SEMKOOPREC: SEMANTIC-CONTROLLED KOOPMAN DYNAMICS WITH CANDIDATE-PRESERVING RANKING CLOSURE

Anonymous authors
Paper under double-blind review

ABSTRACT

Sequential recommendation often entangles three roles: interaction history represents current behavior, the latest event redirects it, and ranking translates a forecast into discrete choices. Pretrained semantics usually remain static item features. We introduce a state–control–decision principle in which event-dependent semantic control drives contextual evolution and a bounded decision interface regulates its ranking effect. SEMKOOPREC realizes this principle with a frozen behavioral anchor, persistent–transient Koopman dynamics, multi-horizon representation closure, and candidate-preserving decision closure. Our analysis follows semantic-event perturbations through persistent and exponentially decaying responses to catalog-wide score error, yielding margin-certified candidate stability or a finite anchor-relative intervention bound. Under the comparison protocol shared by every method in Table 1, SEMKOOPREC achieves the best result on all 12 HR/NDCG metrics across three Amazon benchmarks, with HR@10 gains up to 36.50%. A separate strict history-masked evaluation yields positive gains over the matched anchor across five datasets and alternative semantic representations. Matched ablations support event control, persistent–transient evolution, and bounded decision closure, with fewer than 73K additional trainable parameters.

1 INTRODUCTION

User preference evolves through event-induced changes of intent, not only through the accumulation of similar interactions. Recurrent recommenders model this progression through sequential state updates (Hidasi et al., 2016), while self-attentive models capture longer dependencies across the interaction history (Kang & McAuley, 2018). Bidirectional encoders provide another route to contextual sequence modeling (Sun et al., 2019), and large sequential transducers further increase its capacity (Zhai et al., 2024). In parallel, pretrained language representations enrich items with semantic knowledge (Hou et al., 2022); text-centric pretraining aligns this knowledge with interaction sequences (Li et al., 2023a). Most combinations inject semantics before sequence encoding. This improves what an item representation contains, but leaves implicit how an observed event redirects the subsequent user context. A second gap appears at the output: even a meaningful contextual forecast changes neither candidate membership nor discounted order until it is connected to the discrete full-catalog ranking operation.

We view each observed interaction as both an observation and an event. Its occurrence contributes to the contextual state inferred from history, while its semantic meaning supplies an event-dependent input that redirects subsequent evolution. This yields a state–control–decision design principle: a controlled-evolution operator predicts how an event changes context, and a separate bounded decision-closure operator determines how that forecast may modify the anchor ranking. The two operators are constrained by three requirements. **C1: State–control separation.** Accumulated behavioral context and the semantic role of the latest event must enter evolution through distinct paths. **C2: Stable representation closure.** Repeated forecasts must preserve a non-decaying channel, contract transient discrepancies, and propagate transient error without amplification. **C3: Bounded decision intervention.** Forecast evidence must reach candidate membership and discounted order while preserving a trusted anchor core, limiting cross-boundary admissions, and localizing final order changes to the selected candidate set.

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.

SEMKOOPREC provides a structure-preserving realization of this principle over a frozen Markovian Pre-Trained Transformer (MPT) (Xu et al., 2026a). MPT supplies the history-dependent state, while the latest semantic–collaborative item representation controls a persistent–transient Koopman evolution operator. Its identity and contractive blocks retain cumulative event information while stabilizing transient adaptation, and multi-horizon closure aligns repeated evolution with future contexts. Our analysis connects this continuous evolution to discrete recommendation: an event induces persistent and exponentially vanishing components, rollout discrepancy becomes a uniform catalog-wide score bound, and the resulting candidate decision is either certified by its score margin or confined to a finite neighborhood of the anchor. Candidate-preserving closure realizes the latter guarantee while localizing discounted-order correction. On Beauty, Sports, and Toys, SEMKOOPREC ranks first on all 12 HR/NDCG metrics against 19 baselines and variants under a shared comparison protocol (He et al., 2026; Xu et al., 2026b). Matched comparisons attribute gains to control, transition geometry, and decision closure; positive effects persist across five datasets under a separate strict history-masked protocol and across alternative semantic representations, with fewer than 73K additional trainable parameters.

Our contributions are:

- We formulate sequential recommendation through a state–control–decision principle that separates historical behavioral state, current-event semantic control, and bounded intervention in discrete ranking.
- We establish an event-to-decision analysis: event influence separates into persistent and exponentially decaying components, multi-horizon representation error induces a catalog-wide score bound, and a score-margin condition certifies Top-K stability.
- We formulate dynamics-to-ranking adaptation as bounded decision closure that protects an anchor core and gives an explicit anchor-relative candidate-set bound when the margin certificate is unavailable.
- We establish the best results on all 12 primary metrics under a protocol-matched comparison and confirm matched-anchor gains across five datasets under a separate strict history-masked protocol. Matched ablations and semantic-source transfer attribute the effects of control, transition geometry, and decision closure.

2 RELATED WORK

Sequential recommendation began with first-order Markov factorization and recurrent session modeling (Rendle et al., 2010; Hidasi et al., 2016). Self-attentive and bidirectional encoders then captured longer behavioral dependencies (Kang & McAuley, 2018; Sun et al., 2019). Filter-based sequence models replace attention with frequency-selective mixing (Zhou et al., 2022). Recent systems further exploit side information and inter-session frequency structure (Kim et al., 2025; He et al., 2026). Adaptive wavelet filtering provides a multiscale view of sequential behavior (Xu et al., 2026b). In parallel, universal and text-centric pretraining aligns item language with collaborative sequences (Hou et al., 2022; Li et al., 2023a). E4SRec and LLaRA adapt large language representations to sequential recommendation (Li et al., 2023b; Liao et al., 2024). LLM2Rec directly treats language models as item embedding models (He et al., 2025). MPT instead learns transferable sequence-processing capabilities from synthetic Markov chains (Xu et al., 2026a). These lines provide strong behavioral and semantic representations; SEMKOOPREC joins them by retaining MPT as a frozen behavioral anchor and assigning off-line language representations the distinct role of transition control.

The Koopman operator represents nonlinear dynamics through linear evolution in an observable space (Koopman, 1931). Learned finite-dimensional observables make this principle practical (Brunton et al., 2016; Lusch et al., 2018). Neural identification can learn Koopman-invariant subspaces directly from data (Takeishi et al., 2017). Dynamic mode decomposition with control separates autonomous evolution from external input (Proctor et al., 2016). Within our state–control–decision formulation, persistent–transient Koopman dynamics realizes controlled contextual evolution, while a separate bounded interface maps forecast evidence to candidate membership and discounted order. This coupling gives representation evolution and ranking intervention distinct mathematical roles and complementary closure guarantees.

3 METHOD

3.1 CONTROLLED EVOLUTION AND DECISION CLOSURE

Let $\mathcal{U}$ and $\mathcal{V}$ denote the user and item sets, with $M = |\mathcal{V}|$. User $u \in \mathcal{U}$ has the chronological sequence $\mathcal{S}_u = (i_1^u, \ldots, i_{T_u}^u)$, where $i_t^u \in \mathcal{V}$ is the item observed at event index t and T_u is the sequence length. We write $\mathcal{S}_{u,1:t} = (i_1^u, \ldots, i_t^u)$ for the observed prefix. Given this prefix, the task is to rank every candidate $j \in \mathcal{V}$ so that the next item i_{t+1}^u receives a high rank. We omit superscript u when discussing one user. Event index, not wall-clock duration, is the time variable throughout this paper.

The observed prefix determines a behavioral state $\boldsymbol{h}_t \in \mathbb{R}^{d_h}$, while the latest event supplies a semantic–collaborative control $\boldsymbol{u}_t \in \mathbb{R}^{d_s+d_c}$ that specifies an event-dependent displacement. Let $\boldsymbol{e}_j \in \mathbb{R}^{d_h}$ be the shared representation of item j. For any nonzero vector $\boldsymbol{a}$, write $\bar{\boldsymbol{a}} = \boldsymbol{a}/\|\boldsymbol{a}\|_2$ for its ℓ_2 normalization. We separate contextual evolution from ranking intervention as

$$\boldsymbol{q}_t = \mathcal{T}(\boldsymbol{h}_t, \boldsymbol{u}_t), \qquad \ell_t^b(j) = \langle \bar{\boldsymbol{h}}_t, \bar{\boldsymbol{e}}_j \rangle, \qquad \ell_t^d(j) = \langle \boldsymbol{q}_t, \bar{\boldsymbol{e}}_j \rangle, \qquad \ell_t^f = \mathcal{R}(\ell_t^b, \ell_t^d). \tag{1}$$

Here $\mathcal{T}$ is a controlled-evolution operator, $\boldsymbol{q}_t \in \mathbb{R}^{d_h}$ is its normalized forecast in the anchor space, and $\mathcal{R}$ is a bounded decision-closure operator acting on the anchor and dynamic score fields $\ell_t^b, \ell_t^d : \mathcal{V} \to \mathbb{R}$. This decomposition is a design principle rather than an unconstrained template: evolution must separate state from event control and remain stable under composition, while decision closure must preserve a trusted core, bound membership changes, and localize order correction. SEMKOOPREC realizes $\mathcal{T}$ through persistent–transient Koopman dynamics and realizes $\mathcal{R}$ through candidate-preserving ranking closure, as summarized in Figure 1.

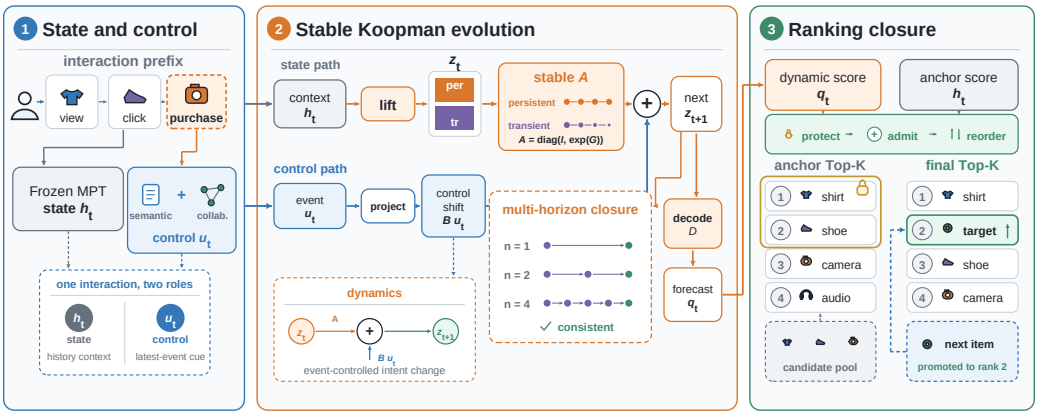

Figure 1: Overview of SEMKOOPREC: frozen-state inference, semantic-controlled persistent–transient Koopman rollouts, and anchor-protected bounded admission with localized Top-K reordering.

3.2 FROZEN SEMANTIC–COLLABORATIVE ANCHOR

For item i, $\boldsymbol{x}_i \in \mathbb{R}^{d_x}$, $\boldsymbol{s}_i \in \mathbb{R}^{d_s}$, and $\boldsymbol{c}_i \in \mathbb{R}^{d_c}$ denote its content, off-line language, and collaborative representations. The anchor maps them to a shared item vector $\boldsymbol{e}_i \in \mathbb{R}^{d_h}$ and applies the frozen causal backbone F_Θ to obtain the behavioral state $\boldsymbol{h}_t = F_\Theta(\boldsymbol{e}_{i_1}, \ldots, \boldsymbol{e}_{i_t}) \in \mathbb{R}^{d_h}$. Thus $\boldsymbol{h}_t$ is the post-interaction context after observing i_t. The subsequent transition is conditioned on a semantic–collaborative control derived from the most recently observed event, so both the lifted state and its control are determined by the observed prefix $i_{1:t}$. The base score is $\ell_t^b(j) = \langle \bar{\boldsymbol{h}}_t, \bar{\boldsymbol{e}}_j \rangle$. All anchor components are fixed before dynamics training, so the full model and no-dynamics control share the same item space, contextual states, and base ranking. Appendix B.1 specifies the semantic–collaborative fusion.

3.3 PERSISTENT–TRANSIENT KOOPMAN REALIZATION

The controlled-evolution operator must distinguish autonomous state transport from event-dependent displacement while remaining stable under repeated composition. We instantiate $\mathcal{T}$ in a learned

observable space. The online lift $\phi_\omega : \mathbb{R}^{d_h} \to \mathbb{R}^{d_z}$, parameterized by ω, converts the detached contextual state into a Koopman observable using a learned matrix $\boldsymbol{W}_\phi \in \mathbb{R}^{d_z \times d_h}$:

$$\boldsymbol{z}_t = \phi_\omega(\mathrm{sg}(\boldsymbol{h}_t)) = \mathrm{LN}(\boldsymbol{W}_\phi \,\mathrm{sg}(\boldsymbol{h}_t)) = [\boldsymbol{z}_t^{\mathrm{per}}; \boldsymbol{z}_t^{\mathrm{tr}}], \tag{2}$$

where $\mathrm{sg}(\cdot)$ stops gradients, LN denotes layer normalization, and $[\cdot;\cdot]$ denotes concatenation. The vectors $\boldsymbol{z}_t^{\mathrm{per}} \in \mathbb{R}^{d_{\mathrm{per}}}$ and $\boldsymbol{z}_t^{\mathrm{tr}} \in \mathbb{R}^{d_{\mathrm{tr}}}$ are respectively the persistent and transient components, with $d_z = d_{\mathrm{per}} + d_{\mathrm{tr}}$. The lifted state evolves under a Koopman operator $\boldsymbol{A} \in \mathbb{R}^{d_z \times d_z}$. The current event provides the normalized semantic–collaborative control $\boldsymbol{u}_t$; a learned matrix $\boldsymbol{B} \in \mathbb{R}^{d_z \times (d_s + d_c)}$ injects it into the transition:

$$\boldsymbol{u}_t = \frac{1}{\sqrt{2}}[\bar{\boldsymbol{s}}_{i_t}; \bar{\boldsymbol{c}}_{i_t}] \in \mathbb{R}^{d_s+d_c}, \qquad \widehat{\boldsymbol{z}}_t^{(1)} = \boldsymbol{A}\boldsymbol{z}_t + \boldsymbol{B}\boldsymbol{u}_t \in \mathbb{R}^{d_z}. \tag{3}$$

Both $\boldsymbol{z}_t$ and $\boldsymbol{u}_t$ are determined by the observed prefix $i_{1:t}$. The operator term $\boldsymbol{A}\boldsymbol{z}_t$ carries the history-dependent latent context, while $\boldsymbol{B}\boldsymbol{u}_t$ contributes an event-dependent semantic–collaborative displacement. Their sum therefore combines history-dependent autonomous evolution with an event-dependent displacement to predict the next contextual state.

We construct $\boldsymbol{A}$ from an unconstrained matrix $\boldsymbol{J} \in \mathbb{R}^{d_{\mathrm{tr}} \times d_{\mathrm{tr}}}$ and an unconstrained vector $\boldsymbol{\xi} \in \mathbb{R}^{d_{\mathrm{tr}}}$. Let $\eta_s > 0$ scale the rotational component and let $\delta_{\mathrm{stab}} > 0$ enforce strict decay:

$$\boldsymbol{\Omega} = \eta_s(\boldsymbol{J} - \boldsymbol{J}^\top) \in \mathbb{R}^{d_{\mathrm{tr}} \times d_{\mathrm{tr}}}, \qquad \boldsymbol{\gamma} = \mathrm{softplus}(\boldsymbol{\xi}) + \delta_{\mathrm{stab}}\mathbf{1} \in \mathbb{R}^{d_{\mathrm{tr}}}_{>0}, \tag{4}$$

$$\boldsymbol{G} = \boldsymbol{\Omega} - \mathrm{diag}(\boldsymbol{\gamma}), \qquad \boldsymbol{A} = \mathrm{blkdiag}(\boldsymbol{I}_{d_{\mathrm{per}}}, \exp(\boldsymbol{G})) \in \mathbb{R}^{d_z \times d_z}. \tag{5}$$

Here $\mathrm{softplus}$ is the standard positive activation, diag constructs a diagonal matrix, $\mathrm{blkdiag}$ constructs a block-diagonal matrix, $\boldsymbol{I}_{d_{\mathrm{per}}}$ is the identity matrix, $\mathbf{1}$ is the all-ones vector, and $\exp(\cdot)$ denotes the matrix exponential. The identity block provides a non-decaying channel under autonomous evolution while permitting new control inputs to accumulate in those coordinates. The transient generator combines rotation through the skew-symmetric $\boldsymbol{\Omega}$ with coordinate-wise dissipation through $\boldsymbol{\gamma}$.

Let $\gamma_{\min} = \min_r \gamma_r > 0$, $\rho = e^{-\gamma_{\min}} < 1$, and partition $\boldsymbol{B} = [\boldsymbol{B}_{\mathrm{per}}; \boldsymbol{B}_{\mathrm{tr}}]$ consistently with the latent coordinates. The block construction determines not only autonomous stability but also how event-dependent control differences propagate.

Theorem 1 (persistent–transient event response). Consider two trajectories with the same pre-event state and identical future controls except that $\boldsymbol{u}_t$ changes by $\Delta\boldsymbol{u}_t$. For every event distance $k \geq 0$,

$$\Delta\boldsymbol{z}_{t+k+1} = \begin{bmatrix} \boldsymbol{B}_{\mathrm{per}}\Delta\boldsymbol{u}_t \\ e^{k\boldsymbol{G}}\boldsymbol{B}_{\mathrm{tr}}\Delta\boldsymbol{u}_t \end{bmatrix}, \qquad \left\|\Delta\boldsymbol{z}^{\mathrm{tr}}_{t+k+1}\right\|_2 \leq \rho^k \left\|\boldsymbol{B}_{\mathrm{tr}}\right\|_2 \left\|\Delta\boldsymbol{u}_t\right\|_2. \tag{6}$$

Thus an interaction induces a retained displacement together with an exponentially vanishing transient response. Appendix A extends this result to arbitrary control-sequence mismatch and derives same-input and bounded-input stability as special cases; the identity block deliberately accumulates event evidence rather than contracting every direction.

3.4 Multi-Horizon Representation Closure

A useful Koopman observable must remain predictive under repeated applications of the same transition. We use a target lift ϕ_{ω^-} to provide stop-gradient successor observables. Let $\mathcal{H} \subset \mathbb{N}_+$ be the rollout-horizon set. Starting from $\widehat{\boldsymbol{z}}_t^{(0)} = \boldsymbol{z}_t$, the n-step representation rollout is

$$\widehat{\boldsymbol{z}}_t^{(m+1)} = \boldsymbol{A}\widehat{\boldsymbol{z}}_t^{(m)} + \boldsymbol{B}\boldsymbol{u}_{t+m}, \quad m \in \{0, \ldots, n-1\}, \qquad \boldsymbol{z}^\star_{t+n} = \phi_{\omega^-}(\mathrm{sg}(\boldsymbol{h}_{t+n})). \tag{7}$$

Let $\cos(\cdot,\cdot)$ denote cosine similarity and let $\mathrm{Huber}(\cdot,\cdot)$ denote the element-wise smooth ℓ_1 loss averaged over valid coordinates. For horizon weights $w_n > 0$ and cosine-distance coefficient $\lambda_{\cos} > 0$, the latent closure loss is

$$\mathcal{L}_{\mathrm{lat}} = \frac{1}{\sum_{n\in\mathcal{H}} w_n} \sum_{n \in \mathcal{H}} w_n \left[\mathrm{Huber}(\widehat{\boldsymbol{z}}_t^{(n)}, \boldsymbol{z}^\star_{t+n}) + \lambda_{\cos}\big(1 - \cos(\widehat{\boldsymbol{z}}_t^{(n)}, \boldsymbol{z}^\star_{t+n})\big)\right]. \tag{8}$$

The decoder $\boldsymbol{D} \in \mathbb{R}^{d_h \times d_z}$ returns and explicitly normalizes the one-step forecast in the anchor space:

$$\boldsymbol{q}_t = \frac{\boldsymbol{D}\widehat{\boldsymbol{z}}_t^{(1)}}{\left\|\boldsymbol{D}\widehat{\boldsymbol{z}}_t^{(1)}\right\|_2} \in \mathbb{R}^{d_h}, \qquad \ell_t^d(j) = \langle \boldsymbol{q}_t, \bar{\boldsymbol{e}}_j \rangle \in \mathbb{R}, \tag{9}$$

where the decoded vector is assumed nonzero and $\ell_t^d(j)$ is the resulting cosine score of candidate j. The decoder aligns the forecast with the next frozen contextual state, while the deployed score uses only the one-step forecast conditioned on the observed prefix $i_{1:t}$. Appendix B.1 gives the decoder alignment term.

Define the target-space residual $\boldsymbol{\varepsilon}_{t+m} = \boldsymbol{z}^\star_{t+m+1} - (\boldsymbol{A}\boldsymbol{z}^\star_{t+m} + \boldsymbol{B}\boldsymbol{u}_{t+m})$ and rollout discrepancy $\boldsymbol{e}_t^{(m)} = \widehat{\boldsymbol{z}}_t^{(m)} - \boldsymbol{z}^\star_{t+m}$. At every horizon, the predicted and target dynamic score fields use the same explicit decoder normalization:

$$\widehat{\ell}^d_{t,n}(j) = \left\langle \frac{\boldsymbol{D}\widehat{\boldsymbol{z}}_t^{(n)}}{\left\|\boldsymbol{D}\widehat{\boldsymbol{z}}_t^{(n)}\right\|_2}, \bar{\boldsymbol{e}}_j \right\rangle, \qquad \ell^{d,\star}_{t,n}(j) = \left\langle \frac{\boldsymbol{D}\boldsymbol{z}^\star_{t+n}}{\left\|\boldsymbol{D}\boldsymbol{z}^\star_{t+n}\right\|_2}, \bar{\boldsymbol{e}}_j \right\rangle.$$

Theorem 2 (multi-horizon representation-to-ranking stability). Let $E_{t,n}$ be the full-state bound in Appendix Equation 26, which combines cumulative persistent residuals with geometrically weighted transient residuals. For every $n \geq 1$, $\left\|\boldsymbol{e}_t^{(n)}\right\|_2 \leq E_{t,n}$. For nonzero decoded forecasts, define $r_{t,n} = \min\{\left\|\boldsymbol{D}\widehat{\boldsymbol{z}}_t^{(n)}\right\|_2, \left\|\boldsymbol{D}\boldsymbol{z}^\star_{t+n}\right\|_2\}$ and let $\Delta^{(K)}_{t,n}$ be the target score gap between ranks K and $K+1$. Then

$$\begin{aligned} \left\|\widehat{\ell}^d_{t,n} - \ell^{d,\star}_{t,n}\right\|_\infty &\leq \varepsilon_{t,n} := 2\left\|\boldsymbol{D}\right\|_2 E_{t,n}/r_{t,n}, \\ \Delta^{(K)}_{t,n} > 2\varepsilon_{t,n} &\Longrightarrow \mathrm{TopK}_K(\widehat{\ell}^d_{t,n}) = \mathrm{TopK}_K(\ell^{d,\star}_{t,n}). \end{aligned} \tag{10}$$

Decoder normalization converts the representation bound into a uniform full-catalog score bound; a sufficiently separated boundary certifies the forecast-induced Top-K set. When the certificate is unavailable, the ranking closure below supplies an unconditional anchor-relative bound. Appendix A proves the result and gives the event-to-score specialization of Theorem 1.

3.5 Candidate-Preserving Decision-Closure Realization

Candidate-preserving decision closure treats forecast-to-ranking adaptation as one constrained decision problem. Its operator $\mathcal{R}$ is defined by three structural invariants: preservation of a trusted anchor core, bounded admission of forecast-supported candidates, and membership-preserving order correction restricted to the selected set. These invariants determine what forecast evidence may change and organize the realization as protect $\rightarrow$ bounded admit $\rightarrow$ local reorder. Write $\mathrm{TopK}_k(f)$ for the k catalog items with largest score under f. Formally, let

$$\mathcal{P}_t = \mathrm{TopK}_{K_b}(\ell_t^b), \qquad \mathcal{A}_t \subseteq \mathcal{V} \setminus \mathcal{P}_t, \qquad |\mathcal{A}_t| = B_t \leq K_c - K_b, \tag{11}$$

where $\mathcal{P}_t$ is the protected anchor core and $\mathcal{A}_t$ is the finite set admitted by the bounded forecast rule. The correction returns $\ell_t^{\mathrm{pre}} : \mathcal{V} \rightarrow \mathbb{R}$ and enforces

$$\begin{aligned} \min_{p \in \mathcal{P}_t} \ell_t^{\mathrm{pre}}(p) &> \max_{j \notin \mathcal{P}_t \cup \mathcal{A}_t} \ell_t^{\mathrm{pre}}(j), \\ \ell_t^{\mathrm{pre}}(a) \geq \ell_t^b(a) \quad (a \in \mathcal{A}_t), &\qquad \ell_t^{\mathrm{pre}}(j) = \ell_t^b(j) \quad (j \notin \mathcal{P}_t \cup \mathcal{A}_t). \end{aligned} \tag{12}$$

Thus forecast evidence may change protected scores and admit selected candidates, while all remaining catalog scores retain their anchor values. Appendices B.2 and B.3 give the bounded admission realization.

Let $\mathcal{C}_t = \mathrm{TopK}_{K_c}(\ell_t^{\mathrm{pre}})$ and let $\mathcal{C}_t^b = \mathrm{TopK}_{K_c}(\ell_t^b)$ be the corresponding anchor set. The admission rule selects the candidate membership; the following order stage acts only inside $\mathcal{C}_t$.

For each $j \in \mathcal{C}_t$, a standardized ranking representation $\boldsymbol{\psi}_t(j) \in \mathbb{R}^{d_\psi}$ aggregates causal score and trajectory evidence; Appendix B.4 gives its feature groups. A listwise scorer $g_{\boldsymbol{\theta}_{\mathrm{ord}}} : \mathbb{R}^{d_\psi} \rightarrow \mathbb{R}$

produces the learned correction. Writing $Z_{\mathcal{C}_t}[\cdot]$ for standardization over $\mathcal{C}_t$, the within-set order priority is

$$\chi_t^{\mathrm{ord}}(j) = Z_{\mathcal{C}_t}[\ell_t^{\mathrm{pre}}](j) + \lambda_{\mathrm{ord}} Z_{\mathcal{C}_t}[g_{\boldsymbol{\theta}_{\mathrm{ord}}}(\boldsymbol{\psi}_t(\cdot))](j), \qquad j \in \mathcal{C}_t, \tag{13}$$

where $\lambda_{\mathrm{ord}} \geq 0$ controls the order correction. The deployed score ℓ_t^f assigns the descending multiset of pre-order values on $\mathcal{C}_t$ to candidates in descending χ_t^{ord} order and leaves scores outside $\mathcal{C}_t$ unchanged.

Theorem 3 (bounded decision intervention). Let $\mathcal{C}_t^f = \mathrm{TopK}_{K_c}(\ell_t^f)$. Under a strict pre-order boundary at K_c, and under member-preserving tie resolution otherwise, bounded admission and score-multiset reassignment satisfy

$$\mathcal{P}_t \subseteq \mathcal{C}_t^f = \mathcal{C}_t, \qquad |\mathcal{C}_t^f \setminus \mathcal{C}_t^b| \leq B_t, \qquad |\mathcal{C}_t^f \triangle \mathcal{C}_t^b| \leq 2B_t. \tag{14}$$

Thus the final set remains within symmetric-difference radius $2B_t$ of the anchor, while the order stage preserves membership and the selected score multiset. In single-target evaluation, reassignment preserves HR@K_c; moving the target earlier cannot decrease NDCG at cutoffs up to K_c. For target $i_{t+1} \in \mathcal{C}_t$, the scorer minimizes

$$\mathcal{L}_{\mathrm{ord}} = -\log \frac{\exp(g_{\boldsymbol{\theta}_{\mathrm{ord}}}(\boldsymbol{\psi}_t(i_{t+1})))}{\sum_{j \in \mathcal{C}_t} \exp(g_{\boldsymbol{\theta}_{\mathrm{ord}}}(\boldsymbol{\psi}_t(j)))}. \tag{15}$$

This likelihood promotes the ground-truth next item within $\mathcal{C}_t$. The admission and reassignment operations together instantiate $\mathcal{R}$. Appendix B.4 gives the feature groups and reassignment.

3.6 Optimization and Complexity

The frozen anchor is not updated. Besides the latent closure loss $\mathcal{L}_{\mathrm{lat}}$ in Equation 8 and the bounded listwise loss $\mathcal{L}_{\mathrm{ord}}$ in Equation 15, let $\mathcal{L}_{\mathrm{dec}}$ denote decoded next-state alignment and $\mathcal{L}_{\mathrm{item}}$ denote next-item cross-entropy under the dynamic score. The complete objective is

$$\mathcal{L} = \mathcal{L}_{\mathrm{lat}} + \lambda_{\mathrm{dec}}\mathcal{L}_{\mathrm{dec}} + \lambda_{\mathrm{item}}\mathcal{L}_{\mathrm{item}} + \lambda_{\mathrm{ord}}\mathcal{L}_{\mathrm{ord}}, \tag{16}$$

where all coefficients are nonnegative.

For one event, latent transition and decoding cost is $O(d_z^2 + d_z(d_s + d_c) + d_h d_z)$. Full-catalog inference adds one $O(Md_h)$ dynamic score product to the anchor score product, and ranking closure costs $O(K_c d_\psi)$ on its bounded candidate set. Computing $\exp(\boldsymbol{G})$ costs $O(d_{\mathrm{tr}}^3)$ only when the operator changes and is cached at evaluation. Theorems 1–3 establish the event-to-decision chain; Appendix A provides the proofs and special cases.

4 Experiments

Our evaluation follows the state–control–decision chain. We test end-to-end ranking under a shared protocol, transfer across stricter masking, domains, and semantic sources, and the roles of aligned control, structured evolution, and bounded decision closure. Together, the experiments ask whether interaction semantics can control stable behavioral evolution and whether its forecast produces supported ranking changes.

4.1 Experimental Setup

Datasets and protocols. We use two evaluation settings for different questions. The *comparison setting* evaluates Amazon Beauty, Sports & Outdoors, and Toys & Games under the shared DIFF/FreqRec convention, which retains historical items in the catalog and enables same-setting comparison with prior methods in Table 1. The separate *strict setting* completely masks pre-endpoint history and compares against matched anchors on these datasets plus Yelp2018-10core and OnlineRetail in Table 2. Both use chronological leave-one-out and full-catalog ranking. Because their candidate masks differ, results are never compared across the two settings. Appendix B gives protocol details.

Baselines. The baselines are GRU4Rec (Hidasi et al., 2016), Caser (Tang & Wang, 2018), SASRec (Kang & McAuley, 2018), BERT4Rec (Sun et al., 2019), DuoRec (Qiu et al., 2022), and ICLRec (Chen et al., 2022); CL4SRec (Xie et al., 2022), MAERec (Ye et al., 2023), SASRecF (Kang & McAuley, 2018; Kim et al., 2025), MSSR (Lin et al., 2024), FamouSRec (Zhang et al., 2025a), and FMLPRec (Zhou et al., 2022); FEARec (Du et al., 2023b), SLIME4Rec (Du et al., 2023a), BSARec (Shin et al., 2024), DIFF (Kim et al., 2025), FreqRec-S (He et al., 2026), FreqRec-P (He et al., 2026), and WEARec (Xu et al., 2026b).

Evaluation. We report full-catalog HR and NDCG at 10 and 20 as test means with sample standard deviations, with paired significance tests on per-user ranks. Frozen MPT provides the behavioral anchor, Qwen3-Embedding-0.6B (Zhang et al., 2025b) supplies off-line semantics, and the complete adaptation adds at most 72,708 trainable parameters.

4.2 Comparison Setting: Same-Protocol Baselines

Table 1: Comparison-setting results; all methods follow the same data, 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	0.0304	0.0147	0.0527	0.0203	0.0187	0.0101	0.0303	0.0131	0.0211	0.0106	0.0348	0.0140
Caser	0.0225	0.0108	0.0403	0.0153	0.0163	0.0080	0.0260	0.0104	0.0161	0.0079	0.0268	0.0106
SASRec	0.0531	0.0283	0.0823	0.0356	0.0298	0.0159	0.0459	0.0200	0.0652	0.0366	0.0929	0.0435
BERT4Rec	0.0705	0.0387	0.1073	0.0480	0.0428	0.0229	0.0649	0.0284	0.0635	0.0353	0.0939	0.0430
DuoRec	0.0965	0.0584	0.1313	0.0671	0.0569	0.0331	0.0791	0.0387	0.0816	0.0515	0.1080	0.0582
ICLRec	0.0726	0.0400	0.1055	0.0483	0.0422	0.0227	0.0632	0.0280	0.0826	0.0473	0.1137	0.0552
CL4SRec	0.0706	0.0417	0.0990	0.0488	0.0375	0.0201	0.0575	0.0251	0.0863	0.0522	0.1143	0.0592
MAERec	0.0789	0.0472	0.1094	0.0548	0.0435	0.0239	0.0645	0.0292	0.0823	0.0499	0.1108	0.0570
SASRecF	0.0849	0.0421	0.1187	0.0506	0.0488	0.0231	0.0711	0.0287	0.0875	0.0571	0.1082	0.0623
MSSR	0.0897	0.0448	0.1281	0.0545	0.0540	0.0255	0.0785	0.0317	0.0982	0.0608	0.1201	0.0684
FMLPRec	0.0559	0.0291	0.0869	0.0369	0.0336	0.0183	0.0525	0.0231	0.0671	0.0365	0.0974	0.0441
FamouSRec	0.0838	0.0497	0.1146	0.0575	0.0424	0.0244	0.0632	0.0297	–	–	–	–
FEARec	0.0982	0.0601	0.1352	0.0694	0.0589	0.0343	0.0836	0.0405	0.1005	0.0606	0.1357	0.0708
SLIME4Rec	0.1006	0.0601	0.1381	0.0696	0.0611	0.0357	0.0869	0.0421	–	–	–	–
BSARec	0.1008	0.0611	0.1373	0.0703	0.0612	0.0360	0.0858	0.0422	0.1023	0.0610	0.1379	0.0713
DIFF	0.0940	0.0526	0.1301	0.0631	0.0580	0.0310	0.0829	0.0381	0.0969	0.0607	0.1210	0.0703
FreqRec-S	0.0962	0.0589	0.1299	0.0674	0.0571	0.0334	0.0825	0.0398	0.1046	0.0655	0.1459	0.0734
FreqRec-P	0.0989	0.0601	0.1359	0.0686	0.0583	0.0342	0.0859	0.0401	0.1044	0.0653	0.1468	0.0735
WEARec	0.1041	0.0614	0.1391	0.0703	0.0631	0.0367	0.0895	0.0433	–	–	–	–
SemKoopRec	**0.1325**	**0.0752**	**0.1515**	**0.0799**	**0.0781**	**0.0421**	**0.0903**	**0.0451**	**0.1428**	**0.0831**	**0.1593**	**0.0873**
Improve	+27.26%	+22.45%	+8.93%	+13.70%	+23.76%	+14.74%	+0.86%	+4.21%	+36.50%	+26.94%	+8.49%	+18.74%

Within the comparison setting, every entry in Table 1 uses the same preprocessing and splits, full-catalog candidate convention, historical-item treatment, and HR/NDCG definitions at identical cutoffs. This is the like-for-like SOTA comparison; no strict-setting result is mixed into the table. Against strong sequential, contrastive, frequency-aware, and side-information recommenders, SemKoopRec ranks first on all 12 metrics. Its HR@10 gains over the strongest reported competitor are 27.26%, 23.76%, and 36.50% on Beauty, Sports, and Toys, and its NDCG@10 gains are 22.45%, 14.74%, and 26.94%. The lead is therefore not tied to one dataset, cutoff, or accuracy notion: hit-based and position-sensitive metrics improve jointly at both reported depths. The larger margins at the tighter cutoff further match the proposed decision geometry, in which the predicted state promotes supported candidates and refines their early order while largely retaining the broader candidate envelope.

4.3 Robustness and Transfer across Protocols, Domains, and Semantic Sources

Table 2: Strict-setting results against matched anchors with complete history masking.

	HR@10			NDCG@10		
Dataset	Anchor	Full	Improve	Anchor	Full	Improve
Beauty	0.1191	**0.1218±0.0005**	+2.25%	0.0700	**0.0716±0.0003**	+2.34%
Sports	0.0668	**0.0726±0.0003**	+8.68%	0.0379	**0.0400±0.0001**	+5.59%
Toys	0.1259	**0.1301±0.0007**	+3.37%	0.0741	**0.0773±0.0003**	+4.32%
Yelp2018-10core	0.0598	**0.0680±0.0004**	+13.78%	0.0307	**0.0331±0.0001**	+7.78%
OnlineRetail	0.2450	**0.2601±0.0026**	+6.13%	0.1456	**0.1500±0.0008**	+2.99%

Table 3: Relative improvement (%) over the matched anchor under strict history masking.

Dataset	HR@5	N@5	HR@10	N@10	HR@20	N@20	MRR
Beauty	**2.97**	**2.77**	**2.25**	**2.34**	**0.00**	**1.19**	**1.84**
Sports	**2.44**	**2.03**	**8.68**	**5.59**	**0.00**	**1.32**	**1.76**
Toys	**1.16**	**3.53**	**3.37**	**4.32**	**0.00**	**2.48**	**4.05**
Yelp2018-10core	**0.00**	**0.00**	**13.78**	**7.78**	**0.00**	**0.41**	**0.49**
OnlineRetail	**0.00**	**0.00**	**6.13**	**2.99**	**0.00**	**0.18**	**0.20**

Table 2 switches to the separate strict setting; it is not a continuation of the baseline comparison in Table 1. Complete history masking removes every previously interacted item before scoring, and the matched anchor keeps the backbone, representations, data, candidate mask, and evaluator fixed. HR@10 improves by 2.25–13.78% and NDCG@10 by 2.34–7.78% on every dataset. The positive changes span review sequences with different densities and the transaction-oriented OnlineRetail corpus, indicating that the adaptation is not specific to one interaction regime. Table 3 localizes the effect: gains concentrate at tighter cutoffs, NDCG@20 and MRR remain positive, and HR@20 is invariant. This metric profile matches a bounded correction that improves the ordering of retained candidates rather than broadly replacing the candidate region. Appendix Table 11 reports all absolute operating points. These values are not compared with Table 1 because the candidate spaces differ.

Table 4: Semantic-source transfer under strict history masking.

Dataset	Semantic source	HR@10	NDCG@10	Δ
Beauty	BGE-M3	**0.1168±0.0021**	**0.0684±0.0005**	**+2.64%**
Beauty	E5-large-v2	**0.1158±0.0028**	**0.0683±0.0009**	**+2.62%**
Beauty	TinyLlama-1.1B-Chat	**0.1174±0.0039**	**0.0689±0.0019**	**+2.54%**
Yelp	BGE-M3	**0.0685±0.0011**	**0.0334±0.0006**	**+7.00%**
Yelp	E5-large-v2	**0.0681±0.0004**	**0.0333±0.0002**	**+7.33%**

Protocol transfer tests whether the decision mechanism survives a harder candidate space; semantic-source transfer tests whether the control formulation survives a different language representation. Table 4 replaces Qwen3-Embedding with BGE-M3 (Chen et al., 2024), E5-large-v2 (Wang et al., 2022), and the causal decoder TinyLlama-1.1B-Chat (Zhang et al., 2024). Every source–dataset combination improves over its source-matched static anchor: NDCG@10 rises by 2.54–2.64% on Beauty and 7.00–7.33% on Yelp. The narrow spread of relative gains within each dataset is notable because the sources vary in scale, training objective, and whether they encode text bidirectionally or causally. Their consistent advantage supports the intended abstraction: language representations supply event-conditioned control coordinates, whereas the transition and decision mechanisms remain shared rather than being specialized to one embedding space.

4.4 Mechanism Attribution: Evolution and Decision Closure

Table 5: Complete pathway versus its matched static anchor.

Dataset	Metric	Anchor	Full	Δ (relative)	95% CI / p_H
Beauty	HR@10	0.1022	**0.1326**	0.0304 (29.70%)	[0.0281,0.0327] / 0.0006
	NDCG@10	0.0472	**0.0753**	0.0281 (59.57%)	[0.0266,0.0295] / 0.0006
Sports	HR@10	0.0595	**0.0781**	0.0186 (31.26%)	[0.0172,0.0201] / 0.0006
	NDCG@10	0.0276	**0.0421**	0.0145 (52.56%)	[0.0137,0.0154] / 0.0006
Toys	HR@10	0.1096	**0.1434**	0.0338 (30.89%)	[0.0313,0.0366] / 0.0006
	NDCG@10	0.0487	**0.0834**	0.0347 (71.28%)	[0.0329,0.0363] / 0.0006

Complete dynamics-to-ranking pathway. Table 5 first isolates the complete pathway by keeping the frozen MPT, item features, fusion, data, and evaluator fixed. HR@10 rises by 29.70–31.26% and NDCG@10 by 52.56–71.28%; all confidence intervals are positive and all six effects remain significant at $p = 0.0006$ after Holm correction. The larger relative gains in NDCG than HR show that the pathway does more than place the target inside the candidate region: it systematically moves relevant items toward positions with higher discounted utility. The comparable HR gains across all three domains, together with the larger NDCG response, identify a repeatable combination of candidate recovery and early-rank refinement. We next decompose this end-to-end gain into control, evolution, and decision effects.

Table 6: Control-signal attribution under the comparison protocol.

Variant	Beauty NDCG@10	Sports NDCG@10	Toys NDCG@10
No control: $\boldsymbol{u}_t = \mathbf{0}$	0.0621	0.0345	0.0743
Shuffled control	0.0594	0.0332	0.0728
Collaborative-only control	0.0708	0.0391	0.0806
Semantic-only control	0.0689	0.0380	0.0792
Full semantic–collaborative control	**0.0753**	**0.0421**	**0.0834**

Control alignment and complementarity. Table 6 identifies what drives the transition. Shuffled control preserves the control distribution but breaks trajectory alignment and is worse than zero control everywhere, showing that the operator responds to event identity; this rules out a distribution-only explanation. Each single view helps, while their combination is strongest, improving NDCG@10 by 12.25–22.03% over zero control and 3.47–7.67% over the stronger single view. The same ordering across all three domains makes the roles identifiable: collaborative evidence supplies plausible movement, while language semantics specifies the event meaning that redirects it.

Table 7: Equal-parameter transition comparison in NDCG@10.

Dataset	Residual MLP	Koopman	Relative Δ	95% CI
Beauty	0.0582	**0.0611**	+4.95%	[0.0017,0.0040]
Sports	0.0329	**0.0341**	+3.74%	[0.0006,0.0019]
Toys	0.0630	**0.0677**	+7.54%	[0.0035,0.0060]

Persistent–transient transition geometry. Table 7 replaces Equation 5 with a residual two-layer SiLU MLP while fixing every other component; the alternatives contain 72,371 and 72,399 trainable parameters. Koopman improves NDCG@10 by 3.74–7.54%, every bootstrap interval excludes zero, and Holm-adjusted $p = 0.0006$, so comparable capacity alone does not recover the benefit. The common ordering across three datasets favors structured evolution over unrestricted capacity. Together with Table 6, this separates signal from evolution: aligned control sets the direction of change, while the persistent–transient operator determines what is retained or contracted.

Table 8: Decision-closure attribution in NDCG@10.

Variant	Beauty	Sports	Toys
w/o recurrence state	0.0694	0.0390	0.0677
Anchor + recurrence only	0.0751	0.0416	0.0804
Full SEMKOOPREC	**0.0753**	**0.0421**	**0.0834**

Candidate-preserving decision closure. Table 8 closes the chain from forecast to ranking. Removing recurrence coordinates lowers NDCG@10 everywhere ($p = 0.0002$), while the full pathway is strongest throughout. The recurrence-only variant confirms complementary causal structure, whereas the complete model adds forecasted evolution. Table 3 and Appendix Table 12 show invariant HR@20 across all five strict-setting datasets and closure variants despite tighter-cutoff gains. This pattern is consistent with Theorem 3: bounded membership intervention followed by within-set correction moves relevant items upward while preserving broad Top-20 inclusion.

5 CONCLUSION

We formulate sequential recommendation through a state–control–decision principle: history determines behavioral state, event semantics redirect contextual evolution, and a bounded interface regulates ranking intervention. SEMKOOPREC realizes this principle with persistent–transient Koopman dynamics and a decision closure that protects an anchor core, bounds forecast-supported admission, and localizes order correction. The analysis links continuous evolution to discrete recommendation: event influence separates into retained and decaying components, rollout error bounds catalog-wide scores, margins certify stable decisions, and the admission budget confines ambiguous cases. The complete model ranks first on all 12 protocol-matched metrics and significantly improves its static anchor. Control, transition, and decision ablations support the three roles separately, while gains remain positive across five strict-setting datasets and alternative semantic representations.

Reproducibility Statement

Appendix A provides the theoretical analysis and proofs. Appendix B documents the model and evaluation protocols, and Appendix B.6 records the architecture, optimization, and decision-closure settings required for reproduction. Appendices C and D provide additional experimental results and efficiency details, respectively.

AI Use Statement

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

References

Steven L. Brunton, Bingni W. Brunton, Joshua L. Proctor, and J. Nathan Kutz. Koopman invariant subspaces and finite linear representations of nonlinear dynamical systems for control. *PLOS ONE*, 11(2):e0150171, 2016. doi: 10.1371/journal.pone.0150171.

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/.

Yongjun Chen, Zhiwei Liu, Jia Li, Julian McAuley, and Caiming Xiong. Intent contrastive learning for sequential recommendation. In *Proceedings of the ACM Web Conference 2022*, pp. 2172–2182, 2022. doi: 10.1145/3485447.3512090.

Xinyu Du, Huanhuan Yuan, Pengpeng Zhao, Junhua Fang, Guanfeng Liu, Yanchi Liu, Victor S. Sheng, and Xiaofang Zhou. Contrastive enhanced slide filter mixer for sequential recommendation. In *39th IEEE International Conference on Data Engineering*, pp. 2673–2685, 2023a. doi: 10.1109/ICDE55515.2023.00205.

Xinyu Du, Huanhuan Yuan, Pengpeng Zhao, Jianfeng Qu, Fuzhen Zhuang, Guanfeng Liu, Yanchi Liu, and Victor S. Sheng. Frequency enhanced hybrid attention network for sequential recommendation. In *Proceedings of the 46th International ACM SIGIR Conference on Research and Development in Information Retrieval*, pp. 78–88, 2023b. doi: 10.1145/3539618.3591689.

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.

Bernard O. Koopman. Hamiltonian systems and transformation in Hilbert space. *Proceedings of the National Academy of Sciences*, 17(5):315–318, 1931. doi: 10.1073/pnas.17.5.315.

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.

Xiaolin Lin, Jinwei Luo, Junwei Pan, Weike Pan, Zhong Ming, Xun Liu, Shudong Huang, and Jie Jiang. Multi-sequence attentive user representation learning for side-information integrated sequential recommendation. In *Proceedings of the 17th ACM International Conference on Web Search and Data Mining*, pp. 414–423, 2024. doi: 10.1145/3616855.3635815.

Bethany Lusch, J. Nathan Kutz, and Steven L. Brunton. Deep learning for universal linear embeddings of nonlinear dynamics. *Nature Communications*, 9:4950, 2018. doi: 10.1038/s41467-018-07210-0.

Joshua L. Proctor, Steven L. Brunton, and J. Nathan Kutz. Dynamic mode decomposition with control. *SIAM Journal on Applied Dynamical Systems*, 15(1):142–161, 2016. doi: 10.1137/15M1013857.

Ruihong Qiu, Zi Huang, Hongzhi Yin, and Zijian Wang. Contrastive learning for representation degeneration problem in sequential recommendation. In *Proceedings of the Fifteenth ACM International Conference on Web Search and Data Mining*, pp. 813–823, 2022. doi: 10.1145/3488560.3498433.

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.

Yehjin Shin, Jeongwhan Choi, Hyowon Wi, and Noseong Park. An attentive inductive bias for sequential recommendation beyond the self-attention. In *Proceedings of the AAAI Conference on Artificial Intelligence*, volume 38, pp. 8984–8992, 2024. doi: 10.1609/aaai.v38i8.28747.

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.

Naoya Takeishi, Yoshinobu Kawahara, and Takehisa Yairi. Learning Koopman invariant subspaces for dynamic mode decomposition. In *Advances in Neural Information Processing Systems*, volume 30, 2017. URL https://proceedings.neurips.cc/paper/2017/hash/3a835d3215755c435ef4fe9965a3f2a0-Abstract.html.

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.

Yaowen Ye, Lianghao Xia, and Chao Huang. Graph masked autoencoder for sequential recommendation. In *Proceedings of the 46th International ACM SIGIR Conference on Research and Development in Information Retrieval*, pp. 321–330, 2023. doi: 10.1145/3539618.3591692.

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.

Junjie Zhang, Ruobing Xie, Hongyu Lu, Wenqi Sun, Wayne Xin Zhao, Yu Chen, and Zhanhui Kang. Frequency-augmented mixture-of-heterogeneous-experts framework for sequential recommendation. In *Proceedings of the ACM on Web Conference 2025*, pp. 2596–2605, 2025a. doi: 10.1145/3696410.3714663.

Peiyuan Zhang, Guangtao Zeng, Tianduo Wang, and Wei Lu. TinyLlama: An open-source small language model. *arXiv preprint arXiv:2401.02385*, 2024.

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*, 2025b. 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 Theoretical Analysis

This section supplies the complete event-to-decision analysis summarized by Theorems 1–3. The argument has three steps: the generator controls event influence in the observable space, decoder normalization transfers representation error to every catalog score, and the decision interface converts that score control into either an exact margin certificate or a finite anchor-relative intervention bound.

A.1 Two Basic Operator Bounds

Lemma 1 (transient semigroup contraction). Let $\boldsymbol{G} = \boldsymbol{\Omega} - \operatorname{diag}(\boldsymbol{\gamma})$, where $\boldsymbol{\Omega}^\top = -\boldsymbol{\Omega}$ and $\gamma_r > 0$. Its Euclidean logarithmic norm is

$$\mu_2(\boldsymbol{G}) = \lambda_{\max}\left(\frac{\boldsymbol{G} + \boldsymbol{G}^\top}{2}\right) = \lambda_{\max}(-\operatorname{diag}(\boldsymbol{\gamma})) = -\gamma_{\min}. \tag{17}$$

Consequently, $\left\|e^{k\boldsymbol{G}}\right\|_2 \le e^{-k\gamma_{\min}} = \rho^k$ for every integer $k \ge 0$.

Lemma 2 (normalized-decoder score transfer). For nonzero vectors $\boldsymbol{x}, \boldsymbol{y}$,

$$\left\| \frac{\boldsymbol{x}}{\|\boldsymbol{x}\|_2} - \frac{\boldsymbol{y}}{\|\boldsymbol{y}\|_2} \right\|_2 \leq \frac{2\,\|\boldsymbol{x}-\boldsymbol{y}\|_2}{\min\{\|\boldsymbol{x}\|_2, \|\boldsymbol{y}\|_2\}}. \tag{18}$$

To see this, add and subtract $\boldsymbol{y}/\|\boldsymbol{x}\|_2$, apply the triangle inequality, and use the reverse triangle inequality to bound $|\,\|\boldsymbol{x}\|_2 - \|\boldsymbol{y}\|_2\,|$. If $\bar{\boldsymbol{e}}_j$ is unit normalized, Cauchy–Schwarz further gives

$$\left| \left\langle \frac{\boldsymbol{x}}{\|\boldsymbol{x}\|_2} - \frac{\boldsymbol{y}}{\|\boldsymbol{y}\|_2}, \bar{\boldsymbol{e}}_j \right\rangle \right| \leq \frac{2\,\|\boldsymbol{x}-\boldsymbol{y}\|_2}{\min\{\|\boldsymbol{x}\|_2, \|\boldsymbol{y}\|_2\}}. \tag{19}$$

A.2 Proof and Consequences of Theorem 1

Subtracting two controlled recurrences gives $\boldsymbol{\delta}_{t+1} = \boldsymbol{A}\boldsymbol{\delta}_t + \boldsymbol{B}\Delta\boldsymbol{u}_t$. Since $\boldsymbol{A} = \mathrm{blkdiag}(\boldsymbol{I}, e^{\boldsymbol{G}})$, recursive expansion yields the exact identities

$$\begin{aligned} \boldsymbol{\delta}^{\mathrm{per}}_{t+n} &= \boldsymbol{\delta}^{\mathrm{per}}_{t} + \sum_{m=0}^{n-1} \boldsymbol{B}_{\mathrm{per}}\Delta\boldsymbol{u}_{t+m}, \\ \boldsymbol{\delta}^{\mathrm{tr}}_{t+n} &= e^{n\boldsymbol{G}}\boldsymbol{\delta}^{\mathrm{tr}}_{t} + \sum_{m=0}^{n-1} e^{(n-1-m)\boldsymbol{G}}\boldsymbol{B}_{\mathrm{tr}}\Delta\boldsymbol{u}_{t+m}. \end{aligned} \tag{20}$$

Lemma 1, submultiplicativity, and the triangle inequality give

$$\left\|\boldsymbol{\delta}^{\mathrm{tr}}_{t+n}\right\|_2 \leq \rho^n \left\|\boldsymbol{\delta}^{\mathrm{tr}}_{t}\right\|_2 + \|\boldsymbol{B}_{\mathrm{tr}}\|_2 \sum_{m=0}^{n-1} \rho^{n-1-m} \left\|\Delta\boldsymbol{u}_{t+m}\right\|_2. \tag{21}$$

If the states agree before event t, $\Delta\boldsymbol{u}_t$ is the only nonzero control difference, and all later controls coincide, Equation 20 reduces to Equation 6. This proves Theorem 1.

Two useful special cases follow immediately. Under identical future controls,

$$\|\boldsymbol{\delta}_{t+n}\|_2^2 \leq \|\boldsymbol{\delta}^{\mathrm{per}}_{t}\|_2^2 + \rho^{2n}\left\|\boldsymbol{\delta}^{\mathrm{tr}}_{t}\right\|_2^2 \leq \|\boldsymbol{\delta}_t\|_2^2, \tag{22}$$

which recovers same-input non-expansiveness. More generally, if $\sup_m \|\Delta\boldsymbol{u}_{t+m}\|_2 \leq \epsilon_u$, then

$$\left\|\boldsymbol{\delta}^{\mathrm{tr}}_{t+n}\right\|_2 \leq \rho^n \left\|\boldsymbol{\delta}^{\mathrm{tr}}_{t}\right\|_2 + \frac{1-\rho^n}{1-\rho}\|\boldsymbol{B}_{\mathrm{tr}}\|_2\,\epsilon_u, \qquad \limsup_{n\to\infty} \left\|\boldsymbol{\delta}^{\mathrm{tr}}_{t+n}\right\|_2 \leq \frac{\|\boldsymbol{B}_{\mathrm{tr}}\|_2}{1-\rho}\epsilon_u. \tag{23}$$

The persistent coordinates are intentionally different: their discrepancy accumulates the explicit sum in Equation 20, rather than being hidden behind a global stability claim.

The event response also reaches the forecast space. Partition $\boldsymbol{D} = [\boldsymbol{D}_{\mathrm{per}}, \boldsymbol{D}_{\mathrm{tr}}]$. Before normalization, a single event perturbation obeys

$$\begin{aligned} \Delta(\boldsymbol{D}\boldsymbol{z}_{t+k+1}) &= \boldsymbol{D}_{\mathrm{per}}\boldsymbol{B}_{\mathrm{per}}\Delta\boldsymbol{u}_t + \boldsymbol{D}_{\mathrm{tr}}e^{k\boldsymbol{G}}\boldsymbol{B}_{\mathrm{tr}}\Delta\boldsymbol{u}_t, \\ \left\|\boldsymbol{D}_{\mathrm{tr}}e^{k\boldsymbol{G}}\boldsymbol{B}_{\mathrm{tr}}\Delta\boldsymbol{u}_t\right\|_2 &\leq \|\boldsymbol{D}_{\mathrm{tr}}\|_2\,\|\boldsymbol{B}_{\mathrm{tr}}\|_2\,\rho^k\,\|\Delta\boldsymbol{u}_t\|_2. \end{aligned} \tag{24}$$

Hence the raw decoded effect approaches its persistent displacement as event distance grows. Applying Lemma 2 to the two decoded trajectories transfers the corresponding finite-horizon latent difference to the normalized dynamic scores used by the recommender.

A.3 Proof of Theorem 2: Representation-to-Ranking Stability

From the residual definition and rollout recurrence,

$$\begin{aligned} \boldsymbol{e}^{(n),\mathrm{per}}_{t} &= \boldsymbol{e}^{(0),\mathrm{per}}_{t} - \sum_{m=0}^{n-1} \boldsymbol{\varepsilon}^{\mathrm{per}}_{t+m}, \\ \boldsymbol{e}^{(n),\mathrm{tr}}_{t} &= e^{n\boldsymbol{G}}\boldsymbol{e}^{(0),\mathrm{tr}}_{t} - \sum_{m=0}^{n-1} e^{(n-1-m)\boldsymbol{G}}\boldsymbol{\varepsilon}^{\mathrm{tr}}_{t+m}. \end{aligned} \tag{25}$$

Define

$$P_{t,n} = \left\| \boldsymbol{e}_t^{(0),\mathrm{per}} \right\|_2 + \sum_{m=0}^{n-1} \left\| \boldsymbol{\varepsilon}_{t+m}^{\mathrm{per}} \right\|_2,$$
$$T_{t,n} = \rho^n \left\| \boldsymbol{e}_t^{(0),\mathrm{tr}} \right\|_2 + \sum_{m=0}^{n-1} \rho^{n-1-m} \left\| \boldsymbol{\varepsilon}_{t+m}^{\mathrm{tr}} \right\|_2,$$
$$E_{t,n} = \sqrt{P_{t,n}^2 + T_{t,n}^2}. \tag{26}$$

The triangle inequality bounds the persistent term in Equation 25 by $P_{t,n}$, while Lemma 1 bounds the transient term by $T_{t,n}$. Because the blocks occupy orthogonal coordinates, their squared bounds add, proving $\left\| \boldsymbol{e}_t^{(n)} \right\|_2 \le E_{t,n}$.

Set $\boldsymbol{x} = \boldsymbol{D}\widehat{\boldsymbol{z}}_t^{(n)}$ and $\boldsymbol{y} = \boldsymbol{D}\boldsymbol{z}_{t+n}^{\star}$. Then $\|\boldsymbol{x}-\boldsymbol{y}\|_2 \le \|\boldsymbol{D}\|_2 \left\| \boldsymbol{e}_t^{(n)} \right\|_2 \le \|\boldsymbol{D}\|_2 E_{t,n}$. Lemma 2 therefore gives the uniform catalog bound in Equation 10.

It remains to prove the margin implication. Let i be any member of the target Top-K set and j any item outside it. By definition of the boundary margin,

$$\widehat{\ell}_{t,n}^{d}(i) - \widehat{\ell}_{t,n}^{d}(j) \ge \ell_{t,n}^{d,\star}(i) - \ell_{t,n}^{d,\star}(j) - 2\varepsilon_{t,n} \ge \Delta_{t,n}^{(K)} - 2\varepsilon_{t,n} > 0. \tag{27}$$

Every target Top-K item consequently remains above every outside item, proving the candidate-set equality. This certificate concerns the forecast-induced dynamic score field; the next result bounds the final decision even when the score margin is too small to certify equality.

A.4 Proof and Metric Consequences of Theorem 3

Equation 12 places every non-admitted item outside $\mathcal{P}_t$ below every protected item. For any $p \in \mathcal{P}_t$, at most $K_b - 1$ protected items and at most B_t admitted items can precede it. Since $K_b + B_t \le K_c$, member-preserving tie resolution gives $\mathcal{P}_t \subseteq \mathcal{C}_t$.

All scores outside $\mathcal{P}_t \cup \mathcal{A}_t$ retain their anchor values, admitted scores do not decrease, and protected items remain selected. Hence a non-admitted item outside $\mathcal{C}_t^b$ cannot displace an unchanged anchor-tail member; every item entering $\mathcal{C}_t$ from outside $\mathcal{C}_t^b$ must belong to $\mathcal{A}_t$. Thus

$$\mathcal{C}_t \setminus \mathcal{C}_t^b \subseteq \mathcal{A}_t, \qquad |\mathcal{C}_t \setminus \mathcal{C}_t^b| \le B_t. \tag{28}$$

The final and anchor sets both have cardinality K_c, so the numbers of entering and leaving items are equal. Consequently,

$$|\mathcal{C}_t \triangle \mathcal{C}_t^b| = 2|\mathcal{C}_t \setminus \mathcal{C}_t^b| \le 2B_t, \qquad \frac{|\mathcal{C}_t \triangle \mathcal{C}_t^b|}{2K_c} \le \frac{B_t}{K_c}. \tag{29}$$

Finally, let $a_1 \ge \cdots \ge a_{K_c}$ be the descending pre-order score values assigned to $\mathcal{C}_t$. The implementation only permutes $\{a_1, \ldots, a_{K_c}\}$ inside $\mathcal{C}_t$ according to Equation 13; all outside scores remain unchanged. Under a strict boundary, every outside score is below a_{K_c}; member-preserving resolution handles equality. Hence

$$\mathcal{C}_t^f = \mathcal{C}_t,$$
$$\{\{\ell_t^f(j) : j \in \mathcal{C}_t\}\} = \{\{\ell_t^{\mathrm{pre}}(j) : j \in \mathcal{C}_t\}\}, \qquad \ell_t^f(j) = \ell_t^{\mathrm{pre}}(j) \text{ for } j \notin \mathcal{C}_t. \tag{30}$$

This completes Theorem 3.

For the single held-out next item used in evaluation, HR@K_c is the indicator that the target belongs to $\mathcal{C}_t$. It is therefore invariant to the within-set reassignment. If the target moves from rank r to $r' \le r$, then for any cutoff $K \le K_c$, its NDCG contribution changes from $\mathbf{1}[r \le K]/\log_2(r+1)$ to $\mathbf{1}[r' \le K]/\log_2(r'+1)$ and cannot decrease. This formalizes why localized order correction can improve tighter-cutoff NDCG while leaving broad-cutoff membership unchanged.

B Additional Model and Protocol Specification

B.1 Frozen Anchor and Decoder Supervision

This subsection gives the anchor construction and decoder supervision summarized in Sections 3.2 and 3.4.

Let $f_{\text{in}} : \mathbb{R}^{d_x} \to \mathbb{R}^{d_h}$, $\boldsymbol{P}_s \in \mathbb{R}^{d_h \times d_s}$, and $\boldsymbol{P}_c \in \mathbb{R}^{d_h \times d_c}$. Writing $\boldsymbol{e}_i^{\text{sem}} = f_{\text{in}}(\boldsymbol{x}_i) + \alpha_s \boldsymbol{P}_s \boldsymbol{s}_i$, the shared anchor representation is

$$\boldsymbol{e}_i = \boldsymbol{e}_i^{\text{sem}} + \alpha_c \big(\boldsymbol{I}_{d_h} - \bar{\boldsymbol{e}}_i^{\text{sem}} (\bar{\boldsymbol{e}}_i^{\text{sem}})^\top\big) \boldsymbol{P}_c \bar{\boldsymbol{c}}_i, \quad [\boldsymbol{h}_1, \dots, \boldsymbol{h}_t] = F_\Theta(\boldsymbol{e}_{i_1}, \dots, \boldsymbol{e}_{i_t}), \quad \ell_t^b(j) = \langle \bar{\boldsymbol{h}}_t, \bar{\boldsymbol{e}}_j \rangle. \tag{31}$$

Here $\alpha_s, \alpha_c \geq 0$ are learned mixing scalars and $\boldsymbol{I}_{d_h}$ is the identity matrix. The decoder auxiliary loss is

$$\mathcal{L}_{\text{dec}} = 1 - \cos(\boldsymbol{q}_t, \bar{\boldsymbol{h}}_{t+1}) + \lambda_{\text{Huber}} \operatorname{Huber}(\boldsymbol{q}_t, \bar{\boldsymbol{h}}_{t+1}), \tag{32}$$

with $\lambda_{\text{Huber}} > 0$.

B.2 Reliability-Gated Bounded Candidate Closure

Let $\operatorname{TopK}_k(f \mid \mathcal{I})$ return the k items in $\mathcal{I}$ with largest score f, with $\mathcal{I} = \mathcal{V}$ when omitted. After applying the protocol-specific candidate mask, define

$$\begin{aligned} \mathcal{I}_t^b &= \operatorname{TopK}_{K_b}(\ell_t^b), & \mathcal{I}_t^{\text{trust}} &= \operatorname{TopK}_{K_{\text{trust}}}(\ell_t^b), \\ \mathcal{I}_t^d &= \operatorname{TopK}_{K_b}(\ell_t^d \mid \mathcal{I}_t^{\text{trust}}), & \mathcal{I}_t^{\text{union}} &= \mathcal{I}_t^b \cup \mathcal{I}_t^d. \end{aligned} \tag{33}$$

Over $\mathcal{I}_t^{\text{union}}$, let (μ_t^b, σ_t^b) and (μ_t^d, σ_t^d) denote the means and stabilized positive standard deviations of the two score fields. The standardized dynamic advantage is

$$\Delta_t(j) = \frac{\ell_t^d(j) - \mu_t^d}{\sigma_t^d} - \frac{\ell_t^b(j) - \mu_t^b}{\sigma_t^b}. \tag{34}$$

A reliability gate $\kappa_t \in [0, 1]$, computed from the contextual state and its forecast, modulates how strongly the dynamic advantage may alter the anchor ranking. Let $\beta_t = \min_{j \in \mathcal{I}_t^b} \ell_t^b(j)$ be the protected boundary, K_{prom} the promotion budget, $\lambda_r \geq 0$ the reranking scale, $m_{\text{cross}} > 0$ the crossing margin, and $m_{\text{floor}} \in (0, m_{\text{cross}})$ the membership floor. Define $\mathcal{I}_t^{\text{prom}} = \operatorname{TopK}_{K_{\text{prom}}}(\Delta_t \mid \mathcal{I}_t^d \setminus \mathcal{I}_t^b)$. The bounded score is

$$\ell_t^{\text{pre}}(j) = \begin{cases} \max\{\ell_t^b(j) + \kappa_t \lambda_r \sigma_t^b \Delta_t(j), \beta_t + m_{\text{floor}} \sigma_t^b\}, & j \in \mathcal{I}_t^b, \\ \ell_t^b(j) + \kappa_t \left(\beta_t - \ell_t^b(j) + m_{\text{cross}} \sigma_t^b\right), & j \in \mathcal{I}_t^{\text{prom}}, \\ \ell_t^b(j), & \text{otherwise.} \end{cases} \tag{35}$$

The first branch changes order without removing anchor members, the second promotes only selected forecast candidates, and the third leaves all other items unchanged.

B.3 Semantic Residual and Outer Proposal

The observed prefix also provides a realized direction in language space. For maximum window $L_o \in \mathbb{N}_+$, available length $L_t = \min\{L_o, t-1\}$, semantic recency factor $\beta_{\text{sem}} \in [0, 1]$, and extrapolation coefficient $\lambda_m \in \mathbb{R}$, define

$$\boldsymbol{m}_t = \frac{\sum_{r=0}^{L_t - 1} \beta_{\text{sem}}^r (\bar{\boldsymbol{s}}_{i_{t-r}} - \bar{\boldsymbol{s}}_{i_{t-r-1}})}{\sum_{r=0}^{L_t - 1} \beta_{\text{sem}}^r} \in \mathbb{R}^{d_s}, \qquad \boldsymbol{s}_t^{\text{res}} = \overline{\bar{\boldsymbol{s}}_{i_t} + \lambda_m \boldsymbol{m}_t} \in \mathbb{R}^{d_s}, \tag{36}$$

with $\boldsymbol{m}_t = \boldsymbol{0}$ when $L_t = 0$. Let $\mathcal{T}_t = \operatorname{TopK}_{K_{\text{trust}}}(\ell_t^{\text{pre}}) \setminus \operatorname{TopK}_{K_b}(\ell_t^{\text{pre}})$ be the editable tail. For outer width $K_o \in \mathbb{N}_+$,

$$\mathcal{O}_t = \operatorname{TopK}_{K_o}\big(\ell_t^d \mid \mathcal{V} \setminus \operatorname{TopK}_{K_{\text{trust}}}(\ell_t^b)\big) \tag{37}$$

is the forecast proposal outside the core trust region. For a finite set $\mathcal{S}$, let $Z_{\mathcal{S}}[f]$ center and variance-normalize scalar field f on $\mathcal{S}$. With $\alpha_o, \lambda_o \in [0, 1]$,

$$\ell_t^{\text{hyb}}(j) = \alpha_o Z_{\mathcal{T}_t \cup \mathcal{O}_t}[\ell_t^b](j) + (1 - \alpha_o) Z_{\mathcal{T}_t \cup \mathcal{O}_t}[\ell_t^d](j), \tag{38}$$

$$\ell_t^{\text{res}}(j) = Z_{\mathcal{O}_t}[\langle \boldsymbol{s}_t^{\text{res}}, \bar{\boldsymbol{s}}_\cdot \rangle](j), \qquad j_t^o = \arg\max_{j \in \mathcal{O}_t} \big((1 - \lambda_o) \ell_t^{\text{hyb}}(j) + \lambda_o \ell_t^{\text{res}}(j)\big). \tag{39}$$

Thus the Koopman score proposes candidates and realized semantic motion resolves their local priority. The bounded admission rule combines this priority with forecast reliability, sets $\mathcal{P}_t = \mathcal{I}_t^b$, and admits at most $K_c - K_b$ items into $\mathcal{A}_t$. The protected separation in Equations 11–12 therefore continues to hold.

B.4 Listwise Order Features

For each $j \in \mathcal{C}_t$, the ranking representation groups five causal evidence sources:

$$\boldsymbol{\psi}_t(j) = [\boldsymbol{\psi}_t^{\mathrm{score}}(j); \boldsymbol{\psi}_t^{\mathrm{sem}}(j); \boldsymbol{\psi}_t^{\mathrm{ctx}}(j); \boldsymbol{\psi}_t^{\mathrm{rank}}(j); \boldsymbol{\psi}_t^{\mathrm{rec}}(j)] \in \mathbb{R}^{d_\psi}. \tag{40}$$

The groups respectively encode anchor and dynamic scores, semantic trajectory, forecast context, incoming rank, and causal recurrence. All are determined by the observed prefix $i_{1:t}$. Equation 13 orders candidates, after which descending pre-order values are reassigned within $\mathcal{C}_t$ and scores outside $\mathcal{C}_t$ remain unchanged.

B.5 Datasets and Evaluation Protocol

Table 9: Dataset statistics for the strict MPT release.

Dataset	Users	Items	Interactions	Avg. length
Beauty	22,363	12,101	198,502	8.88
Toys	19,412	11,924	167,597	8.63
Sports	35,598	18,357	296,337	8.32
Yelp2018-10core	77,277	45,638	2,103,896	27.23
OnlineRetail	16,520	3,469	519,906	31.47

Table 9 covers short Amazon sequences as well as the longer Yelp and retail trajectories, providing variation in catalog size, interaction volume, and sequence length. All datasets use chronological leave-one-out: the last event is test, the penultimate event is validation, and earlier interactions define the observed prefix. Full-catalog evaluation assigns one rank per user. The strict history-masked protocol removes every item in the complete pre-endpoint history after scoring; the comparison protocol applies the no-history-mask convention of the DIFF/FreqRec pipeline. Results are reported separately within each protocol family.

B.6 Reproducibility Configuration

Architecture and dynamics. The frozen MPT anchor uses four Transformer layers, two attention heads, hidden and intermediate dimensions of 256, and maximum sequence length 50. Qwen3-Embedding-0.6B item embeddings are reduced to $d_s = 64$ dimensions by catalog-level PCA; collaborative controls use $d_c = 64$. The Koopman observable has $d_z = 64$ coordinates, split into $d_{\mathrm{per}} = 16$ persistent and $d_{\mathrm{tr}} = 48$ transient dimensions. In Equation 4, $\eta_s = 0.05$ and $\delta_{\mathrm{stab}} = 10^{-4}$. The target-lift exponential-moving-average coefficient is 0.995, the rollout set is $\mathcal{H} = \{1, 2, 4\}$, and $(w_1, w_2, w_4) = (1, 0.5, 0.25)$. We use $\lambda_{\cos} = 0.25$ and $\lambda_{\mathrm{Huber}} = 0.10$. The Koopman core contains 72,371 trainable parameters, and the complete adaptation contains at most 72,708; the MPT anchor and item representations remain frozen.

Optimization and reporting. The adapter uses AdamW with learning rate 10^{-3}, dropout 0.2, temperature 0.07, and 30 epochs. Weight decay is 0.1 on Beauty, Sports, Toys, and Yelp2018-10core, and 1.0 on OnlineRetail. Batch size is 256, except 512 on OnlineRetail. Results aggregate three independent runs and report sample standard deviation with denominator $n - 1$.

Decision closure. All primary comparison runs share $K_b = 10$ and $K_c = 20$. These values define the protected-core size and selected candidate region; the realized admission count $B_t \le K_c - K_b$ determines the anchor-relative intervention radius in Theorem 3. The retrieval widths $K_{\mathrm{trust}} = 20$ and $K_o = 80$, representation width $d_\psi = 19$, and comparison-setting value $K_{\mathrm{prom}} = 3$ specify the realization inside this structural envelope. We use $m_{\mathrm{cross}} = 0.01$, $m_{\mathrm{floor}} = 0.001$, and $\alpha_o = 0.8600$. Table 10 lists the remaining implementation settings; its residual tuple is $(\lambda_m, L_o, \beta_{\mathrm{sem}}, \lambda_o)$, and MLP(16) denotes one hidden layer of width 16. The strict history-masked suite retains the same dimensions and optimizer; Beauty, Sports, and Toys use $\lambda_r = 0.5$, whereas Yelp2018-10core

and OnlineRetail use boundary-only admission with $\lambda_r = 0$. Beauty and Toys admit at most three promoted candidates, while Sports, Yelp2018-10core, and OnlineRetail allow all dynamically retrieved candidates to compete within the bounded editable region. Alternative-semantic experiments replace only the frozen semantic matrix.

Table 10: Decision-layer implementation settings reported for reproducibility.

Dataset	λ_r	Residual tuple	Order scorer	λ_{ord}
Beauty	2.0	$(-0.5000, 3, 1.0000, 0.1800)$	Linear	8
Sports	3.0	$(0.0000, 1, 1.0000, 0.0000)$	MLP(16)	12
Toys	2.0	$(0.7500, 3, 0.7500, 0.1800)$	MLP(16)	12

Table 10 records the remaining dataset-specific implementation settings inside the shared structural envelope. The latent dimensions, frozen anchor, semantic construction, rollout horizons, optimizer, protected-core width, and selected-set width remain fixed. These settings govern how evidence is scored within the admissible region without changing the containment and intervention guarantees of Theorem 3.

C Additional Experimental Results

C.1 Strict History-Masked Results

Table 11: Strict history-masked full-catalog test results.

Dataset	HR@5	NDCG@5	HR@10	NDCG@10	HR@20	NDCG@20	MRR
Beauty	**0.0843±0.0002**	**0.0596±0.0002**	**0.1218±0.0005**	**0.0716±0.0003**	**0.1612±0.0000**	**0.0815±0.0001**	**0.0633±0.0002**
Sports	**0.0463±0.0003**	**0.0316±0.0001**	**0.0726±0.0003**	**0.0400±0.0001**	**0.0959±0.0000**	**0.0458±0.0000**	**0.0354±0.0000**
Toys	**0.0910±0.0001**	**0.0647±0.0002**	**0.1301±0.0007**	**0.0773±0.0003**	**0.1677±0.0000**	**0.0867±0.0002**	**0.0679±0.0003**
Yelp2018-10core	**0.0368±0.0000**	**0.0234±0.0000**	**0.0680±0.0004**	**0.0331±0.0001**	**0.0952±0.0000**	**0.0398±0.0000**	**0.0296±0.0000**
OnlineRetail	**0.1773±0.0000**	**0.1237±0.0000**	**0.2601±0.0026**	**0.1500±0.0008**	**0.3296±0.0000**	**0.1672±0.0000**	**0.1287±0.0000**

Table 11 supplies the complete strict-setting operating points summarized by the matched-anchor and cutoff analyses in the main text.

C.2 Ranking-Closure Ablation

Table 12: Ranking-closure ablation under the comparison protocol.

Dataset	Variant	HR@10	NDCG@10	HR@20	NDCG@20	MRR
Beauty	w/o recurrence state	0.1249	0.0694	0.1516	0.0759	0.0588
	anchor + recurrence only	0.1314	0.0751	0.1516	0.0801	0.0638
	Full SemKoopRec	**0.1326**	**0.0753**	**0.1516**	**0.0800**	**0.0636**
Sports	w/o recurrence state	0.0734	0.0390	0.0901	0.0431	0.0332
	anchor + recurrence only	0.0770	0.0416	0.0901	0.0448	0.0353
	Full SemKoopRec	**0.0781**	**0.0421**	**0.0901**	**0.0451**	**0.0356**
Toys	w/o recurrence state	0.1249	0.0677	0.1592	0.0761	0.0571
	anchor + recurrence only	0.1389	0.0804	**0.1592**	0.0854	0.0682
	Full SemKoopRec	**0.1434**	**0.0834**	**0.1592**	**0.0874**	**0.0705**

Table 12 expands the main-text decision-closure comparison from NDCG@10 to all reported ranking metrics. The complete model is strongest in HR@10 and NDCG@10 on every dataset, while HR@20 remains identical within each dataset. The combination shows that dynamics and recurrence chiefly improve the location and discounted order of relevant items inside a preserved candidate envelope, consistent with the bounded closure construction.

C.3 Main-Table Standard Deviations

Table 13: SemKoopRec main-table results (mean ± sample SD).

Dataset	HR@10	NDCG@10	HR@20	NDCG@20
Beauty	**0.1325±0.0002**	**0.0752±0.0003**	**0.1515±0.0001**	**0.0799±0.0002**
Sports	**0.0781±0.0001**	**0.0421±0.0000**	**0.0903±0.0001**	**0.0451±0.0000**
Toys	**0.1428±0.0007**	**0.0831±0.0003**	**0.1593±0.0001**	**0.0873±0.0001**

Table 13 supplies the uncertainty omitted from the wide protocol-matched comparison table. The small sample deviations show that the reported ordering is stable across independent runs; entries displayed as 0.0000 reflect four-decimal rounding rather than an assumed zero variance.

D Efficiency Details

Table 14: Full-catalog inference efficiency on an RTX 4080 SUPER.

Dataset	System	Test time (s)	Users/s	ms/batch	Peak MiB
	Full closed loop	**1.7350**	**12,889**	**19.716**	**454.04**
Beauty	Koopman core	0.9904	22,580	11.254	435.93
	MLP	0.9676	23,112	10.995	436.88
	Full closed loop	**2.8360**	**12,552**	**20.257**	**696.55**
Sports	Koopman core	1.7005	20,934	12.147	631.01
	MLP	1.6669	21,356	11.907	632.80
	Full closed loop	**1.4946**	**12,989**	**19.665**	**411.07**
Toys	Koopman core	0.8274	23,462	10.886	433.04
	MLP	0.8139	23,851	10.709	434.55

Table 14 reports the complete efficiency analysis. The model confines its additional computation to latent transition and bounded candidate ordering, while full-catalog scoring remains the dominant cost and the stable operator is cached during evaluation. The complete system processes 12,810 users/s (0.0781 ms/user) with 411.07–696.55 MiB peak incremental memory; matched Koopman and MLP training takes 380.25 s and 382.67 s, respectively, showing that the structured operator introduces no material training-time premium over the equal-capacity alternative.