

Department of Biostatistics and Medical Informatics Seminar



Yongyi Guo, PhD

Assistant Professor, Department of Statistics
University of Wisconsin-Madison

<https://sites.google.com/view/yongyi-guo/>

Email: guo98@wisc.edu

Friday, September 19, 2025

12:00-1:00 pm

Morgridge Hall Seminar Room 7560 or

Zoom:

<https://uwmadison.zoom.us/j/99879638765?pwd=wbtxoucEFiIPVCVc9SFbvKB1Av7Xk.1>

Passcode: 343271

Personalizing Digital Health with Reinforcement Learning: Design and Inference in Adaptive Experiments

Abstract: Digital health interventions offer the opportunity to deliver personalized support at scale, but designing and evaluating such interventions raises new statistical and methodological challenges. In this talk, I will describe my work on MiWaves, a mobile health study aimed at reducing cannabis use among emerging adults, where intervention messages were delivered via a just-in-time adaptive intervention (JITAI). To enable personalization, we used a reinforcement learning (RL) framework, where a mixed-effects bandit algorithm adaptively tailored micro-support messages to participants based on their individual histories. While adaptive algorithms improve personalization, they also complicate statistical inference: user trajectories are no longer independent, and classical estimators may fail. I will show how valid inference on key causal quantities, such as causal excursion effects, can still be achieved by leveraging the structure of the bandit algorithm. Specifically, I introduce a statistic–state decomposition of the RL policy, which separates pooled statistics across users from user-specific histories. This decomposition enables asymptotic normality of estimators, construction of hypothesis tests, and a clear connection to the posterior distribution of fixed (pooled) and random (personalized) effects. Beyond providing valid inference, this framework also offers domain scientists interpretable insights into how information is aggregated at both the population and individual levels, allowing them to critique, refine, or redesign policies. This work underscores both the promise and the challenges of applying RL methods in real-world health settings: we can adaptively personalize support while still preserving the ability to draw reliable scientific conclusions.

Bio: Yongyi Guo is an Assistant Professor in the Department of Statistics at the University of Wisconsin–Madison. Her research centers on statistical reinforcement learning, adaptive experimentation, and causal inference, with applications in healthcare and online platforms. She received her Ph.D. in Operations Research and Financial Engineering from Princeton University and was a postdoctoral fellow in the Department of Statistics at Harvard University.



**School of Medicine
and Public Health**

UNIVERSITY OF WISCONSIN-MADISON