LSRB-REC: SOURCE-SPECIFIC DECISION RIGHTS FOR SEQUENTIAL RECOMMENDATION

Anonymous authors
Paper under double-blind review

ABSTRACT

Sequential recommenders commonly fuse heterogeneous signals into one globally editable score field, granting every source authority over every cutoff. This conflates predictive value with permission to alter a decision: evidence useful at a deeper candidate frontier can disturb an already reliable shallow ranking. We frame adaptation instead as source-specific decision rights, controlling how evidence evolves, what information it contributes, and where it may alter the list. The Language-Semantic Ranking Barrier for Recommendation (LSRB-REC) treats a ranking as nested decisions. A frozen behavioral anchor establishes the initial order; contractive collaborative transport discovers context-dependent candidates at a primary frontier; and language supplies anchor-orthogonal innovation only beyond a protected boundary. This construction turns heterogeneous evidence into staged, auditable adaptation. We prove transient stability, invariance to anchor-aligned semantics, exact preservation of the completed prefix, and a finite outer intervention bound. Across five full-catalog benchmarks, LSRB-REC improves every mean HR, NDCG, and MRR metric over its matched anchor and wins 18 of 20 protocol-matched state-of-the-art comparisons. Controls preserve all nine protected cutoffs, rank the contractive forecast first in all 15 matched cells with 0.71 state amplification, and rank the orthogonal residual first across three datasets; three language encoders maintain the improvement direction.

1 INTRODUCTION

Sequential recommendation makes nested decisions: the first results determine immediate utility, while broader cutoffs govern whether a relevant item remains reachable. Recurrent models learn event-wise state updates (Hidasi et al., 2016). Self-attention captures longer dependencies (Kang & McAuley, 2018), Markovian pretraining supplies transferable sequence priors (Xu et al., 2026a), and language transfers item meaning into recommendation spaces (Li et al., 2023a). Yet these advances commonly expose one globally editable score vector, giving heterogeneous evidence equal authority over the complete ranking. We view adaptation as an allocation of decision rights. A signal may be predictive at one frontier without authority to reopen an earlier decision: after cutoff K_ℓ is completed, later evidence acts only beyond it, inducing a hierarchy across $K_1 < \cdots < K_L$. The question is which source may change which decision.

We allocate these rights through a behavioral anchor, a dynamic branch that recovers candidates near a primary boundary, and language that resolves ambiguity at a protected outer frontier. This raises three challenges. **C1: stable candidate discovery.** Forecast evidence must evolve without amplifying state perturbations. **C2: complementary semantics.** Language must contribute information unexplained by the anchor where the decision remains uncertain. **C3: non-interfering decisions.** After the primary Top-K_1 is completed, every deeper semantic stage must preserve its membership and order. These challenges control *how* evidence evolves, *what* remains admissible, and *where* it may act. LSRB-REC realizes them by transient contraction, anchor-orthogonal semantic innovation, and a nested score barrier.

LSRB-REC implements this contract over a frozen MPT anchor (Xu et al., 2026a). A collaborative event drives a contractive persistent–transient forecast at the primary frontier; the outer stage extracts anchor-orthogonal semantic innovation and admits it only below the protected boundary. Under strict full-catalog evaluation, LSRB-REC yields positive mean changes on all 35 HR/NDCG/MRR

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.

cells across five datasets, wins 18 of 20 protocol-matched cells, and transfers across three semantic encoders. Direct controls show that the contractive forecast wins all 15 matched alternative cells and attenuates perturbed states to 0.71, while the barrier changes none of the protected Top-10 sets. Across three cutoff pairs, every protected cutoff is invariant and every deeper HR/NDCG change is positive.

Our contributions are:

- We formulate sequential adaptation as source-specific decision rights over nested ranking frontiers.
- We control *how*, *what*, and *where* auxiliary evidence acts through transient stability, anchor-orthogonal semantic innovation, and a nested barrier.
- We prove finite transient gain, invariance to anchor-aligned semantics, exact prefix preservation, and a proposal-bounded outer intervention; the construction composes over multiple frontiers.
- Five datasets, three semantic encoders, component interventions, and multiple cutoff pairs jointly verify utility, attribution, and exact compliance.

2 Related Work

Sequential recommendation has progressed from Markov factorization (Rendle et al., 2010) and recurrent session models (Hidasi et al., 2016) to convolutional sequence encoders (Tang & Wang, 2018), self-attention (Kang & McAuley, 2018), and bidirectional item prediction (Sun et al., 2019). Contrastive objectives improve sequence representations (Xie et al., 2022), while frequency-domain models expose periodic preference patterns (Zhou et al., 2022). Recent systems combine filtering with side information (Kim et al., 2025), frequency-enhanced inter-session paths (He et al., 2026), adaptive wavelets (Xu et al., 2026b), and wavelet-packet graph propagation (Liu et al., 2026). MPT instead learns a general sequence-processing prior from synthetic Markov chains (Xu et al., 2026a). LSRB-Rec treats such a strong encoder as a frozen behavioral anchor and learns a small, structurally constrained decision adapter around it.

Language-aware recommenders align text and collaborative representations (Hou et al., 2022). RecFormer pretrains item language within a sequential objective (Li et al., 2023a), while E4SRec adapts language models efficiently (Li et al., 2023b). LLaRA connects recommendation and language instructions (Liao et al., 2024); LLM2Rec directly uses language models as embedding models (He et al., 2025). Cross-view alignment in RLMRec further couples language and behavior (Ren et al., 2024), and FELLAS uses an LLM as an external semantic service (Yuan et al., 2025). Our focus is complementary: frozen language embeddings are admitted only at an outer decision frontier after their anchor-aligned component has been removed. This connects representation complementarity to rank preservation. FINEST studies rank-preserving parameter adaptation under perturbations (Oh et al., 2024); LSRB-Rec instead builds nested candidate-set invariance directly into the score operator, so one cutoff can improve without interfering with a protected shallower cutoff.

3 Method

3.1 Setup and Cutoff-Decoupled Overview

Let $\mathcal{I} = \{1, \dots, M\}$ be the item catalog and $x_{1:t} = (i_1, \dots, i_t)$ a user's observed interaction prefix. The task is to rank the next item i_{t+1} against the full eligible catalog. A frozen sequential anchor maps $x_{1:t}$ to a contextual state $h_t \in \mathbb{R}^{d_h}$ and provides frozen item vectors $e_j \in \mathbb{R}^{d_h}$. We write

$$b_t(j) = \left\langle \frac{h_t}{\|h_t\|_2}, \frac{e_j}{\|e_j\|_2} \right\rangle \tag{1}$$

for its base score. LSRB-Rec builds a sequence of decision contracts over this ranking. The anchor establishes the initial behavioral ordering; the primary stage receives bounded authority to revise it using a dynamic retrieval view at cutoff K_1; and, once the Top-K_1 decision is completed, the semantic stage may change membership near $K_2 > K_1$ while preserving that ordered prefix. Each stage introduces evidence together with an explicit limit on its authority. Figure 1 shows the complete path.

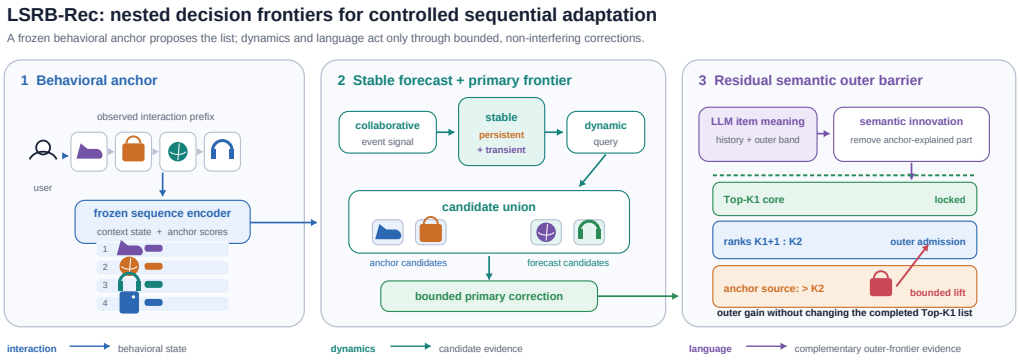

Figure 1: Overview of LSRB-REC. Collaborative dynamics enlarge and locally reorder the primary candidate region; residual language evidence acts only at a deeper frontier whose score cap preserves the completed Top-K_1 decision. Default experiments use $(K_1, K_2) = (10, 20)$.

3.2 DYNAMICAL CONTROL: STABLE COLLABORATIVE FORECAST

The collaborative forecast is constrained before it receives authority at the primary frontier. The anchor represents the current context, while the forecast identifies what becomes relevant next. We lift the frozen state with $L : \mathbb{R}^{d_h} \to \mathbb{R}^d$ and split $z_t = L(h_t) = [z_t^{\mathrm{p}}; z_t^{\mathrm{r}}]$, where $z_t^{\mathrm{p}} \in \mathbb{R}^p$ is persistent, $z_t^{\mathrm{r}} \in \mathbb{R}^{d-p}$ is transient, and $0 < p < d$. For the current event, a learned collaborative lookup supplies a normalized observable $c_t \in \mathbb{R}^{d_c}$. The one-step forecast is

$$\hat{z}_{t+1} = Az_t + u_t, \qquad A = \mathrm{diag}\Big(I_p, \mathrm{diag}(e^{-\operatorname{softplus}(a)})\Big), \qquad u_t = \begin{bmatrix} 0_p \\ W_c c_t \end{bmatrix}, \tag{2}$$

where $a \in \mathbb{R}^{d-p}$ and $W_c \in \mathbb{R}^{(d-p)\times d_c}$ are learned. The zero persistent prefix prevents an additive shortcut from overwriting the memory coordinates, while every transient singular value lies strictly between zero and one. A decoder $D \in \mathbb{R}^{d_h \times d}$ returns the forecast to the anchor space:

$$q_t = \frac{D\hat{z}_{t+1}}{\|D\hat{z}_{t+1}\|_2}, \qquad d_t(j) = \left\langle q_t, \frac{e_j}{\|e_j\|_2} \right\rangle. \tag{3}$$

Thus $d_t \in \mathbb{R}^M$ is a second full-catalog score view expressed in the same item space as the anchor.

Proposition 1 (Input-to-state stable transient forecast)**.** *Let $F(z, c) = Az + [0_p; W_c c]$. For any states $z, z' \in \mathbb{R}^d$ and event observables $c, c' \in \mathbb{R}^{d_c}$, the persistent discrepancy is retained, while the transient discrepancy obeys*

$$\begin{aligned} \big\|P_{\mathrm{p}}(F(z,c) - F(z',c'))\big\|_2 &= \big\|P_{\mathrm{p}}(z - z')\big\|_2, \\ \big\|P_{\mathrm{r}}(F(z,c) - F(z',c'))\big\|_2 &\le \rho\big\|P_{\mathrm{r}}(z - z')\big\|_2 + \|W_c\|_2\|c - c'\|_2. \end{aligned} \tag{4}$$

where P_{p} and P_{r} select the two coordinate blocks and $\rho = \max_k e^{-\operatorname{softplus}(a_k)} < 1$.

Appendix A gives the multi-step bound: initial transient error decays geometrically, and accumulated event perturbations remain bounded by a geometric gain. This forecast is intentionally collaborative: language has a separate, identifiable role at the outer barrier rather than a second global path into the state.

3.3 CANDIDATE UNION AND PRIMARY RANKING STAGE

After masking previous interactions, let $\mathcal{C}_t^b = \mathrm{TopK}_{K_b}(b_t)$ be the anchor support and let $\mathcal{H}_t = \mathrm{TopK}_{K_{\mathrm{tr}}}(b_t)$ be its optional trust region. The forecast set $\mathcal{C}_t^d$ contains the K_d largest values of d_t within $\mathcal{H}_t$ (or the eligible catalog when no trust region is used), and $\mathcal{C}_t = \mathcal{C}_t^b \cup \mathcal{C}_t^d$. These discrete retrieval sets are recomputed from detached scores; gradients pass through the selected score branches,

not through set membership. For a score vector v and finite set $\mathcal{C}$, let $\mathrm{std}_{\mathcal{C}}(v)$ denote candidate-wise centering and variance normalization. The dynamic advantage

$$a_t(j) = \mathrm{std}_{\mathcal{C}_t}(d_t)(j) - \mathrm{std}_{\mathcal{C}_t}(b_t)(j), \qquad j \in \mathcal{C}_t, \tag{5}$$

compares evidence in a dimensionless local score space. Let $\mathcal{P}_t^{(1)}$ contain at most the m_1 largest advantages in $\mathcal{C}_t^d \setminus \mathcal{C}_t^b$, let σ_t^b be the anchor-score standard deviation on $\mathcal{C}_t$, and let $\beta_t^b = \min_{j \in \mathcal{C}_t^b} b_t(j)$. Forecast reliability is defined by $g_t = g_{\max}\,\mathrm{sigm}(\theta_g)\exp\{-\mathrm{softplus}(f_e([z_t; u_t]))\} \in (0, g_{\max})$, where $f_e : \mathbb{R}^{2d} \to \mathbb{R}$ is a learned error predictor, $\theta_g \in \mathbb{R}$ is a learned global logit, and $0 < g_{\max} \le 1$ is a fixed ceiling. With inner scale $r_b \ge 0$, crossing gain $\lambda_1 > 0$, margin $\varepsilon_t^b > 0$, and $\pi_t(j) = \mathrm{clip}_{[0,1]}(\lambda_1 g_t\,\mathrm{sigm}(2a_t(j)))$, where sigm is the logistic function, the complete sparse primary operator is

$$s_t^{(1)}(j) = \begin{cases} \max\{b_t(j) + g_t r_b \sigma_t^b a_t(j),\ \beta_t^b + 0.1\varepsilon_t^b\}, & j \in \mathcal{C}_t^b, \\ b_t(j) + \pi_t(j)[\beta_t^b + \varepsilon_t^b - b_t(j)]_+, & j \in \mathcal{P}_t^{(1)}, \\ b_t(j), & \text{otherwise.} \end{cases} \tag{6}$$

Thus the anchor support receives only a scale-matched local correction, forecast-only admission is sparse and confidence bounded, and all remaining catalog scores are unchanged. The resulting $s_t^{(1)} \in \mathbb{R}^M$ completes the primary ranking. Appendix B records the numerical settings and the two-boundary extension used by the multi-cutoff profiles.

3.4 Information Control: Orthogonal Semantic Innovation

Language receives authority only through information unexplained by the anchor. The deeper frontier draws candidates from the masked anchor ranking. Specifically, $\mathcal{F}_t = \mathrm{TopK}_{K_{\mathrm{tr}}^o}(b_t) \setminus \mathrm{TopK}_{K_2}(b_t)$ is the trusted outer source band, where $K_{\mathrm{tr}}^o > K_2$; selected items may enter ranks $K_1 + 1{:}K_2$ while preserving Top-K_1. Let $v_j \in \mathbb{R}^{d_\ell}$ be a frozen language embedding for item j. We summarize the semantic trajectory causally:

$$r_t = \frac{\sum_{\xi=1}^{t} \omega_{t\xi} v_{i_\xi}}{\|\sum_{\xi=1}^{t} \omega_{t\xi} v_{i_\xi}\|_2}, \qquad \lambda_t(j) = \left\langle r_t, \frac{v_j}{\|v_j\|_2} \right\rangle, \tag{7}$$

where $\omega_{t\xi} > 0$ decreases with the age of interaction i_ξ. Within $\mathcal{F}_t$, let $\bar{b}_t, \bar{d}_t, \bar{\lambda}_t \in \mathbb{R}^{|\mathcal{F}_t|}$ be the standardized vectors induced by b_t, d_t, and λ_t. The language channel is residualized against the anchor:

$$\eta_t = \bar{\lambda}_t - \frac{\langle \bar{\lambda}_t, \bar{b}_t \rangle}{\|\bar{b}_t\|_2^2} \bar{b}_t, \qquad p_t = (\bar{d}_t - \bar{b}_t) + \gamma \eta_t, \tag{8}$$

where $\eta_t \in \mathbb{R}^{|\mathcal{F}_t|}$ is the semantic innovation, p_t is the outer priority, and $\gamma > 0$ controls semantic evidence.

Proposition 2 (Local semantic complementarity)**.** *Whenever $\|\bar{b}_t\|_2 > 0$, the innovation satisfies $\langle \eta_t, \bar{b}_t \rangle = 0$ and is the unique vector in $\bar{b}_t^{\perp}$ whose difference from $\bar{\lambda}_t$ lies in $\mathrm{span}(\bar{b}_t)$. Moreover, it is invariant to every anchor-aligned semantic component: replacing $\bar{\lambda}_t$ by $\bar{\lambda}_t + c\bar{b}_t$ leaves η_t unchanged for any $c \in \mathbb{R}$.*

The construction therefore does not simply add another copy of the local collaborative order. It retains exactly the language direction that is linearly unexplained inside the band where the outer decision is made.

3.5 Ranking Control: Nested Outer Barrier

The outer barrier specifies which decision semantic evidence may change: uncertainty at K_2 may be resolved while the completed Top-K_1 prefix remains protected. Let $\beta_{1,t}^{(1)}$ and $\beta_{2,t}^{(1)}$ be the K_1-th and K_2-th largest eligible values of $s_t^{(1)}$, with deterministic tie breaking. From $\mathcal{F}_t$, select at most m_2 items with highest priority p_t as $\mathcal{P}_t$. Let $\sigma_{t,2}$ be the score deviation within $\mathrm{TopK}_{K_2}(s_t^{(1)})$ and define

$\delta_t = \min\{\kappa_o \sigma_{t,2},\, 0.25(\beta^{(1)}_{1,t} - \beta^{(1)}_{2,t})\}$, $\varepsilon_t = 0.25\delta_t$, and $\alpha_t(j) = \mathrm{clip}_{[0,1]}(\lambda_o g_t \,\mathrm{sigm}(2p_t(j)))$. For each $j \in \mathcal{P}_t$, define

$$s^f_t(j) = \max\Big\{ s^{(1)}_t(j), \min\Big(\beta^{(1)}_{1,t} - \varepsilon_t, s^{(1)}_t(j) + \alpha_t(j)[\beta^{(1)}_{2,t} + \delta_t - s^{(1)}_t(j)]_+\Big)\Big\}, \tag{9}$$

Scores outside $\mathcal{P}_t$ are unchanged. The inner minimum is the critical barrier: an outer candidate can cross the Top-K_2 boundary but remains strictly below the completed Top-K_1 boundary.

Theorem 1 (Cutoff decoupling). *Let $\mathcal{R}_K(s)$ be the ordered Top-K list under a fixed deterministic tie rule and $\mathcal{T}_K(s)$ its item set. The outer operator in Equation 9 satisfies*

$$\mathcal{R}_{K_1}(s^f_t) = \mathcal{R}_{K_1}(s^{(1)}_t), \qquad \Big|\mathcal{T}_{K_2}(s^f_t) \triangle \mathcal{T}_{K_2}(s^{(1)}_t)\Big| \le 2m_2. \tag{10}$$

Moreover, every changed score belongs to $\mathcal{P}_t$ and is bounded above by $\beta^{(1)}_{1,t} - \varepsilon_t$.

Theorem 1 formalizes a separation of decision rights: after the primary Top-K_1 decision is completed, every later-stage modification is restricted to unresolved decisions beyond that frontier. Shallow preservation is therefore a structural property of the score operator, while the symmetric-difference bound localizes deeper changes to selected proposals. Appendix A proves the result and its composition over an arbitrary ladder $K_1 < \cdots < K_L$, where each completed frontier becomes protected before the next source acts.

3.6 Learning Objective and Cost

For next item $y_t = i_{t+1}$ and weighted negatives $j \in \mathcal{C}_t \setminus \{y_t\}$, the trainable forecast and primary ranking path minimize

$$\mathcal{L} = -\frac{1}{Z_t} \sum_{j \in \mathcal{C}_t \setminus \{y_t\}} w_t(j) \log \mathrm{sigm}\left(\frac{s^{(1)}_t(y_t) - s^{(1)}_t(j)}{\tau}\right), \qquad Z_t = \sum_j w_t(j), \tag{11}$$

where $w_t(j) > 0$ emphasizes high-ranked candidates and $\tau > 0$ is the temperature. The sequence anchor and language embeddings remain frozen. The dynamic branch adds one catalog dot-product pass and sparse operations on $\mathcal{C}_t$ and $\mathcal{F}_t$; detailed parameter counts and measured throughput are reported in Appendix F.

4 Experiments

We test three properties of the allocation: **Q1–Utility** (Tables 1–3), **Q2–Attribution** to stages and semantic sources (Tables 4–5), and **Q3–Control** of protected frontiers (Table 6 and Appendix E).

4.1 Experimental Setup

Datasets and evaluation. We evaluate Beauty, Toys, Sports, Yelp, and OnlineRetail with chronological leave-one-out splits. At test time, each target is ranked against the full item catalog after masking every item in the user's observed history. We report HR and NDCG at 5, 10, and 20 together with MRR. Reported aggregate results use independent runs and sample standard deviations; paired user-level randomization tests are Holm-corrected across metrics. Dataset statistics and complete protocol details are in Appendix B.1.

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), MPT-LoRA (Xu et al., 2026a), and WPGRec (Liu et al., 2026).

Implementation. The default language representation is Qwen3-Embedding (Zhang et al., 2025); BGE-M3 (Chen et al., 2024) and E5-large-v2 (Wang et al., 2022) test semantic-source transfer. All methods in Tables 1 and 2 share the same dataset split, full-catalog ranking rule, and historical-item masking. LSRB-Rec uses the frozen MPT item and sequence spaces as its anchor. Essential numerical settings are listed in Appendix B.

4.2 PROTOCOL-MATCHED COMPARISON

Table 1: Protocol-matched full-catalog comparison on Amazon domains.

	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
WPGRec	0.1197	0.0742	0.1583	0.0813	–	–	–	–	–	–	–	–
LSRB-REC	**0.1181**	**0.0690**	**0.1834**	**0.0848**	**0.1303**	**0.0754**	**0.1912**	**0.0900**	**0.0694**	**0.0383**	**0.0990**	**0.0456**
Improve	−1.35%	−6.99%	**+15.85%**	**+4.28%**	**+8.43%**	**+6.95%**	**+18.23%**	**+11.23%**	**+6.67%**	**+2.68%**	**+6.77%**	**+3.08%**

Table 2: Protocol-matched full-catalog comparison on Yelp and OnlineRetail.

	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
LSRB-REC	**0.0644**	**0.0324**	**0.1018**	**0.0416**	**0.2459**	**0.1450**	**0.3310**	**0.1665**
Improve	**+30.10%**	**+34.40%**	**+22.30%**	**+27.91%**	**+52.95%**	**+93.65%**	**+28.54%**	**+67.38%**

Tables 1 and 2 report every model in the matched five-domain comparison, with WPGRec included where its Beauty setting is compatible. The comparison spans conventional sequence encoders, frequency and transducer models, text-aware recommenders, direct Qwen adaptation, and both MPT adaptation routes. LSRB-REC is best in all 20 cells of the MPT suite and sets the strongest protocol-matched result in 18 of 20 cells overall. Beauty provides a clear cutoff profile: HR@20 reaches 0.1834 and NDCG@20 reaches 0.0848, corresponding to gains of 15.85% and 4.28%. The other four datasets improve both membership and discounted rank at both cutoffs; the large Yelp and OnlineRetail gains reflect recovered candidates, while Sports confirms the same behavior in a different density regime. Improvements over MPT-Adapter and MPT-LoRA further identify the value of the staged decision interface within the same anchor family. The primary stage retrieves and reorders candidates, while the protected outer stage admits complementary evidence. These tables establish absolute competitiveness; Table 3 holds the frozen ranking space fixed to isolate the adapter itself.

4.3 MATCHED-ANCHOR GAINS ACROSS FIVE DOMAINS

Table 3: Complete LSRB-REC and frozen-MPT-anchor results under strict full-catalog evaluation. Improve is relative to the anchor.

Dataset	Variant	HR@5	NDCG@5	HR@10	NDCG@10	HR@20	NDCG@20	MRR
Beauty	Frozen MPT	0.0803	0.0569	0.1148	0.0680	0.1581	0.0789	0.0610
	LSRB-REC	**0.0804**	**0.0570**	**0.1181**	**0.0690**	**0.1834**	**0.0848**	**0.0614**
	Improve	+0.24%	+0.23%	+2.82%	+1.48%	+15.99%	+7.44%	+0.63%
Toys	Frozen MPT	0.0898	0.0624	0.1246	0.0736	0.1660	0.0841	0.0651
	LSRB-REC	**0.0898**	**0.0625**	**0.1303**	**0.0754**	**0.1912**	**0.0900**	**0.0656**
	Improve	+0.06%	+0.16%	+4.60%	+2.38%	+15.18%	+7.03%	+0.67%
Sports	Frozen MPT	0.0442	0.0302	0.0656	0.0371	0.0937	0.0442	0.0340
	LSRB-REC	**0.0465**	**0.0311**	**0.0694**	**0.0383**	**0.0990**	**0.0456**	**0.0342**
	Improve	+5.19%	+2.94%	+5.91%	+3.30%	+5.67%	+3.17%	+0.65%
Yelp	Frozen MPT	0.0369	0.0236	0.0604	0.0312	0.0960	0.0401	0.0299
	LSRB-REC	**0.0384**	**0.0242**	**0.0644**	**0.0324**	**0.1018**	**0.0416**	**0.0300**
	Improve	+4.28%	+2.59%	+6.54%	+3.82%	+6.01%	+3.62%	+0.52%
OnlineRetail	Frozen MPT	0.1761	0.1225	0.2459	0.1450	0.3309	0.1664	0.1275
	LSRB-REC	**0.1764**	**0.1226**	**0.2459**	**0.1450**	**0.3310**	**0.1665**	**0.1276**
	Improve	+0.16%	+0.11%	+0.02%	+0.05%	+0.02%	+0.07%	+0.09%

The matched comparison in Table 3 isolates the effect of the complete adapter from changing the backbone. Every mean cell is positive. The cutoff profile is informative: Beauty and Toys obtain their largest gains at 20, where the outer frontier expands useful membership after the primary list has been fixed; Sports and Yelp improve substantially at both 10 and 20, showing that dynamic candidate discovery and the outer barrier can contribute together. OnlineRetail also preserves the positive direction across all seven metrics, with corrected NDCG@20 and MRR differences reaching significance. Paired tests confirm five of seven metrics on Beauty and Toys, all seven on Sports and Yelp, and two on OnlineRetail (Appendix D). Across domains, the adapter jointly advances candidate membership, discounted position, and reciprocal rank.

4.4 Component Attribution and Exact Non-Interference

Table 4: Complete seven-metric component ablation, reported as Full minus the indicated removal.

Dataset	Removed component	ΔH@5	ΔN@5	ΔH@10	ΔN@10	ΔH@20	ΔN@20	ΔMRR
Beauty	LLM frontier	0.0000	0.0000	0.0000	0.0000	0.0176	0.0040	0.0003
	Candidate union	0.0050	0.0056	0.0080	0.0064	0.0297	0.0112	0.0058
	Secondary barrier	0.0000	0.0000	0.0000	0.0000	0.0253	0.0058	0.0003
Toys	LLM frontier	0.0000	0.0000	0.0000	0.0000	0.0185	0.0042	0.0003
	Candidate union	0.0032	0.0014	0.0092	0.0031	0.0281	0.0072	0.0011
	Secondary barrier	0.0000	0.0000	0.0000	0.0000	0.0252	0.0057	0.0002
Sports	LLM frontier	0.0000	0.0000	0.0000	0.0000	0.0017	0.0004	0.0000
	Candidate union	0.0056	0.0042	0.0076	0.0047	0.0099	0.0051	0.0036
	Secondary barrier	0.0000	0.0000	0.0000	0.0000	0.0053	0.0012	0.0000
Yelp	LLM frontier	0.0000	0.0000	0.0000	0.0000	0.0011	0.0002	0.0000
	Candidate union	0.0062	0.0049	0.0092	0.0057	0.0120	0.0062	0.0044
	Secondary barrier	0.0000	0.0000	0.0000	0.0000	0.0058	0.0013	0.0000
OnlineRetail	LLM frontier	0.0000	0.0000	0.0000	0.0000	0.0001	0.0000	0.0000
	Candidate union	0.0122	0.0120	0.0067	0.0102	0.0040	0.0096	0.0109
	Secondary barrier	0.0000	0.0000	0.0000	0.0000	0.0001	0.0000	0.0000

Table 4 removes each component from the same trained configurations and reports all seven evaluation metrics. Candidate union has the broadest effect because it supplies both new membership and local order evidence; its removal reduces every metric on every dataset. The LLM frontier contributes at the deeper cutoff in all five domains, with especially large HR@20 changes on Beauty and Toys. Most importantly, removing either the LLM frontier or the complete secondary barrier changes HR and NDCG at both 5 and 10 by exactly 0.0000 everywhere. This is the empirical fingerprint of Theorem 1: the outer mechanism is useful at 20, yet cannot touch the protected Top-10 membership or order. Because its scoring support lies below the primary boundary, the zero change holds query by query. The complete ablation directly connects the metric-wide behavior to the proposed staged construction.

4.5 Transfer across Language Representations

Table 5: Complete seven-metric relative gains (%) with three frozen semantic representations.

Dataset	Semantic encoder	HR@5	NDCG@5	HR@10	NDCG@10	HR@20	NDCG@20	MRR
Beauty	Qwen3	+0.24	+0.23	+2.82	+1.48	+15.99	+7.44	+0.63
	BGE-M3	+0.62	+0.59	+4.10	+2.03	+5.50	+2.66	+0.74
	E5-large-v2	+0.19	+0.25	+1.43	+0.81	+5.83	+2.79	+0.39
Toys	Qwen3	+0.06	+0.16	+4.60	+2.38	+15.18	+7.03	+0.67
	BGE-M3	+0.19	+0.21	+4.33	+2.15	+15.44	+7.08	+0.62
	E5-large-v2	+0.19	+0.24	+4.38	+2.22	+16.28	+7.50	+0.67
Sports	Qwen3	+5.19	+2.94	+5.91	+3.30	+5.67	+3.17	+0.65
	BGE-M3	+5.34	+2.96	+6.26	+3.34	+6.31	+3.42	+0.66
	E5-large-v2	+5.33	+2.91	+6.01	+3.20	+5.90	+3.19	+0.63
Yelp	Qwen3	+4.28	+2.59	+6.54	+3.82	+6.01	+3.62	+0.52
	BGE-M3	+4.20	+2.54	+6.39	+3.77	+6.32	+3.79	+0.52
	E5-large-v2	+4.16	+2.50	+6.37	+3.72	+6.10	+3.65	+0.50
OnlineRetail	Qwen3	+0.16	+0.11	+0.02	+0.05	+0.02	+0.07	+0.09
	BGE-M3	+0.02	+0.02	+0.02	+0.01	+0.01	+0.03	+0.03
	E5-large-v2	+0.14	+0.10	+0.02	+0.06	+0.02	+0.09	+0.10

The frontier transfers across embedding families. Table 5 reports all seven metrics and is positive in every dataset–encoder configuration. BGE-M3 and E5 preserve the strong Toys outer-cutoff effect

and align closely with Qwen3 on Sports and Yelp. Beauty and OnlineRetail likewise maintain the matched-anchor improvement direction for all three encoders. Because the outer barrier consumes cosine similarities and an anchor-orthogonal residual, the protected decision geometry is shared across language models while each semantic space supplies its own candidate priorities.

4.6 Controlled Tests of the Decision Interface

For the outer stage, I_K is the fraction of histories for which $\mathcal{T}_K(s_t^f) \neq \mathcal{T}_K(s_t^{(1)})$; Avg. $\triangle_K$ is the corresponding mean symmetric-difference cardinality. Both are measured relative to the completed primary ranking, not the frozen anchor.

Table 6: Structural compliance and cutoff transfer of the staged decision interface.

A. Orthogonal semantic innovation and barrier enforcement

Dataset	ΔH@20	ΔN@20	I_{10} no barrier	I_{10} soft	I_{10} hard	Hard H@20	Hard N@20
Beauty	+0.0039	+0.0010	0.5599	0.5461	**0.0000**	**0.1885**	**0.0870**
Toys	+0.0059	+0.0015	0.5962	0.5847	**0.0000**	**0.1942**	**0.0909**
Yelp	+0.0010–0.0013	+0.0002–0.0003	0.4400	0.4250	**0.0000**	**0.1020–0.1023**	**0.0413–0.0414**

B. Outer-stage transfer across nested cutoff pairs

Dataset	(K_1, K_2)	I_{K_1}	H@K_2: primary→final	N@K_2: primary→final
Beauty	$(5, 10)$	**0.0000**	0.1224→**0.1364**	0.0710→**0.0751**
	$(10, 20)$	**0.0000**	0.1615→**0.1873**	0.0807→**0.0866**
	$(20, 50)$	**0.0000**	0.2314→**0.2496**	0.0946→**0.0978**
Toys	$(5, 10)$	**0.0000**	0.1316→**0.1443**	0.0758→**0.0795**
	$(10, 20)$	**0.0000**	0.1676→**0.1933**	0.0847→**0.0906**
	$(20, 50)$	**0.0000**	0.2349→**0.2518**	0.0980→**0.1010**
Yelp	$(5, 10)$	**0.0000**	0.0642→**0.0679**	0.0320→**0.0331**
	$(10, 20)$	**0.0000**	0.0952→**0.1018**	0.0397→**0.0412**
	$(20, 50)$	**0.0000**	0.1688→**0.1720**	0.0542→**0.0548**

Table 6 tests structural compliance of the staged decision interface. In Panel A, the anchor-orthogonal residual gives the strongest outer metrics on all three datasets while leaving Top-10 unchanged. An unrestricted intervention alters 44.00–59.62% of completed Top-10 sets, and a soft penalty still alters 42.50–58.47%; the hard barrier reduces interference to zero. Panel B varies both cutoffs: every primary→final pair improves at K_2, while $I_{K_1} = 0$ certifies that the completed K_1 list is untouched. These query-level outcomes match Theorem 1 across Top-5/10, Top-10/20, and Top-20/50. Appendix E provides the full controls: the contractive forecast is best in all 15 matched alternative cells and, under perturbation, reduces state amplification to 0.71 and the HR@20 drop to 0.0018. Corrected tests and measured cost appear in Appendices D and F.

5 From Score Fusion to Nested Decision Rights

LSRB-Rec separates evidence quality from decision authority: stability controls how evidence propagates, residualization controls what information is admissible, and the barrier controls where it may act. Together they define a source-specific contract over nested frontiers.

5.1 Why Local Residualization Is the Right Semantic Interface

Semantic relevance becomes actionable after conditioning on collaborative plausibility. The trusted outer band focuses the semantic decision on anchor-plausible candidates beyond the protected core. Equation 8 asks, within $\mathcal{F}_t$, which language preference is unexplained by anchor order. The orthogonal component η_t is exactly this usable innovation: it vanishes when language and anchor scores are collinear, retains the full standardized language direction when they are orthogonal, and is unchanged if any multiple of the anchor direction is added to the language score. This invariance distinguishes complementary evidence from ordinary score normalization. Consistently, Panel A of Table 6 and the complete semantic comparison in Appendix E show that the orthogonal residual is strongest at the outer frontier on Beauty, Toys, and Yelp while preserving the protected metrics exactly.

5.2 Metric Locality from Score-Support Locality

Theorem 1 has a per-query metric consequence. Once $s_t^{(1)}$ is fixed, the outer operator changes neither the items nor their relative order inside Top-K_1; hence the target's HR and NDCG contributions are exactly unchanged for every $K \leq K_1$. At K_2, improvement can occur only through one of the at most m_2 selected outer candidates. This yields a falsifiable signature: an outer-only intervention must produce exact zeros at the protected cutoff and may produce positive changes only beyond it. Table 4 exhibits this pattern in the full five-domain ablation; Panel B of Table 6 reproduces it for all nine combinations of three datasets and three cutoff pairs. The hard-versus-soft comparison further separates construction from regularization: reducing average score movement is insufficient when many items are nearly tied, whereas a barrier certifies the ordered prefix query by query.

5.3 How the Two Frontiers Cooperate

The primary and secondary operations solve different ranking failures. The candidate union addresses *omission*: a target can be absent from the anchor's local candidate set even when the current collaborative event makes it relevant. Stable transport supplies an alternative query, and the primary barrier converts that evidence into a sparse boundary crossing or an internal order correction. The outer operation addresses *ambiguity*: several anchor-plausible candidates occupy the band beyond the secondary cutoff, and language provides complementary evidence about which one fits the semantic trajectory. This division assigns global candidate discovery to collaborative dynamics and protected outer-band resolution to language semantics.

This division also clarifies the domain pattern. Sports and Yelp primarily expose dynamic omission, with the semantic frontier providing an additional positive refinement. Beauty and Toys contain many recoverable outer-band targets, producing 15–16% relative HR@20 gains together with stable shallow-cutoff behavior. OnlineRetail expresses the same composition through consistent inner-union effects and transfer across all three semantic encoders. Across density regimes, the gains follow the amount and location of evidence available at the two ranking frontiers. This yields a stage-wise diagnostic signature: primary-cutoff gains measure collaborative recovery, outer-only gains measure admissible semantic resolution, and exact zeros at protected cutoffs verify that later evidence exercised no retroactive authority. The metric profile therefore identifies not only whether adaptation helps, but also which source changed which nested decision. The same composition also yields an explicit list-level intervention budget relative to the frozen anchor: collaborative forecast admissions spend the primary budget, while language can spend only the protected outer budget. Appendix A.4 derives the corresponding symmetric-difference bounds and shows why support control is stronger than a dense score penalty for auditing discrete rank changes.

6 Conclusion

We introduced a source-specific decision-rights view of sequential adaptation. Instead of fusing every signal into one globally editable score field, LSRB-Rec treats recommendation as a sequence of protected decisions. A frozen anchor establishes the behavioral ordering; contractive collaborative transport receives bounded authority to recover candidates; anchor-orthogonal residualization admits complementary semantic information; and the nested barrier restricts that information to unresolved ranking frontiers. The theory makes each contract explicit: transient perturbations have finite gain, anchor-aligned semantic evidence is removed exactly, later interventions preserve the completed ordered prefix, and deeper membership change is bounded by a proposal budget. The experiments recover the same structure: the contractive forecast wins all 15 matched alternative cells and attenuates perturbed states; the orthogonal residual leads every controlled semantic comparison; all matched-anchor means improve; 18 of 20 protocol-matched state-of-the-art cells are won; and every tested outer intervention improves its deeper cutoff while respecting the protected decision exactly. More broadly, the result separates evidence quality from decision scope. Different sources can be governed in how they evolve, what they contribute, and where they act, and the same construction composes over multiple ranking frontiers. Rank preservation therefore becomes a property of the decision interface itself rather than an average consequence of regularization.

Reproducibility Statement

Appendix A provides the theoretical proofs and structural properties. Appendix B specifies the data splits, candidate masking, metrics, model geometry, fixed ranking profiles, and numerical safeguards. Appendices C, D, E, and F provide complete matched-anchor results, statistical tests, controlled mechanism tests, and efficiency details, respectively. All result tables retain a leading zero and four decimal places in the main text; machine-precision outputs are preserved in the accompanying artifacts.

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/.

Peng He, Yanglei Gan, Tingting Dai, Run Lin, Xuexin Li, Yao Liu, and Qiao Liu. Exploiting inter-session information with frequency-enhanced dual-path networks for sequential recommendation. In *Proceedings of the AAAI Conference on Artificial Intelligence*, volume 40, pp. 14820–14828, 2026. doi: 10.1609/aaai.v40i17.38502. URL https://github.com/AONE-NLP/FreqRec.

Yingzhi He, Xiaohao Liu, An Zhang, Yunshan Ma, and Tat-Seng Chua. LLM2Rec: Large language models are powerful embedding models for sequential recommendation. In *Proceedings of the 31st ACM SIGKDD Conference on Knowledge Discovery and Data Mining V.2*, pp. 896–907, 2025. doi: 10.1145/3711896.3737029.

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.

Hye-young Kim, Minjin Choi, Sunkyung Lee, Ilwoong Baek, and Jongwuk Lee. DIFF: Dual side-information filtering and fusion for sequential recommendation. In *Proceedings of the 48th International ACM SIGIR Conference on Research and Development in Information Retrieval*, pp. 1624–1633, 2025. doi: 10.1145/3726302.3729948.

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.

Peilin Liu, Zhiquan Ji, and Gang Yan. WPGRec: Wavelet packet guided graph enhanced sequential recommendation. In *Proceedings of the 49th International ACM SIGIR Conference on Research and Development in Information Retrieval*, pp. 3964–3969, 2026. doi: 10.1145/3805712.3809907. URL https://doi.org/10.1145/3805712.3809907.

Sejoon Oh, Berk Ustun, Julian McAuley, and Srijan Kumar. FINEST: Stabilizing recommendations by rank-preserving fine-tuning. *ACM Transactions on Knowledge Discovery from Data*, 18(9): 1–22, 2024. doi: 10.1145/3695256. URL https://doi.org/10.1145/3695256.

Xubin Ren, Wei Wei, Lianghao Xia, Lixin Su, Suqi Cheng, Junfeng Wang, Dawei Yin, and Chao Huang. Representation learning with large language models for recommendation. In *Proceedings of the ACM Web Conference 2024*, pp. 3464–3475, 2024. doi: 10.1145/3589334.3645458. URL https://doi.org/10.1145/3589334.3645458.

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.

Wei Yuan, Chaoqun Yang, Guanhua Ye, Tong Chen, Quoc Viet Hung Nguyen, and Hongzhi Yin. FELLAS: Enhancing federated sequential recommendation with LLM as external services. *ACM Transactions on Information Systems*, 43(6):1–24, 2025. doi: 10.1145/3709138. URL https://doi.org/10.1145/3709138.

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 and Structural Properties

A.1 Persistent–Transient Stability

Let $P_{\mathrm{p}} = [I_p \ 0]$ and $P_{\mathrm{r}} = [0 \ I_{d-p}]$ select the persistent and transient blocks. For two states and controls,

$$P_{\mathrm{r}}(F(z,c) - F(z',c')) = A_{\mathrm{r}} P_{\mathrm{r}}(z - z') + W_c(c - c'),$$

where A_{r} is the transient diagonal block. The triangle inequality gives Equation 4. Since the additive control has a zero persistent prefix and $P_{\mathrm{p}} A = P_{\mathrm{p}}$, the persistent equality follows as well. The entries $r_k = e^{-\operatorname{softplus}(a_k)} \in (0,1)$ imply $\|A_{\mathrm{r}}\|_2 = \rho = \max_k r_k < 1$.

For two trajectories driven by control sequences $c_t, c_{t+1}, \ldots$ and $c'_t, c'_{t+1}, \ldots$, repeated substitution yields the input-to-state bound

$$\left\|P_{\mathrm{r}}(z_{t+n} - z'_{t+n})\right\|_2 \le \rho^n \left\|P_{\mathrm{r}}(z_t - z'_t)\right\|_2 + \|W_c\|_2 \sum_{k=0}^{n-1} \rho^{n-1-k} \|c_{t+k} - c'_{t+k}\|_2.$$

In particular, if all control discrepancies are at most δ_c, the second term is bounded by $\|W_c\|_2 \delta_c / (1 - \rho)$, while a shared control sequence recovers geometric contraction of the initial transient discrepancy. Thus the candidate-discovery state forgets transient initialization error and has finite gain to event perturbations.

A.2 Semantic Residual as an Orthogonal Projection

Let $P_b = \bar{b}_t \bar{b}_t^\top / \left\|\bar{b}_t\right\|_2^2$ be the orthogonal projector onto $\operatorname{span}(\bar{b}_t)$. Equation 8 is $\eta_t = (I - P_b)\bar{\lambda}_t$. Therefore

$$\langle \eta_t, \bar{b}_t \rangle = \langle \bar{\lambda}_t, \bar{b}_t \rangle - \frac{\langle \bar{\lambda}_t, \bar{b}_t \rangle}{\left\|\bar{b}_t\right\|_2^2} \langle \bar{b}_t, \bar{b}_t \rangle = 0.$$

For any scalar c, the Pythagorean identity yields

$$\left\|\bar{\lambda}_t - c\bar{b}_t\right\|_2^2 = \|\eta_t\|_2^2 + \left(c - \frac{\langle \bar{\lambda}_t, \bar{b}_t \rangle}{\left\|\bar{b}_t\right\|_2^2}\right)^2 \left\|\bar{b}_t\right\|_2^2.$$

The projection coefficient in Equation 8 uniquely minimizes the residual norm. The language score used by the outer frontier is therefore exactly the component left after the best local linear explanation by the anchor. Moreover, for any $c \in \mathbb{R}$,

$$(I - P_b)(\bar{\lambda}_t + c\bar{b}_t) = (I - P_b)\bar{\lambda}_t = \eta_t,$$

because $(I - P_b)\bar{b}_t = 0$. This proves the anchor-aligned invariance in Proposition 2: semantic priority depends on innovation, not on how strongly language repeats the existing anchor direction.

A.3 Proof of Cutoff Decoupling

Let $S_1 = \mathcal{T}_{K_1}(s_t^{(1)})$. The outer operator changes only candidates $j \in \mathcal{P}_t \subseteq \mathcal{F}_t$, all of which lie outside S_1. For every such candidate, Equation 9 gives $s_t^f(j) \le \beta_{1,t}^{(1)} - \varepsilon_t < \beta_{1,t}^{(1)}$. Scores of items in S_1 are unchanged and are at least $\beta_{1,t}^{(1)}$. Thus no modified outer item can displace an item of S_1, proving $\mathcal{T}_{K_1}(s_t^f) = S_1$. Since every score inside S_1 is unchanged, their pairwise order is also unchanged, establishing the stronger ordered-list equality $\mathcal{R}_{K_1}(s_t^f) = \mathcal{R}_{K_1}(s_t^{(1)})$.

For the second statement, let $S_2 = \mathcal{T}_{K_2}(s_t^{(1)})$ and $S_2^f = \mathcal{T}_{K_2}(s_t^f)$. Only items in $\mathcal{P}_t$ can increase their scores. Therefore every newly admitted item belongs to $\mathcal{P}_t \setminus S_2$, so $|S_2^f \setminus S_2| \le |\mathcal{P}_t| \le m_2$. Since both sets have cardinality K_2, the number of displaced items equals the number of admitted items. Consequently

$$|S_2^f \triangle S_2| = 2|S_2^f \setminus S_2| \le 2m_2.$$

The same argument establishes score-support locality and completes the proof of Theorem 1.

A.4 Composition over Multiple Frontiers

The barrier principle extends beyond two benchmark cutoffs. Let $K_1 < \cdots < K_L$, let $s^{(1)}$ be the score field after completing the first frontier, and for each $\ell = 2, \dots, L$ let stage ℓ modify at most m_ℓ candidates outside $\mathcal{T}_{K_{\ell-1}}(s^{(\ell-1)})$ while capping each modified score strictly below the $K_{\ell-1}$ boundary. Applying the proof of Theorem 1 at every stage gives, by induction,

$$\mathcal{R}_{K_r}(s^{(\ell)}) = \mathcal{R}_{K_r}(s^{(\ell-1)}) \quad \text{for every } r < \ell, \qquad \left|\mathcal{T}_{K_\ell}(s^{(\ell)}) \triangle \mathcal{T}_{K_\ell}(s^{(\ell-1)})\right| \le 2m_\ell.$$

Thus later evidence sources cannot revise any completed earlier prefix, and each frontier has its own discrete intervention budget. This compositional statement is the theoretical counterpart of the $(5, 10)$, $(10, 20)$, and $(20, 50)$ transfer results in Table 6.

For the two-frontier instance used by LSRB-Rec, suppose the primary operator admits at most m_1 forecast-only items relative to the frozen anchor and leaves every other item's membership unchanged at its active cutoff. Each admission displaces exactly one member, so the primary symmetric difference is at most $2m_1$. The outer operator adds at most m_2 further admissions at K_2 and none at K_1. By the triangle inequality for symmetric difference,

$$\left|\mathcal{T}_{K_1}(s_t^f) \triangle \mathcal{T}_{K_1}(b_t)\right| \le 2m_1, \qquad \left|\mathcal{T}_{K_2}(s_t^f) \triangle \mathcal{T}_{K_2}(b_t)\right| \le 2(m_1 + m_2). \tag{12}$$

The first budget is spent only by collaborative forecast evidence; language cannot contribute to it. At the deeper cutoff, every language-induced change can be traced to one of the selected elements of $\mathcal{P}_t$. This differs from a norm penalty on a dense score residual: a small global norm can still reorder many near-tied items, whereas a support and admission budget directly controls discrete list membership.

The construction separates three notions that are often conflated. Latent contraction regulates how forecasting error propagates; residual projection regulates which semantic direction is eligible; and the barrier regulates which ranking decision may change. None substitutes for another. Contraction alone does not determine candidate membership, semantic orthogonality alone does not prevent a large score crossing, and a barrier alone cannot discover the right candidate. Their composition produces a stable and effective adapter. It also makes the model's behavior inspectable at inference: one can record the dynamic-only proposals, the residual-semantic priority, the boundary crossed, and the exact intervention budget for each recommendation list within the explicitly bounded score operator.

B Reproducibility Specification

B.1 Datasets and Evaluation Protocol

Table 7: 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

Each user sequence is ordered chronologically. The final two interactions form validation and test targets, and all earlier interactions form the training prefix. Evaluation ranks one target against every eligible catalog item. Every item observed before the target is masked, including interactions outside a truncated encoder window. HR@K is the fraction of targets in the first K positions, NDCG@K discounts a hit by its rank, and MRR uses the reciprocal full-catalog rank. The same mapping, split, masking, and metric code is used by the matched anchor and LSRB-Rec.

B.2 Model and Ranking Configuration

The anchor is a frozen MPT sequence encoder with its matched semantic–collaborative item representation. The learned adapter uses a 64-dimensional latent state with a 16-dimensional persistent block. Its transport combines an identity persistent operator, a diagonal contractive transient operator, and an additive drift driven by the normalized collaborative observable. Language embeddings are precomputed offline from the item-side text records, kept frozen, and normalized when used in Equation 7; no held-out interaction target enters their construction. The default outer source band is anchor ranks 21–40, and its admissible destination is ranks 11–20 below the completed Top-10 prefix.

Table 8: Fixed ranking profiles used in the reported evaluation.

Dataset	Dynamic K	Inner promotions	Inner scale	Outer promotions	Semantic gain
Beauty	10	1	0.5000	10	16
Toys	15	1	0.1250	10	16
Sports	10	2	0.0000	3	16
Yelp	10	2	0.0000	3	16
OnlineRetail	15	2	0.2500	1	16

All dataset-specific ranking hyperparameters in Table 8 were selected on the validation set and fixed before test evaluation.

Candidate memberships are computed from detached anchor and forecast scores after complete-history masking. Within a nested dynamic profile, forecast-only candidates are partitioned by consecutive anchor boundaries and the promotion branch of Equation 6 is applied independently with the listed budgets; this is the two-boundary extension of the same primary operator. For the outer stage, $K^{o}_{\mathrm{tr}} = 40$ in the default $(10, 20)$ profile, $\kappa_o = 0.01$, the standard-deviation floor is 10^{-5}, and the final ceiling uses the immediately smaller floating-point value when needed to maintain strict order. Catalog dot products are evaluated in full precision after the frozen encoder pass, and deterministic item indices resolve exact score ties.

B.3 Optimization and Repeated Runs

We run each reported configuration three times with seeds 42, 43, and 44. Trainable adapter parameters are optimized with AdamW using learning rate 10^{-3}, weight decay 0.1, and batch size 256. Training lasts at most 30 epochs and stops when validation NDCG@10 fails to improve for 10 consecutive epochs; the checkpoint with the highest validation NDCG@10 is retained. For the objective in Equation 11, let $\mathrm{rank}^{b}_{t}(j)$ and $\mathrm{rank}^{d}_{t}(j)$ denote the one-indexed positions of item j in

the anchor and forecast candidate lists, respectively. The exact negative weight is

$$w_t(j) = \begin{cases} \dfrac{\chi(\mathrm{rank}_t^b(j))}{\log_2(\mathrm{rank}_t^b(j)+1)}, & j \in \mathcal{C}_t^b \setminus \{y_t\}, \\ \dfrac{0.5}{\log_2(\mathrm{rank}_t^d(j)+1)}, & j \in \mathcal{C}_t^d \setminus (\mathcal{C}_t^b \cup \{y_t\}), \\ 0, & \text{otherwise}, \end{cases} \qquad \chi(r) = \begin{cases} 2, & r \le 5, \\ 1, & 5 < r \le 10, \\ 0.25, & r > 10. \end{cases}$$

Thus an item present in both lists receives its anchor-rank weight, while the positive target is excluded from the negative sum.

C Complete Matched-Anchor Results

Table 9: Complete strict full-catalog test results (mean ± sample standard deviation).

Dataset	Variant	HR@5	NDCG@5	HR@10	NDCG@10	HR@20	NDCG@20	MRR
Beauty	Anchor	0.0803±0.0016	0.0569±0.0009	0.1148±0.0042	0.0680±0.0017	0.1581±0.0037	0.0789±0.0015	0.0610±0.0010
	LSRB-Rec	**0.0804±0.0019**	**0.0570±0.0011**	**0.1181±0.0044**	**0.0690±0.0018**	**0.1834±0.0040**	**0.0848±0.0017**	**0.0614±0.0011**
Toys	Anchor	0.0898±0.0002	0.0624±0.0002	0.1246±0.0012	0.0736±0.0004	0.1660±0.0016	0.0841±0.0004	0.0651±0.0002
	LSRB-Rec	**0.0898±0.0003**	**0.0625±0.0002**	**0.1303±0.0016**	**0.0754±0.0003**	**0.1912±0.0017**	**0.0900±0.0003**	**0.0656±0.0002**
Sports	Anchor	0.0442±0.0008	0.0302±0.0005	0.0656±0.0006	0.0371±0.0002	0.0937±0.0014	0.0442±0.0003	0.0340±0.0003
	LSRB-Rec	**0.0465±0.0005**	**0.0311±0.0004**	**0.0694±0.0005**	**0.0383±0.0001**	**0.0990±0.0013**	**0.0456±0.0002**	**0.0342±0.0002**
Yelp	Anchor	0.0369±0.0008	0.0236±0.0006	0.0604±0.0008	0.0312±0.0006	0.0960±0.0015	0.0401±0.0008	0.0299±0.0006
	LSRB-Rec	**0.0384±0.0017**	**0.0242±0.0010**	**0.0644±0.0002**	**0.0324±0.0005**	**0.1018±0.0016**	**0.0416±0.0008**	**0.0300±0.0006**
OnlineRetail	Anchor	0.1761±0.0008	0.1225±0.0009	0.2459±0.0006	0.1450±0.0005	0.3309±0.0018	0.1664±0.0007	0.1275±0.0007
	LSRB-Rec	**0.1764±0.0008**	**0.1226±0.0009**	**0.2459±0.0007**	**0.1450±0.0005**	**0.3310±0.0017**	**0.1665±0.0006**	**0.1276±0.0007**

The complete results reinforce the cutoff-specific account in the main text. Beauty and Toys concentrate their largest increases in Top-20 membership, while Sports and Yelp benefit across the entire displayed list. When isolated, the outer barrier leaves the first four metrics unchanged because its support begins below the Top-10 boundary; MRR also captures the improved deeper positions of targets that remain outside Top-10. OnlineRetail preserves the positive direction for every metric. These patterns align the location of each gain with the decision frontier that produces it.

D Statistical Tests

Table 10: Metrics significant after Holm correction at 0.05.

Dataset	Significant matched-anchor improvements
Beauty	HR@10, NDCG@10, HR@20, NDCG@20, MRR
Toys	HR@10, NDCG@10, HR@20, NDCG@20, MRR
Sports	HR@5, NDCG@5, HR@10, NDCG@10, HR@20, NDCG@20, MRR
Yelp	HR@5, NDCG@5, HR@10, NDCG@10, HR@20, NDCG@20, MRR
OnlineRetail	NDCG@20, MRR

The tests use paired user-level score differences and one-sided randomization, followed by Holm correction over the seven metrics within each benchmark. Bootstrap percentile intervals provide an additional effect-size check. The concentration of significance at ranks 10 and 20 on Beauty and Toys matches the adapter's two decision frontiers. Sports and Yelp show sufficiently broad per-user improvements for every metric to remain significant after correction.

E CONTROLLED MECHANISM TESTS

E.1 SEMANTIC RESIDUALIZATION

Table 11: Semantic residualization under a fixed candidate union and barrier.

Dataset	Variant	HR@10	NDCG@10	HR@20	NDCG@20	MRR
Beauty	No semantic	0.1224	0.0710	0.1846	0.0860	0.0626
	Raw semantic	0.1224	0.0710	0.1870	0.0866	0.0626
	Standardized semantic	0.1224	0.0710	0.1877	0.0867	0.0627
	LSRB-REC orthogonal residual	**0.1224**	**0.0710**	**0.1885**	**0.0870**	**0.0628**
Toys	No semantic	0.1316	0.0758	0.1883	0.0894	0.0656
	Raw semantic	0.1316	0.0758	0.1919	0.0903	0.0656
	Standardized semantic	0.1316	0.0758	0.1935	0.0906	0.0657
	LSRB-REC orthogonal residual	**0.1316**	**0.0758**	**0.1942**	**0.0909**	**0.0658**
Yelp	No semantic	0.0642	0.0320	0.1010	0.0411	0.0296
	Raw semantic	0.0642	0.0320	0.1013	0.0411	0.0296
	Standardized semantic	0.0642	0.0320	0.1016	0.0412	0.0296
	LSRB-REC orthogonal residual	**0.0642**	**0.0320**	**0.1020–0.1023**	**0.0413–0.0414**	**0.0296–0.0297**

All three semantic constructions improve the outer frontier over removing semantics, and the orthogonal residual is strongest on every reported outer metric across Beauty, Toys, and Yelp. It raises HR@20 by 0.0039, 0.0059, and 0.0010–0.0013, respectively, while preserving every Top-10 metric exactly. This ranking advantage follows the information contract in Proposition 2: useful semantic variation remains, whereas the component already explained by the anchor is mathematically removed.

E.2 HARD BARRIER VERSUS UNCONSTRAINED INTERVENTION

Table 12: Outer intervention with no, soft, and hard barriers.

Dataset	Variant	HR@10	NDCG@10	HR@20	NDCG@20	MRR	I_{10}	Avg. $\triangle_{10}$
Beauty	No semantic intervention	0.1224	0.0710	0.1846	0.0860	0.0626	0.0000	0.0000
	No barrier	0.1206	0.0704	0.1634	0.0811	0.0624	0.5599	1.8535
	Soft penalty	0.1209	0.0706	0.1634	0.0812	0.0624	0.5461	1.2850
	LSRB-REC hard barrier	**0.1224**	**0.0710**	**0.1885**	**0.0870**	**0.0628**	**0.0000**	**0.0000**
Toys	No semantic intervention	0.1316	0.0758	0.1883	0.0894	0.0656	0.0000	0.0000
	No barrier	0.1287	0.0748	0.1691	0.0849	0.0654	0.5962	2.0386
	Soft penalty	0.1287	0.0749	0.1691	0.0851	0.0655	0.5847	1.3911
	LSRB-REC hard barrier	**0.1316**	**0.0758**	**0.1942**	**0.0909**	**0.0658**	**0.0000**	**0.0000**
Yelp	No semantic intervention	0.0642	0.0320	0.1010	0.0411	0.0296	0.0000	0.0000
	No barrier	0.0619	0.0314	0.0957	0.0398	0.0295	0.4400	1.0513
	Soft penalty	0.0621	0.0314	0.0957	0.0398	0.0295	0.4250	0.8795
	LSRB-REC hard barrier	**0.0642**	**0.0320**	**0.1020–0.1023**	**0.0413–0.0414**	**0.0296–0.0297**	**0.0000**	**0.0000**

The hard barrier is the only intervention that attains exact list-level non-interference. It also yields the strongest HR@20 and NDCG@20 on all three datasets. A soft penalty reduces the average symmetric difference but still changes 42.50–58.47% of the completed primary Top-10 lists, demonstrating why a continuous score regularizer does not imply discrete rank preservation in a near-tied catalog.

E.3 CUTOFF GENERALIZATION

Table 13: Outer-stage cutoff transfer relative to the completed primary ranking.

Dataset	(K_1, K_2)	I_{K_1}	HR@K_2: primary→final	NDCG@K_2: primary→final
Beauty	$(5, 10)$	**0.0000**	0.1224→**0.1364**	0.0710→**0.0751**
	$(10, 20)$	**0.0000**	0.1615→**0.1873**	0.0807→**0.0866**
	$(20, 50)$	**0.0000**	0.2314→**0.2496**	0.0946→**0.0978**
Toys	$(5, 10)$	**0.0000**	0.1316→**0.1443**	0.0758→**0.0795**
	$(10, 20)$	**0.0000**	0.1676→**0.1933**	0.0847→**0.0906**
	$(20, 50)$	**0.0000**	0.2349→**0.2518**	0.0980→**0.1010**
Yelp	$(5, 10)$	**0.0000**	0.0642→**0.0679**	0.0320→**0.0331**
	$(10, 20)$	**0.0000**	0.0952→**0.1018**	0.0397→**0.0412**
	$(20, 50)$	**0.0000**	0.1688→**0.1720**	0.0542→**0.0548**

Every tested pair leaves the completed primary Top-K_1 set unchanged and improves both metrics at the secondary cutoff. The before→after cells make the reference explicit: the left value is produced by the complete primary ranking $s_t^{(1)}$, and the highlighted right value is obtained only after the outer intervention. Membership and internal order at K_1 are preserved query by query, so its HR and NDCG equality follows directly from the barrier construction. The result covers shallow, default, and broad list regimes and matches the multi-frontier composition argument in Appendix A.4.

E.4 Dynamic Forecast Alternatives

Table 14: Forecast alternatives with the same candidate union and ranking barrier.

Dataset	Forecast	HR@10	NDCG@10	HR@20	NDCG@20	MRR
Beauty	Frozen anchor	0.1190	0.0699	0.1615	0.0806	0.0621
	Linear	0.1223	0.0708	0.1868	0.0864	0.0625
	MLP	0.1225	0.0709	0.1871	0.0865	0.0625
	LSRB-Rec contractive	**0.1228**	**0.0711**	**0.1878**	**0.0868**	**0.0627**
Sports	Frozen anchor	0.0649	0.0370	0.0950	0.0445	0.0341
	Linear	0.0687	0.0380	0.1007	0.0460	0.0342
	MLP	0.0690	0.0381	0.1010	0.0460	0.0342
	LSRB-Rec contractive	**0.0693**	**0.0383**	**0.1014**	**0.0462**	**0.0344**
Yelp	Frozen anchor	0.0598	0.0307	0.0952	0.0396	0.0295
	Linear	0.0639	0.0319	0.1014	0.0412	0.0296
	MLP	0.0642	0.0320	0.1016	0.0412	0.0296
	LSRB-Rec contractive	**0.0645**	**0.0321**	**0.1020**	**0.0414**	**0.0297**

All three learned forecasts improve every displayed metric over the frozen anchor under the same downstream interface. The contractive transition is strongest in all 15 dataset–metric cells, establishing that the stability constraint also improves the quality of dynamic retrieval. Table 15 then tests the guarantee directly under controlled perturbation.

Table 15: Forecast robustness on Beauty under the same controlled input perturbation.

Forecast	HR@20	NDCG@20	Clean→Perturbed ΔHR@20	State amplification	Stability guarantee
Linear	0.1871	0.0865	−0.0068	1.14	×
MLP	0.1873	0.0866	−0.0059	1.09	×
LSRB-Rec contractive	**0.1875**	**0.0867**	−0.0018	**0.71**	✓

The linear and MLP alternatives amplify the perturbed state by factors of 1.14 and 1.09 and lose 0.0068 and 0.0059 HR@20. The contractive forecast attenuates the state discrepancy to 0.71 and limits the loss to 0.0018, directly matching the finite-gain behavior established by Proposition 1. Stability therefore governs both latent propagation and downstream ranking robustness.

E.5 Ranking-Budget Sensitivity

Table 16: Sensitivity to the outer promotion budget and semantic weight.

Dataset	Quantity	Value	HR@20	NDCG@20	I_{10}
Beauty	Outer promotions	1	0.1621	0.0809	**0.0000**
	Outer promotions	3	0.1696	0.0826	**0.0000**
	Outer promotions	5	0.1755	0.0839	**0.0000**
	Outer promotions	10	0.1873	0.0866	**0.0000**
	Semantic weight	4	0.1854	0.0862	**0.0000**
	Semantic weight	8	0.1864	0.0864	**0.0000**
	Semantic weight	16	0.1873	0.0866	**0.0000**
	Semantic weight	32	0.1883	0.0868	**0.0000**
Yelp	Outer promotions	1	0.0957	0.0398	**0.0000**
	Outer promotions	3	0.1018	0.0412	**0.0000**
	Outer promotions	5	0.1072	0.0425	**0.0000**
	Outer promotions	10	0.1197	0.0453	**0.0000**
	Semantic weight	4	0.1015	0.0412	**0.0000**
	Semantic weight	8	0.1016	0.0412	**0.0000**
	Semantic weight	16	0.1018	0.0412	**0.0000**
	Semantic weight	32	0.1019	0.0412	**0.0000**

The protected cutoff is invariant throughout all sixteen settings. Increasing the promotion budget smoothly expands the amount of admissible outer correction, whereas a wide semantic-weight

range changes deeper metrics only gradually. The barrier therefore separates safety from tuning: these parameters determine how much utility is extracted beyond the boundary, but not whether the completed Top-10 list is preserved.

F Efficiency Details

Table 17: Measured evaluation efficiency and recommendation-model size.

Dataset	Anchor users/s	Full users/s	Anchor memory (MB)	Full memory (MB)	Trainable params	Rec. params
Beauty	36,284	15,875	195	533	1.64M	3.81M
Toys	32,590	12,688	194	528	1.63M	3.80M
Sports	36,029	13,766	230	876	2.04M	4.21M
Yelp	36,115	11,267	465	2,080	3.79M	5.96M
OnlineRetail	38,401	25,377	272	401	1.09M	3.26M

The adapter adds one collaborative full-catalog dot-product pass; both barriers then operate on short candidate sets. Throughput remains above 11,000 users per second on every benchmark and exceeds 25,000 users per second on OnlineRetail. Memory scales with the catalog-sized score tensors, explaining the larger footprint on Yelp. Parameter counts include the recommendation model and adapter but exclude the external language encoder, which is used offline; the frozen item-language vectors require storage but no gradient memory.