Patrick Bourke

Seattle, WA

pbourke@gmail.com
linkedin.com/in/patrickbourke
github.com/pbourke

SUMMARY

Twenty years of software engineering experience spanning large-scale production systems, quantitative finance and applied AI research. Principal Software Engineer at Microsoft Research on the team behind GraphRAG, focused on graph techniques, large language models and agentic AI systems. Previous roles at Amazon and Two Sigma Investments.

SKILLS

- *Machine Learning*: LLMs, prompt engineering, RAG and GraphRAG, graph theory and graph statistics, dense vector spaces, NLP, search and sentiment analysis.
- Expert Python programmer; extensive experience with data engineering pipelines for machine learning, graph analytics and large-scale data processing.
- Experienced in C#, Java, C++, JavaScript, HTML, CSS and SQL in Linux and Windows environments.
- *Cloud / MLOps*: Azure, Azure Machine Learning, AWS.
- *Software Engineering*: requirements gathering, software design, best-practices implementation — have led the design and implementation of several large projects.
- *Databases*: project experience with SQL Server, PostgreSQL, MySQL, SQLite.
- *Linux*: 20+ years of experience; installation, configuration and maintenance of OS and common applications, multi-tier system administration.

EXPERIENCE

Microsoft

5/2018 – present · Redmond, WA

Principal Software Engineer — Microsoft Research, Agentic Data Memory

- Designed data engineering pipelines and graph models for analysis of changing workplace patterns; starting from anonymized telemetry, builds hundreds of thousands of graphs and conducts temporal and network analysis.
- Developed large-scale feature extraction from millions of graphs, producing dense matrices that support research into cross-graph structural patterns.
- Built LLM evaluation frameworks and analysis pipelines using Azure Machine Learning and Azure OpenAI (GPT-3 through current models), including LLM-based scoring, side-by-side model comparison and automated rubric generation.
- Applied GraphRAG — Microsoft's open-source graph-based RAG framework — to code analysis, building a pipeline combining graph indexing, LLM-generated summaries and semantic search.
- Developed synthetic data generation pipelines using embedding models and LLMs to produce privacy-preserving training datasets.
- Sped up a core component in a reinforcement learning project by 60% through C++ optimization, enabling continued research progress.

Microsoft

6/2015 – 5/2018 · Redmond, WA

Principal Software Engineer — Azure Capacity Management

- Engineering lead for Azure Capacity Planning and Management systems.
- Worked in both the short-term (online resource allocation) and medium-term (hardware purchase/allocation) areas.
- Technologies used: C#/.NET, SQL Server, Azure Cloud Services.

Two Sigma Investments

3/2013 – 4/2015 · New York, NY

Software Engineer, Vice President

- Technical lead for the broker relationship (Alpha Capture) platform of a quantitative hedge fund.
- Changed the platform architecture to be simpler and more maintainable.
- Drove the creation of a company-wide technology solutions manual as a member of the firm's architecture working group.
- Technologies used: Java, SQL Server, RabbitMQ, Apache ZooKeeper, AWS.

Amazon.com, Inc.

10/2009 – 2/2013 · Seattle, WA

Senior Software Development Engineer — Product Imaging

- Led design and implementation of new product imaging platform using Spring, Hibernate, Oracle and AWS.
- Created and maintained software for embedded device project in Dynamic C.

Amazon.com, Inc.

9/2007 – 10/2009 · Seattle, WA

Software Development Engineer — Gift Cards

- Owned software behind worldwide delivery of Amazon gift certificates: email, postal and downloadable (PDF).
- Designed and implemented next-generation gift certificate delivery stack using Spring, Hibernate, Oracle and AWS.
- Improved peak email delivery efficiency across multiple dimensions: throughput, latency, TCO and number of support incidents.

BPS, Inc.

2/2006 – 7/2007 · Toronto, ON

Lead Developer — Product Team

- Led a team of 5 developers in the design and implementation of new features for the company's financial risk-management product.
- Designed and implemented a RESTful interface for the company's product; introduced Spring integration tests, measurably improving code quality.
- Technologies: Tomcat, Hibernate, Spring, Sybase, DB2, Oracle and JBoss Rules.
- First six months included J2EE integration work on a project for a top-tier Wall Street client.

ExtenSys, Inc. 11/2004 – 2/2006 · Toronto, ON

Software Developer

- Designed and implemented an OLTP project for Ontario's electrical utilities (Spring, Tomcat, Hibernate and JBoss).
- Built, tested and deployed against PostgreSQL and Oracle simultaneously.

Semaview, Inc. 5/2003 – 10/2004 · Toronto, ON

Intermediate Software Developer

- Wrote the server portion of the company's eventSherpa calendar product (PHP, Apache, MySQL).
- Introduced high-value features including RSS feeds of calendar data and a mobile version of the site.
- Designed and implemented a J2EE online training application using Tomcat and Spring.

EDUCATION

University of Waterloo 2002 – 2003

Hons. B.Math (Computer Science)

University of Waterloo 1996 – 2001

B.A. (English Literature)

PUBLICATIONS

- Helm, H.S., Basu, A., Athreya, A., Park, Y., Vogelstein, J.T., Priebe, C.E., Winding, M., Zlatic, M., Cardona, A., **Bourke, P.**, et al. (2023). Distance-based positive and unlabeled learning for ranking. *Pattern Recognition*, 134, 109085.
- Zuzul, T., Cox Pahnke, E., Larson, J., **Bourke, P.**, et al. (2024). Dynamic Silos: Increased Modularity and Decreased Stability in Intraorganizational Communication Networks During the COVID-19 Pandemic. *Management Science*, 71(4), 3428–3448.