TESA-REC: EVIDENCE–AUTHORITY SEPARATION FOR SEMANTIC SEQUENTIAL RECOMMENDATION

Anonymous authors
Paper under double-blind review

ABSTRACT

Semantic recommenders commonly allow an auxiliary source to modify the final score whenever that source provides useful information. This conflates two distinct objects: evidence quality and decision authority. We introduce an evidence–authority separation principle and instantiate it as Temporal Entropic Semantic Adaptation for Recommendation (TESA-REC), a zero-trainable-parameter interface around a frozen sequential anchor. The anchor defines a behaviorally admissible set and the finite rank capacity available for intervention. Within that set, temporally partitioned candidate evidence performs a minimum-information update of a recency prior, producing a candidate-conditioned Gibbs belief over temporal explanations. One explanation overtakes another only after crossing a prior-odds evidence barrier, while reversing an anchor preference requires a separate behavioral-gap threshold. A constrained assignment allocates the anchor's existing rank slots according to the resulting priorities. We show universal optimality for nonincreasing position values, maximal invariance to monotone score calibration, exact rank-capacity conservation, and a deterministic global perturbation bound. Across five full-catalog benchmarks, every matched-anchor mean improves and TESA-REC wins all 20 protocol-matched comparisons with the listed baselines; 31 of 35 default metric gains remain significant after correction, and the direction transfers across three language encoders. Matched controls further show that constrained allocation consistently improves upon additive score fusion, while empirical rankings satisfy the predicted invariance and conservation laws exactly.

1 INTRODUCTION

Modern recommenders combine behavioral sequences with language, visual content, knowledge, and contextual signals. Sequential encoders learn behavioral regularities from interaction order (Kang & McAuley, 2018). Language models add item meaning unavailable to identifiers alone (Li et al., 2023a). Existing integration commonly inserts auxiliary evidence into a hidden representation, adds it to a numerical recommendation score, or learns a new reranker. These routes share an implicit assumption: a source that provides useful evidence may also obtain direct authority over the final decision. The two are not equivalent. A semantic model may confidently recognize that two products are related without being entitled to promote a behaviorally implausible item from the catalog tail to the first position.

Semantic similarity and behavioral scores measure different quantities, so direct addition makes intervention depend on numerical calibration. An order-based interface instead asks which behaviorally supported candidate deserves an existing high-value position and what evidence justifies that reassignment. This separates the usefulness of auxiliary information from the scope of its decision authority.

We therefore separate *evidence* from *authority*. A sequential anchor does more than emit scores: its order identifies candidates with sufficient behavioral support to be reconsidered and supplies a finite collection of rank positions. Semantics should explain preference inside this permitted intervention domain rather than reconstruct catalog-wide admissibility. This principle creates three coupled questions. **C1: who is admissible?** The behavioral model must delimit where auxiliary evidence may act. **C2: why is a candidate compatible?** A changing history contains competing immediate,

Reviewers: please read the Reviewer Guidelines (iclr.cc/Conferences/2027/ReviewerGuidelines) and the AI Policy for Reviewers (iclr.cc/Conferences/2027/AIPolicyForReviewers). If you used AI to expand, edit, or polish your review, please provide the input text to the LLM. Better still, consider skipping the LLM and submitting your original text: we, and the authors, are much more interested in your unedited thoughts than in what an LLM has to say. AI-assisted or not, you are putting your name and reputation behind your review: LLM-generated falsehoods, hallucinations or misrepresentations are subject to disciplinary action, which may include desk-rejecting all papers you have authored.

short-term, medium-term, and persistent explanations, and different candidates can favor different explanations. **C3: how much authority may the evidence exercise?** The final decision must allocate existing behavioral capacity without creating new positions or disturbing the ranking outside the permitted domain.

TESA-REC answers these questions as one constrained evidence-to-decision process around a frozen MPT anchor (Xu et al., 2026a). The anchor first defines a behaviorally admissible Top-R set and its fixed rank capacity. Within this set, each temporal block supplies candidate-specific semantic evidence with optional collaborative support. KL-regularized exponential tilting computes the minimum-information departure from a recency prior justified by each candidate. This creates three levels of control: a prior-odds barrier governs when one temporal explanation overtakes another; a behavioral-gap threshold governs whether candidate-conditioned evidence can reverse an anchor preference; and every reversal remains confined to capacity fixed by the anchor. Thus auxiliary authority is earned locally but conserved globally. Experiments follow the same logic: TESA-REC improves every matched-anchor mean across five domains, both evidence specialization and constrained allocation contribute beyond simpler alternatives, three language encoders preserve the direction, and direct audits recover ordinal invariance and capacity conservation on every evaluated ranking.

Our contributions are:

- We introduce evidence–authority separation for semantic sequential recommendation: auxiliary evidence may determine priority among behaviorally admissible candidates but may not create additional ranking capacity.
- We derive candidate-conditioned temporal reasoning as a minimum-information update of a recency prior, with an explicit threshold for one explanation to dominate another.
- We characterize when candidate-conditioned evidence earns enough authority to reverse an anchor preference through a behavioral-gap threshold.
- We cast the final decision as universally optimal allocation of conserved rank capacity and establish ordinal sufficiency, metric locality, and a global intervention bound; matched controls, structural audits, and three semantic encoders validate the complete pathway.

2 Related Work

Sequential recommendation has progressed from Markov factorization (Rendle et al., 2010) and recurrent state updates (Hidasi et al., 2016) to convolutional sequence encoders (Tang & Wang, 2018) and self-attention (Kang & McAuley, 2018). Bidirectional prediction enriches item context (Sun et al., 2019), while contrastive objectives regularize sequential states (Xie et al., 2022). Frequency-domain models expose periodic structure (Zhou et al., 2022), and wavelet architectures separate behavior at multiple resolutions (Xu et al., 2026b). MPT instead learns a transferable behavioral prior from synthetic Markov chains (Xu et al., 2026a). We treat that prior as the mechanism defining behavioral admissibility and ask how a second source should receive authority over its induced order.

Language-aware methods differ in how semantic evidence enters the decision. Representation fusion aligns text with collaborative states before prediction (Hou et al., 2022). Language-centric pretraining places item text directly in the sequential encoder (Li et al., 2023a). Parameter-efficient adaptation modifies a language model toward recommendation objectives (Li et al., 2023b), while frozen embeddings can serve as auxiliary sequence features (Harte et al., 2023). Instruction-based systems connect recommendation states to language reasoning (Liao et al., 2024), and learned rerankers construct a new decision rule over retrieved candidates (Gao et al., 2025). In these families, semantic evidence obtains authority through the fused representation, numerical score, or learned reranker. TESA-REC instead studies a constrained-authority interface. Candidate-conditioned evidence determines relative priority only inside a behaviorally admissible domain, while every final position is inherited from the anchor's fixed rank capacity.

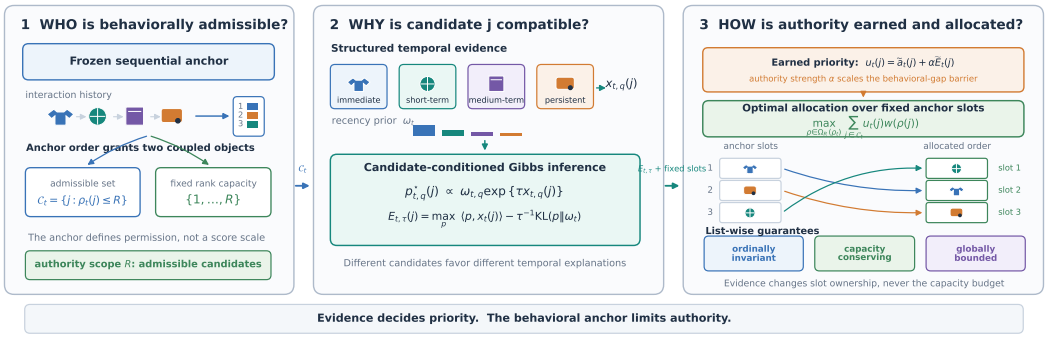


Figure 1: Candidate-conditioned evidence earns priority relative to the ordinal anchor, while admissibility and rank capacity remain fixed.

3 METHOD

3.1 EVIDENCE–AUTHORITY SEPARATION

For the observed prefix $i_{1:t} = (i_1, \ldots, i_t)$, let $\mathcal{I}$ denote the eligible catalog after masking its observed history and let $M = |\mathcal{I}|$. A frozen sequential anchor assigns $b_t : \mathcal{I} \to \mathbb{R}$, where larger scores are preferred. Its score-induced weak order is $i \preceq_t j$ iff $b_t(i) \le b_t(j)$; a fixed tie rule refines $\preceq_t$ into the one-indexed total rank $\rho_t : \mathcal{I} \to \{1, \ldots, M\}$. For an authority budget $1 \le R \le M$, the behaviorally admissible set is $\mathcal{C}_t = \{j \in \mathcal{I} : \rho_t(j) \le R\}$. Letting ρ range over permutations of $\mathcal{I}$, the feasible final rankings are

$$\Omega_R(\rho_t) = \{\rho : \ \rho(j) = \rho_t(j)\ \forall j \notin \mathcal{C}_t, \quad \{\rho(j) : j \in \mathcal{C}_t\} = \{1, \ldots, R\}\}. \tag{1}$$

Thus R is authority scope: the anchor fixes both *who* may be reconsidered and the available slots $\{1, \ldots, R\}$. Each item also has a frozen, normalized language vector $v_j \in \mathbb{R}^{d_\ell}$ and anchor vector $e_j \in \mathbb{R}^{d_c}$; these geometries explain *why* a candidate fits without determining admissibility. TESA-REC then decides *how* to assign the fixed capacity (Figure 1).

3.2 BEHAVIORAL ADMISSIBILITY AND TEMPORAL EVIDENCE

A single history average treats all past events as mutually compatible evidence. We instead partition the interaction positions $\{1, \ldots, t\}$ into B disjoint blocks $\mathcal{B}_{t,q}$, so $\mathcal{B}_{t,q} \cap \mathcal{B}_{t,r} = \varnothing$ for $q \neq r$ and $\bigcup_{q=1}^{B} \mathcal{B}_{t,q} = \{1, \ldots, t\}$. Here ξ indexes a position and i_ξ is the item observed there. Each block is a competing temporal explanation; the implementation uses immediate, short-term, medium-term, and persistent history. Let $\mathcal{Q}_t = \{q \in \{1, \ldots, B\} : \mathcal{B}_{t,q} \neq \varnothing\}$ be the active hypotheses. For a fixed $\varepsilon_{\text{norm}} > 0$, their semantic and collaborative centroids are

$$s_{t,q} = \frac{\sum_{\xi \in \mathcal{B}_{t,q}} v_{i_\xi}}{\max\left\{\left\|\sum_{\xi \in \mathcal{B}_{t,q}} v_{i_\xi}\right\|_2, \varepsilon_{\text{norm}}\right\}}, \qquad c_{t,q} = \frac{\sum_{\xi \in \mathcal{B}_{t,q}} e_{i_\xi}}{\max\left\{\left\|\sum_{\xi \in \mathcal{B}_{t,q}} e_{i_\xi}\right\|_2, \varepsilon_{\text{norm}}\right\}}, \quad q \in \mathcal{Q}_t. \tag{2}$$

Let $\bar{\omega} = (\bar{\omega}_1, \ldots, \bar{\omega}_B) \in \mathbb{R}^B_{>0}$ be a base recency prior. Active weights are $\omega_{t,q} = \bar{\omega}_q / \sum_{r \in \mathcal{Q}_t} \bar{\omega}_r$ for $q \in \mathcal{Q}_t$, which explains the dependence on prefix length. For candidate $j \in \mathcal{C}_t$, block q supplies semantic criterion $x^\ell_{t,q}(j) = \frac{1}{2}(1 + \langle s_{t,q}, v_j \rangle)$ and collaborative criterion $x^c_{t,q}(j) = \frac{1}{2}(1 + \langle c_{t,q}, e_j \rangle)$. Their mixture is

$$x_{t,q}(j) = \lambda x^\ell_{t,q}(j) + (1 - \lambda) x^c_{t,q}(j), \qquad \lambda \in [0, 1]. \tag{3}$$

Here $\lambda = 1$ and 0 recover semantic-only and collaborative-only evidence, respectively. The bounded quantity $x_{t,q}(j) \in [0, 1]$ is candidate support for explanation q, not a catalog score; evaluating it after $\mathcal{C}_t$ is fixed separates evidence from admissibility.

3.3 CANDIDATE-CONDITIONED MINIMUM-INFORMATION INFERENCE

Before observing a candidate, the normalized weights $\omega_t = (\omega_{t,q})_{q \in \mathcal{Q}_t}$ encode a recency prior over nonempty temporal hypotheses. Candidate j then supplies evidence $x_t(j) = (x_{t,q}(j))_{q \in \mathcal{Q}_t}$. For

$\tau > 0$, TESA-REC evaluates their compatibility by

$$E_{t,\tau}(j) = \frac{1}{\tau} \log \sum_{q \in \mathcal{Q}_t} \omega_{t,q} \exp\{\tau x_{t,q}(j)\}\,. \tag{4}$$

At $\tau = 0$, the continuous extension is the prior mean $E_{t,0}(j) = \sum_q \omega_{t,q} x_{t,q}(j)$. The parameter τ controls prior resistance: departing from recency requires stronger evidence when τ is small, whereas larger τ permits greater candidate-specific specialization.

Theorem 1 (Candidate-conditioned minimum-information update)**.** *For $\tau > 0$, let $x_t(j) = (x_{t,q}(j))_{q \in \mathcal{Q}_t}$ and $\Delta(\mathcal{Q}_t) = \{p \in \mathbb{R}^{|\mathcal{Q}_t|}_{\geq 0} : \sum_{q \in \mathcal{Q}_t} p_q = 1\}$. Then*

$$E_{t,\tau}(j) = \max_{p \in \Delta(\mathcal{Q}_t)} \left\{ \langle p, x_t(j) \rangle - \frac{1}{\tau} \mathrm{KL}(p \| \omega_t) \right\}, \quad p^\star_{t,q}(j) = \frac{\omega_{t,q} e^{\tau x_{t,q}(j)}}{\sum_r \omega_{t,r} e^{\tau x_{t,r}(j)}}. \tag{5}$$

Writing $m_{t,\tau}(j) = \langle p^\star_t(j), x_t(j) \rangle$, the same optimizer is equivalently the minimum-information departure achieving at least that evidence level:

$$p^\star_t(j) = \arg \min_{p \in \Delta(\mathcal{Q}_t):\, \langle p, x_t(j) \rangle \geq m_{t,\tau}(j)} \mathrm{KL}(p \| \omega_t).$$

Moreover, $E_{t,\tau}$ is coordinate-wise monotone, translation equivariant, and 1-Lipschitz under $\| \cdot \|_\infty$. It satisfies $\langle \omega_t, x_t(j) \rangle \leq E_{t,\tau}(j) \leq \max_q x_{t,q}(j)$, converges to the maximum as $\tau \to \infty$, and obeys $\partial E_{t,\tau}(j)/\partial \tau = \mathrm{KL}(p^\star_t(j) \| \omega_t)/\tau^2 \geq 0$.

The optimizer $p^\star_t(j)$ is a candidate-conditioned Gibbs belief, not learned attention. The sensitivity of $E_{t,\tau}(j)$ to τ is $\mathrm{KL}(p^\star_t(j) \| \omega_t)/\tau^2$, so the value becomes more sensitive precisely when the candidate-induced belief departs more strongly from the recency prior. The 1-Lipschitz property further bounds its response to perturbations in temporal evidence.

Corollary 1 (Temporal-explanation dominance threshold)**.** *For any active hypotheses q and r,*

$$p^\star_{t,q}(j) > p^\star_{t,r}(j) \quad \Longleftrightarrow \quad x_{t,q}(j) - x_{t,r}(j) > \tau^{-1} \log \frac{\omega_{t,r}}{\omega_{t,q}}.$$

Hence explanation q overtakes r only when its candidate-specific evidence advantage exceeds the corresponding prior-odds penalty.

Appendix A proves the theorem and corollary.

3.4 FROM EVIDENCE TO EARNED PRIORITY

The entropic value supplies evidence, but positive evidence alone does not authorize a behavioral reversal. For any strictly decreasing ordinal potential $\phi : \{1, \ldots, R\} \to \mathbb{R}$, set $a_t(j) = \phi(\rho_t(j))$; we use $\phi_\gamma(r) = -[\log r]^\gamma$, where $\gamma > 0$ controls rank-gap curvature. For any $y : \mathcal{C}_t \to \mathbb{R}$, its candidate-set mean and stabilized population standard deviation are

$$\mu_y = \frac{1}{|\mathcal{C}_t|} \sum_{j \in \mathcal{C}_t} y(j), \qquad s_y = \left[\max \left\{ \frac{1}{|\mathcal{C}_t|} \sum_{j \in \mathcal{C}_t} (y(j) - \mu_y)^2, \varepsilon_s \right\} \right]^{1/2},$$
$$\mathrm{Std}_{\mathcal{C}_t}(y)(j) = \frac{y(j) - \mu_y}{s_y}. \tag{6}$$

where $\varepsilon_s > 0$ is fixed. Let $\widetilde{a}_t = \mathrm{Std}_{\mathcal{C}_t}(a_t)$ and $\widetilde{E}_t = \mathrm{Std}_{\mathcal{C}_t}(E_{t,\tau})$. The single assignment priority used thereafter is

$$u_t(j) = \widetilde{a}_t(j) + \alpha \widetilde{E}_t(j), \qquad \alpha \geq 0, \tag{7}$$

Here α is authority strength inside the scope set by R; u_t never defines unrestricted catalog positions.

Proposition 1 (Behavioral-reversal threshold)**.** *Let $i, j \in \mathcal{C}_t$ satisfy $\rho_t(i) < \rho_t(j)$. For $\alpha > 0$, candidate j receives higher assignment priority than i if and only if*

$$\widetilde{E}_t(j) - \widetilde{E}_t(i) > \frac{\widetilde{a}_t(i) - \widetilde{a}_t(j)}{\alpha}.$$

Thus α inversely scales the evidence advantage required to cross a standardized behavioral gap. For $\alpha = 0$, $u_t = \widetilde{a}_t$ and the anchor order is retained.

3.5 Optimal Allocation under Conserved Rank Capacity

Let $w : \{1, \dots, R\} \to \mathbb{R}_{\geq 0}$ be any position-value function satisfying $w(1) \geq \cdots \geq w(R)$, and define the corresponding rank-capacity budget $B_w(R) = \sum_{k=1}^{R} w(k)$. The final decision allocates this fixed budget by

$$\rho_t^f \in \arg\max_{\rho \in \Omega_R(\rho_t)} \sum_{j \in \mathcal{C}_t} u_t(j) w(\rho(j)). \tag{8}$$

Theorem 2 (Universal optimality under conserved rank capacity). *Let $\pi_t(1), \dots, \pi_t(R)$ order $\mathcal{C}_t$ by decreasing u_t, using anchor rank to break ties. Then the closed-form assignment*

$$\rho_t^f(\pi_t(k)) = k \quad (1 \leq k \leq R), \qquad \rho_t^f(j) = \rho_t(j) \quad (j \notin \mathcal{C}_t) \tag{9}$$

simultaneously solves Equation 8 for every nonincreasing w. If priorities are distinct and w is strictly decreasing, the solution is unique.

The same descending-priority assignment is therefore optimal for every valuation in which higher positions carry no less value; it is not tied to NDCG, MRR, or a particular exposure curve.

Corollary 2 (Exact rank-capacity conservation). *Equation 9 is a permutation of $\{1, \dots, M\}$, conserves the candidate slot multiset $\{\rho_t^f(j) : j \in \mathcal{C}_t\} = \{1, \dots, R\}$, and leaves every outside rank unchanged. Consequently, the Top-K item set is invariant for every $K \geq R$, and for any position-value function w,*

$$\sum_{j \in \mathcal{C}_t} w(\rho_t^f(j)) = \sum_{j \in \mathcal{C}_t} w(\rho_t(j)) = B_w(R). \tag{10}$$

This conservation is enforced exactly by $\Omega_R(\rho_t)$ rather than encouraged through a soft penalty: candidate-conditioned evidence changes who receives rank capacity, never how much exists.

3.6 Ordinal and Intervention Guarantees

Proposition 2 (Monotone calibration invariance). *Replacing b_t by $g \circ b_t$ for any strictly increasing $g : \mathbb{R} \to \mathbb{R}$ leaves $\preceq_t$ and, under the fixed tie rule, ρ_t, $\mathcal{C}_t$, a_t, u_t, and ρ_t^f unchanged.*

Theorem 3 (Maximal ordinal sufficiency). *Let $\mathcal{G}$ be the group of strictly increasing bijections of $\mathbb{R}$. Two finite anchor-score vectors lie in the same $\mathcal{G}$-orbit exactly when they induce the same weak order $\preceq_t$, including ties. Hence, for fixed non-anchor inputs, $\preceq_t$ is a maximal invariant, and every deterministic adapter invariant to $\mathcal{G}$ factors through it. The algorithm uses its fixed total-rank refinement ρ_t.*

Corollary 3 (Bounded and local intervention). *Let d_{KT} count pairwise inversions between the anchor and final orders. Then*

$$d_{\mathrm{KT}}(\rho_t, \rho_t^f) \leq \binom{R}{2}, \qquad \frac{d_{\mathrm{KT}}(\rho_t, \rho_t^f)}{\binom{M}{2}} \leq \frac{R(R-1)}{M(M-1)}.$$

For a target item y_t with $\rho_t(y_t) > R$, its rank and its HR, NDCG, and reciprocal-rank contributions remain unchanged at every cutoff. All reported ranking-metric changes are therefore localized to behaviorally admissible targets.

3.7 End-to-End Decision Closure

The preceding components form one closed decision map rather than independent scoring modules:

$$\begin{aligned} \rho_t &\longrightarrow (\mathcal{C}_t, a_t, \Omega_R(\rho_t)), \\ (\{x_t(j)\}_{j \in \mathcal{C}_t}, \omega_t) &\longrightarrow \{p_t^\star(j), E_{t,\tau}(j)\}_{j \in \mathcal{C}_t}, \\ (a_t, E_{t,\tau}, \alpha) &\longrightarrow u_t, \\ (u_t, \Omega_R(\rho_t)) &\longrightarrow \rho_t^f. \end{aligned} \tag{11}$$

The four maps respectively fix admissibility, ordinal potential, and feasible slots; infer candidate-conditioned evidence; construct the single assignment priority; and realize the closed-form feasible allocation. The final map selects no w: by Theorem 2, the same descending-priority assignment

simultaneously optimizes every nonincreasing position value. Corollary 1 governs when one explanation overtakes another, whereas Proposition 1 governs when a lower-ranked candidate defeats a stronger behavioral preference. Neither can move a candidate across $\mathcal{C}_t$ or produce a ranking outside $\Omega_R(\rho_t)$.

The closure separates explanation updates, order changes, and feasible actions: evidence may change the first without the second, and may change order only across the behavioral gap. Ordinal sufficiency removes nuisance calibration, while Ω_R makes every intervention local and bounded. With frozen vectors and zero trainable parameters, centroid construction, candidate evidence, and assignment require $O((t + RB)(d_\ell + d_c) + R \log R)$ work after anchor ranking; Appendix A gives complete proofs.

4 Experiments

We evaluate five theory-aligned questions. **Q1: Quality.** Does capacity-constrained evidence adaptation improve recommendation under matched protocols and against its exact anchor? **Q2: Inference.** Are gains attributable to candidate-conditioned temporal evidence? **Q3: Authority.** Does fixed-capacity allocation outperform granting the same evidence through score fusion? **Q4: Guarantees.** Are ordinal invariance and decision-capacity conservation realized exactly? **Q5: Transfer.** Does the direction persist across frozen semantic representations?

4.1 Experimental Setup

Evaluation and implementation. We use Beauty, Toys, Sports, Yelp, and OnlineRetail with chronological leave-one-out splits. Each next-item target is ranked against the full eligible catalog after masking the complete observed history; we report HR and NDCG at 5, 10, and 20 together with MRR. Results average three independently trained MPT anchor checkpoints (Xu et al., 2026a); conditional on a frozen anchor and cached item representations, TESA-Rec is deterministic. Qwen3-Embedding is the default language representation (Zhang et al., 2025); BGE-M3 (Chen et al., 2024) and E5-large-v2 (Wang et al., 2022) test transfer. The matched Score Fusion control replaces the behavioral priority potential in Equation 7 with the candidate-standardized anchor score while retaining the same evidence and coefficient. Appendices C–F record profiles, the user-level testing procedure and its 35-test Holm family, complete results, and measured efficiency.

Baselines. GRU4Rec (Hidasi et al., 2016), SASRec (Kang & McAuley, 2018), FMLP-Rec (Zhou et al., 2022), HSTU (Zhai et al., 2024), SASRec$_+$ (Xu et al., 2026a), UniSRec (Hou et al., 2022), RecFormer (Li et al., 2023a), E4SRec (Li et al., 2023b), Qwen2.5 (Yang et al., 2024), MPT-Adapter (Xu et al., 2026a), and MPT-LoRA (Xu et al., 2026a).

4.2 Predictive Quality under Matched Protocols

Table 1: Full-catalog comparison under the same Amazon evaluation setting.

	Beauty				Toys				Sports			
Model	H@10	N@10	H@20	N@20	H@10	N@10	H@20	N@20	H@10	N@10	H@20	N@20
GRU4Rec	0.0524	0.0278	0.0788	0.0344	0.0490	0.0282	0.0704	0.0335	0.0259	0.0135	0.0407	0.0172
SASRec	0.0753	0.0417	0.1060	0.0495	0.0846	0.0493	0.1115	0.0561	0.0405	0.0212	0.0607	0.0263
FMLP-Rec	0.0775	0.0454	0.1072	0.0529	0.0855	0.0521	0.1124	0.0589	0.0413	0.0238	0.0593	0.0283
HSTU	0.0902	0.0535	0.1246	0.0621	0.0981	0.0610	0.1288	0.0687	0.0550	0.0317	0.0788	0.0377
SASRec$_+$	0.0982	0.0591	0.1348	0.0683	0.1040	0.0645	0.1393	0.0733	0.0576	0.0337	0.0824	0.0399
UniSRec	0.0803	0.0382	0.1233	0.0490	0.0829	0.0380	0.1235	0.0483	0.0446	0.0221	0.0699	0.0285
RecFormer	0.0676	0.0333	0.1017	0.0419	0.0842	0.0405	0.1221	0.0501	0.0345	0.0167	0.0555	0.0220
E4SRec	0.0773	0.0469	0.1046	0.0538	0.0798	0.0492	0.1061	0.0558	0.0387	0.0218	0.0563	0.0262
Qwen2.5	0.0826	0.0464	0.1200	0.0558	0.0784	0.0433	0.1114	0.0515	0.0473	0.0256	0.0712	0.0316
MPT-Adapter	0.1118	0.0673	0.1531	0.0777	0.1202	0.0705	0.1617	0.0809	0.0651	0.0373	0.0924	0.0442
MPT-LoRA	0.1100	0.0658	0.1514	0.0763	0.1192	0.0701	0.1611	0.0806	0.0644	0.0365	0.0927	0.0436
TESA-Rec	**0.1184**	**0.0703**	**0.1625**	**0.0815**	**0.1291**	**0.0753**	**0.1728**	**0.0863**	**0.0687**	**0.0383**	**0.0985**	**0.0458**
Improve	**+5.90%**	**+4.46%**	**+6.14%**	**+4.89%**	**+7.40%**	**+6.81%**	**+6.86%**	**+6.67%**	**+5.53%**	**+2.68%**	**+6.26%**	**+3.62%**

Table 2: Full-catalog comparison under the same Yelp and OnlineRetail setting.

	Yelp				OnlineRetail			
Model	H@10	N@10	H@20	N@20	H@10	N@10	H@20	N@20
GRU4Rec	0.0246	0.0116	0.0448	0.0166	0.1175	0.0561	0.1931	0.0751
SASRec	0.0298	0.0123	0.0539	0.0184	0.1245	0.0524	0.2014	0.0718
FMLP-Rec	0.0334	0.0138	0.0606	0.0207	0.1333	0.0548	0.2237	0.0775
HSTU	0.0453	0.0215	0.0773	0.0295	0.1452	0.0656	0.2304	0.0870
SASRec_+	0.0466	0.0209	0.0796	0.0292	0.1589	0.0719	0.2575	0.0968
UniSRec	0.0343	0.0170	0.0567	0.0226	0.1424	0.0662	0.2275	0.0877
RecFormer	0.0247	0.0121	0.0429	0.0166	0.1407	0.0683	0.2194	0.0881
E4SRec	0.0291	0.0139	0.0513	0.0195	0.1551	0.0749	0.2528	0.0995
Qwen2.5	0.0411	0.0204	0.0682	0.0272	0.1503	0.0731	0.2356	0.0945
MPT-Adapter	0.0461	0.0224	0.0783	0.0305	0.1608	0.0741	0.2572	0.0982
MPT-LoRA	0.0495	0.0241	0.0832	0.0325	0.1568	0.0746	0.2502	0.0980
TESA-REC	**0.0576**	**0.0301**	**0.0907**	**0.0384**	**0.2411**	**0.1404**	**0.3267**	**0.1620**
Improve	**+16.36%**	**+24.90%**	**+9.01%**	**+18.15%**	**+49.94%**	**+87.45%**	**+26.87%**	**+62.81%**

Tables 1 and 2 compare TESA-REC with the listed sequential, language-aware, and adapted MPT baselines under the same five-domain setting. TESA-REC leads all 20 comparisons. Beauty HR@10, NDCG@10, HR@20, and NDCG@20 improve by 5.90%, 4.46%, 6.14%, and 4.89%, respectively; Toys and Sports gains range from 2.68% to 7.40%, and all four metrics also improve on Yelp and OnlineRetail. Improve denotes the relative increase over the strongest listed baseline in each column. Gains across membership and discounted-position metrics support assigning high-rank slots through candidate-conditioned evidence within the anchor-defined region. The following matched-anchor evaluation fixes scores, masks, candidates, and representations to isolate evidence specialization and fixed-capacity assignment.

4.3 MATCHED-ANCHOR GAINS ACROSS FIVE DOMAINS

Table 3: Complete matched-anchor results under strict full-catalog evaluation.

Dataset	Variant	HR@5	NDCG@5	HR@10	NDCG@10	HR@20	NDCG@20	MRR
Beauty	Frozen MPT	0.0806	0.0569	0.1134	0.0674	0.1555	0.0780	0.0605
	TESA-REC	**0.0835**	**0.0591**	**0.1184**	**0.0703**	**0.1625**	**0.0815**	**0.0630**
	Improve	**+3.60%**	**+3.96%**	**+4.39%**	**+4.33%**	**+4.53%**	**+4.42%**	**+4.08%**
Toys	Frozen MPT	0.0877	0.0609	0.1236	0.0725	0.1654	0.0830	0.0639
	TESA-REC	**0.0915**	**0.0632**	**0.1291**	**0.0753**	**0.1728**	**0.0863**	**0.0659**
	Improve	**+4.35%**	**+3.70%**	**+4.46%**	**+3.89%**	**+4.52%**	**+3.99%**	**+3.21%**
Sports	Frozen MPT	0.0449	0.0304	0.0674	0.0377	0.0957	0.0448	0.0342
	TESA-REC	**0.0459**	**0.0310**	**0.0687**	**0.0383**	**0.0985**	**0.0458**	**0.0347**
	Improve	**+2.38%**	**+1.81%**	**+1.95%**	**+1.65%**	**+2.94%**	**+2.26%**	**+1.47%**
Yelp	Frozen MPT	0.0355	0.0230	0.0575	0.0300	0.0902	0.0382	0.0287
	TESA-REC	**0.0356**	**0.0230**	**0.0576**	**0.0301**	**0.0907**	**0.0384**	**0.0288**
	Improve	**+0.35%**	**+0.23%**	**+0.29%**	**+0.26%**	**+0.62%**	**+0.46%**	**+0.25%**
OnlineRetail	Frozen MPT	0.1686	0.1167	0.2385	0.1392	0.3234	0.1606	0.1221
	TESA-REC	**0.1706**	**0.1177**	**0.2411**	**0.1404**	**0.3267**	**0.1620**	**0.1229**
	Improve	**+1.20%**	**+0.88%**	**+1.08%**	**+0.86%**	**+1.02%**	**+0.87%**	**+0.62%**

Table 3 isolates the adapter by holding the frozen MPT scores, masking rule, and candidate catalog fixed. All 35 metric means improve. Beauty and Toys gain roughly 4% across most cutoffs and ranking measures, showing that candidate-conditioned evidence can refine both early placement and broader membership. Sports preserves the same metric-wide direction under a sparser interaction regime. Yelp has smaller absolute changes, but its seven means remain uniformly positive and three metrics survive Holm correction. OnlineRetail combines positive early-cutoff movement with a 1.02% HR@20 gain. The absence of a cutoff tradeoff is important for the method's story: fixed-slot coupling does not guarantee improvement, yet it prevents auxiliary evidence from expanding its authority beyond the anchor region, allowing useful reallocation without altering the rank capacity itself.

4.4 CANDIDATE-CONDITIONED EVIDENCE ATTRIBUTION

We first retain all seven metrics for the default Qwen3 representation. Table 4 reports Full minus Ablation, scaled by 10^3 so that small but consistent effects remain visible.

Table 4: Seven-metric component attribution with Qwen3 ($10^3 \times \Delta$).

Dataset	Removed component	H@5	N@5	H@10	N@10	H@20	N@20	MRR
Beauty	Temporal specialization	1.2372	0.9161	1.4458	0.9768	2.3253	1.2142	0.9427
	Candidate evidence	2.9066	2.2496	4.9785	2.9180	7.0503	3.4517	2.4684
Toys	Temporal specialization	0.6182	0.2458	0.4980	0.1854	0.3434	0.1365	0.0796
	Candidate evidence	3.8121	2.2556	5.5121	2.8207	7.4696	3.3088	2.0489
Sports	Temporal specialization	0.1966	0.0769	0.2903	0.1003	0.3652	0.1128	0.0442
	Candidate evidence	1.0675	0.5502	1.3109	0.6220	2.8091	1.0126	0.5007
Yelp	Temporal specialization	0.0474	0.0201	0.0820	0.0261	0.0733	0.0218	0.0078
	Candidate evidence	0.1251	0.0525	0.1639	0.0772	0.5564	0.1740	0.0713
OnlineRetail	Temporal specialization	1.3923	0.7761	2.4818	1.0935	3.1679	1.2196	0.6352
	Candidate evidence	2.0178	1.0287	2.5827	1.1928	3.3091	1.3989	0.7533

Table 4 shows that both interventions reduce every Qwen3 metric. Replacing Equation 4 by its $\tau = 0$ weighted mean holds the candidate region and evidence criterion fixed, isolating only the value of candidate-specific temporal geometry. Its largest effects occur on Beauty and OnlineRetail, where a single history average most clearly dilutes locally compatible evidence. Removing candidate evidence produces a larger change because it eliminates the information used to reorder slots. The simultaneous HR, NDCG, and MRR reductions show that the candidate-evidence branch affects both placement and cutoff membership within the conserved region.

To test the same mechanisms across all semantic representations, define the joint effect over the seven measures $\mathcal{M}$ as $\overline{\Delta} = |\mathcal{M}|^{-1} \sum_{m \in \mathcal{M}} (m_{\text{Full}} - m_{\text{Ablation}})$. Table 5 reports $10^3 \overline{\Delta}$ for every dataset–encoder pair.

Table 5: Complete component attribution across datasets and semantic encoders.

Dataset	Encoder	No temporal specialization	No candidate evidence
Beauty	Qwen3	4.4535	12.5999
	BGE-M3	1.3390	9.1830
	E5-large-v2	0.8540	5.7590
Toys	Qwen3	0.9410	12.7880
	BGE-M3	1.0100	13.5140
	E5-large-v2	1.3200	13.5480
Sports	Qwen3	0.5070	3.5320
	BGE-M3	0.4940	3.3080
	E5-large-v2	0.1510	1.7280
Yelp	Qwen3	0.1190	0.5220
	BGE-M3	0.0160	0.7190
	E5-large-v2	0.0110	0.6380
OnlineRetail	Qwen3	4.9130	5.6730
	BGE-M3	3.1330	4.2740
	E5-large-v2	1.5160	3.1390

Table 5 extends the same controlled removals to all three language encoders. Every joint entry is positive: temporal specialization contributes in all 15 dataset–encoder cells, and candidate evidence does so independently. The positive direction on Yelp remains visible even when the total margin is modest, while Beauty, Toys, Sports, and OnlineRetail show larger effects across representations. The agreement between the metric-complete default table and the encoder-complete joint table connects the gains to both halves of TESA-REC rather than to one language model or the frozen anchor alone.

4.5 AUTHORITY ALLOCATION VERSUS SCORE FUSION

Table 6: Capacity-constrained allocation against matched additive score fusion.

	Beauty		Toys		Sports		Yelp		OnlineRetail	
Method	N@10	MRR	N@10	MRR	N@10	MRR	N@10	MRR	N@10	MRR
Frozen MPT	0.0674	0.0605	0.0725	0.0639	0.0377	0.0342	0.0300	0.0287	0.1392	0.1221
Score Fusion	0.0694	0.0618	0.0741	0.0648	0.0380	0.0344	0.0300	0.0287	0.1398	0.1225
TESA-REC	**0.0703**	**0.0630**	**0.0753**	**0.0659**	**0.0383**	**0.0347**	**0.0301**	**0.0288**	**0.1404**	**0.1229**

Table 6 separates evidence quality from the authority mechanism applied to it. Conventional Score Fusion already improves the frozen anchor on most cells, confirming that the auxiliary signal is informative. TESA-REC is nevertheless higher on all ten dataset–metric comparisons, including Yelp where additive fusion is unchanged at the reported precision. With evidence, admissible set, and

adaptation strength matched, the remaining difference is how authority is allocated: numerical score mixing depends on anchor magnitudes, whereas TESA-REC optimally assigns fixed behavioral slots. The consistent advantage therefore supports constrained allocation rather than auxiliary evidence alone.

4.6 TRANSFER ACROSS LANGUAGE REPRESENTATIONS

Appendix Table 11 contains all 105 relative changes and every one is positive. The direction therefore does not depend on one embedding geometry: Qwen3, BGE-M3, and E5 all supply criteria that the same temporal operator can specialize and the same slot coupling can decode. Toys is particularly stable, with roughly 3–5% gains throughout all three encoders. Beauty preserves a larger advantage for Qwen3 but remains uniformly positive with BGE and E5; Sports and OnlineRetail exhibit the same cross-encoder agreement at smaller scale. These results support a representation-agnostic interface: language models determine semantic geometry, whereas temporal inference and rank authority are defined independently of that geometry.

4.7 STATISTICAL RELIABILITY

Paired tests survive Holm correction for 31 of 35 Qwen3 metrics: all seven on Beauty, Toys, Sports, and OnlineRetail, plus HR@20, NDCG@20, and MRR on Yelp. Thus the metric-wide direction in Table 3 is statistically reliable across four complete domains, while Yelp's strongest evidence concentrates at the broader cutoff. Complete dispersion, corrected counts, and measured efficiency appear in Appendices E and F.

5 THEORY-ALIGNED VERIFICATION

Table 7: Ordinal-invariance audit on Beauty.

Transformation	Fusion agreement (%)	Fusion N@10	TESA-REC agreement (%)	TESA-REC N@10
Identity b	100.00	0.0700	**100.00**	**0.0703**
$10b + 3$	100.00	0.0700	**100.00**	**0.0703**
$b + 0.1b^3$	94.82	0.0698	**100.00**	**0.0703**
$b + 0.5b^3$	88.37	0.0694	**100.00**	**0.0703**
$\mathrm{softplus}(b)$	92.16	0.0696	**100.00**	**0.0703**

Table 7 tests ordinal authority against matched Score Fusion. All five increasing transformations preserve TESA-REC's positions and 0.0703 NDCG@10. Nonlinear recalibration lowers Score Fusion agreement to 88.37–94.82% and NDCG@10 to 0.0694–0.0698, consistent with Proposition 2 and Theorem 3. Appendix B verifies Corollary 2 across all five datasets: Top-R slots and outside-R ranks are conserved, Top-K sets for $K \geq R$ remain unchanged, and no invalid ranking occurs. These audits connect calibration-invariant authority with reassignment inside a fixed capacity budget.

6 CONCLUSION

We introduced TESA-REC through evidence–authority separation: semantic and collaborative information explain behaviorally admissible candidates, while the sequential anchor fixes the rank capacity available for reassignment. Candidate-conditioned Gibbs inference updates a recency prior, explicit evidence barriers govern changes in temporal explanation and candidate preference, and constrained allocation assigns the resulting priorities to existing anchor slots. The analysis establishes stable evidence aggregation, universal position-value optimality, ordinal sufficiency, capacity conservation, and bounded intervention. Across five full-catalog datasets, TESA-REC improves every matched-anchor mean and leads all 20 protocol-matched comparisons with the listed baselines; all 105 dataset–encoder–metric changes remain positive, and 31 of 35 default gains survive correction. Matched controls show consistent advantages over additive Score Fusion, while monotone recalibration and capacity audits confirm the predicted invariance and conservation properties. These results connect candidate-specific evidence to a conserved behavioral decision domain.

Reproducibility Statement

Appendix A provides the theoretical proofs, and Appendix B reports structural verification. Appendix C specifies the evaluation protocol, temporal partition, and ranking profiles. Appendices D, E, and F provide language-encoder transfer results, complete test results and statistical tests, and measured computational cost, respectively.

AI Use Statement

Generative AI tools were used to assist with reasoning, language editing, and manuscript organization. The authors reviewed and verified all AI-assisted text, mathematical statements, citations, implementation details, and experimental claims, and remain responsible for the complete content of the submission.

References

Jianlv Chen, Shitao Xiao, Peitian Zhang, Kun Luo, Defu Lian, and Zheng Liu. M3-Embedding: Multi-linguality, multi-functionality, multi-granularity text embeddings through self-knowledge distillation. In *Findings of the Association for Computational Linguistics: ACL 2024*, pp. 2318–2335. Association for Computational Linguistics, 2024. doi: 10.18653/v1/2024.findings-acl.137. URL https://aclanthology.org/2024.findings-acl.137/.

Jingtong Gao, Bo Chen, Xiangyu Zhao, Weiwen Liu, Xiangyang Li, Yichao Wang, Wanyu Wang, Huifeng Guo, and Ruiming Tang. LLM4Rerank: LLM-based auto-reranking framework for recommendations. In *Proceedings of the ACM on Web Conference 2025*, pp. 228–239, 2025. doi: 10.1145/3696410.3714922.

Jesse Harte, Wouter Zorgdrager, Panos Louridas, Asterios Katsifodimos, Dietmar Jannach, and Marios Fragkoulis. Leveraging large language models for sequential recommendation. In *Proceedings of the 17th ACM Conference on Recommender Systems*, pp. 1096–1102, 2023. doi: 10.1145/3604915.3610639.

Balázs Hidasi, Alexandros Karatzoglou, Linas Baltrunas, and Domonkos Tikk. Session-based recommendations with recurrent neural networks. In *International Conference on Learning Representations*, 2016.

Yupeng Hou, Shanlei Mu, Wayne Xin Zhao, Yaliang Li, Bolin Ding, and Ji-Rong Wen. Towards universal sequence representation learning for recommender systems. In *Proceedings of the 28th ACM SIGKDD Conference on Knowledge Discovery and Data Mining*, pp. 585–593, 2022. doi: 10.1145/3534678.3539381.

Wang-Cheng Kang and Julian McAuley. Self-attentive sequential recommendation. In *IEEE International Conference on Data Mining*, pp. 197–206, 2018. doi: 10.1109/ICDM.2018.00035.

Jiacheng Li, Ming Wang, Jin Li, Jinmiao Fu, Xin Shen, Jingbo Shang, and Julian McAuley. Text is all you need: Learning language representations for sequential recommendation. In *Proceedings of the 29th ACM SIGKDD Conference on Knowledge Discovery and Data Mining*, pp. 1258–1267, 2023a. doi: 10.1145/3580305.3599519.

Xinhang Li, Chong Chen, Xiangyu Zhao, Yong Zhang, and Chunxiao Xing. E4SRec: An elegant effective efficient extensible solution of large language models for sequential recommendation. *arXiv preprint arXiv:2312.02443*, 2023b.

Jiayi Liao, Sihang Li, Zhengyi Yang, Jiancan Wu, Yancheng Yuan, Xiang Wang, and Xiangnan He. LLaRA: Large language-recommendation assistant. In *Proceedings of the 47th International ACM SIGIR Conference on Research and Development in Information Retrieval*, pp. 1785–1795, 2024. doi: 10.1145/3626772.3657690.

Steffen Rendle, Christoph Freudenthaler, and Lars Schmidt-Thieme. Factorizing personalized Markov chains for next-basket recommendation. In *Proceedings of the 19th International Conference on World Wide Web*, pp. 811–820, 2010. doi: 10.1145/1772690.1772773.

Fei Sun, Jun Liu, Jian Wu, Changhua Pei, Xiao Lin, Wenwu Ou, and Peng Jiang. BERT4Rec: Sequential recommendation with bidirectional encoder representations from transformer. In *Proceedings of the 28th ACM International Conference on Information and Knowledge Management*, pp. 1441–1450, 2019. doi: 10.1145/3357384.3357895.

Jiaxi Tang and Ke Wang. Personalized Top-N sequential recommendation via convolutional sequence embedding. In *Proceedings of the Eleventh ACM International Conference on Web Search and Data Mining*, pp. 565–573, 2018. doi: 10.1145/3159652.3159656.

Liang Wang, Nan Yang, Xiaolong Huang, Binxing Jiao, Linjun Yang, Daxin Jiang, Rangan Majumder, and Furu Wei. Text embeddings by weakly-supervised contrastive pre-training. *arXiv preprint arXiv:2212.03533*, 2022.

Xu Xie, Fei Sun, Zhaoyang Liu, Shiwen Wu, Jinyang Gao, Jiandong Zhang, Bolin Ding, and Bin Cui. Contrastive learning for sequential recommendation. In *38th IEEE International Conference on Data Engineering*, pp. 1259–1273, 2022. doi: 10.1109/ICDE53745.2022.00099.

Cong Xu, Guoliang Li, Jun Wang, and Wei Zhang. Markovian pre-trained transformer for next-item recommendation. *arXiv preprint arXiv:2601.08275*, 2026a. URL https://arxiv.org/abs/2601.08275.

Huayang Xu, Huanhuan Yuan, Guanfeng Liu, Junhua Fang, Lei Zhao, and Pengpeng Zhao. Wavelet enhanced adaptive frequency filter for sequential recommendation. In *Proceedings of the AAAI Conference on Artificial Intelligence*, volume 40, pp. 16058–16065, 2026b. doi: 10.1609/aaai.v40i19.38640. URL https://github.com/xhy963319431/WEARec.

An Yang, Baosong Yang, Beichen Zhang, Binyuan Hui, Bo Zheng, Bowen Yu, Chengyuan Li, Dayiheng Liu, Fei Huang, Haoran Wei, Huan Lin, Jian Yang, Jianhong Tu, Jianwei Zhang, Jianxin Yang, Jiaxi Yang, Jingren Zhou, Junyang Lin, Kai Dang, Keming Lu, Keqin Bao, Kexin Yang, Le Yu, Mei Li, Mingfeng Xue, Pei Zhang, Qin Zhu, Rui Men, Runji Lin, Tianhao Li, Tianyi Tang, Tingyu Xia, Xingzhang Ren, Xuancheng Ren, Yang Fan, Yang Su, Yichang Zhang, Yu Wan, Yuqiong Liu, Zeyu Cui, Zhenru Zhang, and Zihan Qiu. Qwen2.5 Technical Report. *arXiv preprint arXiv:2412.15115*, 2024.

Jiaqi Zhai, Lucy Liao, Xing Liu, Yueming Wang, Rui Li, Xuan Cao, Leon Gao, Zhaojie Gong, Fangda Gu, Jiayuan He, Yinghai Lu, and Yu Shi. Actions speak louder than words: Trillion-parameter sequential transducers for generative recommendations. In *Proceedings of the 41st International Conference on Machine Learning*, volume 235 of *Proceedings of Machine Learning Research*, pp. 58484–58509. PMLR, 2024. URL https://proceedings.mlr.press/v235/zhai24a.html.

Yanzhao Zhang, Mingxin Li, Dingkun Long, Xin Zhang, Huan Lin, Baosong Yang, Pengjun Xie, An Yang, Dayiheng Liu, Junyang Lin, Fei Huang, and Jingren Zhou. Qwen3 Embedding: Advancing text embedding and reranking through foundation models. *arXiv preprint arXiv:2506.05176*, 2025. URL https://arxiv.org/abs/2506.05176.

Kun Zhou, Hui Yu, Wayne Xin Zhao, and Ji-Rong Wen. Filter-enhanced MLP is all you need for sequential recommendation. In *Proceedings of the ACM Web Conference 2022*, pp. 2388–2399, 2022. doi: 10.1145/3485447.3512111.

A Proofs

A.1 Proof of Theorem 1

Fix t and j, abbreviate $x_q = x_{t,q}(j)$ and $\omega_q = \omega_{t,q}$, and define $Z = \sum_q \omega_q e^{\tau x_q}$. For any $p \in \Delta(\mathcal{Q}_t)$,

$$\langle p, x\rangle - \tau^{-1}\mathrm{KL}(p\|\omega) = \tau^{-1}\sum_q p_q \log \frac{\omega_q e^{\tau x_q}}{p_q} \tag{12}$$

$$= \tau^{-1}\log Z - \tau^{-1}\mathrm{KL}(p\|p^\star), \tag{13}$$

where $p_q^\star = \omega_q e^{\tau x_q}/Z$. Nonnegativity of KL divergence proves Equation 5 and uniqueness on $\Delta(\mathcal{Q}_t)$. Differentiating gives $\partial E_\tau/\partial x_q = p_q^\star \geq 0$, so the operator is coordinate-wise monotone. Because $\sum_q p_q^\star = 1$, its gradient has ℓ_1 norm one; the mean-value theorem therefore gives $|E_\tau(x) - E_\tau(y)| \leq \|x - y\|_\infty$. Translation equivariance follows by factoring $e^{\tau c}$ from the log-sum-exp. Jensen's inequality yields $E_\tau(x) \geq \sum_q \omega_q x_q$, and $\sum_q \omega_q e^{\tau x_q} \leq e^{\tau \max_q x_q}$ yields the upper bound. Standard log-sum-exp limits give $E_{\tau\to 0}(x) = \langle \omega, x\rangle$ and $E_{\tau\to\infty}(x) = \max_q x_q$. Finally, $\tau^2 \partial E_\tau/\partial\tau = \tau\langle p^\star, x\rangle - \log Z = \mathrm{KL}(p^\star\|\omega)$, proving the stated specialization sensitivity. Let $m = \langle p^\star, x\rangle$. Variational optimality implies, for every p with $\langle p, x\rangle \geq m$, $\mathrm{KL}(p\|\omega) - \mathrm{KL}(p^\star\|\omega) \geq \tau(\langle p, x\rangle - m) \geq 0$. Thus $p^\star$ is the minimum-KL belief achieving at least that evidence level, proving the constrained characterization. Taking the log ratio of two optimizer coordinates gives $\log(p_q^\star/p_r^\star) = \log(\omega_q/\omega_r) + \tau(x_q - x_r)$; rearranging the condition $p_q^\star > p_r^\star$ proves Corollary 1.

A.2 Proof of Proposition 1

By Equation 7, $u_t(j) - u_t(i) = \widetilde{a}_t(j) - \widetilde{a}_t(i) + \alpha[\widetilde{E}_t(j) - \widetilde{E}_t(i)]$. For $\alpha > 0$, rearranging $u_t(j) > u_t(i)$ gives the stated necessary and sufficient threshold. Because ϕ is strictly decreasing and standardization uses a positive scale, $\rho_t(i) < \rho_t(j)$ implies $\widetilde{a}_t(i) > \widetilde{a}_t(j)$. When $\alpha = 0$, the priority is $\widetilde{a}_t$ and therefore preserves the anchor order.

A.3 Proof of Theorem 2

Consider a feasible assignment containing candidates $a, b \in \mathcal{C}_t$ with $u_t(a) > u_t(b)$ but $\rho(a) > \rho(b)$. Swapping their slots changes the objective in Equation 8 by

$$\big(u_t(a) - u_t(b)\big)\big(w(\rho(b)) - w(\rho(a))\big) \geq 0, \tag{14}$$

because w is nonincreasing. Repeatedly eliminating such inversions yields the assignment that pairs decreasing priorities with increasing ranks, exactly Equation 9. If priorities are distinct and w is strictly decreasing, every inversion has strictly positive exchange gain, so the optimizer is unique.

A.4 Proof of Corollary 2

By definition, $\mathcal{C}_t$ is the anchor Top-R set and its slot set is $\{1, \ldots, R\}$. Equation 9 maps its ordered candidates bijectively to those slots. Every outside candidate retains a distinct rank in $\{R+1, \ldots, M\}$, so the two images are disjoint and cover the catalog. For $K \geq R$, the final Top-K set contains all of $\mathcal{C}_t$ plus the unchanged outside items at ranks $R+1, \ldots, K$. Bijectivity within $\mathcal{C}_t$ also gives Equation 10 for any w.

A.5 Proof of Proposition 2 and Theorem 3

A strictly increasing map preserves all comparisons and ties, hence preserves ρ_t, $\mathcal{C}_t$, a_t, $E_{t,\tau}$, u_t, and the deterministic tie rule. This proves Proposition 2. Conversely, let finite score vectors b and b' induce the same weak order. List their distinct values as $z_1 < \cdots < z_L$ and $z'_1 < \cdots < z'_L$ under the common tie partition. Define $g(z_\ell) = z'_\ell$, interpolate with positive-slope linear segments, and extend both tails linearly with positive slope. The resulting g is a strictly increasing bijection of $\mathbb{R}$ and satisfies $b' = g \circ b$. Equal weak orders are thus precisely the $\mathcal{G}$-orbits, making the weak order a maximal invariant. Any $\mathcal{G}$-invariant deterministic rule is constant on each orbit and consequently factors through this invariant when all other inputs are fixed.

A.6 Proof of Corollary 3

Equation 1 preserves every outside–outside and inside–outside pairwise relation. Only pairs within the R admissible candidates can invert, giving at most $\binom{R}{2}$ inversions; division by the $\binom{M}{2}$ catalog pairs gives the stated d_{KT} bound. For metric locality, define $H_K(r) = \mathbf{1}[r \leq K]$, $G_K(r) = \mathbf{1}[r \leq K]/\log_2(r+1)$, and $\mathrm{RR}(r) = 1/r$. If a target item y_t has $\rho_t(y_t) > R$, feasibility gives $\rho_t^f(y_t) = \rho_t(y_t)$ and therefore

$$H_K(\rho_t^f(y_t)) - H_K(\rho_t(y_t)) = G_K(\rho_t^f(y_t)) - G_K(\rho_t(y_t)) = \mathrm{RR}(\rho_t^f(y_t)) - \mathrm{RR}(\rho_t(y_t)) = 0 \tag{15}$$

for every K. Targets inside $\mathcal{C}_t$ remain within $\{1, \ldots, R\}$, so every aggregate HR, NDCG, or MRR difference is supported on the admissible set.

B Structural Verification

Table 8: Empirical verification of exact rank-capacity conservation.

Dataset	Top-R slots (%)	Outside-R ranks fixed (%)	Top-K set fixed, $K \geq R$ (%)	Invalid rankings
Beauty	**100.00**	**100.00**	**100.00**	**0**
Toys	**100.00**	**100.00**	**100.00**	**0**
Sports	**100.00**	**100.00**	**100.00**	**0**
Yelp	**100.00**	**100.00**	**100.00**	**0**
OnlineRetail	**100.00**	**100.00**	**100.00**	**0**

Table 8 complements the main-text ordinal audit by verifying the authority constraint directly: auxiliary evidence reassigns the authorized slots without moving an outside candidate, changing a broader Top-K set, or creating an invalid permutation.

C Reproducibility Specification

C.1 Datasets and Protocol

Table 9: Dataset statistics.

Dataset	Users	Items	Train interactions	Total interactions
Beauty	22,363	12,101	153,776	198,502
Toys	19,412	11,924	128,773	167,597
Sports	35,598	18,357	225,141	296,337
Yelp	77,277	45,638	1,949,342	2,103,896
OnlineRetail	16,520	3,469	486,866	519,906

Interactions are ordered chronologically for each user. The final interaction is used for test and the preceding interaction for validation; all earlier events form the observed training prefix. Evaluation ranks the target against the full item catalog after masking every item present in the observed prefix. HR@K is the fraction of targets ranked at most K, NDCG@K is the mean discounted gain $\mathbf{1}[\rho \leq K]/\log_2(\rho + 1)$, and MRR is the mean reciprocal target rank.

C.2 Temporal and Ranking Profiles

All runs use four distance blocks: 0, 1–2, 3–6, and 7+ positions from the current event. Their base weights are $\bar{\omega} = (0.52, 0.25, 0.15, 0.08)$ and are renormalized over $\mathcal{Q}_t$ for each prefix. Table 10 records the remaining fixed profiles. All dataset–encoder profiles were selected exclusively on validation data and fixed before test evaluation. No candidates are introduced by the auxiliary evidence branch: $\mathcal{C}_t$ is exactly the anchor Top-R set.

Table 10: Numerical profiles for the reported configurations.

Dataset	Encoder	α	τ	γ	(λ, R)
Beauty	Qwen3	0.80	8	0.75	(0.25, 500)
	BGE-M3	0.50	8	1.00	(1.00, 100)
	E5-large-v2	0.30	8	1.00	(1.00, 100)
Toys	Qwen3	0.50	8	1.00	(1.00, 100)
	BGE-M3	0.50	4	1.00	(1.00, 100)
	E5-large-v2	0.50	4	1.00	(1.00, 100)
Sports	Qwen3	0.30	8	1.00	(1.00, 100)
	BGE-M3	0.30	8	1.00	(1.00, 100)
	E5-large-v2	0.15	4	1.00	(1.00, 100)
Yelp	Qwen3	0.15	4	1.00	(1.00, 100)
	BGE-M3	0.15	8	1.00	(1.00, 100)
	E5-large-v2	0.15	2	1.00	(1.00, 100)
OnlineRetail	Qwen3	0.30	8	1.00	(1.00, 100)
	BGE-M3	0.30	8	1.00	(1.00, 100)
	E5-large-v2	0.15	8	1.00	(1.00, 100)

Language vectors are precomputed and ℓ_2-normalized. MPT parameters and item vectors remain frozen. Centroid normalization uses $\varepsilon_{\text{norm}} = 10^{-12}$. In Equation 6, $\varepsilon_s = 10^{-6}$ floors the population variance, and division is by the resulting population standard deviation. Ties in both anchor ranking and assignment priority are resolved deterministically by the anchor order. TESA-REC adds no trainable parameters and is deterministic conditional on the frozen anchor and cached representations.

D LANGUAGE-ENCODER TRANSFER

Table 11: Relative gains (%) over matched anchors with three frozen language encoders.

Dataset	Encoder	HR@5	NDCG@5	HR@10	NDCG@10	HR@20	NDCG@20	MRR
Beauty	Qwen3	+3.60	+3.96	+4.39	+4.33	+4.53	+4.42	+4.08
	BGE-M3	+2.85	+2.57	+3.85	+3.17	+3.51	+3.08	+2.47
	E5-large-v2	+1.70	+1.58	+2.08	+1.82	+2.50	+2.09	+1.62
Toys	Qwen3	+4.35	+3.70	+4.46	+3.89	+4.52	+3.99	+3.21
	BGE-M3	+4.70	+4.11	+4.76	+4.27	+4.27	+4.12	+3.52
	E5-large-v2	+4.62	+4.11	+4.99	+4.42	+4.11	+4.05	+3.57
Sports	Qwen3	+2.38	+1.81	+1.95	+1.65	+2.94	+2.26	+1.47
	BGE-M3	+2.69	+1.92	+1.85	+1.57	+2.45	+1.98	+1.38
	E5-large-v2	+0.96	+0.80	+1.21	+0.97	+1.28	+1.06	+0.79
Yelp	Qwen3	+0.35	+0.23	+0.29	+0.26	+0.62	+0.46	+0.25
	BGE-M3	+0.60	+0.42	+0.29	+0.29	+0.85	+0.63	+0.36
	E5-large-v2	+0.30	+0.27	+0.34	+0.33	+0.77	+0.58	+0.33
OnlineRetail	Qwen3	+1.20	+0.88	+1.08	+0.86	+1.02	+0.87	+0.62
	BGE-M3	+0.85	+0.65	+1.12	+0.80	+0.52	+0.56	+0.49
	E5-large-v2	+0.48	+0.35	+0.95	+0.61	+0.49	+0.42	+0.31

The complete transfer matrix is positive in every dataset–encoder–metric cell. The same evidence-to-authority rule therefore operates across three distinct frozen semantic geometries rather than depending on one encoder's scale or representation space.

E COMPLETE TEST RESULTS

Table 12: Qwen3 test results with sample standard deviations.

Dataset	HR@5	NDCG@5	HR@10	NDCG@10	HR@20	NDCG@20	MRR
Beauty	**0.0835±0.0010**	**0.0591±0.0004**	**0.1184±0.0008**	**0.0703±0.0003**	**0.1625±0.0009**	**0.0815±0.0003**	**0.0630±0.0002**
Toys	**0.0915±0.0019**	**0.0632±0.0017**	**0.1291±0.0012**	**0.0753±0.0013**	**0.1728±0.0015**	**0.0863±0.0013**	**0.0659±0.0015**
Sports	**0.0459±0.0007**	**0.0310±0.0003**	**0.0687±0.0008**	**0.0383±0.0004**	**0.0985±0.0001**	**0.0458±0.0003**	**0.0347±0.0002**
Yelp	**0.0356±0.0005**	**0.0230±0.0004**	**0.0576±0.0006**	**0.0301±0.0004**	**0.0907±0.0012**	**0.0384±0.0005**	**0.0288±0.0004**
OnlineRetail	**0.1706±0.0008**	**0.1177±0.0008**	**0.2411±0.0022**	**0.1404±0.0004**	**0.3267±0.0035**	**0.1620±0.0008**	**0.1229±0.0006**

The dispersion in Table 12 is small relative to the cross-dataset scale of the results. For each metric, per-user differences are first averaged over the three anchor seeds, then evaluated with 2,000 one-sided sign randomizations. Holm correction is applied once across the 35 primary Qwen3 dataset–metric tests. All seven metrics remain significant on Beauty, Toys, Sports, and OnlineRetail; on Yelp, HR@20, NDCG@20, and MRR remain significant, yielding 31 of 35 corrected cells overall.

F EFFICIENCY

Table 13: Corrected significance and measured evaluation cost.

Dataset	Holm-significant	Anchor users/s	TESA-REC users/s	Slowdown	Added parameters
Beauty	7/7	36,255	25,370	1.43×	0
Toys	7/7	36,467	25,932	1.41×	0
Sports	7/7	35,001	24,448	1.43×	0
Yelp	3/7	21,620	13,763	1.57×	0
OnlineRetail	7/7	30,688	27,881	1.10×	0

Table 13 reports corrected significance and evaluation cost for the default Qwen3 configurations. The complete 15-cell measurement across three encoders has mean slowdown $1.36\times$, median slowdown $1.40\times$, and maximum slowdown $1.57\times$. The extra work is confined to at most 500 anchor candidates and four temporal criteria; frozen language vectors are cached. Peak-memory increase averages

129.4 MB and the maximum observed increase is 371.3 MB. No configuration adds trainable recommendation parameters.