Title: Key Methods Powering ChatGPT, Bard and LLMs: Instruction Tuning, Prompting and Beyond

Abstract of the talk: Humans have the remarkable ability to solve different tasks by simply reading textual, instructions that define the tasks and looking at a few examples. NLP models built with the conventional machine learning paradigm, however, often struggles to generalize, across tasks (e.g. a question-answering system cannot solve classification tasks) despite training with lots of examples. A long-standing challenge in AI is to build a model that learns a new task by understanding the human-readable instructions that define it. To study this, we build NATURAL INSTRUCTIONS and SUPERNATURAL INSTRUCTIONS, large-scale datasets of diverse tasks, their human-authored instructions, and instances. We adopt generative pre-trained language models to encode task-specific instructions along with input and generate task output. Our results indicate that the instruction-tuning helps models achieve cross-task generalization. This leads to the question: how to write good instructions? Backed by extensive empirical analysis on large language models, we observe important attributes for successful instructional prompts and propose several reframing techniques for model designers to create such prompts. Our results show that reframing notably improves few-shot learning performance; this is particularly important on large language models, such as GPT3 where tuning models or prompts on large datasets is expensive. We also observe that representing a chain of thought instruction of mathematical reasoning questions as a program improves model performance significantly. We leverage our observation to build a large-scale mathematical reasoning model BHASKAR and a unified benchmark LILA. In case of program synthesis tasks, we observe that summarizing a question (instead of expanding as in chain of thought) helps models significantly. We study,,instruction-example equivalence, power of decomposition instruction to replace the need for new models and origination of dataset bias from crowdsourcing instructions to better understand the advantages and disadvantages of the instruction paradigm. We apply the instruction paradigm to match real user needs and introduce a new prompting technique HELP ME THINK to help humans perform various tasks by asking questions. Finally, we develop SELF-INSTRUCT to automatically generate instructions and improve instruction following capability of models automatically in an iterative fashion.

About the Speaker: Dr. Swaroop Mishra is a Research Scientist in Google DeepMind (Formerly Google Brain). He primarily works on instruction-tuning for Bard, he also collaborates with various teams, for prompting and AI self-improvement related consultancies.

His key research contributions, are instruction-tuning, various prompt-engineering methods e.g. reframing, question-decomposition, program representation to improve mathematical reasoning, stepback prompting and instruction-bias. His paper won the outstanding paper award at EACL 2023. He has published regularly in top AI conferences such as ACL, EMNLP, NeurIPS, ICML and AAAI. He received his M.Tech from IIT Kanpur in 2016 and PhD from Arizona State University in 2023. He has worked with Google Research, Microsoft Research and Allen Institute for AI.