

AgentTailor

Dual-Gate LLM-Assisted Re-ranking for
Personalized Long-Tail Recommendation

Long-tail items need
personal preference support.


Case: An Awfully Big Adventure

Genre mismatch and a disliked director.
The item keeps its original low rank.

Jinpeng Chen* · Zhenye Yang · Huan Li
Kaimin Wei · Senzhang Wang · Hongbo Gao

Action and spy stories.
No Mike Newell, please.


A hidden gem?
Let's check the fit.



Illustrated case, paraphrased dialogue

Retrieve, assess, then place selectively

Two history-based permissions separate evidence assessment from ranking changes.



No activation or no placement permission: preserve the Stage-1 order.

History controls permission. Evidence guides ordering.

$$e_{ui} = \text{rel}_{ui} \cdot \text{tailfit}_{ui}$$

Evidence = relevance $\times$ tail fit.

Placement adds a weak Stage-1 rank prior.

Tail pressure λ is smoothed historical tail engagement.

$$\lambda_u < \delta_a$$

$$\delta_a \leq \lambda_u < \delta_p$$

$$\lambda_u \geq \delta_p$$

Keep the order

No evidence assessment is activated.

Assess evidence

Keep the Stage-1 order.

Permit placement

Use evidence scores and the rank prior for stable placement.

Fixed thresholds satisfy activation $\leq$ placement. Permission can still yield no ranking change.

Stable placement retains the relative order of non-routed candidates.

Stronger overall and long-tail recommendation

Comparison with LLM4Rerank using the same Stage-1 candidate pools.

MovieLens-1M

+17.4%

Tail HR@10 relative gain

	LLM4Rerank	AgentTailor
Overall	0.091	0.099
Tail	0.069	0.081

Overall gain: +8.8%

BookCrossing

+13.3%

Tail HR@100 relative gain



	LLM4Rerank	AgentTailor
Overall	0.097	0.109
Tail	0.090	0.102

Overall gain: +12.4%

Five runs. Gains use reported rounded scores. Tail HR measures hits on held-out tail targets.

Evidence synthesis and stable placement both matter

Stage-2 ablations with the candidate pools held fixed.



Removing SPA raises head HR, but lowers overall and tail HR.

Make long-tail placement personal

Retrieve relevant items.

Rank with personal evidence.

History sets the permissions.


Evidence and the rank prior guide ordering.

Preserve the relative order of non-routed items.

Paper



Code



github.com/Alex-bupt/AgentTailor

* Corresponding author

jpchen@bupt.edu.cn

Higher overall and tail HR on both evaluated datasets. Preserved relative order applies to non-routed candidates.