

# Ashmit Mukherjee

New York University Abu Dhabi  
asm8879@nyu.edu | ashmitmukherjee.com

## Education

---

### New York University Abu Dhabi

*Expected May 2027*

B.S. in Computer Science

*Studied at:* NYU Abu Dhabi; NYU Paris; NYU New York

**Relevant Coursework:** Natural Language Processing; Principles of Data Science; Statistics for Social & Behavioral Sciences; Data Structures & Algorithms; Computer Systems Architecture; Operating Systems

## Research Interests

---

I work at the intersection of machine learning and computational social science, with a particular focus on how artificial agents reshape human cooperation, coordination, and norm formation. My current work models social dilemmas in mixed Human-AI populations, asking when AI-mediated interventions stabilize cooperative equilibria and when they disrupt them. My broader portfolio — spanning multilingual and code-mixed NLP and protein representation learning — reflects a methodological interest in how models learn meaningful structure from heterogeneous data, and how we evaluate that learning rigorously.

## Publications

---

**A. Mukherjee** (co-first author). “Bio-Informed LoRA for Signal Peptide Prediction.” *Under review at EMNLP 2026 Workshop*, 2026.

## Research Experience

---

### Research Assistant — Human-AI Collaboration, NYU Abu Dhabi

*Feb 2026–Present*

- Developing a simulation-based research platform for studying **AI-mediated coordination in multi-party social dilemmas**, using global supply chains as the domain (advisors: Prof. Hanan Salam, Prof. Benjamin Rosche).
- Built a calibrated game-theoretic environment with five buyers and three suppliers across repeated rounds, integrating public goods support, strategic ignorance, threshold remediation, exit-versus-engagement, and auditing into a single domain-grounded model; main buyer policy uses quantal response equilibrium with a Stage 2 belief solver.
- Reproduced the target inefficient-equilibrium pattern (~22–26% welfare loss vs cooperative ceiling, ~99% supplier defection, ~7% remediation) across a 12-cell calibration sweep with 50 simulations per cell; backend has 101 passing tests with three completed audit rounds.
- Designing the next phase as an experimental platform comparing algorithmic baselines against human-participant sessions in which AI agents act as advisors, coordinators, monitors, or commitment-support systems.

### Research Assistant, eBRAIN Lab — NYU Abu Dhabi

*Jan 2026–Present*

- Co-developed **Bio-Informed LoRA**, a parameter-efficient fine-tuning method that injects residue-level biological priors (BLOSUM62, hydrophobicity, Grantham) into LoRA adapters for ESM-2 protein language models on signal peptide prediction.
- Showed BLOSUM62-guided LoRA matches or exceeds full fine-tuning on SignalP6 benchmarks across SP MCC, cleavage-site exactness, and residue-level F1, while training only ~3.6% of parameters and using ~43% less peak GPU memory at the 3B-parameter scale.
- Designed multi-seed evaluation protocols and a one-factor sensitivity screen demonstrating that single-seed improvements frequently fail to replicate — a methodological finding central to the paper’s contribution.
- Manuscript under review (see Publications).

## Research Projects

---

### Hinglish Named Entity Recognition Benchmark

*Fall 2024*

- Fine-tuned multilingual transformer models (mBERT, XLM-RoBERTa) on the COMI-LINGUA dataset for Hindi-English code-mixed NER.

- Achieved 78% entity-level F1, outperforming zero-shot GPT-4o and LLaMA 3.1 baselines under controlled evaluation.
- Analyzed trade-offs between model scale, task-specific fine-tuning, generalization, and computational cost.

### Data Science Salary Prediction Platform (Methodological Study)

*Fall 2025*

- Conducted large-scale regression benchmarking using AutoML (PyCaret) across 15+ models on a dataset of 3,000+ job postings.
- Analyzed feature contributions using SHAP to study labor market dynamics across geography, firm size, and skill requirements.
- Tracked experiments with MLflow; reported MAE, RMSE, and  $R^2$  across conditions.

### Crime Data Analysis and Prediction

*Spring 2025*

- Studied relationships between socioeconomic indicators and violent crime rates using regression modeling on data from 1,994 communities.
- Evaluated model performance via cross-validation ( $R^2 = 0.85$ ) and exploratory statistical analysis; built interactive visualizations for interpretability.

### CAMP — Campus Asset Management Platform

*Fall 2024*

- Designed and implemented a multi-user system for resource allocation and scheduling under shared constraints, with role-based access and conflict resolution.
- Developed RESTful APIs and automated CI/CD pipelines using Docker and GitHub Actions for reproducible deployment.

## Industry Experience

---

### Machine Learning Engineer Intern, Zeek

*Jun–Aug 2025*

- Developed and evaluated segmentation models on noisy real-world data, iterating on preprocessing and model selection to meet production reliability constraints.
- Applied model interpretability techniques to study robustness and bias sensitivity, informing decisions on which model variants to deploy.
- Contributed to the team's evaluation workflow, integrating standardized metrics and validation procedures into the model development cycle.

## Technical Skills

---

**Languages:** Python, C++, C#, JavaScript

**ML & NLP:** PyTorch, Transformers (Hugging Face), Scikit-Learn, XGBoost, PyCaret, SHAP

**Multi-Agent & RL:** NumPy, NetworkX, reinforcement learning pipelines

**Data Analysis:** Pandas, SciPy, Matplotlib, Seaborn, Plotly

**Experiment Tracking:** MLflow, Git, Docker, GitHub Actions

**Other:** Streamlit, MongoDB,  $\text{\LaTeX}$