

PATRICK FARRELL

ANN ARBOR, MI

317-281-3881 | pat@patf.net

linkedin.com/in/patrickfarrell | github.com/pfarrell | patf.net

SUMMARY

Experienced software engineer and leader with a deep background in getting data to power businesses. Extensive experience identifying critical business issues and opportunities, organizing ideas into projects, and assembling and leading teams to deliver those projects as shipped software. Recognized for the ability to clearly communicate at many different levels of understanding, from customers to system architects. Demonstrated ability to assemble and lead teams to solve difficult problems by extending existing codebases or starting new projects from scratch.

Languages:	Python, Kotlin, Java, Ruby, C#, JavaScript, SQL, Regex, Shell
Environments:	GNU/Linux, Windows, Solaris, OSX, Amazon Web Services, Google Cloud, Microsoft Azure, Kubernetes, Docker
Databases and Frameworks:	MSSQL Server, Postgres, MySQL, Oracle, SQLite, Elasticsearch, Redis, MongoDB, Cassandra, Hadoop, DynamoDB, Redshift, RabbitMQ, Kafka, Zookeeper, Apache, Nginx, Jetty, JBoss, Spring, Ruby on Rails, Dropwizard, Pandas
Verticals:	Automated machine learning (ML), database integrated development environments (IDEs), health care, fintech, online video, advertising, email marketing, credit card and ACH, auto industry, eGov, document imaging systems, and optical and intelligent character recognition (OCR/ICR)

EXPERIENCE

DATAROBOT, Boston, MA

2/2019 – 9/2022

Artificial intelligence (AI) cloud platform.

Director of Engineering

- Led a development team to create a collaborative data catalog as a new pillar in the product ([AI Catalog](#)). Designed the application programming interface (API) and integrated Elasticsearch into software-as-a-service (SaaS) and on-premises deployments to support multi-language indexing and advanced search. Project was a major sales driver for expanding DataRobot to support data science teams (Python, Elasticsearch, and MongoDB).
- Created and led a data engineering team in support of a 200+ person project using machine learning to model the spread of COVID-19 ([Decision Intelligence Platform](#)). Project made crucial recommendations to Moderna and the Department of Health and Human Services influencing the US pandemic response. System landed a major contract with the US government, and received a [patent](#) in June 2022 (Python, Postgres, Jenkins, and ML).
- Identified major scaling bug in Google's BigQuery JDBC driver. Worked with Google to formulate a fix ([GBQJ-581](#)). Designed and deployed custom workaround, saving \$1.5M in recurring subscriber revenue (Python and Java).
- Designed and coded a project to allow DataRobot to safely store encrypted data with support for key rotation and expiration. Project became an enabling factor and building block for transparent encryption services throughout the product (Python and MongoDB).

CURSOR, San Francisco, CA

8/2017 – 2/2019

Collaborative SQL tool.

Co-Founder and Head of Engineering

- Technical lead of a startup developing an application capable of capturing, organizing, and recommending queries and code. Responsibilities included architecture, coding, deployment, project management, hiring, customer support, and any of the myriad things that happen in a startup.
- Architected and developed desktop application to connect to data stores and capture metadata about the interactions while respecting data security boundaries (Kotlin, Elasticsearch, Postgres, SQLite, and React).
- Startup was acquired by DataRobot while raising Series A round.

GRAND ROUNDS, San Francisco, CA

1/2016 – 8/2017

Health care navigator.

Engineering Manager

- Organized data engineering team to take data preparation load off the data science team.
- Refactored legacy codebase to allow the import of millions of patient records, immediately allowing a 2x increase in data (Ruby and Java).
- Implemented data pipelines to perform entity resolution in DynamoDB (AWS, PySpark, and Ruby).

PROSPER MARKETPLACE, San Francisco, CA

6/2014 – 1/2016

Peer to peer lending platform.

Architect

- Designed and implemented pub/sub service to provide event-based synchronization between two systems previously synced via batch processes. Major step in dropping business accounting processes from days to hours. (C#, RabbitMQ, and MSSQL Server)
- Architect on a full system rewrite in Java designed to move system to microservices architecture. (Java, JBoss, Spring)
- Designed and led team to code geolocation, income verification, and a proprietary confidence scoring as microservices. (Java, Spring, and ML)

REALGRAVITY, San Francisco, CA

9/2012 – 6/2014

Ulive.com, online video syndicator.

Principal Engineer

- Hired and led a five-person data engineering team. Replaced homegrown batch processing system with one based on the Lambda architecture to handle streaming and batch workflows. (Ruby, Hadoop, and Redis)
- Designed and coded a REST API to execute queries across various streaming and historical data stores, improving efficiency of support team and speeding ticket resolution. (Ruby, and MySQL)
- Designed and built page level video analytics pipeline which became a core direction of the business (Ruby, Hadoop, Hive, d3, JavaScript, and MySQL).

EXACTTARGET, Indianapolis, IN

4/2007 – 9/2012

Email service provider.

Principal Engineer

Senior Software Engineer

- Sharded a live single point of failure (SPOF) database (DB) with zero downtime. Data included all authentication and authorization data used by every subsystem. Major factor in getting system to 99.98% actual uptime in 2011 and reducing continuous CPU load on DB server from 80% to 20%, and was recognized with an award from the Chief Executive Officer (CEO) (C#, Visual Basic 6, and MSSQL Server).
- Led an eight-person team to add major security enhancements to the product including SAML and SSO support, two-factor authentication, IP Whitelists, and machine fingerprinting (Java, C#, C, and JavaScript).
- Designed and implemented a rules engine to apply arbitrary functions against incoming data sets to identify suspicious system behavior. Processed over 20 million data points per day (C#, SQL, and Hadoop).
- Designed and coded signature system for encrypted coupon application in response to a decryption flaw, stopping a multi-million dollar revenue loss (C#, HMAC, and AES).
- Member of a six-person team to create [DataExtensions](#), which added the ability for customers to import data to custom tables and have those entities exposed as first class system objects. Major driver of sales and system scalability. Still in use as part of the Salesforce Marketing Cloud (C# and SQL).

Employment information prior to 2007 available on [LinkedIn](#) or on request

EDUCATION

Bachelor of Arts, Mathematics, [INDIANA UNIVERSITY](#), Bloomington, IN

Bachelor of Arts, Psychology, [INDIANA UNIVERSITY](#), Bloomington, IN

Graduate Computer Science work, [PURDUE UNIVERSITY](#), Indianapolis, IN

INVENTIONS

- US Pat. [20220199266](#), June 23, 2022: Systems and methods for using machine learning with epidemiological modeling.
- Created [DBMulticast](#), a multiplexing database query tool that reduced data investigations at ExactTarget from tens of minutes to seconds.
- Created [PShare](#), a streaming music website in continuous development since 2002.
- Numerous [personal projects](#) dealing with search engines, web crawlers, machine learning, photography, streaming data, diagramming, personal productivity, and legacy data systems.