

# AJ Djalali

Athens, GA · 312.576.4525 · alex.djalali@gmail.com · alexdjalali.github.io

[LinkedIn](#) · [Google Scholar](#) · [GitHub](#)

## SUMMARY

Research faculty and engineering leader with 12 years building language technology that ships—from formal semantics research at Stanford through two acquisitions (ClearGraph by Tableau, PerceptivePanda by Zapier) and into production systems serving hundreds of thousands of users. PhD in Linguistics, 9 patents in NL interfaces, 15+ peer-reviewed publications. Designed the NLP architecture behind Tableau's Ask Data before transformer models existed, built deterministic control layers over LLMs grounded in discourse theory, and currently build production AI systems at Georgia Tech—including vector search over 400M images and RAG pipelines on AWS Bedrock. Scaled teams from founding engineer to 30+, led AI strategy across Salesforce's cloud portfolio, and consult on AI agents and retrieval-augmented generation for startups in maritime and financial services.

## EXPERIENCE

### Georgia Institute of Technology

Research Faculty & Senior Research Engineer

Atlanta, GA

Oct 2025 – Present

- Embedded with research teams to build production AI systems from scratch—hands-on engineering, not advisory
- Built enterprise search & ingestion platform: vector search over 400M+ images via CLIP/Qdrant at sub-200ms latency, distributed GPU ingestion via Ray/Spark on Databricks, zero-error completion across multi-million-image datasets
- Built congressional text analysis platform: 181M+ tokens, Kafka-decoupled ingestion, Elasticsearch + Neo4j entity graphs, spaCy NLP pipeline, full Prometheus/Grafana observability
- Designing graduate curriculum on software engineering for AI and distributed systems—15 weeks covering ADRs, trunk-based dev, SRE, and AI-augmented workflows

### Gridline

AI Consultant

Atlanta, GA

Oct 2025 – Present

- Built production RAG system for financial advising on AWS Bedrock with S3-backed vector storage

### VesselAware

Software Engineering Advisor

Remote

Jan 2024 – Present

- Building AI agents for maritime intelligence; guiding architecture and technical strategy

### PerceptivePanda (acquired by Zapier, Jan 2026)

Co-Founder & CTO

San Francisco, CA

May 2023 – Jul 2025

- Built AI-native customer research platform replacing human-led interviews with LLM-driven micro-interviews at scale
- Core innovation: deterministic control layer wrapping LLMs using dialogue state management grounded in the Questions Under Discussion model from my Stanford research—tracked addressed questions, identified threads for probing, orchestrated LLM calls through a structured state machine
- Raised institutional pre-seed; accepted into StartX '24 summer cohort

### Salesforce

Software Engineering Architect

San Francisco, CA

Jun 2020 – May 2023

Principal Member of Technical Staff

Jun 2020 – Mar 2022

- Led integration of Tableau Ask Data across multiple Salesforce clouds (Sales, Service, Marketing, Commerce)
- Elected to Salesforce AI Northstar Council (2022, 2023)—cross-organization body setting company-wide AI strategy
- Architected next-generation analytics platform spanning CRM Analytics query engine and multiple cloud products
- Led technical integration of Tableau-acquired Narrative Science into the Salesforce ecosystem
- Recognized for Most Patents Filed and Granted across the organization (2022)

### Tableau (via ClearGraph acquisition)

Staff Software Engineer

Seattle, WA / Palo Alto, CA

Aug 2017 – May 2020

- Created the NLP technology behind Tableau's Ask Data feature, deployed to hundreds of thousands of users
- Scaled semantic indexing infrastructure using Elasticsearch to support enterprise-scale data catalogs
- Published "Inferencing underspecified natural language utterances in visual analysis" at IUI 2019
- Filed 9 patents covering NLQ architecture, intent inference, cascading edits, table calculations, and incremental visual feedback

### ClearGraph (acquired by Tableau, 2017)

VP, Natural Language Technology

San Francisco, CA

Nov 2014 – Jul 2017

- Built the core NLP pipeline translating natural language into structured database queries—using Montague Grammar and compositional semantics from my doctoral research, years before BERT existed
- Grew the engineering team from founding stage to 30+ engineers post-acquisition
- Awarded the Celent Model Bank Award (2017) for innovation in financial services technology

## EDUCATION

---

### Stanford University

PhD, Linguistics

Stanford, CA  
2009 – 2015

- Dissertation: *The syntax and semantics of ordinary comparative constructions in English*
- Developed compositional semantic framework using continuation semantics and Montague Grammar
- 5 peer-reviewed publications; Phi Beta Kappa Fellowship (2014); collaborated with Christopher Potts on natural logic

### University of Amsterdam

Fulbright Scholar — Institute of Logic, Language & Computation

Amsterdam, NL  
2008 – 2009

### Northwestern University

BA, Linguistics — Cum Laude with Honors; Phi Beta Kappa

Evanston, IL  
2004 – 2008

## PUBLICATIONS & PATENTS

---

15+ publications ([Google Scholar](#)) · 9 patents ([full listing](#))

- **Djalali, AJ** et al. "Inferencing underspecified natural language utterances in visual analysis." *IUI*, 2019.
- **Djalali, AJ**. "Synthetic logic." *Linguistic Issues in Language Technology*, 2014.
- **Djalali, AJ** & C. Potts. "A Formal Proof System for Natural Logic." Stanford, 2011.
- 9 patents (7 granted + 2 pending) covering NLQ architecture, intent inference, underspecification resolution, cascading edits, table calculations, and incremental visual feedback

## TECHNICAL PROJECTS

---

- **Enterprise Semantic Search Platform** — Production vector search over 400M+ images. CLIP embedding generation on distributed Ray/GPU clusters via Databricks, Qdrant nearest-neighbor retrieval with hybrid dense+keyword scoring at sub-200ms latency. Circuit breakers, exponential backoff, Redis caching for fault tolerance. Batch-level fault isolation and dead-letter tracking for zero-error ingestion across multi-million-image datasets. Full Prometheus/Grafana observability.
- **House Proceedings Corpus** — Multi-service NLP platform for 181M+ tokens of congressional text. FastAPI gateway, Kafka-decoupled ingestion, Elasticsearch full-text + faceted search, Neo4j entity graphs linking speakers/bills/committees. spaCy pipeline extracting named entities, noun phrases, and dependency structures at ingest time. Clean layered architecture with repository pattern and DI. [github.com/alexddjalali/hpc](https://github.com/alexddjalali/hpc)
- **LLM-Driven Interview Platform** — AI agent system replacing human-led customer research with autonomous LLM micro-interviews at scale. Deterministic control layer orchestrating LLM calls via dialogue state machine grounded in Questions Under Discussion theory—tracked addressed questions, identified threads for probing, enforced analytical coherence while preserving conversational fluidity. (PerceptivePanda, acquired by Zapier)
- **Financial RAG System** — Production retrieval-augmented generation pipeline on AWS Bedrock for financial advisory use cases. S3-backed vector storage, document chunking with semantic boundaries, contextual retrieval with source attribution. (Gridline)

## TECHNICAL SKILLS

---

**AI / ML** LLM orchestration & prompt engineering, RAG pipelines, AI agent architectures, dialogue state machines, vector search & embeddings (CLIP, sentence-transformers), semantic retrieval, few-shot / chain-of-thought prompting, evaluation frameworks

**NLP** Formal semantics, Montague Grammar, compositional semantics, natural logic, discourse theory (QUD), dependency parsing, NER, spaCy, pre-transformer NLQ systems

**Distributed Systems** Ray (distributed GPU clusters), Apache Spark, Databricks, Kafka, event-driven architectures, batch & stream processing at scale

**Languages** Python, Go, TypeScript/JavaScript, Bash, Lua, Java

**Data & Search** Qdrant, Elasticsearch, Neo4j, PostgreSQL, Redis, MongoDB, S3, vector databases

**Cloud & Infra** AWS (Bedrock, S3, EC2, Lambda), Docker, Kubernetes, Terraform, Helm, GitHub Actions

**Frameworks** FastAPI, React, Next.js, Node.js, GraphQL

**Observability** Prometheus, Grafana, OpenTelemetry, structured logging, SLO-driven monitoring

## HONORS & AWARDS

---

- StartX '24 Summer Cohort
- Salesforce Most Patents Filed & Granted (2022)
- Salesforce AI Northstar Council (2022, 2023)
- Celent Model Bank Award (2017)
- Phi Beta Kappa Fellowship, Stanford (2014)
- Fulbright Fellowship (2008)
- Phi Beta Kappa, Northwestern
- Demoz Paper Award, Northwestern
- National Merit Scholar Semi-Finalist