DISSERTATION PROPOSAL

H. Satyam Verma

"Smart Decisions in Healthcare: From Transparency to Personalization"

Tuesday, November 25, 2025
10:00am
Tepper 4242

Chapter 1: Unpacking Transparency: Information, Bargaining, and AI-Enabled Search in Healthcare Markets

Healthcare remains one of the largest financial burdens for U.S. consumers, yet prices have historically been opaque. The federal Price Transparency Act aims to improve price and quality transparency by requiring hospitals (providers) and insurers (payers) to disclose negotiated rates for all covered services. The policy is expected to enhance competition and inform consumer choice, though its actual effects remain uncertain. The main challenge is that transparency operates through several connected layers: it can shift bargaining power between payers and providers, affect how payers set prices, and change how consumers choose among providers. This chapter develops a stylized game-theory model to study how price transparency alters equilibrium outcomes in the healthcare market. The framework focuses on a single representative medical service to isolate the core strategic interactions. Payers and providers bargain bilaterally over reimbursement rates, while consumers choose among providers based on price and quality. Transparency enters as an information shock that strengthens payer bargaining leverage and increases the salience of prices in consumer decision-making. AI-enabled information tools are modeled as a complementary mechanism that reduces consumer search costs, making price information more accessible and amplifying the behavioral effects of transparency. To illustrate these mechanisms, the model is paired with numerical what-if analyses using real, publicly available hospital data on negotiated rates. The simulations explore how improvements in transparency, together with easier access to information through AI-assisted search, propagate through bargaining and consumer choice to reshape equilibrium prices, incentives, and welfare.

Chapter 2 — Data-to-Dose: Efficient Synthetic Data Generation with Expert Guidance for Personalized Dosing

Personalized dosing aims to determine medication levels that are optimal for each patient's unique physiological and clinical characteristics. Achieving this objective requires learning reliable relationships between dose, efficacy, and toxicity. In practice, however, data suitable for this task are often scarce relative to the diversity of patient populations, and existing datasets may underrepresent key subgroups defined by genotype, comorbidities, or treatment history. Motivated by this challenge, we develop GenEx, a framework that integrates expert-guided preference learning with generative synthetic data augmentation to learn personalized dosing policies when clinical evidence is limited or unevenly distributed. GenEx represents clinician reasoning through pairwise comparisons of candidate doses, such as "A is more reasonable than B," which update a Gaussian-process surrogate representing the efficacy–toxicity balance. In parallel, a conditional generative model learns to simulate plausible patient–dose–outcome triplets conditioned on clinical covariates. When the surrogate's implied rankings and the generator's outputs align, a trust-gating mechanism selectively injects synthetic samples into regions where observations are sparse but uncertainty is high, expanding informative coverage while maintaining consistency with expert intent. We analyze the learning procedure and establish sublinear regret bounds together with bounded bias under trust-gated synthetic augmentation. We evaluate the framework in two complementary settings. First, we study learning dynamics on simulated datasets that emulate realistic clinical dosing environments. These datasets are constructed with clinician input to capture variability in pharmacokinetic response and feedback behavior, enabling controlled analysis of data efficiency, bias, and regret under different scarcity regimes. Second, we extend the evaluation to real clinical trajectories from kidney-transplant patients, reformulated into independent one-shot scenarios to test generalization to factual data. These studies examine performance relative to purely generative and purely expert-based baselines and assess how the trust-gated mechanism supports safe and data-efficient learning under scarcity. Collectively, this chapter develops and analyzes a principled approach for integrating expert feedback and generative modeling in personalized dosing, supported by theoretical regret guarantees.

Chapter 3 — Safe Sequences: AI Co-pilots for Clinical Decision-Making

Building on the principles of Chapter 2, this chapter extends our framework to sequential decision-making, where observations unfold over time in the form of patient trajectories. Available trajectories are often limited, incomplete, irregular, or missing for certain profiles, making it difficult to directly learn personalized dosing policies from observational data. This motivates a more adaptive learning approach that can operate effectively even under partial information. To address this challenge, we introduce a generative modeling and human-feedback subsystem that enables the algorithm to reason and learn under scarcity. Unlike Chapter 2, we do not rely solely on preference feedback but experiment with different forms of expert input to better capture the sequential nature of decisions. The generative model constructs short-horizon world representations that infer plausible near-term outcomes from limited histories, while human feedback provides expert guidance to calibrate and refine these representations. In this framework, the expert supplements the lack of reliable data while also contributing domain knowledge—clinical reasoning, safety heuristics, and contextual insights—that are not easily captured in the data itself. The two components together provide a structured way to learn effectively when information is incomplete or uncertain.

We further examine this second role of expert input in contexts with richer sequential information, where it continues to shape model behavior by articulating principles that remain implicit even when more data are available. This progression allows the system to evolve from using expert feedback primarily to address data scarcity toward leveraging it as an alignment signal that refines model reasoning and interpretability. Finally, we explore a bidirectional learning dynamic in which both the algorithm and the expert improve through interaction. The algorithm continues to generate uncertainty-aware dosing recommendations, while the expert learns from these recommendations—refining clinical intuition, recognizing counterfactual possibilities, and improving decision consistency over time. In turn, the algorithm reduces its uncertainty in areas where expert feedback remains strong, creating a complementary exchange that promotes mutual learning. This formulation positions the generative model as an AI co-pilot for clinical decision-making, supporting an adaptive, interpretable, and continuously improving human–AI collaboration.

Proposed Committee: Sridhar Tayur (Chair), Holly Wiberg (Heinz College of Information Systems and Public Policy), Shixiang (Woody) Zhu (Heinz College of Information Systems and Public Policy, and Tinglong Dai (Johns Hopkins Carey Business School)