

The Second Workshop on Applied Machine Learning Management @ KDD 2023

Title

Incorporating Advances in LLMs into the Machine Learning Lifecycle

Round Table Proposal

New advances in large language models (LLMs) have the potential to be one of the most disruptive technological advancements in recent history. The pace at which new innovations are being created & shared in the machine learning community far outstrips the pace of other innovative advancements in technology. There is an urgent need amongst leaders in the machine learning community to start building best practices around the creation & use of large language models as there is no precedent for the pace of innovation currently being witnessed.

Having spoken to ML leaders across organizations of various sizes, in this roundtable, we will discuss some of the pressing topics on the minds of machine learning leaders & practitioners who are trying to incorporate advances in LLMs into the machine learning lifecycle. The goal of the round table would be to gather insights and share learnings from the successes and failures of recent experimentation.

1. How do you trade-off short term product commitments while trying to experiment with disruptive advancements in LLM technologies?

Discussion starting points:

- a. How much effort are you dedicating to incorporating LLMs into product use cases?
- b. Are you actively building new Proof of Concepts (PoCs)?
- c. How are you managing expectations with your leadership?

2. How are you thinking about the trade-off between fine tuning open source LLMs internally versus using APIs from companies like OpenAI?

Discussion starting points:

- a. Have you successfully fine tuned open source LLMs for product / research?
- b. Have you successfully used APIs from companies like OpenAI for product / research?
- c. What challenges did you face when trying to do one vs the other?

3. How are you dealing with licensing & legal concerns over using LLMs trained by various organizations?

Scenarios:

- a. When they have commercially usable licenses (Apache V2 / MIT)
- b. When they have research only licenses
- c. When they have commercially usable licenses with additional clauses around responsible AI (<https://bigscience.huggingface.co/blog/the-bigscience-rail-license>)
- d. When they have commercially usable licenses but you're unsure of the data lineage, do you train your own models

4. How are you handling the hallucination problem?

Discussion starting points:

- a. Using LLMs for internal-only use cases
- b. Incorporating LLMs into product use cases

5. How are you internally sharing best practices to avoid duplication of efforts?

Discussion starting points:

- a. For prompt creation
- b. For fine-tuning model development

Presenter Bio

Navneet Rao is a senior engineering manager at Thumbtack, a late-stage Sequoia backed consumer facing marketplace for local services. His org is responsible for matching customers on Thumbtack to professionals & local businesses in their neighborhood using data science & machine learning. They work on challenges in areas like search, ranking & marketplace efficiency. Recently, he's also been leading efforts around experimenting with LLMs for product use cases. Before Thumbtack, he was a tech lead for the conversational AI team at Watson Assistant, IBM's largest AI product. He focused on machine learning challenges like language understanding, intent & entity detection, transfer learning & machine bias detection. He also holds several patents in this space. He is an alumni of the Language Technologies Institute (LTI), at the School of Computer Science at Carnegie Mellon University.

Presenter Company

Thumbtack (www.thumbtack.com) is a local services marketplace where customers find and hire skilled professionals. Our app intelligently matches customers to electricians, landscapers, photographers and more with the right expertise, availability, and pricing. Headquartered in San Francisco, Thumbtack has raised more than 400 million USD from Baillie Gifford, Capital G, Javelin Venture Partners, Sequoia Capital, and Tiger Global Management among others. Available in all 3,143 U.S. counties, more than 4M customers have used Thumbtack in the last 12 months to find and hire professionals and local businesses.