



# Ranking with Long-Term Constraints

Kianté Brantley  
Cornell University  
Ithaca, NY, USA  
kdb82@cornell.edu

Zhichong Fang  
Cornell University  
Ithaca, NY, USA  
zf94@cornell.edu

Sarah Dean  
Cornell University  
Ithaca, NY, USA  
sdean@cornell.edu

Thorsten Joachims  
Cornell University  
Ithaca, NY, USA  
tj@cs.cornell.edu

## ABSTRACT

The feedback that users provide through their choices (e.g., clicks, purchases) is one of the most common types of data readily available for training search and recommendation algorithms. However, myopically training systems based on choice data may only improve short-term engagement, but not the long-term sustainability of the platform and the long-term benefits to its users, content providers, and other stakeholders. In this paper, we thus develop a new framework in which decision makers (e.g., platform operators, regulators, users) can express long-term goals for the behavior of the platform (e.g., fairness, revenue distribution, legal requirements). These goals take the form of exposure or impact targets that go well beyond individual sessions, and we provide new control-based algorithms to achieve these goals. In particular, the controllers are designed to achieve the stated long-term goals with minimum impact on short-term engagement. Beyond the principled theoretical derivation of the controllers, we evaluate the algorithms on both synthetic and real-world data. While all controllers perform well, we find that they provide interesting trade-offs in efficiency, robustness, and the ability to plan ahead.

## CCS CONCEPTS

• Information systems → Recommender systems.

## KEYWORDS

ranking, exposure allocation, long-term objectives

### ACM Reference Format:

Kianté Brantley, Zhichong Fang, Sarah Dean, and Thorsten Joachims. 2024. Ranking with Long-Term Constraints. In *Proceedings of the 17th ACM International Conference on Web Search and Data Mining (WSDM '24), March 4–8, 2024, Merida, Mexico*. ACM, New York, NY, USA, 10 pages. <https://doi.org/10.1145/3616855.3635819>

## 1 INTRODUCTION

Optimizing search and recommendation platforms based on feedback that users provide through their choices (e.g., clicks, purchases) has led to great improvements in ranking quality. However, myopically training systems based on choice data may only improve short-term engagement, but not the long-term sustainability of the platform and the long-term benefits to its users, content providers, and other stakeholders [21]. In particular, platforms operate as part of a complex socio-technical system, and many have argued how



This work is licensed under a Creative Commons Attribution International 4.0 License.

WSDM '24, March 4–8, 2024, Merida, Mexico  
© 2024 Copyright held by the owner/author(s).  
ACM ISBN 979-8-4007-0371-3/24/03.  
<https://doi.org/10.1145/3616855.3635819>

such AI systems can amplify misinformation [16], harm supply through rich-get-richer dynamics [32], incentivize spam [24], or perpetuate human biases [31].

In this complex space of problems and competing interests, we argue that improved tools for explicitly steering the long-term dynamics of the platform are needed. These tools should enable decision-makers to specify long-term goals for the search and recommendation algorithms beyond short-term engagement maximization. While approaches using reinforcement learning have the potential to directly optimize long-term goals, it is challenging to apply them to complex and large-scale information retrieval settings [1, 29]. End-to-end frameworks obscure important decision points [19], leading to problems like a lack of reproducibility [4, 12, 20], reward hacking [15, 36–38, 41], and user tampering and manipulation [8, 13, 14, 28]. Instead, we argue that providing designers with a novel *macroscopic* view will enable strategic reasoning about long-term platform dynamics and that it will enable new tools for steering the platform. The key algorithmic challenge lies in bridging the gap between long-term goals at a macro-level that span many requests, and the micro-level goal of maximizing engagement for each individual request.

In this paper, we develop a new class of macro-level interventions for steering the long-term dynamics of AI platforms, as well as the mechanisms for optimally executing these macro-level interventions. In our framework, further described in Sections 1.1 and 2, macro-level interventions take the form of exposure or impact targets over substantial periods of time (e.g., days, weeks, months). The macro-level interventions can come from various decision-makers, including the users themselves (e.g., "I want to buy at least 30% local products next month" on an e-commerce platform), the platform operator (e.g., "promote local music communities by serving at least 50% local artists on average" on the Localify! music platform [30]), or regulators (e.g., the recent settlement between Meta and the Department of Housing and Urban Development (HUD) [43] that requires Meta to ensure that each housing-related ad is shown with demographic parity to all protected groups over the ad's lifetime). All such macro-level goals steer the behavior of the system over the course of many requests. This creates complex interactions between individual requests, their short-term metrics, and the long-term goals [34].

We address the key technical problem of designing algorithms which break down macro-level goals into a sequence of individual rankings that least hurt the micro-level metric (e.g., engagement). We view these algorithms as controllers which drive the value of macro-level metrics towards specified targets while responding to incoming requests in real time. In Section 3, we rigorously derive three controllers. The first is a baseline approach that satisfies the macro-level goals at a high cost to the micro-level utility. The second enables a finer trade-off between macro- and micro-level

objectives. The final controller incorporates planning to handle requests coming from non-stationary distributions with temporal patterns. To clarify the design of these controllers and their affordances and limitations, we make interesting novel connections between ranking and concepts from online stochastic optimization [3] and model-predictive control [7]. Furthermore, we evaluate the controllers on a number of synthetic and real-world datasets in Section 4, which provides practical guidance on when the use of each controller is most appropriate.

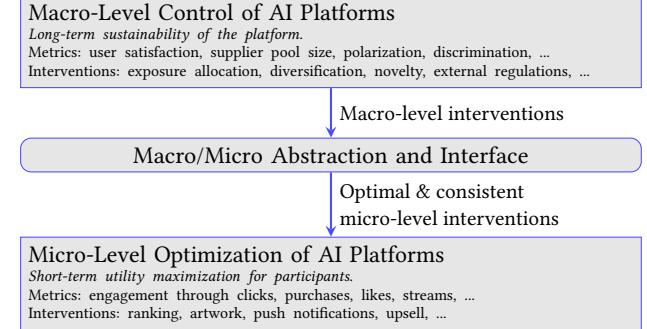
## 1.1 Motivation & Related Work

We argue that one of the key challenges in steering the long-term dynamics of AI platforms results from a mismatch in time scales. Algorithms on these platforms typically optimize metrics pertaining to individual requests or sessions, while the dynamics we aim to control play out over weeks or months of repeated interactions. Optimization on a per-request basis is ill-suited for even expressing long-term objectives, much less for steering their dynamics. Instead, we argue that we need a novel *macroscopic* view to enable strategic reasoning about the long-term dynamics of the platform in addition to the *microscopic* view that our methods currently focus on.

Existing work on incorporating long-term goals in search and recommendation has largely focused on fairness. Early works defining fairness criteria posed them as constraints on impact or exposure to be fulfilled within a single ranking [10, 40, 50, 50]. Later work introduced a temporal perspective, including Celis et al. [9] who develop an online algorithm for recommending diverse viewpoints and Morik et al. [35], Usunier et al. [44] who present algorithms for satisfying fairness cumulatively over multiple rankings. More recently, reinforcement learning algorithms have been applied to long-term fairness in order to handle endogenous dynamics—i.e., the impact of a ranking decision on future utilities and constraints [18, 48, 52]. Beyond fairness, long-term exposure constraints have been motivated as an optimal strategy under the endogenous dynamics of content-provider viability [33, 51]. Many of these settings fit into the framework that we propose. However, unlike approaches which attempt to directly handle dynamics, we argue for elevating such strategic concerns to the definition of interventions.

Partitioning into microscopic and macroscopic views has proven essential in the control of other complex systems. For example, macroeconomic metrics like gross domestic product and unemployment rate describe the state of our economy as a whole, and we use these metrics to reason about its long-term dynamics. Macro-level interventions are used to influence these metrics, like the interest rates set by the Federal Reserve. Similar examples are also widespread in engineered systems, where, for example, the macroscopic control intervention of a self-driving car (e.g., turn left 10 degrees) hides the microscopic execution of this command (e.g., voltages going to the steering motors) behind a control system.

Figure 1 illustrates an analogous micro-macro view of AI platforms, where the macro-level metrics we aim to optimize are quantities like customer satisfaction, retention, polarization, or the size of the supplier pool. It is not hard to think of possible macro-level interventions either, like the rate with which the service is interrupting users with push notifications, how aggressively to prune clickbait,



**Figure 1: We propose to separate macro-level control used for steering the long-term dynamics of the platform from its micro-level engagement optimization. The interface layer provides an abstraction by optimally translating strategic macro-level interventions into a sequence of micro-level actions with minimal impact on short-term metrics.**

or how much exposure to give to smaller suppliers. Enabling reasoning at the level of macro-level metrics and interventions opens the door for future investigations of platform dynamics. At the macro level, it will be far more tractable to understand how interventions affect the long-term metrics we aim to optimize for. For example, establishing a causal model of how exposure allocation to small suppliers (a scalar) relates to the size of the supplier pool (another scalar) is considerably less complex than estimating a causal link between millions of rankings and supplier-pool size.

In this paper, we focus on macro-level interventions that represent constraints on exposure or impact. Exposure can be quantified by models like the position-based model (PBM) [11] which assigns a score to each position in a ranking, representing the probability of being viewed by the user. The impact can be directly measured by clicks or purchases. Adding such long-term constraints provides a rich new language for guiding system behavior, as illustrated by the following examples that can be implemented in our framework:

- (1) Give local artists at least  $X$  percent of the overall exposure over the next month. (Item Group Exposure)
- (2) Show new artist  $A$  to at least  $X$  users over the next week. (Single Item Exposure)
- (3) Given a well-calibrated but imperfect spam filter, ensure that the expected exposure to spam across all users is less than  $X$ . (Item Group Exposure with uncertain Group Membership)
- (4) Do not send more than  $X$  push messages on average per week to user  $U$ . (Single User Exposure)
- (5) Show each housing-related ad to protected groups by demographic parity over the lifetime of the ad. (Single Item / User Group Exposure)
- (6) Support goal of user  $U$  to buy at least 30% of products from local suppliers. (Single User Impact)

These examples show that macroscopic interventions can shape the aggregate experience of items over a given time span (first three), but also the aggregate experience of users (last three). The macroscopic interventions can provide constraints on the experience of a single user or item (2, 4, 6), on the collective experience of

item or user groups (1, 3), or on the complex interaction between item and user groups (5). These interventions may be exact when we have precise knowledge of class membership, or they may be approximate and fulfilled only in expectation (e.g., based on the probability of an article being spam). Finally, in some cases, the constraints directly act upon exposure (first five), while in example six the user asks the system to support a particular impact goal [40] that includes the reactions of the users (e.g., clicks, purchases). Similar macro-level constraints are also relevant to other aspects of AI platforms like ads [45, 47, 53, 54] and ad-pacing [2, 46] in particular.

## 2 MACRO/MICRO CONTROL PROBLEM

We now formalize the problem of translating a set of macro-level interventions into a sequence of micro-level actions. We model this translation as a problem of optimal control which seeks to ensure that the micro-level actions achieve the desired macro-level interventions in aggregate. This is analogous to how control is used in mechanical systems to provide layers of abstraction with clear semantics. Controllers are used to keep a system in desirable states even under external perturbations and incomplete knowledge of the dynamics (e.g., keep a plane in level flight). Under bounds on worst-case conditions, controllers can be proven to be stable, safe, and performant [5, 39]. Furthermore, the control perspective has already proven useful for other aspects of online systems [22, 26, 46, 47, 49, 54].

The macro/micro control problem takes the following form. At each time step  $t$  from 1 to the final time step  $T$ , a new context  $x_t$  arrives. Each context  $x_t$  is drawn independently from some unknown and possibly shifting sequence of distributions

$$x_t \sim P_t.$$

At the micro-level, our goal is to derive a ranking policy that selects an action  $a_t$  (i.e., ranking) for each context that maximizes a micro-level ranking metric  $u(a_t|x_t)$  (e.g., Discounted Cumulative Gain (DCG) [23]). Thus, over all contexts  $x_1, \dots, x_T$  our policy should choose actions  $a_1, \dots, a_T$  which achieve a large cumulative value:

$$\sum_{t=1}^T u(a_t|x_t). \quad (1)$$

However, unlike conventional ranking policies, ours must also consider macro-level goals in addition to the micro-level utility. In particular, the policy needs to fulfill  $m$  constraints that range over all time steps from 1 to  $T$ :

$$\begin{aligned} \sum_{t=1}^T c_1(a_t|x_t) &\geq \tau_1 \\ &\vdots \\ \sum_{t=1}^T c_m(a_t|x_t) &\geq \tau_m \end{aligned} \quad (2)$$

Each of the  $m$  constraints represents one macro-level intervention, where each  $c_i(a_t|x_t)$  represents a metric (e.g., exposure of an ad to demographic user groups, purchase of a local product), and each corresponding  $\tau_i$  represents the cumulative goal over the time period from 1 to  $T$ . For a more compact presentation, we use the

---

### Algorithm 1: Macro/Micro Control Loop

---

```

Input: target  $\tau$ , violation cost  $\phi$ , final time  $T$ , controller  $\Pi$ 
set initial state  $s_0 = 0$ 
foreach  $t$  from 1 to  $T$  do
  observe  $x_t \sim P_t$ 
  select action  $a_t = \Pi(x_t, s_{t-1}, t)$ 
  update state  $s_t = s_{t-1} + c(a_t|x_t)$ 
Compute objective:  $\sum_{t=1}^T u(a_t|x_t) - \phi^\top (\tau - s_T)_+$ 

```

---

following vector notation to represent all  $m$  constraints:

$$\sum_{t=1}^T c(a_t|x_t) \geq \tau. \quad (3)$$

While maximizing Equation (1) subject to Equation (3) may look like a straightforward optimization problem at first glance, it cannot be solved using standard optimizers. Crucially, contexts arrive sequentially, and the policy must immediately choose an action  $a_t$  at time step  $t$  upon receiving  $x_t$ . However, at time  $t$ , we only know the contexts  $x_1, \dots, x_t$ , but not the future contexts  $X_t := (x_{t+1}, \dots, x_T)$ . This means that we cannot directly evaluate the constraints, and must pick the current action under partial information.

To address this problem, we first introduce a state  $s_t$  that reflects the progress made towards fulfilling the constraints up to time  $t$ :

$$s_t = \sum_{t'=1}^t c(a_{t'}|x_{t'}). \quad (4)$$

Then, we consider ranking controllers of the form  $a_t = \Pi(x_t, s_{t-1}, t)$ . The ranking  $a_t$  is chosen based on the context  $x_t$ , the state  $s_{t-1}$ , and the time  $t$ . The macro-level intervention is achieved when the constraints are fulfilled in the final time step  $T$ . This is equivalent to ensuring that the terminal state  $s_T \geq \tau$ . The controller aims to reach this target state but must make decisions only on the basis of the current state and context without exact knowledge of the future contexts.

Fulfilling all constraints may not always be possible and, furthermore, it may not be desirable to arbitrarily sacrifice the micro-level objective from Equation (1). We thus consider soft constraints and denote the macro-violation cost as

$$\phi^\top (\tau - s_T)_+, \quad (5)$$

where  $\phi \geq 0$  is an  $m$  dimensional parameter vector that expresses how costly it is to violate each constraint. The “hinge loss”  $(\cdot)_+$  sets all negative components of the input vector to zero, so any dimension of the terminal state  $s_T$  that is above its target  $\tau$  contributes zero to the macro-violation cost.

The overall objective is the sum of the micro-level utilities minus the macro-violation cost at the final time step  $T$ . For a given sequence of contexts, this objective is:

$$\sum_{t=1}^T u(a_t|x_t) - \phi^\top \left( \tau - \sum_{t=1}^T c(a_t|x_t) \right)_+. \quad (6)$$

Note that this final objective can only be computed after time step  $T$ . Since the actions must be chosen sequentially, this setting has the form of a closed-loop control problem, where the controller

can react to the state  $s_{t-1}$ . This control loop is summarized in Algorithm 1.

We conclude our general setup with a discussion of metrics which depend on *modeled* vs. *observed* feedback. When metrics are defined by a model (e.g., the position-based model of exposure), the result of any arbitrary action  $a$  can be anticipated once the context  $x_t$  is observed. In contrast, metrics defined by observed feedback (e.g., clicks or hover time) are known only for the chosen action  $a_t$  and are observed only after the action is taken. The control loop illustrated in Algorithm 1 is valid whether the metrics defining  $u$  and  $c$  are modeled or observed. However, the action selection step defined by  $\Pi$  benefits from the ability to anticipate the effect of arbitrary actions. We, therefore, focus on modeled metrics in the controller development below—in other words, we take  $u(\cdot|x_t)$  and  $c(\cdot|x_t)$  to be known functions. We call this the *full information* setting, referring to the fact that the context provides sufficient information. In practice, controllers may operate on the basis of imperfect models (e.g., using a learned relevance predictor instead of true relevances), even if the closed-loop logic proceeds according to observed feedback, and better models may be learned interactively by using this feedback. However, for clarity of exposition, we leave such a scenario to future work.

## 2.1 Linear Utilities and Constraints

We now define a specific class of models for describing the relationship between an arbitrary action  $a$  and the macro- and micro-level objectives for a given context  $x$ . For ease of exposition, we will refer to the micro-level objective  $u(a|x)$  as “utility” and each macro-level  $c_i(a|x)$  as “progress towards intervention  $i$ ”.

The utility depends on both the relevance of the items and their ranked positions. In the full information setting, the relevance of the items can be determined from context  $x_t$ . Denote by  $\mathbf{r}_{t,j}$  the relevance of item  $j$  at time  $t$ . Further define a position-dependent weight  $\mathbf{u}_k$  for each position for  $k \in [n]$  (e.g., Discounted Cumulative Gain [23] is  $\mathbf{u}_k = \log_2(k+1)^{-1}$ ). Then an item  $j$  ranked in position  $k$  contributes  $\mathbf{r}_{t,j}\mathbf{u}_k$  to the utility. Denoting by  $\text{rank}(j|a_t)$  the position of item  $j$  under the ranking specified by  $a_t$ , the utility has a linear form:

$$u(a_t|x_t) = \sum_{j=1}^n \mathbf{r}_{t,j} \mathbf{u}_{\text{rank}(j|a_t)}.$$

Without loss of generality, we assume that  $\mathbf{u}_k$  is non-increasing in position  $k$ , meaning that the higher an item is ranked, the more its relevance contributes to the utility.

In the absence of macro-interventions, a utility-maximizing action sorts the items in order of their relevance scores:  $a_t = \text{argsort } \mathbf{r}_t$ . However, our ranking controllers also consider progress towards the macro-level interventions of the following linear form:

$$c_i(a_t|x_t) = \sum_{j=1}^n W_{t,ij} \mathbf{e}_{\text{rank}(j|a_t)} \quad i \in \{1, \dots, m\}.$$

Above,  $W_{t,ij}$  determines the contribution of item  $j$  to intervention  $i$  (e.g., an indicator of group membership). As with relevance, we assume the full information setting, so that the context  $x_t$  contains enough information to determine this quantity. Denoted by  $\mathbf{e}_k$  is

another position-dependent weight. It is not necessary to assume that it is equal to  $\mathbf{u}_k$  or even that it is also non-increasing in  $k$ .

From here forward, we denote the parameters of the utility and progress functions by vectors and matrices:  $\mathbf{r}_t \in \mathbb{R}^n$ ,  $\mathbf{u} \in \mathbb{R}^n$ ,  $W_t \in \mathbb{R}^{m \times n}$ , and  $\mathbf{e} \in \mathbb{R}^n$ . We assume that position weights  $\mathbf{u}$  and  $\mathbf{e}$  are known<sup>1</sup> and that the context provides full information about utility and interventions so that  $x_t = (\mathbf{r}_t, W_t)$ .

For the ranking controllers that we develop, it is convenient to use permutation matrices for representing rankings. A permutation matrix has exactly one entry equal to 1 in each row and column, and 0 elsewhere. If  $\Sigma \in \mathbb{R}^{n \times n}$  represents a ranking, then  $\Sigma_{kj} = 1$  means that item  $j$  will be placed in position  $k$ . Using this notation, the utility and exposure quantities can be written in a compact matrix-vector notation, identifying  $a_t$  with the corresponding permutation matrix  $\Sigma_t$ :

$$u(a_t|x_t) = \mathbf{r}_t^\top \Sigma_t \mathbf{u}, \quad c(a_t|x_t) = W_t \Sigma_t \mathbf{e}. \quad (7)$$

Since searching in the discrete space of permutations can be computationally challenging, it is sometimes convenient to search over ranking distributions. This corresponds to considering policies represented by doubly stochastic matrices rather than permutation matrices. The set of doubly stochastic matrices is defined as

$$\Delta = \left\{ \Sigma \in \mathbb{R}_+^{n \times n} \mid \sum_{k=1}^n \Sigma_{kj} = 1 \text{ and } \sum_{j=1}^n \Sigma_{kj} = 1 \forall j, k \in \{1, \dots, n\} \right\}.$$

Given a doubly stochastic  $\Sigma$ , a ranking can be sampled via the Birkhoff-von Neumann decomposition [6, 40].

## 3 CONTROLLERS FOR RANKING

We now introduce three controllers to address the macro/micro control problem. We begin with a baseline *myopic* controller that has a high reduction in micro-level engagement. Next, we use the lens of online optimization to introduce a controller appropriate for *stationary* context distributions and draw the connection to a previously proposed P-controller for ranking under fairness constraints [35]. Finally, we develop a more sophisticated *predictive* controller that can anticipate and plan for non-stationarities in the context distribution.

### 3.1 Myopic Controller (MC)

Actions must be chosen at every time step  $t$  without knowledge of future contexts. As a result, the controller cannot exactly optimize (6). A simple idea to address this issue is to define an intermediate objective at each time  $t$ :

$$\sum_{t'=1}^t u(a_{t'}|x_{t'}) - \phi^\top \left( \frac{t}{T} \tau - \sum_{t'=1}^t c(a_{t'}|x_{t'}) \right)_+. \quad (8)$$

This intermediate objective scales the target  $\tau$  linearly by  $\frac{t}{T}$  and removes the effect of future timesteps. Note that this is not equivalent to the original objective due to the nonlinearity of the hinge loss. Effectively, this objective treats every timestep as if it were the final timestep (albeit scaling the target value). Since there are no future timesteps to consider under this simplified objective, maximizing the objective for selecting the current action  $a_t$  is well-defined.

<sup>1</sup>It is trivial to extend our controllers to the case that position-dependent weights vary with context, so long as the context provides full information on them.

This leads to the following controller, which we call the Myopic Controller (MC):

$$\Pi_{MC}(x_t, s_{t-1}, t) = \operatorname{argmax}_a u(a|x_t) - \phi^\top \left( \frac{t}{T} \tau - s_{t-1} - c(a|x_t) \right)_+$$

This expression contains only the terms from the objective (8) which affect the argmax. Notice that the maximizing action depends on the past actions and contexts only through the state  $s_{t-1}$ . Algorithm 2 presents the linear program (LP) implementation.

---

**Algorithm 2:** Myopic Controller LP

---

**Input:**  $x_t = (\mathbf{r}_t, W_t), s_{t-1}, t$

**Parameters:**  $\tau \in \mathbb{R}^m, \phi \in \mathbb{R}^m$

$$\Sigma_t = \operatorname{argmax}_{\Sigma \in \Delta} \mathbf{r}_t^\top \Sigma \mathbf{u} - \phi^\top \left( \frac{t}{T} \tau - s_{t-1} - W_t \Sigma \mathbf{e} \right)_+$$

**Return:**  $a_t \sim \Sigma_t$

---

While the intermediate objective at each time step is simple, it is overly strict. Specifically, it charges the full violation cost at the current time step if the controller is unable to reach the scaled target  $\frac{t}{T}$ . It thus ignores that the full violation cost is truly incurred only in the final time step, and that the intermediate violations may cancel out before the final time step. We can thus expect this controller to perform very conservatively.

### 3.2 Stationary Controller (SC)

To address the inappropriate strictness of the Myopic Controller, we turn to ideas from online convex programming. Algorithms developed for this setting select optimization variables, in our case actions, at each time step based on streaming optimization parameters, in our case contexts [3]. As a first step, consider the following objective,

$$\min_{0 \leq \lambda \leq \phi} \frac{1}{T} \sum_{t=1}^T u(a_t|x_t) - \lambda^\top \left( \frac{1}{T} \tau - \frac{1}{T} \sum_{t=1}^T c(a_t|x_t) \right), \quad (9)$$

where the inequality constraints on the Lagrange multiplier vector  $\lambda \in \mathbb{R}^m$  are defined elementwise. Besides rescaling by  $\frac{1}{T}$ , this is equal to the original objective (6): minimizing over the multiplier  $\lambda$  constrained to  $[0, \phi]$  is an exact reformulation of the constraints implicit in the hinge loss appearing in the macro-violation cost.

So far, the reformulation does not solve the problem of intermediate objectives, since minimizing the multiplier  $\lambda$  requires summing over the entire horizon. However, if the value of  $\lambda$  were fixed, then the objective would be separable over timesteps, and maximizing with respect to the action  $a_t$  at time  $t$  would no longer depend on the future. But we would have the same issues as the fixed violation cost in the MC. How should a value of  $\lambda$  be selected? The key insight from the online optimization literature is to alternate between maximizing the objective while holding  $\lambda$  fixed and updating  $\lambda$  to iteratively minimize the objective. Updating the multiplier in this way is like learning a dynamic violation cost. Concretely, actions are chosen according to

$$a_t = \operatorname{argmax}_a u(a|x_t) - \lambda_{t-1}^\top \left( \frac{1}{T} \tau - c(a|x_t) \right). \quad (10)$$

This expression is simplified to contain only the terms which affect the argmax. It can be interpreted as approximating the average

utility and macro-level progress over time as the utility or progress at each time step. This is well motivated for i.i.d. contexts [3], and we therefore call this the Stationary Controller (SC).

It remains to specify the multiplier updates. In general,  $\lambda_t$  is defined based on  $\lambda_{t-1}$  and the gradient of the objective with respect to the multiplier:  $\frac{1}{T} \tau - c(a_t|x_t)$ . In experiments, we use a variant of online gradient descent with adaptive step size. For the sake of exposition, we derive a closed-form expression for the controller in the simpler case of gradient descent with fixed step size  $\gamma > 0$  and initialization  $\lambda_0 = 0$ . In this case,

$$\lambda_t = \lambda_{t-1} + \gamma \left( \frac{1}{T} \tau - c(a_t|x_t) \right) = \gamma \left( \frac{t}{T} \tau - s_t \right). \quad (11)$$

The multiplier  $\lambda_t$  is exactly the tracking error between the linearly scaled target and the current state, scaled by the step size  $\gamma$ . Accounting for the bounds on  $\lambda$ , the closed-loop control law can be written as

$$\Pi_{SC}(x_t, s_{t-1}, t) = \operatorname{argmax}_a u(a|x_t) + \gamma \left( \frac{t-1}{T} \tau - s_{t-1} \right)_{[0, \phi]}^\top c(a|x_t).$$

where  $(z)_{[0, \phi]} = \max\{0, \min\{\phi, z\}\}$  is an elementwise clipping operation. This closed-form expression illustrates that the Stationary Controller adapts the weight on macro-level interventions depending on how much progress it has made towards the goal. Algorithm 3 presents the linear program (LP) implementation.

---

**Algorithm 3:** Stationary Controller LP

---

**Input:**  $x_t = (\mathbf{r}_t, W_t), s_{t-1}, t$

**Parameters:**  $\tau \in \mathbb{R}^m, \phi \in \mathbb{R}^m, \gamma \in \mathbb{R}$

$$\Sigma_t = \operatorname{argmax}_{\Sigma \in \Delta} \mathbf{r}_t^\top \Sigma \mathbf{u} - (\lambda_{t-1}^\top)_{[0, \phi]} W_t \Sigma \mathbf{e}$$

update  $\lambda_t$  from  $\lambda_{t-1}$  with  $\gamma$  and gradient  $\frac{1}{T} \tau - W_t \Sigma_t \mathbf{e}$

**Return:**  $a_t \sim \Sigma_t$

---

**3.2.1 Proportional (P) control.** We briefly outline the connection between the Stationary Controller and a (seemingly) heuristic method for boosting the position of certain items within a ranking. Proportional (or simply “P”) control is a general control technique which applies a correction proportional to the size of a tracking error. In the context of ranking, P-control makes direct adjustments to relevance scores and was first proposed by Morik et al. [35] for achieving long-term fairness constraints.

To draw this connection, we write the control law  $\Pi_{SC}(x_t, s_{t-1}, t)$  as a linear optimization problem and assume that  $\mathbf{u} = \mathbf{e}$ :

$$\begin{aligned} & \operatorname{argmax}_{\Sigma \in \Delta} \left( \mathbf{r}_t^\top + \gamma \left( \frac{t-1}{T} \tau - s_{t-1} \right)_{[0, \phi]}^\top W_t \right) \Sigma \mathbf{u} \\ &= \operatorname{argsort} \mathbf{r}_t + \gamma W_t^\top \left( \frac{t-1}{T} \tau - s_{t-1} \right)_{[0, \phi]} \end{aligned} \quad (12)$$

The equality holds because  $\mathbf{u}$  is non-increasing. Therefore, in this special case, the control law simply sorts items by adjusted relevance scores. The adjustments are proportional to the tracking error, where the transpose matrix  $W^\top$  can be understood as translating from macro-level goals to individual items. Therefore, items associated with lagging macro-level metrics will be boosted. The step size parameter  $\gamma$  can be interpreted as the “gain” of the P-controller, determining its sensitivity to tracking errors.

The P-controller is usually thought of as a heuristic, derived without reference to an overall objective. Its score adjustment does not usually take into account the macro-violation cost parameter  $\phi$  and it cannot account for possible differences between  $u$  and  $e$ . Despite this cost-obliviousness, the derivation above shows that P-control arises as a special case of our Stationary Controller.

### 3.3 Predictive Controller (PC)

All controllers presented so far attempt to make progress towards macro-level goals at a constant rate over the horizon: at step  $t$ , the target is determined to be  $\frac{t}{T}\tau$ . This does not account for non-stationarity in the context distribution. For example, certain types of items may be relevant only on weekends or evenings. Attempting to progress on macro-level goals at a linear rate fails to take into account variable underlying demand.

We therefore derive a predictive controller which accounts for the entire time horizon. Denote future actions as  $A_h = (a_{h+1}, \dots, a_T)$  and similarly for contexts  $X_h$ . The total progress can be written as

$$\sum_{t=1}^T c(a_t|x_t) = s_{h-1} + c(a_h|x_h) + C(A_h|X_h) \quad (13)$$

where  $C(A_h|X_h)$  is the “progress-to-go” at  $h$  defined as the sum of  $c(a_t|x_t)$  for  $t$  from  $h+1$  to  $T$ . The first term in the above expression is the state: the accumulated progress so far. The middle term is the contribution at time  $h$ , and the final term is the cumulative progress to come. This expression explicitly separates the contribution of the past (known), present (current decision), and future (unknown).

The portion of the optimization objective (6) depending on the action  $a_t$  at time  $t$  can be written as

$$u(a_t|x_t) - \phi^\top (\tau - (s_{t-1} + c(a_t|x_t) + C(A_t|X_t)))_+ \quad (14)$$

Notice that because the utility is separable over time, the contributions of past and future utility do not impact the decision at time  $t$ . Due to the hinge loss, the macro-level goal is not separable over time. The progress-to-go  $C(A_t|X_t)$  depends on the future contexts, which are unknown, and future actions, which are hard to choose without knowing the contexts. Instead, we propose using predicted values denoted by  $\widehat{C}_t$ . In Appendix ??, we present methods for forecasting the progress-to-go from historical data.

The following develops a multi-forecast predictive controller that can make use of such progress-to-go estimates. Given a bootstrap sample of  $B$  forecasts of the progress-to-go, the predictive controller selects an action which maximizes the average objective over these  $B$  possible futures. The multi-forecast objective at time  $t$  is represented by the following optimization problem:

$$\max_a \frac{1}{B} \sum_{b=1}^B u(a|x_t) - \phi^\top (\tau - s_{t-1} - c(a|x_t) - \widehat{C}_t^b)_+ \quad (15)$$

Finally, we introduce multipliers for each forecast and use the same alternating online optimization approach as developed in the previous section for the Stationary Controller. Putting together all the pieces, the predictive controller  $\Pi_{PC}(x_t, s_{t-1}, t)$  selects actions according to:

$$a_t = \operatorname{argmax}_a u(a|x_t) + \frac{1}{B} \sum_{b=1}^B \left( \lambda_{t-1}^b \right)_{[0, \phi]}^\top c(a|x_t), \quad (16)$$

and updates each multiplier by defining  $\lambda_t^b$  based on  $\lambda_{t-1}^b$  and the gradient of the objective with respect to the multiplier:  $\tau - s_{t-1} - c(a_t|x_t) - \widehat{C}_t^b$  for  $b \in [B]$ . For the simple case of online gradient descent, this update takes the form

$$\lambda_t^b = \lambda_{t-1}^b + \gamma \left( \tau - s_{t-1} - c(a_t|x_t) - \widehat{C}_t^b \right) \quad b = 1, \dots, B.$$

Similar to the Stationary Controller, actions are chosen according to a weighted objective of utility and progress towards the macro-level targets. However, while the Stationary Controller updates the multiplier to target a linear rate of progress, the predictive controller updates the multiplier depending on potentially non-stationary forecasts of progress-to-go (e.g., expected higher demand on the weekend). Algorithm 4 presents the LP implementation.

## 4 EXPERIMENTS

While each controller comes with a strong conceptual and theoretical motivation, we now evaluate how far these arguments translate into improved empirical performance. In particular, we evaluate the controllers on real-world datasets to assess their differences on realistic data. Furthermore, we present experiments on synthetic data to explore in which situations PC outperforms SC. Implementations of the controllers and code for reproducing the experiments are available at [https://github.com/xkianteb/ranking\\_constraints](https://github.com/xkianteb/ranking_constraints).

### 4.1 Experiment Setup

In addition to the controllers discussed in Section 3, we include results for two additional controllers for comparison. As an (unachievable) skyline, we report the performance of an oracle controller that has access to the whole sequence of test time contexts and directly optimizes the overall objective. For further comparison, we also include the unconstrained controller (MC w/o constraints), which only optimizes utility without enforcing any macro-level interventions.

We conduct experiments on three datasets: The first dataset is KuaiRec [17], which is a fully observed dataset collected from the recommendation logs of the video-sharing mobile app Kuaishou. The KuaiRec dataset consists of 1,411 users, 3,327 items, 4,676,570 interactions, and has a density of 99.6%. We filter to include only items that every user has interacted with, which reduces the number of items from 3,327 to 2,062. We consider the task of ranking videos for sequentially arriving users. Since the KuaiRec dataset does not provide relevance scores, we define the relevance score as half of the normalized watch ratio:  $play\_duration/(2 * video\_duration)$ , capped at 1. This is the recommended relevance signal provided

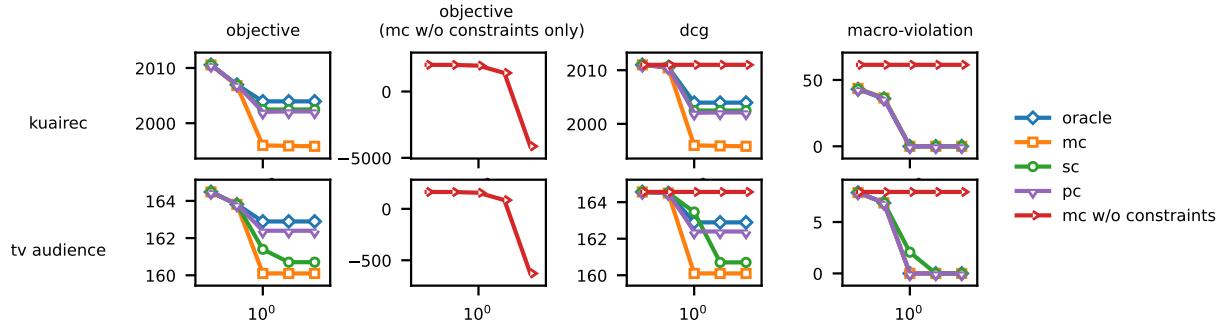
---

#### Algorithm 4: Predictive Controller LP

---

**Input:**  $x_t = (r_t, W_t), s_{t-1}, t$   
**Parameters:**  $\tau \in \mathbb{R}^m, \phi \in \mathbb{R}^m, \gamma \in \mathbb{R}, \{(\widehat{C}_t^b)_{t=1}^T\}_{b=1}^B$   
 $\Sigma_t = \operatorname{argmax}_{\Sigma \in \Delta} r_t^\top \Sigma u + \frac{1}{B} \sum_{b=1}^B \left( \lambda_{t-1}^b \right)_{[0, \phi]}^\top W_t \Sigma e$   
 update  $\lambda_t^b$  from  $\lambda_{t-1}^b$  with  $\gamma$  and gradient  
 $\tau - s_{t-1} - W_t \Sigma_t e - \widehat{C}_t^b$  for  $b = 1, \dots, B$   
**Return:**  $a_t \sim \Sigma_t$

---



**Figure 2: Experiment results comparing all controllers across two datasets KuaiRec and Tv Audience.** The x-axis is  $\phi$  on a log-scale in all plots. The first column is the final objective (6) value, the middle column is the the utility metric (DCG), and the final column is the macro-violation. The oracle has access to the test time contexts and directly optimizes the original objective (6). The MC w/o constraints is an unconstrained utility maximizing controller.

by the dataset publishers<sup>2</sup>. We define an exposure intervention on two arbitrarily chosen groups to evaluate the performance on multi-group constrained ranking. Each group contained two videos, one of which overlapped. In particular, we set the exposure targets to be 1.1 times and 3 times the exposure of the unconstrained controller (MC w/o constraints).

The next dataset we consider is the linear television dataset Tv Audience [42]. This dataset contains temporal television watching behavior for 13k users: the watch duration consists of 217 channels over 19 weeks with an hourly time resolution. For our experiment, we only use the first 12 weeks and ignored the remaining weeks, which results in a total of 288 timeslots. We consider the task of ranking channels over time. The relevance score of a channel during a particular hour is defined as the number of viewers normalized by the channel’s maximum viewers per hour over the past several weeks. To evaluate the temporal prediction capabilities of controllers, we define an exposure intervention on a group of one arbitrarily selected late-night channel that users mostly watch during the night or late evening hours. The intervention is a 100% exposure boost, which is equivalent to setting the exposure target to be twice that of the unconstrained controller (MC w/o constraints).

For our final dataset, we created a fully synthetic dataset to better understand the situations where PC outperforms SC. This dataset consists of temporal patterns of relevance scores for eight items over a horizon of 400 steps. The first four items are always relevant and have the highest relevance scores. The remaining four items are used to define exposure interventions of two disjoint groups. Two of the remaining items form one group and are most relevant during the first half of the time horizon, while the other two items form the second group and are most relevant during the second half of the time horizon. Unlike the previous two datasets, all controllers are trained and evaluated on the same exact relevance scores. This means that we can assume accurate knowledge of the future context distribution when forecasting  $\hat{C}_t$ , which ensures that bad forecasts do not confound the evaluation of the controllers.

**Metrics.** For our experiments, we use discounted cumulative gain (DCG) [23] as  $u$  for the utility metric and reciprocal rank (RR)  $e_k = 1/k$  as our exposure curve [25].

<sup>2</sup>see the data description section on <https://kuairec.com/>

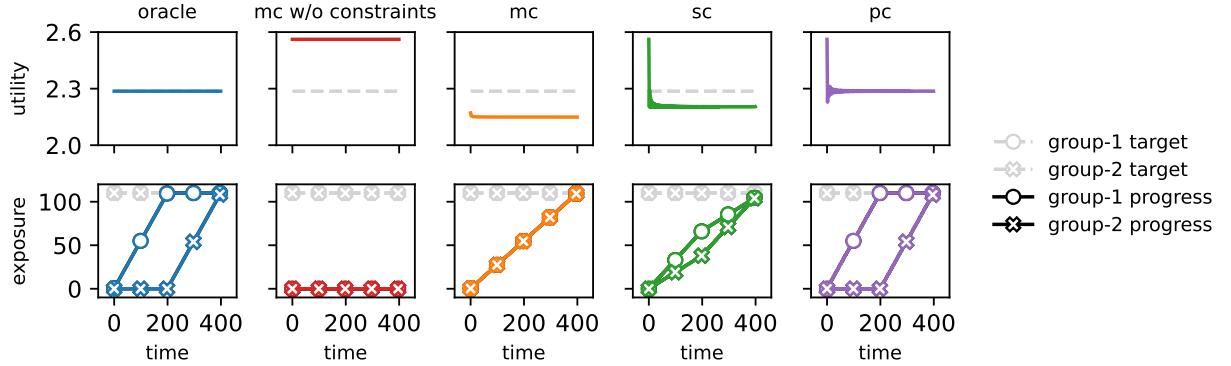
**Hyperparameters.** SC and PC have parameters that must be tuned and estimated to achieve good performance. We tuned these parameters by dividing the data into three sets: train, development, and test. For the PC, we used the train set to estimate the forecast of the progress-to-go and the development set to simulate online contexts as described in Appendix ???. We performed a grid search with the train and development sets to select the best forecast parameter based on the overall objective (6). To update the multiplier  $\lambda$ , we utilized the Adam optimizer [27] for both SC and PC. Additionally, we performed another grid search with the train and development datasets to pick the best Adam hyperparameters based on the overall objective (6). In our experiments, we investigated the performance of each controller for different macro-violation penalty values. We selected the best-performing hyperparameters separately for each controller and penalty value, and we repeated this process for all datasets. Detailed hyperparameter ranges are given in Appendix ???.

## 4.2 Experiment Results

In this section, we evaluate key properties of the controllers using the datasets discussed in the previous section. For all experiments, the x-axis represents the varying macro-violation vector  $\phi$  (ranging from  $1 \times 10^{-2}$  to  $1 \times 10^2$ ), and the y-axis shows the performance of the controllers in terms of the utility (1), violation cost (5) or the overall objective (6).

### 4.2.1 Which controller achieves the highest overall objective?

Figure 2 compares all controllers on both KuaiRec, a stationary dataset, and Tv Audience, a non-stationary dataset with temporal shifts in the context distribution. The performance of the oracle policy provides a skyline because it has (unrealistic) knowledge of all future contexts. We first note that for both datasets, all algorithms are largely equivalent when the macro-violation cost factor  $\phi$  is small since a small cost factor implies little influence on the objective compared to short-term utility maximization. As the importance of the long-term constraints increases with increasing macro-violation cost factor  $\phi$ , the MC performs substantially worse than the other controllers, because its actions are chosen at every time step without consideration of how increasing violations in the current state interact with future contexts. The right middle plots



**Figure 3: Comparison of all controllers on a synthetic dataset to showcase when the PC should be preferred. The o’s and x’s represent two groups of items. The top row is the average utility over time and the dashed grey line represents the highest achievable utility under the exposure constraint. The bottom row displays the exposure over time of both item groups. The grey x’s and o’s represent the target exposure for the groups.**

show, across both datasets, that under an increasing violation cost, all controllers eventually treat the macro-constraints as hard constraints and reduce their violation cost close to zero. However, both SC and PC perform substantially better than MC in terms of the overall objective. Comparing SC and PC on the KuaiRec dataset, we see that both perform about the same. This is to be expected since there is no temporal pattern that PC could learn to exploit with its progress-to-go estimates. Instead, it can only learn to predict the average exposure, which is precisely what the SC is optimizing. On the Tv Audience dataset, which has non-stationary temporal patterns, we see that the PC can plan for the temporal pattern and perform better than all other controllers.

**4.2.2 How sensitive is PC to the number  $B$  of forecast samples?**  
 While the PC performs well on temporal datasets and matches the performance of SC on non-temporal datasets, it has the number of forecast samples  $B$  as an additional parameter that needs to be selected. ?? in the appendix shows the performance of PC depending on the choice of  $B$  for different cost vectors. Note that the results on Tv audience dataset are the median of 20 independent runs since the length of the test set is only 48 introducing a great amount of noise. When the number of forecasts is greater than 20, then PC performs well on both datasets, giving a reference point for selecting  $B$ . However, additional savings in computation time are possible for some datasets since we find that on the KuaiRec dataset, smaller values of  $B$  can suffice to get good performance.

#### 4.2.3 When is it advantageous to use a predictive controller?

From the experiments in Figure 2, we see that PC performs better than other controllers on the Tv Audience dataset, which has a non-stationary temporal pattern. To further illustrate this point, Figure 3 is a salient example that explores when the PC should be preferred over other controllers that are oblivious to any temporal patterns. In the figure, each column represents the controller utility at the top and exposure at the bottom. There are eight items in total and one item in each group. We compute nDCG@4 and RR@4; which means only the top-4 items receive utility and exposure. When no macro-level exposure goals are enforced, as seen in the MC w/o constraints, the exposure for both groups is zero because

they are not among the top-4 ranked items. As illustrated by the oracle controller, the first group is more relevant during the first half of the time horizon, while the second group is more relevant during the second half. The PC performs similarly to the oracle controller because it can leverage this temporal pattern. The SC is better compared to MC due to tuning the gain parameters which allows for additional flexibility. However, both MC and SC follow a linear exposure target and are thus unable to boost the two groups separately.

## 5 CONCLUSIONS

We formalize and address the problem of how to design algorithms that convert macro-level goals into a sequence of individual rankings that have the least impact on micro-level metrics. The algorithms we introduce are analogous to how control is used in mechanical systems to provide layers of abstraction with clear semantics. By introducing three new controllers, we cover a range of application scenarios. Furthermore, we provide rigorous justification for proportional controllers for ranking. Of the three controllers we introduce, we find that controllers based on online optimization (i.e., SC and PC) outperform the more naive MC controller. Furthermore, we find that the predictive controller (PC) performs better than the stationary controller (SC) in non-stationary settings.

This paper opens up a wide range of future work. By making new technical connections between ranking and control theory, it opens up a new set of tools for designing adaptive ranking policies. Furthermore, we anticipate that the macroscopic view of ranking platforms we introduce will provide a conceptual framework for making these platforms more steerable.

## 6 ACKNOWLEDGMENTS

This research was supported in part by NSF Awards IIS-2312865, IIS-2008139, and CCF 2312774; a CI Fellowship; and a gift from Wayfair. We thank Ali Vanderveld and Benjamin Schroeder for their valuable input. All content represents the opinion of the authors, which is not necessarily shared or endorsed by their respective employers and/or sponsors.

## 7 ETHICAL CONSIDERATIONS

Understanding that search and recommendation AI platforms are socio-technical systems means that platform designers must consider the systems' technical and social aspects. In particular, when recommendation systems are being optimized for short-term engagement, it is potentially at the cost of the long-term sustainability of the platform. Not considering long-term sustainability could have social implications like amplifying misinformation or providing disparate utility to different groups. This could affect users' long-term satisfaction with a platform. Our work focuses on developing algorithms that incorporate a platform designer's long-term goals while maintaining a platform's short-term goals. Our intent is that this will enable platform designers to better incorporate all of the socio-technical aspects of a system into the algorithmic decision-making of a search and recommendation AI platform. However, the increased ability to control long-term platform behavior does not automatically lead to better behavior, and responsible governance in setting long-term goals is of crucial importance.

## REFERENCES

[1] M Mehdi Afsar, Trafford Crump, and Behrouz Far. 2022. Reinforcement learning based recommender systems: A survey. *Comput. Surveys* 55, 7 (2022), 1–38.

[2] Deepak Agarwal, Souvik Ghosh, Kai Wei, and Siyu You. 2014. Budget Pacing for Targeted Online Advertisements at LinkedIn. In *Proceedings of the 20th ACM SIGKDD International Conference on Knowledge Discovery and Data Mining* (New York, New York, USA) (KDD '14). ACM, New York, NY, USA, 1613–1619. <https://doi.org/10.1145/2623330.2623366>

[3] Shipra Agrawal and Nikhil R Devanur. 2014. Fast algorithms for online stochastic convex programming. In *Proceedings of the twenty-sixth annual ACM-SIAM symposium on Discrete algorithms*. SIAM, 1405–1424.

[4] Marcin Andrychowicz, Anton Raichuk, Piotr Stanczyk, Manu Orsini, Sertan Girgin, Raphaël Marinier, Leonard Hussetin, Matthieu Geist, Olivier Pietquin, Marcin Michalski, et al. 2020. What matters for on-policy deep actor-critic methods? a large-scale study. In *International conference on learning representations*.

[5] Karl Johan Astroem and Richard Murray. 2020. *Feedback Systems: An Introduction for Scientists and Engineers*. Princeton University Press.

[6] Garrett Birkhoff. 1946. Tres observaciones sobre el álgebra lineal. *Univ. Nac. Tucumán, Ser. A* 5 (1946), 147–154.

[7] Francesco Borrelli, Alberto Bemporad, and Manfred Morari. 2017. *Predictive control for linear and hybrid systems*. Cambridge University Press.

[8] Micah Carroll, Dylan Hadfield-Menell, Stuart Russell, and Anca Dragan. 2021. Estimating and Penalizing Preference Shift in Recommender Systems. In *Fifteenth ACM Conference on Recommender Systems*. 661–667.

[9] L Elisa Celis, Sayash Kapoor, Farnood Salehi, and Nisheeth Vishnoi. 2019. Controlling polarization in personalization: An algorithmic framework. In *Proceedings of the conference on fairness, accountability, and transparency*. 160–169.

[10] L Elisa Celis, Damian Straszak, and Nisheeth K Vishnoi. 2017. Ranking with fairness constraints. *arXiv preprint arXiv:1704.06840* (2017).

[11] Nick Craswell, Onno Zoeter, Michael Taylor, and Bill Ramsey. 2008. An Experimental Comparison of Click Position-bias Models. In *International Conference on Web Search and Data Mining* (WSDM). ACM, 87–94.

[12] Logan Engstrom, Andrew Ilyas, Shibani Santurkar, Dimitris Tsipras, Firdaus Janoos, Larry Rudolph, and Aleksander Madry. 2020. Implementation matters in deep policy gradients: A case study on ppo and trpo. *arXiv preprint arXiv:2005.12729* (2020).

[13] Charles Evans and Atoosa Kasirzadeh. 2021. User Tampering in Reinforcement Learning Recommender Systems. *arXiv preprint arXiv:2109.04083* (2021).

[14] Tom Everitt, Marcus Hutter, Ramana Kumar, and Victoria Krakovna. 2021. Reward tampering problems and solutions in reinforcement learning: A causal influence diagram perspective. *Synthese* 198, Suppl 27 (2021), 6435–6467.

[15] Tom Everitt, Victoria Krakovna, Laurent Orseau, Marcus Hutter, and Shane Legg. 2017. Reinforcement learning with a corrupted reward channel. *arXiv preprint arXiv:1705.08417* (2017).

[16] Miriam Fernández and Alejandro Bellogín. 2020. Recommender Systems and Misinformation: The Problem or the Solution?. In *RecSys 2020 Workshop on Online Misinformation- and Harm-Aware Recommender Systems (CEUR Workshop Proceedings, Vol. 2758)*. Antonela Tommasel, Daniela Godoy, and Arkaitz Zubiaga (Eds.). CEUR-WS.org, 40–50. <http://ceur-ws.org/Vol-2758/OHARS-paper3.pdf>

[17] Chongming Gao, Shijun Li, Wenqiang Lei, Jiawei Chen, Biao Li, Peng Jiang, Xiangnan He, Jiaxin Mao, and Tat-Seng Chua. 2022. KuaiRec: A Fully-observed Dataset and Insights for Evaluating Recommender Systems. In *Proceedings of the 31st ACM International Conference on Information and Knowledge Management* (Atlanta, GA, USA) (CIKM '22). 11 pages. <https://doi.org/10.1145/3511808.3557220>

[18] Yingqiang Ge, Shuchang Liu, Ruoyuan Gao, Yikun Xian, Yunqi Li, Xiangyu Zhao, Changhua Pei, Fei Sun, Junfeng Ge, Wenwu Ou, et al. 2021. Towards long-term fairness in recommendation. In *Proceedings of the 14th ACM International Conference on Web Search and Data Mining*. 445–453.

[19] Thomas Krendl Gilbert, Nathan Lambert, Sarah Dean, Tom Zick, and Aaron Snoswell. 2022. Reward reports for reinforcement learning. *arXiv preprint arXiv:2204.10817* (2022).

[20] Peter Henderson, Riashat Islam, Philip Bachman, Joelle Pineau, Doina Precup, and David Meger. 2018. Deep reinforcement learning that matters. In *Proceedings of the AAAI conference on artificial intelligence*, Vol. 32.

[21] Henning Hohnhold, Deirdre O'Brien, and Diane Tang. 2015. Focusing on the Long-Term: It's Good for Users and Business. In *Proceedings of the ACM SIGKDD International Conference on Knowledge Discovery and Data Mining* (Sydney, NSW, Australia) (KDD '15). ACM, New York, NY, USA, 1849–1858. <https://doi.org/10.1145/2783258.2788583>

[22] Tamas Jambor, Jun Wang, and Neal Lathia. 2012. Using Control Theory for Stable and Efficient Recommender Systems. In *Proceedings of the 21st International Conference on World Wide Web* (Lyon, France) (WWW '12). ACM, New York, NY, USA, 11–20. <https://doi.org/10.1145/2187836.2187839>

[23] Kalervo Järvelin and Jaana Kekäläinen. 2002. Cumulated gain-based evaluation of IR techniques. *ACM Transactions on Information Systems (TOIS)* 20, 4 (2002), 422–446.

[24] Nitin Jindal and Bing Liu. 2007. Analyzing and Detecting Review Spam. In *Seventh IEEE International Conference on Data Mining* (ICDM 2007). 547–552. <https://doi.org/10.1109/ICDM.2007.68>

[25] Thorsten Joachims, Adith Swaminathan, and Tobias Schnabel. 2017. Unbiased learning-to-rank with biased feedback. In *ACM Conference on Web Search and Data Mining (WSDM)*. 781–789.

[26] Niklas Karlsson and Jianlong Zhang. 2013. Applications of feedback control in online advertising. *2013 American Control Conference* (2013), 6008–6013.

[27] Diederik P Kingma and Jimmy Ba. 2014. Adam: A method for stochastic optimization. *arXiv preprint arXiv:1412.6980* (2014).

[28] David Krueger, Tegan Maharaj, and Jan Leike. 2020. Hidden Incentives for Auto-Induced Distributional Shift. *arXiv preprint arXiv:2009.09153* (2020).

[29] Yuanguo Lin, Yong Liu, Fan Lin, Lixin Zou, Pengcheng Wu, Wenhua Zeng, Huan-huan Chen, and Chunyan Miao. 2023. A survey on reinforcement learning for recommender systems. *IEEE Transactions on Neural Networks and Learning Systems* (2023).

[30] localify [n. d.]. Locality: Local Artist and Music Event Recommendation Application. <http://www.localify.org/>. <http://www.localify.org/>

[31] Masoud Mansouri, Hima Abdollahpouri, Mykola Pechenizkiy, Bamshad Mobasher, and Robin Burke. 2020. Feedback Loop and Bias Amplification in Recommender Systems. In *Proceedings of the 29th ACM International Conference on Information and Knowledge Management* (Virtual Event, Ireland) (CIKM '20). ACM, New York, NY, USA, 2145–2148. <https://doi.org/10.1145/3340531.3412152>

[32] Rishabh Mehrotra, James McInerney, Hugues Bouchard, Mounia Lamas, and Fernando Diaz. 2018. Towards a Fair Marketplace: Counterfactual Evaluation of the trade-off between Relevance, Fairness and Satisfaction in Recommendation Systems. In *ACM Conference on Information and Knowledge Management* (CIKM).

[33] Martin Mladenov, Elliot Creager, Omer Ben-Porat, Kevin Swersky, Richard Zemel, and Craig Boutilier. 2020. Optimizing long-term social welfare in recommender systems: A constrained matching approach. In *International Conference on Machine Learning*. PMLR, 6987–6998.

[34] Jamie Morgenstern. 2021. Ensuring equity in online advertising for employment, housing, and credit. Medium, <https://jamiemmt-cs.medium.com/ensuring-equity-in-online-advertising-for-employment-housing-and-credit-82931668c420>.

[35] Marco Morik, Ashudeep Singh, Jessica Hong, and Thorsten Joachims. 2020. Controlling fairness and bias in dynamic learning-to-rank. In *Proceedings of the 43rd International ACM SIGIR Conference on Research and Development in Information Retrieval*. 429–438.

[36] Andrew Y Ng, Daishi Harada, and Stuart Russell. 1999. Policy invariance under reward transformations: Theory and application to reward shaping. In *Icm*, Vol. 99. 278–287.

[37] Andrew Y Ng, Stuart Russell, et al. 2000. Algorithms for inverse reinforcement learning. In *Icm*, Vol. 1. 2.

[38] Alexander Pan, Kush Bhatia, and Jacob Steinhardt. 2022. The effects of reward misspecification: Mapping and mitigating misaligned models. *arXiv preprint arXiv:2201.03544* (2022).

[39] Suresh Sethi and Gerald Thompson. 2022. *Optimal Control Theory: Applications to Management Science and Economics*. Springer. <https://doi.org/10.1007/0-387-29903-3>

[40] Ashudeep Singh and Thorsten Joachims. 2018. Fairness of exposure in rankings. In *Proceedings of the 24th ACM SIGKDD International Conference on Knowledge Discovery & Data Mining*. 2219–2228.

[41] Joar Skalse, Nikolas Howe, Dmitrii Krasheninnikov, and David Krueger. 2022. Defining and characterizing reward gaming. *Advances in Neural Information Processing Systems* 35 (2022), 9460–9471.

[42] Roberto Turrin, Andrea Condorelli, Paolo Cremonesi, and Roberto Pagano. 2014. Time-based TV programs prediction. In *1st Workshop on Recommender Systems for Television and Online Video at ACM RecSys*, Vol. 14.

[43] Southern District of New York United States District Court. 2022. Settlement on housing advertisements between Meta and HUD. <https://www.justice.gov/opa/press-release/file/1514031/download>

[44] Nicolas Usunier, Virginie Do, and Elvise Dohmatob. 2022. Fast online ranking with fairness of exposure. In *Proceedings of the 2022 ACM Conference on Fairness, Accountability, and Transparency*. 2157–2167.

[45] Di Wu, Xiujun Chen, Xun Yang, Hao Wang, Qing Tan, Xiaoxun Zhang, Jian Xu, and Kun Gai. 2018. Budget Constrained Bidding by Model-Free Reinforcement Learning in Display Advertising. In *International Conference on Information and Knowledge Management (CIKM)*. 1443–1451.

[46] Jian Xu, Kuang-chih Lee, Wentong Li, Hang Qi, and Quan Lu. 2015. Smart Pacing for Effective Online Ad Campaign Optimization. In *Proceedings of the 21th ACM SIGKDD International Conference on Knowledge Discovery and Data Mining* (Sydney, NSW, Australia) (KDD '15). ACM, New York, NY, USA, 2217–2226. <https://doi.org/10.1145/2783258.2788615>

[47] Xun Yang, Yasong Li, Hao Wang, Di Wu, Qing Tan, Jian Xu, and Kun Gai. 2019. Bid Optimization by Multivariable Control in Display Advertising. In *Proceedings of the 25th ACM SIGKDD International Conference on Knowledge Discovery and Data Mining* (Anchorage, AK, USA) (KDD '19). ACM, New York, NY, USA, 1966–1974. <https://doi.org/10.1145/3292500.3330681>

[48] Eric Yang Yu, Zhizhen Qin, Min Kyung Lee, and Sicun Gao. 2022. Policy Optimization with Advantage Regularization for Long-Term Fairness in Decision Systems. *arXiv preprint arXiv:2210.12546* (2022).

[49] A. Zarezade, A. De, U. Upadhyay, H. Rabiee, and M. Gomez-Rodriguez. 2018. Steering Social Activity: A Stochastic Optimal Control Point of View. *Journal of Machine Learning Research* (2018).

[50] Meike Zehlike, Francesco Bonchi, Carlos Castillo, Sara Hajian, Mohamed Megahed, and Ricardo Baeza-Yates. 2017. Fa\* ir: A fair top-k ranking algorithm. In *Proceedings of the 2017 ACM on Conference on Information and Knowledge Management*. 1569–1578.

[51] Ruohan Zhan, Konstantina Christakopoulou, Ya Le, Jayden Ooi, Martin Mladenov, Alex Beutel, Craig Boutilier, Ed Chi, and Minmin Chen. 2021. Towards content provider aware recommender systems: A simulation study on the interplay between user and provider utilities. In *Proceedings of the Web Conference 2021*. 3872–3883.

[52] Dell Zhang and Jun Wang. 2021. Recommendation Fairness: From Static to Dynamic. *arXiv preprint arXiv:2109.03150* (2021).

[53] Haoqi Zhang, Junqi Jin, Zhenzhe Zheng, Fan Wu, Haiyang Xu, and Jian Xu. 2022. Control-based Bidding for Mobile Livestreaming Ads with Exposure Guarantee. In *International Conference on Information and Knowledge Management (CIKM)*. 2539–2548.

[54] Weinan Zhang, Yifei Rong, Jun Wang, Tianchi Zhu, and Xiaofan Wang. 2016. Feedback Control of Real-Time Display Advertising. In *Proceedings of the Ninth ACM International Conference on Web Search and Data Mining* (San Francisco, California, USA) (WSDM '16). ACM, New York, NY, USA, 407–416. <https://doi.org/10.1145/2835776.2835843>