

Recommender Systems: A Market-Based Design ^{*}

Yan Zheng Wei [†], Luc Moreau and Nicholas R. Jennings
Department of Electronics and Computer Science
University of Southampton
Southampton SO17 1BJ, UK
{ yzw01r, L.Moreau, nrj }@ecs.soton.ac.uk

ABSTRACT

Recommender systems have been widely advocated as a way of coping with the problem of information overload for knowledge workers. Given this, multiple recommendation methods have been developed. However, it has been shown that no one technique is best for all users in all situations. Thus we believe that effective recommender systems should incorporate a wide variety of such techniques and that some form of overarching framework should be put in place to coordinate the various recommendations so that only the best of them (from whatever source) are presented to the user. To this end, we show that a marketplace, in which the various recommendation methods compete to offer their recommendations to the user, can be used in this role. Specifically, this paper presents the *principled* design of such a marketplace; detailing the auction protocol and reward mechanism and analyzing the rational bidding strategies of the individual recommendation agents.

Categories and Subject Descriptors

I.2.11 [Distributed Artificial Intelligence]: Multiagent systems; H.3.3 [Information Search and Retrieval]: Information filtering

General Terms

Algorithms, Design, Economics

Keywords

Recommender System, Mechanism Design, Auctions

1. INTRODUCTION

The receipt of undesirable or non-relevant information that results in an economic loss for the recipient [11], generally referred to as *information overload*, is a major problem for

^{*}This research is funded in part by QinetiQ and the EPSRC Magnitude project (reference GR/N35816).

[†]The first author of this paper is a student.

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. To copy otherwise, to republish, to post on servers or to redistribute to lists, requires prior specific permission and/or a fee.

AAMAS'03, July 14–18, 2003, Melbourne, Australia.
Copyright 2003 ACM 1-58113-683-8/03/0007 ...\$5.00.

many individuals. For this reason, significant research endeavor is being invested in building support tools that ensure the right information is delivered to the right people at the right time. While search engines and information filtering tools can assist in this endeavor, they are typically not personalized to individual users or their prevailing context and they tend not to deliver an appropriate amount of information [18]. To overcome these limitations, *recommender systems* have been advocated. Such systems help make choices among recommendations from all kinds of sources without having sufficient personal experience of all these alternatives [13]. In a typical recommender system, recommendations are provided as inputs and the system then aggregates and directs them to appropriate recipients. Thus a recommender system's main value lies in information aggregation and its ability to match the recommenders with those seeking recommendations.

Recommender systems have been applied in many application domains (including music albums [15], video [9] and Web navigation [7]) and many different techniques have been used to make the recommendations. For example, some are based on the correlation between the item contents (such as term frequency inverse document frequency [12] and weighting [7]), while others are based on the correlation between users' interests (such as votes [8] and trails [5]). However, there is no universally best method for all users in all situations [3] and we believe this situation is likely to continue as ever more methods are developed. Moreover, the ranking of relevance produced by the different methods can vary dramatically from one another. Given this situation, we believe the best way forward in this area is to allow the multiple recommendation methods to co-exist and to provide an overarching system that coordinates their outputs such that only the best recommendations (from whatever source or method) are presented to the user. To this end, in [12] we developed a system that provides recommendations about "where shall I read next?" as the user browses web pages (see Figure 1). We showed that a *market-based approach* is an efficient means of achieving such coordination because the problem of selecting appropriate recommendations to place in the limited sidebar space can be viewed as one of scarce resource allocation and markets are an effective solution for this class of problems [4]. In our system, the various recommendation techniques (represented as economic agents) compete with one another to advertise their recommendations to the user. Those agents that make recommendations that are selected by the user as being good are rewarded and those agents that make poor recommen-



Figure 1: Browser with Recommendations

dations make losses (since they have to pay to advertise their recommendations). Thus, over the longer-term, those agents that make good recommendations become richer and so are able to get their recommendations advertised more frequently than the methods whose recommendations are infrequently chosen by the user.

While the marketplace we developed for [12] worked efficiently much of the time, in certain cases it exhibited undesirable behaviour (for example, agents that were rewarded initially had sufficient funds to continuously outbid the other recommenders even when they had poorer recommendations and the agents were not economically efficient because of the way the marketplace limited the rewards). Thus, we need to develop a more principled approach to mechanism design to ensure the ensuing marketplace does not have such undesirable pathologies.

Against this background, this paper advances the state of the art in the following ways. Firstly, we outline a method for coordinating the behaviour of multiple recommendation methods with diverse measures of similarity (no other recommendation system has attempted to incorporate multiple approaches). Secondly, we design a marketplace that is Pareto-optimal, maximizes social welfare, stable and fair to all the recommending agents for an important real world problem. Thirdly, our market design relates an agent's rational reasoning to the user's perceived quality of recommendations.

The remainder of this paper is structured in the following manner. Section 2 details the market mechanism design with respect to the auction protocol, the reward mechanism, the bidding strategy and an analysis of the market equilibrium. Section 3 evaluates this design in terms of its economic efficiency criteria. Section 4 outlines related work in terms of reducing information overload and market-based systems. Section 5 concludes and points to future work.

2. DESIGNING THE MARKETPLACE

Our marketplace operates according to the following metaphor. A user agent acting on behalf of the user is selling sidebar space where recommendations may be displayed. The number of such slots is fixed and limited. Information providers (the component recommending agents)¹ want to

¹We assume that each recommending agent is self-interested and is unaware of other agents' valuations of its recommen-

get their recommendations advertised in the user's browser and so compete in the marketplace to maximize their individual gain by purchasing this advertising space when they have what they believe are good recommendations. Their bids indicate how much they are willing to pay for such slots. The recommender system acts as the auctioneer and selects the most valuable items (highest bids) which it then displays as its recommendations (those agents that provided these *shortlisted* items are then charged according to their bids). The user then chooses some of these recommendations (or not) according to their interests. The agents that provided the user-selected recommendations receive some reward (since these recommendations are deemed useful), while those not chosen receive no reward.

Ideally, we would like to use one of the standard auction protocols for our marketplace. However, this is not possible because of the peculiarities of our scenario. Specifically, standard auctions could probably deal with the shortlisting phase, but they do not consider the subsequent reward phase. This means a bespoke mechanism is needed.

To evaluate our mechanism we will use some standard economic metrics [14]:

Pareto Efficiency: This is important from the point of view of the individual agents since it enables us to compare alternative mechanisms. A solution x is Pareto efficient if there is no other solution x' such that at least one agent is better off in x' than in x and no agent is worse off in x' than in x .

Social Welfare Maximization: In our context, social welfare is a combination of all the agents' utilities. This measure provides a way of ranking different distributions of utility among agents and of indicating what is best for the group as a whole.

Individual Rationality: Participation in an auction is individually rational to an agent if its payoff in the auction is no less than what it would get by not participating. A mechanism is individually rational if participation is individually rational for all agents. Individually rational protocols are essential because without them, there is no incentive for agents to participate.

Stability: A protocol is stable if it provides all agents with an incentive to behave in a particular way over time. Protocols should be designed to be stable because if a self-interested agent is better off behaving in some other manner than desired, it will do so. Stability is important because without it the system behaviour is unpredictable.

Fairness: In our context, a market is fair if it gives all recommendations equal opportunity of being shortlisted (irrespective of the agent or method that makes the recommendation). This is important because we want the system to shortlist the best recommendations in an unbiased manner.

With these metrics in place, the rest of the section details the auction protocol we designed, the reward mechanism we established and the bidding strategies of the individual agents. Section 3 then evaluates these components against the above criteria.

This private-value hypothesis is justified because all recommending agents operate using different kinds of information sources and they are unaware of the existence of each other.

2.1 The Auction Protocol

To ensure recommendations are provided in a timely and computationally efficient manner, we choose a *generalized first-price sealed-bid auction* in which all agents whose recommendations are shortlisted pay an amount equal to their valuation of the advertisement (meaning we have price differentiation). We choose a sealed bid auction (in which agents will typically make a single bid) to minimize the time for running the auction and the amount of communication generated. We choose a first price auction with price differentiation because the relative ordering of the recommendations effects the likelihood of them being selected by the user.

In more detail, the market operates in the following manner. Each time the user browses a new page the auction is activated. In each such activation, the auctioneer agent calls for a number of bids, say M ($M > 0$) equal to the number of recommendations it is seeking. After a fixed time, the auctioneer agent ranks all the bids it received by their bidding price, and directs the M bids with the highest prices to the user's browser. Those bidding agents whose recommendations are shortlisted pay the auctioneer agent according to how much they bid. Those bidding agents whose recommendations are not shortlisted do not pay anything. The user may then take up a number of these shortlisted recommendations in which case the agent that supplied them is rewarded.

More formally, the variables representing the different entities and values in each auction round are:

S : the number of recommending agents ($S \gg 1$ ²);
 $A_{b1}, A_{b2}, \dots, A_{bS}$: S bidding agents;
 A_B : complete set of bidding agents, i.e., $A_{b1}, A_{b2}, \dots, A_{bS}$;
 A_a : auctioneer agent;
 A_u : user agent;
 T_b : duration of the auction;
 M : number of recommendations that A_u requests from A_a ;
 $b_{ij} = \langle A_{bi}, rec_j, price_j \rangle$: bid provided by A_{bi} , containing the j^{th} recommendation with bidding price $price_j$ ($i \in [1..S], j \in [1..M]$);
 B^{ALL} : all bids submitted to A_a ;
 B^M : shortlisted bids recommended to A_u ;
 B^R : bids selected by the user (will be rewarded by A_a);
 SU : recommendations displayed in user's sidebar (B^M ignoring the prices);
 SU^R : recommendations selected by user (B^R ignoring the prices);
 N : number of user-selected recommendations;
 b_l, b_h : two bids for temporary use ($l, h \in [1..M]$);
 R_h : reward to h^{th} user-selected recommendation.

The protocol for each auction round is formally defined in Figure 2. It should be noted that: (i) function *GenerateBid* ($A_{bi}, rec_j, price_j$) relates to the bidding strategy and will be discussed in section 2.3; (ii) function *UserSelectsRecs* (SU) concerns the user's behaviour of making choices among the shortlisted recommendations; and (iii) function *ComputeReward* (b_h) concerns the reward mechanism and will be discussed in section 2.2.

2.2 The Reward Mechanism

With the auction protocol in place, we now turn to the reward mechanism. According to our protocol, the user may select multiple recommendations from the shortlist.

²We assume the number of recommenders is sufficiently large with respect to the number of sidebar slots such that there is sufficient competition to make the marketplace operate efficiently.

```

i ∈ [1..S]    j, l, h ∈ [1..M]
BALL = φ;
BM = φ;
BR = φ;
CallForBids(AB, M, Tb);
repeat during the duration of auction Tb    // receiving bids
{
    bij = GenerateBid(Abi, recj, pricej);
    BALL = BALL ∪ {bij};
}
for l = 1 to M do    // shortlisting
{
    bl = FindBidWithLthTopPrice(BALL, l);
    BM = BM ∪ {bl};
}
SU = { ⟨Abi, recj⟩ | ⟨Abi, recj, pricej⟩ ∈ BM };
SUR = UserSelectsRecs(SU);    // SUR ⊆ SU
BR = { ⟨Abi, recj, pricej⟩ |
    ⟨Abi, recj⟩ ∈ SUR and ⟨Abi, recj, pricej⟩ ∈ BM };
N = |BR|;
for h = 1 to N do    // rewarding
{
    bh = FindHthBid(BR, h);
    Rh = ComputeReward(bh);
}

```

Figure 2: The Auction Protocol

For each such user-selected recommendation, the suggesting agent is given a reward³. In defining the *ComputeReward* function, our aim is to ensure that it is both Pareto efficient and social welfare maximizing. To this end, this subsection addresses the following issues: (i) How is one reward mechanism judged to be better than another? (ii) Does there exist a reward mechanism that is the best amongst all possible mechanisms? First, however, a complete set of reward mechanisms is introduced.

2.2.1 The Complete Set of Reward Mechanisms

Let us assume we have N (defined in section 2.1) user-selected recommendations to be rewarded and the auctioneer has an amount of payoff, R^T , to be distributed to the relevant agents. The problem is then how to best split R^T into parts and distribute them to each of the rewarded recommending agents such that we cannot find any other more optimal allocation solutions.

To this end, we define the complete set of reward mechanisms as follows: Suppose the h^{th} ($h \in [1..N]$) user-selected recommendation receives an amount of payoff R_h . Then, all possible reward mechanisms are such that the sum of each payoff is less than or equal to R^T . That is, $\sum_{h=1}^N R_h \leq R^T$. Therefore, we have a complete set of reward mechanisms, $\hat{\mathfrak{R}}$, such that:

$$\hat{\mathfrak{R}} = \{ (R_1, R_2, \dots, R_N) \mid \sum_{h=1}^N R_h \leq R^T \}$$

Now each element of $\hat{\mathfrak{R}}$ is a possible allocation of R^T and $\hat{\mathfrak{R}}$ can be split into two complementary subsets: $\hat{\mathfrak{R}}_1$ that does not completely allocate all of R^T (called a *With Surplus Mechanism* (WSM)) and $\hat{\mathfrak{R}}_2$ that does allocate all of R^T (called a *No Surplus Mechanism* (NSM)):

$$\hat{\mathfrak{R}}_1 = \{ (R_1, R_2, \dots, R_N) \mid \sum_{h=1}^N R_h < R^T \} \text{ (WSM)}$$

$$\hat{\mathfrak{R}}_2 = \{ (R_1, R_2, \dots, R_N) \mid \sum_{h=1}^N R_h = R^T \} \text{ (NSM)}$$

From these two subsets, we want to identify those that are both Pareto efficient and social welfare maximizing.

³A given agent may have multiple recommendations selected in a given auction in which case it receives multiple rewards.

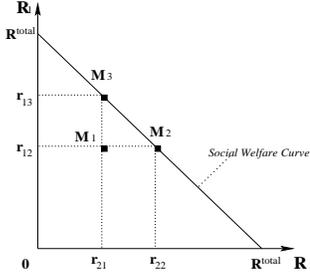


Figure 3: Pareto Optimization

2.2.2 Pareto Optimal Reward Mechanisms

If there is only one recommendation to be rewarded, it is trivially true that awarding all of R^T to this recommendation is the Pareto optimal solution. However, when there is more than one recommendation to be rewarded, the allocation is more complicated. To simplify the presentation, we discuss the case where two recommendations are rewarded (i.e., $N = 2$). This is chosen since it can easily be depicted and it gives us a direct impression of allocation. The general case with multiple recommendations rewarded ($N \geq 2$) can be analyzed in the same way. To this end, Figure 3 depicts the case where there are two recommendations to be rewarded (R_1 and R_2). The axes represent the payoff allocated to each recommendation. We define the budget payoff curve as the line joining $(0, R^T)$ and $(R^T, 0)$ and it represents the payoffs whose sum is R^T . The triangle formed by the budget payoff curve and the axes, including the edges, contains all possible allocations of R^T . The outcome of each possible reward mechanism corresponds to a point within this area⁴. Those points on the budget payoff curve represent the elements of NSM since each of these points allocates the total amount of R^T . For example, for point M_2 , $R_1 = r_{12}, R_2 = r_{22}$ and $R_1 + R_2 = r_{12} + r_{22} = R^T$. Generally speaking, therefore, any mechanism in the set NSM maximizes the total payoff and it is impossible to distinguish between any of these points. A mechanism that produces a reward in the triangle, but not on the budget payoff curve, is by definition in the WSM set. For example, for point M_1 , $R_1 = r_{12}, R_2 = r_{21}$ and $R_1 + R_2 = r_{12} + r_{21} < R^T$.

In terms of Pareto efficiency, for any point representing a WSM outcome, at least one Pareto optimal point can be found representing a related NSM. For example, in Figure 3, point M_1 (in WSM) can straightforwardly be transformed into M_2 (in NSM) by giving R_2 the extra amount of reward ($r_{22} - r_{21}$). However, those points on the budget curve cannot be improved upon since giving extra reward to either recommendation necessarily results in a loss to the other. Therefore, all NSM outcomes are Pareto efficient.

2.2.3 Social Welfare Maximizing Reward Mechanisms

Pareto efficiency has nothing to say about the distribution of welfare across agents. Thus, given two mechanisms that produce outcomes that are both Pareto efficient, it is not possible to say which is better. Thus, we need a further means of differentiation. To this end, we seek to define a

⁴One point in this area may represent multiple reward mechanisms since different mechanisms may result in the same outcome. In this case, our concern is how much the reward to one recommendation is related to the reward of another. Hence, we are concerned only with the outcome and ignore what specific reward mechanism the outcome comes from.

social welfare function that is able to assign a ranking to all Pareto efficient mechanisms. This ranking specifies the “social preference” [17]p590 of a distribution of overall welfare to different rewarded recommendations and should ensure that recommendations are rewarded according to how good they are. However, in our system, there are two different views on the quality of a recommendation. Firstly, each recommending agent has an *internal quality* measure of its recommendation that is generated from the specific method it uses. This value is used to compute the agents’ bid price — the higher the quality, the higher its bid price. Secondly, a user of the recommender system also has a view of the quality of the recommendations (here termed *user perceived quality*) that indicates which of them will be taken up and which will be ignored. This user perceived quality can be defined as $Q_h \in [1..100]$ ($h \in [1..N]$). Given this, our aim is to provide a marketplace that shortlists the most valuable recommendations in decreasing order of user perceived quality. To do this, we can segment our set of potential reward mechanisms that are Pareto efficient (i.e. \mathfrak{R}_2) into two complementary subsets: those that allocate reward in a manner proportional to the user perceived quality and those that do not:

Proportional Reward Mechanism (PRM)

$$\mathfrak{R}_P = \{ (R_1, R_2, \dots, R_N) \mid R_h = \frac{Q_h}{\sum_{i=1}^N Q_i} \times R^T, \text{ where } h \in [1..N] \}$$

*Non-Proportional Reward Mechanism (NPM)*⁵

$$\hat{\mathfrak{R}}_N = \hat{\mathfrak{R}}_2 - \hat{\mathfrak{R}}_P$$

Given these two sets, we can now define our social welfare function in terms of utility. As noted above, we want the system to prefer a reward mechanism that distributes the welfare to the user-selected recommendations according to how well they satisfy the user. Therefore, a Cobb-Douglas utility function [17] is introduced. This function shows preferences of the inputs in a manner proportional to the value of their powers:

$$U(R_1, R_2, \dots, R_N) = R_1^{Q_1} \cdot R_2^{Q_2} \cdot \dots \cdot R_N^{Q_N} \quad (1)$$

In this function, the powers, Q_1, Q_2, \dots, Q_N , describe how important each rewarded agent’s utility is to the overall social welfare. Specifically, a reward mechanism, $M_i = (R_{1,i}, R_{2,i}, \dots, R_{N,i})$ is better than (or more socially-preferred to) $M_j = (R_{1,j}, R_{2,j}, \dots, R_{N,j})$, if $U(M_i) > U(M_j)$ and $i \neq j$.

Our objective now is to find if there exists a best mechanism within $\hat{\mathfrak{R}}_2$. Thus, we need to determine if there is a mechanism that has the maximum utility value, given a total amount of reward R^T . That is:

Proposition:

Does there exist an $M' \in \hat{\mathfrak{R}}_2$, such that $\forall M \in \hat{\mathfrak{R}}_2$, if $M \neq M'$, $U(M') > U(M)$?

Conditions:

$$N \text{ is a natural number} \quad (2)$$

$$Q_i > 0 \text{ and is constant } (i \in [1..N]) \quad (3)$$

$$R^T > 0 \text{ and is constant} \quad (4)$$

$$R_1 + R_2 + \dots + R_N = R^T \quad (5)$$

$$R_i > 0 \quad (i \in [1..N]) \quad (6)$$

⁵Note that $\hat{\mathfrak{R}}_P$ contains only one element (given an R^T and a set of Q_i , whose values are fixed, there is only one solution for $\hat{\mathfrak{R}}_P$). While $\hat{\mathfrak{R}}_N$ contains multiple elements.

⁶We do not consider the case of $R_i = 0$, $i \in [1..N]$, since this case must result in $U = 0$ and any mechanism with a positive utility is better than this solution.

Proof:

Because of the limited space, we just outline the key steps. In case of $N = 1$, $R_1 = R^T$ ensures the maximal value of U and this is the solution that we want. We now turn to the case of $N > 1$. Based on the given conditions, a monotonic transformation, $V = \ln U$, simplifies the problem.

$$V(R_1, R_2, \dots, R_N) = \sum_{i=1}^N Q_i \ln R_i \quad (7)$$

Hence, finding the maximum value of U is equivalent to finding that of V . Function (7) has one constraint (condition (5)) on the N input variables. Thus, only $N - 1$ variables remain independent. Let us consider that R_1 is dependent of the other $N - 1$ variables

$$R_1 = R^T - (R_2 + R_3 + \dots + R_N) \quad (8)$$

Substituting equation (8) for R_1 in function (7), we get:

$$V = Q_1 \ln[R^T - (R_2 + \dots + R_N)] + Q_2 \ln R_2 + \dots + Q_N \ln R_N \quad (9)$$

Therefore, in (9), $R_2 \dots R_N$ are independent of each other. The necessary condition for V reaching extrema is:

$$\begin{cases} \frac{\partial V}{\partial R_2} = \frac{-Q_1}{R^T - (R_2 + R_3 + \dots + R_N)} + \frac{Q_2}{R_2} = 0 \\ \frac{\partial V}{\partial R_3} = \frac{-Q_1}{R^T - (R_2 + R_3 + \dots + R_N)} + \frac{Q_3}{R_3} = 0 \\ \vdots \\ \frac{\partial V}{\partial R_N} = \frac{-Q_1}{R^T - (R_2 + R_3 + \dots + R_N)} + \frac{Q_N}{R_N} = 0 \end{cases} \quad (10)$$

Now (10) has $N - 1$ equations and $N - 1$ variables and is nonsimplified. Its unique solution is $R_j = R^T \frac{Q_j}{\sum_{i=1}^N Q_i}$, ($j \in [2 \cdot N]$). Substituting this for R_2 to R_N in equation (8), we get:

$$R_h = R^T \frac{Q_h}{\sum_{i=1}^N Q_i}, \text{ where } h \in [1 \cdot N]. \quad (11)$$

We record this extremum, (11), as M_{PRM} and note that it represents the PRM by its definition.

We now need to verify whether point M_{PRM} is a maximum or a minimum. From (11), we know that $\frac{Q_1}{R_1} = \frac{Q_2}{R_2} = \dots = \frac{Q_N}{R_N}$. We assume $K = \frac{Q_1}{R_1}$. From condition (3) and (6), we know $K > 0$. The second derivative of V is:

$$d^2V = -\left(\frac{Q_1}{R_1^2} dR_1^2 + \frac{Q_2}{R_2^2} dR_2^2 + \dots + \frac{Q_N}{R_N^2} dR_N^2\right). \quad (12)$$

At point M_{PRM} , there is a constraint on dR_1, dR_2, \dots, dR_N . This is, by differentiating condition (5) on both sides, $dR_1 + dR_2 + \dots + dR_N = 0$. So, $dR_1 = -(dR_2 + \dots + dR_N)$. The second derivative of V at point M_{PRM} is,

$$d^2V_{M_{PRM}} = -K^2 \left[\frac{(dR_2 + \dots + dR_N)^2}{Q_1} + \frac{dR_2^2}{Q_2} + \dots + \frac{dR_N^2}{Q_N} \right]. \quad (13)$$

Since $Q_i > 0$ ($i \in [1..N]$), $d^2V_{PRM} < 0$. Therefore, V and U get maximum value at solution (11). Hence, M_{PRM} is the unique maximum point and \mathfrak{R}_P , represented by M_{PRM} , is the best mechanism in $\hat{\mathfrak{R}}_2$. ■

We now illustrate this outcome with an example with two recommendations being rewarded (Figure 4). Here, the axes represent the payoffs allocated to the two recommendations. M_{PRM} and M_{NPM} represent the PRM and an element of the NPM set. The utility curves defined by $U(R_1, R_2) = R_1^{Q_1} \cdot R_2^{Q_2}$ (as per function (1)) are depicted in Figure 4 and they give us a direct impression of the comparison of the different mechanisms. In Figure 4, the mechanisms represented by points on the same utility indifference curve are as good as each other since they produce the same utility. However, mechanisms represented by points

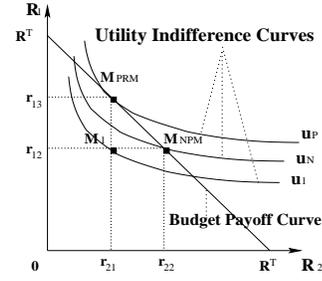


Figure 4: Social Welfare Maximization

on the outer utility curves are better (or more preferred) than those on the inner curves. This is because the outer curves bear higher utility than the inner ones. Thus, in Figure 4, $u_P > u_N > u_1$. So, the mechanism represented by M_{PRM} is better than the one represented by M_{NPM} , which, in turn, is better than M_1 . This discussion tells us that, by providing utility function (1) for the reward mechanism, the unique element of \mathfrak{R}_P represents the best possible mechanism. Therefore, this is the one we should use.

2.2.4 Designing the Reward Mechanism

Having identified $\hat{\mathfrak{R}}_P$ as the best reward mechanism for our protocol, we now need to define the total payoff R^T . In fact, the absolute value of R^T is not important. What the analysis in section 2.2.3 essentially tells us is how a reward to one recommendation should be related to that of another. In addition, it is difficult to determine the actual value of R^T without delving into the specifics of a particular marketplace. To this end, we adjust the reward mechanism of (11) to an equivalent one that does not rely on R^T and is, therefore, easier to compute. In our revised mechanism, all user-selected recommendations are ordered in decreasing rank of user perceived quality (such that $Q_1 > Q_2 > \dots > Q_N$) and each reward is based on the $(M + 1)^{th}$ price P_{M+1} (the highest not shortlisted bid) instead of R^T :

$$R_h = \delta \cdot Q_h \cdot P_{M+1} \quad (14)$$

where $h \in [1..N]$, δ is the reward coefficient and $\delta > 0$. This new mechanism also ensures recommendations are rewarded proportionally to their user perceived quality and is therefore also ideal from the perspective of maximizing social welfare⁷. We base the reward on P_{M+1} (whose value is not known by the bidding agents) so that the market cannot be manipulable by the participants [17]p289. If the reward is based on the prices from the rewarded recommendations, the rewarded agents might be able to affect the market through their prices since they are aware of the history of both rewards and bid prices. Our approach also reduces the possibility of bidding collusions because the reward is based on something that the rewarded agents are unaware of and cannot control.

However, reward mechanism (14), as it currently stands, does not satisfy the system objective of shortlisting the most valuable recommendations in decreasing order of user perceived quality. This is because all individually rational agents will bid the same price (marginally higher than P_{M+1}) to

⁷As each R_h ($h \in [1..N]$) is known, the value of the total payoff $\sum_{h=1}^N R_h$ is also known. Among all possible allocations for this total payoff, mechanism (14) ensures the maximal social welfare according to section 2.2.3.

maximize their revenue. This is because a bidder's revenue is the reward obtained minus the bidding price that has been paid and, hence, a rational bidder should bid as low as possible to be shortlisted. When all shortlisted recommendations have the same bidding price, the system cannot differentiate and rank them by price. Therefore, we need a mechanism that can relate and regulate the bidding price according to user perceived quality (i.e. higher quality means a higher price).

To achieve this, we involve two other variables: P_h ($h \in [1..N]$) and P_m^* ($m \in [1..M]$). P_h is the bidding price of the h^{th} rewarded recommendation (user-selected recommendation with the h^{th} highest user perceived quality). P_m^* is the historical average bidding price of the m^{th} shortlisted recommendation during the system's lifetime (note the bidding agents do not actually know this value). By this definition, P_m^* indicates the price that the majority of bidders are willing to pay for the m^{th} advertisement displayed in the user's browser sidebar.

With this additional information, we can now fine-tune the reward mechanism towards the system objective. Instead of (14), we adjust the reward to the h^{th} rewarded recommendation to:

$$R_h = \delta \cdot Q_h \cdot P_{M+1} - \alpha \cdot |P_h^* - P_h| \quad (15)$$

where α is another system coefficient and $\alpha > 1$. The specific values of δ and α are not yet defined and their values will depend upon the specifics of the application.

The reward mechanism in (15), compared with (14), gives recommending agents the incentive to adjust their bids to different levels according to their belief about the corresponding user perceived quality. With (15), the market can differentiate shortlisted recommendations by price so that the marketplace can shortlist good recommendations in decreasing order of user perceived quality. Moreover, under certain conditions, mechanism (15) will tend to be the same as mechanism (14) (to be discussed in section 2.4).

2.3 Recommender Bidding Strategies

Rational bidders seek to maximize their revenue and they do this by bidding sensibly for recommendations that they believe are valuable to the user. The outcome of such bids is that the corresponding recommendation is: not shortlisted, shortlisted but not rewarded, or rewarded. Depending on what happened to its previous bid for the given recommendation, a rational bidder should base the bidding price of its next bid (P^{next}) for that recommendation on (i) the internal quality, (ii) the last bid price (P^{last}) and (iii) the previous rewards to this recommendation. Assuming the internal quality for the specific recommendation is unchanged, we need only consider the bidding strategies with respect to price and reward.

2.3.1 Bid Not Shortlisted

This leaves the agent's revenue unchanged since it neither has to pay for its advertising, nor does it receive a reward. The only way to increase revenue is to get the recommendation shortlisted (since this might bring a reward). Therefore, the agent will increase its bidding price for the same recommendation:

$$P^{next} = Y \cdot P^{last} \quad (Y > 1)$$

This is the dominant strategy in this case since being shortlisted is the only way of increasing revenue.

2.3.2 Bid Shortlisted But Not Rewarded

These agents lose revenue since they pay for the advertising but receive no reward. This means the agent overrated its internal quality with respect to the user perceived quality and so the agent should decrease its price in subsequent rounds so as to lose less:

$$P^{next} = Z \cdot P^{last} \quad (0 < Z < 1)$$

This is the dominant strategy in this case since keeping the same price or even raising it will result in further losses.

2.3.3 Bid Rewarded

These agents have a good correlation between their internal quality for a recommendation and that of the user perceived quality. Therefore, these agents have a chance of increasing their revenue. The profit made by the h^{th} rewarded recommendation is:

$$\xi_h = \delta \cdot Q_h \cdot P_{M+1} - \alpha \cdot |P_h^* - P_h| - P_h$$

However, since the agent is unaware of P_h^* , it does not know whether ξ_h has been maximized. Hence, what it could do is to minimize ($\alpha \cdot |P_h^* - P_h| + P_h$) so as to maximize ξ_h . Furthermore, the agent does not know whether P_h is higher or lower than P_h^* . In either case, however, the agent will definitely make a loss if P_h is not close to P_h^* (proof below).

Assume the set of recommending agents remains unchanged between successive auctions (we discuss what happens when this situation does not hold in section 2.4). The user perceived quality for the h^{th} rewarded recommendation will remain in the h^{th} place in subsequent auctions. Given this, P_h^* is related to Q_h , such that the agent with the h^{th} rewarded recommendation is able to estimate the value of P_h^* . Now consider the design of the strategy for the h^{th} rewarded recommendation. We find that the h^{th} rewarded agent can always be aware of whether its price is closer to or farther from the h^{th} historical average market price, P_h^* , by adjusting its bidding prices. In this way, the agent can minimize its loss. The proof is given below.

Assumptions [static marketplace]:

(i) The h^{th} rewarded recommendation remains the h^{th} highest user perceived quality in subsequent bids. (ii) P_h^* remains stable in subsequent bids. (iii) There are sufficient bidders in the market with not-shortlisted increasing prices and shortlisted but not rewarded decreasing prices to ensure P_{M+1} remains stable. (iv) $\Delta P > 0$, which represents an increment or a decrement of bidding prices.

Proposition:

If the h^{th} rewarded recommendation's current bidding price is below the historical average market price ($P_h < P_h^*$), increasing the price by ΔP and still being below the average price ($P_h + \Delta P$, with $P_h + \Delta P < P_h^*$) results in an increase in profit; decreasing the price by ΔP , ($P_h - \Delta P$) results in a decrease in profit. If the h^{th} rewarded recommendation's current bidding price is above the historical average ($P_h > P_h^*$), increasing the price by ΔP , ($P_h + \Delta P$) results in a decrease in profit; decreasing the price by ΔP and still remaining above the average price ($P_h - \Delta P$, with $P_h - \Delta P > P_h^*$) results in an increase in profit.

Proof:

According to assumptions (i) and (ii), P_h^* is unchanged with respect to the h^{th} rewarded recommendation. So the corresponding agent can estimate the value of P_h^* .

When $P_h < P_h^*$, its profit in the current bid is:

$$\begin{aligned} \xi_{hl} &= \delta \cdot Q_h \cdot P_{M+1} - \alpha |P_h^* - P_h| - P_h \\ &= \delta \cdot Q_h \cdot P_{M+1} - \alpha P_h^* + (\alpha - 1)P_h \end{aligned}$$

current price	adjustment	$ P_h^* - P_h $	$\Delta\xi$
$P_h < P_h^*$	$+\Delta P$	\searrow	> 0
	$-\Delta P$	\nearrow	< 0
$P_h > P_h^*$	$+\Delta P$	\nearrow	< 0
	$-\Delta P$	\searrow	> 0

Table 1: Price Adjustment and Results

Given that P_{M+1} is stable (assumption (iii)), if the agent raises the price by ΔP in the next bid, its profit in the next bid will be:

$$\begin{aligned}\xi_{hli} &= \delta \cdot Q_h \cdot P_{M+1} - \alpha |P_h^* - (P_h + \Delta P)| - (P_h + \Delta P) \\ &= \delta \cdot Q_h \cdot P_{M+1} - \alpha P_h^* + (\alpha - 1)P_h + (\alpha - 1)\Delta P\end{aligned}$$

Since $\alpha > 1$ and $\Delta P > 0$, $\xi_{hli} - \xi_{hl} = (\alpha - 1)\Delta P > 0$.

When $P_h < P_h^*$, if the agent decreases the price by ΔP in the next bid, its profit will be:

$$\begin{aligned}\xi_{hld} &= \delta \cdot Q_h \cdot P_{M+1} - \alpha |P_h^* - (P_h - \Delta P)| - (P_h - \Delta P) \\ &= \delta \cdot Q_h \cdot P_{M+1} - \alpha P_h^* + (\alpha - 1)P_h - (\alpha - 1)\Delta P\end{aligned}$$

Therefore, $\xi_{hld} - \xi_{hl} = -(\alpha - 1)\Delta P < 0$.

The case when $P_h > P_h^*$ can be proven in the same way. ■

This proof tells us that a rational rewarded bidder will adjust its price to the corresponding average market price to maximize its profit. The proof also indicates that, whatever its current price is with respect to the historical average, when adjusting the bid price, if the adjustment results in making less profit, it indicates the action is wrong and $(P_h \pm \Delta P)$ is farther from P_h^* ; if it results in making more profit, it indicates the action is right and $(P_h \pm \Delta P)$ is closer to P_h^* . This phenomenon is listed in Table 1 ($\Delta\xi$ represents the possible profit of the next bid compared to that of the current bid). Table 1 also specifies the strategy for the rewarded agents. This strategy (to bid closer to the corresponding historical average market price) is the dominant strategy for the rewarded agents since otherwise they will definitely receive a loss of revenue. The actual value of ΔP will be defined in an application specific manner.

2.4 Market Equilibrium

According to the strategy for rewarded bidders, such bidders must bid in a manner that aligns their internal view of quality with that of the user. Thus, over time, each individual recommending agent improves its correspondence between its bid price and the user's preferences for recommendations. Only by achieving this can an agent maximize its profit. How quickly this convergence occurs depends on the adjustment of price ΔP .

Under the assumption of a static marketplace (section 2.3), the market reaches an equilibrium. The h^{th} historical average market price reflects the market equilibrium price: thus, at a certain price, the quantity of demand of the h^{th} advertisement slot equals the quantity of the supply (see Figure 5(a)⁸). In the long run, however, these assumptions will not hold and the equilibrium will tend to be broken. However, this new market situation will gradually tend towards another equilibrium and will reach it as long as the changes in the recommending agents are not too frequent

⁸Strictly speaking, the demand curve should be discrete in this case. And the quantity of supply is 1 since we differentiate between each of the M slots and there is only one h^{th} slot. To simplify the discussion, however, we use a continuous demand curve in this context.

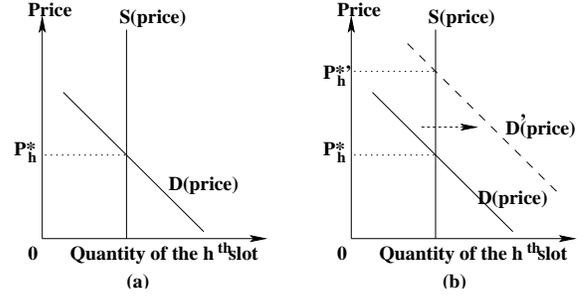


Figure 5: Market Equilibrium and Its Change

(a) The supply curve S is vertical indicating that whatever the deal is, the supply of the h^{th} advertisement slot is constant. The demand curve D has a slope indicating that more agents are willing to pay a low price and few agents are willing to pay a high price for the same slot. The cross indicates that at a certain price level the quantity of demand equals that of supply. This cross point represents the market equilibrium. (b) At each price level, more recommendations become available and the demand curve shifts to the right.

with respect to convergence times (see Figure 5(b)). If, for example, there is more demand in the system, the demand curve will shift right compared to (a). This means that at each price level, there are more bidders willing to pay for the same advertisement slot (because, for example, more better recommendations are being produced).

At equilibrium, since the bidding prices are aligned with the user perceived quality, the system can produce a shortlist of recommendations in decreasing order of user perceived quality which is precisely the objective of the recommender system.

3. EVALUATING THE MARKETPLACE

This section evaluates the market mechanism design with respect to the desiderata of section 2.

- *Pareto Efficiency:* With the reward mechanism defined in (15), the historical average market price, P_h^* , reflects how the majority of bidders value a given advertisement slot and this price becomes the expected equilibrium price. With such a reward mechanism, each bidder iterates itself to the corresponding expected equilibrium price. Therefore, the market has a tendency to converge to the equilibrium. With the market tending to equilibrium, the second term in reward mechanism (15) tends to zero. Therefore, this mechanism tends to be the same as mechanism (14), which is the ideal Pareto efficient mechanism.

- *Social Welfare Maximization:* With the market tending to equilibrium, reward mechanism (15) tends to be the same as mechanism (14). Thus, this reward mechanism tends to reward all user selected recommendations in a manner that is proportional to their user perceived quality. Therefore, (15) maximizes social welfare and is the most socially-preferred.

- *Individual Rationality and Stability:* According to the analysis in section 2.3, the market mechanism produces individually rational dominant strategies for the cases in which the bids are not shortlisted, shortlisted but not rewarded and rewarded. With all agents taking their dominant strategy, the market will dynamically reach the equilibrium and this equilibrium is stable since the market always tends towards it (section 2.4).

- *Fairness*: With our reward mechanism, no greedy strategies (bid as high as possible in order to be shortlisted irrespective of the agent's internal quality evaluation) can survive in this market. Although they can guarantee to be shortlisted, they cannot guarantee being selected if their recommendation is of insufficient quality. Thus bidding higher than its internal quality measure means the agent will make a loss in revenue. Therefore, the prices of the rewarded recommendations cannot continue to rise indefinitely with respect to P_{M+1} . In return, P_{M+1} remains relatively stable so all bidders have equal opportunity of being shortlisted (they simply have to bid higher than P_{M+1}). Hence, the system is fair to all recommending agents.

4. RELATED WORK

A number of information filtering tools (e.g. [18], [16], [1]) have been developed to cope with the problem of information overload. However, these systems tend to filter based on document content and in many cases in our Web browsing domain issues such as quality, style and other machine unparseable properties are the key to giving good recommendations [15]. Thus, recommender systems have been advocated. Developed systems include GroupLens [10] (collaborative filtering of Usenet news to help people find articles they will like in the huge stream of available articles), Ringo [15] (automates word-of-mouth recommendations by weighting user votes to recommend music albums and artists), and Memior [5] (uses trails to support users in finding colleagues with similar interests). While such recommender systems tackle the weaknesses of content-based filtering techniques, each system employs a variety of techniques that are more or less successful for particular users in particular contexts. For this reason, in [12] we developed an extensible multiagent recommendation system for Web documents that incorporates multiple recommendation methods into a single system and for the reason discussed in section 1 we used a market-based approach to coordinate the different methods. However, as also noted in section 1, this system had a number of limitations. Related to this, [2] developed a system that used a competitive market-based allocation of consumer attention space as a means of investigating the user's behaviour in making choices when faced with multiple items and [6] used a variety of marketplaces to moderate a range of digital library services. However, neither of these systems are specifically targetted at making recommendations.

5. CONCLUSIONS AND FUTURE WORK

This paper has outlined an architecture for a recommender system that incorporates multiple heterogeneous recommendation methods. Each recommending agent competes in a marketplace to have their recommendation displayed to the user. Specifically, we designed the auction protocol and the reward mechanism that should be deployed in our application and analyzed the strategies that the individual bidding agent should employ. Our design was shown to be Pareto efficient, social welfare maximizing, stable and fair to all participants. Using our marketplace, the recommender system should be able to put forward the best recommendations to the user. For the future, however, we need to incorporate this mechanism design into our existing recommendation system to determine whether the theoretical properties of our mechanism hold in practice.

6. REFERENCES

- [1] Autonomy. Agentware i3. Technical report, Autonomy, Inc., US, 1997.
- [2] S. M. Bohte, E. Gerding, and H. L. Poutre. Competitive market-based allocation of consumer attention space. In *Proc. of the 3rd ACM Conf on Electronic Commerce*, pages 202–205, US, 2001.
- [3] J. S. Breese, D. Heckerman, and C. Kadie. Empirical analysis of predictive algorithms for collaborative filtering. In *Proc. of the 14th Conf on Uncertainty in Artificial Intelligence*, pages 43–52, USA, 1998.
- [4] S. H. Clearwater, editor. *Market-Based Control. A Paradigm for Distributed Resource Allocation*. World Scientific, 1996.
- [5] D. De Roure, W. Hall, S. Reich, and A. Pikrakis *et al*. Memoir - an open framework for enhanced navigation of distributed information. *Information Processing and Management*, 37:53–74, 2001.
- [6] E. H. Durfee, D. L. Kiskis, and W. P. Birmingham. The agent architecture of the university of michigan digital library. *IEE Proc on Software Engineering*, 144(1):61–71, 1997.
- [7] S. El-Beltagy, W. Hall, D. De Roure, and L. Carr. Linking in context. In *Proc. of the 12th ACM Conf on Hypertext and Hypermedia*, pages 151–160, Denmark, 2001. ACM Press.
- [8] D. Goldberg, D. Nichols, B. Oki, and D. Terry. Using collaborative filtering to weave an information tapestry. *Comm of the ACM*, 35(12):61–70, 1992.
- [9] W. Hill, L. Stead, M. Rosenstein, and G. Furnas. Recommending and evaluating choices in a virtual community of use. In *Proc. on Human Factors in Computing Systems*, pages 194–201, 1995.
- [10] J. A. Konstan, B. N. Miller, D. Maltz, and J. L. Herlocker *et al*. Grouplens: applying collaborative filtering to usenet news. *Comm of the ACM*, 40(3):77–87, 1997.
- [11] R. M. Losee. Minimizing information overload: The ranking of electronic messages. *Journal of Information Science*, 15(3):179–189, 1989.
- [12] L. Moreau, N. Zaini, J. Zhou, and N. R. Jennings *et al*. A market-based recommender system. In *Proc. of the AOIS02 Workshop*, Bologna, 2002.
- [13] P. Resnick and H. R. Varian. Recommender Systems. *Comm of the ACM*, 40(3):56–58, 1997.
- [14] T. W. Sandholm. Distributed rational decision making. In G. Weiss, editor, *Multiagent Systems: A Modern Approach to Distributed Artificial Intelligence*, pages 201–258. MIT Press, US, 1999.
- [15] U. Shardanand and P. Maes. Social information filtering: algorithms for automating word of mouth. In *Proc. on Human Factors in Computing Systems*, pages 210–217, 1995.
- [16] B. Sheth and P. Maes. Evolving agents for personalized information filtering. In *Proc. of the 9th Conf on Artificial Intelligence for Applications (CAIA '93)*, pages 345–352, Orlando, 1993.
- [17] H. R. Varian. *Intermediate Microeconomics: A Modern Approach*. Norton, NY, 6th edition, 2003.
- [18] T. Yan and H. Garcia-Molina. SIFT—A tool for wide-area information dissemination. In *Proc. 1995 USENIX Technical Conf*, pages 177–186, US, 1995.