# Wayne Krug

# **Technical Skills**

- Languages: C, C++, Java, Scala, Python
- Libraries and frameworks: Hadoop, MPI, Apache Lucene, SciPy, NumPy
- Operating systems: Linux (Red Hat Enterprise, CentOS, Fedora, Ubuntu, Debian)
- Domains: radar signal processing, real-time systems, distributed systems, natural language processing, machine learning
- Hardware: embedded development, production support and test, field support

# Experience

# Senior ML Software Engineer, Twitter, New York, NY

Software engineer on the Natural Language Processing Signals (NLPS) team, part of the Cortex machine learning organization. This team focuses on extracting signals from text using ML models and providing these signals to other teams within Twitter to use in their products. Signals are provided via services running in the Twitter data centers. It is composed of both machine learning engineers (MLEs) and machine learning software engineers (MLSWEs).

- Designed and implemented a service to link named entities found in Tweets to real-world entities in the Wikidata knowledge graph using a combination of online and offline ML models
- Designed and implemented a service to host a next-generation named entity recognition model
- Designed a system to handle high request rates (>1MM RPS) to NLPS services while maintaining low request latency (<20 ms)
- Led the team's efforts to comply with regulatory requirements
- Collaborated with the Recommendations team to integrate NER into their Entity Real Graph system
- Mentored two new-to-Twitter MLSWEs (one on NLPS, one on another team)
- Mentored MLE team members on Twitter internal systems and best practices for production systems
- Contributed small patches to core internal systems when needed to unblock NLPS projects
- Developed dashboards for monitoring and serviceability of NLPS services
- Configured and managed load testing configurations for production services
- Wrote and maintained production batch processing jobs in Scala and Scalding
- Developed best practices for synchronization of data between on-premises storage and Google Cloud Platform

Principal Engineer/Lead Backend Engineer, Jetblack, New York, NY Jul 2018 – Jan 2019

Lead a team of 3 responsible for the development and maintenance of the company's customer-related systems (profiles, payments, and messaging) and logistics systems (source management, inventory management, price monitoring, and delivery estimation).

- Lead the design and development of a system for streamlining the workflow of customer agents and logistics sourcers for acquiring items ordered by the customer and providing estimated delivery dates
- Helped define company-wide coding and documentation standards
- Managed the transition of engineering responsibility from outside contractors to the in-house team

# Senior Software Engineer, Bloomberg, LP, New York, NY

Developed various pieces of C++ and Java software as part of the machine learning (ML) platform team of Bloomberg Law. Software included both ML components built in to products made by other teams and internal services that other teams could use for various ML tasks.

- Lead engineer for Draft Analyzer, a hybrid C++/Java system for characterizing the language of legal documents.
- Planned and initiated a long-term redesign of the Draft Analyzer system to take it from the MVP state to a stable, resilient production system.
- Led the transition of the ML team's Java services from manual build and deployment to an automated CI/CD system.

#### Senior Software Engineer, Language Computer Corporation, Richardson, TX April 2012 - Oct 2016

- Developed natural language processing (NLP) applications in Java for research, prototype, and production. • Architect/lead engineer for Estesa document summarization system built with Java and AngularJS. Demon-
- strated product at a Navy trade show in June 2015.

Jan 2019 - Feb 2023

Oct 2016 - Jul 2018

- Architected common Java system for REST applications and a simple yet fast Bayesian profiling system, both implemented across several company projects.
- Created a Java research application "Ghostwriter" to profile and identify document authors.
- Conducted/published research on machine learning models used to identify non-content words in documents.
- Developed "Capable", a Java application to compile cultural profiles of Twitter hashtag usage.
- Built Java/Python tools to automatically download and index data from Twitter and the Google News RSS feed.
- Developed JNI adapters to connect Java application code to C-based machine learning libraries.

Various Engineer Roles, Raytheon Company, Dallas, TX	Sep 2002 - April 2012
Senior Systems Engineer	May 2008 – Apr 2012
Systems Engineer	July 2006 – May 2008
Electrical Engineer	Sep 2002 – July 2006

#### Project highlights:

### P-8 Advanced Airborne Sensor (and its predecessor AN/APG-149 surveillance radar)

Key contributor on the development and testing of signal processing algorithms. Development (in C) targeted realtime embedded Linux on POWER and x86-64 commercial server platforms. Additional work in Python, Fortran, and Matlab.

- Modified and tuned the Red Hat's Enterprise Linux kernel to add real-time capabilities.
- Refactored legacy C code increasing throughput by >300% and reducing memory use by >90%.
- Drove a well-received trade study of clustered SAN filesystems (GPFS, Lustre, GlusterFS).
- Architected a subsystem to efficiently pack/unpack variable-length binary system messages.
- Researched and benchmarked various processor architectures and wrote findings for AAS proposals.
- Developed platform-specific optimizations for the signal processor's common math library.
- Created a job management framework to speed the creation of cluster processing jobs used to test new algorithm iterations.
- Authored/coauthored multiple internal whitepapers and presented at 2007 annual meeting.
- Made regular presentations on work at quarterly analyst meetings.
- Provided real-time data QC and technical support for a two-week field experiment

#### F-22A Radar

Supported the production and field deployment of the power supply for the F-22A radar. Performed hardware failure analysis and test, evaluation of replacement parts during production, and test software development.

- Identified and fixed a long-standing bug in the power supply controller software that eliminated some intermittent false test failures and reduced unnecessary retesting.
- Significantly reduced hardware testing process cycle time saving >\$185K
- Participated in a high-profile field failure investigation.

# Education

Bachelor of Science in Electrical Engineering, Florida Atlantic University - Boca Raton, FL,

# Publications

- W. Krug and M. Tomlinson, Automated Non-Content Word List Generation Using hLDA, in The 26th Annual Florida Artificial Intelligence Research Society (FLAIRS) Conference, 2013
- M. Tomlinson, W. Krug, D. Hinote, and D. Bracewell, *#impressme: The Language of Motivation in User Generated Content*, in The 15th International Conference on Intelligent Text Processing and Computational Linguistics (CICLing), 2014 (accepted, to be published in April 2014)
- M. Tomlinson, D. Bracewell, W, Krug and D. Hinote *#mygoal: Finding Motivations on Twitter*, in the Proceedings of the 9th Language Resources and Evaluation Conference (LREC), 2014 (accepted, to be published in May 2014)
- L. Barrett, W. Krug, Z. Lu, K. Martin, R. Martin, A. Ortan, A. Pradhan, A. Sherman, M. Sherman, R. Smey, and T. Wenzel Civil Asset Forfeiture: A Judicial Perspective, Data For Good Exchange, 2017