Bridging Recommendation and Marketing via Recurrent Intensity Modeling

Yifei Ma, Ge Liu, Anoop Deoras
AWS AI Labs

Popularity Bias Effects in User-Centric Recommender Systems

- Popular items often get exposed to more users, who in turn drive the popularity of the popular items
- Must design mechanisms to explore the underexposed “cold-start” items
- Practical solution: offer every item a free marketing budget for exploration
- Question: how to integrate marketing into an existing recommender system?

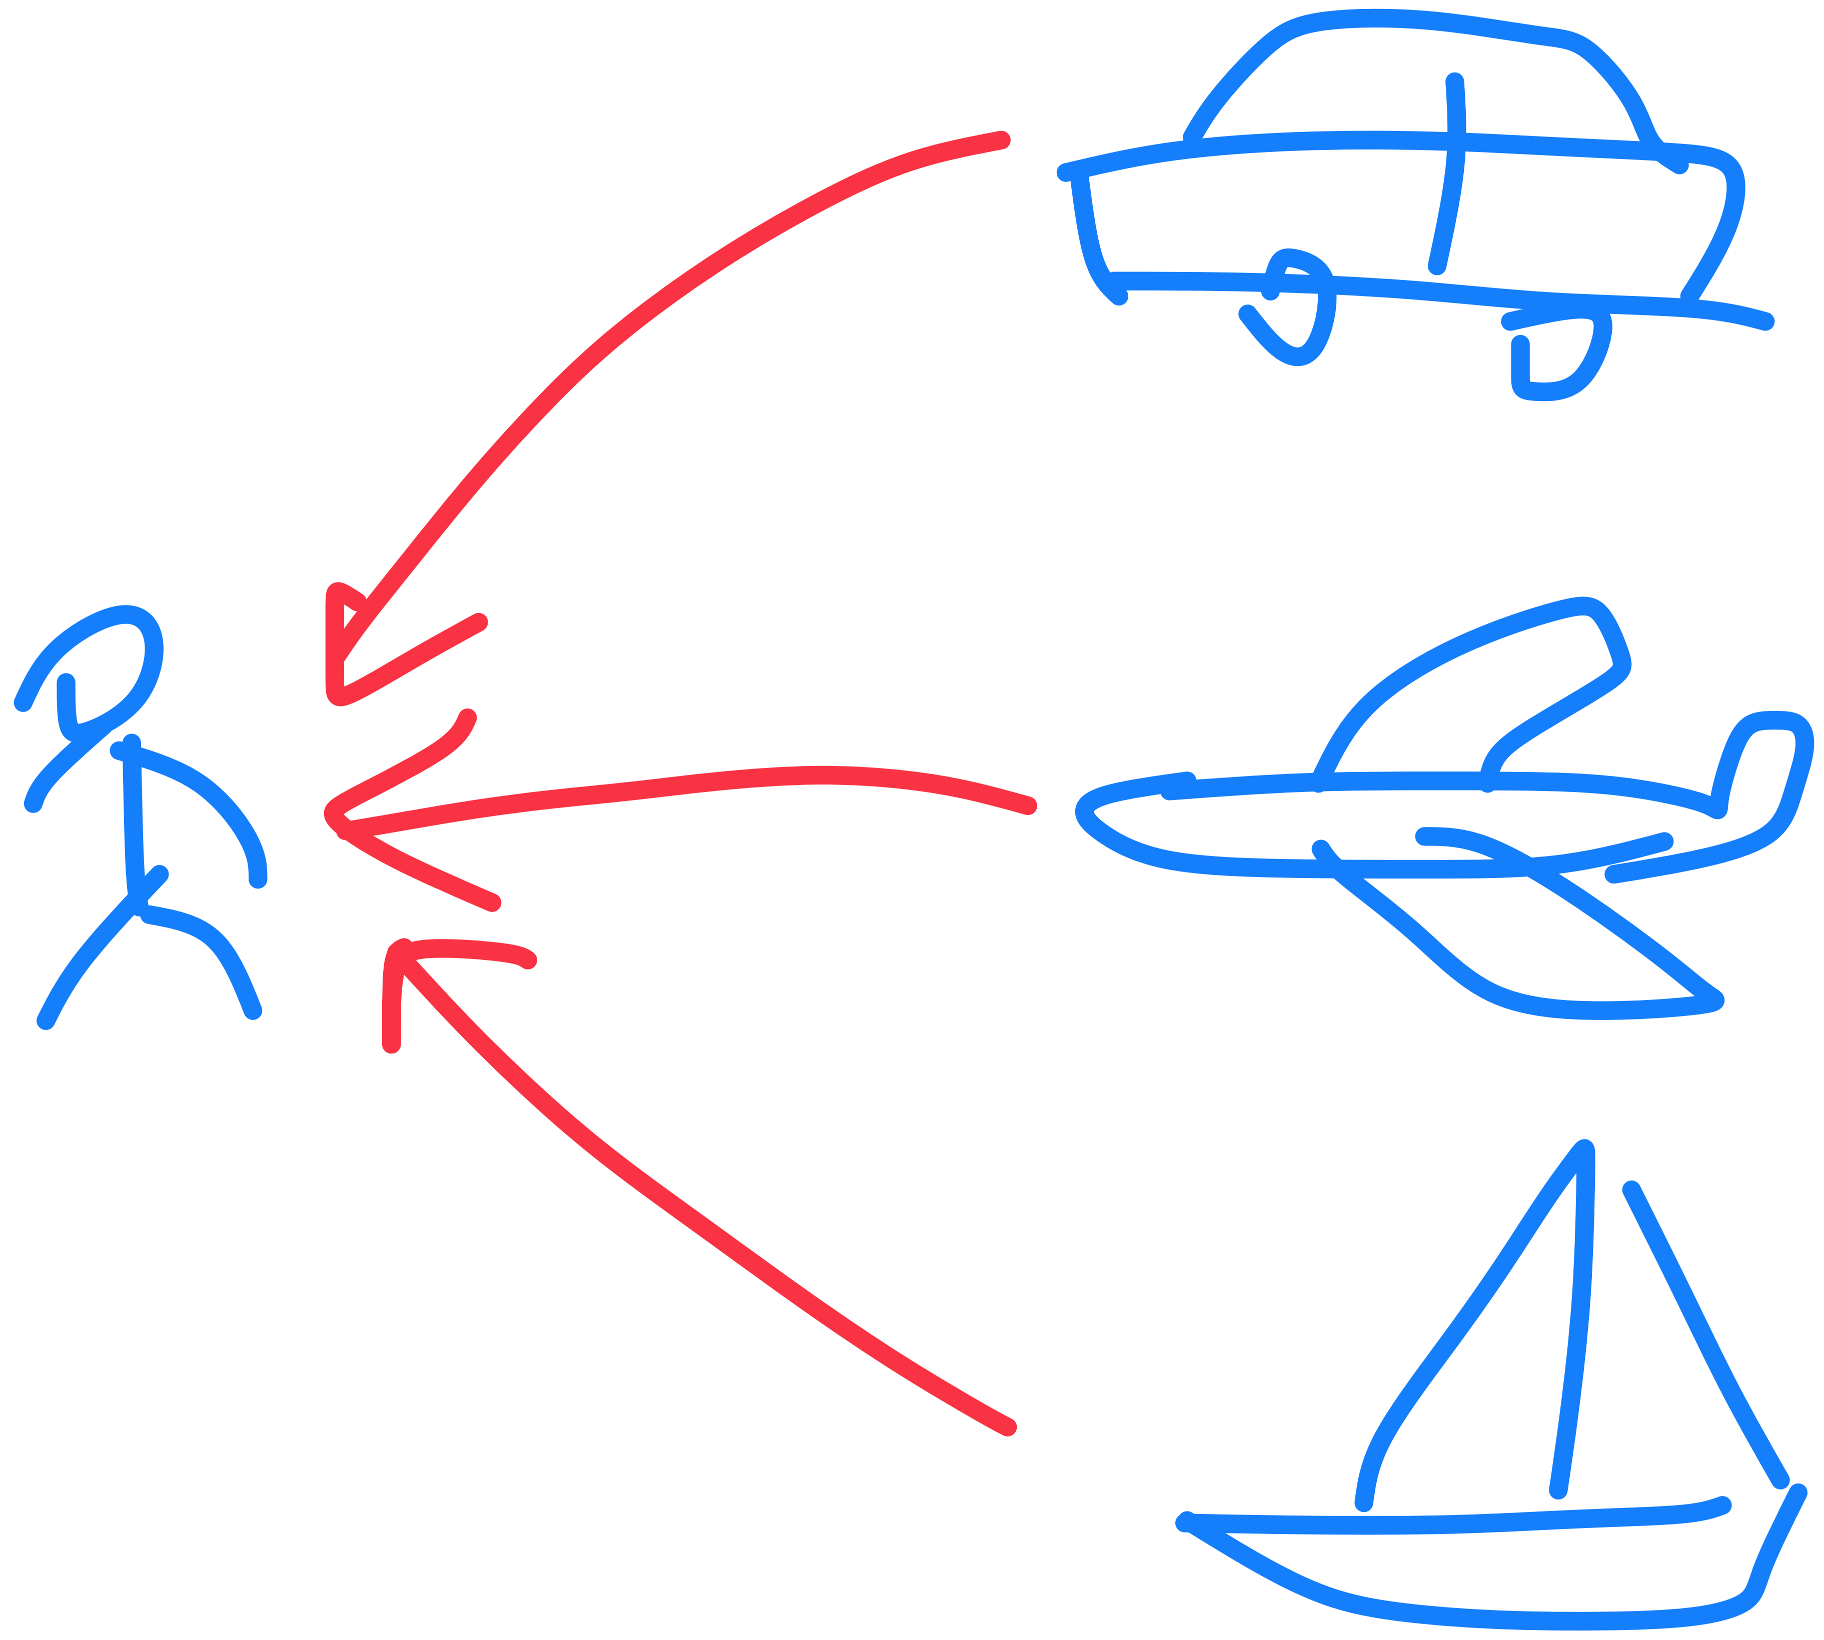

Main Contributions of Our Work

- UserRec: recommend users for item marketers
- OnlnMtch: match multiple users and items in an online real-time environment
- Show results with improved item diversity with least compromises in recommendation relevance

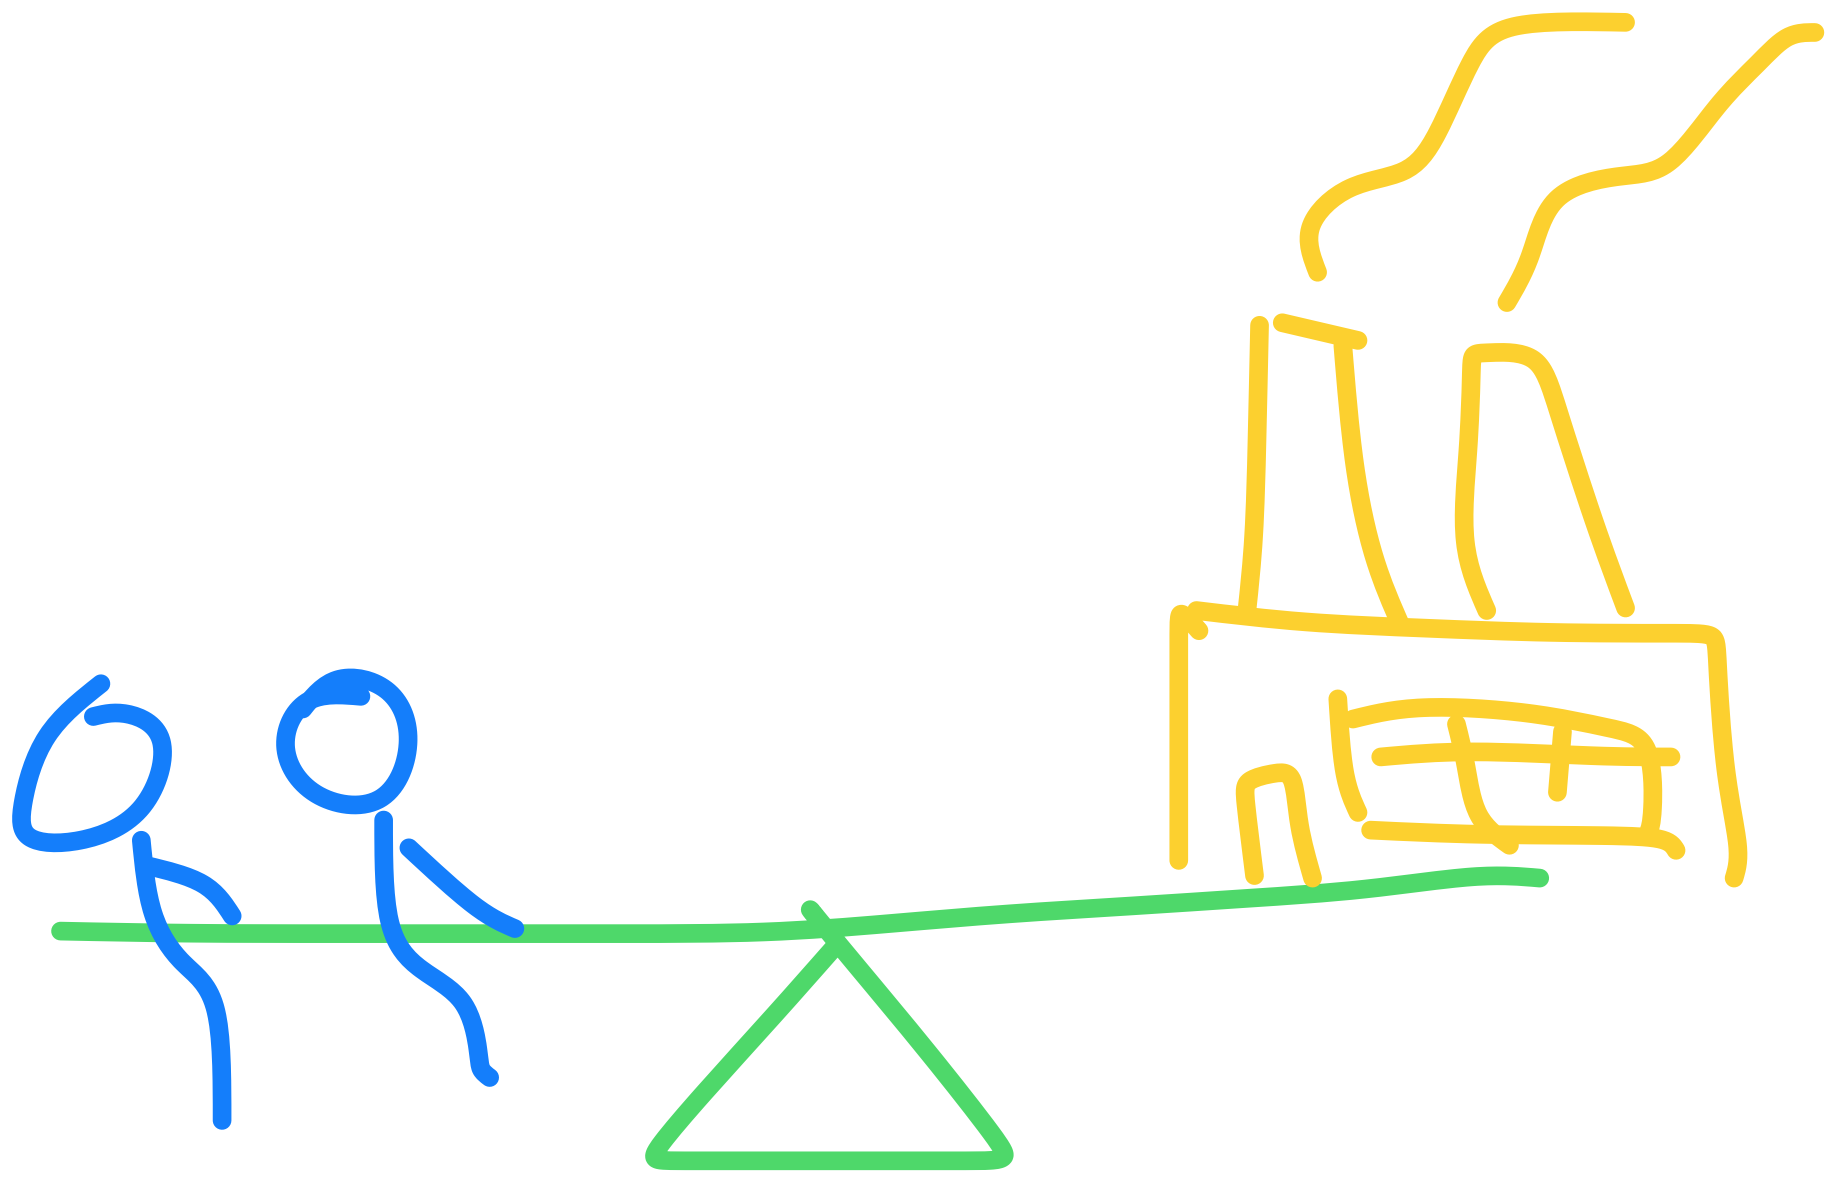

UserRec Generalizes Related-Product Advertising

- Place advertisements on related items, such as frequently co-viewed products
  - Increase exposure to users with potential interests based on a single event in their browsing histories
- UserRec considers the users' full-history information (time-series data)
- Challenge: the time series data cannot be compared just like random variables
  - E.g., compare between a matching but less active user and an active but less matching user?
  - Not feasible with probability density estimation based on the users' hidden-state variables learned for next-item recommendation

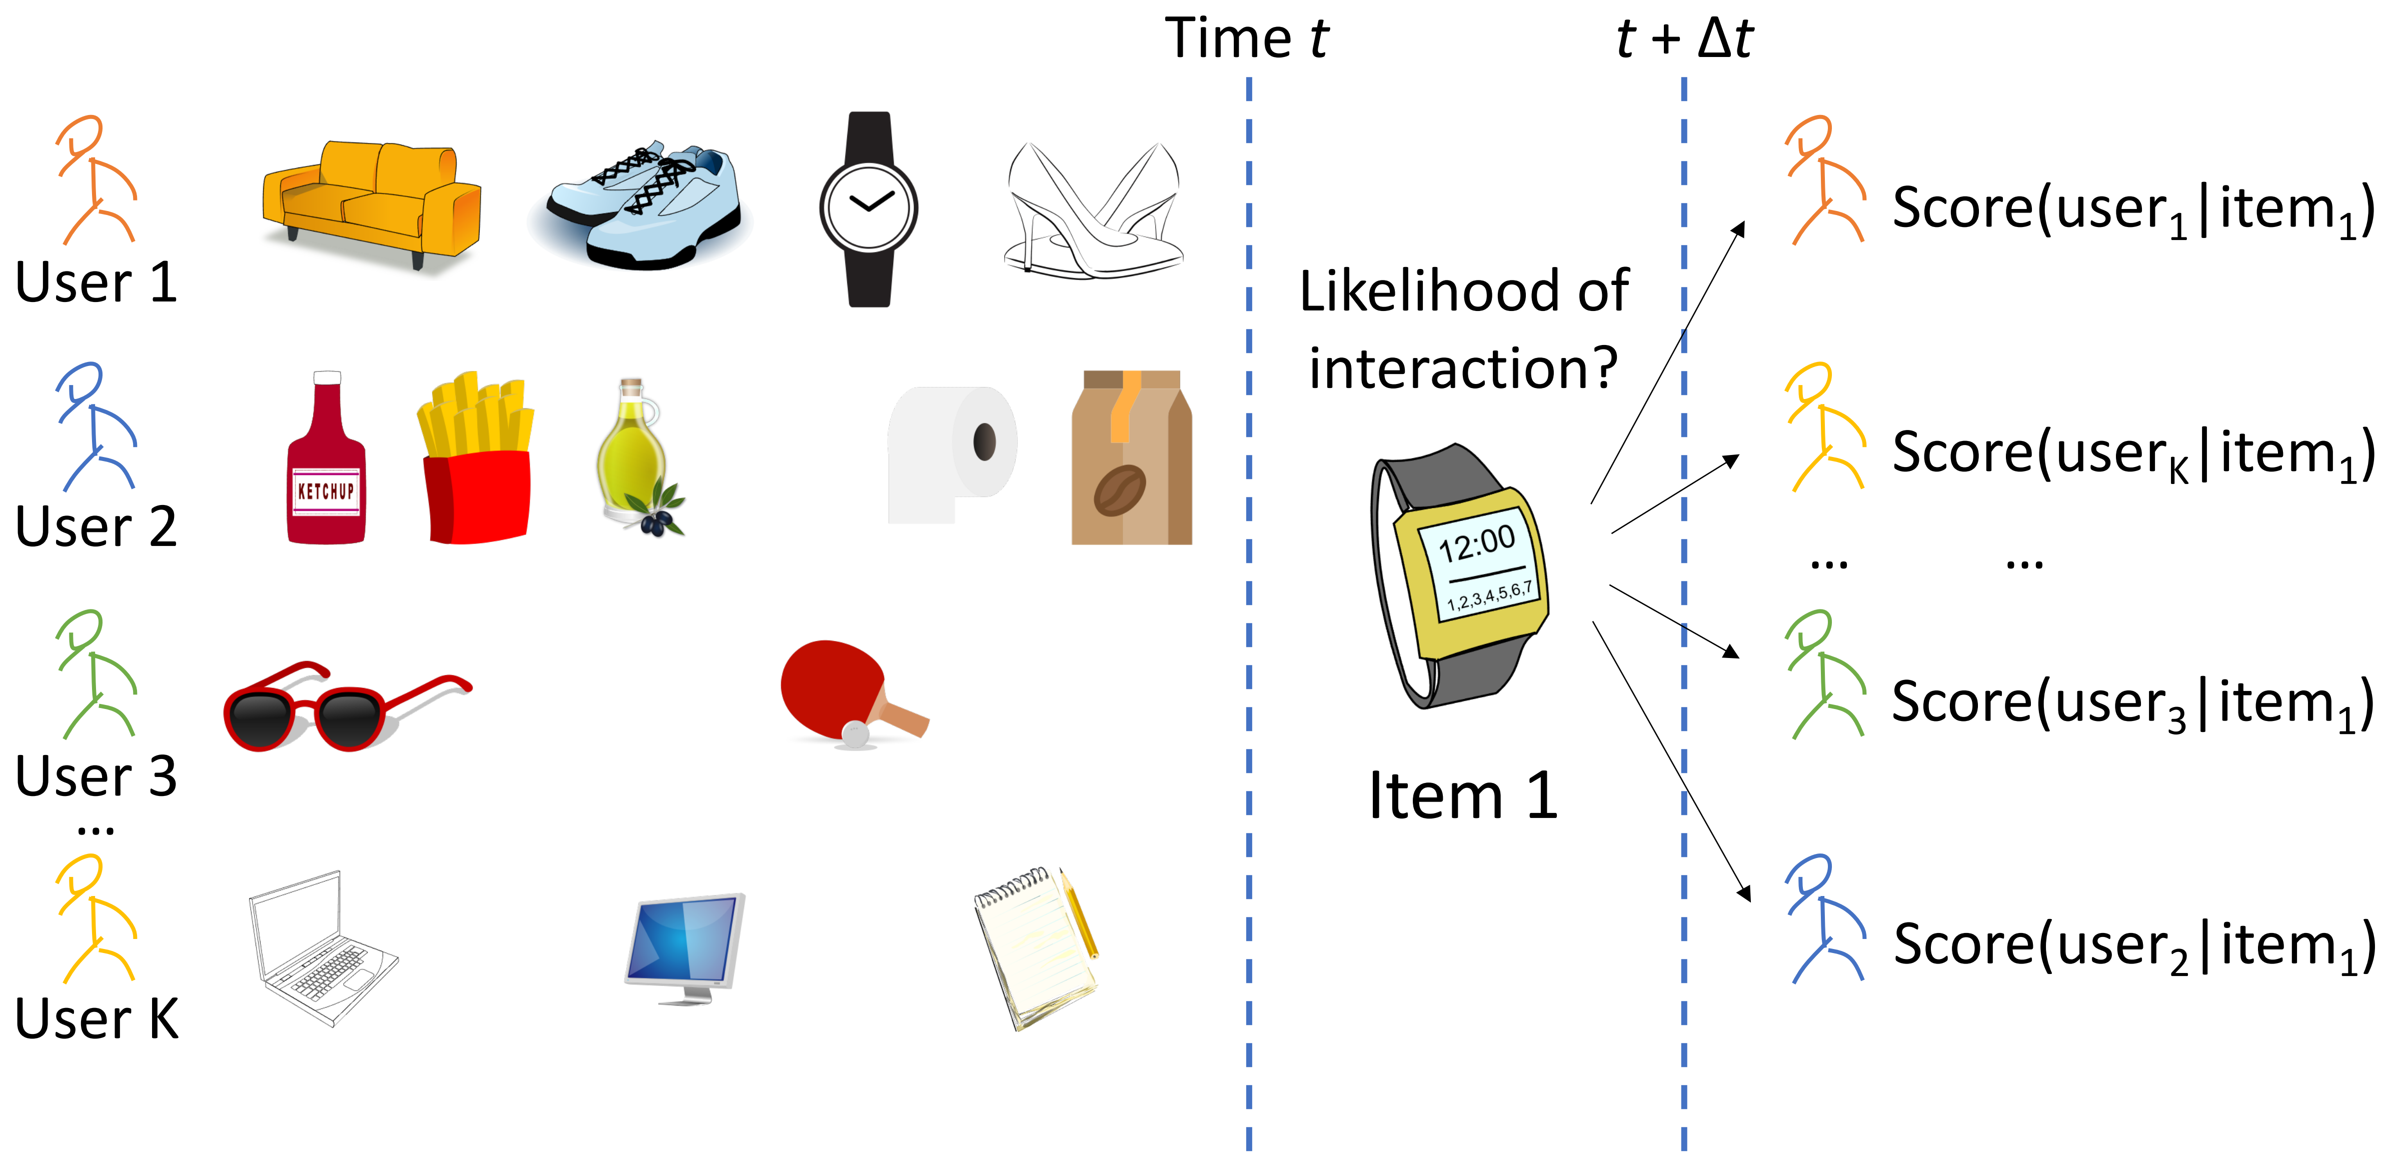

Introducing Recurrent Intensity Modeling (RIM)

- Key observation: find users who are most likely to interact with the given item in *the next period of time*
- Formally defined by the *intensity parameter λ(x, y)* in Temporal Point Processes, where x is a user-state and y is an item:

$$x = s(t) \qquad \lambda(x, y) = \lim_{\Delta t \to 0} \mathbb{E}\left[\frac{\Delta N}{\Delta t} \middle| x, y \right]$$

- MTPP Decomposition λ(x, y) = λ(x) p(y | x), similar to Bayes' rule
  - Reuse RNN/Transformer models from recommender systems for p(y | x) for simplicity

Three Marketing Scenarios

	UserRec	Offline Match	Online Match
Pick top users among current users for one marketed item	Y	Y	Y
Market for multiple items at the same time		Y	Y
• Set minimal-exposure constraints for every item for exploration, combined with max-capacity constraints for every user		Y	Y
Responsive marketing when a user appears just like RecSys			Y
• Distribute marketing materials in a future period of time			Y

Constrained Optimization and Dual Algorithm

$$\max_{0 \leq \pi \leq 1} \mathbb{E}_{xy}(\lambda_{xy}\pi_{xy})$$

Binary matching decision

$$\text{s.t.} \begin{cases} \mathbb{E}_y(\pi_{xy}) \leq \alpha, \forall x \\ \mathbb{E}_x(\pi_{xy}) \geq \beta, \forall y \end{cases}$$

Every user can receive <= α portion of all items

Every item must reach >= β portion of all users

Expectations over a distribution of users $x \in \mathcal{X}$ and a finite set of items $y \in \mathcal{Y}$

- Let v_y be the dual variable for item y, learned from historical data
- When a user appears online with a random state x,
  - return top items based on the modified scores $\{\lambda_{xy} - v_y : y \in \mathcal{Y}\}$
- $v_y \leq 0$ corresponds to positive biases for cold-start exploration

Featured OnlnMtch Experiment

- Take user-item interaction data from MovieLens* dataset
- Pick test-start time T and train RIMs from past data
- Define item min-exposure constraint 0<β<1
- Learn Dual planner based on user-state distribution at T' to simulate OnlnMtch at T
- Vary β to generate a trade-off curve between relevance & diversity

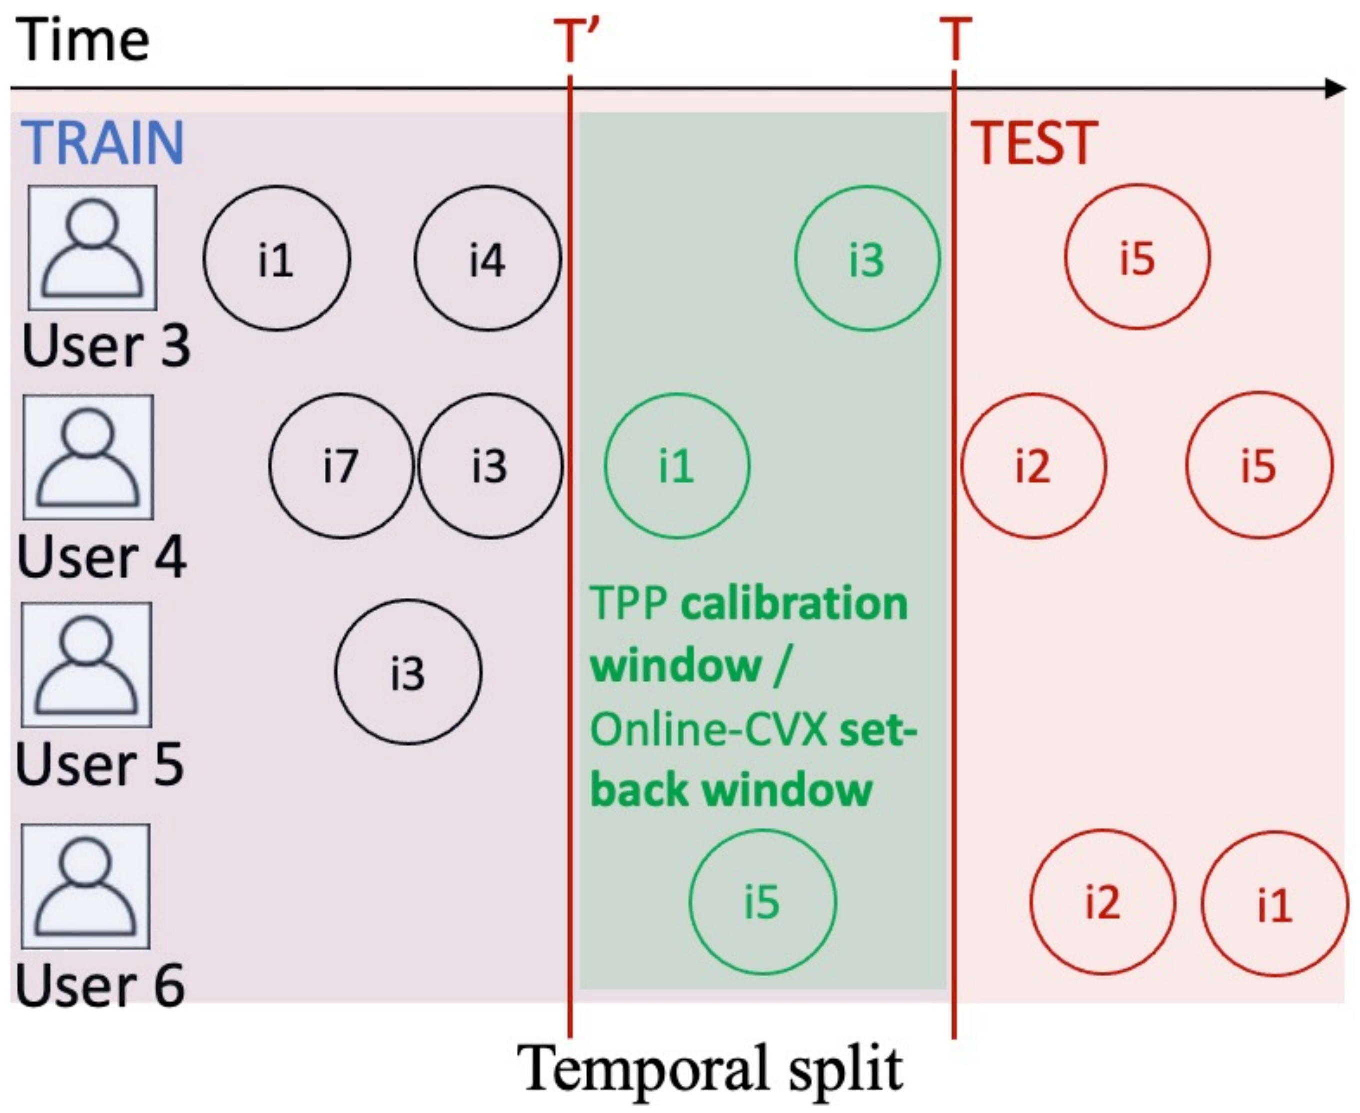

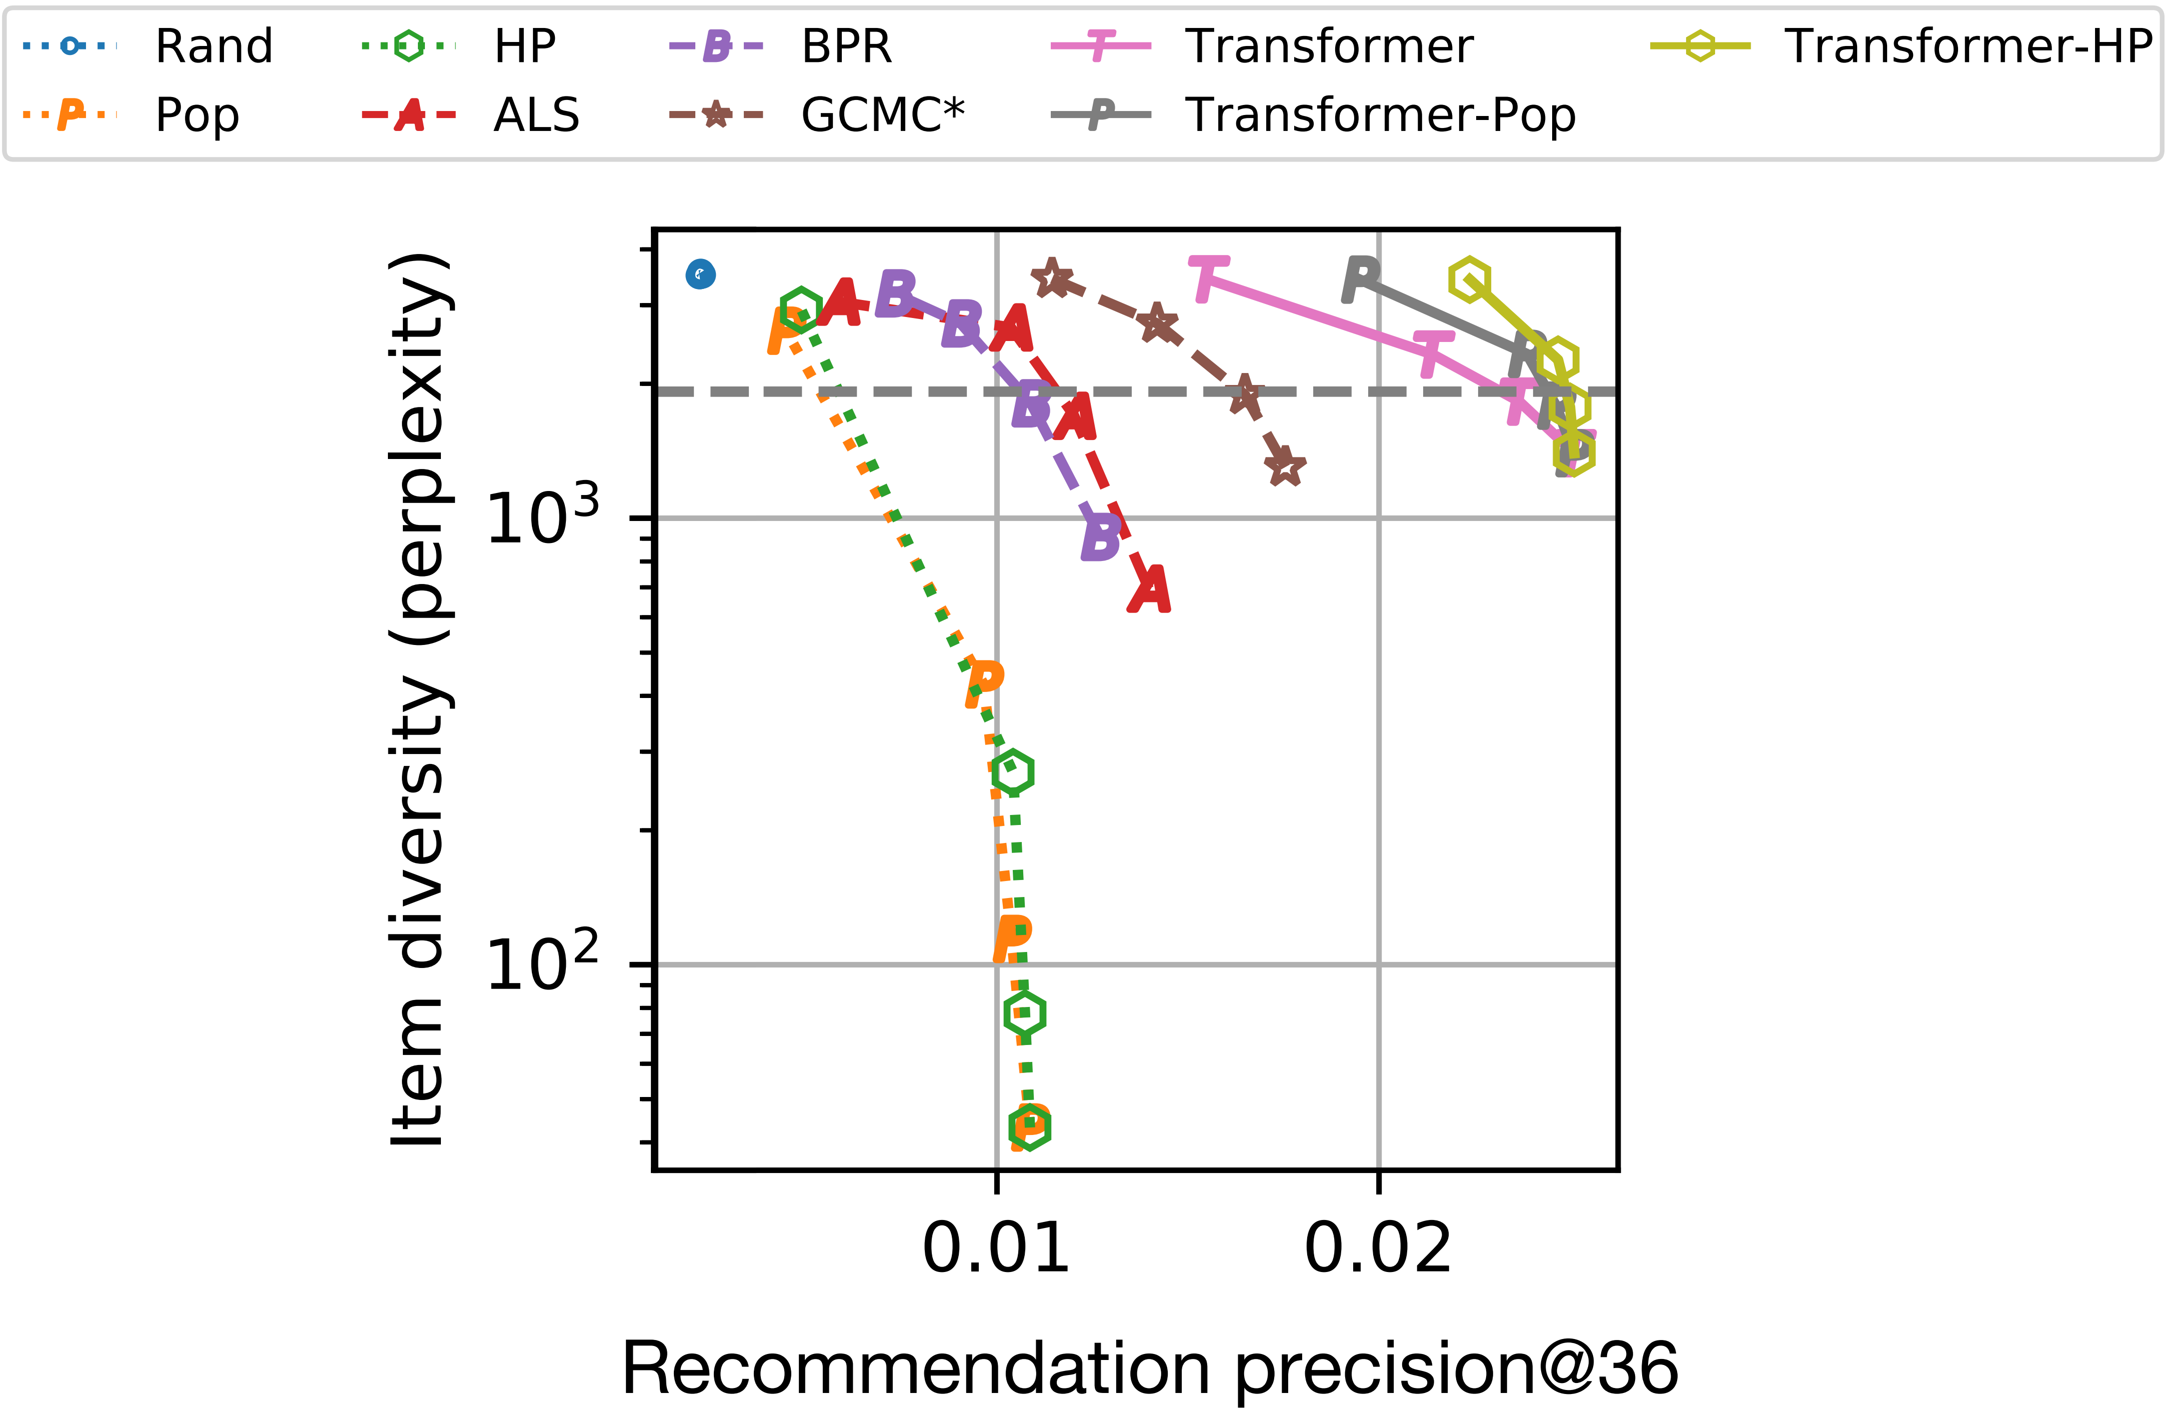

- Transformer-HP (a variant of RIMs) holds the Pareto front
- GitHub repository: https://github.com/awslabs/recurrent-intensity-model-experiments

* F. Maxwell Harper and Joseph A. Konstan. 2015. The MovieLens Datasets: History and Context. ACM Transactions on Interactive Intelligent Systems (TiiS) 5, 4, Article 19 (December 2015), 19 pages. DOI=http://dx.doi.org/10.1145/2827872

Conclusions

- Recommender systems are often user-centric, but they also consider item exploration as a key challenge
- We borrow ideas from marketing to meet the exploration challenge
  - RIM extends sequential recommendation models for personalized marketing
  - Dual modifies rec scores to distribute items across user & time space
- Show favorable relevance-diversity trade-offs using back-test time periods
- Future work: close the loop with human feedback to study the long-term exploration effects after multiple rounds of matching iterations