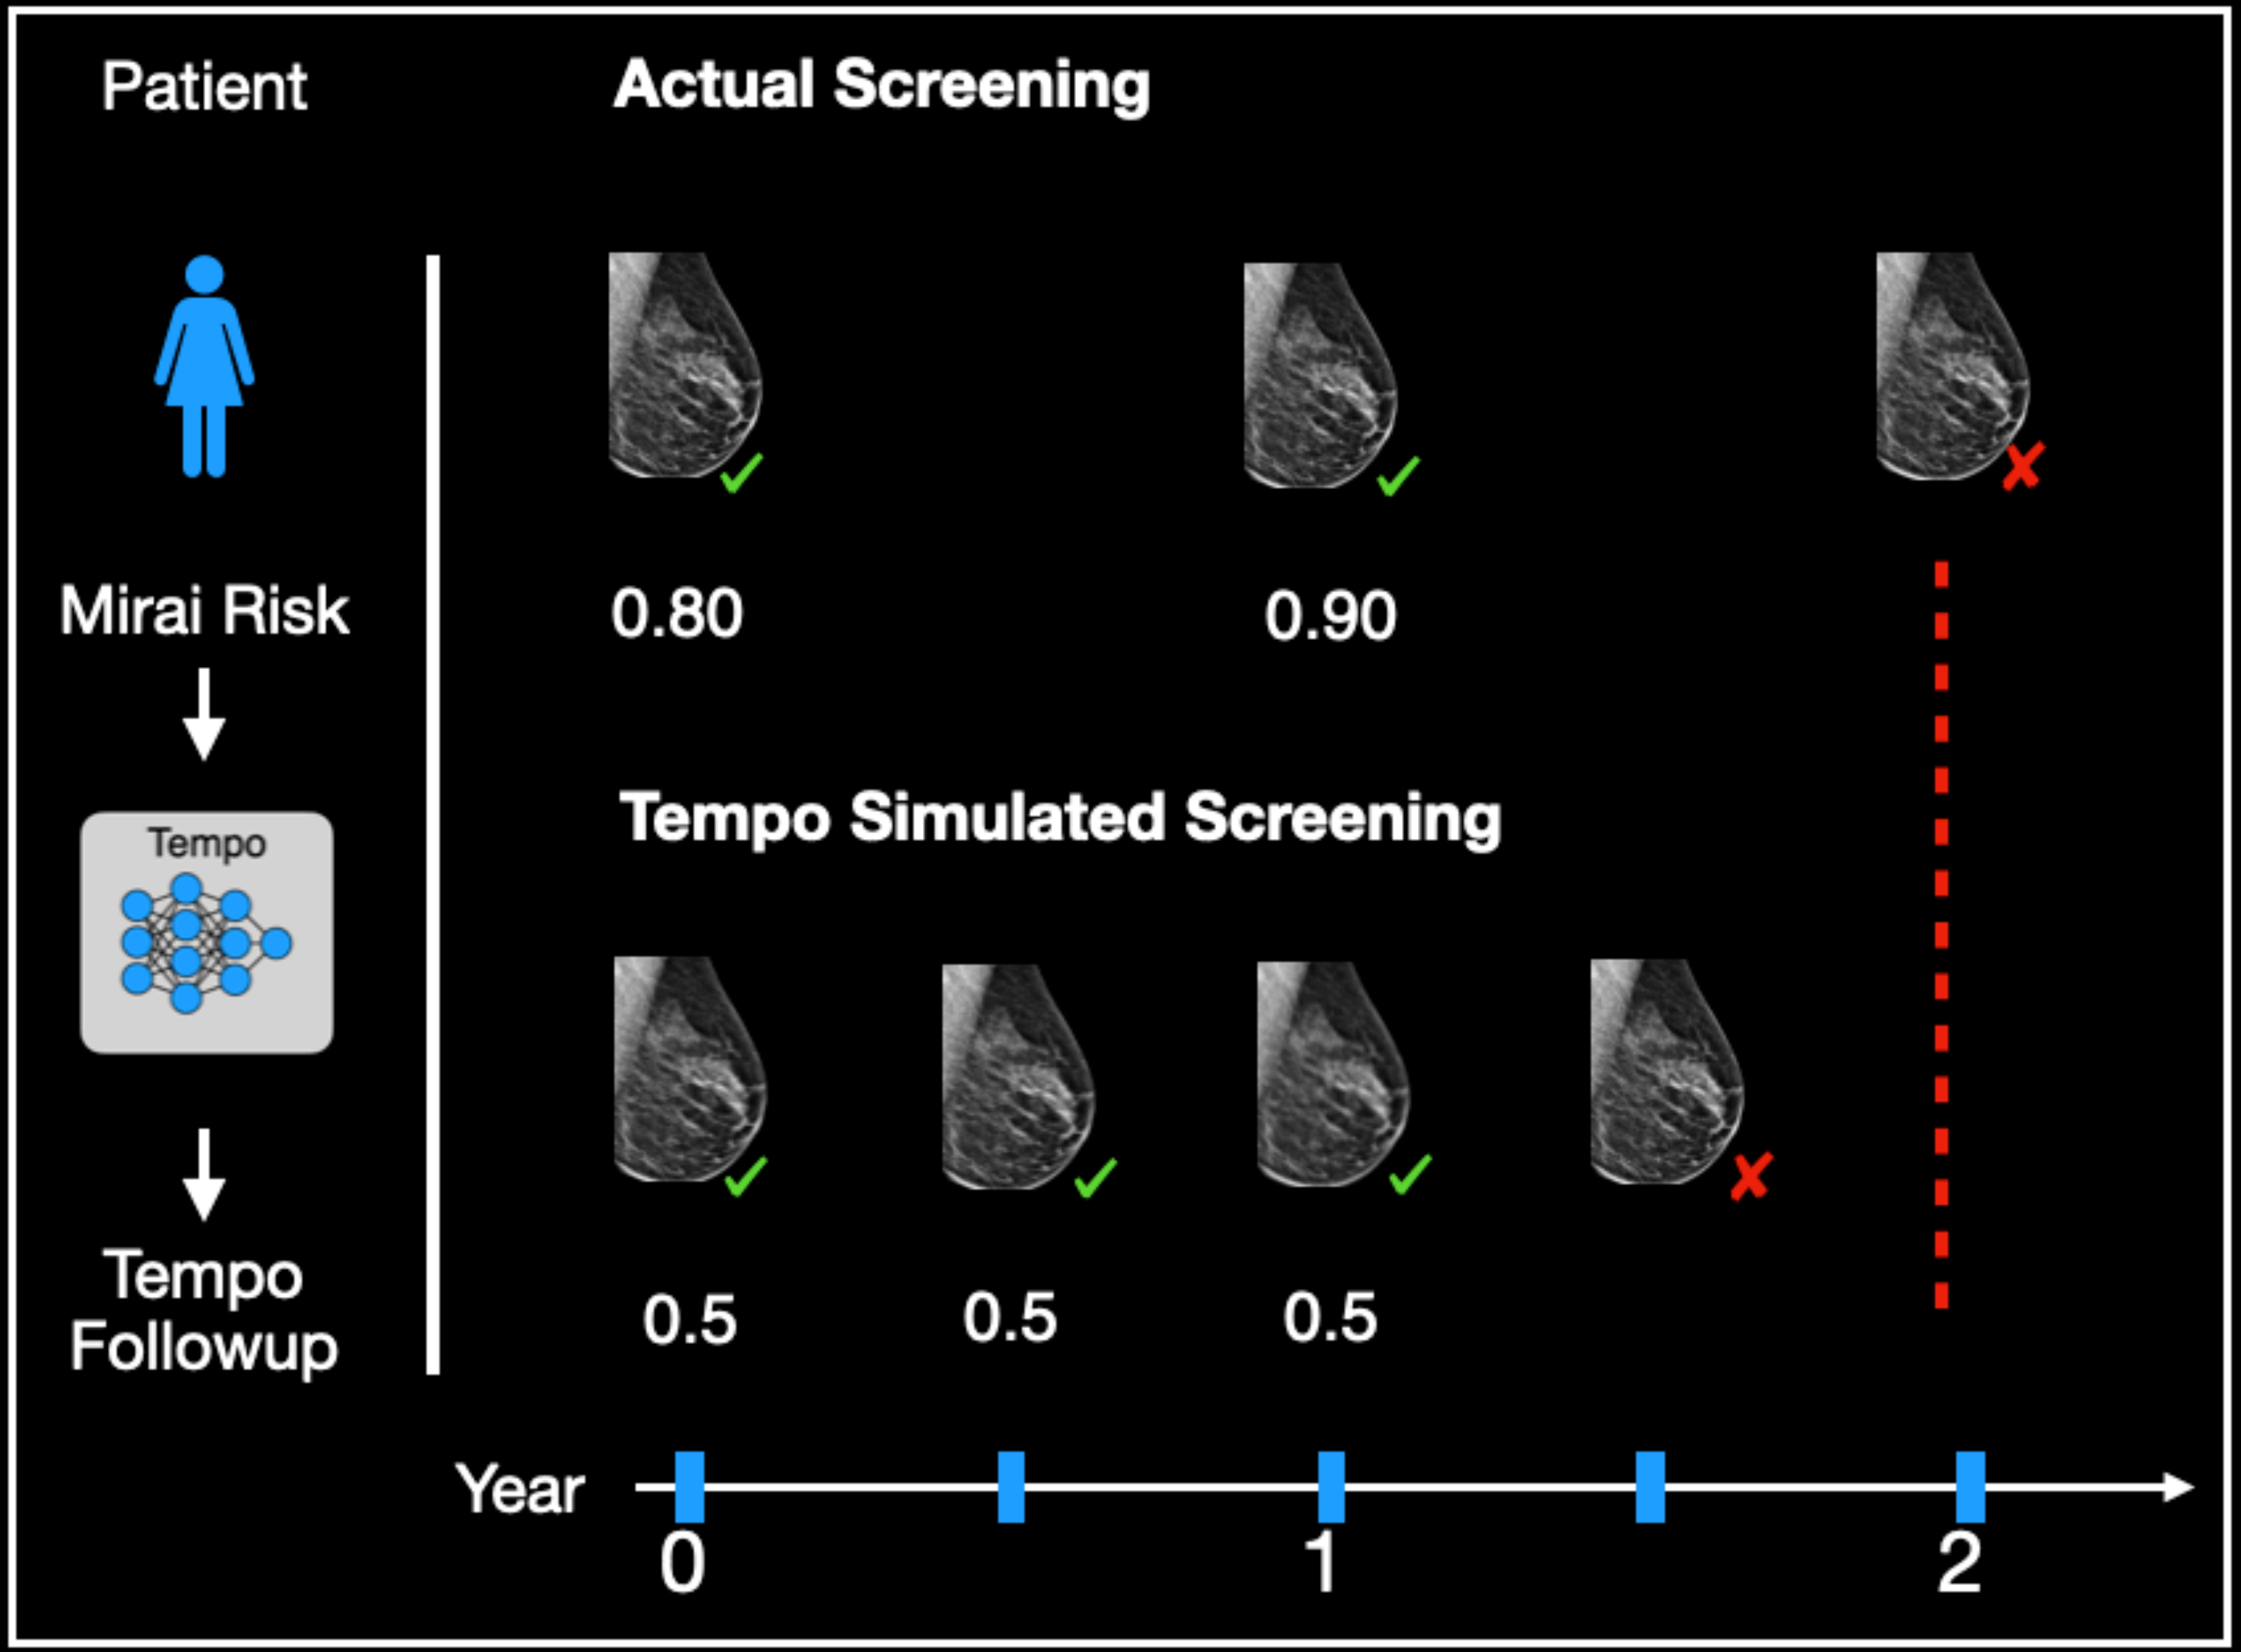

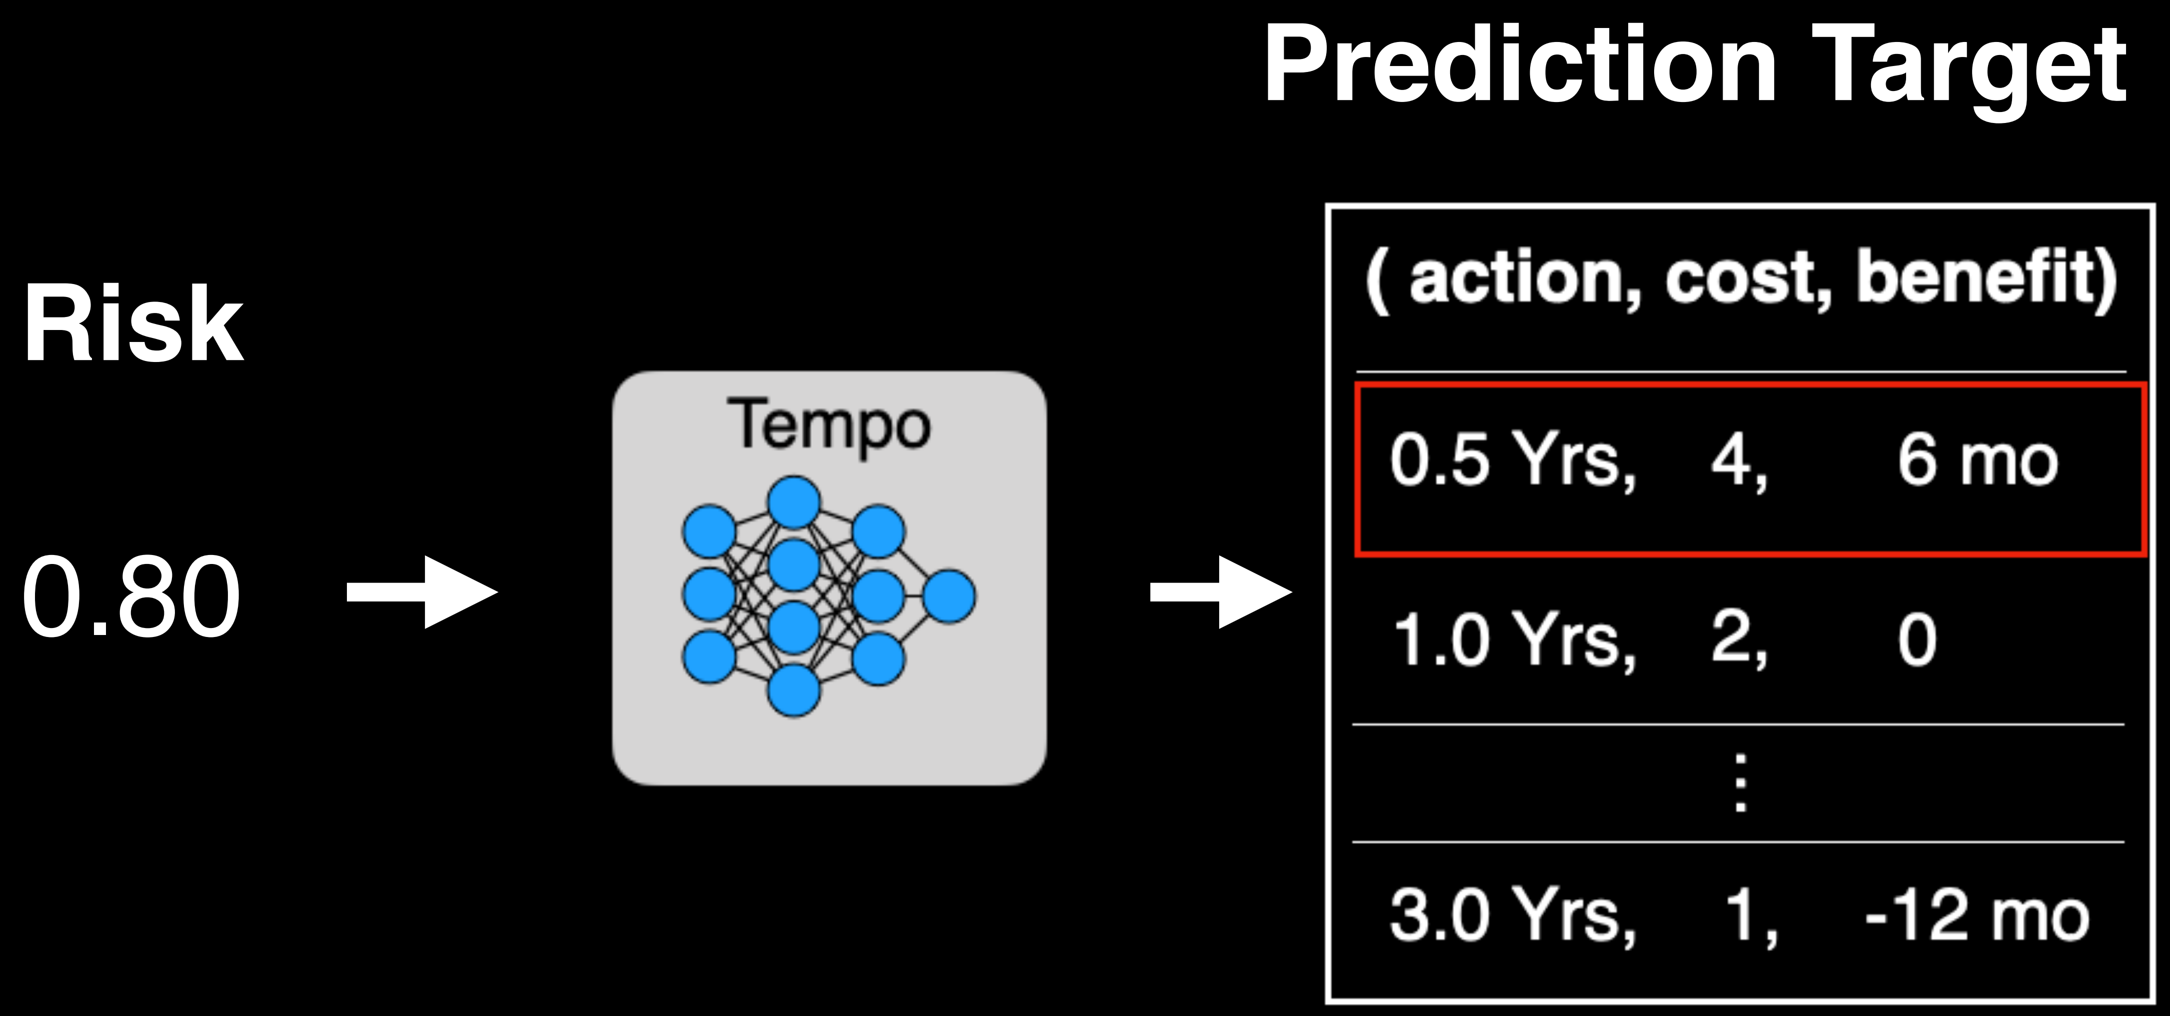

Reward = λ_1 **Benefit -** λ_2**Cost**

Learning for a fixed preference

Patient | **Risk**

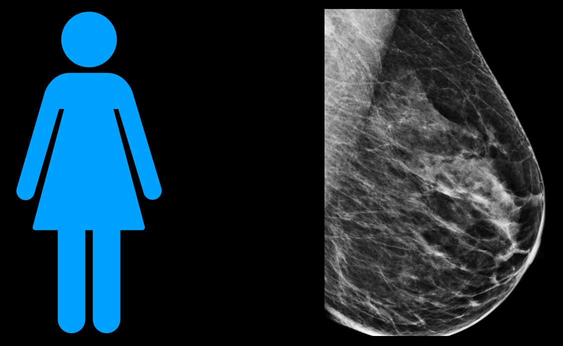

→ 0.80 →

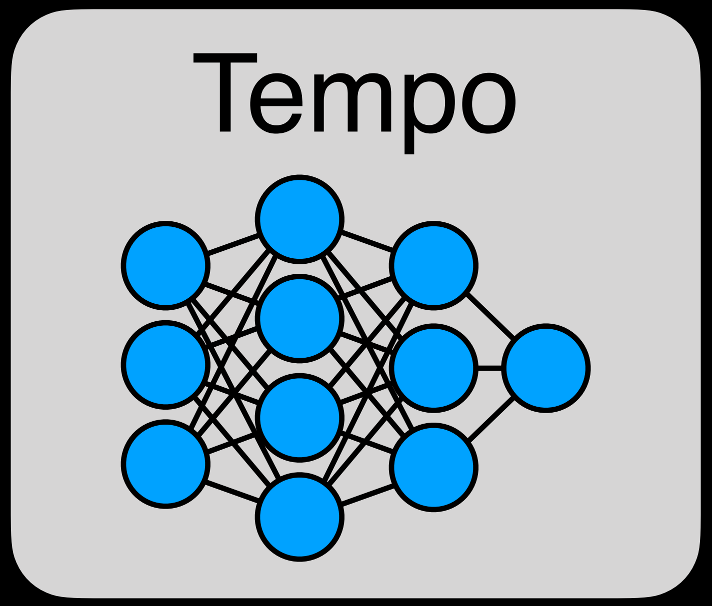

→ **1 year followup**

Reward = λ_1 Early Detection Benefit - λ_2Screening Cost

Learning for a fixed preference: Q Learning

Patient

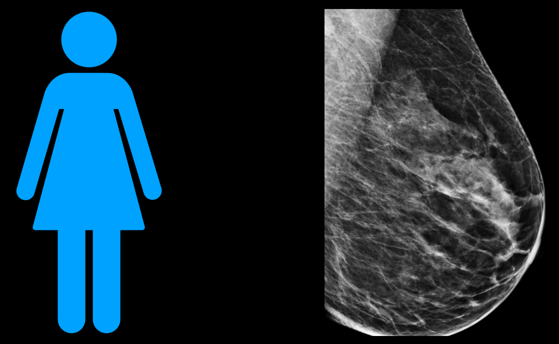

→

Risk

0.80

→

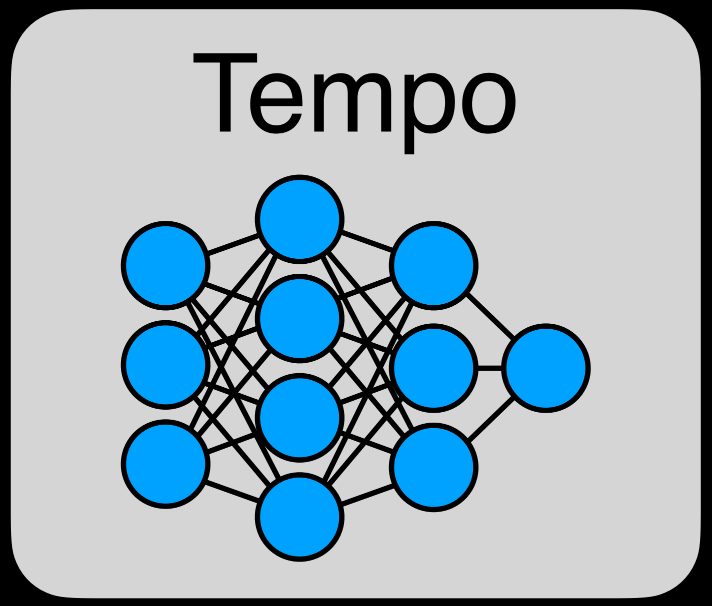

→ **1 year followup**

Reward = λ_1 Early Detection Benefit - λ_2Screening Cost

$$Q(s,a) = R(s,a) + \gamma \text{max}_a Q(s',a)$$

$$\mathcal{L}(s,a) = ||R(s,a) + \text{max}_a Q(s',a) - Q(s,a)||^2$$

Tricks for stable training

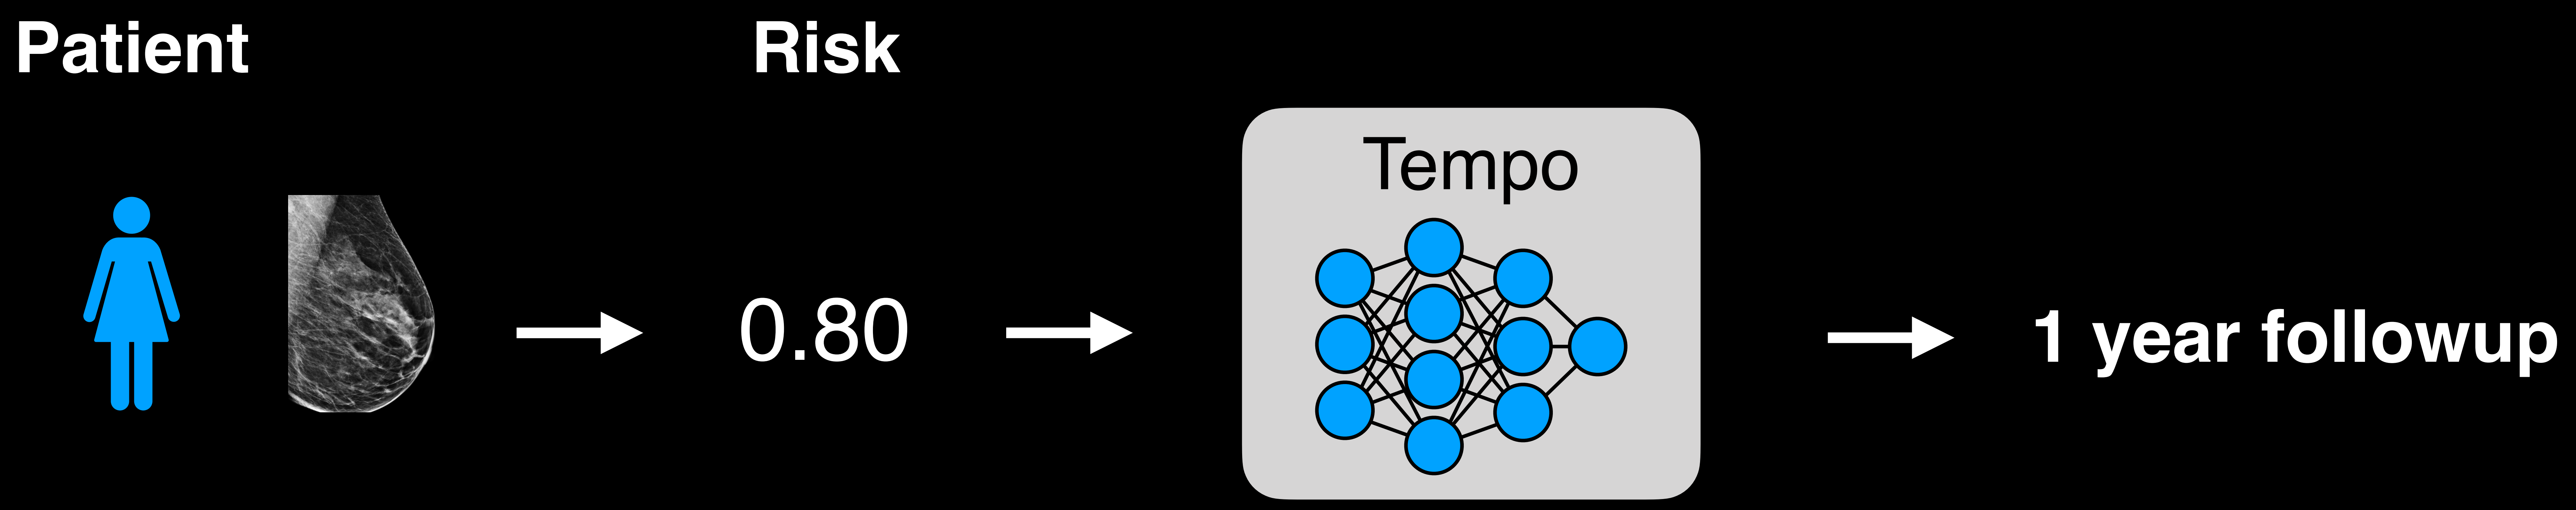

Randomly sample transitions from *Experience Replay Buffers*

Use slowly updated target network (i.e. copy Q every 100~ steps)

$$\mathscr{L}(s,a) = ||R(s,a) + \text{max}_a Q_{\boxed{target}}(s',a) - Q(s,a)||^2$$

Scaling to unknown preferences

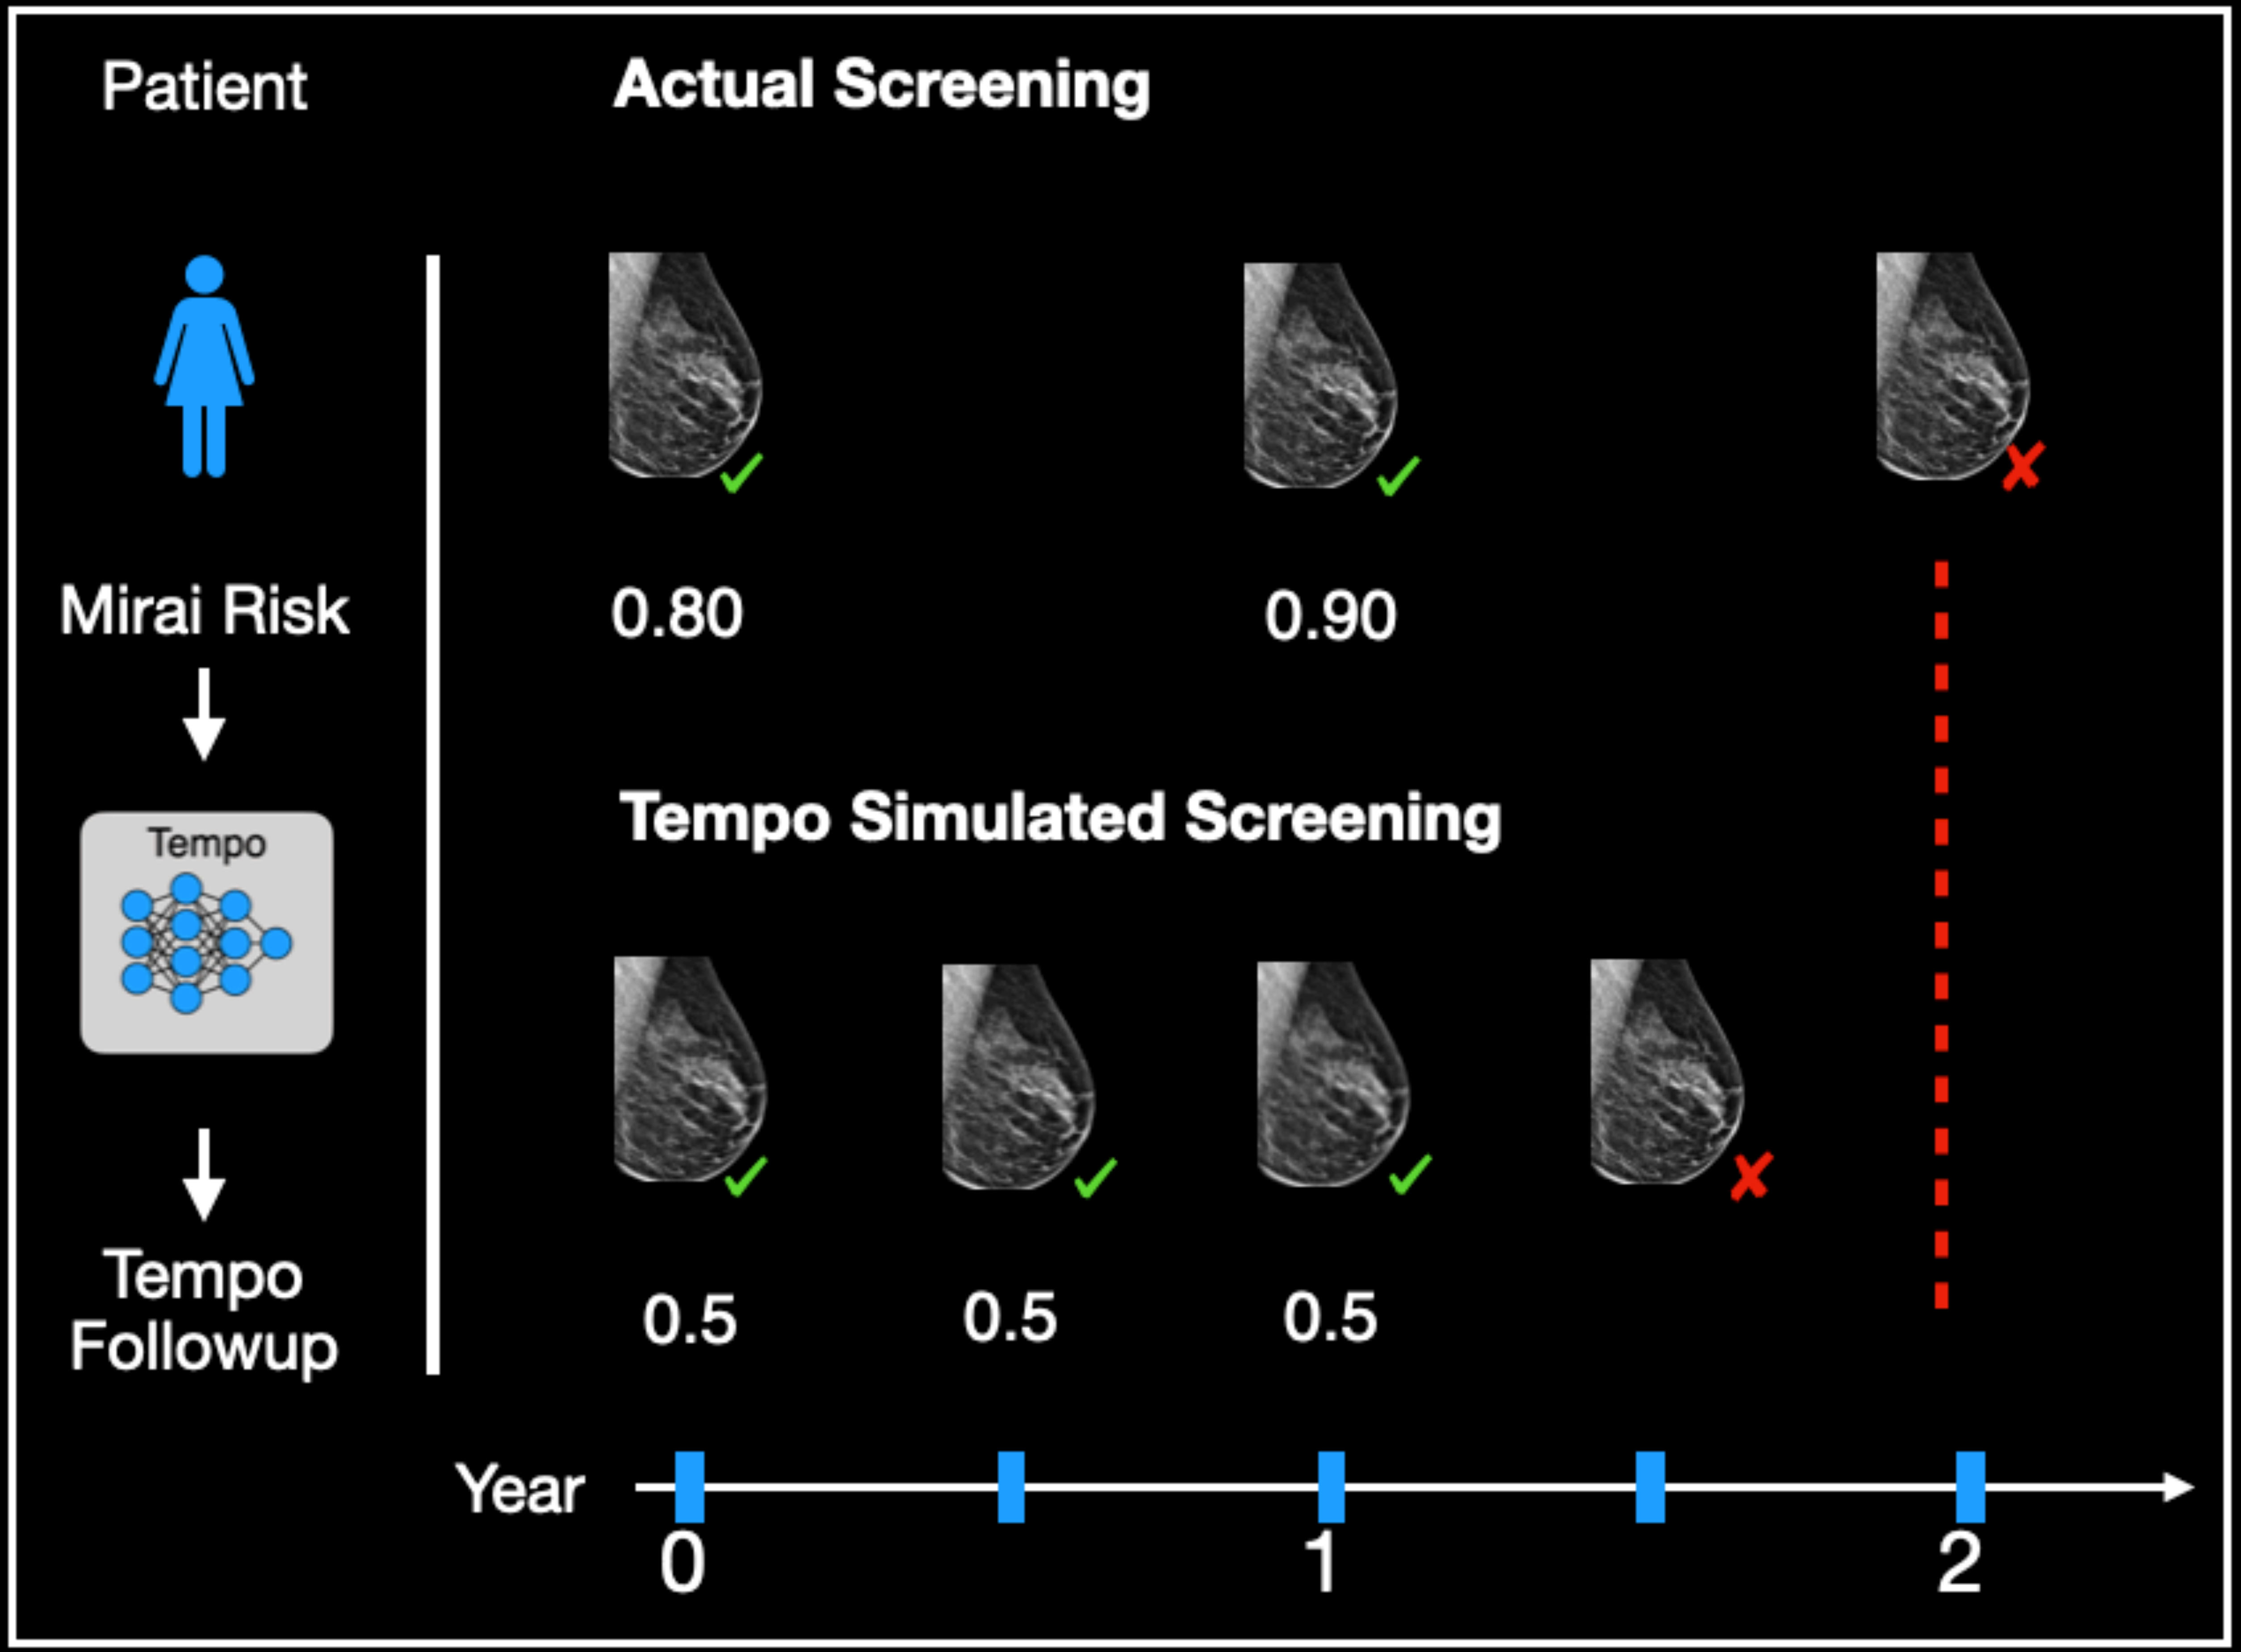

Risk

0.80 → 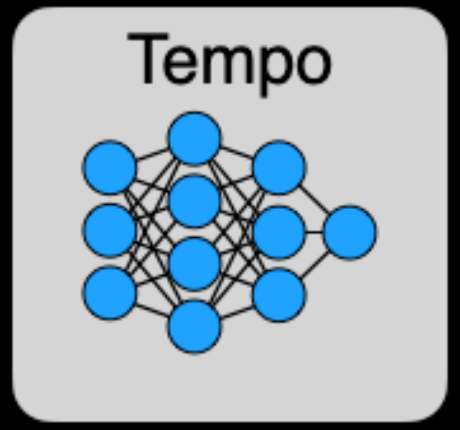 →

Prediction Target

(action, cost, benefit)		
0.5 Yrs,	4,	6 mo
1.0 Yrs,	2,	0
⋮		
3.0 Yrs,	1,	-12 mo

Reward = λ_1 Benefit - λ_2Cost

- λ unknown at training time

\+ We have access to [Benefit, Cost]

Supporting diverse clinical requirements

Patient

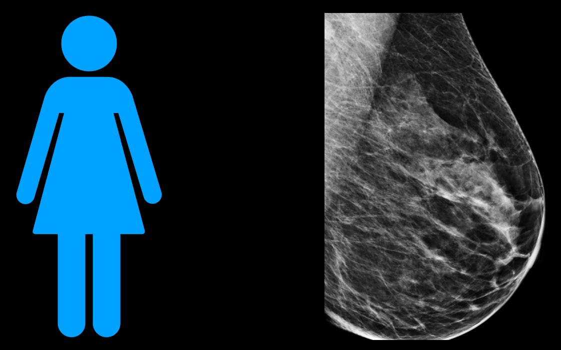

Risk

0.80

Tempo

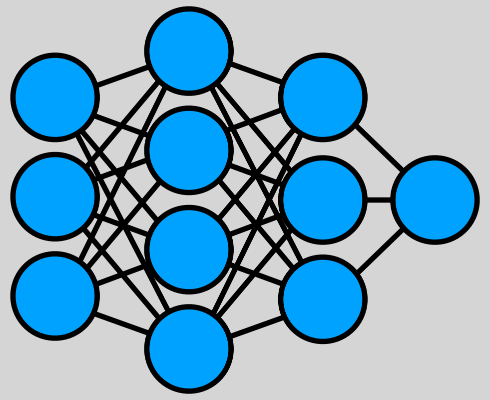

1 year followup

preference: $\lambda = \begin{bmatrix} 0.9 \\ 0.1 \end{bmatrix}$

$\vec{r}(s, a) =$ [Early Detection, Screening Cost]

Trained across possible (λ_1, λ_2) to maximize:

$\lambda \cdot \vec{r} = \lambda_1$ Early Detection Benefit - λ_2Screening Cost

Towards multi-objective RL: Scalarized updates

Patient **Risk**

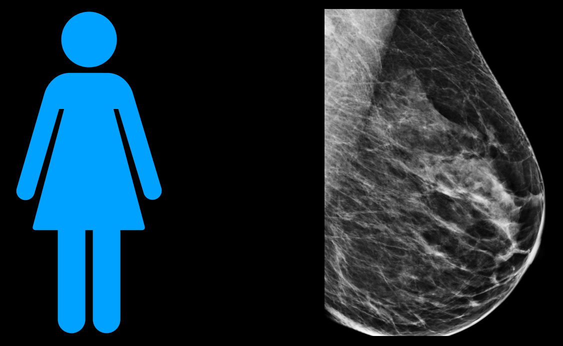

0.80

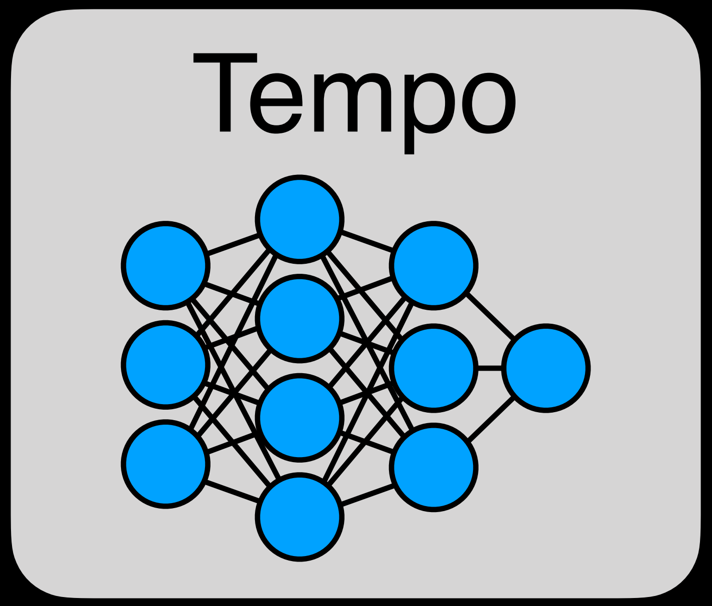

1 year followup

preference: $\lambda = \begin{bmatrix} 0.9 \\ 0.1 \end{bmatrix}$

$\vec{r}(s,a) = [\text{Early Detection, Screening Cost}]$

$Q(s,a,\lambda) = \lambda \vec{r}(s,a) + \gamma \lambda \text{max}_a Q(s',a,\lambda)$

Towards multi-objective RL: Scalarized updates

Patient **Risk**

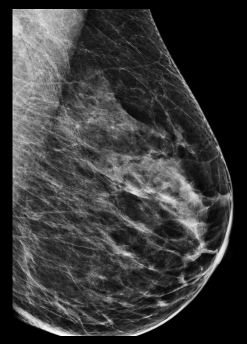

→ 0.80 →

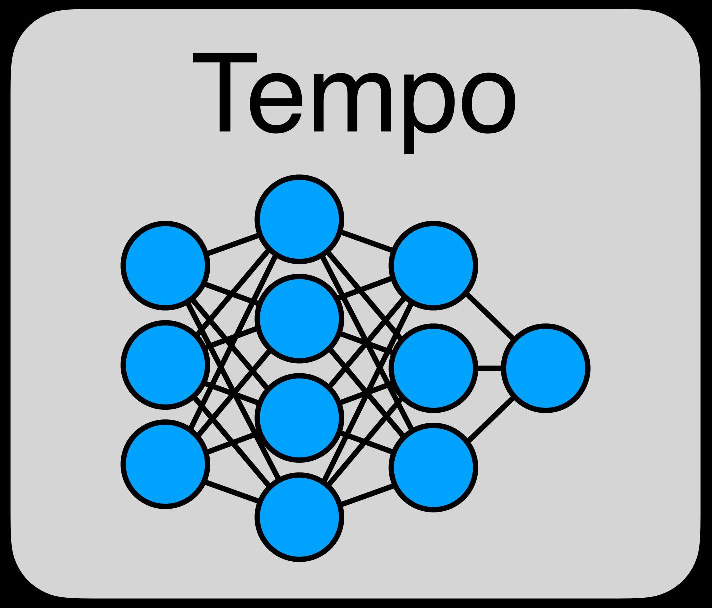

→ **1 year followup**

preference: $\lambda = \begin{bmatrix} 0.9 \\ 0.1 \end{bmatrix}$

$$\vec{r}(s, a) = [\text{Early Detection, Screening Cost}]$$

$$Q(s, a, \lambda) = \lambda \vec{r}(s, a) + \gamma \lambda \text{max}_a Q(s', a, \lambda)$$

Doesn't use relationship between λ

Envelope Q-Learning

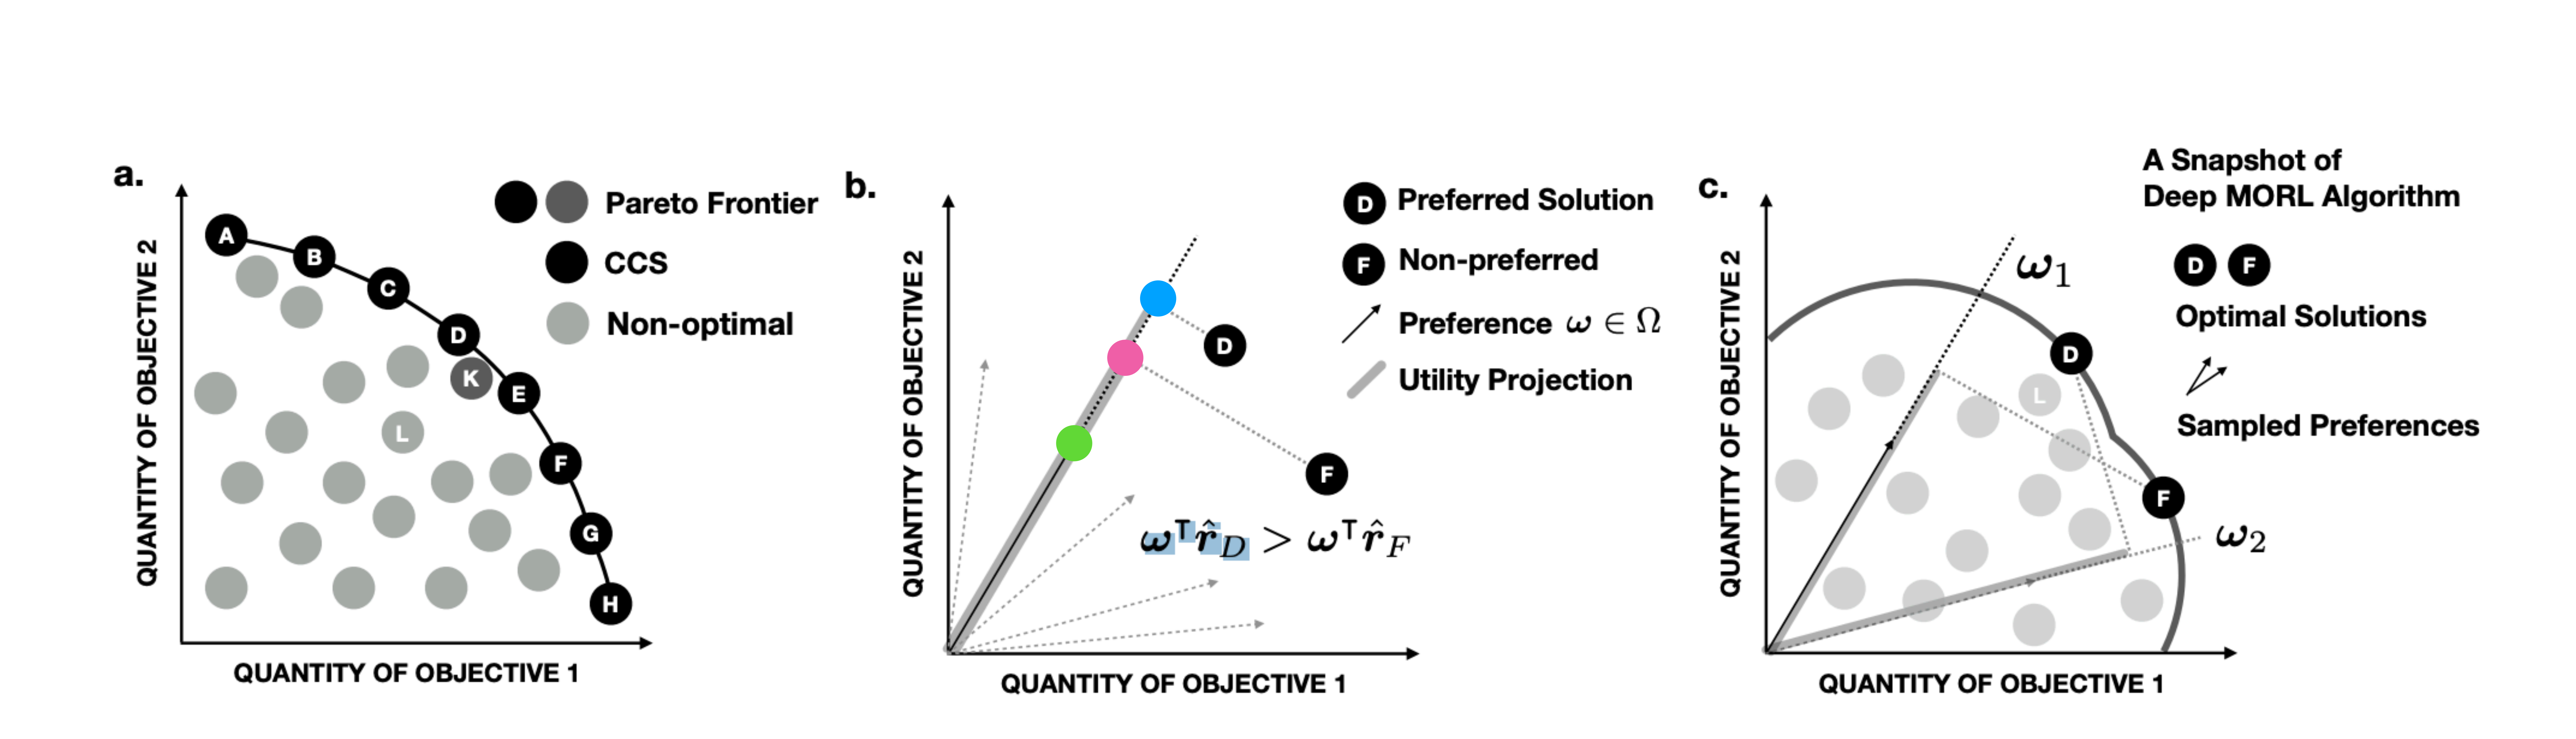

Search *envelope* of policy:

Identify λ' more effective for true target λ

A Generalized Algorithm for Multi-Objective Reinforcement Learning and Policy Adaptation

Runzhe Yang, Xingyuan Sun, Karthik Narasimhan

Envelope Q-Learning

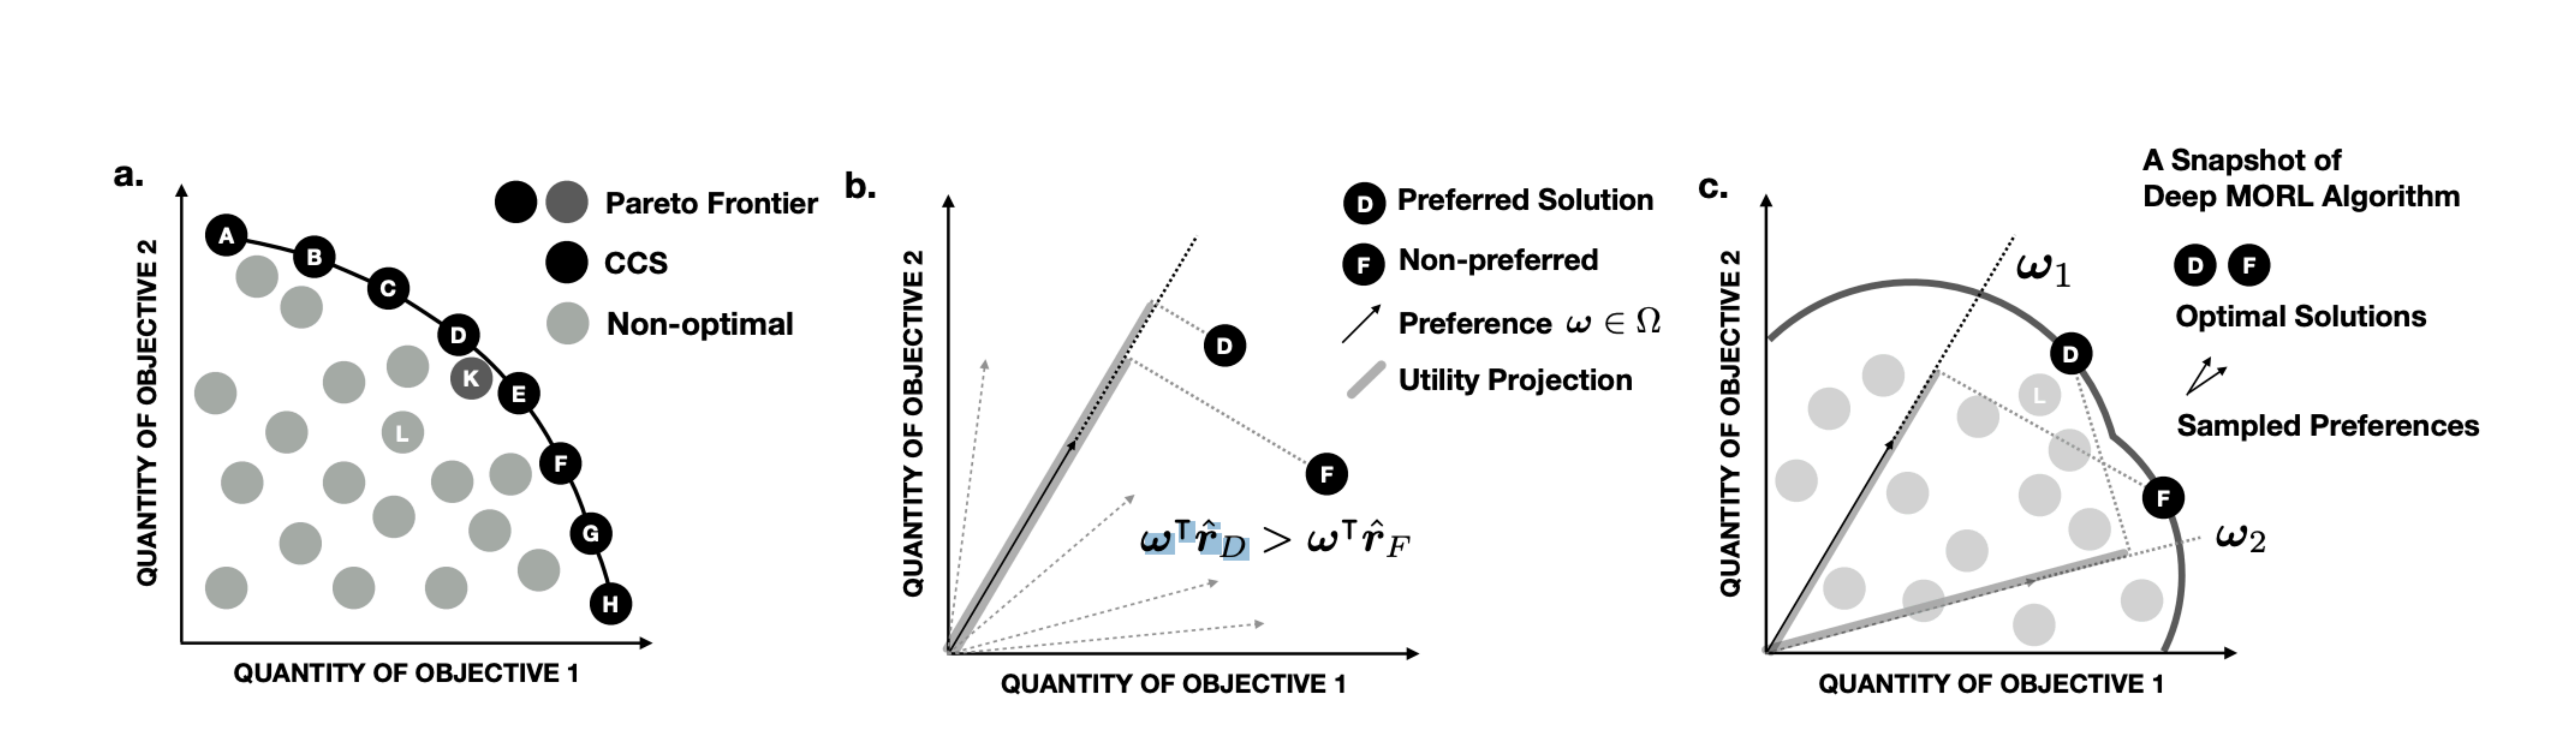

$$y = \vec{r}(s,a) + \gamma \text{arg}_Q \text{max}_{a,\lambda'} \lambda^t \; Q(s', a, \lambda')$$

$$\mathcal{L}(s,a,\lambda) = ||y - Q(s,a,\lambda)||^2$$

A Generalized Algorithm for Multi-Objective Reinforcement Learning and Policy Adaptation

Runzhe Yang, Xingyuan Sun, Karthik Narasimhan

Supporting diverse clinical requirements

Patient | **Risk**

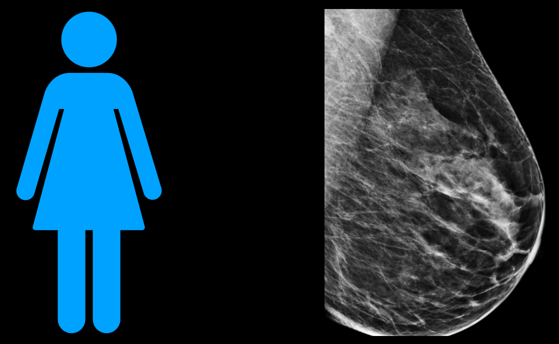

→ 0.80 → Tempo

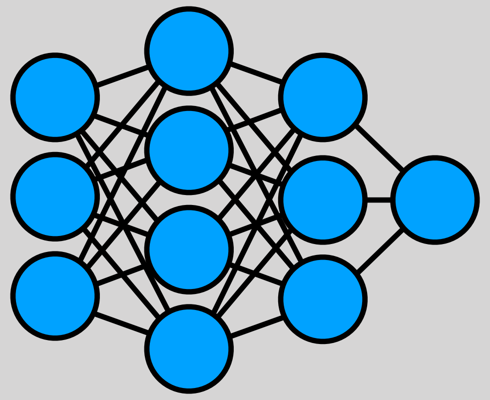

→ **1 year followup**

↑

preference: $\lambda = \begin{bmatrix} 0.9 \\ 0.1 \end{bmatrix}$

Trained across possible (λ_1, λ_2) to maximize:

λ_1 Early Detection Benefit - λ_2Screening Cost

Experimental Setup

Train all models on MGH training set

Test on MGH, Emory, Karolinska, and CGMH

Evaluate Screening Efficiency: $\dfrac{\textbf{Early Detection}}{\textbf{Avg Mammo per Year}}$

Results: MGH Screening Efficiency

Supporting diverse clinical needs

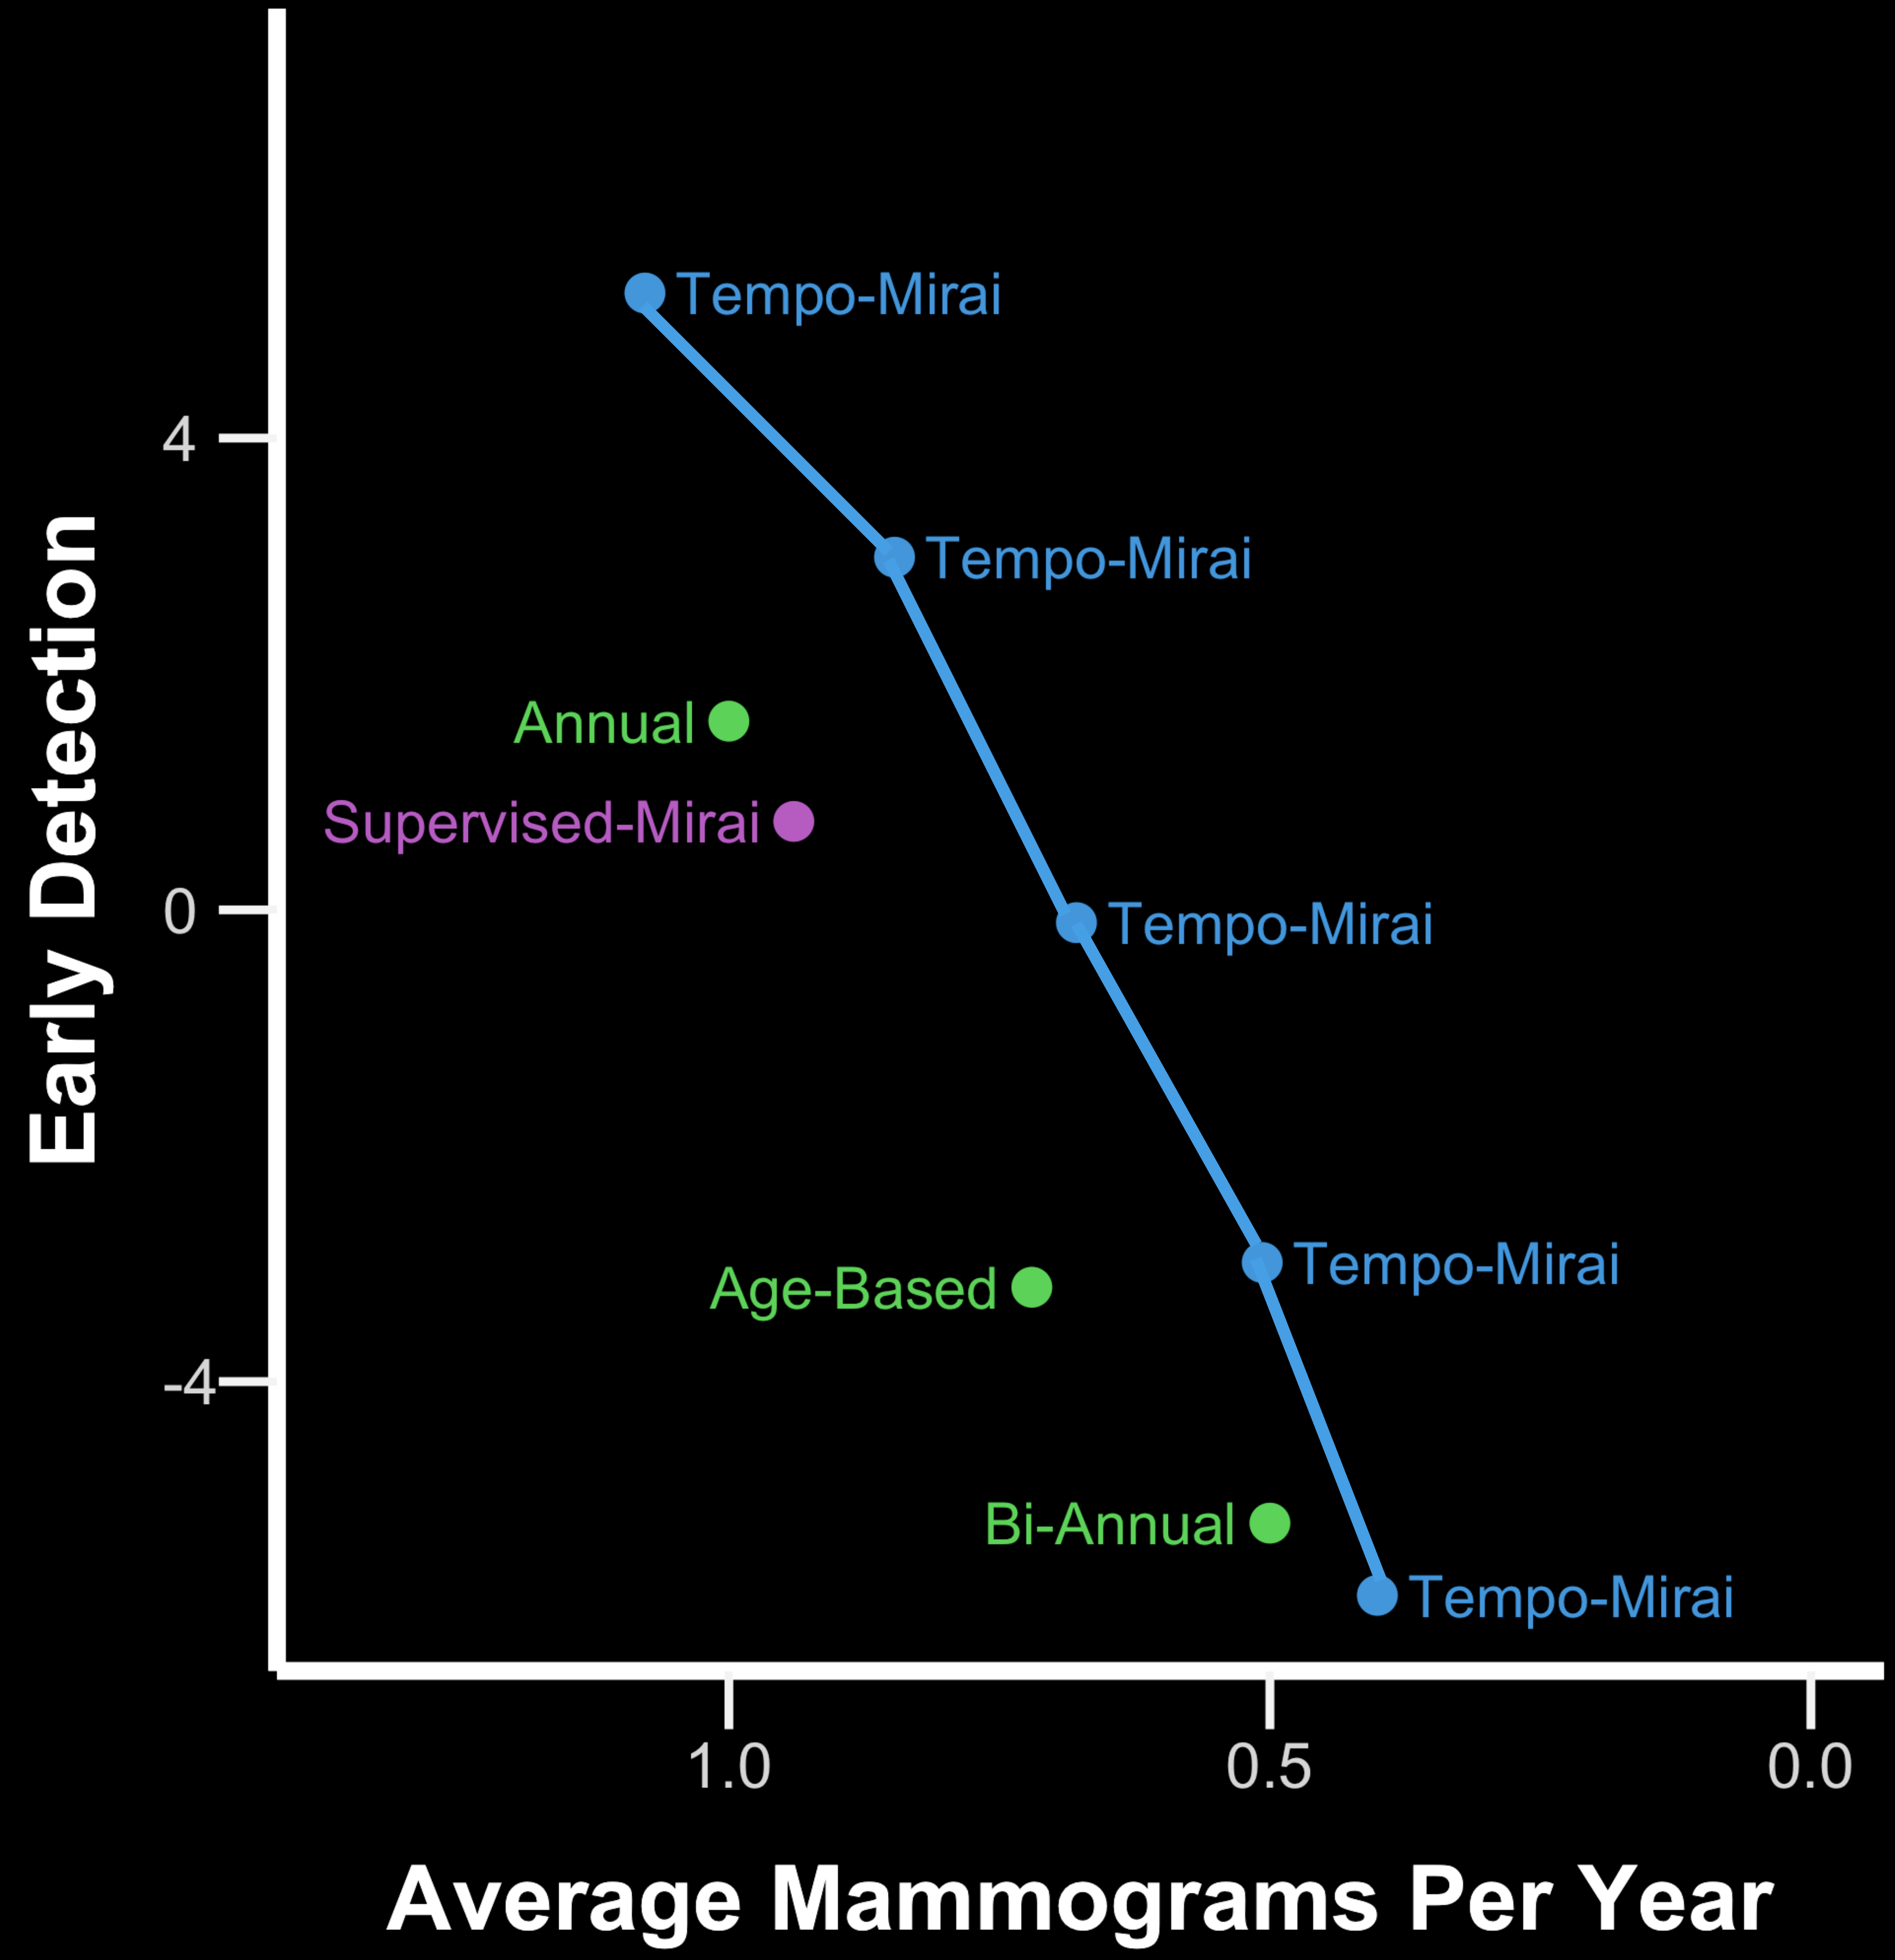

MGH Test Set

How do the policy behave?

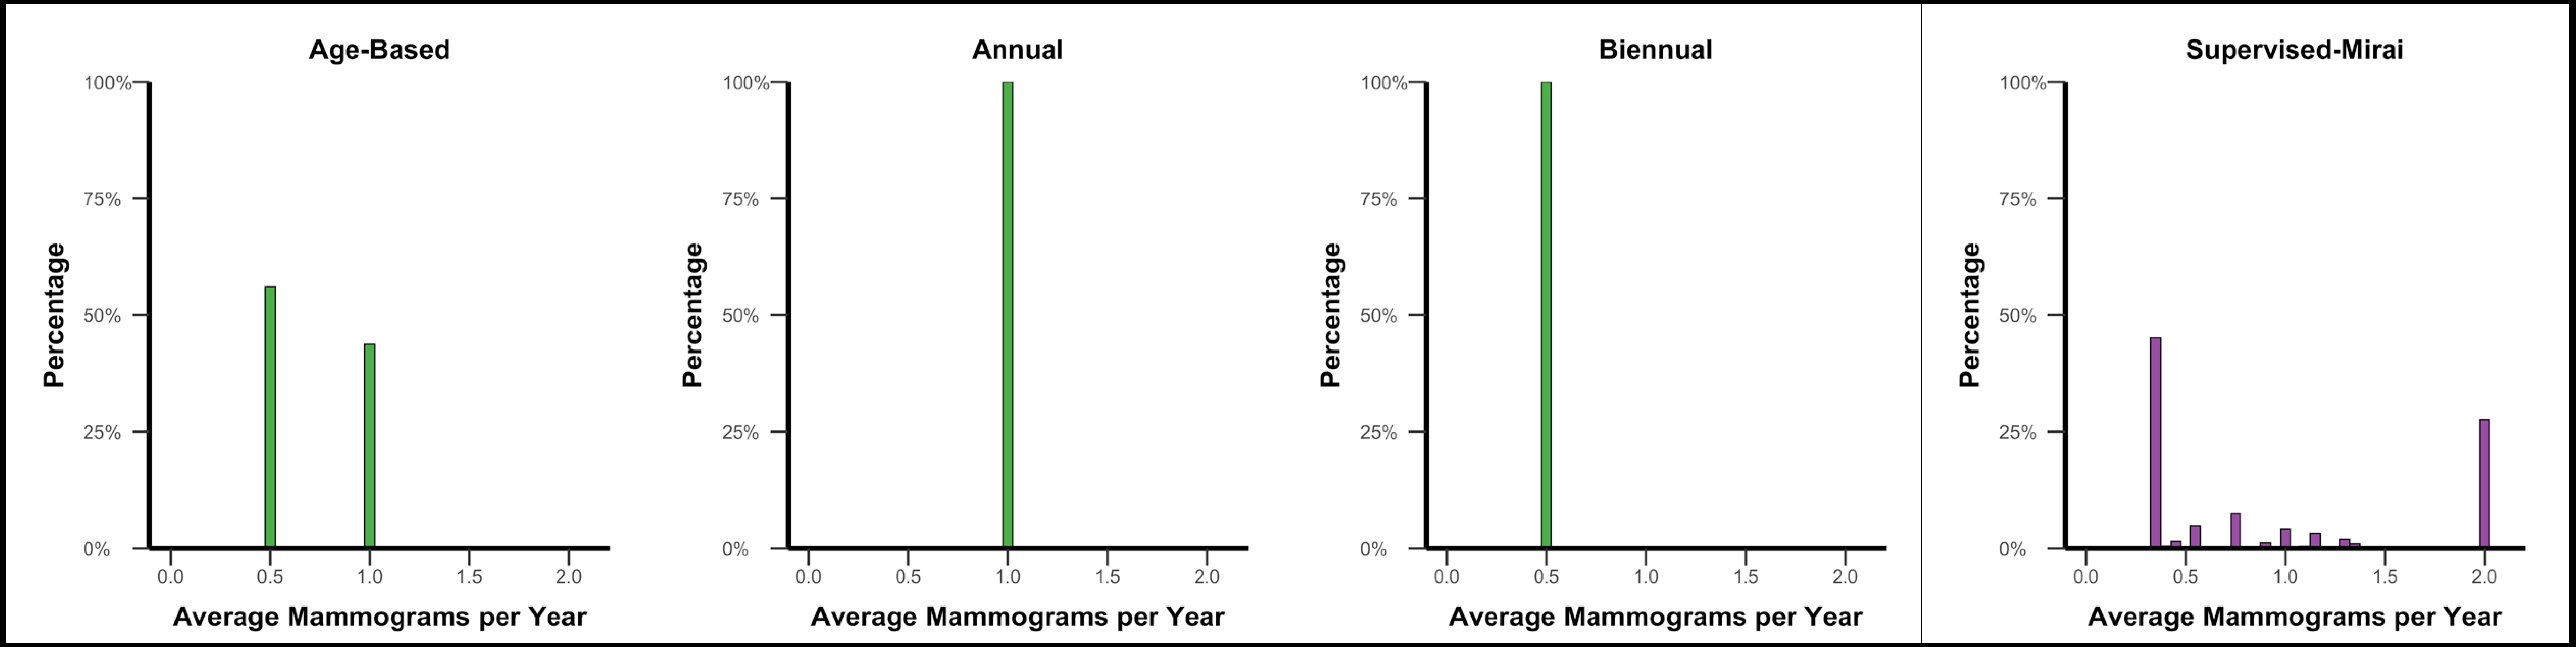

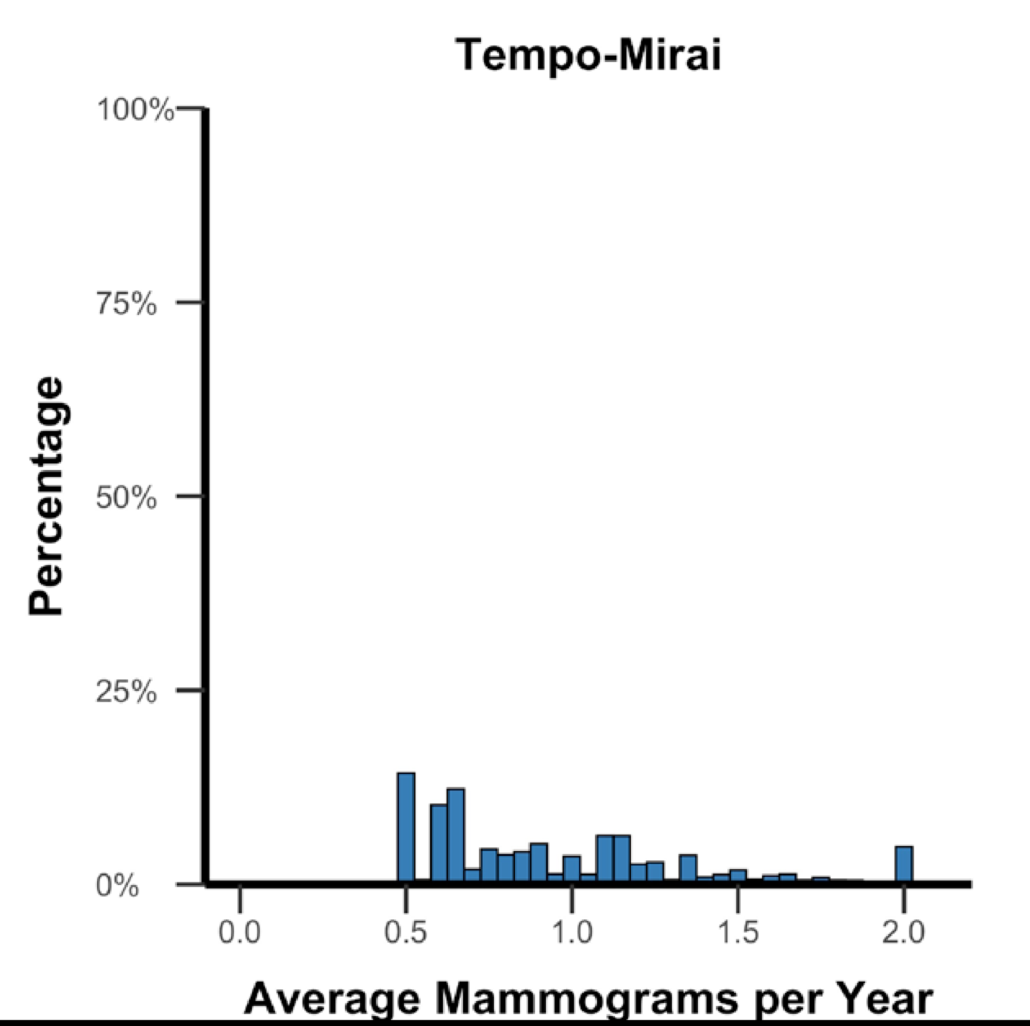

How to Deploy?

Validate on **retrospective data**

Choose desired operating point

Run prospective trials

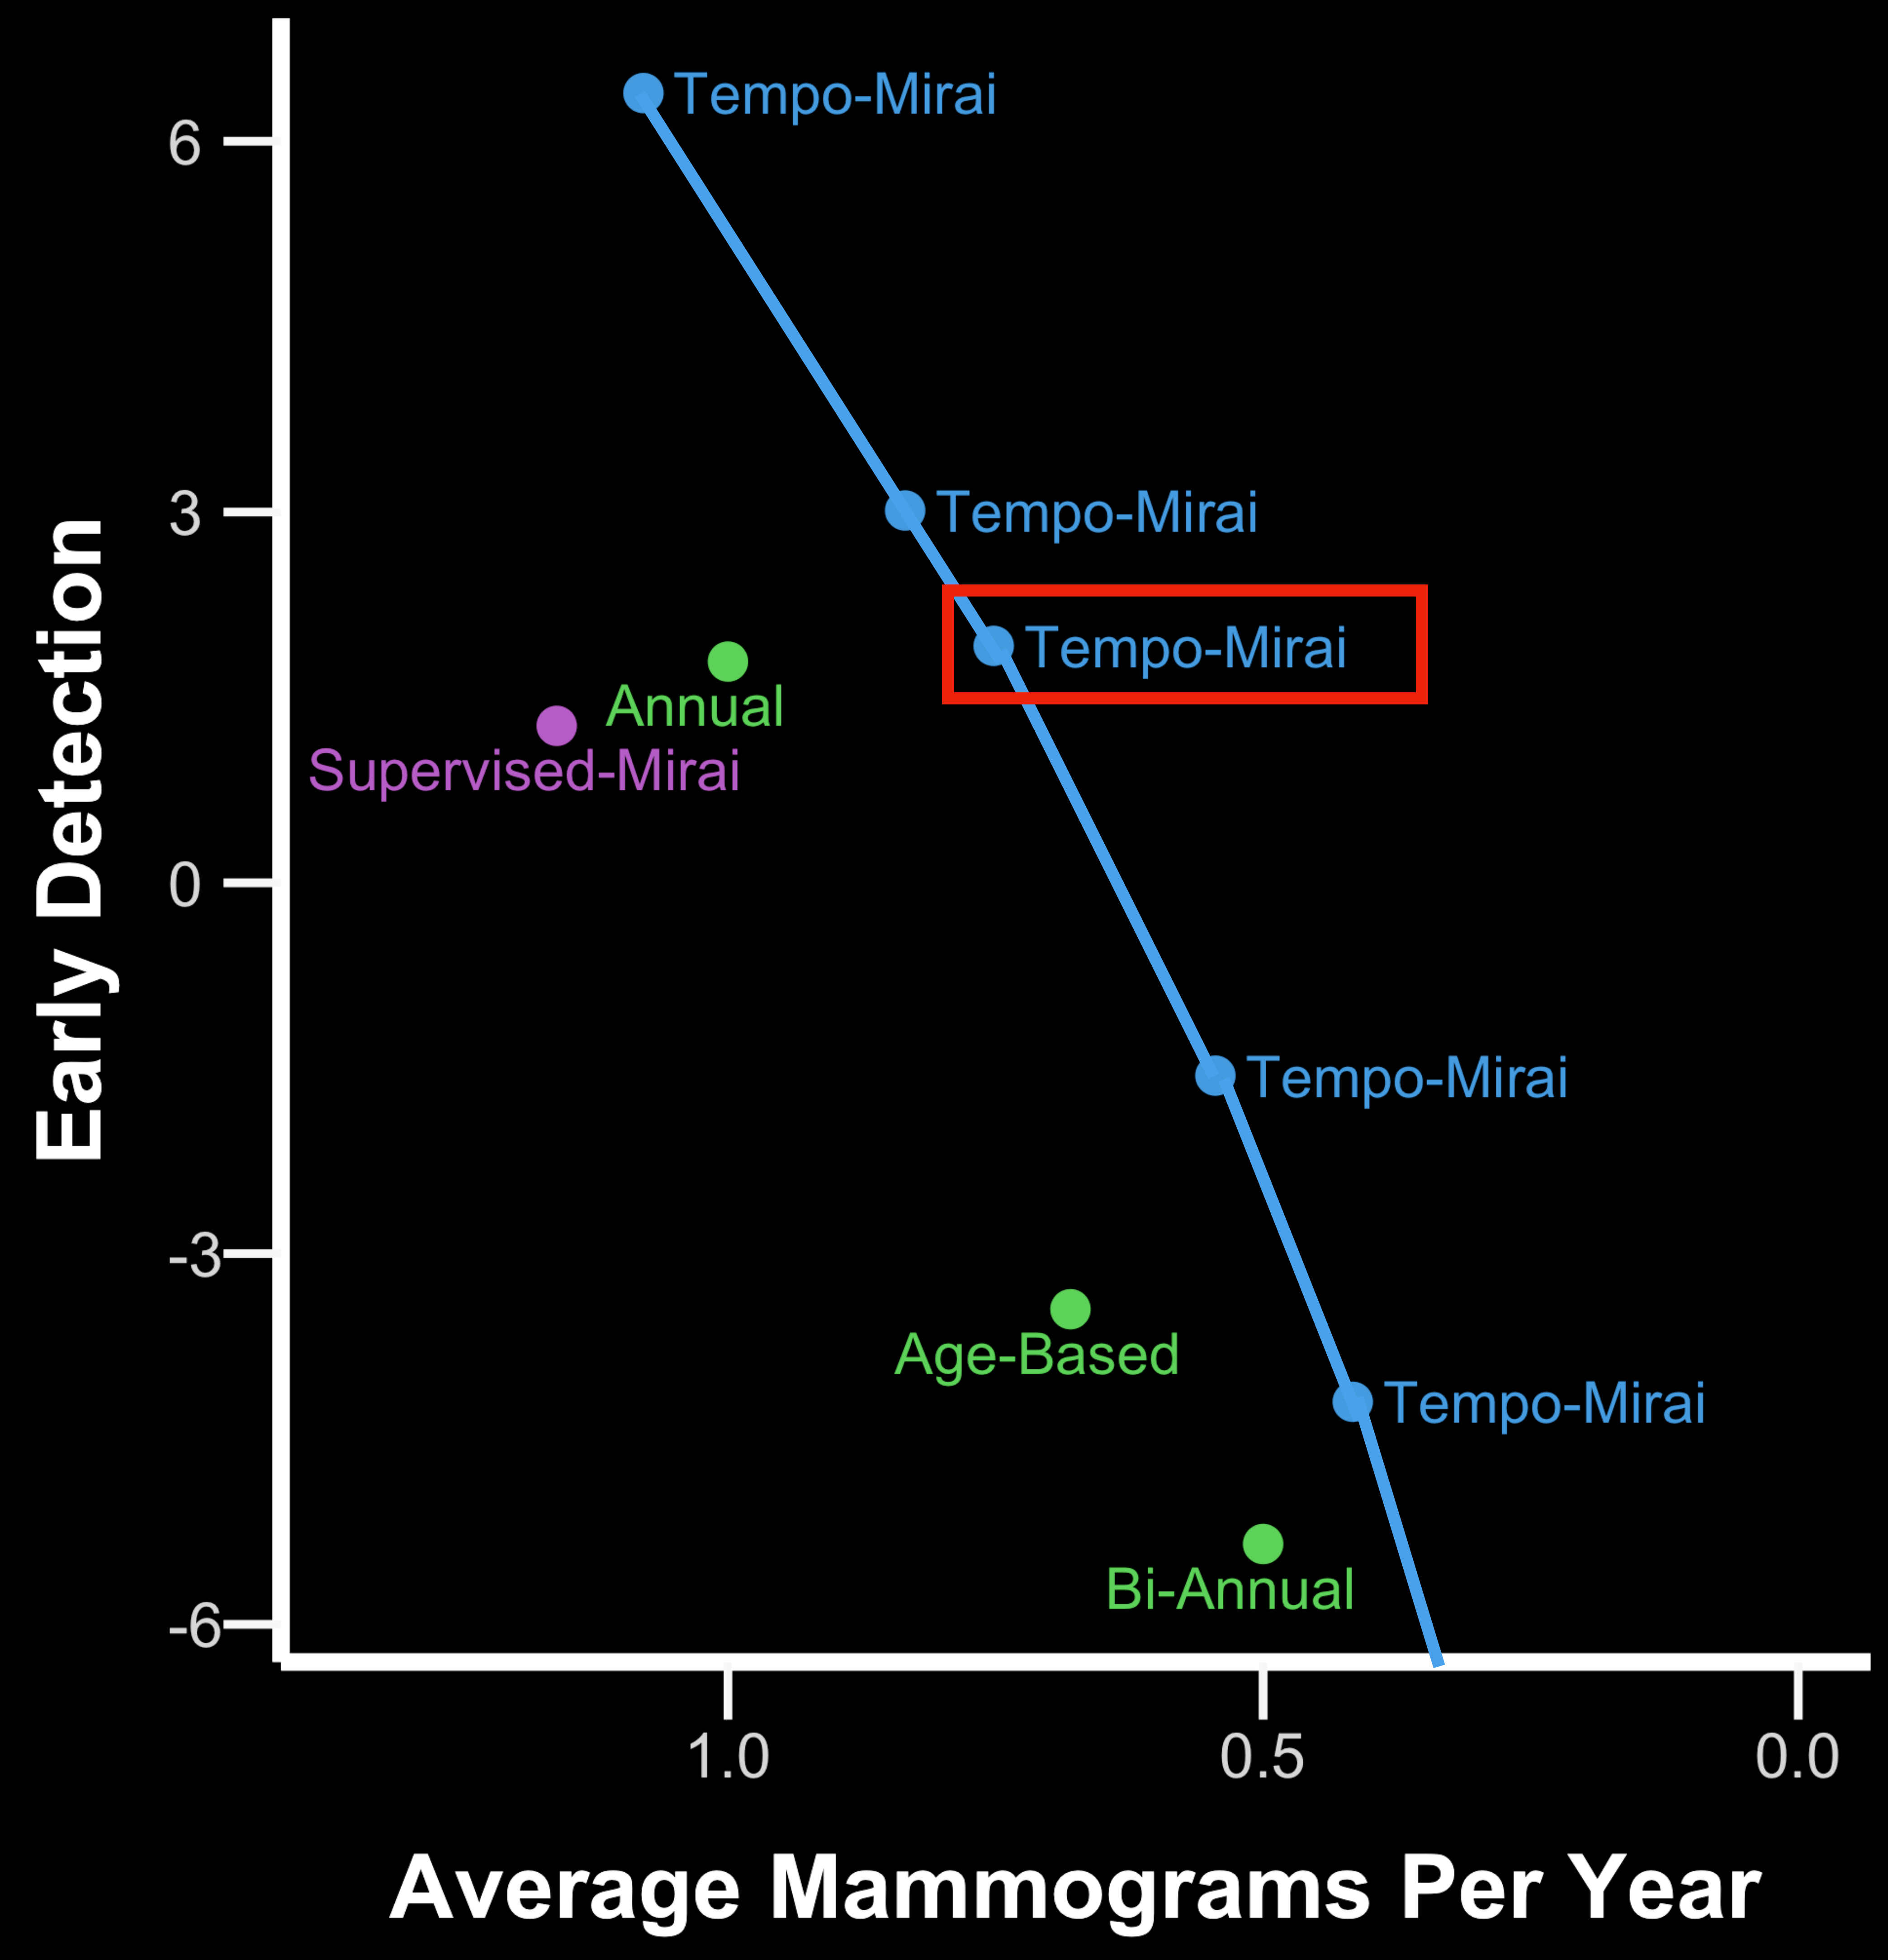

Emory Test Set

Summary

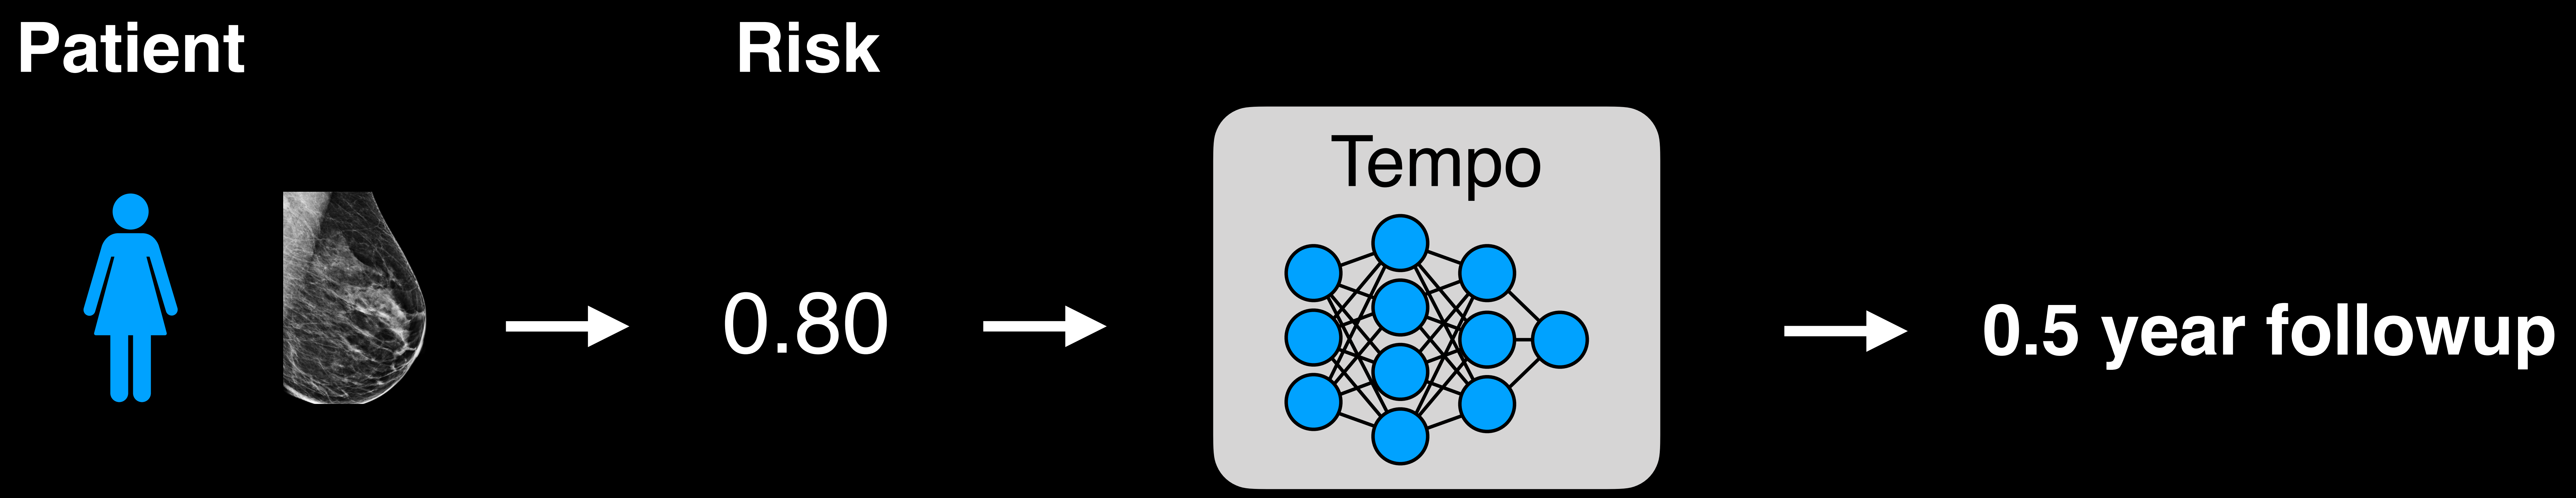

Learn personalized screening policies by modeling individuals

Applicable with arbitrary reward design / choice of risk model

Better **early detection and less overscreening**

Ongoing Work: AI to start cancer screening

Patient **EHR**

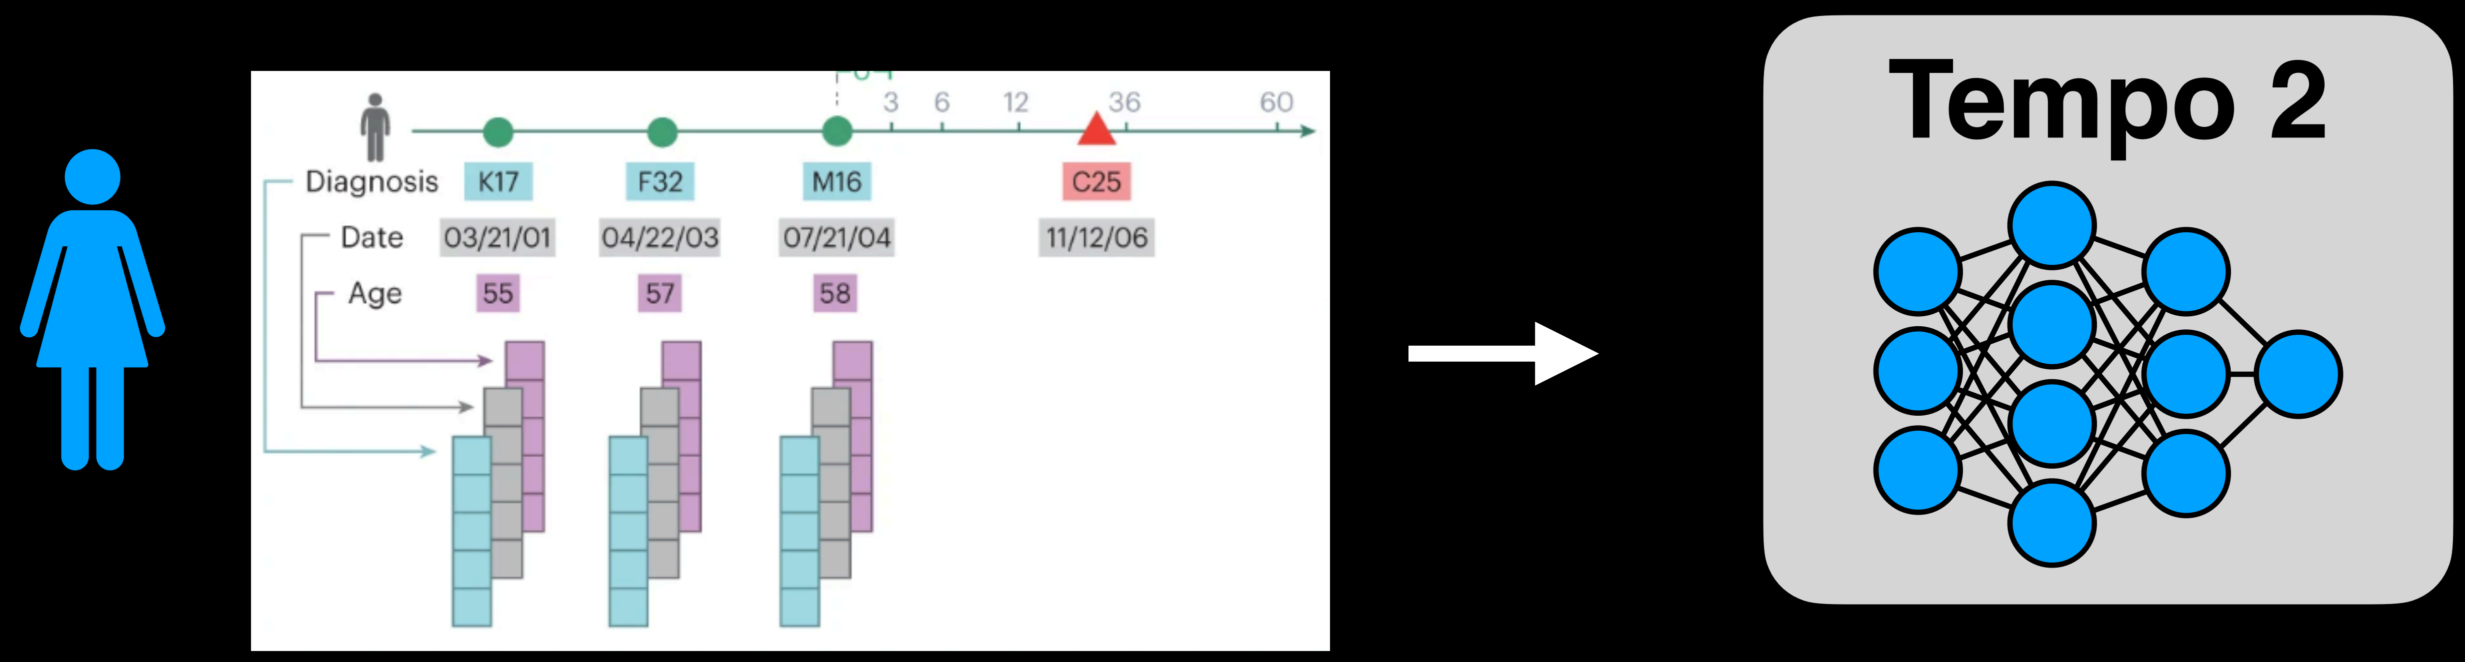

Mammo at 35

Model all disease codes in EHR + **EHR of Parents**

At each year, AI predicts if to screen for cancer [mammo, LDCT, etc]

Goal: Help women < 40, non-smokers who get lung cancer, etc.

Led by: Mikkel Odgaard

Ongoing Work: AI to start cancer screening

EHR Records + Info in Danish Registries

Patient

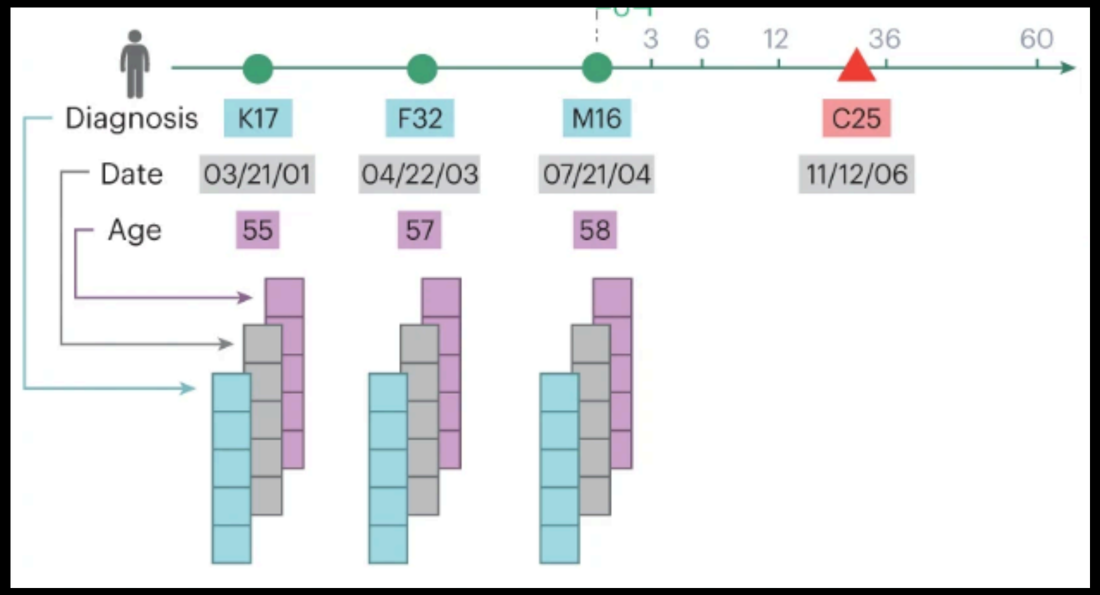

Tempo 2

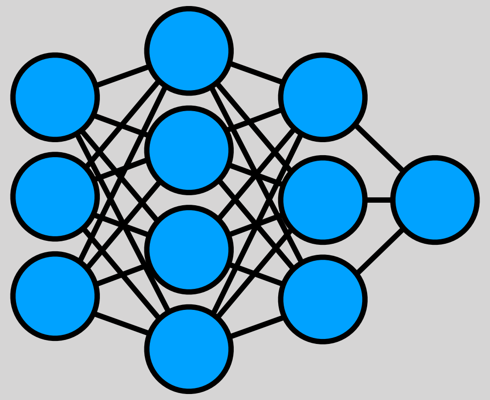

Cancer screening this year?

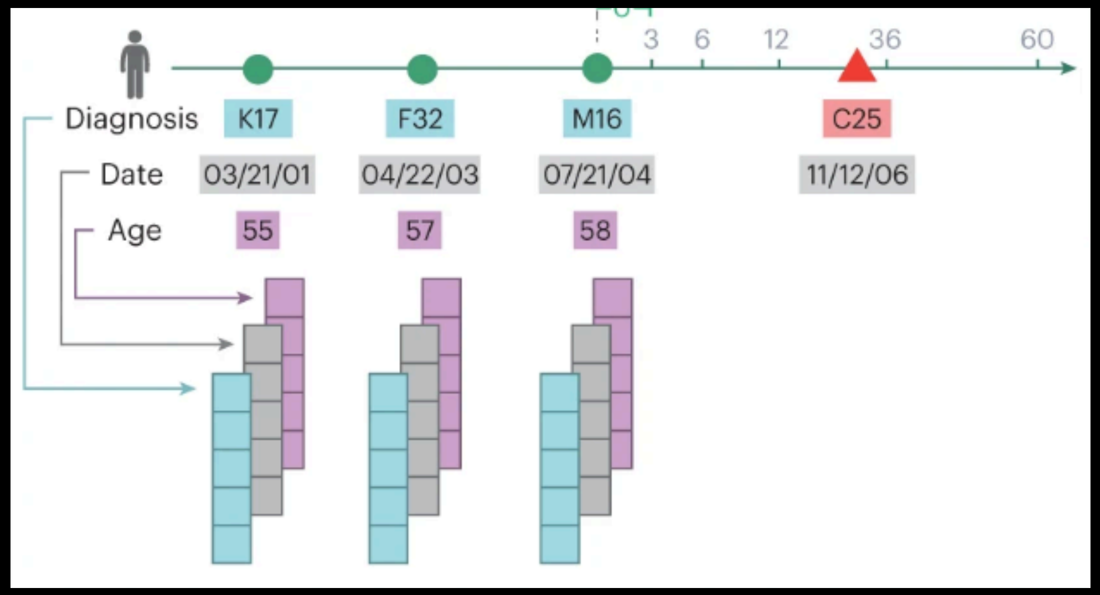

Parents

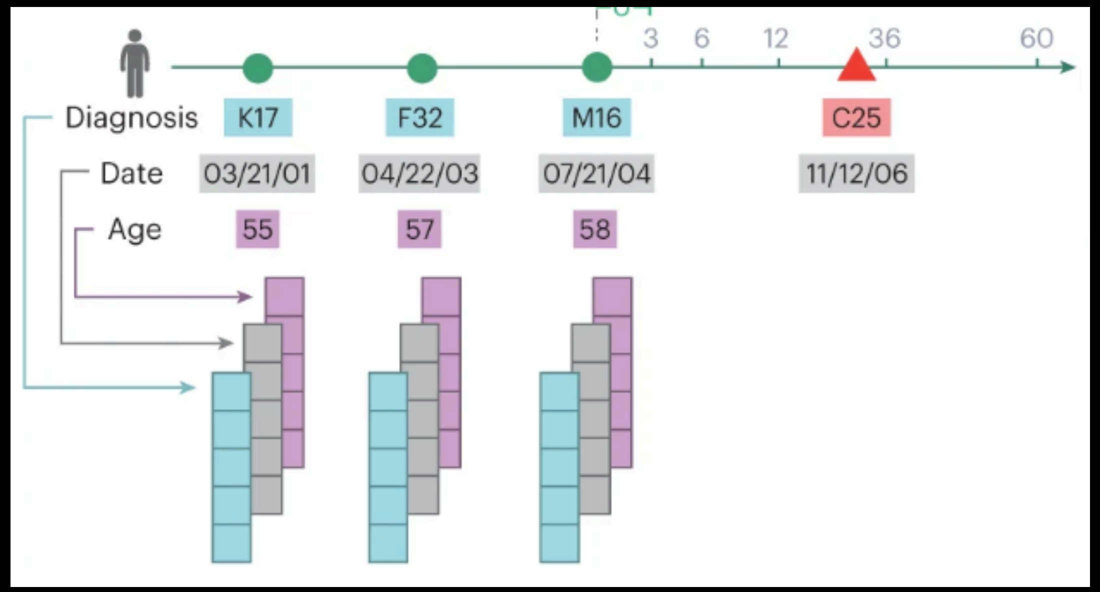

Today: Towards AI-driven care

Prediction

Control

Translation

Today: Towards AI-driven care

Translation

Ongoing Prospective Trials: Mirai-MRI

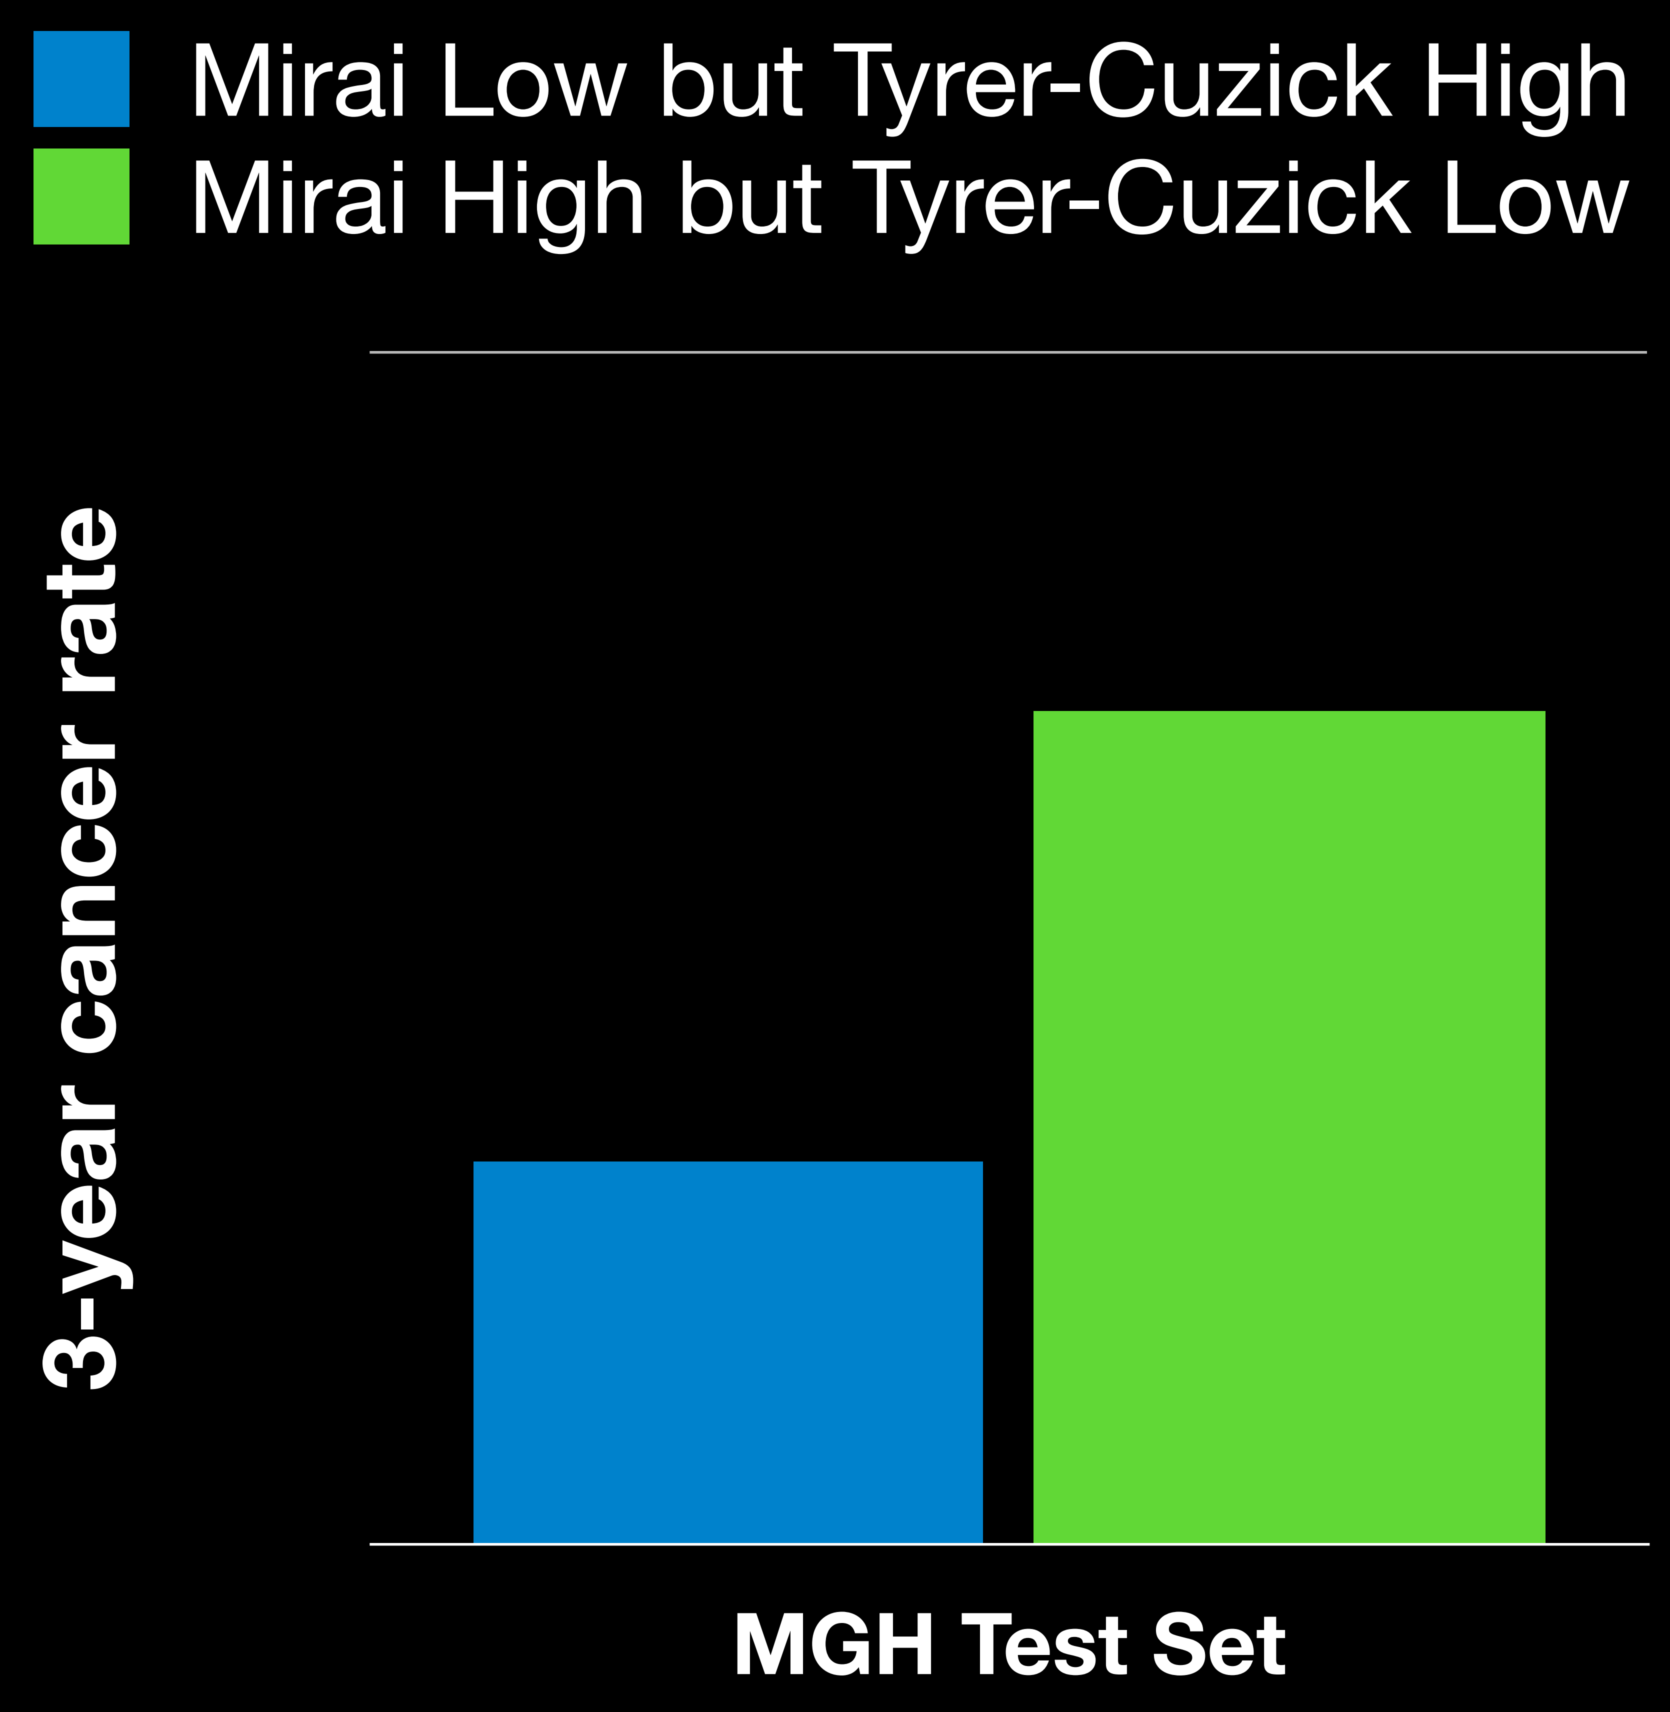

Retrospective analysis

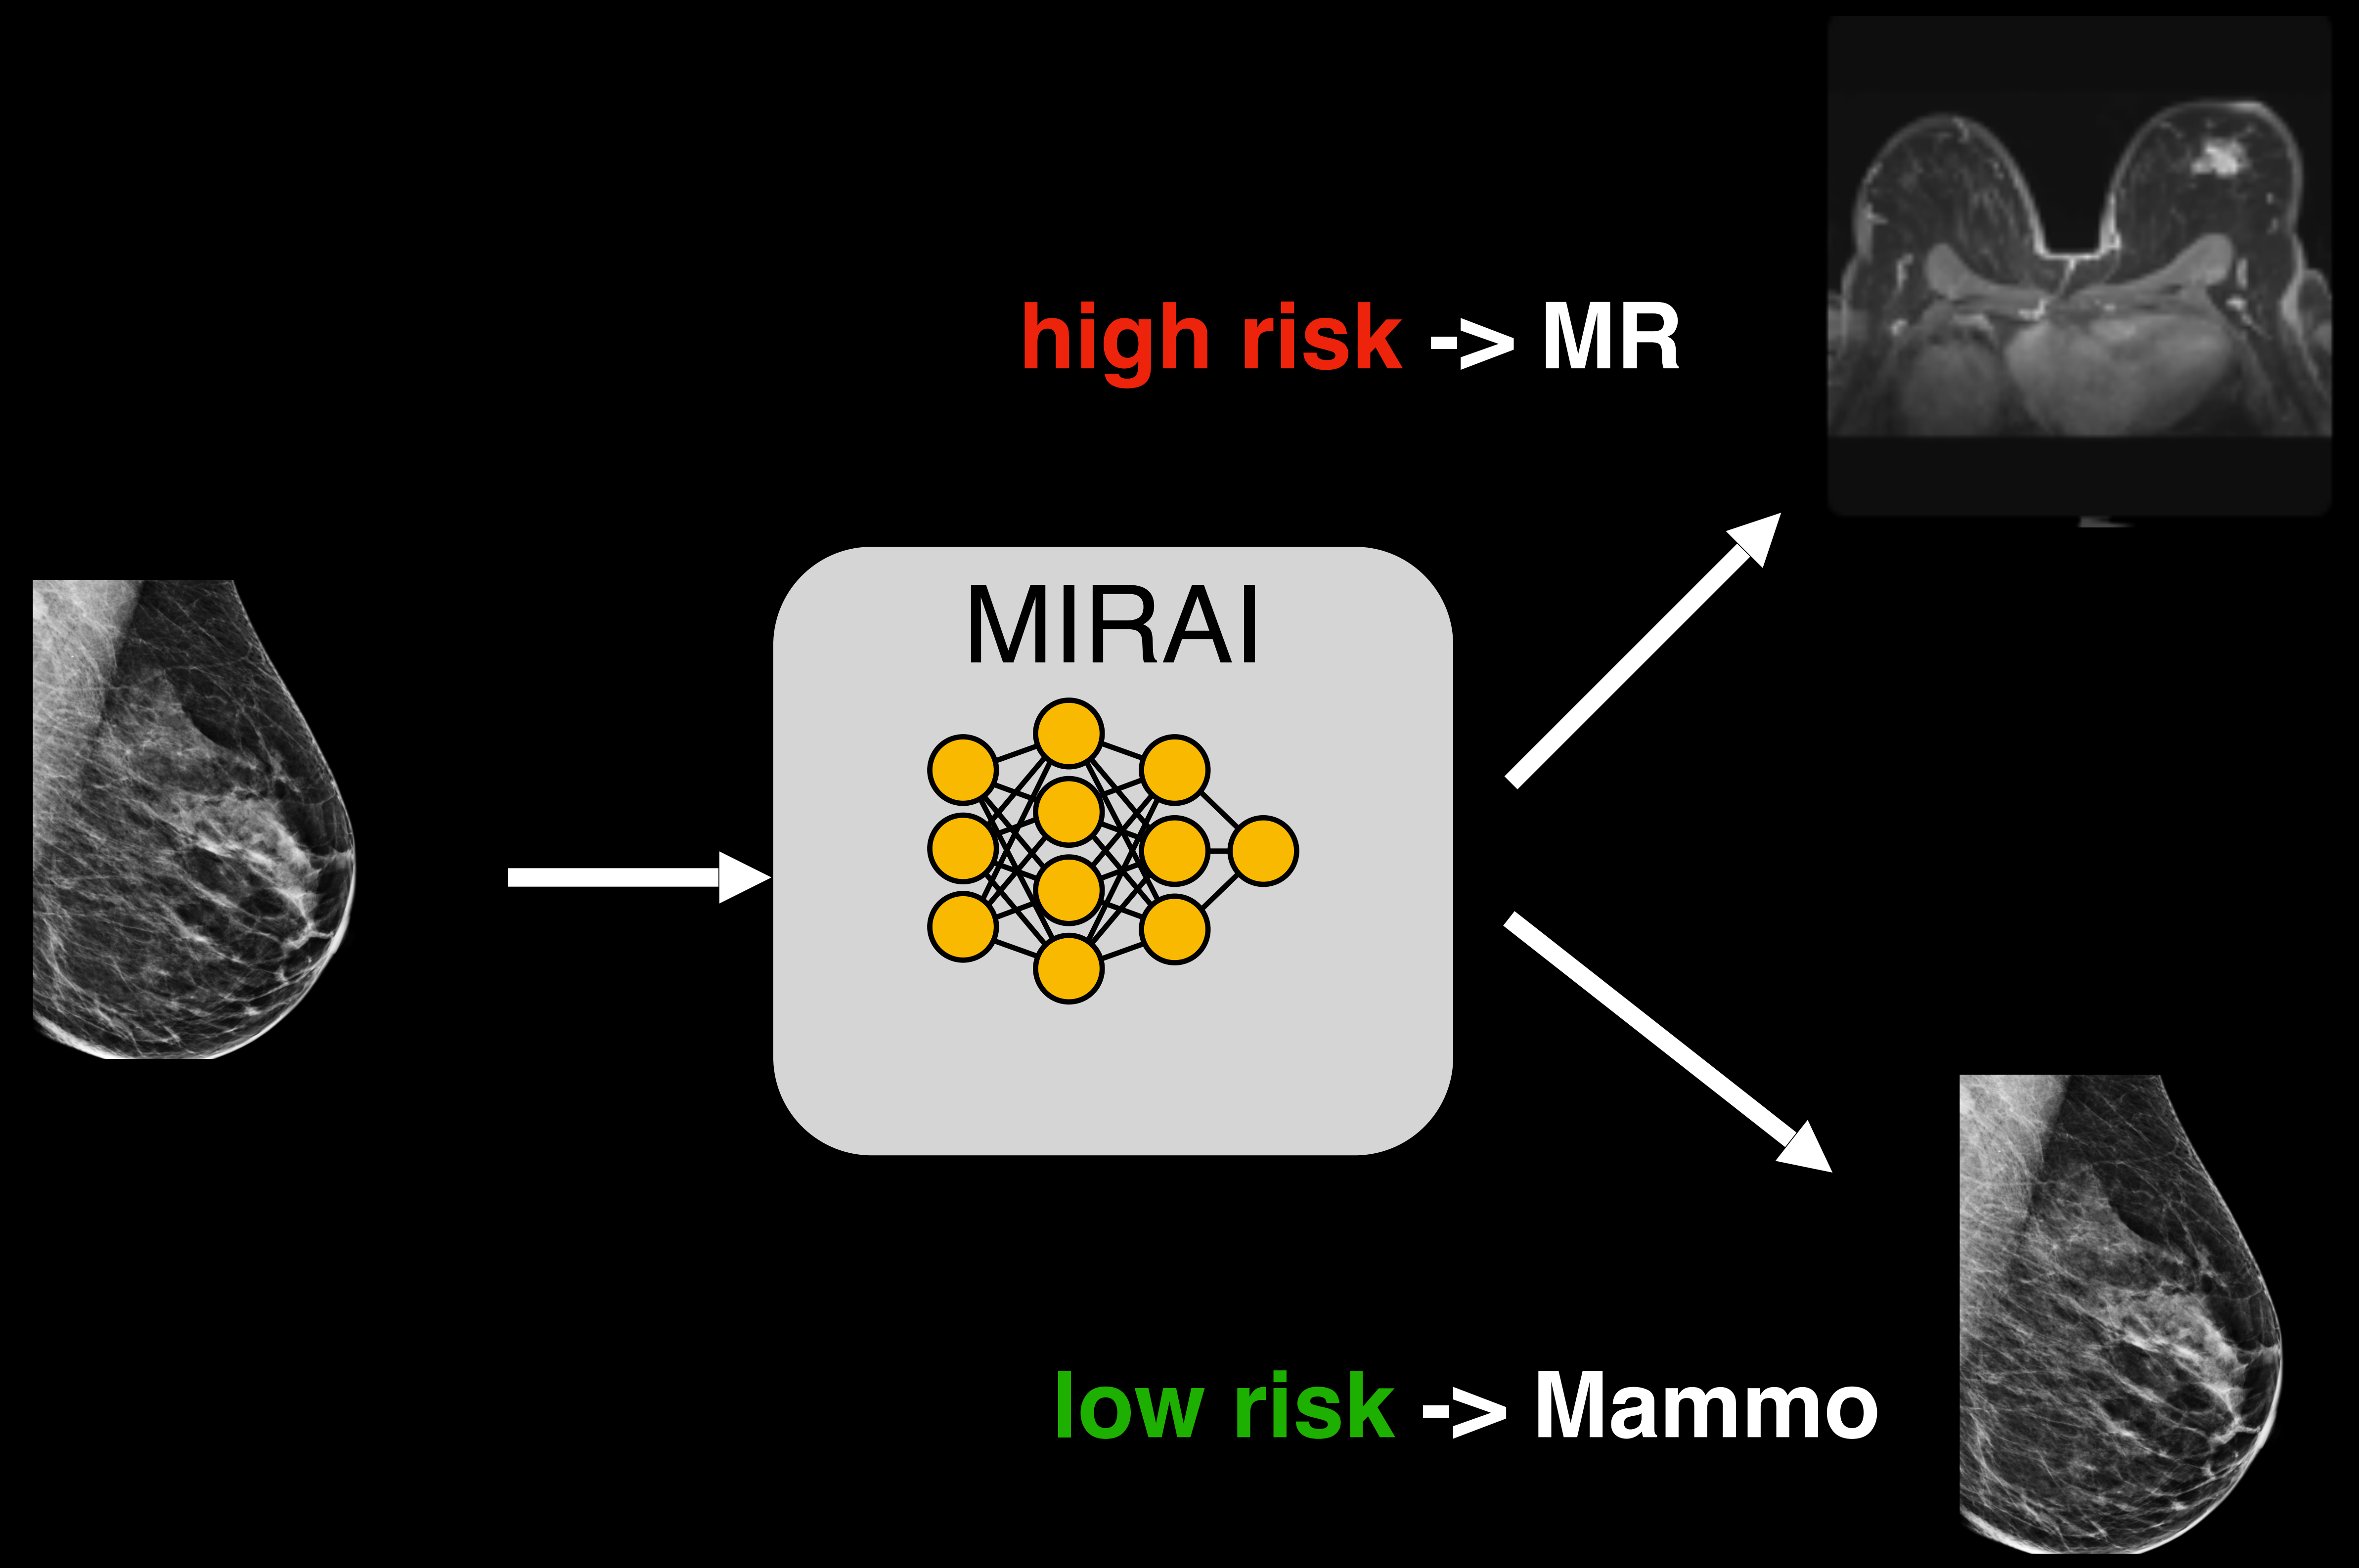

Mirai-based Supplemental Imaging

NCT 05968157

Ongoing Prospective Trials: Mirai-SDA

SDA Workflow:

- **Realtime** AI-based cancer risk assessment
- Invite high risk patient for **same-day** diagnostic exam
- Expectation: >50% of cancer cases will receive same-day diagnostic
- **100%** enrollment so far!

Screen

<5 seconds

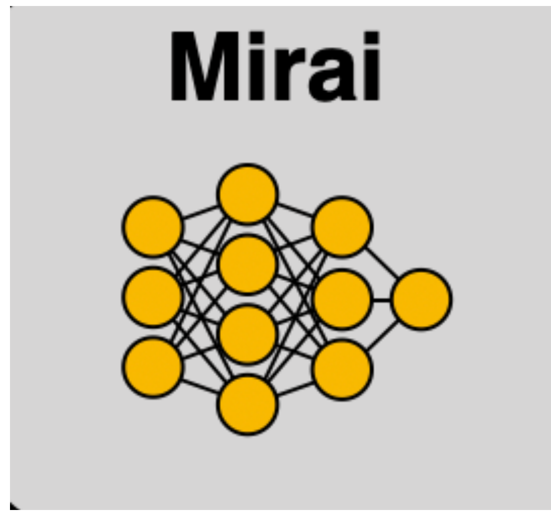

AI Risk

Top 10% Risk

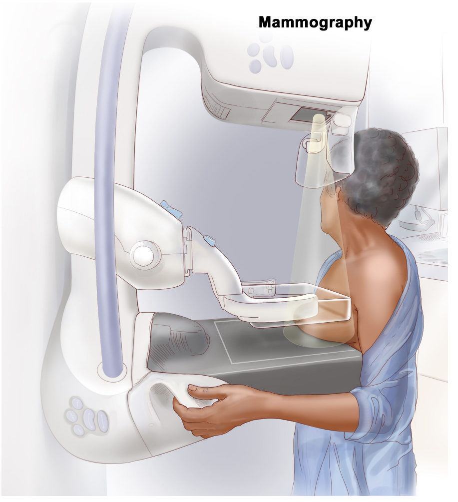

Same Day Diagnostic

How do we evaluate constant evolving AI tools?

Evaluation

Traditional Randomized Trials

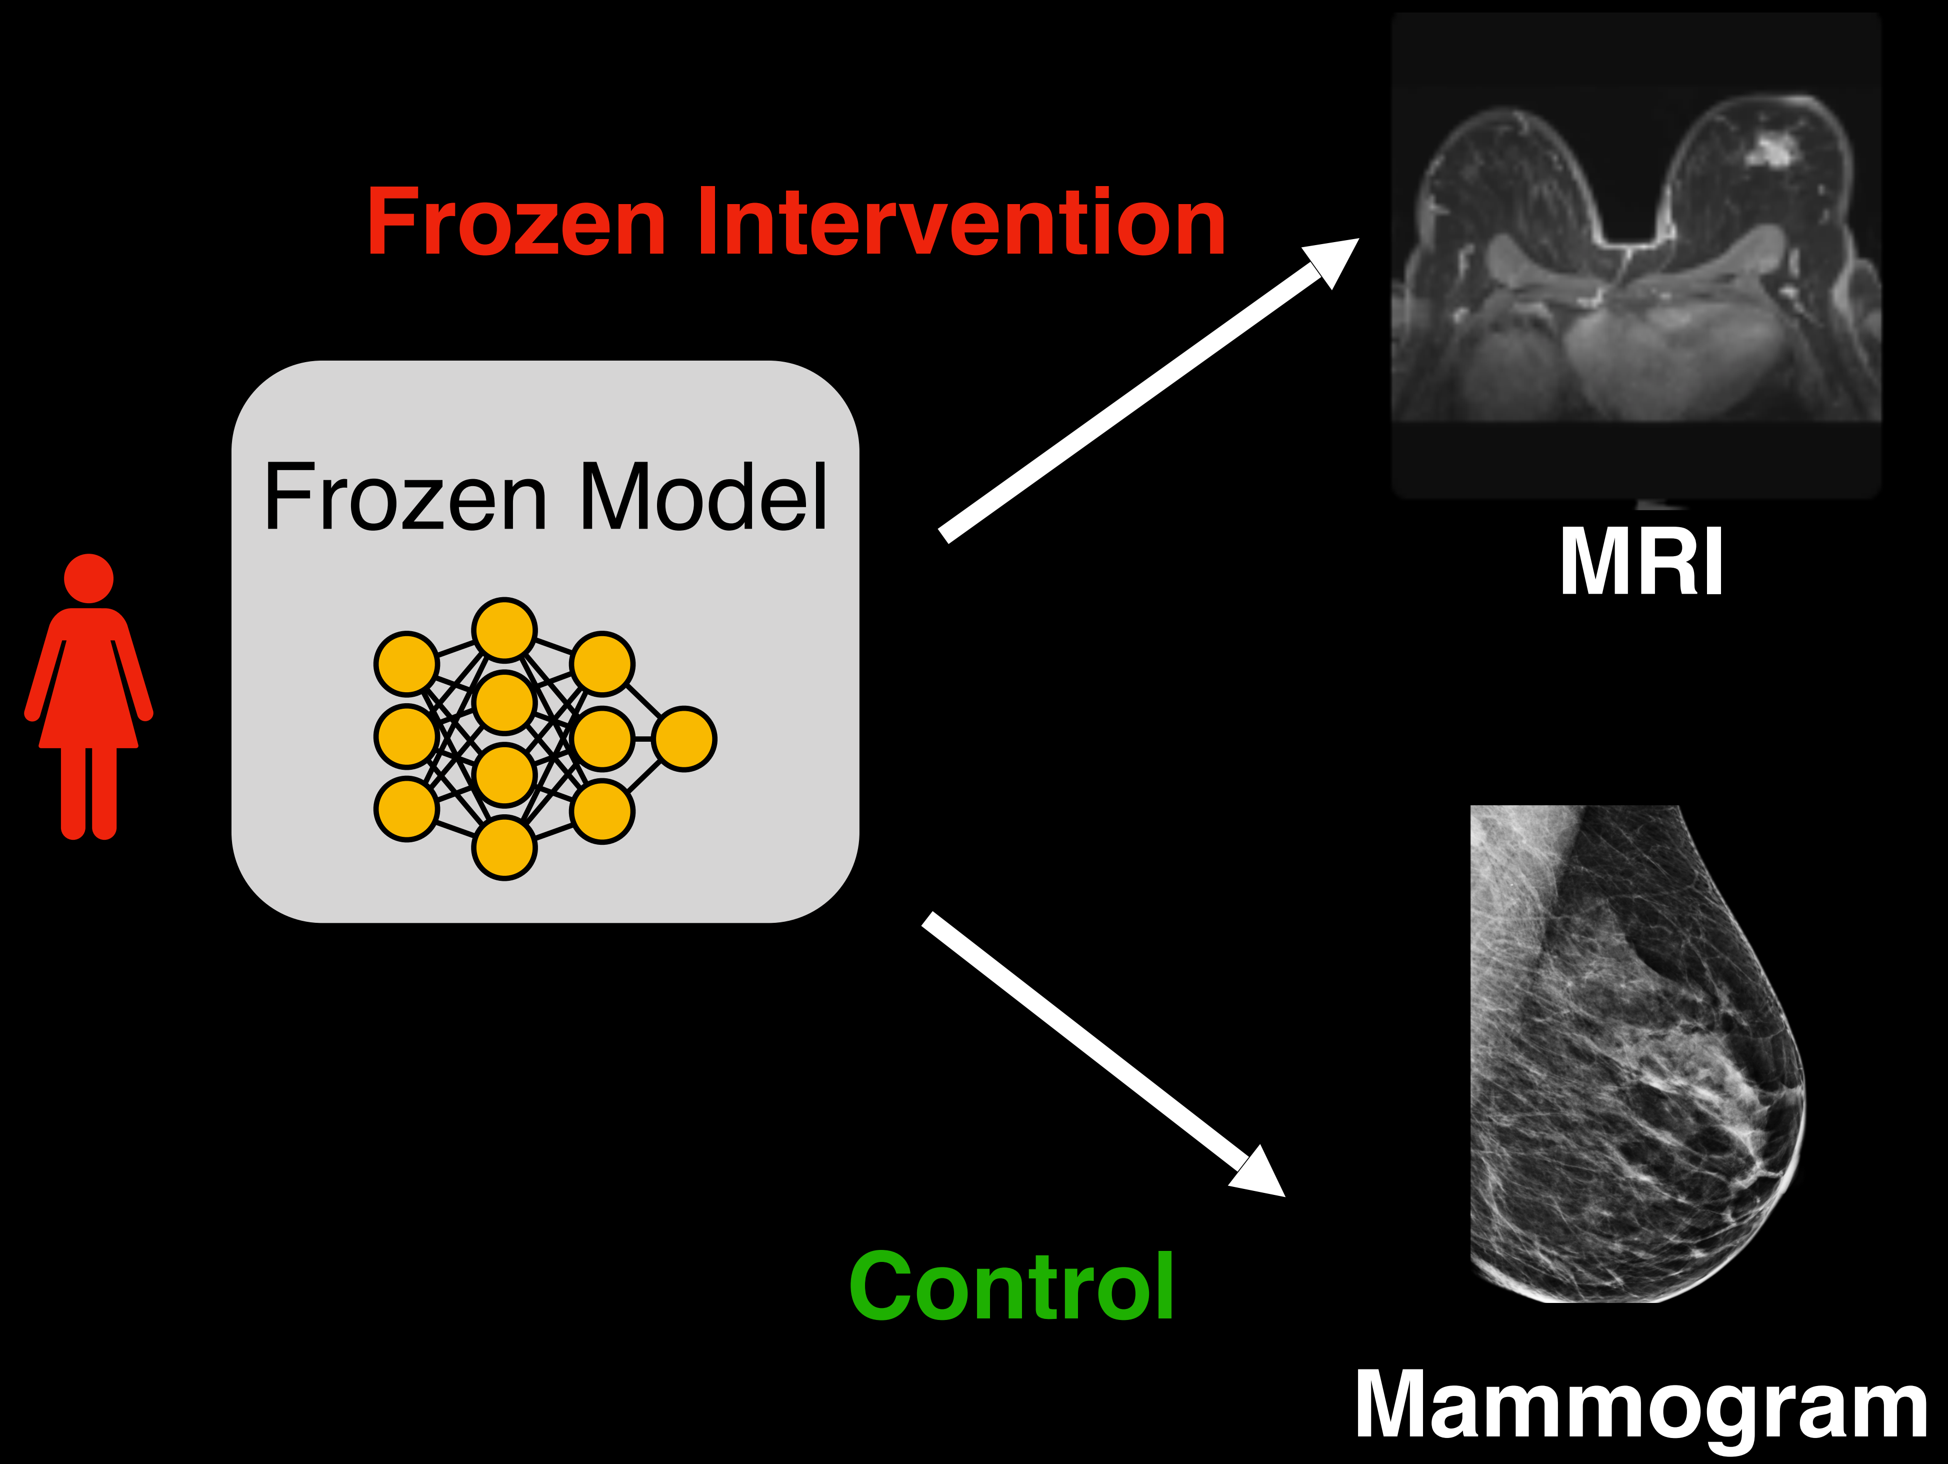

AI **obsolete** by end of trial..
Incompatible with rapid model innovation

Led by: Wenxin Zhang

Ongoing work: Reusable and AI-Adaptive RCTs

Adaptive AI-Platform Trials + RWE

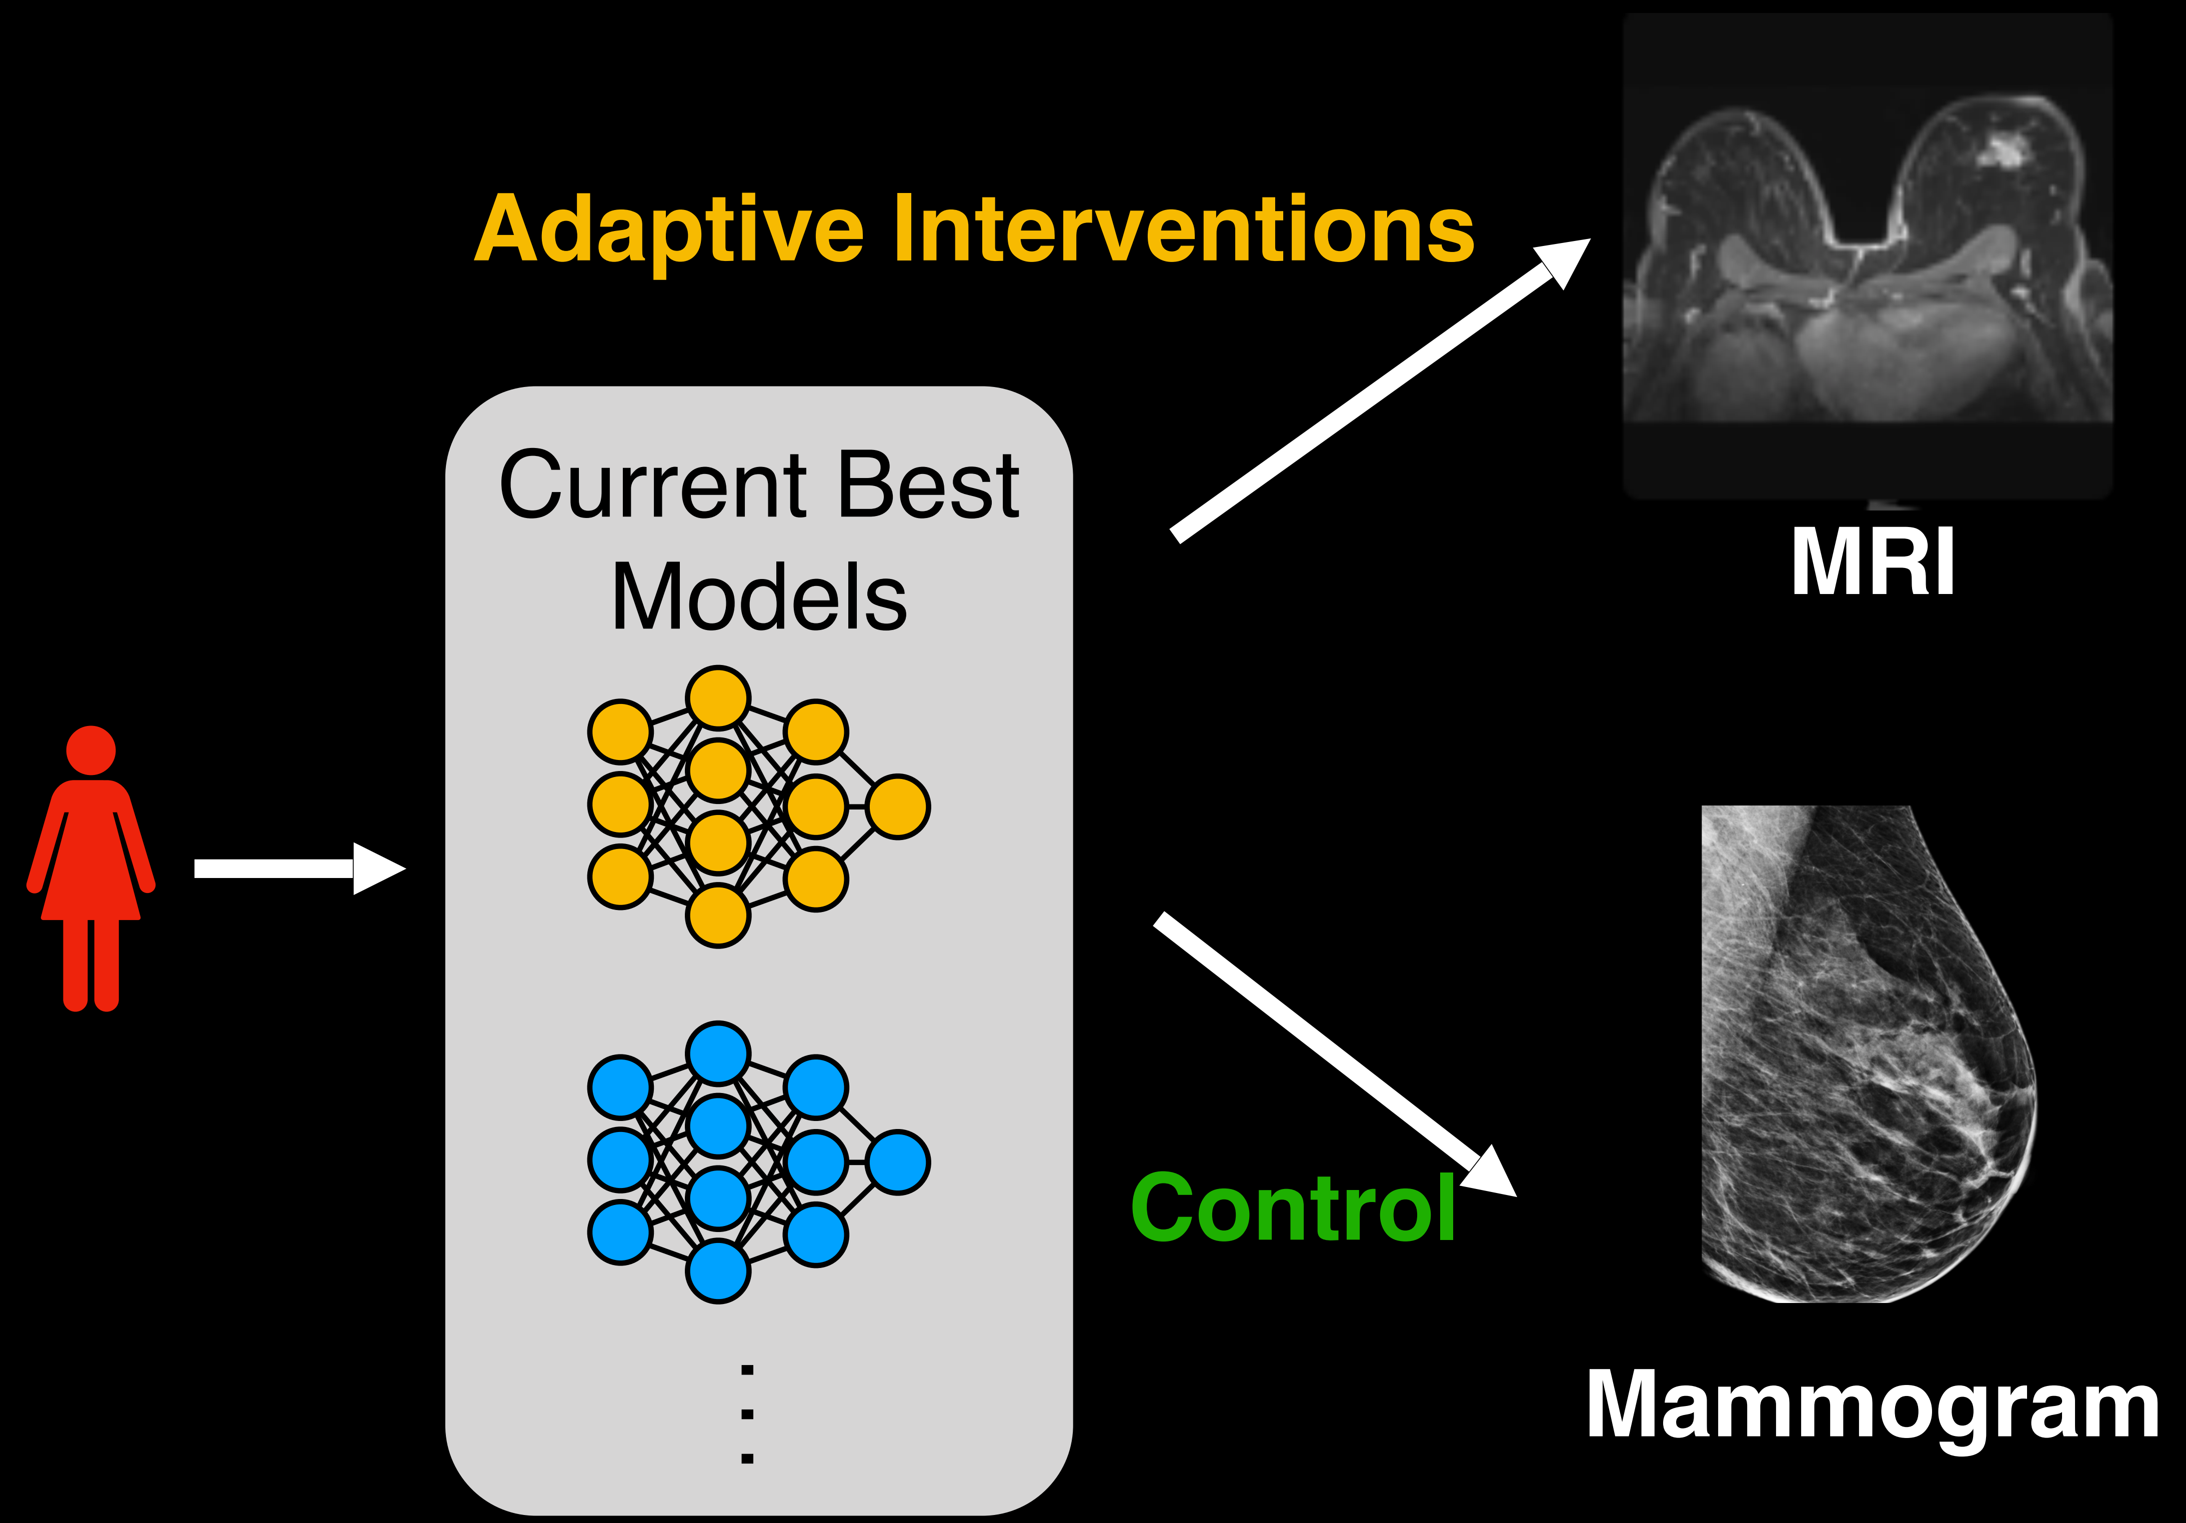

Continuous **platform trials** for improving AI
Combine evidence across all models + historical data (RWE). **Fast**

Key Ideas:

- AI model **allocates** intervention
- Models generations **overlap** in decisions
- Re-use data and do adaptive enrollment

Simulation: Mirai-SDA

Adaptive AI-Platform Trials + RWE

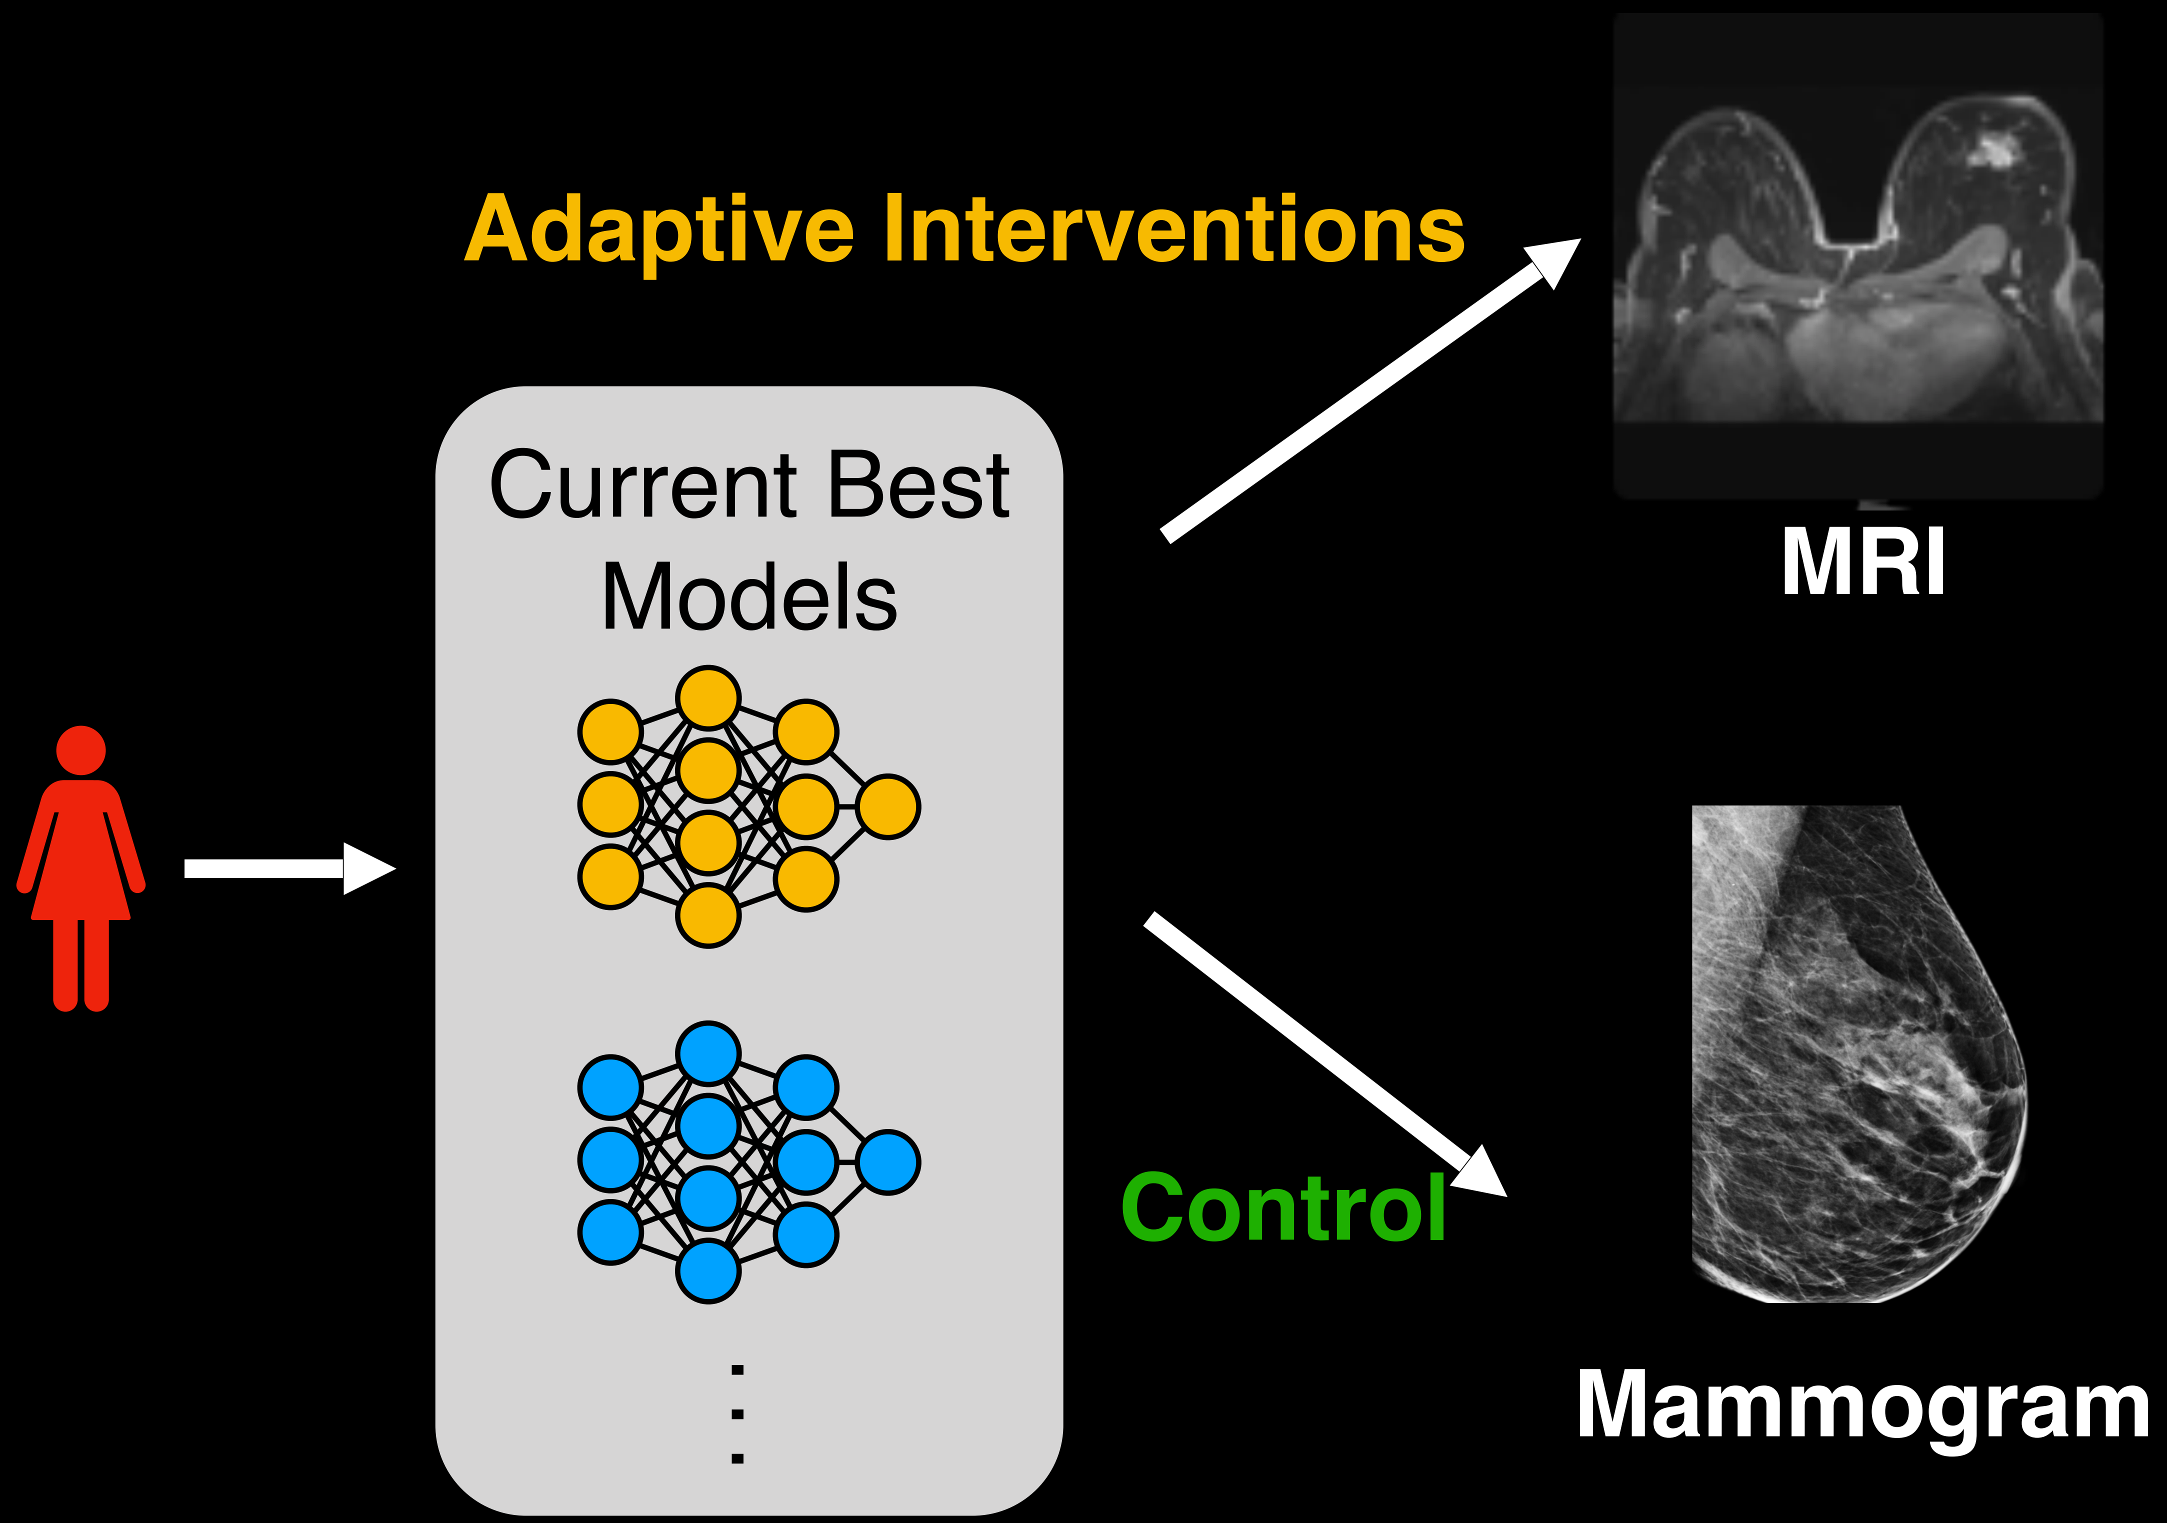

Continuous **platform trials** for improving AI
Combine evidence across all models + historical data (RWE). **Fast**

Mirai-1 Trial:

- Enroll 100 patients

Mirai-1.5 Trial

- Naive trial: 100 new patients
- **70 %** patients high risk by M1.5 and M1.0
- With data-reuse: 30 new patients

Trials will only get faster and easier

Ongoing work: BRIDGE Adaptive RCTs

Clinical domain	Legacy model	New model	Same training dataset?	Same input features?	Same model endpoint?	Top-5% Overlap (%)
Breast cancer	AI-Density	Mirai	✓	✓	—	2.5
	ImgOnly DL	Mirai	✓	✓	✓	46.6
Cardiovascular disease	SEER	S4-ECG	—	✓	—	14.2
	ResNet	S4-ECG	✓	✓	✓	49.6
Sepsis	LSTM-Dynamic	LSTM-Full	✓	—	✓	49.3
	LSTM-Full	GRU-Full	✓	✓	✓	52.3

Overlaps are common!

Examples across :

- **breast cancer (Mammo)**
- **cardiovascular disease (ECG)**
- **sepsis (EMR)**

Ongoing work: BRIDGE Adaptive RCTs

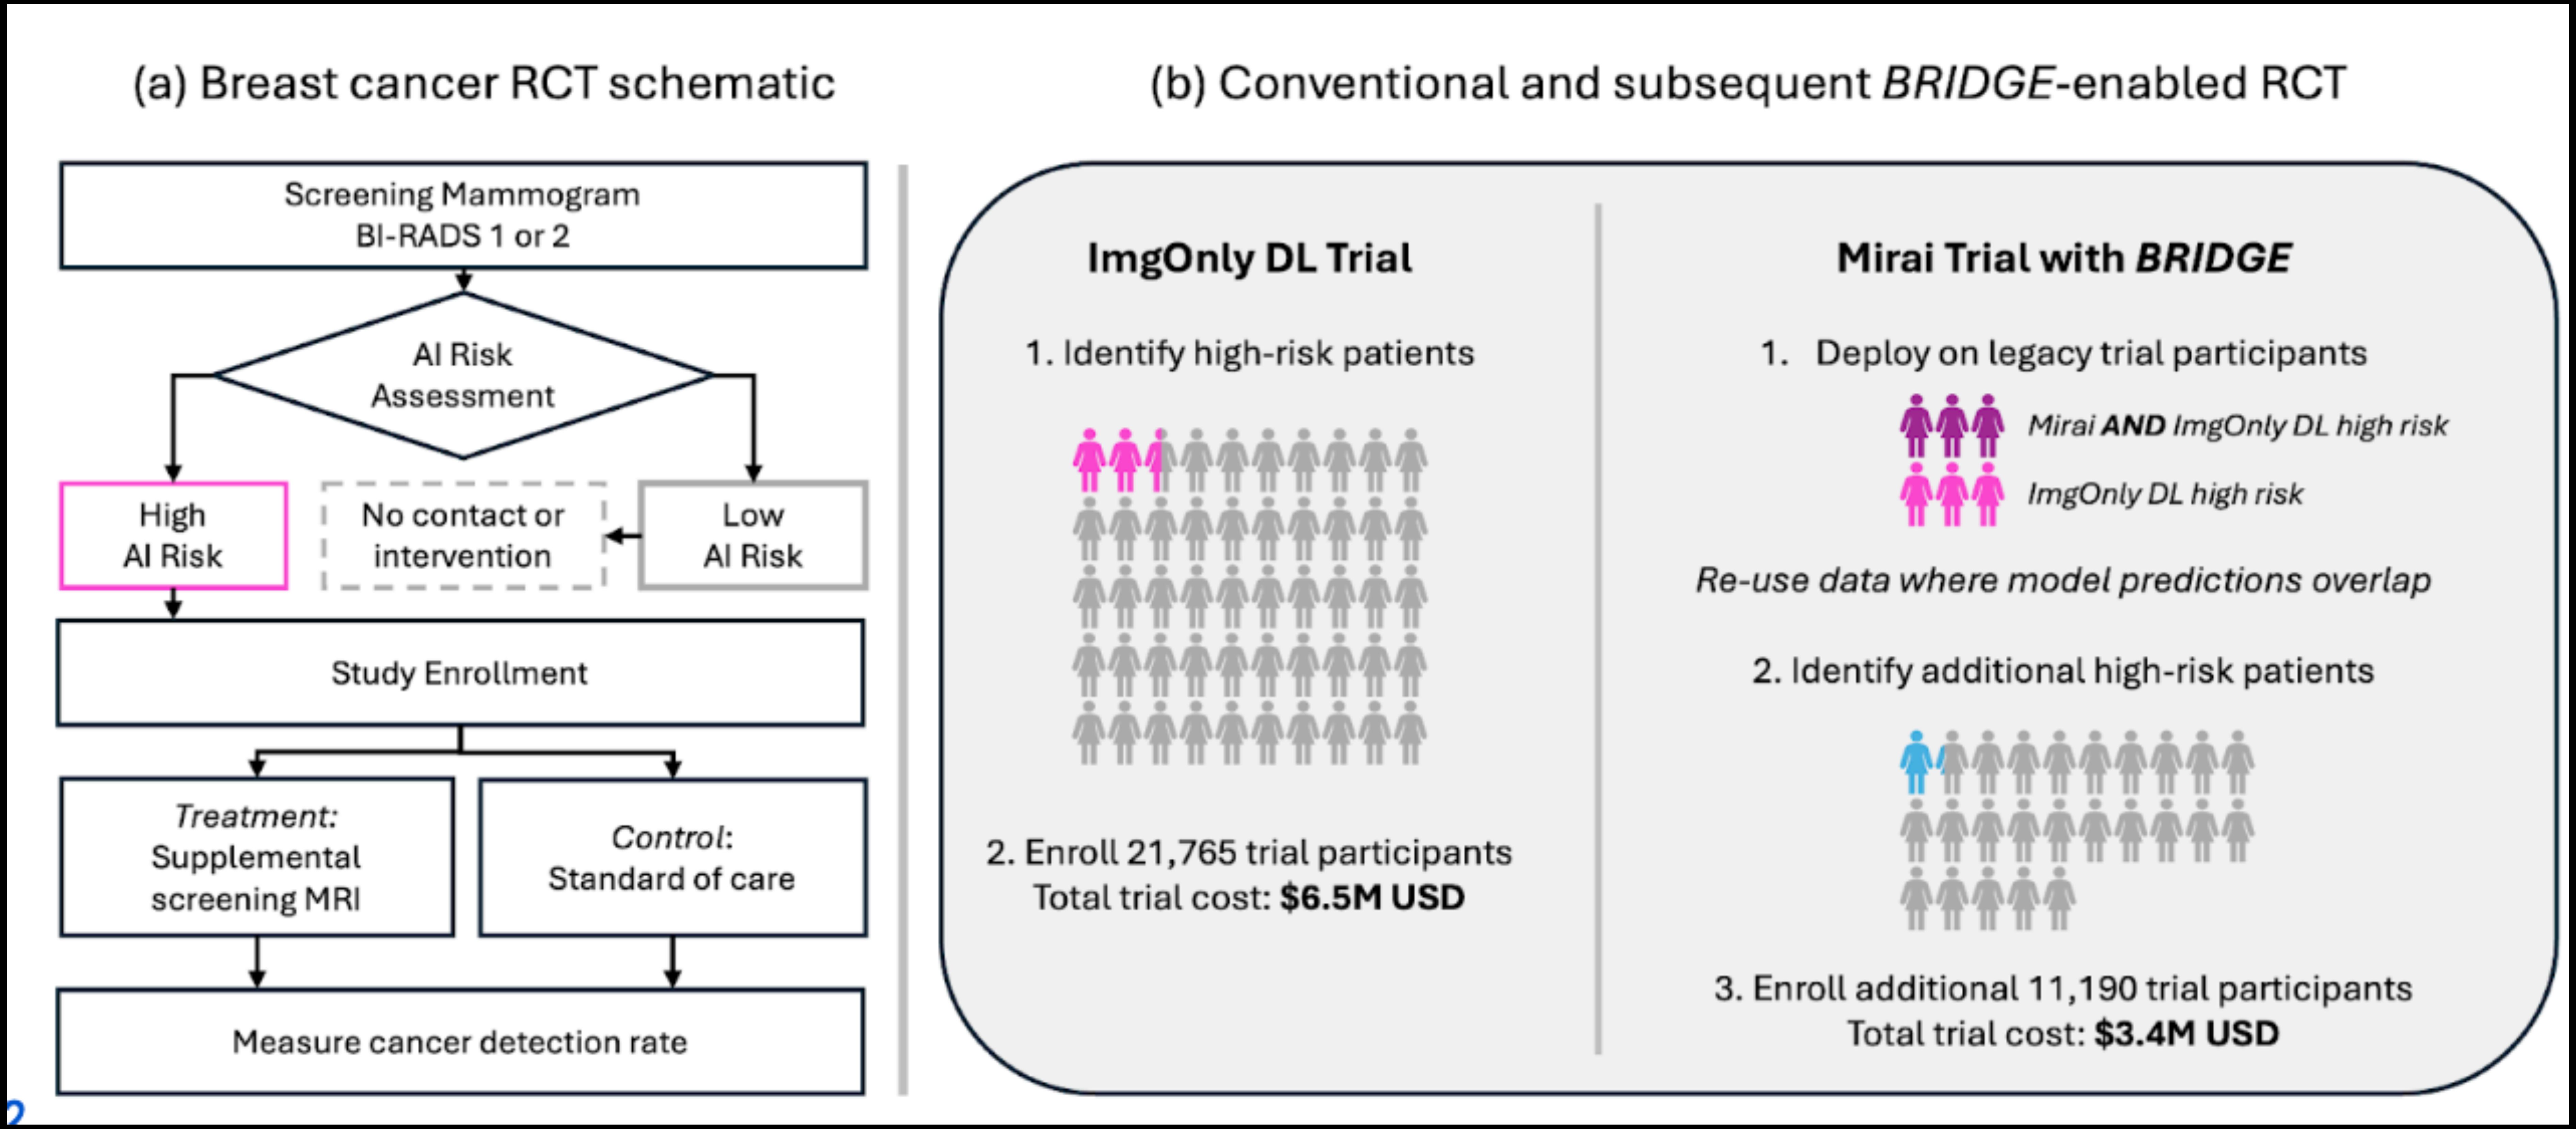

Recap: Towards AI-driven care

Prediction

Control

Translation

What if it all works?

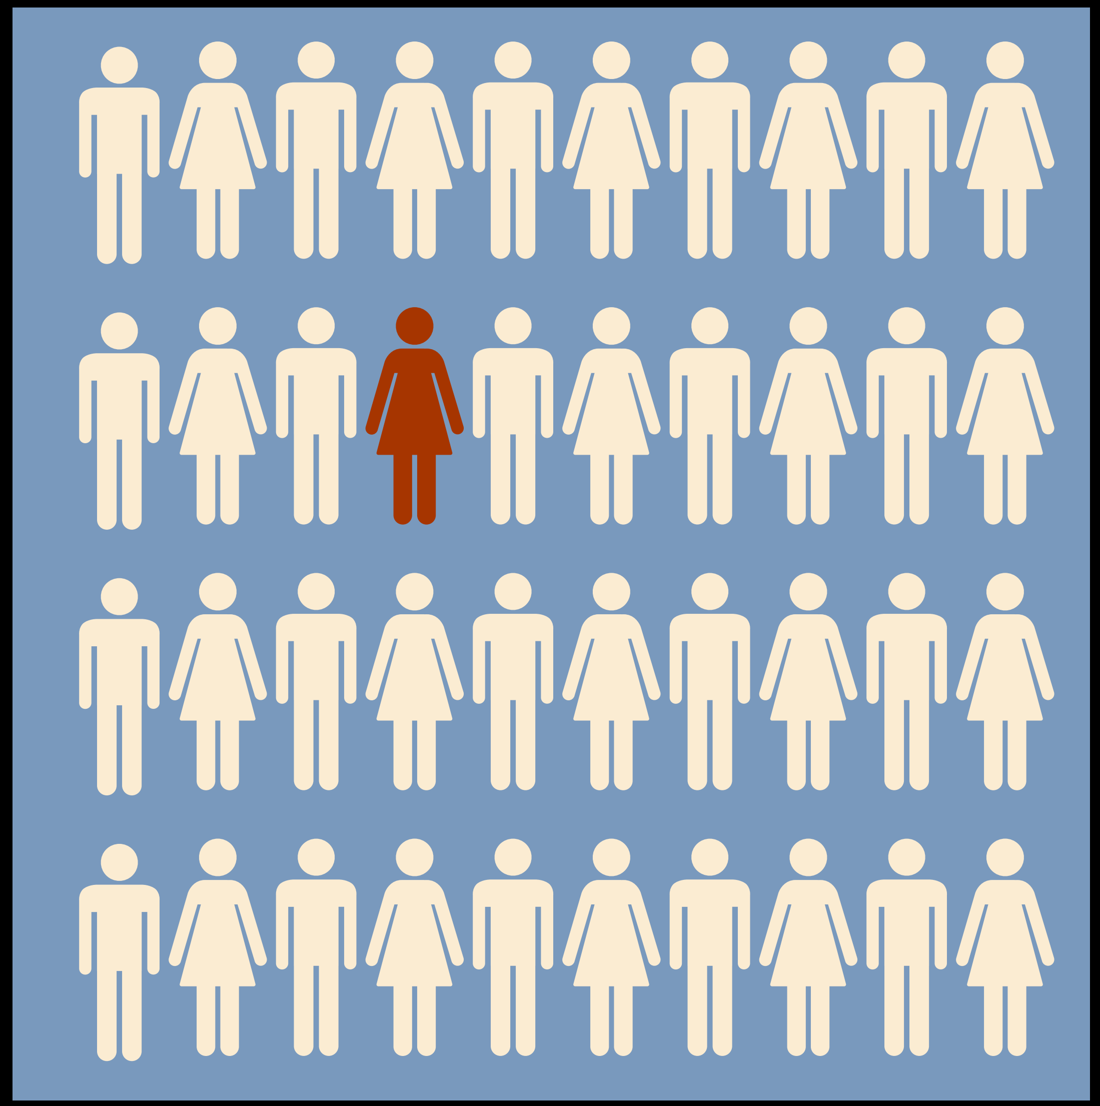

Rethink screening criteria / guidelines across diseases

New doors for prevention and therapeutic development

Questions?