# A Data-driven Optimization Approach to Designing Parsimonious Treatment Guidelines

## BIOGRAPHY

Gian-Gabriel Garcia is an Assistant Professor in the Department of Industrial & Systems Engineering (ISE) at the University of Washington. His research interests are in the design, analysis, and optimization of data-driven frameworks at the intersection of optimization, machine learning, and artificial intelligence as motivated by high-impact problems in medical decision-making, health policy, and healthcare operations. His recent work emphasizes the design of frameworks that consider interpretability, fairness, and robustness, and his work has found application in various disease areas including: chronic diseases, concussion, opioid-related overdose, mental health, and maternal health. Dr. Garcia has received several scholarly recognitions for his scholarship, including the IISE Transactions Best Paper Award in Operations Engineering and Analytics, INFORMS Minority Issues Forum Best Paper Award, INFORMS Bonder Scholarship, and the Society for Medical Decision Making's Lee B. Lusted Prize in Quantitative Methods and Theoretical Developments. He has also been awarded external funding support by the National Institutes for Health, Agency for Healthcare Research and Quality, Department of Defense, and Children's Healthcare of Atlanta.

## ABSTRACT

Clinical treatment guidelines play an important role in managing chronic diseases at the population level by providing standardized recommendations to guide medical decision-making. While many existing guidelines include some degree of patient stratification, these groupings may not be optimized to maximize population-level health outcomes. In contrast, fully personalized treatment policies can theoretically optimize population-level health, but can be impractical to implement at a wide scale due to their complexity and potential technological requirements. Accordingly, we propose the Treatment Guideline Design Problem (TGDP), a new data-driven optimization framework that jointly partitions a population into a fixed number of subgroups and determines an optimal treatment policy for each subgroup. In the TGDP, we model each patient as a Markov Decision Process to capture heterogeneity in disease progression and treatment response. We prove that the TGDP is NP-hard and that the expected population-level utility increases monotonically with the number of stratifications. To address computational intractability, we develop both exact and heuristic algorithms to solve the TGDP. Numerical experiments show that our heuristic methods achieve near-optimal outcomes while scaling efficiently to large populations. We then apply our methods to a case study on hypertension management using a nationally representative U.S. dataset. We find that even a small number of optimized stratifications can yield substantial improvements in population-level outcomes over current clinical guidelines. Overall, our framework provides a principled and practical approach for designing treatment guidelines that balances personalization and implementability, helping clinicians, managers, and health policymakers make advances toward optimal population health outcomes.