Phone: 832-767-8651
Email: ankur.p.parikh@gmail.com
http://www.cs.cmu.edu/~apparikh

Ankur Parikh

Research Summary

My research focuses on developing effective **machine learning** paradigms to tackle complex applications in the context of modern data/computational resources, particularly in **natural language processing**. My style balances between broad abstractions and narrow domain-specific solutions, and I strive to develop theoretically principled, scalable, and easily reproducible methods that give strong empirical results on real applications.

Employment

Google New York - Research Scientist
Aug 2015 - present
Research new methodologies for natural language understanding, and deploy these technologies to improve Google products.

Education

Carnegie Mellon Machine Learning Department - Ph.D. in Machine Learning
Aug 2009 - May 2015, M.S. degree received in 2012

- *GPA*: 4.0 / 4.0 (A+ = 4.33)
- *Research Advisor*: Prof. Eric P. Xing
- *Thesis Title*: Spectral Probabilistic Modeling and Applications to Natural Language Processing
- *Thesis Committee*: Drs. Eric Xing (chair), Geoff Gordon, John Platt (Microsoft Research), Noah Smith, Le Song (Georgia Tech)

Princeton University, B.S.E. in Computer Science with Certificate in Applied and Computational Mathematics
Sept 2005 - June 2009

- *GPA*: 3.85 / 4.0, (CS classes: 3.90 / 4.0) **summa cum laude**
- *Research Advisors*: Prof. Robert Schapire and Prof. Sharad Malik
- Phi Beta Kappa and Tau Beta Pi

Awards

- **Best Paper Runner Up**, Conference on Empirical Methods in Natural Language Processing (EMNLP 2014)
- **Best Paper in Translational Bioinformatics**, International Conference on Intelligence Systems for Molecular Biology (ISMB 2011)
- **NSF Graduate Research Fellowship 2011**
- Nominated for MLD graduate student teaching award
- Phillip Goldman Senior Prize in Computer Science

Internships

- **Microsoft Research**– Research Intern, Natural Language Processing Group
  Summer 2013
  - *Mentors*: Drs. Hoifung Poon and Kristina Toutanova

  - *Topic*: Weakly Supervised Complex Knowledge Extraction for Biomedical Literature

- **Microsoft Research** - Research Intern, Machine Learning Group
  Summer 2011
  - *Mentors*: Drs. Asela Gunawardana and Christopher Meek
  - *Topic*: Scalable Modeling of Event Streams for Large Structured Label Spaces

- **Microsoft** - Software Development Engineer Intern, Search Platform
  Summer 2009
  - *Mentors*: Puneet Sahni and Dr. Balaji Shyamkumar
  - *Topic*: Improving web crawling/indexing algorithms

- **Google** - Intern Engineer, Java Platforms Group
  Summer 2008
  - *Mentor*: Dan Grove
  - *Topic*: Exposing garbage collection inside Java Virtual Machine

Teaching

- **Guest Lecturer at New York University (NYU)**
  Fall 2016, Statistical Natural Language Processing (1 lecture)

- **Spectral Learning Tutorial at Microsoft Research**
  Summer 2013
  - *Title*: Spectral Learning for Graphical Models: An Intuitive Introduction with a Focus on NLP

- **Guest Lecturer** - Probabilistic Graphical Models
  Spring 2012/2013
  - 3-lecture unit on Kernel Graphical Models and Spectral Learning
  - Designed all instructional material, since area was very nascent at the time and few (if any) teaching materials existed
  - Prof. David Sontag at NYU later used a subset of the slides for his class

- **Teaching Assistant** - Probabilistic Graphical Models (advanced graduate class)
  Spring 2012

- **Teaching Assistant** - Great Theoretical Ideas in Computer Science (undergraduate class)
  Fall 2011

Mentoring

- Micol Marchetti-Bowick (PhD student) - machine learning/computational biology
  Nov 2013 - current
  - *Results*: Hybrid Multiview Subspace Learning for High Dimensional Data (under review)

- Mrinmaya Sachan (PhD student) - machine learning/natural language processing
  Oct 2014 - current

Community

- Primary organizer of NIPS 2014 Modern Machine Learning and Natural Language Processing Workshop (with Avneesh Saluja, Chris Dyer, and Eric Xing)

- Co-organizer of ICML 2013 Spectral Learning Workshop
  (with Byron Boots, Borja Balle, and Daniel Hsu)
- Primary organizer of NIPS 2012 Spectral Algorithms for Latent Variables Workshop
  (with Le Song and Eric Xing)
- CMU ML PhD Admissions Committee Spring 2013 / 2014
  1 of 9 members (2 faculty and 7 students) that evaluate 300+ applicants
- Co-organizer of Carnegie Mellon Machine Learning Lunch Seminar 2010, 2012

Publications

Conference/ Journal Papers

∗ = equal contributing author

[1] Ankur P Parikh, Oscar Täckström, Dipanjan Das, and Jakob Uszkoreit. A Decomposable Attention Model for Natural Language Inference. In *Conference on Empirical Methods in Natural Language Processing (EMNLP)*, 2016, short paper.

[2] Ankur P Parikh, Hoifung Poon, and Kristina Toutanova. Grounded semantic parsing for complex knowledge extraction. In *Conference of the North American Chapter of the Association for Computational Linguistics (NAACL)*, 2015

[3] Ankur P Parikh, Avneesh Saluja, Chris Dyer, and Eric P Xing. Language modeling with power low rank ensembles. In *Conference on Empirical Methods in Natural Language Processing (EMNLP)*, 2014, **Best Paper Runner Up**.

[4] Ankur P Parikh, Shay B Cohen, and Eric P Xing. Spectral unsupervised parsing with additive tree metrics. In *Proceedings of the 52nd Annual Meeting of the Association for Computational Linguistics (ACL)*, 2014.

[5] Ankur P Parikh, Ross E Curtis, Irene Kuhn, Sabine Becker-Weimann, Mina Bissell, Eric P Xing, and Wei Wu. Network analysis of breast cancer progression and reversal using a tree-evolving network algorithm. *PLoS computational biology*, 2014.

[6] Ankur P Parikh, Wei Wu, and Eric P Xing. Robust reverse engineering of dynamic gene networks under sample size heterogeneity. In *Pacific Symposium on Biocomputing (PSB)*, 2014.

[7] Le Song, Mariya Ishteva, Ankur P Parikh, Haesun Park, and Eric P Xing. Hierarchical tensor decomposition of latent tree graphical models. In *Proceedings of the 30th International Conference on Machine Learning (ICML)*, 2013.

[8] Ankur P Parikh, Le Song, Mariya Ishteva, Gabi Teodoru, and Eric P Xing. A spectral algorithm for latent junction trees. *The 28th Conference on Uncertainty in Artificial Intelligence (UAI)*, 2012.

[9] Qirong Ho, Ankur P Parikh, and Eric P Xing. A multiscale community blockmodel for network exploration. *Journal of the American Statistical Association*, 2012 (earlier version in AISTAT 2011).

[10] Ross E Curtis, Jing Xiang, Ankur P Parikh, Peter Kinnaird, and Eric P Xing. Enabling dynamic network analysis through visualization in tvnviewer. *BMC bioinformatics*, 2012.

[11] Le Song, Ankur P Parikh, and Eric P Xing. Kernel embeddings of latent tree graphical models. In *Advances in Neural Information Processing Systems (NIPS)*, 2011.

[12] Ankur P Parikh*, Wei Wu*, Ross E Curtis, and Eric P Xing. Treegl: reverse engineering tree-evolving gene networks underlying developing biological lineages. In *the Nineteenth International Conference on Intelligence Systems for Molecular Biology (ISMB), Published in Bioinformatics*, 2011, **Best Paper in Translational Bioinformatics**.

[13] Ankur P Parikh, Le Song, and Eric P Xing. A spectral algorithm for latent tree graphical models. In *Proceedings of the 28th International Conference on Machine Learning (ICML)*, 2011.

[14] Qirong Ho, Ankur P Parikh, Le Song, and Eric P Xing. Multiscale community blockmodel for network exploration. In *Proceedings of the 14th International Conference on Artificial Intelligence and Statistics (AISTAT)*, 2011.

[15] Mladen Kolar, Ankur P Parikh, and Eric P Xing. On sparse nonparametric conditional covariance selection. In *Proceedings of the 27th International Conference on Machine Learning (ICML)*, 2010.

Patents

[16] Asela JR Gunawardana, Christopher A Meek, and Ankur P Parikh. Forecasting a future event in an event stream, October 3 2013. US Pub. No. US 20130262369 A1.

Talks

- A Decomposable Attention Model for Natural Language Inference
  - *11/2016-* Conference on Empirical Methods in Natural Language Processing (EMNLP), short paper talk
- Language Modeling with Power Low Rank Ensembles
  - *10/2014-* Conference on Empirical Methods in Natural Language Processing (EMNLP), best paper runner up talk
  - *10/2014-* University of Pennsylvania CLUNCH
  - *10/2014* - MIT CSAIL Natural Language Processing Group
- Spectral Unsupervised Parsing with Additive Tree Metrics
  - *10/2014-* Harvard Computer Science
  - *6/2014-* Meeting of the Association of Computational Linguistics (ACL), paper talk
  - *2/2014* - Carnegie Mellon Machine Learning Lunch Seminar
- Robust Reverse Engineering of Dynamic Gene Networks under Sample Size Heterogeneity
  - *1/2014-* Pacific Symposium on Biocomputing (PSB), paper talk
- Spectral Learning for Graphical Models: An Intuitive Introduction with a Focus on NLP
  - *7/2013-* Tutorial at Microsoft Research
- Reverse Engineering Tree-Evolving Gene Networks Underlying Developing Biological Lineages
  - *7/2011-* International Conference on Intelligence Systems for Molecular Biology (ISMB), paper talk
- A Spectral Algorithm for Latent Tree Graphical Models
  - *6/2011-* International Conference on Machine Learning (ICML), paper talk
- On Sparse Nonparametric Conditional Covariance Selection

- *9/2010* - Carnegie Mellon Machine Learning Lunch Seminar
- *6/2010*- International Conference on Machine Learning (ICML), paper talk

References

- **Dipanjan Das**, Staff Research Scientist, Google New York, dipanjand@google.com
- **Slav Petrov**, Senior Staff Research Scientist, Google New York, slav@google.com
- **Eric Xing**, Professor, Carnegie Mellon, epxing@cs.cmu.edu
- **Chris Dyer**, Assistant Professor, Carnegie Mellon (on leave at Google Deepmind), cdyer@cs.cmu.edu
- **Le Song**, Assistant Professor, Georgia Tech, lsong@cc.gatech.edu