

Bowen Fang

bf2504@columbia.edu | (917)293-3109

LinkedIn: www.linkedin.com/in/bowenfang/ | Personal website: <https://bwfbowen.github.io/>

EDUCATION

Columbia University in the City of New York

New York, NY

M.S. in Operations Research. GPA: 4.11/4.0

Expected Dec 2023

- Academic Mentor: [Cédric Josz](#)
- Relevant coursework:
PhD-level: Advanced Big Data and AI, Big Data Analysis, ML and High-dimensional Analysis, Robot Learning
Master-level: Optimization, Stochastic Models

Peking University

Beijing, China

B.S. in Big Data Management and Application, minor in Economics

Jul 2022

- Thesis title: A Trading Strategy Based on Transformer for Simulation. Advisor: [Meng Fan](#)
- Relevant coursework: Database System, Deep Learning, Data Structure and Algorithm, Graph Theory, Game Theory, Econometrics

RESEARCH EXPERIENCE

Dept. of Civil Engineering, Columbia University

NY, United States

Advisor: [Sharon Di](#)

Dec 2022 - present

- My research focuses on planning with RL.
- Developed heuristic operators and RL framework for PDTSP.
- Developed a learned model for simultaneous planning and mapping.

Knowledge Discovery Lab, Peking University

Beijing, China

Advisor: [Meng Fan](#)

Spring 2021 - Spring 2022

- My research focuses on RL application on Finance.
- Developed a temporal Transformer model for short-term market movement prediction. Experimentally investigated DQN trading agent on agent-based simulation.

PUBLICATION

(Submitted to ICAPS) Bowen, F., Xu C. and Sharon D., Zhengkun P. Plan and learn to Map for Joint SLAM and Navigation

- Proposed end-to-end MCTS with encoder-decoder architecture that generates and works on interpretable hidden state.
- Demonstrated effectiveness by conducting Navigation, mapping and ablation experiments on Gibson dataset.

(Submitted to AAMAS) Bowen, F., Xu C. and Sharon D., Learn to Tour: Operator Design For Solution Feasibility Mapping in Pickup-and-delivery Traveling Salesman Problem

- Develop operators that are proved to always generate feasible Hamiltonian cycles in PDTSP.
- We integrate the proposed operators into an RL framework, and demonstrate the effectiveness by conducting experiments on various problem scales. We provide a comparative analysis of computational efficiency and solution quality between our method and off-the-shelf solvers.

TALK

A Learning-Based Approach for High Capacity Ridesharing, INFORMS 2023

AZ, United States

RESEARCH PROJECTS

Autonomous Learning of Physical Environment through Neural Tree Search[\[link\]](#)

NY, United States

- Proposed a MCTS-based reinforcement learning algorithm to perform active slam.
- Combined MuZero loss function with SLAM loss to enable end-to-end building map while planning.

Temporal Graph Attention Network Prediction on Ethereum Transaction Cost[\[link\]](#)

NY, United States

- Proposed a GNN model based on temporal transaction network that outperforms LSTM, GRU.
- Enabled automatic train and tune with Airflow and deployed on GCP to provide real-time predictions.

Light Attention Vision Modules for DRL on Atari Games[\[link\]](#) NY, United States

- Implemented fast attention mechanism, which requires linear time and space instead of quadratic.
- Experimentally explored DQN agent with convolution, regular and fast attention on Atari games.

OPEN SOURCE EXPERIENCE

Author of Python Open-source Library MCTS-based Reinforcement Learning [\[link\]](#) NY, United States

- Implemented easy-to-modify and light weight MuZero and its variants that can be seamlessly fitted into RL pipeline and provide APIs that are familiar to RL practitioners compared to other implementations.
- Implemented recent extensions of MuZero within the RL framework Acme to enrich the toolbox. [\[link\]](#)
- Achieved that Muax is now an example project listed in DeepMind's library mctx.

TEACHING EXPERIENCE

Graduate Optimization Models and Methods NY, United States

Teaching assistant

Spring 2023/ Fall 2023

- Teaching assistant for the graduate Optimization Models and Methods class. Topics include linear programming, the simplex method, duality, nonlinear, integer and dynamic programming.
- Graded homework and course project, provided detailed feedback, revised solutions and improved the final project moving object detection coding part. The detailed summary can be found at [\[link\]](#)

INDUSTRY EXPERIENCE

MathWorks NY, United States

Part-time Data Scientist

Aug 2022 - Dec 2022

- Developed novel RL algorithms with MATLAB to solve Goal-based Wealth Management which make better decisions in all scenarios and longer horizons. [Blog](#).
- Proposed reward engineering methods to enable the agents to approximate to optimal behavior generated from dynamic programming on synthetic data with limited access to 10% of all cases.
- Realized an increase from 41% to 61% average success rate compared to Q-Learning for portfolios with practical constraints and the result is supported by extensive experiments.

AI TOPIA

Beijing, China

ML Researcher Intern

Sep 2021 - Jun 2022

- Researched on order placement optimization so to better control execution cost for our trading strategies.
- Developed a new tick-level buy/sell matching program that incorporates Indexed Priority Queue to handle large raw data parallelly on cluster to provide the team data with higher quality and faster.
- Optimized tick-level order placement strategy with 2 researchers using DRL so it would place on average 2-ticks better price and remain similar inference time through distillation.

Nomura Securities Orient International

Shanghai, China

Software Engineering Intern

Jul 2021 - Sep 2021

- Perfected our website so customers can get a more thorough understanding of fund products compared to the old version. Built dashboard and SQL database for metrics include style and systematic risk.
- Realized a new feature in MongoDB and a web crawler so that data from target URLs are stored into SQL databases and updated in front-end visualizations on a daily basis.

BLOGS AND TUTORIALS[\[LINK\]](#)

Some Recent Advancement Around MuZero Oct 2023

Adding MuZero into RL Toolkits at Ease Jun 2023

“Hindsight” – An easy yet effective RL Technique HER with Pytorch implementation May 2022

What are the Effective Deep Learning Models for Tabular Data? Mar 2022

SKILLS

- Technical: Airflow, AWS, Docker, GCP, Git, MATLAB, MongoDB, Spark, SQL, Java, Python
- Language: Fluent in English and Mandarin