ATTRIBUTE-STRUCTURED TEMPORAL RULE REPRESENTATIONS FOR SEQUENTIAL RECOMMENDATION

Anonymous authors
Paper under double-blind review

ABSTRACT

Logic-based sequential recommenders introduce logical operators to impose structured reasoning for better recommendation, but few of them explicitly explore the impact of logical structure on this complex preference learning. Item attributes offer a natural grounding for making the role of logical structure explicit and empirically testable. yet existing methods rarely construct logic at this level. To address this gap, we introduce **TAPR**, a **T**emporal **A**ttribute **P**robabilistic **R**ule model that grounds logical preference rules in explicit item attributes derived from a user's interaction history. Probabilistic conjunction then composes this evidence into a structured rule that directly shapes the final ranking score. Implemented in a Beta representation space, this formulation captures soft preference evidence while remaining closed under probabilistic conjunction, allowing composed rules to stay in the same representation space for ranking. Experiments on three real-world datasets show that TAPR is competitive with strong sequential recommendation baselines. Synthetic experiments further verify that the explicit logical structure contributes to ranking behavior rather than serving only as a descriptive representation. Additional empirical analyses show that the learned rules are behaviorally faithful and that their utility depends on data availability and attribute-combination difficulty.

1 INTRODUCTION

Sequential recommendation seeks to model users' evolving preferences from historical interactions, where both interacted items and their dependencies across time provide evidence for future choices. Most sequential recommenders capture such evidence through latent sequence representations (Kang & McAuley, 2018; Yuan et al., 2019; Fan et al., 2022), which are effective at learning behavioral patterns but provide limited support for explicit reasoning. This has motivated a line of research that introduces logical structure into recommendation reasoning. Existing efforts along this direction can be broadly grouped into two categories: neural-symbolic recommenders use neural logical operators to compose preference evidence (Chen et al., 2021; Yuan et al., 2023), while rule-based collaborative filtering models organize attribute-level evidence into explicit rules that can support intrinsic explanations (Zhang et al., 2022; 2024).

However, attribute-level rule structure remains largely absent from sequential preference modeling. Existing neural-symbolic sequential recommenders typically perform logical reasoning over item-level representations (Yuan et al., 2023), leaving the semantic constituents of their logical expressions ungrounded in the attributes of individual interactions. Attribute-enhanced sequential recommenders, in contrast, incorporate attribute information into sequence modeling, but usually absorb its effects into latent representations rather than treating attributes as identifiable units of rule composition (Zhang et al., 2019; Xie et al., 2022). Consequently, existing sequential models do not explicitly represent how attribute evidence from different interaction events is combined into a preference rule. Item attributes provide a natural grounding for addressing this gap: they offer semantically identifiable evidence that can be associated with specific events, composed into rules, and related across the interaction history. Such a formulation also makes the resulting structure directly inspectable, enabling analysis of not only whether it improves ranking, but how it participates in the ranking process. This suggests a sequential rule representation in which attribute evidence remains explicit, while temporal structure determines how that evidence contributes to the current preference rule.

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.

Realizing such a representation raises three technical challenges. First, attribute information should be represented as identifiable pieces of logical evidence that preserve fine-grained information from each interaction. Second, the logical operator must be context-dependent, so that the contribution of each piece of evidence is determined by the user's interaction history rather than by context-free set composition. Third, the operator must explicitly capture the temporal structure of sequential recommendation, allowing relations across interaction events to directly shape how grounded evidence is combined. We address these challenges with TAPR, a sequential recommendation model that constructs a structured probabilistic rule representation from attribute-grounded logical evidence. First, item attributes are grounded as logical evidence, with item identities retained as complementary evidence. Each piece of evidence is represented by a factorized Beta distribution, and a weighted product-of-experts provides a closed-form conjunction that keeps individual evidence and their compositions in the same representation space. Second, a causal convolution estimates the local contextual relevance of each piece of evidence. Third, same-type and strict ordered relation channels condition the weights with which evidence enters the user rule, allowing temporal structure to shape composition without changing the conjunction operator itself. In this way, TAPRturns rule formation into an identifiable computational process rather than an implicit consequence of sequence encoding. This makes the imposed structure directly testable: its contribution can be traced, intervened on, and evaluated against ranking behavior, allowing us to study explicit rules as an inductive bias rather than merely as an architectural component.

Our work makes three main contributions:

- **Attribute-grounded rule representation.** We introduce a probabilistic representation that organizes event-bound attribute evidence into explicit preference rules for sequential recommendation.
- **Temporally conditioned rule composition.** We develop a composition mechanism in which local context and inter-event relations determine how historical evidence contributes to the user rule.
- **Structurally testable rule formation.** TAPR preserves an explicit evidence-to-rule path, enabling targeted interventions and behavioral analysis of how the learned structure affects ranking.

Accordingly, we evaluate TAPR on three real-world datasets and a synthetic dataset. TAPR remains competitive with strong sequential recommendation baselines, while controlled interventions show that its explicit rule structure participates in ranking decisions. On MovieLens, the data-efficiency analysis further suggests that this inductive bias is most consistently useful for compositionally difficult targets: improvements appear first in catalog-wide rank and extend to top-K retrieval as supervision increases.

2 Related Work

Sequential and Logic-based Recommendation. Sequential recommenders summarize ordered histories as latent states, such as using self-attention in SASRec (Kang & McAuley, 2018), causal convolutions in NextItNet (Yuan et al., 2019), attribute-aware attention in FDSA and DIF-SR (Zhang et al., 2019; Xie et al., 2022), or stochastic representations in STOSA (Fan et al., 2022). Logic-based approaches instead compose explicit logical variables for cognition capacities: NCR learns neural logical operators over user–item interactions (Chen et al., 2021), while SR-PLR represents sequential preference evidence with Beta distributions (Yuan et al., 2023), following related uses of Beta embeddings for logical query composition (Ren & Leskovec, 2020). Latent sequential models retain evidence primarily in a hidden state, whereas existing logic-based sequential models compose item-level variables without grounding them in typed attributes at particular events. TAPR combines these concerns by grounding probabilistic evidence in event–attribute pairs and using temporal relations to condition their composition into a history-level user rule.

Explicit Rule-based Recommendation. Rule-based neuro-symbolic recommenders expose an intermediate rule, expression, or reasoning path so that a prediction can be associated with semantically identifiable evidence. NS-ICF learns attribute-level logical structures for interpretable collaborative filtering (Zhang et al., 2022), and FENCR augments neural collaborative reasoning with feature-aware

expressions for explainable recommendation (Zhang et al., 2024). NLQ4Rec expresses recommendation as a neural logic query (Wu et al., 2024), while LOGER constructs neural reasoning paths for faithful explanations (Zhu et al., 2021). This line differs from logic-based sequential recommendation in its primary emphasis: the former exposes a semantic intermediate object for explanation, whereas the latter focuses on composing logical preference evidence over a sequence. Existing rule-based models, however, do not explicitly represent how attribute evidence attached to different interaction events is combined through temporal relations. TAPR connects the two directions by constructing an attribute-grounded, history-level rule for sequential prediction and evaluating whether its components have measurable effects on ranking behavior.

Relational Inductive Biases. Relational inductive-bias methods make relations between objects an explicit part of the computation rather than leaving them entirely implicit in object representations. The Abstractor architecture introduces relational cross-attention to separate relational information from object-level features in Transformer reasoning (Altabaa et al., 2024). TAPR shares the goal of explicit relational modeling but uses a different mechanism and target: its relation channels operate on typed evidence bound to historical events and use learned lags to condition rule composition. This specialization makes event order part of rule construction and permits targeted interventions on individual relation channels.

3 METHODOLOGY

3.1 OVERVIEW

Problem formulation. Let $\mathcal{U}$ and $\mathcal{I}$ denote the user and item sets. For a user $u \in \mathcal{U}$, the interaction history available when predicting the item at step t is the chronologically ordered prefix

$$\mathcal{H}_{u,t} = (i_{u,1}, i_{u,2}, \ldots, i_{u,t-1}), \qquad i_{u,s} \in \mathcal{I}, \tag{1}$$

where s records interaction order. We retain at most the L most recent events. Each item i is associated with values from G attribute types: $\mathcal{V}_g$ denotes the vocabulary of type g, and $\mathcal{A}_{i,g} \subseteq \mathcal{V}_g$ denotes the values carried by i. This definition covers both single- and multi-valued attributes. For a uniform treatment of item identities and attributes, we reserve $g = 0$ for identity, set $\mathcal{A}_{i,0} = \{i\}$, and write $\mathcal{G}_i = \{g : \mathcal{A}_{i,g} \neq \varnothing\}$ for the types available for item i.

Given $\mathcal{H}_{u,t}$ and a candidate set $\mathcal{C}_{u,t} \subseteq \mathcal{I}$ containing the observed next item $i_{u,t}$, the task is to assign each candidate $c \in \mathcal{C}_{u,t}$ a score $s_{u,t}(c)$. Candidates are ranked in descending order of this score, with the objective of placing $i_{u,t}$ ahead of non-interacted items. TAPR constructs $s_{u,t}(c)$ by first representing the retained history as a rule over event-bound identity and attribute evidence, and then measuring how well the candidate agrees with that rule.

Architecture overview. As shown in Figure 1, TAPR follows three stages. First, the attribute-structured probabilistic conjunction maps item identities and typed attribute values to factorized Beta distributions. Values belonging to the same attribute type are combined into event-bound evidence while remaining in the same distribution family. Second, the temporal relation-conditioned rule representation summarizes this evidence and estimates its contribution to the current preference rule. A causal convolution captures local historical context, while separate same-type and strict THEN channels encode relations within an attribute type and across ordered interaction events. Their outputs determine the composition weights used to form the history-level user rule Q_u^R. Third, rule–candidate matching composes a candidate distribution Q_c^C from its identity and attributes and scores it by the negative KL divergence to Q_u^R. The three stages are detailed in Sections 3.2–3.4.

3.2 ATTRIBUTE-STRUCTURED PROBABILISTIC CONJUNCTION

Probabilistic logical evidence. For an event at position s, attribute type $g \in \mathcal{G}_{i_{u,s}}$, and value $v \in \mathcal{A}_{i_{u,s},g}$, we represent the logical evidence $A_{s,g,v}$ by a product of D independent Beta distributions:

$$\boldsymbol{h}_{s,g,v} = \boldsymbol{e}_v + \boldsymbol{e}_g^{\text{type}} + \boldsymbol{p}_{t-s},$$
$$[\tilde{\boldsymbol{\alpha}}_{s,g,v}; \tilde{\boldsymbol{\beta}}_{s,g,v}] = W_B \boldsymbol{h}_{s,g,v} + \boldsymbol{b}_B.$$

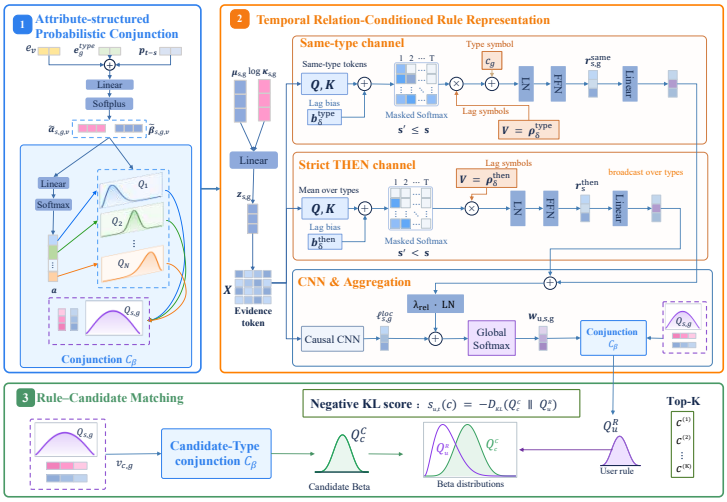


Figure 1: Overview of TAPR.

$$\begin{aligned}\boldsymbol{\alpha}_{s,g,v} &= 1 + \mathrm{softplus}(\tilde{\boldsymbol{\alpha}}_{s,g,v}) + \epsilon,\\ \boldsymbol{\beta}_{s,g,v} &= 1 + \mathrm{softplus}(\tilde{\boldsymbol{\beta}}_{s,g,v}) + \epsilon,\\ Q_{s,g,v} &= \mathrm{Beta}^{\otimes D}(\boldsymbol{\alpha}_{s,g,v}, \boldsymbol{\beta}_{s,g,v}).\end{aligned} \tag{2}$$

Here $\boldsymbol{e}_v$, $\boldsymbol{e}_g^{\mathrm{type}}$, and $\boldsymbol{p}_{t-s}$ encode the value, attribute type, and recency, respectively, and $\epsilon > 0$ is a numerical constant. The same map without $\boldsymbol{p}_{t-s}$ produces position-independent candidate evidence $Q_{c,g,v}$. In each latent dimension, the mean $\mu = \alpha/(\alpha+\beta)$ represents the direction of the learned evidence and the concentration $\kappa = \alpha + \beta$ represents its strength; neither is interpreted as a calibrated truth probability.

Probabilistic conjunction. Following the probabilistic intersection operator of BetaE (Ren & Leskovec, 2020), we combine Beta evidence through a normalized weighted product. For non-negative weights a_n, this operator remains in the same distribution family:

$$\mathcal{C}_\beta(\{Q_n\}, \{a_n\}) \equiv \mathrm{Normalize}\left[\prod_{n=1}^{N} Q_n^{a_n}\right]. \tag{3}$$

For $Q_n = \mathrm{Beta}^{\otimes D}(\boldsymbol{\alpha}_n, \boldsymbol{\beta}_n)$, its closed-form parameters are

$$\begin{aligned}\boldsymbol{\alpha}_\cap &= 1 + \sum_n a_n(\boldsymbol{\alpha}_n - \mathbf{1}),\\ \boldsymbol{\beta}_\cap &= 1 + \sum_n a_n(\boldsymbol{\beta}_n - \mathbf{1}),\\ \mathcal{C}_\beta(\{Q_n\}, \{a_n\}) &= \mathrm{Beta}^{\otimes D}(\boldsymbol{\alpha}_\cap, \boldsymbol{\beta}_\cap).\end{aligned} \tag{4}$$

We use this operator as a probabilistic conjunction, without claiming Boolean equivalence; Appendix B provides the normalization proof and interpretation limits.

For a multi-valued type g, learnable simplex weights combine its value-level distributions through $\mathcal{C}_\beta$ into $Q_{s,g}$, with the same operation producing $Q_{c,g}$ for a candidate.

3.3 Temporal Relation-Conditioned Rule Representation

Local evidence summary. We summarize each Beta evidence distribution by retaining its mean and concentration:

$$\boldsymbol{\mu}_{s,g} = \frac{\boldsymbol{\alpha}_{s,g}}{\boldsymbol{\alpha}_{s,g} + \boldsymbol{\beta}_{s,g}}, \qquad \boldsymbol{\kappa}_{s,g} = \boldsymbol{\alpha}_{s,g} + \boldsymbol{\beta}_{s,g}.$$

$$\boldsymbol{z}_{s,g} = W_{\text{tok}}[\boldsymbol{\mu}_{s,g}; \log \boldsymbol{\kappa}_{s,g}] + \boldsymbol{b}_{\text{tok}}. \tag{5}$$

Tokens from one event are concatenated in a fixed type order, with zeros for absent types, to form $\boldsymbol{x}_s$. A causal convolution then produces a local relevance logit without accessing future events:

$$\boldsymbol{h}_{1:L} = \text{CausalCNN}(\boldsymbol{x}_{1:L}), \qquad \ell^{\text{loc}}_{s,g} = \boldsymbol{w}_g^\top \boldsymbol{h}_s + b_g. \tag{6}$$

Temporally conditioned relations. Inspired by the relational cross-attention of Abstractors (Altabaa et al., 2024), we use content similarity to select related evidence while aggregating learned relation symbols rather than content values. The same-type channel relates $\boldsymbol{z}_{s,g}$ to earlier or current evidence of type g. The THEN channel relates the event summary $\bar{\boldsymbol{z}}_s = |\mathcal{G}_{i_{u,s}}|^{-1} \sum_g \boldsymbol{z}_{s,g}$ only to earlier events. For channel $\chi \in \{\text{same}, \text{then}\}$, let j denote its target, $\mathcal{S}_\chi(j)$ its causal source set, and $\delta(j,k)$ the lag from source k to target j. Multi-head relation aggregation is

$$e^{\chi,m}_{j,k} = \frac{(W_Q^{\chi,m} \boldsymbol{y}_j)^\top (W_K^{\chi,m} \boldsymbol{y}_k)}{\sqrt{d_k}} + b^{\chi,m}_{\delta(j,k)}, \tag{7a}$$

$$a^{\chi,m}_{j,k} = \text{softmax}_{k \in \mathcal{S}_\chi(j)}(e^{\chi,m}_{j,k}). \tag{7b}$$

$$\boldsymbol{r}^\chi_j = f_\chi \left(\boldsymbol{c}^\chi_j + \text{Concat}_{m=1}^{H} \sum_{k \in \mathcal{S}_\chi(j)} a^{\chi,m}_{j,k} \boldsymbol{\rho}^{\chi,m}_{\delta(j,k)} \right). \tag{7c}$$

The two channels instantiate the generic notation as follows:

Channel	Target and representation	Source set $\mathcal{S}_\chi(j)$	Context
Same-type	$j = (s,g),\ \boldsymbol{y}_j = \boldsymbol{z}_{s,g}$	$\{(s',g) : s' \le s\}$	$\boldsymbol{c}_g$
THEN	$j = s,\ \boldsymbol{y}_j = \bar{\boldsymbol{z}}_s$	$\{s' : s' < s\}$	$\boldsymbol{0}$

The same-type source set includes only retained positions at which type g is present. Here $\boldsymbol{\rho}^{\chi,m}_\delta$ and $b^{\chi,m}_\delta$ are a learned lag symbol and bias, $\boldsymbol{c}_g$ is a type embedding, and f_χ is a learned projection. Thus the query–key similarity selects relation instances, while the lag symbols carry the relation-specific output. Unlike the general Abstractor, TAPR specializes this bottleneck to typed historical evidence and separates same-type from ordered inter-event relations.

The two relation summaries give the relation-conditioned logit

$$\begin{aligned} \boldsymbol{h}^{\text{rel}}_{s,g} &= \text{LN}\big(W_{\text{same}} \boldsymbol{r}^{\text{same}}_{s,g} + W_{\text{then}} \boldsymbol{r}^{\text{then}}_s + \boldsymbol{b}_{\text{rel}}\big), \\ \ell^{\text{rel}}_{s,g} &= \boldsymbol{w}^\top_{\text{rel}} \boldsymbol{h}^{\text{rel}}_{s,g} + b_{\text{rel}}. \end{aligned} \tag{8}$$

We use THEN only as the name of this ordered relation channel; it is not a temporal-logic operator.

Finally, let $\mathcal{E}_{u,t} = \{(s,g) : s \text{ is retained in } \mathcal{H}_{u,t},\ g \in \mathcal{G}_{i_{u,s}}\}$ denote the valid event–type pairs. The local and relational logits are normalized over this set, and the resulting weights instantiate the history-level user rule:

$$\begin{aligned} r_{u,s,g} &= \ell^{\text{loc}}_{s,g} + \lambda_{\text{rel}} \ell^{\text{rel}}_{s,g}, \\ w_{u,s,g} &= \text{softmax}_{(s,g) \in \mathcal{E}_{u,t}}(r_{u,s,g}). \end{aligned} \tag{9}$$

$$Q^R_u = \mathcal{C}_\beta(\{Q_{s,g}\}_{(s,g) \in \mathcal{E}_{u,t}}, \{w_{u,s,g}\}_{(s,g) \in \mathcal{E}_{u,t}}). \tag{10}$$

Here $\lambda_{\text{rel}} \ge 0$ controls the relational residual. Q^R_u is the only history-level object we call a user rule; $Q_{s,g,v}$ and $Q_{s,g}$ remain evidence distributions.

3.4 RULE–CANDIDATE MATCHING

Learned candidate-type logits $\xi_{c,g}$ define weights $\nu_{c,g} = \text{softmax}_{g \in \mathcal{G}_c}(\xi_{c,g})$ that compose the position-independent candidate evidence into

$$Q^C_c = \mathcal{C}_\beta(\{Q_{c,g}\}_{g \in \mathcal{G}_c}, \{\nu_{c,g}\}_{g \in \mathcal{G}_c}). \tag{11}$$

STOSA provides a precedent for ranking items by distances between stochastic sequence and candidate representations (Fan et al., 2022). TAPR adopts this distribution-matching view but

Table 1: Statistics of three public datasets.

Dataset	Users	Items	Train events	Attribute Types
MovieLens	2,113	10,109	851,372	5
H&M	52,880	20,769	4,032,210	11
Yelp	68,008	21,637	1,673,656	15

replaces STOSA's Gaussian/Wasserstein geometry with a directed KL divergence between factorized Beta distributions. Let $Q^C_{c,d}$ and $Q^R_{u,d}$ denote the one-dimensional marginals in dimension d. Then

$$\begin{aligned} s_{u,t}(c) &= -D_{\mathrm{KL}}(Q^C_c \,\|\, Q^R_u) \\ &= -\sum_{d=1}^{D} D_{\mathrm{KL}}\big(Q^C_{c,d} \,\|\, Q^R_{u,d}\big)\,. \end{aligned} \tag{12}$$

This direction measures how costly it is to encode the candidate evidence under the history-specific user rule. Appendix B gives the one-dimensional closed form. The dimension-level terms, composition weights, and relation logits expose internal quantities for controlled analysis, but they do not by themselves establish faithful input attribution; Section 4.4 therefore evaluates faithfulness through interventions.

Training objective. For each valid target $i_{u,t}$, we sample one negative item $c^-_{u,t}$ that does not occur in user u's observed interactions. We optimize the pairwise BPR objective (Rendle et al., 2009):

$$\begin{aligned} \mathcal{L}_{\mathrm{BPR}} &= -\sum_{(u,t)} \log \sigma\big(s_{u,t}(i_{u,t}) - s_{u,t}(c^-_{u,t})\big) \\ &= \sum_{(u,t)} \mathrm{softplus}\big(s_{u,t}(c^-_{u,t}) - s_{u,t}(i_{u,t})\big)\,. \end{aligned} \tag{13}$$

4 Experiments

We evaluate predictive performance and examine how the proposed rule structure shapes ranking behavior. Beyond full-catalog next-item ranking, the experiments test the effect of the explicit rule structure, the behavioral faithfulness of extracted rules, and when logic-structured representation provides a useful inductive bias.

4.1 Experimental Setup

Datasets. We use three public datasets with item-side attributes: MovieLens (Cantador et al., 2011)[1] contains movie interactions and metadata such as genres and cast; H&M (García Ling et al., 2022)[2] contains customer purchase histories and structured product descriptors; and Yelp (Yelp, 2015)[3] contains user–business interactions with business categories and attributes. We order interactions within each user and split them chronologically into training, validation, and test prefixes. For H&M, repeated customer–item interactions on the same day are deduplicated. Yelp retains repeated user–business interactions, so its sequence order reflects the observed order within each user. Table 1 summarizes the processed data; Appendix C provides preprocessing and attribute-construction details.

For mechanism tests, we construct a synthetic dataset with known generative rules and oracle next items. Each factual sequence is paired with counterfactual sequences created by assignment, order, or substitution interventions. Strict-binding and same-type-binding proof sets isolate the two relation families in TAPR. Appendix E specifies the data-generation process and intervention protocol.

[1] https://grouplens.org/datasets/hetrec-2011/

[2] https://www.kaggle.com/competitions/h-and-m-personalized-fashion-recommendations

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

Table 2: Overall performance on all datasets.

	MovieLens		Yelp		H&M	
Model	NDCG@10	HR@10	NDCG@10	HR@10	NDCG@10	HR@10
SASRec	0.0292 ± 0.0019	0.0550 ± 0.0028	0.0266 ± 0.0025	0.0504 ± 0.0047	0.0426 ± 0.0011	0.0791 ± 0.0017
NextItNet	0.0276 ± 0.0025	0.0540 ± 0.0040	0.0274 ± 0.0004	0.0513 ± 0.0009	0.0418 ± 0.0011	0.0779 ± 0.0008
STOSA	0.0300 ± 0.0030	0.0534 ± 0.0023	0.0330 ± 0.0001	0.0604 ± 0.0007	0.0540 ± 0.0013	0.0914 ± 0.0030
SR-PLR	$\underline{0.0325 \pm 0.0019}$	0.0538 ± 0.0015	0.0295 ± 0.0016	0.0567 ± 0.0021	0.0436 ± 0.0011	0.0808 ± 0.0009
Attr-SASRec	0.0279 ± 0.0020	0.0571 ± 0.0035	0.0195 ± 0.0010	0.0405 ± 0.0036	0.0373 ± 0.0024	0.0714 ± 0.0033
Attr-NextItNet	0.0270 ± 0.0026	0.0576 ± 0.0040	0.0280 ± 0.0014	0.0539 ± 0.0024	0.0402 ± 0.0018	0.0766 ± 0.0032
Attr-SR-PLR	0.0302 ± 0.0027	$\underline{0.0589 \pm 0.0039}$	0.0198 ± 0.0028	0.0396 ± 0.0067	0.0382 ± 0.0030	0.0753 ± 0.0058
FDSA	0.0316 ± 0.0026	0.0583 ± 0.0016	0.0197 ± 0.0013	0.0395 ± 0.0018	0.0317 ± 0.0010	0.0631 ± 0.0020
DIF-SR	0.0263 ± 0.0015	0.0522 ± 0.0011	$\underline{0.0399 \pm 0.0004}$	$\underline{0.0721 \pm 0.0010}$	$\underline{0.0593 \pm 0.0019}$	$\underline{0.1057 \pm 0.0050}$
NCR	0.0250 ± 0.0025	0.0510 ± 0.0041	0.0185 ± 0.0006	0.0382 ± 0.0014	0.0196 ± 0.0036	0.0400 ± 0.0057
TAPR	$\mathbf{0.0452 \pm 0.0031}$	$\mathbf{0.0823 \pm 0.0042}$	$\mathbf{0.0401 \pm 0.0008}$	$\mathbf{0.0760 \pm 0.0005}$	$\mathbf{0.0631 \pm 0.0006}$	$\mathbf{0.1120 \pm 0.0024}$
Gain	+39.1%	+39.7%	+0.5%	+5.4%	+6.4%	+6.0%

Baselines. We compare with four ID-based sequential recommenders: SASRec (Kang & McAuley, 2018), NextItNet (Yuan et al., 2019), SR-PLR (Yuan et al., 2023), and STOSA (Fan et al., 2022). Attr-SASRec, Attr-NextItNet, and Attr-SR-PLR extend their ID-only counterparts with the same attribute input: attribute-value embeddings are mean-pooled into an attribute representation, which is then averaged with the item-ID representation before entering the original sequence encoder. FDSA (Zhang et al., 2019) and DIF-SR (Xie et al., 2022) provide additional attribute-based comparisons, while NCR (Chen et al., 2021) represents logic-based recommendation.

Evaluation metrics. We use full-catalog HR@10 and NDCG@10 as the primary ranking metrics. The inductive-bias analysis additionally uses HR@100 and mean rank percentile, where lower values are better; the structural and faithfulness analyses use their respective mechanism-specific metrics.

Implementation details. Model selection is based on validation performance, followed by evaluation on the held-out test suffix. Appendix D reports the primary MovieLens configuration and model-selection protocol. The extraction thresholds used in the faithfulness analysis are selected on validation data.

4.2 Overall Recommendation Performance

Table 2 reports full-catalog next-item ranking on the three real-world datasets. On MovieLens, TAPR improves over the strongest baseline by 39.1% in NDCG@10 and 39.7% in HR@10. It also leads DIF-SR on H&M by 6.4% and 6.0%, respectively. On Yelp, its NDCG@10 is comparable to DIF-SR (0.0401 versus 0.0399), while its HR@10 is 5.4% higher. Thus, the clearest advantage appears on MovieLens, with smaller gains on H&M and Yelp, which means our model is competitive for the best baseline. The subsequent analyses examine the contribution of the rule structure in these ranking tasks.

4.3 Ablation Study

Table 3 separates attribute evidence, probabilistic rule construction, and explicit relation modeling. The ID-only ablation removes attribute evidence. strict-no-rule does not model attributes as Beta distributions and therefore removes the probabilistic rule representation. CNN-only retains Beta conjunction and rule matching but removes the explicit same-type and THEN relation channels, deriving evidence weights only from the local causal encoder. no-THEN, used in the synthetic analysis, removes the THEN relation while retaining the remaining rule structure. TAPR outperforms ID-only and strict-no-rule on both metrics across all three datasets, while CNN-only is lower on MovieLens and H&M but slightly higher on Yelp. The rule representation is therefore supported more consistently than the explicit relation channels, whose contribution is isolated below.

The synthetic dataset supplies oracle targets under paired assignment and order interventions. As Figure 2(a) shows, TAPR, no-THEN, and CNN-only retain the Beta-conjunction rule path and achieve nonzero bidirectional validity on both proof sets, whereas strict-no-rule collapses. This establishes structural responsiveness of the rule path rather than superiority of a particular relation channel. To isolate THEN, we disable it at inference in a frozen TAPR checkpoint and compare the effect with a

Table 3: Ablations on the real-world datasets (mean ± sample standard deviation). Bold marks the best result in each column. "Full vs. best ablation" is the relative difference between TAPR and the strongest ablation; a negative value indicates that an ablation is higher.

	MovieLens		Yelp		H&M	
Variant	NDCG@10	HR@10	NDCG@10	HR@10	NDCG@10	HR@10
TAPR	$\mathbf{0.0452 \pm 0.0031}$	$\mathbf{0.0823 \pm 0.0042}$	0.0401 ± 0.0008	0.0760 ± 0.0005	$\mathbf{0.0631 \pm 0.0006}$	$\mathbf{0.1120 \pm 0.0024}$
CNN-only	0.0398 ± 0.0087	0.0729 ± 0.0122	$\mathbf{0.0408 \pm 0.0013}$	$\mathbf{0.0773 \pm 0.0013}$	0.0612 ± 0.0009	0.1105 ± 0.0024
ID-only	0.0370 ± 0.0031	0.0714 ± 0.0032	0.0371 ± 0.0010	0.0706 ± 0.0026	0.0581 ± 0.0028	0.0984 ± 0.0050
strict-no-rule	0.0385 ± 0.0051	0.0694 ± 0.0074	0.0274 ± 0.0003	0.0542 ± 0.0009	0.0592 ± 0.0014	0.1011 ± 0.0025
Full vs. best ablation	+13.6%	+12.9%	−1.7%	−1.7%	+3.1%	+1.4%

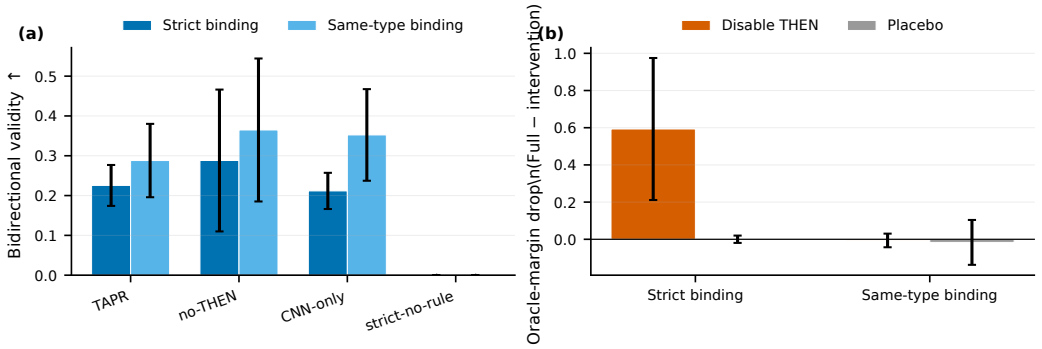


Figure 2: Mechanism evidence on the synthetic dataset. (a) Bidirectional validity across independently trained variants; TAPR, no-THEN, and CNN-only retain the Beta-conjunction rule path, whereas strict-no-rule removes it. (b) Frozen-TAPR intervention, reported as the full-model bidirectional oracle margin minus the margin after disabling THEN or applying a placebo perturbation; positive values indicate that the intervention weakens the oracle margin. Error bars show one sample standard deviation across four runs.

placebo perturbation. On strict-binding pairs, disabling THEN reduces the bidirectional oracle margin by 0.5930, compared with 0.0002 for the placebo, while neither produces a comparable change on same-type-binding pairs (Figure 2(b)). Appendix E reports the complete metrics.

4.4 Behavioral Faithfulness of Extracted Rules

Using validation-selected thresholds, we compare extracted rules with controls matched by event positions and retained component counts. Keep Top-1 Preservation tests whether retained evidence reconstructs the original decision, Hierarchy Completeness tests its event–attribute organization, and Faithful Coverage additionally requires a positive deletion effect. Across all three datasets, extracted rules outperform their matched controls on each metric (Figure 3), indicating that both the selected evidence and its organization are more closely associated with the ranking decision than equally sized alternatives.

Appendix F provides the full protocol and representative success, boundary, and failure cases.

4.5 When Does Logic-Structured Representation Provide a Useful Inductive Bias?

We vary the retained training fraction over $\{5, 10, 25, 50, 100\}\%$ on each dataset and partition targets by attribute-combination difficulty using training-prefix statistics. We compare TAPR with the attribute-based FDSA and DIF-SR and the logic-based NCR. Figure 4 reports mean rank percentile on the hard slice, measuring where the target falls in the full-catalog ranking; lower values are better. Appendix G details the slice construction and complementary MovieLens results.

Across all three datasets, TAPR attains the best mean rank percentile at every training fraction. Its advantage over the strongest comparator is largest at 5% training data on H&M and Yelp, reducing mean rank percentile by 40.7% and 39.2%, respectively; the gains narrow to 1.4% and 2.5% with full

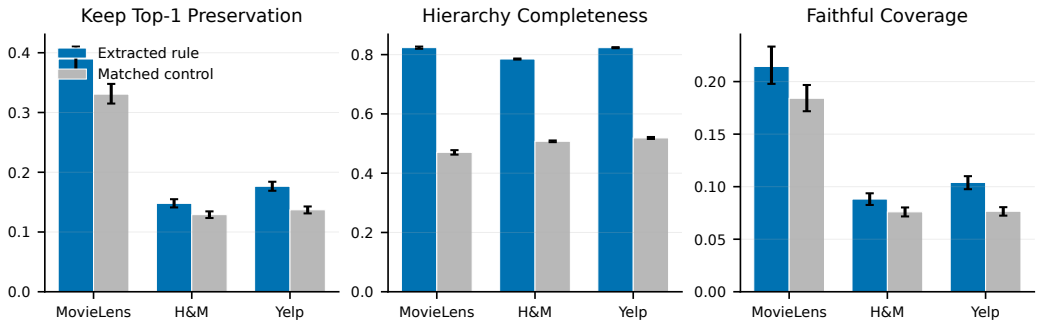


Figure 3: Faithfulness of extracted rules across three datasets. The panels report Keep Top-1 Preservation, Hierarchy Completeness, and Faithful Coverage; higher values are better. Extracted rules are compared with controls matched by event positions and retained component counts, and error bars show 95% user-level bootstrap confidence intervals.

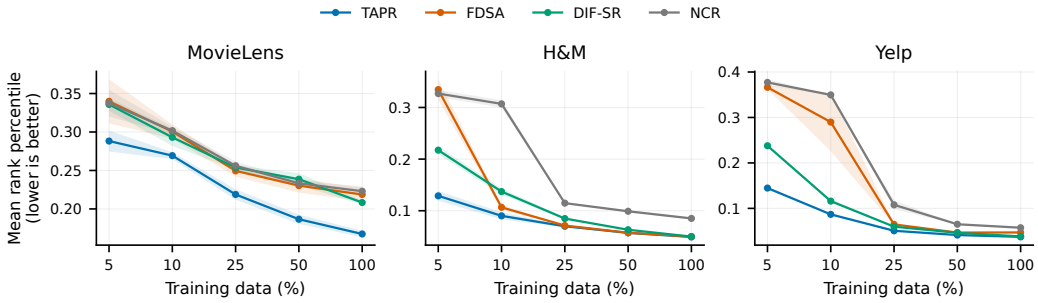


Figure 4: Mean rank percentile on the hard compositional slice across training-data fractions. Lower is better. Shaded bands show one sample standard deviation across runs; all cells contain three runs except H&M NCR at 10%, which contains two valid runs.

training data. On MovieLens, TAPR also remains best across all five fractions. These results show that the rule representation improves catalog-wide ranking for difficult attribute compositions across datasets, with its clearest low-data advantage on H&M and Yelp.

5 Conclusion

We introduced TAPR, a sequential recommender that grounds item identities and typed attributes as probabilistic evidence, uses temporally conditioned relations to compose a history-specific user rule, and ranks candidates by distributional agreement. Across three real-world datasets, TAPR is competitive with strong baselines and achieves the best mean rank percentile on hard compositional slices at every evaluated training fraction. Its largest margins on H&M and Yelp occur with the least training data, providing evidence that the logic-structured representation offers a useful inductive bias under limited supervision. Matched-control and synthetic interventions further indicate that selected evidence, its event-level organization, and the THEN relation make distinct contributions to ranking. Together, these findings support explicit rule structure as a testable modeling bias rather than a universal replacement for latent sequence representations. Future work should study adaptive relation selection and rule construction with incomplete or learned attributes, while continuing to evaluate whether intermediate symbolic structures make reproducible contributions to model decisions.

AI USE STATEMENT

Generative AI tools were used to assist with drafting and revising manuscript prose, organizing the presentation of experimental results, and preparing figure-generation scripts. The authors independently checked all AI-assisted text, code, numerical claims, and figures against the cited sources and stored experimental outputs. Generative AI tools were not used to generate experimental data, select reported results, or replace author judgment in designing or evaluating the proposed method. The authors take responsibility for the final content of this work.

REPRODUCIBILITY STATEMENT

We provide the mathematical definition of TAPR and the closure proof for its probabilistic conjunction in Sections 3 and Appendix B. Appendix C documents dataset preprocessing, chronological splits, leakage controls, and attribute construction; Appendix D records the primary MovieLens configuration, model-selection rules, and randomness protocols. The synthetic dataset construction, intervention protocol, and complete supporting results are specified in Appendix E; Appendices F and G provide the faithfulness protocol, low-data definitions, and full result tables. Each reported experiment is associated with a resolved configuration and evaluation artifact, enabling the reported analyses to be reproduced without using test targets for configuration selection.

REFERENCES

Awni Altabaa, Taylor W. Webb, Jonathan D. Cohen, and John Lafferty. Abstractors and relational cross-attention: An inductive bias for explicit relational reasoning in Transformers. In *The Twelfth International Conference on Learning Representations*, 2024. URL https://openreview.net/forum?id=XNa6r6ZjoB.

Iván Cantador, Peter Brusilovsky, and Tsvi Kuflik. Second workshop on information heterogeneity and fusion in recommender systems (HetRec2011). In *Proceedings of the Fifth ACM Conference on Recommender Systems*, pp. 387–388. ACM, 2011.

Hanxiong Chen, Shaoyun Shi, Yunqi Li, and Yongfeng Zhang. Neural collaborative reasoning. In *Proceedings of the Web Conference 2021*, pp. 1516–1527, 2021. doi: 10.1145/3442381.3449973.

Ziwei Fan, Zhiwei Liu, Yu Wang, Alice Wang, Zahra Nazari, Lei Zheng, Hao Peng, and Philip S. Yu. Sequential recommendation via stochastic self-attention. In *Proceedings of the ACM Web Conference 2022*, pp. 2036–2047, 2022. doi: 10.1145/3485447.3512077.

Carlos García Ling, ElizabethHMGroup, FridaRim, inversion, Jaime Ferrando, Maggie, neuraloverflow, and xlsrln. H&M Personalized Fashion Recommendations. Kaggle competition, 2022. URL https://www.kaggle.com/competitions/h-and-m-personalized-fashion-recommendations.

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

Hongyu Ren and Jure Leskovec. Beta embeddings for multi-hop logical reasoning in knowledge graphs. In *Advances in Neural Information Processing Systems*, volume 33, pp. 19716–19726, 2020.

Steffen Rendle, Christoph Freudenthaler, Zeno Gantner, and Lars Schmidt-Thieme. BPR: Bayesian personalized ranking from implicit feedback. In *Proceedings of the Twenty-Fifth Conference on Uncertainty in Artificial Intelligence*, pp. 452–461. AUAI Press, 2009. URL https://www.auai.org/uai2009/papers/UAI2009_0139_48141db02b9f0b02bc7158819ebfa2c7.pdf.

Maonian Wu, Bang Chen, Shaojun Zhu, Bo Zheng, Wei Peng, and Mingyi Zhang. Neuro-symbolic recommendation model based on logic query. *Knowledge-Based Systems*, 284:111311, 2024. doi: 10.1016/j.knosys.2023.111311.

Yueqi Xie, Peilin Zhou, and Sunghun Kim. Decoupled side information fusion for sequential recommendation. In *Proceedings of the 45th International ACM SIGIR Conference on Research and Development in Information Retrieval*, pp. 1611–1621, 2022. doi: 10.1145/3477495.3531963.

Yelp. Yelp open dataset. Yelp data resource, 2015. URL https://business.yelp.com/data/resources/open-dataset/.

Fajie Yuan, Alexandros Karatzoglou, Ioannis Arapakis, Joemon M. Jose, and Xiangnan He. A simple convolutional generative network for next item recommendation. In *Proceedings of the Twelfth ACM International Conference on Web Search and Data Mining*, pp. 582–590, 2019. doi: 10.1145/3289600.3290975.

Huanhuan Yuan, Pengpeng Zhao, Xuefeng Xian, Guanfeng Liu, Yanchi Liu, Victor S. Sheng, and Lei Zhao. Sequential recommendation with probabilistic logical reasoning. In *Proceedings of the Thirty-Second International Joint Conference on Artificial Intelligence*, pp. 2432–2440, 2023. doi: 10.24963/ijcai.2023/270.

Tingting Zhang, Pengpeng Zhao, Yanchi Liu, Victor S. Sheng, Jiajie Xu, Deqing Wang, Guanfeng Liu, and Xiaofang Zhou. Feature-level deeper self-attention network for sequential recommendation. In *Proceedings of the Twenty-Eighth International Joint Conference on Artificial Intelligence*, pp. 4320–4326, 2019. doi: 10.24963/ijcai.2019/600.

Wei Zhang, Junbing Yan, Zhuo Wang, and Jianyong Wang. Neuro-symbolic interpretable collaborative filtering for attribute-based recommendation. In *Proceedings of the ACM Web Conference 2022*, pp. 3229–3238, 2022. doi: 10.1145/3485447.3512042.

Xiaoyu Zhang, Shaoyun Shi, Yishan Li, Weizhi Ma, Peijie Sun, and Min Zhang. Feature-enhanced neural collaborative reasoning for explainable recommendation. *ACM Transactions on Information Systems*, 43(1):7:1–7:33, 2024. doi: 10.1145/3690381.

Yaxin Zhu, Yikun Xian, Zuohui Fu, Gerard de Melo, and Yongfeng Zhang. Faithfully explainable recommendation via neural logic reasoning. In *Proceedings of the 2021 Conference of the North American Chapter of the Association for Computational Linguistics: Human Language Technologies*, pp. 3083–3090, 2021. doi: 10.18653/v1/2021.naacl-main.245.

A Appendix Roadmap

The appendix separates mathematical justification, reproducibility details, and supporting evidence. Appendix B proves the closure used by the probabilistic conjunction and states its interpretation limits. Appendix C documents real-data preprocessing and Appendix D records implementation choices. Appendix E gives the complete synthetic dataset construction, metrics, and results. Appendix F specifies the faithfulness protocol and case selection. Appendix G gives the low-data slice definition and full results, and Appendix H reports sensitivity analyses.

B Probabilistic Conjunction: Closure and Interpretation

BetaE introduced a learned probabilistic intersection by taking a normalized weighted product of Beta densities, and SR-PLR subsequently used the same construction for sequential recommendation (Ren & Leskovec, 2020; Yuan et al., 2023). This section gives the complete derivation for the parameterization used by TAPR and makes explicit the conditions under which the simplified parameter average is valid.

B.1 One-dimensional closure

Proposition 1 (closure of the weighted geometric pool). Let $q_n(x) = \mathrm{Beta}(x; \alpha_n, \beta_n)$ for $n = 1, \ldots, N$, where $\alpha_n, \beta_n > 0$. Let $a_n \geq 0$, and define

$$\tilde{\alpha} = 1 + \sum_{n=1}^{N} a_n(\alpha_n - 1), \qquad \tilde{\beta} = 1 + \sum_{n=1}^{N} a_n(\beta_n - 1). \tag{14}$$

If $\tilde{\alpha} > 0$ and $\tilde{\beta} > 0$, then the normalized weighted product

$$q_\wedge(x) = \frac{\prod_{n=1}^{N} q_n(x)^{a_n}}{\int_0^1 \prod_{n=1}^{N} q_n(t)^{a_n}\, dt} \tag{15}$$

is exactly $\mathrm{Beta}(x; \tilde{\alpha}, \tilde{\beta})$.

Proof. The density of the nth factor is

$$q_n(x) = \frac{x^{\alpha_n - 1}(1-x)^{\beta_n - 1}}{B(\alpha_n, \beta_n)}. \tag{16}$$

Raising each factor to a_n and collecting powers gives

$$\prod_{n=1}^{N} q_n(x)^{a_n} = \left[\prod_{n=1}^{N} B(\alpha_n, \beta_n)^{-a_n}\right] x^{\sum_n a_n(\alpha_n - 1)}(1-x)^{\sum_n a_n(\beta_n - 1)} \tag{17}$$

$$= C\, x^{\tilde{\alpha} - 1}(1-x)^{\tilde{\beta} - 1}, \tag{18}$$

where $C = \prod_n B(\alpha_n, \beta_n)^{-a_n}$ is independent of x. Because $\tilde{\alpha}, \tilde{\beta} > 0$, the kernel is integrable on $(0, 1)$ and

$$\int_0^1 C\, t^{\tilde{\alpha} - 1}(1-t)^{\tilde{\beta} - 1}\, dt = CB(\tilde{\alpha}, \tilde{\beta}). \tag{19}$$

Substituting this normalizer into Equation 15 cancels C and yields

$$q_\wedge(x) = \frac{x^{\tilde{\alpha} - 1}(1-x)^{\tilde{\beta} - 1}}{B(\tilde{\alpha}, \tilde{\beta})}, \tag{20}$$

which is the density of $\mathrm{Beta}(\tilde{\alpha}, \tilde{\beta})$. □

In TAPR, Equation 2 enforces $\alpha_n, \beta_n > 1$ and all conjunction weights are non-negative. Hence $\tilde{\alpha}, \tilde{\beta} \geq 1$, so the integrability condition is satisfied by construction.

Corollary 1 (simplex form). If the weights additionally satisfy $\sum_n a_n = 1$, then

$$\tilde{\alpha} = \sum_n a_n \alpha_n, \qquad \tilde{\beta} = \sum_n a_n \beta_n. \tag{21}$$

This follows because $1 - \sum_n a_n = 0$. Thus the parameter interpolation written by BetaE and SR-PLR is the simplex-specialized form of Proposition 1. For example, $0.6\ \mathrm{Beta}(3,2)$ and $0.4\ \mathrm{Beta}(2,5)$ combine to $\mathrm{Beta}(2.6, 3.2)$.

B.2 Factorized closure and basic properties

For $Q_n = \mathrm{Beta}^{\otimes D}(\boldsymbol{\alpha}_n, \boldsymbol{\beta}_n)$ and scalar simplex weights $a_n \geq 0$ with $\sum_n a_n = 1$, the same argument applies independently in every dimension:

$$\mathcal{C}_\beta(\{Q_n\}, \{a_n\}) = \mathrm{Beta}^{\otimes D}\left(\sum_n a_n \boldsymbol{\alpha}_n, \sum_n a_n \boldsymbol{\beta}_n\right). \tag{22}$$

The operator is closed and invariant to a simultaneous permutation of inputs and weights. It is also idempotent under simplex weights: combining repeated copies of the same Q returns Q. These are properties of the probabilistic operator; they do not imply that a learned attention function assigning the weights is itself permutation invariant.

B.3 Why we call the operator a probabilistic conjunction

Equation 15 is a normalized weighted geometric pool, equivalently a weighted product-of-experts. Its exact comparative statement is that, for any interior points x and y,

$$\log \frac{q_\wedge(x)}{q_\wedge(y)} = \sum_n a_n \log \frac{q_n(x)}{q_n(y)}. \tag{23}$$

Thus every factor with nonzero weight contributes additively to output log-density ratios. This identity does *not* guarantee that a high-density output region is a high-density region of every input: normalization and sharply conflicting factors can make their compromise the output mode. We use *probabilistic conjunction* following the learned soft-intersection construction of BetaE and SR-PLR (Ren & Leskovec, 2020; Yuan et al., 2023), not as a theorem of set intersection or Boolean truth. The output is not a calibrated probability that a proposition is true, and learned weights can attenuate an input. Under simplex weights, concentration parameters are interpolated rather than accumulated, so the operator should not be described as unweighted Bayesian evidence multiplication.

This distinction matters for our claim. TAPR contributes typed, position-bound atoms and temporally conditioned weights; it does not claim a new Beta closure theorem or exact Boolean semantics. The controlled interventions in Appendix E test whether the learned composition behaves as the intended ordered mechanism.

B.4 Factorized Beta KL divergence

For completeness, the score in Equation 12 follows from the one-dimensional identity

$$\begin{aligned} D_{\mathrm{KL}}(\mathrm{Beta}(a,b) \| \mathrm{Beta}(c,d)) = {} & \log \frac{B(c,d)}{B(a,b)} + (a-c)\psi(a) + (b-d)\psi(b) && (24) \\ & + (c+d-a-b)\psi(a+b). && (25) \end{aligned}$$

For independent factors, the log-density ratio is additive, and therefore

$$D_{\mathrm{KL}}(Q_c^C \| Q_u^R) = \sum_{j=1}^{D} D_{\mathrm{KL}}\left(\mathrm{Beta}(\alpha_{c,j}^C, \beta_{c,j}^C) \middle\| \mathrm{Beta}(\alpha_{u,j}^R, \beta_{u,j}^R)\right). \tag{26}$$

This exact dimension-level decomposition is useful for tracing scores, but it is not by itself a faithful atom-level attribution because the rule parameters depend jointly on all atoms through learned weights. Appendix F therefore evaluates explanations through keep, remove, and hierarchy interventions.

Table 4: Detailed post-processing statistics. Sequence quantiles are computed before truncation to the model history window. Each dataset contributes one validation and one test target per retained user.

Dataset	Users	Items	Total events	Train events	Min	P10	Median	P90	Max
MovieLens	2,113	10,109	855,598	851,372	20	44	273	969	3,410
H&M	52,880	20,769	4,137,970	4,032,210	50	52	67	117	741
Yelp	68,008	21,637	1,809,672	1,673,656	10	10	16	51	1,416

C Real-World Data and Preprocessing

All datasets use a chronological leave-two-out protocol: the last interaction of each retained user is the test target, the penultimate interaction is the validation target, and preceding interactions form the training prefix. Training uses shifted prefixes, whereas validation and test use a fixed history ending immediately before their target. Full-catalog evaluation ranks every catalog item after masking previously observed items according to the same protocol for all models.

C.1 MovieLens

We use the HetRec MovieLens 2K release and retain all 2,113 users. Movie attributes are typed as genre, director, actor, country, and coarse filming location. The final vocabularies contain 20 genre values, 4,024 directors, 4,000 actors, 71 countries, and 1,389 locations. An item has 12.19 retained values on average (minimum 2, maximum 63). The final test set contains 12 cold-item events involving 11 distinct items, corresponding to a 0.568% cold-event rate. Cold targets remain in the full catalog rather than being silently removed.

C.2 H&M

The H&M source contains transaction dates rather than within-day event times. We therefore collapse repeated customer–article events on the same calendar day and order the remaining events by date with a deterministic tie-breaker. This removes 3,212,929 duplicate rows (10.11% of raw transaction rows). We retain eleven structured article fields: product type, product group, graphical appearance, color group, perceived color value, perceived color master, department, index, index group, section, and garment group. Item frequency is recomputed from distinct users in the training prefixes only. Joint fixed-point filtering enforces a minimum retained sequence length of 50, at least 20 training users per item, and an attribute-signature uniqueness rate of at least 0.90 with maximum multiplicity four. The final observed signature-uniqueness fraction is 0.9001.

C.3 Yelp

We restrict Yelp businesses to the restaurant/food catalog and preserve repeated user–business interactions; timestamps induce only within-user ordering and are not interpreted as calibrated time gaps. Each item receives up to five leaf-category values and values from fourteen structured families: Alcohol, casual ambience, bike parking, lot parking, street parking, catering, good-for-kids, television, noise level, outdoor seating, delivery, price range, reservations, and WiFi. Together with category, this gives fifteen attribute types. Values and items are filtered using training prefixes only, with minimum sequence length 10, at least 20 distinct training users per item, minimum attribute-family coverage 0.80, and the same signature multiplicity constraint as H&M. The final signature-uniqueness fraction is 0.9283.

Leakage controls. Attribute vocabularies, value-support thresholds, item-frequency filters, and composition-difficulty counts are estimated only from training prefixes. Validation targets select model checkpoints and explanation thresholds. Test targets are used only after these choices are frozen. Dataset builders store source hashes, output checksums, filter trajectories, and split manifests; generated directories are immutable rather than overwritten in place.

Table 5: Primary TAPR implementation settings on MovieLens. Dataset-specific batch sizes may be reduced to fit memory, while the evaluation protocol and selection rule remain unchanged.

Setting	Value	Setting	Value
Latent/hidden dimension	64	Maximum history	50
Base heads / layers	2 / 2	Relation heads	4
Causal kernel	3	Dropout	0.2
Learning rate	10^{-3}	Weight decay	0
Train / eval batch	64 / 32	Item score chunk	2,048
Maximum epochs	50	Selection metric	validation NDCG@10
Rule count	1	Candidate score	directed Beta KL

D Implementation and Reproducibility Details

The primary MovieLens configuration uses $D = 64$, two encoder layers, two base attention heads, four heads in each relation branch, a causal convolution kernel of width 3, dropout 0.2, one user rule, and maximum history length 50. We optimize with learning rate 10^{-3}, batch size 64, evaluation batch size 32, no weight decay, and at most 50 epochs. Full-catalog scores are computed in chunks of 2,048 items. Checkpoints are selected by validation NDCG@10; test metrics are read only after selection.

Randomness and uncertainty. Standard real-data comparisons use independently trained seeds and report the mean and sample standard deviation. Low-data results use seeds 2024, 2025, and 2026. Synthetic evidence uses four independent configurations: $(2026, 2024)$, $(2026, 2025)$, $(2027, 2024)$, and $(2028, 2024)$, where each pair denotes the world seed and model seed. Faithfulness confidence intervals use a user-level bootstrap so that multiple counterfactual evaluations of the same user are not treated as independent.

Reproducibility artifacts. For each reported result, we retain the dataset builder, split manifest, resolved configuration, and evaluation entry point. Each dataset build and result summary is associated with a manifest and checksum, so preprocessing, model selection, and evaluation can be reproduced without using test targets to choose a configuration.

E Synthetic Dataset

The synthetic dataset asks whether a model recovers and uses a known semantic–temporal mechanism rather than an item-ID, attribute-bag, position, sequence-length, or null-class shortcut. Causal language in this section is restricted to this known synthetic structural causal model (SCM); it does not identify causal effects in observed human behavior.

E.1 World, catalog, and oracle

Each generated world independently permutes item IDs and visible attribute IDs. Four semantic roles are connected by eight directed edges,

$$(0,1),(1,0),(1,2),(2,1),(2,3),(3,2),(3,0),(0,3). \tag{27}$$

There are three source namespaces: *active sources*, which may trigger an overriding rule; *fallback sources*, which always define a valid default target; and role-less *distractors*. If payload cardinality is P, each target pool is indexed by a directed edge e and two payloads,

$$g = eP^2 + p_a P + p_b. \tag{28}$$

With $P = 8$, the benchmark contains 512 semantic target pools. Every pool has four exchangeable target replicas; source roles also have four replicas. Target items expose compositional first-role, second-role, left-payload, and right-payload attributes, so relation and entity OOD evaluation does not require an arbitrary item-ID lookup.

Every history contains one complete fallback pair and either one complete active pair or one incomplete active pair containing exactly one active event. There is exactly one valid fallback relation and at most one valid active relation. The oracle is

$$G(H) = \begin{cases} G_{\text{active}}(H), & \text{if a complete active rule is present,} \\ G_{\text{fallback}}(H), & \text{otherwise.} \end{cases} \tag{29}$$

An intervention pair is accepted only if the active and fallback targets differ. The generator stores the selected mechanism, relation, payloads, relevant positions, and oracle pool, permitting an independent oracle recomputation.

E.2 Support and split construction

For each train-supported target pool, active-selected and fallback-selected episodes are balanced. Thus both sides of insertion and removal are marginally supported in training rather than being separated by a null class. Within each payload pair, two relations are held out for validation- and test-composition OOD; two additional relations are held out world-wide for relation OOD. Entity OOD changes source and target replicas without changing the semantic target pool. Held-out pools are excluded from sampled training negatives.

The formal configuration uses maximum history length 12 with realized lengths 8, 10, and 12; 32,768 training episodes; 4,096 episodes in each observational evaluation split; 1,024 validation pairs and 4,096 test pairs per intervention family. Test family IDs are disjoint from training, validation, observational splits, and other test families. A generated world is the independent replicate; paired histories within a world are dependent observations.

E.3 Matched interventions

All pairs preserve history length, positions, the complete fallback pair, nuisance schedule, and candidate catalog.

Semantic-value intervention. One active event's payload changes while its role, nuisance attributes, replica, and position are fixed. Both resulting target pools are distinct and train-supported.

Temporal-order intervention. The two active items swap positions. The item multiset is exactly preserved, while the directed relation and oracle pool change; both directions are train-supported.

Antecedent removal and rule insertion. Removal replaces one active event with a role-less distractor that preserves payload, nuisance, replica, and position, switching the oracle from the active pool to the already present fallback pool. Insertion is the exact inverse matched pair. Both targets are non-null and train-supported.

Irrelevant change. A payload or nuisance field of one role-less filler changes while the oracle pool, selected mechanism, active/fallback pairs, and relevant positions remain fixed. This intervention supplies the placebo/invariance control.

Strict-binding proof sets use matched order-sensitive pairs to test TAPR's THEN relation. Same-type-binding proof sets analogously isolate relations among repeated observations of the same attribute type. Each proof set is evaluated with the same frozen candidates and paired oracle alternatives.

E.4 Metrics

Let H_n and H_n^{do} denote a factual/intervened pair, and let G_n and G_n^{do} be their oracle pools. Let $S(g \mid H)$ be the pool score, computed as the log-mean-exp of scores of exchangeable replicas in pool g.

For changing pairs, paired interventional joint correctness is

$$\text{PIJC} = \frac{1}{N} \sum_{n=1}^{N} \mathbf{1}[\hat{G}_n = G_n] \mathbf{1}[\hat{G}_n^{do} = G_n^{do}]. \tag{30}$$

Table 6: Cross-model structural capability on the synthetic dataset (mean ± sample standard deviation across four independent runs). BV is bidirectional validity/PEDC.

Proof set	Model	CF-NDCG@10	BV / PEDC	Oracle margin
Strict	TAPR	0.3388 ± 0.0628	0.2255 ± 0.0514	2.7591 ± 0.2443
Strict	no-THEN	0.3231 ± 0.0378	0.2881 ± 0.1781	3.2978 ± 2.2802
Strict	CNN-only	0.3369 ± 0.0634	0.2117 ± 0.0455	3.0258 ± 0.5682
Strict	strict-no-rule	0.0136 ± 0.0017	0.0000 ± 0.0000	0.0000 ± 0.0000
Same-type	TAPR	0.3367 ± 0.0628	0.2880 ± 0.0921	1.0107 ± 0.5361
Same-type	no-THEN	0.3184 ± 0.0323	0.3647 ± 0.1796	1.3788 ± 0.6572
Same-type	CNN-only	0.3333 ± 0.0552	0.3523 ± 0.1151	1.3619 ± 0.7191
Same-type	strict-no-rule	0.0128 ± 0.0016	0.0000 ± 0.0000	0.0000 ± 0.0000

Paired effect-direction correctness, called *bidirectional validity* in the main text, is

$$\mathrm{PEDC} = \frac{1}{N}\sum_{n=1}^{N} \mathbf{1}[S(G_n \mid H_n) > S(G_n^{do} \mid H_n)]\mathbf{1}[S(G_n^{do} \mid H_n^{do}) > S(G_n \mid H_n^{do})]. \tag{31}$$

Define the two directional margins

$$m_n^F = S(G_n \mid H_n) - S(G_n^{do} \mid H_n), \tag{32}$$

$$m_n^I = S(G_n^{do} \mid H_n^{do}) - S(G_n \mid H_n^{do}). \tag{33}$$

The bidirectional oracle margin is $N^{-1}\sum_n (m_n^F + m_n^I)/2$. CF-NDCG@10 is the ordinary NDCG@10 of G_n^{do} in the intervened pool ranking. We additionally report factual/intervened pool accuracy, their absolute gap, and balanced-side accuracy

$$\mathrm{BSA} = \tfrac{1}{2}(\mathrm{Acc}_F + \mathrm{Acc}_I). \tag{34}$$

For invariant pairs, joint interventional invariance correctness is

$$\mathrm{JIIC} = \frac{1}{N_0}\sum_{n:G_n = G_n^{do}} \mathbf{1}[\hat{G}_n = G_n]\mathbf{1}[\hat{G}_n^{do} = G_n]\mathbf{1}[\hat{G}_n = \hat{G}_n^{do}]. \tag{35}$$

Prediction-change rate is diagnostic only because a changed prediction need not be the correct causal response.

Checkpoint selection. The implemented protocol selects the model state using only IID observational validation-pool NDCG@10. Intervention test files are not read during selection. This is the criterion recorded by the released training code and result manifests.

E.5 Structural capability results

Table 6 reports all four independent runs aggregated by proof set. The strict-no-rule control collapses on both proof sets. The other three models respond to the interventions, but their confidence intervals overlap and their relative ordering varies; this is evidence for the rule path, not a claim of universal superiority for TAPR's THEN branch.

E.6 Frozen-model mechanism attribution

For a frozen TAPR instance, let M_{full} be the bidirectional oracle margin under full inference and $M_{-\chi}$ the margin after disabling relation branch χ. The branch effect is $\Delta_\chi = M_{\mathrm{full}} - M_{-\chi}$. The placebo applies the same branch intervention on the non-matching proof set. The specificity contrast is the corresponding difference in differences.

The oracle-margin effect and specificity contrast for TAPR's THEN branch are positive in all four runs. PEDC specificity is also positive in all four runs. PIJC is much smaller because it requires both pool-level top-1 decisions to be exactly correct; it is positive in three of four runs. The same-type intervention does not show an analogous margin effect, which is a useful negative result rather than evidence for a generic benefit of disabling any branch.

Table 7: Frozen-model mechanism attribution across four runs. "Effect" is full minus branch-off; "placebo" uses the non-matching proof set; "specificity" is their difference.

Target branch	Metric	Full	Effect	Placebo	Specificity
TAPR: THEN	Oracle margin	2.7591	0.5930	0.0002	0.5928
TAPR: THEN	PEDC	0.2255	0.0483	−0.0007	0.0491
TAPR: THEN	PIJC	0.0220	0.0044	0.0023	0.0021
Same-type	Oracle margin	1.0107	−0.0062	−0.0167	0.0105
Same-type	PEDC	0.2880	0.0005	−0.0028	0.0033
Same-type	PIJC	0.0521	0.0022	0.0004	0.0018

E.7 GENERATOR VALIDATORS AND SHORTCUT SENTINELS

A formal world is rejected unless: (i) an independent oracle reproduces every stored target; (ii) every history has exactly one fallback pair and at most one active rule; (iii) target-by-mechanism training cells are balanced; (iv) both sides of every changing pair are distinct, non-null, and train-supported; (v) intervention edit counts and temporal multiset preservation are exact; (vi) irrelevant changes preserve the oracle and mechanism; (vii) family IDs are disjoint across splits; (viii) held-out pools never appear as positives or sampled negatives in training; and (ix) seeded regeneration is byte-identical. Item-ID, visible-attribute bag, absolute-position, gap, length, and mechanism-only lookup controls are recorded as shortcut sentinels. Oracle scoring must obtain 1.0 on factual accuracy, PIJC, PEDC, and JIIC, while a seeded random scorer must remain near chance.

F FAITHFULNESS PROTOCOL AND CASE STUDIES

We evaluate the MovieLens TAPR checkpoint trained with seed 2024 on all 2,113 test users under full-catalog ranking. Event and attribute thresholds are selected on validation data and frozen at $(e, a) = (0.25, 0.25)$. The test set is not used to select thresholds, users, or controls. Let $s_o(c)$, $s_{-R}(c)$, and $s_R(c)$ denote candidate scores under the original history, after removing selected rule evidence, and after keeping only the selected rule evidence plus the always-retained ID baseline. Let $\hat{c}_o$ and $\hat{c}_R$ be the respective top-1 candidates. If y is the observed target, define its margin as its score minus the maximum score among non-target candidates.

F.1 METRICS AND MATCHED CONTROLS

Keep Top-1 Preservation. This metric asks whether the retained explanation reconstructs the model's decision:

$$\text{KeepTop1} = \frac{1}{|\mathcal{U}|} \sum_{u \in \mathcal{U}} \mathbf{1}[\hat{c}_{o,u} = \hat{c}_{R,u}]. \tag{36}$$

Faithful Coverage. For user u, necessity is the change in target margin after removing the selected evidence,

$$N_u = \text{margin}_{o,u}(y_u) - \text{margin}_{-R,u}(y_u). \tag{37}$$

The rule is counted as locally faithful when it is extractable, $N_u > 0$, and keeping it preserves the original top-1 decision. Faithful Coverage is the fraction of users satisfying all three conditions.

Hierarchy Completeness. For selected event s, let p_s be the parent-event deletion effect, $d_{s,g}$ the deletion effect of selected child component g, and $D_s = \sum_g d_{s,g}$. Event-level completeness is

$$C_s = \left[1 - \frac{|p_s - D_s|}{|p_s| + \sum_g |d_{s,g}| + \epsilon} \right]_{[0,1]}. \tag{38}$$

We average C_s across selected events and then across users. The hierarchy control cyclically reassigns every selected child set to a different selected parent event with no fixed point. It preserves the number of events and components while breaking the learned parent–child association.

Table 8: MovieLens full-catalog faithfulness results. Intervals are 95% user-level bootstrap confidence intervals. Hierarchy metrics use the 1,950 users with at least two selected events; all other rows use 2,113 users.

Metric	Extracted rule	Matched control	Paired lift
Keep Top-1 Preservation	0.3895 [0.3687, 0.4103]	0.3307 [0.3150, 0.3478]	0.0588 [0.0457, 0.0724]
Hierarchy Completeness	0.8235 [0.8197, 0.8271]	0.4703 [0.4626, 0.4777]	0.3533 [0.3442, 0.3623]
Faithful Coverage	0.2144 [0.1978, 0.2333]	0.1841 [0.1718, 0.1968]	0.0303 [0.0172, 0.0445]

Table 9: Faithfulness lifts on the additional full-catalog panels. Values are extracted-rule performance minus the relevant matched control, with 95% user-level bootstrap intervals. H&M and Yelp each use a pre-specified panel of 10,000 test users; hierarchy rows use 9,937 and 9,446 eligible users, respectively.

Dataset	Metric	Paired lift
H&M	Keep Top-1 Preservation	0.0189 [0.0142, 0.0235]
H&M	Hierarchy Completeness	0.2776 [0.2744, 0.2808]
H&M	Faithful Coverage	0.0122 [0.0083, 0.0162]
Yelp	Keep Top-1 Preservation	0.0394 [0.0343, 0.0443]
Yelp	Hierarchy Completeness	0.3045 [0.3007, 0.3084]
Yelp	Faithful Coverage	0.0274 [0.0223, 0.0321]

Matched random evidence. For each selected event, the random control samples the same number of eligible attribute components from that event. Thus it preserves event positions, temporal availability, rule size, and per-event component counts. Four random replicates are averaged with seed 2026. This is a stringent control: it overlaps the extracted rule by 79.95% on average.

F.2 Full quantitative results

The three comparisons answer different questions. Keep Top-1 tests sufficiency for reconstructing the original decision; Faithful Coverage additionally requires a positive deletion effect; Hierarchy Completeness tests whether atom-to-event ownership carries information beyond the selected atoms themselves. The controls are matched in complexity, so the lifts are not attributable merely to retaining more evidence.

F.3 Pre-specified case panel

Cases are selected from four pre-specified strata with selection seed 2026; selection does not maximize test effect magnitude. Table 10 records the chosen instances. The complete release includes each history, typed attributes, selected events/components, atom weights, candidate scores, KL terms, and original/keep/remove rank trajectories.

The typical and complex cases show successful decision reconstruction with increasingly distributed evidence. The threshold-boundary case makes visible sensitivity to the frozen extraction threshold. The failure case is retained to show that a high hierarchy score does not guarantee top-1 sufficiency; selected evidence can be internally well organized while omitting decision-relevant context.

G Low-Data Composition-Difficulty Analysis

G.1 Nested subsets and fixed difficulty slices

Training subsets retain $\{5, 10, 25, 50, 100\}\%$ of the MovieLens training users using seeded nested sampling; validation and test users, targets, histories, candidate catalog, and difficulty labels remain fixed. Every model is retrained separately for each fraction and seed. Results below aggregate seeds 2024, 2025, and 2026 with the sample standard deviation.

Table 10: Pre-specified MovieLens case panel. "Components" counts selected event–attribute components.

Case type	User	Faithful	Events	Components	Original top-1
Typical success	68947	yes	4	16	*The Matrix*
Complex reasoning	48076	yes	5	20	*WALL-E*
Threshold boundary	70878	no	4	15	*Pulp Fiction*
Faithfulness failure	52328	no	3	11	*The Matrix*

Table 11: Mean rank percentile by training fraction and difficulty (lower is better; mean ± sample standard deviation across three seeds).

Train	Model	Easy	Middle	Hard
5%	TAPR	$\mathbf{0.1043} \pm 0.0108$	$\mathbf{0.1373} \pm 0.0102$	$\mathbf{0.2882} \pm 0.0137$
	FDSA	0.1134 ± 0.0034	0.1588 ± 0.0196	0.3399 ± 0.0286
	NCR	0.1152 ± 0.0048	0.1623 ± 0.0100	0.3375 ± 0.0174
	DIF-SR	0.1083 ± 0.0120	0.1538 ± 0.0162	0.3358 ± 0.0083
10%	TAPR	$\mathbf{0.0831} \pm 0.0075$	$\mathbf{0.1132} \pm 0.0067$	$\mathbf{0.2693} \pm 0.0046$
	FDSA	0.0872 ± 0.0056	0.1249 ± 0.0128	0.3006 ± 0.0093
	NCR	0.0875 ± 0.0079	0.1232 ± 0.0088	0.3017 ± 0.0051
	DIF-SR	0.0847 ± 0.0044	0.1202 ± 0.0046	0.2931 ± 0.0112
25%	TAPR	0.0685 ± 0.0034	$\mathbf{0.0907} \pm 0.0020$	$\mathbf{0.2188} \pm 0.0068$
	FDSA	$\mathbf{0.0671} \pm 0.0052$	0.0953 ± 0.0044	0.2495 ± 0.0076
	NCR	0.0677 ± 0.0057	0.0984 ± 0.0042	0.2561 ± 0.0057
	DIF-SR	0.0692 ± 0.0035	0.0951 ± 0.0016	0.2538 ± 0.0029
50%	TAPR	0.0659 ± 0.0042	$\mathbf{0.0826} \pm 0.0012$	$\mathbf{0.1866} \pm 0.0060$
	FDSA	0.0645 ± 0.0061	0.0885 ± 0.0096	0.2303 ± 0.0091
	NCR	$\mathbf{0.0625} \pm 0.0013$	0.0862 ± 0.0004	0.2333 ± 0.0035
	DIF-SR	0.0696 ± 0.0017	0.0959 ± 0.0008	0.2386 ± 0.0079
100%	TAPR	$\mathbf{0.0541} \pm 0.0032$	$\mathbf{0.0725} \pm 0.0029$	$\mathbf{0.1674} \pm 0.0042$
	FDSA	0.0668 ± 0.0034	0.0903 ± 0.0016	0.2187 ± 0.0087
	NCR	0.0605 ± 0.0017	0.0851 ± 0.0044	0.2230 ± 0.0062
	DIF-SR	0.0632 ± 0.0023	0.0864 ± 0.0009	0.2083 ± 0.0030

Difficulty uses only training-prefix attribute statistics. For a target containing cross-type value pair (a, b), let c_a, c_b, and c_{ab} be their training-event counts and let N be the number of training interactions. Pair surprise is

$$\text{surprise}(a, b) = \log \frac{(c_a + \alpha)(c_b + \alpha)}{(c_{ab} + \alpha)(N + \alpha)}, \qquad \alpha = 1. \tag{39}$$

Pairs are eligible only when both marginals have support at least $\max(20, 0.001N) = 852$. An item's score is the mean of the largest 25% of its eligible pair surprises. Easy, middle, and hard are defined by validation-target tertiles, fixed at 0.05196 and 0.34493. Test targets are never used to estimate supports or thresholds. This yields 628 easy, 693 middle, 785 hard, and 7 ineligible test users.

G.2 Full mean-rank-percentile results

G.3 Full HR@100 results

The rank-percentile table shows the most stable interaction: TAPR ranks hard targets ahead of every comparator at every fraction. HR@100 identifies the boundary more sharply. At 5% and 10%, the hard-target rank improvement has not yet reached the top 100 consistently. From 25% onward, TAPR is best on hard-target HR@100, with all three seeds agreeing in direction.

H Sensitivity Analyses

H.1 Three-seed one-factor sensitivity

We evaluate three implementation axes on MovieLens under full-catalog ranking: the residual relation-logit coefficient λ_{rel}, latent/hidden dimension D, and maximum history length L. All runs use seeds 2024, 2025, and 2026, the same chronological split, and the same training and evaluation

Table 12: HR@100 by training fraction and difficulty (higher is better; mean ± sample standard deviation across three seeds). Bold marks the best of the four compared models within each cell.

Train	Model	Easy	Middle	Hard
5%	TAPR	$\mathbf{0.3169} \pm 0.0245$	$\mathbf{0.2429} \pm 0.0079$	0.0437 ± 0.0195
	FDSA	0.2861 ± 0.0330	0.2145 ± 0.0334	$\mathbf{0.0467} \pm 0.0116$
	NCR	0.2861 ± 0.0223	0.1881 ± 0.0108	0.0348 ± 0.0085
	DIF-SR	0.3073 ± 0.0130	0.2174 ± 0.0217	0.0463 ± 0.0141
10%	TAPR	$\mathbf{0.3970} \pm 0.0106$	0.2482 ± 0.0175	0.0238 ± 0.0051
	FDSA	0.3429 ± 0.0212	$\mathbf{0.2487} \pm 0.0216$	0.0348 ± 0.0045
	NCR	0.3508 ± 0.0169	0.2285 ± 0.0124	0.0251 ± 0.0099
	DIF-SR	0.3583 ± 0.0111	0.2400 ± 0.0068	$\mathbf{0.0429} \pm 0.0027$
25%	TAPR	0.3917 ± 0.0240	$\mathbf{0.3040} \pm 0.0167$	$\mathbf{0.0688} \pm 0.0193$
	FDSA	0.4087 ± 0.0093	0.2703 ± 0.0182	0.0335 ± 0.0096
	NCR	$\mathbf{0.4130} \pm 0.0080$	0.2535 ± 0.0140	0.0289 ± 0.0045
	DIF-SR	0.4023 ± 0.0148	0.2621 ± 0.0188	0.0361 ± 0.0072
50%	TAPR	$\mathbf{0.4209} \pm 0.0285$	$\mathbf{0.3410} \pm 0.0022$	$\mathbf{0.1291} \pm 0.0209$
	FDSA	0.4018 ± 0.0398	0.2684 ± 0.0188	0.0437 ± 0.0051
	NCR	0.4172 ± 0.0146	0.2684 ± 0.0109	0.0297 ± 0.0129
	DIF-SR	0.3636 ± 0.0144	0.2559 ± 0.0030	0.0764 ± 0.0084
100%	TAPR	$\mathbf{0.4623} \pm 0.0155$	$\mathbf{0.3583} \pm 0.0150$	$\mathbf{0.1503} \pm 0.0157$
	FDSA	0.3646 ± 0.0256	0.2915 ± 0.0177	0.0917 ± 0.0297
	NCR	0.4236 ± 0.0256	0.2689 ± 0.0079	0.0293 ± 0.0056
	DIF-SR	0.4034 ± 0.0130	0.3006 ± 0.0096	0.1244 ± 0.0116

Table 13: Three-seed one-factor sensitivity on MovieLens under full-catalog evaluation (mean ± sample standard deviation). Each block varies one factor around the shared anchor $(\lambda_{\text{rel}}, D, L) = (1, 64, 50)$. Validation NDCG@10 selects checkpoints; test metrics are reported after selection.

Axis	Value	Val. NDCG@10	Test NDCG@10	Test HR@10	Test MRR
λ_{rel}	0.0	0.0477 ± 0.0013	0.0473 ± 0.0011	0.0844 ± 0.0024	0.0447 ± 0.0013
	0.5	0.0472 ± 0.0019	0.0465 ± 0.0014	0.0817 ± 0.0017	0.0445 ± 0.0013
	1.0	0.0473 ± 0.0007	0.0463 ± 0.0011	0.0828 ± 0.0021	0.0440 ± 0.0020
	2.0	0.0479 ± 0.0022	0.0437 ± 0.0021	0.0773 ± 0.0014	0.0424 ± 0.0027
D	32	0.0358 ± 0.0026	0.0323 ± 0.0028	0.0604 ± 0.0022	0.0319 ± 0.0034
	64	0.0473 ± 0.0007	0.0463 ± 0.0011	0.0828 ± 0.0021	0.0440 ± 0.0020
L	5	0.0348 ± 0.0038	0.0302 ± 0.0060	0.0568 ± 0.0036	0.0290 ± 0.0070
	20	0.0442 ± 0.0004	0.0381 ± 0.0024	0.0715 ± 0.0019	0.0362 ± 0.0025
	50	0.0473 ± 0.0007	0.0463 ± 0.0011	0.0828 ± 0.0021	0.0440 ± 0.0020

protocol. The shared anchor is $(\lambda_{\text{rel}}, D, L) = (1, 64, 50)$; each row changes one axis while retaining the other two anchor values. Each run selects its checkpoint by validation NDCG@10, after which the held-out test metrics are recorded. Thus, Table 13 is a sensitivity analysis rather than a test-set model-selection procedure.

In the residual relation encoder, λ_{rel} rescales the normalized relation logit before it is added to the causal-CNN logit. The coefficient therefore controls the magnitude of the relational residual; it is not a loss weight. In particular, $\lambda_{\text{rel}} = 0$ suppresses this residual while retaining the causal-CNN and rule-matching components.

For $\lambda_{\text{rel}} \in \{0, 0.5, 1,$ the reported test metrics are close under this protocol; doubling the residual coefficient to 2.0 lowers all three reported test metrics. This pattern indicates that the model does not require an amplified relational residual on this dataset. It does not test the necessity of the full rule representation, because setting $\lambda_{\text{rel}} = 0$ leaves the remaining model components unchanged.

The $D = 64$ configuration outperforms $D = 32$ in this sweep. Since D is implemented through the shared hidden width, this comparison jointly changes representational dimension and model capacity; it should not be interpreted as an isolated causal effect of Beta dimensionality. Finally, performance rises from $L = 5$ to $L = 50$ across the reported metrics, supporting the use of the 50-event context in the reported MovieLens configuration. These one-factor results describe the local configuration behavior and do not establish dataset-independent optima.