Jiaqi (Charles) Shang

jqshang@cs.toronto.edu | jqshang.github.io | 647-895-1571 | Google Scholar | charlesshang11 | jqshang

EDUCATION

University of Toronto, *Toronto, Ontario, Canada* Sept. 2025 - Dec. 2026 (Expected)
Master of Science in Applied Computing

- GPA: 4.0/4.0 (All A+)
- **Coursework:** Computational Imaging, Machine Learning, Causal Inference, Computational Models of Semantic Change.

University of Waterloo, *Waterloo, Ontario, Canada* Sept. 2020 - Apr. 2025
Bachelor of Mathematics in Computer Science & Statistics, Minor in Computational Mathematics

- GPA: 89.51%; Excellent Standing and Term Distinction in every term.
- **Coursework:** Artificial Intelligence, Machine Learning, Algorithms and Data Structures, Computational Inference.

Research Affiliations

University of California, Riverside, *Riverside, California, USA* Apr. 2026 - Present
Visiting Graduate Student, Department of Mechanical Engineering

California State University, Fullerton, *Fullerton, California, USA* Jan. 2026 - Present
Volunteer Research Assistant, Department of Engineering

PUBLICATION

Predict by Pointing: A Task-Agnostic Output Head for Tabular In-Context Learning
Qiuhao Zeng, Jerry Huang, Younwoo Choi, **Jiaqi Shang**, Rahul G Krishnan
Submitted to ICLR 2027, under review

Learning to Represent Social Information in Multi-Agent Reinforcement Learning for Self-Organized Systems
Jiaqi Shang, Evan Grzeszkiewicz-O'Neill, Hao Ji, Bingling Huang
IDETC-CIE 2026, oral presentation

ERD: Extended RAW-Diffusion Framework for De-rendering sRGB Images
Jiaqi Shang, *Yifan Qu*, Jianbo Qi
ISPRS 2026, oral presentation

Basis Transformers for Multi-Task Tabular Regression
Wei Min Loh, **Jiaqi Shang**, Pascal Poupart
Preprint

RESEARCH EXPERIENCE

Student Research Assistant, *Remote* Jan. 2026 - present
Supervisor: Bingling Huang, California State University, Fullerton, California, USA

- Built a multi-agent assembly task environment in **Python** (**Pygame**, **Pymunk**) with rigid-body physics, sensor-based observations, and shaped reward functions to support cooperative **multi-agent reinforcement learning (MARL)** training.
- Designed three **social learning** information encoding schemes for MARL agents, enabling permutation-aware reasoning over identifiable teammates.
- Analyzed emergent role specialization via **PCA** and **K-means clustering** on agents' force and torque contributions, revealing a hierarchical role structure correlated with team performance.
- Building a modular 2D fish-swimming **RL** environment in **Python** with modular flow fields, body-to-world transforms, and composable forces.
- Training agents with **PyTorch** to navigate time-varying vortex flows, design reward functions to discover energy-optimal routes that exploit vorticity-induced forces.

Machine Learning Researcher (Part-time), *Orka, Remote* Oct. 2025 - Sept. 2026

- Built a **GAN**-based **speech enhancement** system pairing a **causal transformer-based U-Net** generator with Multi-Period and Multi-Scale discriminators for **adversarial training**.
- Implemented a composite training objective combining multi-resolution STFT loss, adversarial loss, and feature-matching loss to improve perceptual quality over magnitude-only baselines.

- Developed end-to-end training and inference pipelines on the VoiceBank-DEMAND and DNS Challenge 3 corpora, including data augmentation.
- Evaluated enhanced speech with standard objective metrics (PESQ-WB, STOI/ESTOI, SI-SDR) to benchmark denoising quality and validate architectural choices.

Student Research Assistant, *University of Toronto, Toronto, Ontario, Canada* Oct. 2025 - present
Supervisor: Rahul G. Krishnan, University of Toronto

- Cleaned and preprocessed single-cell Perturb-seq datasets; ran model training experiments with generalization across biological context for the Arc Institute's Virtual Cell Challenge.
- Expanded a **scaling-law** framework for foundation models CausalPFN, by swapping transformer backbones with FlashAttention.
- Extended TraCEE, a transformer for **causal effect estimation**, with a Tab-ICL v2 mixture prior and an ACIC 2026 **calibration** pipeline to detect prior distribution mismatch.

Undergraduate Research Assistant, *University of Waterloo, Waterloo, Ontario, Canada* Jan. 2024 - Apr. 2025
Supervisor: Pascal Poupart, University of Waterloo

- Co-designed a **transformer-based** tabular architecture ("basis transformers") to handle heterogeneous columns, text features, and missing or noisy values; lifted median R^2 by +0.338 with the lowest variance across 34 OpenML-CTR23 tasks while using 5x fewer parameters than the best baseline.
- Performed end-to-end **EDA** on **T4** (1.34TB) and **OpenML benchmarks** (**PCA**, **leakage checks**, **type inference**); encoded invariances (hierarchical structure, numeric representation) that stabilized results across datasets.
- Implemented robust preprocessing and augmentation for mixed-type tables (numeric tokenization, text column embeddings, few-shot reference sampling); improved generalization to unseen schemas without metadata beyond column names.
- Built and maintained a rigorous experimentation pipeline using **Python** with **PyTorch** for baselines; reported reproducible R^2 and F1 metrics with best-checkpoint selection.
- Scaled training and evaluation on A100 GPUs with **HuggingFace Accelerate** and **Slurm**; reduced run-to-run variance and enabled zero-shot and few-shot experiments across seeds and datasets for reliable comparisons.

Undergraduate Research Scientist, *Vector Institute, Toronto, Ontario, Canada* May. 2024 - Aug. 2024
Supervisor: Pascal Poupart, University of Waterloo

- Developed a Contextual RBF-Histogram algorithm using **Python** for **contextual bandits** (RBF-kernel–weighted reward bins) to select the best arm per context, achieving 55.6% correct arm selection on a test set of 2K data points.
- Formulated and implemented a **contextual multi-armed bandit (CMAB)** procedure for sequential decision-making, using an RBF reward for greedy arm selection and **XGBoost** for label generation, with online data aggregation to improve the policy.
- Evaluated on datasets over 10K rows, producing diagnostic plots to measure uncertainty and explainability.

WORK EXPERIENCE

Data AI Platform Developer (Intern), *Geotab, Toronto, Ontario, Canada* May. 2026 - present

- Building the **GenAI** Platform and **Agentic AI** solutions to empower developers and serve 3500+ users.
- Building and maintaining a **GenAI Gateway** that provides a unified interface to various upstream LLM services with **FastAPI**.
- Maintaining an internal Chat application that allows seamless experience between various **LLM models**, empowering thousands of Geotabbers with SOTA AI technologies with **LiteLLM**.
- Serving the **MCP client** and Agentic AI solution that allows developers to integrate their own service and scaling of the Agent's ability to serve the users better.
- Ensuring the high scalability and availability of these services while keeping a low maintenance effort by engineering sophisticated **CI/CD** pipelines leveraging **Kubernetes**, **Helm**, and **Terraform**.

Full Stack Developer (Co-op), *Aterlo Networks, Waterloo, Ontario, Canada* May. 2023 - Dec. 2023

- Engineered OOP **Python** data models and a cloud ingestion pipeline for multi-agent network telemetry, supporting 10M+ fixed-wireless/fiber devices across 20+ customer sites and reducing network load by 15%.
- Automated unit-test generation with **Python unittest and mock** (diff detection and docs) for thousands of models, accelerating regression coverage and team adoption; tool was integrated into workflows.
- Designed and shipped a responsive device information card UI (**JS/TS/Node, CSS**) surfacing port stats and connected-device counts across all supported routers/APs, validated via simulation on phone/tablet/desktop.
- Implemented semaphore-based session locking for **SSH** device tooling to prevent concurrent writes, stabilizing **APIs** under load and cutting failure modes from overlapping sessions.
- Built a **Protobuf** to **UML** generator and wired it into **Bitbucket CI**, visualizing internal **API message** relationships to speed schema reviews and clarify service contracts across **Python/JS** codebases.
- Reverse-engineered legacy devices (no manuals) using **cURL** and vendor CLIs; instrumented **Google Cloud queries** and dashboards to isolate model-specific logic defects and shorten debug cycles.

Data Engineer (Co-op), *Tesla, Shanghai, China* Mar. 2022 - Aug. 2022

- Built an end-to-end **Django** and **MySQL** portal for recruiting data (1K+ rows, 40+ fields, 20 users); added **SSO**, **RBAC**, and row-level access to eliminate overwrites and lift team throughput 50%.
- Implemented **CRUD**, advanced **SQL** filters, and **Bokeh** interactive tables/charts; automated sub-team weekly summaries with email digests, saving hours of manual reporting each week across 5 groups.
- Containerized with **Docker** and deployed on **Kubernetes**; delivered reproducible dev/prod environments and reliable releases while coordinating access/compliance with 3+ departments.
- Built **PySpark** and **Airflow ELT** joining vision and telematics for seat-belt detection analysis; disentangled true unbuckled behavior from detector failures, flagging 20% of 10M+ events as potential model misses.
- Published curated **MySQL** gold tables and SQL-ready reports powering safety analytics and offline model evaluation; accelerated root-cause triage and focused model iteration on highest-impact scenarios.
- Bootstrapped **Excel/VBA macros** (data validation, dropdowns, summary generation) to stabilize the legacy sheet pre-migration, reducing data-entry errors and transition risk.

Machine Learning Engineer (Co-op), *iFlytek, Beijing, China* Sept. 2021 - Dec. 2021

- Engineered an ensemble of seven Chinese paraphrase-generation algorithms (**Python, jieba, LTP**), delivering 95.8% precision and 99.2% recall while preserving fluency and maximizing semantic diversity.
- Built a **BERT back-translation** pipeline to augment training data and diversify candidates, strengthening model generalization across domains.
- Preprocessed 1M+ sentences with generative models via a scalable, reproducible data pipeline, enabling large-scale training and evaluation.
- Applied semantic role labeling, POS/NER, and dependency parsing (**LTP**) to enforce linguistic constraints during generation, improving readability and grammaticality of outputs.
- Established an evaluation framework tracking precision and recall, creating objective quality gates for release decisions.
- Synthesized and presented state-of-the-art research on paraphrase generation and semantic similarity, informing algorithm selection and prioritizing high-impact experiments.

SELECTED PROJECTS

Multi-Agent Reinforcement Learning for Self-Propelled Fish Navigation in Hydrodynamic Flow Apr. 2026 - present

- Built a **multi-agent reinforcement learning** framework for **self-propelled fish navigation** in flow, with PPO training pipelines for single- and multi-agent settings.
- Designed a controlled observation-feature ablation to improve the **explainability** of which flow/state features drive navigation.

Determine Protein Information Flow Via Granger Causality Analysis Sept. 2025 - Dec. 2025

- Formulated **protein allostery mapping** as a causal inference problem on **molecular dynamics (MD) trajectories** of Adenylate Kinase (AdK) to infer **directed information flow** under **cyclic dependencies**.
- Developed a **neural Granger causality** framework to estimate residue-level information flow, combining a **Bayesian prior** to improve biological plausibility. Validated on **synthetic ground-truth simulations** and **real AdK trajectories**; recovered clean and directed signalling pathways consistent with known AdK mechanics and outperformed correlation-based baselines.

AWARD, SCHOLARSHIP, and GRANT

- ($30,000) Mitacs Accelerate Award 2026
- ($7,500) Vector Institute Research Grant 2024
- ($1,500) University of Waterloo President's Research Award 2024
- ($3,000) University of Waterloo CS International Student Upper-Year Scholarship 2022
- ($2,000) University of Waterloo President's Scholarship of Distinction 2021

TECHNICAL SKILL

- Languages and Libraries: Python (PyTorch, NumPy, pandas, scikit-learn, SciPy; unittest/mock), R, C++, SQL, Bash.
- Platforms and DevOps Tools: Django, MySQL; Git and Bitbucket Pipelines (CI/CD); Docker, Kubernetes; AWS EC2.
- Machine Learning and Data: Transformers; NLP paraphrasing; safe/explainable RL (contextual bandits, Bayesian optimization); EDA and feature engineering; data augmentation.