

Alex Beutel

alexbeutel.com

alex@beu.tel

EDUCATION	Carnegie Mellon University August 2011–May 2016 <i>Ph.D., Computer Science</i> May 2016 <i>Masters of Science, Computer Science</i> December 2013 Thesis title: “Understanding User Behavior through Large-Scale Graph Analysis” Committee: Christos Faloutsos, Alex Smola, Geoff Gordon, Phillip Yu
	Duke University August 2007–May 2011 <i>Bachelor of Science, Quantitative Studies in Computer Science and Physics</i> GPA: 3.858/4.0; Dean’s List (FA08, FA09) with Distinction (FA07, SP08, SP10) Graduated <i>Magna cum Laude</i> and with <i>Highest Distinction</i> in Computer Science
INDUSTRY EXPERIENCE	Google , Research Senior Research Scientist Aug. 2016–Present Research Intern, hosted by Ed H. Chi May 2015–Aug. 2015 Marianas Labs , Deep learning research consultant Feb. 2015–May 2015 Microsoft , Research Intern Cloud and Information Service Lab, hosted by Markus Weimer June 2014–Aug. 2014 Facebook , Software Engineering Intern Site Integrity, hosted by Wanhong Xu, Chris Palow May 2012–Aug. 2013 Site Integrity (remote part-time) Sept. 2012–May 2013 News Feed Ranking, hosted by Wanhong Xu, Lars Backstrom May 2013–Aug. 2013
ACADEMIC EXPERIENCE	Carnegie Mellon University, Computer Science Department Sept. 2011–Aug. 2016 Advised by Professor Christos Faloutsos and Professor Alex Smola. My research focuses on large-scale user behavior modeling, including fraud detection, recommendation systems, and scalable machine learning. Duke University, Department of Computer Science Jan. 2010–Aug. 2011 Research assistant for Prof. Pankaj K. Agarwal in computational geometry. Duke University, Department of Computer Science Oct. 2009–Dec. 2010 Research assistant for Prof. Xiaowei Yang in networks and distributed systems. Duke University, Department of Physics April 2009–Aug. 2009 Research assistant for Prof. Chris Walter in the high energy physics, neutrino group
HONORS	SIGKDD Doctoral Dissertation Award Runner-up , 2017 Best Paper Award , <i>ACM KDD</i> 2016 Young Researcher at the Heidelberg Laureate Forum 2015 , Heidelberg, Germany. Best Paper Finalist , <i>ACM KDD</i> 2014 Phi Beta Kappa Honor Society , 2012 Alex Vasilos Memorial Award , Duke University Computer Science, 2011 Best Paper Award , <i>ACM GIS</i> 2010 Computer Science Undergraduate Research Fellow , Duke University 2010
FUNDING AWARDS	Facebook Graduate Fellowship , 2013–2014 \$79,202 NSF Graduate Research Fellowship , 2011–2016 \$132,000 Yahoo! Faculty Research and Engagement Program , 2014 \$10,000 Aided Professor Christos Faloutsos in writing the research proposal NSF Collaborative Grant (Award No. IIS-1408924) , 2014 \$307,908 Helped multiple professors with the research proposal ACM GIS Student Travel Grant Award , 2011 \$1,000

- C30. **Towards Neural Mixture Recommender for Long Range Dependent User Sequences**
Jiaxi Tang, Francois Belletti, Sagar Jain, Minmin Chen, Alex Beutel, Can Xu, Ed H. Chi. *WWW 2019: The 2019 Web Conference*, 2019.
- C29. **Top-K Off-Policy Correction for a REINFORCE Recommender System**
Minmin Chen*, Alex Beutel*, Paul Covington*, Sagar Jain, Francois Belletti, Ed H. Chi. *Twelfth ACM International Conference Web Search and Data Mining (WSDM)*, 2019.
- C28. **Putting Fairness Principles into Practice: Challenges, Metrics, and Improvements**
Alex Beutel, Jilin Chen, Tulsee Doshi, Hai Qian, Allison Woodruff, Christine Luu, Pierre Kreitmann, Jonathan Bischof, Ed H. Chi. *AAAI/ACM Conference on Artificial Intelligence, Ethics, and Society (AIES)*, 2019.
- C27. **Counterfactual Fairness in Text Classification through Robustness**
Sahaj Garg, Vincent Perot, Nicole Limtiaco, Ankur Taly, Ed H. Chi, Alex Beutel. *AAAI/ACM Conference on Artificial Intelligence, Ethics, and Society (AIES)*, 2019.
- C26. **SageDB: A Learned Database System**
Tim Kraska, Mohammad Alizadeh, Alex Beutel, Ed H. Chi, Jialin Ding, Ani Kristo, Guillaume Leclerc, Samuel Madden, Hongzi Mao, Vikram Nathan. *Ninth Biennial Conference on Innovative Data Systems Research (CIDR)*, 2019.
- C25. **Categorical-Attributes-Based Item Classification for Recommender Systems**
Qian Zhao, Jilin Chen, Minmin Chen, Sagar Jain, Alex Beutel, Francois Belletti, Ed H. Chi. *Proceedings of the 12th ACM Conference on Recommender Systems (RecSys)*, 2018.
- C24. **Q&R: A Two-Stage Approach Toward Interactive Recommendation**
Konstantina Christakopoulou, Alex Beutel, Rui Li, Sagar Jain, Ed H. Chi. *Proceedings of the 24th ACM SIGKDD International Conference on Knowledge Discovery and Data Mining (KDD Applied Data Science)*, 2018.
- C23. **The Case for Learned Index Structures**
Tim Kraska, Alex Beutel, Ed H. Chi, Jeffrey Dean, Neoklis Polyzotis. *ACM SIGMOD International Conference on Management of Data (SIGMOD)*, 2018.
- C22. **Factorized Recurrent Neural Architectures for Longer Range Dependence**
Francois Belletti, Alex Beutel, Sagar Jain, Ed H. Chi. *21st International Conference on Artificial Intelligence and Statistics (AISTATS)*, 2018.
- C21. **Latent Cross: Making Use of Context in Recurrent Recommender Systems**
Alex Beutel, Paul Covington, Sagar Jain, Can Xu, Jia Li, Vince Gatto, Ed H. Chi. *Eleventh ACM International Conference Web Search and Data Mining (WSDM)*, 2018.
- C20. **The Many Faces of Link Fraud**
Neil Shah, Hemank Lamba, Alex Beutel, Christos Faloutsos. *IEEE International Conference on Data Mining (ICDM)*, 2017.
- C19. **Beyond Globally Optimal: Focused Learning for Improved Recommendations**
Alex Beutel, Ed H. Chi, Derek Zhiyuan Cheng, Hubert Pham, John Anderson. *Proceedings of the 26th International Conference on World Wide Web (WWW)*, 2017.
- C18. **Recurrent Recommender Networks**
Chao-Yuan Wu, Amr Ahmed, Alex Beutel, Alex Smola, How Jing. *Tenth ACM International Conference Web Search and Data Mining (WSDM)*, 2017.
- C17. **FRAUDAR: Bounding Graph Fraud in the Face of Camouflage**
Bryan Hooi, Hyun Ah Song, Alex Beutel, Neil Shah, Kijung Shin, Christos Faloutsos. *Proceedings of the 22nd ACM SIGKDD International Conference on Knowledge Discovery and Data Mining (KDD)*, 2016.

- C16. **BIRDNEST: Bayesian Inference for Ratings-Fraud Detection**
Bryan Hooi, Neil Shah, Alex Beutel, Stephan Gunnemann, Leman Akoglu, Mohit Kumar, Disha Makhija, Christos Faloutsos. *2016 SIAM International Conference on Data Mining (SDM)*, 2016.
- C15. **A General Suspiciousness Metric for Dense Blocks in Multimodal Data**
Meng Jiang, Alex Beutel, Peng Cui, Bryan Hooi, Shiqiang Yang, Christos Faloutsos. *IEEE International Conference on Data Mining (ICDM)*, 2015.
- C14. **ACCAMS: Additive Co-Clustering to Approximate Matrices Succinctly**
Alex Beutel, Amr Ahmed, Alexander J. Smola. *Proceedings of the 24th International Conference on World Wide Web (WWW)*, 2015.
- C13. **ND-SYNC: Detecting Synchronized Fraud Activities**
Maria Giatsoglou, Despoina Chatzakou, Neil Shah, Alex Beutel, Christos Faloutsos, Athena Vakali. *19th Pacific-Asia Conference on Knowledge Discovery and Data Mining (PAKDD)*, 2015.
- C12. **Spotting Suspicious Link Behavior with fBox: An Adversarial Perspective**
Neil Shah, Alex Beutel, Brian Gallagher, Christos Faloutsos. *IEEE International Conference on Data Mining (ICDM)*, 2014.
- C11. **CatchSync: Catching Synchronized Behavior in Large Directed Graphs**
Meng Jiang, Peng Cui, Alex Beutel, Christos Faloutsos, Shiqiang Yang. *Proceedings of the 20th ACM SIGKDD International Conference on Knowledge Discovery and Data Mining (KDD)*, 2014.
- C10. **Inferring Strange Behavior from Connectivity Pattern in Social Networks**
Meng Jiang, Peng Cui, Alex Beutel, Christos Faloutsos, Shiqiang Yang. *18th Pacific-Asia Conference on Knowledge Discovery and Data Mining (PAKDD)*, 2014.
- C9. **Fugue: Slow-Worker-Agnostic Distributed Learning for Big Models**
Abhimanu Kumar, Alex Beutel, Qirong Ho, Eric P. Xing. *17th International Conference on Artificial Intelligence and Statistics (AISTATS)*, 2014.
- C8. **FlexiFaCT: Scalable Flexible Factorization of Coupled Tensors on Hadoop**
Alex Beutel, Abhimanu Kumar, Evangelos E. Papalexakis, Partha Pratim Talukdar, Christos Faloutsos, Eric P. Xing. *2014 SIAM International Conference on Data Mining (SDM)*, 2014.
- C7. **CoBaFi: Collaborative Bayesian Filtering**
Alex Beutel, Kenton Murray, Christos Faloutsos, Alexander J. Smola. *Proceedings of the 23rd International Conference on World Wide Web (WWW)*, 2014.
- C6. **CopyCatch: Stopping Group Attacks by Spotting Lockstep Behavior in Social Networks**
Alex Beutel, Wanhong Xu, Venkatesan Guruswami, Christopher Palow, Christos Faloutsos. *Proceedings of the 22nd International Conference on World Wide Web (WWW)*, 2013.
- C5. **Network Anomaly Detection using Co-clustering**
Evangelos E. Papalexakis, Alex Beutel, Peter Steenkiste. *Proceedings of the 2012 International Conference on Advances in Social Networks Analysis and Mining (ASONAM)*, 2012.
- C4. **Interacting Viruses on a Network: Can both survive?**
Alex Beutel, B. Aditya Prakash, Roni Rosenfeld, Christos Faloutsos. *Proceedings of the 18th ACM SIGKDD International Conference on Knowledge Discovery and Data Mining (KDD)*, 2012.
- C3. **Winner-takes-all: Competing Viruses on fair-play networks**
B. Aditya Prakash, Alex Beutel, Roni Rosenfeld, Christos Faloutsos. *Proceedings of the 21st International Conference on World Wide Web (WWW)*, 2012.

- C2. **TerraNNI: Natural Neighbor Interpolation on a 3D Grid Using a GPU**
Alex Beutel, Thomas Moelhave, Pankaj K. Agarwal, Arnold P. Boedihardjo, James A. Shine. *Proceedings of the 19th International Symposium on Advances in Geographic Information Systems (ACM GIS)*, 2011.
- C1. **Natural Neighbor Interpolation Based Grid DEM Construction Using a GPU**
Alex Beutel, Thomas Moelhave, Pankaj K. Agarwal. *Proceedings of the 18th International Symposium on Advances in Geographic Information Systems (ACM GIS)*, 2010.
- W9. **Lifting the Curse of Multidimensional Data with Learned Existence Indexes**
Stephen Macke, Alex Beutel, Tim Kraska, Maheswaran Sathiamoorthy, Derek Zhiyuan Cheng, Ed H. Chi. *ML for Systems workshop at NeurIPS*, 2018.
- W8. **Data Decisions and Theoretical Implications when Adversarially Learning Fair Representations**
Alex Beutel, Jilin Chen, Zhe Zhao, Ed H. Chi. *Workshop on Fairness, Accountability, and Transparency in Machine Learning*, 2017.
- W7. **Joint Training of Ratings and Reviews with Recurrent Recommender Networks**
Chao-Yuan Wu, Amr Ahmed, Alex Beutel, Alex Smola. *Workshop track at 5th International Conference on Learning Representations (ICLR)*, 2017.
- W6. **EdgeCentric: Anomaly Detection in Edge-Attributed Networks**
Neil Shah, Alex Beutel, Bryan Hooi, Leman Akoglu, Stephan Gunnemann, Disha Makhija, Mohit Kumar, Christos Faloutsos. *IEEE International Conference on Data Mining (ICDM) Workshop on Data Mining for Cyber Security*, 2016.
- W5. **Additive Co-Clustering of Gaussians and Poissons for Joint Modeling of Ratings and Reviews**
Chao-Yuan Wu, Alex Beutel, Amr Ahmed, Alexander J. Smola. *NeurIPS workshop on Nonparametric Methods for Large Scale Representation Learning*, 2015.
- W4. **Collaborative Bayesian Filtering: Patterns and Methods**
Alex Beutel, Kenton Murray, Christos Faloutsos, Alexander J. Smola. *Workshop on Information Networks (WIN)*, 2015.
- W3. **Elastic Distributed Bayesian Collaborative Filtering**
Alex Beutel, Markus Weimer, Tom Minka, Yordan Zaykov, Vijay Narayanan. *NeurIPS Distributed Machine Learning and Matrix Computations workshop*, 2014.
- W2. **FlexiFaCT: Scalable Flexible Factorization of Coupled Tensors on Hadoop**
Alex Beutel, Abhimanu Kumar, Evangelos E. Papalexakis, Partha Pratim Talukdar, Christos Faloutsos, Eric P. Xing. *NeurIPS Big Learning Workshop*, 2013.
- W1. **Volumetric Grid Construction using 3D Natural Neighbor Interpolation on the GPU**
Alex Beutel, Thomas Moelhave, Pankaj K. Agarwal. *MASSIVE '11: Proceedings of the Workshop on Massive Data Algorithmics*, 2011.

REFEREED
WORKSHOP
PAPERS

- J5. **Graph-Based Fraud Detection in the Face of Camouflage**
Bryan Hooi, Kijung Shin, Hyun Ah Song, Alex Beutel, Neil Shah, Christos Faloutsos. *ACM Transactions on Knowledge Discovery from Data (TKDD)*, 2017.
- J4. **Spotting Suspicious Behaviors in Multimodal Data: A General Metric and Algorithms**
Meng Jiang, Alex Beutel, Peng Cui, Bryan Hooi, Shiqiang Yang, Christos Faloutsos. *Transactions on Knowledge and Data Engineering (TKDE)*, 2016.
- J3. **Catching Synchronized Behaviors in Large Networks: A Graph Mining Approach**
Meng Jiang, Peng Cui, Alex Beutel, Christos Faloutsos, Shiqiang Yang. *ACM Transactions on Knowledge Discovery from Data (TKDD)*, 2016.

REFEREED
JOURNAL
PAPERS

- J2. **TerraNNI: Natural Neighbor Interpolation on 2D and 3D Grids using a GPU**
Pankaj K. Agarwal, Alex Beutel, Thomas Moelhave. *ACM Transactions on Spatial Algorithms and Systems (TSAS)*, 2016.
- J1. **Inferring Lockstep Behavior from Connectivity Pattern in Large Graphs**
Meng Jiang, Peng Cui, Alex Beutel, Christos Faloutsos, Shiqiang Yang. *Knowledge and Information Systems (KAIS)*, 2015.

BOOK	Network Anomaly Detection using Co-clustering
CHAPTERS	Evangelos E. Papalexakis, Alex Beutel, Peter Steenkiste. <i>Springer Encyclopedia of Social Network Analysis and Mining</i> , 2012.
OTHER PAPERS	User Behavior Modeling with Large-Scale Graph Analysis Alex Beutel. <i>Ph.D. Thesis, Carnegie Mellon University</i> , 2016.
	User Behavior Modeling and Fraud Detection Alex Beutel, Christos Faloutsos. <i>IEEE Intelligent Systems: Trends and Controversies</i> , 2016.
	Explaining reviews and ratings with PACO: Poisson Additive Co-Clustering Chao-Yuan Wu, Alex Beutel, Amr Ahmed, Alexander J. Smola. <i>Companion Proceedings of the 25th International Conference on World Wide Web (WWW)</i> , 2016.
	Detecting Suspicious Following Behavior in Multimillion-Node Social Networks Meng Jiang, Peng Cui, Alex Beutel, Christos Faloutsos, Shiqiang Yang. <i>Companion Proceedings of the 23rd International Conference on World Wide Web (WWW)</i> , 2014.
	From Point Cloud to 2D and 3D Grids: A Natural Neighbor Interpolation Algorithm using the GPU Alex Beutel. <i>Senior Thesis - Graduation with Highest Distinction, Duke University</i> , 2011.
TUTORIALS	Graph-Based User Behavior Modeling: From Prediction to Fraud Detection Alex Beutel, Leman Akoglu, Christos Faloutsos. <i>KDD</i> , 2015
	Fraud Detection through Graph-Based User Behavior Modeling Alex Beutel, Leman Akoglu, Christos Faloutsos. <i>ACM CCS</i> , 2015
TALKS	ML for Data Systems Stanford EE380 Colloquium, Palo Alto, CA, October 2018
	Q&R: A Two-Stage Approach Toward Interactive Recommendation <i>KDD</i> , London, UK, August 2018
	Dynamics and Context in Neural Recommender Systems <i>LearnIR Workshop at WSDM</i> , Los Angeles, CA, February 2018
	Latent Cross: Making Use of Context in Recurrent Recommender Systems <i>WSDM</i> , Los Angeles, CA, February 2018
	Dynamics and Context in Neural Recommender Systems Pinterest, San Francisco, CA, February 2018
	A Machine Learning Approach to Databases Indexes <i>ML Systems at NeurIPS</i> , Long Beach, CA, December 2017
	Using Context when Modeling User Behavior: Improving Fraud Detection, Neural Recommenders, and Fairness M.I.T., Cambridge, MA, November 2017
	Using Context when Modeling User Behavior: Improving Fraud Detection, Neural Recommenders, and Fairness Brown University, Providence, RI, November 2017
	Beyond Globally Optimal: Focused Learning for Improved Recommendations Google Student Research Summit, Mountain View, CA, September 2017
	Beyond Globally Optimal: Focused Learning for Improved Recommendations <i>WWW</i> , Perth, Australia, April 2017

Beyond Who and What: Answering How and Why by Modeling Large Graphs
 Northeastern University, Boston, MA, March 2016

Beyond Who and What: Answering How and Why by Modeling Large Graphs
 Arnhold Institute for Global Health, Mount Sinai School of Medicine, New York, NY, March 2016

Beyond Who and What: Answering How and Why by Modeling Large Graphs
 IOMS, Stern School of Business, New York University, New York, NY, March 2016

Beyond Who and What: Answering How and Why by Modeling Large Graphs
 Google Research, Mountain View, CA, March 2016

Beyond Who and What: Answering How and Why by Modeling Large Graphs
 Microsoft, Redmond, WA, March 2016

Beyond Who and What: Answering How and Why by Modeling Large Graphs
 Georgia Institute of Technology, Atlanta, GA, February 2016

Beyond Who and What: Answering How and Why by Modeling Large Graphs
 New York University, Courant Institute, New York, NY, February 2016

ACCAMS: Additive Co-Clustering to Approximate Matrices Succinctly
 University of Pennsylvania, Philadelphia, PA, November 2015

Collaborative Bayesian Filtering: Patterns and Methods
 WIN, New York, NY, October 2015

ACCAMS: Additive Co-Clustering to Approximate Matrices Succinctly
 WWW, Florence, Italy, May 2015

Distributed Machine Learning for User Behavior Modeling
 Facebook, New York, NY, May 2015

Distributed Machine Learning for User Behavior Modeling
 Google Research, New York, NY, May 2015

CoBaFi: Collaborative Bayesian Filtering
 WWW, Seoul, South Korea, April 2014

SGD on Hadoop for Big Data and Huge Models
 Duke University, Durham, NC, 2014

CopyCatch: Stopping Group Attacks by Spotting Lockstep Behavior in Social Networks
 WWW, Rio de Janeiro, Brazil, May 2013

Interacting Viruses on a Network: Can both survive?
 KDD, Beijing, China, August 2012

TerraNNI: Natural Neighbor Interpolation on a 3D Grid Using a GPU
 ACM GIS, Chicago, IL, November 2011

Natural Neighbor Interpolation Based Grid DEM Construction Using a GPU
 ACM GIS, San Jose, CA, November 2010

TEACHING
 EXPERIENCE

Guest Lecture: Intro to Data & Computational Science (Brown DATA 1030) **Fall 2017**
 “Building Blocks of Neural Networks and Research Applications of RNNs”

Guest Lecture: Machine Learning with Large Datasets (CMU 10-805) **Spring 2015**
 “SGD on Hadoop for Big Data and Huge Models”

Guest Lecture: Machine Learning with Large Datasets (CMU 10-805) **Spring 2014**
 “SGD on Hadoop for Big Data and Huge Models”

Teaching Assistant: Database Applications (CMU 15-415/615) **Spring 2014**

Teaching Assistant: Multimedia DB & Data Mining (CMU 15-826) **Fall 2013**

PATENTS

Elastic multi-resolution model-serving to compute inferences, Christopher Olston, Noah Fiedel, Ed H. Chi, Alexander Beutel. Defensive Publication 668.

Detection of Lockstep Behavior, Alex Beutel and Wanhong Xu.
Patent number 9077744; issued July 7, 2015.

STUDENTS
MENTORED

Francois Belletti (2017, UC Berkeley; next position: Google Research)
Konstantina Christakopoulou (2017, UMN; next position: Google Research)
Stephen Macke (2018, UIUC)
Candice Schumann (2018, UMD)
Sahaj Garg (2018, Stanford undergraduate)

SERVICE

PC Member: WSDM 2018, 2019
PC Member: KDD 2017, 2018, 2019
PC Member: SDM 2017, 2018, 2019
PC Member: WWW 2017, 2018, 2019
PC Member: FAT 2019
PC Member: SIGMOD 2019
PC Member: ORSUM 2018
PC Member: SocInfo 2016
PC Member: IEEE DSAA Special Session on Big Behavioral Data Analytics 2016
PC Member: ACM/IEEE ASONAM 2016
Co-organizer: Workshop on Machine Learning Systems, NeruIPS 2015
PhD Forum Committee, ICDM 2015
Mentor at Doctoral Consortium, WSDM 2018
PC Member: Special Session on Big Behavioral Data Analytics, IEEE DSAA 2015
PC Member: Web Information System Engineering (WISE), 2014
PC Member: Diffusion Networks and Cascade Analytics Workshop, WSDM 2014
Reviewer: TKDD, TKDE, NeruIPS, ICML, ICLR, INFORMS Journal on Computing, Neurocomputing