RISHAV SEN

Nashville, TN · +1-615-955-7951 · rishav.s3n@gmail.com · rishavsen1.github.io · LinkedIn · GitHub · Google Scholar

TECHNICAL SKILLS

Control, Planning & Time-Series: Rolling-horizon MPC / MC-MPC, MCTS, scenario generation, time-series forecasting and foundation models.
Deep Learning: DDPG, PPO, GAN, VAE, action masking, policy guidance, RLlib.
Optimization & ML: MILP (CPLEX, Gurobi), Neuro-symbolic AI, physics-informed forecasting, digital twins (SUMO, GridLAB-D).
Generative & Agentic AI: LLM APIs (Gemini, Claude, Codex), Claude Code, LLM explainability, LLM-orchestrated optimization, RAG.
Tools: PyTorch, Scikit-learn, Python, C++, SQL, PySpark, AWS (S3/SageMaker), Docker, Git.
Visualization & Monitoring: Plotly, Dash, Grafana, D3.js.

EXPERIENCE

Vanderbilt University, SCOPE Lab | Graduate Research Assistant *Nashville, TN · Aug 2021 – Present*

Applied machine learning and decision systems for energy and mobility under uncertainty; developed through field validation.

- **Online decision-making at scale:** Developed rolling-horizon controllers (MC-MPC, MCTS) and safe reinforcement learning (DDPG) within 1.6% of the optimal MILP oracle while maintaining 100% constraint satisfaction across twelve months of live data, with fifteen-minute decision epochs, sub-second inference, and scaling beyond 200 vehicles.
- **Neuro-symbolic AI:** Embedded a learned value-function surrogate into a MILP using McCormick envelopes and SOS2 piecewise-linear constraints, bridging short-horizon control with multi-day demand objectives and delivering 3.5% lower peak power with zero state-of-charge violations.
- **Foundation-model forecasting:** Built day-ahead building-load forecasting with XGBoost and the zero-shot Salesforce Moirai time-series foundation model, calibrated against DOE EnergyPlus simulations, ASHRAE occupancy schedules, and NASA POWER weather data.
- **Safe reinforcement learning:** Designed policy guidance for DDPG with domain-specific action masking to enforce hard operational constraints (state-of-charge bounds, charger limits) during training, achieving zero constraint violations on field-scale building data.
- **Multi-agent design and LLM explainability:** Formulated a strategy-proof negotiation mechanism as a Semi-Markov Decision Process that raised V2B participation by 8%, and began developing an LLM explainability layer that translates grid constraints and time-varying pricing into natural-language cost and flexibility trade-offs for users.

OpenMinds Inc. | NextGen Leader, Energy & Climate Fellowship *Remote, United States · May 2025 – Jan 2026*

- **Scenario modeling and decision support:** Quantified $100M in projected ten-year ERCOT savings for battery-enabled HVAC deployment using XGBoost and Gaussian Mixture Models with Carrier and CenterPoint Energy, and built a decision-support dashboard (Grafana, Dash) translating model outputs into operational insight for energy planners.

Nissan North America Inc. | Research Intern, Connected Energy & Data Analytics *Santa Clara, CA · Jun – Sep 2023*

- **Cost optimization and behavior modeling:** Engineered a MILP-based V2B power-flow optimization pipeline that reduced peak-demand charges by 15% on real commercial building loads, and modeled electric-vehicle user behavior from more than 500,000 telematics records (ZINB, Beta regression, Weibull AFT) for integration into a digital twin used for policy validation.
- **Digital twin:** Extended the V2B simulation framework with vehicle-to-home capability for residential grid-services modeling.

EDUCATION

Ph.D., Electrical & Computer Engineering, Vanderbilt University *Nashville, TN · Aug 2021 – Expected Jul 2026*

AWARDS & RECOGNITION

- Best Paper Award Finalist, AAMAS 2025 (top 5 of 900-plus submissions, multi-agent systems).
- Second Place, Poster Competition, INFORMS 2025 Annual Meeting (top 2 of 300-plus submissions).
- Best PhD Forum Poster Award, ACM/IEEE ICCPS 2025.
- Russell G. Hamilton Scholar, Vanderbilt University (merit fellowship awarded to under 2% of graduates).

Key Projects & Publications

Neuro-Symbolic V2B Controller (P-V2B) *MPC + RL · ICCPS 2026*
Embedded a learned value-function into MPC, bridging short-horizon dispatch with long-horizon, multi-day objectives while capturing recurring behavior; achieved 3.5% lower power peaks and 2.5% cost savings with 100% constraint satisfaction over twelve months of live EV + building data.

CONSENT: Strategy-Proof Negotiation Framework *Mechanism Design + MPC · AAMAS 2026*
Formulated V2B negotiation as an SMDP with MPC scheduling and a discrete-choice user model, providing operational guarantees on strategy-proofness; cut user charging costs 22% and building costs 3.5% versus smart-charging baselines on twenty-two months of Nissan field data.

Agentic AutoML for RL *LLM agents + RL*
Built an LLM agent for autonomous RL architecture and hyperparameter search across three benchmarks; its two-phase workflow (LLM discovery then Optuna tuning) raised mean Acrobot reward 5x over the Optuna-only baseline (+260 vs +51), matching tuned PPO-Zoo performance.

MoveOD: Synthetic Demand and Mobility Pipeline *Statistics + MILP · ICCPS 2026 · INFORMS 2025 (2nd Place)*
First open-source pipeline generating commuter origin-destination (OD) matrices for every U.S. county; fuses and statistically calibrates Census, building-load, and traffic data at county scale using MILP and PySpark, producing reproducible demand inputs for transit and energy simulation.