

---

---

## DR. SOFIA ATHENIKOS — PhD, Information Science x PhD, Philosophy

---

---

**I ship. I own. I deliver**

Website: <https://athenikos.dev> Email: [sathenikos@gmail.com](mailto:sathenikos@gmail.com)

LinkedIn: <https://www.linkedin.com/in/sofiaphdphd> GitHub: <https://github.com/aifoss>

---

### SUMMARY

---

Versatile senior software engineer with 13+ years of shipping high-impact systems — often end-to-end, often as the sole owning engineer. Strong expertise in backend development, distributed systems, and large-scale data platforms. Track record of solo-owned projects and initiatives across Amazon, Twitter, Bank of America Merrill Lynch, Flipboard, Vericast, and RR Donnelley. Currently building AI-driven applications. Seeking Senior / Staff / Principal Software Engineer roles — AI-driven, with maximum personal space and latitude to ship, own, and deliver.

---

### SIGNAL STATS

---

13+ Years | 9 Companies | 6 Domains | 6+ Projects Owned | 100+ Teams Onboarded | 3,000+ Serviceable Daily Jobs

---

### SELECTED OWNERSHIP PROJECTS

---

- nCountR — Retail Data Pipeline & Analytics Suite RR Donnelley · 2025
  - Reactive Platform — Event-Driven Workflow Orchestration Vericast → RR Donnelley · 2023-2025
  - HDFS Image Analyzer — Storage Analysis & Archival Tool Vericast · 2023
  - Dark Traffic Proxy — Production Traffic Shadowing Service Twitter · 2022-2023
  - Entitlement Data Migration — Hierarchical Restructuring BoA Merrill Lynch · 2017-2018
  - Content Extraction — Core Platform Refactoring Flipboard · 2015-2016
- 

### CAREER SNAPSHOT

---

2024 - 2026	RR Donnelley	Principal Software Engineer
2023 - 2024	Vericast	Principal Software Engineer
2019 - 2023	Twitter	Sr. Software Engineer
2018 - 2019	Morgan Stanley	Associate (Sr. SWE)
2017 - 2018	BoA Merrill Lynch	Consultant (Sr. SWE)
2016 - 2017	Bloomberg	Consultant (Sr. SWE)
2015 - 2016	Flipboard	Software Engineer
2013 - 2015	IPsoft	R&D Engineer (Amelia AI)
2012 - 2013	Amazon	Software Development Engineer

---

### EDUCATION

---

Drexel University	Ph.D., Information Science	GPA 3.85
Drexel University	M.S., Computer Science	GPA 3.85
Temple University	Ph.D., Philosophy	GPA 3.92
Temple University	M.A., Religion	GPA 3.92

---

---

→ Full CV follows on next page

---

---

---

---

**DR. SOFIA ATHENIKOS**  
**PhD, Information Science**  
**PhD, Philosophy**

---

---

**I ship. I own. I deliver.**

Website: <https://athenikos.dev>  
Email: [sathenikos@gmail.com](mailto:sathenikos@gmail.com)  
LinkedIn: <https://www.linkedin.com/in/sofiaphdphd>  
GitHub: <https://github.com/aifoss>

---

---

---

---

## SUMMARY

---

---

Versatile senior software engineer with 13+ years of shipping high-impact systems — often end-to-end, often as the sole owning engineer. Strong expertise in backend development, distributed systems, and large-scale data platforms. Track record of solo-owned projects and initiatives across major employers including Amazon, Twitter, Bank of America Merrill Lynch, Flipboard, Vericast, and RR Donnelley. Currently building AI-driven applications.

Seeking Senior / Staff / Principal Software Engineer roles — AI-driven, with maximum personal space and latitude to ship, own, and deliver.

---

---

---

---

## SIGNAL STATS

---

---

- 13+ years of professional software engineering experience
  - 9 companies (Amazon, IPsoft, Flipboard, Bloomberg, BoA Merrill Lynch, Morgan Stanley, Twitter, Vericast, RRD)
  - 6 industry domains (e-commerce, AI/NLP, media, finance, social media, AdTech)
  - 6+ projects owned end-to-end
  - 100+ internal engineering teams onboarded (Twitter DTP)
  - 3,000+ serviceable daily jobs (Vericast Reactive Platform)
- 
- 

---

---

## WORK EXPERIENCE

---

---

---

---

**RR Donnelley · 2024 - 2026**  
**Principal Software Engineer · Full Time · Remote**

---

---

Continued as a core member of the 3-member R&D (NERD) team (with the addition of a support engineer) after Vericast's Valassis subsidiary was acquired by RR Donnelley, reporting to the (same) VP of Engineering and delivering high-priority, high-visibility projects. Worked cross-functionally to shape the product roadmap and accelerate project delivery.

Key contributions:

- Continued sole engineering ownership of the Reactive Platform — the event-driven workflow orchestration and execution service originally started at Vericast — through full implementation, testing, validation with the first prospective client team, and readiness for production rollout.

- Added comprehensive observability to the Reactive Platform using Grafana, Graphite, and OpenTelemetry, improving system reliability and reducing debugging time.

- Conducted R&D for the nCountR retail analytics project as the sole engineer — from proof-of-concept through full implementation of the end-to-end data pipeline, the journey-transaction correlation algorithm, and the customer journey visualization.

Stack: Java, Scala, Python, SQL, Kafka, Snowflake, S3, Grafana, Graphite, OpenTelemetry

---

## **Vericast · 2023 - 2024**

### **Principal Software Engineer · Full Time · Remote**

---

Joined as one of two principal-level engineers on the core R&D (NERD) team at Valassis Digital Corp, a Vericast subsidiary, reporting directly to the VP of Engineering. Conducted high-priority, high-visibility R&D — establishing proof-of-concepts that guided the product roadmap, then driving the systems and services to full implementation.

Key contributions:

- Developed the HDFS Image Analyzer command-line utility as the sole engineer — a tool that scans the HDFS image (~8M files), applies user-defined filtering predicates, and generates archival recommendation reports. Deployed and used in production by the Platform team to manage storage at scale, contributing to storage cost savings while eliminating arduous manual cleanup effort.

- Engineered the Reactive Platform as the sole engineer — an event-driven workflow orchestration and execution service comprising three core services (Event Generation, Batching, Processing), an Overseer service, client utilities, and a configuration UI. Designed to replace the legacy batch-processing system and reduce data processing latency for the 3,000+ daily jobs run by internal engineering teams.

Stack: Java, Kafka, HDFS

---

## **Twitter · 2019 - 2023**

### **Sr. Software Engineer · Full Time · Remote**

---

Member of the Capacity Correctness & Reliability team within Core Infra | Performance Engineering. Team focus: Implementing critical services for Twitter service reliability and efficiency, development velocity, and capacity planning.

Key contributions:

- Refactored the "Perfy" codebase (Scala) — Twitter's performance/load/regression testing service — and added a broad set of new capabilities: health-check client for target service instances, encryption of API connections, dynamic target-instance lookup via platform specs, custom target mapping, metric collection and storage across test types, a metrics dashboard, email/Phabricator notifications with test details and results, and real-time metric push.

- Refactored the "RedCurve" codebase (Scala) and added significantly enhanced dashboards — including an upgraded taskrun dashboard, a service-data dashboard, a historical dashboard for visualizing service stats over time, and the ability to extract and store specs for target instances and taskruns.

- Refined the "PagerDuck" codebase (Go) and implemented new features.

- Created shared libraries by refactoring common code out of Perfy and RedCurve — reducing duplication and improving maintainability.

- Built and rolled out Dark Traffic Proxy (DTP) — a new production traffic shadowing service (Scala) that replaced the legacy Diffy Proxy. Solo-designed and implemented the full service, then onboarded 100+ internal engineering teams to production. DTP was officially designated as Twitter's recommended traffic-shadowing service.

Team recognition: "Improving Our Focus" Award at Twitter OneTeam 2020.

Stack: Scala, Go, TypeScript, React, Finagle, Finatra, Aurora, Mesos, Grafana

---

**Morgan Stanley · 2018 - 2019**  
**Associate (Senior Software Engineer) · Full Time · New York, NY**

---

Member of the Quantitative and Structured Products group within Institutional Securities Technology | ESTAR. Team focus: Automated pricing system for equity derivative trading.

Key contributions (built solo across all components listed below):

- Implemented a modified pricing-context re-allocation algorithm — improving how the pricing system reuses and reassigns computational context across pricing runs.
- Devised a pricing-result diagnostics data exchange format (XML/JSON) and implemented the data mapping classes that convert pricing-result objects into diagnostics objects serializable/deserializable to and from the format.
- Implemented data mapping classes for converting pricing-result objects for consumption by an internal third-party application, with JSON serialization/deserialization.
- Built a KDB client service — with associated data mapper, logger, reader, writer, and factory classes — for performing read/write operations against KDB (the time-series database used widely in finance).
- Implemented REST API endpoints for storing and retrieving serialized diagnostics, pricing, and testing data in KDB, along with DAO classes that process pricing and test results and call those endpoints.
- Implemented a parameter override mechanism for repricing, plus numerous additional features that enhanced the automated pricing system overall.

Stack: Java, Spring Boot, KDB, JSON/XML, Swagger

---

**Bank of America Merrill Lynch · 2017 - 2018**  
**Consultant (Senior Software Engineer) · Contract · New York, NY**

---

Sole-engineer engagement for the Equity Linked Technology (ELT) group at BoA Merrill Lynch. Took on a long-standing, high-stakes, irreversible data restructuring and migration project that the group had previously been unable to staff.

Key contributions:

- Restructured ELT's role-based access (RBA) entitlement data — used by thousands of traders, engineers, and executives — from a flat, unstructured list into a hierarchical model, while building the bidirectional sync pipeline that eliminated long-standing inconsistency risk between the UI/service and the database.
- Devised and executed a 50+ step, two-day production migration with zero data loss and no production disruption. The change was irreversible — once executed, no rollback was possible — and required step-by-step validation throughout.

- Delivered the full project in 10 months — two months ahead of the originally projected 12-month timeline — while incorporating additional enhancements beyond the original scope at the request of the primary business stakeholder. Recognized by ELT executives as a successful outcome on a project that had previously resisted resolution.

Stack: Java, JAX, Spring Boot, iBATIS, Oracle DB, Toad

---

**Bloomberg · 2016 - 2017**  
**Consultant (Senior Software Engineer) · Contract · New York, NY**

---

Member of the Data Intelligence team within the R&D Consumer Media Services group. Worked on Bloomberg's entity tagging, personalization, and recommendation systems — the backend systems that power content tagging and personalized content delivery for consumer-facing media products.

Key contributions:

- Revised and extended the data ingestion pipeline to ingest data from new data sources, broadening the input surface available to downstream tagging and recommendation components.
- Revised and extended the tagging, personalization, and recommendation systems to accommodate new tag types, enabling richer classification of ingested content.
- Refactored the tagging system code to use new terminologies — accommodating terminological evolution and extension while improving clarity and consistency across the codebase.

Stack: Java, Spring, Spring Boot, Tomcat, Hadoop, HBase

---

**Flipboard · 2015 - 2016**  
**Software Engineer · Full Time · New York, NY**

---

Member of the Content Extraction team within the Platform Engineering group — working on the system that defines Flipboard's product: extracting content from numerous sources and reformatting it for delivery across web, mobile, and other devices.

Key contributions:

- Solo-built or rewrote multiple core components of the content extraction stack — including the central extraction microservice (~57 classes), the core extraction package (~81 classes), a new schema and validation microservice (~39 classes), and a rewritten config and management microservice (~46 classes) with a Git-based file management system newly implemented using JGit.
- Implemented supporting tooling for the extraction system: a PDF content handler that converts PDF sources to well-formed HTML; a concurrent asynchronous test client for realistic load and behavior validation; and cross-version result validators that catch regressions between releases.
- Drove deep refactoring across the content extraction codebase proactively — not as an assignment, but because the existing code was getting in the way of safely extending the most critical part of the product.
- In numbers (across the entire work): 1,200+ commits, 250K+ lines added, 154K+ lines deleted.

Stack: Java, Spring, Tomcat, S3, XML/DOM

---

**IPsoft · 2013 - 2015**  
**R&D Engineer · Full Time · New York, NY**

---

Member of the core R&D team (~10 engineers) reporting directly to the CEO, building Amelia — an AI/NLP-based virtual customer service agent and the company's emerging flagship product. Worked primarily on the engineering infrastructure that powered and supported Amelia's conversational behavior and reasoning: the conversation workflow execution module and the semantic knowledge management module.

Key contributions:

- Developed and maintained the BPM module — the engine responsible for capturing, executing, and maintaining real-time conversational issue-resolution workflows between Amelia and the user. Each conversation was represented as a graph, with the module managing its creation, execution, and management lifecycle. Collaborated with one other engineer.
- Devised and implemented algorithms supporting conversational workflow execution and management, including answer-type handling, multiple-choice handling, and graph merging (for consolidating structurally similar conversation flows).
- Also contributed to the semantic knowledge management module — the framework used to create and operate against a semantic knowledge base storing entities, relationships, and inference rules. Authored inference rules over the knowledge base, improving the system's ability to derive and infer facts during conversations.
- Collaborated directly with the CEO on R&D direction — including studying and adapting research papers (e.g., on graph merging) into concrete enhancements implemented for and integrated into Amelia.

Stack: Java, Spring, Tomcat, Activiti (BPM), SQL, Hibernate, Cassandra, Hazelcast, ActiveMQ

---

**Amazon · 2012 - 2013**  
**Software Development Engineer · Full Time · New York, NY**

---

Founding member of the New York team built to launch the Amazon Marketing Services (AMS) Advertising Platform — a new self-service product geared toward Amazon sellers (not buyers), enabling them to create brand pages, post advertising messages, and analyze their marketing campaigns. The platform was co-developed by an existing San Francisco team and the newly created NY team; later, ownership was divided such that SF focused on Amazon Pages while NY focused on Amazon Posts and Amazon Analytics.

Key contributions:

- Contributed across the full stack, with Java/Spring-based backend services, AWS-backed storage (S3, RDS), and JavaScript/jQuery on the frontend — typical of e-commerce self-service platforms of the period. Built web services according to multi-tier architecture for Amazon Posts and Amazon Analytics, supporting the successful launch of the AMS platform.
- Restructured and refactored the codebase proactively according to the principle of clean separation of responsibilities, improving understandability, maintainability, and extensibility.
- Later belonged to a 3-engineer API focus subgroup within the NY team, designing, implementing, and maintaining public-facing REST APIs.

Stack: Java, Spring, REST, AWS (S3, RDS), SQL, JavaScript, jQuery

---

---

## EDUCATION

---

---

Drexel University · Philadelphia, PA  
Ph.D., Information Science  
GPA: 3.85 / 4.0

Drexel University · Philadelphia, PA  
M.S., Computer Science  
GPA: 3.85 / 4.0

Temple University · Philadelphia, PA  
Ph.D., Philosophy  
GPA: 3.92 / 4.0

Temple University · Philadelphia, PA  
M.A., Religion  
GPA: 3.92 / 4.0

---

---

## TOOLKIT

---

---

### AI / Gen AI / ML / DL

LLM (Claude, GPT) · LangChain4J · Vapi · Deepgram · 11Labs · Agentic Architecture · Prompt Engineering · Python (Scikit-Learn, NumPy, SciPy, Pandas) · TensorFlow · Keras · DL4J · FloydHub · Dato · MATLAB · Octave · Weka

### PROGRAMMING LANGUAGES

Java · Scala · Python · Go · TypeScript · Kotlin · Rust · R · C++ · C · Solidity · Lisp · Scheme

### BACKEND & MICROSERVICES

REST API · Spring · Spring Boot · Spring Data · Spring Cloud · Netflix OSS (Eureka, Hystrix, Ribbon, Zuul) · Tomcat · Servlet · JSP · Java Handler

### DATA PLATFORMS & STREAMING

Kafka · Spark · Hadoop · MapReduce · HDFS · HBase · ElasticSearch

### DATA STORAGE, ORM, & CACHING

SQL · PostgreSQL · MySQL · Oracle · KDB · Cassandra · MongoDB · Snowflake · Supabase · iBATIS · JDBC · JPA · Hibernate · Quill · Toad · Hazelcast · H2 · Redis

### ARCHITECTURE & DISTRIBUTED SYSTEMS

Event-Driven Architecture · Agentic Architecture · SOA · Multi-Tier Server-Client Architecture · Microservices · REST · Message-Driven Systems · Workflow Orchestration · BPM (Activiti)

### MESSAGING & STREAMING

Kafka · ActiveMQ · RabbitMQ

### CLOUD & CONTAINERIZATION

AWS (S3, RDS) · Docker · Kubernetes

### OBSERVABILITY & MONITORING

Grafana · Graphite · OpenTelemetry

#### DATA EXTRACTION & PROCESSING

XML · DOM · SAX · JAXB · JSON · Jackson · Gson

#### FRONTEND & FULL-STACK WEB DEVELOPMENT

HTML · CSS · JavaScript · TypeScript · React · jQuery · AJAX · Bootstrap · KnockoutJS · ExtJS · Swagger · Node.js · Express.js · Zod · Ngrok · Railway · Next.js · Tailwind CSS · Framer Motion · Lucide React · JetBrains Mono · Geist · React Hook Form · Formspree · Vercel

#### TESTING & QA

JUnit · ScalaTest · Mockito · PowerMock · EasyMock

#### BUILD & CI/CD

Maven · Gradle · Ant · Ivy · Pants · Bazel · Cargo · Jenkins · Hudson · GitLab

#### VERSION CONTROL & CODE REVIEW

Git · GitHub · GitLab · BitBucket · SVN · Perforce · Phabricator · ReviewBoard

#### STATISTICS & DATA SCIENCE

R · RStudio · SPSS

#### SEMANTIC WEB & RULE ENGINES

OWL · RDF · SPARQL · Protégé · Sesame · PowerLooms · Drools

---

---

**END OF CV**

---

---