Shunyu Wu

Ph.D. Candidate —Decision-Focused Learning & Optimization

shunyuwu@sjtu.edu.cn shuanyuwu.github.io Google Scholar +86 137 7667 7189

Research Summary

I build learning-to-optimize systems that are stable, safety-first, and deployable. My work centers on predict-then-optimize (PtO) and decision-focused learning (DFL) with learned objectives/constraints. I design deployment-aligned training (dual-gap views), light-weight proximal updates with *inexact* surrogates, and graph/physics-aware models that connect to LP/MILP online optimization. These pipelines run in production at thermal power stations and water utilities with configuration-driven services, feasible-set tightening, and two-layer protections.

Education

Apr. 2022–Mar. 2026(Expected)	**Ph.D.** in Control Science and Engineering	Shanghai Jiao Tong University
Sep. 2019–Mar. 2022	**M.Sc.** in Control Engineering	Shanghai Jiao Tong University
Sep. 2015–Jun. 2019	**B.Eng.** in Automation	Nanjing University of Science and Technology

Research Agenda

My future research focuses on building next-generation decision systems that are both theoretically sound and practically deployable. I aim to extend my work in decision-focused learning along three directions:

- **Unsupervised Neural Solvers** Develop label-free solvers for constrained optimization that remain differentiable and can be trained end-to-end. These solvers will enable rapid problem-solving without requiring ground-truth decisions.
- **Differentiable Decision Pipelines** Design modular optimization pipelines where objectives and constraints are embedded as trainable blocks. This will make decision models easier to train, verify, and adapt to new tasks.
- **Large Models for Decision Intelligence** Explore the use of large language models (LLMs) for extracting constraints and guiding decision processes, combined with differentiable decision pipelines for closed-loop adaptation in real environments.

The long-term vision is to apply these methods in social, energy, and scientific domains, creating AI systems that deliver reliable, safe, and explainable decisions at scale.

Research Program —Three Highlights

1. Dynamics of End-to-End DFL Training **Lead** 2024–2026

- Concept: Analyzed unstable training dynamics of DFL when the argmin map is piecewise-constant. Borrowed ideas from computational physics (energy landscapes, damping, relaxation) to characterize instability and guide stabilization.
- Contributions: Proposed variational free-energy regularization across separation regimes and proximal descent with *inexact* surrogate gradients.
- Key works: *Crossing the Separation Point: Stabilizing DFL with Variational Free-Energy* (NeurIPS under review); *Proximal Descent for Stable Decision-Focused Training with Inexact Surrogate Gradients* (NeurIPS under review).

2. DFL for Constrained Optimization —Regret and Unsupervised Solvers **Lead** 2024–2026

- Concept: Extended DFL to handle learned *constraints*, beyond objective-only uncertainty. Developed dual-gap regret objectives and explored unsupervised neural solvers that generate feasible decisions without labels.

- Contributions: Introduced parametric dual-gap regret as a general principle for constrained PtO; designed neural "projection-to-envelope" modules for feasibility-preserving optimization.
- Key works: *Beyond Objective Parameters: Parametric Dual-Gap Regret* (ICLR planned); *Projection-to-Envelopes Solver for Constrained Optimization* (in preparation).

3. DFL Pipelines Coupled with Physical Systems **Lead/Co-lead** 2019–Present

- Concept: Coupled graph/physics-informed surrogates with LP/MILP optimizers to build end-to-end deployable pipelines for water and energy systems. Emphasized feasible-set tightening, two-layer protections, and configuration-driven services.
- Contributions: Developed PtO pipelines for sensor placement in water networks and water–energy scheduling. Pioneered end-to-end training of graph regret surrogates directly tied to anomaly coverage and cost.
- Key works: *End-to-End Stochastic Predict-then-Optimize for Water-Energy Scheduling* (IEEE TSG, 2025); *Planning the Invisible: End-to-End Sensor Deployment for Anomaly-Aware WDNs* (AAAI under review); *CritiCoder: End-to-End Uncertain Regression for WDN* (IEEE TCSS, 2023).

Field Deployments and Collaborations

Thermal Power Stations —Closed-loop Efficiency Optimization **Project Lead** 2025.05–Present

- Context: Deployed with a national energy provider across 330 MW (2 units) and 660 MW (3 units) coal-fired units, using 200–300 sensors with 10 s sampling.
- Problem: Efficiency tuning under varying coal quality and load disturbances, subject to strict NOx and reheater safety constraints.
- Solution: Integrated SVR-based surrogates with MILP optimization and two-layer protection; configuration-driven templates allowed rapid cross-site transfer.
- Impact: Efficiency **+1.6–2.0%**, coal use **-1–3 g/kWh**, with **stable 3+ month operation** since May 2025.
- Role: Led technical route and safety strategy; coordinated **SJTU (3–4)** with **XATU/NUC (5–7)** teams in remote collaboration.

Thermal Power Stations —Intelligent Soot-blowing **Co-Lead** 2024–Present

- Context: Targeted slagging/fouling in rapidly varying load regimes.
- Solution: Developed cleaning-factor models and constrained optimization for soot-blowing schedule (mutual exclusion, windows, min-interval).
- Impact: Reduced flue-gas temp volatility by **25%**, cut steam consumption **10%**, improved transparency of operation.
- Role: Designed models and indices, interfaced with efficiency loop, managed deployment and operator training.

Urban Water Utilities —Pump Scheduling and Anomaly Control **Algorithm Lead** 2019–2022

- Context: Shanghai demonstration zone "One Plant–Two Stations," connected with SCADA system.
- Solution: Predict–optimize pipeline: demand forecasting + pressure/flow constrained MILP scheduling; hot updates and rollback mechanisms integrated into SCADA.
- Impact: Peak energy **-10%**, pump on/off **-20%**, pressure compliance **+3%**.
- Role: Led design of end-to-end model, deployment acceptance tests, and operator training manuals.

Urban Drainage —Resilience and Anomaly Management **Co-Algorithm Lead** 2022–2024

- Context: Combined sewer overflow control in Shanghai, facing unexpected inflow and sensor failures.
- Solution: Robust forecasting + redundancy-based thresholds with voting; template library of anomaly strategies integrated with SCADA.
- Impact: Overflow reduced by **≈21%** in 5-year rainfall simulations; operational risk significantly lowered.
- Role: Co-led anomaly-strategy design, integration testing, and knowledge base creation.

Publications —First-Author (selected)

- **Wu, S.**, Wang, J., Xu, H., *et al.*. End-to-End Stochastic Predict-then-Optimize for Cost-efficient Water-Energy Resource Scheduling. *IEEE Transactions on Smart Grid*, accepted, to appear 2025
- **Wu, S.**, Wang, J., Xu, H., *et al.*. Knowledge-based Bi-correction for Lag-free Daily Urban Water Demand Forecasting. *Expert Systems with Applications*, 2024, 255: 124508
- **Wu, S.**, Wang, J., Xu, H., Zhao, S., Xu, J.. CritiCoder: End-to-End Uncertain Regression for Macroscopic Pressure Models in WDN. *IEEE Trans. on Computational Social Systems*, 2023, 11(2): 2222–2233

Publications —Co-authored (selected)

- Xu, J., **Wu, S.**, *et al.*. Reinforcement learning controller design for discrete-time-constrained nonlinear systems with weight initialization. *IEEE Transactions on Systems, Man, and Cybernetics: Systems*, 2024, 54(4): 2368–2378
- Zhong, Y., **Wu, S.**, *et al.*. Prediction of energy consumption in hot rolling roughing based on TDADE. *IEEE Transactions on Automation Science and Engineering*, 2023, 21(1): 555–568
- Xu, J., **Wu, S.**, *et al.*. Adaptive dynamic programming for optimal control of discrete-time nonlinear systems. *IEEE Transactions on Systems, Man, and Cybernetics: Systems*, 2023, 54(3): 1489–1501
- Zhong, Y., **Wu, S.**, *et al.*. A series-concatenation hybrid prediction model of energy consumption in roughing process. *IEEE Transactions on Automation Science and Engineering*, 2023, 21(3): 4585–4598
- Kan, T., Wang, J., **Wu, S.**. The CEEMDAN-ps-gcGRU Model on Water Pressure Prediction With Strong Irregularity. *IFAC-PapersOnLine*, 2023, 56(2): 7108–7113

Preprints & Under Review

- **Wu, S.**, Wang, J., Geng, H.. Crossing the Separation Point: Stabilizing Decision-Focused Learning with Variational Free-Energy. *NeurIPS (under review)*, 2025
- **Wu, S.**, Wang, J., Rao, J., Xiao, X., Cai, H.. Proximal Descent for Stable Decision-Focused Training with Inexact Surrogate Gradients. *NeurIPS (under review)*, 2025
- **Wu, S.**, Wang, J., Rao, J., Wang, H., Xiao, X.. Planning the Invisible: End-to-End Sensor Deployment Regret Minimization for Anomaly-Aware Water Networks. *AAAI (under review)*, 2026
- **Wu, S.**, Wang, J.. Beyond Objective Parameters: Parametric Dual-Gap Regret for Predict-then-Optimize under Unknown Constraint Parameters. *ICLR (planned; in revision)*, 2026
- **Wu, S.**. Projection-to-Envelopes Solver for Constrained Optimization. *in preparation*, ( )

Awards

2023	ABB Intelligent Technology Innovation Competition—National 1st Prize	**1/194**
2023	SJTU *Excellent Graduate Scholarship*	6/208
2023	SJTU *Merit Student*	Top 3%
2022	SJTU *Outstanding Graduate*	Top 15%
2018	Jiangsu Province *Advanced Class Collective*	Top 10%

Academic Service

2021–Present	**Conference reviewing**	ICLR, AAAI, CDC, IFAC
	Learning-to-optimize, decision-focused pipelines, control/optimization under uncertainty.	
2021–Present	**Journal reviewing**	IEEE TSMC (Systems), TASE, TIV
	Methodological soundness, reproducibility, deployability, and clear reporting standards.	
Ongoing	**Industry collaboration**	Utilities & energy
	Acceptance tests, A/B, gray rollout, rollback, post-mortems—bridging prototypes to production.	

Teaching & Mentoring

2023–2024 **Teaching Assistant**, Advanced Process Control — SJTU
Lectures, assignments, recitations.

2020–Present **Mentoring & Supervision** — SJTU / XATU / NUC
Co-mentored 1 Ph.D. and 13 M.Sc.; research pipelines and engineering replication.

Skills

- **Languages**: Python, MATLAB, C
- **ML/Modeling**: PyTorch, scikit-learn, XGBoost/LightGBM, NumPy/Pandas, PyTorch Geometric
- **Optimization**: LP/MILP, robust/stochastic, bi-level/bi-parametric, setpoint optimization; Gurobi, Pyomo
- **Systems**: configuration-driven services, decoupled pipelines, hot updates, gray release; EPANET/Toolkit, SCADA, Redis, Docker, Git, Linux