

Wild Wild Tests

Keynote

Jacopo Tagliabue^{1,2}

¹*Bauplan Labs, New York, USA*

²*New York University, New York, USA*

Abstract

Recommender systems are typically evaluated through performance metrics computed over held-out data points. However, real-world behavior is more nuanced: we explore the shortcomings of a simplistic, one-metric-fits-all approach and fight the false testing dichotomy quantitative-and-scalable vs. qualitative-and-manual. We introduce RecList to address the engineering challenge of a rounded evaluation of recommender systems, and discuss the motivation and opportunities behind two recent data challenges, EvalRS 2022 and 2023.


Speaker biography

Educated in several acronyms across the globe (UNISR, SFI, MIT), Jacopo was co-founder and CTO of Tooso, a NLP startup in San Francisco acquired by TSX:CVO. He led Coveo's AI roadmap from scale-up to IPO, and built out Coveo Labs, an applied R&D practice rooted in collaboration (with Stanford, Netflix, Microsoft, NVIDIA) open source and open science. Co-creator of the RecList library for trustworthy evaluation of recommender systems, his work appeared in venues such as NAACL, RecSys, WWW, SIGIR. While building his new company, Bauplan, he is teaching ML Systems at NYU, which is mostly notable because it is the only job he ever had that his parents understand.

ORSUM@ACM RecSys 2023: 6th Workshop on Online Recommender Systems and User Modeling, jointly with the 17th ACM Conference on Recommender Systems, September 19th, 2023, Singapore



© 2022 Copyright for this paper by its authors. Use permitted under Creative Commons License Attribution 4.0 International (CC BY 4.0).

 CEUR Workshop Proceedings (CEUR-WS.org)