

POUYA GHAHRAMANIAN

Ankara, Türkiye • +90 536 832 8546 • pouyaghahramanian@gmail.com •
github.com/pouyaghahramanian • linkedin.com/in/pouyaghahramanian

EDUCATION

Ph.D. in Computer Engineering | *Bilkent University, Ankara, Turkey* 2022 – Expected 2027

- Research: machine learning on large-scale data streams — classification, anomaly detection, concept-drift adaptation, and LLM / time-series foundation model adaptation.
- Member of the Bilkent Information Retrieval Group (BilIR).

M.Sc. in Computer Engineering | *Bilkent University, Ankara, Turkey* 2019 – 2022

- Thesis: [Evolving Text Stream Classification with a Novel Neural Ensemble Architecture](#) — published in ACM TKDD.

B.Sc. in Computer Engineering | *Iran University of Science and Technology, Tehran, Iran* 2013 – 2018

EXPERIENCE

Data Scientist | *Invent.ai, Ankara, Turkey* September 2025 – Present

- Cut **lost sales by 10%** for an enterprise retail client by improving demand forecast accuracy and the replenishment decisions driven from it.
- Own **demand forecasting** and **replenishment** models in production for **three** enterprise retail clients, covering roughly **1M SKU-store pairs** — from feature pipeline through model release.
- Build and operate large-scale ML pipelines on **Databricks**, **PySpark**, and **Airflow**: feature generation, training, backtesting, and scheduled inference.
- Design the offline and online evaluation behind model releases — backtesting setup, holdout construction, and drift monitoring on live series.
- Deliver inventory-visibility models and “next best action” recommendations for client planners, working with product and engineering to ship them.

Senior AI/ML Researcher (part-time) | *Bilkent University, Ankara, Turkey* April 2019 – Present

- Named senior researcher on three government-funded R&D programs (TÜBİTAK 117E870, 120E103, 125E060); lead the ML workstream on classification, anomaly detection, and concept-drift adaptation over large-scale streaming and event data.
- Designed **AdaNEN**, a neural ensemble for evolving data streams, improving classification accuracy by up to **8.8%** across **13** benchmark datasets (ACM TKDD).
- Co-designed **LACE**, an unsupervised concept-drift detector for multi-label streams that cuts detection delay by **63.5%** versus the previous best unsupervised method at a **0%** missed-detection rate (ACM CIKM 2026).
- Proposed the “Adaptimizer” optimizer and an **On-the-Fly Adaptation (OFA)** framework for continual **LLM** adaptation under temporal drift, evaluated across two decades of news (ACM CIKM 2025); implemented in **PyTorch** and **HuggingFace**.
- Own projects end-to-end — dataset construction, training, evaluation metrics, and offline/online experiments — turning research into reusable, well-evaluated models.

Teaching Assistant (part-time) | *Bilkent University, Ankara, Turkey* February 2019 – Present

- Teach and mentor students in Information Retrieval, Algorithms, and Computer Organization, explaining complex ML concepts to technical and non-technical audiences; supervise student projects on online learning, neural retrieval, and ranking systems.

Earlier | *Software and mobile engineering*

2017 – 2018

- **Android Developer Intern**, Petanux GmbH, Bonn, Germany (2018) — built an Android app for online price comparison across multiple vendors.
- **Software Engineer**, Shams Clinic, Tehran, Iran (2017) — built a clinic management system for appointment scheduling, staff coordination, and financial tracking, used in daily operations.

PUBLICATIONS

- Gofralilar, M. K., **Ghahramanian, P.**, & Can, F. (2026). LACE: Unsupervised Concept Drift Detection in Multi-Label Data Streams Through Label Cluster Evolution. **ACM CIKM**. [Code](#)
- **Ghahramanian, P.**, Bakhshi, S., & Can, F. (2025). LLM-OFA: On-the-Fly Adaptation of Large Language Models to Address Temporal Drift Across Two Decades of News. **ACM CIKM**
- **Ghahramanian, P.**, Bakhshi, S., Bonab, H., & Can, F. (2024). A Novel Neural Ensemble Architecture for On-The-Fly Classification of Evolving Text Streams. **ACM TKDD**
- Bakhshi, S., **Ghahramanian, P.**, Bonab, H., & Can, F. (2023). A Broad Ensemble Learning System for Drifting Stream Classification. **IEEE Access**

Under review: **Ghahramanian, P.**, & Can, F. CoSign: Dual-Timescale Co-Signed Consensus for Online Adaptation of Time-Series Foundation Models. **AAAI 2027**. *A streaming optimizer that adapts a TSFM's LoRA weights under delayed supervision; ranks first against 14 streaming optimizers across 17 datasets.*

Full list and citations: [Google Scholar](#).

TECHNICAL SKILLS

Machine Learning & LLMs: PyTorch, HuggingFace Transformers, Scikit-learn, NumPy, pandas, Deep Learning, Neural Networks, Large Language Models (LLMs), Fine-tuning (LoRA / PEFT), Retrieval-Augmented Generation (RAG), Vector Databases

ML Methods: Time-Series Forecasting & Demand Prediction, Classification, Anomaly Detection, Recommendation, Ensemble Learning, Online & Continual Learning, Concept-Drift Adaptation, Time-Series Foundation Models, Feature Engineering, Experimentation (offline/online), Statistical Hypothesis Testing (Friedman, Wilcoxon), Model Evaluation

Data & MLOps: SQL, Databricks, PySpark, Airflow, FastAPI, Docker, Linux, Git, CI/CD, AWS (EC2, S3, Bedrock), Azure ML

Programming: Python (primary), SQL, C++, Java, C#

AWARDS & SCHOLARSHIPS

Gold Medal (x2), 2010: Iranian National Student Olympiads — Physics and Research

TÜBİTAK Research Scholarships: funded researcher on three national R&D programs

PROFESSIONAL SERVICE

Conference Reviewer: ACM SIGIR (2023–2025), ACM CIKM (2023–2026), SIGIR-AP (2024–2026)

LANGUAGES

English (Advanced, TOEFL iBT: 104), **Turkish** (Fluent), **Azerbaijani** (Native), **Persian** (Native)