

# Current Advances in LLM Reasoning

Akhil Arora, Vishrav Chaudhary, Julia Kreutzer, Nearchos Potamitis, Nouha Dziri, Niket Tandon

As large language models (LLMs) increasingly tackle reasoning-heavy tasks, from mathematics to commonsense to multilingual understanding, researchers face three pressing questions: How well do models reason? How can we make them reason better? What are the next frontiers in LLM reasoning? This tutorial answers these questions through a unified view of LLM reasoning. This tutorial explores comprehensive evaluation strategies to assess the reasoning abilities of models and discusses two types of methods to improve models' reasoning: advanced inference time methods, such as structured and self-improvement inference methods, and (ii) post-training methods, such as RLHF, DPO, and GRPO that aim to make LLMs think more like humans. The tutorial explores these technical discussions while maintaining a practical outlook through illustrative demos and short guided hands-on exercises. The tutorial is designed for both researchers and practitioners seeking practical insights into LLM reasoning.

---

## Akhil Arora

email: [akhil.arora@cs.au.dk](mailto:akhil.arora@cs.au.dk)

website: <https://au-clan.github.io/people/aarora/>

Akhil Arora is a professor of computer science at Aarhus University (AU), where he leads the CLAN for AI Research on Language and Networks (or “CLAN” for short). He is a fellow of the Copenhagen Center for Social Data Science, an affiliate of the Pioneer Centre for AI, and a formal collaborator of the Wikimedia Foundation. Akhil’s research lies broadly in human-centered AI with a focus on improving human knowledge-seeking, bridging knowledge gaps, and promoting knowledge equity on the Web.

## Vishrav Chaudhary

email: [vishrav@meta.com](mailto:vishrav@meta.com)

Vishrav Chaudhary is a Research Scientist working on post-training research and reasoning capabilities in large and small LMs. His research spans efficient multilingual modeling, reinforcement learning from human feedback, and cross-lingual reasoning. Vishrav has over a decade of experience leading major multilingual modeling efforts at Meta, Microsoft Turing, and FAIR, including contributions to the LLaMA-4, Phi3, and Phi4 model families.

## Julia Kreutzer

email: [juliakreutzer@cohere.com](mailto:juliakreutzer@cohere.com)

website: <https://juliakreutzer.github.io/>

Julia Kreutzer is a Senior Research Scientist focusing on multilingual reasoning and evaluation for LLMs. She has previously worked at Google Translate and obtained her PhD from the University of Heidelberg. Julia has led research at the intersection of multilingual evaluation, post-training, and inclusivity in language technologies.

## Nearchos Potamitis

email: [nearchos.potamitis@cs.au.dk](mailto:nearchos.potamitis@cs.au.dk)

website: <https://potamitisn.github.io/>

Nearchos Potamitis is a PhD student in Computer Science at Aarhus University (AU), advised by Akhil Arora. His research explores reasoning, efficiency, and evaluation in LMs, focusing on practical

systems for robust and affordable inference. Nearchos has prior experience developing educational materials and interactive demos for graduate NLP courses at AU and for international summer schools.

**Nouha Dziri**

email: [nouha.dziri@gmail.com](mailto:nouha.dziri@gmail.com)

website: <https://nouhadziri.github.io/>

Nouha Dziri is a Research Scientist at the Allen Institute for AI (AI2) working with Yejin Choi. Her research spans post-training, reasoning robustness, and LLM interpretability. Nouha has co-led the OLMo safety and post-training teams and previously worked at Google Research, Microsoft Research, and Mila.

**Niket Tandon**

email: [nikett@gmail.com](mailto:nikett@gmail.com)

website: <https://niket.tandon.info/>

Niket Tandon is a Principal Research Scientist focusing on feedback-guided and self-reflective reasoning in LLMs. Niket previously led reasoning research at AI2's Aristo team and earned his PhD from MPI for Informatics under Gerhard Weikum. His work has appeared in ACL, NeurIPS, and EMNLP, and includes the creation of large-scale commonsense resources like WebChild.