

# NEST: Simulating Pandemic-like Events for Collaborative Filtering by Modeling User Needs Evolution

Chenglong Ma  
chenglong.ma@student.rmit.edu.au  
RMIT University  
Melbourne, Australia

Pablo Castells  
pablo.castells@uam.es  
Universidad Autónoma de Madrid  
Madrid, Spain

Yongli Ren  
yongli.ren@rmit.edu.au  
RMIT University  
Melbourne, Australia

Mark Sanderson  
mark.sanderson@rmit.edu.au  
RMIT University  
Melbourne, Australia

## ABSTRACT

We outline a simulation-based study of the effect rapid population-scale concept drifts have on Collaborative Filtering (CF) models. We create a framework for analyzing the effects of macro-trends in population dynamics on the behavior of such models. Our framework characterizes population-scale concept drifts in item preferences and provides a lens to understand the influence events, such as a pandemic, have on CF models. Our experimental results show the initial impact on CF performance at the initial stage of such events, followed by an aggravated population herding effect during the event. The herding introduces a popularity bias that may benefit affected users, but which comes at the expense of a normal user experience. We propose an adaptive ensemble method that can effectively apply optimal algorithms to cope with the change brought about by different stages of the event.

## CCS CONCEPTS

• Information systems → Recommender systems.

## KEYWORDS

Collaborative Filtering, pandemic-like events, simulation, human needs, herding behavior, popularity bias, concept drift

### ACM Reference Format:

Chenglong Ma, Yongli Ren, Pablo Castells, and Mark Sanderson. 2022. NEST: Simulating Pandemic-like Events for Collaborative Filtering by Modeling User Needs Evolution. In *Proceedings of the 31st ACM International Conference on Information and Knowledge Management (CIKM '22)*, October 17–21, 2022, Atlanta, GA, USA. ACM, New York, NY, USA, 11 pages. <https://doi.org/10.1145/3511808.3557407>

## 1 INTRODUCTION

People react to changes in the conditions of their life [25]. In response to certain shifts, individuals may alter their purchasing

Permission to make digital or hard copies of all or part of this work for personal or classroom use is granted without fee provided that copies are not made or distributed for profit or commercial advantage and that copies bear this notice and the full citation on the first page. Copyrights for components of this work owned by others than ACM must be honored. Abstracting with credit is permitted. To copy otherwise, or republish, to post on servers or to redistribute to lists, requires prior specific permission and/or a fee. Request permissions from [permissions@acm.org](mailto:permissions@acm.org).  
*CIKM '22, October 17–21, 2022, Atlanta, GA, USA*

© 2022 Association for Computing Machinery.  
ACM ISBN 978-1-4503-9236-5/22/10...\$15.00  
<https://doi.org/10.1145/3511808.3557407>

habits, which could be copied by others through positive feedback [29, 69]. Examples of such changes include the stockpiling of essentials during uncertain times [49].

Since Recommender Systems (RSs) – CF in particular – draw on behavioral data, changes in the external environment may re-frame human behavior impacting the performance of CF because the altered behavior may not match what the system has learned up to that point. Typically, CF models involve two phases:

- *knowledge learning*: the models learn user traits from past interaction data;
- *recommendation generation*: the models yield a subset of items as recommendations to match users' predicted tastes.

Change in human behavior in the context of RSs is referred to as concept drift [17, 36]. While such drift is well studied, population-scale drifts as occurred during the COVID pandemic [13, 61] are less understood. Many people changed their routine, buying different products, or volumes of products, challenging the agility of commercial supply chains [32]. Such changes can disrupt a CF model's ability to identify meta-users (i.e., coherent user profiles during training). These changes require remodeling of the dynamics of intra- and inter-user analysis over time. We examine the impact of population-scale drifts on CF models. As it is difficult to find datasets reflecting such drifts, therefore, we take a simulator-based framework that we call **Need Evolution SimulaTor** (NEST).

Based on Maslow's theory of user needs [43, 44, 64], we extend a Markov chain model [11] with a hidden layer to model both users' observable preferences and inner need states. We embed the model into a reinforcement learning (RL) based OpenAI Gym framework [4, 30] to simulate individuals' behavior at different stages of a population-scale event. NEST disentangles CF training from the discrepancy of user preferences by modularizing various CF agents. By setting different types of "what if" scenarios, NEST creates a sandbox environment to provide analysis and suggestions for creators of RSs. We focus our evaluation on CF models within e-commerce domains, i.e., users' purchasing behavior. While such a simulator cannot mirror user behavior with full fidelity, our goal is to replicate psychological need-driven user behaviors found in real RSs and to evaluate their macro-trend impact on RSs in a parameterized interacting environment with a set of counterfactual experiments<sup>1</sup>. The contributions of the paper are as follows:

<sup>1</sup>Code is available at <https://github.com/ChenglongMa/NEST>

- An RL-based interaction model NEST for simulating user need and behavior dynamics in a non-stationary context. The model connects user needs, behavior, and RSs uniformly.
- A study of the multi-faceted impact of pandemic-like events on classic CF models, including how models cope with such changes based on a set of counterfactual experiments.
- An adaptive ensemble method to optimize the overall effectiveness of the CF models assembly, which maximizes the adaptability of RSs in a non-stationary environment.

## 2 RELATED WORK

We consider “pandemic-like” events as events occurring unpredictably by ordinary people, that have substantial impact on everyday life, and result in population-scale effects on behavior. While major events such as the pandemics of 1918 [28] and 2019 [68] are examples of such events, our focus goes beyond such “once-in-a-century” occurrences to events that occur several times in a life span, that have a significant impact on people’s behavior patterns. For example, events that are statistically fat-tailed, such as financial crises [45, 65], are often unpredictable but have a broad and deep impact on global population with a great likelihood of occurrence.

Many researchers have studied the impact of the recent pandemic on people’s needs and behaviors, e.g., diet, psychology, work, entertainment. Di Renzo et al. [15] find that people feel more stress and alter their habitual behaviors due to the fear of the virus. Yuen et al. [69] identify psychological causes and social influence of panic buying. Fink et al. [16] notice the popularity of “coronamusical” which is widely used as solitary emotional regulation to meet people’s socio-emotional needs.

### 2.1 Human Needs and Behavior

Maslow developed a hierarchy of 5-levels of human needs [43, 44] classified into **Physiological** (Ph), **Safety** (Sf), **Love & Belongingness** (LB), **Cognitive** (Co) and **Self-Actualization** needs (SA)<sup>2</sup>. Behavior is motivated to achieve certain such needs; once the need is satisfied, it becomes no longer active, until it is, eventually, brought back into active state. Maslow argues that lower level needs must be met prior to higher needs. However, extensive research [e.g., 11, 61] show that the movement of needs is flexible and multi-directional, people can better satisfy lower-level needs by improved lifestyle [11]. Alderfer [2] argues that multiple needs may be active synchronously and if the satisfaction of higher-level needs is suppressed, the desire for lower-level needs will become intense.

The theory of social identity [62] suggests that individuals are inclined to recognize themselves as belonging to a social group, wherein they share similar values and preferences with other members [29]. Chung [11] maps the 5-level model of needs to five states in a latent space, in which the movements among needs are multi-directional. We extend Chung’s idea to a hidden Markov Model (HMM) that further considers the probability of specific emitted actions from people. In the context of e-commerce, for instance, actions are user purchases, with probabilities driven by user preferences on different products.

Changes in behavior can impair the performance of systems, such as search engines, RSs, and decision-making. Suh et al. [61]

<sup>2</sup>There are different definitions of need at each level, we use those from Suh et al. [61].

carry out a population-scale study in search systems across all levels of Maslow’s needs during the recent pandemic, finding that expression of basic needs (e.g., food, water, shelter) exponentially increased, while higher-level needs (e.g., fashion, education) relatively declined, thus describing population-scale drifts with data-backed evidence. Škare et al. [58] measure the impact of the COVID-19 pandemic on tourism systems, emphasizing the sharp decline in the tourism economy and the slower than expected recovery. Gu et al. [20] investigate the impact of the pandemic on purchasing behavior in e-commerce systems finding the pandemic impels a constancy and promptness of online buying behavior.

### 2.2 Collaborative Filtering

We distinguish two types of learning in CF: offline CF and online CF. Offline CF (e.g.,  $k$  nearest neighbor ( $k$ NN) and matrix factorization (MF)) learns from “static” training data, which must be fully accessible at a *knowledge learning* phase. Only then can the subsequent *recommendation generation* proceed into predicting good choices to be suggested to users [1, 17]. In contrast, online CF treats historical user interaction data as sequences or streams. It does not require complete training data at the *knowledge learning* phase; instead, it constantly updates models as new observations arrive and performs *recommendation generation* asynchronously [1, 17].


The changes caused by pandemic-like events not only affect people’s behavior, but also affect the performance of CF models. Compared to online CF, conventional offline CF models often fail to capture the drift of interest from users [51]. Although some offline algorithms are developed to incorporate temporal information in data [e.g., 36], they are incapable of fitting unseen sudden changes in user behavior until retrained on new observations.

### 2.3 Temporal Dynamics in CF

Changes in user behavior are a form of *concept drift* [66, 67], which introduces new challenges to temporal dynamic modeling. Koren [36] identifies different paradigms for temporal effects in user dynamics, and proposes a variant of SVD++ (TimeSVD) to capture both persistent and transient signals in such behavior, balancing long and short-term drifts. This approach does not consider item-specific drift and suffers from cold start and data sparsity problems. Rendle et al. [51] combine MF and Markov chain (MC) models to build a personalized transition graph for prediction. The method can learn the sequential evolution in user behavior by an extension of BPR [50]. He and McAuley [22] model dynamic user preferences for visual factors of fashion products based on a convolutional neural network. Kang and McAuley [34] balance long-term vs. short-term preferences and sparse vs. dense data by incorporating MCs and recurrent neural networks (RNNs) in a self-attentive sequential model. Recent research [e.g., 18, 39] considers temporal drift by analyzing similarities among user trajectories (i.e., modeling users’ dynamic portraits in *compatible* long/short-term preferences).

### 2.4 Simulation in Recommendation

The sudden and unexpected changes arising from pandemic-like events are rarely studied. People’s abrupt adaptive behavior is less relevant to their past actions. At different development stages of the event, individuals may show disparate behavior patterns. Shifts



**Figure 1: a) Illustration of interactions between needs, behavior and CF models. b) Sequential interactions between needs, behavior and the CF models in 5 bottom-up layers.**

occur at the level of the entire population, where individuals imitate and influence each other rather than moving in isolation. Such events raise difficulties for researchers. The rarity of such events results in a scarcity of data, making the collection of behavioral patterns nearly infeasible. We therefore consider simulation as a feasible approach to overcome these difficulties.

Simulation is common in RSs. Any offline recommendation experiment is a simulation of a production RS; the better the experiment, the more useful the observations will be. Additional layers of simulation are often added. For instance, null hypotheses can be simulated by randomization of variables of interest, e.g., by randomly swapping rating values between user-item pairs in the data, we may simulate probabilistic independence between the rating count and the average rating value of items [5, 7, 46]. The depiction of RSs as a one-shot recommendation round captures only a highly partial side of reality. The field is moving beyond this limited view, to an understanding of recommendation as a continued cyclic relationship between RSs and users, where feedback at each step adds to the inputs available to the RS in the next step [52].

Ciampaglia et al. [12] investigate the impact of social conformity on the quality of user choices in a simulation where users, their tastes, attention, decisions, are synthetic along with items and item quality. Based on a variation of multi-armed bandits (MAB), Schmit and Riquelme [55] propose a model to simulate users' interaction with RSs, evaluating RSs with respect to the *feedback loop*. Based on this, Chaney et al. [10] examine how feedback loops in RSs cause high homogeneity and low utility of users. However, both of them consider static user preference only and ignore change over time. Recently, Rohde et al. [52] and Ie et al. [30] proposed RL-based simulation environments (i.e., RecoGym, RecSim) wrapped in the novel OpenAI Gym framework [4], which facilitates the RL algorithm design and analysis of user behaviors.

In this research, we consider the reverse effect described above, where macro-trends in population dynamics influence the reactions of RSs and the qualities of recommendations.

### 3 TERMINOLOGY AND DEFINITIONS

We start from the relationship among user needs, user behaviors, and user interactions with CF models. As shown in Fig. 1a, users' latent need state is an independent variable. The observable user behavior (e.g., purchasing and feedback) acts as a mediating variable while the recommendations of CF models are the dependent variable. We deduce our conclusion from the premise that users' needs, behaviors, and preferences on items are strongly correlated, as RSs can only estimate preferences of users based on the behavior they explicitly exhibited. Fig. 1b is a detailed example of Fig. 1a. During the pre-event stage, the CF model can fit users' preferences effectively. However, the occurrence of a sudden unexpected event first disrupts the latent need model, and a degrading cascade reaction follows in the interactions between users and the RS. We explain this in greater detail in the next subsections.

#### 3.1 User Needs and Preferences

Given a RS, there are a finite user set  $\mathcal{U}$  and an item set  $\mathcal{I}$ . The average user  $u \in \mathcal{U}$  will explore their surroundings. In the context of RSs, this can be characterized as:

- (1) *human need system*  $\mathcal{N}$ , a latent discrete space of needs states, that we represent using Maslow's 5-level needs, i.e.,  $\mathcal{N} \subseteq \{Ph, Sf, LB, Co, SA\}$ . The system evolves as a holism, in which a new need  $\eta_t \in \mathcal{N}$  is active or dormant depending on the previously satisfied need  $\eta_{t-1}$ , i.e., need  $\eta \sim \Pr(\eta_t | \eta_{t-1})$ ;
- (2) *behavior pattern*  $\mathcal{P}$ , an action space representing preferences from users in  $\mathcal{U}$  for items in  $\mathcal{I}$ . A *motivation policy*  $\pi_m : \mathcal{N} \times (\mathcal{U}, \mathcal{I}, t) \rightarrow \mathcal{P}$  maps a need state to a prior distribution of preferences for items in  $\mathcal{I}$  at time  $t$ . Then the next observable preference will be conditional on the current latent need state, i.e., preference  $p \sim \Pr(p_t | \eta_t)$ .

The private preferences of users are their *inner* incentives of decision-making while their circumstances and the opinions of others constitute *external* incentives. Furthermore, the expressions of users' private preferences in the form of recommendations will in turn affect others' decisions via a feedback loop.

#### 3.2 User Behavior

A user  $u$  reveals his/her preferences  $\mathcal{P}_u$  to a RS through implicit (e.g., view, purchasing) and explicit (e.g., rating) actions. Meanwhile, the system perceives and learns their actual latent needs. The system verifies its assumptions by presenting slates of personalized recommendations and observing users' feedback on the slates. Let  $c_t \in \mathcal{I}$  be a recommended item consumed (purchased, clicked, etc.) by user  $u$  at time  $t$ , and let  $\vec{H}(t) = [h_1, h_2, \dots, h_t]$ , with  $h_t = (u_t, \mathcal{P}_t, c_t)$ , be a history of user engagement with recommended items (note that need states  $\mathcal{N}$  are unobservable to RSs). A *recommendation policy*  $\pi_r : \vec{H}(t-1) \times (\mathcal{U}, \pi_m, t) \rightarrow \mathcal{S}_t$  maps a history of *past* preferences to a slate of recommended items [30], where  $\mathcal{S}_t \sim \Pr(\mathcal{S}_t | \mathcal{I}, \vec{H}(t-1))$ .

The behaviors in the logged data can be viewed as a signal of users' private preference subjoining their implicit perception of product quality. And the quality of a certain item can be the convergence of others' beliefs on its attributes [31, 55]. This coincides with the findings of Smith and Sørensen [59] that individuals make

decisions subject to their endowed characteristics and observation of others' behaviors. Technically, the resulting feedback of such behavior is a combination of an individual preference  $p_t$  and the perceived quality of the selected item  $q_t$  at time  $t$ , i.e.,  $r_t = f(p_t, q_t)$ , for a certain function  $f$  – in our experiments, we take the convex combination  $f(p_t, q_t) = p_t + q_t + \epsilon_t$ , where  $\epsilon_t$  represents random noise, but this is a modular choice in our framework. Note that, as we focus on the changes from users' perspective, we assume the quality of items does not change over time, i.e.,  $q_t = q$ .

### 3.3 User-CF Interactions

Suh et al. [61] identify boundaries between changes in the development of the COVID-19 pandemic. Inspired by this, we identify three stages in the development of a pandemic-like event: **pre-event** to represent the normal days prior to the event; **during-event** to represent the time range from the beginning to the end of the event; **post-event** to represent the period after the event subsides. The setting can be further subdivided into more stages to depict different phases of an event.

Fig. 1b illustrates with a toy example how the policies  $\pi_m$  and  $\pi_r$  operate over time. The *Events* layer describes the stages in the evolution of the outlier event. Correspondingly, the *Need States* layer shows people's dominant needs at different event stages, i.e., growth-based needs (e.g., Co, SA needs) dominate the pre-event stage while they yield to basic needs (e.g., Ph, Sf needs) to react to the surprise situation. The *Preferences* layer represents the preferred items for the user to meet his/her needs shown in the *Need States* layer driven by policy  $\pi_m$ . The *Recommendations* layer at the top demonstrates the slate of recommendations generated by policy  $\pi_r$ . Finally, back to the *Choice* layer, it shows that the user makes the final decision combining the recommendations and his/her own preference. Therefore, the estimated recommendation score  $\hat{r}_{uit}$  of item  $i$  at time  $t$  is a real-valued function of  $\{(r_{ui}, t) : t \in T\}$  based on the policy  $\pi_r$ ;  $r_{ui}$  is the integrated rating (e.g., feedback) of user  $u$  on item  $i$ . Finally, the decision  $c_t$  of the active user at time  $t$  is drawn from the probabilistic combination of external recommendation scores and the inner dynamic user preference, i.e.,<sup>3</sup>

$$c_t \sim \text{Pr}(\hat{r}_t + p_t). \tag{1}$$

User decisions and feedback further participate in the cycle of the *recommendation generation*, and their need states  $\mathcal{N}$  update once they complete the transaction and gain the utilities accordingly.

If the output of policy  $\pi_m$  is stable or only mildly drifting, policy  $\pi_r$  will generate expected recommendations. Otherwise, if users are in a high-velocity  $\mathcal{N}$  space, the RS may not keep up with the speed of change as it can only observe the past behavior that the user explicitly exhibited. In that case, recommendations will lag behind the user's current need state. During a pandemic-like event, this deviation is more obvious and thus the performance of CF models can be heavily impacted.

## 4 METHODOLOGY

As there are no representative behavior samples in existing datasets (e.g., MovieLens [21], Amazon reviews [48]), we propose the NEST model based on the OpenAI Gym framework [4], which supports

<sup>3</sup>We may omit the subscripts of  $u$  and  $i$  for readability when there is no ambiguity.




Figure 2: Modules in NEST, including an item model (blue), a need-HMM (green), a user-choice model (yellow), a user-feedback model (orange) and a recommender model (gray).




Figure 3: The illustration of need-HMM. The yellow trajectory shows that one's need shifts from SA to Sf at time  $t_1$  and thus the sanitizer becomes his/her preferred item.

sequential evaluation for online CF models and synthetic data generation for offline CF models. Our models represent the foregoing interaction problem within an RL environment. The main components of our model are shown in Fig. 2. The pluggable *recommender model* supports various Gym compatible RL-based agents. It interacts with the other modules by learning the observable user preferences and the item quality, recommending slates of items to users.

### 4.1 Item Model

The item model samples items drawn from a prior distribution characterizing their qualities. The quality feature can be interpreted as a fusion of consumer reviews for an item [31]. We assume the recommender agents can collect the quality information from statistics, which is thus observable to the agents and available for recommendation generation. In NEST, the quality is modeled and drawn from a normal distribution, i.e.,  $q_i \sim N(\mu_i, \sigma_i^2)$ . The number of items  $n = |Z|$  is also a parameter of the item model. The item model parameters are thus a triplet  $(\mu_i, \sigma_i, n)$ . Particularly, we set  $\mu_i = 0$  and  $\sigma_i = 1$  in our experiments.

### 4.2 Need-HMM

Based on the work of Chung [11], we propose a user need transition model (need-HMM) to describe the evolution of individuals'

needs and the dynamic preferences given certain need state. Need-HMM is a time-inhomogeneous hidden Markov process that generates a set of discrete probabilistic preferences on recommended items given the user's hidden need state at each time point. It is characterized by a tuple  $(N, \mathcal{S}, A, E)$  where, again,  $N$  is the hidden need system;  $\mathcal{S}$  is a slate of recommended items,  $\mathcal{S} \subseteq \mathcal{I}$ ;  $A$  is a transition probability matrix describing the movements of the need system. The probabilities may vary over time;  $E$  is an emission probability matrix representing the dynamic preferences of a user on slate  $\mathcal{S}$ .

*Human Need System.* Need-HMM treats the human need system  $N$  as a probabilistic holism consisting of different levels of need states,  $\eta \in N \subseteq \{Ph, Sf, LB, Co, SA\}$ , which are unobservable to recommender agents. According to the representative theories in behavioral sciences [2, 43], the evolution within the need system are holistic and changes in any state can cause concomitant reactions in other states. For example, once a user's basic needs are met, their drive towards other needs changes accordingly (e.g., keen on knowledge acquisition and less inclined towards grocery shopping).

*Transition Probability Matrix.* The movements among need states are probabilistic following a statistical pattern, which is described in a prior transition probability matrix  $A$ :

$$A[x, y] = \Pr(\eta_{t+1} = y \mid \eta_t = x), \quad (2)$$

where  $\eta_t$  (resp.,  $\eta_{t+1}$ ) is the need state at time  $t$  (resp.,  $t + 1$ ). The need states are equipped with the Markov property as people focus on their predominant need and the most urgent desire. Thus, their current state mainly depends – as a simplification, only depends – on the immediately preceding one, not earlier ones. Besides, the movements are multi-directional and people can potentially switch to lower needs even they are in a higher need. Meanwhile, the transition probabilities in  $A$  are inhomogeneous in time as people's inner status may be suddenly shifted due to an upheaval in the external environment, e.g., the occurrence of COVID-19 pandemic. We set different transition matrices for different event stages, e.g.,  $A_{pre}$ ,  $A_{during}$  and  $A_{post}$ . According a neuroscience theory [42],  $A$  is drawn from normal distributions. Specifically, the initial parameters of  $A_{pre}$  and  $A_{post}$  are estimated based on user behaviors shown in Amazon reviews data [48], while  $A_{during}$  is skewed to basic needs according to the findings of Suh et al. [61].

*Observable Preferences on Slates.* People's preferences on items are driven by their inner need states. We premise users are unable to collect all information of items in RSs before making a decision. People's behavior tends to be rationally bounded, in that a satisfactory solution is sought, rather than the optimal one [57]. Self-assessment of inner needs is usually imperfect and inaccurate, whereby the output of need-HMM (i.e., user preferences on items in the slate  $\mathcal{S}$ ) follows probabilistic patterns rather than deterministic assertion.

*Emission Probability Matrix.* We use an emission matrix  $E$  to describe the output of need-HMM. Conditional on certain need state, the model emits statistical preference distributions for items in  $\mathcal{S}$ . It facilitates the decision-making process in the simulation that the preference is drawn from the emission probabilities:

$$p_{it} \sim E[\eta, i] = \Pr(c_t = i \mid \eta_t = \eta). \quad (3)$$

As shown in Fig. 3, the model contains 5 hidden states corresponding to the 5 levels of needs defined in [43, 61]. If the initial need state  $\pi_{m0}$  is given, need-HMM can predict the need states for successive time periods. Accordingly, we can dynamically and sequentially generate probability distributions of users' preferences.

Need-HMM is the core module in NEST upon which we can simulate the dynamics of user preference during different event stages by manipulating the transition matrix, and observe the reactions of CF models behind the sudden changes.

### 4.3 User Choice and Feedback Model

When interacting with the simulation environment, the plugged recommender will recommend a  $K$ -sized slate of items to a user at each epoch. The consumption decision of the user is determined by the choice model. When the recommendation slate is delivered, the user chooses items based on their private need (which appears as preferences for items) and recommendation scores. It is worth mentioning that, to make the simulation robust and consistent with the need-HMM emission, the choice policy is probabilistic rather than deterministic [40], i.e., Eq. (1).

Once an item is consumed by a user, they produce feedback on the consumed item. The feedback model postulates the rating is a convex combination of the user preference at that time and the perceived quality of the selected item, i.e.,

$$r_{uit} = f(p_{uit}, q_{ui}) = \alpha \cdot p_{uit} + (1 - \alpha) \cdot q_{ui} + \varepsilon_t, \quad (4)$$

where  $\alpha$  balances self-need-driven and crowd-opinion-driven feedback depending of how self-sufficient or herd-influenced the user is. In our experiments, we set  $\alpha = 0.5$  to avoid overly subjective hypotheses, but this can be further explored as future work on NEST;  $\varepsilon_t$  is an additional random error drawn independently from a Gaussian noise distribution  $E \sim N(\mu, \sigma^2)$ , where  $\mu = 0$  and  $\sigma \in \mathbb{R}^+$ . Finally, the user need state updates after each transaction.

### 4.4 Recommender Model

NEST can also be treated as a dynamic Bayesian network over trajectories of user need evolution, recommendations and interaction with RSs [30]. It can factorize over time horizon  $T$  as:

$$\begin{aligned} & \Pr(p_1, \dots, p_T, c_1, \dots, c_T, \mathcal{S}_1, \dots, \mathcal{S}_T) \\ &= \sum_{(\eta_1, \dots, \eta_T)} \left[ \Pr(\eta_0) \Pr(\mathcal{S}_0) \Pr(c_0 \mid \mathcal{S}_0, \eta_0) \right. \\ & \quad \left. \prod_{t=1}^T \Pr(p_t \mid \eta_t) \Pr(\eta_t \mid \eta_{t-1}) \Pr(c_t \mid \mathcal{S}_t, p_{t-1}) \Pr(\mathcal{S}_t \mid \mathcal{I}, \vec{H}(t-1)) \Pr(\mathcal{I}) \right], \end{aligned}$$

where  $\Pr(p_t \mid \eta_t) \Pr(\eta_t \mid \eta_{t-1})$  is the need-HMM with motivation policy  $\pi_m$ ,  $\Pr(c_t \mid \mathcal{S}_t, p_{t-1})$  is the choice model,  $\Pr(\mathcal{S}_t \mid \mathcal{I}, \vec{H}(t-1))$  is the recommendation policy  $\pi_r$  as aforementioned.

Since NEST is built on top of the OpenAI Gym framework, it can work seamlessly with various RL agents designed for this framework, such as Deep Q-Learning [47], or Implicit Quantile Networks [14]. However, as we focus on the impact of macro-trends in population dynamics on RSs, we apply vanilla RL agents in this paper. Besides, to restrain the feedback loop issue [55], we let the




Figure 4: The Adaptive Ensembles Model

agent take the empirical averages over time horizon  $T$  as the recommendation score, i.e.,

$$\hat{r}_{it} = \frac{1}{T} \sum_{t=1}^T r_{it}. \quad (5)$$

#### 4.5 Synthetic Data Generation

To enable offline CF model evaluation within NEST, we set up a static synthetic dataset generation strategy. We add a logging layer in the simulator, which can aggregate relevant statistics pertaining to these interactions for traditional offline CF training. We take snapshots of the running logs of NEST as the training data for offline CF models. For example, when users interact with NEST in each epoch, we log user id, the consumed item, its corresponding rating, timestamp of consumption and their current need state label. Like in the real system, we store the static dataset and share it with offline CF models. To alleviate the feedback loop or other effects on user behavior caused by RSs, we purposely keep the recommendations “vacuous”, i.e., the recommender model generates *random* recommendations in slates. Therefore, users choose items based on their preferences only, and Eq. (1) of the choice model becomes:

$$c_t \sim \text{Pr}(p_t + \varepsilon_t), \quad (6)$$

where  $\varepsilon_t$  is an additional noise variable.

#### 4.6 Adaptive Ensembles

In the context of the pandemic-like event, population-scale concept drifts are widely observed in CF-based RSs. Generalizing and updating the learning algorithms from the time-evolving data present a challenge to CF models. The evolving data can be interpreted as a mixture distribution of target concepts. Ensemble learning has been proved as an effective mechanism in this case [17], in which each candidate algorithm models a target concept respectively. To adapt to the drifts, we propose an Adaptive Ensembles (AE) A/B testing environment inspired by Cañamares’ work in [6] to optimize the overall effectiveness of the assembly. The ensembles are built upon a Thompson Sampling based MAB environment and integrate several widely used CF methods (e.g.,  $k$ NN, MF) and non-personalized methods (e.g., popularity-based methods, average-rating-based methods) into the candidate pool  $C_a$ . Each individual model serves as an arm

Table 1: Need Categories and Amazon Review Datasets

Need Category	Amazon Datasets Examples
SA	Arts Crafts and Sewing; Musical Instruments; Sports and Outdoors
Co	Books; Industrial and Scientific
LB	Gift Cards; Pet Supplies
Sf	Tools and Home Improvement; Patio Lawn and Garden
Ph	Grocery and Gourmet Food; Clothing Shoes and Jewelry




Figure 5: The illustration of Iterative Evaluation. Each column is one iteration (i.e., one snapshot of NEST) of training data and test data. From top to bottom, it logs in chronological order. Each iteration contains 100 epochs in training data and 20 epochs in test data.

and the bandit learns their contribution progressively in the non-stationary environment (Fig. 4). Specifically, the bandit maintains a beta distribution  $B(\alpha_a, \beta_a)$  as the posterior reward for each arm. The parameters  $\alpha_a$  and  $\beta_a$  indicate the number of times to achieve the best and worst moving recall for CF model  $a$ . The bandit pulls the arm  $C_a^*(t)$  probabilistically drawn from their beta distribution  $\rho_a \sim B(\alpha_a, \beta_a)$  and outputs a weighted average of the individual predictions. We evaluate the moving precision of its output as the reward and give feedback to the bandit. The bandit update the beta distribution for next epoch from an iterative training process in the context of pandemic-like events.

### 5 EXPERIMENT

We use a semi-parametric synthetic data generated by the NEST framework. For different stages of an event, we sample from different data sources. Particularly, for pre-event and post-event stages, we take the seeds from various categories of Amazon reviews datasets [48] and match them to user need categories [61], as shown in Table 1. We purify and pooling 10,000 users who have revealed interactions across all 5-level needs; and then construct the initial need states and transition matrices  $A_{pre}$  and  $A_{post}$  for sampled users in the synthetic datasets. For during-event stage, we extract user need distributions from Suh et al. [61]’s outcome and estimate the transition matrix  $A_{during}$ .

#### 5.1 Evaluation Mechanism

By setting different hyperparameters in NEST, we simulate scenarios in a non-stationary context for both online and offline CF models.




Figure 6: Counterfactual results of ranking accuracy of CF models running in different events

We design a set of counterfactual experiments to validate the output of the NEST framework and verify the following hypotheses:

- H1. CF models fail to generate high-quality recommendations in pandemic-like events of different intensities.
- H2. Temporal based models outperform other CF models in pandemic-like events.
- H3. Users who are not affected by the event will receive better recommendations.

*Sequential Evaluation.* The online CF models are trained in a sequential environment. We simulate 1000 users and 1000 items<sup>4</sup> with 30,000 epochs in the environment. The epochs are divided into three 10,000 event stages. The recommender agent collects ratings from users as rewards, which combine the qualities of items and preferences of users, i.e., Eq. (4). Then it generates slates of recommendations ranking by the scoring function Eq. (5). We apply the proposed preference-based natural exploration strategy [3] into NEST following Eq. (1). After each transaction, users will update their need states. User preferences may shift in different event stages. Therefore, users are motivated to explore new items and find the suitable choice synchronously. The impact of changes in user behavior on CF models can be demonstrated by plotting the real-time precision and popularity of selected item (Fig. 9).

*Iterative Evaluation.* As offline CF models are trained on static rating matrix data and can only be retrained regularly, we deploy an iterative evaluation strategy, as shown in Fig. 5.

We define the dissemination magnitude of the event in different iterations. We set the timing of occurrences of the outlier event, which divides the horizon into two event stages (i.e., pre- and during-event) and varies the proportion of normal behavior and affected behavior in the training dataset. At iteration 0, we fit the algorithms on users’ pure normal behavioral data (i.e., pre-event data) and measure the qualities of recommendations on their during-event data. In the subsequent 10 iterations, we iteratively bring the event forward by 10 epochs but keep the total data volume the same so that the testee model can feed on different proportional behavior data. At the appointed epoch, we simulate the intensity of the pandemic-like event from different perspectives:

- (1) We alter the direction of the skewness of the time inhomogeneous transition matrix  $A_{during}$  of users to control the number and size of herds formed.

<sup>4</sup>We also test different sizes of user and item sets. They show similar results, thus we demonstrate only one set of experimental results.




Figure 7: Diversity of recommendations

- (2) We vary population sizes affected by the event to control the scale of herds formed. The adaptation of CF models may have different influence on different types of users.
- (3) We hypothesize that users have different rates of adaptation to the event. Some users may be more sensitive to changes in the external environment, while others less so.

We report the evolution of the recommendation quality across the iterations upon different settings. As a counterfactual comparison, we set 10 anterior iterations (labelled between [-10, -1]) where the training set and test set are both *stationary* data.

We use 20 CF baselines from both rating prediction and item ranking tasks: *Temporal Aware*: TimeSVD [36], RNN4Rec [26], Caser [63]; *Conventional*: User-kNN [60], Item-kNN [54], SlopeOne [38], SVD [37], SVD++ [35], BPR [50], LDA [19], RBM [53], pLSA [27]; *Advanced*: RSTE [41], SocialMF [33], NeuMF [24], CFGAN[9], LightGCN [23]; *Non-personalized*: MostPopular, ItemMean, Random.

We use metrics to examine multiple aspects of the quality of recommendations: *Predictive Accuracy Metrics*: Mean Absolute Error (MAE) and Root Mean Square Error (RMSE) [56]. *Rank Accuracy Metrics*: Precision@k, Recall@k, F1@k and nDCG@k [56]. *Diversity Metrics*: We define two dimensions of the diversity [8]: *intra-user* diversity (e.g., Mean Intra User Distance(MIUD)) and *inter-user* diversity (e.g., Gini index).




Figure 8: Evaluation of baselines on different user groups

## 5.2 Experiment Results and Analysis

As we evaluated plenty of baseline algorithms, to show their performance changes more clearly, we applied a Gaussian filter transformation ( $g(x) = \frac{1}{\sqrt{2\pi}\sigma} e^{-x^2/(2\sigma^2)}$ , where  $\sigma = 2$ ) on their results. Owing to spatial confined, we combine some different metrics that can lead to the same conclusion, e.g., we use Precision@10 to represent predictive and ranking accuracy performance.

**5.2.1 Quality of Recommendations in Events of Varying Intensities.** To verify Hypothesis 1 and 2, we simulate events from 3 perspectives: 1) the typical pandemic-like events, i.e., users show *population-scale sudden drifts* in their behaviors; 2) the events of *non-population scale sudden drifts*, where users shift to various dominant need states at during-event stage; 3) the events of *population scale gradual drifts*, where population shift to new need states in different adaptation rates. The results (Fig. 6) show that

- (1) Most baselines perform solidly and personalized CF models outperform non-personalized ones when users' needs and preference are stable in normal days (i.e., pre-event);
- (2) In events of population-scale drifts, most models (except Random) achieve better precision at during-event stage (Fig. 6a, 6c), where Non-personalized (ItemMean, MostPopular) models are extraordinarily conspicuous. This contrary to our hypothesis suggests that the herding effect seems to increase the homogeneity of users and make them more predictable as Non-personalized models only work if everyone is in the same cluster and only interested in popular items. Also, the sudden shift of users to new states accelerates the formation of herds, which allows algorithms to acquire new user data faster and more than gradual drifts; and
- (3) In a non-population scale event (Fig. 6b), users shift different needs at during-event stage, which reduces the size of herds among users. Temporal Aware and Neighborhood based CF models achieve better performance while other models fail in this scenario. In such event, users are biased towards specific needs according to their own interests instead of imitating others, which forms *rational* herds in the system. It suggests users abandon old preferences and form

new personalities, which breaks the stereotype of users in algorithms. However, neighborhood based models do not require a prior-knowledge of users and only make neighbor estimates from existing data. Temporal aware models maintain good performance due to their good adaptability.<sup>5</sup>

**5.2.2 Diversity of Recommendations.** Diversity metrics are generally used to measure the ability of CF models to help users explore new items. MIUD@10 measures the intra-list diversity of recommendations. As shown in Fig. 7, some baselines (e.g., SVD++, TimeSVD, non-personalized models) generate more divergent recommendations. However, since iteration 5, all baselines are generating more homogeneous recommendations. The diversity of recommendations becomes worse than normal days. Albeit the homogenization of user behaviors during the pandemic-like event “sharpens” the accuracy of recommendations, it aggravates the emergence of echo chambers and narrows user attention.

**5.2.3 Impact of Pandemic-like Events on Different Users.** To verify Hypothesis 3, we divide the entire population evenly into *Normal* group (users with unshakable faith) and *Affected* group (adaptive users). Fig. 8 illustrates the result of nDCG@10 of all baselines on the two groups. At pre-event stage, CF models generate comparable recommendations on both groups. When the outlier event occurs, the affected users may get poor recommendations, but the normal users even get better results as they are traceable and conform to their old impressions. After iteration 3, however, the performance on the affected users increases linearly, but the normal users start to be “affected” that their experience is worse than before. Therefore, we reject the original hypothesis that the “better” overall performance comes at the expense of a normal user experience.

**5.2.4 Performance of Online CF Models.** For online CF models, we set 10,000 epochs per stage and measure Precision@10 with a 100-epoch window. Fig. 9a shows the evolution where the model achieves stable increasing trend at pre-event stage. However, when the pandemic-like event occurs, it experiences a steep drop, and then increase rapidly until the end of the during-event stage. Then, we observe another steep drop when the event subsides.

It seems that the sequential model works better during the pandemic-like event. However, when shifting our viewing angle to popularity of selected items at each epoch, we observe a skewed popularity bias in Fig. 9b. At pre-event stage, the selected items do not show a high popularity once the model is well-trained, which indicates that people have different needs on items. At the beginning of during-event, the overall popularity plummets down because users shift their needs and choose items that they do not buy often before but more urgently needed at this time. After that, the popularity shows a linear increase during the pandemic-like event. It means that more users are prone to purchasing popular items instead of what they used to buy. In other words, individuals feel safer when following the public opinions. This trend indicates the herding effect is spreading in the system. It also explains the rapid increasing trend in Fig. 9a, because users become more predictable. However, this trend gets alleviated once the pandemic-like event ends, because users calm down and have developed new lifestyle.

<sup>5</sup>RNN4Rec achieve its best performance with large data or additional features




Figure 9: Evaluation of online CF models




Figure 10: Results of AE model in Iterative Evaluation

Furthermore, with the help of a built-in logging layer in NEST, which can access statistics of any observable and latent features, we can collect users’ need expressions as they make decisions. We investigate the quantity of user expression in different needs (in Fig. 9c). It is observed a delicate changes that expressions of needs distribute evenly at pre-event and post-event stages, but the Ph need suddenly rises across the during-event stage.

**5.2.5 Adaptive Ensembles.** Both the evaluation results of offline CF models and online CF models manifest a herding effect among user interactions, which brings challenges to CF models whereas some non-personalized methods may achieve a better recommendation. We examine the proposed adaptive ensembles to integrate different methods and learn the optimal algorithm in certain event stage. Specifically, we put four common-used algorithms into the candidate pool  $C$ , which are  $k$ NN, MF, MostPopular and Random algorithms. Fig. 10 shows the Precision@ $k = 1$  of AE and its candidates in the iterative evaluation process where the proposed AE model can effectively and learn the best candidate algorithm. Especially at the during-event stage, AE beats its candidates observably. Therein, 1) at iteration 1, AE model generate comparable recommendations as in normal days while others are frightened by the pandemic-like event; 2) During the entire during-event stage, AE can match users’ dynamic preferences effectively. Moreover, the candidate algorithms largely affects the performance of the

A/B testing environment model. Therefore, enriching the candidate pool will lead to better performance.

## 6 CONCLUSION

We investigate the impacts of “pandemic-like” events on collaborative filtering based on a simulation study. Research and surveys in other fields have shown that user behavior has undergone tremendous changes since the occurrence of such events. The imitation behavior and information cascades between individuals can make the herd effect prominent. We propose a NEST framework to examine the characteristics of concept drifts in user behavior, which is widely applicable in various scenarios: pandemic-like events, gradual population-scale events (e.g., the annual Black Friday/Double 11 shopping festival) and sudden non-population scale events (e.g., local wildfire or flood). We further provide an elementary understanding of the macro-trends of population dynamics on CF models. Our preliminary work show that the CF models are liable to be shocked by the sudden changes initially. However, the irrational herd behavior assimilates user preferences which makes users more predictable. The CF models can learn the changed profiles quickly, but concurrently amplify the popularity bias of items.

## 7 ACKNOWLEDGMENTS

This research was conducted in part by the Australian Research Council Discovery Grant – DP190101485 – and by the ARC Centre of Excellence for Automated Decision-Making and Society, which was funded partially by the Australian Government through the Australian Research Council – CE200100005.

## REFERENCES

- [1] Marie Al-Ghossein, Talel Abdesslem, and Anthony BARRÉ. 2021. A Survey on Stream-Based Recommender Systems. *ACM Computing Surveys (CSUR)* 54, 5 (2021), 1–36.
- [2] Clayton P. Alderfer. 1969. An Empirical Test of a New Theory of Human Needs. *Organizational Behavior and Human Performance* 4, 2 (May 1969), 142–175. [https://doi.org/10.1016/0030-5073\(69\)90004-X](https://doi.org/10.1016/0030-5073(69)90004-X)
- [3] Hamsa Bastani, Mohsen Bayati, and Khashayar Khosravi. 2021. Mostly exploration-free algorithms for contextual bandits. *Management Science* 67, 3 (2021), 1329–1349.
- [4] Greg Brockman, Vicki Cheung, Ludwig Pettersson, Jonas Schneider, John Schulman, Jie Tang, and Wojciech Zaremba. 2016. Openai gym. *arXiv preprint arXiv:1606.01540* (2016).
- [5] R. Cañameres and P. Castells. 2018. Should I follow the crowd? A probabilistic analysis of the effectiveness of popularity in recommender systems. In *Proceedings of the 41st Annual International ACM SIGIR Conference on Research and Development in Information Retrieval (SIGIR 2018)*. ACM, New York, NY, USA, 415–424.

- [6] Rocío Cañamares, Marcos Redondo, and Pablo Castells. 2019. Multi-armed recommender system bandit ensembles. In *Proceedings of the 13th ACM Conference on Recommender Systems*. 432–436.
- [7] P. Castells and R. Cañamares. 2018. Characterization of fair experiments for recommender system evaluation: A formal analysis. In *Proceedings of the Workshop on Offline Evaluation for Recommender Systems (REVEAL 2018) at the 12th ACM Conference on Recommender Systems (RecSys 2018)*.
- [8] Pablo Castells, Neil Hurley, and Saul Vargas. 2022. Novelty and diversity in recommender systems. In *Recommender systems handbook*. Springer, 603–646.
- [9] Dong-Kyu Chae, Jin-Soo Kang, Sang-Wook Kim, and Jung-Tae Lee. 2018. CF-GAN: A Generic Collaborative Filtering Framework Based on Generative Adversarial Networks. In *Proceedings of the 27th ACM International Conference on Information and Knowledge Management (Torino, Italy) (CIKM '18)*. Association for Computing Machinery, New York, NY, USA, 137–146. <https://doi.org/10.1145/3269206.3271743>
- [10] Allison J. B. Chaney, Brandon M. Stewart, and Barbara E. Engelhardt. 2018. How Algorithmic Confounding in Recommendation Systems Increases Homogeneity and Decreases Utility. In *Proceedings of the 12th ACM Conference on Recommender Systems*. ACM, Vancouver British Columbia Canada, 224–232. <https://doi.org/10.1145/3240323.3240370>
- [11] Kae H. Chung. 1969. A Markov Chain Model of Human Needs: An Extension of Maslow's Need Theory. *AMJ* 12, 2 (June 1969), 223–234. <https://doi.org/10.5465/254818>
- [12] G. L. Ciampaglia, A. Nematzadeh, F. Menczer, and A. Flammini. 2018. How algorithmic popularity bias hinders or promotes quality. *Scientific Reports* 8 (2018).
- [13] Olivier Coibion, Yuriy Gorodnichenko, and Michael Weber. 2020. *Labor Markets During the COVID-19 Crisis: A Preliminary View*. Technical Report w27017. National Bureau of Economic Research. <https://doi.org/10.3386/w27017>
- [14] Will Dabney, Georg Ostrovski, David Silver, and Rémi Munos. 2018. Implicit quantile networks for distributional reinforcement learning. In *International conference on machine learning*. PMLR, 1096–1105.
- [15] Laura Di Renzo, Paola Gualtieri, Francesca Pivari, Laura Soldati, Alda Attinà, Giulia Cinelli, Claudia Leggeri, Giovanna Caparelli, Luigi Barrea, Francesco Scerbo, et al. 2020. Eating habits and lifestyle changes during COVID-19 lockdown: an Italian survey. *Journal of translational medicine* 18 (2020), 1–15.
- [16] Lauren K Fink, Lindsay A Warrenburg, Claire Howlin, William M Randall, Niels Chr Hansen, and Melanie Wald-Fuhrmann. 2021. Viral tunes: changes in musical behaviours and interest in coronamusic predict socio-emotional coping during COVID-19 lockdown. *Humanities and Social Sciences Communications* 8, 1 (2021), 1–11.
- [17] João Gama, Indrè Žliobaitė, Albert Bifet, Mykola Pechenizkiy, and Abdelhamid Bouchachia. 2014. A survey on concept drift adaptation. *ACM computing surveys (CSUR)* 46, 4 (2014), 1–37.
- [18] Honghao Gao, Li Kuang, Yuyu Yin, Bin Guo, and Kai Dou. 2020. Mining consuming behaviors with temporal evolution for personalized recommendation in mobile marketing apps. *Mobile Networks and Applications* 25 (2020), 1233–1248.
- [19] Tom Griffiths. 2002. Gibbs sampling in the generative model of latent dirichlet allocation. (2002).
- [20] Shengyu Gu, Beata Ślusarczyk, Sevda Hajizada, Irina Kovalyova, and Amina Sakhbieva. 2021. Impact of the covid-19 pandemic on online consumer purchasing behavior. *Journal of Theoretical and Applied Electronic Commerce Research* 16, 6 (2021), 2263–2281.
- [21] F. Maxwell Harper and Joseph A. Konstan. 2015. The MovieLens Datasets: History and Context. *ACM Trans. Interact. Intell. Syst.* 5, 4, Article 19 (Dec. 2015), 19 pages. <https://doi.org/10.1145/2827872>
- [22] Ruining He and Julian McAuley. 2016. Ups and Downs: Modeling the Visual Evolution of Fashion Trends with One-Class Collaborative Filtering. In *Proceedings of the 25th International Conference on World Wide Web (Montréal, Québec, Canada) (WWW '16)*. International World Wide Web Conferences Steering Committee, Republic and Canton of Geneva, CHE, 507–517. <https://doi.org/10.1145/2872427.2883037>
- [23] Xiangnan He, Kuan Deng, Xiang Wang, Yan Li, YongDong Zhang, and Meng Wang. 2020. LightGCN: Simplifying and Powering Graph Convolution Network for Recommendation. In *Proceedings of the 43rd International ACM SIGIR Conference on Research and Development in Information Retrieval (Virtual Event, China) (SIGIR '20)*. Association for Computing Machinery, New York, NY, USA, 639–648. <https://doi.org/10.1145/3397271.3401063>
- [24] Xiangnan He, Lizi Liao, Hanwang Zhang, Liqiang Nie, Xia Hu, and Tat-Seng Chua. 2017. Neural Collaborative Filtering. In *Proceedings of the 26th International Conference on World Wide Web (Perth, Australia) (WWW '17)*. International World Wide Web Conferences Steering Committee, Republic and Canton of Geneva, CHE, 173–182. <https://doi.org/10.1145/3038912.3052569>
- [25] Harry Helson. 1964. Adaptation-level theory: an experimental and systematic approach to behavior. (1964).
- [26] Balázs Hidasi, Alexandros Karatzoglou, Linas Baltrunas, and Domonkos Tikk. 2015. Session-based recommendations with recurrent neural networks. *arXiv preprint arXiv:1511.06939* (2015).
- [27] Thomas Hofmann. 2004. Latent semantic models for collaborative filtering. *ACM Transactions on Information Systems (TOIS)* 22, 1 (2004), 89–115.
- [28] Mark Honigsbaum. 2013. Regulating the 1918–19 pandemic: Flu, stoicism and the Northcliffe press. *Medical history* 57, 2 (2013), 165–185.
- [29] Jen-Hung Huang and Yi-Fen Chen. 2006. Herding in online product choice. *Psychology & Marketing* 23, 5 (2006), 413–428.
- [30] Eugene Ie, Chih-wei Hsu, Martin Mladenov, Vihan Jain, Sanmit Narvekar, Jing Wang, Rui Wu, and Craig Boutilier. 2019. Recsim: A configurable simulation platform for recommender systems. *arXiv preprint arXiv:1909.04847* (2019).
- [31] Bar Ifrach, Costis Maglaras, and Marco Scarsini. 2014. Bayesian social learning with consumer reviews. *ACM SIGMETRICS Performance Evaluation Review* 41, 4 (2014), 28–28.
- [32] Dmitry Ivanov. 2020. Viable supply chain model: integrating agility, resilience and sustainability perspectives—lessons from and thinking beyond the COVID-19 pandemic. *Annals of operations research* (2020), 1–21.
- [33] Mohsen Jamali and Martin Ester. 2010. A Matrix Factorization Technique with Trust Propagation for Recommendation in Social Networks. In *Proceedings of the Fourth ACM Conference on Recommender Systems (Barcelona, Spain) (RecSys '10)*. Association for Computing Machinery, New York, NY, USA, 135–142. <https://doi.org/10.1145/1864708.1864736>
- [34] Wang-Cheng Kang and Julian McAuley. 2018. Self-Attentive Sequential Recommendation. In *2018 IEEE International Conference on Data Mining (ICDM)*. 197–206. <https://doi.org/10.1109/ICDM.2018.00035>
- [35] Yehuda Koren. 2008. Factorization Meets the Neighborhood: A Multifaceted Collaborative Filtering Model. In *Proceedings of the 14th ACM SIGKDD International Conference on Knowledge Discovery and Data Mining (Las Vegas, Nevada, USA) (KDD '08)*. Association for Computing Machinery, New York, NY, USA, 426–434. <https://doi.org/10.1145/1401890.1401944>
- [36] Yehuda Koren. 2009. Collaborative filtering with temporal dynamics. In *Proceedings of the 15th ACM SIGKDD international conference on Knowledge discovery and data mining*. 447–456.
- [37] Yehuda Koren, Robert Bell, and Chris Volinsky. 2009. Matrix factorization techniques for recommender systems. *Computer* 42, 8 (2009), 30–37.
- [38] Daniel Lemire and Anna Maclachlan. 2005. Slope one predictors for online rating-based collaborative filtering. In *Proceedings of the 2005 SIAM International Conference on Data Mining*. SIAM, 471–475.
- [39] Jiacheng Li, Yujie Wang, and Julian McAuley. 2020. Time Interval Aware Self-Attention for Sequential Recommendation. In *Proceedings of the 13th International Conference on Web Search and Data Mining (Houston, TX, USA) (WSDM '20)*. Association for Computing Machinery, New York, NY, USA, 322–330. <https://doi.org/10.1145/3336191.3371786>
- [40] Jordan J Louviere, David A Hensher, and Joffre D Swait. 2000. *Stated choice methods: analysis and applications*. Cambridge university press.
- [41] Hao Ma, Irwin King, and Michael R Lyu. 2009. Learning to recommend with social trust ensemble. In *Proceedings of the 32nd international ACM SIGIR conference on Research and development in information retrieval*. 203–210.
- [42] Wei Ji Ma, Jeffrey M Beck, Peter E Latham, and Alexandre Pouget. 2006. Bayesian inference with probabilistic population codes. *Nature neuroscience* 9, 11 (2006), 1432–1438.
- [43] A. H. Maslow. 1943. A Theory of Human Motivation. *Psychological Review* 50, 4 (1943), 370–396. <https://doi.org/10.1037/h0054346>
- [44] Abraham H. Maslow. 1970. New Introduction: Religions, Values, and Peak-Experiences. *Journal of Transpersonal Psychology* 2, 2 (1970), 83–90.
- [45] Mieszko Mazur, Man Dang, and Miguel Vega. 2021. COVID-19 and the march 2020 stock market crash. Evidence from S&P1500. *Finance research letters* 38 (2021), 101690.
- [46] E. Mena-Maldonado, R. Cañamares, P. Castells, Y. Ren, and M. Sanderson. 2021. Popularity Bias in False-Positive Metrics for Recommender Systems Evaluation. *ACM Transactions on Information Systems* 39, 3 (2021).
- [47] Volodymyr Mnih, Koray Kavukcuoglu, David Silver, Andrei A Rusu, Joel Veness, Marc G Bellemare, Alex Graves, Martin Riedmiller, Andreas K Fidjeland, Georg Ostrovski, et al. 2015. Human-level control through deep reinforcement learning. *nature* 518, 7540 (2015), 529–533.
- [48] Jianmo Ni, Jiacheng Li, and Julian McAuley. 2019. Justifying recommendations using distantly-labeled reviews and fine-grained aspects. In *Proceedings of the 2019 Conference on Empirical Methods in Natural Language Processing and the 9th International Joint Conference on Natural Language Processing (EMNLP-IJCNLP)*. 188–197.
- [49] Ravi Philip Rajkumar and SM Arafat. 2021. Model Driven Causal Factors of Panic Buying and Their Implications for Prevention: A Systematic Review. *Psychiatry International* 2, 3 (2021), 325–343.
- [50] Steffen Rendle, Christoph Freudenthaler, Zeno Gantner, and Lars Schmidt-Thieme. 2009. BPR: Bayesian Personalized Ranking from Implicit Feedback. In *Proceedings of the Twenty-Fifth Conference on Uncertainty in Artificial Intelligence (Montreal, Quebec, Canada) (UAI '09)*. AUAI Press, Arlington, Virginia, USA, 452–461.
- [51] Steffen Rendle, Christoph Freudenthaler, and Lars Schmidt-Thieme. 2010. Factorizing Personalized Markov Chains for Next-Basket Recommendation. In *Proceedings of the 19th International Conference on World Wide Web - WWW '10*. ACM Press,

- Raleigh, North Carolina, USA, 811. <https://doi.org/10.1145/1772690.1772773>
- [52] David Rohde, Stephen Bonner, Travis Dunlop, Flavian Vasile, and Alexandros Karatzoglou. 2018. RecoGym: A Reinforcement Learning Environment for the problem of Product Recommendation in Online Advertising. In *Proceedings of the Workshop on Offline Evaluation for Recommender Systems (REVEAL 2018) at the 12th ACM Conference on Recommender Systems (RecSys 2018)*.
- [53] Ruslan Salakhutdinov, Andriy Mnih, and Geoffrey Hinton. 2007. Restricted Boltzmann machines for collaborative filtering. In *Proceedings of the 24th international conference on Machine learning*. 791–798.
- [54] Badrul Sarwar, George Karypis, Joseph Konstan, and John Riedl. 2001. Item-based collaborative filtering recommendation algorithms. In *Proceedings of the 10th international conference on World Wide Web*. 285–295.
- [55] Sven Schmit and Carlos Riquelme. 2018. Human interaction with recommendation systems. In *International Conference on Artificial Intelligence and Statistics*. PMLR, 862–870.
- [56] Guy Shani and Asela Gunawardana. 2011. *Evaluating Recommendation Systems*. Springer US, Boston, MA, 257–297. [https://doi.org/10.1007/978-0-387-85820-3\\_8](https://doi.org/10.1007/978-0-387-85820-3_8)
- [57] Herbert A Simon. 1956. Rational choice and the structure of the environment. *Psychological review* 63, 2 (1956), 129.
- [58] Marinko Škare, Domingo Riberio Soriano, and Małgorzata Porada-Rochoń. 2021. Impact of COVID-19 on the Travel and Tourism Industry. *Technological Forecasting and Social Change* 163 (Feb. 2021), 120469. <https://doi.org/10.1016/j.techfore.2020.120469>
- [59] Lones Smith and Peter Sørensen. 2000. Pathological outcomes of observational learning. *Econometrica* 68, 2 (2000), 371–398.
- [60] Xiaoyuan Su and Taghi M. Khoshgoftaar. 2009. A Survey of Collaborative Filtering Techniques. *Advances in Artificial Intelligence 2009* (2009), 1–19. <https://doi.org/10.1155/2009/421425>
- [61] Jina Suh, Eric Horvitz, Ryen W. White, and Tim Althoff. 2021. Population-Scale Study of Human Needs During the COVID-19 Pandemic: Analysis and Implications (WSDM '21). Association for Computing Machinery, New York, NY, USA, 4–12. <https://doi.org/10.1145/3437963.3441788>
- [62] Henri Tajfel, John C Turner, William G Austin, and Stephen Worchel. 1979. An integrative theory of intergroup conflict. *Organizational identity: A reader* 56, 65 (1979), 9780203505984–16.
- [63] Jiayi Tang and Ke Wang. 2018. Personalized top-n sequential recommendation via convolutional sequence embedding. In *Proceedings of the Eleventh ACM International Conference on Web Search and Data Mining*. 565–573.
- [64] Louis Tay and Ed Diener. 2011. Needs and Subjective Well-Being around the World. *Journal of Personality and Social Psychology* 101, 2 (Aug. 2011), 354–365. <https://doi.org/10.1037/a0023779>
- [65] John B Taylor and John C Williams. 2009. A black swan in the money market. *American Economic Journal: Macroeconomics* 1, 1 (2009), 58–83.
- [66] Alexey Tsymbal. 2004. The problem of concept drift: definitions and related work. *Computer Science Department, Trinity College Dublin* 106, 2 (2004), 58.
- [67] Gerhard Widmer and Miroslav Kubat. 1996. Learning in the presence of concept drift and hidden contexts. *Machine learning* 23, 1 (1996), 69–101.
- [68] Eugene Koh Boon Yau, Nicholas Pang Tze Ping, Wendy Diana Shoosmith, Sandi James, Noor Melissa Nor Hadi, and Jiann Lin Loo. 2020. The behaviour changes in response to COVID-19 pandemic within Malaysia. *The Malaysian journal of medical sciences: MJMS* 27, 2 (2020), 45.
- [69] Kum Fai Yuen, Xueqin Wang, Fei Ma, and Kevin X Li. 2020. The psychological causes of panic buying following a health crisis. *International journal of environmental research and public health* 17, 10 (2020), 3513.