

Arun Tejasvi Chaganty

Senior Research Scientist

 arunchaganty@gmail.com

 arun.chagantys.org

 github.com/arunchaganty

EXPERIENCE

Google

Senior Research Scientist

Dec 2020–(today)

- Created **dialog inpainting** (co-first author), a technique to generate millions of information-seeking conversations from documents using language models (T5 S–XXL). Implemented the entire bulk inference pipeline (average throughput of ~3k inference calls/s) using Apache Beam. Led human evaluation and safety analysis. Trained masked language models and retrieval models.
- Created the **Conversational Playlist Curation Dataset** (first author; PI), one of the first resources for conversational recommendation with multiple item ratings per-turn. Designed and implemented human-human methodology, including all annotation interfaces.
- Developed **Talk the Walk** (PI), a recipe to generate millions of (music) recommendation-seeking conversations from existing playlists using a combination of random walks and language models. Bootstrapped an end-to-end conversation recommendation system that significantly outperforms baselines in live experiments.
- Defined task and evaluation methodology for **RARR**, a post-hoc attribution and revision method for large language models (PaLM-540B).

Research Intern

Summer 2014

- Explored multi-sentence relation extraction for knowledge bases.

Square

AI Lead

May 2019–Dec 2020

- Led a small team of AI engineers that built **Square Assistant** —a chatbot we launched in October 2019 that helps customers book and reschedule appointments with Square merchants.
- Designed and shipped **conversational rescheduling** feature that increased booking and rescheduling success rates by helping customers find a concrete time for their appointment; the feature understands temporal constraints in user utterances using a model-based semantic parser.
- Developed a **type-safe domain-specific language** to describe asynchrony and interruptions in dialog flows using coroutines. Implemented Java-to-Java compiler. DSL reduced feature code 10–20x and fixed subtle asynchrony bugs.
- Developed most of the AI model deployment, logging and data annotation infrastructure.

EDUCATION

Stanford University

PhD (Computer Science) 2012-2018

Advised by Percy Liang

Indian Institute of Technology,
Madras

MTech. (Computer Science)

BTech. (Computer 2011-2012

Science) 2007-2011

Minor in Physics GPA: 9.24/10

AWARDS

- Stanford Graduate Fellow ('14–'17)
- Robert Padovani Scholar ('09)
- Google Summer of Code ('08)
- Kishore Vaigyanik Protsahan Yojana Scholar ('06–'07)

SKILLS

Natural Language Processing

Conversational AI • Recommendation

Systems • Synthetic Data Generation

• Evaluation • Retrieval •

Crowdsourcing • Semantic Parsing •

Information Extraction

Machine Learning

Deep Learning • Latent Variable

Models • Probabilistic Programming

Programming

Python (PyTorch, Tensorflow) •

Typescript (Angular, React) • SQL •

Bash • Java • C++

Computer Science

Compilers • Operating Systems •

Computer Networks • Cloud

Computing

Eloquent Labs

Head of AI

June 2018–May 2019

- Led a small team of AI engineers that built a conversational AI system for enterprise customer service. Interfaced with clients directly.
- Developed a **human-in-the-loop system to fine-tune question similarity models** for particular clients; led to 2–3x increases in precision and recall for each client.
- Startup acquired by Square in May 2019.

Stanford University

PhD Candidate

Sept 2012–Sept 2018

- Led / part of the Stanford team at TAC-KBP 2013, 2015–17. Our entry was the top-ranked at the TAC-KBP 2015--17 Cold Start tracks.
- Co-author of **CoreNLP Server**, an extremely popular API server for the Stanford CoreNLP package.
- 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 human-in-the-loop solution for knowledge-base population evaluation that eliminates *pooling bias* using a novel importance-reweighted estimator that decreases annotation costs by a factor of 4 (**EMNLP 2017**).
- Numeric comparisons, while common in the news, are 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 semantic parser to identify such frames in text (**EMNLP 2018**).
- People best understand concepts through comparisons: we provide a system to generate compositional 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**).
- Can we efficiently learn latent variable models with guarantees? We show that this is possible for a variety of models satisfying a 'uniformly bottlenecked' assumption including discriminative mixtures of linear experts (**ICML 2013**), high tree-width models, log-linear models and multi-view Markov random fields (**ICML 2014**). In later work, we show guaranteed recovery for any mixture model with polynomial moments is possible via reduction to the generalized moment problem (**NIPS 2015**). All of these methods require tensor factorization, which we show can be more efficiently performed by reduction to simultaneous matrix diagonalization using random productions (**AISTATS 2015**).

Microsoft Research India

Intern

Summer 2009, 2010, 2011

- 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**).

PUBLICATIONS

1. **A. Chaganty, M. Leszczynski, R. Ganti, S. Zhang, K. Balog, and F. Radlinski (2023).** Beyond Single Items: Exploring User Preferences in Item Sets with the Conversational Playlist Curation Dataset. *In Submission.*
2. **M. Leszczynski, R. Ganti, S. Zhang, K. Balog, F. Radlinski, F. Pereira, and A. Chaganty (2023).** Generating Synthetic Data for Conversational Music Recommendation Using Random Walks and Language Models. *arXiv.*
3. **L. Gao, Z. Dai, P. Pasupat, A. Chen, A. Chaganty, Y. Fan, V. Zhao, N. Lao, H. Lee, D. Juan, and K. Guu (2022).** RARR: Researching and Revising What Language Models Say, Using Language Models. *arXiv.*
4. **Z. Dai, A. Chaganty, V. Zhao, A. Amini, Q. Rashid, M. Green, and K. Guu (2022).** Conformal Retrofitting via Riemannian Manifolds: distilling task-specific graphs into pretrained embeddings. *In submission; Advances in Neural Information Processing Systems (NeurIPS).*
5. **J. Dieter, and A. Chaganty (2020).** Conformal Retrofitting via Riemannian Manifolds: distilling task-specific graphs into pretrained embeddings. *arXiv.*
6. **J. Dieter, T. Wang, G. Angeli, A. Chang, and A. Chaganty (2019).** Mimic and Rephrase: Reflective Listening in Open-Ended Dialogue. *Computational Natural Language Learning (CoNLL).*
7. **M. Lamm, A. Chaganty, C. D. Manning, D. Jurafsky, and P. Liang (2018).** Textual Analogy Parsing: What's Shared and What's Compared among Analogous Facts. *Empirical Methods in Natural Language Processing (EMNLP).*
8. **A. Chaganty, S. Mussmann, and P. Liang (2018).** The price of debiasing automatic metrics in natural language evaluation. *Association for Computational Linguistics (ACL).*
9. **A. T. Chaganty, A. Paranjape, J. Bolton, M. Lamm, J. Lei, A. See, K. Clark, Y. Zhang, P. Qi, and C. D. Manning (2017).** Stanford at TAC KBP 2017: Building a Trilingual Relational Knowledge Graph. *Text Analytics Conference (TAC).*
10. **A. Chaganty, A. Paranjape, P. Liang, and C. Manning (2017).** Importance sampling for unbiased on-demand evaluation of knowledge base population. *Empirical Methods in Natural Language Processing (EMNLP).*
11. **A. T. Chaganty, and P. Liang (2016).** How Much is 131 Million Dollars? Putting Numbers in Perspective with Compositional Descriptions. *Association for Computational Linguistics (ACL).*
12. **Y. Zhang, A. Chaganty, A. Paranjape, D. Chen, J. Bolton, P. Qi, and C. D. Manning (2016).** Stanford at TAC KBP 2016: Sealing Pipeline Leaks and Understanding Chinese. *Text Analytics Conference (TAC).*
13. **S. I. Wang, A. Chaganty, and P. Liang (2015).** Estimating Mixture Models via Mixture of Polynomials. *Advances in Neural Information Processing Systems (NeurIPS).*
14. **K. Werling, A. Chaganty, P. Liang, and C. Manning (2015).** On-the-Job Learning with Bayesian Decision Theory. *Advances in Neural Information Processing Systems (NeurIPS).*
15. **V. Kuleshov, A. Chaganty, and P. Liang (2015).** Tensor factorization via matrix factorization. *Artificial Intelligence and Statistics (AISTATS).*
16. **G. Angeli, V. Zhong, D. Chen, A. Chaganty, J. Bolton, M. Premkumar, P. Pasupat, S. Gupta, and C. D. Manning (2015).** Stanford at TAC KBP 2015: Bootstrapped Self Training for Knowledge Base Population. *Text Analytics Conference (TAC).*
17. **A. Chaganty, and P. Liang (2014).** Estimating Latent-Variable Graphical Models using Moments and Likelihoods. *International Conference on Machine Learning (ICML).*
18. **A. Chaganty, and P. Liang (2013).** Spectral Experts for Estimating Mixtures of Linear Regressions. *International Conference on Machine Learning (ICML).*
19. **G. Angeli, A. Chaganty, A. Chang, K. Reschke, J. Tibshirani, J. Wu, O. Bastani, K. Siilats, and C. D. Manning (2013).** Stanford's 2013 KBP System. *Text Analytics Conference (TAC).*
20. **A. T. Chaganty, A. Lal, A. Nori, and S. Rajamani (2013).** Probabilistic model approximation for statistical relational learning. *US Patent App. 13/308,571.*
21. **A. Chaganty, A. Lal, A. Nori, and S. Rajamani (2013).** Combining Relational Learning with SMT Solvers using CEGAR. *Computer Aided Verification (CAV).*
22. **A. Chaganty, A. Nori, and S. Rajamani (2013).** Efficiently Sampling Probabilistic Programs via Program Analysis. *Artificial Intelligence and Statistics (AISTATS).*
23. **A. Chaganty, P. Gaur, and B. Ravindran (2012).** Learning in a Small World. *Autonomous Agents and Multi-Agent Systems (AAMAS).*
24. **A. Chaganty (2012).** Inter-task Learning with Spatio-Temporal Abstractions. *Master's Thesis.*