Tree of Preferences for Diversified Recommendation

Hanyang Yuan[1†], *Ning Tang*[2†], *Tongya Zheng*[3*], *Jiarong Xu*[2*], *Xintong Hu*[1]
Renhong Huang[1], *Shunyu Liu*[4], *Jiacong Hu*[1], *Jiawei Chen*[1], *Mingli Song*[1]

1 Zhejiang University, 2 Fudan University,
3 Hangzhou City University, 4 Nanyang Technological University
†Equal contribution ∗ Corresponding authors

NeurIPS 25: Tree of Preferences for Diversified Recommendation

The key to *diversity*

- Capture user **latent preferences under data bias**

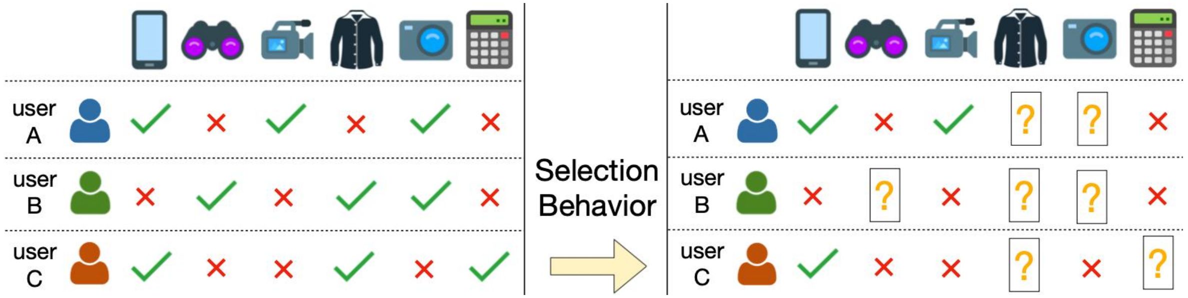

Selection bias

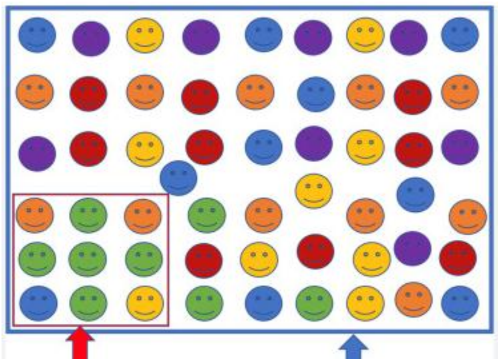

Exposure bias

Limitation of exiting work

- Via deep learning: rely on **observed data** only
- Via LLM: focus on **coarse category level** only

NeurIPS 25: Tree of Preferences for Diversified Recommendation

Our motivation

- From **user behaviors** to **latent** preferences
- With LLM’s expertise

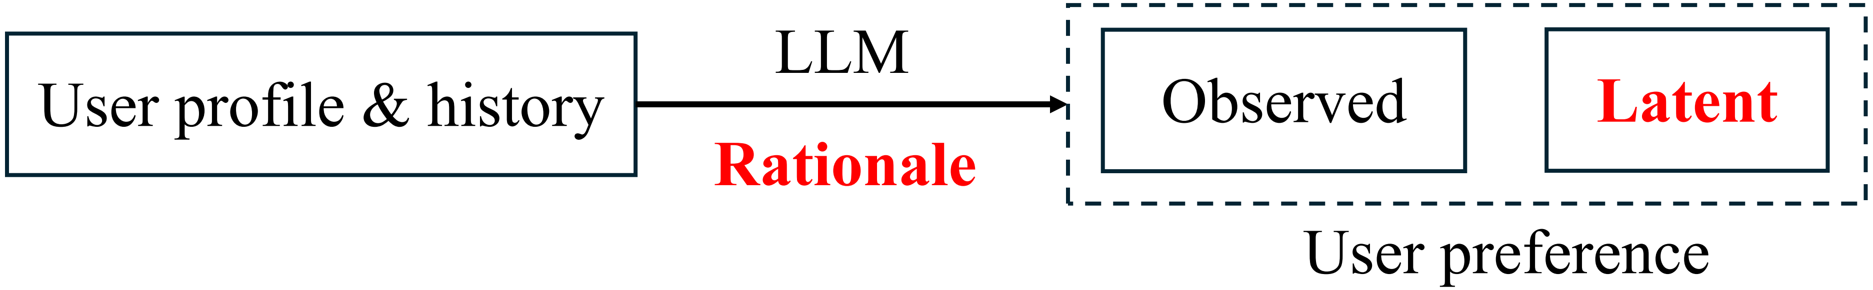

- Example: travel, transportation, local cuisines, and photography

NeurIPS 25: Tree of Preferences for Diversified Recommendation

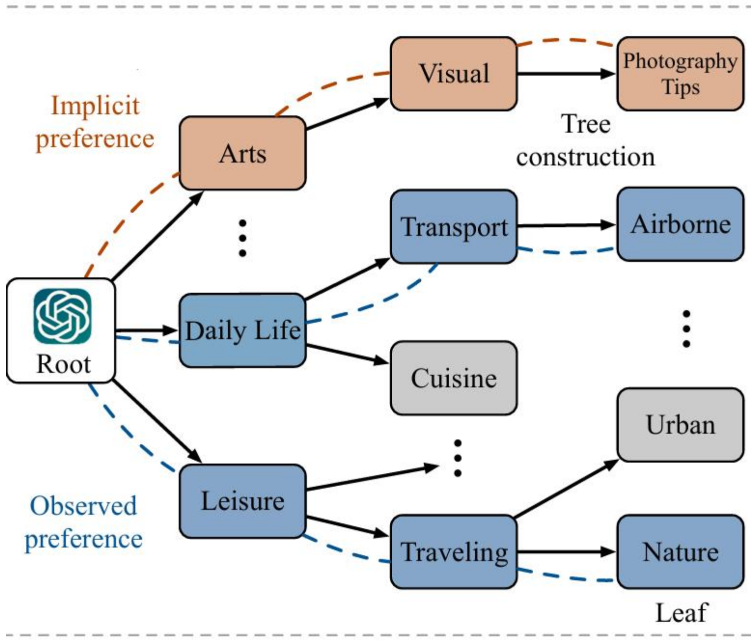

Notaion: $\mathsf{Prompt}_{\mathsf{ToP}}$.
Usage: ToP construction.
Content: You are an expert system designed to classify and organize user preferences based on the following information of item examples. <u>Your task is to generate a multi-level tree structure for user preferences, called Tree of Preference (ToP), progressively refining them from broad to specific preferences.</u> Each node represents a type of user preference, and each edge signifies a finer preference division from its parent node.
Constraints:
- Each node should be divided into 3-5 finer preferences (branches), except for leaf nodes.
- Use diverse preference-dividing criteria at each level.
- Nodes must represent clear, actionable preferences.
...
Items samples: {sampled_items_information}.
Return ToP in the following format: {ToP_format}.

Uncovering latent preferences

- Step 1: Constructing **Tree of Preferences** over items

$$\mathcal{T}(\mathcal{V},\mathcal{E}) = \mathsf{LLM}(\mathsf{Prompt}_{\mathsf{ToP}}(\mathcal{S}))$$

NeurIPS 25: Tree of Preferences for Diversified Recommendation

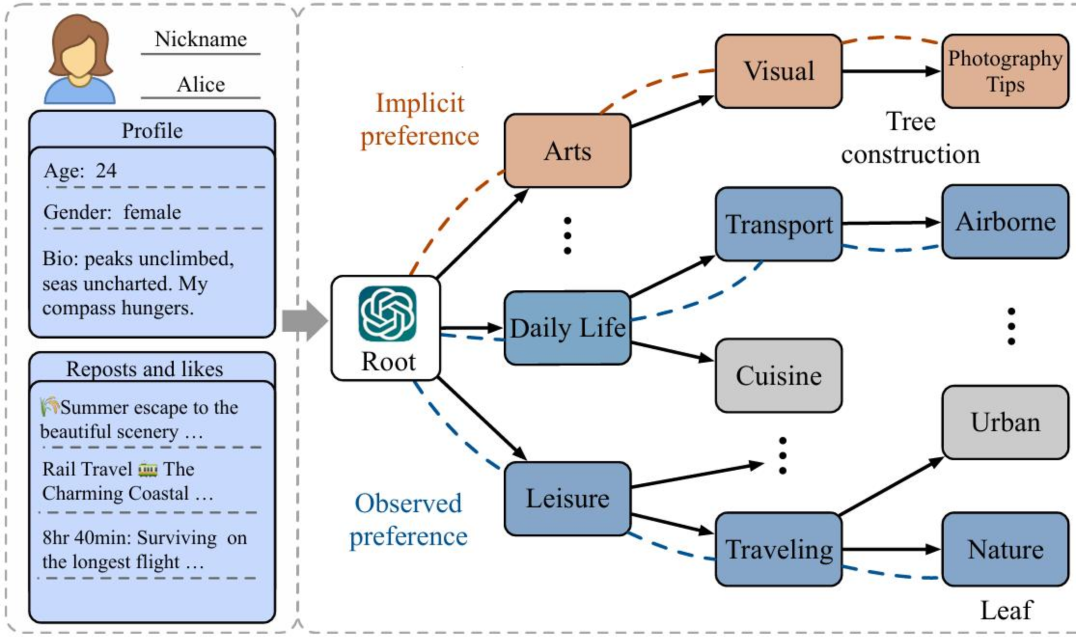

Notaion: $\mathsf{Prompt}_{\mathsf{PR}}$.
Usage: Preference reasoning.
Content: You are an expert system designed to capture user latent preferences through rationale reasoning. You will be provided with the user's profile, observed interactions with items, and a multi-level tree structure representing different user preferences, organized from broad to specific preferences (Tree of Preference, ToP). Using the profile and interactions, your task is to identify the user's latent preferences by exploring the ToP top-down, following the coarse-to-fine paths that best match the user's behavior and reasoning the rationale behind their actions, while identifying any unobserved preferences not reflected in the history.
Constraints:
- Perform a breadth-first search. At each level, select the preferences that best match the user behavior, storing the corresponding nodes.
- For subsequent levels, activate only the child nodes of the stored nodes at the previous level, continuing the selection until reaching final level.
- Reevaluate the exploration to identify any unobserved preferences, adding them if found, following the root-to-leaf path.
- Control the total number of selected paths as {number_of_paths}. The final output should be the leaf nodes of all selected paths.
. . .
User profile: {user_profile}.
User historical interactions: {user_interactions}.
ToP: {ToP_content}.
Return selected leaf nodes in the following format: {leaf_nodes_format}, along with a concise explanation of each selection in the format of {reason_format}.

Uncovering latent preferences

- Step 2: Capturing latent preferences via **rationale reasoning**.

$$\{v_1, ...v_n\} = \mathsf{LLM}(\mathsf{Prompt}_{\mathsf{PR}}(\mathcal{A}_u, \mathcal{R}_u)), \{v_1, ...v_n\} \subset \mathcal{V}_{\mathsf{Leaf}}$$

NeurIPS 25: Tree of Preferences for Diversified Recommendation

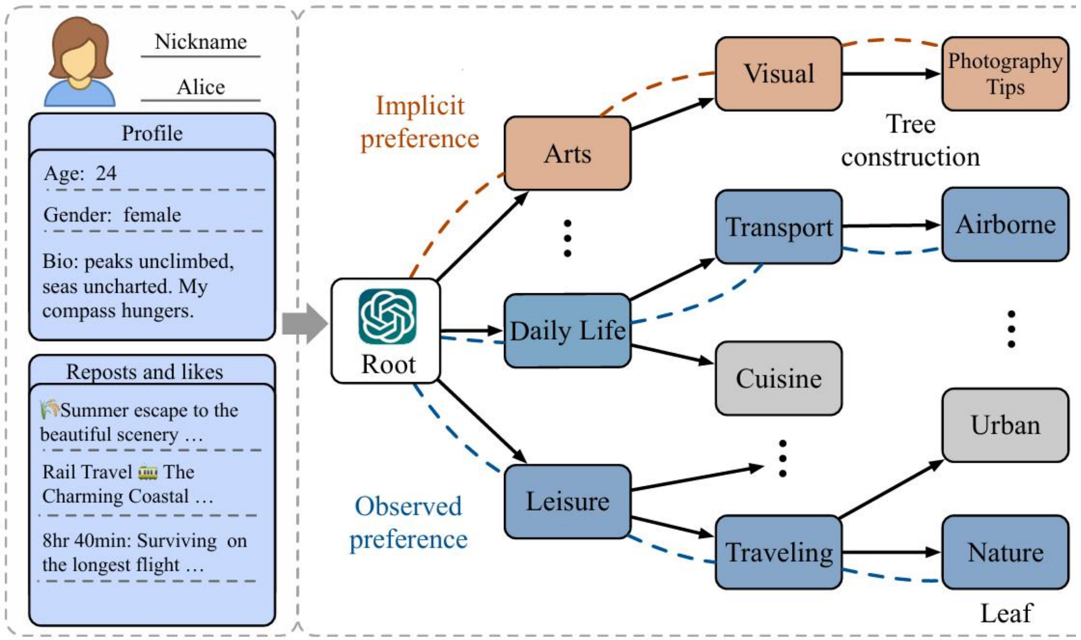

Notaion: $\text{Prompt}_{\text{PR}}$.
Usage: Preference reasoning.
Content: You are an expert system designed to capture user latent preferences through rationale reasoning. You will be provided with the user's profile, observed interactions with items, and a multi-level tree structure representing different user preferences, organized from broad to specific preferences (Tree of Preference, ToP). <u>Using the profile and interactions, your task is to identify the user's latent preferences by exploring the ToP top-down, following the coarse-to-fine paths that best match the user's behavior and reasoning the rationale behind their actions, while identifying any unobserved preferences not reflected in the history.</u>
Constraints:
- Perform a breadth-first search. At each level, select the preferences that best match the user behavior, storing the corresponding nodes.
- For subsequent levels, activate only the child nodes of the stored nodes at the previous level, continuing the selection until reaching final level.
- Reevaluate the exploration to identify any unobserved preferences, adding them if found, following the root-to-leaf path.
- Control the total number of selected paths as {number_of_paths}. The final output should be the leaf nodes of all selected paths.
. . .
User profile: {user_profile}.
User historical interactions: {user_interactions}.
ToP: {ToP_content}.
Return selected leaf nodes in the following format: {leaf_nodes_format}, along with a concise explanation of each selection in the format of {reason_format}.

Uncovering latent preferences

- Step 2: Capturing latent preferences via **rationale reasoning**.
- A variant with Lora fine-tuning:

User profile & history + ToP → User preference

NeurIPS 25: Tree of Preferences for Diversified Recommendation

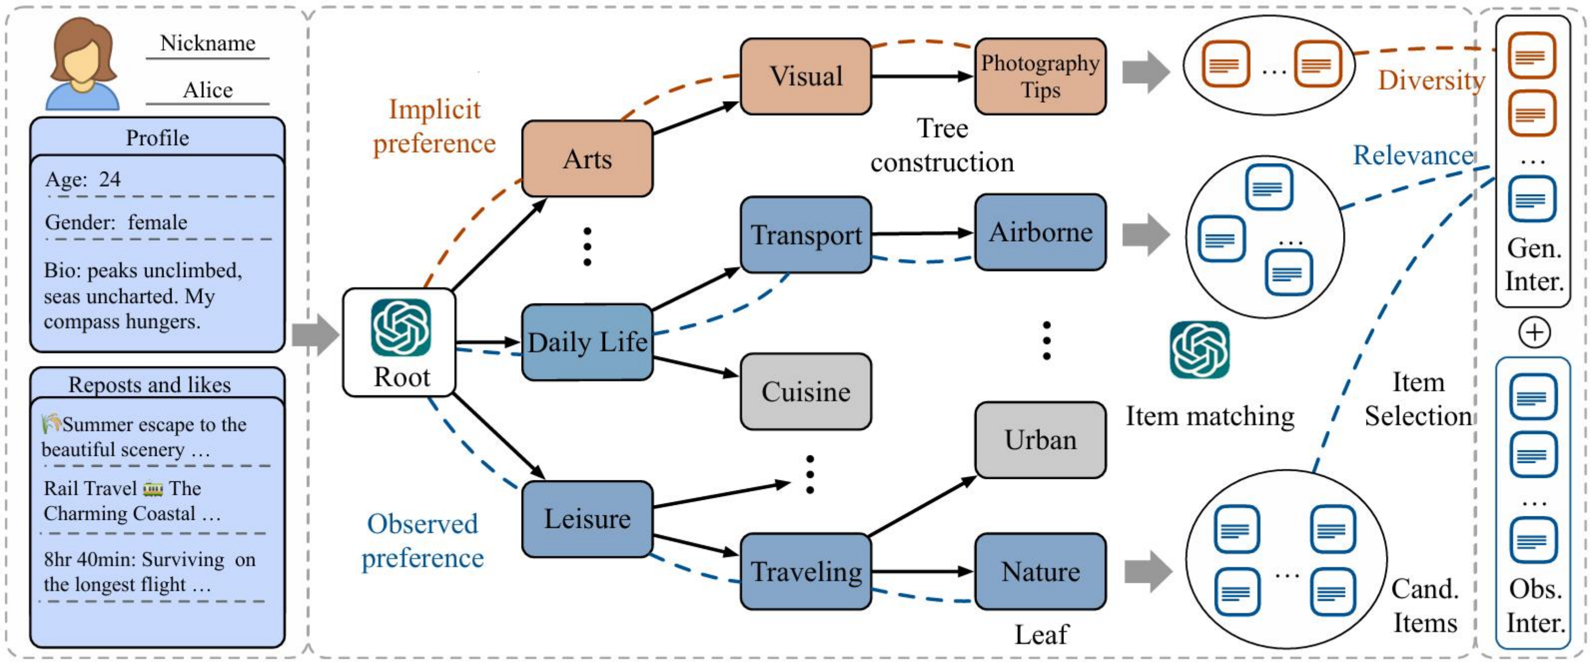

Guiding diverse and relevant recommendations

- Matching candidate items with preferences

$$v_i = \mathsf{LLM}(\mathsf{Prompt}_{\mathsf{IM}}(\mathcal{A}_i)), v_i \in \mathcal{V}_{\mathsf{Leaf}}$$

NeurIPS 25: Tree of Preferences for Diversified Recommendation

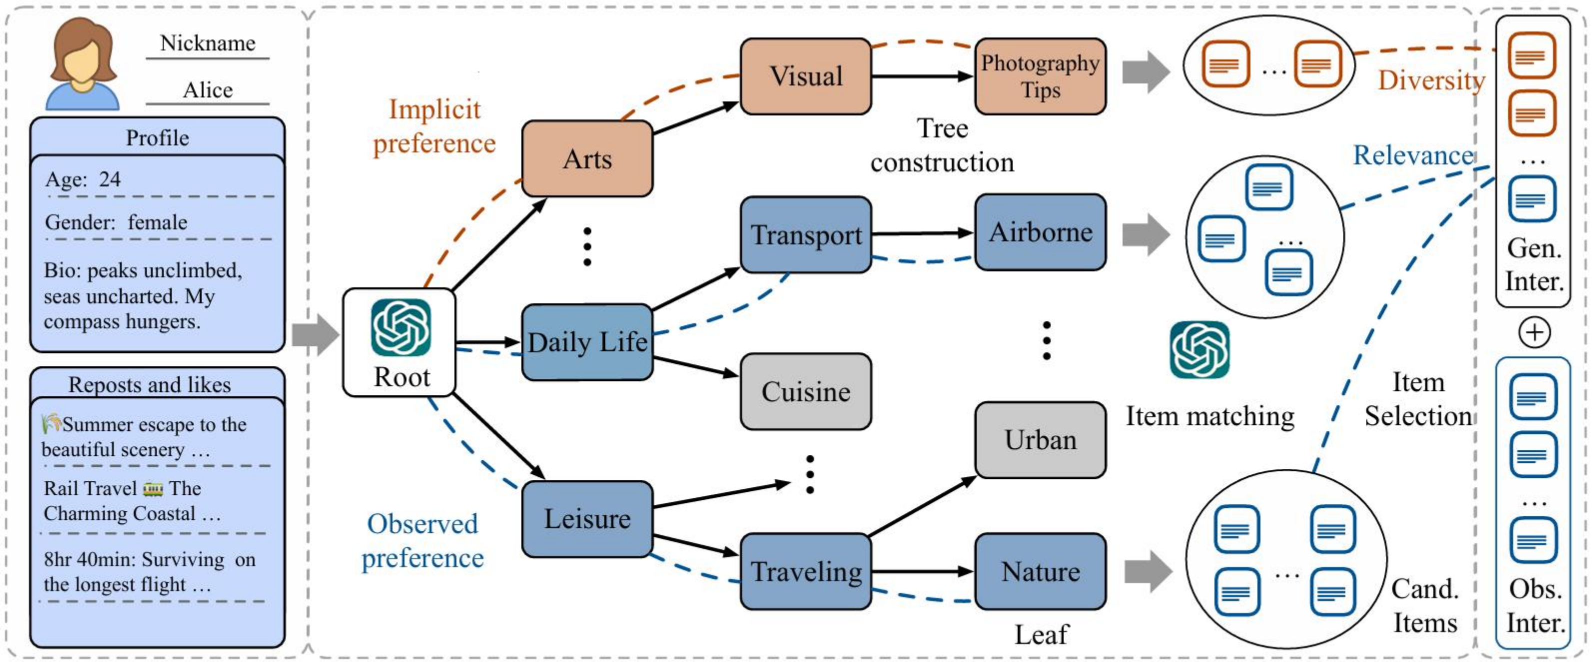

Guiding diverse and relevant recommendations

- Data generation for debiasing user interactions.

$$s(u,i) = (1-\lambda)\cdot s_{\text{rel}}(u,i) + \lambda \cdot s_{\text{div}}(u,i)$$

- Dynamic user selection: reduce token burden

NeurIPS 25: Tree of Preferences for Diversified Recommendation

Experiments

- Twitter, Weibo, Amazon
- Metric
  - Relevance: recall
  - Diversity: category-entropy

- Baselines
  - Heuristics: random, MMR…
  - Traditional: LCD-UC, CDM…
  - LLM-based: LLMRec, LLM4Rerank…

	Twitter				**Weibo**				**Amazon**			
	R@50	R@100	CE@50	CE@100	R@50	R@100	CE@50	CE@100	R@50	R@100	CE@50	CE@100
LightGCN*	0.0567	0.0830	1.2841	1.3413	0.1052	0.1669	0.9905	1.0763	0.1362	0.2105	0.5004	0.5609
Random	0.0494^{-}	0.0730^{-}	1.2954^{+}	1.3475^{-}	0.0988^{-}	0.1577^{-}	0.9995^{+}	1.0801^{+}	0.1356^{-}	0.2062^{-}	0.5184^{+}	0.5792^{+}
MMR	0.0540^{-}	0.0790^{-}	1.3078^{+}	1.3550^{-}	0.0990^{-}	0.1578^{-}	1.0081^{+}	1.1005^{+}	0.1371^{+}	0.2115^{+}	0.5141^{+}	0.5755^{+}
DPP	0.0467^{-}	0.0765^{-}	1.3048^{+}	1.3532^{-}	0.0963^{-}	0.1530^{-}	**1.0362**$^{+}$	1.1150^{+}	0.1283^{-}	0.2035^{-}	0.5181^{+}	0.5748^{+}
CDM	0.0562^{-}	0.0814^{-}	1.2986^{+}	1.3461^{-}	0.1014^{-}	0.1620^{-}	1.0018^{+}	1.0912^{+}	0.1349^{-}	0.2103^{-}	0.5228^{+}	0.5816^{+}
Box	0.0527^{-}	0.0741^{-}	1.2844^{+}	1.3407^{-}	0.0996^{-}	0.1587^{-}	1.0238^{+}	1.1034^{+}	0.1228^{-}	0.2019^{-}	0.5186^{+}	0.5844^{+}
LCD-UC	0.0517^{-}	0.0768^{-}	1.3154^{+}	1.3784^{-}	0.1038^{-}	0.1625^{-}	1.0211^{+}	1.0956^{+}	0.1295^{-}	0.2065^{-}	0.5202^{+}	0.5842^{+}
LLMRec-MMR	0.0558^{-}	0.0820^{-}	1.3056^{+}	1.3551^{-}	0.1041^{-}	0.1662^{-}	1.0246^{+}	1.1182^{+}	0.1363^{+}	0.2113^{+}	0.5177^{+}	0.5836^{+}
LLM4Re-A	0.0562^{-}	0.0827^{-}	1.2855^{+}	1.3424^{-}	0.1032^{-}	0.1656^{-}	0.9891^{-}	1.0745^{-}	0.1359^{-}	0.2049^{-}	0.5028^{+}	0.5863^{+}
LLM4Re-AD	0.0560^{-}	0.0822^{-}	1.2864^{+}	1.3466^{-}	0.1044^{-}	0.1652^{-}	1.0001^{+}	1.0823^{+}	0.1332^{-}	0.2042^{-}	0.5131^{+}	0.5827^{+}
ToP-Rec	**0.0586**$^{+}$	**0.0841**$^{+}$	**1.3275**$^{+}$	**1.3852**$^{+}$	**0.1054**$^{+}$	**0.1667**$^{-}$	1.0333^{+}	**1.1369**$^{+}$	**0.1380**$^{+}$	**0.2120**$^{+}$	**0.5298**$^{+}$	**0.5902**$^{+}$

NeurIPS 25: Tree of Preferences for Diversified Recommendation

Experiments
- Twitter, Weibo, Amazon
- Metric
  - Relevance: recall
  - Diversity: category-entropy

- Baselines
  - Heuristics: random, MMR…
  - Traditional: LCD-UC, CDM…
  - LLM-based: LLMRec, LLM4Rerank…