# ATLASREC: LANGUAGE-GROUNDED MULTI-HORIZON SUCCESSOR DYNAMICS FOR SEQUENTIAL RECOMMENDATION

**Anonymous authors**
Paper under double-blind review

## ABSTRACT

Language-enhanced sequential recommenders improve how items and histories are represented, but next-item prediction also requires direction: knowing what a sequence means is not the same as knowing where it is going. We formulate this task as a *state–dynamics–decision* problem and introduce ATLASREC. Frozen language representations estimate a recency-weighted semantic state; training interactions estimate sparse, popularity-relative successor dynamics; and a pretrained recommender provides the behavioral decision reference. An event $r$ positions back queries the $(r+1)$-step operator, so every source is interpreted at its exact distance from the target. We show that repeating a one-step operator assumes horizon-invariant dynamics and otherwise incurs a discounted conditional-KL mismatch. A minimum-information closure incorporates the two evidence fields while deviating minimally from the behavioral prior, and finite-sample successor uncertainty yields an explicit robust ranking condition. On three same-setting Amazon benchmarks, ATLASREC is best in 11 of 12 comparisons with prior methods. Against its exact anchor, it improves all 21 full-catalog metrics, with relative gains up to 18.64%. Component, horizon, state, popularity, encoder-transfer, and efficiency results validate the complete formulation.

## 1 INTRODUCTION

Sequential recommendation predicts the next interaction from an ordered history. Recurrent models summarize histories through event-wise state updates (Hidasi et al., 2016). Self-attentive models expose longer behavioral dependencies (Kang & McAuley, 2018). Language-pretrained recommenders add transferable item meaning beyond interaction identifiers (Li et al., 2023a). Most language-enhanced methods therefore strengthen the representation of an item or sequence. Yet semantic representation answers what the current history means, whereas recommendation must predict how that state evolves. Semantic similarity is compatible in either direction; the next interaction is not. *Semantics describes the current state, but recommendation requires its dynamics.*

This distinction becomes concrete once the prediction horizon is made explicit. For the target $i_{t+1}$, the latest event $i_t$ is one step away, $i_{t-1}$ is two steps away, and $i_{t-r}$ is $r+1$ steps away. Applying the same one-step relation to every position asks what immediately follows each historical item, which is the wrong temporal question for older evidence. The correct relation is distance dependent: what tends to occur $r+1$ steps after this source? Language geometry can locate the semantic region compatible with the current intent, interaction data can identify direction within that region, and a pretrained sequential model can supply the behavioral reference that resolves the final order. Accordingly, we formulate language-enhanced sequential recommendation as a *state–dynamics–decision* problem: language estimates the current state, horizon-specific successor statistics characterize its evolution, and the behavioral prior determines how that evidence changes the recommendation.

We propose ATLASREC, a language-grounded atlas of successor dynamics around a dataset-adapted MPT behavioral prior (Xu et al., 2026a). A frozen language encoder estimates semantic state from recent events. A family of popularity-relative successor operators describes one-step through multi-step evolution, with each historical event querying the operator indexed by its exact distance from the target. Repeating the one-step law is exact only under horizon invariance; otherwise its expected



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.

mismatch is a discounted sum of conditional KL divergences. A minimum-information closure then incorporates semantic and successor evidence by making the smallest KL-geometric update of the behavioral prior, while a finite-sample bound carries successor uncertainty into robust ranking conditions. The empirical design mirrors this decomposition through matched-anchor, component, horizon, state, popularity, cross-domain, encoder-transfer, sensitivity, and efficiency evaluations.

Our contributions are:

- We formulate language-enhanced sequential recommendation as state–dynamics–decision, separating semantic state, directional evolution, and behavioral decision.
- We construct target-distance-aligned successor operators, so each historical event queries the law for its exact distance from the target.
- We characterize one-step replacement by a discounted conditional-KL mismatch and derive minimum-information closure, finite-sample concentration, and robust ranking guarantees.
- Across five datasets and four language encoders, ATLASREC improves all 21 matched-anchor metrics and leads 11 of 12 state-of-the-art metrics under the shared Amazon setting.

## 2 RELATED WORK

Sequential recommendation began with Markov factorization (Rendle et al., 2010). Recurrent networks introduced learned session states (Hidasi et al., 2016). Convolutional encoders captured local high-order patterns (Tang & Wang, 2018). SASRec brought causal self-attention to item sequences (Kang & McAuley, 2018). BERT4Rec used bidirectional masking to enrich sequential context (Sun et al., 2019). Contrastive objectives improve state discrimination (Xie et al., 2022). Frequency-selective MLPs provide an alternative to attention (Zhou et al., 2022). HSTU scales sequential transduction to industrial regimes (Zhai et al., 2024). Recent wavelet models resolve behavior over multiple temporal bands (Xu et al., 2026b). WPGRec couples scale-aligned graph propagation with a stationary wavelet packet tree (Liu et al., 2026). MPT instead pretrains a transferable sequence prior on synthetic Markov chains (Xu et al., 2026a). ATLASREC complements these encoders by estimating an explicit, training-only successor field around their behavioral scores.

Language-aware recommenders transfer semantic structure in several ways. UniSRec aligns transferable sequence representations across domains (Hou et al., 2022). RecFormer places item text inside a language-pretrained sequence model (Li et al., 2023a). E4SRec adapts language representations through an efficient sequential interface (Li et al., 2023b). LLM2Rec treats LLMs as item embedding models (He et al., 2025). DIFF filters and fuses dual sources of side information (Kim et al., 2025). ACE controls anisotropy before adapting LLM embeddings to a recommender (Lee et al., 2026). Classical pointwise mutual information measures association relative to a marginal reference (Church & Hanks, 1990). Existing language-enhanced methods primarily enrich item or history representations, whereas transition methods primarily model behavioral association. ATLASREC connects these views: language estimates the current semantic state, horizon-indexed interaction statistics estimate its directional evolution, and the behavioral prior closes both into the next-event decision.

## 3 METHOD

### 3.1 STATE–DYNAMICS–DECISION FORMULATION

Let $\mathcal{I} = \{1, \dots, M\}$ be the item catalog and $i_{1:t} = (i_1, \dots, i_t)$ the observed prefix. A pretrained sequential anchor returns a normalized state $\mathbf{h}_t \in \mathbb{R}^{d_c}$ and normalized collaborative item vectors $\mathbf{e}_j \in \mathbb{R}^{d_c}$. Its logit and distribution are

$$b_t(j) = \langle \mathbf{h}_t, \mathbf{e}_j \rangle / \tau_b, \qquad \boldsymbol{\pi}_t^b(j) = \frac{\exp b_t(j)}{\sum_{k \in \mathcal{I}} \exp b_t(k)}, \tag{1}$$

where $\tau_b > 0$ is the anchor temperature and $\boldsymbol{\pi}_t^b \in \mathbb{R}^M$ is the behavioral prior over the catalog. ATLASREC adds a semantic potential $a_t : \mathcal{I} \to \mathbb{R}$ for current-state compatibility and a target-distance-aligned successor potential $d_t : \mathcal{I} \to \mathbb{R}$ for future direction; their minimum-information closure changes the behavioral order only when the combined evidence exceeds its anchor margin (Figure 1).

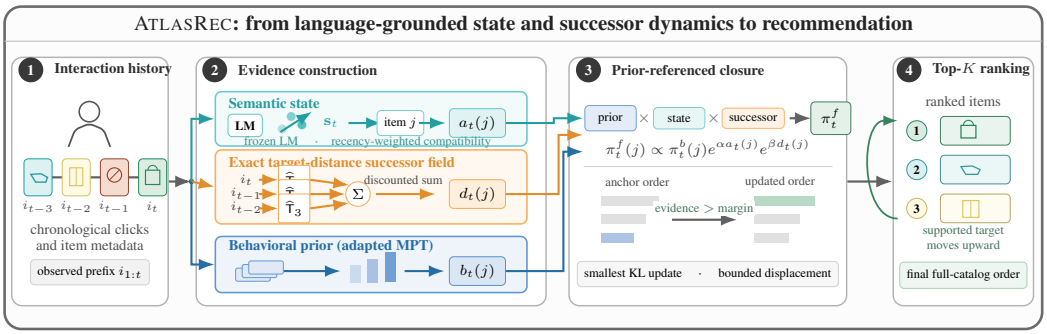



Figure 1: The ATLASREC pathway. The observed prefix produces a language-grounded semantic state, an exactly horizon-aligned successor field, and a behavioral prior. Minimum-information closure multiplies the two evidence experts into the prior, changing only candidate pairs whose combined evidence exceeds the anchor margin.

### 3.2 LANGUAGE-GROUNDED SEMANTIC STATE

Let $x_j$ denote the metadata text associated with item $j$. A frozen language encoder produces $\mathbf{u}_j = \mathrm{Enc}_{\mathrm{LM}}(x_j) \in \mathbb{R}^{d_0}$. For the encoders considered here, $d_0 \geq d_s$; we map once into the semantic space using a fixed row-orthonormal projection $\mathbf{W}_{\mathrm{sem}} \in \mathbb{R}^{d_s \times d_0}$ and normalize:

$$\mathbf{v}_j = \frac{\mathbf{W}_{\mathrm{sem}} \mathbf{u}_j}{\max\{\|\mathbf{W}_{\mathrm{sem}} \mathbf{u}_j\|_2, \varepsilon_v\}} \in \mathbb{R}^{d_s}, \tag{2}$$

where $\varepsilon_v > 0$ prevents division by zero. Both the encoder and $\mathbf{W}_{\mathrm{sem}}$ remain fixed.

The last item alone is a local observation of intent, whereas an unweighted history average treats stale and recent evidence identically. For a history depth $H \geq 1$ and semantic decay $\kappa \in (0, 1]$, let $H_t = \min\{H, t\}$. We define the language-grounded state and candidate potential as

$$\mathbf{s}_t = \frac{\sum_{r=0}^{H_t-1} \kappa^r \mathbf{v}_{i_{t-r}}}{\max\left\{\left\|\sum_{r=0}^{H_t-1} \kappa^r \mathbf{v}_{i_{t-r}}\right\|_2, \varepsilon_s\right\}}, \qquad a_t(j) = \langle \mathbf{s}_t, \mathbf{v}_j \rangle. \tag{3}$$

Here $\varepsilon_s > 0$ is a numerical safeguard. Thus $\mathbf{s}_t \in \mathbb{R}^{d_s}$ is a moving state in a frozen language geometry, $\|\mathbf{s}_t\|_2 \leq 1$, and $a_t(j) \in [-1, 1]$. It is computed from observed events only. Equation 3 is deliberately a lightweight state estimator: language locates the candidate region compatible with current intent, while the successor operators below determine the behaviorally supported direction within that region.

**Proposition 1** (Temporal semantic consensus)**.** *Let* $\mathbf{m}_t = \sum_{r=0}^{H_t-1} \kappa^r \mathbf{v}_{i_{t-r}}$ *and suppose* $\|\mathbf{m}_t\|_2 \geq \varepsilon_s$. *Then* $\mathbf{s}_t$ *is the maximizer*

$$\mathbf{s}_t = \arg\max_{\|\mathbf{s}\|_2=1} \sum_{r=0}^{H_t-1} \kappa^r \langle \mathbf{s}, \mathbf{v}_{i_{t-r}} \rangle,$$

*equivalently the unit direction minimizing the same recency-weighted sum of squared semantic distances. The last-item state is the boundary case* $H = 1$, *while uniform history is the boundary case* $\kappa = 1$.

Proposition 1 makes the temporal role of language explicit: the state is the maximum-agreement direction of the observed semantic trajectory, with recent events contributing more to the direction rather than merely changing a downstream score coefficient.

### 3.3 POPULARITY-RELATIVE SUCCESSOR OPERATORS

Semantic compatibility is non-directional. We obtain direction from training sequences. Let $H_{\mathrm{op}}$ be the largest constructed transition horizon and let $n \in \{1, \ldots, H_{\mathrm{op}}\}$. For items $i, j \in \mathcal{I}$, $C_n(i, j)$ counts training occurrences in which $j$ appears exactly $n$ events after $i$, and

$$C_n^{\mathrm{src}}(i) = \sum_{k \in \mathcal{I}} C_n(i, k)$$

is the number of observed $n$-step successors of source $i$. The implementation uses one shared popularity reference: $C^{\mathrm{pop}}(j)$ is the add-one-smoothed training frequency with which $j$ occurs as the next event, and

$$p(j) = \frac{C^{\mathrm{pop}}(j)}{\sum_{k\in\mathcal{I}} C^{\mathrm{pop}}(k)}.$$

For $C_n^{\mathrm{src}}(i) > 0$, define $P_n(j \mid i) = C_n(i,j)/C_n^{\mathrm{src}}(i)$; a source with no $n$-step observation has an all-zero row. Given a cap $L > 0$, the popularity-relative positive transition information is

$$\tau_n(i,j) = \begin{cases} \mathrm{clip}_{[0,L]}\left(\log \frac{P_n(j|i)}{p(j)}\right), & C_n(i,j) > 0, \\ 0, & C_n(i,j) = 0, \end{cases} \qquad \mathrm{clip}_{[0,L]}(x) = \min\{L, \max\{0, x\}\}. \quad (4)$$

The piecewise definition avoids evaluating $\log 0$. It also distinguishes a missing transition from negative evidence: both receive zero positive lift.

The actual operator retains at most $K_s$ observed successors per source. Let $\mathcal{S}_n(i) \subseteq \mathcal{I}$ contain the $K_s$ candidates with largest $C_n(i,j)$, or all observed candidates when fewer are available. For $j \in \mathcal{S}_n(i)$, let $q_n(i,j) \in \{1, \ldots, |\mathcal{S}_n(i)|\}$ be its descending count rank. A support-preserving rank-based reliability modulation gives the stored weight

$$\widehat{\tau}_n(i,j) = \begin{cases} (1-\mu)\tau_n(i,j) + \mu \dfrac{L}{\log_2(q_n(i,j)+1)}, & j \in \mathcal{S}_n(i), \\ 0, & j \notin \mathcal{S}_n(i), \end{cases} \qquad \widehat{\mathsf{T}}_n = [\widehat{\tau}_n(i,j)]_{i,j=1}^{M}, \quad (5)$$

where $\mu \in [0,1)$ mixes popularity-relative lift with the empirical support-rank signal. Hence $\widehat{\mathsf{T}}_n \in \mathbb{R}_{\geq 0}^{M\times M}$ has at most $MK_s$ nonzeros and $0 \leq \widehat{\tau}_n(i,j) \leq L$. The rank term is defined only on retained observations, strengthening evidence within empirical support without creating a new source–target edge.

**Proposition 2** (Popularity-null identification)**.** *For an observed pair, $\tau_n(i,j) > 0$ exactly when $P_n(j \mid i) > p(j)$ before upper clipping. Let $u(j) = 1/M$ and define the signed, unclipped contrasts $\lambda_n^{\mathrm{pop}}(i,j) = \log(P_n(j \mid i)/p(j))$ and $\lambda_n^{\mathrm{uni}}(i,j) = \log(P_n(j \mid i)/u(j))$. Then*

$$\lambda_n^{\mathrm{uni}}(i,j) = \lambda_n^{\mathrm{pop}}(i,j) + \log\big(Mp(j)\big).$$

*Consequently, under the popularity null $P_n(\cdot \mid i) = p(\cdot)$, the popularity-relative contrast is identically zero, whereas a uniform reference retains the source-independent catalog-frequency term $\log(Mp(j))$.*

The decomposition shows what the correction removes before clipping: a globally popular candidate with $p(j) > 1/M$ receives a positive offset under a uniform reference even when the source supplies no candidate-specific transition information. Equation 4 removes that background, and Equation 5 combines source-specific excess transition information with bounded support-rank reliability on observed edges. Popularity-relative lift therefore calibrates the statistical meaning of successor evidence rather than serving as an additional preference channel.

### 3.4 Horizon-Aligned Multi-Step Successor Dynamics

The defining property of the successor field is not merely that it contains multi-step transitions, but that the source event's exact target distance selects its operator: $i_t \mapsto \widehat{\mathsf{T}}_1$, $i_{t-1} \mapsto \widehat{\mathsf{T}}_2$, and $i_{t-2} \mapsto \widehat{\mathsf{T}}_3$. A static transition feature uses the latest item or reuses one relation independently of source position. AtlasRec instead constructs a target-distance-indexed field whose statistical meaning changes with every event's temporal role. Choose $H \leq H_{\mathrm{op}}$, retain $H_t = \min\{H, t\}$ events, and let $\eta \in (0,1]$ be the dynamic decay. The scalar potential of candidate $j$ and corresponding catalog vector are

$$d_t(j) = \sum_{r=0}^{H_t-1} \eta^r \widehat{\tau}_{r+1}(i_{t-r}, j), \qquad \mathbf{d}_t = [d_t(1), \ldots, d_t(M)]^\top = \sum_{r=0}^{H_t-1} \eta^r \widehat{\mathsf{T}}_{r+1}^\top \boldsymbol{\delta}_{i_{t-r}}, \quad (6)$$

where the scalar expression uses the stored weight $\widehat{\tau}_{r+1}$ from Equation 5, and $\boldsymbol{\delta}_i \in \mathbb{R}^M$ is the one-hot vector of item $i$. Thus $d_t(j) \in \mathbb{R}$ and $\mathbf{d}_t \in \mathbb{R}^M$ are no longer overloaded. The construction is temporally valid: every $\widehat{\mathsf{T}}_n$ is estimated from training transitions, and online inference uses only the

observed prefix. Each summand identifies both the source event and the horizon supporting candidate $j$.

This construction raises the central design question: why should an older event query a different successor law rather than reuse the one-step relation? To answer it, let $P_n^\star(\cdot \mid i)$ be the population $n$-step successor law and $p^\star$ the population next-item marginal. On their common positive support define $\lambda_n^\star(i,j) = \log[P_n^\star(j \mid i)/p^\star(j)]$. Let $\widetilde{d}_t^\star$ use $\lambda_{r+1}^\star$ at history position $r$, and let $\widetilde{d}_t^{(1)}$ replace every such term by $\lambda_1^\star$.

**Theorem 1** (Horizon-mismatch characterization)**.** *For every observed prefix and candidate in the common support,*

$$\widetilde{d}_t^\star(j) - \widetilde{d}_t^{(1)}(j) = \sum_{r=1}^{H_t-1} \eta^r \log \frac{P_{r+1}^\star(j \mid i_{t-r})}{P_1^\star(j \mid i_{t-r})}. \tag{7}$$

*Moreover, if $P_{r+1}^\star(\cdot \mid i_{t-r}) \ll P_1^\star(\cdot \mid i_{t-r})$, the discounted expected componentwise cost of the repeated one-step replacement is*

$$\mathcal{M}_t = \sum_{r=1}^{H_t-1} \eta^r \mathrm{KL}\big(P_{r+1}^\star(\cdot \mid i_{t-r}) \,\|\, P_1^\star(\cdot \mid i_{t-r})\big) \geq 0.$$

*It vanishes exactly when every queried long-horizon law equals its one-step counterpart. Thus repeated one-step propagation is uniformly exact only under horizon-invariant successor dynamics.*

Theorem 1 identifies target-distance alignment as the exact way to avoid a conditional-KL mismatch whenever successor laws vary across horizons. Equations 4–5 provide its bounded, sparse empirical realization. Proposition 3 in Appendix A bounds finite-prefix truncation, Theorem 4 propagates estimation uncertainty through the same discounted field, and the one-step control in Section 4.5 changes only the operator index to test the theorem's horizon-invariance condition.

### 3.5 A Bounded Multi-Horizon Successor-Evidence Expert

The additive field in Equation 6 admits an exact multiplicative interpretation. Define the bounded evidence factor $\widehat{R}_n(i,j) = \exp\{\widehat{\tau}_n(i,j)\} \in [1, e^L]$. It equals one outside the retained transition support. Exponentiating the dynamic correction yields

$$\Gamma_t(j) := \exp\{\beta d_t(j)\} = \prod_{r=0}^{H_t-1} \widehat{R}_{r+1}(i_{t-r}, j)^{\beta\eta^r}. \tag{8}$$

Thus every historical event contributes a discounted successor-evidence expert at exactly the horizon separating it from the target. The $\tau_n$ component is a clipped positive likelihood ratio, while the rank component records bounded confidence among retained observations. Their neutral value is the multiplicative identity rather than negative evidence, and repeated horizon-aligned support compounds across the prefix. For candidates $j$ and $k$, $\Gamma_t(j)/\Gamma_t(k)$ decomposes into event-wise evidence contrasts, so every dynamic preference can be traced to specific history positions. This differs from placing a single co-occurrence vector beside the anchor: the operator family represents a path of distance-indexed evidence, while Equation 6 closes that path at the current decision.

### 3.6 Behavioral-Prior Evidence Closure

The behavioral anchor is the decision reference rather than another interchangeable feature channel. The semantic state and successor field provide bounded evidence for updating that reference. With nonnegative strengths $\alpha$ and $\beta$, the final score is

$$f_t(j) = b_t(j) + \alpha a_t(j) + \beta d_t(j), \qquad \boldsymbol{\pi}_t^f = \mathrm{softmax}(\mathbf{f}_t), \tag{9}$$

Here $\mathbf{f}_t = [f_t(1), \ldots, f_t(M)]^\top \in \mathbb{R}^M$ and $\boldsymbol{\pi}_t^f \in \mathbb{R}^M$ is the final catalog distribution. Language compatibility and successor dynamics act as log-potentials, producing the closest anchor-referenced distribution that realizes their rewards.

**Theorem 2** (Minimum-information closure)**.** *Let $\Delta^{M-1} = \{\boldsymbol{\pi} \in \mathbb{R}^M_{\geq 0} : \sum_{j\in\mathcal{I}} \boldsymbol{\pi}(j) = 1\}$ be the catalog probability simplex. The distribution in Equation 9 is the unique solution of*

$$\boldsymbol{\pi}_t^f = \arg\min_{\boldsymbol{\pi}\in\Delta^{M-1}} \left\{ \mathrm{KL}(\boldsymbol{\pi}\|\boldsymbol{\pi}_t^b) - \alpha\mathbb{E}_{j\sim\boldsymbol{\pi}}[a_t(j)] - \beta\mathbb{E}_{j\sim\boldsymbol{\pi}}[d_t(j)] \right\}, \tag{10}$$

*and equivalently*

$$\boldsymbol{\pi}_t^f(j) = \frac{\boldsymbol{\pi}_t^b(j)\exp\{\alpha a_t(j)\}\exp\{\beta d_t(j)\}}{\sum_{k\in\mathcal{I}} \boldsymbol{\pi}_t^b(k)\exp\{\alpha a_t(k)\}\exp\{\beta d_t(k)\}}.$$

The Lagrangian proof is in Appendix A. The displayed factorization has a direct functional reading: $\boldsymbol{\pi}_t^b$ is the behavioral prior, $\exp\{\alpha a_t\}$ is semantic-state evidence, and $\exp\{\beta d_t\}$ is successor evidence. The closure is the unique distribution that incorporates both evidence fields while deviating minimally from the pretrained prior in KL geometry. Normalization cancels in every pairwise decision:

$$\log\frac{\pi_t^f(i)}{\pi_t^f(j)} = \log\frac{\pi_t^b(i)}{\pi_t^b(j)} + \alpha\big(a_t(i) - a_t(j)\big) + \beta\big(d_t(i) - d_t(j)\big).$$

Thus a rank change has an exact evidence account: it occurs only when the semantic and successor contrasts jointly overcome the corresponding anchor log-odds margin.

**Corollary 1** (Bounded displacement and margin stability)**.** *Define the time-dependent evidence radius*

$$B_t = \alpha + \beta L \begin{cases} (1-\eta^{H_t})/(1-\eta), & 0 < \eta < 1, \\ H_t, & \eta = 1. \end{cases}$$

*Then $\mathrm{KL}(\boldsymbol{\pi}_t^f\|\boldsymbol{\pi}_t^b) \leq 2B_t$. If two items satisfy $b_t(i) - b_t(j) > 2B_t$, then $f_t(i) > f_t(j)$. Hence the closure remains in a bounded information-geometric neighborhood of the behavioral prior and can reverse only anchor comparisons whose margins lie inside a deterministic evidence band.*

For a target item, the pairwise identity above also gives an exact recovery criterion: semantic and successor evidence correct an anchor error precisely when their combined contrast exceeds the erroneous anchor margin. Theorem 3 in Appendix A shows that such a correction survives empirical operator error whenever its oracle margin exceeds twice the propagated dynamic uncertainty.

Since $|a_t(j)| \leq 1$ and $0 \leq d_t(j) \leq L\sum_{r=0}^{H_t-1}\eta^r$, the correction magnitude is at most $B_t$; the result follows by the triangle inequality. The bound partitions candidate pairs into a large-margin region that preserves the behavioral order and an evidence-responsive region in which semantic and successor contrasts determine the update.

**Optimization and complexity.** The MPT anchor is adapted with standard next-item cross-entropy, $\mathcal{L} = -\sum_t \log\boldsymbol{\pi}_t^b(i_{t+1})$, and is then fixed while applying the evidence layer. Language vectors and successor operators are also fixed at inference. For one prefix, constructing $\mathbf{s}_t$ costs $O(H_t d_s)$, computing its catalog compatibility costs $O(Md_s)$, accumulating $\mathbf{d}_t$ costs $O(H_t K_s)$, and combining catalog scores costs $O(M)$. The sparse operator family requires $O(H_{\mathrm{op}}MK_s)$ storage, not $O(H_{\mathrm{op}}M^2)$.

## 4 Experiments

The evaluation asks: **Q1** does state–dynamics evidence improve a strong prior (Table 2); **Q2** are the two fields complementary and **Q3** is exact horizon alignment necessary (Table 4); **Q4** do state estimation and popularity calibration behave as predicted (Table 5); and **Q5** does the formulation transfer across domains and encoders at practical cost (Tables 3 and 6, Appendix C.2)? Table 1 gives the same-setting comparison.

### 4.1 Experimental Setup

**Evaluation and implementation.** We use Amazon Beauty, Sports & Outdoors, Toys & Games, Yelp2018-10core, and OnlineRetail with chronological leave-one-out splits. Each target is ranked against the full eligible catalog after masking the complete observed history; HR and NDCG are reported at 5, 10, and 20 together with MRR. The main estimates average three anchor seeds.

Controlled component, horizon, state, and popularity comparisons share one reference checkpoint so that each row changes only the named factor. MPT supplies the behavioral anchor (Xu et al., 2026a). Sentence-T5-XL supplies the default frozen language vectors (Ni et al., 2022). Qwen3-Embedding-0.6B (Zhang et al., 2025), BGE-base-en-v1.5 (Xiao et al., 2024), and E5-base-v2 (Wang et al., 2022) test representation transfer. Appendix B gives dataset statistics and the numerical configuration.

**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), WPGRec (Liu et al., 2026), and ACE (Lee et al., 2026).

### 4.2 Same-Setting Comparison with Prior Methods

Table 1: Full-catalog results under the same Amazon data, split, candidate, history, and metric definitions.

| | Beauty | | | | Sports | | | | Toys | | | |
|---|---|---|---|---|---|---|---|---|---|---|---|---|
| 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 | 0.0259 | 0.0135 | 0.0407 | 0.0172 | 0.0490 | 0.0282 | 0.0704 | 0.0335 |
| SASRec | 0.0846 | 0.0493 | 0.1151 | 0.0571 | 0.0485 | 0.0275 | 0.0691 | 0.0327 | 0.0872 | 0.0535 | 0.1149 | 0.0605 |
| FMLP-Rec | 0.0775 | 0.0454 | 0.1072 | 0.0529 | 0.0413 | 0.0238 | 0.0593 | 0.0283 | 0.0855 | 0.0521 | 0.1124 | 0.0589 |
| HSTU | 0.1035 | 0.0621 | 0.1413 | 0.0715 | 0.0601 | 0.0341 | 0.0858 | 0.0406 | 0.1103 | 0.0659 | 0.1474 | 0.0753 |
| SASRec$_+$ | 0.1093 | 0.0648 | 0.1508 | 0.0752 | 0.0647 | 0.0369 | 0.0914 | 0.0437 | 0.1173 | 0.0697 | 0.1584 | 0.0801 |
| UniSRec | 0.0807 | 0.0471 | 0.1115 | 0.0549 | 0.0458 | 0.0260 | 0.0647 | 0.0307 | 0.0916 | 0.0558 | 0.1190 | 0.0627 |
| RecFormer | 0.0700 | 0.0413 | 0.0959 | 0.0478 | 0.0399 | 0.0234 | 0.0558 | 0.0275 | 0.0805 | 0.0504 | 0.1036 | 0.0563 |
| E4SRec | 0.0879 | 0.0526 | 0.1196 | 0.0605 | 0.0518 | 0.0303 | 0.0734 | 0.0358 | 0.0980 | 0.0592 | 0.1291 | 0.0671 |
| Qwen2.5 | 0.0826 | 0.0464 | 0.1200 | 0.0558 | 0.0473 | 0.0256 | 0.0712 | 0.0316 | 0.0784 | 0.0433 | 0.1114 | 0.0515 |
| MPT-Adapter | 0.1118 | 0.0673 | 0.1531 | 0.0777 | 0.0651 | 0.0373 | 0.0924 | 0.0442 | 0.1202 | 0.0705 | 0.1617 | 0.0809 |
| MPT-LoRA | 0.1100 | 0.0658 | 0.1514 | 0.0763 | 0.0644 | 0.0365 | 0.0927 | 0.0436 | 0.1192 | 0.0701 | 0.1611 | 0.0806 |
| ACE | 0.0886 | 0.0433 | 0.1288 | 0.0534 | – | – | – | – | 0.0977 | 0.0468 | 0.1342 | 0.0560 |
| WPGRec | 0.1197 | 0.0742 | 0.1583 | 0.0813 | 0.0704 | 0.0401 | 0.0991 | 0.0439 | – | – | – | – |
| AtlasRec | **0.1205** | **0.0730** | **0.1628** | **0.0836** | **0.0725** | **0.0425** | **0.1009** | **0.0496** | **0.1353** | **0.0832** | **0.1763** | **0.0936** |
| Improve | **+0.67%** | −1.62% | **+2.84%** | **+2.83%** | **+2.98%** | **+5.99%** | **+1.82%** | **+12.98%** | **+12.56%** | **+18.01%** | **+9.03%** | **+15.70%** |

Table 1 compares all methods under the same Amazon data, split, candidate mask, history mask, and metric definitions. AtlasRec leads 11 of 12 HR/NDCG cells: it wins three Beauty metrics, all four Sports metrics, and all four available Toys metrics, with 9.03–18.01% gains on Toys; the remaining Beauty NDCG@10 is 0.0730 versus 0.0742. Its advantage over MPT-Adapter and MPT-LoRA isolates the value of an explicit state–dynamics layer around the same pretrained model family. The comparison with WPGRec further separates frequency-scale modeling from target-distance modeling: wavelets organize behavioral bands, whereas AtlasRec assigns each historical source the successor law for its actual target distance. Improvements at both 10 and 20 show that this alignment raises early order quality without sacrificing broader-cutoff retrieval.

### 4.3 Matched-Anchor Gains

Table 2: Matched-anchor performance and paired significance under strict full-catalog evaluation.

| Dataset | Variant | H@5 | H@10 | H@20 | N@5 | N@10 | N@20 | MRR |
|---|---|---|---|---|---|---|---|---|
| Beauty | Anchor | 0.0800±0.0008 | 0.1116±0.0005 | 0.1536±0.0013 | 0.0571±0.0006 | 0.0672±0.0005 | 0.0778±0.0006 | 0.0608±0.0006 |
| | AtlasRec | 0.0865±0.0005 | 0.1205±0.0007 | 0.1628±0.0011 | 0.0621±0.0006 | 0.0730±0.0004 | 0.0836±0.0002 | 0.0657±0.0004 |
| | Improve | +8.06% | +7.97% | +6.02% | +8.72% | +8.64% | +7.57% | +8.03% |
| Sports | Anchor | 0.0451±0.0004 | 0.0661±0.0002 | 0.0939±0.0009 | 0.0305±0.0004 | 0.0373±0.0003 | 0.0443±0.0004 | 0.0338±0.0003 |
| | AtlasRec | 0.0511±0.0002 | 0.0725±0.0001 | 0.1009±0.0007 | 0.0356±0.0002 | 0.0425±0.0002 | 0.0496±0.0003 | 0.0388±0.0003 |
| | Improve | +13.17% | +9.59% | +7.43% | +16.53% | +13.94% | +12.17% | +14.69% |
| Toys | Anchor | 0.0867±0.0011 | 0.1217±0.0030 | 0.1636±0.0028 | 0.0604±0.0006 | 0.0716±0.0012 | 0.0821±0.0011 | 0.0632±0.0006 |
| | AtlasRec | 0.0993±0.0016 | 0.1353±0.0010 | 0.1763±0.0020 | 0.0717±0.0008 | 0.0832±0.0008 | 0.0936±0.0010 | 0.0742±0.0008 |
| | Improve | +14.53% | +11.16% | +7.78% | +18.64% | +16.23% | +13.93% | +17.25% |

| Dataset | Relative-gain range | Positive metrics | Paired test |
|---|---|---|---|
| Beauty | 6.02%–8.72% | 7/7 | all $p < 0.01$ |
| Sports | 7.43%–16.53% | 7/7 | all $p < 0.01$ |
| Toys | 7.78%–18.64% | 7/7 | all $p < 0.01$ |

**Q1—State–dynamics increment.** Table 2 holds the MPT checkpoint, semantic source, split, candidate mask, and evaluator fixed; the anchor removes only $a_t$ and $d_t$. All 21 metrics improve, by

6.02–8.72% on Beauty, 7.43–16.53% on Sports, and 7.78–18.64% on Toys, with HR, NDCG, and MRR rising together. The largest gains occur at tighter cutoffs, where bounded evidence can resolve uncertain anchor comparisons most directly, while positive rank-20 gains show that the update also brings missed targets into the broader evaluated region. Because behavioral capacity is unchanged, the improvement identifies the contribution of the state–dynamics evidence layer rather than a stronger sequence backbone; the paired summary confirms that every matched-seed difference is positive with $p < 0.01$.

### 4.4 Cross-Domain Evaluation

Table 3: Full AtlasRec results on all five datasets.

| Dataset | H@10 | N@10 | H@20 | N@20 |
|---|---|---|---|---|
| Beauty | **0.1205** | **0.0730** | **0.1628** | **0.0836** |
| Sports | **0.0725** | **0.0425** | **0.1009** | **0.0496** |
| Toys | **0.1353** | **0.0832** | **0.1763** | **0.0936** |
| Yelp | **0.0523** | **0.0264** | **0.0867** | **0.0351** |
| OnlineRetail | **0.1692** | **0.0864** | **0.2590** | **0.1090** |

**Q5—Cross-domain transfer.** Table 3 extends the complete model from Amazon product-review sequences to two structurally different domains. Yelp couples a larger catalog with sparse user histories, so its result emphasizes semantic regularity when direct successor support is limited. OnlineRetail contains denser transactional recurrence and repeat-purchase structure, allowing horizon-specific rows to accumulate stronger destination evidence across earlier purchases. Together with the three Amazon categories, these results cover different catalog sizes, interaction densities, transition sparsities, and recurrence patterns under the same state–dynamics–decision interface. Appendix C.2 further holds the OnlineRetail evidence configuration fixed while changing the frozen language geometry; all 28 anchor comparisons remain positive across Sentence-T5, Qwen3, BGE, and E5.

### 4.5 Component and Horizon Attribution

Table 4: Core ablation with the anchor and evaluator held fixed.

| Dataset | Variant | H@10 | N@10 | H@20 | N@20 | MRR |
|---|---|---|---|---|---|---|
| Beauty | Anchor | 0.1119 | 0.0671 | 0.1549 | 0.0780 | 0.0607 |
| | Semantic state only | 0.1189 | 0.0698 | 0.1612 | 0.0805 | 0.0621 |
| | Successor dynamics only | 0.1148 | 0.0692 | 0.1569 | 0.0798 | 0.0624 |
| | One-step history | 0.1183 | 0.0719 | 0.1612 | 0.0827 | 0.0649 |
| | **Full multi-horizon** | **0.1210** | **0.0727** | **0.1641** | **0.0836** | **0.0652** |
| Sports | Anchor | 0.0663 | 0.0373 | 0.0948 | 0.0445 | 0.0339 |
| | Semantic state only | 0.0694 | 0.0386 | 0.0978 | 0.0458 | 0.0348 |
| | Successor dynamics only | 0.0696 | 0.0409 | 0.0986 | 0.0482 | 0.0375 |
| | One-step history | 0.0706 | 0.0409 | 0.0990 | 0.0480 | 0.0373 |
| | **Full multi-horizon** | **0.0726** | **0.0427** | **0.1016** | **0.0500** | **0.0391** |
| Toys | Anchor | 0.1222 | 0.0715 | 0.1643 | 0.0821 | 0.0629 |
| | Semantic state only | 0.1304 | 0.0762 | 0.1731 | 0.0869 | 0.0667 |
| | Successor dynamics only | 0.1270 | 0.0782 | 0.1667 | 0.0882 | 0.0700 |
| | One-step history | 0.1326 | 0.0820 | 0.1734 | 0.0923 | 0.0732 |
| | **Full multi-horizon** | **0.1347** | **0.0829** | **0.1762** | **0.0933** | **0.0738** |

**Q2–Q3—Complementarity and horizon alignment.** Table 4 shows that semantic state and successor dynamics each improve every displayed anchor metric, while their combination exceeds either field on all three datasets. Beauty receives the larger standalone HR@10 gain from semantic compatibility, whereas Sports and Toys receive the larger standalone NDCG@10 gain from directional successor evidence, revealing dataset-dependent roles rather than interchangeable scores. Replacing the aligned field with history-aggregated one-step evidence reduces all 15 metrics, including NDCG@10 by 1.11%, 4.40%, and 1.10% on Beauty, Sports, and Toys. Since history, support budget, semantic vectors, and anchor scores remain fixed, this control isolates the operator index and directly tests Theorem 1: older events contain useful transition meaning that is lost when every source is forced through $\widehat{\mathsf{T}}_1$. Appendix C.3 further shows that the full model exceeds both the stronger single field and the one-step control in every one of the 15 HR, NDCG, and MRR comparisons.

### 4.6 State Construction and Popularity Reference

Table 5: Semantic-state and popularity-reference controls under the matched full-catalog protocol.

| Control | Dataset | Variant | H@10 | N@10 | MRR |
|---|---|---|---|---|---|
| State construction | Beauty | Last item | 0.1183 | 0.0719 | 0.0649 |
| | | Uniform history | 0.1198 | 0.0722 | 0.0649 |
| | | **Recency-weighted history** | **0.1210** | **0.0727** | **0.0652** |
| | Sports | Last item | 0.0706 | 0.0409 | 0.0373 |
| | | Uniform history | 0.0727 | 0.0425 | 0.0389 |
| | | **Recency-weighted history** | **0.0726** | **0.0427** | **0.0391** |
| Popularity reference | Beauty | Uniform reference (no correction) | 0.1205 | 0.0723 | 0.0649 |
| | | **Popularity-relative (AtlasRec)** | **0.1210** | **0.0727** | **0.0652** |
| | Sports | Uniform reference (no correction) | 0.0727 | 0.0426 | 0.0390 |
| | | **Popularity-relative (AtlasRec)** | **0.0726** | **0.0427** | **0.0391** |

**Q4—Evidence estimation.** Table 5 shows that replacing the last-item state with a short semantic trajectory improves Beauty HR@10/NDCG@10 by 2.28%/1.11% and Sports by 2.83%/4.40%; recency weighting then raises NDCG@10 and MRR over uniform history on both datasets. This ordering matches Proposition 1: temporal aggregation estimates the maximum-agreement direction of recent intent, while recency determines which intent dominates. Popularity-relative lift improves every shown Beauty metric and the position-sensitive Sports metrics. In the role identified by Proposition 2, this reference calibrates successor evidence by removing catalog-wide frequency before source-specific support is evaluated. The complete-cutoff controls in Appendix C.4 show that recency weighting improves every NDCG and MRR entry over uniform history, while Appendix C.2 confirms that the resulting roles persist across frozen language geometries.

### 4.7 Sensitivity, Transfer, and Evaluation Cost

Table 6: Sensitivity and measured full-catalog evaluation cost.

| Sweep | Setting and corresponding NDCG@10 | | | | | | | |
|---|---|---|---|---|---|---|---|---|
| $\alpha$ setting | 1.6 | 2.0 | 2.4 | 2.8 | 3.2 | 3.6 | 4.0 | 4.8 |
| NDCG@10 | 0.0725 | 0.0725 | 0.0727 | 0.0729 | 0.0730 | 0.0730 | 0.0730 | 0.0727 |
| $\beta$ setting | 0.12 | 0.18 | 0.24 | 0.30 | 0.36 | 0.42 | – | – |
| NDCG@10 | 0.0722 | 0.0726 | 0.0730 | 0.0728 | 0.0724 | 0.0718 | – | – |
| $H$ setting | 1 | 2 | 3 | 5 | 8 | – | – | – |
| NDCG@10 | 0.0719 | 0.0726 | 0.0730 | 0.0725 | 0.0724 | – | – | – |

| Variant | Time (s) | Mem. (MiB) |
|---|---|---|
| Anchor | $4.442 \pm 0.396$ | 1868 |
| AtlasRec | $4.613 \pm 0.521$ | 1878 |

**Q5—Sensitivity, transfer, and cost.** Table 6 links the operating regime to the theory. Performance is stable over a broad $\alpha$ interval, while the interior $\beta$ optimum reflects the closure trade-off: successor evidence first corrects uncertain anchor margins, but excessive strength admits weaker contrasts. The peak at $H = 3$ matches Theorem 4, since additional horizons reduce truncation while drawing on progressively sparser transition rows. Under one fixed OnlineRetail configuration, changing only the frozen encoder preserves 10.78–13.45% NDCG@10 gains and 14.31–16.79% MRR gains. With language vectors and sparse operators cached, the complete update adds 0.171 seconds (3.85%) and 10 MiB (0.54%); its online work is sparse evidence accumulation rather than another sequence encoder.

## 5 Conclusion

We presented AtlasRec, a state–dynamics–decision framework that treats frozen language geometry as an estimate of current intent rather than another item feature. Target-distance-aligned successor operators supply the missing direction of evolution, while minimum-information closure combines semantic and transition evidence with a pretrained behavioral prior. The analysis identifies the conditional-KL cost of repeated one-step dynamics, propagates finite-sample uncertainty, and establishes bounded ranking stability. Across five datasets and four encoders, AtlasRec leads 11 of 12 same-setting comparisons and improves all 21 matched-anchor metrics, with ablations confirming the roles of state estimation, horizon alignment, and popularity-relative evidence. Future work will learn smoother cross-horizon operators, update the successor atlas under evolving catalogs, and extend the framework to longer-term trajectory prediction.

### Reproducibility Statement

Appendix A provides the theoretical proofs. Appendix B specifies the split, masking rule, metrics, model geometry, optimization, successor construction, evidence parameters, and random seeds. Appendix C reports additional experimental results, including seed stability, encoder transfer, and complete ablation controls. Source code and machine-precision result files will accompany the final version.

### 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

Kenneth Ward Church and Patrick Hanks. Word association norms, mutual information, and lexicography. *Computational Linguistics*, 16(1):22–29, 1990. URL https://aclanthology.org/J90-1003/.

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.

Dongcheol Lee, Hye-young Kim, and Jongwuk Lee. ACE: Anisotropy-controllable embedding for LLM-enhanced sequential recommendation. In *Proceedings of the 49th International ACM SIGIR Conference on Research and Development in Information Retrieval*, pp. 3919–3923, 2026. doi: 10.1145/3805712.3809911.

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.

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.

Jianmo Ni, Gustavo Hernández Ábrego, Noah Constant, Ji Ma, Keith B. Hall, Daniel Cer, and Yinfei Yang. Sentence-T5: Scalable sentence encoders from pre-trained text-to-text models. In *Findings of the Association for Computational Linguistics: ACL 2022*, pp. 1864–1874. Association for Computational Linguistics, 2022. doi: 10.18653/v1/2022.findings-acl.146. URL https://aclanthology.org/2022.findings-acl.146/.

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.

Shitao Xiao, Zheng Liu, Peitian Zhang, Niklas Muennighoff, Defu Lian, and Jian-Yun Nie. C-Pack: Packed resources for general Chinese embeddings. In *Proceedings of the 47th International ACM SIGIR Conference on Research and Development in Information Retrieval*, pp. 641–649. ACM, 2024. doi: 10.1145/3626772.3657878.

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 PROPOSITION 1

For any unit vector $\mathbf{s}$ and normalized item vectors $\mathbf{v}_{i_{t-r}}$,

$$\sum_{r=0}^{H_t-1} \kappa^r \|\mathbf{s} - \mathbf{v}_{i_{t-r}}\|_2^2 = 2 \sum_{r=0}^{H_t-1} \kappa^r - 2\langle \mathbf{s}, \mathbf{m}_t \rangle.$$

Minimizing the left side is therefore equivalent to maximizing $\langle \mathbf{s}, \mathbf{m}_t \rangle$. By Cauchy–Schwarz, the maximum is $\|\mathbf{m}_t\|_2$ and is attained by $\mathbf{s} = \mathbf{m}_t / \|\mathbf{m}_t\|_2$, which equals Equation 3 when $\|\mathbf{m}_t\|_2 \geq \varepsilon_s$. Setting $H = 1$ or $\kappa = 1$ gives the stated boundary cases.

## A.2 PROOF OF PROPOSITION 2

For an observed pair, monotonicity of the logarithm gives $\log(P_n(j \mid i)/p(j)) > 0$ exactly when $P_n(j \mid i) > p(j)$; positive clipping preserves this condition until the upper cap. Moreover,

$$\log \frac{P_n(j \mid i)}{u(j)} = \log \frac{P_n(j \mid i)}{p(j)} + \log \frac{p(j)}{1/M},$$

which is the claimed decomposition. Substituting $P_n(\cdot \mid i) = p(\cdot)$ proves the null statement.

## A.3 FINITE OBSERVED-PREFIX TRUNCATION

**Proposition 3** (Finite observed-prefix truncation)**.** *Let* $H_t^{\mathrm{op}} = \min\{H_{\mathrm{op}}, t\}$ *and let* $\mathbf{d}_t^{\mathrm{op}} = \sum_{r=0}^{H_t^{\mathrm{op}}-1} \eta^r \widehat{\mathsf{T}}_{r+1}^{\top} \boldsymbol{\delta}_{i_{t-r}}$. *For* $\eta < 1$,

$$\|\mathbf{d}_t^{\mathrm{op}} - \mathbf{d}_t\|_\infty \leq \frac{L\eta^{H_t}(1 - \eta^{H_t^{\mathrm{op}} - H_t})}{1-\eta} \leq \frac{L\eta^{H_t}}{1-\eta}.$$

Every coordinate of an omitted vector term is nonnegative and at most $L\eta^r$. Therefore

$$\|\mathbf{d}_t^{\mathrm{op}} - \mathbf{d}_t\|_\infty \leq \sum_{r=H_t}^{H_t^{\mathrm{op}}-1} L\eta^r = \frac{L\eta^{H_t}(1 - \eta^{H_t^{\mathrm{op}} - H_t})}{1-\eta} \leq \frac{L\eta^{H_t}}{1-\eta}.$$

This is the claimed finite-prefix bound.

## A.4 PROOF OF THEOREM 2

Write $r_t(j) = \alpha a_t(j) + \beta d_t(j)$. The objective in Equation 10 is strictly convex on the support of $\boldsymbol{\pi}_t^b$. Introducing a multiplier $\lambda$ for $\sum_{j \in \mathcal{I}} \boldsymbol{\pi}(j) = 1$, stationarity gives

$$\log \frac{\boldsymbol{\pi}(j)}{\boldsymbol{\pi}_t^b(j)} + 1 - r_t(j) + \lambda = 0,$$

so $\boldsymbol{\pi}(j) = \boldsymbol{\pi}_t^b(j) e^{r_t(j)} e^{-1-\lambda}$. Normalization yields the product-of-experts expression in Theorem 2. Substituting $\boldsymbol{\pi}_t^b(j) \propto e^{b_t(j)}$ gives $\boldsymbol{\pi}_t^f = \mathrm{softmax}(\mathbf{f}_t)$, exactly Equation 9. Strict convexity establishes uniqueness.

## A.5 PROOF OF COROLLARY 1

From unit normalization, $|a_t(j)| \leq 1$. Equation 5 gives $0 \leq \widehat{\tau}_n(i,j) \leq L$, so Equation 6 implies $0 \leq d_t(j) \leq L \sum_{r=0}^{H_t-1} \eta^r$. Consequently $|f_t(j) - b_t(j)| \leq B_t$. For any $i, j$,

$$f_t(i) - f_t(j) \geq b_t(i) - b_t(j) - 2B_t.$$

An anchor margin greater than $2B_t$ therefore remains positive. For $r_t(j) = \alpha a_t(j) + \beta d_t(j)$ and $Z_t = \mathbb{E}_{j \sim \boldsymbol{\pi}_t^b}[e^{r_t(j)}]$, Theorem 2 gives

$$\mathrm{KL}(\boldsymbol{\pi}_t^f \| \boldsymbol{\pi}_t^b) = \mathbb{E}_{j \sim \boldsymbol{\pi}_t^f}[r_t(j)] - \log Z_t.$$

Since $|r_t(j)| \leq B_t$, the expectation is at most $B_t$ and $\log Z_t \geq -B_t$. Hence the KL displacement is at most $2B_t$.

### A.6 Pairwise Geometry of the Evidence Closure

For a candidate pair $(i, j)$, define the anchor margin $m_t^b(i,j) = b_t(i) - b_t(j)$, semantic contrast $\Delta_t^a(i,j) = a_t(i) - a_t(j)$, and successor contrast $\Delta_t^d(i,j) = d_t(i) - d_t(j)$. Equation 9 gives the exact identity

$$f_t(i) - f_t(j) = m_t^b(i,j) + \alpha\Delta_t^a(i,j) + \beta\Delta_t^d(i,j).$$

The pair reverses relative to the anchor precisely when the combined evidence crosses $-m_t^b(i,j)$. This partitions the catalog into comparisons already settled by the behavioral prior and comparisons for which the observed semantic trajectory supplies enough evidence to change the decision. The same identity also separates the two experimental interventions: removing semantic state sets $\Delta_t^a = 0$, removing successor dynamics sets $\Delta_t^d = 0$, and the complete model evaluates their joint displacement of the anchor margin.

The variational form adds a distributional interpretation. If $Z_t$ denotes the normalizer in Theorem 2, then

$$\log \frac{\pi_t^f(j)}{\pi_t^b(j)} = \alpha a_t(j) + \beta d_t(j) - \log Z_t.$$

Thus the update is a catalog-wide exponential tilt of the behavioral reference. The normalizer changes probability mass but cancels from all pairwise rankings, so the score-space and probability-space accounts are exactly consistent.

For held-out target $y$, define its candidate set $\mathcal{C}_t$ as in Appendix B. Its final rank has the exact form

$$\mathrm{rank}_f(y) = 1 + \sum_{j \in \mathcal{C}_t \setminus \{y\}} \mathbf{1}\left[m_t^b(y,j) + \alpha\Delta_t^a(y,j) + \beta\Delta_t^d(y,j) < 0\right].$$

Consequently, $\mathrm{rank}_f(y) \le K$ exactly when at most $K-1$ candidate margins remain negative. This connects evidence closure directly to HR@$K$, NDCG@$K$, and MRR rather than only to score perturbation.

**Theorem 3** (Robust pairwise ranking recovery)**.** *Let $\mathbf{d}_t^\circ$ be an oracle successor field on the same sparse support and suppose $\|\mathbf{d}_t - \mathbf{d}_t^\circ\|_\infty \le \epsilon_t^d$. For target $y$ and competitor $j$, define the oracle closed margin*

$$m_t^\circ(y,j) = m_t^b(y,j) + \alpha\Delta_t^a(y,j) + \beta\big(d_t^\circ(y) - d_t^\circ(j)\big).$$

*If $m_t^\circ(y,j) > 2\beta\epsilon_t^d$, then $f_t(y) > f_t(j)$. If this condition holds for all but at most $K-1$ candidates, then* $\mathrm{rank}_f(y) \le K$.

Indeed, each candidate's dynamic error is at most $\epsilon_t^d$, so the empirical pairwise dynamic contrast differs from its oracle counterpart by at most $2\epsilon_t^d$. Subtracting this worst-case deviation proves both claims.

### A.7 What Horizon Alignment Changes

For each $r$, subtracting the one-step signed information from the aligned signed information cancels the shared marginal:

$$\lambda_{r+1}^\star(i,j) - \lambda_1^\star(i,j) = \log \frac{P_{r+1}^\star(j \mid i)}{P_1^\star(j \mid i)}.$$

The $r = 0$ term is zero; summing the remaining discounted contrasts proves Equation 7. If $J \sim P_{r+1}^\star(\cdot \mid i)$, then

$$\mathbb{E}\left[\log \frac{P_{r+1}^\star(J \mid i)}{P_1^\star(J \mid i)}\right] = \mathrm{KL}\big(P_{r+1}^\star(\cdot \mid i) \| P_1^\star(\cdot \mid i)\big).$$

Discounting and summing gives $\mathcal{M}_t$. Nonnegativity and the equality condition follow from Gibbs' inequality. If absolute continuity fails, the corresponding component cost is infinite, so horizon invariance still characterizes zero mismatch.

For the bounded empirical operators actually used by AtlasRec, let $\mathbf{d}_t^{(1)}$ be the one-step control obtained by replacing every $\widehat{\mathsf{T}}_{r+1}$ in Equation 6 with $\widehat{\mathsf{T}}_1$. Their difference is

$$\mathbf{d}_t - \mathbf{d}_t^{(1)} = \sum_{r=1}^{H_t - 1} \eta^r \big(\widehat{\mathsf{T}}_{r+1} - \widehat{\mathsf{T}}_1\big)^\top \boldsymbol{\delta}_{i_{t-r}}.$$

Because every stored coordinate lies in $[0, L]$,

$$\|\mathbf{d}_t - \mathbf{d}_t^{(1)}\|_\infty \le L \sum_{r=1}^{H_t-1} \eta^r.$$

The discrepancy is therefore bounded, vanishes for a one-event history, and is supported entirely by older events whose correct target distance exceeds one. The horizon control changes no item history, semantic vector, anchor score, or successor support budget; it isolates whether distance-indexed transition meaning improves the same decision problem.

### A.8 Finite-Sample Successor Uncertainty

We now relate training support to the dynamic error used in Theorem 3. Fix a source item $i$ and horizon $n$. Let $N_{i,n} = C_n^{\mathrm{src}}(i)$ be the observations in that source–horizon row, $P_n^\star(\cdot \mid i)$ its conditional successor law, $N_{\mathrm{pop}}$ the number of next-event observations before add-one smoothing, and $p^\star$ the catalog marginal. Under the conditional multinomial model, the row counts sample $P_n^\star(\cdot \mid i)$ and the next-event counts sample $p^\star$. Define

$$\epsilon_{i,n} = \sqrt{\frac{\log(4M^2 H_{\mathrm{op}}/\delta)}{2N_{i,n}}}, \qquad \epsilon_{\mathrm{pop}} = \sqrt{\frac{\log(4M/\delta)}{2N_{\mathrm{pop}}}} + \frac{M}{N_{\mathrm{pop}} + M}.$$

The final term accounts for Laplace smoothing of the catalog marginal. For the empirical support $\mathcal{S}_n(i)$ and support ranks $q_n(i,j)$, define an oracle-on-support weight $\widehat{\tau}_n^\circ(i,j)$ by Equation 5 with the same $\mathcal{S}_n(i)$ and $q_n(i,j)$, but with the clipped population lift $\mathrm{clip}_{[0,L]} \log[P_n^\star(j \mid i)/p^\star(j)]$. The result is conditional on this observed support; $\rho > 0$ denotes a common probability floor over the queried retained tuples.

**Theorem 4** (Successor-field concentration)**.** *With probability at least $1 - \delta$, all rows with $N_{i,n} > 0$ and all candidates simultaneously satisfy*

$$|P_n(j \mid i) - P_n^\star(j \mid i)| \le \epsilon_{i,n}, \qquad |p(j) - p^\star(j)| \le \epsilon_{\mathrm{pop}}.$$

*Suppose that, on every queried retained tuple, $P_n^\star(j \mid i), p^\star(j) \ge \rho > 0$ and both deviations are at most $\rho/2$. Then*

$$|\widehat{\tau}_n(i,j) - \widehat{\tau}_n^\circ(i,j)| \le \frac{2(1-\mu)}{\rho}(\epsilon_{i,n} + \epsilon_{\mathrm{pop}}).$$

*For $d_t^\circ(j) = \sum_{r=0}^{H_t-1} \eta^r \widehat{\tau}_{r+1}^\circ(i_{t-r}, j)$, this gives*

$$\|\mathbf{d}_t - \mathbf{d}_t^\circ\|_\infty \le \epsilon_t^d := \frac{2(1-\mu)}{\rho} \sum_{r=0}^{H_t-1} \eta^r (\epsilon_{i_{t-r},r+1} + \epsilon_{\mathrm{pop}}). \tag{11}$$

*If $\mathbf{d}_t^{\circ,\mathrm{op}}$ additionally retains every available horizon through $H_t^{\mathrm{op}}$, then*

$$\|\mathbf{d}_t - \mathbf{d}_t^{\circ,\mathrm{op}}\|_\infty \le \epsilon_t^d + L \sum_{r=H_t}^{H_t^{\mathrm{op}}-1} \eta^r.$$

Hoeffding's inequality followed by a union bound gives simultaneous probability deviations. On the stated event, empirical and population probabilities are at least $\rho/2$; the logarithm is therefore $2/\rho$-Lipschitz. Clipping is nonexpansive, and the shared rank term cancels, proving the weight bound. Discounted summation proves Equation 11; Proposition 3 supplies the final omitted-horizon term. The last display exposes the finite-horizon trade-off: increasing $H$ removes bounded successor mass from the truncation term while accumulating uncertainty from additional, typically sparser horizon-specific rows. The interior optimum in Table 6 is consistent with this prediction.

### A.9 Sparse Construction and Scaling

Let $N_{\mathrm{tr}}$ be the number of training interactions. Counting all source–successor pairs through horizon $H_{\mathrm{op}}$ requires $O(H_{\mathrm{op}} N_{\mathrm{tr}})$ additions. Retaining the largest $K_s$ counts for each source and horizon

yields at most $H_{\mathrm{op}} M K_s$ stored weights. A streaming top-$K_s$ construction costs $O(E \log K_s)$ over the observed transition edges $E$, while the dense $M \times M$ matrices in the notation never need to be materialized.

At inference, semantic-state construction costs $O(H_t d_s)$, catalog compatibility costs $O(M d_s)$, and horizon-aligned successor accumulation touches at most $H_t K_s$ entries. The final catalog update is $O(M)$. The evidence layer therefore adds $O(M d_s + H_t K_s)$ computation per prefix and $O(H_{\mathrm{op}} M K_s)$ offline storage, with no quadratic catalog term.

# B Implementation and Reproducibility

## B.1 Datasets and Evaluation

Table 7: Dataset statistics.

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

Interactions are ordered by timestamp per user. The final event is the test target and the preceding event is the validation target. Full-catalog ranking masks items already present in the observed prefix while retaining the held-out target. For target rank $r$, HR@$K$ is $\mathbf{1}[r \le K]$, NDCG@$K$ is $\mathbf{1}[r \le K]/\log_2(r+1)$, and reciprocal rank is $1/r$; dataset metrics average these quantities over users.

More explicitly, for catalog $\mathcal{I}$, observed prefix $\mathcal{H}_t = (i_1, \ldots, i_t)$, target $i_{t+1}$, and final score $f_t$, the evaluated set and target rank are

$$\mathcal{C}_t = \{i_{t+1}\} \cup \big(\mathcal{I} \setminus \{i_1, \ldots, i_t\}\big), \qquad r_t = 1 + \sum_{j \in \mathcal{C}_t \setminus \{i_{t+1}\}} \mathbf{1}[f_t(j) > f_t(i_{t+1})].$$

The same $\mathcal{C}_t$ and rank computation are used for every matched anchor and AtlasRec run. Thus the reported differences change only the ordering induced by the evidence layer, not catalog coverage or historical-item eligibility.

## B.2 Numerical Specification

Table 8: Numerical configuration used by the reported model.

| Component | Quantity | Value |
|---|---|---|
| Behavioral anchor | layers / width / heads | $4/256/2$ |
| Behavioral anchor | maximum length / dropout / $\tau_b$ | $50/0.2/0.07$ |
| Optimization | optimizer / learning rate / weight decay | AdamW / $10^{-3}$ / 0.1 |
| Optimization | batch size / runs | $64/3$ |
| Semantic state | $d_s/H/\kappa$ | $64/3/0.65$ |
| Successor family | $H_{\mathrm{op}}/K_s/L$ | $4/32/6$ |
| Successor family | $\mu/\eta$ | $0.25/0.8$ |
| Evidence closure | $\alpha/\beta$ | $2.4/0.24$ |

The three runs use seeds 42, 43, and 44. In Equation 2, $x_j$ concatenates title, categories, and brand for Amazon; item name, categories, and city for Yelp; and the product-text field for OnlineRetail. The fixed seeded projection $\mathbf{W}_{\mathrm{sem}}$ maps every frozen language representation into the common $d_s$-dimensional geometry before normalization.

### B.3 Operator Construction and Evaluation Path

The successor family is built once from the training split. For each horizon $n$, training prefixes contribute counts $C_n(i, j)$ and the shared next-event marginal supplies $p(j)$. Each observed edge receives the clipped popularity-relative lift in Equation 4; the $K_s$ largest-count successors per source are retained, ranked within that support, and stored with Equation 5. Sparse entries are accumulated by scatter addition. No validation or test transition contributes to the operators.

For an observed prefix, inference follows four deterministic steps: (i) obtain the anchor logits $\mathbf{b}_t$; (ii) form the normalized recency-weighted semantic state $\mathbf{s}_t$ and compatibility field $\mathbf{a}_t$; (iii) gather one sparse successor row for every available source–horizon pair and accumulate $\mathbf{d}_t$; and (iv) evaluate Equation 9. The shared evaluator subsequently masks the complete observed history from both the anchor and AtlasRec and ranks the remaining catalog. This ordering makes the candidate set, masking rule, and metric computation identical for the matched comparison.

The language encoder, semantic projection, popularity marginal, and successor operators are fixed during recommendation. Consequently, all item-side semantic vectors and transition weights can be cached. The online state depends only on the observed prefix, and every non-anchor contribution to a candidate score can be traced to either a semantic inner product or a stored source–horizon edge.

## C Additional Experimental Results

### C.1 Matched-Seed Stability and Significance

The paired unit is the matched anchor seed: each full-model run and its anchor share the trained checkpoint and evaluator, so the test directly measures the evidence-layer increment at fixed behavioral capacity. All 21 metric differences are positive and every paired two-sided test is below $0.01$. The largest gains occur on Sports and Toys NDCG@5, while the positive range on Beauty remains consistent across recall, discounted gain, and reciprocal rank.

### C.2 Complete Encoder-Transfer Results

Table 9: Complete OnlineRetail encoder-transfer results.

*(a) Hit rate*

| Encoder | Variant | H@5 | H@10 | H@20 |
|---|---|---|---|---|
| Sentence-T5-XL | Anchor | $0.0830 \pm 0.0018$ | $0.1548 \pm 0.0007$ | $0.2459 \pm 0.0053$ |
| | AtlasRec | $0.0929 \pm 0.0010$ | $0.1598 \pm 0.0024$ | $0.2504 \pm 0.0052$ |
| Qwen3-Embedding-0.6B | Anchor | $0.0813 \pm 0.0018$ | $0.1564 \pm 0.0009$ | $0.2520 \pm 0.0021$ |
| | AtlasRec | $0.0908 \pm 0.0014$ | $0.1598 \pm 0.0005$ | $0.2523 \pm 0.0028$ |
| BGE-base-en-v1.5 | Anchor | $0.0839 \pm 0.0017$ | $0.1559 \pm 0.0006$ | $0.2499 \pm 0.0038$ |
| | AtlasRec | $0.0920 \pm 0.0004$ | $0.1597 \pm 0.0016$ | $0.2510 \pm 0.0027$ |
| E5-base-v2 | Anchor | $0.0842 \pm 0.0011$ | $0.1566 \pm 0.0006$ | $0.2513 \pm 0.0005$ |
| | AtlasRec | $0.0955 \pm 0.0009$ | $0.1637 \pm 0.0002$ | $0.2559 \pm 0.0007$ |

*(b) Discounted gain and reciprocal rank*

| Encoder | Variant | N@5 | N@10 | N@20 | MRR |
|---|---|---|---|---|---|
| Sentence-T5-XL | Anchor | $0.0472 \pm 0.0012$ | $0.0704 \pm 0.0008$ | $0.0932 \pm 0.0020$ | $0.0595 \pm 0.0013$ |
| | AtlasRec | $0.0578 \pm 0.0007$ | $0.0793 \pm 0.0006$ | $0.1021 \pm 0.0009$ | $0.0695 \pm 0.0001$ |
| Qwen3-Embedding-0.6B | Anchor | $0.0461 \pm 0.0007$ | $0.0702 \pm 0.0008$ | $0.0942 \pm 0.0003$ | $0.0594 \pm 0.0010$ |
| | AtlasRec | $0.0556 \pm 0.0007$ | $0.0778 \pm 0.0003$ | $0.1011 \pm 0.0005$ | $0.0679 \pm 0.0004$ |
| BGE-base-en-v1.5 | Anchor | $0.0472 \pm 0.0012$ | $0.0704 \pm 0.0007$ | $0.0941 \pm 0.0013$ | $0.0595 \pm 0.0008$ |
| | AtlasRec | $0.0564 \pm 0.0005$ | $0.0782 \pm 0.0003$ | $0.1012 \pm 0.0006$ | $0.0683 \pm 0.0004$ |
| E5-base-v2 | Anchor | $0.0477 \pm 0.0010$ | $0.0710 \pm 0.0008$ | $0.0948 \pm 0.0007$ | $0.0602 \pm 0.0009$ |
| | AtlasRec | $0.0587 \pm 0.0011$ | $0.0806 \pm 0.0009$ | $0.1038 \pm 0.0008$ | $0.0703 \pm 0.0012$ |

Every cell in a AtlasRec row exceeds its source-matched anchor. The complete table also shows why the main text reports several ranking measures: the largest relative gains occur at early, position-sensitive metrics, but the positive direction remains at HR@20 for all four encoders.

This transfer experiment holds one OnlineRetail evidence configuration fixed while changing only the frozen language geometry. Its purpose is representation transfer, whereas Table 3 uses the validated configuration assigned to each dataset. Under the fixed transfer configuration, NDCG@10 gains are

12.57%, 10.78%, 11.02%, and 13.45% for Sentence-T5, Qwen3, BGE, and E5; the corresponding MRR gains are 16.79%, 14.31%, 14.71%, and 16.74%. The stable direction across all 28 cells demonstrates encoder-agnostic transfer of the state–dynamics interface.

### C.3 Complementarity and Horizon Attribution

Table 10: Absolute gains of the joint model over its strongest single field and over repeated one-step history.

| | Full − best single field | | | | | Full − one-step history | | | | |
|---|---|---|---|---|---|---|---|---|---|---|
| Dataset | H@10 | N@10 | H@20 | N@20 | MRR | H@10 | N@10 | H@20 | N@20 | MRR |
| Beauty | +0.0021 | +0.0029 | +0.0029 | +0.0031 | +0.0028 | +0.0027 | +0.0008 | +0.0029 | +0.0009 | +0.0003 |
| Sports | +0.0030 | +0.0018 | +0.0030 | +0.0018 | +0.0016 | +0.0020 | +0.0018 | +0.0026 | +0.0020 | +0.0018 |
| Toys | +0.0043 | +0.0047 | +0.0031 | +0.0051 | +0.0038 | +0.0021 | +0.0009 | +0.0028 | +0.0010 | +0.0006 |

The left block measures complementarity rather than the gain over a weak reference: every entry remains positive after comparing the complete model with the stronger of semantic-only and successor-only variants. The right block isolates temporal indexing with the same history and candidate support. Its 15 positive entries show that exact horizon alignment contributes beyond simply adding more historical events. The two blocks together match the model factorization: semantic state and successor direction are individually useful, jointly stronger, and most effective when each older source is queried at its actual distance from the target.

### C.4 Complete State and Popularity Controls

Table 11: Complete metrics for the semantic-state and popularity-reference controls.

| Control | Dataset | Variant | H@5 | N@5 | H@10 | N@10 | H@20 | N@20 | MRR |
|---|---|---|---|---|---|---|---|---|---|
| State | Beauty | Last item | 0.0851 | 0.0612 | 0.1183 | 0.0719 | 0.1612 | 0.0827 | 0.0649 |
| | | Uniform history | 0.0859 | 0.0613 | 0.1198 | 0.0722 | 0.1633 | 0.0831 | 0.0649 |
| | | **Recency-weighted** | **0.0859** | **0.0614** | **0.1210** | **0.0727** | **0.1641** | **0.0836** | **0.0652** |
| | Sports | Last item | 0.0488 | 0.0339 | 0.0706 | 0.0409 | 0.0990 | 0.0480 | 0.0373 |
| | | Uniform history | 0.0508 | 0.0355 | 0.0727 | 0.0425 | 0.1015 | 0.0498 | 0.0389 |
| | | **Recency-weighted** | **0.0509** | **0.0357** | **0.0726** | **0.0427** | **0.1016** | **0.0500** | **0.0391** |
| Popularity | Beauty | Uniform reference | 0.0854 | 0.0610 | 0.1205 | 0.0723 | 0.1638 | 0.0832 | 0.0649 |
| | | **Popularity-relative** | **0.0859** | **0.0614** | **0.1210** | **0.0727** | **0.1641** | **0.0836** | **0.0652** |
| | Sports | Uniform reference | 0.0504 | 0.0354 | 0.0727 | 0.0426 | 0.1015 | 0.0499 | 0.0390 |
| | | **Popularity-relative** | **0.0509** | **0.0357** | **0.0726** | **0.0427** | **0.1016** | **0.0500** | **0.0391** |

The complete metrics separate recall from order quality. Recency weighting dominates last-item state throughout and improves every NDCG and MRR entry over uniform history. Popularity correction improves all seven Beauty metrics and six of seven Sports metrics; the remaining Sports HR@10 difference is 0.0001, while NDCG@5, NDCG@10, NDCG@20, and MRR all increase. This pattern is consistent with a calibration term whose principal effect is where supported candidates appear, not expansion of the candidate support.

### C.5 Full-Catalog Effect Map

Table 12 consolidates absolute score changes rather than repeating complete result rows. The matched-dataset block shows that the evidence closure raises every ranking criterion on Beauty, Sports, and Toys. The encoder-transfer block holds the OnlineRetail evidence configuration fixed and changes the frozen language geometry; all 28 entries remain positive. Across both blocks, the largest increments generally occur in HR@5, NDCG@5, NDCG@10, and MRR, while HR@20 also remains positive. This profile is consistent with a decision layer that improves early ordering through language compatibility and horizon-aligned successors while retaining useful broad-cutoff candidates.

Table 12: Absolute ATLASREC-over-anchor changes from the reported means.

| Setting | Dataset / encoder | H@5 | H@10 | H@20 | N@5 | N@10 | N@20 | MRR |
|---|---|---|---|---|---|---|---|---|
| Matched dataset | Beauty | **+.0065** | **+.0089** | **+.0092** | **+.0050** | **+.0058** | **+.0058** | **+.0049** |
| | Sports | **+.0060** | **+.0064** | **+.0070** | **+.0051** | **+.0052** | **+.0053** | **+.0050** |
| | Toys | **+.0126** | **+.0136** | **+.0127** | **+.0113** | **+.0116** | **+.0115** | **+.0110** |
| Encoder transfer | Sentence-T5-XL | **+.0099** | **+.0050** | **+.0045** | **+.0106** | **+.0089** | **+.0089** | **+.0100** |
| | Qwen3-Embedding-0.6B | **+.0095** | **+.0034** | **+.0003** | **+.0095** | **+.0076** | **+.0069** | **+.0085** |
| | BGE-base-en-v1.5 | **+.0081** | **+.0038** | **+.0011** | **+.0092** | **+.0078** | **+.0071** | **+.0088** |
| | E5-base-v2 | **+.0113** | **+.0071** | **+.0046** | **+.0110** | **+.0096** | **+.0090** | **+.0101** |