

Rethinking Search Engines and Recommendation Systems: A Game Theoretic Perspective

Moshe Tennenholtz

Technion — Israel Institute of Technology
Haifa, Israel
moshet@ie.technion.ac.il

Oren Kurland

Technion — Israel Institute of Technology
Haifa, Israel
kurland@ie.technion.ac.il

ABSTRACT

Search engines and recommendation systems present users with data items such as Web pages, news articles, songs, advertisements, etc. The items are often produced by incentivized parties. For example, many publishers of Web pages want the pages to be highly ranked by search engines in response to queries of interest. Moreover, these parties often respond to decisions made by the system; e.g., changing a Web page so as to promote it in rankings induced by the search engine. Users of the system can also act strategically, intentionally providing the system with misinformation or deliberately adopting specific recommendations so as to alter the system’s future recommendations. The system itself might also be in competition with others (e.g., other search engines) and accordingly become “incentivized” to avoid users’ churn. Thus, settings in which modern search engines and recommendation systems operate are highly dynamic with much of the dynamics driven by strategic players (users, data creators, competitors). Interestingly, the vast majority of search and recommendation algorithms ignore this strategic-behavior based dynamics. We discuss our recent work that shows how this reality results in algorithms which are sub-optimal by design. This fundamental design deficiency calls for a completely new treatment of the involved challenges in the face of strategic environments. We also describe our work that addresses some of these challenges using completely novel approaches, many of which are based on algorithmic game theory and mechanism design.

1 INTRODUCTION

In her popular book, *Weapons of Math Destruction*, data scientist Cathy O’Neil elegantly describes to the general population the danger of the data science revolution in decision making. She describes how the USNews ranking of universities, which orders universities based on 15 measured properties, created new dynamics in university behavior, as they adapted to these measures, ultimately resulting in decreased social welfare. Unfortunately, the idea that data science related algorithms, such as ranking, cause changes in behavior, and that this dynamics may lead to socially inferior outcomes, is dominant in our new on-line economy.

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.

Conference’17, July 2017, Washington, DC, USA
© 2018 Association for Computing Machinery.

Ranking also plays a crucial role in search engines and recommendation systems — two prominent data science applications that we focus on in this paper. Search engines (e.g., Google and Bing) rank Web pages, images and other items in response to a query. Recommendation systems recommend items by ranking them using information induced from some context — e.g., the Web page a user is currently browsing, a specific application the user is running on her mobile phone, the time of day, etc.

In the industrial search and information retrieval community, the practice of adapting content to satisfy a search engine’s ranking function or that of a recommendation engine has been typically referred to as search-engine-optimization (SEO); SEO has become one of the most profitable professions in the data science era. Accordingly, there has been much work on devising algorithms that block spammers and guarantee that content presented to users is of high quality. Yet, virtually all retrieval and recommendation algorithms ignore post-ranking/recommendation effects (i.e., manipulations applied to items) on the corpus. As it turns out, this reality results in under-optimized fundamental paradigms for devising search and recommendation algorithms as we discuss below. For example, a publisher of a Web page can try and mimic other Web pages to promote her page in rankings, thereby potentially causing content changes that are not for the better — e.g., reducing content breadth in terms of topical coverage.

The observations just described call for action. We believe that the design of search and recommendation methods which is central to the Web and to on-line media services, needs to be revolutionized to account for the strategic incentives of the various parties that affect (or affected by) the system. The parties can be publishers/-content providers, the systems themselves (e.g., search engines) which might be in competition with each other, and/or users of the systems.

We are in the process of establishing an entirely new repertoire of incentive-compatible search and recommendation algorithms. These are devised through pioneering the application of game-theoretic mechanism design to these tasks.

Game theory is the branch of mathematics that models multi-agent interactions. Mechanism design is the part of game theory that designs protocols/algorithms for environments consisting of self-motivated participants. Mechanism design has been central in bridging computer science and game theory [28], including wide application to electronic commerce (e.g., [10]), advertising (e.g., [38]) and routing networks (e.g., [33]).

In this paper we survey our recent work on creating theoretical foundations and developing empirically evaluated algorithms to build a fundamental bridge between game theory, and more specifically mechanism design, and search and recommendation.

We argue that it is only natural to model the competitive search and recommendation settings using game theory. For example, considering the system (search engine or recommendation system) as a mediator and those who produce items to be searched or recommended as players entails a suite of specific game settings. Items being highly ranked or recommended amount to a high “payoff” for the owner of the item. Game-theoretic modeling allows to reason about the strategic behavior of players and the situations the setting can converge to (namely, equilibria). Furthermore, the task of devising search or recommendation algorithms becomes a mechanism design (or more specifically, mediator design) problem — a completely new view of these two important tasks.

The incentivized players we mainly focus on in this paper are content providers/publishers who create and manipulate the retrieved/recommended items. In addition, we describe our work on (i) addressing the competition between search engines — i.e., where the engines are the incentivized players — and more generally, competition between prediction algorithms, focusing on regression; and (ii) addressing some of the implications of the strategic behavior of users of recommendation systems in social networks.

The rest of the paper is structured as follows. In Section 2 we discuss some of the fundamental, and far reaching, modeling and algorithmic implications of the competitive and incentivized settings in which search engines operate. Specifically, we discuss post-ranking corpus effects and survey our work on a game-theoretic analysis of the strategic behavior of publishers (content providers) and on addressing ranking robustness. We also briefly describe the challenges of empirical evaluation and analysis in competitive search settings. In Section 3 we focus on our theoretical work on the sub-optimality in competitive settings of the most basic search (ranking) and recommendation principle. In Sections 2 and 3 the incentivized parties in the competitive search and recommendation settings are publishers. In Section 4 we survey our work on settings where the incentivized parties are ranking algorithms and prediction algorithms (regression) that compete with each other. In Section 5 the incentivized players are users of a recommendation system in social networks.

2 SEARCH ENGINES

The main goal of search engines is to rank documents in a corpus by their presumed relevance to the information need expressed by a query posted by a user. There has been a huge body of work on devising relevance estimates for documents. For example, the similarity between term-based vectors representing documents and queries serves as a relevance estimate in the well known vector space model [34]. Modern relevance ranking functions are learned (trained) using a training set which contains queries and corresponding relevance judgments of documents [24, 27]. These approaches allow to integrate various types of relevance estimates.

The theoretical foundation of virtually all ad hoc retrieval paradigms, classical and modern, is the probability ranking principle (*PRP*) [32]: documents are ranked by the probability that they are relevant to the query wherein the probability is estimated using all the information available to the search system. The PRP is the direct basis of the probabilistic retrieval approach [37] and was shown to underlie the basic language modeling approach [22]. Feature-based

learning-to-rank methods can also be viewed as obeying the PRP, even if relevance probability is not directly estimated. That is, these approaches are trained to minimize ranking loss with respect to ground truth ranking and they essentially integrate various relevance estimates so as to improve overall relevance estimation which guides ranking. Neural-network-based retrieval methods estimate relevance (often probability) by learning document and query representations and/or integrating multiple relevance signals [27]. As feature-based approaches they can be viewed as obeying the PRP. The PRP was shown to be optimal under some assumptions; e.g., the independence of document relevance from that of others [32].

2.1 Post-Ranking Dynamics

Careful examination of the PRP, and retrieval methods based on the PRP, reveals a major gap: post-ranking effects are not accounted for; specifically, changes in the corpus documents made by incentivized publishers that respond to the induced ranking so as to improve the ranking of their documents. We note that throughout Sections 2 and 3, we focus on publishers (content providers) as those affected by, and affecting, the dynamics. We do not consider the dynamics driven by users of the systems, e.g., via clickthrough actions. In Section 6 we further discuss the importance, challenges and potential future directions of accounting simultaneously for publisher and user driven dynamics.

To highlight the importance of accounting for post-ranking effects on corpus content consider the following example [2]. A publisher writes a document which contains both common information and information which is unique to the document; i.e., this unique information cannot be found in any other document in the corpus. The publisher, for some reason, is more interested in her document being highly ranked for queries that touch on the non-unique information. However, the ranking function “penalizes” documents having a mixture of information (i.e., not being focused on the information sought for in these queries). In terms of the PRP and current retrieval paradigms the approach of the ranking function is justifiable: one can assume that the user who posted the query would prefer reading documents that are focused on information related to the query rather than having to wade through potentially long documents and distilling out the sought information. Now, suppose that the publisher decides to remove the unique information from the document so as to obtain higher ranking for queries pertaining to the common information. This means that valuable information in the corpus is lost, and search effectiveness for queries which target the unique information will consequently degrade. Overall, the potential satisfaction of users’ information need with respect to the corpus will decrease.

Indeed, we showed, using a game theoretical analysis, that the PRP is sub-optimal as it does not promote content breadth in the corpus in terms of topical coverage [2]. In other words, retrieval methods based on the PRP do not provide a strong enough incentive for publishers to produce diversified content. It also turns out [2] that introducing randomness to ranking functions can help to promote content breadth in the corpus, and hence, increase the overall attainable utility for search engine users. We further discuss the sub-optimality of the PRP and the merits of using non-deterministic ranking functions in Section 3.

Given the significant impact of induced rankings on the content in the corpus, due to the actions employed by incentivized publishers who respond to these rankings, the natural challenge that rises is analyzing the *strategic behavior* of publishers. Previous work has characterized search engine optimization (SEO) techniques intended to promote documents in rankings [15]. We also note that content dynamics on the Web was studied and analyzed (e.g., [30, 35]). However, the post-ranking perspective has not been actually modeled or analyzed; i.e., the specific types of responses of publishers to rankings, and more generally, their ranking-motivated strategies in terms of document manipulation, were not studied.

2.1.1 Strategic Publishers. Since the early days of the Web, different types of SEO techniques have been identified [15]. For example, publishers can “stuff” keywords in their Web pages so as to promote them in rankings induced for queries that include these keywords. The underlying (often correct) assumption is that increased similarity between the query and the Web page increases the retrieval score of the page, and hence improves its potential rank position. On the other hand, compromising the quality of the page by filling it with potentially query relevant keywords can also reduce the retrieval score as ranking functions often utilize page-quality estimates [6].

There is a fundamentally important question which goes beyond the actual *general* actions that publishers use as part of their SEO efforts: what is the strategic behavior of the publishers with respect to induced rankings? In other words, given that they do not know what the ranking function is¹, but they can observe past rankings induced for queries of interest, what would be an effective response strategy to rankings?² We have recently addressed this question using a game theoretic analysis [31]. Our main theoretical result was that a “worthwhile” strategy for publishers who want to promote their documents in rankings induced for a specific query is to make their documents become more similar to those highly ranked in the past for the query. By “worthwhile” we refer to the fact that this strategy results in a certain equilibrium in a game-theoretic modeling of the retrieval setting wherein publishers are players and the ranking function is a mediator, not exposed to the players, which induces rankings for queries.

The intuitive theoretical finding that mimicking the “winner” from previous rankings is worthwhile was supported by analyzing the strategic publishing behavior of students who served as publishers in a content-based ranking competition we organized. In Section 2.3 we provide details of this competition.

2.2 Addressing Unwarranted Effects of the Dynamics

As discussed, a major part of the dynamics of a corpus in competitive retrieval settings is driven by incentivized publishers. The dynamics can have undesirable effects, specifically, in terms of degrading retrieval effectiveness. For example, some Web pages could be spam that is intended to be promoted in rankings and to attract

clicks — i.e., *black hat* SEO [15]. At the same time, corpus dynamics is driven to a major extent by *white-hat* SEO efforts which need not necessarily hurt retrieval effectiveness; these are legitimate actions applied to Web pages so as to promote them in rankings. However, even white-hat SEO can lead to undesirable effects; e.g., rapid changes to the relative ranking of documents due to indiscernible document manipulation. As a result, users might respond by consistently reformulating their queries or simply losing faith in the search engine.

2.2.1 Ranking Robustness: A Blessing or a Curse? Following the arguments just posed, one should presumably opt for *ranking robustness*; that is, small indiscernible changes of documents should not result in major changes to induced rankings [14]. Some support for this argument can be drawn using one of the most fundamental hypotheses in the field of information retrieval, namely, the *cluster hypothesis* [18] as we recently described [14]. Additional support can be drawn from arguments made in work on adversarial classification, specifically, in the vision domain. The main premise in this line of work was that small indiscernible (adversarial) changes of objects should not result in changes to classification decisions (cf., [11, 13]). Thus, while changes of rankings in our setting that reflect discernible changes of, or differences between, documents are naturally very important, changes of rankings that are “harder to explain” (a.k.a, explainable IR) are less warranted.

On the other hand, there are arguments for hurting ranking robustness, to some extent, for special purposes. For example, recent work showed that introducing randomization to induced rankings along time, which hurts robustness, can increase fairness with respect to publishers of Web pages [7]. Specifically, given that users browsing the search results page mainly pay attention to the top ranked results [19], allowing Web pages to be positioned at the highest ranks even if they are not among the most relevant, can increase the attention given to other publishers and hence promote fairness. Another justification for hurting ranking robustness is promoting content breadth in the corpus, specifically, by introducing randomization to ranking functions, as we recently showed using a game theoretic analysis [2].

Considering the findings described above, there is a tradeoff between ranking robustness and content breadth and/or ranking fairness. Accounting for all these aspects simultaneously is an intriguing avenue for future work.

2.2.2 Robustness of Ranking Functions: Regularization and Beyond. In recent work [14], we presented initial theoretical foundations for the analysis of the robustness of ranking functions to (adversarial) document manipulations intended for promoting the documents in rankings. While the goal was indeed to study adversarial effects on robustness, the formal treatment was more general and accounted for any type of a change.

An important theoretical result to which we provided empirical support was the connection between *regularization* — a suite of approaches intended to improve the effectiveness of a ranking function over queries not seen during the training phase — and ranking robustness. We showed that the “stronger” the regularization, the more robust the induced rankings are to document manipulations.

¹There are efforts to reverse engineer ranking functions but these are often of quite limited success.

²A response in terms of manipulating document content can be, for example, at the “micro-level”: selecting the terms to add or remove, or at the “macro-level”: making a document more similar to another document [31].

Controlling ranking robustness by tuning regularization is only a first initial step. Exploring other fundamental approaches to improving, or more generally, controlling ranking robustness is a completely open research avenue. While existing relevance ranking functions are trained to improve effectiveness or minimize reductions in effectiveness with respect to that of known ranking functions [39], ranking robustness as a manifestation of post-ranking effects is a missing component in such approaches. The treatment of robustness of ranking functions is also important, as mentioned above, given the emergence of recent work on improving ranking fairness by damaging robustness [7]. One would strive to simultaneously maximize fairness and minimize ranking instability.

2.3 Empirical Analysis and Evaluation

The Web is the canonical example of a competitive retrieval setting. Specifically, many publishers of Web pages have an incentive (e.g., financial) to have their pages highly ranked for queries they care about. Isolating Web dynamics associated with incentive driven changes to documents is difficult since it occurs on a backdrop of constant Web page changes and dynamics that are not driven by ranking incentives.

To allow the study of corpus dynamics that results from the strategic behavior of publishers in response to rankings, as well as to devise new retrieval paradigms that account for this type of dynamics and its effects, controlled experiments are called for. For example, we recently reported a ranking competition held between students in a class [31]. The students were instructed to write plain text documents with the aim of being highly ranked for given queries. That is, the competition was focused on content-based manipulation. The incentive of students to take part in the competition and produce documents that would presumably be ranked high was bonus points for the course grade if they were to win many rounds of the competition; i.e., if the documents they created were highly ranked. At the end of each round of the competition, the students were shown the rankings induced for queries over the documents they have created/changed. That is, the only information available for “publishers” in this competition, as is the case over the Web, is the query, the induced ranking and the ranked documents.

The competition was not large scale and focused on a specific (yet important) aspect of manipulation — namely, content manipulation. Nevertheless, as mentioned in Section 2.1, it provided empirical support to a game theoretic result we presented about the strategic behavior of publishers in response to induced rankings [31]: publishers make their documents similar to those which were highly ranked for the same queries in previous rounds of the competition.

Organizing ranking competitions at large scale is extremely difficult, specifically, if various aspects are to be studied simultaneously — e.g., the manipulation of both content and hyperlink structure. Furthermore, to develop novel ranking functions, one has to run A/B testing between old and new functions in real time — i.e., to have publishers participating in the game respond, by manipulating their documents, to the ranker which is applied. Thus, analysis of ranking competitions between incentivized publishers, as well as evaluation of novel ranking methods that address the resulting corpus dynamics, is a challenge at its own right.

3 NON-OPTIMALITY OF CLASSICAL SEARCH AND RECOMMENDATION METHODS

As already noted above, the probability ranking principle (PRP) is the theoretical foundation of most retrieval approaches that rank documents in response to a query. It is, in essence, also the theoretical foundation of recommendation methods that rank and recommend more appropriate items. Accordingly, aside for potential exploration intended to collect data for training, search and recommendation methods based on the PRP are deterministic. As it turns out, when the publishers of items to be searched or recommended are strategic, this classical deterministic approach has to be replaced. In Sections 3.1 and 3.2 we consider the PRP and approaches devised based on the PRP for search (retrieval) and recommendation, respectively, in settings with strategic publishers.

3.1 Beyond PRP: Randomness Increases Social Welfare in Search

One of the important goals of analysis of a competitive retrieval setting with corpus dynamics driven by incentivized (strategic) selfish players — publishers in our case — is to derive insightful statements about the steady states that the setting can converge to, if at all. It is important to note that here, dynamics refers to that of documents in the corpus upon which search is performed. Dynamics of clickthrough patterns³ and that of users’ queries [40] is outside the scope of this paper. To address this goal, it is only natural to take a game theoretic approach which allows to reason about the different equilibria the search setting can converge to. Equilibrium is a steady state wherein players (publishers) have no incentive to deviate from the current actions (strategies) they employ⁴. The fundamental question then becomes how to model a retrieval setting as a game.

We have recently modeled competitive retrieval settings as games [2, 3]. Publishers are players whose strategies are the types of documents they can write: either single-topic or multi-topic documents. The search engine’s ranking function is the mediator which induces ranking. A publisher is rewarded if her document is the highest ranked for a query. To simplify the analysis, we assumed that each query is about a single topic and that the distribution over queries is known to the publishers. We then examined two settings: (i) all publishers have the same equal writing quality over all topics, and (ii) publishers have differential writing quality for different topics.

The type of documents publishers can write and the assumption regarding their writing quality entails a specific game. We assumed that the ranking function has full knowledge of whether a document is relevant to a query (specifically, by topic matching) and analyzed the equilibria of the game, i.e. publishers’ behavior in which modification of behavior will not be any unilaterally beneficial. We then analyzed the *social welfare* attained in the game, specifically, in various equilibria. Social welfare was defined as the sum of utilities attained by publishers. We assumed that users’ utilities and publishers’ utilities were aligned and those were determined by whether the highest ranked document was relevant. A simple utility function was used: 1 if the highest ranked document

³Clickthrough refers to clicks of users on the results presented by the search engine.

⁴These strategies can be probability distributions (a.k.a., mixed strategies) over pure strategies.

is relevant and 0 otherwise [2, 3]. This utility function corresponds to a setting where a user examines only the top ranked document and her utility solely depends on the relevance of this document. The alignment between publishers’ and users’ utilities facilitates the formal treatment and corresponds to the assumption that if a user is satisfied by seeing a relevant document then the publisher of the document is rewarded to the same extent. In practice, however, publishers’ and users’ utilities are not necessarily aligned and more evolved models are thereby called for [2, 3].

The games were further analyzed by computing the *price of anarchy* [20, 33]: the ratio between the maximal social welfare that can be attained in a game and the minimal social welfare attained in any equilibrium of a game. In other words, price of anarchy is the price the echo system “pays” for the selfish behavior of players. Appendix A describes an example of a game and its analysis.

We found [2, 3] that if publishers write single-topic documents and they have equal writing qualities for all topics then the PRP is indeed an optimal criterion for ranking. But, if publishers write multi-topic documents or have differential writing qualities the PRP becomes sub-optimal in terms of price of anarchy. Appendix A provides a simple example that demonstrates the sub-optimality of the PRP. Indeed, the PRP motivates publishers to write single-topic documents, or documents only on the topics they are most qualified to write, as their resultant ranking would be higher. Thus, the PRP essentially does not promote content breadth in the corpus — in terms of topical coverage — in these natural settings. In Section 2.1 we described a canonical example of this situation: the publisher of a multi-topic document, with information unique to the document, is driven by the PRP to remove the unique information so as to better focus the document for higher ranking for queries of interest.

It turns out that introducing randomization to the ranking functions can improve the price of anarchy with respect to that attained by using the PRP as a ranking criterion, as randomization can help to promote content diversity [2, 3]. A case in point, promoting in rankings a multi-topic document with respect to a single topic document, in case the retrieval score of the latter is a bit higher than that of the former, does not significantly harm per-query retrieval effectiveness but does result in improved price of anarchy (and social welfare) due to increased content breadth in the corpus.

These findings are especially important because they imply that the most fundamental principle used by existing retrieval methods and search engines for ranking (the PRP) is sub-optimal in competitive retrieval settings. Furthermore, these findings motivate a new view of how ranking functions should be learned: not (only) for minimizing a “local” loss function but rather maximizing the resultant social welfare attained in equilibria. This is a brand new challenge for the information retrieval community, and other communities interested in ranking, as the task involves estimating post-ranking effects. Our findings about the merits of applying non-deterministic ranking functions is a first small step in this direction.

3.2 Characterizing a Desired Fair Recommendation System

Heretofore, we have mainly focused on search (retrieval) systems. Recommendation systems (RSs hereinafter) have also rapidly developed over the past decade. By predicting a user preference for

an item, RSs have been successfully applied in a variety of applications. However, in order to address significant applications, where information is integrated from many sources, such systems have to generate recommendations that satisfy the needs of both the end users and other parties, and in particular content providers (publishers). Many of the content providers provide sponsored or partially sponsored content that may be relevant to a user, and an RS that will not be fair with content providers, will not survive.

Hence, an RS should be able to deal in a fair way with strategic content providers, each of which aims at maximizing its exposure. In addition, if an RS is not designed properly, then content providers may reach a fluctuating system where content providers rapidly change the type of content they provide only to defeat their competitors. Therefore, the RS need not only be fair but to also yield stability, i.e., convergence to equilibrium of content providers’ strategies. These requirements from an RS were only very recently formulated in a rigorous manner [5]. It has been shown that classical RSs, which operate in the PRP style (i.e., always show the most relevant content) fail to satisfy the above requirements (i.e., will be unfair or lead to fluctuations). Indeed, the latter work also shows a particular RS, selecting among content providers in a well defined probabilistic manner, that does satisfy the related requirements, and in a sense is the only recommendation system that satisfies them. Technically, the results use a classical concept in cooperative game theory, the Shapley value [36], which measures the contribution of parties to a cooperative act, and adapt it to serve as an efficient mechanism for probabilistic selection among content providers.

These findings complement the observation about the failure of the most fundamental principle used by existing retrieval methods and search engines for ranking (the PRP), by showing that the corresponding method also fails in the context of recommendation systems. It is sub-optimal when considering strategic content providers. Indeed, well-defined and well-grounded probabilistic substitutes are offered.

4 DUELING

So far we have dealt with the design of systems, a.k.a. mechanisms or mediators, while taking into account that participants in these systems (specifically, publishers/content generators) are self-motivated. Indeed, the role of a search engine, as well as the role of a recommendation system, which integrate information from different content providers, is to serve as a mediator among selfish participants. However, such mechanisms may also face a competition with other mechanisms as we discuss below. Thus, we now turn to describe models that address such competitions (specifically, between search engines/ranking functions and between prediction algorithms). In these settings, the systems themselves can be thought of the players, in contrast to the settings discussed thus far where publishers were the players.

Consider two competing search engines. A user thinks of a desired Web page and submits a query intended to target such a page to both search engines. The engine that ranks the desired page higher is chosen by the user as the “winner.” Given a probability distribution over the users as for their desired pages given that query, the optimal algorithm ranks the pages, $\omega_1, \omega_2, \dots, \omega_n$, in

a decreasing order of the corresponding probabilities; this is essentially the probability ranking principle (PRP) discussed above. However, in a competitive setting the ranking $\omega_2, \omega_3, \dots, \omega_n, \omega_1$ beats (assuming the proportion of users interested in ω_1 is less than 0.5) on every item that may be searched except for ω_1 . This is a highly simplified example of the fact that different data science mechanisms are in competition with one another, and therefore the optimal way to design them should be re-considered.

Such setting was formalized first in [16], under the title *dueling algorithms*. The most basic building block in such settings is to consider a state-of-the-art algorithm (e.g., a ranker in the information retrieval setting, a regressor in a prediction setting) and see if it can be defeated when put in a competitive setting; i.e., whether one can provide an algorithm that most users will find more compelling than the classical one, when they have to choose. We now discuss two achievements that have been obtained in this regard.

4.1 Best-Response Regression

An interesting extension of the above general discussion is in the context of Machine Learning, and in particular regression tasks. In a regression task, a predictor is given a set of instances (points) along with a real value for each point. Subsequently, she has to identify the value of a new instance as accurately as possible. In [4] the authors consider a regression task tackled by two players, where the payoff of each player is the proportion of the points she predicts more accurately than the other player. On the technical side, in order to do so, the authors revise the probably approximately correct (PAC) learning framework to deal with the case of a duel between two predictors. Then they devise an algorithm which finds a linear regression predictor that is a best response to any (not necessarily linear) regression algorithm. This algorithm has linearithmic sample complexity, and polynomial time complexity when the dimension of the instances domain is fixed. The approach has been also tested in a high-dimensional setting, and it has been shown to significantly defeat classical regression algorithms in the prediction duel.

4.2 Responding to a Strong Ranker

Returning back to the search and ranking motivation, here is a fundamental interesting question. Say that a search engine has a relatively weak (somewhat ineffective) relevance ranking function; specifically, with respect to that of another (strong) search engine. This could be for example due to a relatively small user base: indeed, user engagement signals are important features in relevance ranking functions.

Thus, an interesting duel between the two search engines emerges. The duel entails an interesting bi-modal utility function for the weak search engine: the goal is to produce in response to a query a document result list whose effectiveness is not significantly lower than that of the strong search engine; and, the result list should also be quite different than that produced by the strong engine. In [17] we presented a per-query algorithmic approach that leverages fundamental retrieval principles such as pseudo-feedback-based relevance modeling [23]. That is, the weak ranker observes the documents most highly ranked by the strong ranker, uses them to induce a model of relevance for the query, and uses this model for

ranking. We demonstrated the merits of our approach in improving the search effectiveness of the weak ranker using TREC data⁵ — a suite of datasets used for evaluation of retrieval methods.

5 STRATEGIC USERS

The previous sections emphasized the need to account for the incentives of strategic parties when devising major data science applications such as search (ranking), recommendation, and prediction (specifically, regression). The strategic parties were associated with stakeholders such as content providers or prediction experts. The users of the systems were not strategic; they were the ones posting a query, asking for recommendation, or interested in prediction. In a sense, the users were the products the strategic parties were aiming to attract or to serve better than their competitors. However, users may have their own incentives which might cause them to potentially not follow recommendations by the system or to provide dishonest inputs to the system. We next briefly describe one of our recent works that addresses such aspects.

5.1 Incentive-Compatible Explore and Exploit

Recommendation systems that are based on the interleaving of exploration and exploitation paradigm have a two way relationship with their customers — on the one hand they provide recommendations while on the other hand they use customers as their source of information. This reality leads to an important challenge: a user might not accept a recommendation if he/she believes the recommendation is done to benefit exploration rather than be the optimal one given current information. The challenge is to devise a system which will behave close to optimal, but will be also incentive-compatible, i.e., rational users who care only for their expected utility will accept the system’s recommendations.

We have recently addressed this challenge in designing recommendation systems, specifically for social networks, where users can (partly) observe each other [1]⁶. In particular, we investigated the conditions on the social network which allow for asymptotically optimal outcomes. Our results show that for reasonable networks, where each user can observe many of the other users, but where still most users can not see most of the other users, an incentive-compatible and approximately optimal recommendation system does exist.

The literature on incentivizing exploration relevant to our study can be viewed as extension of the celebrated multi-armed bandit problems ([8]) to deal with settings where exploration can only be done by self motivated myopic agents and a central planner must incentivize exploration. The tension between the objective of the mediator (the recommendation engine) and the individual agents in a multi-armed bandit context was first introduced in [21] who study this in a very simple setting. They identify an incentive compatible scheme with which a central mediator with commitment power can asymptotically steer the users towards taking the optimal action. This exciting news has been extended in [25] and [26] to several more elaborate bandit settings and to additional optimization criteria such as regret minimization. Whereas both of these papers account for agents’ incentives and in particular the

⁵<https://trec.nist.gov/>

⁶The pioneering work on this challenge assumed no user communication [9, 21].

misalignment of incentives of the agents and the mediator they ignore other societal aspects. In particular, these papers make an implicit assumption that agents cannot see nor communicate with any other agent. Needless to say, the assumption that agents have no knowledge of each other, and cannot see the actions chosen by their neighbors, is unrealistic. Che and Horner [9] also study mechanisms for social learning through a mediator. Similar to [21], they assume agents are short lived and do not observe each other and the mechanism must incentivize them to explore. In contrast with [21] the model in [9] is one of continuous time.

Unfortunately, the results obtained in [21], [25] and [26] are not robust to changes in such network assumptions and their implicit assumption of no visibility turns out to be critical. In fact, even with very little observability, for example when each agent just sees the action chosen by his immediate predecessor, the schemes proposed in both works cease to be incentive compatible and lead to market failure. The challenge of approximately optimal incentive compatible schemes subject to partial observability among the agents was addressed in our work [1]. In that paper it is assumed that agents can view actions chosen by some of their predecessors. We characterize structural properties that allow for the design of optimal schemes and provide explicit constructions. In contrast with the recent literature relevant to search and recommendation where payments are avoided, other works study payments to agents as a means to incentivize exploration (e.g., [12])

6 CONCLUSIONS AND LOOKING FORWARD

Nowadays, many search engines and recommendation systems operate in dynamic settings wherein dynamics is driven by interested parties. The parties can be those that generate items to be retrieved or recommended, the users themselves or competing systems.

It is thus somewhat surprising that there has been little to no work on modeling the incentives of parties involved, and more generally, on devising algorithms that account for the incentive-driven dynamics of the ecosystem they operate in. We discussed our results regarding the consequences of ignoring the dynamics in several cases: the attained solutions are, by design, sub-optimal. In particular, we showed that the most basic principle underlying most query-based ranking (retrieval) algorithms and recommendation methods is sub-optimal in competitive settings.

Accordingly, we argued for a need for revolution in the way search and recommendation algorithms are devised. We suggested game theory, and more specifically mechanism design, as a framework for devising the algorithms. The basic idea is to treat interested/incentivized parties as players and the data system/algorithm as a mediator. The challenge is that of mechanism (mediator) design: devising an algorithm that accounts for the environment which consists of self-motivated parties.

An interesting question is whether the game theoretic models constitute effective approximations for real user behavior. In our recent work we demonstrated such an alignment [31]: the publishing strategies of publishers were aligned with the game-theoretic results; namely, mimicking documents which were highly ranked in the past. Studying the robustness of game-theoretic solutions in the face of irrational behavior of users is a direction we intend to explore in future work.

Another important aspect is the overhead in terms of system and algorithmic complexity when accounting for incentives and applying game theoretic principles. We argue that the overhead need not be significant. For example, the non-deterministic ranking functions we proposed in Ben-Basat et al. [2, 3] improve social welfare with respect to the PRP and practically post no algorithmic or system-implementation overhead. Similarly, our approaches for improving ranking robustness [14] and for dueling of search engines [17] incur no significant overhead.

Our on-going and ultimate operational goal is to devise new, improved, search and recommendation algorithms via mechanism design. We surveyed a few such examples, but there is much more work to be done in that respect. We envision a modular development of game-theoretic-based search and recommendation solutions, as was and is the case for current solutions which ignore incentives. First, fundamental relevance estimates for search and recommendation will be improved by mechanism design and these can be integrated — e.g., via learning-to-rank — with other estimates for which incentive issues are not accounted for. At the macro level, one can devise loss functions for search and recommendation functions that account for incentives. These can be devised in some cases, as we recently showed for ranking robustness [14], regardless of the features used. Overall, such modular development allows, in many cases, to separate the development of none-game-theoretic aspects from those which rely on game theory.

The work we surveyed (e.g., on non-deterministic ranking functions [2, 3]) provides some initial guidelines about how one would devise solutions that account for incentives. We believe that the suite of guidelines will broaden with more work on these directions, as is the case, for example, for neuro IR: current state-of-the-art neuro IR algorithms are significantly different from those proposed a few years ago given the experience which has been gained.

The spectrum of challenges remained to be addressed is huge. For example, accounting for the interactions between different types of incentivized players involved — e.g., publishers and users of a search engine. Indeed, users of the search engine are not only incentivized consumers of retrieved items, but also interested parties that affect the search engine’s behavior (specifically, its ranking function) via their interactions. That is, in commercial search engines, users’ clicks on retrieved results are treated as implicit relevance feedback [19]. Query reformulations are an additional implicit relevance feedback signal. Thus, modeling user interactions with the search system is an important part of the analysis of any search engine and the design of retrieval algorithms. In our setting, the emerging challenge is modeling simultaneously, using a game theoretic approach, the effects of the strategic behavior of interested users with that of interested publishers. For example, to contrast ranking mechanisms in terms of utility and social welfare, we can potentially use Bayesian games, rather than perfect information games which we used in our initial work [2, 3], wherein clicks and their dynamics are used to devise relevance estimates.

There are additional numerous “micro-level” challenges. Examples include (i) devising content-based and other (e.g., hyperlink-based) relevance estimates that account for incentivized publishers, (ii) moving from list-based effectiveness evaluation to amortized and equilibrium based evaluation (cf. [3, 7]) given the dynamics of the retrieval setting, and (iii) adapting specialized retrieval approaches

(e.g., those based on fusion or federation) to an incentives-based setting.

Our study is part of a more general attack on applying mechanism design for data science (see <http://mdds.net,technion.ac.il>). Indeed, while algorithmic game theory had remarkable success in addressing incentives in variety of settings, such as a variety of fundamental resource allocation settings (e.g. auctions), and congestion control, among others [28], both the repertoire of games and the techniques to be employed need to be expanded in order to deal with data science algorithms. While in this paper we referred mainly to search and recommendation systems, the above mentioned site shows applications to prediction, regression, segmentation, diffusion, and others. Two of the major techniques used are around approximate mechanism design without money [29] and dueling algorithms [16], which their general discussion goes beyond the scope of this paper.

Acknowledgments. We thank the reviewers of CACM for their helpful comments. The support of EU project 740435 - MDDS is gratefully acknowledged.

REFERENCES

- [1] Gal Bahar, Rann Smorodinsky, and Moshe Tennenholtz. 2016. Economic Recommendation Systems: One Page Abstract. In *Proceedings of EC*. 757–757.
- [2] Ran Ben-Basat, Moshe Tennenholtz, and Oren Kurland. 2015. The Probability Ranking Principle is Not Optimal in Adversarial Retrieval Settings. In *Proceedings of ICTIR*. 51–60.
- [3] Ran Ben-Basat, Moshe Tennenholtz, and Oren Kurland. 2017. A Game Theoretic Analysis of the Adversarial Retrieval Setting. *J. Artif. Intell. Res.* 60 (2017), 1127–1164.
- [4] Omer Ben-Porat and Moshe Tennenholtz. 2017. Best Response Regression. In *Proceedings of NIPS*. 1498–1507.
- [5] Omer Ben-Porat and Moshe Tennenholtz. 2018. A Game-Theoretic Approach to Recommendation Systems with Strategic Content Providers. In *Proceedings NeurIPS*. 1118–1128.
- [6] Michael Bendersky, W. Bruce Croft, and Yanlei Diao. 2011. Quality-biased ranking of web documents. In *Proceedings of WSDM*. 95–104.
- [7] Asia J. Biega, Krishna P. Gummadi, and Gerhard Weikum. 2018. Equity of Attention: Amortizing Individual Fairness in Rankings. In *Proc. of SIGIR*. 405–414.
- [8] Sébastien Bubeck and Nicolò Cesa-Bianchi. 2012. Regret Analysis of Stochastic and Nonstochastic Multi-armed Bandit Problems. *Foundations and Trends in Machine Learning* 5, 1 (2012), 1–122. <https://doi.org/10.1561/22000000024>
- [9] Y.K Che and J.O. Horner. 2015. Optimal Design for Social Learning. (2015). Cowles Foundation Discussion Paper No. 2000.
- [10] Peter Cramton, Yoav Shoham, and Richard Steinberg. 2007. An overview of combinatorial auctions. *SIGecom Exchanges* 7, 1 (2007), 3–14.
- [11] Nilesh Dalvi, Pedro Domingos, Mausam, Sumit Sanghai, and Deepak Verma. 2004. Adversarial Classification. In *Proceedings of KDD*. 99–108.
- [12] Peter I. Frazier, David Kempe, Jon M. Kleinberg, and Robert Kleinberg. 2014. Incentivizing exploration. In *ACM Conference on Economics and Computation, EC '14, Stanford, CA, USA, June 8-12, 2014*. 5–22.
- [13] Ian J. Goodfellow, Jonathon Shlens, and Christian Szegedy. 2015. Explaining and Harnessing Adversarial Examples. In *Proc. of ICLR*.
- [14] Gregory Goren, Oren Kurland, Moshe Tennenholtz, and Fiana Raiber. 2018. Ranking Robustness Under Adversarial Document Manipulations. In *Proceedings of SIGIR*. 395–404.
- [15] Zoltán Gyöngyi and Hector Garcia-Molina. 2005. Web Spam Taxonomy. In *Proceedings of AIRWeb 2005*. 39–47.
- [16] Nicole Immerlica, Adam Tauman Kalai, Brendan Lucier, Ankur Moitra, Andrew Postlewaite, and Moshe Tennenholtz. 2011. Dueling algorithms. In *STOC*. 215–224.
- [17] Peter Izsak, Fiana Raiber, Oren Kurland, and Moshe Tennenholtz. 2014. The search duel: a response to a strong ranker. In *Proceedings of SIGIR*. 919–922.
- [18] N. Jardine and C. J. van Rijsbergen. 1971. The use of hierarchic clustering in information retrieval. *Information Storage and Retrieval* 7, 5 (1971), 217–240.
- [19] Thorsten Joachims, Laura Granka, Bing Pan, Helene Hembrooke, and Geri Gay. 2005. Accurately interpreting clickthrough data as implicit feedback. In *Proceedings of SIGIR*. 154–161.
- [20] E. Koutsoupias and C. Papadimitriou. 1999. Worst-Case Equilibria. In *STACS*.
- [21] Ilan Kremer, Yishay Mansour, and Motty Perry. 2014. Implementing the wisdom of the crowd. *Journal of Political Economy* 122 (2014), 988–1012.
- [22] John Lafferty and Chengxiang Zhai. 2003. Probabilistic Relevance Models Based on Document and Query Generation. In *Language Modeling and Information Retrieval*. Kluwer Academic Publishers, 1–10.
- [23] Victor Lavrenko and W. Bruce Croft. 2001. Relevance-Based Language Models. In *Proceedings of SIGIR*. 120–127.
- [24] Tie-Yan Liu. 2011. *Learning to Rank for Information Retrieval*. Springer. I–XVII, 1–285 pages.
- [25] Yishay Mansour, Aleksandrs Slivkins, and Vasilis Syrgkanis. 2015. Bayesian Incentive-Compatible Bandit Exploration. In *Proceedings of EC*.
- [26] Yishay Mansour, Aleksandrs Slivkins, Vasilis Syrgkanis, and Zhiwei Steven Wu. 2016. Bayesian Exploration: Incentivizing Exploration in Bayesian Games. *CoRR abs/1602.07570* (2016). arXiv:1602.07570 <http://arxiv.org/abs/1602.07570>
- [27] Bhaskar Mitra and Nick Craswell. 2017. Neural Models for Information Retrieval. *CoRR abs/1705.01509* (2017). <http://arxiv.org/abs/1705.01509>
- [28] Noam Nisan, Tim Roughgarden, Eva Tardos, and Vijay V. Vazirani. 2007. *Algorithmic Game Theory*. Cambridge University Press.
- [29] A. Procaccia and M. Tennenholtz. 2009. Approximate Mechanism Design without Money. In *EC-09*.
- [30] Kira Radinsky and Paul N. Bennett. 2013. Predicting content change on the Web. In *Proceedings of WSDM*. 415–424.
- [31] Nimrod Raifer, Fiana Raiber, Moshe Tennenholtz, and Oren Kurland. 2017. Information Retrieval Meets Game Theory: The Ranking Competition Between Documents' Authors. In *Proceedings of SIGIR*. 465–474.
- [32] Stephen E. Robertson. 1977. The Probability Ranking Principle in IR. *Journal of Documentation* (1977), 294–304. Reprinted in K. Sparck Jones and P. Willett (eds), *Readings in Information Retrieval*, pp. 281–286, 1997.
- [33] T. Roughgarden and E. Tardos. 2002. How bad is selfish routing? *J. ACM* 49, 2 (April 2002), 236–259.
- [34] Jerard Salton, Anita Wong, and Chung Shu Yang. 1975. A vector space model for automatic indexing. *Commun. ACM* 18, 11 (1975), 613–620.
- [35] Aécio S. R. Santos, Bruno Pasini, and Juliana Freire. 2016. A First Study on Temporal Dynamics of Topics on the Web. In *Proceedings of WWW*. 849–854.
- [36] Lloyd S Shapley. 1953. A Value for n-Person Games. In *Contributions to the Theory of Games II*, Harold W. Kuhn and Albert W. Tucker (Eds.). Princeton University Press, Princeton, 307–317.
- [37] Karen Sparck Jones, Steve Walker, and Stephen E. Robertson. 2000. A probabilistic model of information retrieval: development and comparative experiments - Part 1. *Information Processing and Management* 36, 6 (2000), 779–808.
- [38] Hal R. Varian. 2009. Online Ad Auctions. *American Economic Review* 99, 2 (2009).
- [39] Lidian Wang, Paul N. Bennett, and Kevyn Collins-Thompson. 2012. Robust ranking models via risk-sensitive optimization. In *Proceedings of SIGIR*. 761–770.
- [40] Grace Hui Yang, Marc Sloan, and Jun Wang. 2016. *Dynamic Information Retrieval Modeling*. Morgan & Claypool Publishers.

A A SIMPLE EXAMPLE OF THE NON-OPTIMALITY OF THE PRP

A simple way to demonstrate the sub-optimality of the probability ranking principle (PRP) in a game-theoretic setting is as follows. Assume two authors (players 1 and 2) who can write on two topics – i.e., each player chooses between two strategies: writing on topic *A* or writing on topic *B*. There is big user demand, $U(A)$, to topic *A*, and a slightly lower user demand $U(B) = U(A) - \epsilon$, to topic *B*. (We assume a query is about a single topic and the demand represents the distribution over queries.) Player 1 can write with optimal quality (1) on both topics, while player 2 can write with close to optimal quality $(1 - \epsilon)$ on topic *A* but with really poor quality (ϵ) on topic *B*. The authors' gain (utility) is the quality of their writing times the demand they get on the topic for which they are ranked the highest. The total users' satisfaction (social welfare) though is the overall total quality over both topics. According to the PRP, the higher quality document is ranked first for a query for each topic. In this game the only equilibrium is for player 1 to write on topic *A*, and for player 2 to write on topic *B*; in this case the total users' satisfaction is not much more than half of the total social welfare (satisfaction) the users could get (i.e., the price of anarchy is slightly less than 2). The PRP simply gives no chance to player 2, who could be writing so well on topic *A*, although this is the only topic he writes well about.