Pervasive digitization and affordable sensors have enabled personalized data collection at the user level. These novel data streams coupled with machine learning (ML) enable decision-making about what and when to deliver, that is personalized to users, by considering their behaviors and contexts. This development is useful in many domains like medicine, mobile health, public policy, and e-commerce. However, for reliable **personalized decision-making**, we need **user-level predictions along with uncertainty quantification for actionable quantities**, e.g., confidence intervals for user-level treatment effects of a drug in a clinical trial. Learning these quantities involves two fundamental tasks: (1) **Estimation and inference** from data when a good mechanistic model of how the decision affects a user is unavailable, like in medicine or computer systems; and (2) **model-based simulations** when there is a known computational or stochastic model for the decision's effect, e.g., in cardiology or engineering systems. My research builds machine learning solutions for decision-making that address the algorithmic, statistical and computational challenges associated with learning personalized quantities.

Overall, I take a **multi-disciplinary approach** to **data science** and bring together ideas from computer science, electrical engineering, and statistics in **collaboration with domain experts** to develop statistical ML solutions for personalized decision-making in real-world problems.[(i)] My research spans across *algorithms in optimization and random sampling, causal inference, reinforcement learning (RL), Bayesian inference, and high-dimensional statistics*. On the one hand, my background in computer science and electrical engineering is critical in **designing and analyzing efficient algorithms** and methods that work on a range of problems with various scales. On the other hand, my training in statistics helps me **design statistical models** for the problem at hand, reason about the amount of data needed, and **quantify uncertainty** in the predictions from these algorithms. I now provide an overview of the challenges in the two tasks that my research addresses.

**(1) Estimation and inference.** In settings where we do not have a good mechanistic model of how the decision/treatment affects the user, *I design algorithms that provide user-level estimates with an accurate measure of statistical uncertainty, even with a small sample size and are also computationally efficient thereby well suited for big data settings.* My research provides these estimates for data is collected in different ways, including randomized experiments (like clinical trials or A/B testing), observational studies (like with medical records), and sequential experiments on digital platforms (like adaptive trials in mobile health). For personalized inference with such data, conventional ML approaches face multiple **statistical challenges**. In real-world clinical trials, sample size is small by design due to monetary and other risk constraints. With small sample sizes, complex methods overfit while simple methods do not provide a good user-level fit. In observational studies, unknown treatment mechanism (e.g., how the medicine was chosen) introduces non-trivial biases for user-level estimates. In sequential experiments, online adaptive algorithms like bandits enable the personalization of treatments to user behavior and context. However, the feedback from these algorithms is known to render classical approaches, e.g., even least squares with adaptively collected data with linear models, unreliable due to **inaccurate uncertainty estimates**. My research builds provable and practical solutions to tackle these challenges.

**(2) Model-based simulations.** With a known mechanistic model for how a certain decision/treatment affects a user, **computational challenges** arise when learning and quantifying and propagating uncertainty in the estimates via high-dimensional computer simulations. *I provide provable guarantees*

[(i)] Research across Ph.D. in EECS, UC Berkeley ('15-'21) and postdoc in CS & Statistics, Harvard & EECS, MIT (since '21).

*on popular uncertainty quantification methods (Markov chain Monte Carlo (MCMC)) in high dimensions, design new computationally efficient variants, and new compression algorithms that provide provably huge speed ups for uncertainty propagation to expensive downstream tasks.* For example, computational cardiologists want to predict personalized disease progression and therapy response for a given user via non-invasive digital simulations (without running a trial) using multi-scale models. Such simulations often proceed in two stages. First, one estimates smaller-scale (e.g., single-cell) models for a given user along with uncertainty quantification via **random sampling** with MCMC. Second, this **uncertainty is propagated** to the higher-scale (e.g., tissue or heart) models using complex simulators. Both stages are known to be computationally daunting due to challenges (a) with MCMC convergence in high dimensions and (b) the very high dimensionality of the simulators. E.g., to model the effect of calcium signaling dysregulation on heartbeats, MCMC for generating a million cell-model samples might take 2 CPU weeks, while a *single* heart-level simulation might take 4 CPU weeks. Such challenges also arise in system design via complex multi-scale simulations in many engineering applications like aerospace system, autonomous driving, power plants, and the upcoming digital twin technology. My research provides solutions to the computational bottlenecks in both stages.

Here is a brief overview of my research for personalized decision-making in three different settings: sequentially adaptive experiments, randomized controlled trials (RCT), and model-based simulations.

(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** how effective is the treatment at user-level, referred to as individual treatment effects (ITE) and assessing the personalization achieved by the RL algorithm used in the trial.[16, 17, 18]

(S2) **Heterogeneity in RCT**: As noted earlier, ML models provide unreliable estimates of ITE for real-world RCTs in medicine due to limited sample size (and hence low signal-to-noise ratio). I introduce a calibration and stability based discovery procedure to reliably **estimate** the heterogeneity in treatment's effect in various subgroups in the study.[15]

(S3) **Simulations with computational models**: My work builds new bridges between optimization and sampling for design and analysis for fast MCMC methods[2, 8, 3] and discrepancy and compression to design effective uncertainty propagation methods[12, 13, 7] that collectively speed up computations in complex high-dimensional simulation systems.

**Applications and impact.** My research involves collaborations with specific goals for real-world decision-making and has immediate impact in healthcare. 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 personalized 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 engineering systems, recommender systems, and problems in public policy and social sciences. 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: data efficient estimation and inference and computation efficient model-based simulations.

## 1. Data 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.

**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 treatments 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 inference guarantee for user-level effects in sequential experiments**, namely a non-asymptotic error bound and an asymptotic confidence interval. We use the popular ML approach of nearest neighbors to build our estimate and prove the guarantee under mild assumptions and a very flexible statistical model (non-parametric mixed-effect/latent factor model). We prove that our estimate admits a squared error of $O(T^{-\frac{1}{2}}+N^{-1})$ for each user at each time[16]. We improve it further to $O(T^{-1}+N^{-1})$ with a new improved variant of nearest neighbors that we refer to as doubly robust nearest neighbors.[17] We use these estimators in mobile health studies to estimate if the mobile notification was useful in helping users become physically more active.

For (**ii**), i.e., to estimate if the PA is in fact personalizing treatments to the user, 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 in terms of value function difference, 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 ML-based 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. Computation efficient model-based simulations

In many applications in engineering, healthcare, and other domains, complex multi-scale simulations are used for system design. For example, in a digital twin experiment in cardiology, scientists use a computational model for heart to simulate a heartbeat to better understand the disease 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. MCMC is a common choice for Bayesian inference in stage I, but its **slow convergence** is a known bottleneck in high dimensions. For example, in the 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 coresets, quadrature methods, quasi-Monte Carlo (including my work[9]), 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 $\delta$ 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. This goal requires the development of algorithms that take into account multiple objectives, e.g., personalize the decisions to users during the trial (low regret during study) as well as provide a reliable estimate of whether the decisions are working (small confidence intervals after study). Moreover, these algorithms and inference methods need to plan and account for the dynamic nature of users caused by the decisions (effectively a transfer learning task), e.g., in mobile health, the user behavior changes due to repeated mobile notifications. 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 in both engineering and healthcare.

**Speeding up empirical risk minimization (ERM).** Compressed datasets would be directly useful for speeding up ML algorithms for ERM. For example, runtime of kernel-based ERM using gradient descent with $N$ points can be reduces from order $N^3$ to order $N^{3/2}$ if the $N$ data points are compressed to $N^{1/2}$ points. To avoid the larger generalization error caused by naive compression, I am excited to extend the procedures like kernel thinning and compress++[12, 13, 7] to supervised learning settings. Such effective extensions for compressing labeled datasets would also be useful in transfer learning tasks where a compressed dataset for earlier tasks is retained while fine tuning neural networks for future tasks.

**Algorithm design via affordable simulators.** Computational stability, autonomous operation, and learning speed (bias-variance tradeoff) make bandit algorithms the default choice of PA in many experimental studies. However this choice is often misspecified when the underlying environment is not a bandit and can lead to sub-optimal performance. This sub-optimality can be further exacerbated by the non-stationarity of the environment due to real-time data and sensor constraints. However, even characterization of the optimal algorithm in idealized settings is a non-trivial task when dealing with multiple objectives. For example in adaptive trials in healthcare, the two goals can be **during-study personalization and after-study causal inference**. And these goals can appear at odds with each other. Typically inference requires continued exploration, while personalization translates to low regret which in turn requires exploitation of the information accumulated thus far. For a principled algorithm design, first I plan to characterize the Pareto frontier curve for this trade-off when dealing with multiple goals in sequential experiments and design algorithms that can achieve it. To design a robust algorithm that performs well even under various misspecifications issues that arise in practice, typically extensive simulations with multiple PA candidates and data-inspired environments need to be conducted. I plan to make this step computationally feasible by **building an affordable simulator** via prototype simulations rather than exhaustive search.

**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 nearest neighbors would be a natural starting point. Furthermore, I am also excited to extend my work on personalization[18] to unit-level PA diagnostics using statistical approaches like bootstrap to more general RL environments like longitudinal data that make minimal assumptions.

## 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+.