From Preference into Decision Making

Modeling User Interactions in Recommender Systems

Qian Zhao (Bloomberg, work done during Ph.D. at GroupLens)
Martijn C. Willemsen (Eindhoven University of Technology)
Gediminas Adomavicius (Carlson School of Management, University of Minnesota)
F. Maxwell Harper (Amazon)
Joseph A. Konstan (GroupLens, University of Minnesota)

Recommender Systems

- Most prior research focuses on
  - Learning from user action feedback (e.g., ratings or clicks etc.)
- This work is about
  - Learning from all user browsing (both user action and inaction) activities

Classical Approach: From Explicit/Implicit Feedback to Preference

Explicit/Implicit Feedback and Preference

What about other (inaction) items?

Key Missing Factor

- Within page comparison
- Decision of action vs. inaction

Zhao et al. *RecSys'18*

Explicit/Implicit Feedback and Preference

What if this is the 3rd time the user sees this page?

Key Missing Factor

- Temporal dependency to prior browsing

Zhao et al. *CSCW'17*

Example Impact on User Experience

MovieLens

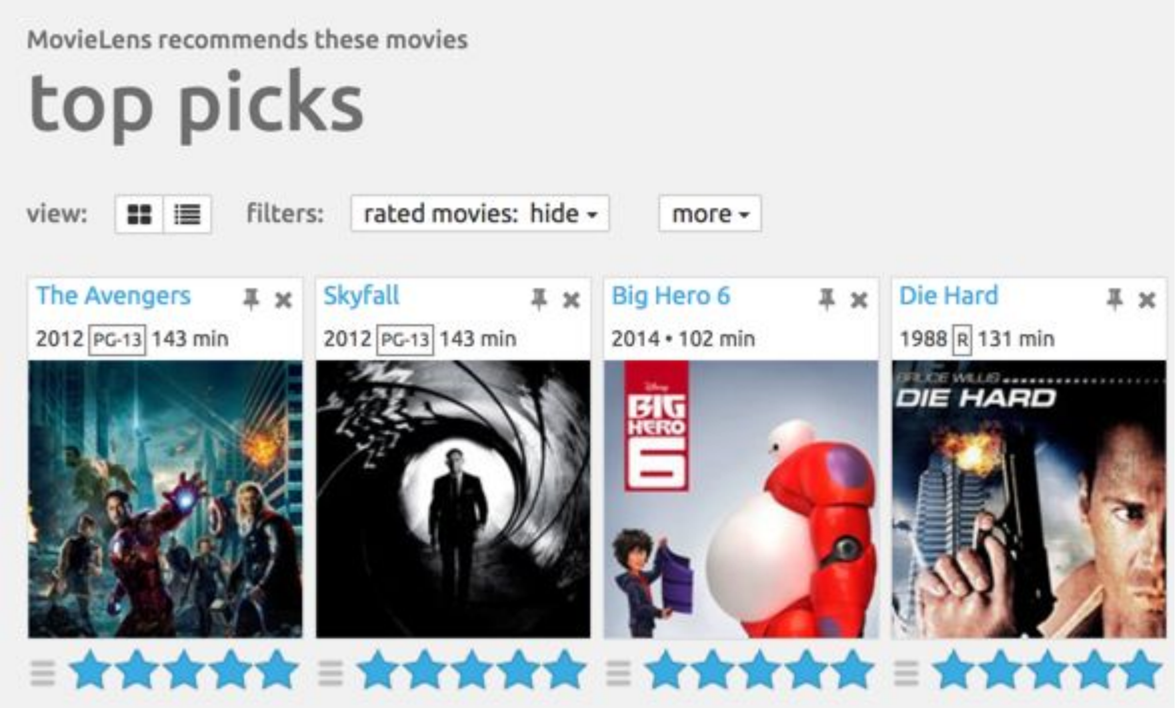

Zhao et al. *CSCW'17*

Again and Again

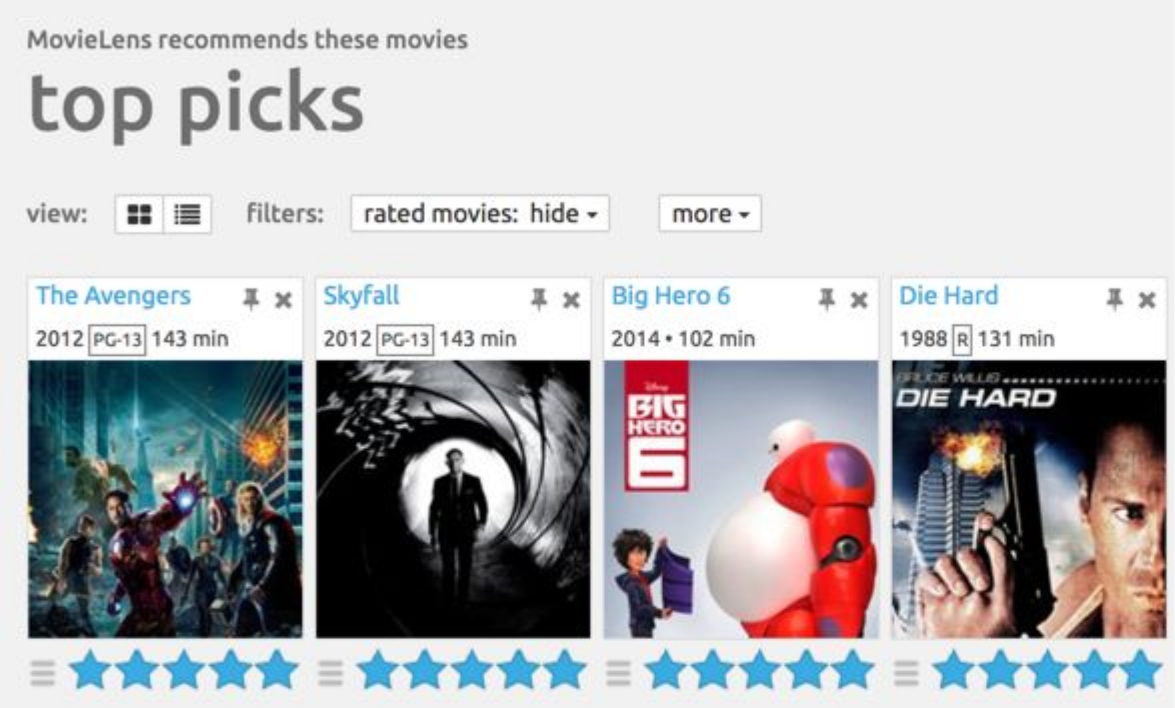

Zhao et al. *CSCW'17*

Youtube

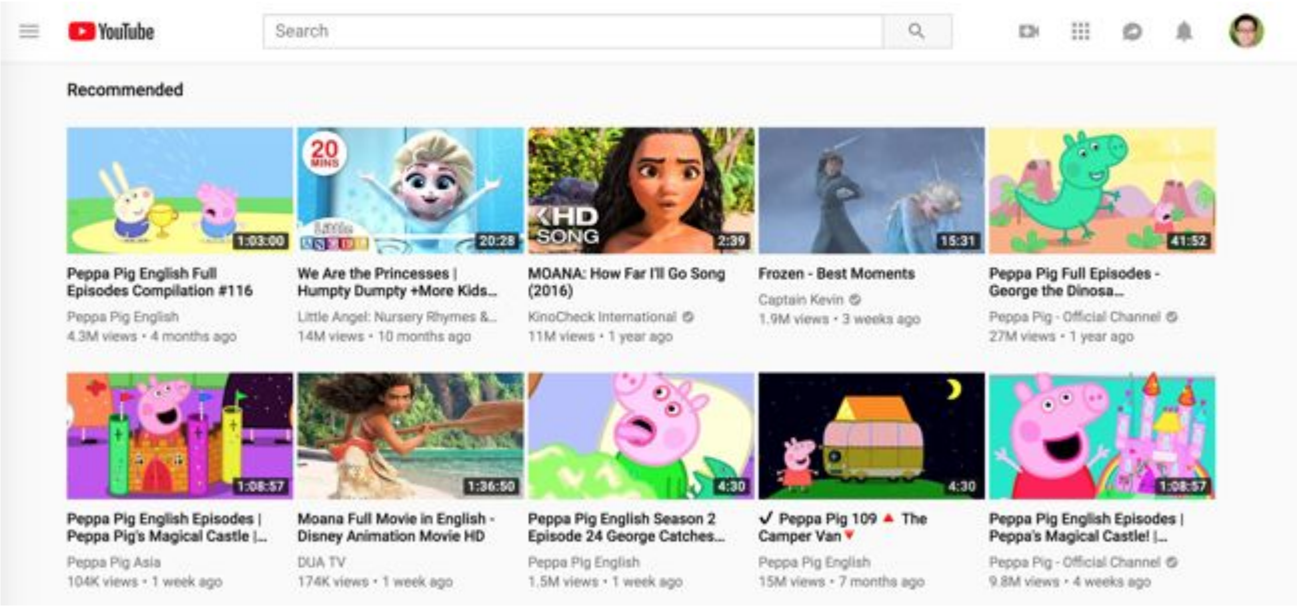

Zhao et al. *RecSys'18*

Interested in two of them

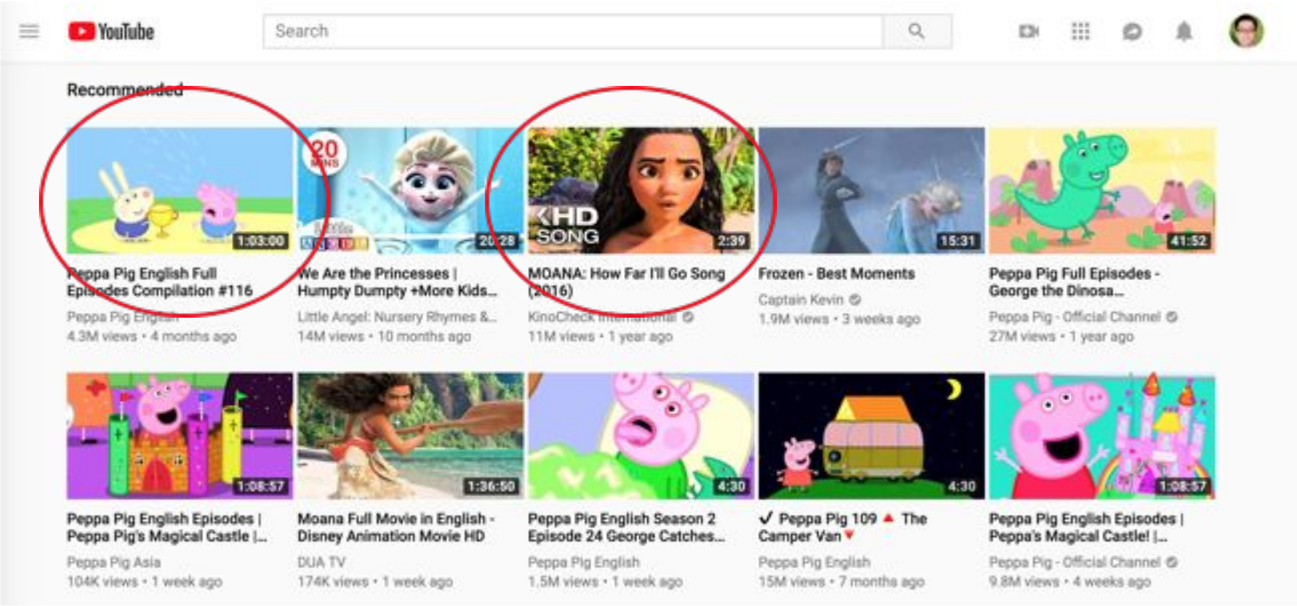

Zhao et al. *RecSys'18*

Watch one of the two

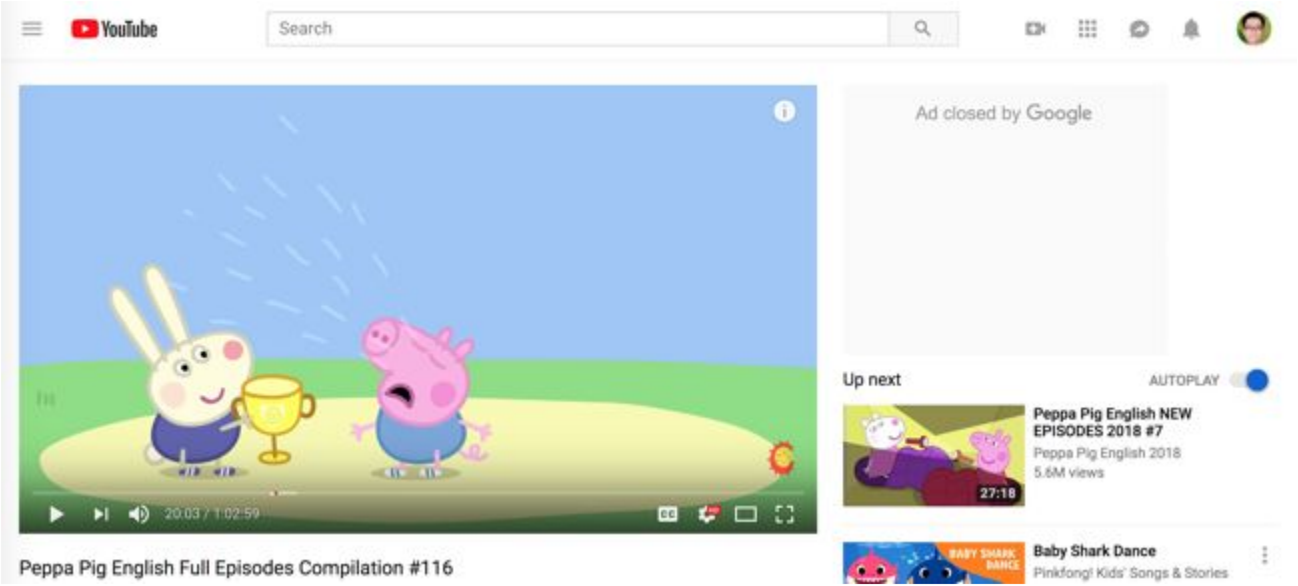

Zhao et al. *RecSys'18*

Back to home page…?

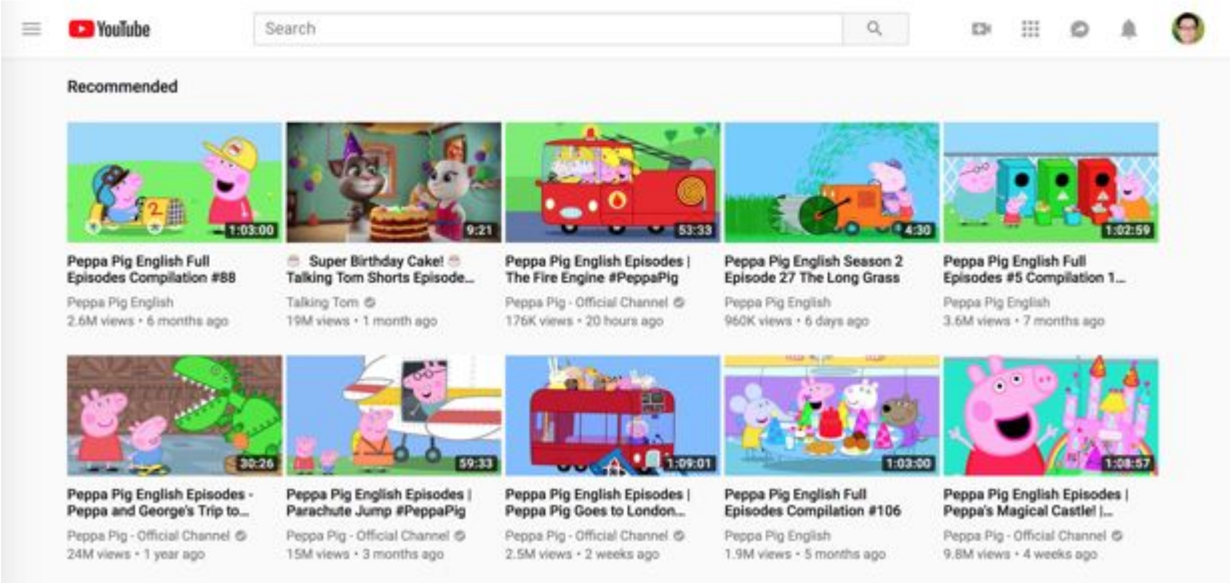

Zhao et al. *RecSys'18*

Problem

- How to model all user browsing activities jointly?
  - Temporal browsing (viewing recommendation lists and/or searching/filtering) page by page
  - Action (performing actions on recommended items, e.g., clicking, consuming) on a page
  - Inaction (neglecting or skipping recommended items) on a page

Decision Field Theory

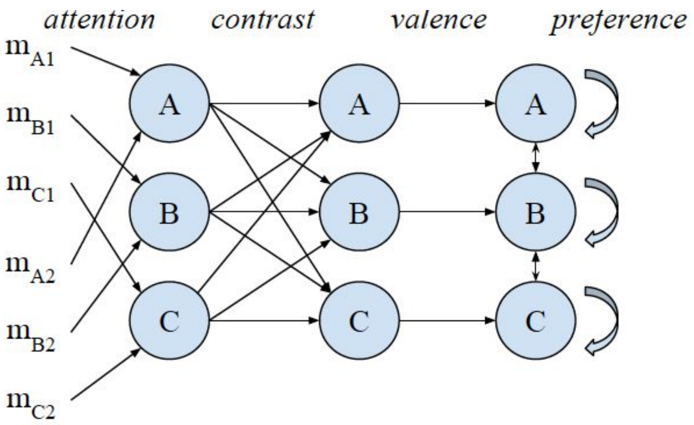

Figure 1: The connectionist network representation of DFT.

Busemeyer et al. *Psychological Review 1993*

Decision Field Theory

Preference is actually the end result of micro-level decision making processes.

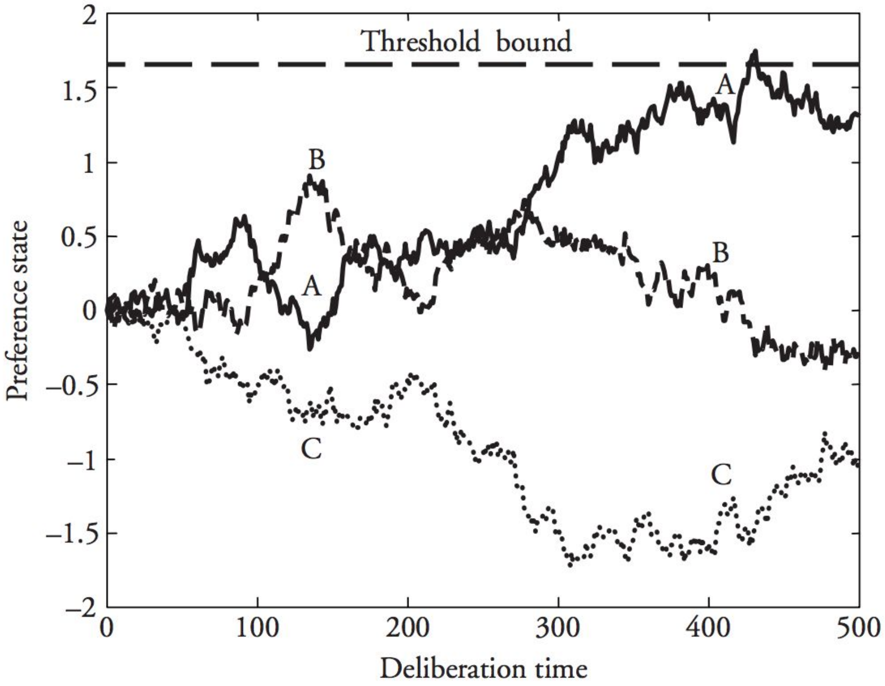

Figure 7.1 The decision process for a choice among three actions

Busemeyer *et al.* 2004

The Page-Level RNN Model

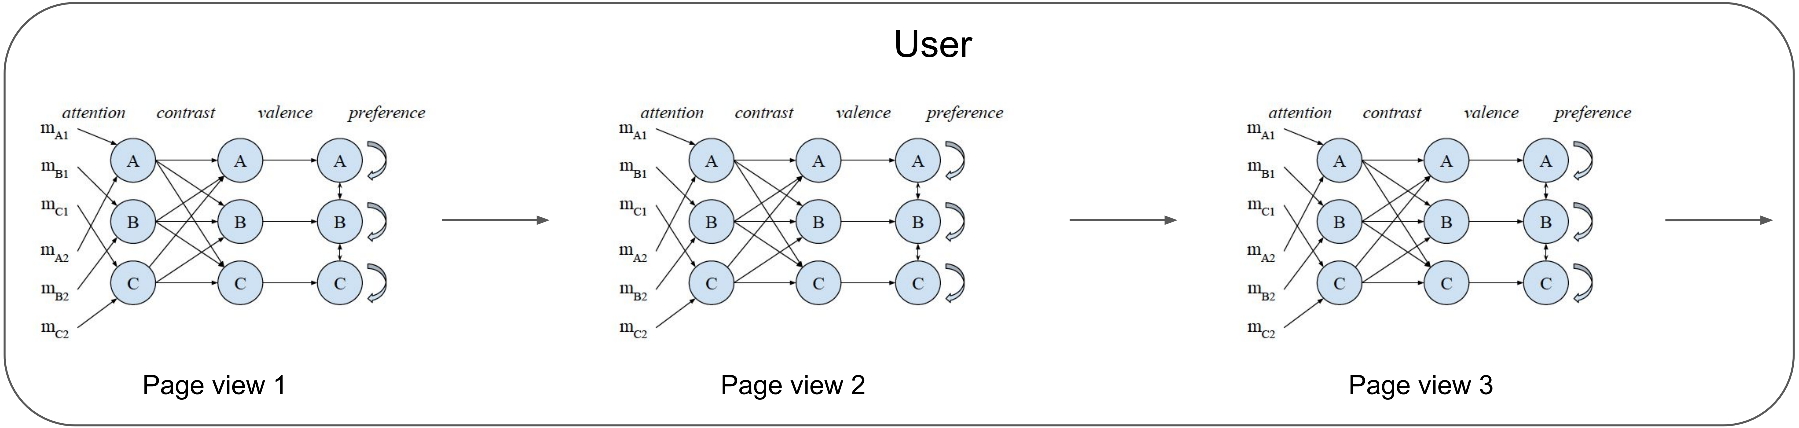

- Learn embeddings instead of using attributes
- Weights are estimated from data
- Page view sequence as the deliberation process
- There are actually 24 items because of the page size in MovieLens.
- Contrast layer is a ReLU layer
- Preference accumulation corresponds to a RNN layer

The Page-Level RNN Model

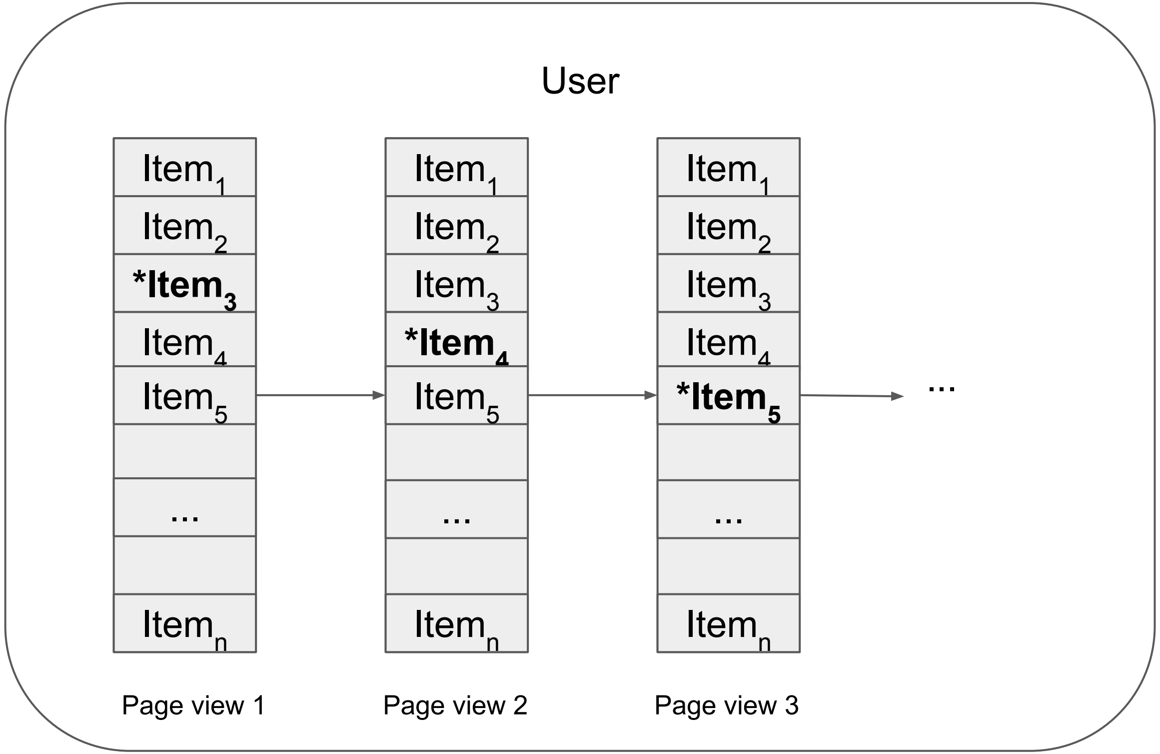

We used n=24 because of the page size in MovieLens.

The Page-Level RNN Model

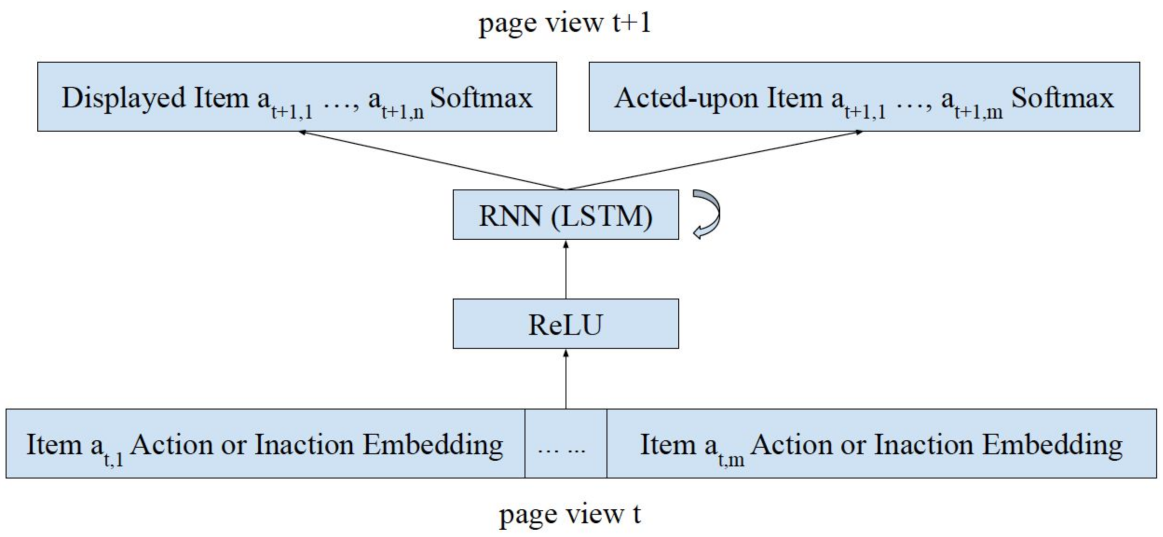

Dataset: MovieLens System Logs

- Jan. 12, 2017 to Jan. 14, 2018
- 60K movies
- 22K users
- 45M movie displays and 1.16M (positive) actions
- Temporal splitting
  - Training: Jan. 12, 2017 to Oct. 31, 2017
  - Testing: Nov. 1, 2017 to Jan. 14, 2018

Classical ***User*** & Preference Models

- SVD (Koren et al. 2009):

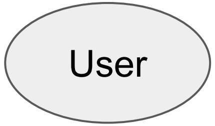

- SVD++ (Koren et al. 2010):

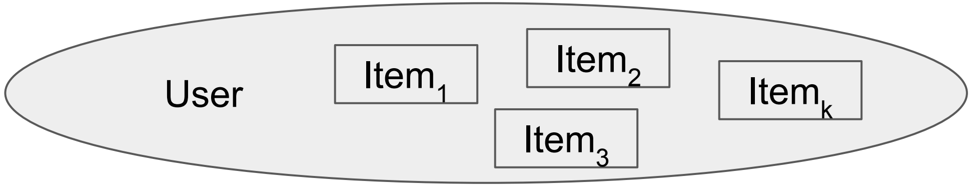

- RNN (Wu et al. 2017):

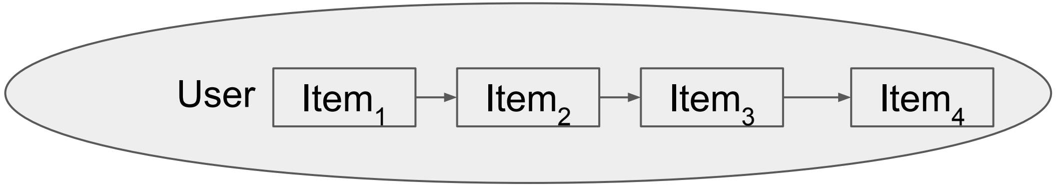

Results

User Model	MAP@8
SVD	*0.097*
SVD++	*0.106*
RNN	*0.119*
PL-RNN	***0.141***

Modeling user actions as a sequence through RNN seems to learn a better user representation compared with SVD or SVD++

Page-level **RNN gains substantially better accuracy than RNN.**

Classical User & ***Preference*** Models

- Independent binary (logistic; Hu et al. 2008): $\widehat{p(r|u,v)} = g^r(1-g)^{1-r}$

- Competitive (softmax; Yang et al. 2011, *CCF*): $\widehat{p(a|u,v)} = \widehat{p_a}(a_i = 1|u,v) = \frac{exp(f(s,o_i))}{\sum_{k=1}^{\alpha+1} exp(f(s,o_k))}$

- Relative (pairwise ranking; Rendle et al. 2009): $p(r_a > \widehat{r_b|u,v}_a, v_b) = g(f(s, o(v_a) - o(v_b)))$

“*Negative*” Items

- Negative sampling (Hu et al. 2009)

- Displayed but inaction (Yang et al. 2011; Zhao et al. 2018)

Results (MAP@8)

User Model	Preference Model	Displayed but Inaction "Negative"	Sampled "Negative"
SVD	logistic	*0.0006*	*0.104*
	softmax	*0.0006*	*0.097*
	pairwise ranking	*0.114*	*0.106*

Softmax and logistic preference models are sensitive to the choice of negative items.

Pairwise preference models can learn equally well on both types of negative items.

Messages

- Jointly modeling the three aspects (temporal browsing, action and inaction) of user-system interaction in recommender systems has benefits in
    - (this work) offline recommendation accuracy
    - (maybe, future work) online user experience
- Go from simplified preference assumptions into modeling the complex micro-level user decision making processes.

Thanks! Questions?

- Title: “From Preference into Decision Making: Modeling User Interactions in Recommender Systems”
- Authors: Qian Zhao, Martijn C. Willemsen, Gediminas Adomavicius, F. Maxwell Harper, and Joseph A. Konstan.
- Contact: qzhao2018@gmail.com

Hybrid Based on *Retrieving & Ranking*

Trained with
- Action items as positive
- Randomly sampled as negative

Trained with
- Action items as positive
- Displayed but inaction as negative

Results

User Model	Preference Model	“Negative” Items	MAP@8
SVD	softmax	sampled	*0.097*
SVD++			*0.106*
RNN			*0.119*
RNN & SVD	softmax & pairwise	sampled & displayed but inaction	*0.140*
PL-RNN	softmax	sampled & displayed but inaction	*0.141*