Patrick Farrell

Ann Arbor, MI
317-281-3881 | mr.patfarrell@gmail.com
linkedin.com/in/patrickfarrell | github.com/pfarrell

Summary

Experienced software engineer and leader with a deep background in getting data to power businesses, building and leading teams, and crafting software from idea to delivery and beyond. Proven history identifying critical business issues and opportunities, organizing ideas into projects, and delivering on time. Recognized for clearly communicating at different levels of understanding.

Languages: Python, Kotlin, Java, JavaScript, Node.js, Ruby, C#, C, SQL, Bash, PowerShell

Environments: GNU/Linux, Windows, OSX, Amazon Web Services, Google Cloud, Microsoft Azure, Kubernetes, Docker

DBs and Tools: PostgreSQL, MSSQL Server, DynamoDB, SQLite, ElasticSearch, Redis, MongoDB, Cassandra, Redshift, RabbitMQ, Kafka, Nginx, Ruby On Rails, Pandas, AWS Lambda, Terraform, dbt, Github

Experience

INDIGO INSURANCE, Nashville, TN 03/2023 - 01/2024
Medical Malpractice Insurance Carrier
Founding Engineer

- Architected and delivered the MVP, established project management, CI/CD, documentation and automated release management. Hired the engineering team creating the new business (*React, Python, AWS, Jira, Github*)
- Identified and fixed issues in a proprietary risk assessment ML algorithm. Automated underwriting decisions, launched Indigo on schedule, and processed new business 14x faster (*React, AWS Lambda, DynamoDB, ML*)
- Created an entity resolution system to identify physicians from minimal user-supplied information. Achieved 95% accuracy in exact matches (*Python, AWS Lambda*)

CAREER BREAK 10/2022 - 03/2023
Personal Sabbatical

DATAROBOT, Boston, MA 02/2019 - 10/2022
Automated Machine Learning Platform
Director of Engineering

- Invented and produced a collaborative data catalog. Major advancement to the existing product to support data science teams generating over $1M in new revenue (*Python, Elasticsearch, RabbitMQ, MongoDB*)
- Founded and led a 20 person data team for a project modeling the pandemic and simulating policy changes. Project made major recommendations and influenced the delivery of the vaccine (*Python, Postgres, Jenkins, ML*)
- Identified a major scaling bug in Google's BigQuery JDBC driver. Advised Google on how to formulate a fix (GBQJ-581). Saved $1.5M in recurring subscriber revenue (*Python*, *Java*)

CURSOR (acquired by DataRobot), San Francisco, CA 08/2017 - 02/2019
Collaborative coding product
Co–Founder, **Head of Engineering**

- Technical co-founder of the startup. Designed and delivered the product and infrastructure, organized the company roadmap, hired the team, and handled customer support to create the company
- Conceived and built an application to automatically organize data analyst's work, increasing productivity by 40% (*Electron*, *Kotlin, Elasticsearch, Postgres, SQLite, React*)
- Devised automated local dataset generation integrated with data capture speeding common data analyst tasks by 20% (*Kotlin, Sqlite*)

GRAND ROUNDS, San Francisco, CA 01/2016 - 08/2017

Health care navigator

Engineering Manager

- Started a data engineering team reducing data science involvement in data preparation by 75%
- Rewrote data processing to support the import of millions of patient records, allowing a 2x increase in data (*Ruby, Java*)
- Coded data pipelines to support entity resolution in DynamoDB enabling scalable data queries that were 3x faster (*AWS, PySpark, Ruby*)

PROSPER MARKETPLACE, San Francisco, CA 06/2014 - 01/2016

Peer-to-peer lending platform

Architect

- Built a pub/sub service providing event-based synchronization of loan servicing systems. Major step in making business accounting processes 4x faster (*C#, RabbitMQ, MSSQL Server*)
- Conducted a major system evolution to move to a microservices architecture and deliver feature releases 10x faster with continuous deployment (*Java, JBoss, Spring*)
- Crafted geolocation, income verification, and a proprietary confidence score as microservices increasing the number of serviceable loans by 30% (*Java, Spring, ML*)

REALGRAVITY, San Francisco, CA 9/2012 - 06/2014

Ulive.com, online video syndicator

Principal Engineer

- Revamped a video traffic processor handling 4 terabytes of data per day and supporting real-time reporting (AWS, *Ruby, Hadoop, Redis*)
- Implemented a query tool to marry streaming and historical data stores into an easy to use product, improving the efficiency of the support team and speeding customer issue resolution by 65% (*Ruby, MySql*)
- Generated custom page level video analytics pipeline creating new business opportunities (*Ruby, Hadoop, Hive, d3, JavaScript, MySQL*)

EXACTTARGET, Indianapolis, IN 04/2006 - 09/2012

Email marketing

Principal Engineer

- Sharded application's central database with zero downtime or data loss. Major factor in getting system to 99.98% uptime and reducing CPU load 4x (*C#, MSSQL Server*)
- Led an eight-person team to add major security enhancements to the product including SAML and SSO support, two-factor authentication, and machine fingerprinting reducing security incidents by over 50% (*Java, C#, C, JavaScript*)
- Produced a rules engine to identify suspicious system behavior. Processed over 20M data points per day and exposed malicious customer behavior (*C#, SQL, Hadoop*)

Employment information prior to 2006 available on LinkedIn or on request

Education

- Bachelor of Arts, Mathematics, Indiana University
- Bachelor of Arts, Psychology, Indiana University
- Graduate Computer Science work, Purdue University

Projects

- US Patent: Systems and methods for using machine learning with epidemiological modeling (Patent #20220199266)
- Created DBMulticast, a multiplexing database app which reduced data investigations at ExactTarget by 600% (*C#, SQL*)
- Wealth of side projects dealing with search engines, music streaming, web crawling, machine learning, photography, streaming data, data visualization, personal productivity, and data querying