

# Addressing Present Bias in Movie Recommender Systems and Beyond

Kai Lukoff<sup>a</sup>

<sup>a</sup>University of Washington, Seattle, WA, USA

## Abstract

Present bias leads people to choose smaller immediate rewards over larger rewards in the future. Recommender systems often reinforce present bias because they rely predominantly upon what people have done in the past to recommend what they should do in the future. How can recommender systems overcome this present bias to recommend items in ways that match with users' aspirations? Our workshop position paper presents the motivation and design for a user study to address this question in the domain of movies. We plan to ask Netflix users to rate movies that they have watched in the past for the long-term rewards that these movies provided (e.g., memorable or meaningful experiences). We will then evaluate how well long-term rewards can be predicted using existing data (e.g., movie critic ratings). We hope to receive feedback on this study design from other participants at the HUMANIZE workshop and spark conversations about ways to address present bias in recommender systems.

## Keywords

present bias, cognitive bias, algorithmic bias, recommender systems, digital wellbeing, movies

## 1. Introduction

People often select smaller immediate rewards over larger rewards in the future, a phenomenon that is known as *present bias* or *time discounting*. This applies to decisions such as what snack to eat [1, 2], how much to save for retirement [3], or which movies to watch [2]. For example, when people choose a movie to watch this evening they often choose guilty pleasures like *The Fast and The Furious*, which are enjoyable in-the-moment, but then quickly forgotten. By contrast, when they choose a movie to watch next week, they are more likely to choose films that are challenging but meaningful,

such as *Schindler's List* [2].

Recommender systems (RS), algorithmic systems that predict the preference a user would give to an item, often reinforce present bias. Today, the dominant paradigm of recommender systems is behaviorism: recommendations are selected based on behavior traces ("what users do") and they largely neglect to capture explicit preferences ("what users say") [4]. Since "what users do" reflects a present bias, RS that rely upon such actions to train their recommendations will prioritize items that offer high short-term rewards but low long-term rewards. In this way, recommender systems may reinforce what the current self wants rather than helping people reach their ideal self [5].

This position paper for the HUMANIZE workshop proposes a study design to address these topics in the domain of movies. In Study 1, a survey of Netflix users, we investigate: How should a RS make recommendations by asking ordinary users about the rather academic concept of "long-term

Joint Proceedings of the ACM IUI 2021 Workshops, April 13-17, 2021, College Station, USA

✉ kai1@uw.edu (K. Lukoff)

🌐 <https://kailukoff.com> (K. Lukoff)

🆔 0000-0001-5069-6817 (K. Lukoff)

© 2020 Copyright for this paper by its authors. Use permitted under Creative Commons License Attribution 4.0 International (CC BY 4.0).

CEUR Workshop Proceedings  
(CEUR-WS.org)



rewards"? And can long-term rewards be predicted based on existing data (e.g., movie critic ratings)? In Study 2, a participatory design exercise with a movie RS, we ask: How do users want a RS to balance short-term and long-term rewards? And what controls would users like to have over such a RS?

We expect that our eventual findings will also inform the design of recommender systems that address present bias in other domains such as news, food, and fitness.

## 2. Related Work

The social psychologist Daniel Kahneman describes people as having both an experiencing self, who prefers short-term rewards like pleasure, and a remembering self, who prefers long-term rewards like meaningful experiences [6]. Lyngs et al. describe three different approaches to the thorny question of how to measure a user's "true preferences" [5]. The first approach aligns with the experiencing self, the second with the remembering self, and the third with the wisdom of the crowd.

The first approach follows the experiencing self and asserts that *what users do is what they really want*, which many in the Silicon Valley push one step further to *what we can get users to do is what they really want* [7]. Social media that are financed by advertising are "compelled to find ways to keep users engaged for as long as possible" [8]. To achieve this, social media services often give the experiencing self exactly what it wants, knowing that it will override the preferences of the remembering self and lead the user to stay engaged for longer than they had intended.

The second approach prioritizes the remembering self, calling for systems to prompt the user to reflect on their ideal self. In this vein, Slovak et al. propose to design to help users to reflect upon how they wish

to transform their behavior [9]. Lukoff et al. previously explored how experience sampling can be used to measure how meaningful people find their interactions with smartphone apps immediately after use [10]. However, building such reflection into RSs remains a major challenge because it is unclear *how* and *when* a system should ask a user about the "long-term rewards" of an experience. It may be that the common approach of asking users to rate items on a "5-star" scale reflects a combination of short-term and long-term rewards, and that a different prompt is required to capture evaluations of long-term rewards more specifically. It is also an open question how well such long-term rewards can be inferred from existing data.

The third perspective leverages the wisdom of the crowd, by using the collective *elicited preferences* of similar users with more experience to make recommendations. Recommender systems today tend to use the "behavior of the crowd" as input into their models, in the form of behavioral data of similar users, but largely neglect elicited preferences [4].

Finally, Ekstrand and Willemsen propose participatory design as a general corrective to the behaviorist bias of recommender systems [4]. Harambam et al. explored using participatory methods to evaluate a recommender system for news, suggesting that giving users control might mitigate filter bubbles in news consumption [11]. Participatory design is a promising way to investigate how users want a RS to balance short-term and long-term rewards and the controls they would like to have.

## 3. Proposed Study Design

In what follows, we propose a study design to better understand how to measure the long-

term rewards of items in the context of movie recommendations. We hope to receive feedback on this study design from other participants at the HUMANIZE workshop and prompt conversations about ways to address present bias in recommender systems more broadly.

### 3.1. Study 1: Eliciting user preferences for long-term rewards

Our first study is a survey of Netflix users that we are currently piloting, which addresses two research questions:

- **RQ 1a:** What wording should be used to ask users to rate the “short-term rewards” and “long-term rewards” of a movie? In other words, what wording captures the right construct and makes sense to users?
- **RQ 1b:** How well can a recommender system predict the long-term rewards of a movie for an individual using data other than explicit user ratings?

#### 3.1.1. Study 1 Methods

We will ask Netflix users who watch at least one movie per month to download their past viewing history and share it with us. We will ask them to rate 30 movies: 10 watched in the past year, 10 watched 1-2 years ago, and 10 watched 2-3 years ago. Participants will rate each movie for short-term reward, long-term reward, and other constructs that might be correlated with these rewards (e.g., meaningfulness, memorability).

The current wording of our questions is:

- For short-term rewards: How rewarding was this movie **while** you were watching it?

- For long-term rewards: How rewarding was this movie **after** you watched it?

Participants will rate all questions on a 1-5 scale, from "Not at all" to "Very."

We are also interested in understanding what other constructs are correlated with the both short-term and long-term rewards. To this end, we are also asking about related constructs, such as:

- How enjoyable was this movie **while** you were watching it?
- How interesting was this movie **while** you were watching it?
- How meaningful was this movie **after** you watched it?
- How memorable was this movie **after** you watched it?

We are currently piloting all questions using a talk-aloud protocol in which participants explain their thinking to us as they complete the survey. We are checking to make sure that the wording makes sense to participants and to identify constructs that are related to short-term and long-term rewards. The constructs that are most closely related to these rewards in the piloting will be included in the final survey, so that each movie will be rated for a cluster of constructs all related to short-term and long-term rewards. For the final survey, we plan to recruit about 50 Netflix users to generate a total of 1,500 movie ratings.

#### 3.1.2. Study 1 Planned Analysis

To address **RQ 1a**, we will report the qualitative results of our talk-aloud piloting of the survey wording. We will also report the correlation between how participants rated

our measures of short-term and long-term rewards and related constructs (e.g., meaningfulness, memorability).

To address **RQ 1b**, first we will test how well existing data correlates with long-term rewards. The existing data we plan to test includes: user ratings (from others), critic ratings, box office earnings, genre, and the day of the week the movie was watched. One notable limitation of our study design here is that we will not have access to behavioral data like clicks, views, or time spent about movies.

Second, we will create machine learning models to test how well we can predict the long-term rewards of a movie might provide for an individual, as assessed by metrics like precision and recall. Specifically, we will create:

- A generalized model that makes the same predictions for each movie for all participants;
- A personalized model that makes individualized predictions for each movie for each participant;

Finally, we will also create generalized and personalized models that predict short-term rewards too, and compare their performance against the models that predict long-term rewards. Our suspicion is that existing data may be more predictive of short-term rewards than long-term rewards, because long-term rewards may require a form of reflection that most existing data do not capture.

### **3.2. Study 2: Addressing present bias via user control mechanisms**

Study 2 is currently planned as a participatory design exercise with a movie recommender system, in which we ask:

- **RQ 2a:** What preferences do users have for how a movie RS should weight the short-term versus long-term rewards of the movies it recommends? In what contexts would users prefer what weights?
- **RQ 2b:** How would users like to control how a movie RS weighs the short-term versus long-term rewards of the movies it recommends?

At the workshop, we hope to elicit feedback on how Study 2 might be revised to best answer our research questions.

#### **3.2.1. Study 2 Methods**

Our exercise will begin by showing participants a set of recommended movies that is heavily weighted towards short-term rewards, based on the ratings that we obtained in Study 1. Then we will show them a set of recommendations that is heavily weighted towards long-term rewards. We will ask participants to describe which recommendations they would prefer and why. We will also ask about which contexts (e.g., mood, day of the week) affect which types of rewards they would prefer.

Next, we will solicit feedback from users on a paper prototype of a RS that offers users control over recommendations at the input, process, and output stages. For instance, at the input stage, users could indicate their general preferences for short-term versus long-term rewards. At the process stage, users could choose from different "algorithmic personas" to filter their recommendations, e.g., "the guilty pleasure watcher" or "the classic movie snob." At the output stage, users might control the order in which recommendations are sorted. This exercise draws from the study design in Harambam et al., in which users participants evaluated and described the control mechanisms they

would like to have based on a prototype of a news recommender system, with a focus on addressing the bias of filter bubbles [11].

## 4. Workshop relevance

Today’s recommender systems often prioritize “what users do” and neglect “what users say.” As a result, they tend to reinforce the current self rather than foster the ideal self. This study design proposes to study this with regards to movies. But the same problem also applies to other domains such as online groceries, where the current self might want cookies while the ideal self wants blueberries, or digital news, where the current self might want to read stories that agree with their worldview while the ideal self wants to be challenged by different perspectives [12]. The methods we propose in this study design are relevant beyond just movies.

We expect that all workshop participants will benefit from a lively discussion of how to conceptualize and measure user preferences in ways that go beyond the current behaviorist paradigm of prioritizing what users do over explicit preferences. Our proposal to ask users for their explicit ratings and correlate these with other data is just one possible approach, and we would like to discuss what other methods workshop participants would suggest and how well these apply to other domains such as groceries and news. Addressing present bias also raises philosophical issues: Is it always irrational to pursue short-term rewards over long-term rewards? Are users in a position to judge their own long-term rewards? How far should computing systems go in nudging or shoving users towards long-term rewards?

Finally, present bias is just one of many cognitive biases. We hope that our submission will also contribute to the growing conversation on how to use psychological theory

to address cognitive biases in intelligent user interfaces [13].

## Acknowledgments

Thank you to Minkyong Kim, Ulrik Lyngs, David McDonald, Sean Munson, and Alexis Hiniker for feedback on this research agenda and drafts of the position paper.

## References

- [1] M. K. Lee, S. Kiesler, J. Forlizzi, Mining behavioral economics to design persuasive technology for healthy choices, in: Proceedings of the SIGCHI Conference on Human Factors in Computing Systems, CHI '11, ACM, New York, NY, USA, 2011, pp. 325–334.
- [2] D. Read, van Leeuwen B, Predicting hunger: The effects of appetite and delay on choice, *Organ. Behav. Hum. Decis. Process.* 76 (1998) 189–205.
- [3] H. E. Hershfield, D. G. Goldstein, W. F. Sharpe, J. Fox, L. Yeykelis, L. L. Carstensen, J. N. Bailenson, Increasing saving behavior through age-progressed renderings of the future self, *J. Mark. Res.* 48 (2011) S23–S37.
- [4] M. D. Ekstrand, M. C. Willemsen, Behaviorism is not enough: Better recommendations through listening to users, in: Proceedings of the 10th ACM Conference on Recommender Systems, RecSys '16, ACM, New York, NY, USA, 2016, pp. 221–224.
- [5] U. Lyngs, R. Binns, M. Van Kleek, others, So, tell me what users want, what they really, really want!, *Extended Abstracts of the (2018)*.
- [6] D. Kahneman, *Thinking, Fast and Slow*, Farrar, Straus and Giroux, 2011.
- [7] N. Seaver, *Captivating algorithms: Rec-*

- ommender systems as traps, *Journal of Material Culture* (2018). URL: <https://doi.org/10.1177/1359183518820366>.
- [8] T. Dingler, B. Tag, E. Karapanos, K. Kise, A. Dengel, Detection and design for cognitive biases in people and computing systems, <http://critical-media.org/cobi/background.html>, 2020.
  - [9] P. Slovák, C. Frauenberger, G. Fitzpatrick, Reflective practicum: A framework of sensitising concepts to design for transformative reflection, in: *Proceedings of the 2017 CHI Conference on Human Factors in Computing Systems*, dl.acm.org, 2017, pp. 2696–2707.
  - [10] K. Lukoff, C. Yu, J. Kientz, A. Hiniker, What makes smartphone use meaningful or meaningless?, *Proc. ACM Interact. Mob. Wearable Ubiquitous Technol.* 2 (2018) 22:1–22:26.
  - [11] J. Harambam, D. Bountouridis, M. Makhortykh, J. Van Hoboken, Designing for the better by taking users into account: a qualitative evaluation of user control mechanisms in (news) recommender systems, in: *Proceedings of the 13th ACM Conference on Recommender Systems*, dl.acm.org, 2019, pp. 69–77.
  - [12] S. A. Munson, P. Resnick, Presenting diverse political opinions: how and how much, in: *Proceedings of the SIGCHI conference on human factors in computing systems*, dl.acm.org, 2010, pp. 1457–1466.
  - [13] D. Wang, Q. Yang, A. Abdul, B. Y. Lim, Designing theory-driven user-centric explainable AI, of the 2019 CHI Conference on ... (2019). URL: <https://dl.acm.org/doi/abs/10.1145/3290605.3300831>.