Elisabeth (Lissa) Moore

(413) 345 2944 • lissa@lissalytics.com • http://www.lissalytics.com

Summary

Machine learning researcher with skills in advanced machine learning, graph analysis, and computational social science. Currently leading machine learning research projects for the high performance computing division at Los Alamos National Laboratory. Seeking a research position in machine learning and data science.

Technical Skills

- Relational learning, ensemble methods, graphical models
- Agile software development and code testing
- Proficient: Python, Ruby on Rails, Java
- Familiar: C, C++, Clojure, Eclipse, Haskell, HTML/CSS, JavaScript, MIPS, Pixar's Renderman, Prolog, Matlab, Mathematica, R, x86
- Analyst-centric data mining tool development
- Presentation and communication of scientific results
- Databases: PostgreSQL, MySQL, MonetDB
- Amazon Web Services: Redshift, S3, SQS, SNS

Education

M.S. in Computer Science, University of Massachusetts Amherst

"Exploring Collective Behavior in Social Computation Through Relational Statistical Models" *2011 - 2015*

Advisor: David Jensen, Knowledge Discovery Laboratory

- **Research highlights**: relational and temporal machine learning, graph analysis, explanatory anomaly detection
- Coursework: Probabilistic Graphical Models, Statistics, Artificial Intelligence, Modern Computer Architecture, Research Methods in Empirical Computer Science, Advanced Algorithms, Advanced Compiler Design, Robotics

B.A. in Computer Science, Summa cum Laude with Distinction, Amherst College

Senior honors thesis: "Computing with Quantum Physics" *2007 - 2011*

Advisor: Lyle A. McGeoch, Computer Science

Select Projects

Explainable Anomaly Detection
LANL: Ongoing
Adapting classifier-adjusted density estimation to include an interpretable component by aggregating across a random forest, and provided mechanism to modify learned model to avoid false positives.

System Log Anomaly Detection
LANL: Ongoing
Developing fine-grained anomaly detection framework for complex computing system logs using methods from social network analysis, including statistical relational learning and Infomap for community detection.

HPC Job Outcome Prediction
LANL: Ongoing
Applying statistical topic modeling, including latent Dirichlet allocation, and other natural language processing techniques for early prediction of HPC job outcomes via related computer-generated textual logs.

Hardware Fault Identification
LANL: Ongoing
Modernizing memory fault characterization through supervised methods. Random forests predict fault modes faster and more accurately than current state-of-the-art, a crucial improvement for exascale machines.

Twitter User Classification
MIT-LL, LANL: 2013, 2015
Accurately classified Twitter account types without use of profile information. Created novel hypergraph representation for use in relational probability trees, and compared content to context features.

Strategic Collective Behavior
MIT-LL, UMass Amherst: 2015
Designed a relational temporal model of human behavior in game theory experiments with human subjects. Compared validity across varying strategies based on inherent versus contextual features.

Experience

Los Alamos National Laboratory – High Performance Computing Design **Los Alamos, NM**

Research Scientist, Ultrascale Systems Research Center *September 2015 - present*

- Leading division effort on machine learning for high performance computing problems
- Developing novel interpretable machine learning methods
- Improving detection and characterization of memory faults (DRAM, SRAM, and SSD)
- Designing and implementing anomaly detection for system monitoring and user behavior
- Organizing large-scale collaborations with academia

Fiksu, Inc. **Northampton, MA**
Software Engineer *September 2014 - July 2015*

- Led technical design and specification writing for predictive analytics project
- Developed, tested, and deployed quality code in a high-impact production environment
- Maintained Amazon Web Service infrastructure

Los Alamos National Laboratory – Center for Nonlinear Studies **Los Alamos, NM**
Graduate Research Assistant, Quantum Computing for Machine Learning *Summer 2014*
Supervisor: Rolando Somma

- Contributed to design and theoretical analysis of quantum computing algorithm for machine learning
- Simulated quantum computing system for experimentation

MIT Lincoln Laboratory – Human Language Technology **Lexington, MA**
Graduate Research Assistant, Relational Learning for Big Data Exploration *Summer 2013 - August 2014*
Supervisor: William Campbell

- Applied relational learning techniques to microfinance (Kiva) and microblogging (Twitter) data
  - Detected anomalies in Kiva that corresponded to real-world fraud
  - Classified Twitter user type (celebrity, business organization, etc.) using only a few tweets
- Engineered, with collaborators, a novel hypergraph analytical method for Twitter
- Developed slides on relational learning presented to DARPA sponsor
- Secured graduate research funding to continue work at UMass Amherst

University of Massachusetts Amherst **Amherst, MA**
Graduate Research Assistant, Knowledge Discovery Laboratory *Fall 2011 - Spring 2014*
Advisor: David Jensen

- Researched strategic behavior of individuals and groups in social networks
- Contributed to DARPA ADAMS and XDATA projects
- Enhanced statistical relational learning techniques for classification and anomaly detection

Publications

- **Baseman, Elisabeth**, Nathan DeBardeleben, Sean Blanchard, Juston Moore, Olena Tkachenko, Kurt Ferreira, Taniya Siddiqua, and Vilas Sridharan. *Physics-Informed Machine Learning for DRAM Error Modeling*. IEEE International Symposium on Defect and Fault Tolerance in VLSI and Nanotechnology Systems, 2018.
- Hickman, Megan, Dakota Fulp, **Elisabeth Baseman**, Sean Blanchard, Hugh Greenberg, William Jones, and Nathan DeBardeleben. *Enhancing HPC System Log Analysis by Identifying Message Origin in Source Code*. IEEE Software Reliability Engineering (Industry Track), 2018.
- Amvrosiadis, George, Jun Woo Park, Gregory R. Granger, Garth A. Gibson, **Elisabeth Baseman**, and Nathan DeBardeleben. *On the Diversity of Cluster Workloads and its Impact on Research Results*. Usenix Annual Technical Conference, 2018.
- DeLucia, Alexandra, and **Elisabeth Baseman**. *Work in Progress: Topic Modeling for HPC Job State Prediction*. HPDC: Workshop on Machine Learning for Computing Systems, 2018.
- Goetting, Ian, **Elisabeth Baseman**, and Huiping Cao. *Work in Progress: Causal Relationships amongst Sensors in the Trinity Supercomputer*. HPDC: Workshop on Machine Learning for Computing Systems, 2018.
- Haque, Abida, Alexandra DeLucia, and **Elisabeth Baseman**. *Markov Chain Modeling for Anomaly Detection in High Performance Computing System Logs*. SC: Fourth International Workshop on HPC User Support Tools, 2017.
- Siddiqua, Taniya, Vilas Sridharan, Steven Raasch, Nathan DeBardeleben, Kurt Ferreira, Scott Levy, **Elisabeth Baseman**, and Qiang Guan. *Lifetime Memory Reliability Data from the Field*. Defect and Fault Tolerance in VLSI and Nanotechnology Systems 2017. Best Paper Nominee.
- **Baseman, Elisabeth**, Nathan DeBardeleben, Kurt Ferreira, Vilas Sridharan, Taniya Siddiqua, and Olena Tkachenko. *Automating DRAM Fault Mitigation By Learning From Experience*. Dependable Systems and Networks 2017.
- **Baseman, Elisabeth**, Sean Blanchard, Zongze Li, and Song Fu. *Relational Synthesis of Text and Numeric*

Data for Anomaly Detection on Computing System Logs. ICMLA 2016.

- Morrow, Adam, **Elisabeth Baseman**, and Sean Blanchard. *Ranking Anomalous High Performance Computing Sensor Data Using Unsupervised Clustering*. Computational Science and Computational Intelligence: Symposium on Parallel and Distributed Computing and Computational Science 2016.
- **Baseman, Elisabeth**, Sean Blanchard, Nathan DeBardeleben, Amanda Bonnie, and Adam Morrow. *Interpretable Anomaly Detection for Monitoring of High Performance Computing Systems*. Outlier Definition, Detection, and Description on Demand: KDD 2016 workshop paper.
- **Baseman, Elisabeth**, Nathan DeBardeleben, Kurt Ferriera, Scott Levy, Steven Raasch, Vilas Sridharan, Taniya Siddiqua, Qiang Guan. *Improving DRAM Fault Characterization Through Machine Learning*. Dependable Systems and Networks 2016.
- Siddiqua, Taniya, Vilas Sridharan, Nathan DeBardeleben, **Elisabeth Baseman**, Qiang Guan, Devesh Tiwari, Christian Engelmann, and Saurabh Gupta. *Memory Error Analysis and Lessons Learned from Large-scale Field Data*. 2016, Los Alamos National Laboratory and AMD internal document.
- Guan, Chung, Nathan DeBardeleben, Panruo Wu, Stephan Eidenbenz, Sean Blanchard, Laura Monroe, **Elisabeth Baseman**, and Li Tan. *Design, Use, and Evaluation of P-FSEFI: A Parallel Soft Error Fault Injection Framework for Emulating Soft Errors in Parallel Applications*. SIMUTOOLS 2016.
- **Baseman, Elisabeth**, and David Jensen. *Collaborative Behavior in Social Networks: A Relational Approach*. Networks in the Social and Information Sciences. NIPS 2015 workshop paper.
- **Baseman, Elisabeth** and David Jensen. *Exploring Collective Behavior in Social Computation Through Relational Statistical Models*. Computational Social Science Society of the Americas 2015.
- Campbell, William, **Elisabeth Baseman**, and Kara Greenfield. *Content + Context = Classification: Examining the Roles of Social Interactions and Linguist Content in Twitter User Classification*. Social NLP. COLING 2014 workshop paper.
- Campbell, William, **Elisabeth Baseman**, and Kara Greenfield. *Content + Context Networks for User Classification in Twitter*. Frontiers of Network Analysis. NIPS 2013 workshop paper.
- **Baseman, Elisabeth**. *Computing with Quantum Physics*. Amherst College honors thesis, 2011.

Presentations

- **Baseman, Elisabeth**. *"Explainable Machine Learning... with an Application to High Performance Computing."* IEEE Women in Engineering Tech Summit Los Alamos, August 7, 2018.
- **Baseman, Elisabeth**. *"Machine Learning Meets High Performance Computing: Modernizing Monitoring"*. National Energy Research Scientific Computing Center invited talk, July 16, 2018.
- **Baseman, Elisabeth**. *"Interpretable Anomaly Detection for High Performance Computing Centers: Monitoring System Logs"*. Chesapeake Large-Scale Analytics Conference invited talk. October 18, 2017.
- **Baseman, Elisabeth**. *"Machine Learning for High Performance Computing: Modernizing Monitoring"*. New Mexico State University invited talk. October 11, 2017.
- **Baseman, Elisabeth**. *"Interpretable, Context-Aware Anomaly Detection for High Performance Computing Systems: Monitoring Syslog"*. Carnegie Mellon University invited talk. August 29, 2017.
- **Baseman, Elisabeth**. *"Machine Learning for High Performance Computing"*. Southern Data Science Conference invited talk. April 7, 2017.
- **Baseman, Elisabeth**. *"Don't Be Pipelined"*. HPC Pipeline Workshop: Diversifying the HPC Workforce invited talk. January 26, 2017.
- **Baseman, Elisabeth**. *"Machine Learning for Detection and Diagnosis: From Computational Social Science to High Performance Computing"*. United States Department of Defense invited talk. September 15, 2016.
- **Baseman, Elisabeth**.*"Little Machines Helping Big Machines: Data Science for High Performance Computing"*. Los Alamos National Laboratory Ultrascale Systems Research Center 1st Annual Symposium keynote. August 4, 2016.
- **Baseman, Elisabeth**, Nathan DeBardeleben, Kurt Ferreira, Scott Levy, Steven Raasch, Vilas Sridharan, Taniya Siddiqua, Qiang Guan. *"Machine Learning for Automatic Memory Fault Mode Characterization"*. Silicon Errors in Logic — System Effects 2016 random access talk.

- **Baseman, Elisabeth**. *"Relational Learning for Fraud Detection"*. Fiksu engineering talk. April 2015.
- **Baseman, Elisabeth**. *"An Introduction to Quantum Computing"*. Fiksu engineering talk. November 2014.
- **Baseman, Elisabeth**. *"Applications of Relational Learning"*. Invited talk for MIT Lincoln Laboratory Computing and Analytics group. March 12, 2014.
- **Baseman, Elisabeth**. *"Relational Learning for XDATA"*. Talk for MIT Lincoln Laboratory Human Language Technology group. Summer 2013.

Posters

- DeLucia, Alexandra, and **Elisabeth Baseman**. *High Performance Computing Job Outcome Prediction By Mining System Logs*. Southern Data Science Conference, 2018.
- **Baseman, Elisabeth**. *Helping Exascale Computers Help Us: Machine Learning for High Performance Computing*. NIPS: Women in Machine Learning Workshop, 2017.
- **Baseman, Elisabeth**, Nathan DeBardeleben, Kurt Ferreira, Scott Levy, Steven Raasch, Vilas Sridharan, Taniya Siddiqua, and Qiang Guan. *A Machine Learning Approach for Automatic Characterization of Memory Faults*. Conference on Data Analysis 2016 poster.
- **Baseman, Elisabeth** and David Jensen. *Relational Statistical Models of Collaborative Behavior in Social Networks*. Women in Machine Learning 2015 workshop poster.
- **Baseman, Elisabeth**, Michael Kearns, Stephen Judd, and David Jensen. *Dynamic Statistical Models of Collective Social Network Behavior*. New England Machine Learning Day 2014 poster.
- **Baseman, Elisabeth**, Michael Kearns, Stephen Judd, and David Jensen. *Statistical Models of Collective Social Network Behavior*. Women in Machine Learning 2013 workshop poster.

Teaching Experience

Los Alamos National Laboratory

Lecturer: Applied Machine Learning Tutorial — *February 2016 - Present*

- Lecturing weekly on current machine learning techniques
- Tailoring topics to high performance computing problems
- Demonstrating use of Python for machine learning
- Providing advice on machine learning for HPC researchers' current projects

University of Massachusetts Amherst

Teaching Assistant: Introduction to Problem Solving with Computers (Java Programming) — *Spring 2013*

- Lectured twice a week
- Held weekly office hours
- Graded midterm and final exams

Amherst College

Teaching Assistant, Laboratory Teaching Assistant, Grader — *Spring 2009 - Spring 2011*

- Introduction to Computer Science II
- Computer Systems II
- Electromagnetism and Optics
- Modern Physics

Advising

Los Alamos National Laboratory

Machine Learning Research Advisor/Mentor

- Alexandra DeLucia, Rollins College, Undergraduate Honors Thesis, Post-Baccalaureate Student, 2017-present
- Michael Kuchnik, Carnegie Mellon University, Graduate Summer Student, Summer 2018
- Emily Porter, University of Texas at Austin, Undergraduate Summer Student, Summer 2018
- David Huff, New Mexico Institute of Mining and Technology, Post-Baccalaureate Summer Student, Summer 2018
- Randall Woodall, New Mexico State University, Undergraduate Summer Student, Summer 2018
- Ian Goetting, New Mexico State University, Undergraduate Honors Thesis, 2017-2018
- Ashley Michalenko, New Mexico State University, Masters Student, 2017
- Abida Haque, Georgia Tech, Post-Masters Student, 2017
- Olena Tchakenko, Florida International University, Post-Baccalaureate Student, 2017
- Adam Morrow, Brigham Young University, Undergraduate Summer Student, Summer 2016
- Zongze Li, University of North Texas, Graduate Summer Student, Summer 2016

Fiksu, Inc.

Career Mentor

- Evelyn Ting, Amherst College, Undergraduate Student, Spring 2015
- Elizabeth Lefever, Amherst College, Undergraduate Student, Spring 2015

Awards

Laboratory Pathfinder Grant: LANL, High Performance Computing, 2017 FY

Laboratory Directed Research & Development grant: LANL, High Performance Computing, 2016 FY

Summer GRA funding grant: LANL, Theoretical Division, Center for Nonlinear Studies, Summer 2014

CISCO grant to attend CRA-W Grad Cohort: University of Massachusetts Amherst, Spring 2014

RA funding grant: MIT Lincoln Laboratory, Fall 2013 - September 2014

Grant for UMass CS Women: Women for UMass, Fall 2012

Graduate School Fellowship: University of Massachusetts Amherst, 2011

Computer Science Prize: Amherst College, Spring 2011

Scholarship to Quantum Information Science for Undergraduates: Massachusetts Institute of Technology, 2010

Amherst College Schupf Scholar: Fall 2007 - Spring 2011

IBM Thomas J Watson Memorial Scholarship: Fall 2007 - Spring 2011

Memberships, Activities, and Service

Academic

Leadership Board Member, Los Alamos National Laboratory Atomic Women: 2017-present

Chair, Workshop on Machine Learning for Computing Systems (MLCS): HPDC 2018

Workshop Committee Member, Supercomputing 2018 (SC'18): 2017-present

Advisory Board Member, Southern Data Science Conference (SDSC): 2017-present

Reviewer, Women in Machine Learning Workshop (WiML): 2017-present

Program Committee Member, Computational Social Science Society of the Americas (CSSSA): 2016-present

Reviewer, Statistical Analysis and Data Mining (SAM): 2016

Technical Program Committee Member, 2nd International Workshop of Fault Tolerant Systems (FTS): 2016

Subreviewer, High-Performance Parallel and Distributed Computing (HPDC): 2016

Subreviewer, International Symposium on Cluster, Cloud, and Grid Computing (CCGrid): 2016

Session Chair, Computational Social Science Society of the Americas: October 2015

UMass Amherst Graduate Women in STEM: Spring 2013 - Spring 2014

Student Member, Association for Computing Machinery: Fall 2011 - present

Co-Chair, UMass CS Women: 2012

Associate Member, Sigma Xi: 2011 - present

President and Founder, Amherst College Women in Technology and Science: 2010 - 2011

Outreach

Workshop Presenter, Expand Your Horizons Conference, Santa Fe, NM: March 2016, March 2017, March 2018

Science Fair Judge, Hampden Charter School of Science, Chicopee, MA: Spring 2012