Joe Bylund

joseph.bylund@gmail.com
791 Tremont Street, E509, Boston, MA 02118
347-829-5863

Moat New York, NY

2013 — Present
Senior Data Scientist & Back-end Engineer

- Designed and implemented distributed, fault tolerant ETL reducing cost by an order of magnitude, increasing reliability and reducing processing time from ~10 hours to ~1 hour, making data available to clients far earlier in the day (using python, SQS, Redis, PostgreSQL).
- Numerous data driven API improvements which lead to 3-4x improvement in API latency as well as maximum request size (PHP, CakePHP).
- Contributed improvements to ORM (CakePHP) and core PHP in order to reduce the number of queries necessary to render a page by 5x (decreasing page load time by ~3x) (CakePHP, c).
- Standardized deployment framework used to deploy thousands of servers of ~30 different roles to AWS (AWS, EC2, boto3).
- Migrated primary non-statistical database (users, accounts, metadata) from MySQL to PostgreSQL improving uptime & flexibility (MySQL, PostgreSQL, foreign data wrappers).
- Migrated primary statistical database (500 million rows/day) from non-first normal form to first normal form schema improving query latency, reducing storage demands and increasing throughput.
- Architected and implemented sophisticated message routing system which is responsible for moving ~40 billion events per day from our pixel servers to our real time processing servers while balancing CPU and memory constraints (c++).
- Architected and prototyped massively parallel decentralized data lake using AWS lambda and S3 for cost effective storage and low latency and cost effective queries (AWS lambda, python, PostgreSQL).

Columbia University New York, NY

2007 — 2013, GPA 3.9/4.0
Integrated Program In Cellular, Molecular and Biomedical Studies
Doctor of Philosophy - 2013 - Monte-Carlo Sampling of Protein-Ligand Interactions and Computational Improvements to Implicit Solvent Models
Master of Philosophy - 2010
Master of Arts - 2009

- Lead developer and maintainer (~80% of source commits) of Protein Local Optimization Program (PLOP) project, a molecular mechanics library developed at Columbia University, University of California at San Francisco, and Schrödinger (fortran).
- Designed and implemented the computational mutation scanning module of PLOP.
- Rehabilitated project from non-compiling state on arrival.
- Redesigned build system to automatically determine dependencies and take advantage of parallel compilation, reducing build time from ~30 minutes to ~3 minutes and greatly accelerating development.
- Created a small molecule database representing 95%+ of small molecules in the Protein Data Bank, extending PLOP from a protein-only program to a general molecular mechanics toolkit.
- Designed and implemented a Perl based automated regression testing framework, which accelerated development while minimizing bugs and regressions.

- Created a project wiki, combining scattered documentation and completing missing documentation.

Rice University Houston, TX

2003 — 2007, GPA 3.7/4.0
Bachelor of Arts - Mathematics
Bachelor of Science - Biology (Ecology and Evolutionary Biology)
Relevant coursework: Machine Learning, Ordinary and Partial Differential Equations, Real and Complex Analysis, Combinatorics, Number Theory, Mathematical Logic, Modern Algebra, Euclidean and Non-Euclidean Geometry.

- Designed and implemented two methods of combining information from multiple crystal structures into a single ”composite motif”. These combined motifs increased sensitivity and specificity of motif matching algorithms.
- Completed senior thesis project identifying homologous pseudogenes in human and chimpanzee, and determining differential mutation rates.

Shotwell Project Yorba Foundation

2012 — 2013
Community Contributor

- Contributed a number of new features and bug fixes to the Shotwell photo organizer program.
- Recursively included contained files in the folder browser.
- Added support of panoramic images as event thumbnails.
- Updated searches to search comments and robustly treat accented characters.
- Fixed a number of UI experiences such as adding icons to buttons and windows, and correcting misleading text.

Other

- Extensive experience with AWS services and AWS api (EC2, dynamodb, kinesis, RDS...)
- Python, C++, shell (and previously FORTRAN, vala)
- Postgresql, mysql, Vertica, redis, git

Last updated August, 2018