

Vahid Faraji (jobehdar)

SENIOR APPLIED AI SPECIALIST

Istanbul, Turkiye, last update (June 2026)

vfaraji89@gmail.com

+90 544 392 1474

LinkedIn | GitHub

Personal Web

Summary

Applied AI Specialist with 9+ years building data products and enterprise platforms, progressing into LLM systems, context engineering, and complex agentic workflows. Expert in designing and scaling multi-agent architectures using the Model Context Protocol (MCP) to bridge LLMs with enterprise data sources. Experienced in establishing operational guardrails on Azure Foundry, Google Cloud (Vertex AI), OpenAI, and Anthropic SDKs with a strict focus on system observability, token budget management, and cost optimization.

Multilingual: Turkish, English, Spanish.

Experience

Senior Applied AI

Kariyer.net (ilab group) | Mar 2025 – Present

- Design and implement production-ready agentic workflows (LangGraph, MCP, Agno) featuring custom enterprise routing, safety fallbacks, and real-time observability telemetry.
- Developed token-efficient context engineering frameworks adopted as an organizational standard, significantly lowering total LLM token burn while boosting response factual accuracy.
- Lead automation, flow, and LLM Agent initiatives across Sales, Finance, and R&D, driving text-to-SQL reporting engines from initial scoping through to production hand-off.
- Built AI-powered feedback automation loops serving over 100 internal enterprise users, cutting manual operational workload by 60% via Python and automated workflows.
- Active contributor to group-wide AI transformation, securing a 40K USD Perplexity enterprise grant and establishing baseline standards for agentic infrastructure evaluations.
- Instructor for Applied AI technical courses covering custom tools, contextual design, and multi-agent systems; AI Career Mentor at TEV (Turkish Education Foundation).

Product Manager – Data

Kariyer.net | Oct 2022 – Mar 2025

- Supervised and designed metadata schemas serving over 300M+ text and numerical data assets, aligning cross-functional product roadmaps directly with core engineering teams.
- Led internal training and adoption initiatives for C-level executives to accelerate AI-native product delivery and lifecycle management.
- Designed real-time analytical reporting pipelines tracking high-velocity labor statistics, establishing data quality matrices, strict governance protocols, and executive dashboards.

Business Analyst – Data Analytics

WorqCompany | Sep 2021 – Oct 2022

- Built analytical financial forecast matrices and predictive revenue simulation streams used directly by leadership for strategic mid-range planning and scenario analysis.

People Data (Co-Founder)

Defaro.io | 2019 – 2021

- Launched a specialized labor-market compensation and analytics startup tracking localized talent supply signals, incubated and backed by Özyegin University.

Technical Skills

Protocol & Architecture: Model Context Protocol (MCP) server/client design, Claude Architect patterns, tool-augmented systems, LangGraph, Agno, multi-agent coordination frameworks

AI & Context Engineering: Context rot management, token budgeting, advanced RAG optimization, embedding engineering, prompt composition, performance benchmarking, Arize observability

Languages & APIs: Python (Async, FastAPI), TypeScript, SQL, Pydantic data validation, RESTful API design

Search & Infrastructure: PostgreSQL/pgvector, hybrid semantic search, custom re-ranking, token-aware chunking, Azure (Foundry, DevOps, Synapse), Google Cloud (BigQuery, Vertex AI, Agent Garden), Docker

Selected Projects

Weekend Advisor – Baachee.ai (Agent AI Studio)

Serve as strategic technical advisor architecting rapid-deployment agent frameworks, defining best practices for custom tool integration and open-standard protocols. Guided the structural design of specialized data agents, including end-to-end text2report systems, autonomous text-to-SQL semantic parsers, and node-graph canvas flow architectures.

Keywords: Strategic Architecture, Framework Design, Protocol Integration

Tokalator – Context Engineering Toolkit (Open Source)

VS Code extension and web platform for real-time token budgeting, tab relevance scoring, and context optimization for AI coding assistants.

Keywords: VS Code Extension, TypeScript, Token Economics

Multi-Agent Text Data Classification (Enterprise Data)

MCP-based multi-agent system for cross-lingual normalization of free-form job titles to O*NET and ESCO taxonomies. 72.5% reduction in classification time and up to 95% segment-level accuracy.

Keywords: MCP, Grounding, AI search, Caching, Multi-Agent, Hooks and MCP

Enterprise AI Transformation (Sales, Finance & R&D)

Design and delivery of AI transformation initiatives across multiple business units, using agentic workflows and LLMs for sales reporting, financial analytics, and R&D automation. Included PoC design, MLOps-ready APIs, security reviews, and production hand-off.

Keywords: FastAPI, Agno, Pydantic, Azure/GCP

Text-to-SQL (“Text to Report”)

Natural-language to SQL pipeline that turns business questions into structured reports via RAG and query-generation agents, enabling self-service analytics and reducing analytics backlog by 40%.

Keywords: DataHub, Azure Synapse, RAG, Docker

Job Matching & Workforce Canvas Graphs

AI-assisted matching pipelines for workforce programs. Developed node-graph components for visual canvas tracking of position clusters, processing complex geographic and department variables.

Keywords: Claude Code, Python, Node-Graph Analytics

Bank POS Commission Reconciliation (Finance AI Transformation)

Automated pipeline reading POS commission statements from 4 major Turkish banks, comparing con-

tracted vs. applied rates, and surfacing discrepancies through analytical dashboards.

Keywords: Azure OpenAI, Streamlit, Python, NumPy, Pandas

Education

M.A. in Economics

Ibn Haldun University | 2019 – 2023

Thesis: Real-Time Labor Market Analysis: A Data-Driven Approach

Teaching Assistant for Econometrics I; Research Fellowship in labor-market analytics

Publications

Tool-Augmented Multi-Agent Systems for Job Position Normalisation, accepted at *CAIS '26: ACM Conference on AI and Agentic Systems*, San Jose, 2026.

[ACM DL] | [alphaXiv]

Leveraging LLMs For Turkish Skill Extraction | [arXiv:2601.22885](https://arxiv.org/abs/2601.22885) (January 2026)

A large-scale empirical study on extracting structured skill signals for the Turkish labor market using modern LLM-based pipelines.

Tokalator: A Context Engineering Toolkit for AI Coding Assistants | [arXiv:2604.08290](https://arxiv.org/abs/2604.08290) (April 2026)

A technical paper introducing an open-source toolkit for token-aware context management, cost modeling, and workflow optimization in AI-assisted software development.

Talks & Community

Trainer & Mentor: *Applied AI Course – Agents, MCP, and AI-Native Development* (18-hour technical training tracking system context design and production tool execution) 2026

AI Career Mentor: TEV (Turkish Education Foundation) 2026

Invited Speaker:

- *Context Engineering for AI-Native Dev – Only Builders Club* 2026
- *Deep Agent Frameworks for SaaS – Bridge Summit '25* 2025
- *Enterprise Prompt Engineering and Tool Design – Kayseri University & Van University* 2025

Certifications & Additional

Certifications: Anthropic Claude Certified Architect (Active Practitioner), Google GenAI Leader

Languages: Turkish (Native), English (C1), Spanish (Instructor Level), Portuguese (Basic)