

# Bojun Feng

(773) 980-3464 | bojundf@gmail.com | github.com/Bojun-Feng | linkedin.com/in/Bojun-Feng

## WORK EXPERIENCE

### Computational Social Science Lab, University of Pennsylvania

Dec 2023 – Present

#### Research Assistant

Philadelphia, PA

- Slashed dataset labeling time from 400 to 3 minutes by leveraging LLMs and parallel programming techniques on a multimodal dataset of 300+ behavioral research experiments
- Boosted accuracy by over 15% and F1 score by over 10% on a domain-specific labeling task by implementing human-in-the-loop outlier detection and iteratively refining a GPT-4o model with 2+ million tokens of data
- Optimized deployment speed by 90% through CI/CD with GitHub Actions, developing and deploying 4 internal tools with React.js, Node.js, Docker, and Postgres on AWS EC2 and RDS
- Spearheaded evaluation of 10+ open-source multimodal models by designing internal benchmarks to assess LLM performance

### EMK Lab, University of Chicago

Jan 2022 – Present

#### Data Science Intern

Chicago, IL

- Accelerated video labeling by 50% through training a Residual Neural Network to analyze rodent memory formation from behavioral video data, leveraging clustering and topological techniques for efficient processing
- Architected data pipeline processing 3 million rows of experimental data, uncovering relationships between rodents' learning processes and 60+ behavioral patterns
- Secured the Computing division award and presented poster at the 2023 Mind Bytes Research Symposium, topping a list of 50+ contenders

### XorbitsAI

Jul 2023 – Sep 2023

#### Software Engineer Intern

Beijing, China

- Integrated 10 open source LLMs into the Python object-oriented backend for deployment and inference, added support for 4 different levels of quantization, integrated library into other open source projects like Llama Index
- Authored 20+ pull requests and reviewed code on Git with continuous integration, unit tests, and proper documentation, played a key role in the progression of Xinference from 0 to 800+ stars on GitHub in 2 months
- Engineered dashboard with 6k+ users with scalable RESTful API to track, create, and delete running models

## PROJECT / ACHIEVEMENT

### Meta Hacker Cup

Oct 2024 – Nov 2024

- Placed global top 50 (top 0.01%) among 22k+ contestants in Meta Hacker Cup 2024 Coding Competition

### Uchicago Offer Acceptance Prediction Model

Nov 2023 – Feb 2024

- Achieved 90% accuracy in predicting University of Chicago applicants' offer acceptance by training XGBoost and Neural Network classifiers on 3K+ data points spanning 10 years
- Enhanced model performance by over 10% by deploying the Mixtral 8x7B LLM model with Docker and Kubernetes in the data processing pipeline to securely extract privileged information from raw text inputs

## EDUCATION

### University of Chicago

June 2025

#### Master, Computer Science

3.55/4.0 GPA

Classes and Honors: Generative AI, Cloud Computing, Optimization, Theory of Algorithms, Dean's List

### University of Chicago

June 2025

#### Bachelor, Computational and Applied Mathematics

3.55/4.0 GPA

## SKILLS

Programming Language / Frameworks: Python, JavaScript, TypeScript, SQL, Node.js, React.js, Go

Data Science / Machine Learning: PyTorch, Transformers, Llama.cpp, Numpy, Pandas, Matplotlib, Hadoop, Llamafire

Other Tools and Technologies: Git, Docker, Kubernetes, Jupyter Notebook, Azure, AWS, S3, EC2, RDS, Lambda, Lightsail