

Research Interests

- How can we efficiently evaluate open-domain language tasks like information extraction or summarization with humans in the loop?
- How can we leverage natural language to make it easier for *people* to consume information?

Education

- **Stanford University** California, USA
PhD. in Computer Science Aug. 2012 — present
 - Focus Areas: Statistical Learning Theory, Machine Learning, Natural Language Processing.
- **Indian Institute of Technology, Madras** Chennai, India
BTech./MTech. in Computer Science and Engineering; **CGPA: 9.24/10** Aug. 2007 — June 2012
 - Major: **Computer Science and Engineering**; Minor: **Physics**

Publications

- (Upcoming) A. Chaganty, A. Paranjape, K. Clark, J. Lei, A. See, M. Lamm, J. Bolton, S. Reddy, Y. Zhang, P. Qi, C. D. Manning, “*Stanford’s 2017 KBP System*”. Text Analysis Conference 2017.
- A. Chaganty, P. Liang “*Importance sampling for on-demand evaluation of knowledge base population*”. Empirical Methods in Natural Language Processing (EMNLP) 2017.
- Y. Zhang, A. Paranjape, D. Chen, A. Chaganty, J. Bolton, P. Qi, C. D. Manning, “*Stanford’s 2016 KBP System*”. Text Analysis Conference 2016.
- A. Chaganty, P. Liang “*How Much is 131 Million Dollars? Putting Numbers in Perspective with Compositional Descriptions*”. Association for Computational Linguistics (ACL) 2016.
- G. Angeli, D. Chen, V. Zhong, A. Chaganty, K. Clark, C. D. Manning, “*Stanford’s 2015 KBP System*”. Text Analysis Conference 2015.
- K. Werling, A. Chaganty, P. Liang, C. Manning “*On the Job Learning with Bayesian Decision Theory*”. Neural Information Processing Systems (NIPS) 2015.
- S. Wang, A. Chaganty, P. Liang, “*Estimating Mixture Models via Mixtures of Polynomials*”. Neural Information Processing Systems (NIPS) 2015.
- V. Kuleshov*, A. Chaganty*, P. Liang, “*Tensor Factorization via Matrix Factorization*”. Artificial Intelligence and Statistics (AISTATS) 2015.
- A. Chaganty, P. Liang, “*Estimating Latent Variable Graphical Models with Moments and Likelihoods*”. International Conference on Machine Learning (ICML) 2014.
- G. Angeli, A. Chaganty, A. Chang, K. Reschke, J. Tibshirani, J. Wu, O. Bastani, K. Siilats, C. D. Manning, “*Stanford’s 2013 KBP System*”. Text Analysis Conference 2013.
- A. Chaganty, P. Liang, “*Spectral Experts for Estimating Mixtures of Linear Regressions*”. International Conference on Machine Learning (ICML) 2013.
- A. Chaganty, A. Lal, A. Nori, S. Rajamani, “*Combining Relational Learning with SMT Solvers using CEGAR*”. Computer Aided Verification (CAV) 2013.
- A. Chaganty, A. Nori, S. Rajamani, “*Efficiently Sampling Probabilistic Programs via Program Analysis*”. AI & Statistics (AISTATS) 2013.

- A. Chaganty, B. Ravindran, “*Learning in a Small World*”. Autonomous Agents and Multi-Agent Systems (AAMAS) 2012.
- A. Chaganty “*Inter-task Learning with Spatio-Temporal Abstractions*”. Master’s thesis (IIT Madras) 2012.

Awards

- **Stanford Graduate Fellowship (2014 – 2017):** Fellowship sponsored by *J. Hewes Crispin and Marjorie Holmes Crispin*.
- **Roberto Padovani Scholarship (2009):** for “stellar performance” on my internship assignment at Qualcomm 2009.
- **Kishore Vaigyanik Protsahan Yojana (2006–08):** Fellowship awarded by the Department of Science and Technology, Govt. of India to promote interest in the basic sciences.

Research Experience

<http://arun.chagantys.org/research>

- **On-demand evaluation** Stanford University
with Ashwin Paranjape, Percy Liang, Christopher Manning January 2017 – present
 - How can we scalably introduce human feedback into when evaluating open-ended language tasks like information extraction or summarization?
 - Proposed framework and solution for knowledge base population based on crowdsourcing and importance-reweighted estimators. (**EMNLP 2017**).
 - Currently working on a solution to evaluate automatic summarization.
- **Natural language numeracy** Stanford University
with Matt Lamm, Percy Liang, Christopher Manning January 2016 – present
 - People best understand concepts through comparisons: we provide a system to generate these comparisons for numerical expressions in text (**ACL 2016**).
 - Numeric comparisons also are very common in the news, but hard to identify because their definition emerges only in context. We define an explicit representation, called a *contrast frames*, for such comparisons and build a system to identify such frames in text.
- **On-the-job Learning** Stanford University
with Keenon Werling, Percy Liang May 2015 – October 2015
 - How can we produce useful predictive output starting without any data, by allowing our systems to query the crowd at test time?
 - Proposed framework and a solution based on Bayesian Decision Theory. (**NIPS 2015**).
 - How do we make complex tasks amenable to a lay-crowd-worker?
- **Learning Latent Variable Models** Stanford University
with Percy Liang September 2012 – September 2015
 - Can we efficiently learn latent variable models with guarantees (as opposed to EM)?
 - Used linear regression and the method of moments to develop a statistically consistent algorithm for a mixture of linear experts model with polynomial sample and computational complexity (**ICML 2013**).
 - An algorithm to learn parameters for any discrete graphical model satisfying a ‘uniformly bottlenecked’ assumption; the family includes models with high treewidth as well as log-linear models (**ICML 2014**).
 - An improved algorithm for tensor factorization through random projections and simultaneous matrix diagonalization (with Volodymyr Kuleshov at **AISTATS 2015**).
 - Recovery for any mixture model with polynomial moments via reduction to the generalized moment problem (with Sida Wang at **NIPS 2015**).
- **Relation Extraction for Knowledge Base Population** Stanford University
with Christopher Manning January 2013 – September 2015
 - Using logical inference in a distantly supervised setting to populate relations in a knowledge base.
 - Our entry was among the *top 5* at the **TAC-KBP 2013** General Slot Filling competition.
 - Our entry was the top entry at the **TAC-KBP 2015** and **TAC-KBP 2016** Cold Start competitions.

- **Program Analysis meets Probabilistic Programs** Microsoft Research India
with Aditya Nori, Sriram Rajamani May 2011 – July 2012
 - Used dynamic analysis and concolic execution to efficiently sample from probabilistic programs by avoiding invalid states in both an importance sampling and Metropolis-Hastings setting (**AISTATS 2013**).
 - Applied Counter-Example Guided Abstraction Refinement, and generalization (from program analysis) to the Markov Logic Network framework, with significant performance improvements over prior art (**CAV 2013**) .
 - Filed for patent: *“Probabilistic Model Approximation for Statistical Relational Learning”*.

Teaching and Leadership Experience

- **President, Asha for Education – Stanford** California, USA
Stanford University Aug 2014 – Aug 2016
 - Non-profit focused on improving education for children in India.
- **Teaching Assistant for CS224D (Deep learning for NLP)** California, USA
Stanford University January 2017 – Apr 2017
- **Teaching Assistant for CS221 (Artificial Intelligence)** California, USA
Stanford University Sept 2013 – Dec 2013

Development Experience

- **Dynamic Firewalls for Micro-budgeting Connectivity** San Diego, CA, USA
Qualcomm Corporate R&D May 2009 – July 2009
 - Compared the performance of `iptables` on Linux, and Windows Filtering Platform on Windows, and wrote a WFP driver to implement a dynamic firewall on Windows. Was awarded the **Roberto Padovani Scholarship** from Qualcomm in 2009 for “stellar performance” on my internship assignment.
- **Integrating Vim with the Anjuta IDE** Hyderabad, India
Google Summer of Code '08 May 2008 – Aug 2008
 - Integrating the popular Vim text editor into Anjuta, an open-source IDE in collaboration with the GNOME Foundation, under the Google Summer of Code Program.