

Felix Parker

flixpar@gmail.com felixparker.com github.com/flixpar linkedin.com/in/felix-parker

Experience

Postdoctoral Researcher, Center for Systems Science and Engineering, Johns Hopkins University Feb. 2026 – Present

- Developing methods for training LLMs to reason explicitly about numerical time series: built a residual VQ-VAE time series tokenizer, then post-trained an LLM with SFT followed by GRPO-based RL with verifiable rewards, consistently improving over supervised baselines on tasks requiring contextual and multi-step temporal reasoning (under review).
- Building agentic LLM systems that perform expert-level analysis of clinical time series with textual and tabular context, applied to clinical problems at the Johns Hopkins Hospital.
- Also training SimplexRL, an RL-based foundation model for pivot selection in the simplex algorithm, and building a goal-oriented LLM recommender system that helps patients with hypertension adhere to the DASH diet.

Ph.D. Researcher, Center for Systems Science and Engineering, Johns Hopkins University Sept. 2020 – Feb. 2026

- Designed and trained TsLLM, a multimodal LLM that natively perceives and generates time series interleaved with text (ICML 2026): pretrained a patch-based β -VAE time series encoder–decoder on 50M sequences spanning 88 domains, and fused it with a pretrained 7B LLM through learned cross-attention adapter layers.
- Built the full TsLLM training pipeline: curated a 27B-token multimodal corpus with large-scale synthetic data generation, and trained in three stages on H100 GPUs; outperformed frontier LLMs and specialized models on ECG question answering and contextual forecasting, without degrading text-only performance.
- Developed MedTsLLM, an LLM framework for multimodal analysis of physiological signals (semantic segmentation, boundary detection, and anomaly detection) that outperformed state-of-the-art deep learning and clinical baselines on ECG and respiratory waveforms (MLHC 2024).
- Created robust mixed-integer optimization models for hospital capacity management and patient/resource redistribution during demand surges (*Health Care Management Science*); retrospective COVID-19 case studies of optimized redistribution showed $\geq 85\%$ reductions in required surge capacity.
- Built and deployed an interactive decision-support dashboard that the Johns Hopkins Health System used for capacity management during the COVID-19 pandemic; work supported by a Rockefeller Foundation grant.

Undergraduate Researcher, Malone Center for Engineering in Healthcare, Johns Hopkins University 2018 – 2020

- Developed deep learning systems for medical imaging and video: sickle-cell retinopathy detection from retinal fundus photographs (*JAMA Ophthalmology*), laryngeal lesion segmentation in endoscopic video, and surgical phase recognition in cataract surgery (*JAMA Network Open*).

Education

Johns Hopkins University – Ph.D., Systems Engineering (Advisor: Kimia Ghobadi) Feb. 2026

Johns Hopkins University – B.S., Computer Science; Applied Mathematics and Statistics May 2020

Awards: ICML Gold Reviewer Award (2026); Gordon Croft Endowed Fellowship (2020–21)

Selected Publications

Full list: [Google Scholar](#)

- **F. Parker**, N. Chan, C. Zhang, K. Ghobadi. TsLLM: Augmenting LLMs for General Time Series Understanding and Prediction. *ICML*, 2026.
- **F. Parker**, N. Chan, C. Zhang, K. Ghobadi. Eliciting Chain-of-Thought Reasoning for Time Series Analysis using Reinforcement Learning. Under review, 2025.
- H. Bui, **F. Parker**, K. Ghobadi, A. Liu. Q-Learning with Shift-Aware Upper Confidence Bound in Non-Stationary Reinforcement Learning. *AISTATS*, 2026.
- N. Chan*, **F. Parker***, W. Bennett, T. Wu, M.Y. Jia, J. Fackler, K. Ghobadi. MedTsLLM: Leveraging LLMs for Multimodal Medical Time Series Analysis. *Machine Learning for Healthcare (MLHC)*, 2024. (*equal contribution)
- **F. Parker**, F. Ganjkanloo, D.A. Martínez, K. Ghobadi. Optimal Hospital Capacity Management During Demand Surges. *Health Care Management Science*, 2026.

Skills

Machine learning: LLM post-training and fine-tuning (SFT, RL with verifiable rewards, GRPO), LLM agents, synthetic data generation, evals and benchmarking, multimodal model design, time series modeling, deep RL, computer vision

Engineering: Python, Julia, C++, JavaScript; PyTorch, JAX, Hugging Face Transformers; multi-GPU model training, data pipelines

Optimization: mixed-integer, convex, and robust optimization; Gurobi, JuMP