

Brian Martin

2016A Grant St
Berkeley, CA 94703

☎ 815.814.1392

✉ brianmartin@gmail.com

📄 brianmartin.github.io

Education

- 2011–2013 **University of Massachusetts**, Amherst, MA.
MSc Computer Science, AI / Natural Language Processing
- 2008–2011 **Hampshire College**, Amherst, MA.
BA Computer Science, Chinese
- 2007–2008 **New York University**, New York, NY.
Economics, Math
- 2004–2007 **Illinois Mathematics and Science Academy**, Aurora, IL.

Experience

Employment

- 2016 - **Data/ML Engineer, Platform Infra Lead**, *Opendoor*, San Francisco, CA.
- Present
- Scaling engineering and data science infrastructure and tooling (from 30 to 120+ engineers).
 - Worked on the Data Infrastructure and Valuation Modeling teams before starting the Platform Infra team, reporting to the VP Eng.
 - Successfully introduced and drove org-wide adoption of Kubernetes.
 - Gained experience in administrative concerns: reduced engineering software and hosting costs by 30%; implemented org-wide incident management as well as project management tooling.
- 2012 - 2016 **Software Engineer, Technical Lead**, *Premise Data*, San Francisco, CA.
- Second engineering hire. Learned a lot from seed to series C growth.
 - Worked on data infrastructure and products, core backend, ops, and mobile (Android).
 - Ton of fun hiring, mentoring, and generally obviating myself.
- 2011 - 2013 **Software Engineer: Information Extraction and Synthesis Lab**, *University of Massachusetts*, Amherst, MA.
- Contributed to Factorie, a "toolkit for deployable probabilistic modeling, implemented in Scala."
 - Implemented major portions of IESL's NLP package, including: data loading, sentence and word segmentation, part-of-speech tagging, and named entity recognition.
- Summer 2012 **Engineering Intern: Data Science**, *Cloudera*, San Francisco, CA.
- 2010–2011 **Research Assistant: Evolutionary Computing**, *Hampshire College*, Amherst, MA.
Developed a job distribution, logging, and statistics aggregation system for a 160-core cluster (written in Clojure using AMQP).
- Summer 2010 **Google Summer of Code: Distributed Systems**, *Open Source Lab, Oregon State University*.
Implementation of several features for Pydra, a distributed and parallel computing framework for Python. This included heterogenous provisioning across providers (EC2, Rackspace, Linode) and runtime statistics collection to improve the job scheduler when facing network bottlenecks.

Other

- Summers 2008+2009 **Intensive Chinese language study**, *University of Chicago and Anhui Agricultural University, Hefei, China*, respectively.
- 2006–2007 **Mentorship: Graph Theory**, *Mathematics Department, Northwestern University*, Evanston, IL, under Prof. Ezra Getzler.
- 2005–2006 **Mentorship: Animation, 3D Modeling**, *Electronic Visualization Lab, University of Illinois*, Chicago, IL, under Director Jason Leigh.

Publications

Machine Learning (natural language, graphical models)

- Workshop **Joint inference of entities, relations, and coreference** S. Singh, S. Riedel, B. Martin, J. Zheng, A. McCallum *Proceedings of the 2013 workshop on Automated knowledge base construction, 1-6*
- Inducing Value Sparsity for Parallel Inference in Tree-shaped Models** S. Singh, B. Martin, A. McCallum *NIPS 2011 Workshop on Computational Trade-offs in Statistical Learning*.
- Competition **Universal Schema for Slot Filling and Cold Start: UMass IESL at TACKBP 2013** S. Singh, L. Yao, D. Belanger, A. Kobren, S. Anzaroot, M. Wick, A. Passos, H. Pandya, J. Choi, B. Martin, A. McCallum
- Undergraduate Thesis **Variants of the Belief Propagation Algorithm for Graphical Models** B. Martin advised by A. McCallum, L. Spector, and P. Dickson *Hampshire College*.
- ### Evolutionary Computing
- Book Chapter **What's in an Evolved Name? The Evolution of Modularity via Tag-Based Reference** L. Spector, K. Harrington, B. Martin, T. Helmuth *In Genetic Programming Theory and Practice IX. New York: Springer. In press*.
- Conference **Tag-based modules in genetic programming** L. Spector, B. Martin, K. Harrington, T. Helmuth *In Proceedings of the Genetic and Evolutionary Computation Conference (GECCO-2011). ACM Press. pp. 1419-1426*
- Workshop **Size-Based Tournaments for Node Selection** T. Helmuth, L. Spector, B. Martin *Genetic and Evolutionary Computation Conference: Graduate Student Workshop. Dublin, Ireland. ACM*.

Other

- Magazine **Big Data and Internships at Cloudera** Y. Chen, A. Ferguson, B. Martin, A. Wang, P. Wendell *XRDS: The ACM Magazine for Students. Fall 2012. 10.1145/2331042.2331054*

Selected Coursework

- Graduate Streaming Algorithms, Machine Learning, Graphical Models, Bayesian Text Modeling, Automated Knowledge Base Construction, Applied Information Theory, Foundations of Natural Language Processing, Systems, Mathematical Optimization, Readings in Big Data
- Undergraduate **Computer Science** Research in AI, Applied Cryptography, Genetic Programming, Creative Programming Workshop, Programming Techniques for the Interactive Arts, Data Structures and Algorithms, Foundations of CS
- Math** Advanced Topics in Continuous Applied Math: Dynamical Systems and Chaos, Reasoning about Uncertainty, Statistics, Multivariate Calculus, Graph Theory, Linear Algebra

Technical Proficiencies

- Languages Python, Scala, Java, Bash, Clojure
- Technologies AWS, Kubernetes, Docker, Linux, Postgres, Spark, Redis, Elastic, Git, Maven, L^AT_EX, ..., keywords forever

Other Achievements

Class President, Illinois Mathematics and Science Academy
Eagle Scout, Boy Scouts of America

Personal interests

Cooking, piano, reading, running, woodworking/boatbuilding, my cat, growing berries and figs