Designing Recommender Systems with Reachability in Mind

Sarah Dean [1] **Mihaela Curmei** [1] **Benjamin Recht** [1]

Abstract

Access to digital content is often mediated by recommendations, which are primarily designed to predict user preferences. We take an alternate view in this work, exploring mechanisms which determine the availability of content from the perspective of user agency. We introduce a notion of *reachability* and show how it is determined by preference scoring functions, item selection rules, and user interface decisions. Simulated experiments illustrate these insights.

1. Introduction

From algorithmically sorted newsfeeds to personalized search, recommender systems mediate access to information in our digital world. The recommendations aim to surface content that will be enjoyed, highly rated, or consumed, and thus they are primarily designed to accurately predict user preferences. This singular focus has lead to concerns of unintended consequences like polarization and radicalization (Dandekar et al., 2013; Faddoul et al., 2020). The recommender systems literature has long proposed a variety of other metrics for evaluation, including concepts of novelty, serendipity, diversity, and coverage (Herlocker et al., 2004; Kaminskas & Bridge, 2016). Seldom are such metrics incorporated into the first-principles design process of the recommendation algorithms; instead, they are secondary objectives used in post-hoc approaches like reranking.

Unintended consequences arise primarily from the interactive nature of recommender systems. Many of the problems attributed to personalization online occur because of feedback effects, and thus a static perspective is insufficient for capturing important system properties. While these dynamic effects can be explored under human behavior models (Chaney et al., 2018), we favor a view based on user agency. Inspired by work on actionable recourse for consequential binary decisions (Ustun et al., 2019) and its recent extension to the recommender setting (Dean et al., 2020), we focus on understanding *reachability* within a recommender system. By considering all possible outcomes of an interaction with the recommender, reachability quantifies the influence a user has over their recommendations.

In this work, we take initial steps towards understanding how to design recommender systems that provide reachability. We introduce a stochastic version of the reachability metric, and view a recommender system as the combination of a preference model, a (possibly randomized) selection rule, and a user update. We show how each component affects reachability properties of the recommender system, and illustrate with simulated experiments.

[1]Department of Electrical Engineering and Computer Science, University of California at Berkeley, Berkeley, USA. Correspondence to: Sarah Dean <dean_sarah@berkeley.edu>.

2. Problem Setting and Definitions

Assume there are n individuals as well as a collection of k pieces of content. For consistency with the recommender systems literature, we refer to individuals as users, pieces of content as items, and expressed preferences as ratings. We will denote a rating by user u of item i as $r_{ui} \in \mathcal{R}$, where $\mathcal{R} \subseteq \mathbb{R}$ denotes the values which ratings can take. The number of *observed ratings* will generally be much smaller that the total number of possible ratings. The goal of a recommendation system is to understand the preferences of users and recommend relevant content.

2.1. Scoring functions and selection rules

In this work, we consider recommenders which are the composition of a *scoring function* ϕ with *selection rule* π. The scoring function ϕ takes as input historical data (e.g. observed ratings, user and item features) and returns a score for each user and item pair. For a given user u and item i, we denote s_{ui} to be the associated score, and for user u we will denote by $\mathbf{s}_u \in \mathbb{R}^k$ the vector of scores for all items. A common example of a scoring function arises by predicting future ratings based on historical data.

The selection rule is a policy which, for given user u and scores $\mathbf{s}_u$, selects one or more items from a set of specified target items Ω_u^t as the next recommendation. The simplest selection rule is a Top-1 policy, which is a deterministic rule that selects the item with the highest score for each user. In this work, we are primarily interested in stochastic policies which select items according to a probability distribution parametrized by scores $\mathbf{s}_u$. A stochastic selection rule recommends an item i according to $\mathbb{P}\left(\pi(\mathbf{s}_u, \Omega_u^t) = i\right)$, which is 0 for all items $i \notin \Omega_u^t$. For example, to select among items that have not yet been seen by the user, we set $\Omega_u^t = \Omega_u^c$, where Ω_u denotes the set of items seen by the user u. Note that deterministic policies are a special case of stochastic policies, with a degenerate probability distribution.

We will focus on how scores are updated after a round of user interaction. For example, if a user consumes and rates several new items, the recommender system should update the score model in response. Therefore, we parameterize the score function by an update rule, so that the new score vector is $\mathbf{s}_u^+ = \phi_u(\mathbf{a})$, where $\mathbf{a} \in \mathcal{A}_u$ represents actions taken by user u and $\mathcal{A}_u$ represents the set of all possible actions. Thus ϕ_u encodes the historical data, the preference model class, and the update algorithm. We define the form of this function in more detail in Section 3.

2.2. Reachability problem

First defined by Dean et al. (2020), an item i is *deterministically reachable* by a user u if there is some allowable modification to the user's ratings $\mathbf{r}_u$ that causes item to be recommended. Allowable modifications can include history edits, such as removing or changing ratings of previously rated items. They can also include future looking modifications which assign ratings to a subset of unseen items.

In our setting where recommendations are made stochastically, an item i is ρ *reachable* by a user u if there is some allowable action $\mathbf{a}$ such that $\mathbb{P}(\pi(\phi_u(\mathbf{a}), \Omega_u^t) = i) \geq \rho$. The maximum ρ reachability for a user-item pair is defined as the solution to the following optimization problem:

$$\rho^*(u,i) = \max_{\mathbf{a}\in\mathcal{A}_u} \; P(\pi(\phi_u(\mathbf{a}), \Omega_u^t) = i). \tag{1}$$

For example, a simple stochastic extension of the Top-1 policy is the ε-greedy policy which selects the top scoring item with probability $1-\varepsilon$, and with probability ε chooses uniformly among the remaining items. Then $\rho^*(u,i) = 1-\varepsilon$ if item i is deterministically reachable by user u, and is $\varepsilon/(|\Omega_u^t|-1)$ otherwise.

3. Solving the stochastic reachability problem

In what follows, we will show how to compute the maximum ρ reachability for user-item pairs within a system. We first specify the forms of the user action model, the score function, and the selection rule.

User action model We will consider actions that are updates to the user rating vector $\mathbf{r}_u \in \mathcal{R}^k$, a sparse vector of observed ratings. For each user, we will distinguish between *mutable* and *immutable* ratings within the vector $\mathbf{r}_u$.

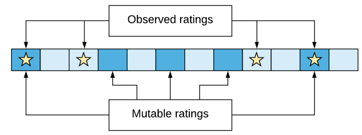

Figure 1. Mutable vs. observed ratings

Let Ω_u^m denote the set of items with mutable ratings. Then the action set $\mathcal{A}_u = \mathcal{R}^{|\Omega_u^m|}$ corresponds to changing or setting the value of these ratings. The updated rating vector $\mathbf{r}_u^+ \in \mathcal{R}^m$ is equal to $\mathbf{r}_u$ at the indices corresponding to immutable ratings and equal to the action $\mathbf{a}$ for all mutable ratings. Note the partition into mutable and immutable is distinct from earlier partition of ratings into observed and unobserved. In our general framework, mutable ratings can be both seen (history edits) and unseen (future reactions), as illustrated in Figure 1. For the reachability problem, we will consider a set of target items Ω_u^t that does not intersect with the mutable items Ω_u^m or the seen items Ω_u.

We remark that additional user or item features used for scoring and thus recommendations could be incorporated into this framework as either mutable or immutable features. The only computational difficulty arises when mutable features are discrete or categorical.

Recommender model The recommender model is comprised of a scoring function and a selection function, which we now specify. We consider *affine score update functions* where for each user, scores are determined by some affine function of the action: $\phi_u(\mathbf{a}) = B_u\mathbf{a} + \mathbf{c}_u$ where $B_u \in \mathbb{R}^{k\times|\Omega_u^m|}$ and $\mathbf{c}_u \in \mathbb{R}^k$ are model parameters determined in part by historical data. Such a scoring model arises from many different recommenders, but in this work we will primarily focus on the predictions arising from matrix factorization models.

Example 1. Matrix factorization models compute scores as rating predictions so that $S = PQ^\top$, where $P \in \mathbb{R}^{n\times d}$ and $Q \in \mathbb{R}^{k\times d}$ are respectively user and item factors for some latent dimension d. They are learned via the optimization

$$\min_{P,Q} \sum_u \sum_{i\in\Omega_u} \left\|\mathbf{p}_u^\top \mathbf{q}_i - r_{ui}\right\|_2^2 .$$

Under a stochastic gradient descent minimization scheme with step size α, the one-step update rule for a user factor is

$$\mathbf{p}_u^+ = \mathbf{p}_u - \alpha \sum_{i\in\Omega_u^m} \left(\mathbf{q}_i\mathbf{q}_i^\top \mathbf{p}_u - \mathbf{q}_i r_{ui}\right),$$

Notice that this expression is affine in the mutable ratings. Therefore, we have an affine score function:

$$\phi_u(\mathbf{a}) = Q\mathbf{p}_u^+ = Q\left(\mathbf{p}_u - \alpha Q_m^\top Q_m \mathbf{p}_u - \alpha Q_m^\top \mathbf{a}\right)$$

where we define $Q_m = Q_{\Omega_u^m} \in \mathbb{R}^{|\Omega_u^m|\times p}$.

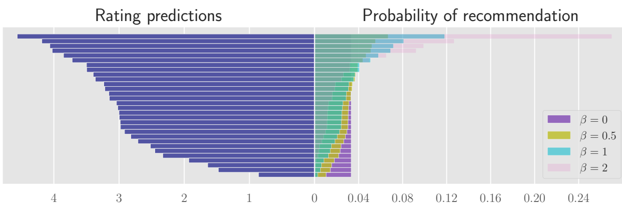

Figure 2. Example predicted ratings (left) and resulting recommendation distributions (right) under a β softmax selection rule.

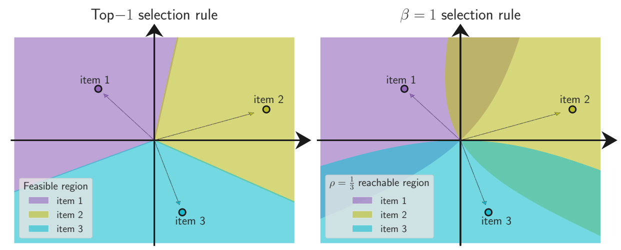

Figure 3. Feasible sets of top-1 selection rule (left) and $\rho = \frac{1}{3}$-reachable sets of $\pi_{\beta=1}$ (right). The item factors are vectors represented by the dots, the shaded regions indicate the location of user factors for which an item is feasible (left plot) or has probability larger than $\frac{1}{\texttt{num\_items}}$ to be recommended next (right plot).

We now turn to the selection component of the recommender, which translates the score $\mathbf{s}_u$ into a probability distribution over unseen items. The stochastic policy we consider is:

Definition 1. Soft-max selection
For $i \in \Omega_u^t$, the probability of item selection is given by

$$P(\pi_\beta(\mathbf{s}_u, \Omega_u^t) = i) = \frac{e^{\beta s_{ui}}}{\sum_{j \in \Omega_u^t} e^{\beta s_{uj}}} .$$

Figure 2 illustrates the distributions induced by this rule for various values of β.

Result 1. *Let $\mathbf{b}_{ui}$ denote the ith row of the action matrix B_u and consider the β-softmax selection rule. Then the stochastic reachability problem for user u and item i:*

$$\begin{aligned} \gamma^\star(u,i) = \min_{\gamma, \mathbf{a} \in \mathcal{A}_u} \quad & \gamma \\ \text{s.t.} \quad & \beta(\mathbf{b}_{ui}^\top \mathbf{a} + c_{ui}) + \gamma \\ & \geq \underset{\substack{j \in \Omega_u^t \\ j \neq i}}{\mathrm{LSE}} \left(\beta(\mathbf{b}_{uj}^\top \mathbf{a} + c_{uj})\right) . \end{aligned} \tag{2}$$

where LSE *is the log-sum-exp function. The resulting reachability parameter is $\rho^\star(u,i) = 1/(1 + e^{\gamma^\star(u,i)})$.*

Proof sketch. Consider the probability statement $P(\pi(\phi(\mathbf{a}), \Omega_u) = i) \geq \rho$. For the softmax selection rule, this is equivalent to

$$e^{\beta s_{ui}} \geq \rho \sum_{j \in \Omega_u^t} e^{\beta s_{uj}} .$$

Rearranging the expression, taking the logarithm of both sides, and substituting the affine expression for ϕ_u results in the constraint

$$\beta(\mathbf{b}_{ui}^\top \mathbf{a} + c_{ui}) + \log\left(\frac{1-\rho}{\rho}\right) \geq \underset{\substack{j \in \Omega_u^t \\ j \neq i}}{\mathrm{LSE}} \left(\beta(\mathbf{b}_{uj}^\top \mathbf{a} + c_{uj})\right)$$

The result follows by a change of variables to $\gamma = \log(\frac{1-\rho}{\rho})$, noting that maximizing ρ is equivalent to minimizing γ. □

4. Stochastic Reachability As A Relaxation

In this section, we explore how the softmax style selection rule is a relaxation of a top-1 recommendation. As illustrated in Figure 2, for larger values of β, the selection rule distribution becomes closer to the deterministic top-1 rule. This also means that the stochastic reachability problem can be viewed as a relaxation of the top-1 reachability problem. Figure 3 illustrates the regions of latent space in which items are reachable for a toy matrix factorization example with three items. We see that the region within which an item has a higher than average probability of being stochastically recommended is a geometric relaxation of the feasible set for deterministic reachability under a top-1 rule.

Result 2. *Consider the stochastic reachability problem for a β-softmax selection rule as $\beta \to \infty$. Then if an item i is top-1 reachable by user u, $\rho^*(u,i) \to 1$. In the opposite case that item i* is not *top-1 reachable, we have that $\rho^*(u,i) \to 0$.*

Proof sketch. Define

$$\tau_\beta = \beta(\mathbf{b}_{ui}^\top \mathbf{a} + c_{ui}) - \underset{\substack{j \notin \Omega_u \\ j \neq i}}{\mathrm{LSE}} \left(\beta(\mathbf{b}_{uj}^\top \mathbf{a} + c_{uj})\right)$$

and see that by the constraints of the optimization problem, $\gamma^\star = -\tau_\beta$. We can immediately notice that

$$\lim_{\beta \to \infty} \frac{1}{\beta} \tau_\beta = (\mathbf{b}_{ui}^\top \mathbf{a} + c_{ui}) - \max_{\substack{j \notin \Omega_u \\ j \neq i}} \left(\mathbf{b}_{uj}^\top \mathbf{a} + c_{uj}\right)$$

yields a top-1 constraint. If an item i is top-1 reachable for user u, then there is some $\mathbf{a}$ such that the above expression is positive. Therefore, as $\beta \to \infty$, $\gamma^\star = -\tau_\beta \to -\infty$, hence $\rho^\star \to 1$. In the opposite case when an item *is not* top-1 reachable we have that $\gamma^\star \to \infty$, hence $\rho^\star \to 0$. □

This connection leads to two interesting insights. First, for fixed $\beta \gg 1$, we see that top-1 reachability ensures that $\rho^\star$

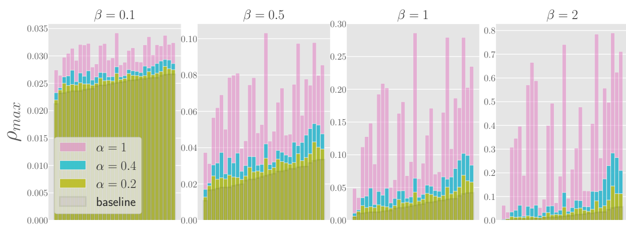

Figure 4. Values $\rho_\star$ for all target items of a user u for various values of selection parameter β and action parameter α.

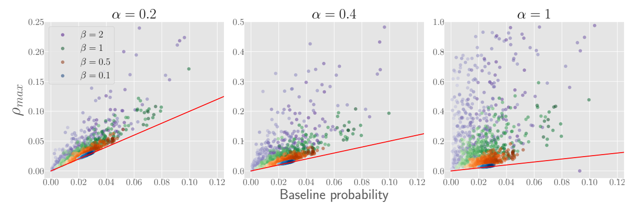

Figure 5. Each figure plots the baseline probabilities vs. $\rho_\star$ for varying β values and a fixed user action model parametrized by α. The red diagonal line represents `baseline`$= \rho_\star$. The hue of the scatter points represents the predicted rating, with darker points denoting higher predictions.

will be close to 1. Therefore, the results about model and action space geometry from Dean et al. (2020) carry over to this stochastic setting. Second, for items which are top-1 reachable, larger values of β result in larger $\rho^\star$. However, if β is too large, then items which are not top-1 reachable will have small $\rho^\star$. This illustrates a delicate balance for using randomness to ensure user agency.

5. Experiments

Setup We generate an underlying rating matrix with $n = 50$ users and $k = 100$ items and randomly select 20% of the entries to be observed. Based on the observed entries, we fit a matrix factorization model based using libFM (Rendle, 2012). We consider a range of selection rules parametrized by $\beta \in \{0.1, 0.5, 1, 2\}$. For each user u we randomly select half of the unobserved items to be target items Ω_u^t. We consider a simple user model where the mutable ratings for each user are α fraction of the unobserved items that are not target items. We vary the value of $\alpha \in \{0.2, 0.4, 1\}$, so small values correspond to restrictive (small) user action spaces and large values correspond to permissive (large) user action spaces. For each variation of the above parameters we solve the stochastic reachability problem (2) for all users and their target items using CVXPY (Diamond & Boyd, 2016). We compare against the baseline selection probabilities, which represent the default selection rates in the absence of any update to the user scores.

Results Across settings, we see that larger action spaces correspond to larger values of $\rho_\star$. Figure 4 plots $\rho_\star$ for all target items of a single user, and it shows that stochastic reachability monotonically increases in the fraction of mutable items α. Figure 4 also demonstrates that for a near uniform selection rule ($\beta = 0.1$), the reachability is minimal with respect to baseline selection probabilities. Reachability is more pronounced for skewed selection policies that prefer items with high predicted ratings. These gains are not uniform across items or baseline selection probabilities, and thus rating predictions are not indicative of how reachable or unreachable an item is in our setting.

These findings are reinforced by Figure 5. The blue points correspond to user-item pairs under a nearly uniform selection policy ($\beta = 0.1$). In this case, regardless of the user agency parameter α, users do not achieve meaningful reachability as the blue dots remain close to the diagonal red line. However, for skewed selection policies (displayed by purple points) larger values of α can dramatically increase reachability. The green and purple points corresponding to $\beta = 1$ and $\beta = 2$ respectively are much more scattered along the y-axis, suggesting that reachability can only partially be explained in terms of item selection policies and user selection models.

6. Conclusion and Discussion

In this paper, we build off work by Dean et al. (2020) to propose a generalization of item reachability which incorporates stochastic recommendation policies. We see that introducing randomness into recommendations allows for a more flexible definition of reachability, but has mixed effects on user agency. By performing this analysis from the perspective of all possible user actions, we implicitly assume that users can strategically realign their position in latent space by modifying their ratings for a set of action items. Though idealized, this setup allows us to compute an upper bound on the likelihood of an item being recommended without making assumptions about user behaviors. It can be viewed as a necessary but not sufficient coverage requirement, and is thus useful for system design and audit.

This work presents a partial characterization of the relationship between stochastic reachability, properties of observed data, item scoring functions, and item selection policies. We believe that expanding upon these is a fruitful area of both theoretical and experimental research. We also leave for future work extensions that consider more realistic user action models that factor in the cost of strategic modifications of user preferences.

References

Chaney, A. J., Stewart, B. M., and Engelhardt, B. E. How algorithmic confounding in recommendation systems increases homogeneity and decreases utility. In *Proceedings of the 12th ACM Conference on Recommender Systems*, pp. 224–232, 2018.

Dandekar, P., Goel, A., and Lee, D. Biased assimilation, homophily, and the dynamics of polarization. In *Proceedings of the National Academy of Sciences*, pp. 5791–5796, 2013.

Dean, S., Rich, S., and Recht, B. Recommendations and user agency: the reachability of collaboratively-filtered information. In *Proceedings of the 2020 Conference on Fairness, Accountability, and Transparency*, pp. 436–445, 2020.

Diamond, S. and Boyd, S. CVXPY: A Python-embedded modeling language for convex optimization. *Journal of Machine Learning Research*, 17(83):1–5, 2016.

Faddoul, M., Chaslot, G., and Farid, H. A longitudinal analysis of youtube's promotion of conspiracy videos. *arXiv preprint arXiv:2003.03318*, 2020.

Herlocker, J., Konstan, J., Terveen, L., and Riedl, J. Evaluating collaborative filtering recommender systems. *ACM transactions on information systems*, 22(1):5–53, 2004.

Kaminskas, M. and Bridge, D. Diversity, serendipity, novelty, and coverage: A survey and empirical analysis of beyond-accuracy objectives in recommender systems. *ACM transactions on interactive intelligent systems*, 7 (1), 2016.

Rendle, S. Factorization machines with libFM. *ACM Trans. Intell. Syst. Technol.*, 3(3):57:1–57:22, May 2012. ISSN 2157-6904.

Ustun, B., Spangher, A., and Liu, Y. Actionable recourse in linear classification. In *Proceedings of the Conference on Fairness, Accountability, and Transparency*, pp. 10–19, 2019.