Toward the Development of Richer Properties for Recommender Systems

David Shriver

Problem

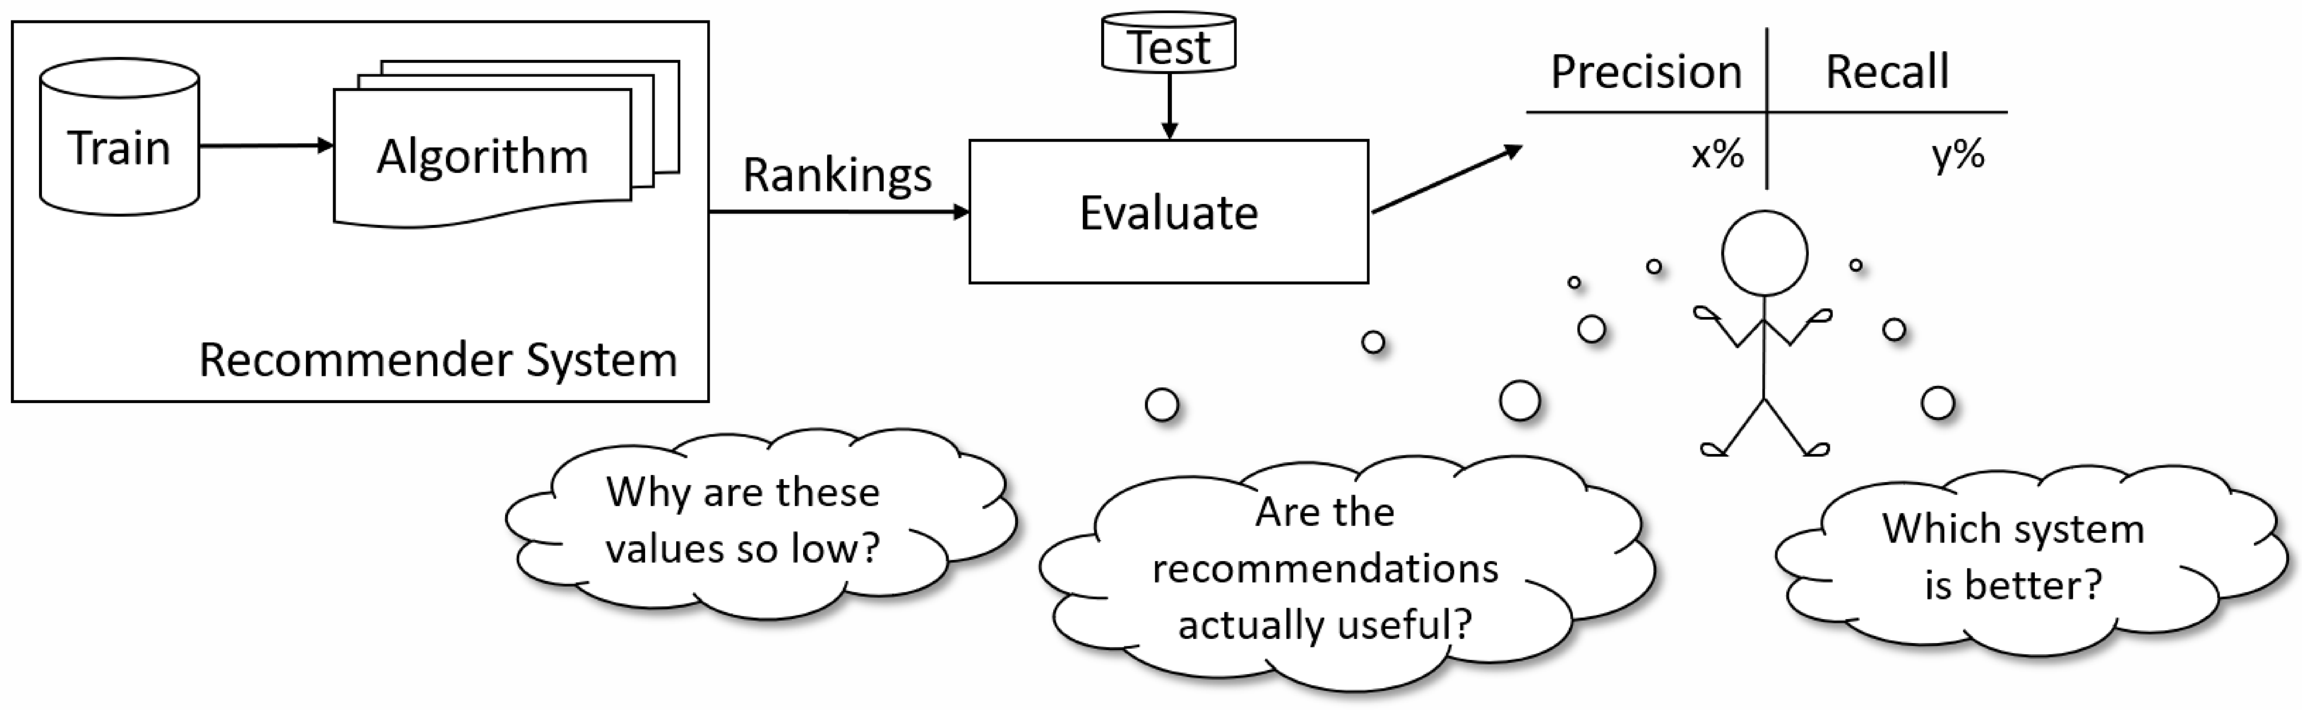

Precision and recall:

- Provide little information or intuition about the usefulness of recommendations
- Depend on the number of relevant items relative to the size of the dataset
- Can be misleading if recommendations are anomalous or counterintuitive

Deriving Recommender Properties

View RS as a set of Users and Items, related through ratings and ranks.

Define property templates that capture useful recommender system properties.

Two such templates are:

1. Number of Recommended Items (NRI)

$$\#\{i \mid i \in I \land \exists u \in U : rank(i,u) \neq \perp\} \geq n$$

2. Number of Recommendation Sets (NRS)

$$\#\{\{i \mid i \in I \land rank(i,u) \leq k\}\ u \in U\} \geq n$$

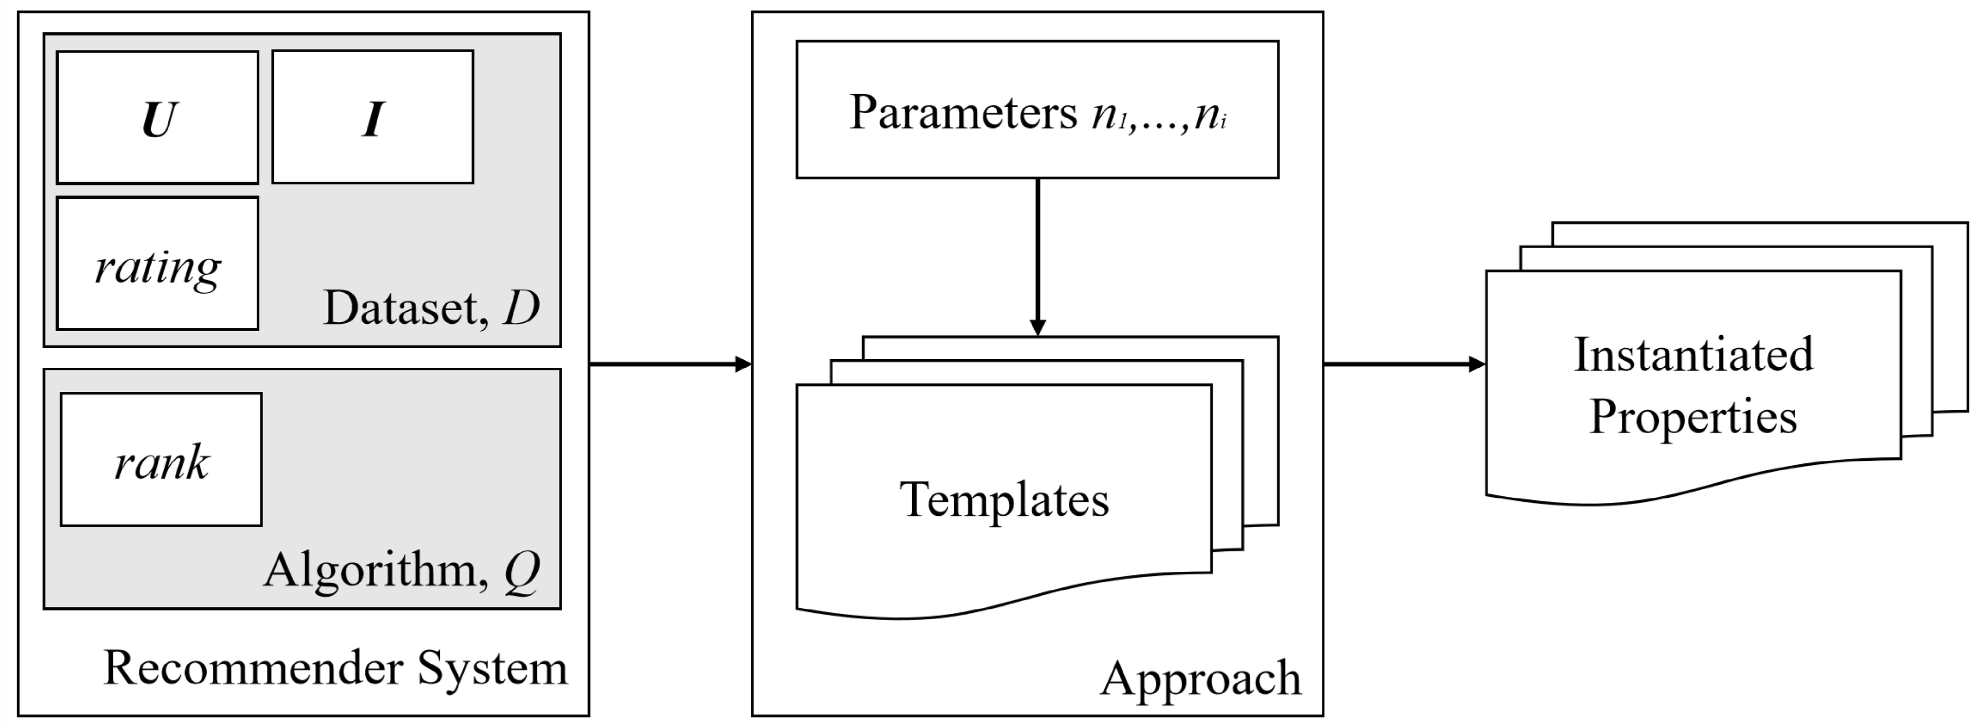

Instantiate parameters $n_1 \ldots n_i$ to cause the property templates to evaluate to true, based on ratings provided to a RS and the rankings produced.

Initial Exploration

We instantiate the NRI and NRS templates for the following algorithms on the MovieLens-20M dataset:

- Item-Item, a model-based collaborative filtering algorithm
- LightFM, a hybrid, content-based recommender algorithm

Algorithm	Precision	Recall	NRI	NRS
Item-Item	0.00179	0.000341	2791	138417
LightFM	0.00520	0.00213	786	85784

Precision and recall:

- Offer little intuition about behavior
- Are miseadingly small
- Indicate LightFM may be superior

Item-Item recommended 2000 more unique items than LightFM.

Indicates that LightFM produced more conservative recommendations.

Item-Item recommended 50000 more unique recommendation sets.

Indicates that Item-Item produced more personalized recommendations.

Using Instantiated Properties

- *Algorithm developers*: perform a richer comparison of the behavior of their proposed algorithm in relation to others
- *RS developers*: ensure that a RS meets some specification. They can also be used as assertions as a dataset evolves
- *Recommender clients*: compare RS and guide the selection of a recommender for a particular application

Future Work

- A more exhaustive exploration of the space of templates that can be generated using our model
- Define templates that capture learned influence relationships
- Use instantiated properties to assist in the explanation of anomalous recommendation behavior

Acknowledgements

This work has been supported in part by National Science Foundation Award #1526652. I would also like to thank Sebastian Elbaum, David S. Rosenblum, and Matt Dwyer for their guidance during this work.