

MICHAEL YAO

michaelsyao.com

michael.yao@pennmedicine.upenn.edu

RESEARCH STATEMENT

My research focuses on trustworthiness and robustness for deep learning, offline optimization, meta-learning, and bandit problem formulations. I am broadly interested in developing methods that leverage prior knowledge to help algorithms better generalize to new distributions. I explore these problems in the setting of generative design, AI4Science, minority health disparities, and medical imaging.

EDUCATION

University of Pennsylvania, MD-PhD Program

NIH F30 NRSA Fellow, HHMI-NIBIB Interfaces Fellow | 2021 - Present

Advised by [Osbert Bastani](#) and [James Gee](#)

MD-PhD, Bioengineering (in progress)

MSE, Computer Science (in progress)

California Institute of Technology

Salutatorian | 2017 - 2021

Advised by [Mikhail Shapiro](#)

BS, Applied Physics

EXPERIENCE

Health VC Fellow, [25madison](#)

New York City | 2024

Bridged clinical and engineering teams to help drive clinical operations for stealth healthtech incubation. Led exploratory investment research into digital health sectors.

AI Clinical Fellow, [Glass Health](#)

Remote | 2023-2024

Released and assessed clinical guideline articles as knowledge sources for large language models (LLM). Investigated applications of LLMs for medical education.

Medical LLM Consultant, [Scale AI](#)

Remote | 2023

Evaluated use cases of LLMs for healthcare. Red team tested LLMs for accuracy and trustworthiness in real-world clinical workflows.

PhD Research Intern, [Microsoft](#)

Redmond, WA | 2022

Developed ML methods for accelerated MRI imaging. Proposed novel techniques for better generalization of MRI image reconstruction models.

Software Engineer, [Hyperfine Research](#)

Guilford, CT | 2021

Implemented and validated algorithms for more robust MRI signal acquisition and image post-processing in MR software across 25+ hospital sites.

SELECTED
PUBLICATIONS

- [1] **Yao MS**, Gee JC, Bastani O. Diversity by design: Leveraging distribution matching for offline model-based optimization. Under peer review. (2024). [Link](#)
- [2] Gee JC, **Yao MS**. Effective structured information extraction from chest

- radiography reports using open-weights large language models. Radiology (Editorial). (2025). [Link](#)
- [3] **Yao MS**, Chae A, Kahn CE, Witschey WR, Gee JC, Sagreiya H, Bastani O. Evidence is all you need: Ordering imaging studies via language model alignment with the ACR Appropriateness Criteria. Under peer review. (2024). [Link](#)
- [4] **Yao MS**, Zeng Y, Bastani H, Gardner J, Gee JC, Bastani O. Generative adversarial model-based optimization via source critic regularization. NeurIPS. (2024). [Link](#)
- [5] Wu Y, Liu Y, Yang Y, **Yao MS**, Yang W, Shi X, Yang L, Li D, Liu Y, Gee JC, Yang X, Wei W, Gu S. A concept-based interpretable model for the diagnosis of choroid neoplasias using multimodal data. arXiv Preprint. (2024). [Link](#)
- [6] Yang Y, Gandhi M, Wang Y, Wu Y, **Yao MS**, Callison-Burch C, Gee JC, Yatskar M. A textbook remedy for domain shifts: Knowledge priors for medical image analysis. NeurIPS (Spotlight). (2024). [Link](#)
- [7] Chae A†, **Yao MS**†, Sagreiya H, Goldberg AD, Chatterjee N, MacLean MT, Duda J, Elahi A, Borthakur A, Ritchie MD, Rader D, Kahn CE, Witschey WR, Gee JC. Strategies for implementing machine learning algorithms in the clinical practice of radiology. Radiology. (2024). [Link](#)
- [8] **Yao MS**†, Chae A†, MacLean MT, Verma A, Duda J, Gee JC, Torigian DA, Rader D, Kahn CE, Witschey WR, Sagreiya H. SynthA1c: Towards clinically interpretable patient representations for diabetes risk stratification. PRIME MICCAI. (2023). [Link](#)
- [9] **Yao MS**, Hansen MS. A path towards clinical adaptation of accelerated MRI. Proc ML4H. (2022). [Link](#)

TEACHING

Instructor and Curriculum Lead, Ethical Algorithms for the Modern Clinician | [Link](#)

TA, Distributed Systems (CIS 5050, Penn) | Spring 2025

TA, Principles of Deep Learning (ESE 5460, Penn) | Fall 2024

TA, Imaging Informatics (EAS 5850, Penn) | Spring 2024, Summer 2024

Head TA, Healthcare and Technology (CIS 7000, Penn) | Fall 2023, Fall 2024

TA, Diagnostic Ultrasound for Medical Students (Penn) | 2023 - Present

TA, Pre-Clinical Medicine (Penn) | 2023 - Present

Head TA, Applied Mathematics (ACM 95a, Caltech) | Winter 2021

TA, Graduate Classical Physics (Ph 106a, Caltech) | Fall 2020

TA, Applied Mathematics (ACM 95b, Caltech) | Spring 2020

TA, Quantum Physics (Ph 12b, Caltech) | Winter 2020

TA, Electrodynamics and Magnetism (Ph 1c, Caltech) | Spring 2019

TA, Operating Systems (CS 24, Caltech) | Spring 2019

TA, Waves and Oscillations (Ph 12a, Caltech) | Fall 2019

TA, Electrodynamics and Magnetism (Ph 1c, Caltech) | Spring 2019

TA, Special Relativity and Electrostatics (Ph 1b, Caltech) | Winter 2019

SERVICE

Referee

RSNA Radiology

RSNA Radiology: Artificial Intelligence

International Conference on Learning Representations (ICLR)

International Conference on Machine Learning (ICML)

AHLI Conference on Health, Inference, and Learning (CHIL)

Ongoing

Anti-Racism Curriculum Lead, Medical Education, University of Pennsylvania SOM

AI Curriculum Task Force Member, University of Pennsylvania SOM

Board Member, Radiology Interest Group, University of Pennsylvania SOM

Admissions Committee, University of Pennsylvania SOM

Peer Tutor, University of Pennsylvania SOM

Technology Committee Vice-Chair, [American Physician Scientists Association](#)

Director of Data Science & AI, [MDplus](#)

Peer Mentor, University of Pennsylvania [Step-Ahead Mentorship Program](#)

Prior Service

Editor-in-Chief, [Caltech Undergraduate Research Journal](#)

Volunteer Tutor, [Caltech RISE Tutoring Program](#)

Peer Tutor, Caltech Deans' Office

Student Body Representative, [Caltech Academics and Research Committee](#)