Digital Clinical Trials for Oncology Patients with Novel Machine Learning and AI Architectures

Pratik Shah. Ph.D.
Principal Investigator

Disclosures: none

Pratik Shah. Ph.D.
pratiks@media.mit.edu

Part 1: Limitations and applications of machine learning (and AI)

(Automation vs. Knowledge vs. Intelligence)

mit media lab
Pratik Shah. Ph.D.
pratiks@media.mit.edu

Machine Learning

www.machinelearningmastery.com/a-tour-of-machine-learning-algorithms/

mit media lab
Pratik Shah. Ph.D.
pratiks@media.mit.edu

Research Areas

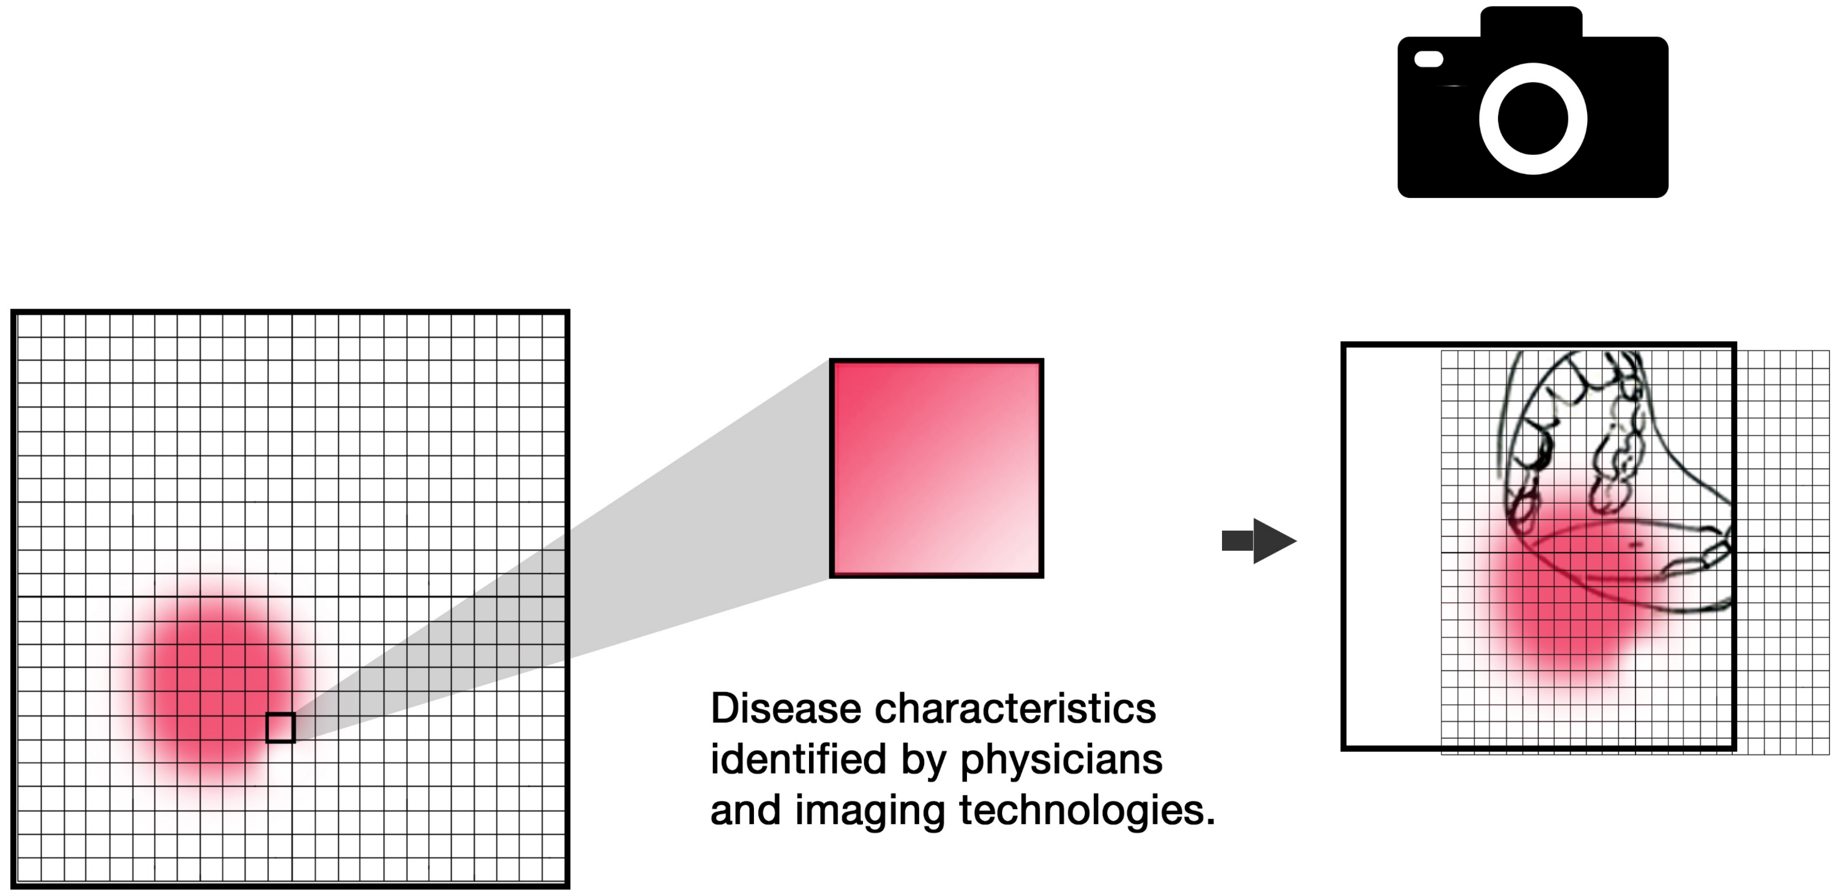

Machine saliency and unorthodox imaging of biomarkers

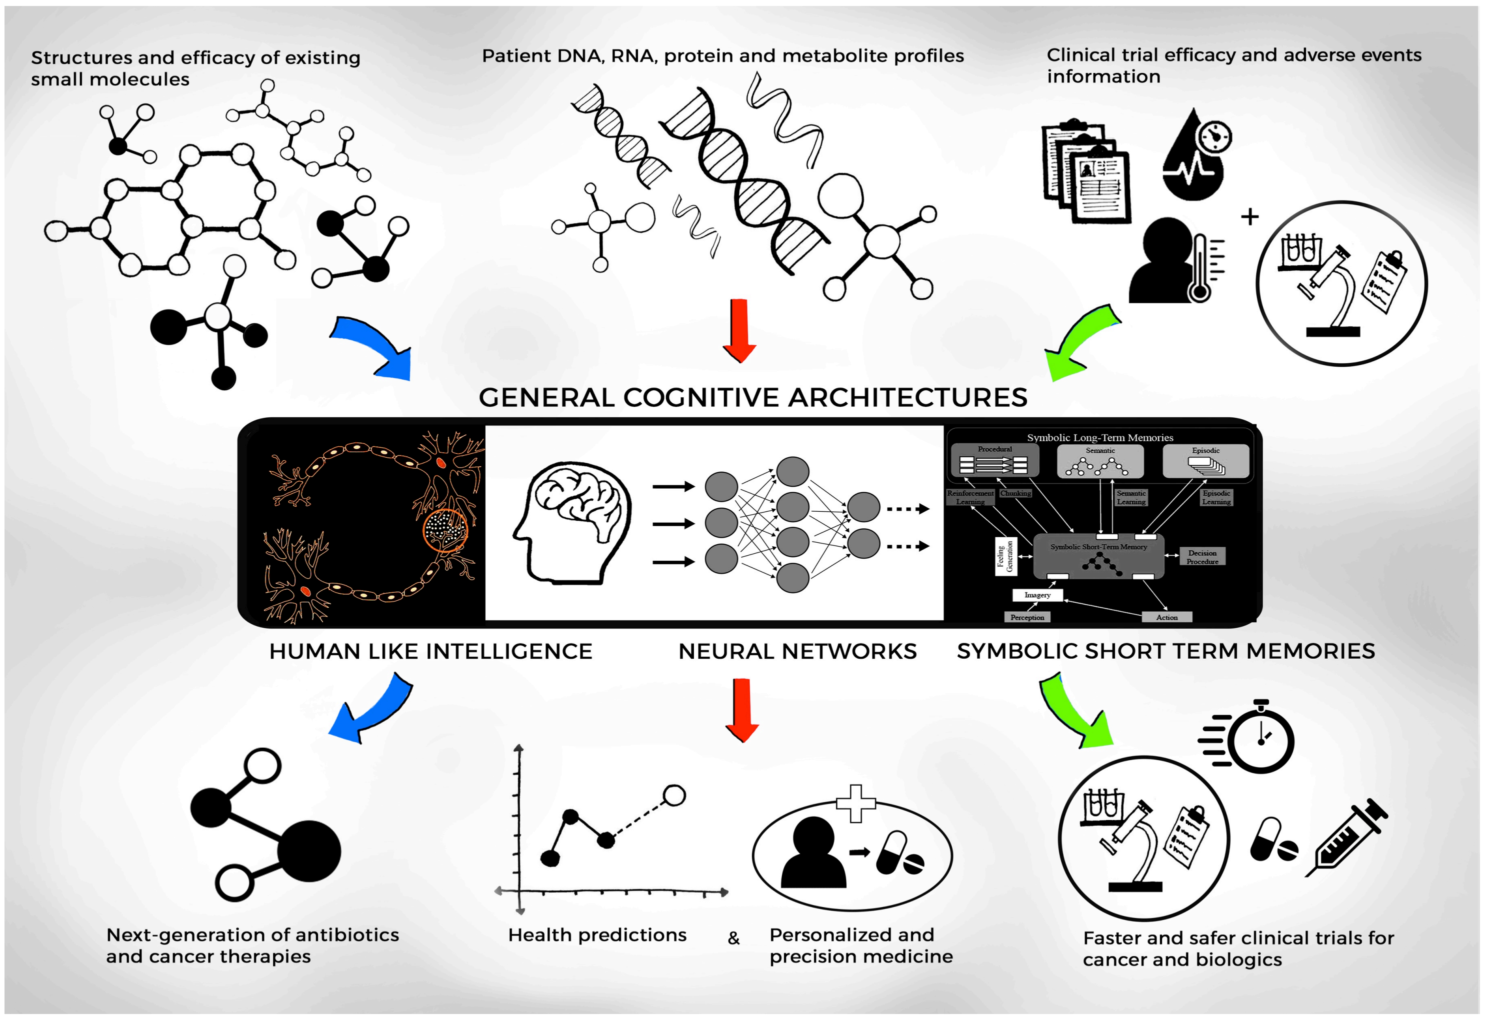

Novel clinical trial platforms

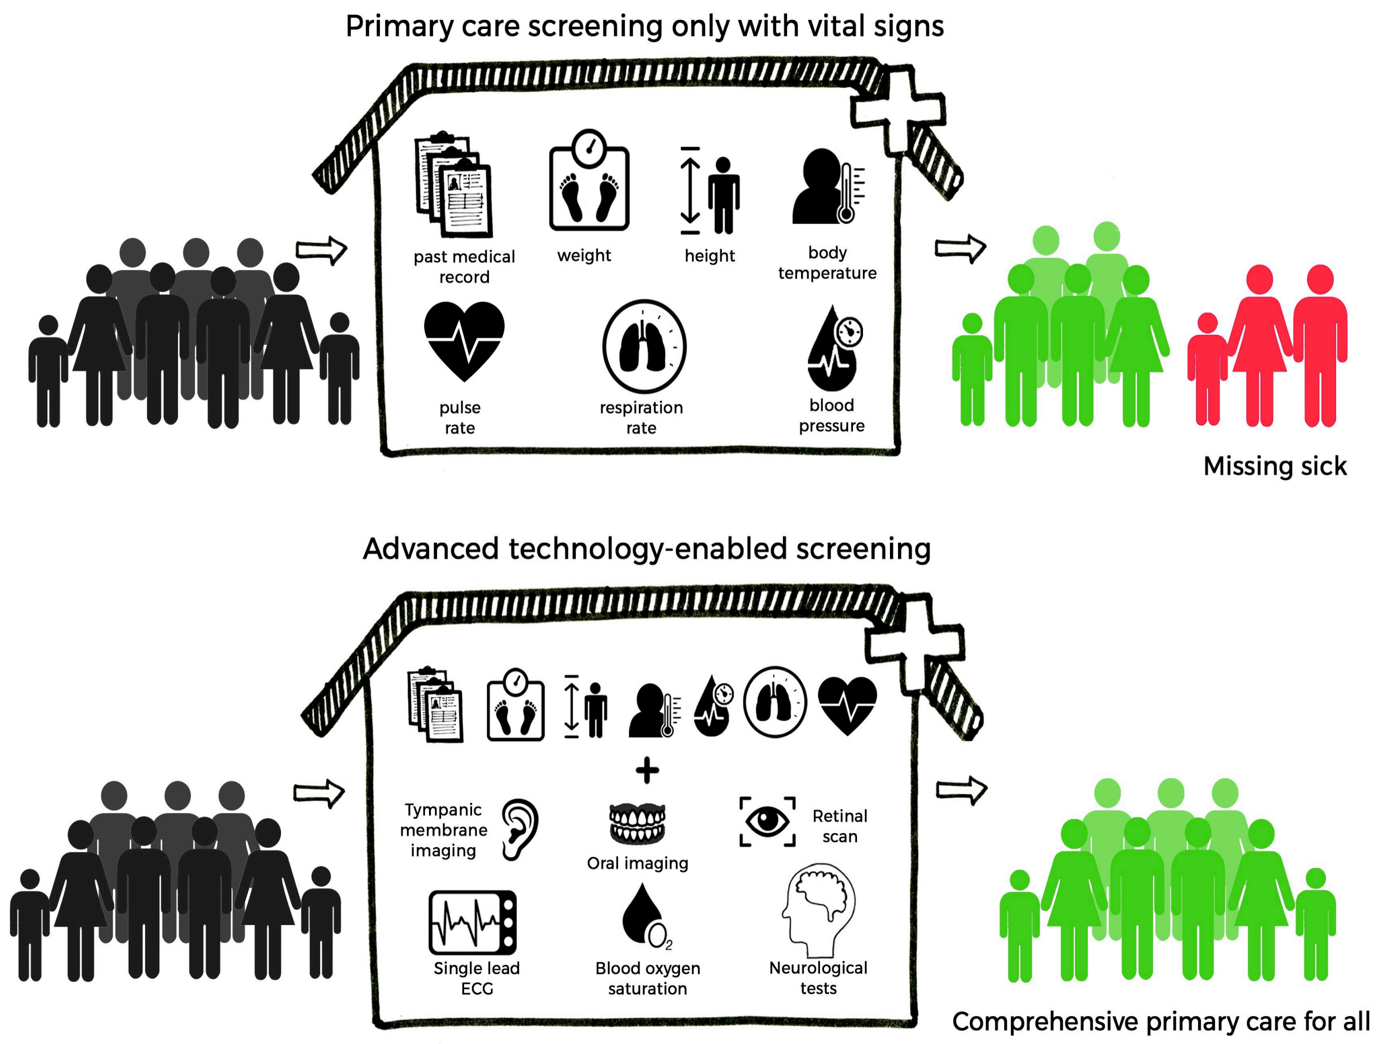

Real world data and evidence

Pratik Shah. Ph.D.
pratiks@media.mit.edu

Oral and Prostate Cancer

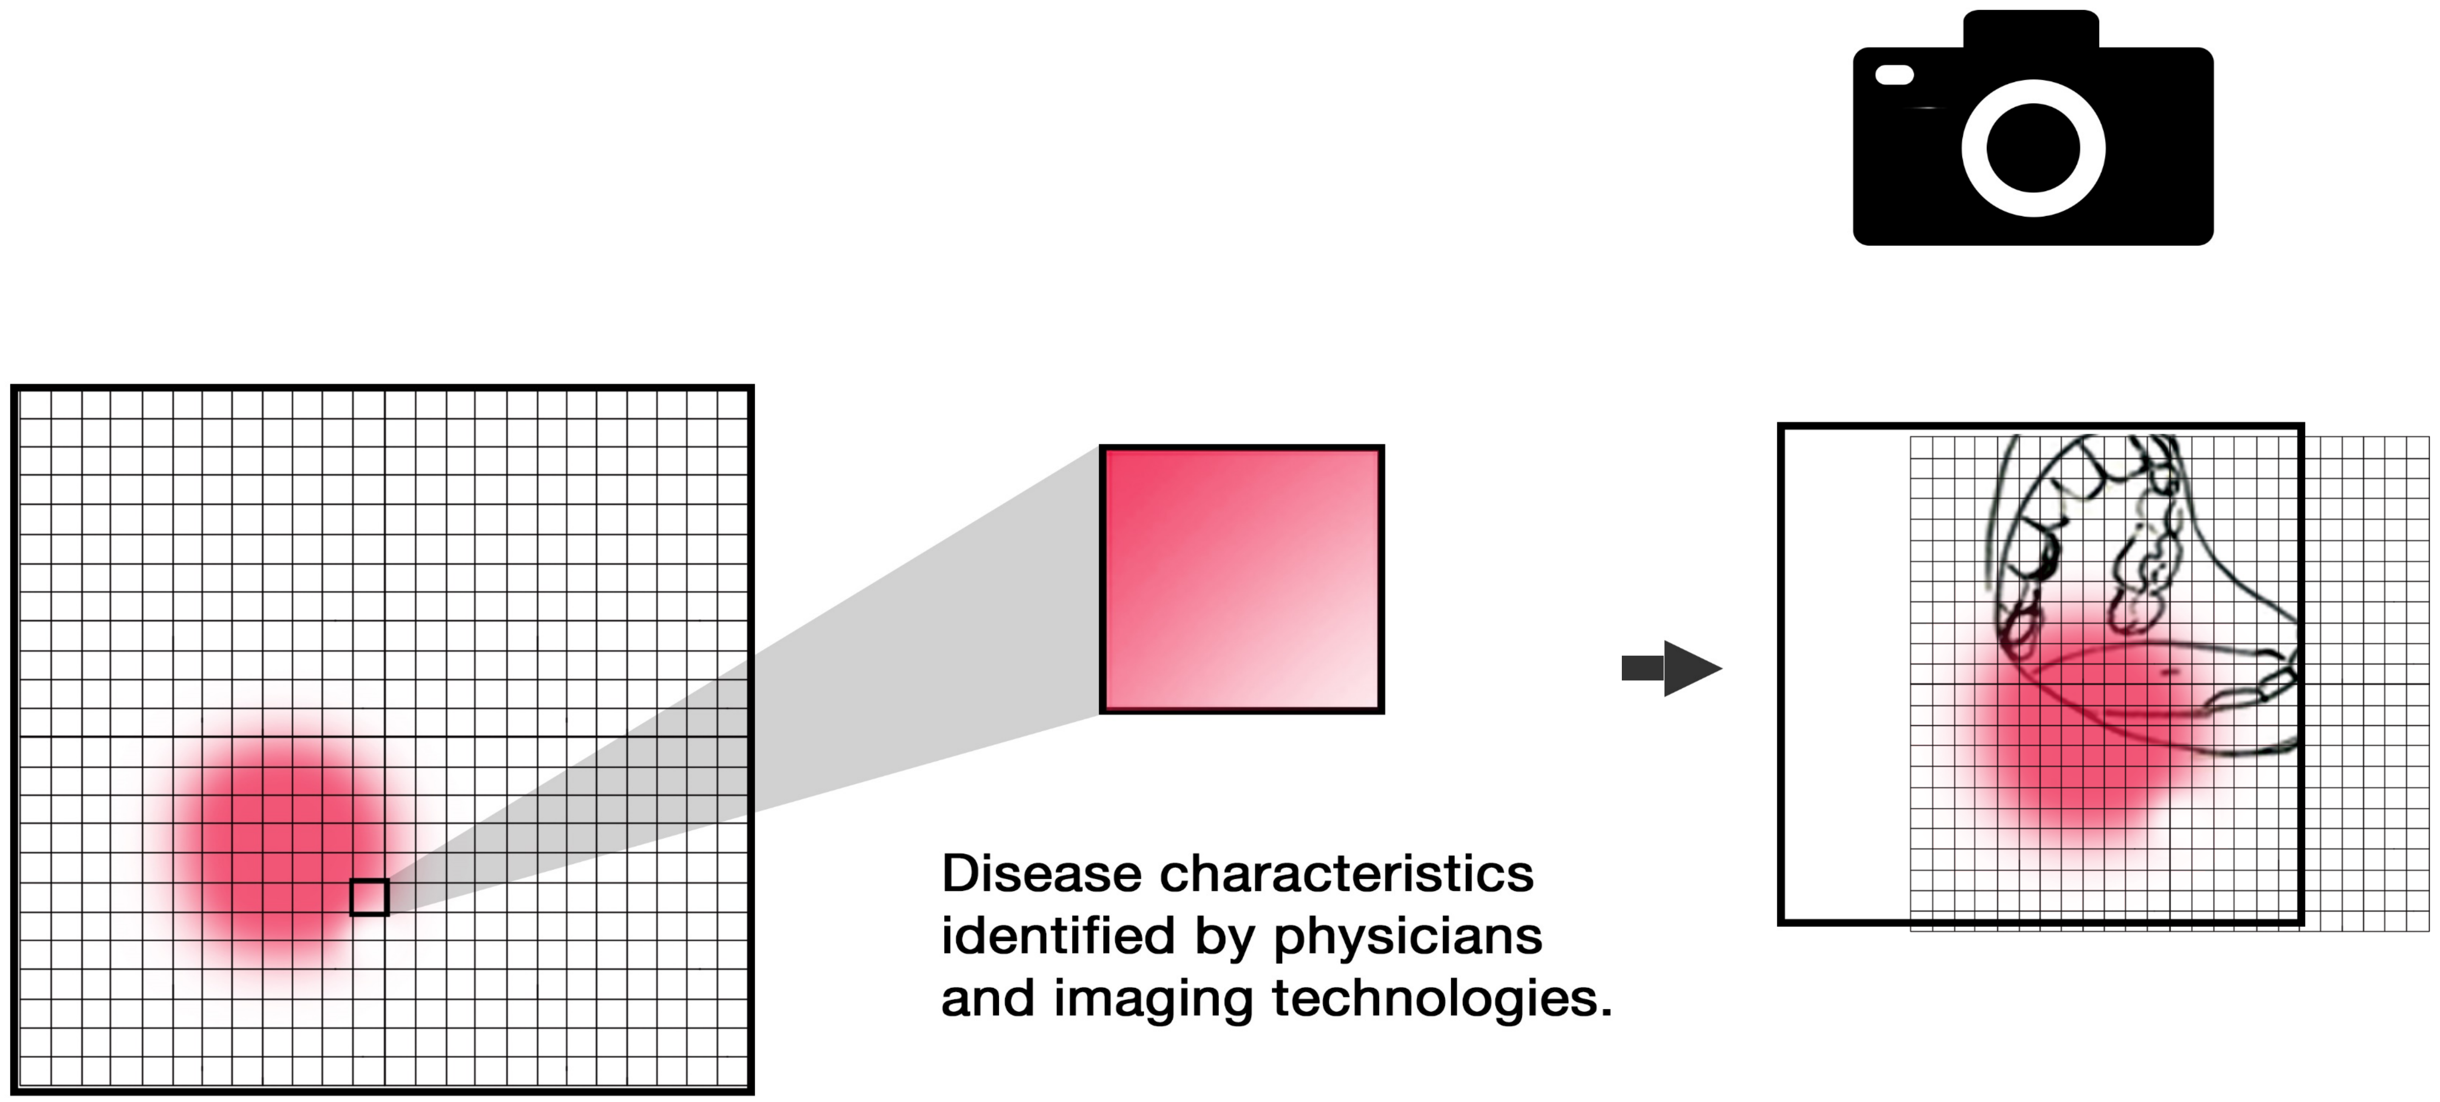

Machine saliency and unorthodox imaging of biomarkers

Pratik Shah. Ph.D.
pratiks@media.mit.edu

Neural Networks with Human and Device Cognition

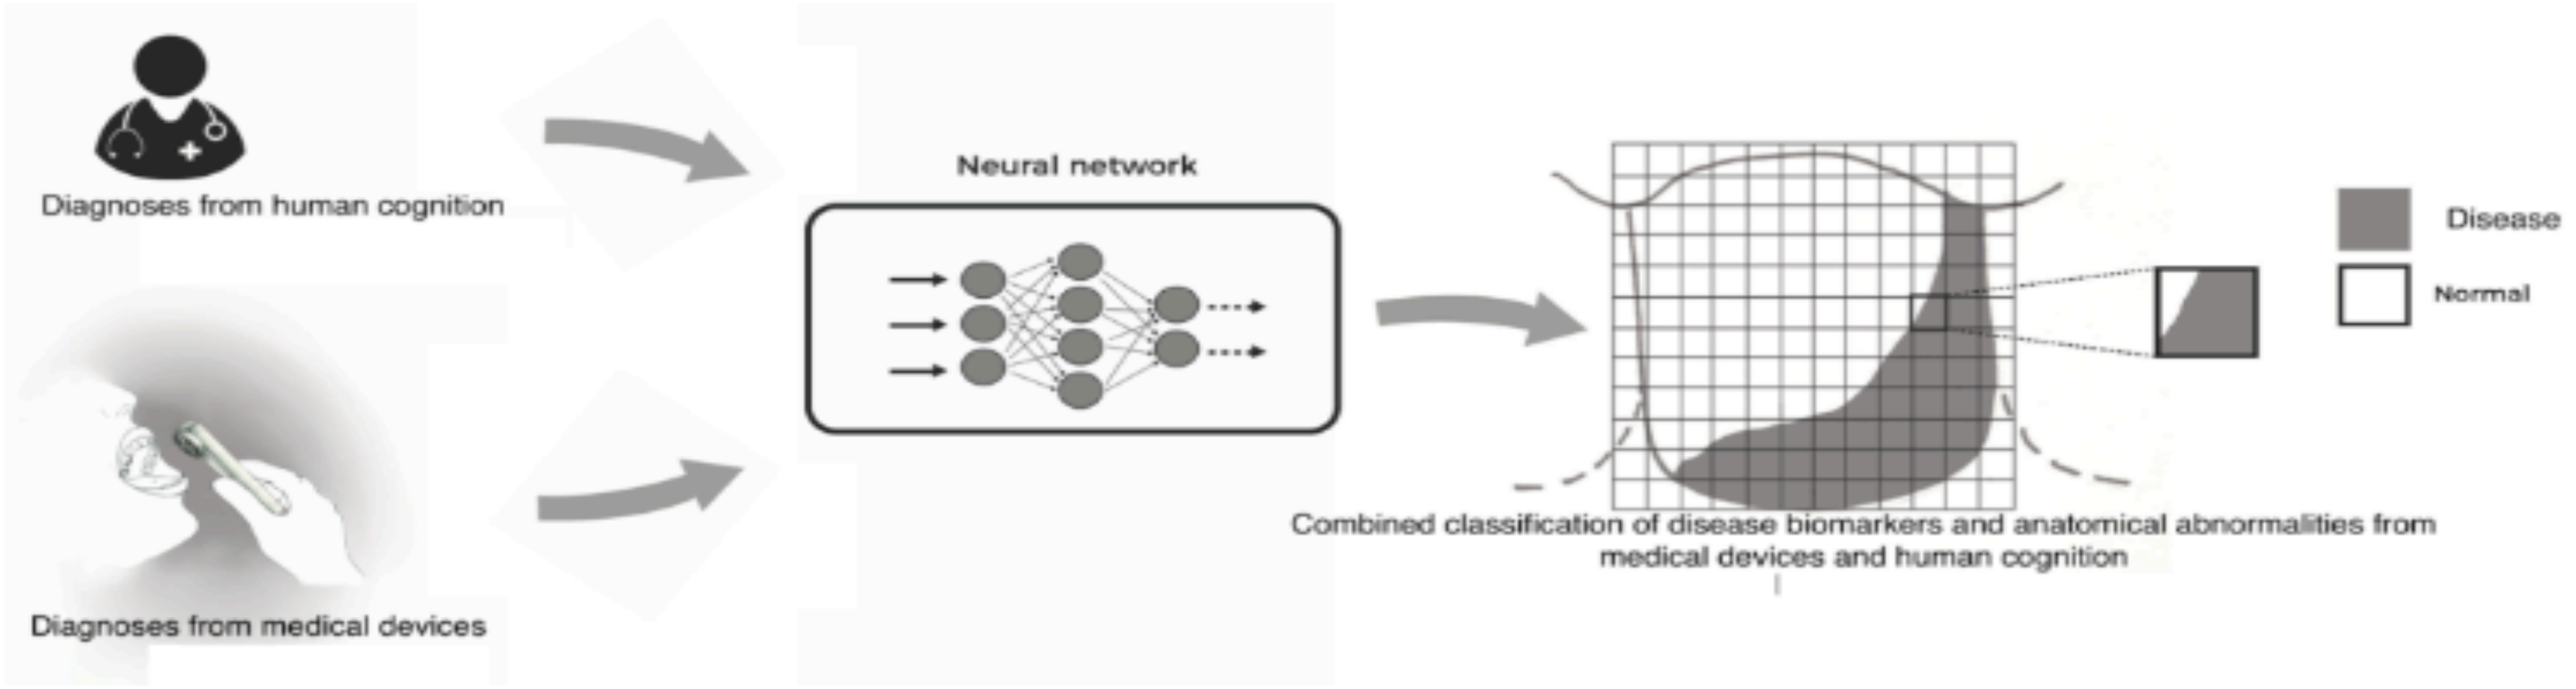

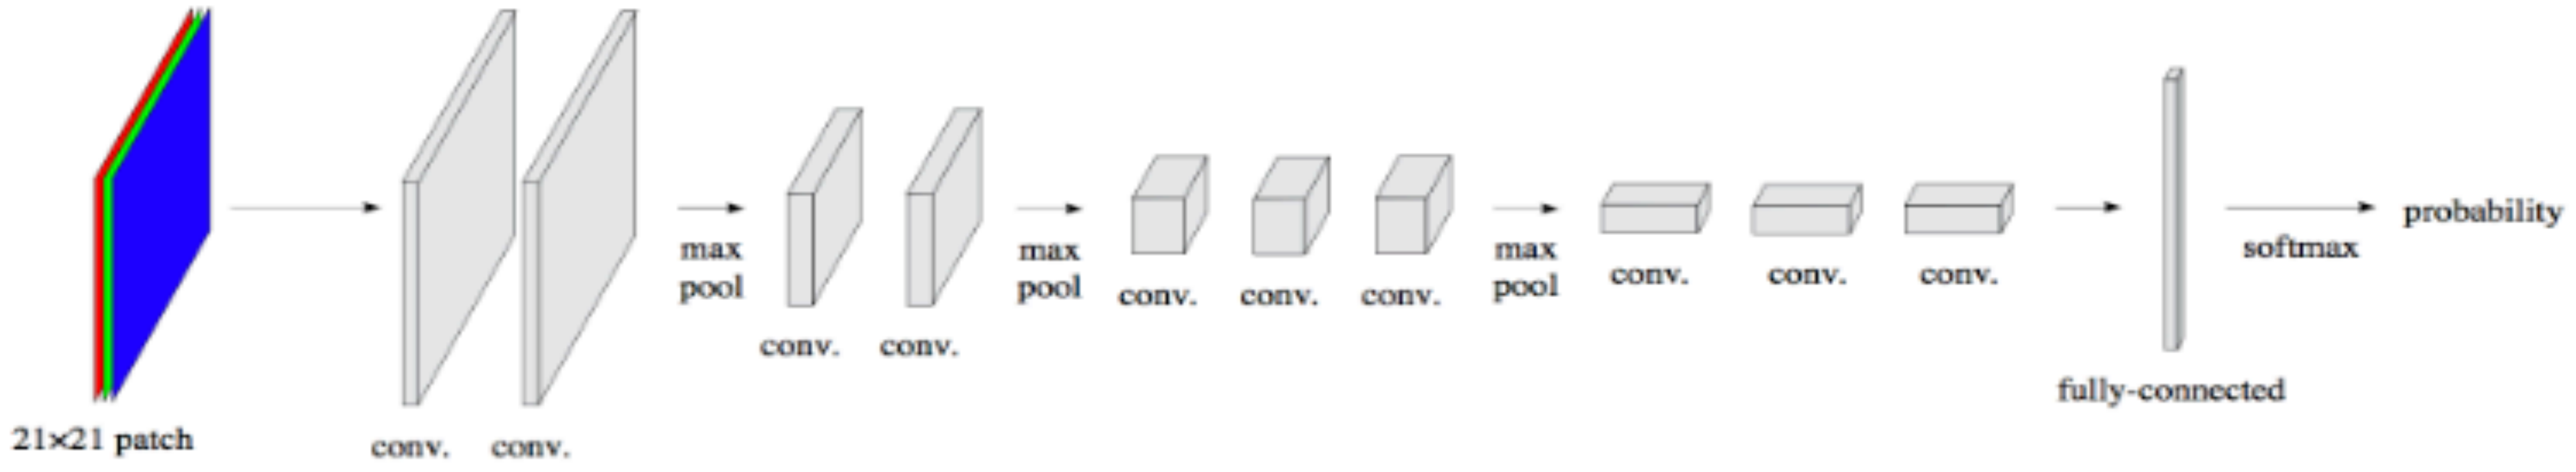

Deep learning architecture.

Convolutional neural networks for combined classification of biomarkers and expert annotations using white light images
G.Yauney, K. Angelino, D. Edlund and **Pratik Shah**
2017 17th International Conference on Bioinformatics and Bioengineering. DOI: 10.1109/BIBE.2017.00-37

Pratik Shah. Ph.D.
pratiks@media.mit.edu

Unorthodox Image Processing and Biomarker Generation

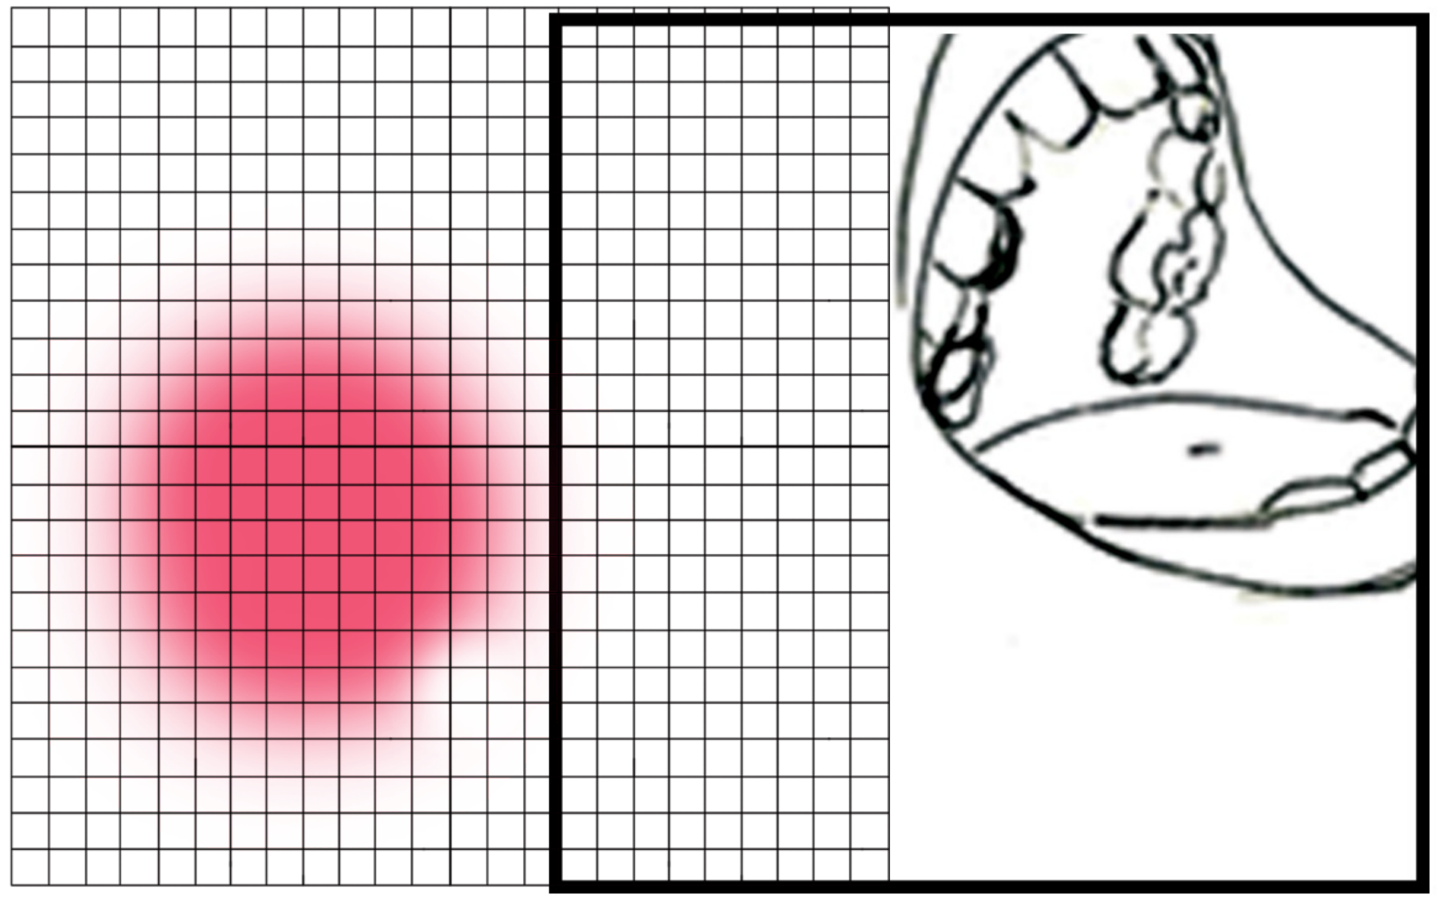

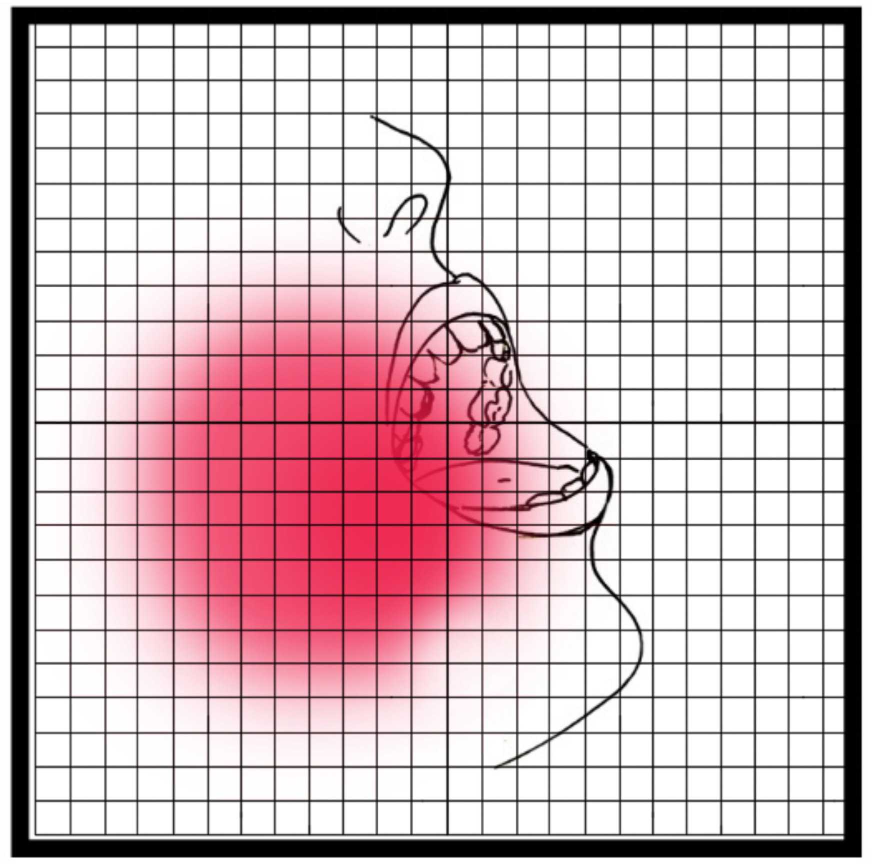

Pratik Shah. Ph.D.
pratiks@media.mit.edu

Computational Histological Staining and Destaining of Prostate Core Biopsy RGB Images with Generative Adversarial Neural Networks

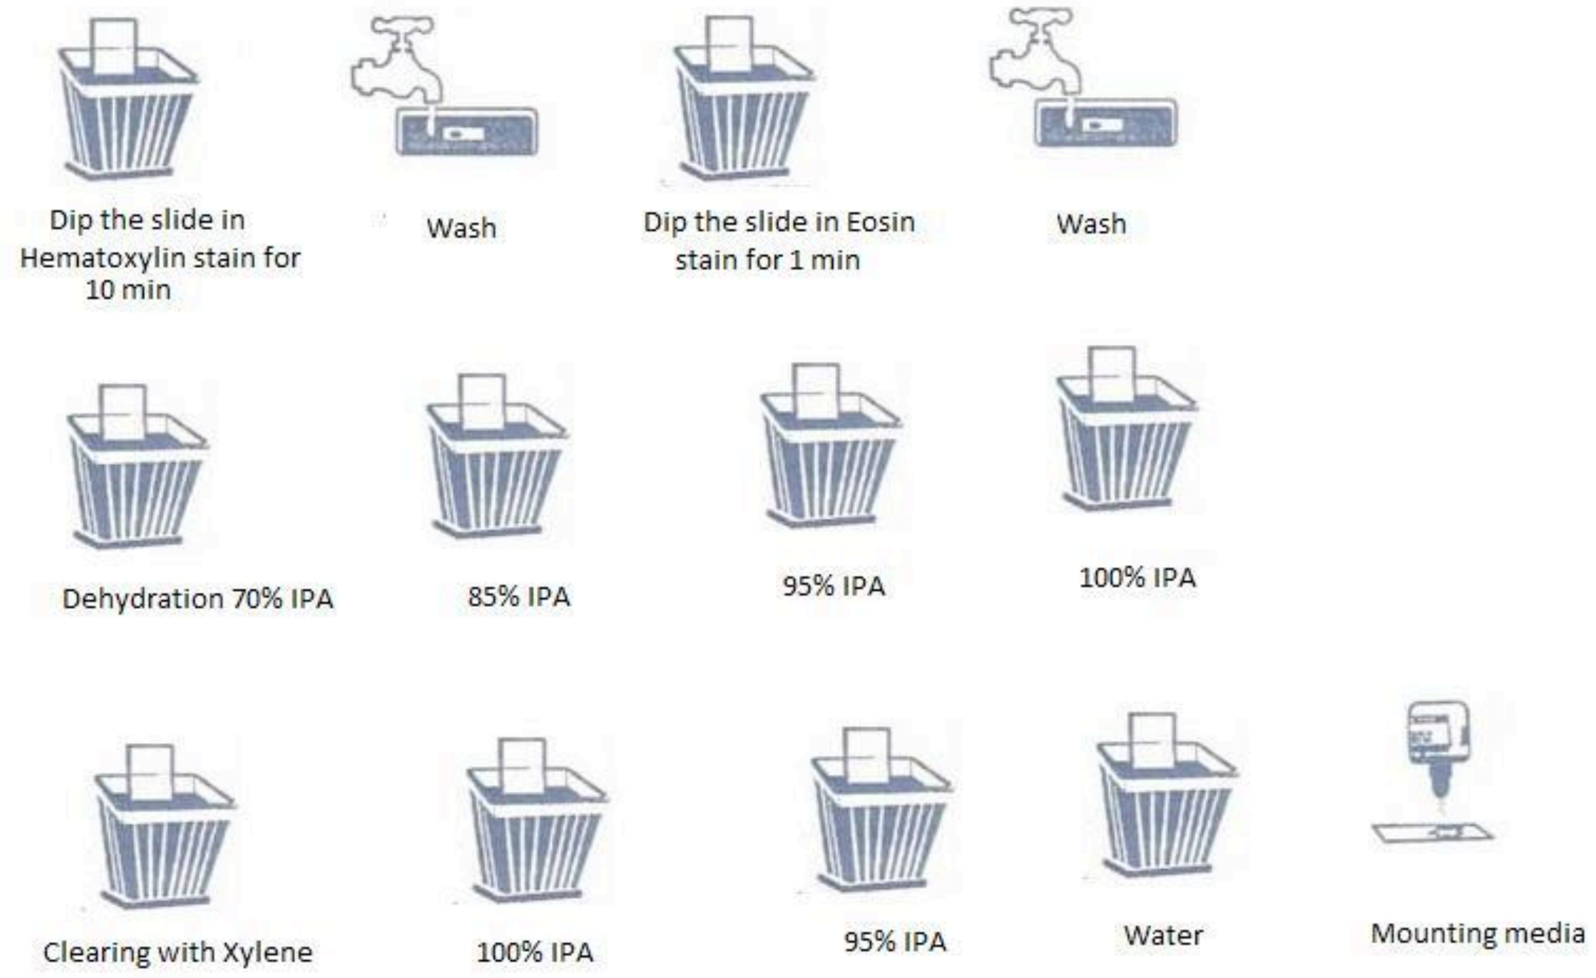

Collaboration with Brigham and Womens Hospital

mit media lab
Pratik Shah. Ph.D.
pratiks@media.mit.edu

Generative Adversarial Neural Networks for New Medical Knowledge

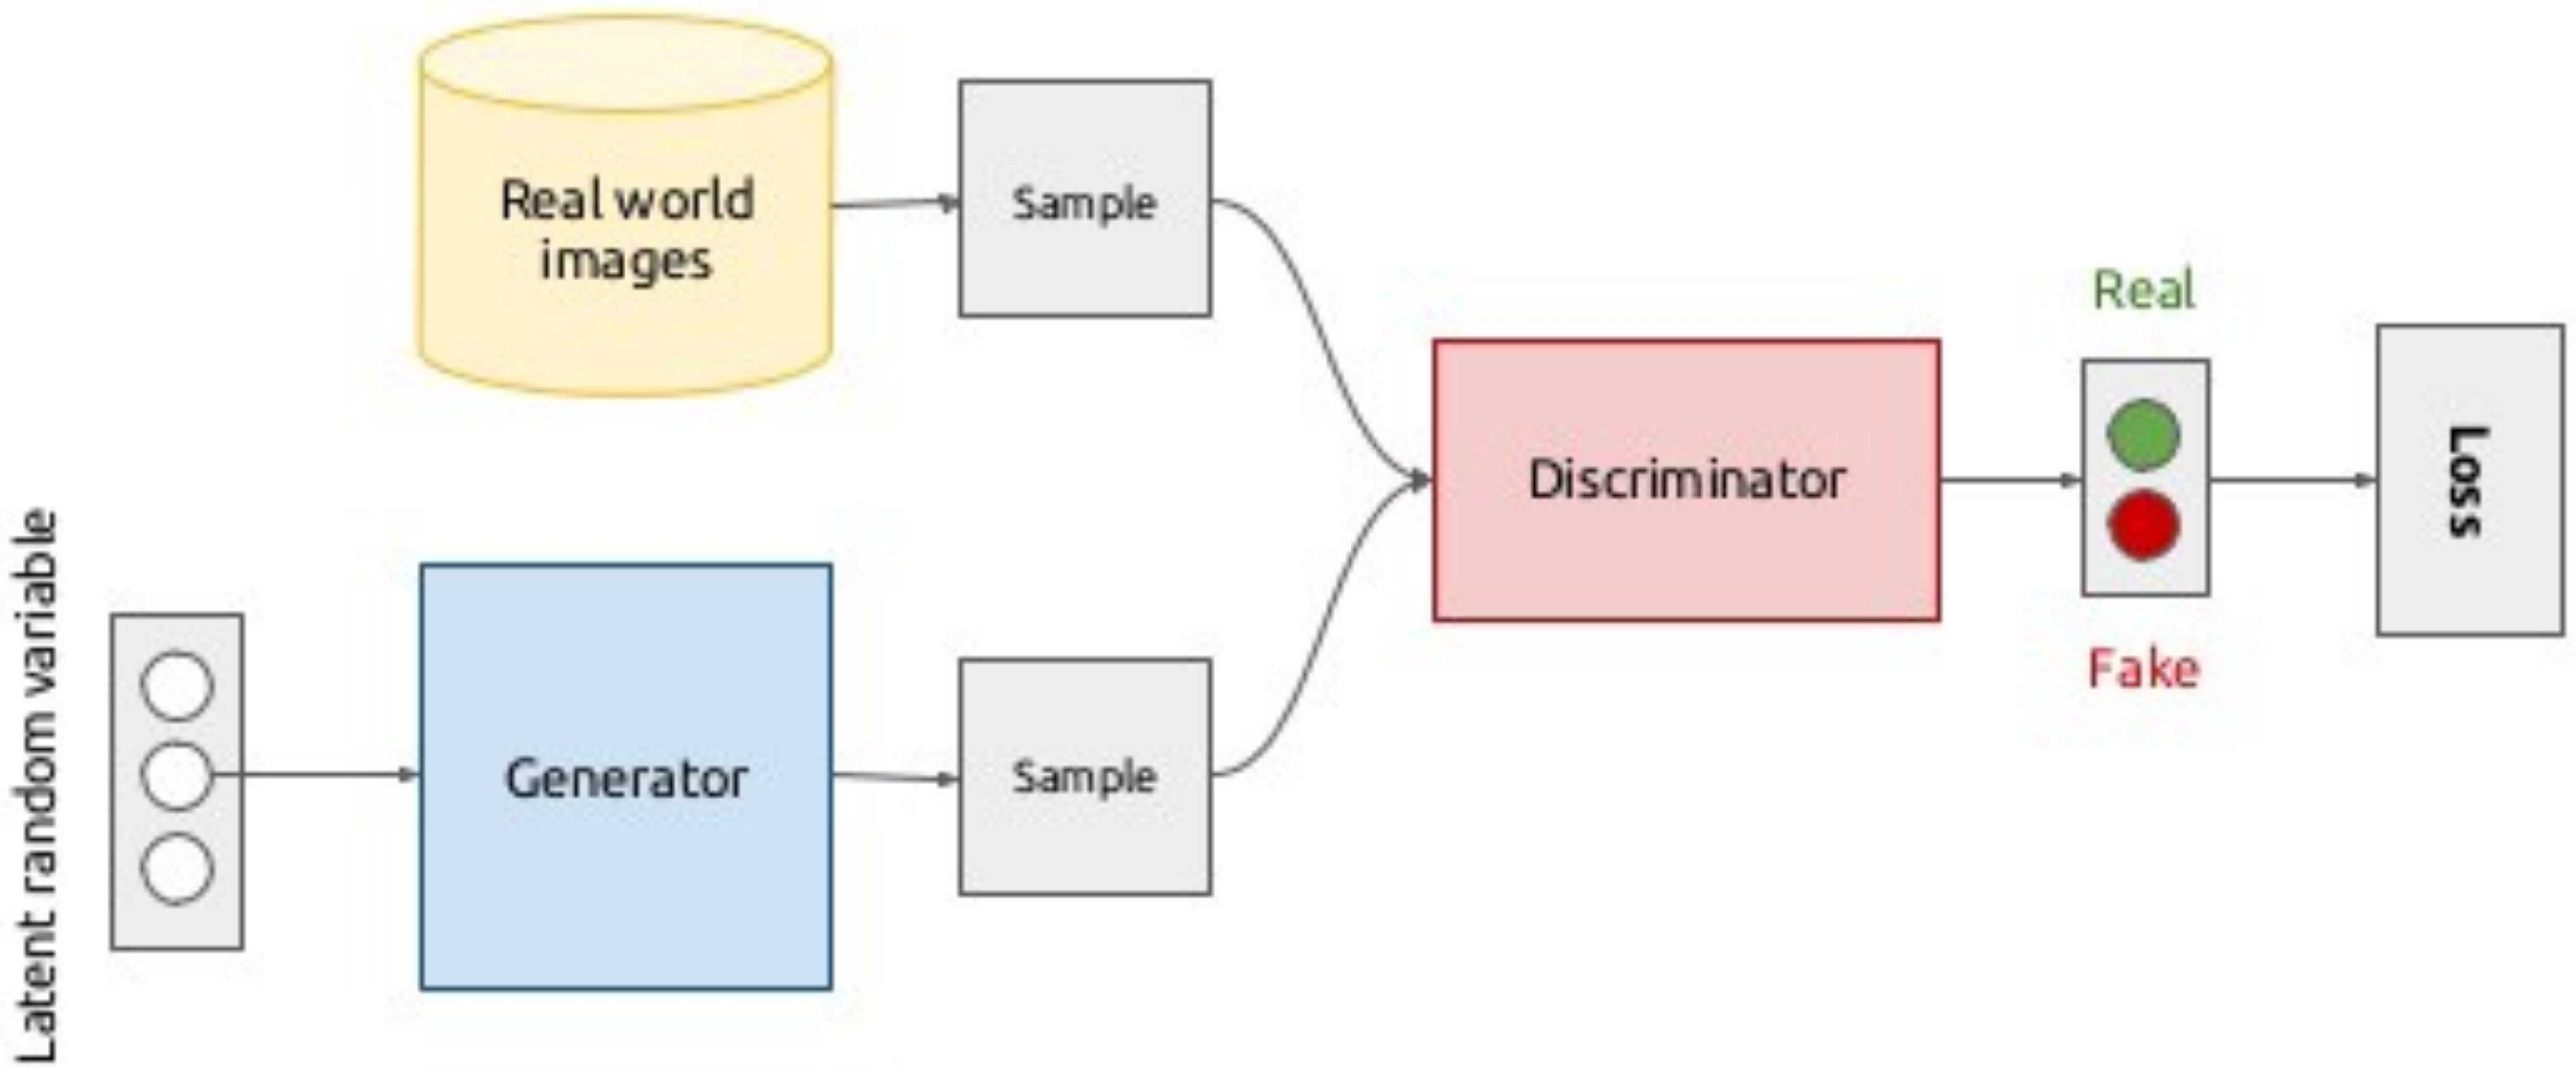

synthetic data/images

realistic looking high resolution images

$$\min_G \max_D V(D,G) = \mathbb{E}_{\boldsymbol{x}\sim p_{\text{data}}(\boldsymbol{x})}[\log D(\boldsymbol{x})] + \mathbb{E}_{\boldsymbol{z}\sim p_{\boldsymbol{z}}(\boldsymbol{z})}[\log(1 - D(G(\boldsymbol{z})))].$$

H & E Staining and Destaining of Prostate Core Biospy Images

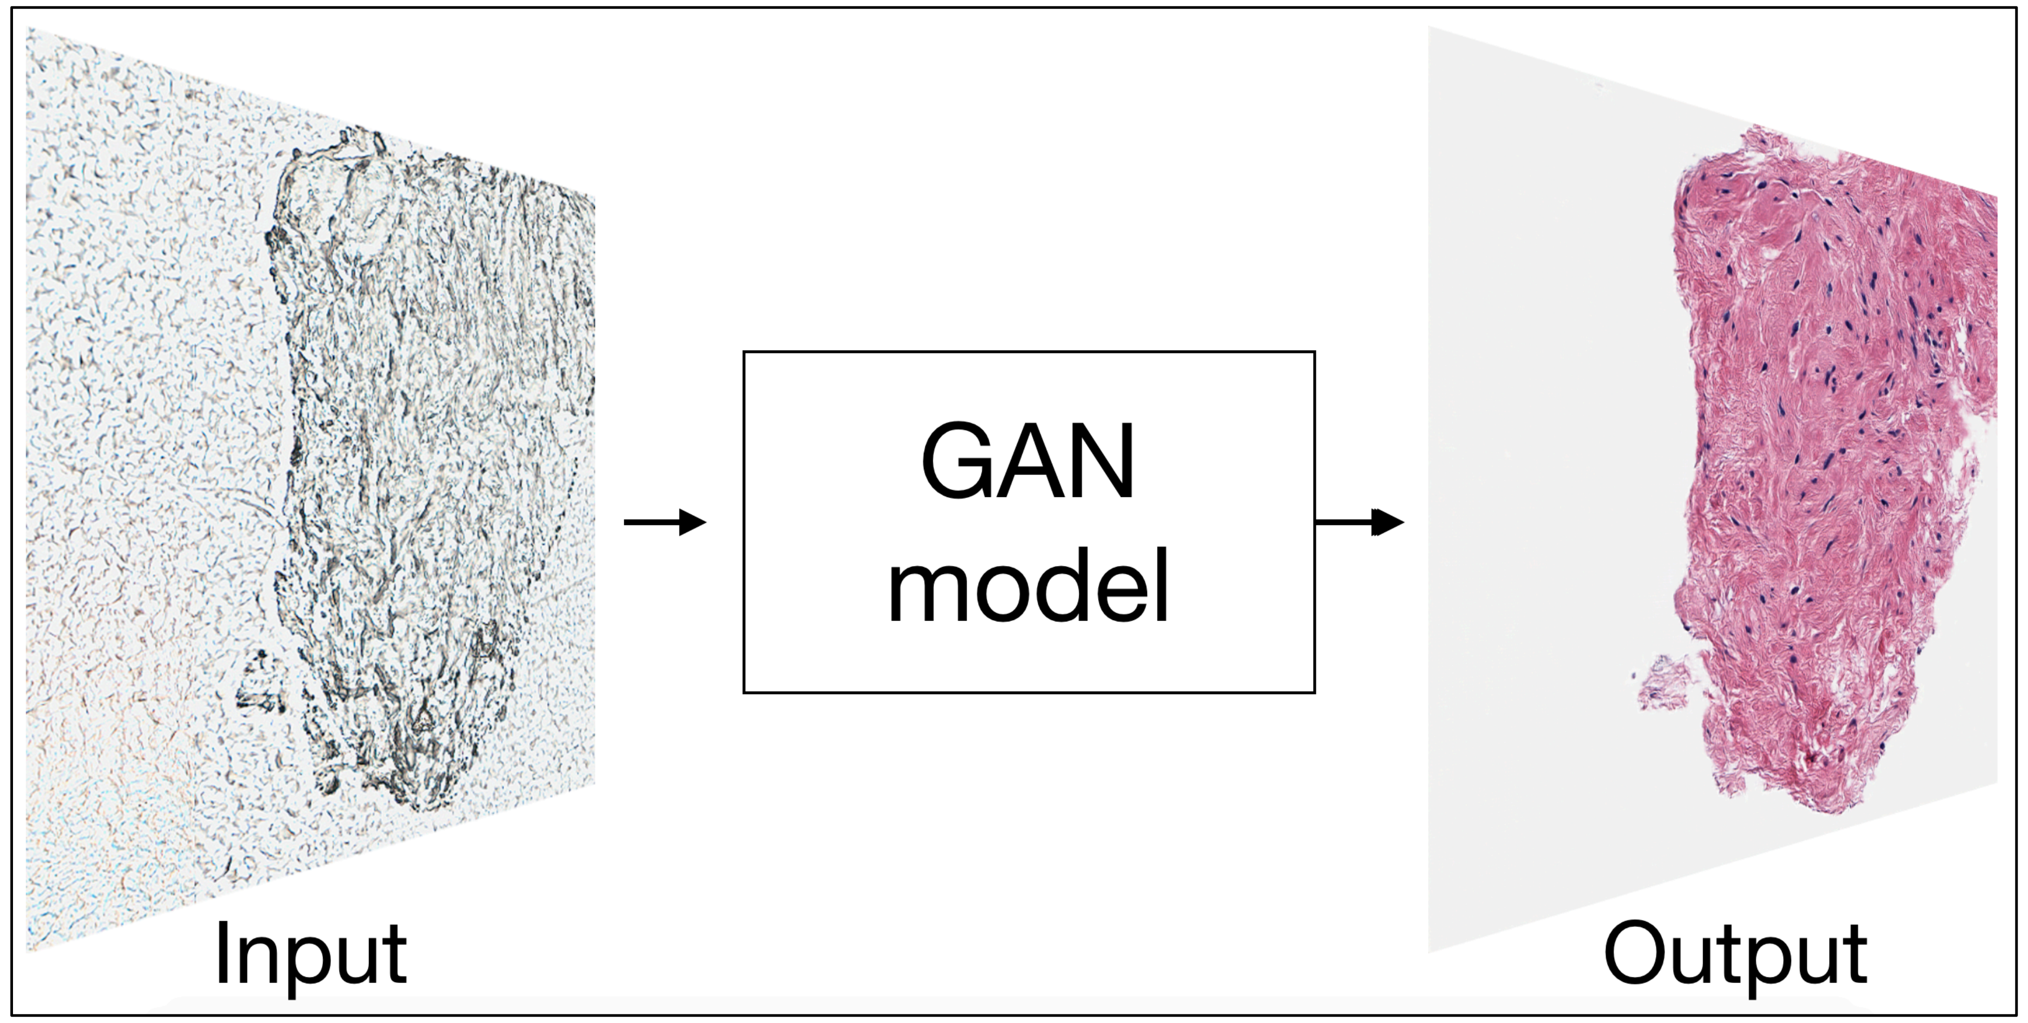

Computational staining

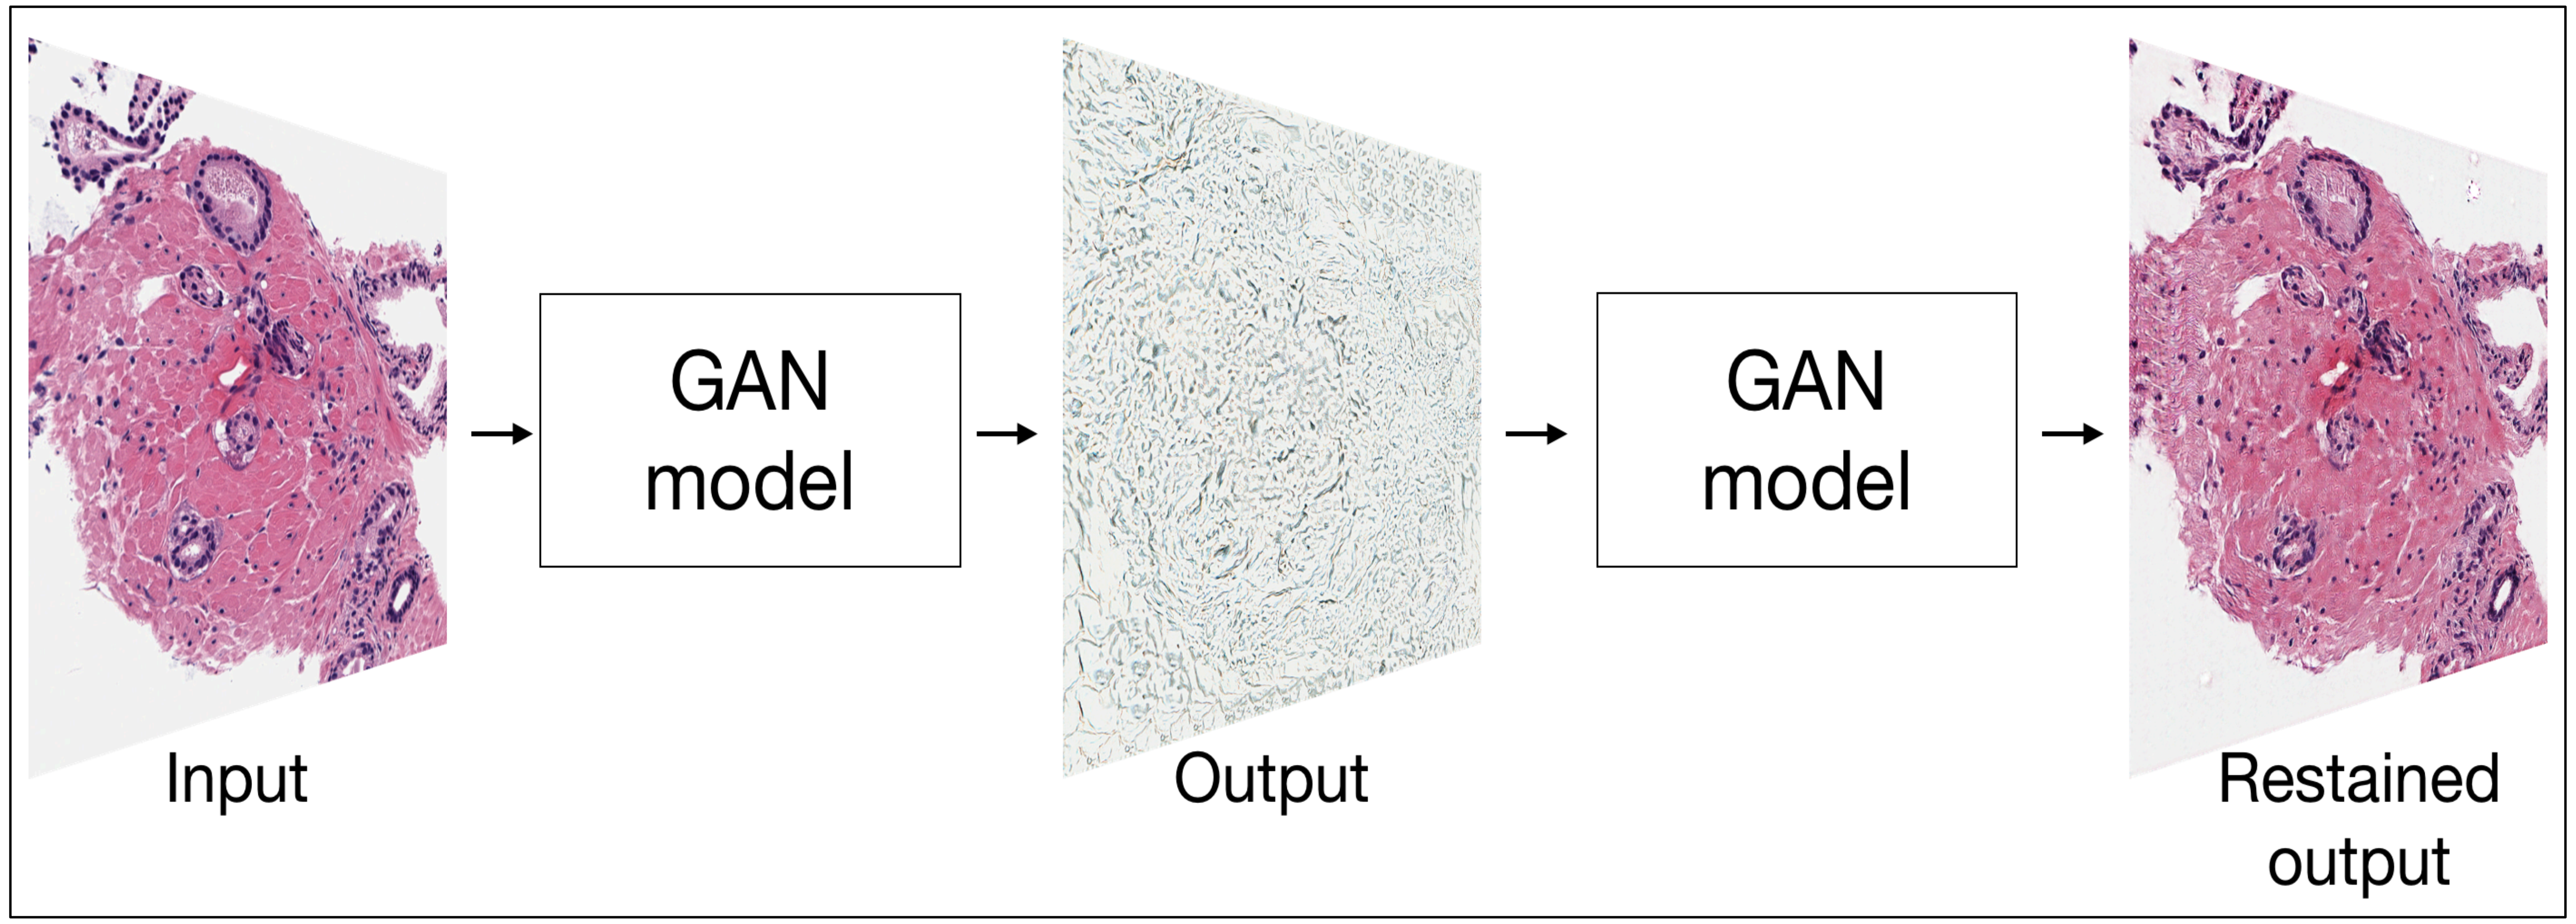

Computational destaining

Collaboration with Brigham and Womens Hospital

Computational histological staining and destaining of prostate core biopsy RGB images with generative adversarial neural networks
A. Rana, G. Yauney, A. Lowe and Pratik Shah
2018 17th International Conference on Machine Learning and application.

mit media lab
Pratik Shah. Ph.D.
pratiks@media.mit.edu

Computational Staining and Destaining of Prostate Cancer Biopsy

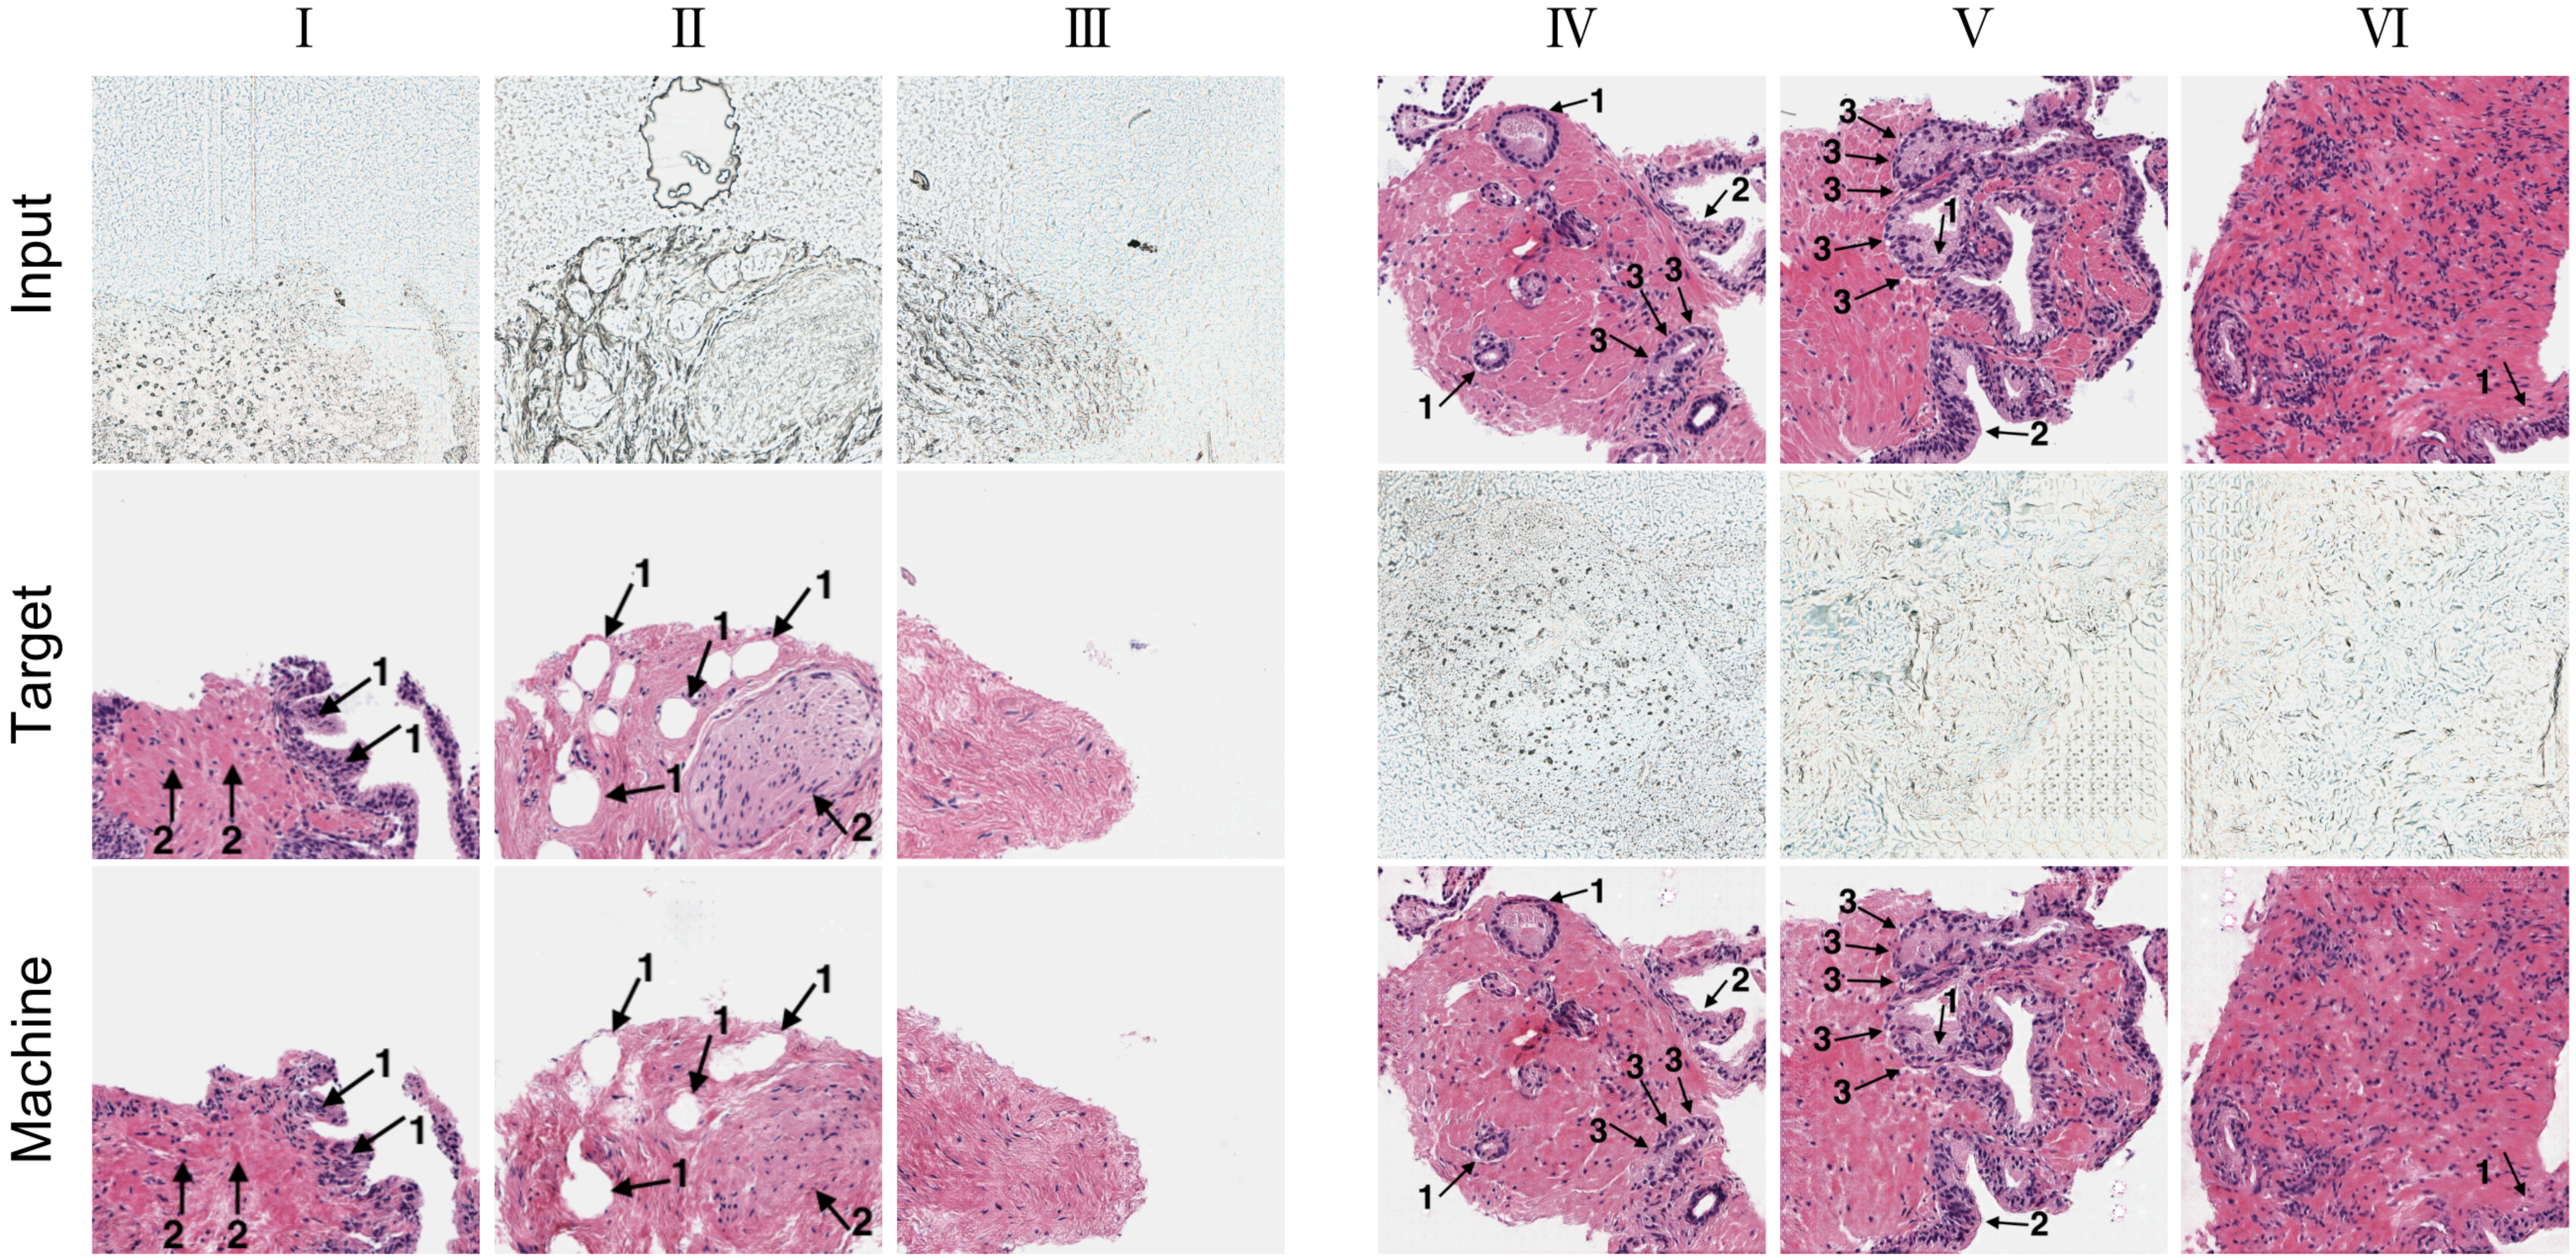

Computational staining

Computational destaining

Collaboration with Brigham and Womens Hospital

Computational histological staining and destaining of prostate core biopsy RGB images with generative adversarial neural networks
A. Rana, G. Yauney, A. Lowe and Pratik Shah
2018 17th International Conference on Machine Learning and application.

Pratik Shah. Ph.D.
pratiks@media.mit.edu

Glioblastomas, Prostate, Colorectal and Small Cell Lung Cancer

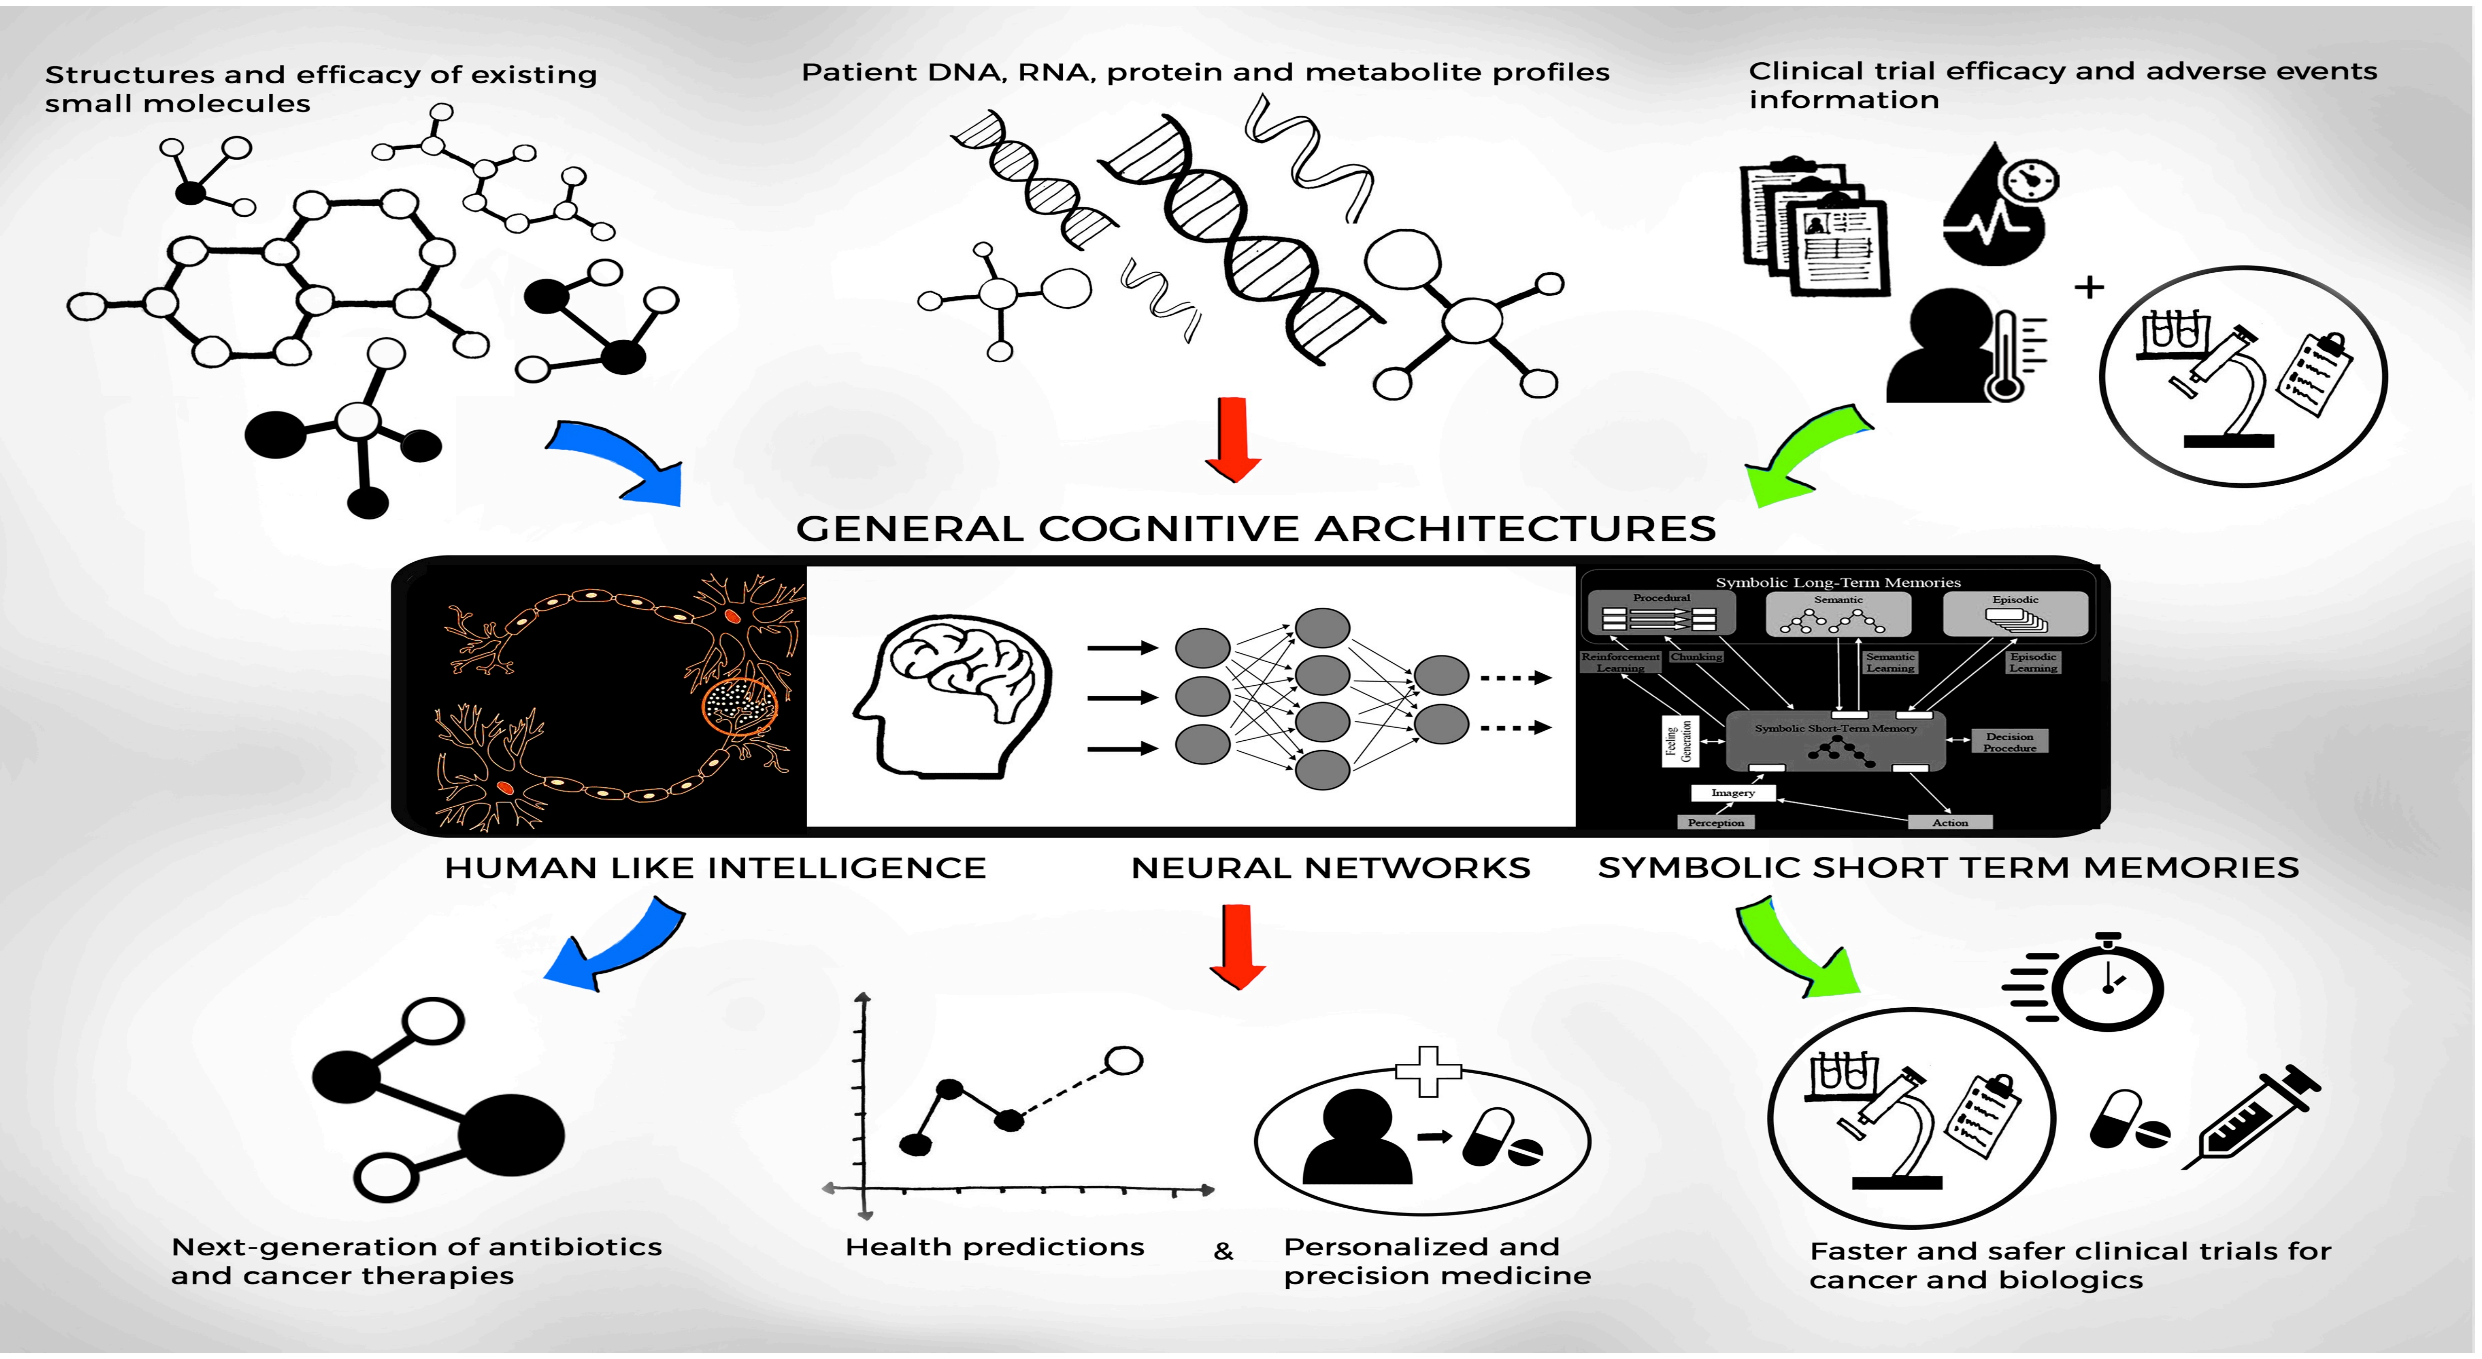

Novel clinical trials

Pratik Shah. Ph.D.
pratiks@media.mit.edu

Glioblastoma Treatment and Toxicity with Chemo-and-Radiotherapy

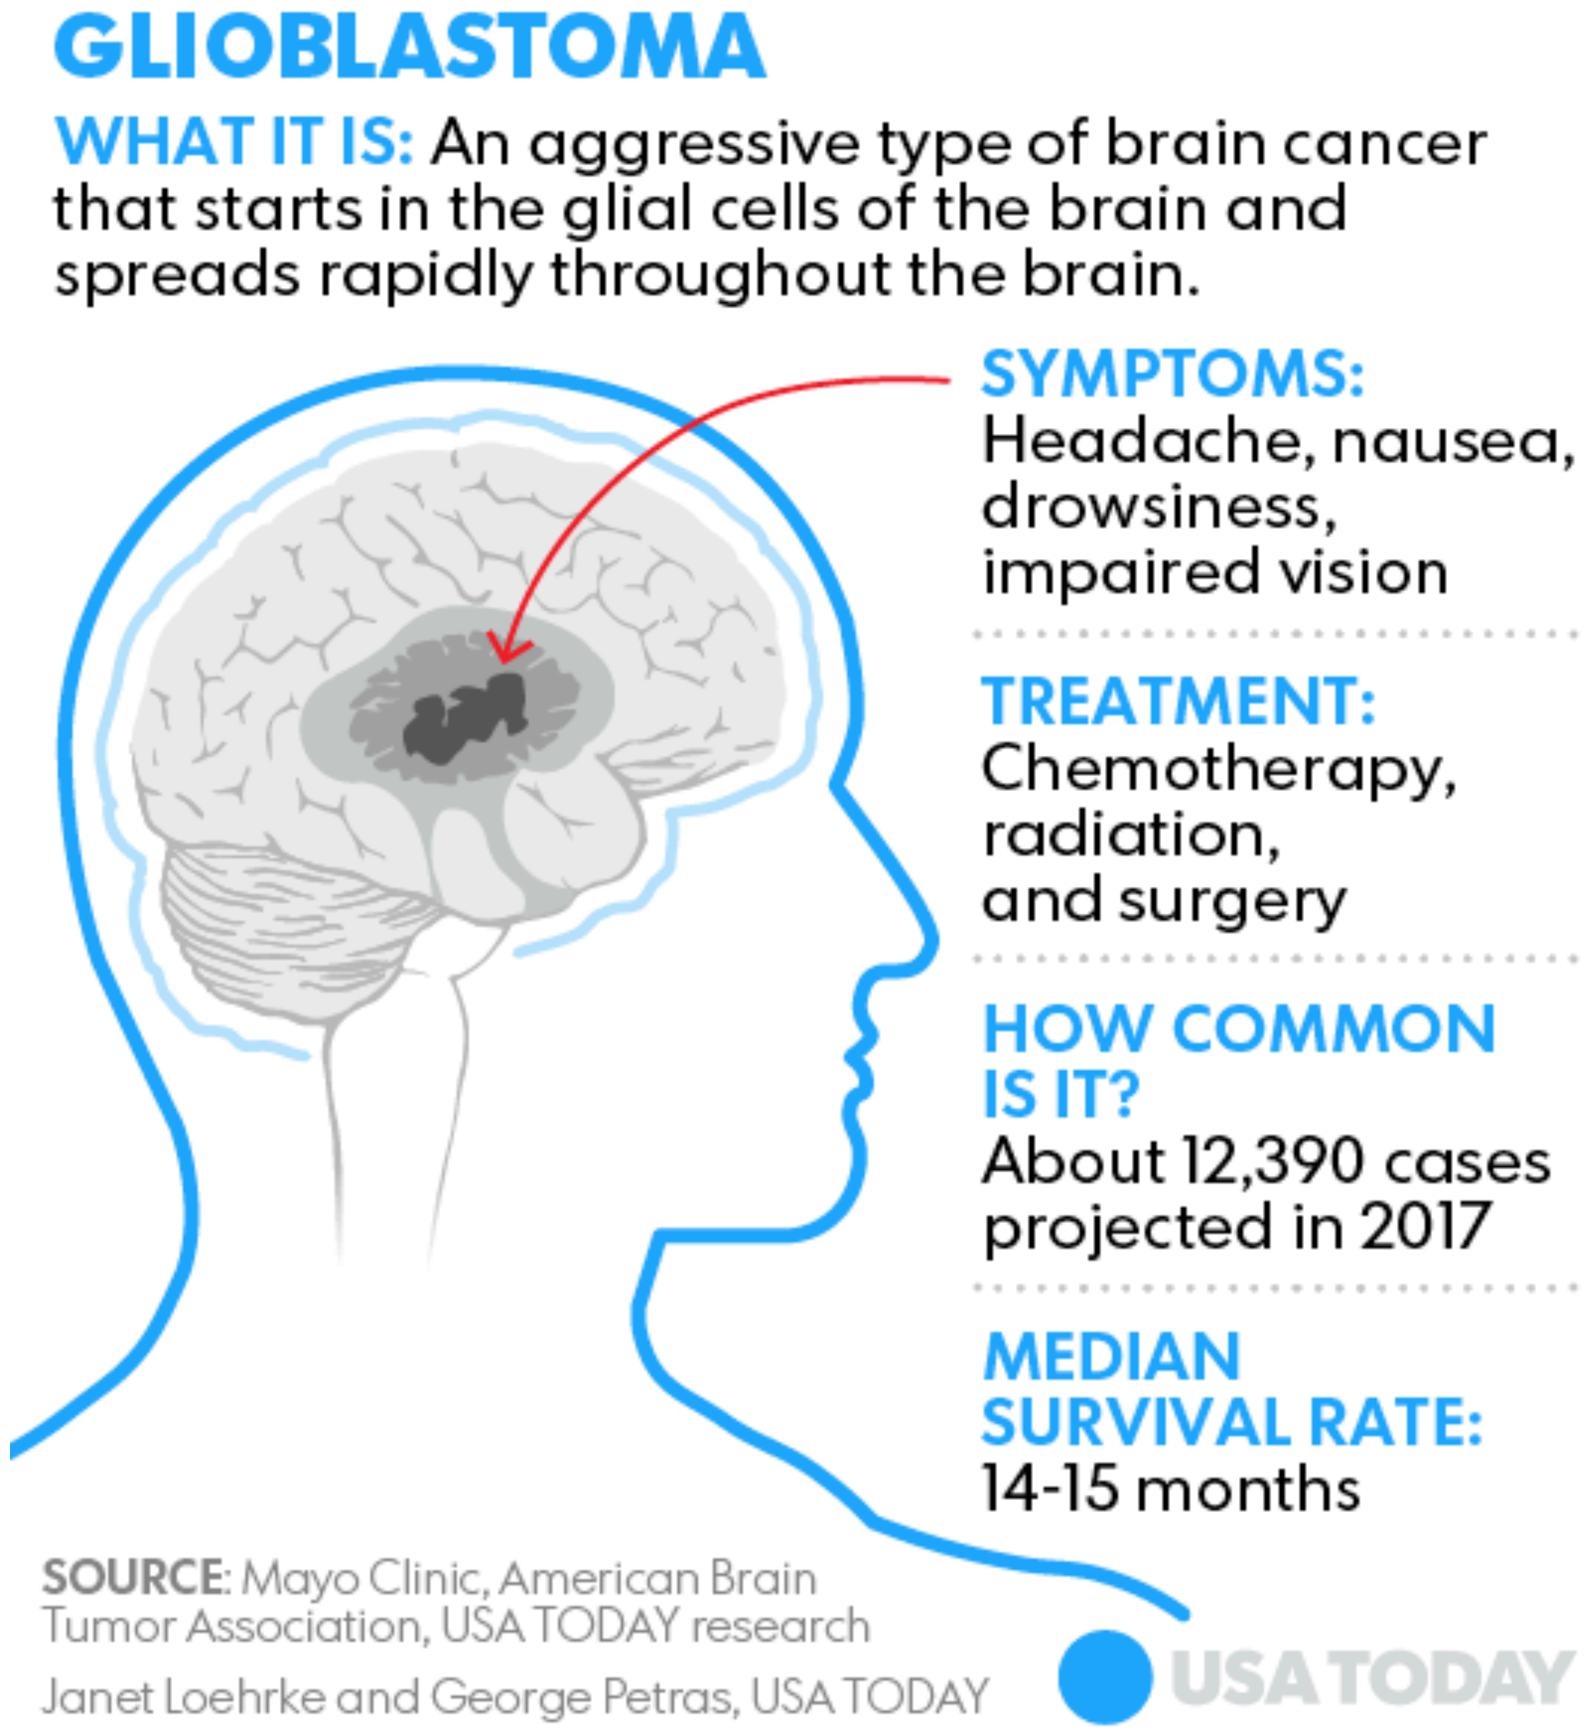

Patient tumor size
Dosing bounds
Final observation
Dose administered

Human expert Temozolomide dosing regimen

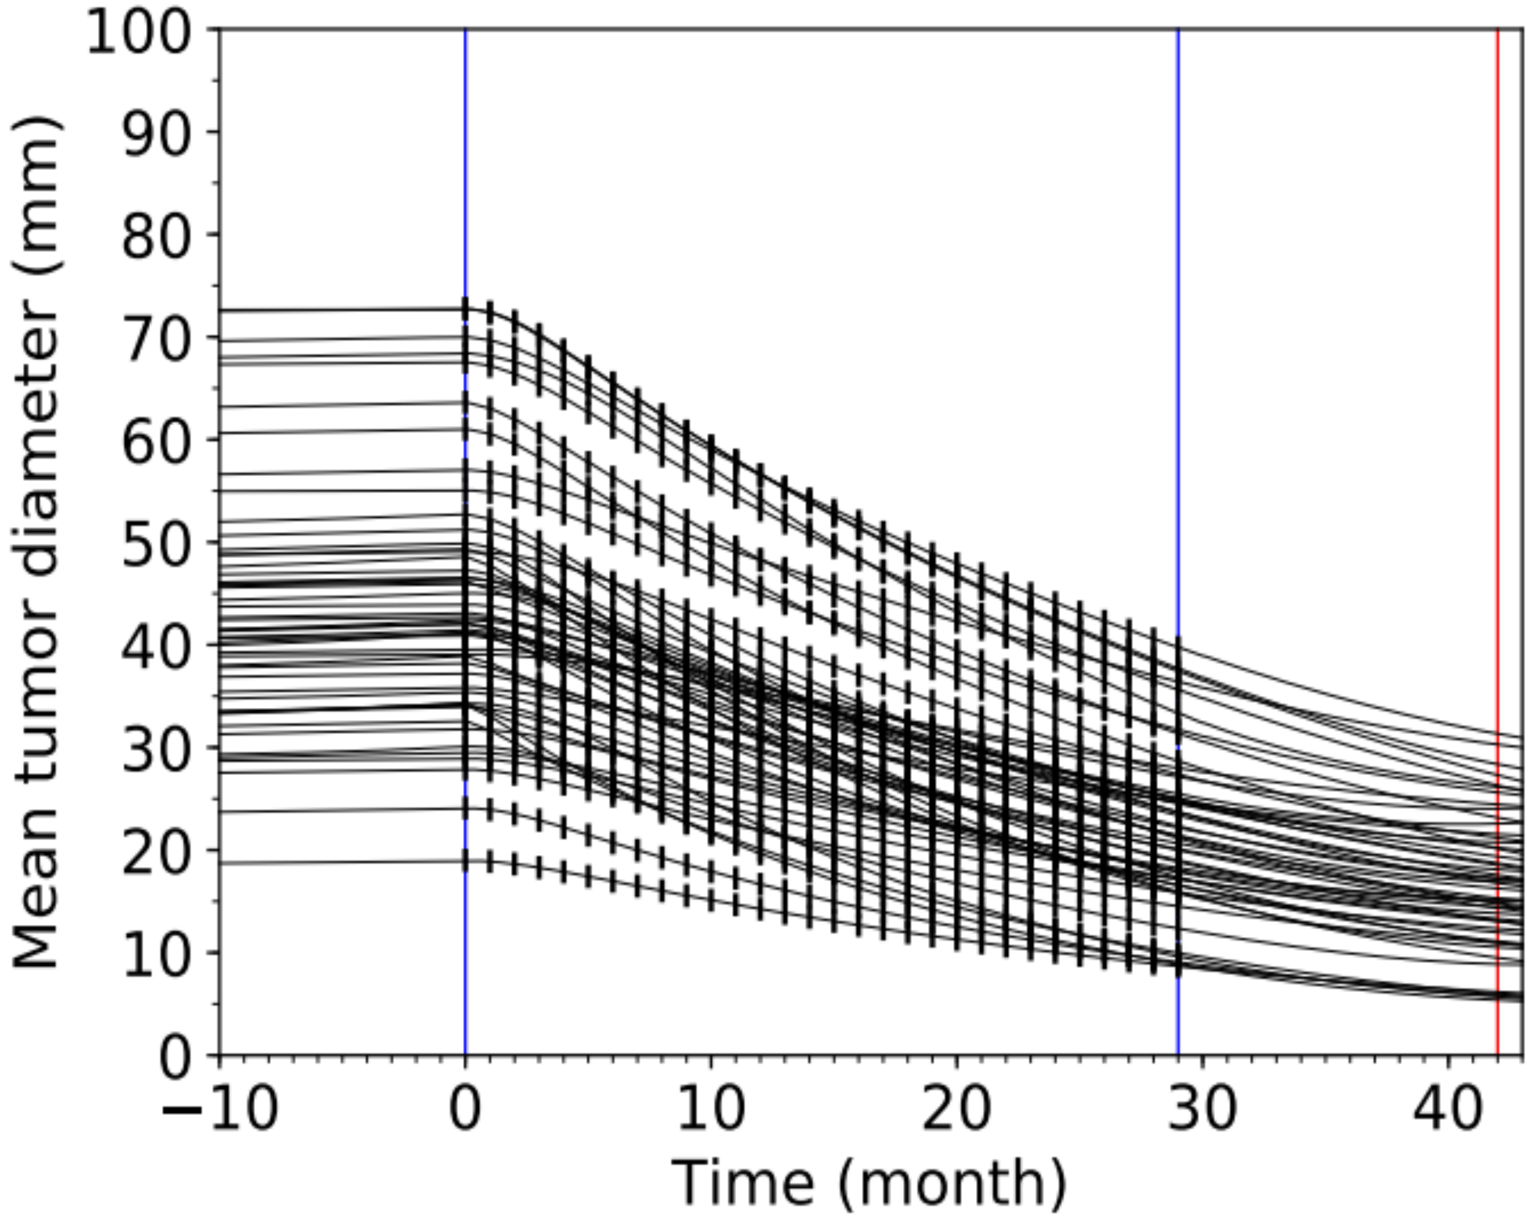

Human expert PCV dosing regimen

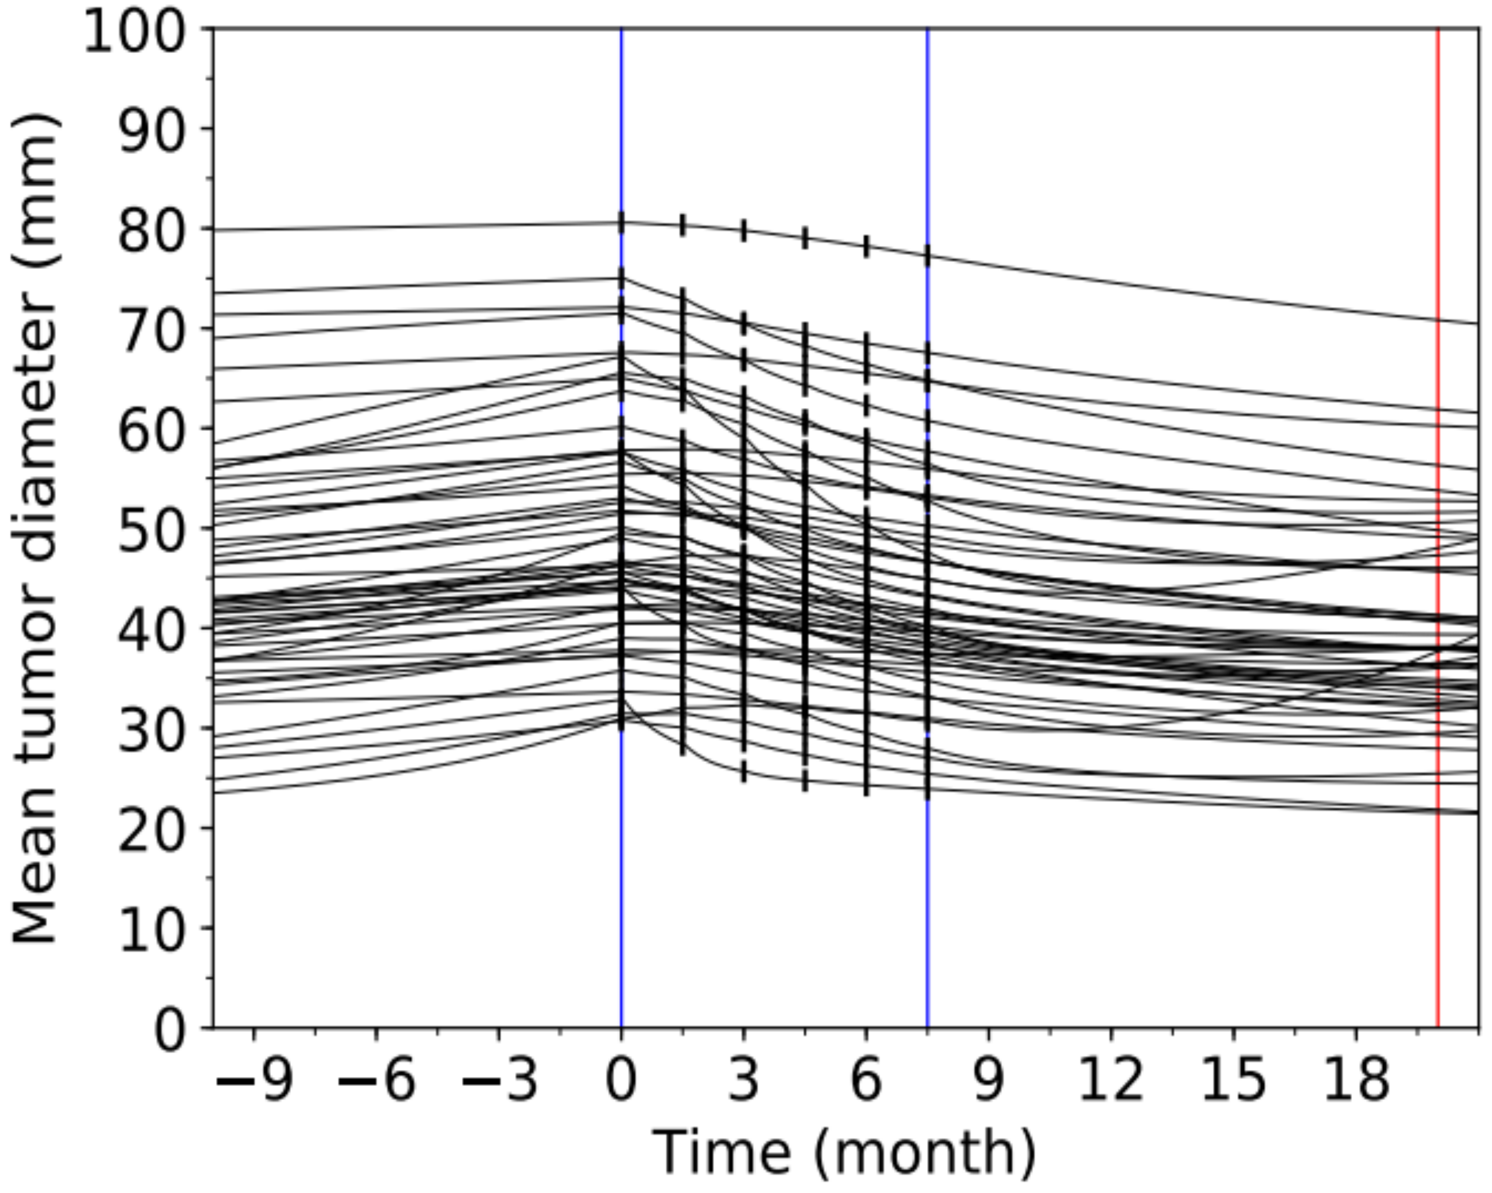

Pratik Shah. Ph.D.
pratiks@media.mit.edu

Reinforcement Learning

Formal definition:

- States
- Actions
- Transition
- Reward
- Discount factor

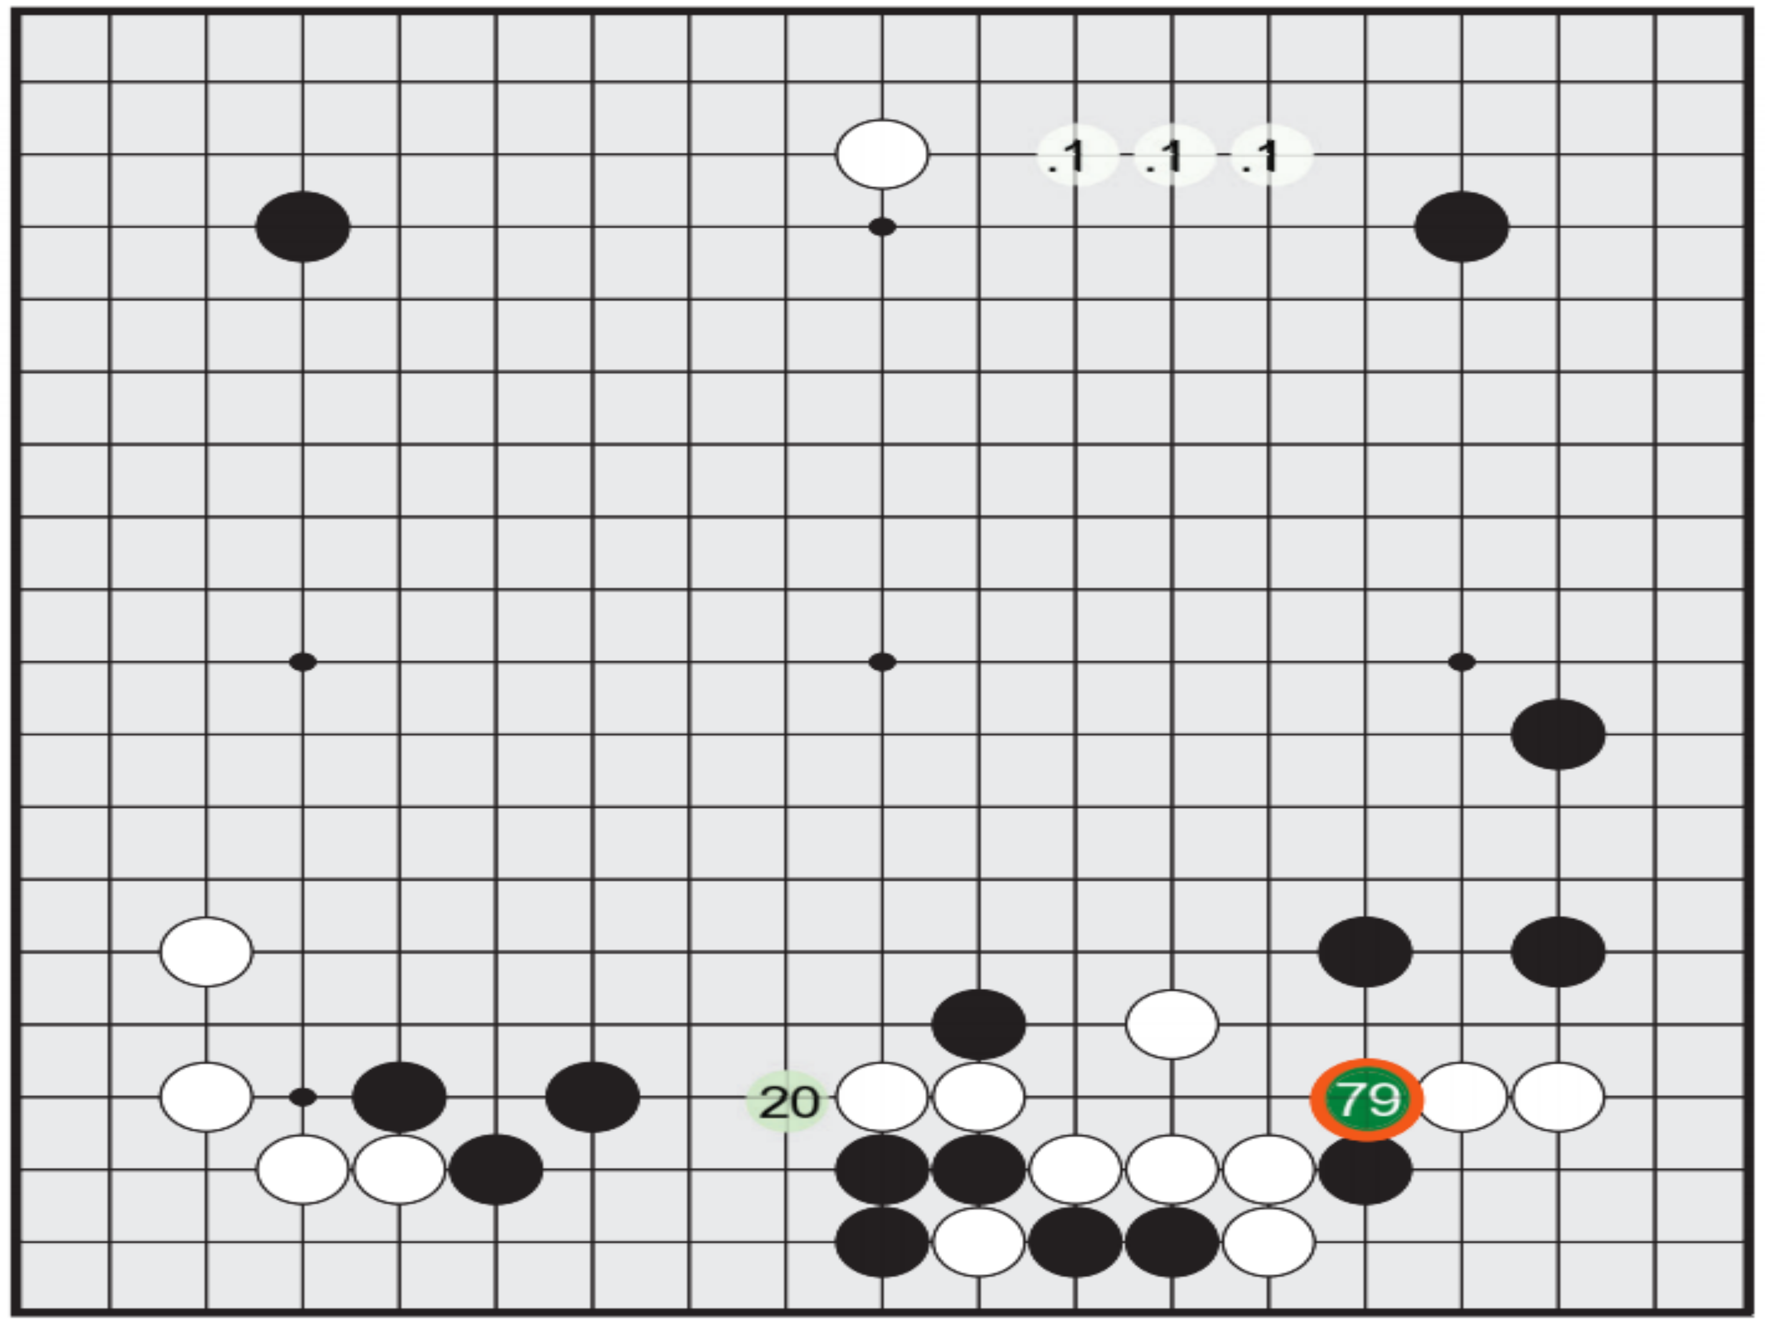

Silver, David, Aja Huang, Chris J. Maddison, Arthur Guez, Laurent Sifre, George Van Den Driessche, Julian Schrittwieser et al. "Mastering the game of Go with deep neural networks and tree search." *Nature* 529, no. 7587 (2016): 484-489.

Sutton & Barto, *Reinforcement Learning: an Introduction.* Cambridge: MIT Press: 1996

Mnih, Volodymyr, Koray Kavukcuoglu, David Silver, Andrei A. Rusu, Joel Veness, Marc G. Bellemare, Alex Graves et al. "Human-level control through deep reinforcement learning." *Nature* 518, no. 7540 (2015): 529-533.

Novel Dosing and Clinical Trials with Action-Derived reward policies

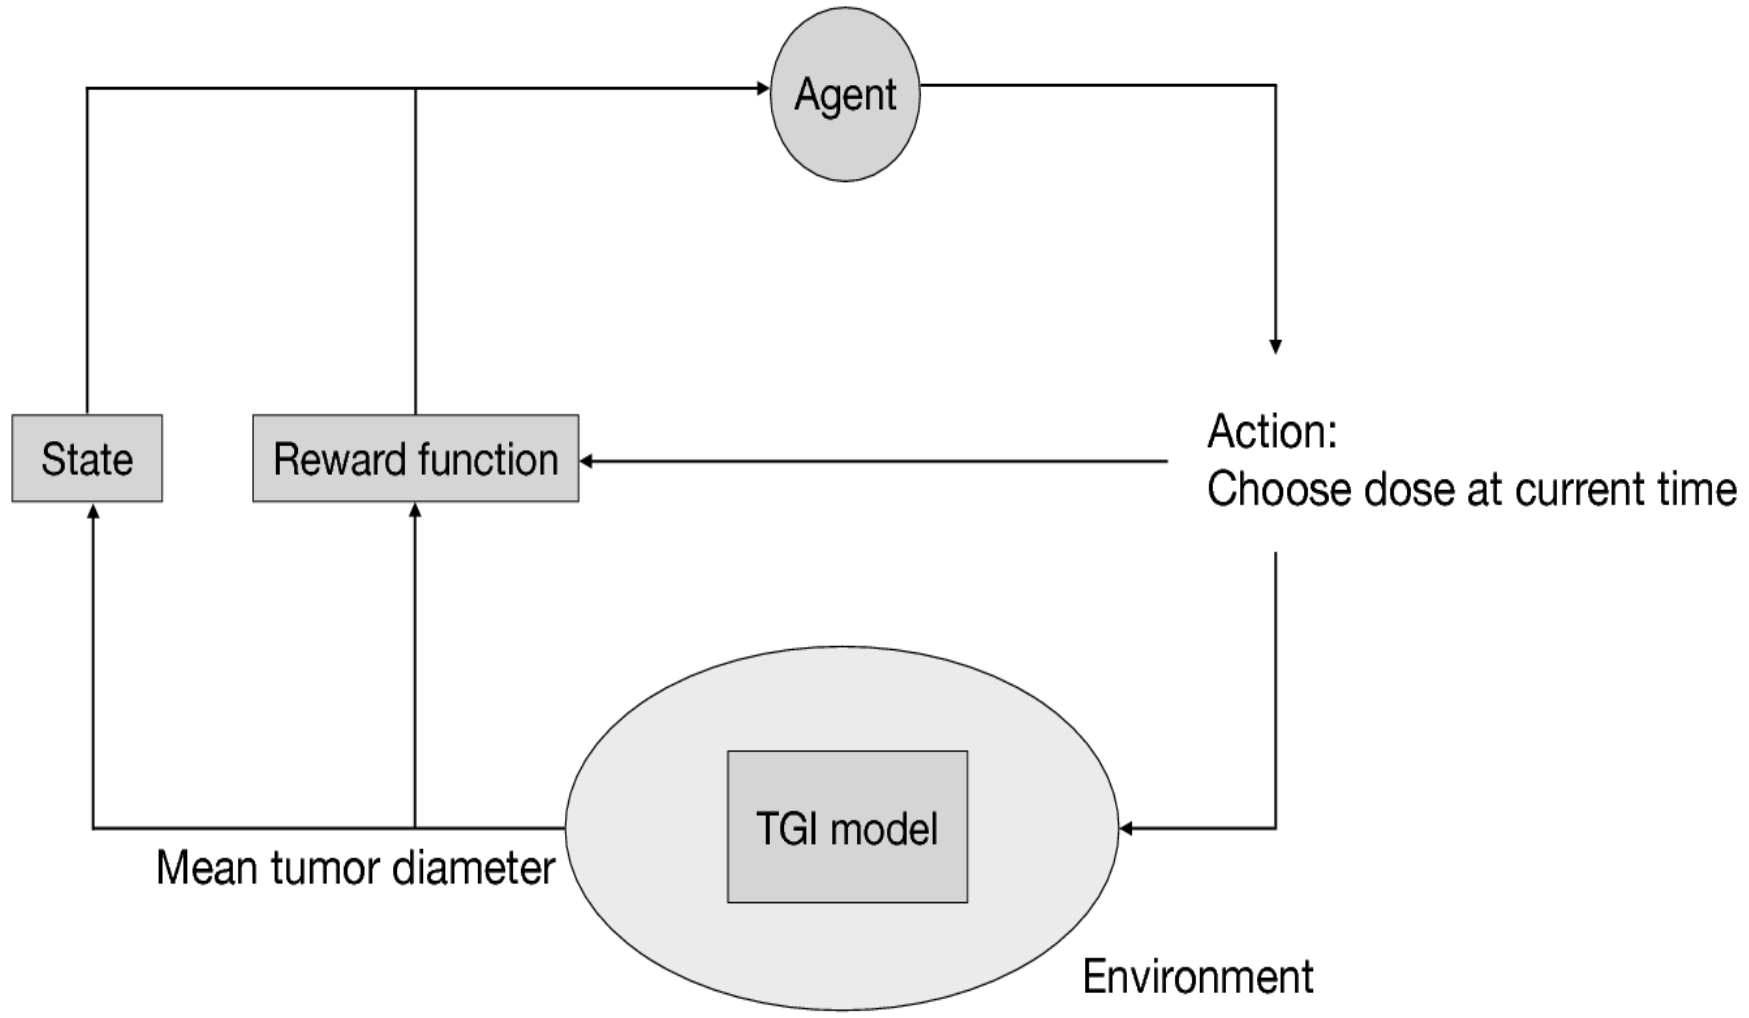

The reinforcement learning agent interacts with an environment containing a tumor growth inhibition (TGI) model. The reward is determined in part by the values used for the reinforcement learning model's state and the agent's most recent action.

Off-Policy deep Q learning

$$Q^*(s,a) = R(s,a,s') + \gamma \sum_{s' \in S} T(s,a,s') \max_{a'} Q^*(s',a')$$

Actions are sampled according to an epsilon-greedy strategy, after which the optimal action from each state can be found:

$$\pi^*(s) = \text{argmax}_a Q^*(s,a)$$

$R = c\ (MTD_t - MTD_{t'})$ - penalty . concentration

$R_{final} = c_{final}\ (MTD_{initial} - MTD_{final})$

$C= 1$, $C_{final}=10$

TMZ penalties= 1 and 5

PCV penalties= 1 and 10

Digital Therapeutics and Algorithmic Design of Clinical Trials

Goal: Reduce number of doses for all patients

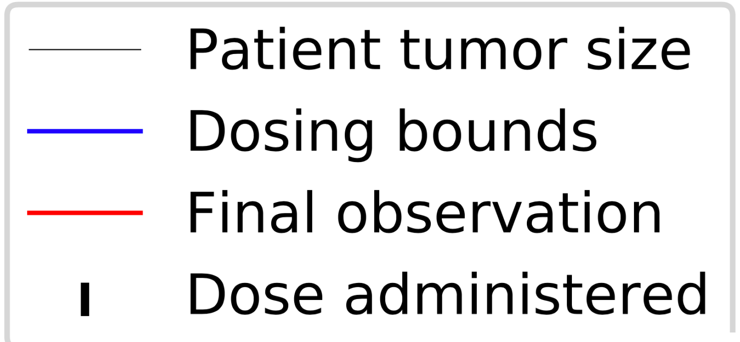

Human experts

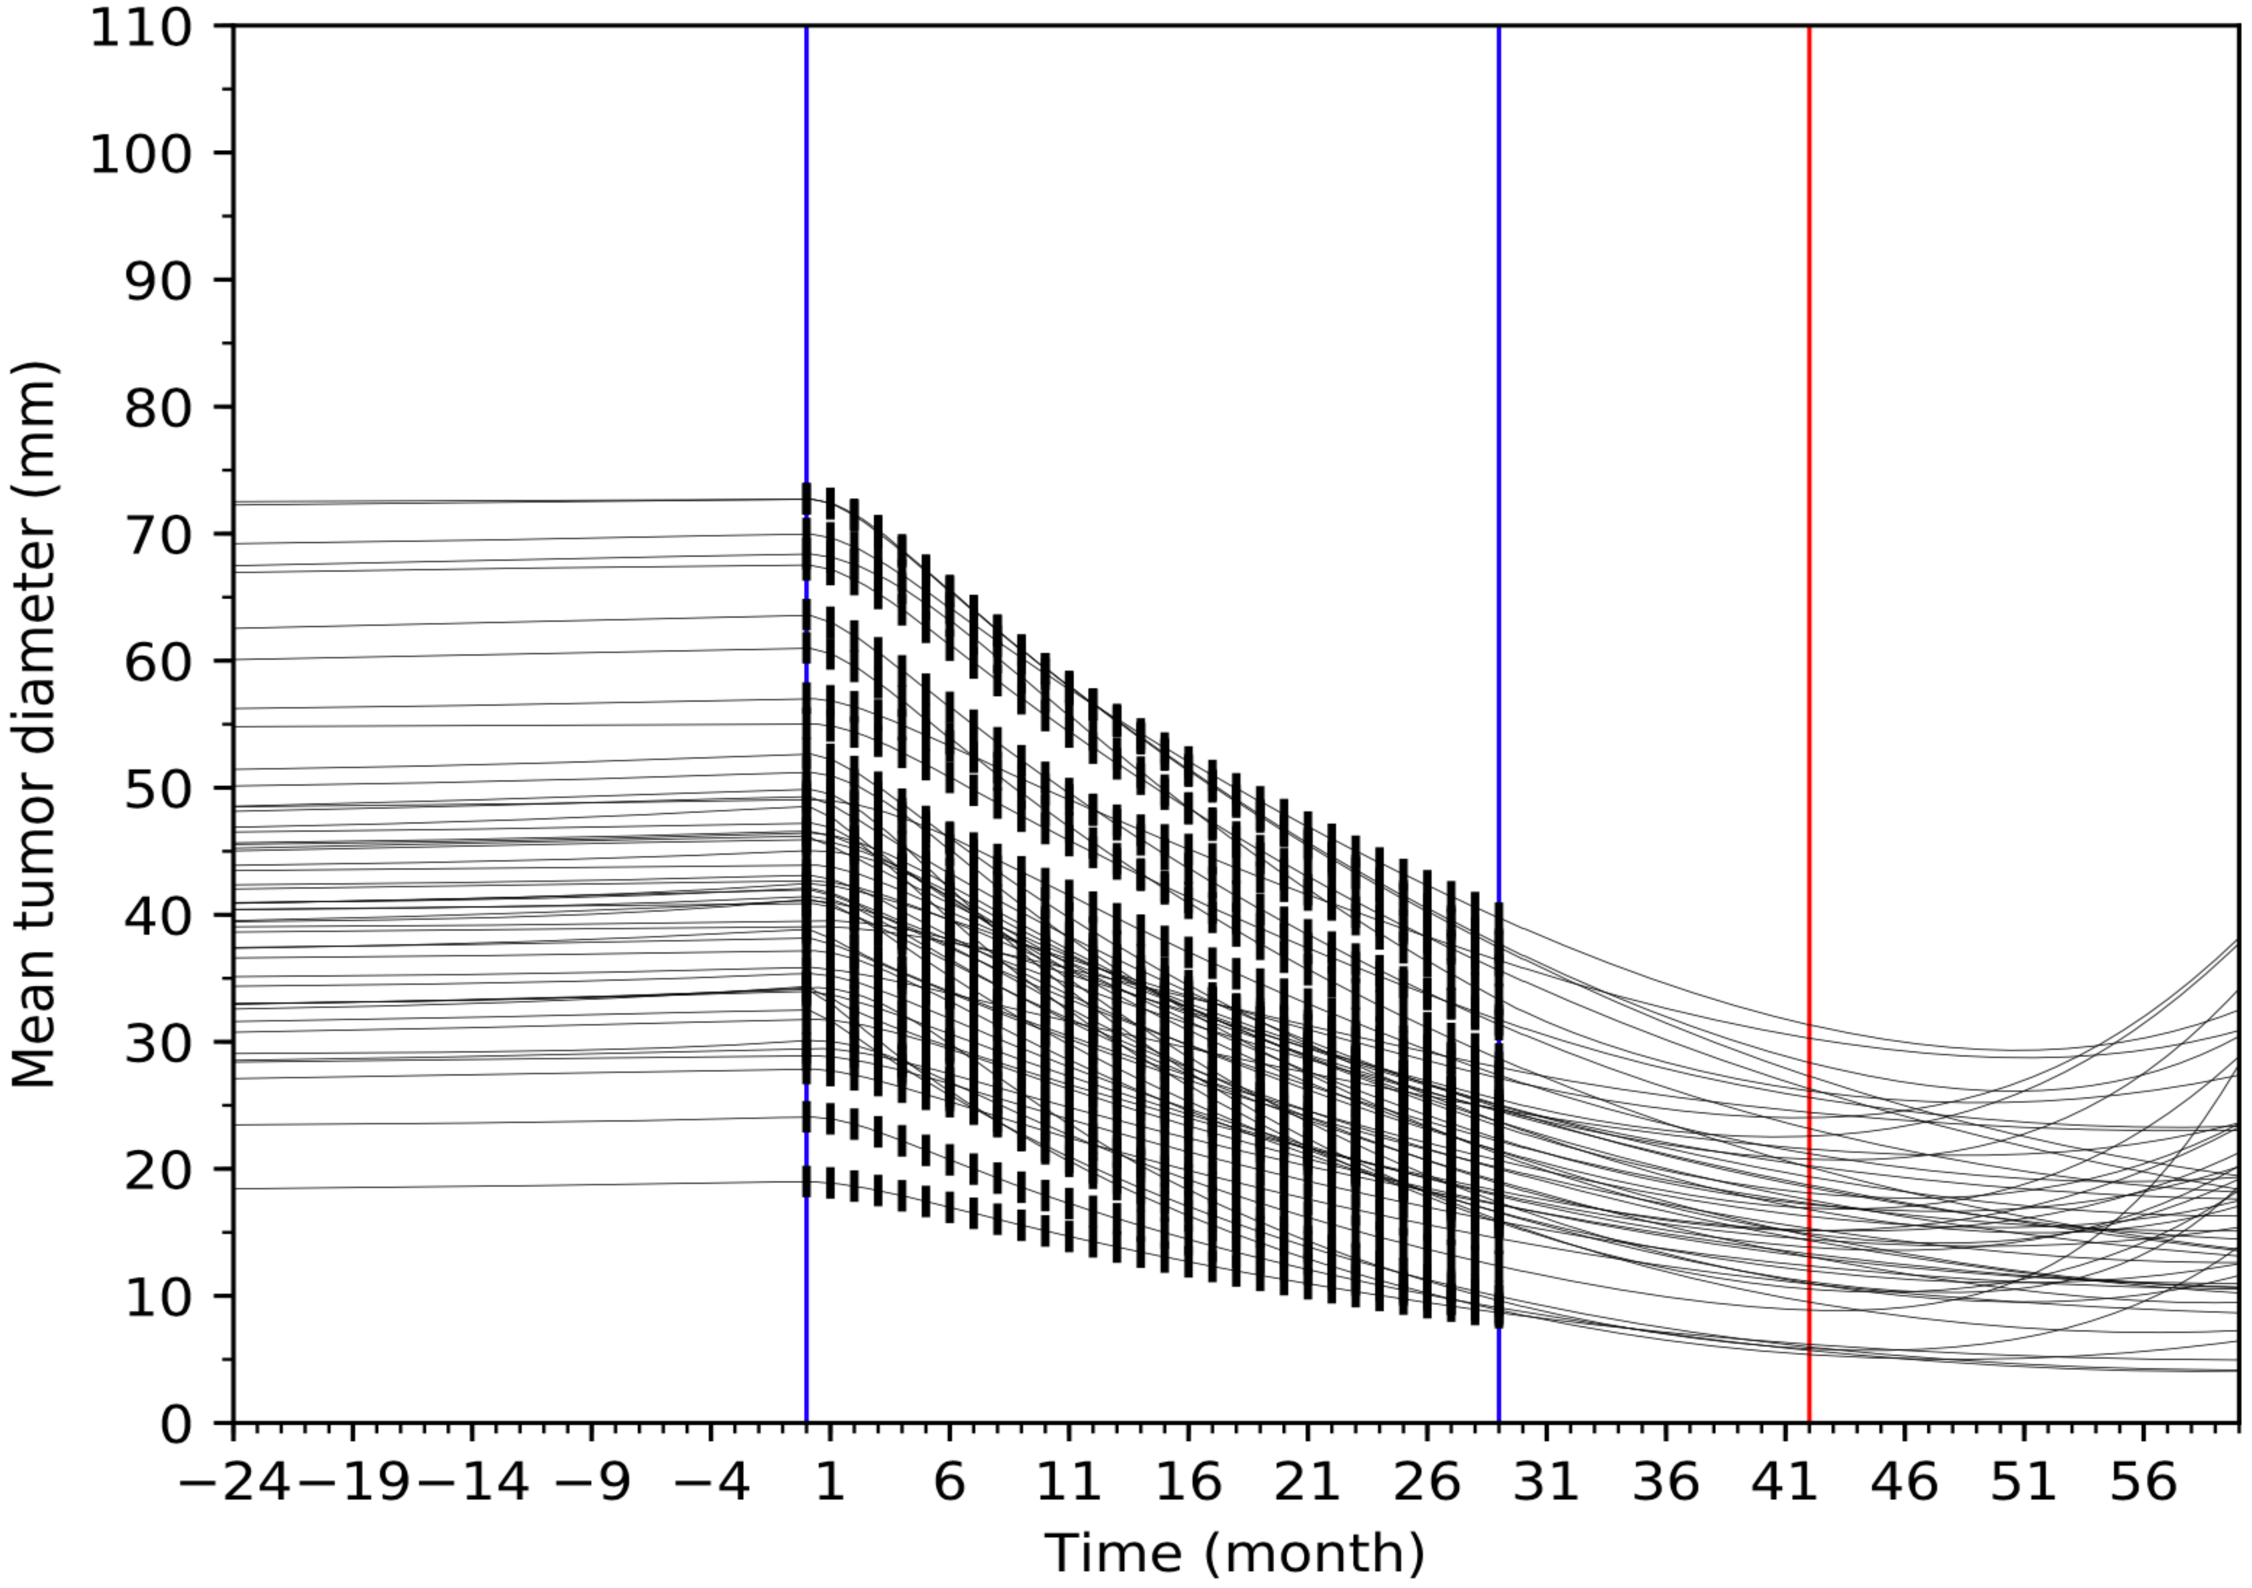

Machine policy

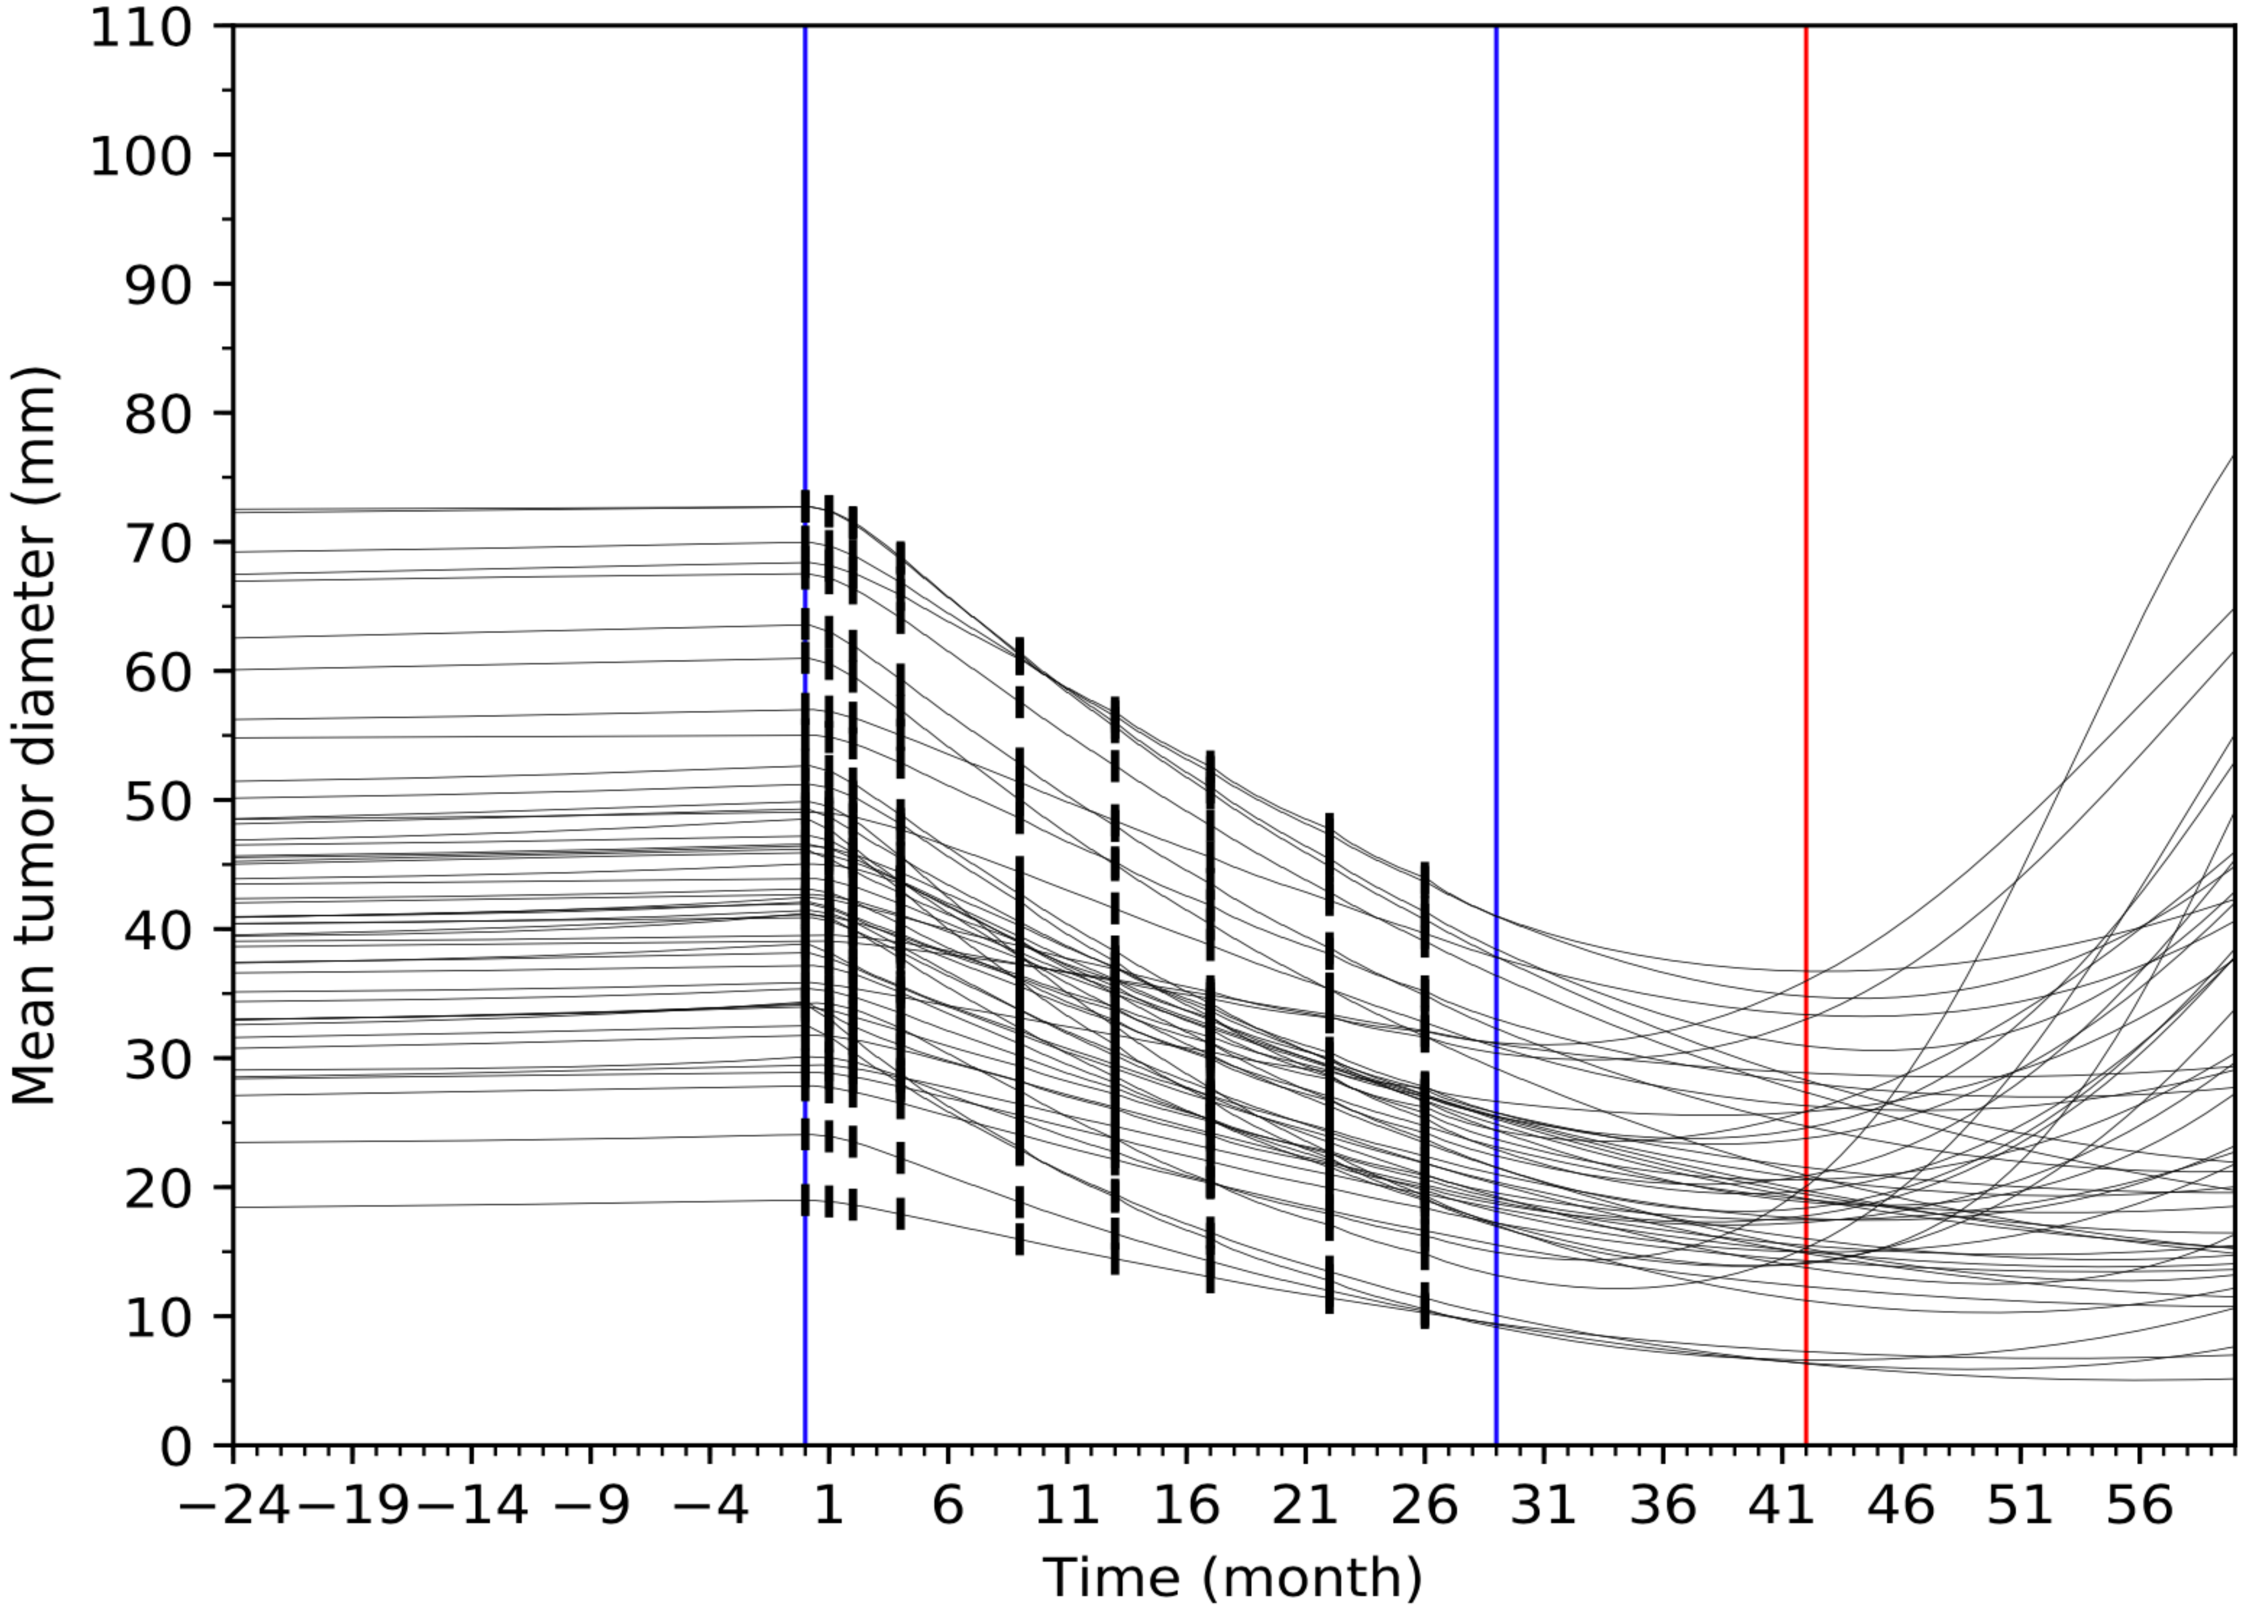

Pratik Shah. Ph.D.
pratiks@media.mit.edu

Digital Therapeutics and Algorithmic Design of Clinical Trials

Goal: Reduce number of doses for all patients

	Month of trial																													
Dosing regimen	0	1	2	3	4	5	6	7	8	9	10	11	12	13	14	15	16	17	18	19	20	21	22	23	24	25	26	27	28	29
Expert	✓	✓	✓	✓	✓	✓	✓	✓	✓	✓	✓	✓	✓	✓	✓	✓	✓	✓	✓	✓	✓	✓	✓	✓	✓	✓	✓	✓	✓	✓
Machines	✓	✓	✓	.	✓	.	.	.	.	✓	.	.	.	✓	.	.	.	✓	.	.	.	.	✓	.	.	.	✓	.	.	.

2018- Proceedings of Machine Learning Research-MLHC-2018
G. Yauney and Pratik Shah

Pratik Shah. Ph.D.
pratiks@media.mit.edu

Digital Therapeutics and Algorithmic Design of Clinical Trials

Goal: Reduce number of interventions for individual patients

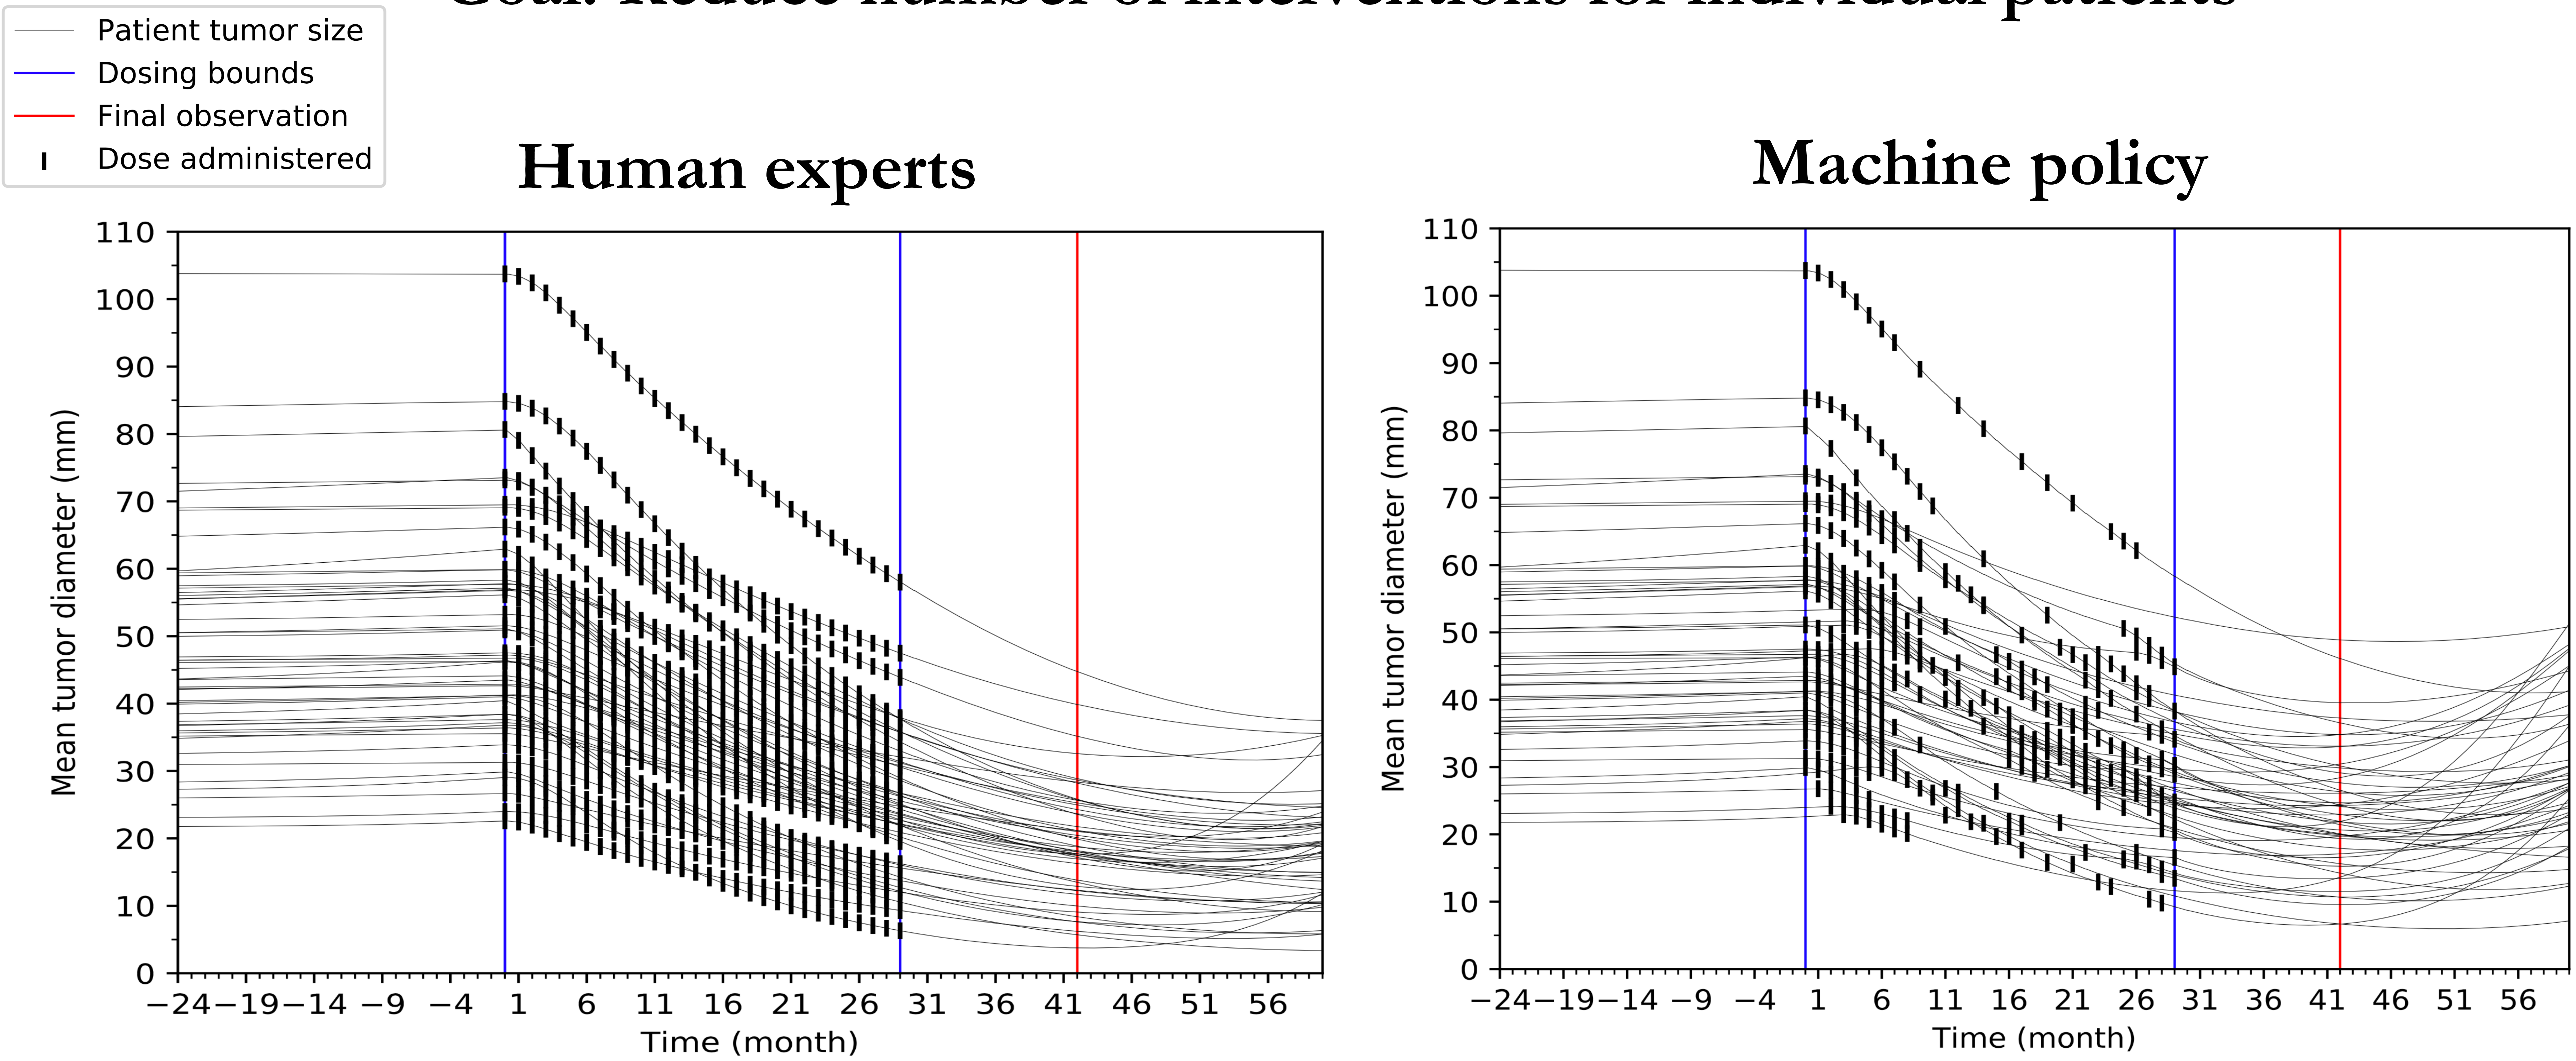

2018- Proceedings of Machine Learning Research-MLHC-2018
G. Yauney and Pratik Shah

Pratik Shah. Ph.D.
pratiks@media.mit.edu

Learned policies with penalties and MTD reductions in TMZ cohort trials

			Month of trial																													
	Conc.	Penalty	0	1	2	3	4	5	6	7	8	9	10	11	12	13	14	15	16	17	18	19	20	21	22	23	24	25	26	27	28	29
A	Fixed	None	1	1	1	1	1	1	1	1	1	1	0	1	1	1	1	1	1	1	1	1	1	1	1	1	1	1	1	1	1	1
B		Small	1	1	1	0	1	0	0	0	0	1	0	0	0	1	0	0	0	1	0	0	0	0	1	0	0	0	1	0	0	0
C		Large	1	0	0	0	0	0	1	0	0	0	0	0	0	0	0	1	0	0	0	0	0	0	1	0	0	0	0	0	0	1
D	Variable	None	1	1	1	1	.25	.75	1	1	.75	.75	.75	1	1	1	1	1	1	1	1	1	1	1	1	1	1	1	1	1	1	1
E		Small	1	1	.75	.75	0	0	.25	.5	0	.5	0	0	.5	0	.5	0	0	.5	0	0	.5	0	.5	0	0	.25	.5	0	0	.5
F		Large	1	.5	0	0	0	0	0	0	.25	.25	0	.25	.25	0	.25	0	.25	0	.25	.25	0	.25	0	.25	.25	0	.25	.25	0	.25

Table 2: Learned policies for trial-based TMZ experiments where all simulated patients received the same dose each month. Agents with a fixed concentration could give a unit dose, whereas agents with a variable concentration could give doses at 25%, 50%, 75%, and 100% of the unit dose's concentration. Penalty: size of the dose penalty the agent tried to minimize while maximizing mean tumor diameter reduction. Conc: concentration.

	Trial parameters			Average MTD change			
	Treatment	Conc.	Type	Expert policy	No penalty	Small penalty	Large penalty
A	TMZ	Fixed	Patient	-61.04% ± 11.63%	-60.95% ± 11.64%	-51.89% ± 15.31%	-35.97% ± 17.05%
B	TMZ	Fixed	Trial	-62.18% ± 10.50%	-62.17% ± 10.51%	-54.09% ± 14.05%	-46.27% ± 17.93%
C	TMZ	Variable	Patient	-60.51% ± 10.68%	-60.23% ± 10.80%	-39.35% ± 46.02%	-7.15% ± 40.80%
D	TMZ	Variable	Trial	-62.86% ± 11.41%	-62.72% ± 11.43%	-54.03% ± 15.21%	-45.69% ± 18.06%

mit media lab
Pratik Shah. Ph.D.
pratiks@media.mit.edu

Research topic III

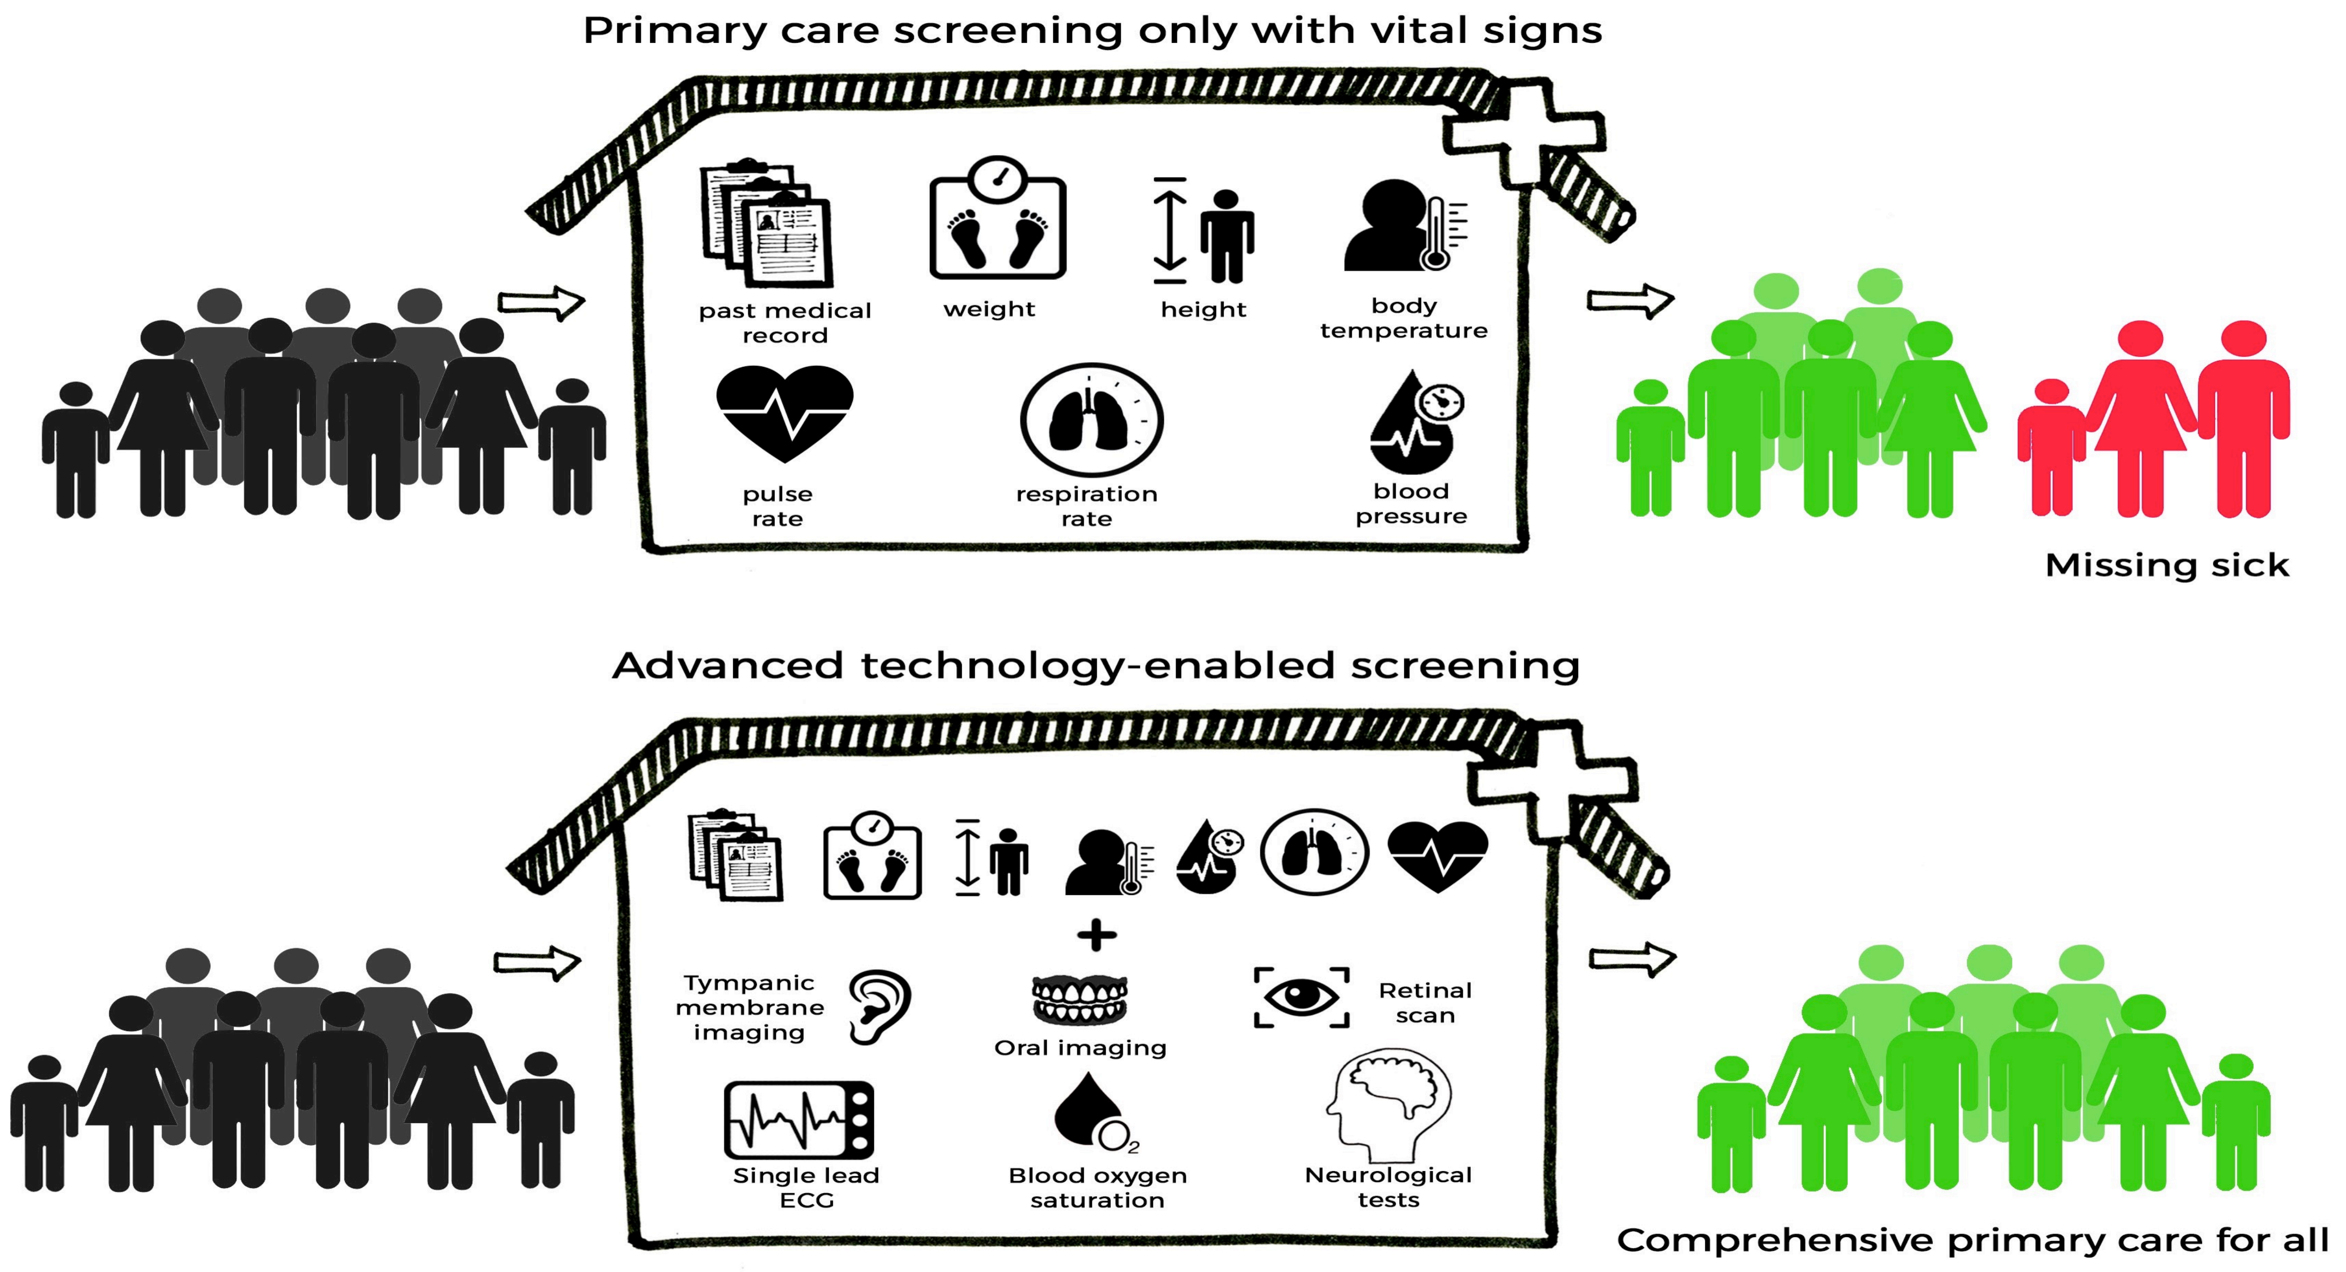

Real world data and evidence

Pratik Shah. Ph.D.
pratiks@media.mit.edu

Novel end points and Point of Care Evidence

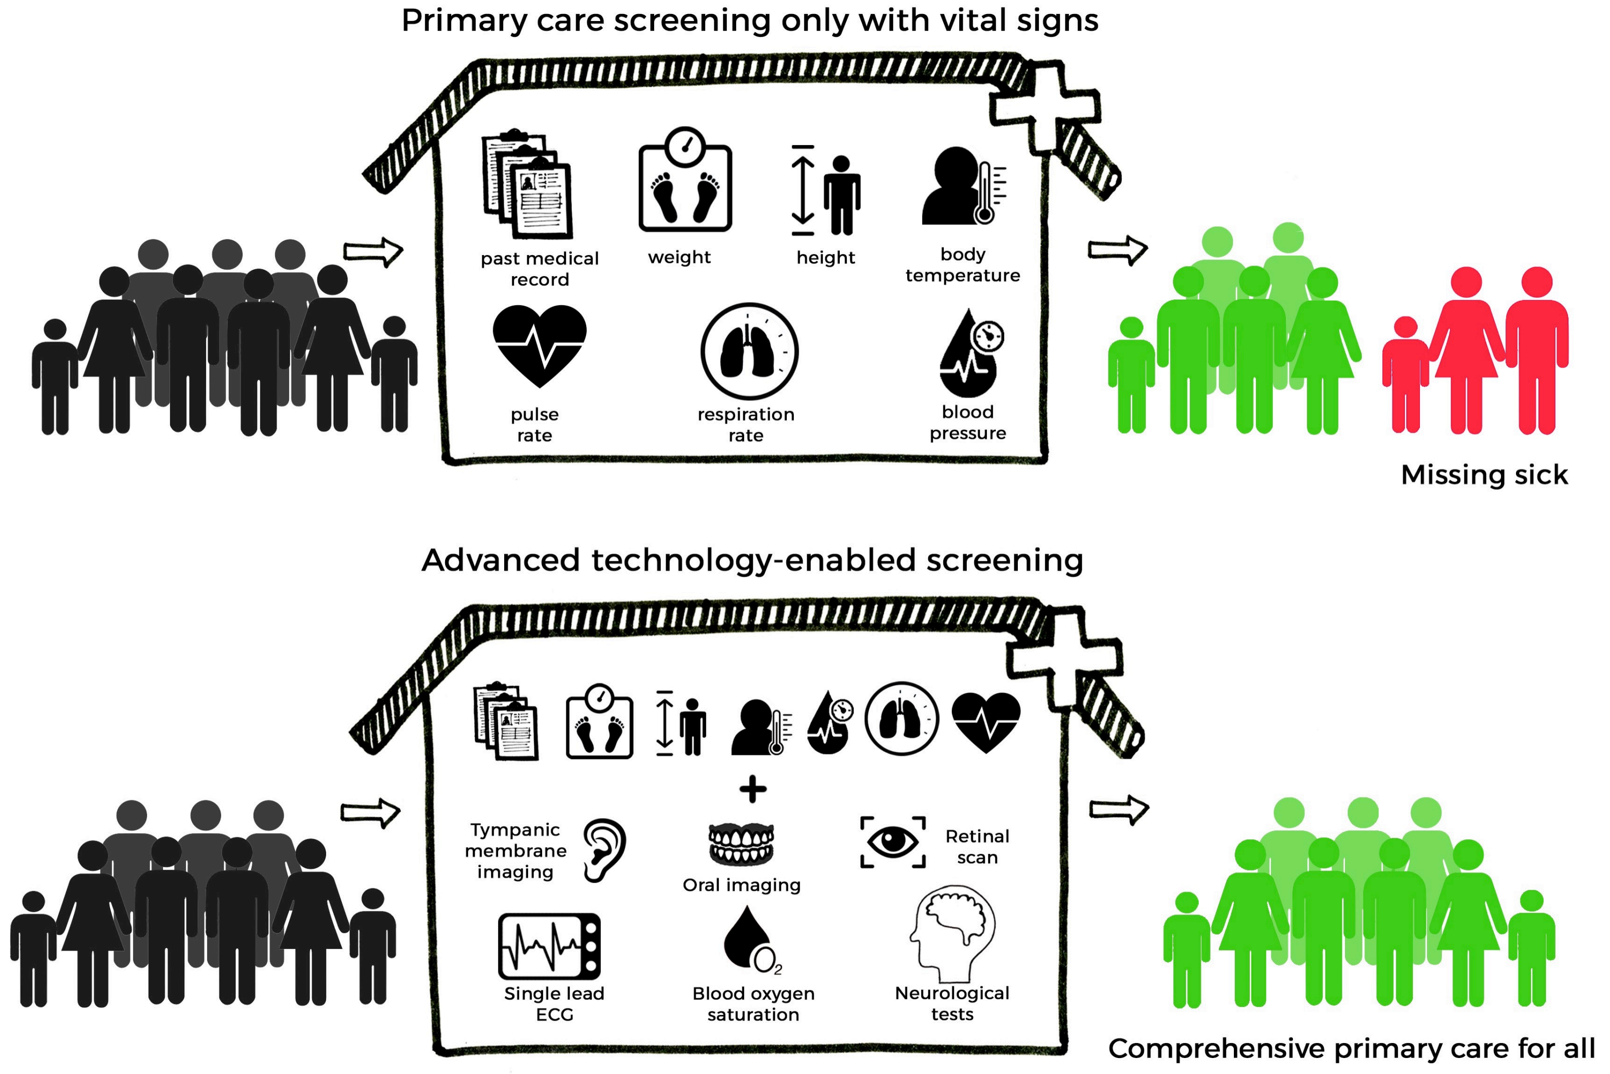

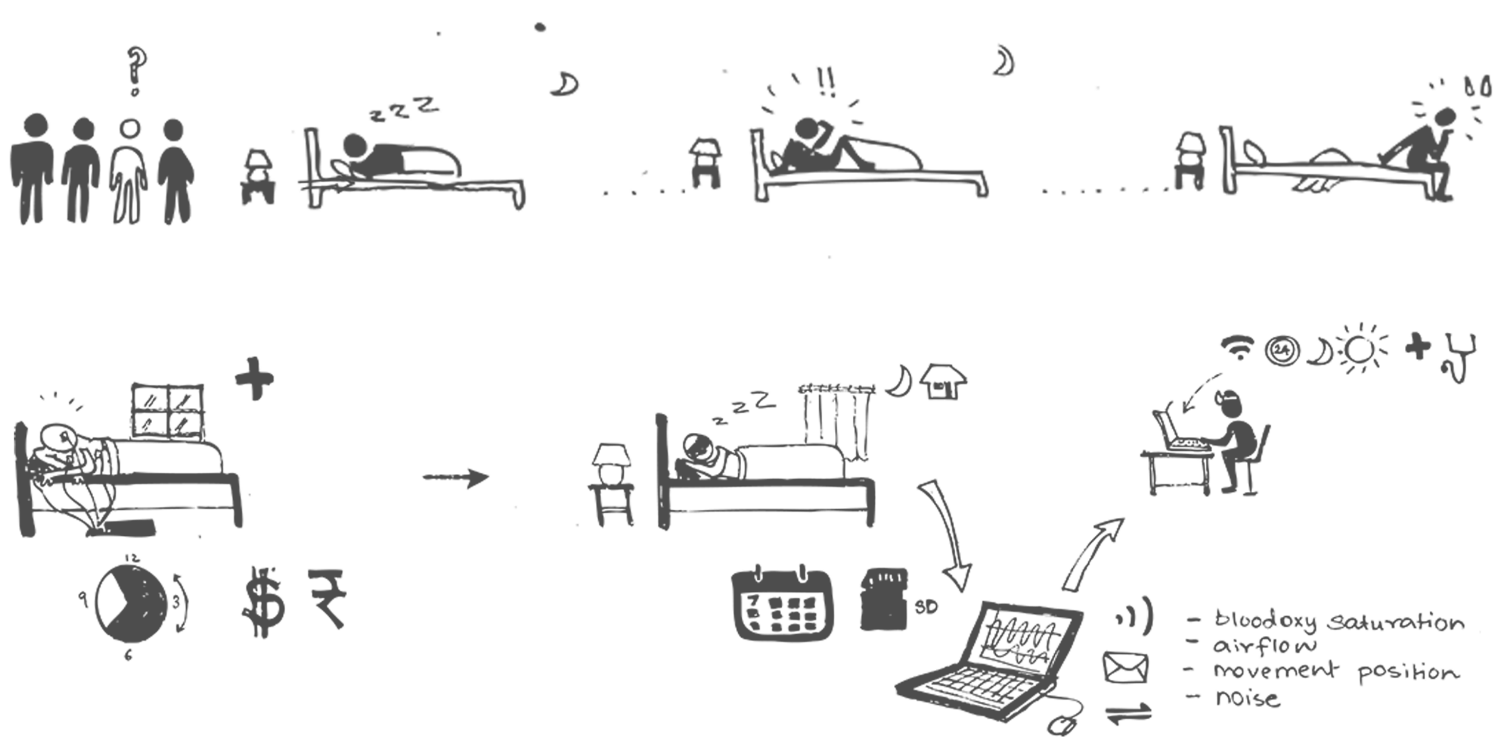

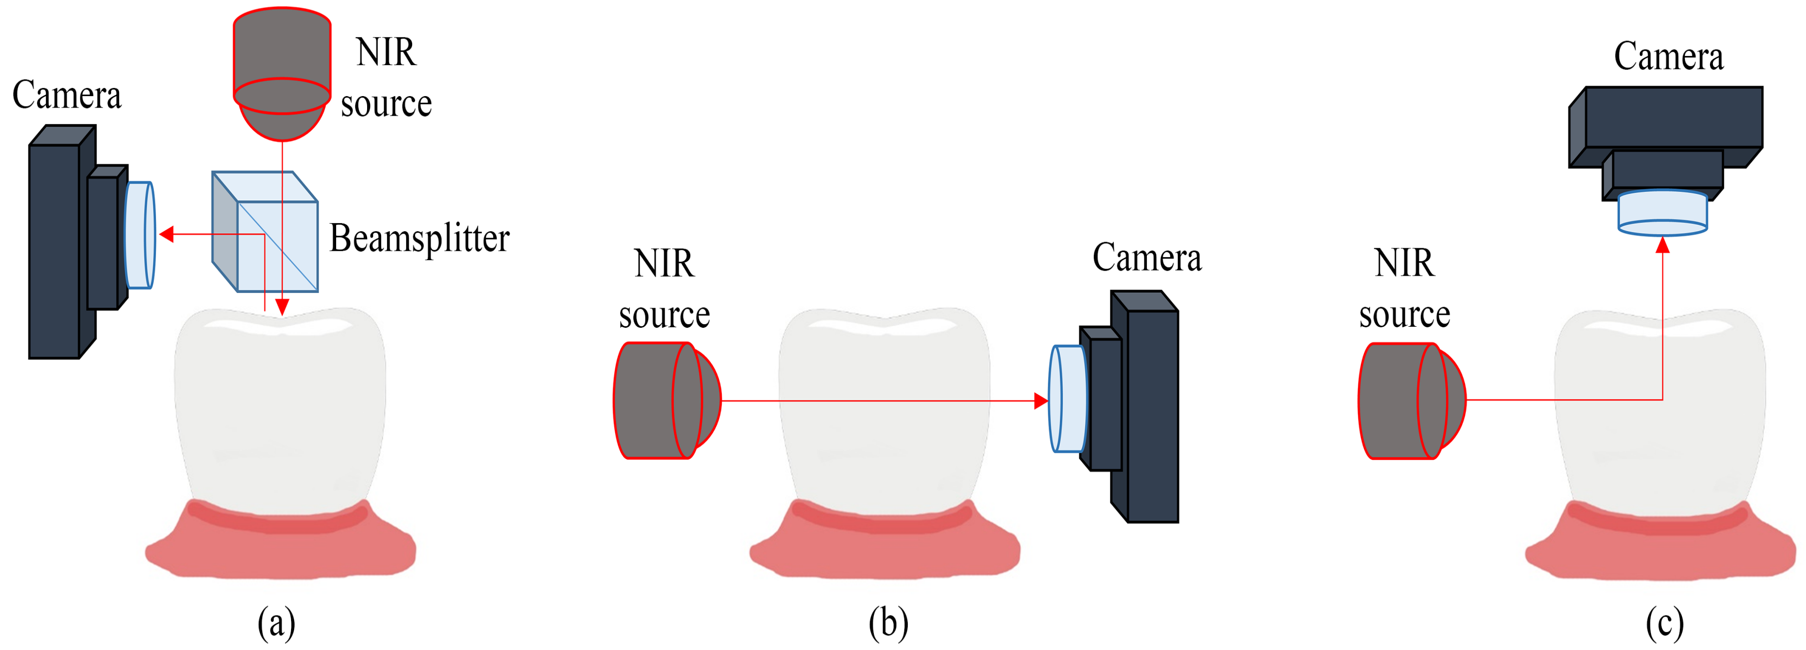

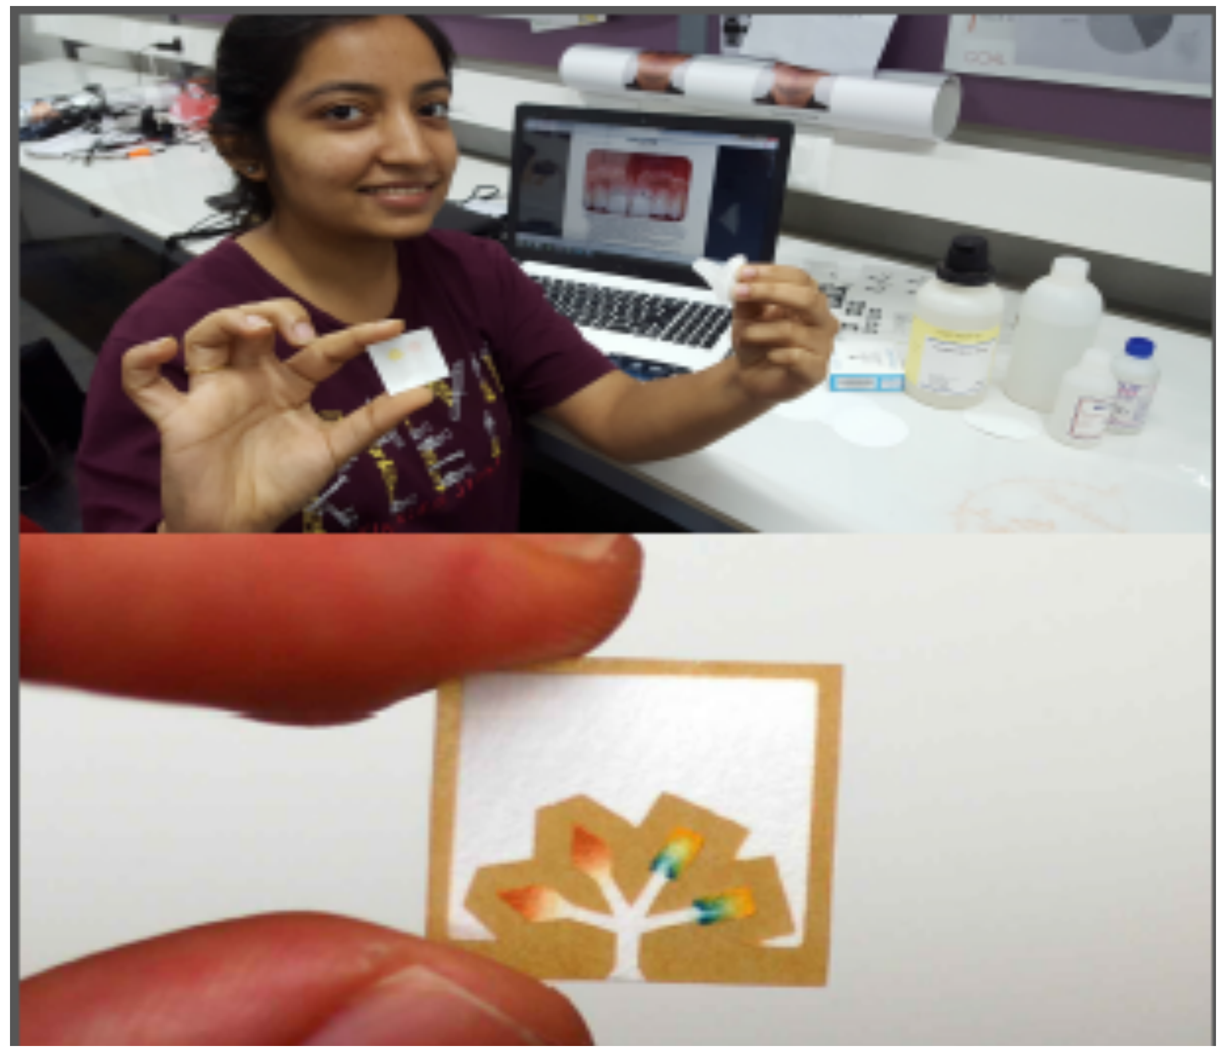

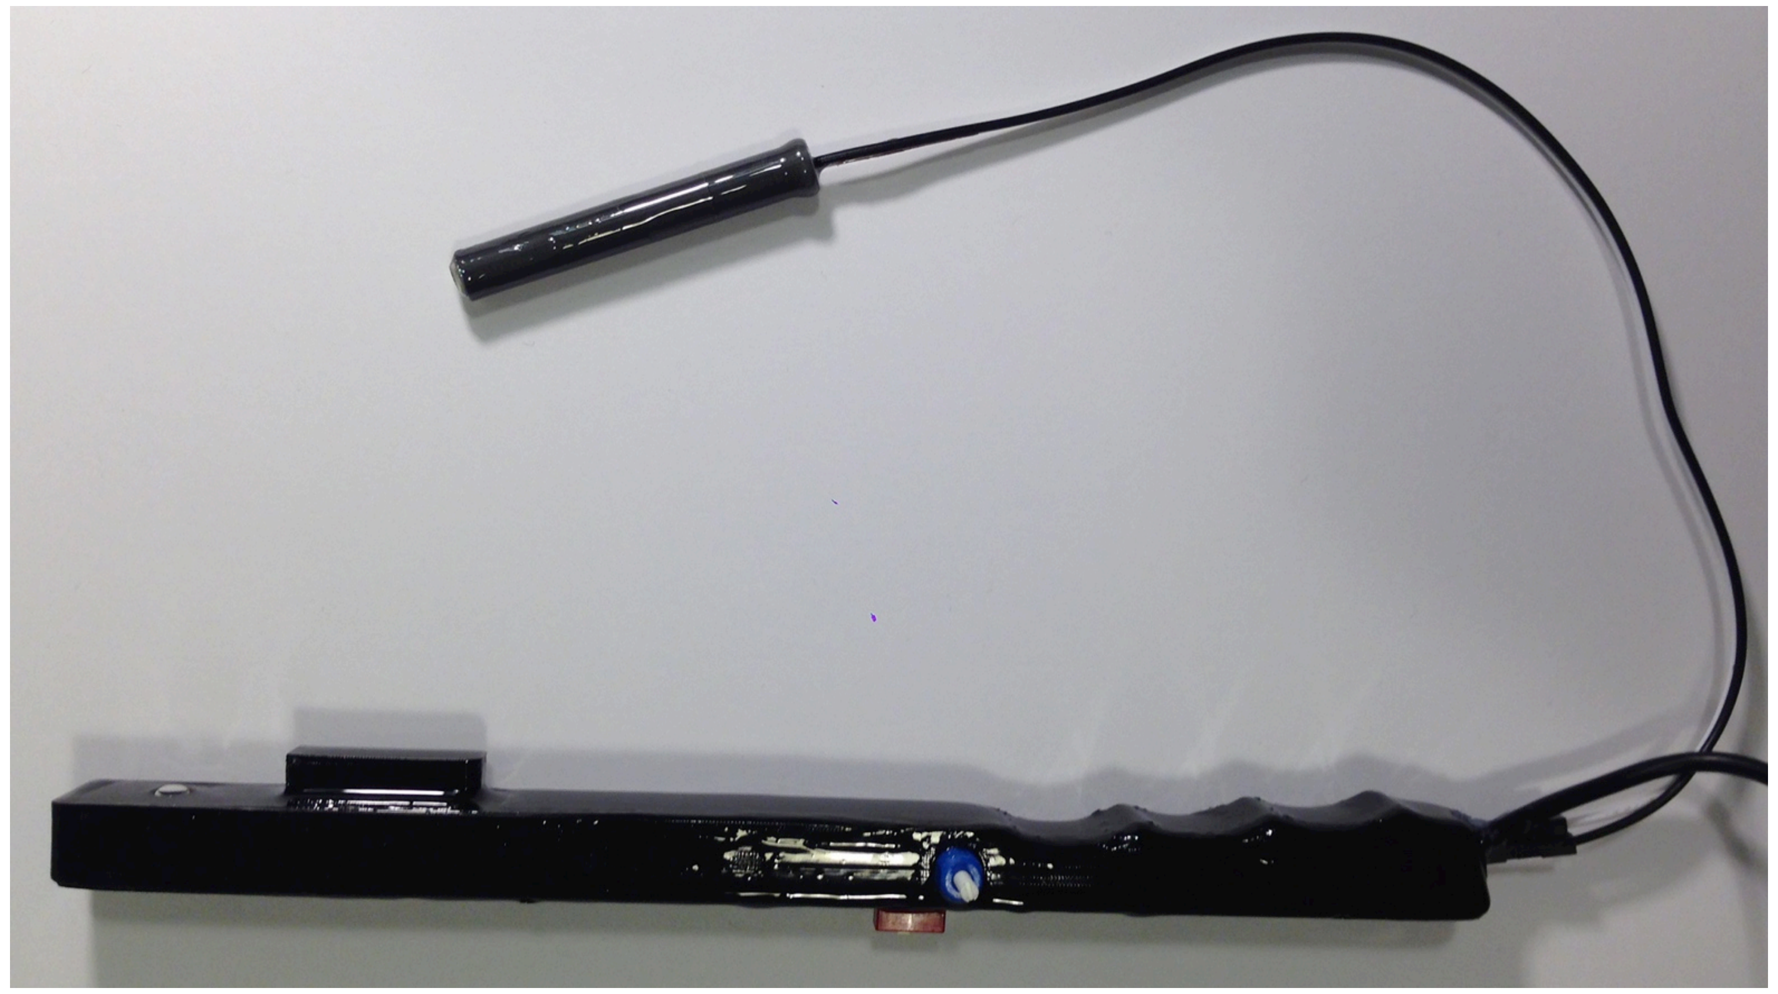

Pratik Shah. Ph.D.
pratiks@media.mit.edu

Real World Data and Evidence using Point-of-Care Devices

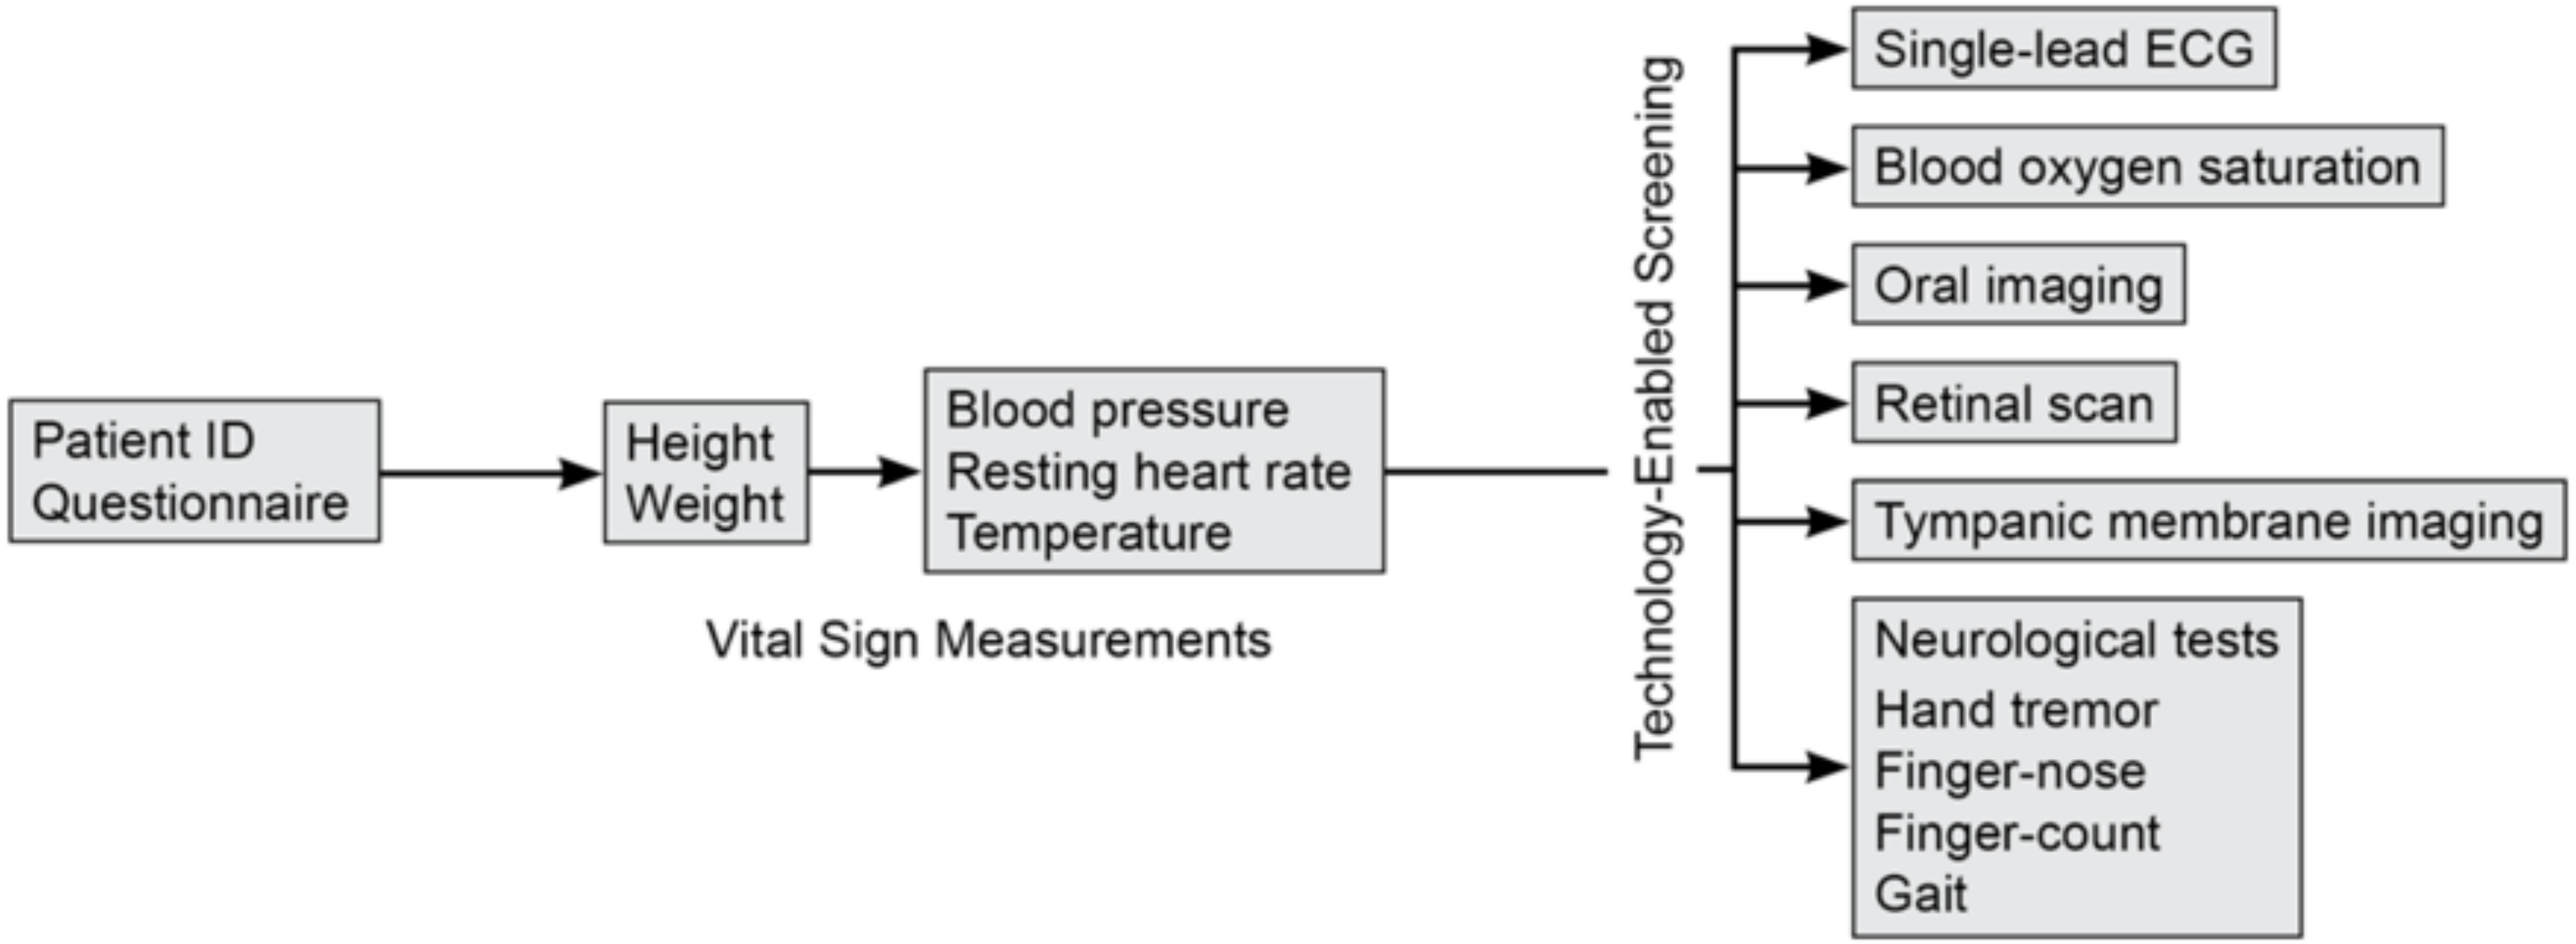

Technology-Enabled Non-Invasive Examinations Augment Primary Care
Pratik Shah, Gregory Yauney, et.al.
The BMJ Open Apr 2018, 8 (4) e018774; DOI: 10.1136/bmjopen-2017-018774.

Pratik Shah. Ph.D.
pratiks@media.mit.edu

Artificial Intelligence in Clinical Development to Improve Public Health:

Key recommendations

- Data and research sandbox that addresses key challenges and leverages opportunities to hosts confidential and high value health data and facilitates collaborations;

- Research collaborations academia, government, life sciences, biotechnology, foundations, universities, and technology corporations

- New models and technologies for health research, early discovery, safer/faster clinical trials, reduced healthcare costs, and digitally empowered researchers, clinicians, regulators, and patients;

Key recommendations: Artificial Intelligence in Clinical Development to Improve Public Health
2018, Under review. Nature Drug Discovery

Pratik Shah. Ph.D.
pratiks@media.mit.edu

Machine Learning Toolbox for Health

Tasks / Approaches	Training with less data	Training with few labels	Learning to learn	Maintaining data Privacy	Prediction of future events
One shot learning					
Transfer learning					
Distributed learning					
Secure learning					
Generative adversarial networks					
Reinforcement learning					

mit media lab

Pratik Shah. Ph.D.
pratiks@media.mit.edu

More information: Health 0.0 Research Program

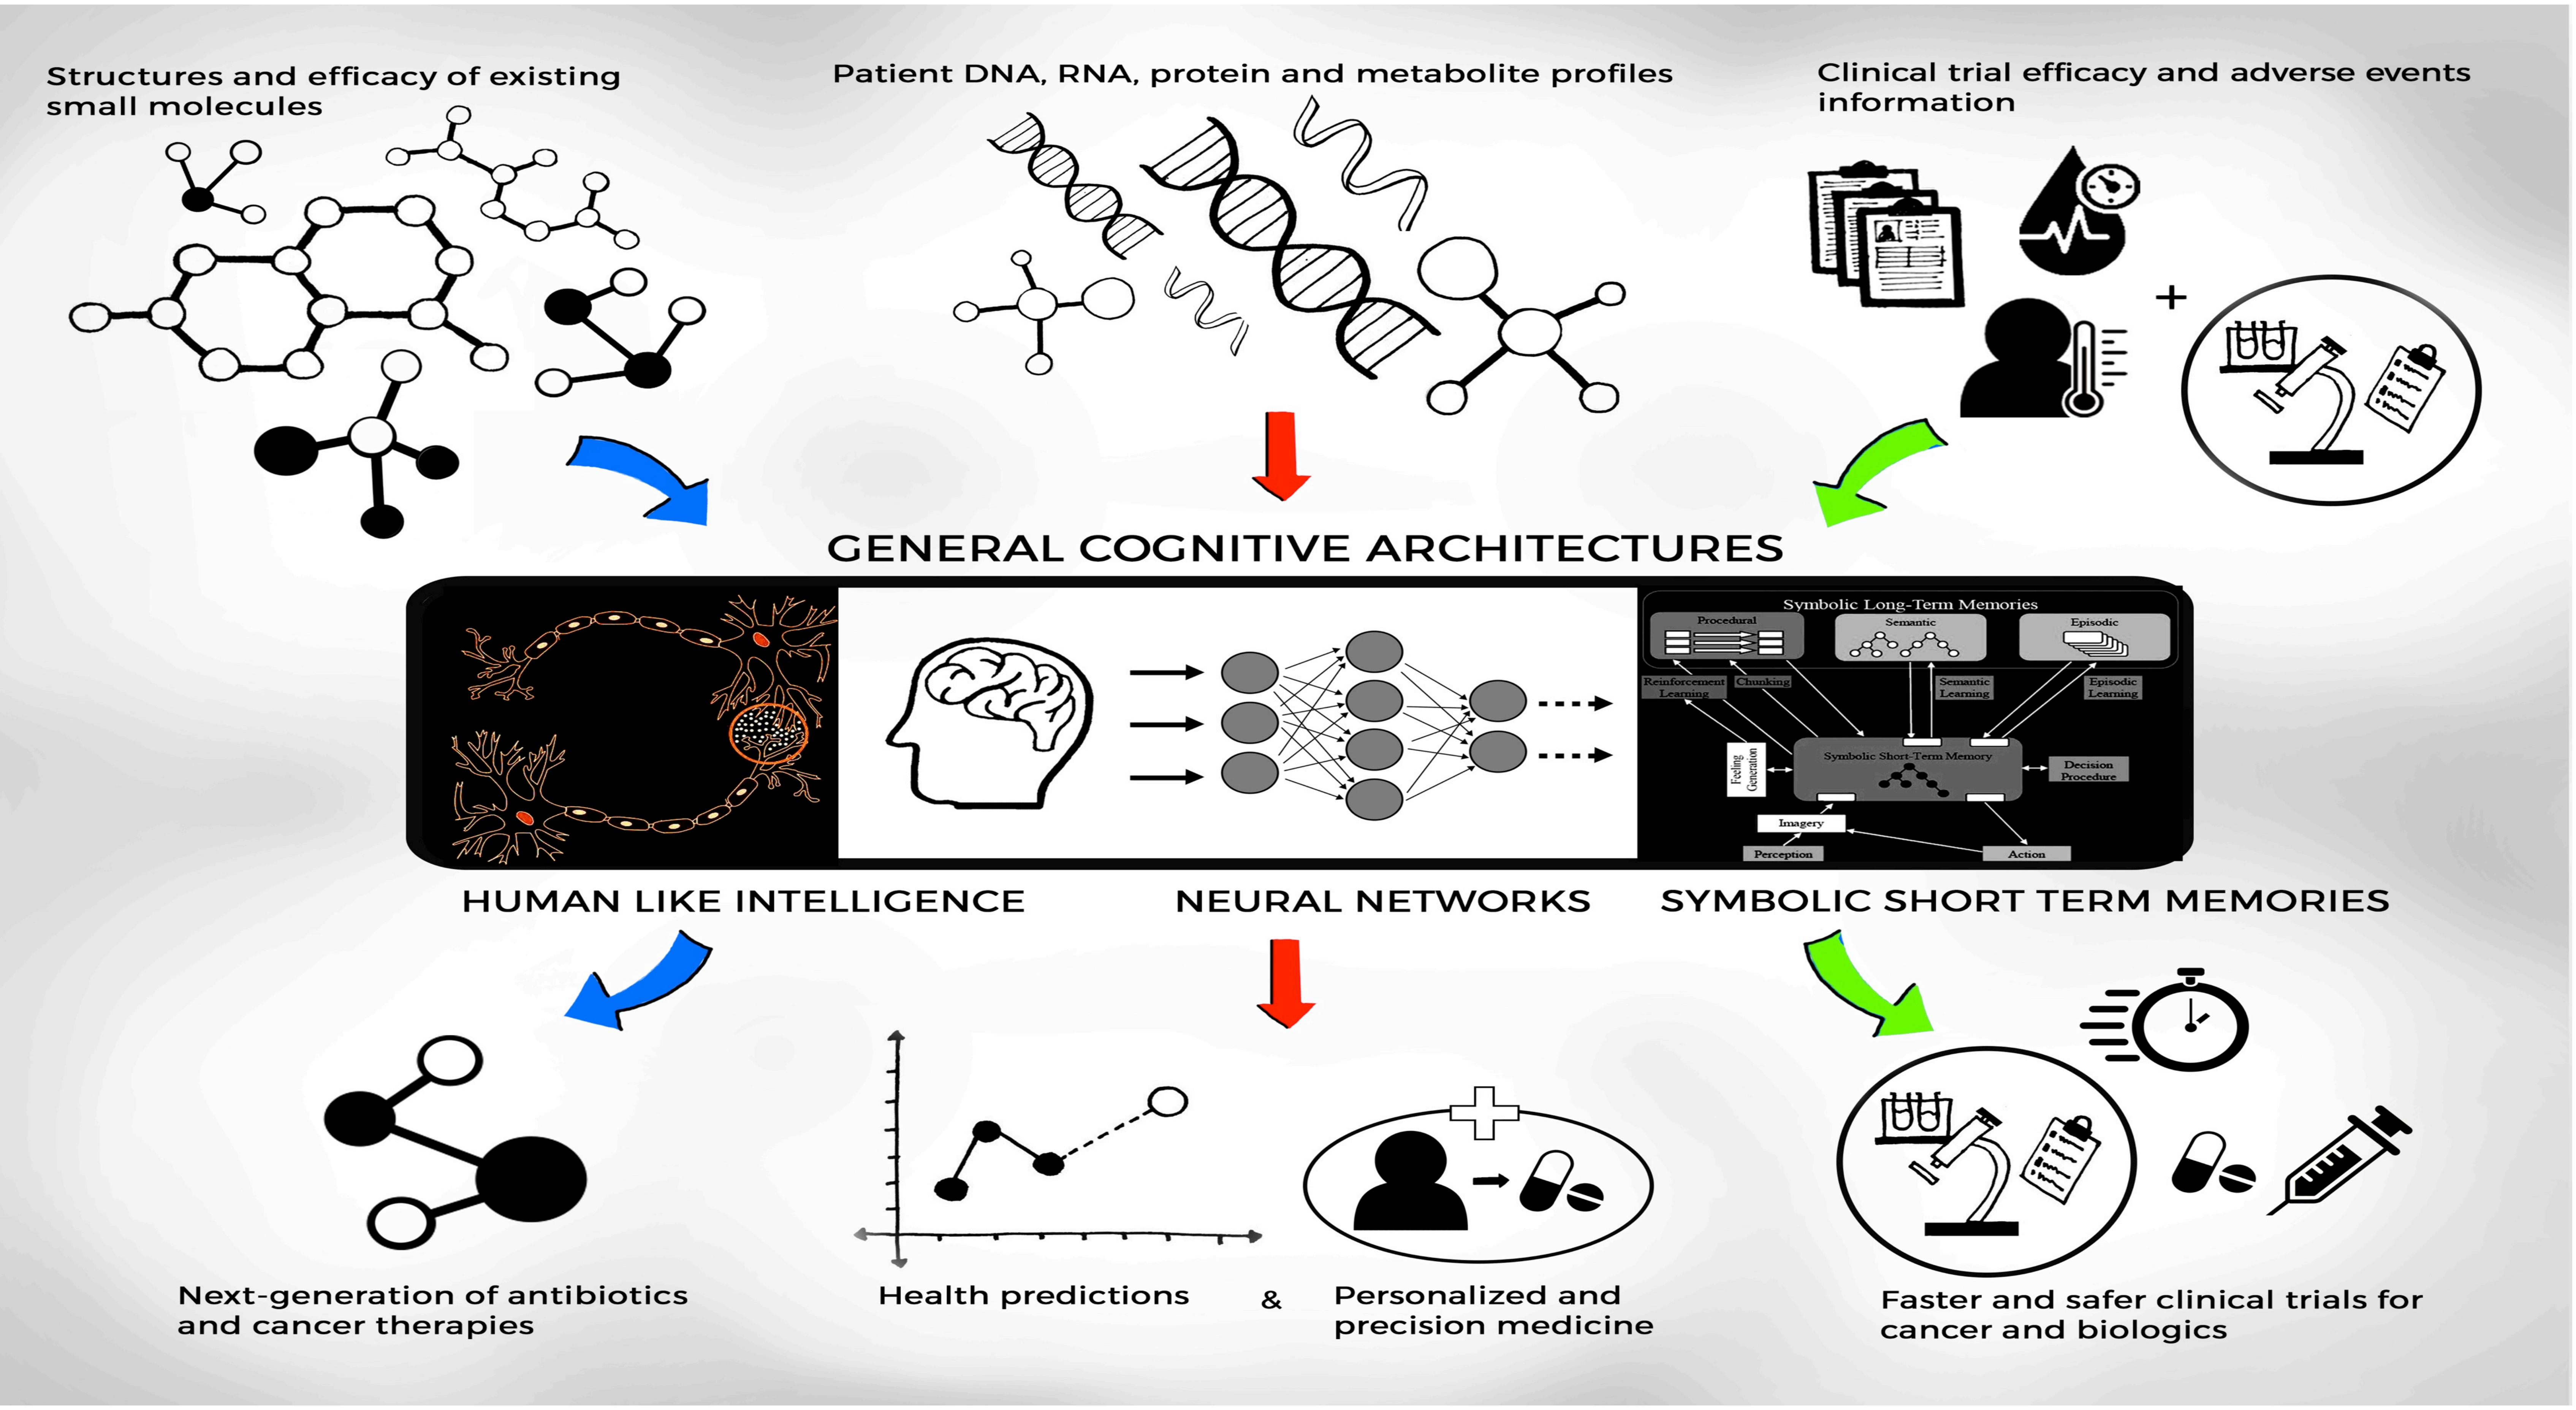

mit media lab

Pratik Shah. Ph.D.
pratiks@media.mit.edu

Language and Literacy for Machine Learning Research

- Inclusion and diversity in datasets are key to prevent a biased clinical development process driven by algorithms
- Causal inferences and data standards should be implemented
- Machine learning and humans: synergy not competition
- Black box: Saliency, explicability and explainability is important but not rate-limiting
- Deployment and testing: New clinical trials for algorithms to be piloted
- Regulation and communication with patients and physicians to de-risk technology
- Automation vs. Knowledge vs. Intelligence
- It is a new field. We can do it right together vs. operating in silos

mit media lab
Pratik Shah. Ph.D.
pratiks@media.mit.edu