

INTERESTS	<b>Natural Language Processing, Deep Learning, Conversational AI, Grounded Language Models</b>	
EXPERIENCE	<b>Founding Engineer @ Samaya AI</b>	(Oct 22 - Ongoing)
	<ul style="list-style-type: none"><li>• <b>ML:</b> Built and deployed multiple Large Language Models (LLMs) for knowledge discovery</li><li>• <b>SWE:</b> Contribute toward frontend (Vue), backend (Flask), infrastructure (k8s, redis, postgres) and DevOps (ArgoCD, GH CI, pants) as needed. Set up documentation, testing and typing best practices.</li><li>• <b>Hiring:</b> Set up interviewing and onboarding processes, conducted 50+ interviews.</li><li>• <b>Product:</b> Identified gaps in product via user interviews and solved them with novel ML components.</li><li>• <b>Leadership:</b> Lead a team of 5 to build ML models &amp; UI to increase response comprehensiveness.</li></ul>	
EDUCATION	<b>Stanford University</b> , Stanford, CA	(2016 - 22)
	<i>PhD in CS advised by Prof. Christopher D. Manning</i>	
	<b>Stanford University</b> , Stanford, CA	(2014 - 16)
	<i>Masters in CS with depth in AI advised by Prof. Jure Leskovec</i> 3.76/4	
	<b>Indian Institute of Technology, Bombay (IITB)</b> , Mumbai, India	(2010 - 14)
	<i>Bachelors (with honors) in CS and Minor in EE advised by Prof. Sunita Sarawagi</i> 9.69/10 (Ranked 2 <sup>nd</sup> in CS)	
RESEARCH OVERVIEW AND PUBLICATIONS	<b>Neural Systems for Informative Conversations</b>	Advisor: Christopher Manning [2020 - now]
	<ul style="list-style-type: none"><li>• We identify human strategies for informative conversations. For one of the strategies, acknowledgements, we improve chatbot utterances using conditional mutual information (<b>NAACL 2021</b>).</li><li>• For open-ended generation, we supervise a neural retriever and a generator with a posterior guide-retriever leading. After training the retriever has higher relevance and the generator is more grounded (with Omar Khattab, Matei Zaharia, Christopher Potts in <b>ICLR 2022</b>).</li><li>• How do large, pretrained language models fail when conversing with humans? We identify failure modes and quantify with robust crowdsourcing (ongoing with Stanford CRFM)</li></ul>	
	<b>Increasing user initiative in social conversations</b>	Advisor: Christopher Manning [2021 - now]
	<ul style="list-style-type: none"><li>• We define user initiative in the context of a social chatbot, propose automatic metrics that correlate with it and identify some strategies to improve it (with Amelia Hardy in <b>SIGDIAL 2021</b>).</li><li>• Abrupt and unnatural turn taking in voice assistants prevents quality social conversations. We define and implement a “speech manager” that replaces ASR &amp; TTS and adds support for backchannel production and handling (ongoing with Siyan Li, under submission to <b>SIGDIAL 2022</b>)</li></ul>	
	<b>Knowledge Base Population</b>	Advisor: Christopher Manning [2016-17]
	<ul style="list-style-type: none"><li>• We built KBP systems as part of the winning team in in TAC-KBP shared task for 2016, 2017</li><li>• Proposed an importance-reweighted estimator for KBP evaluation and implemented a crowdsourced leaderboard (with Arun Chaganty &amp; Percy Liang in <b>EMNLP 2017</b>)</li></ul>	
	<b>Motifs in Temporal Networks</b>	Advisor: Jure Leskovec [2014-16]
	<ul style="list-style-type: none"><li>• Devised new algorithms to count temporal motifs efficiently (on timestamped edges in a graph) and used them to find patterns in various communication datasets (with Austin Benson in <b>WSDM 2017</b>)</li></ul>	
	<b>Automated Link Prediction</b>	Advisor: Jure Leskovec [2014-16]
	<ul style="list-style-type: none"><li>• Predicted missing links between web pages by learning from human navigational patterns: actively collected paths (with Robert West; <b>WWW 2015</b>) and passively collected server logs (<b>WSDM 2016</b>)</li></ul>	
LEADERSHIP	<b>Co-led Stanford’s Alexa Prize 2020 Team, Chirpy Cardinal</b>	[2019-20]
	Led our team to create a social chatbot from scratch in 9 months. It was deployed to tens of thousands of users and was rated 3.6/5 on average with top-10% conversations lasting 12 minutes or more. We	

won 2<sup>nd</sup> place (and \$100,000) against experienced teams and created a robust research program. We were one of the first teams to use neural models in practical settings and open-source the entire system.

**Team:** Recruited, onboarded and led a team of 10 students with complementary skills.

**Planning:** Scoped projects with concrete milestones on a year-long roadmap toward the vision.

**Architecture:** Designed a flexible and modular framework for high reliability and low latency.

**Process:** Set up weekly sprints, 1-1s, rotating on-calls, issue tracking and communication briefs.

**Best practices:** Implemented code-reviews, integration tests, style guidelines, code docs + wiki.

**Tracking:** Created metrics to indicate performance and dashboards for root-cause analysis.

**Mentorship:** Mentored junior students in tooling, implementation, analysis and research.

MENTORSHIP	<b>Stanford's Alexa Prize 2021 Team</b> [2020-21] Mentored new team lead and new members to build upon the existing architecture, pipelines and best-practices, and developed research-based solutions. Won 2 <sup>nd</sup> place (and \$100,000) again. <b>CS Undergrad Research (CURIS) Mentor</b> for 2 undergrad students: Jillian Tang and Avanika Narayan <b>MS CS Research Mentor</b> for 2 students: Amelia Hardy and Siyan Li
ACHIEVEMENTS	<ul style="list-style-type: none"><li>• Our team came 2<sup>nd</sup> in <b>Alexa Prize Competition</b> twice (2020 &amp; 2021) both with a \$100,000 award.</li><li>• Received <b>Narotam Sekhsaria Postgraduate Scholarship 2014</b> (12 out of 20,000 applicants)</li><li>• <b>All India Rank 81</b> in IIT-Joint Entrance Examination (2010) from around 450,000 entrants</li><li>• Awarded <b>Gold Medal in Indian National Physics Olympiad 2010</b> (top 35 from all over India) and attended <b>Orientation cum Selection Camp (OCSC) for International Physics Olympiad</b>.</li><li>• <b>All India Rank 3</b> in National Science (2008), <b>4</b> in Cyber (2010), <b>7</b> in Mathematics (2010) Olympiads</li></ul>
INVITED TALKS	<b>Google's N2Formal Reading Group</b> Sept 2021 Hindsight: Posterior-guided training of retrievers for improved open-ended generation <b>L3-AI: The Level 3 AI Assistant Conference</b> June 2021 Neural Generation for Social Conversations <b>23rd International Conference on Text, Speech and Dialogue: Amazon Alexa Webinar</b> Sept 2020 Behind the scenes of building a socialbot
GRANTS	<ul style="list-style-type: none"><li>• Wrote a research proposal for <b>Alexa Prize 2021</b> and was granted <b>\$250,000</b> to support students.</li><li>• Wrote a research proposal for <b>Alexa Prize 2020</b> and was granted <b>\$250,000</b> to support students.</li></ul>
INTERNSHIPS	<b>Facebook Inc.</b>   <i>Natural Language Understanding</i> Sonal Gupta [Summer '19] Investigated label crowding for intent classification and methods to factorize label prediction <b>Microsoft Research, Redmond</b>   <i>Deep Learning For Graphics</i> Brian Guenter [Summer '13] Designed and implemented neural nets with convolution layer and domain specific objectives <b>IST Austria</b>   <i>Formal Methods</i> Prof. K. Chatterjee [Summer '12] Realizing Unrealizable specifications using edit distances
PROGRAMMING PROFICIENCY	<b>Current:</b> Python   PyTorch   Numpy & Scipy   Pandas   SQL   Flask   Javascript   HTML   L <sup>A</sup> T <sub>E</sub> X <b>Past:</b> C++   Rust   Java   C   C#   PHP   Django   Scrapy   Hadoop   Spark   Tensorflow   Keras   Boost
ADVANCED COURSES	Machine Learning   Probabilistic Graphical Models   Convex and Linear Optimization Natural Language Processing   Deep Learning for NLP   Computer Vision   CNNs for Vision Social, Information Networks   Experimental Robotics Complexity Analysis   Randomized Algorithms   Neuromorphic Engineering
TEACHING	<i>Stanford:</i> <b>Head TA</b> for <b>Information Retrieval (CS276)</b>   <b>Natural Language Processing (CS224N)</b> <i>IIT Bombay:</i> Operating Systems and Lab   Computer Networks   Modern Physics   Introductory CS

POSITIONS OF RESPONSIBILITY	Stanford NLP <i>Seminar Organizer</i> (organized 24 talks) <i>Research mentor</i> to Stanford's Alexa Prize Team Stanford NLP Dialogue <i>Reading Group Organizer</i> <i>Co-Lead</i> of Stanford's Alexa Prize Team Stanford NLP Group <i>Social Climate (DEI) Committee</i> member Institute <i>Student Mentor</i> , IIT Bombay Manager, Technovation (Student org. for technical innovation), IIT Bombay <i>Awarded IIT Bombay Institute Organisational Color 2013 (awarded to 14 students out of 9000).</i>	2021 - 22 2020 - 21 2020 - 21 2019 - 20 2019 - 20 2013 - 14 2012 - 13
SERVICE	Reviewer for TKDE 2018, AAAI 2018, ICML 2019, NeurIPS 2019, Repl4NLP 2020	
CONFERENCE PUBLICATIONS	<b>Hindsight: Posterior-guided training of retrievers for improved open-ended generation</b> <i>Ashwin Paranjape</i> , Omar Khattab, Christopher Potts, Matei Zaharia Christopher D. Manning (ICLR 2022) <b>Effective Social Chatbot Strategies for Increasing User Initiative</b> Amelia Hardy, <i>Ashwin Paranjape</i> and Christopher D. Manning (SIGDIAL 2021) <b>Human-like informative conversations via conditional mutual information</b> <i>Ashwin Paranjape</i> , Christopher D. Manning (NAACL 2021) <b>Importance sampling for unbiased on-demand evaluation of knowledge base population</b> Arun Chaganty*, <i>Ashwin Paranjape</i> *, Percy Liang and Christopher Manning (EMNLP 2017) <b>Motifs in Temporal Networks</b> <i>Ashwin Paranjape</i> *, Austin Benson* and Jure Leskovec (WSDM 2017) <b>Improving Website Hyperlink Structure Using Server Logs</b> <i>Ashwin Paranjape</i> *, Robert West*, Jure Leskovec and Leila Zia (WSDM 2016) <b>Mining Missing Hyperlinks from Human Navigation Traces: A Case Study of Wikipedia</b> Robert West, <i>Ashwin Paranjape</i> , and Jure Leskovec: (WWW 2015) <b>Unsupervised Word Sense Disambiguation Using Markov Random Field and Dependency Parser</b> Devendra S. Chaplot, Pushpak Bhattacharya and <i>Ashwin Paranjape</i> (AAAI 2015)	
		* Equal contribution
TECHNICAL REPORTS	<b>Neural, Neural Everywhere: Controlled Generation Meets Scaffolded, Structured Dialogue</b> Ethan A Chi, Caleb Chiam, Trenton Chang, Swee Kiat Lim, Chetanya Rastogi, Alexander Iyabor, Yutong He, Hari Sowrirajan, Avanika Narayan, Jillian Tang, Haojun Li, <i>Ashwin Paranjape</i> , Christopher D. Manning (Alexa Prize Proceedings 2021) <b>Neural Generation Meets Real People: Towards Emotionally Engaging Mixed-Initiative Conversations</b> <i>Ashwin Paranjape</i> , Abigail See, Kathleen Kenealy, Haojun Li, Amelia Hardy, Peng Qi, Kaushik Ram Sadagopan, Nguyet Minh Phu, Dilara Soyulu, Christopher D. Manning (Alexa Prize Proceedings 2020) <b>Stanford at TAC KBP 2017: Building a Trilingual Relational Knowledge Graph</b> Arun Tejasvi Chaganty, <i>Ashwin Paranjape</i> , Jason Bolton, Matthew Lamm, Jinhao Lei, Abigail See, Kevin Clark, Yuhao Zhang, Peng Qi, Christopher D. Manning (TAC 2017) <b>Stanford at TAC KBP 2016: Sealing Pipeline Leaks and Understanding Chinese</b> Yuhao Zhang, Arun Tejasvi Chaganty, <i>Ashwin Paranjape</i> , Danqi Chen, Jason Bolton, Peng Qi, Christopher D. Manning (TAC 2016)	