

Aman Choudhri

amanchoudhri@gmail.com | San Francisco, CA | github.com/amanchoudhri | amanchoudhri.com

EDUCATION

Columbia University, Columbia College New York, NY
BA in Mathematics and Statistics | GPA: 4.05 | Phi Beta Kappa, Magna Cum Laude May 2025
• **Courses** (* for PhD-level): Agentic Systems*; High Performance ML*; Bayesian Optimization*; Econometrics III*; Causal Inference*; Interpretable ML*; Multilevel Regression Modeling; Advanced Probability Theory; Probabilistic ML*

PUBLICATIONS

Peterson R., Tanelus A., **Choudhri A.**, Ivan V., Prasad A., Schneider D., Sanes D., Williams A. (2023). "Benchmarks and deep learning models for localizing rodent vocalizations in social interactions." *ML for Audio Workshop, NeurIPS*.
Peterson R., **Choudhri A.**, Mitelut C., Tanelus A., Capo-Battaglia A., Williams A., Schneider D., Sanes D. (2023). "Unsupervised discovery of family specific vocal usage in the Mongolian gerbil." *eLife*, 12:RP89892.
Shahn Z.*, **Choudhri A.***, Jung B., Talmor D., Lehman L., Baedorf-Kassis E. (2023). "Effects of aggressive and conservative strategies for mechanical ventilation liberation." *Journal of Critical Care*, 76, 154275.

EXPERIENCE

Exa Labs San Francisco, CA
ML Engineer Aug 2025 – Present
• Training and serving models for search.
Columbia University, Gelman Lab New York, NY
Research Assistant May 2024 – May 2025
• Designed and conducted a 6000+ respondent survey studying how people's social circles relate to their political beliefs
• Analyzed resulting data using multilevel (random-effects) linear and logistic regression modeling in R and Stan
Public Policy Lab (PPL) New York, NY
Summer Intern Jun 2024 – Aug 2024
• Developed natural language processing pipeline analyzing transcripts of interviews with social service recipients, saving 2+ weeks of PPL researcher time per project (sentence transformer embeddings and k-means clustering)
Flatiron Institute, Williams Lab New York, NY
Research Assistant Jun 2022 – Aug 2023
• Designed custom deep learning model architecture for gerbil vocal data that provides accurate estimates of its own uncertainty after synthesizing research on approaches to neural net uncertainty quantification (NeurIPS '23)
• Built parallelized Slurm pipeline and uncertainty calibration toolkit for scalable model training/evaluation in Pytorch
• Uncovered family-specific gerbil "dialects" by clustering 500k+ vocal calls using variational Gaussian mixtures in Stan
CUNY School of Public Health, Shahn Lab New York, NY
Volunteer Researcher Jan 2021 – Jun 2022
• Applied causal inference to estimate the effectiveness of care strategies for ICU patients on mechanical ventilation

PROJECTS

Federal Disaster Relief Agent Jan 2024 – Present
• Building an agentic system to simplify federal disaster relief application for state/local governments (code available here)
• Constructed VLM pipeline to extract structured information from PDFs of past grant applications with 97.5% accuracy
Wind Farm Power Optimization Sep 2024 – Present
• Applying Bayesian optimization to design wind turbine layouts for maximal power production (code available here)
• Implemented multi-fidelity strategy that efficiently combines information from cheap approximate models and expensive high-fidelity simulations, identifying optimal turbine configurations with 65% fewer high-cost simulation runs
• Built SLURM orchestration API to queue/manage fluid dynamics simulations parallelized across 192 HPC CPU cores
Bloggregator: LLM-Powered Blog Aggregator May 2025
• Built a monitor for blogs without RSS feeds, generating webscraping schemas using LLMs with 100% parsing success
• Created Github Actions workflow to process 10+ blogs and 400+ posts daily, generating summaries/topics for new posts

ACTIVITIES

President, Columbia Bhangra: Led 30-person competitive dance team. Managed \$15,000+ budget; hosted and secured \$5000 in corporate sponsorships for a 6-team competition yielding \$6000+ in revenue from 400+ attendees

SKILLS

Programming Languages: Python, R, Stan, SQL, CUDA
Python Libraries: PyTorch, JAX, Transformers, Scikit-Learn, NumPy, Pandas, Plotly
Infrastructure and Tools: Linux, Bash, Git, Docker, GCP, Slurm, Cron, PostgreSQL