Accessible technology and extensive digitization is fueling detailed data collection at the user level. These novel data streams enable decision-making about what and when to deliver that is personalized to users by taking their behaviors and contexts into account. This development is useful in many domains like medicine, mobile health, public policy, and e-commerce. Such **personalized decision-making** involves two fundamental tasks: (1) **estimation and inference** from data when there is no model for the decision's effect on a user and (2) **simulations** when there is a known model for the decision's effect on the user. Here we must overcome the issues facing classical approaches, namely **statistical biases** due to adaptively collected data and **computational bottlenecks** caused by high-dimensional models.

I tackle these statistical and computational challenges by working on theories and methods across *causal inference, reinforcement learning (RL), Bayesian inference, optimization, and high-dimensional statistics*. My research provides strategies for personalized decision-making in three different settings: sequentially adaptive experiments, randomized controlled trials (RCT), and model-based simulations with immediate impact on healthcare. Below I provide a brief summary of my contributions.

(S1) **Sequentially adaptive experiments with RL**: I develop strategies to adjust the biases inherent to adaptive trials in mobile health. My work designs algorithms for **estimating** individual treatment effects (ITE) and assessing the personalization achieved by the RL algorithm used in the trial.[16, 17, 18]

(S2) **Heterogeneity in RCT**: Machine learning (ML) models provide unreliable estimates of ITE for typical RCTs in medicine due to low signal-to-noise ratio. I introduce a data-driven approach for RCTs to reliably **estimate** the heterogeneity in treatment effect.[15]

(S3) **Simulations with computational models**: My work on fast sampling with Markov chain Monte Carlo (MCMC)[2, 8, 3] and uncertainty propagation using compression[12, 13, 7] speed up computationally challenging high-dimensional personalized **simulations**, e.g., in cardiology.

Applications. My research involves collaborations with specific goals for real-world decision-making. Methods from (S1) help **health scientists** design effective mobile apps to assist users in achieving physical fitness and managing stress and addiction. Approaches from (S2) help **doctors** estimate if a drug is very effective for some clinically interpretable patient subgroup. My work in (S3) speeds up non-invasive therapy response simulations in digital twin **heart experiments**. Finally, I helped build forecasting models in a non-profit collaboration for **limited PPE allocation** to needy hospitals during the COVID pandemic.[1]

Besides healthcare, my research is also applicable for decision-making in public policy, social science, and recommender systems in e-commerce. I intend to continue cutting-edge interdisciplinary research and develop principled approaches to solve real-world problems. I now summarize my research in two broad threads: sample-efficient estimation and inference and computationally-efficient simulations.[(i)]

1. Sample-efficient estimation and inference

My first research thread introduces *theory and methods for reliable and personalized recommendations with novel data streams* and tackle several known challenges in sequential experiments and RCTs.

[(i)] Research across Ph.D. co-advised by Prof. Martin Wainwright & Prof. Bin Yu ('15-'21, UC Berkeley), postdoc co-advised by Prof. Susan Murphy & Prof. Devavrat Shah ('21–, Harvard & MIT), and collaborations with Lester Mackey & many others.

Counterfactual inference and personalization. Sequential experiments in mobile health and digital platforms aim to provide personalized treatments to N users over T time points using online personalizing algorithms (PAs), e.g., bandit or RL algorithms. The PA's role is to learn user treatment effects and assign personalized treatment over time. The following projects aim to **infer if (i) the digital interventions and (ii) the PAs used to assign them are effective** at the individual level in sequential experiments. These tasks are **challenging** due to *heterogeneity* across users and time, *biases* due to the PA's sequential adaptivity, the *lack of mechanistic models* for treatment effect, and data's *noisiness*. For (**i**), we provide the **first guarantee for ITE in sequential experiments** using a non-parametric mixed-effect/latent factor model for the potential outcomes and nearest neighbor (NN) estimators. We prove an ITE error of $O(T^{-\frac{1}{2}}+N^{-1})$ for each user at each time[16] and improve it further to $O(T^{-1}+N^{-1})$ with a doubly robust NN for non-adaptive experiments.[17] For (**ii**), we introduce a methodology for assessing the **personalization achieved by an online PA in non-Markovian environments**. We compare the rewards yielded by the online PA to those yielded by a baseline PA, and provide its fundamental decomposition across three practical axes—time, covariates used by the PA (states), and pre-study covariates.[18]

Next, I discuss estimation of heterogeneous treatment effects in RCTs.[15] Via a case study on a drug trial, we highlight the **poor generalization of popular conditional average treatment effect models (CATEm) due to low signal-to-noise ratio**. We introduce a **CATEm-based discovery** procedure for subgroups with heterogeneous treatment effects. The heterogeneity in subgroups discovered in the case study surprisingly **generalized to another independent trial**, i.e., there were statistically significant heterogeneous subgroup treatment effects after adjustment for multiple testing. In another work,[6] we tackle unmeasured confounding in sequential observational studies (unknown treatment policy) via exponential family models. We introduce a method to **estimate unit-specific parameter with a single p-dimensional sample per unit** despite unobserved confounding. We provide a parameter error of $O((s\log k)/p)$ when the parameters are s-sparse combination of k known vectors.

Statistical-computational trade-offs in overparameterized models My research also includes work at the forefront of statistical learning theory. For missing and heterogeneous data, mixture models and expectation-maximization (EM) are the default choice. EM is known to work well for correctly specified models and provide $\Theta(N^{-\frac{1}{2}})$ error in $O(\log N)$ steps with N samples. We provide the **first guarantee that EM behaves poorly with overparametrization**: $\Theta(N^{-\frac{1}{4}})$ error in $O(N^{\frac{1}{2}})$ steps.[11, 10] We show the promise of the usually ignored Newton EM, which achieves the same error in $O(\log N)$ steps.[5] On the other hand, the good generalization of overparameterized models in supervised learning seems to contradict the bias-variance trade-off. We prove that the minimum description length **complexity undergoes a phase change** from $O(d)$ to $O(\log d)$ with d features as d grows larger than N—and hence a good generalization is possible. Moreover, we show that the surprises with bias-variance trade-off occur with poor estimators and **disappear** with regularized estimators.[14]

2. Computationally efficient simulation methodologies

In a digital twin experiment in cardiology, scientists use a computational model for heart to simulate personalized therapy response for arrhythmia.[12] Such simulations often proceed in two stages: (I) estimating smaller-scale (e.g., single-cell) models for each user via Bayesian inference followed by (II) uncertainty propagation to the higher-scale (e.g., tissue or heart) models using complex simulators.

Similar simulations also arise in digital twin technology experiments for system design in energy sector and internet of things. MCMC is a common choice for Bayesian inference in stage I, but its **slow convergence** is a known bottleneck in high dimensions. In cardiology experiments, MCMC is typically run for $\geq 10^6$ iterations (2 CPU weeks) making stage I computationally expensive. In these experiments, stage II is even more expensive as a *single* heart-level simulation takes $\geq$ 5 CPU weeks. A common approach for stage II is to thin/compress the long MCMC chain from stage I. But, standard compression procedures are known to **degrade in high dimensions for integration and testing**. I now summarize my work that addresses these challenges using theory and methods for both MCMC and compression.

Fast compression algorithms in high dimensions. IID sampling and fast mixing MCMC suffer from the poor error rate of $\Theta(N^{-\frac{1}{2}})$ for approximating function expectations with N points. Approaches from quasi-Monte Carlo (including my work[9]), quadrature methods, coresets, and prototype search provide improved error rates for low dimensional distributions, especially uniform on $[0,1]^d$. However, an effective solution for high dimensional distributions with unbounded support (like typical posterior distributions) was unknown. We provide the **first provable and practical strategy, kernel thinning** (KT),[12, 13] a procedure that thins N points to $N^{\frac{1}{2}}$ points with near-optimal $\widetilde{O}_d(N^{-\frac{1}{2}})$ integration error in reproducing kernel Hilbert spaces—a near-quadratic improvement over the $\Theta(N^{-\frac{1}{4}})$ error with $N^{\frac{1}{2}}$ points—for distributions with sub-exponential tails on $\mathbb{R}^d$. We introduce a meta-procedure, *Compress++*[7] that speeds up generic thinning methods and makes KT's runtime $\widetilde{O}(N)$. We build a *compress-then-test* strategy using KT that is **100-200x faster than state of the art** non-parametric hypothesis tests.[4]

Fast MCMC algorithms in high dimensions. A thread of my research establishes the number of iterations (mixing time) needed by popular MCMC algorithms to converge within error δ for a given approximately log-concave distribution in $\mathbb{R}^d$. We provide the **first mixing time bound for state-of-the-art** Metropolized Hamiltonian Monte Carlo (HMC), introduce new MCMC algorithms, and provide several insights en route: (i) We show that **mixing time improves with the usage of gradients** as we move from metropolis random walk, to Langevin algorithms (LA) to HMC.[8, 3] (ii) We prove that **the accept-reject correction step** improves the mixing time of LA from $O(\frac{1}{\delta})$ to $O(\log \frac{1}{\delta})$ (typically used in practice, this correction step is was often ignored in prior theoretical works).[8] (iii) We illustrate that state-of-the-art interior point methods can be suitably adapted to design **state-of-the-art MCMC for constrained sampling**, demonstrating that sampling benefits from faster optimization methods.[2] (iv) Finally, we provide simultaneous improvements to the mixing times of many MCMC algorithms with poor initialization.[3]

3. Future directions

A principled end-to-end pipeline for sequential decision-making is one of my goals, toward which my prior and ongoing research constitute important first steps. Below I outline some future research directions that extend my research trajectory toward this goal. En route I envision many opportunities to build new scientific collaborations across medicine, mobile health, and e-commerce.

Trade-offs between during-study and after-study objectives. The two goals of sequential experiments—**during-study personalization and after-study causal inference**—appear at odds with each other. The latter requires continued exploration, while the former requires exploitation of

the information accumulated thus far. I plan to characterize the Pareto frontier curve for this trade-off in the context of the two goals of sequential experiments as a critical step towards an end-to-end pipeline.

Inference in real-world environments. My work on counterfactual inference[16] in sequential experiments provides a stepping stone to **ITE inference and diagnostics in more complex settings**. It would be interesting to extend it to settings with temporal features (contextual bandits) and delayed effects (Markov decision processes (MDPs)). A two-staged procedure for **contextual bandits and MDPs** that involves parametric modeling refined using NN would be a natural starting point. I am also excited to extend my work on personalization[18] to unit-level PA diagnostics using statistical approaches like bootstrap to more general environments like longitudinal data that make minimal assumptions.

Algorithm design via affordable simulators. Computational stability, autonomous operation, and learning speed make bandit algorithms the default choice of PA. However, this choice of PA is generally misspecified. This misspecification is further exacerbated by the non-stationarity of the environment due to real-time data and sensor constraints. Planning for these issues requires extensive simulations with multiple PA candidates and data-inspired environments with varying misspecifications. I am excited to tackle these challenges by (i) **characterizing the optimal algorithm** that meets the during- and after-study objectives under the above constraints and (ii) **building an affordable simulator** to design a robust algorithm via prototype simulations rather than exhaustive search.

References

[1] N. Altieri, R. L. Barter, J. Duncan, **R. Dwivedi**, K. Kumbier, X. Li, R. Netzorg, B. Park, C. Singh, Y. S. Tan, T. Tang, Y. Wang, C. Zhang, and B. Yu. Curating a COVID-19 Data Repository and Forecasting County-Level Death Counts in the United States. *Harvard Data Science Review*, (Special Issue 1), 2021.

[2] Y. Chen, **R. Dwivedi**, M. J. Wainwright, and B. Yu. Fast MCMC sampling algorithms on polytopes. *Journal of Machine Learning Research*, 19(1), 2018.

[3] Y. Chen, **R. Dwivedi**, M. J. Wainwright, and B. Yu. Fast mixing of Metropolized Hamiltonian Monte Carlo: Benefits of multi-step gradients. *Journal of Machine Learning Research*, 2020.

[4] C. Domingo-Enrich, **R. Dwivedi**, and L. Mackey. Compress then test: Powerful kernel testing in near-linear time. *In conference submission*, 2022.

[5] N. Ho, K. Khamaru, **R. Dwivedi**, M. J. Wainwright, M. I. Jordan, and B. Yu. Instability, computational efficiency and statistical accuracy. *arXiv 2005.11411, accepted with minor revision at JMLR*, 2022+.

[6] A. Shah, **R. Dwivedi**, D. Shah, and G. Wornell. On counterfactual inference with unobserved confounding. *Neurips Workshop 2022, full version to be submitted to Annals of Statistics*.

[7] A. Shetty, **R. Dwivedi**, and L. Mackey. Distribution compression in near-linear time. In *ICLR*, 2022.

[8] **R. Dwivedi**, Y. Chen, M. J. Wainwright, and B. Yu. Log-concave sampling: Metropolis-Hastings algorithms are fast. *Journal of Machine Learning Research*, 20(183):1–42, 2019.

[9] **R. Dwivedi**, O. N. Feldheim, O. Gurel-Gurevich, and A. Ramdas. The power of online thinning in reducing discrepancy. *Probability Theory and Related Fields*, 174(1):103–131, 2019.

[10] **R. Dwivedi**, N. Ho, K. Khamaru, M. Wainwright, M. Jordan, and B. Yu. Sharp analysis of Expectation-Maximization for weakly identifiable models. In *AISTATS*. PMLR, 2020.

[11] **R. Dwivedi**, N. Ho, K. Khamaru, M. J. Wainwright, M. I. Jordan, and B. Yu. Singularity, misspecification and the convergence rate of EM. *The Annals of Statistics*, 48(6), 2020.

[12] **R. Dwivedi** and L. Mackey. Kernel thinning. In *COLT 2021*, Extended abstract. *arXiv, under revision at JMLR*.

[13] **R. Dwivedi** and L. Mackey. Generalized kernel thinning. In *ICLR*, 2022.

[14] **R. Dwivedi**, C. Singh, B. Yu, and M. J. Wainwright. Revisiting minimum description length complexity in overparameterized models. *arXiv 2006.10189, accepted with minor revision at JMLR*, 2022+.

[15] **R. Dwivedi**, Y. S. Tan, B. Park, M. Wei, K. Horgan, D. Madigan, and B. Yu. Stable discovery of interpretable subgroups via calibration in causal studies. *International Statistical Review*, 88:S135–S178, 2020.

[16] **R. Dwivedi**, K. Tian, S. Tomkins, P. Klasnja, S. Murphy, and D. Shah. Counterfactual inference for sequential experiments. *arXiv preprint 2202.06891, to be submitted to Annals of Statistics*, 2022.

[17] **R. Dwivedi**, K. Tian, S. Tomkins, P. Klasnja, S. Murphy, and D. Shah. Doubly robust nearest neighbors for latent factor models. *arxiv preprint 2211.14297, Working paper*, 2022.

[18] **R. Dwivedi**, K. Zhang, P. Chhabaria, P. Klasnja, and S. Murphy. Assessing personalization by a reinforcement learning algorithm. *Working paper*, 2022+.