

From M Passes to One


A MapReduce-Style Bootstrap for Pointwise Off-Policy Evaluation in Production Recommenders

Ghulam Ahmed Ansari · Cheng Ding

LinkedIn Corporation


01 / MAKE FULL-SLATE OPE USEFUL ON FEED

From dedicated random sessions to production logs



02 / THE ONE-PASS TILED CONSTRUCTION

Make uncertainty practical: 50 passes become one



REDUCE TO ESTIMATES

50 masked accumulators

Keep replicas separate; sum, then normalize

M estimates

↓

Sample variance

0 1 0 1 → session weight 2 (illustrative).

Masks and updates still scale with M.

Within 1% relative variance error vs. same-mask multi-pass reference.


No measured 50x speedup claim.

03 / TEST THE SLATE ASSUMPTION

Low pointwise error through K = 60

Pointwise assumes no preceding-item dependence.
Feed composite engagement; on-policy reference (Table 1).

Normalized mean squared error (log; lower is better)




Pointwise has lower error here; cause not isolated.
Unstable sequential weights are one possible reason.

04 / ALIGN SCORES ACROSS RANKER VERSIONS

Calibrate scores before computing propensities

Re-weighting can change score scales; raw softmax can confuse scale with preference.




Online-direction agreement: 1/2 → 2/2

Three ranker versions, two consecutive comparisons.
Calibration restored agreement after re-weighting.

05 / VALIDATE AGAINST ONLINE OUTCOMES

Does the complete system agree with online outcomes?

Second-pass ranker: same 11 model comparisons



Lines: Wilson 95% confidence intervals.

Combined offline-online agreement (two surfaces)

70% (16/23) 95% CI 49-84%

Suggested Post: 75% (9/12), separate model set.

Agreement: same online direction and significance.
Kendall rank correlation +0.29; p=0.04; broader 26 pairs.