

Chao Peter Yang

Tel: (+1) 773 - 901 - 0080 Email: chaopeter.yang@gmail.com LinkedIn Personal Website

EDUCATION

- Duke University** Aug 2024 - May 2026
♦ *M.S. in Interdisciplinary Data Science, Merit Scholarship Recipient (70%)* GPA: 3.96 / 4.0
♦ Relevant Courses: Natural Language Processing, Statistical Modeling, Data Engineering, Theory of Machine Learning
- University of Michigan - Ann Arbor** Aug 2018 - May 2021
♦ *B.S. in Data Science and Mathematics* GPA: 3.65 / 4.0
♦ Honors: **Highest Honors** in Data Science (**One of Two** in Department), **University Honors** (2019, 2021)

TECHNICAL SKILLS

- Programming** Python (PySpark, PyTorch, sklearn, SciPy Pandas, Numpy), R (SparkR, dplyr), SQL, Rust, Scala, C++, Javascript, MATLAB, Git
- Machine Learning** Regressions, Boosted Trees, RLHF, Deep Learning, LSTM, CNN, GNN, Transformers, Diffusion
- Gen AI/Agents** Langgraph, ReAct Agents, Amazon Bedrock, Langfuse, RAG, MCP Server
- Big Data/Databases** Spark, Databricks, AWS, SQL Server, PostgreSQL, MongoDB, SQLite
- Visualization** matplotlib, ggplot, Tableau, Excel Dashboard
- Languages** **Fluent:** English, Hungarian, Mandarin Chinese — **Conversational:** German, Japanese

EXPERIENCES

- Amazon.com, Inc. – Amazon Robotics** May 2025 – Aug 2025
Data Scientist Intern Boston, MA
- ♦ Researched and developed an advanced **AI Agent** for root-cause investigation integrating multiple data sources and **MCP servers**, reducing average warehouse troubleshooting time from several days to **2.5 minutes (75% success rate)**.
 - ♦ Built a scalable agentic framework that unified internal agent development using **LangGraph** and **Amazon Bedrock**, streamlining multi-agent workflows and accelerating deployment cycles.
 - ♦ Engineered a dynamic, production-ready evaluation pipeline that enabled rapid setup and large-scale benchmarking of agent performance, leveraging **LLM-as-a-Judge** techniques with **Langfuse** integration.

- Interpretable Machine Learning Lab, Duke University** Aug 2024 – Current
Research Assistant, Co-advised by Prof. Cynthia Rudin and postdoc Dr. Stephen Ni-Hahn Durham, NC
- ♦ Co-first authored **ProGress**—structured symbolic music generation via rule-guided **Discrete Diffusion**; 45-subject study beat baselines on quality with only ~3M params; **Accepted for NeurIPS 2025**.
 - ♦ Researched and developed a custom implementation of **DiffPool** for **Heterogeneous GNN** used in musical analysis in **PyTorch**, reducing validation **cross-entropy loss** by 60% with additional **hyperparameter tuning**.
 - ♦ Developed **Proximal Policy Optimization** for **Graphical Neural Network** to enable **Reinforcement Learning from Human Feedback**, enabling automated, personalized musical analysis. *Under Review ACMHCI '26*

- Curinos, Inc. (Fintech 100, Banking Consulting)** Aug 2021 – June 2024
Senior Data Science Analyst Chicago, IL
- ♦ Researched and developed nonlinear price elasticity models for **asset-liability management (ALM)** to predict acquisition and balance flow for retail banks, improving prediction vs. legacy models by $\approx 23\%$ R^2 in testing.
 - ♦ Led ML engineering team of 4 to migrate legacy modeling pipeline from using **Cloudera** to **Databricks**, coordinating across teams to schedule testing, promotion, and release plans, leading to more than \$100k in annual savings for platform expenses, and a 30% decrease in data processing time on average.
 - ♦ Created automated ad-hoc regression notebooks with **PySpark** for creating, testing, and validating models with different configurations, reducing the time to build proof-of-concept models by 50%.
 - ♦ Tuned parameters for price elasticity models en masse for major US banks, each with 10,000+ model segments, achieved $\approx 12\%$ faster convergence with higher R^2 and precise **Model Risk Management** documentation.

PROJECTS

- AI Engineering — Duke ProfMatch** Project Link Fall 2024
- ♦ Developed LLM-based professor **recommendation system** for Duke students using **GPT-4o-mini** with state-of-the-art **Graph-based Retrieval Augmented Generation** system, **LightRAG**, allowing for personalized recommendations.
 - ♦ Designed and implemented **Streamlit**-based chat-bot-like interface for ProfMatch with full NetworkX visualization.
- Modeling — Muscribe** Repo Link Fall 2023
- ♦ Developed and implemented both a **Transformer** and a **Convolutional Recurrent Neural Network** to transcribe audio music files to sheet music with **Pytorch**, while achieving accuracy comparable to SOTA models, despite limited training resources.