


← Go to **ACM SIGIR 2026 Short Papers Track** homepage (/group?id=ACM.org/SIGIR/2026/Short_Papers_Track)

2 Versions ▾

ZoRRO: A Zero-Weight Personalized Recommender System for Scalable News Recommendation



*Johannes Kruse (/profile?id=~Johannes_Kruse1),
Ryotaro Shimizu (/profile?id=~Ryotaro_Shimizu1),
Kasper Lindskow (/profile?id=~Kasper_Lindskow1),
Jon Tofteskov (/profile?id=~Jon_Tofteskov1),
Michael Riis Andersen (/profile?id=~Michael_Riis_Andersen2),
Julian McAuley (/profile?id=~Julian_McAuley1),
Jes Frellsen (/profile?id=~Jes_Frellsen1)*

Published: 02 Apr 2026, Last Modified: 02 Apr 2026 SIGIR 2026 short papers
 Short Papers Track, Area Chairs, Reviewers, Authors Revisions (/revisions?id=mkMpySMUdw)
 BibTeX CC BY 4.0 (<https://creativecommons.org/licenses/by/4.0/>)

Confirm That All Author Profiles Are Complete: The authors confirm that they have read and understood this passage from the Call for Papers, and that all authors have completed their profile accordingly

Common Submission Requirements: The authors confirm that they have read and understood requirements in the Call for Papers, and that their submission complies with it

Keywords: recommender systems, news recommendation, scalability, online A/B testing

Abstract:

In this paper, we introduce ZoRRO (Zero-Weight Personalized Recommender System), a lightweight and computationally efficient framework for personalized news recommendations that models user interests with minimal overhead through a zero-weight, training-free design. We show that ZoRRO outperforms strong neural baselines in offline ranking evaluations and delivers click-through rate performance in online A/B testing that is nearly on par with a state-of-the-art deep learning model, while operating more than 600x faster. Our experiments reveal gaps between offline and online performance and metrics, and show that models with similar ranking or CTR outcomes can still generate markedly different recommendation distributions, influencing the overall news flow. Our code is available at <https://anonymous.4open.science/r/zorro-B556> (<https://anonymous.4open.science/r/zorro-B556>).

Submission Number: 41

Filter by reply type... ▾

Filter by author... ▾

Search keywords...

Sort: Newest First

-

=

≡

Everyone

Program Chairs

Submission41 Area...

Submission41 Reviewers

5 / 5 replies shown

Submission41 Authors

Submission41...

Submission41...

Submission41...

Submission41...

Add: **Withdrawal**

Paper Decision

Decision by Program Chairs 31 Mar 2026, 16:21 (modified: 02 Apr 2026, 07:16)

Program Chairs, Area Chairs, Reviewers, Authors Revisions (/revisions?id=4zdMViG4a)

Decision: Accept

Meta Review of Submission41 by Area Chair 8DN8

Meta Review by Area Chair 8DN8 20 Mar 2026, 14:39 (modified: 02 Apr 2026, 07:28)

Area Chairs, Authors, Reviewers Submitted, Program Chairs Revisions (/revisions?id=YnkDMj0LJD)

Metareview:

This is the meta-review by the SPC assigned to this paper.

The reviewers agree that this work is relevant, practical, and well evaluated for a short paper. Some issues were identified (categorical similarity term, more concise writing, strengthening the online-offline gap discussion), but they were deemed minor and hopefully addressable for the final version.

Recommendation: Accept

Confidence: 4: I am confident but not absolutely certain that the evaluation is correct

review

Official Review by Reviewer bb2w 20 Mar 2026, 14:36 (modified: 02 Apr 2026, 07:32)

Program Chairs, Area Chairs, Reviewers Submitted, Reviewer bb2w, Authors

Revisions (/revisions?id=YMPq9JBHda)

Relevance To Conference: Recsys and SIGIR

Relevance Score: 2 excellent

Novelty: -1 poor

Technical Soundness: 1 good

Quality Of Presentation: 1 good

Strengths:

- extremely efficient in time, memory and parameterization
- strong evaluation, including beyond accuracy metrics and online testing
- explicit addressing of the offline-online mismatch

Weaknesses:

- limited novelty
- over NRMS: small offline gain, slight online loss; this discussion is mostly qualitative
- performance probably hinges on the choice of embedding model and recency settings, so ablations would've been nice

Overall Recommendation: A solid short paper with simple yet effective idea. Perhaps there is a bit of overclaiming when saying that it outperforms neural methods.

Overall Recommendation Score: 1 weak accept

Nominate For Best Paper: 0 no

Reviewer Certification: 1 yes

Reviewer Confidence: I am fairly confident that my evaluation is correct, but I have significant doubts.

A solid short paper on news recommendation

Official Review by Reviewer dhHP 📅 14 Mar 2026, 17:15 (modified: 02 Apr 2026, 07:32)

👁 Program Chairs, Area Chairs, Reviewers Submitted, Reviewer dhHP, Authors

📄 Revisions (/revisions?id=w3c5INSbBp)

Relevance To Conference: Of great interest to a large group of audience

Relevance Score: 2 excellent

Novelty: 0 fair

Technical Soundness: 1 good

Quality Of Presentation: 1 good

Strengths:

1. Potential good impact. The conclusion of ZoRRO is that a well-designed zero-weight system can even outperform heavy DLRS-based methods when taking into account business attributes, item characteristics, and simple but comprehensive representations. Simple yet effective methods can potentially bring good impacts to the community
2. In-depth analysis on both accuracy-oriented and beyond-accuracy metrics.
3. Online A/B test in addition to effectiveness in offline evaluations.

Weaknesses:

1. The design choice of a cosine similarity between one-hot vectors needs a second thought.
2. There is an online-offline gap across the performance, which is missing an analysis or discussion based on quantitative results.

Overall Recommendation: This paper is focused on resolving the low efficiency of traditional DLRS-based news recommenders by providing a weight-free model, ZoRRO, and disentangling the relevance into intrinsic and relational. With no weights updated via complex ML models, ZoRRO could largely improve the efficiency while tallying or even outperforming other state-of-the-art methods. The paper is well written and with extensive analysis on both online and offline performance, demonstrated by both accuracy and beyond-accuracy superiority. However, there are also some (moderate) weaknesses to point out as follows:

1. The design choice of a cosine similarity between one-hot vectors needs a second thought. In Eq. (4), the cosine similarity between two one-hot encoded vectors is also modeled as part of the distance. However, the similarity is in fact only binary, with a value of 1 acquired when two vectors share the same category and 0 otherwise. This is a shallow and possibly suboptimal representation. I would suggest authors to consider some other methods, such as projecting categories as embeddings, or do an embedding discretization so that the numbers are no longer binary.
2. There is an online-offline gap across the performance, which is missing an analysis or discussion based on quantitative results. Though there are some speculations stated in Sec 4.2 on the possible explanations of such a gap, it is still missing more quantitative deep dives on whether the mentioned explanations truly stand. It would be of help if related discussions could go even deeper.
3. As a side comment, refs [31] and [32] seem to be the same but just with different bib entries.

Overall speaking, a good paper on an important topic in news recommendation, but there miss some in-depth discussions.

Overall Recommendation Score: 1 weak accept

Nominate For Best Paper: 0 no

Reviewer Certification: 1 yes

Reviewer Confidence: I am confident, but not absolutely certain that my evaluation is correct.

Review for ZoRRO

Official Review by Reviewer 4rKP 📅 13 Mar 2026, 02:25 (modified: 02 Apr 2026, 07:32)

👁 Program Chairs, Area Chairs, Reviewers Submitted, Reviewer 4rKP, Authors

 [Revisions \(/revisions?id=p9bz3cfpSo\)](/revisions?id=p9bz3cfpSo)

Relevance To Conference: The paper presents a personalized news recommender system, which is consistent with the topic of SIGIR.

Relevance Score: 2 excellent

Novelty: 1 good

Technical Soundness: 1 good

Quality Of Presentation: 0 fair

Strengths:

(1) The zero-weight design of ZoRRO significantly reduces the computational complexity and inference cost, while maintaining high recommendation accuracy. (2) Detailed experiments have demonstrated the advantages of ZoRRO in terms of inference speed and resource consumption. (3) In addition to traditional evaluations, the paper also conducted an analysis beyond accuracy, including diversity, serendipity, coverage, novelty, etc., providing a more comprehensive evaluation of the model performance.

Weaknesses:

(1) The figures and tables should be arranged in the order in which they appear in the paper. The arrangement of the figures and tables should be placed close to the description in the text. The arrangement of Table 1, Table 2, Figure 1 and Figure 2 is inappropriate. (2) There are writing errors in the text. For instance, in section 4.1.3, "With a fixed history length of 20 and an embedding dimension of 768, it achieves nearly 2,000 requests per second (RPS) with an average latency of 0.5 ms." (3) The paper points out the performance gap of ZoRRO in offline evaluations and online A/B testing, but fails to provide a reasonable explanation for this phenomenon. It is suggested to further analyze the impact of factors such as user behavior differences and changes in data distribution on the model performance.

Overall Recommendation: The proposed ZoRRO recommender system strikes a good balance between computational complexity and recommendation accuracy through zero-weight design. While there are some drawbacks, the overall work of the paper is of significant importance and practical application value. It is recommended that the authors further revise and improve the paper.

Overall Recommendation Score: 1 weak accept

Nominate For Best Paper: 0 no

Reviewer Certification: 1 yes

Reviewer Confidence: I am confident, but not absolutely certain that my evaluation is correct.

[About OpenReview \(/about\)](/about)

[Hosting a Venue \(/group?id=OpenReview.net/Support\)](/group?id=OpenReview.net/Support)

[All Venues \(/venues\)](/venues)

[Sponsors \(/sponsors\)](/sponsors)

[News \(/group?id=OpenReview.net/News&referrer=\[Homepage\] \(/\)\)](/group?id=OpenReview.net/News&referrer=[Homepage] (/))

[FAQ \(https://docs.openreview.net/getting-started/frequently-asked-questions\)](https://docs.openreview.net/getting-started/frequently-asked-questions)

[Contact \(/contact\)](/contact)

Donate (/donate)

[Terms of Use \(/legal/terms\)](/legal/terms)

[Privacy Policy \(/legal/privacy\)](/legal/privacy)

[OpenReview \(/about\)](/about) is a long-term project to advance science through improved peer review with legal nonprofit status. We gratefully acknowledge the support of the [OpenReview Sponsors \(/sponsors\)](/sponsors). © 2026 OpenReview