Arun Tejasvi Chaganty

chaganty@cs.stanford.edu
http://arun.chagantys.org

Work Experience

- **Square Inc.** — California, USA
  Artificial Intelligence Lead — May 2019 —
  - Responsible for the design and implementation of conversational AI for the Conversations team. This includes but is not limited to the collection and overseeing the annotation of data, design and training of machine learning models, and the development of a conversational dialogue engine.
  - Responsible for coordinating work for a team of 3–4 AI engineers to execute the above tasks.
  - Responsible for communicating the capabilities and limitations of the conversational AI to other groups at Square and the Conversations team.
- **Eloquent Labs Inc.** — California, USA
  Head of Artificial Intelligence — November 2018 — May 2019
  - Lead research and development on conversational artificial intelligence (AI). This includes but is not limited to the collection and overseeing the annotation of data, design and training of machine learning models, and the development of a conversational dialogue engine.
  - Responsible for coordinating work for a team of 2–3 AI engineers to execute the above tasks.
  - Responsible for client-facing communication regarding the capabilities and limitations of Eloquent's conversational AI.

Education

- **Stanford University** — California, USA
  PhD. in Computer Science — Aug. 2012 — September 2018
  - Focus Areas: Natural Language Processing, Statistical Learning Theory, Machine Learning.
- **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**

Awards

- **Stanford Graduate Fellowship (2014–17)**
- **Roberto Padovani Scholarship (2009):** one of about seven awardees for "stellar performance" on my internship assignment at Qualcomm Research.

Research Overview

- **Mimic rephrasal for empathetic listening** — Stanford University
  with Justin Dieter, Tian Wang, Angel Chang, Gabor Angeli — June 2018 – June 2019
  - How can automated assistant (chatbots) perform empathetic listening when talking to users who ask it questions beyond its capabilities? We show that adopting a mimic and rephrase approach using a sequence-to-sequence neural network is able to generate responses on par with humans. **(CoNLL 2019)**
- **On-demand evaluation** — Stanford University
  with Ashwin Paranjape, Steve Mussmann, Percy Liang, Christopher Manning — January 2017 – present
  - Can we scalably evaluate open-ended language tasks like information extraction or summarization with human feedback? We show fundamental limitations with existing automatic metrics. **(ACL 2018)**
  - Proposed a framework and solution for knowledge base population based on crowdsourcing and importance-reweighted estimators that decreased annotation costs by a factor of 4. **(EMNLP 2017)**.
- **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, such as describing Cristiano Ronaldo's signing fee of $131 million as roughly the amount it would take to pay everyone in Kansas City the median salary for a week. **(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 *textual analogy frames*, for such comparisons and build a system to identify such frames in text. **(LREC 2018; EMNLP 2018)**

- **On-the-job Learning** Stanford University
  with Keenon Werling, Percy Liang May 2015 – October 2015
  - How do we produce reliable predictive models starting with zero-training data? We propose a framework that allows our systems to query the crowd at test time using Bayesian Decision Theory. **(NIPS 2015)**.

- **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** Stanford University (and Google)
  with Christopher Manning January 2013 – September 2017
  - Using logical inference in a distantly supervised setting to populate relations in a knowledge base.
  - Part of / led Stanford team at **TAC-KBP 2013, 2015–17**. Our entry was the top-ranked at the **TAC-KBP 2015–17** Cold Start tracks.
  - Using collective inference over the graph of extracted relations to improve precision and recall (*at Google*).

- **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)** .
  - Awarded patent: *"Probabilistic Model Approximation for Statistical Relational Learning"*.

Leadership Experience

- **Leader, Stanford TAC-KBP** California, USA
  Stanford University May 2017 – Sept 2017
  - Organized team while participating the TAC-KBP competition (for knowledge base population).
  - We expanded our system to three languages (English, Chinese and Spanish) and were the top ranked system in the slotfilling track.

- **President, Asha for Education – Stanford** California, USA
  Stanford University Aug 2014 – Aug 2016
  - Non-profit focused on improving education for underprivileged children in India.
  - Improved chapter health by recruiting new students, introduced a new regular talk series, fostered collaboration with other non-profits and helped start a new undergraduate chapter.

Publications (19 publications; in total, cited at least 425 times as of November 2019)

- J. Dieter, T. Wang, G. Angeli, A. Chang, A. Chaganty, *"Mimic and Rephrase: Reflective Listening in Open-Ended Dialogue" "*. Computational Natural Language Learning (CoNLL) 2019.

- M. Lamm, A. Chaganty, D. Jurafsky, P. Liang, C. D. Manning, *"Textual Analogy Parsing: What's Shared and What's Compared among Analogous Facts" "*. Empirical Methods in Natural Language Processing (EMNLP) 2018.
- A. Chaganty, S. Mussmann, P. Liang, *"The price of debiasing automatic metrics in natural language evaluation"*. Association for Computation Linguistics (ACL) 2018.
- M. Lamm, A. Chaganty, D. Jurafsky, P. Liang, C. D. Manning, *"QSRL: A semantic role-labeling schema for quantitative facts"*. Financial Narrative Processing Workshop, LREC 2018.
- 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.