

H. Satyam Verma

PhD Candidate in Operations Management
Tepper School of Business · Carnegie Mellon University

hverma@andrew.cmu.edu · +1 (412) 954-7406 · [Google Scholar Profile](#) · hsatyamverma.com

Research Interests

Healthcare Operations · Sequential Decision-Making Under Uncertainty · Optimization and Machine Learning · Healthcare Markets and Pricing · Business Analytics

Education

- 2022 – Present **Ph.D., Operations Management**
Tepper School of Business, Carnegie Mellon University, Pittsburgh, PA
Dissertation: *Smart Decisions in Healthcare: From Transparency to Personalization*
Advisor: Sridhar R. Tayur
Committee: Sridhar R. Tayur (Chair), Tinglong Dai, Holly Wiberg, Shixiang (Woody) Zhu
Expected Graduation: Spring 2027
- 2019 – 2021 **M.Sc., Operations Research**
Indian Institute of Technology (IIT) Bombay, Mumbai, India
- 2015 – 2018 **B.Sc. (Hons.), Mathematics**
Kirori Mal College, University of Delhi, Delhi, India
First Division with Honors

Working Papers

- [1] **Verma, H. S.**, Wiberg, H., Zhu, S., and Tayur, S. R. “GenEx: Generative Augmentation of Scarce Data through Expert Guidance for Personalized Dosing.”
Manufacturing & Service Operations Management (M&SOM) — **Under review (Major Revision)**, 2026.
★ **Runner-Up**, INFORMS HAS Best Student Paper Competition, 2025.
Abstract: Personalized treatment optimization is constrained by data scarcity, especially for rare patient profiles, limiting individualized and safe dosing recommendations. We introduce GenEx, a hybrid framework integrating preference-based Bayesian optimization with a generative model fine-tuned via expert feedback. A trust-gated component proposes synthetic data in uncertain regions, while pairwise expert rankings guide updates to both surrogate and generator. Our main theoretical contribution is a bias-exposed confidence bound separating statistical uncertainty from synthetic-label corruption; building on it, we derive an explicit $\tilde{O}(\sqrt{T})$ regret rate, with bias-decay conditions derived rather than assumed. Experiments across $m \in \{1, \dots, 5\}$ drug dimensions and non-i.i.d. covariate distributions show that GenEx significantly outperforms expert-only and synthetic-only baselines. A retrospective tacrolimus study in kidney transplant recipients demonstrates a mean 25.9% improvement in treatment utility, with gains retained under noisy feedback. A core novelty is reconceptualizing the surrogate as a co-designed component rather than a passive predictor: by jointly refining it via expert preferences and

uncertainty-directed synthetic augmentation, the framework actively shapes its own training distribution. GenEx is humans augmenting AI as much as it is AI augmenting humans. Available at [SSRN](#).

- [2] **Verma, H. S.**, Agrawal, D., Hosseininasab, A., and Tayur, S. R. “Price Transparency and Hospital–Insurer Bargaining: A Three-Channel Model with Evidence from Negotiated Rate Data.” *Working paper*.

Refereed Publications

- [1] Dhingra, R., **Verma, H. S.**, Reiffers-Masson, A., and Kavitha, V. “Multi-page Menu Recommendation in Cascade Model with Externalities.” *60th IEEE Conference on Decision and Control (CDC)*, pp. 6614–6619, 2021. [10.1109/CDC45484.2021.9682971](#)
- [2] Meisheri, H. *et al.* (including **Verma, H. S.**) “Scalable Multi-Product Inventory Control with Lead Time Constraints Using Reinforcement Learning.” *Neural Computing and Applications*, vol. 34(3), pp. 1735–1757, 2022. [10.1007/s00521-021-06129-w](#)

Presentations

- **INFORMS Healthcare Conference, 2026.**
GenEx: Generative Augmentation of Scarce Data through Expert Guidance for Personalized Dosing
Price Transparency and Hospital–Insurer Bargaining: A Three-Channel Model with Evidence from Negotiated Rate Data
- **POMS Conference, 2026.**
Price Transparency and Hospital–Insurer Bargaining: A Three-Channel Model with Evidence from Negotiated Rate Data
- **NeurIPS Workshop on ML × OR (MLxOR), 2025. (Poster)**
Data to Dose: Efficient Synthetic Data Generation with Expert Guidance for Personalized Dosing
- **INFORMS Annual Meeting, Atlanta, Oct 2025.**
Data to Dose: Efficient Synthetic Data Generation with Expert Guidance for Personalized Dosing
- **INFORMS Workshop on Data Science, Oct 2025.**
Data to Dose: Efficient Synthetic Data Generation with Expert Guidance for Personalized Dosing
- **Cornell ORIE Young Researchers Workshop, Oct 2025.**
Data to Dose: Efficient Synthetic Data Generation with Expert Guidance for Personalized Dosing
- **POMS Conference, May 2025.**
Improved Personalized Drug Dosing through Data Generation and Expert Feedback
- **IEEE Conference on Decision and Control (CDC), Dec 2021.**
Multi-page Menu Recommendation in Cascade Model with Externalities

Teaching Experience

- 2024 – 2026 **Teaching Assistant** — Carnegie Mellon University
- *Operations Management*: MBA & Undergrad Core (Fall 2024, Spring/Fall 2025)
 - *Operations and Supply Chain Analytics*: MBA/Undergrad Electives (Spring 2024, Spring 2026)
 - *Supply Chain Management*: MBA/Undergrad Electives (Spring 2026)
 - *Service Management: Go-to-Market Strategy*: MBA Elective (Fall 2025)
- 2020 – 2021 **Teaching Assistant** — Indian Institute of Technology (IIT) Bombay
- *Probability and Stochastic Models I* (Fall 2020)
 - *Probability and Stochastic Processes II* (Spring 2021)

Teaching Interests: Operations Management (core MBA/undergrad), Healthcare Operations, Stochastic Models, Machine Learning for Business, Data Analytics & Optimization.

Honors and Awards

- 2025 **Runner-Up, INFORMS HAS Best Student Paper Competition**
- 2024 **Generative AI Fellowship**
Center for Intelligent Business, Tepper School of Business, CMU
- 2022 **William Larimer Mellon Fellowship**
Tepper School of Business, Carnegie Mellon University
- 2018 **All-India Rank 8 — Joint Admission to M.Sc. Programs (JAM) Mathematics**

Professional Service

- **Webmaster and Marketing Chair**, YinzOR Student Conference, CMU INFORMS Student Chapter, 2023–2026.
- **Webmaster**, CMU INFORMS Student Chapter, 2024–2025.

Prior Industry Experience

- 2021 – 2022 **Data Scientist**, Delhivery
Optimization team. Built and tested large-scale optimization models for logistics network operations. Developed a totally-unimodular LP formulation for order allocation.
- 2020 **Research Intern**, Tata Consultancy Services (TCS) Research
Developed a reinforcement learning framework for scalable multi-product inventory control with lead time constraints. This work resulted in a publication in *Neural Computing and Applications*.

Technical Skills

Programming	Python (PyTorch, GPyTorch, scikit-learn, pandas, NumPy), R, MATLAB
Optimization	Gurobi, CPLEX, CVXPY, AMPL
Typesetting	$\text{\LaTeX}$, Overleaf, Git / GitHub

References

Sridhar R. Tayur (Primary Advisor)
University Professor of Operations Management
Tepper School of Business, CMU
tayur@cmu.edu

Holly Wiberg
Assistant Professor of Operations Research
Heinz College, CMU
hwiberg@andrew.cmu.edu

Tinglong Dai
Bernard T. Ferrari Professor of Business
Johns Hopkins Carey Business School
dai@jhu.edu

Shixiang (Woody) Zhu
Assistant Professor of Data Analytics
Heinz College, CMU
shixianz@andrew.cmu.edu