

Matrix factorization techniques for recommender systems (Koren et al., 2009)

Satvik Saha (ss7228)

September 16, 2025

This paper describes a recommender system, predicting *users'* ratings for *items* (in this instance, films). By assigning each user u a latent factor $p_u \in \mathbb{R}^k$ and each item i a latent factor $q_i \in \mathbb{R}^k$, the rating r_{ui} user u has for i is modeled as approximately $q_i^\top p_u$. If all user-item pairs (u, i) were observed, or if we choose to impute missing entries, recovering the latent factors could be presented as a matrix factorization problem of the form $R = Q^\top P$. Here, with only pairs from $\mathcal{K}$ observed, we may instead minimize a squared error loss $\sum_{(u,i) \in \mathcal{K}} (r_{ui} - q_i^\top p_u)^2$. This model can be extended by adding regularization terms, including user and item biases, user and item attributes, and even temporal variations.

This has the usual topic modeling interpretation, with each item being a q_i -mixture of k topics, and each user having preferences for those topics proportional to p_u . The better aligned these vectors are, the higher the score $q_i^\top p_u$ for the pair (u, i) .

Here, user attributes Y and item attributes X have been incorporated by replacing the user latent vector p_u with $p_u + (\text{item attributes for items preferred by user } u) + (\text{user attributes})$. A probabilistic model (such as a modified version of FLAIR) might use these attributes in the following manner:

$$r_{ui} \sim f(q_i^\top p_u + \lambda_u^\top x_i + \eta_i^\top y_u),$$

introducing more latent factors λ_u, η_i , or by introducing hierarchy

$$\begin{aligned} r_{ui} \mid p_u, q_i &\sim f(\cdot \mid q_i^\top p_u), \\ p_u &\sim \pi_1(\cdot; y_u), \\ q_i &\sim \pi_2(\cdot; x_i). \end{aligned}$$

Probabilistic models like these would perhaps offer more information (say about uncertainty) than point estimates obtained by direct matrix factorization.

References

- Yehuda Koren, Robert Bell, and Chris Volinsky. Matrix factorization techniques for recommender systems. *Computer*, 42(8):30–37, 2009. doi: 10.1109/MC.2009.263.
- Lorenzo Mauri and David B. Dunson. Factor pre-training in bayesian multivariate logistic models, 2025. URL <https://arxiv.org/abs/2409.17441>.
- Sohrab Salehi, Achille Nazaret, Sohrab P Shah, and David M. Blei. Population Priors for Matrix Factorization.