Impacts of Architectural Enhancements on Sequential Recommendation Models

Mufhumudzi Muthivhi[1], Terence L. van Zyl[2], and Hairong Wang[1]

[1] University of the Witwatersrand, Johannesburg GT, 2000, South Africa
1599695@students.wits.ac.za, hairongwng@gmail.com
[2] University of Johannesburg, Johannesburg GT, 2092, South Africa
tvanzyl@uj.ac.za

Abstract. Improving the architecture of sequential recommendation models has long been associated with enhanced accuracy metrics. However, new evaluation methods reveal that these improvements often favour a specific user cohort. The study aims to show that their popularity bias limits the performance of sequential recommendation models despite the numerous architectural enhancements adopted from NLP and Computer Vision. We propose a novel evaluation methodology to reflect users' preferences for popular and unpopular items accurately. We vary the threshold across the power law distribution to obtain two item subsets. This process sheds light on the extent of bias and performance discrepancies across the user spectrum. Our analysis of the experimental results reveals sequential recommendation models are limited to performing only as well as the epsilon-greedy algorithm. Consequently, the enhanced accuracy metrics such as Hit Rate and Normalized Discounted Cumulative Gain, frequently highlighted in the research, tend to stem primarily from the user group positioned at the short head of a power-law distribution.

Keywords: Sequential recommendation model · SASRec · Popularity bias.