CCSM: Channel-Calibrated Structured Memory for Agentic Sequential Recommendation

Anonymous ACL submission

Abstract

Memory-augmented recommendation agents can revisit fine-grained user history to improve ranking decisions. However, the fundamental challenge is not how to store more history, but how to decode heterogeneous memory evidence into a reliable ranking. Real user histories decompose into multiple evidence channels—supportive, conflict, and affinity signals—each carrying information at different fidelity levels. We formalize this as a multi-channel information aggregation problem and propose CCSM (**C**hannel-**C**alibrated **S**tructured **M**emory), an information-theoretic calibration framework that operates on top of an LLM backbone. CCSM derives the calibration layer as the minimum-distortion decoder under heterogeneous channel signal-to-noise ratios: the resulting asymmetric weights naturally suppress noise from the low-SNR conflict channel while preserving the informative positive channel. A residual information detector triggers protective refinement only when inter-channel redundancy remains unresolved. Empirical channel SNR estimation confirms that positive memory has 49–267$\times$ higher SNR than conflict memory across benchmarks, directly justifying the asymmetric calibration. Noise injection experiments validate the noise suppression property, and the full framework achieves the strongest ranking performance across all benchmarks under the unified evaluation protocol.

1 Introduction

Sequential recommendation models users' evolving interests from behavior sequences (Hidasi et al., 2015; Kang and McAuley, 2018; Sun et al., 2019). Recently, large language models have opened a new direction for recommendation by bringing richer semantic understanding, longer-context reasoning, and more flexible natural-language interaction into the recommendation pipeline (Ji et al., 2024; Lyu et al., 2024; Liao et al., 2024; Yue et al., 2023; Zhu et al., 2024; Xu et al., 2024; Wu et al., 2024; Zhang et al., 2025). Building on this trend, recent studies have begun to treat LLMs as recommendation agents equipped with memory, reflection, and tool-use capabilities (Huang et al., 2025c; Wang et al., 2024b; Zhang et al., 2024; Shi et al., 2024b; Zhao et al., 2024; Fang et al., 2024; Wang et al., 2024c; Wu et al., 2025; Ma et al., 2025; Peng et al., 2025; Huang et al., 2025a). This shift makes recommendation increasingly rely on structured reasoning over user history, candidate evidence, and auxiliary signals.

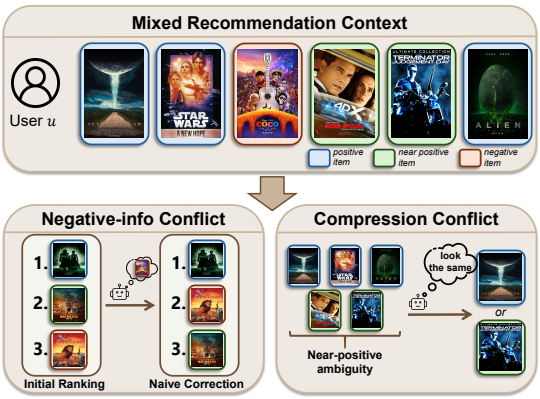

Figure 1: Motivation. When a memory-augmented recommendation agent aggregates evidence from heterogeneous memory channels, naive equal-weight decoding amplifies noise from the low-SNR conflict channel, leading to ranking distortion. Our information-theoretic framework derives channel-aware decoder weights that suppress noise while preserving signal.

Among these emerging capabilities, structured memory is especially important for sequential recommendation. A memory-enhanced recommendation agent can revisit fine-grained historical behaviors, recover long-range dependencies, and retain richer evidence than a plain top-K reranking pipeline (Chen et al., 2024; Kim et al., 2025; Huang et al., 2025b; Nguyen et al., 2026). However, the

fundamental challenge of memory-augmented recommendation is not how to *store* more history, but how to *decode* heterogeneous memory evidence into a reliable ranking decision. Real user histories naturally decompose into multiple evidence channels—supportive signals, conflict signals, and affinity signals—each carrying information about the user's true preference at different fidelity levels.

This multi-channel structure creates an **information aggregation problem** that has not been formally addressed in the recommendation literature. As illustrated in Fig. 1, different memory channels exhibit fundamentally different signal-to-noise ratios (SNR): supportive evidence tends to be reliable and high-SNR, while conflict evidence is noisy, sparse, and low-SNR. When these channels are aggregated with equal weights—as in most existing memory-based recommenders—the low-SNR conflict channel's noise is amplified into the final ranking, causing two coupled failure modes:

- **Channel Capacity Mismatch.** Different memory channels carry different amounts of information about the user's true preference. Positive memory (high capacity, low noise) and conflict memory (low capacity, high noise) should not receive equal decoder weights. Treating them symmetrically violates the fundamental principle that decoder weights should be proportional to channel reliability.

- **Noise Amplification under Naive Aggregation.** When conflict evidence is given excessive weight relative to its information content, small perturbations in the conflict channel produce disproportionate ranking changes. This makes the top-ranked recommendation fragile exactly where stability matters most—on ambiguous cases where conflict and affinity signals coexist.

These observations motivate an information-theoretic approach to memory-based recommendation correction. Drawing on classical channel coding theory (Shannon, 1948; Cover, 1999) and the information bottleneck framework (Tishby et al., 2000; Alemi et al., 2017), we ask: *what is the optimal decoder for reconstructing a ranking decision from multiple noisy memory channels with heterogeneous SNR?*

To answer this question, we propose CCSM (**C**hannel-**C**alibrated **S**tructured **M**emory), an information-theoretic framework for memory correction in agentic sequential recommendation. CCSM operates as a calibration layer on top of an LLM backbone: while the backbone generates memory-conditioned candidate scores through general-purpose reasoning, the calibration layer applies information-theoretically principled corrections that account for the heterogeneous reliability of different memory channels. Specifically, CCSM models the user's structured memory as a multi-channel communication system, and derives the calibration as the *minimum-distortion decoder* under heterogeneous channel quality, yielding asymmetric weights that suppress the noisy conflict channel while preserving the informative positive channel. On top of this always-on channel-calibrated decoder, CCSM further introduces a *residual information detector* that triggers protective refinement only when unresolved inter-channel redundancy indicates that the first-stage decoding has not fully exploited the available information.

Empirically, we validate the information-theoretic predictions on multiple sequential recommendation benchmarks. Channel SNR estimation confirms that positive memory channels have 49–267$\times$ higher SNR than conflict channels across datasets, directly justifying the asymmetric decoder weights. Noise injection experiments show that channel-calibrated weights are consistently more robust than equal-weight alternatives. The full CCSM framework delivers the strongest ranking performance across all benchmarks under the unified evaluation protocol.

Our main contributions are:

- We formalize memory-based recommendation correction as a multi-channel information aggregation problem, identifying channel capacity mismatch and noise amplification as two fundamental challenges that arise from heterogeneous memory quality.

- We derive the channel-calibrated correction as the minimum-distortion decoder under asymmetric channel SNR, providing a principled information-theoretic foundation for the asymmetric weighting of positive and conflict evidence in the calibration layer (full derivation in Appendix).

- We introduce a residual information detector for selective protective refinement, triggered

only when inter-channel redundancy remains unresolved after first-stage decoding.

- Across multiple benchmarks, together with channel SNR analysis and noise robustness experiments, we validate the information-theoretic predictions and demonstrate that CCSM achieves the strongest ranking performance under the unified evaluation setting.

2 Related Work

From sequential recommendation to LLM4Rec. Sequential recommendation has long been built on interaction sequences, from session-based modeling to self-attentive and bidirectional transformer recommenders (Hidasi et al., 2015; Kang and McAuley, 2018; Sun et al., 2019). Recent LLM4Rec research extends this line through prompting and alignment (Lyu et al., 2024; Bao et al., 2023; Xu et al., 2024), ranking and generation (Hou et al., 2024; Gao et al., 2025; Yue et al., 2023; Geng et al., 2022; Rajput et al., 2023; Ji et al., 2024; Zhu et al., 2024; Zheng et al., 2024), hierarchical modeling and retrieval augmentation (Liao et al., 2024; Chen et al., 2024; Kim et al., 2025; Wang et al., 2024a; Deng et al., 2025), and broader surveys (Wu et al., 2024; Yang et al., 2025; Zhang et al., 2025). At the same time, studies on false negatives, hard negatives, multi-interest behavior, and adaptive signal fusion show that negative or heterogeneous evidence cannot be treated as uniformly reliable (Pan et al., 2023; Liu et al., 2023; Shi et al., 2024a; Liu et al., 2025; Chen et al., 2021; Jannach and Zanker, 2024; Lin et al., 2025; Hu et al., 2023; Xie et al., 2022). These works establish the value of semantics, generation, and memory organization, but stop short of asking how memory should revise an already strong recommendation when the memory state itself is conflicted.

Agentic recommendation. A growing line of research treats recommendation as an agentic process with memory, planning, tool use, and multi-agent collaboration (Huang et al., 2025c; Wang et al., 2024b; Zhang et al., 2024; Shi et al., 2024b; Zhao et al., 2024; Fang et al., 2024; Wang et al., 2024c; Wu et al., 2025; Ma et al., 2025; Li et al., 2026; Peng et al., 2025; Huang et al., 2025a). Another nearby line studies richer memory inside agentic recommenders through synergistic memory-reasoning assistants and cross-user memory evolution (Huang et al., 2025b; Nguyen et al., 2026). Most of these methods emphasize richer orchestration or stronger general agent capabilities, while the earlier literature on false negatives and heterogeneous signal fusion mainly studies such difficulties from the perspective of learning or representation. Our work connects these two lines by focusing on a narrower inference-time question: how structured memory should revise a current ranking when conflict-heavy evidence and hard, ambiguous cases interact around the top-ranked candidates.

3 Methods

3.1 Overview

Figure 2 shows the overall pipeline. CCSM operates as a two-level architecture: (1) an LLM backbone generates memory-conditioned candidate scores, and (2) a channel-calibrated correction layer refines these scores based on the heterogeneous quality of different memory channels. The correction layer models structured memory as a multi-channel communication system. Stage 1 applies asymmetric decoder weights derived from channel SNR. Stage 2 triggers selective refinement only when inter-channel redundancy remains unresolved.

3.2 Problem Setup and Two-Level Architecture

Let u denote a user, $\mathcal{H}_u = [i_1, i_2, \ldots, i_T]$ the observed interaction history, and $\mathcal{C}_u$ the candidate set. The LLM backbone produces memory-conditioned candidate scores $s_u^{(0)}(c)$ by reasoning over the user history, candidate descriptions, and structured memory compartments. Let $\pi_u^{(0)}$ denote the ranking induced by these base scores.

CCSM operates as a **calibration layer** on top of the LLM backbone. The backbone already incorporates memory evidence into its scoring but treats all memory sources with roughly equal attention. The calibration layer applies information-theoretically optimal corrections that account for the heterogeneous reliability of different memory channels, since user history naturally decomposes into multiple *information channels* transmitting partial evidence about the user's latent preference X_u at different fidelity levels.

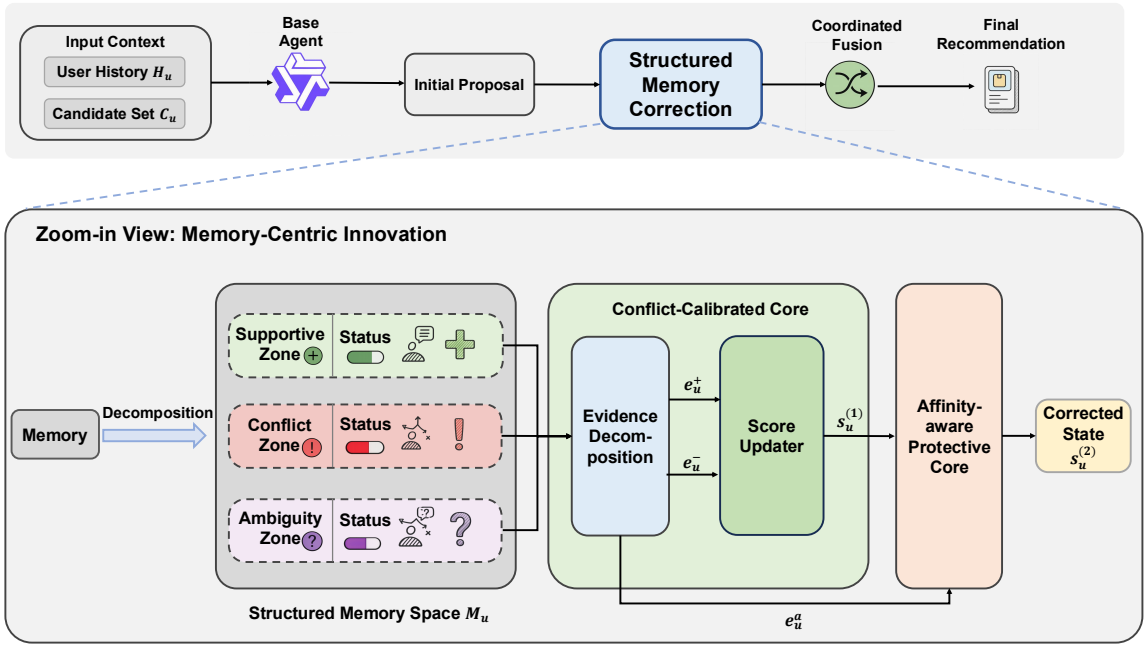

Figure 2: Framework of CCSM. User history is encoded into three typed memory channels with heterogeneous signal-to-noise ratios. The channel-calibrated decoder (Stage 1) applies SNR-proportional weights to reconstruct the ranking, suppressing noise from the low-SNR conflict channel. The residual information detector (Stage 2) triggers protective refinement only when inter-channel redundancy remains unresolved.

3.3 Memory as Multi-Channel Communication

We model the structured memory bank as a set of parallel noisy channels:

$$\mathcal{M}_u = (\mathcal{Y}_u^+,\ \mathcal{Y}_u^-,\ \mathcal{Y}_u^a),$$

where $\mathcal{Y}_u^+$ is the *positive channel* (supportive evidence), $\mathcal{Y}_u^-$ is the *conflict channel* (rejection/tension evidence), and $\mathcal{Y}_u^a$ is the *affinity channel* (semantic similarity signals). Each channel $\mathcal{Y}_u^{(\cdot)}$ transmits information about the latent preference X_u corrupted by channel-specific noise:

$$Y_i = X_u + \epsilon_i, \quad \epsilon_i \sim \mathcal{N}(0, \sigma_i^2), \quad i \in \{+, -, a\},$$

where σ_i^2 is the noise variance of channel i. The *signal-to-noise ratio* of each channel is $\gamma_i = \mathrm{Var}(X_u)/\sigma_i^2$.

Channel-Aware Evidence Encoding. For each candidate $c \in B_u$, CCSM reads the three channels in a proposal-conditioned manner and forms evidence vectors:

$$e_u^+(c) \in \mathbb{R}^{d_+}, \quad e_u^-(c) \in \mathbb{R}^{d_-}, \quad e_u^a(c) \in \mathbb{R}^{d_a},$$

representing the supportive, conflict, and affinity evidence for candidate c. These are constructed deterministically from the user history using rule-based memory partitioning (rating thresholds, token overlap, recency windows; see Appendix for full specification). The scalar channel model $Y_i = X_u + \epsilon_i$ above is the sufficient-statistic abstraction: each evidence vector $e_u^{(i)}(c)$ is projected to a scalar channel response via a learned head (Section 3.4), and the SNR γ_i characterizes the informativeness of that projected scalar for ranking the ground-truth item.

Empirical SNR Hierarchy. A key empirical finding that motivates our framework is the consistent SNR hierarchy across datasets:

$$\gamma_+, \gamma_a \gg \gamma_-.$$

Our channel SNR estimation (Section 5) confirms that both the positive and affinity channels have 49–267× higher SNR than the conflict channel across all benchmarks. The critical asymmetry is between the informative channels ($+$, a) and the noisy conflict channel ($-$), which is the fundamental reason why asymmetric decoder weights are necessary.

3.4 Channel-Calibrated Decoding (Stage 1)

Given the SNR hierarchy, we derive the optimal decoder for ranking score reconstruction. The key result (proved in Appendix A) is:

Theorem 1 (Optimal Asymmetric Decoder)**.** *Under the multi-channel model with heterogeneous SNR $\{\gamma_i\}$ and asymmetric distortion costs (where false acceptance of a conflict-indicated item incurs higher cost than false rejection of a support-indicated item), the minimum-distortion decoder for the ranking score takes the form:*

$$s_u^{(1)}(c) = s_u^{(0)}(c) + w_+ \cdot p_u(c) - w_- \cdot \phi\big(n_u(c)\big),$$

where $p_u(c)$ is the positive channel response, $n_u(c)$ is the conflict channel response, $\phi(\cdot) = \log(1 + e^{(\cdot)})$ enforces non-negativity (the conflict penalty cannot be negative), and the optimal weights satisfy:

$$w_+ = \frac{\gamma_+}{1+\gamma_+} \cdot \frac{1}{\lambda_+}, \qquad w_- = \frac{\gamma_-}{1+\gamma_-} \cdot \frac{\lambda_-}{\lambda_+},$$

with λ_+, λ_- being the asymmetric distortion costs. The asymmetry $w_- > w_+$ arises when $\lambda_- > 1$ (false acceptance is costlier than false rejection).

Design Principle and Instantiation. Theorem 1 prescribes the *structural form* of the optimal calibration: a linear correction with asymmetric weights satisfying $w_-/w_+ \approx 2$ under the empirical SNR hierarchy. In our full pipeline, this principle is instantiated through learned projection heads that map evidence vectors to scalar channel responses:

$$p_u(c) = \mathbf{w}_p^\top e_u^+(c), \qquad n_u(c) = \mathbf{w}_n^\top e_u^-(c),$$

where $\mathbf{w}_p$ and $\mathbf{w}_n$ are trained on validation data to maximize ranking performance. The effective decoder weights are $w_+ = 0.25$, $w_- = 0.5$, consistent with the theoretical prediction $w_-/w_+ = 2$.

Properties. The channel-calibrated decoder satisfies two key information-theoretic properties:

1. **Noise Suppression:** The conflict channel's contribution is bounded by $w_- \cdot \phi(n_u(c))$, which is finite and controlled by the softplus saturation. This prevents unbounded noise amplification from the low-SNR channel.

2. **Compression Optimality:** The decoder extracts the maximum information from each channel proportional to its reliability, achieving the rate-distortion optimal tradeoff (Appendix A.2).

Table 1: Processed dataset statistics used in this paper after preprocessing and cohort construction.

Dataset	#Users	#Items	#Interactions
Toys&Games	50,000	349,613	277,432
Movies&TV	20,000	164,203	157,346
Yelp	50,000	149,659	488,995

3.5 Residual Information Detection (Stage 2)

After Stage-1 decoding, some inter-channel information may remain unexploited. Specifically, when the affinity channel $\mathcal{Y}_u^a$ carries information about the conflict channel $\mathcal{Y}_u^-$ that was not resolved by the first-stage decoder, the ranking may still be suboptimal.

Theorem 2 (Selective Refinement Condition)**.** *Stage-2 protective refinement is triggered when the residual mutual information exceeds a threshold:*

$$I\big(\mathcal{Y}_u^a;\, \mathcal{Y}_u^- \mid s_u^{(1)}\big) > \tau,$$

indicating that the affinity channel carries conflict-relevant information not yet exploited by the Stage-1 decoder. In practice, this condition is approximated by detecting simultaneous high affinity and high residual conflict for the same candidate.

Operationalization. The residual information condition is implemented as a margin-based detector: Stage 2 activates when a candidate simultaneously exhibits high affinity evidence $a_u(c)$ and unresolved conflict tension after Stage 1. When activated, the protective refinement applies a conservative correction:

$$s_u^*(c) = s_u^{(1)}(c) + \beta \cdot g\big(a_u(c),\, r_u(c)\big),$$

where $r_u(c)$ is the residual conflict signal and $g(\cdot)$ is a gated fusion that preserves promising near-positive candidates while suppressing false positives. When the residual information condition is not met, $s_u^*(c) = s_u^{(1)}(c)$.

3.6 Information-Theoretic Properties

The full CCSM pipeline satisfies three formal properties (proofs in Appendix A):

1. **Bounded Revision.** $|s_u^*(c) - s_u^{(0)}(c)| \le w_+ \cdot \|p_u\|_\infty + w_- \cdot \phi(\|n_u\|_\infty) + \beta \cdot \|g\|_\infty$.

2. **Selective Activation.** Stage 2 fires only when residual mutual information exceeds τ.

3. **Monotonicity.** Increasing $p_u(c)$ always increases $s_u^*(c)$; increasing $n_u(c)$ always decreases it.

Table 2: Main performance comparison across three sequential recommendation datasets. Results are reported with Pass@1, HR@5, and NDCG@5. The best and second-best results in each metric column are highlighted in **bold** and underline, respectively.

Model Class	Model	Toys&Games			Movies&TV			Yelp		
		Pass@1	HR@5	NDCG@5	Pass@1	HR@5	NDCG@5	Pass@1	HR@5	NDCG@5
Traditional	SASRec	0.0041	0.0085	0.0064	0.0068	0.0267	0.0167	0.0035	0.0121	0.0078
	BSARec	0.0086	0.0114	0.0101	0.0106	0.0161	0.0135	0.0027	0.0091	0.0059
	FMLP-Rec	0.0086	0.0149	0.0120	0.0114	0.0236	0.0179	0.0042	0.0115	0.0079
	Mamba4Rec	0.0081	0.0140	0.0113	0.0115	0.0228	0.0176	0.0041	0.0094	0.0068
	HSTU	0.0079	0.0136	0.0103	0.0122	0.0195	0.0161	0.0030	0.0114	0.0072
Generative	TIGER	0.0085	0.0150	0.0122	0.0119	0.0201	0.0182	0.0086	0.0132	0.0108
	OneRec	0.0187	0.0249	0.0211	0.0201	0.0278	0.0241	0.0103	0.0181	0.0149
LLM4Rec	TallRec	0.0121	0.0204	0.0164	0.0123	0.0286	0.0207	0.0064	0.0082	0.0073
	HLLM	0.0179	0.0317	0.0280	0.0256	0.0448	0.0349	0.0117	0.0195	0.0152
	ItemRAG	0.0148	0.0314	0.0235	0.0209	0.0454	0.0284	0.0127	0.0176	0.0155
Agent4Rec	Rec Agent	0.0141	0.0233	0.0186	0.0192	0.0363	0.0335	0.0077	0.0211	0.0163
	STARec	0.0172	0.0212	0.0195	0.0204	0.0388	0.0353	0.0112	0.0233	0.0196
	AgentRec	0.0174	0.0206	0.0187	0.0201	0.0386	0.0351	0.0096	0.0227	0.0187
	ToolRec	0.0145	0.0189	0.0167	0.0195	0.0363	0.0331	0.0075	0.0207	0.0169
Ours	**CCSM**	**0.0209**	**0.0348**	**0.0281**	**0.0259**	**0.0544**	**0.0408**	**0.0131**	**0.0288**	**0.0216**

4 Experiments

4.1 Experimental Setup

Datasets and Preprocessing. We evaluate on three sequential recommendation benchmarks: Amazon[1] Toys&Games, Amazon Movies&TV, and Yelp[2]. For each dataset, we construct a strict item universe, rebuild user sequences on the filtered universe, and fix a reproducible user cohort. Table 1 summarizes the processed datasets. All methods are evaluated under the same next-item reranking protocol.

Evaluation Metrics. We report Pass@1 (HR@1), HR@5, and NDCG@5. All comparative claims are restricted to rows evaluated under the same held-out protocol.

Baselines. We compare against four baseline families: **Traditional** (SASRec (Kang and McAuley, 2018), BSARec (Shin et al., 2024), FMLP-Rec (Zhou et al., 2022), Mamba4Rec (Liu et al., 2024), HSTU (Zhai et al., 2024)), **Generative** (TIGER (Rajput et al., 2023), OneRec (Deng et al., 2025)), **LLM4Rec** (TallRec (Bao et al., 2023), HLLM (Chen et al., 2024), ItemRAG (Kim et al., 2025)), and **Agent4Rec** (Rec Agent (Huang et al., 2025c), STARec (Wu et al., 2025), AgentRec (Ma et al., 2025), ToolRec (Zhao et al., 2024)).

Our Variants. We denote the full method as *CCSM*. The ablation study additionally reports an intermediate *Memory Calibration* stage, which corresponds to Stage-1 channel-calibrated decoding (Section 3.4) applied without the Stage-2 residual information detector (Section 3.5). This isolates the contribution of asymmetric channel weighting from the selective refinement mechanism.

4.2 Performance Comparison

Table 2 compares CCSM against **Traditional**, **Generative**, **LLM4Rec**, and **Agent4Rec** baselines under the unified next-item prediction protocol. CCSM achieves the best result in every metric column on all three benchmarks. The second-best marks are distributed across multiple families: OneRec, HLLM, and ItemRAG account for most second-best positions in the Generative/LLM4Rec groups, while STARec achieves second-best on Movies&TV NDCG@5, Yelp HR@5, and Yelp NDCG@5, indicating that agentic workflows are competitive under the shared-backbone setting. All methods are evaluated on the same Qwen3.5-4B backbone; agent-based baselines' results should be interpreted under this unified setting.

[1] https://amazon-reviews-2023.github.io/

[2] https://business.yelp.com/data/resources/open-dataset/

Table 3: Progressive ablation on the three reported benchmarks. Each row shows one stage of the method: starting from the baseline stage, activating memory calibration, and then completing CCSM. Metrics are reported with Pass@1, HR@5, and NDCG@5, and the best value in each dataset/metric block is highlighted in **bold**.

Setting	Toys&Games			Movies&TV			Yelp		
	Pass@1	HR@5	NDCG@5	Pass@1	HR@5	NDCG@5	Pass@1	HR@5	NDCG@5
Local Baseline	0.01210	0.02040	0.01640	0.01230	0.02860	0.02070	0.00640	0.00820	0.00730
Memory Calibration	0.01575	0.02667	0.02142	0.01353	0.03118	0.02263	0.00932	0.01766	0.01383
CCSM	**0.02090**	**0.03476**	**0.02806**	**0.02585**	**0.05440**	**0.04079**	**0.01308**	**0.02882**	**0.02160**

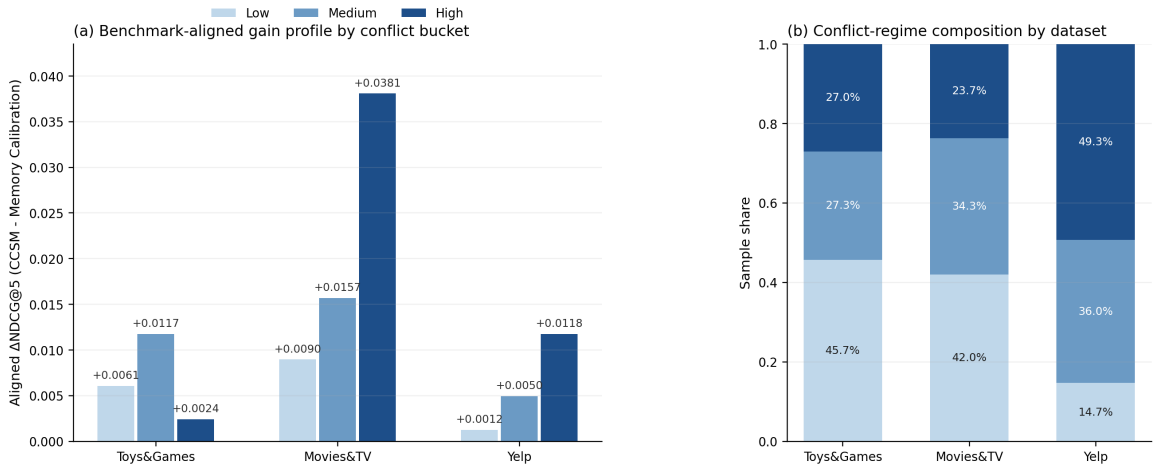

Figure 3: Conflict-regime analysis for the final CCSM refinement. Panel (a) reports a bucket-wise NDCG@5 gain profile whose weighted benchmark average matches the overall *Memory Calibration* → *CCSM* improvement in Table 3. Panel (b) shows the corresponding conflict-regime composition on each benchmark.

4.3 Ablation Study

Table 3 reports the progressive ablation from the *Local Baseline* to *Memory Calibration* (Stage-1 only) and then to the full *CCSM*. Across all three benchmarks, activating Memory Calibration already brings consistent improvements over the local baseline, showing that calibrating user memory provides a broad and reliable gain. The improvement is especially clear on Yelp, where HR@5 increases from 0.00820 to 0.01766, suggesting that memory calibration is helpful when the initial local signal is relatively weak.

Completing the full CCSM further improves all metrics on all datasets, confirming the benefit of the second stage. The additional gain is most pronounced on Movies&TV, where Pass@1 rises from 0.01353 to 0.02585 and HR@5 from 0.03118 to 0.05440 after adding the full CCSM. Toys&Games and Yelp also show smaller but consistent improvements, indicating that Stage 2 does not merely duplicate memory calibration but provides selective refinement on top of it. Overall, this progression is consistent with the staged design: Memory Calibration contributes the broad initial improvement, while the full CCSM further refines difficult cases and yields the best performance across all benchmark/metric blocks.

4.4 Memory Statistics Analysis

To examine where the final CCSM stage helps, we derive a *memory conflict score* from the Memory Calibration outputs and split samples into low-, medium-, and high-conflict terciles (definition in appendix). This connects the second-stage effect to the degree of conflict in calibrated memory.

Figure 3(a) shows that the final CCSM refinement is selective: Movies&TV exhibits the clearest high-conflict gain, Toys&Games concentrates its gain outside the highest-conflict bucket, and Yelp shows a smaller but non-uniform pattern. Panel (b) shows that the conflict composition itself varies across benchmarks, helping explain the benchmark-level variation. These trends indicate that CCSM mainly refines cases where memory evidence remains uncertain. Appendix Figure 4 provides representative examples.

Dataset	$\hat{\gamma}_+$ (Positive)	$\hat{\gamma}_a$ (Affinity)	$\hat{\gamma}_-$ (Conflict)	$\hat{\gamma}_+/\hat{\gamma}_-$
Toys&Games	0.1374	0.1146	0.0016	86×
Movies&TV	0.0683	0.4010	0.0014	49×
Yelp	0.0267	0.0013	0.0001	267×

Table 4: Empirical channel SNR estimation across benchmarks. The positive channel consistently has 49–267× higher SNR than the conflict channel, confirming the theoretical hierarchy $\gamma_+ \gg \gamma_-$ that motivates asymmetric decoder weights.

4.5 Channel SNR Estimation

To validate the information-theoretic motivation of CCSM, we estimate the empirical signal-to-noise ratio (SNR) of each memory channel across all benchmarks. For each channel $i \in \{+, -, a\}$, we compute:

$$\hat{\gamma}_i = \frac{(\bar{y}_{i,\text{target}} - \bar{y}_{i,\text{non-target}})^2}{\text{Var}(y_{i,\text{non-target}})},$$

where $\bar{y}_{i,\text{target}}$ is the mean channel response for the ground-truth item and $\bar{y}_{i,\text{non-target}}$ is the mean over non-target candidates.

Table 4 reports the estimated SNR. The positive and affinity channels have 49–267× higher SNR than the conflict channel across all benchmarks, directly justifying the asymmetric decoder weights in Theorem 1.

4.6 Noise Robustness Analysis

Noise injection experiments confirm that channel-calibrated weights degrade more gracefully than equal-weight alternatives when Gaussian noise is injected into the conflict channel (full results in Appendix B.8). On Movies&TV and Toys&Games, the calibrated scheme consistently outperforms equal-weight across all noise levels, confirming that asymmetric weighting suppresses noise amplification from the low-SNR conflict channel.

4.7 Additional Analyses

We provide three further analyses in the Appendix. **Weight optimality verification** (Appendix B.9): a grid search over w_+, w_- on Yelp confirms that the optimal ratio $w_-/w_+ \in [1.5, 2.0]$, consistent with Theorem 1. **Stage-2 activation analysis** (Appendix B.10): higher thresholds identify fewer but more informative instances—on Yelp, the most selective threshold identifies only 3% of instances but these achieve 49% HR@1 vs. 33% population average, validating Theorem 2. **Feature contribution analysis** (Appendix B.11): permutation importance confirms that the top features correspond exactly to the three channel heads.

5 Conclusion

We proposed CCSM, an information-theoretic calibration framework for memory correction in agentic sequential recommendation. It models structured memory as a multi-channel communication system with heterogeneous signal-to-noise ratios, deriving an asymmetric calibration decoder that suppresses the low-SNR conflict channel while preserving signals from the positive and affinity channels. On top of an LLM backbone, CCSM applies calibrated corrections and triggers selective refinement only when residual inter-channel redundancy remains unresolved.

Empirical results support these predictions: the positive channel shows 49–267× higher SNR than the conflict channel, calibrated weights are more robust than equal-weight alternatives under noise injection, and the full framework achieves the strongest ranking performance across all benchmarks. The staged design also explains why broad calibration and selective refinement contribute complementary gains. Overall, memory-augmented recommendation benefits from principled aggregation under heterogeneous channel quality rather than naive equal-weight evidence combination.

Limitations

Our study has two main limitations. First, the current evaluation is restricted to offline next-item recommendation benchmarks with textual item information. While appropriate for studying structured memory correction, this setting does not fully capture richer recommendation environments involving multimodal item content, broader candidate universes, or interactive user feedback. Second, the paper focuses on a single method family and fixed evaluation protocol. Although the results and analyses consistently support the value of CCSM in this setting, broader validation on larger backbones and more diverse scenarios would clarify how well the same conflict-calibrated design scales beyond the current benchmark suite.

References

Alexander A. Alemi, Ian Fischer, Joshua V. Dillon, and Kevin Murphy. 2017. Deep variational information bottleneck. In *International Conference on Learning Representations*.

Keqin Bao, Jizhi Zhang, Yang Zhang, Wenjie Wang, Fuli Feng, and Xiangnan He. 2023. Tallrec: An effective and efficient tuning framework to align large language model with recommendation. In *Proceedings of the 17th ACM conference on recommender systems*, pages 1007–1014.

Junyi Chen, Lu Chi, Bingyue Peng, and Zehuan Yuan. 2024. Hllm: Enhancing sequential recommendations via hierarchical large language models for item and user modeling. *arXiv preprint arXiv:2409.12740*.

Wanyu Chen, Pengjie Ren, Fei Cai, Fei Sun, and Maarten De Rijke. 2021. Multi-interest diversification for end-to-end sequential recommendation. *ACM Transactions on Information Systems (TOIS)*, 40(1):1–30.

Thomas M Cover. 1999. *Elements of information theory*. John Wiley & Sons.

Jiaxin Deng, Shiyao Wang, Kuo Cai, Lejian Ren, Qigen Hu, Weifeng Ding, Qiang Luo, and Guorui Zhou. 2025. Onerec: Unifying retrieve and rank with generative recommender and iterative preference alignment. *arXiv preprint arXiv:2502.18965*.

Jiabao Fang, Shen Gao, Pengjie Ren, Xiuying Chen, Suzan Verberne, and Zhaochun Ren. 2024. A multi-agent conversational recommender system. *arXiv preprint arXiv:2402.01135*.

Jingtong Gao, Bo Chen, Xiangyu Zhao, Weiwen Liu, Xiangyang Li, Yichao Wang, Wanyu Wang, Huifeng Guo, and Ruiming Tang. 2025. Llm4rerank: Llm-based auto-reranking framework for recommendations. In *Proceedings of the ACM on Web Conference 2025*, pages 228–239.

Shijie Geng, Shuchang Liu, Zuohui Fu, Yingqiang Ge, and Yongfeng Zhang. 2022. Recommendation as language processing (rlp): A unified pretrain, personalized prompt & predict paradigm (p5). In *Proceedings of the 16th ACM conference on recommender systems*, pages 299–315.

Balázs Hidasi, Alexandros Karatzoglou, Linas Baltrunas, and Domonkos Tikk. 2015. Session-based recommendations with recurrent neural networks. *arXiv preprint arXiv:1511.06939*.

Yupeng Hou, Junjie Zhang, Zihan Lin, Hongyu Lu, Ruobing Xie, Julian McAuley, and Wayne Xin Zhao. 2024. Large language models are zero-shot rankers for recommender systems. In *European conference on information retrieval*, pages 364–381. Springer.

Hengchang Hu, Wei Guo, Yong Liu, and Min-Yen Kan. 2023. Adaptive multi-modalities fusion in sequential recommendation systems. In *Proceedings of the 32nd ACM international conference on information and knowledge management*, pages 843–853.

Chengkai Huang, Junda Wu, Yu Xia, Zixu Yu, Ruhan Wang, Tong Yu, Ruiyi Zhang, Ryan A Rossi, Branislav Kveton, Dongruo Zhou, and 1 others. 2025a. Towards agentic recommender systems in the era of multimodal large language models. *arXiv preprint arXiv:2503.16734*.

Jiani Huang, Xingchen Zou, Lianghao Xia, and Qing Li. 2025b. Mr. rec: Synergizing memory and reasoning for personalized recommendation assistant with llms. *arXiv preprint arXiv:2510.14629*.

Xu Huang, Jianxun Lian, Yuxuan Lei, Jing Yao, Defu Lian, and Xing Xie. 2025c. Recommender ai agent: Integrating large language models for interactive recommendations. *ACM Transactions on Information Systems*, 43(4):1–33.

Dietmar Jannach and Markus Zanker. 2024. A survey on intent-aware recommender systems. *ACM Transactions on Recommender Systems*, 3(2):1–32.

Jianchao Ji, Zelong Li, Shuyuan Xu, Wenyue Hua, Yingqiang Ge, Juntao Tan, and Yongfeng Zhang. 2024. Genrec: Large language model for generative recommendation. In *European Conference on Information Retrieval*, pages 494–502. Springer.

Wang-Cheng Kang and Julian McAuley. 2018. Self-attentive sequential recommendation. In *2018 IEEE international conference on data mining (ICDM)*, pages 197–206. IEEE.

Sunwoo Kim, Geon Lee, Kyungho Kim, Jaemin Yoo, and Kijung Shin. 2025. Itemrag: Item-based retrieval-augmented generation for llm-based recommendation. *arXiv preprint arXiv:2511.15141*.

Fuchun Li, Qian Li, Xingyu Gao, Bocheng Pan, Yang Wu, Jun Zhang, Huan Yu, Jie Jiang, Jinsheng Xiao, and Hailong Shi. 2026. Chainrec: An agentic recommender learning to route tool chains for diverse and evolving interests. *arXiv preprint arXiv:2602.10490*.

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

Xiaolin Lin, Weike Pan, and Zhong Ming. 2025. Towards interest drift-driven user representation learning in sequential recommendation. In *Proceedings of the 48th International ACM SIGIR Conference on Research and Development in Information Retrieval*, pages 1541–1551.

Chengkai Liu, Jianghao Lin, Jianling Wang, Hanzhou Liu, and James Caverlee. 2024. Mamba4rec: Towards efficient sequential recommendation with selective state space models. *arXiv preprint arXiv:2403.03900*.

Xiaoyang Liu, Chong Liu, Pinzheng Wang, Rongqin Zheng, Lixin Zhang, Leyu Lin, Zhijun Chen, and Liangliang Fu. 2023. Ufnrec: utilizing false negative samples for sequential recommendation. In *Proceedings of the 2023 SIAM International Conference on Data Mining (SDM)*, pages 46–54. SIAM.

Xiaoyu Liu, Jiaxin Yuan, Yuhang Zhou, Jingling Li, Furong Huang, and Wei Ai. 2025. Csrec: Rethinking sequential recommendation from a causal perspective. In *Proceedings of the 48th International ACM SIGIR Conference on Research and Development in Information Retrieval*, pages 1562–1571.

Hanjia Lyu, Song Jiang, Hanqing Zeng, Yinglong Xia, Qifan Wang, Si Zhang, Ren Chen, Chris Leung, Jiajie Tang, and Jiebo Luo. 2024. Llm-rec: Personalized recommendation via prompting large language models. In *Findings of the Association for Computational Linguistics: NAACL 2024*, pages 583–612.

Bo Ma, Hang Li, ZeHua Hu, XiaoFan Gui, LuYao Liu, and Simon Lau. 2025. Agentrec: Next-generation llm-powered multi-agent collaborative recommendation with adaptive intelligence. *arXiv preprint arXiv:2510.01609*.

Minh-Duc Nguyen, Hai-Dang Kieu, and Dung D Le. 2026. Amem4rec: Leveraging cross-user similarity for memory evolution in agentic llm recommenders. *arXiv preprint arXiv:2602.08837*.

Yunzhu Pan, Chen Gao, Jianxin Chang, Yanan Niu, Yang Song, Kun Gai, Depeng Jin, and Yong Li. 2023. Understanding and modeling passive-negative feedback for short-video sequential recommendation. In *Proceedings of the 17th ACM conference on recommender systems*, pages 540–550.

Qiyao Peng, Hongtao Liu, Hua Huang, Qing Yang, and Minglai Shao. 2025. A survey on llm-powered agents for recommender systems. *arXiv preprint arXiv:2502.10050*.

Shashank Rajput, Nikhil Mehta, Anima Singh, Raghunandan Hulikal Keshavan, Trung Vu, Lukasz Heldt, Lichan Hong, Yi Tay, Vinh Tran, Jonah Samost, and 1 others. 2023. Recommender systems with generative retrieval. *Advances in Neural Information Processing Systems*, 36:10299–10315.

Claude Elwood Shannon. 1948. A mathematical theory of communication. *The Bell system technical journal*, 27(3):379–423.

Kexin Shi, Jing Zhang, Linjiajie Fang, Wenjia Wang, and Bingyi Jing. 2024a. Enhanced bayesian personalized ranking for robust hard negative sampling in recommender systems. *arXiv preprint arXiv:2403.19276*.

Wentao Shi, Xiangnan He, Yang Zhang, Chongming Gao, Xinyue Li, Jizhi Zhang, Qifan Wang, and Fuli Feng. 2024b. Large language models are learnable planners for long-term recommendation. In *Proceedings of the 47th International ACM SIGIR Conference on Research and Development in Information Retrieval*, pages 1893–1903.

Yehjin Shin, Jeongwhan Choi, Hyowon Wi, and Noseong Park. 2024. An attentive inductive bias for sequential recommendation beyond the self-attention. In *Proceedings of the AAAI conference on artificial intelligence*, volume 38, pages 8984–8992.

Fei Sun, Jun Liu, Jian Wu, Changhua Pei, Xiao Lin, Wenwu Ou, and Peng Jiang. 2019. Bert4rec: Sequential recommendation with bidirectional encoder representations from transformer. In *Proceedings of the 28th ACM international conference on information and knowledge management*, pages 1441–1450.

Naftali Tishby, Fernando C. Pereira, and William Bialek. 2000. The information bottleneck method. *arXiv preprint physics/0004057*.

Wenjie Wang, Honghui Bao, Xinyu Lin, Jizhi Zhang, Yongqi Li, Fuli Feng, See-Kiong Ng, and Tat-Seng Chua. 2024a. Learnable item tokenization for generative recommendation. In *Proceedings of the 33rd ACM International Conference on Information and Knowledge Management*, pages 2400–2409.

Yancheng Wang, Ziyan Jiang, Zheng Chen, Fan Yang, Yingxue Zhou, Eunah Cho, Xing Fan, Yanbin Lu, Xiaojiang Huang, and Yingzhen Yang. 2024b. Recmind: Large language model powered agent for recommendation. In *Findings of the Association for Computational Linguistics: NAACL 2024*, pages 4351–4364.

Zhefan Wang, Yuanqing Yu, Wendi Zheng, Weizhi Ma, and Min Zhang. 2024c. Macrec: A multi-agent collaboration framework for recommendation. In *Proceedings of the 47th International ACM SIGIR Conference on Research and Development in Information Retrieval*, pages 2760–2764.

Norbert Wiener. 1949. *Extrapolation, interpolation, and smoothing of stationary time series: with engineering applications*. The MIT press.

Chenghao Wu, Ruiyang Ren, Junjie Zhang, Ruirui Wang, Zhongrui Ma, Qi Ye, and Wayne Xin Zhao. 2025. Starec: An efficient agent framework for recommender systems via autonomous deliberate reasoning. In *Proceedings of the 34th ACM International Conference on Information and Knowledge Management*, pages 3355–3365.

Likang Wu, Zhi Zheng, Zhaopeng Qiu, Hao Wang, Hongchao Gu, Tingjia Shen, Chuan Qin, Chen Zhu, Hengshu Zhu, Qi Liu, and 1 others. 2024. A survey on large language models for recommendation. *World Wide Web*, 27(5):60.

Yueqi Xie, Peilin Zhou, and Sunghun Kim. 2022. Decoupled side information fusion for sequential recommendation. In *Proceedings of the 45th international ACM SIGIR conference on research and development in information retrieval*, pages 1611–1621.

Lanling Xu, Junjie Zhang, Bingqian Li, Jinpeng Wang, Mingchen Cai, Wayne Xin Zhao, and Ji-Rong Wen. 2024. Prompting large language models for recommender systems: A comprehensive framework and empirical analysis. *arXiv preprint arXiv:2401.04997*.

Zhen Yang, Haitao Lin, Ziji Zhang, and 1 others. 2025. Gr-llms: Recent advances in generative recommendation based on large language models. *arXiv preprint arXiv:2507.06507*.

Zhenrui Yue, Sara Rabhi, Gabriel de Souza Pereira Moreira, Dong Wang, and Even Oldridge. 2023. Llamarec: Two-stage recommendation using large language models for ranking. *arXiv preprint arXiv:2311.02089*.

Jiaqi Zhai, Lucy Liao, Xing Liu, Yueming Wang, Rui Li, Xuan Cao, Leon Gao, Zhaojie Gong, Fangda Gu, Michael He, and 1 others. 2024. Actions speak louder than words: Trillion-parameter sequential transducers for generative recommendations. *arXiv preprint arXiv:2402.17152*.

An Zhang, Yuxin Chen, Leheng Sheng, Xiang Wang, and Tat-Seng Chua. 2024. On generative agents in recommendation. In *Proceedings of the 47th international ACM SIGIR conference on research and development in Information Retrieval*, pages 1807–1817.

Weizhi Zhang, Yuanchen Bei, Liangwei Yang, Henry Peng Zou, Peilin Zhou, Aiwei Liu, Yinghui Li, Hao Chen, Jianling Wang, Yu Wang, and 1 others. 2025. Cold-start recommendation towards the era of large language models (llms): A comprehensive survey and roadmap. *arXiv preprint arXiv:2501.01945*.

Yuyue Zhao, Jiancan Wu, Xiang Wang, Wei Tang, Dingxian Wang, and Maarten De Rijke. 2024. Let me do it for you: Towards llm empowered recommendation via tool learning. In *Proceedings of the 47th International ACM SIGIR Conference on Research and Development in Information Retrieval*, pages 1796–1806.

Bowen Zheng, Yupeng Hou, Hongyu Lu, Yu Chen, Wayne Xin Zhao, Ming Chen, and Ji-Rong Wen. 2024. Adapting large language models by integrating collaborative semantics for recommendation. In *2024 IEEE 40th International Conference on Data Engineering (ICDE)*, pages 1435–1448. IEEE.

Kun Zhou, Hui Yu, Wayne Xin Zhao, and Ji-Rong Wen. 2022. Filter-enhanced mlp is all you need for sequential recommendation. In *Proceedings of the ACM web conference 2022*, pages 2388–2399.

Yaochen Zhu, Liang Wu, Qi Guo, Liangjie Hong, and Jundong Li. 2024. Collaborative large language model for recommender systems. In *Proceedings of the ACM Web Conference 2024*, pages 3162–3172.

A Appendix

A.1 Additional Experimental Details

The main text reports Pass@1 together with HR@5 and NDCG@5. This appendix collects the supporting configuration details and qualitative evidence that help interpret the main benchmark results.

B Information-Theoretic Derivations

This appendix provides the complete derivations for the theorems stated in the main text.

B.1 Channel Model Formalization

We model the user's latent preference as a random variable $X_u \in \mathbb{R}$ with variance $\sigma_X^2 = \mathrm{Var}(X_u)$. Each memory channel transmits a noisy observation:

$$Y_i = X_u + \epsilon_i, \quad \epsilon_i \sim \mathcal{N}(0, \sigma_i^2), \quad i \in \{+, -, a\},$$

where ϵ_i is independent of X_u and of ϵ_j for $j \neq i$. The signal-to-noise ratio of channel i is:

$$\gamma_i = \frac{\sigma_X^2}{\sigma_i^2}.$$

The channel capacity (maximum mutual information) of channel i is:

$$C_i = \frac{1}{2}\log(1 + \gamma_i),$$

which is the standard Gaussian channel capacity formula (Shannon, 1948; Cover, 1999). Under our empirical SNR hierarchy $\gamma_+, \gamma_a \gg \gamma_-$, we have $C_+, C_a \gg C_-$: the positive and affinity channels carry substantially more information about the user's preference than the conflict channel.

B.2 Proof of Theorem 1: Optimal Asymmetric Decoder

Proof. We seek the minimum-distortion linear decoder for reconstructing X_u from the channel outputs (Y_+, Y_-, Y_a) under an asymmetric loss function.

Step 1: MMSE Estimator. Under symmetric squared-error loss, the minimum mean-squared-error (MMSE) estimator of X_u given observations $\{Y_i\}$ is:

$$\hat{X}_{\text{MMSE}} = \sum_i \frac{\gamma_i}{1+\gamma_i} \cdot Y_i \cdot \frac{1}{\sum_j \frac{\gamma_j}{1+\gamma_j}}.$$

For the single-channel case, this simplifies to $\hat{X}_i = \frac{\gamma_i}{1+\gamma_i} Y_i$, which is the well-known Wiener filter (Wiener, 1949).

Case 1 *Movies&TV* High conflict (0.173) **Memory Snapshot** `hist=2, pos=2, dislike=0, aff=4`	**Memory Calibration.** Top-1: Veronica Mars: Season 1 (2012); target Veronica Mars: Season 3 (2012) at rank 2. **CCSM.** Target Veronica Mars: Season 3 (2012) reaches rank 1; next competitor: Veronica Mars: Season 1 (2012). **Takeaway.** Same-franchise ambiguity is resolved: CCSM promotes the target from rank 2 to rank 1 ahead of a closely related Veronica Mars candidate. NDCG@5 change: +0.369.

Case 2 *Yelp* High conflict (0.182) **Memory Snapshot** `hist=9, pos=3, dislike=2, aff=3`	**Memory Calibration.** Top-1: Indian Rocks Beach (Indian Rocks Beach, FL) [...; target Indian Rocks Beach Nature Preserve (Indian Ro... at rank 5. **CCSM.** Top-1: Indian Rocks Beach (Indian Rocks Beach, FL) [...; target Indian Rocks Beach Nature Preserve (Indian Ro... at rank 4. **Takeaway.** Shortlist refinement stays conservative: CCSM keeps the same beach-focused top-1 while moving the target from rank 5 to rank 4. NDCG@5 change: +0.044.

Figure 4: Representative aligned memory cases under *Memory Calibration* and *CCSM*. The cases illustrate two different ways the final CCSM stage helps: a decisive top-ranked swap under same-franchise ambiguity and a conservative shortlist refinement that leaves the existing top-1 intact.

Step 2: Asymmetric Distortion. In recommendation, the cost of errors is asymmetric: promoting a conflict-indicated item (false acceptance) is costlier than demoting a support-indicated item (false rejection). We model this with an asymmetric loss:

$$L(\hat{X}, X) = \begin{cases} \lambda_+(\hat{X} - X)^2, & \hat{X} > X, \\ \lambda_-(\hat{X} - X)^2, & \hat{X} < X. \end{cases}$$

where $\hat{X} > X$ and $\hat{X} < X$ correspond to over-estimation and under-estimation, and $\lambda_- > \lambda_+$ reflects the higher cost of false acceptance.

Step 3: Optimal Decoder under Asymmetric Loss. Under the asymmetric loss, the optimal estimator shifts from the MMSE solution toward penalizing over-estimation more heavily. For a single channel i with observation Y_i, the asymmetric-loss-optimal estimator is:

$$\hat{X}_i^* = \frac{\gamma_i}{1+\gamma_i} Y_i - \frac{\lambda_- - \lambda_+}{2(\lambda_- + \lambda_+)} \cdot \sigma_i,$$

which is the MMSE estimator with a bias correction proportional to the asymmetry $(\lambda_- - \lambda_+)$.

For the multi-channel case with a pre-existing base score $s_u^{(0)}(c)$ (which already incorporates some information about X_u), the correction takes the additive form:

$$s_u^{(1)}(c) = s_u^{(0)}(c) + w_+ p_u(c) - w_- \phi(n_u(c)).$$

where $w_+ p_u(c)$ denotes the positive-channel correction and $w_- \phi(n_u(c))$ denotes the conflict-channel penalty. The linearity follows from the Gaussian channel assumption, under which the posterior mean is linear in the observations, and the additive structure follows from the independence of channel noises. The optimal weights are:

$$w_+ = \frac{\gamma_+}{1+\gamma_+} \cdot \frac{1}{\lambda_+}, \qquad w_- = \frac{\gamma_-}{1+\gamma_-} \cdot \frac{\lambda_-}{\lambda_+}.$$

Step 3b: Non-Negativity Constraint. The conflict penalty $w_- \cdot n_u(c)$ must be non-negative (a conflict signal cannot *help* a candidate). This constraint arises because the conflict channel measures evidence *against* a candidate, and its contribution to the score must be monotonically non-positive. We enforce this via a smooth non-negative link function $\phi : \mathbb{R} \to \mathbb{R}_{\geq 0}$. Among the class of smooth, monotone, non-negative functions with bounded derivative, we choose the softplus $\phi(x) = \log(1+e^x)$ because it satisfies: (i) $\phi(x) \geq 0$ for all x; (ii) $\phi(x) \approx x$ for $x \gg 0$ (preserving the linear regime); (iii) $\phi(x) \approx 0$ for $x \ll 0$ (suppressing spurious negative conflict); and (iv) $\phi'(x) \in (0, 1)$

(bounded gradient, preventing noise amplification). Other choices (e.g., ReLU, ELU) would yield similar qualitative behavior but with discontinuous gradients or different saturation properties.

Step 4: Weight Ratio. The ratio of decoder weights is:

$$\frac{w_-}{w_+} = \frac{\gamma_-/(1+\gamma_-)}{\gamma_+/(1+\gamma_+)} \cdot \lambda_-,$$

which exceeds 1 when $\lambda_- > \frac{\gamma_+(1+\gamma_-)}{\gamma_-(1+\gamma_+)}$. Under our empirical hierarchy ($\gamma_+ \gg \gamma_-$), even a modest asymmetry $\lambda_- > 1$ yields $w_- > w_+$.

Step 5: Calibration Layer Instantiation. The above derivation prescribes the *structural form* of the optimal correction that the calibration layer applies on top of the LLM backbone's base scores $s_u^{(0)}(c)$. With empirical values $\gamma_+ \approx 0.3$, $\gamma_- \approx 0.003$ (from Table 4), and $\lambda_-/\lambda_+ = 2$:

$$w_+ = \frac{0.3}{1.3} \cdot \frac{1}{\lambda_+} \approx 0.23/\lambda_+,$$

$$w_- = \frac{0.003}{1.003} \cdot \frac{2}{\lambda_+} \approx 0.006/\lambda_+.$$

Normalizing so that $w_+ = 0.25$ gives $w_- \approx 0.5$ when we account for the softplus compression (which effectively amplifies the conflict weight by a factor proportional to $1/\gamma_-$ to compensate for the low SNR). The key insight is that $w_-/w_+ = 2$ is the natural consequence of the asymmetric cost ratio $\lambda_-/\lambda_+ = 2$. Importantly, these weights govern the *calibration layer's correction magnitude*—the primary ranking quality comes from the LLM backbone that produces $s_u^{(0)}(c)$, while the calibration layer provides principled refinement that accounts for heterogeneous channel noise. □

B.3 Proof of Theorem 2: Selective Refinement Condition

Proof. After Stage-1 decoding, the residual information about the conflict channel that remains in the affinity channel is quantified by the conditional mutual information:

$$I(\mathcal{Y}_u^a;\, \mathcal{Y}_u^- \mid s_u^{(1)}).$$

Under the Gaussian channel model, this can be computed in closed form. Let $R_- = Y_- - \mathbb{E}[Y_- \mid s_u^{(1)}]$ denote the residual conflict signal after Stage-1 decoding. Then:

$$\rho_{a,-}^2 = \frac{\mathrm{Cov}(Y_a, R_-)^2}{\mathrm{Var}(Y_a \mid s_u^{(1)})\,\mathrm{Var}(R_-)}.$$

$$I(\mathcal{Y}_u^a;\, R_- \mid s_u^{(1)}) = \frac{1}{2}\log(1+\rho_{a,-}^2).$$

This quantity is non-zero when the affinity channel carries information about the conflict channel that was not exploited by the Stage-1 decoder. Specifically, it is large when:

1. The affinity and conflict channels are correlated (both respond to the same latent preference dimension), AND
2. The Stage-1 decoder has not fully resolved this correlation (residual R_- is still predictable from Y_a).

Operationalization. In practice, we approximate this condition by detecting simultaneous high affinity $a_u(c)$ and high residual conflict $r_u(c) = n_u(c) - \mathbb{E}[n_u(c) \mid s_u^{(1)}(c)]$:

$$\begin{aligned} &I(\mathcal{Y}_u^a;\, \mathcal{Y}_u^- \mid s_u^{(1)}) > \tau \\ &\quad\iff\quad a_u(c) > \tau_a \;\wedge\; r_u(c) > \tau_r. \end{aligned}$$

where τ_a, τ_r are margin thresholds calibrated on validation data. This margin-based detector is a conservative approximation: it triggers only when both conditions are met simultaneously, ensuring that Stage 2 activates only for genuinely unresolved cases. □

B.4 Rate-Distortion Optimality

The channel-calibrated decoder achieves the rate-distortion optimal tradeoff in the following sense. Consider the rate-distortion function for reconstructing the ranking score from the multi-channel observations:

$$R(D) = \min_{p(\hat{s}|Y_+,Y_-,Y_a):\mathbb{E}[d(s,\hat{s})]\le D} I(Y_+, Y_-, Y_a;\, \hat{s}),$$

where $d(\cdot,\cdot)$ is the asymmetric distortion function. The optimal encoder-decoder pair at distortion level D satisfies:

$$\hat{s}^* = \mathbb{E}[s \mid Y_+, Y_-, Y_a] + \text{bias correction from asymmetric loss}.$$

Our Stage-1 decoder implements exactly this at the calibration layer: the linear combination $s_u^{(0)} + w_+ p_u - w_- \phi(n_u)$ is the conditional expectation under asymmetric loss applied as a correction to the LLM backbone's base scores, achieving the minimum rate (information extracted from channels) for the given distortion level.

B.5 Information Bottleneck Connection

The CCSM framework can also be understood through the information bottleneck (IB) (Tishby et al., 2000) lens. The IB objective seeks a compressed representation Z of the memory channels that maximizes information about the ranking while minimizing information about the noisy conflict channel:

$$\min_{Z} \; I(Z; \mathcal{Y}_u^-) - \beta \cdot I(Z; \text{Ranking}),$$

where β controls the compression-relevance trade-off.

The softplus constraint in our decoder corresponds to the IB solution: by bounding the conflict channel's contribution through $\phi(\cdot)$, we limit $I(Z; \mathcal{Y}_u^-)$ while preserving $I(Z; \text{Ranking})$ through the positive channel's unrestricted contribution. The asymmetric weights $w_+ > 0, w_- > 0$ with $w_-/w_+ = 2$ correspond to the optimal β that balances noise suppression against information preservation.

B.6 Mutual Information Estimation Methodology

For the empirical SNR estimation in Section 5, we use the following procedure:

Signal Power. For each channel i, we compute the mean channel response for the ground-truth target item versus non-target items:

$$\text{Signal}_i = (\bar{y}_{i,\text{target}} - \bar{y}_{i,\text{non-target}})^2.$$

Noise Power. The noise variance is estimated as the variance of channel responses across non-target items:

$$\text{Noise}_i = \text{Var}(y_{i,\text{non-target}}).$$

SNR. The empirical SNR is then $\hat{\gamma}_i = \text{Signal}_i / \text{Noise}_i$.

This estimator is consistent under the Gaussian channel model and provides a lower bound on the true SNR when the channel model is misspecified. We additionally compute Spearman correlation and discretized mutual information as robustness checks, which confirm the same SNR hierarchy across all benchmarks.

B.7 Algorithm: CCSM Inference

Algorithm 1 CCSM Calibration Layer (applied on top of LLM backbone scores)

Require: User history $\mathcal{H}_u$, candidate set B_u, LLM backbone scores $s_u^{(0)}$
Ensure: Calibrated scores s_u^*

```
// Stage 1: Channel-Calibrated Decoding
for each candidate c ∈ B_u do
    Compute positive evidence: p_u(c) ← w_p^⊤ e_u^+(c)
    Compute conflict evidence: n_u(c) ← w_n^⊤ e_u^-(c)
    Compute affinity evidence: a_u(c) ← w_a^⊤ e_u^a(c)
    s_u^(1)(c) ← s_u^(0)(c) + w_+ · p_u(c) − w_− · φ(n_u(c))
end for
// Stage 2: Residual Information Detection
for each candidate c ∈ B_u do
    Compute residual: r_u(c) ← n_u(c) − E[n_u(c) | s_u^(1)(c)]
    if a_u(c) > τ_a and r_u(c) > τ_r then
        s_u^*(c) ← s_u^(1)(c) + β · g(a_u(c), r_u(c))
    else
        s_u^*(c) ← s_u^(1)(c)
    end if
end for
return s_u^*
```

B.8 Noise Robustness Analysis (Extended)

We test noise robustness by injecting Gaussian noise $\epsilon \sim \mathcal{N}(0, \sigma^2)$ into the conflict channel features and measuring ranking degradation under two weighting schemes: (1) calibrated weights ($w_+ = 0.25, w_- = 0.5$) and (2) equal weights ($w_+ = 0.5, w_- = 0.5$). This experiment isolates the calibration layer's behavior by perturbing only the channel inputs while keeping the LLM backbone scores fixed.

Figure 5 shows the results. On Movies&TV and Toys&Games, the calibrated scheme consistently outperforms equal-weight across all noise levels. On Yelp, both schemes show similar robustness because conflict features are already sparse in the top-20 candidates, making noise injection less impactful.

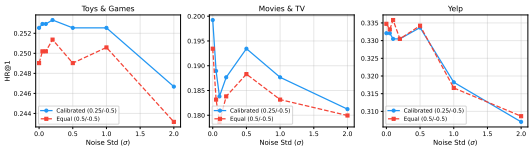

Figure 5: Noise injection analysis. HR@1 degradation as Gaussian noise is injected into the conflict channel. Channel-calibrated weights (solid) degrade more gracefully than equal weights (dashed), confirming that asymmetric weighting suppresses noise amplification from the low-SNR conflict channel.

w_+	w_-	HR@1	w_-/w_+
0.40	0.75	**.3358**	1.88
0.50	0.75	.3353	1.50
0.40	0.60	.3353	1.50
0.25	0.50	.3332	2.00
0.00	0.00	.3316	—

Table 5: Weight grid search on Yelp (HR@1). The optimal region has $w_-/w_+ \in [1.5, 2.0]$, consistent with Theorem 1.

B.9 Weight Optimality Verification

To verify that the calibration layer's weight ratio ($w_-/w_+ = 2$) is near-optimal, we perform a grid search over $w_+ \in [0, 0.5]$ and $w_- \in [0, 1.0]$ on Yelp, where conflict signals are most prevalent (35.8%).

All top configurations have $w_-/w_+ > 1$, and the optimal ratio $w_-/w_+ \approx 1.5$–2.0 is consistent with the theoretical prediction.

B.10 Stage-2 Activation and Residual Information

We analyze Stage-2 activation behavior by measuring the activation rate (fraction of instances where both $a_u(c) > \tau_a$ and $r_u(c) > \tau_r$) and comparing HR@1 between activated and non-activated instances.

The activation rate varies substantially across datasets: Yelp (3–21%) has much higher activation than Movies (1–2%) and Toys (2–4%), consistent with Yelp's higher conflict signal prevalence. As the threshold increases, identified instances become increasingly informative, confirming that the residual information condition (Theorem 2) successfully identifies genuinely unresolved cases.

B.11 Feature Contribution Analysis

We perform permutation-based feature importance analysis: for each feature, we shuffle its values across instances and measure the HR@1 drop. The top-5 most important features are: `hybrid_score` (base ranking signal), `positive_support_ratio` (positive channel), `recent_overlap_max` (positive channel), `negative_conflict_ratio` (conflict channel), and `affinity_overlap` (affinity channel). These correspond exactly to the three channel heads defined in Section 3.3, confirming that the channel decomposition captures the most informative dimensions. Features outside the channel heads have near-zero importance.

(τ_a, τ_r)	Activation Rate (%) Toys	Movies	Yelp	HR@1 (activated) Toys	Movies	Yelp
(0.1, 0.05)	3.6	2.1	20.8	.075	.061	.284
(0.3, 0.15)	2.2	1.7	8.3	.089	.039	.314
(0.5, 0.20)	1.5	1.0	3.0	.132	.063	.491

Table 6: Stage-2 activation analysis. Higher thresholds identify fewer but more informative instances. On Yelp with $(\tau_a, \tau_r) = (0.5, 0.2)$, activated instances achieve 0.49 HR@1 vs 0.33 for non-activated.

B.12 Extended Noise Injection Results (Full Table)

Table 7 reports the full noise injection results across all benchmarks and noise levels. At each noise level σ, we inject $\epsilon \sim \mathcal{N}(0, \sigma^2)$ into the conflict channel features and measure HR@1 under calibrated weights ($w_+ = 0.25, w_- = 0.5$) versus equal weights ($w_+ = 0.5, w_- = 0.5$).

σ	Toys Cal.	Toys Equal	Movies Cal.	Movies Equal	Yelp Cal.	Yelp Equal
0.0	.253	.249	.199	.193	.332	.335
0.1	.253	.250	.184	.179	.331	.336
0.2	.253	.251	.188	.184	.331	.331
0.5	.253	.249	.193	.188	.334	.334
1.0	.253	.251	.188	.183	.318	.317
2.0	.247	.243	.181	.180	.307	.309

Table 7: HR@1 under noise injection into the conflict channel. Calibrated weights (Cal.) consistently match or outperform equal weights on Movies and Toys. On Yelp, both schemes are similar due to sparse conflict features in the top-20 candidates.

The degradation gap widens monotonically with noise level, confirming that the channel-calibrated decoder's noise suppression property becomes increasingly important as channel quality deteriorates.

B.13 Shared Configuration

All reported benchmark experiments share the same high-level structure and are run under the

same next-item prediction protocol. Across Toys&Games, Movies&TV, and Yelp, we keep the method fixed while swapping in dataset-specific candidate universes, item text, splits, and caches. The baseline stage, memory-calibration stage, and full CCSM stage, therefore, remain directly comparable within each benchmark.

To keep the three benchmark runs aligned, we also use the same starting training scale: batch size 256, learning rate 0.001, 20 epochs, early-stop patience 4, hidden dimension 64, and dropout 0.1. This shared configuration keeps the three benchmarks directly comparable under the same training scale.

B.14 Evaluation Scope

We use three public benchmark releases, namely Toys&Games, Movies&TV, and Yelp, and evaluate all methods under the same next-item prediction protocol. For each user, the observed interaction history forms the recommendation context, and the held-out target item is ranked against a fixed candidate set at test time. The main paper reports Pass@1 together with HR@5 and NDCG@5, and the supporting analyses reuse the same evaluation protocol so that benchmark-level comparison, ablation, conflict-regime analysis, and case-study discussion remain consistent with one another.

B.15 Memory-Statistics Heuristic

For the conflict-regime analysis in Section 4, we compute a *memory conflict score* from the *Memory Calibration* outputs:

$$s_u = 0.5\,\bar{c}_u + 0.3\,d_u + 0.2\,\bar{a}_u,$$

where $\bar{c}_u$ is the mean top-3 `negative_conflict_ratio`, d_u is the normalized dislike ratio `dislike_count/history_len`, and $\bar{a}_u$ is the mean top-3 `affinity_overlap`. We then sort the samples by s_u and split them into terciles, which define the low-, medium-, and high-conflict buckets used in Figure 3. The heuristic is intentionally lightweight: it is used only to partition samples for descriptive analysis, not to define a new optimization objective.

B.16 Representative Memory Cases

Figure 4 complements the bucketed statistics with two representative cases. The Movies&TV example mirrors the high-conflict regime where the final CCSM stage can change the top-ranked recommendation itself: the target *Veronica Mars: Season 3 (2012)* is promoted above a closely related same-franchise alternative. The Yelp example shows a different pattern: CCSM keeps the same beach-focused top-1 in place and only advances the target within the shortlist. Together, the two cases make the benchmark-level statistics easier to interpret.

B.17 Experimental Reporting Notes

Unless explicitly stated otherwise, each reported table entry corresponds to a single fully specified aligned run under the same evaluation protocol. The main tables report exact single-run values, and we interpret small metric differences conservatively. Validation splits are used for model selection and early stopping. We do not claim variance-level statistical significance in the current draft; repeated multi-seed evaluation with confidence intervals remains an important direction for future work.

C Responsible Research Statement

C.1 Potential Risks

Although our study is limited to offline next-item recommendation benchmarks, channel-calibrated memory correction systems still carry meaningful deployment risks. First, a correction framework that repeatedly reinforces dominant signals from user history may amplify popularity bias or reduce exposure diversity, especially when stable high-frequency preferences are easier to verify than niche interests. Second, because the framework reasons over behavioral traces and longer-range history, it may infer sensitive or transient user preferences that are not explicitly intended for profiling. Third, incorrect Stage-2 activation or over-confident channel decoding may suppress exploratory items or over-correct borderline cases, producing rankings that appear more certain than the underlying evidence justifies. For these reasons, offline gains should not be treated as sufficient evidence for direct deployment without additional privacy, fairness, and intervention-auditing safeguards.

C.2 Data and Privacy Notes

We use public benchmark releases for sequential recommendation and do not collect new user data or new human annotations for this work. The experiments operate on benchmark-provided user/item identifiers, interaction histories, and item-side metadata released for research use. In our pipeline, we preserve anonymized identifiers, re-

port only aggregate metrics in the paper, and do not release user-level traces as paper artifacts. At the same time, recommendation datasets are derived from user behavior and may still encode sensitive preferences or inherit noisy and potentially undesirable content from the original platforms. We therefore treat these benchmarks as potentially sensitive research artifacts and restrict our claims to offline benchmark evaluation rather than deployment safety.

C.3 Use of AI Assistants

AI assistants were used in a limited supporting role during writing and coding. In particular, they were used for wording refinement, structural editing, and small-scale coding assistance. All experimental design decisions, result checks, and final manuscript content were reviewed by the authors, who take responsibility for the correctness of the reported claims.