

THE ECONOMICS OF ATTENTION: MAXIMIZING USER VALUE IN INFORMATION-RICH ENVIRONMENTS

BERNARDO A. HUBERMAN* and FANG WU†

HP Labs, Palo Alto, CA 94306, USA

**bernardo.huberman@hp.com*

†fang.wu@hp.com

Received 6 December 2007

Revised 19 February 2008

A distinguishing feature of the information era is the saliency of people's attention as a scarce resource. Unlike an earlier time when information was a valuable resource, its easy availability has shifted the focus to the limited bandwidth that people can devote to ubiquitous media and news. With this in mind, we introduce an automatic configuration mechanism that generates the most relevant information to be presented to limited attention users. It also guarantees to maximize their total expected utility from the information they receive. A computationally efficient algorithm is used to assign an index value to each information item, which then determines whether or not a given item appears in the top list presented to users at a given time.

Keywords: Economics of attention; search engine; dual-speed restless bandit problem.

1. Introduction

An interesting but daunting consequence of the prevalence of the Web and digital media is that information, which used to be scarce and therefore valuable, is now so ubiquitous as to be almost devoid of monetary value. Search engines, billions of websites, targeted advertisement and easy access to digital content provide us with myriad ways of taking care of our most complex informational and entertainment needs. What is now scarce, and therefore valuable, is the user's attention, which explains the intense efforts made to obtain it through focused advertising, pop-ups, short videos embedded in news portals and, most disheartening, spam.

The interaction between attention and information, first articulated by Simon in 1971 [18], has been studied in fields from psychology [12] to economics [6, 9, 13–15]. Recently, Falkinger developed an equilibrium model of the economics of attention that elucidates the interplay between the intensity of the information available to consumers and the amount of attention that they can devote to it [4, 5]. Specifically, he showed that when the information exposure of individuals is low an information-poor economy ensues in which there is no scarcity of attention. Conversely, an information-rich economy with a consequent scarcity of attention is bound to appear

whenever the technology makes it easy to reach a large number of individuals without raising costs, or when individual wealth increases.

An information-rich regime is characterized by keen competition for the user's attention, resulting in a flood of information from which people often find it hard to sort out the most relevant and useful pieces. In addition, the law of surfing, which states that the probability of a user accessing a number of items in a single session markedly decays with the number of items [11], puts a strong constraint on the amount of information that ever gets explored in a single surfing session [1].

Rather than leaving it to the consumer to cope with this distracting overload, providers often try to present first the most salient items in their inventory while taking into account the visual real estate available on a given device. Search engines such as Google or Yahoo do not exhibit all their search results on one webpage, but rather prioritize and display them on consecutive pages whose value is assumed to be decreasingly lower to the user. Google has even implemented a function called "I'm feeling lucky" that will automatically choose the "best" search result for the user, given a set of search terms. The same applies to large recommendation sites such as CNET, where items or stores are ranked and displayed according to the number of positive rankings they receive.

These approaches suffer from two problems. The first one lies with the content provider, who needs to decide what to prioritize in order to get the user's attention. This decision can be made on the basis of some objective criterion (page rank in search, number of recommendations for software, popularity of a site, saliency of news) or some heuristic rule that the content provider develops. In either case it is not clear that such procedures maximize the user's value. For example, while an algorithm like page rank inserts the most-linked-to pages in the first page of a query result, other links in other pages often contain incipiently valuable information that is not available to the user.

The second problem stems from the finite number of items that a user can attend to in a given time interval. Because of this, a user is more likely to explore the first few items presented to him. For example, there is empirical evidence that a typical user seldom visits pages beyond the first one in a search result, so that a page ranked at the bottom by a search engine is unlikely to be viewed by many users [3,10]. This behavior tends to reinforce the leading position of those top items and further increase their popularity, which in turn penalizes new content that is not well known yet. Thus it is easy for an item to get locked in a top ranking and hard for other bottom items to surface, even though the latter can often be more valuable. It is therefore important to break this distorting reinforcement process by encouraging users to explore more items, thus increasing on average the value they obtain [17].

In this paper we solve these problems by presenting an automatic configuration mechanism that maximizes user value in information-rich environments. It does so by mapping the problem of optimizing the information one gets from exploring

Web pages or any other digital content to that of the optimal allocation of effort to a number of competing projects. We then formulate our problem as a *dual-speed restless bandit problem*, which is a special case of the celebrated *restless bandit problem*, first solved by Gittins and Whittle a number of years ago [7, 19]. By using a computationally efficient algorithm recently developed by Bertsimas and Niño-Mora [2], we calculate an *index* for each item state, which is then used to decide which item goes into the top list at a given time.

In what follows we first formulate the problem of choosing a subset of items that maximize the informational value for users, and show how it corresponds to a discounted dual-speed multiarm bandit problem. Next, we use the adaptive greedy algorithm of Bertsimas and Niño-Mora to compute the index for each item state and proceed to illustrate the use of this mechanism by resorting to a concrete example. As we show, this approach yields useful and nontrivial results, for in several instances the rank order of the indices suggests the presentation in the top page of items that at first inspection would appear to have little value to the user. We conclude by summarizing our results, describing possible applications, and their implementation within several scenarios.

2. Deciding What to Display

Consider a system that wishes to present n different items to a large number of users but can only display up to k items ($k < n$) at any given time. Since an item displayed in front of a user has a higher probability of being chosen than when it is not displayed, we will call these k items the “top list.” We will also assume that the system can update its top list at discrete times $t = 0, 1, 2, \dots$

If the system can track a certain set of properties for each item, such as its reputation, history or age, we say that the item is in a “state” defined by those properties. Let E be the set of all possible states, i.e. all possible combinations of those trackable properties. In general, the state of an item may change as time goes on. As an example, on a software download site the number of downloads, or the average rating of a particular package, may vary from week to week. We assume that the state of each item changes according to a Markov process independent of the state of other items, with transition probabilities $\{P_{ij}^1 : i, j \in E\}$ if the item is on the top list, and $\{P_{ij}^0 : i, j \in E\}$ if it is not. We also make the assumption that having an item on the top list encourages more users to try it out and consequently accelerates its transitions from one state to the other. Conversely, when an item transitions away from the top list it slows down its rate of change by an amount ϵ_i which is less than 1. This *dual-speed assumption* can be stated as

$$P_{ij}^0 = \begin{cases} \epsilon_i P_{ij}^1, & i \neq j, \\ (1 - \epsilon_i) + \epsilon_i P_{ii}^1, & i = j, \end{cases} \quad (1)$$

where $\epsilon_i \in [0, 1]$.

Consider the total expected utility r_i obtained at one time step by those users who decide to try an item on the top list which has state i . This utility may depend on many factors, such as the total expected number of users choosing the item at a given time step, or the expected quality of the item. Since we can always enlarge the definition of “state” to include these factors, the utility r_i is uniquely determined by the item state i . In other words, we can assume that $r = (r_i)_{i \in E}$ is an $|E|$ -dimensional constant vector known by the system.

Our goal then is to design a system that maximizes the total expected utility of all users:

$$\max_{u \in \mathcal{U}} \mathbb{E}_u \left[\sum_{t=0}^{\infty} \sum_{m=1}^n \beta^t r_{i_m(t)} I_m(t) \right], \quad (2)$$

where $0 < \beta \leq 1$ is the future discount factor, $i_m(t)$ is the state of item m at time t , and

$$I_m(t) = \begin{cases} 1 & \text{if item } m \text{ is displayed at time } t, \\ 0 & \text{otherwise.} \end{cases} \quad (3)$$

We seek to find the optimal strategy, u , in the space \mathcal{U} of stationary strategies (strategies that depend on current item states only). This optimal strategy can then get translated into the set of offerings that should appear in the top list.

3. The Solution

The model we just described is essentially a *dual-speed restless bandit* problem, as defined in Refs. 19 and 8. It is restless because changes of state can also occur when the items are not displayed in the top list, and dual-speed because those changes do happen at a different speed than those on the top list. Bertsimas and Niño-mora have shown that the optimal solution to the dual-speed restless bandit problem has the following simple form [8, 16]. It is possible to attach an *index* to each item state, so that the top list is the one including those items with the largest indices. This way the user value gets maximized. It is worth remarking that it is not obvious why the states’ relative importance can be measured by one independent index. In fact, for a general restless bandit problem without the dual-speed assumption, such a set of indices may not exist [19].

Nevertheless, Bertsimas and Niño-mora showed that a relaxed version of the dual-speed problem is always *indexable* (i.e. such indices always exist) and also proposed an efficient adaptive greedy algorithm to compute these indices. By “relaxed” we mean that instead of displaying *exactly* k items at each time, k items *on average* are displayed [19]. For this relaxed problem, it can be shown that there exists a set of indices $\{G_i\}_{i \in E}$ and a Lagrange multiplier γ such that the optimal strategy is to always display those items whose G index is greater than γ [16, 19]. Note that in situations where the top list can have variations in the number of items, the relaxed situation is the one that applies. In the limit of no variations, while

the solution is known to be suboptimal [8] suggests that it is a good approximation to the optimal one.

In order to use their algorithm, we first calculate a set of constants A_i^S , which we now define. Assume that E is finite. For any subset $S \subseteq E$, we define the S -active policy u_S to be the strategy that recommends all items whose state is in S . Now consider an item that starts from an initial state $X(0) = i$. Under the action implied by strategy u_S , its total occupancy time in S is given by

$$V_i^S = \mathbb{E}_{u_S} \left[\sum_{t=0}^{\infty} \beta^t I_S(t) | X(0) = i \right], \tag{4}$$

where

$$I_S(t) = \begin{cases} 1 & \text{if } X(t) \in S, \\ 0 & \text{otherwise.} \end{cases} \tag{5}$$

We have

$$V_i^S = \begin{cases} 1 + \beta \sum_{j \in E} P_{ij}^1 V_j^S, & i \in S, \\ \beta \sum_{j \in E} P_{ij}^0 V_j^S, & i \in S^c. \end{cases} \tag{6}$$

The variables $\{V_i^S\}_{i \in E}$ can be solved from the set of linear equations above.

A matrix of constants $\{A_i^S\}_{i \in E, S \subseteq E}$ is defined by means of V_i^S , as follows:

$$A_i^S = 1 + \beta \sum_{j \in E} P_{ij}^1 V_j^{S^c} - \beta \sum_{j \in E} P_{ij}^0 V_j^{S^c}. \tag{7}$$

The constants $\{A_i^S\}$ are then used in the Bertsimas–Niño–Mora algorithm (Algorithm 1), as indicated in the table.

Algorithm 1: Bertsimas–Niño–Mora adaptive greedy algorithm

Step 1. Set $S_{|E|} = E$ and

$$y^{S_{|E|}} = \max \left\{ \frac{r_i}{A_i^E} : i \in E \right\}. \tag{8}$$

Select $\pi_{|E|}$ as any maximizer and set $G_{\pi_{|E|}} = y^{S_{|E|}}$.

Step 2. For $k = 2, 3, \dots, |E|$, set $S_{|E|-k+1} = S_{|E|-k+2} \setminus \{\pi_{|E|-k+2}\}$ and

$$y^{S_{|E|-k+1}} = \max \left\{ \frac{r_i - \sum_{j=1}^{k-1} A_i^{S_{|E|-j+1}} y^{S_{|E|-j+1}}}{A_i^{S_{|E|-k+1}}} : i \in S_{|E|-k+1} \right\}. \tag{9}$$

Select $\pi_{|E|-k+1}$ as any maximizer and set $G_{\pi_{|E|-k+1}} = G_{\pi_{|E|-k+2}} + y^{S_{|E|-k+1}}$.

Once we compute the G index for each state using Algorithm 1, the strategy is to always display the k items whose states have the largest G indices. For our dual-speed restless bandit problem, it follows from Ref. 8 that $A_i^S > 0$ for all $i \in E$ and $S \subseteq E$, so that the relaxed version of our problem is indexable and Algorithm 1 provides a good heuristic for the unrelaxed problem.

4. An Example

4.1. Description

We now show how our mechanism works by presenting a simple example. Consider a website that can display k out of n items at any given time. Each item can have a star rating among 1–5, with 5 stars denoting the highest rating and 1 the lowest, aggregated from users who volunteer to rate the item. Each item also has an access level from 1 to 5 indicating its click rate in one time step, with 5 denoting the largest number of clicks received. The access levels can be set rather arbitrarily, e.g. level 1 for 1000 clicks, level 2 for 2000 clicks, level 3 for 5000 clicks, etc. Hence each state can be represented as a 2-vector $(s, a) \in \{1, 2, 3, 4, 5\}^2$, where s is the star rating and a is the access level.

In addition to those 25 states there is one more state, 0, which we call the “unknown” state. Each item initially starts in this state, as it has never been either accessed or rated. We assume that occasionally an item will “die,” and if that happens it is immediately replaced by a new item. This is equivalent to assuming that there is a small transition probability from each of the 25 states to the unknown state, entering which implies starting over. State 0 thus serves as both the sink and the source.

The transition probabilities are assumed to be as follows:

$$\begin{cases} P^1(s, a; s + 1, a) = 0.1, & 1 \leq s \leq 4, 1 \leq a \leq 5, \\ P^1(s, a; s - 1, a) = 0.1, & 2 \leq s \leq 5, 1 \leq a \leq 5, \\ P^1(s, a; s, a + 1) = 0.2, & 1 \leq s \leq 5, 1 \leq a \leq 4, \\ P^1(s, a; s, a - 1) = 0.1, & 1 \leq s \leq 5, 2 \leq a \leq 5, \\ P^1(s, a; 0) = 0.01, & 1 \leq s \leq 5, 1 \leq a \leq 5, \\ P^1(0; s, 1) = 0.1, & 1 \leq s \leq 5, \end{cases} \quad (10)$$

and

$$\epsilon_i = 0.1 \quad \text{for all } i \in E, \quad (11)$$

which expresses the fact that displaying an item on the top list accelerates its transition speed by ten times. Note that we have assumed that an item’s access level tends to increase more than to decrease. The states and the transition probabilities are illustrated in Fig. 1.

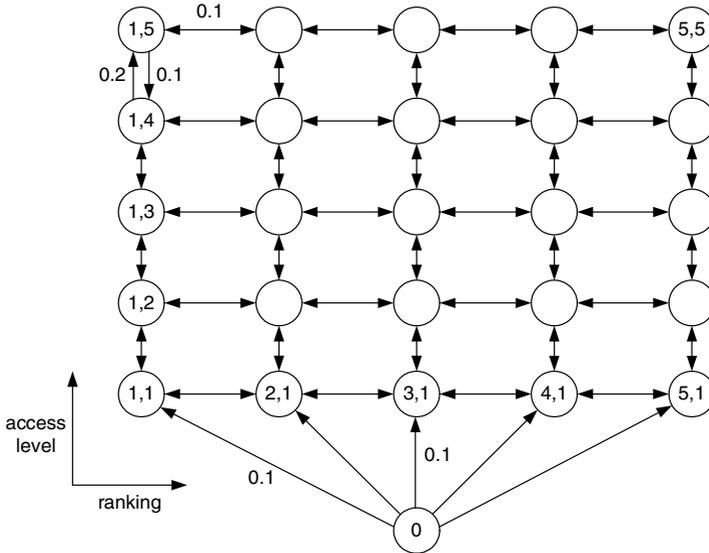


Fig. 1. An illustration of the 26 states and the transition probabilities P^1 between them. The horizontal axis denotes the ratings, and the vertical axis the five possible access levels. What is not shown in the figure is that every state in the 5×5 grid transits to state 0 with a small probability, 0.01.

We naturally set the reward of each state to be

$$\begin{cases} r(s, a) = s \cdot a, \\ r(0) = 0. \end{cases} \tag{12}$$

That is, items on the website are rewarded in proportion to the number of clicks they receive and their ratings.

4.2. Solution

The G index rankings of the 26 states are calculated using the Bertsimas–Niño-Mora heuristic described in Sec. 3. The result is shown in Fig. 2. As can be seen, state (5,5) has the largest G index, state (5,4) the second-largest, and so on. The absolute values of the indices are not as important as their relative orders. Those items that acquire the largest G index value are the ones to be displayed to the user.^a

The result of this example is by no means trivial. For example, it is not obvious that the unknown state which gives no reward should have higher display priority

^aAt the very beginning, when there is no previous information for any of the items (all “unknown”), the initial list could be randomly composed. Since what we are maximizing is the long-term utility of the site owner, the initial display of information is not that critical.

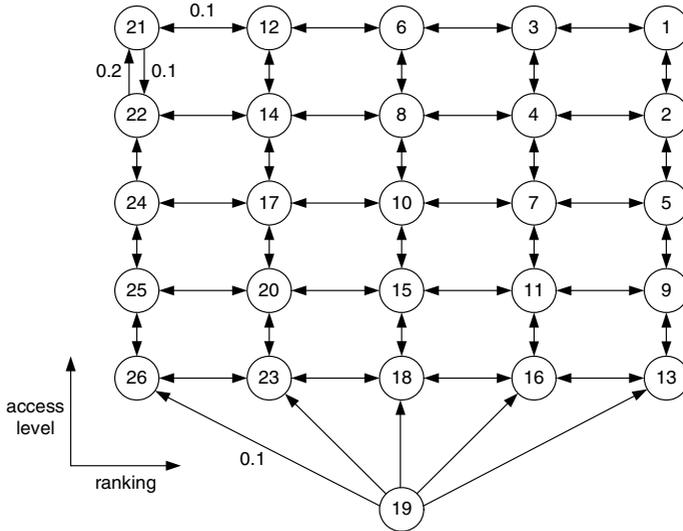


Fig. 2. The 26 states ranked by their G indices, from highest to lowest. For example, the state (5,5) in the top right corner has the highest G index and the state (1,1) in the lower left corner has the lowest G index.

than state (2,2), but lower priority than (3,1). This effect is due to the fact that the algorithm gives high index values to potentially valuable states.

It is rather straightforward to extend this example to larger systems and to compute the transition probabilities from actual data from a portal.

5. Conclusion

The information-rich environment which digital technology and the Web have launched upon us is both a panacea and a source of bewilderment. Sponsored advertisements, thousands of items to be explored when posing a query to a search engine or a portal, and recommendation mechanisms that are hard to evaluate are now part and parcel of what users of the Web face on a daily basis.

While the reasons for this information onslaught are easily explainable in terms of an economy of attention where the latter is the scarce resource, a solution to the problems facing the user is wanting. This paper provides a mechanism for increasing, and often optimizing, the value people can get out of information-rich environments. The mechanism consists in an algorithm that automatically generates the top candidate items to be displayed on any device which are guaranteed to provide a maximum discounted value to most users.

This algorithm relies on the computation of a set of indices, each allocated to each item state on a list, which can be computed by accessing the rates at which items are visited and the rankings they receive from users. These rates determine the transition probabilities, which are then used as inputs into the actual computation of the index for each state.

We showed that the actual computation of these indices can be performed by mapping the problem of optimizing the information one gets from any other digital content to that of the optimal allocation of effort to a number of competing projects. This allowed us to formulate our problem as a *dual-speed restless bandit problem*, which is a special case of the restless multiarm bandit problem. By using a computationally efficient algorithm recently developed by Bertsimas and Niño-Mora, we were then able to calculate an index for each item state.

This mechanism can be used to solve a multiplicity of problems, ranging from the decision on which search results to display on the top page of a search engine, to the menu of items that a portal decides to present to users and the order in which a journal presents its contents to users. Other applications that are worth noting are the decision on what to display in devices with a small visual real estate, the relevant information that should be presented to analysts confronted with mountains of data, how to sort through blogs and other forms of user-generated media, and the decision on how to best optimize movie and video directories. Another area is instrumentation, whose sole purpose is to inform the user about the state of the world in which it is embedded. Last but not least, advertising is another potential beneficiary of this technology, for it could simply use click patterns from visitors to given portals to decide on which ads to present at given times.

Given the flood of information that most people encounter in their daily interaction with the Web, we believe that this mechanism will alleviate information overload while maximizing the value that users obtain from such a medium.

References

- [1] Adar, E. and Huberman, B. A., The economics of surfing, *Quart. J. Electron. Commer.* **1** (2000) 203–214.
- [2] Bertsimas, D. and Nino-Mora, J., Conservation laws, extended polymatroids and multiarmed bandit problems: A unified approach to indexable systems, *Math. Operat. Res.* **21** (1996) 257–306.
- [3] Cho, J. and Roy, S., Impact of search engines on page popularity, in *Proc. World-Wide Web Conference* (2004), Vol. 33, pp. 76–98.
- [4] Falkinger, J., Limited attention as the scarce resource in an information-rich economy, IZA Discussion Papers 1538 (2005).
- [5] Falkinger, J., Attention economies, *J. Econ. Theory* (2006), forthcoming.
- [6] Franck, G., Science communication, a vanity fair, *Science* **286** (1999) 53–55.
- [7] Gittins, J. C., Bandit processes and dynamic allocation indices, *J. R. Stat. Soc. Ser. B* **41** (1979) 148–164.
- [8] Glazebrook, K. D., Nino-Mora, J. and Ansell, P. S., Index policies for a class of discounted restless bandits, *Adv. Appl. Prob.* **34** (2002) 754–774.
- [9] Goldhaber, M. H., Attention economics and the net, *First Monday* 2 1998.
- [10] Halvey, M., Keane, M. T. and Smyth, B., Mobile web surfing is the same as web surfing, *Commun. ACM* **49**(3) (2006).
- [11] Huberman, B. A., Pirolli, P., Pitkow, J. and Lukose, R., Strong regularities in World Wide Web surfing, *Science* **33** (1998) 76–98.
- [12] Kahneman, D., *Attention and Effort* (Prentice-Hall, 1973).

- [13] Klamer, A. and Dalen, H. P. V., Attention and the art of scientific publishing, *J. Econ. Methodol.* **9** (2002) 289–315.
- [14] Lanham, R. A., *The Economics of Attention: Style and Substance in the Age of Information* (University of Chicago Press, 2006).
- [15] Levy, D. M., The market for fame and fortune, *Hist. Polit. Econ.* **33** (1988) 615–625.
- [16] Nino-Mora, J., Restless bandits, partial conservation laws and indexability, *Adv. Appl. Prob.* **33** (2001) 76–98.
- [17] Pandey, S., Roy, S., Olston, C., Cho, J. and Chakrabarti, S., Shuffling a stacked deck: The case for partially randomized ranking of search engine results, in *Proc. VLDB* (2005).
- [18] Simon, H. A., Designing organizations for an information rich world. In: M. Greenberger (ed.), *Computers, Communications and the Public Interest* (Johns Hopkins Press, 1971), pp. 38–52.
- [19] Whittle, P., Restless bandits: Activity allocation in a changing world, *J. Appl. Prob.* **25A** (1998) 287–298.