
OBJECTIVE	Seeking a full-time software engineering position with emphasis on data mining and machine learning.	
TECHNICAL SKILLS	<ul style="list-style-type: none"> • Languages: Python, Java, C, C++, Matlab, MySQL, DB2, PHP, JavaScript, Haskell. • Modeling: Keras, Tensorflow, NLTK, spaCy, scikit-learn. • Infrastructure + Scaling: Docker, Kubernetes, AWS, Spark, Hadoop, Airflow. • Interests: Machine Learning, NLP, Heterogeneous Information Networks, Graph Theory, Artificial Intelligence. 	
EDUCATION	University of Illinois at Urbana-Champaign , Urbana, IL (Currently on leave) Ph.D. in Computer Science Area of study: Data Mining Advisor: Prof. Jiawei Han	Sep. 2012 - May. 2015
	Boston University , Boston, MA B.A./M.A. in Computer Science, <i>Summa Cum Laude</i> . GPA: 3.95/4.00 in Computer Science, 3.85/4.00 overall Trustee Scholar (recipient of a merit-based, four-year, full tuition scholarship)	Sep. 2008 - May 2012
INDUSTRY EXPERIENCE	Co-founder and CTO , Migacore Technologies <ul style="list-style-type: none"> • ML + Engineering: Designing, building, and refining a system to (1) crawl and extract events worldwide, and (2) learn what types of events are most impactful to demand for air travel (<i>Python, NLTK, spaCy, Keras, scikit-learn, Spark+Dask</i>) • Infrastructure: Set up and maintain continuous deployment (<i>CircleCI, Docker</i>) • Team: Building, developing, and growing the machine learning and engineering teams. 	May 2017 – present
	Software Engineer , Palantir Technologies <ul style="list-style-type: none"> • Built an internal deployability tool, which allows customers to describe their product stack in a custom DSL (Java, Groovy). • Built AtlasDB, a transactional layer on top of a key-value store (Java, distributed systems). 	Jun. 2015 – Nov. 2016
	Software Engineering Intern , Google <ul style="list-style-type: none"> • Built and evaluated a classification system to help solve a (non-public) internal task. • The application required high model interpretability. The models generated by my framework are more interpretable than all baselines, without a loss in performance. 	May. 2014 – Aug. 2014
	Software Engineering Intern , American Express <ul style="list-style-type: none"> • Explored strategies for recommending merchants to cardholders. • Ongoing work on recommending merchants away from a user's home city 	Jun. 2013 – Aug. 2013
	Software Engineering Intern , Groupon <ul style="list-style-type: none"> • Improved the algorithm that personalizes deal recommendations for users by using a more accurate estimate for the price of each deal. • Observed a 3.3% increase in revenue, as validated by A/B tests, with high statistical confidence. • Wrote new code for the recommendation engine in Java, and analysis tools in Hadoop. 	Jun. 2012 – Aug. 2012

Financial Software Developer Intern, Bloomberg L.P. May 2011 – Aug. 2011

- Redesigned the Fair Value Detail function on the Bloomberg terminal, which computes and compares the fair values of futures contracts.
- This function is used by financial professionals over 2,000 times daily.
- Wrote the backend functionality in C++ and the user interface in Javascript.

Financial Software Developer Intern, Bloomberg L.P. May 2010 – Aug. 2010

- Designed and built an analysis tool for Bloomberg Tradebook's Pair Trading algorithm.

RESEARCH
EXPERIENCE

Research Assistant, University of Illinois Aug. 2012 – present

- Studying the spread of information in heterogeneous information networks.

Research Assistant (Authority Influence), Boston University Sep. 2010 – May 2011

- Defined a model to discern between and quantify the strength of influence exerted by peer and authority relationships in a social network, and applied the model to two real-world datasets to extract meaningful conclusions about the underlying social processes.

Research Assistant (Query by Humming), Boston University May 2009 – May 2010

- Designed and implemented an efficient and accurate query-by-humming system in C++

PUBLICATIONS

- Honglei Zhuang, Jing Zhang, **George Brova**, Jie Tang, Hasan Cam, Xifeng Yan, Jiawei Han: *"Mining Query-Based Subnetwork Outliers in Heterogeneous Information Networks"*, ICDM 2014.
- Huan Gui, Yizhou Sun, Jiawei Han, **George Brova**: *"Modeling Topic Diffusion in Multi-Relational Bibliographic Information Networks"*, CIKM 2014.
- Fangbo Tao, **George Brova**, Jiawei Han, Heng Ji, Chi Wang, Brandon Norick, Ahmed El-Kishky, Jialu Liu, Xiang Ren, Yizhou Sun: *"NewsNetExplorer: automatic construction and exploration of news information networks"*, SIGMOD Conference 2014.
- Fangbo Tao, Xiao Yu, Kin Hou Lei, **George Brova**, Xiao Cheng, Jiawei Han, Rucha Kanade, Yizhou Sun, Chi Wang, Lidan Wang, Tim Weninger: *"Research-insight: providing insight on research by publication network analysis"*, SIGMOD Conference 2013.
- Aris Anagnostopoulos, **George Brova**, Evimaria Terzi: *"Peer and authority pressure in information-propagation models"*, Principles and Practice of Knowledge Discovery in Databases (ECML-PKDD) 2011.

TEACHING

- Teaching Assistant for CS411: Databases, Fall 2014
- Teaching Assistant for CS412: Data Mining, Fall 2013

HONORS AND
AWARDS

- Honorable mention, Computing Research Association Outstanding Undergraduate Researcher Award, a nationwide competition (2011)
- UROP Travel Award to present paper at the ECML-PKDD conference (2011)
- UROP funding for work on query-by-humming system (2009-2010)