CSRec: Rethinking Sequential Recommendation from A Causal Perspective

Presenter: Dr. Qiaozhu Mei

Xiaoyu Liu*, Jiaxin Yuan*, Yuhang Zhou, Jingling Li, Furong Huang, and Wei Ai
University of Maryland, College Park,
{xliu1231, jyuan98, tonyzhou, jingling, furongh, aiwei} @umd.edu
*equal contribution

Sequential Recsys: Predicts User Preference For Next Item

User interests & previous positive decisions

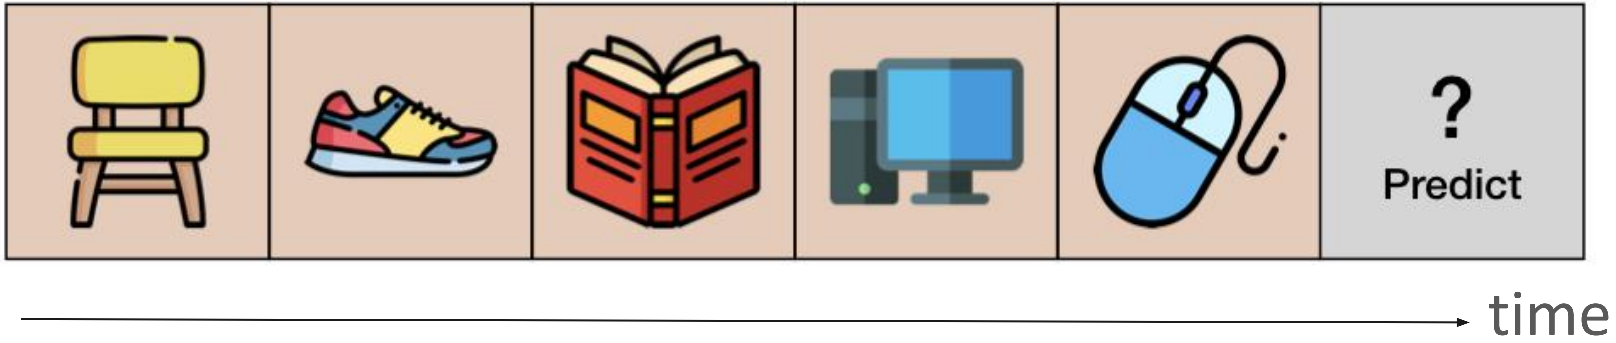

Traditional Sequential Recsys - A Visualization

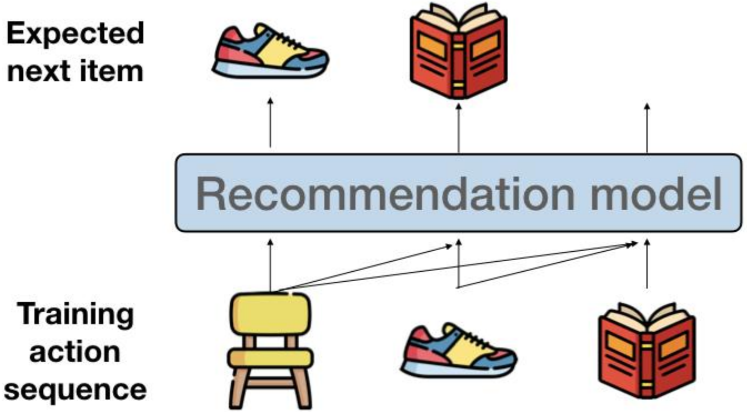

H: historical data

P: user preferences at time t

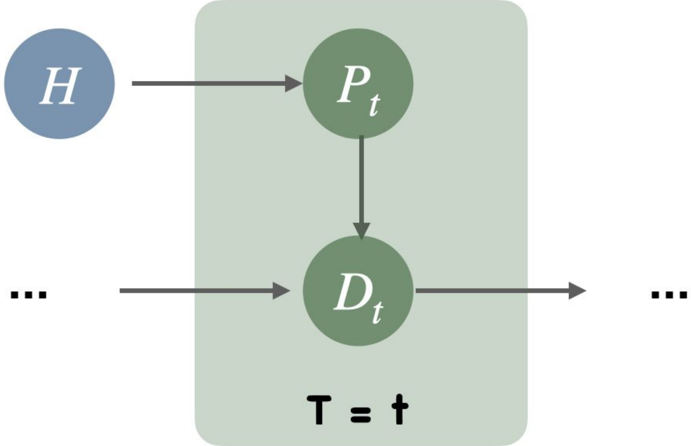

D: user decision at time t

Current preference is derived from previous preference.

Decision is equivalent to preference.

Modeling Preference Is Not Enough

Echo Chamber → **low recommendation quality**

Existing recommender systems assume a user's decision directly reflects their preference. In sequential settings, this leads to repeatedly suggesting similar high-scoring items—even when the user has already purchased one—reinforcing a narrow view of their interests.

Biased Strategies → **missing negative samples**

Users may reject recommendations during the sequence, resulting in both positive and negative decisions. However, existing frameworks often overlook these negative signals because they are missing from observational data, leading to biased learning of recommendation strategies.

Changed Decision → **Interaction not considered**

Recommender systems can influence user decisions, shaping what users choose over time. However, this influence is not accounted for in existing frameworks.

Recommendation: More Than Just User Preference

system influence is missing

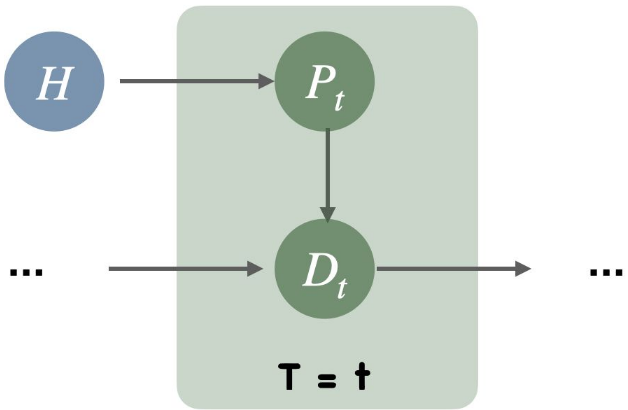

VS

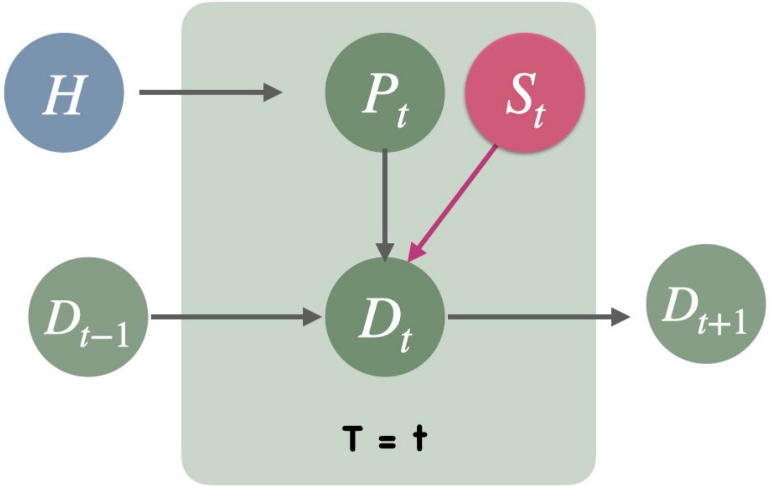

Recsys differentiates user decision from user preference

Recommendation: More Than Just User Preference

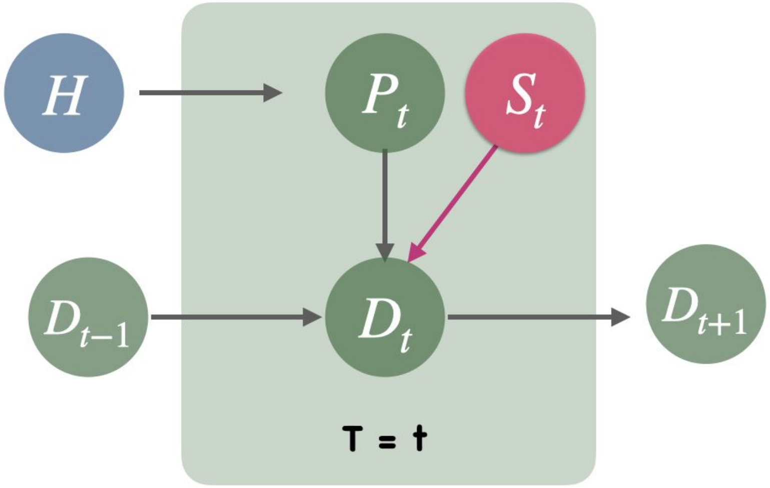

- How recsys,previous decision and preference **collaboratively** affect the user's decision.
- How **recsys itself** affects user's decision.

Requires **maintaining** the causal relationship and also **isolating** the causal relationship

Understanding Do Notation - Intervention

We aim to understand how the **changes in recommended items S (intervention)** affect the user's decision D, with user preference P

$$P(D|\underbrace{do(S = s)}_{\text{Intervention}}, P)$$

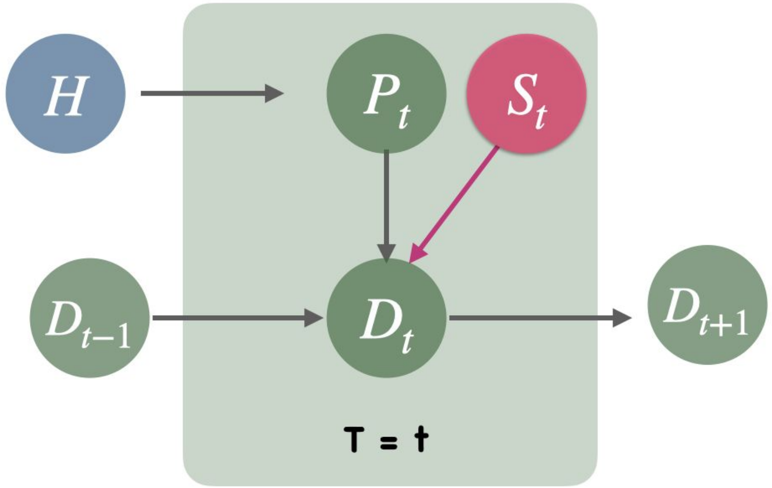

Maintaining Causal Relationship: Theoretical Insights

Maximize the **probability of positive decision** under recommendation.

$$\max_{\mathrm{do}(S_t)\in\mathcal{S}} \mathbb{P}(D_t|\mathrm{do}(S_t, ..., S_1), P_t)$$

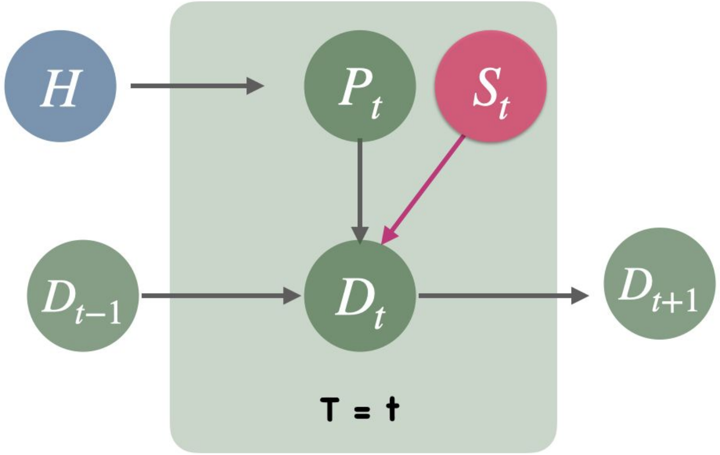

Current decision is a weighted sum of **previous decisions** and **recommendations**

$$\text{s.t.} \quad \mathbb{P}(D_t|\mathrm{do}(S_t, ..., S_1), P_t) = \sum_{D_{t-1}} \mathbb{P}(D_t|S_t, D_{t-1}, P_t)\cdot\mathbb{P}(D_{t-1}|\mathrm{do}(S_{t-1}, ..., S_1), P_t)$$

Current Decision Being A Weighted Sum of Previous Decisions

decision towards the second recommendation

$$
\begin{aligned}
&\mathbb{P}(D_2|\mathrm{do}(S_2 = s_2, S_1 = s_1), P_2) \\
&= \sum_{D_1 \in \{0,1\}} \mathbb{P}(D_2|S_2 = s_2, D_1, , P_1)\mathbb{P}(D_1|\mathrm{do}(S_1 = s_1), P_1)
\end{aligned}
$$

prob of successful first recommendation

$$
\begin{aligned}
&= \mathbb{P}(D_2|D_1 = 1, S_2 = s_2, P_2)\boxed{\mathbb{P}(D_1 = 1|\mathrm{do}(S_1 = s_1), P_1)} + \\
&\quad \mathbb{P}(D_2|D_1 = 0, S_2 = s_2, P_2)\boxed{\mathbb{P}(D_1 = 0|\mathrm{do}(S_1 = s_1), P_1)}.
\end{aligned}
$$

prob of unsuccessful first recommendation

users' natural decision s based on previous decision and temporal preference

Isolating Causal Relationship: How Recsys Makes A Difference

Compliance → **Treatment effect** → For a given item, the recommender system may make a user more likely—or less likely—to purchase it.

A new perspective in design → Compliance-based designs, i.e., detect echo chamber, maximizing the sales of an item

System Influence as Treatment Effect - Evaluation

Compliance rate (CR)

$$\frac{1}{|C|}\sum_{i\in C}[\mathbb{P}(D^i = 1|do(S = s), P^i) - \mathbb{P}(D^i = 0|S, P^i)]$$

Non-compliance rate (NCR)

$$\frac{1}{|C|}\sum_{i\in C}[\mathbb{P}(D^i = 0|do(S = s), P^i) - \mathbb{P}(D^i = 1|S, P^i)]$$

	Organic Decision	**Intervened Decision**
Compliance	Negative	Positive
Non-compliance	Positive	Negative

→ **higher compliance: higher user satisfaction**

→ **Lower compliance: decline in user satisfaction**

C represents the set of items that the user's decision changed after intervention

System Influence as Treatment Effect - New Design Directions

The system can be trained to maximize the **probability of a user accepting a recommendation (compliance rate)**

this item is accepted by the user

$$\text{argmax}_S \sum_{c \in C} [\mathbb{P}(D^c = 1 | do(S = c), P^i) - \mathbb{P}(D^c = 0 | S = c, P^i)]$$

this item is declined by the user

for each item

Results: Music Streaming Sessions Dataset

MSSD: first introduced at the WSDM Cup 2019 Spotify Sequential Skip Prediction Challenge

60 million listening sessions, 3.7 million tracks with acoustic features

		General SOTA Sequential Recommenders							WSDM winning model	
Data	Metric	GRU4Rec	Caser	SASRec	BERT4Rec	FMLPRec	BSARec	CSRec	Few-Shot	CSRec+Few-Shot
X_{obs}	↑ HR@10	0.1311	0.1012	0.1082	0.0407	0.1952	0.1901	**0.2007**	0.2235	**0.2364**
	↑ NDCG@10	0.0312	0.0457	0.0981	0.0771	0.1031	0.1487	**0.2117**	**0.1674**	0.1572
X_{intv}	↑ AHR@0.1	0.0047	0.1019	0.1293	0.1107	0.0968	0.1263	**0.2084**	0.3115	**0.5279**
	↑ AHR@0.2	0.1117	0.1461	0.1926	0.1468	0.1412	0.1692	**0.3509**	0.4127	**0.6764**
	↑ AHR@0.5	0.2403	0.2467	0.2956	0.3401	0.3009	0.4323	**0.5371**	0.5061	**0.7418**
	↑ AUC	0.2774	0.2998	0.3762	0.4239	0.4003	0.4572	**0.5651**	0.5424	**0.7851**

Results: Synthetic Book Dataset

Synthetic dataset using GPT-4 as agents to mimic 10000 users', each with unique preference

Amazon book dataset as a list of potential items

Repeat the generation process for a set of different temperatures

Data	Metric	GRU4Rec	Caser	SASRec	BERT4Rec	FMLPRec	BSARec	CSRec
X_{obs}	↑ HR@10	**0.3980**	0.3857	0.3865	0.3951	0.3893	0.3888	0.3933
	↑ NDCG@10	**0.1797**	0.1778	0.1753	0.1773	0.1777	0.1741	0.1716
X_{intv}	↑ AHR@0.1	0.619 ± 0.0206	0.626 ± 0.0154	0.622 ± 0.0157	0.624 ± 0.0162	0.625 ± 0.0153	0.621 ± 0.0167	**0.635** ± 0.0142
	↑ AHR@0.2	0.583 ± 0.02	0.606 ± 0.0163	0.614 ± 0.0153	0.591 ± 0.0197	0.601 ± 0.0176	0.0587 ± 0.0158	**0.641** ± 0.0146
	↑ AHR@0.5	0.831 ± 0.0134	0.834 ± 0.0157	0.833 ± 0.0098	0.823 ± 0.0133	0.834 ± 0.0108	0.826 ± 0.0164	**0.884** ± 0.0100
	↑ AUC	0.875 ± 0.0157	0.910 ± 0.0140	0.921 ± 0.0117	0.856 ± 0.0168	0.899 ± 0.0144	0.872 ± 0.0140	**0.930** ± 0.0086

* AHR: averaged hit rate

CSRec Enables Using Compliance Rate to Evaluate User Satisfaction

In another synthetic dataset, we simulated an echo chamber scenario by asking the system to **repeatedly recommend highly similar items** that user prefers.

CR provides additional signals beyond traditional AUC metrics.

Eval Metric	CSRec+BSARec	CSRec+SASRec
↑ AUC	0.621	0.617
↑ CR	0.081	0.149
↓ NCR	0.252	0.107

AUC shows over-positive signals while CR indicates that the user may not be satisfied with the recommendation quality.

Questions

Thank you!