Risk-aware Language Models of Collaborative Filtering

Jun Wang
Department of Computer Science, University College London
Malet Place, London, WC1E 6BT, UK
Wang.Jun@acm.org

Background

Collaborative filtering is a major technique to make personalized recommendations about information items (movies, books, webpages etc) to individual users. In the literature, a common research objective is to predict unknown ratings of items for a user, on the condition that the user has explicitly rated a certain amount of items.

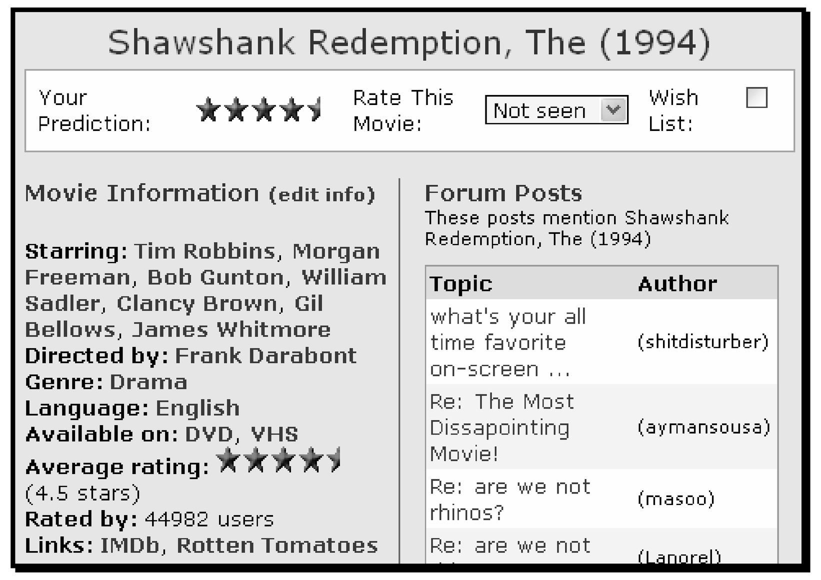

FIGURE 1: the MovieLens Recommendation Scenario

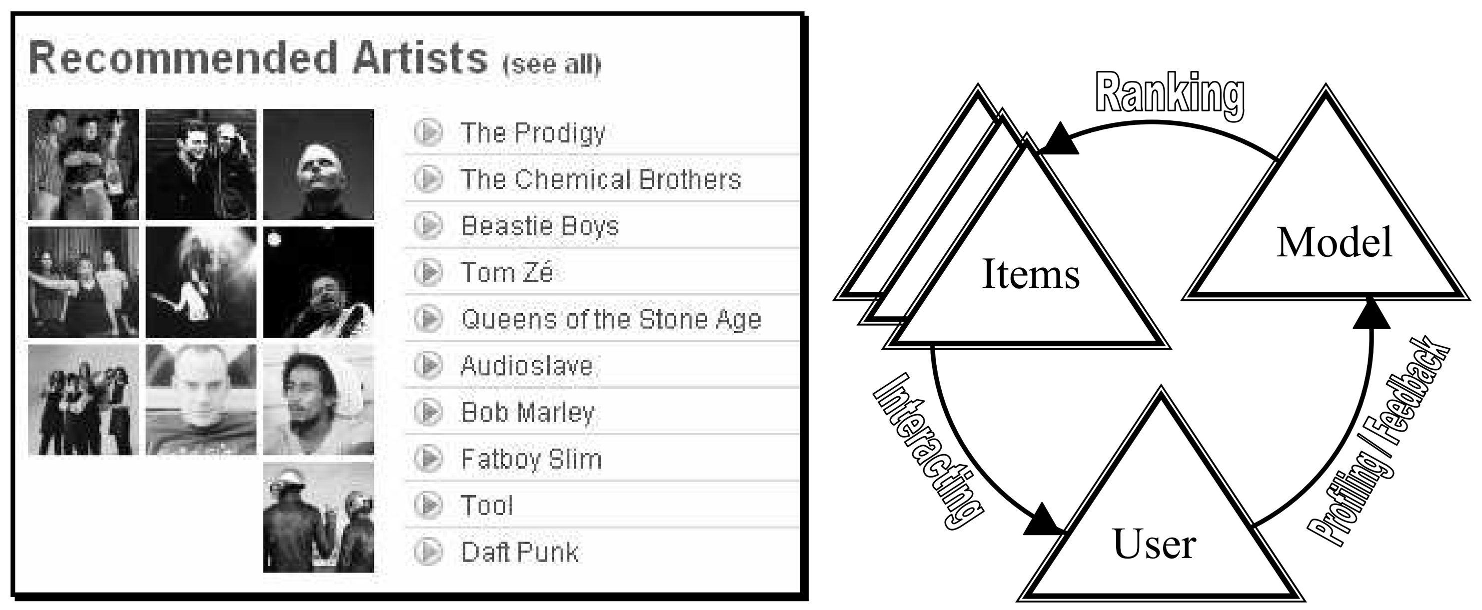

Figure 3: the Last.FM Recommendation Scenario

However, we may only have *implicit* evidence of user preferences, such as "playback times of a music file" or "visiting frequency of a web-site"; A more practical view of the recommendation task is to directly generate a top-N ranked list of items that the user is most likely to like [1].

Risk-aware Modelling

We measure how probable an item (denoted as i) is to be suggested to a given user (denoted as u): $p(i|u)$ [2]

$$o_u(i) \equiv p(i|u) \propto_i \log p(u|i) + \log p(i)$$

Two parts: its likelihood towards the user preference $p(u|i)$ and its popularity $p(i)$. Applying *linear smoothing* gives [3]:

$$o_u(i) \equiv \sum_{i' \in L_u} \ln \left(\lambda \frac{c(i',i)}{c(i)} + (1-\lambda) \frac{c(i')}{\sum_{i'} c(i')} \right) + \ln \frac{c(i)}{\sum_{i'} c(i')}$$

The ranking score of a target item i is a combination of its popularity and its co-occurrence with the items $i' \in L_u$ in the preference list of the user. $\lambda \in [0,1]$ is a linear smoothing parameter to further smooth it from a background model $(P(i'))$. $c(i',i)$ denotes the number of user preferences where both items i and i' occur, and $c(i)$ denotes the number of user preferences where item i occurs.

Alternatively, using *Bayes-smoothing* gives:

$$o_u(i) \equiv \sum_{i' \in L_u} \ln \left(\frac{c(i',i) + \mu \cdot \frac{c(i')}{\sum_{i'} c(i')}}{c(i) + \mu} \right) + \ln \frac{c(i)}{\sum_i c(i)}$$

Unreliably-estimated items may be ranked highly in the ranked list. large variance is unreliable and its rank score should be penalized [4].

$$\hat{\theta}_i \equiv Mean(\theta_i) - \frac{b}{2} Var(\theta_i)$$

where $Mean(\theta_i)$ is the mean of θ_i while $Var(\theta_i)$ denotes its variance. $b > 0$, and it is a parameter that adjusts the risk preference and can be tuned from data.

Some Experiments

Figure 4: Recommendation Precision in Last.FM. Model 1: the linear smoothing method; Model 2: risk-aware linear smoothing model ; and Model 3: risk-aware Maximum Likelihood Estimate model. A Wilcoxon signed-rank test is conducted and the significant ones (P-value < 0.05) over Model 1 the linear smoothing model are marked as *.

Observed items	Method	Parameter	Top-10	Top-5	Top-1
5	①	λ=0.9949	0.4696	0.5081	0.5988
	②	λ=0.989 b=22.9	**0.4888***	**0.5281***	**0.6279***
	③	b =130.8	0.4786	0.5243	0.5925
10	①	λ=0.99718	0.5131	0.5530	0.6029
	②	λ=0.9969 b=16	**0.5378***	**0.5938***	**0.6694***
	③	b =170.8	0.5356	0.5925	0.6590
15	①	λ=0.998363	0.5632	0.6146	0.6757
	②	λ=0.99712 b=17.01	**0.5969***	**0.6674***	**0.7526***
	③	b =190.8	0.5784	0.6482	0.7339

References

[1] J. Wang. *Relevance Models for Collaborative Filtering*. Delft University of Technology, ISBN 978-90-9022932-4, 2008. http://web4.cs.ucl.ac.uk/staff/jun.wang/papers/phdthesis.pdf.

[2] J. Wang, A. P. de Vries, and M. J. Reinders. A user-item relevance model for log-based collaborative filtering. In *Proc. of ECIR06, London, UK*, pages 37–48, Berlin, Germany, 2006. Springer Berlin / Heidelberg.

[3] C. Zhai and J. Lafferty. A study of smoothing methods for language models applied to ad hoc information retrieval. In *SIGIR '01*, 2001.

[4] J. Zhu, J. Wang, M. Taylor, and I. Cox. Risky business: Modeling and exploiting uncertainty in information retrieval. In *SIGIR09*, 2009.