

RISHABH JOSHI

CONTACT INFORMATION	Google DeepMind Mountain View, CA, USA	E-mail : rishabhjoshi@google.com Links : Webpage , Google Scholar , LinkedIn
EDUCATION	Carnegie Mellon University - School of Computer Science ; Pittsburgh, PA <i>Master of Science in Language Technologies</i> Courses : Machine Learning, Neural Networks for NLP, Dialog Seminar, Deep RL, Multilingual NLP	Aug 2021 CGPA : 3.94
	Birla Institute of Technology and Science (BITS), Pilani ; Pilani, India <i>Bachelor of Engineering (Hons.) in Computer Science</i> Thesis on "Distantly Supervised Relation Extraction using Side Information" at IISc, Bangalore	June 2018 CGPA : 9.28/10
	The Mother's International School ; Delhi, India Class XII, CBSE Class X, CBSE	May 2014 95.6% CGPA : 10/10
EXPERIENCE	Google Deepmind <i>Senior Research Engineer</i> <i>Research Engineer (Formerly Brain Team)</i>	Mountain View, CA Apr 2024 - Present Oct 2021 - Apr 2024
	<ul style="list-style-type: none">Proposed multiple research works aimed towards improving alignment to human feedback of large language models using preference optimization that ultimately end up in Google's LLMs.Core-contributor to Bard, Search, Gemini and Gemini 1.5 LLM post-training specifically, SFT, RM training and RLHF / preference optimization, and improving instruction tuning.Performed research in self-improving language models and scaling up RL training.Contributed to research and deployment of internal summarization models which were used in multiple Google products and had SOTA summarization performance across all models available.	
	Carnegie Mellon University <i>Graduate Research Assistant</i> (Advisors : Yulia Tsvetkov , Alan Black , Alex Rudnicky)	Pittsburgh, PA Sept 2019 - Aug 2021
	<ul style="list-style-type: none">Proposed an interpretable approach for knowledge proliferation in academic research articles by identifying key phrases using importance attribution.Performed a linguistic driven analysis on the characteristics of good negotiators and incorporated explicit strategy-sequence structure using Graph Neural Networks to improve non-collaborative dialogue systems.Collaborated with team Tartan to develop a conversational system in Amazon Alexa Challenge 2019; developed mini-bots and analyzed conversational story structure, reached semi-finals.Modeled the structure of dialogue by analyzing the flow of conversational topics and detecting non-coherence. Submitted proposal for Alexa Challenge 2020.	
	Samsung Research Institute <i>Software Engineer - Natural Language Understanding</i>	Bangalore, India Aug 2018 - Aug 2019
	<ul style="list-style-type: none">Improved open domain dialogue systems using side information and contextual knowledge.Constructed a low resource intent classification and speaker recognition solution using GMMs.	
	Indian Institute of Science <i>Undergraduate Thesis</i> (Advisor : Partha Talukdar)	Bangalore, India Jan 2018 - July 2018
	<ul style="list-style-type: none">Constructed a unique India-Centric Knowledge Graph based on the Never-Ending Language Learning Paradigm.Improved Distantly-Supervised Relation Extraction using Side Information achieving SOTA results.	
	Samsung Research Institute <i>Summer Intern - 5G Team</i>	Bangalore, India May 2017 - July 2017
	<ul style="list-style-type: none">Developed tools for packet generation and distributed analysis for the data link layer of 5G protocol.	
	Indian Institute of Remote Sensing, ISRO <i>Research Intern</i> (Advisor : Sameer Saran)	Dehradun, India May 2016 - July 2016
	<ul style="list-style-type: none">Developed core API and execution engine of the DataCube for the effective storage, retrieval and analysis of large earth observation datasets using Python and ideas from distributed and parallel computing.	
PUBLICATIONS	Pierre Harvey Richemond, Yunhao Tang, Daniel Guo, Daniele Calandriello, Mohammad Gheshlaghi Azar, Rafael Rafailov, Bernardo Avila Pires, Eugene Tarassov, Lucas Spangher, Will Ellsworth, Aliaksei Severyn, Jonathan Mallinson, Lior Shani, Gil Shamir, Rishabh Joshi , Tianqi Liu, Remi Munos, Bilal Piot. Offline Regularised Reinforcement Learning for Large Language Models Alignment.	[Paper]

Daniele Calandriello, Daniel Guo, Remi Munos, Mark Rowland, Yunhao Tang, Bernardo Avila Pires, Pierre Harvey Richemond, Charline Le Lan, Michal Valko, Tianqi Liu, **Rishabh Joshi**, Zeyu Zheng, Bilal Piot. Human alignment of large language models through online preference optimisation. *In Proc. ICML '24* [Paper]

Gemini Team, **Rishabh Joshi**. Gemini 1.5 : Unlocking multimodal understanding across millions of tokens of context. [Paper]

Tianqi Liu, Zhen Qin, Junru Wu, Jiaming Shen, Misha Khalman, **Rishabh Joshi**, Yao Zhao, Mohammad Saleh, Simon Baumgartner, Jialu Liu, Peter J Liu, Xuanhui Wang. Lipo : Listwise preference optimization through learning-to-rank. [Paper]

Gemini Team, **Rishabh Joshi**. Gemini : a family of highly capable multimodal models. [Paper]

Polina Zablotskaia, Misha Khalman, **Rishabh Joshi**, Livio Baldini Soares, Shoshana Jakobovits, Joshua Maynez, Shashi Narayan. Calibrating Likelihoods towards Consistency in Summarization Models. [Paper]

Tianqi Liu, Yao Zhao, **Rishabh Joshi**, Misha Khalman, Mohammad Saleh, Peter J Liu, Jialu Liu. Statistical rejection sampling improves preference optimization. *In Proc. ICLR '24*. [Paper]

Yao Zhao, **Rishabh Joshi**, Tianqi Liu, Misha Khalman, Mohammad Saleh, and Peter J. Liu. SLiC-HF : Sequence Likelihood Calibration with Human Feedback. [Paper]

Yao Zhao, Mikhail Khalman, **Rishabh Joshi**, Shashi Narayan, Mohammad Saleh, Peter J Liu. Calibrating Sequence likelihood Improves Conditional Language Generation. *In Proc. ICLR '23*. [Paper]

Rishabh Joshi*, Vidhisha Balachandran*, Emily Saldanha, Maria Glenski, Svitlana Volkova, Yulia Tsvetkov. Unsupervised Keyphrase Extraction via Interpretable Neural Networks. *In Findings EACL '23*. [Paper]

Rishabh Joshi, Vidhisha Balachandran, Shikhar Vashishth, Alan W Black and Yulia Tsvetkov. DialoGraph : Incorporating Interpretable Strategy-Graph Networks into Negotiation Dialogues. *In Proc. ICLR '21*. [Paper]

Ritam Dutt, Sayan Sinha, **Rishabh Joshi**, Surya Shekhar Chakraborty, Meredith Riggs, Xinru Yan, Haogang Bao and Carolyn Rose. ResPer : Computationally Modelling Resisting Strategies in Persuasive Conversations. *In Proc. EACL '21*. [Paper]

Ritam Dutt, **Rishabh Joshi** and Carolyn Rosé. Keeping Up Appearances : Computational Modeling of Face Acts in Persuasion Oriented Discussions. *In Proc. EMNLP '20*. [Paper]

Fanglin Chen, Ta-Chung Chi, Shiyang Lyu, Jiachen Gong, Tanmay Parekh, **Rishabh Joshi**, Anant Kaushik and Alexander Rudnicky. Tartan : A Two-Tiered Dialog Framework for Multi-Domain Social Chitchat. *In Proc. 3rd Alexa Prize '19-'20*. [Paper]

Shikhar Vashishth, **Rishabh Joshi**, Ritam Dutt, Denis Newman-Griffis and Carolyn Rose. MedType : Improving Medical Entity Linking with Semantic Type Prediction. *In Proc. JMR '21*. [Paper] [Demo]

Rishabh Joshi*, Sopan Khosla*, Ritam Dutt*, Alan W Black and Yulia Tsvetkov. LTIatCMU at SemEval-2020 Task 11 : Incorporating Multi-Level Features for Multi-Granular Propaganda Span Identification. *In Proc. SemEval '20*. [Paper]

Rishabh Joshi*, Gaurav Kumar*, Jaspreet Singh* and Promod Yenigalla. AMUSED : A Multi-Stream Vector Representation Method for Use in Natural Dialogue. *In Proc. LREC '20*. [Paper]

Rakesh Bal*, Sayan Sinha*, Swastika Dutta, **Rishabh Joshi**, Sayan Ghosh and Ritam Dutt. Analysing the Extent of Misinformation in Cancer Related Tweets. *In Proc. ICWSM '20*. [Paper]

Shikhar Vashishth, **Rishabh Joshi**, Sai Suman Prayaga, Chiranjib Bhattacharyya and Partha Talukdar. RE-SIDE : Improving Distantly-Supervised Neural Relation Extraction using Side Information. *In Proc. EMNLP '18*. [Paper] [Code]

PROJECTS

Exploring DataMaps for better QA Data Augmentation, CMU Jan 2021 - May 2021

- Proposed effective data augmentation techniques for Machine Reading by utilizing data maps of augmented data. Showed that it is most effective to utilize amniguous data points.

Improving Multilingual Text Summarization using Adapters, CMU Sept 2020 - Jan 2021

- Improved few-shot multilingual abstractive summarization using adapter networks.

Coherence Aware Curious Bot, Advisor : Alex Rudnicky, CMU June - Oct 2020

- Built a curious bot system which learns to ask better questions to keep conversations engaging.
- Leveraged Knowledge Graphs to form novel questions and keep track of dialogue coherence.

Propaganda Identification in News, Advisor : Yulia and Alan, CMU Jan - May 2020

- Proposed a multi-granular framework to incorporate multi-level features for the SemEval 2020 task of detecting propaganda spans in news articles. Performed 4th in span identification task.

TEACHING	<p>Teaching Assistant for Applied Machine Learning 11-663, CMU Jan - Dec 2020</p> <ul style="list-style-type: none"> Assisted Carolyn Rosé by taking office hours for around 100 students and assisting them in their group projects during Spring and Fall 2020. <p>Teaching Assistant for Data Structures and Algorithms, BITS Pilani Jan - May 2017</p> <ul style="list-style-type: none"> Held two lab sections and was the jury for the online judge with the responsibility of assisting 200+ students. The course was taught by Sundar S Balasubramaniam.
HONORS AND AWARDS	<ul style="list-style-type: none"> Honourable Mention, Won a cash prize of \$200 for an oral presentation at CMU LTI SRS 2019 1st, Campus ML Hackathon organized by MapMyIndia on Kaggle in APOGEE, BITS Pilani 2017 3rd, ML Hackathon organized by D.E. Shaw on HackerEarth 2017 2nd, E-Yantra, a robotics competition organized by MHRD and IIT Bombay 2016 Merit, Recipient of Institute Merit Scholarship for academic excellence 2015 1574, All India Rank (AIR), IIT-JEE (out of around 1,300,000) 2014 Inspire, Recipient of Govt. of India Inspire Scholarship (top 1% in CBSE) 2014 All India Topper, Chemistry, CBSE Class XII Board (100/100) School Topper, Computer Science, CBSE Class XII Board (99/100)
POSITIONS OF RESPONSIBILITY	<p>Executive Coordinator, Embryo Club, BITS Pilani Apr 2016 - Apr 2017</p> <ul style="list-style-type: none"> Handled the finances of the club and raised club treasury from Rs. 20,000 to Rs. 180,000. Organized on campus talk by Amish Tripathi, handling all expenses and negotiations. <p>Joint Activities Coordinator, ACM BITS Chapter Aug 2015 - Aug 2016</p> <ul style="list-style-type: none"> Conducted more than 8 events for ACM BITS chapter throughout the year. BITS-ACM won the best ACM chapter award for that year in India.
EXTRA CURRICULAR ACTIVITIES	<p>Volunteered for National Service Scheme, BITS Pilani chapter from Aug 2014 to Aug 2017. Helped underprivileged people to become acquainted in computers. Organized Junoon, sports festival for physically challenged. Lectured junior year students on transition trends of Campus-to-Corporate (C2C) under Sushila Shekhawat.</p>