

Achieving ML Reliability At Scale - Learnings and Opportunities

Machine Learning (ML) has revolutionized the overall tech ecosystem enabling complex scenarios with tremendous accuracy. The total [market size](#) for ML at the end of 2022 was estimated to be \$20 billion, and is expected to grow 10x (to ~\$225 billion) by 2030. However, at this scale, ML quality - especially reliability of ML artifacts - becomes extremely critical in order to continue delivering high-quality output, both for new product entrants as well as established players.

At Meta, where ML innovations help both the end users and customers get the most out of our offerings, we have built ML reliability for Ads from the ground up, in the process setting up sufficient guardrails throughout the whole ML lifecycle, i.e. data preparation, model development, model training, and evaluation. This has come a long way over the last couple of years when ML reliability was considered an afterthought to now being treated as a first-class citizen in Ads ML development, and continues to evolve with rapid, diverse, and more complex ML innovations.

Through this talk, we wish to share our learnings on the topic by providing real life examples and usecases around the following high-level themes:

1. Importance of ML reliability and its breakdown into different ML entities - for e.g., bad snapshot and staleness for models, coverage drop and value distribution shift for features
2. How to introduce ML reliability into a new product offering
3. Take reliability to next level and include it in all stages of ML lifecycle
4. Future direction of ML reliability and its opportunities/challenges

Presenters

Ritu Singh is a Technical Program Manager at Meta. She is responsible for leading machine learning reliability for the Meta Ads platform. Ritu's team has developed numerous reliability solutions for complex machine learning problems, many of which have been adopted by different ML product teams within Meta. Before joining Meta, Ritu has worked on various data and machine learning products at Microsoft.

Namit Gupta is a Technical Program Manager at Meta. He owns the reliability space for Ads ML predictions in order to help Meta's end users and advertisers derive maximum value out of the product offerings. Prior to Meta, he was a Product Manager in the Microsoft Office Identity division responsible for building identity experiences at scale for Microsoft 365 suite of apps.