

# Bowen Fang

[bf2504@columbia.edu](mailto:bf2504@columbia.edu) | Tel: 9172933109 | Personal Website: [www.github.com/bwfbowen](http://www.github.com/bwfbowen) |  
Google Scholar: [scholar.google.com/citations?user=bEebg80AAAAJ](https://scholar.google.com/citations?user=bEebg80AAAAJ) | LinkedIn: [www.linkedin.com/in/bowenfang/](http://www.linkedin.com/in/bowenfang/) |

## EDUCATION

---

Ph.D. Student, School of Engineering and Applied Science, **Columbia University**, New York, NY 2024-present

- Focus: Reinforcement Learning
- GPA: 4.1
- Advisor: Sharon Di
- Relevant courses: Diffusion Models AI & RL, Probabilistic Models and Machine Learning

M.S. in Operations Research, **Columbia University**, New York, NY Dec 2023

- GPA: 4.13
- Relevant courses: Advanced Big Data and AI, ML and High-dimensional Analysis, Robot Learning, Optimization, Stochastic Models

B.Mgt in Big Data Management and Application, **Peking University**, Beijing China Sep 2018 - Jul 2022

- GPA: 3.5
- Relevant courses: Algorithms and Data Structure, Database

## INDUSTRY EXPERIENCE

---

Applied Scientist, AWS Deep Engine-Science, New York, NY May 2025 – Aug 2025

- Developed a novel RLVR algorithm for mathematical reasoning, boosting pass@1 accuracy by 9% over the base Llama and QWen and outperforming the state-of-the-art DAPO by 7% on pass@k.
- Implemented a custom off-policy rollout mechanism in verl to maintain model entropy at ~6.0 (vs. <1.0 in standard methods) to ensure training stability and effectiveness.
- Automated the deployment and management of a multi-node (3x8 A100 GPUs) distributed training cluster on AWS, leveraging vllm for efficient inference.

Part-time Data Scientist, **MathWorks**, New York, NY Aug 2022 - Dec 2022

- Reproduced **Dynamic Programming** for Goal-based Wealth Management, developed 4 reward signals, devised experiments with 6 model-free agents to discover the best algorithm
- Achieved an increase of success rate from 41% to 61%, compared to **Q-Learning** baseline, the success rate is -0.07 from theoretical best probability. A blog with code was published officially and **added into MATLAB toolbox** [\[blog\]](#)

Machine Learning Research Engineer Intern, **AI TOPIA**, Beijing, China Sep 2021 - Jun 2022

- Proposed learning algorithms to optimize execution cost for intraday order placement optimization. Implemented 5 **model-free algorithms** and a high frequency trading environment
- Utilized advanced models including **DeepLOB** and XGBoost to improve the short-term prediction. Realized on average 2-ticks better price for tick-level order placement

Quant Analyst Intern, **Nomura Orient International Securities**, Shanghai, China Jul 2021 - Sep 2021

- Perfected our **FOF website platform**. Built **dashboard and SQL database** for metrics include style and systematic risk
- Realized a new feature in **MongoDB** and a **web crawler**, which has been **added to the department's core codebase**

Data Scientist, Intern, **Deloitte**, Beijing, China Aug 2020 - Oct 2020

- Completed a **web crawler** project and automatically produced formatted excel spreadsheet weekly
- Collaborated with data scientists on the **interactive data visualization**, preprocessed fundamental data for **business intelligence dashboard** products

## PUBLICATIONS AND PREPRINTS

---

**Decaying Budget Forcing: A Simple and Effective Reinforcement Learning Approach for Balancing Accuracy and Capacity in Mathematical Reasoning**

**Bowen Fang**, Hengzhi Pei, Leonard Lausen (in submission to ACL)

## Efficient Consistency Model Training for Policy Distillation in Reinforcement Learning

**Bowen Fang**, Xuan Di *ICLR 2025 Workshop on Deep Generative Model in Machine Learning: Theory, Principle and Efficacy*  
[Paper]

## SLAMuZero: Plan and learn to Map for Joint SLAM and Navigation

**Bowen Fang**, Xu Chen, Zhengkun Pan, Xuan Di *Proceedings of the International Conference on Automated Planning and Scheduling, ICAPS 2024*  
[Paper], [Code]

## Learn to Tour: Operator Design to Unlock Feasibility Mapping

**Bowen Fang**, Xu Chen, and Xuan Di *ITSC 2025*  
[arxiv]

## TravelLM: Could you plan my new public transit route in face of a network disruption

**Bowen Fang**, Yang Zixiao, and Xuan Di *ITSC 2025*

## MetaLLMetro: LLM-Centric Virtual Reality Navigation Platform for Metro Stations

Yongjie Fu, **Bowen Fang**, Mengxuan Liu, Xuan Di, *ITSC 2025*

## PROJECT EXPERIENCE

---

### SINA: Seamless, Intelligent Navigation Anywhere

NY, United States

- Awarded first place in the CS3 VALIDATE accelerator, securing grant funding to develop a next-generation emergency response and navigation platform.
- Pioneered a modern emergency response method, now U.S. Patent Pending, that intelligently routes users based on tiered environmental and network signals.

### Python Open-source Library Muax[code]

NY, United States

- Implemented variants of **MuZero** with **JAX/TensorFlow**, supported **distributed RL** and **distributed training**
- Integrated with DeepMind's **Acme** framework, compared **MCTS** search policies on **Open Spiel** multi-agent games (e.g., Go)

### SLAMuZero[code]

NY, United States

- Enhanced **MuZero** perception with **SLAM** for **embodied AI**, experiment on Meta's **Habitat simulator**
- Realized comparable performance for indoor **exploration** tasks on Gibson dataset, with better sample efficiency

### LLM Agent Navigation in VR Unity

NY, United States

## TECHNICAL SKILLS

---

- Languages: Python, SQL, C++, C#, Java, JavaScript, HTML/CSS, MATLAB
- Machine Learning:
  - Deep Learning: JAX, PyTorch, TensorFlow, Hugging Face Transformers, vLLM
  - General ML & Vision: Scikit-learn, XGBoost, OpenCV
  - Reinforcement Learning: Acme, mctx (MuZero), verl, RLlib (Ray), stable baselines, OpenSpiel
  - Data Analysis & Visualization: Pandas, NumPy, SciPy, Matplotlib, Seaborn, Plotly
- Data & MLOps: Apache Spark, Docker, Airflow, MongoDB
- Cloud Platforms: Amazon Web Services (AWS), Google Cloud Platform (GCP)
- Developer Tools: Git, Linux, Unity, XCode (iOS)

## HONORS AND AWARDS

---

NSF Center for Smart Streetscapes Validator Award  
Calatrava Family Fellowship

2025  
2024

## PRESENTATIONS

---

- Learn to Tour: Operator Design for Solution Feasibility Mapping, INFORMS, Phoenix, United States, 2023

- SLAMuZero: Plan and learn to Map for Joint SLAM and Navigation, ICAPS, Banff, Canada, 2024

## PROFESSIONAL SERVICE

---

### Reviewer

- Transportation Research Board
- Transportation Research Part C: Emerging Technologies

### Workshop Organization

- Student Assistant, Urban Mobility AI Workshop, Columbia University, 2024
  - Developed and maintained workshop website with **HTML/CSS/JS**, Assisted in organization and coordination
  - Workshop URL: [[Homepage](#)]

## TEACHING EXPERIENCE

---

Teaching Assistant, IEOR department, **Columbia University**, New York, NY

2023

- Graduate Optimization Models and Methods (IEOR E4004) instructed by Cedric Jozs
- Self-motivated to enhance final project on **Moving Object Detection**, created sample videos and code templates for quick start[[summary](#)]