

Maximizing profit using recommender systems

Aparna Das^{*}
Brown University
Providence, RI
aparna@cs.brown.edu

Claire Mathieu[†]
Brown University
Providence, RI
claire@cs.brown.edu

Daniel Ricketts
Brown University
Providence, RI
daniel.bmore.ricketts@gmail.com

ABSTRACT

Traditional recommendation systems make recommendations based solely on the customer's past purchases, product ratings and demographic data without considering the profitability of the items being recommended. In this work we consider the question of how a vendor can directly incorporate the profitability of items into its recommendation system so as to maximize expected profit while still providing accurate recommendations. Our approach uses the output of any traditional recommender system and adjusts it according to item profitability. Our approach is parametrized so the vendor can control the amount of deviation between the recommendation incorporating profits and the traditional recommendation. We study our approach under two settings and show that it can achieve significantly more profit than traditional recommendations.

Categories and Subject Descriptors

H.3.5 [Information Systems]: Information storage and retrieval—*Online-information services, Commercial services*

Keywords

Recommender systems, Product profitability, Algorithms, Electronic commerce.

1. INTRODUCTION

Recommendation Systems are important tools for major companies such as Amazon, Netflix and Pandora. Recommendation Systems use a customer's demographic data, past purchases and past product ratings to predict how the customer will rate new products [1, 11, 9, 13]. They have been shown to help customers become aware of new products, increase sales and encourage customers to return to the business for future purchases [6, 15]. Designing recommendation systems that accurately predict customer ratings has generated much research and interest both in the academic and business communities. The Netflix prize was a manifestation of this interest [12]. However, the majority of the work on recommender systems has not explicitly considered how the profitability of products could be incorporated

into the recommendations. An article published in Knowledge@Wharton claims that the actual Netflix recommendation system modifies its ratings to encourage consumers to order the more obscure movies which are presumably cheaper for Netflix to supply than major blockbusters[18]. While Netflix does not, publicly reveal whether it uses such methods, it seems natural for a business to incorporate the profitability of products into its recommendations. From the viewpoint of the customer, recommendations are helpful suggestions, but from the viewpoint of the vendor they are extremely targeted advertisements, and the explicit goal of advertising is to increase profit.

In this paper, we study the question of how a vendor might incorporate the profitability of items into its recommendations. A naive approach is to give the most profitable items the highest recommendations. Then these items would presumably be bought more often and the business would make more money. However this tactic has some obvious flaws. While the customer may initially follow the vendor's recommendation, she may find that she does not like the items as much as the vendor predicted. After only a few such experiences, she would realize that the vendor's recommendations do not accurately reflect her tastes. In the best case for the vendor, the customer would ignore the vendor's recommendations and continue her natural purchasing behavior while in the worst case she would lose trust not only in the vendor's recommendations but also in the vendor as a whole and take her business elsewhere. Thus incorporating the profitability of items into recommendations must be done carefully so that customer's trust is not compromised.

A reasonable assumption is that as long as the vendor consistently presents recommendations that are similar enough to the customer's own ratings, the customer will maintain a high level of trust in the accuracy of the vendor's recommendations. For this reason, we use an established *similarity measure* as a measure of trust. We assume that the vendor has access to a vector \vec{c} giving the consumer's true ratings for items. The vendor's objective is to present a recommendation vector \vec{r} to the customer which is within a certain threshold of similarity to \vec{c} , and maximizes the vendor's expected profit (by incorporating the profitability of items into the recommendation). Section 2 gives details of the model.

One question that might arise is how the vendor can determine the customer's true ratings (i.e \vec{c}) if the ratings are for products that the consumer has not yet rated herself. There is a large amount of research and activity in developing highly accurate recommendation systems that solve this problem and we assume the predictions from these systems

^{*}Supported by NSF grant CCF-0728816.

[†]Supported by NSF grant CCF-0728816.

for well established customers are good approximations to the customer's true ratings.

The idea that the customer will maintain high trust as long as the vendor's recommendations are similar to her own ratings has empirical support. Hill et al. showed that users asked to rate the same item at different times supply different ratings [8]. Amatriain et al. report similar findings in their study which evaluates noise in user ratings of movies [2]. Thus there is some natural variability in the ratings that users supply, and so slight differences between the predicted ratings and the actual customer ratings should be insignificant to the customer.

Chen et al. also considered using the profitability of items in recommender systems [5], but they do not explicitly require a level of accuracy in their system.

The outline of our paper is: section 2 defines our model and problem, section 3 describes the similarity measure and gives justification for its use as a measure of customer trust and section 4 describes our approach for maximizing expected profit in different scenarios.

2. MODEL

Let n be the number of items being sold by the vendor. Let \vec{c} and \vec{r} be the vectors of length n where the i th components denoted c_i and r_i gives a rating for item i . All items are rated using numbers between zero and some maximum rating m .

We focus on the scenario where the vendor is interacting with an established customer who the vendor would like to continue to do business with in the long term. We assume the vendor uses a recommendation system which makes good predictions about how such a customer rates items and that vector \vec{c} gives these predictions¹. The vendor presents the customer with recommendation vector \vec{r} to help her decide which item to purchase. In many real applications, the vendor only wants to provide the customer with the top n recommended items. In this case the vendor would display the n highest rated items in \vec{r} .

The customer has a certain level of *trust* in the accuracy of the vendor's recommendation which she has developed from experience with the vendor's past recommendations. When choosing an item to purchase, the customer considers \vec{r} based on how much trust she has in the vendor's recommendations. The exact influence that a customer's level of trust in the vendor's recommendations has on her purchasing behavior is too complex to model precisely. We make the following simplifying assumptions to allow for a model of that behavior: First we assume that the consumer's trust is closely tied to how similar the vendor's recommendations are to the ratings she would give items. Second we assume that as long the similarity constraint is consistently met by the vendor, the consumer will use the recommendations when deciding what items to purchase. Specifically we assume there is a function $T(\vec{r})$ that assigns a scalar value to the notion of similarity between \vec{r} and \vec{c} where higher values indicate greater similarity. Second, we assume that if for every \vec{r} that the vendor presents to the customer, $T(\vec{r})$ meets or exceeds some threshold value τ , then the customer's trust in the recommendation system will remain at a constant, significantly high level. Finally, we assume that if the cus-

¹The customer does not necessarily know all the entries of \vec{c} herself as she has not purchased all items.

tomers trust remains at this level, her purchasing decision will be solely a function of \vec{r} . The customer's level of trust and her subsequent purchasing behavior is unknown if the vendor presents \vec{r} such that $T(\vec{r}) < \tau$.

The intuition behind these assumptions is that, as long as the vendor recommendation consistently predicts ratings that are similar enough to how the customer would rate the items, the customer will maintain a certain level of trust in the accuracy of the recommendation system. As a result, she will use the vendor's recommendation as information when deciding which items to purchase. However if \vec{r} is too far from \vec{c} , the customer loses trust the vendor's recommendations and no longer considers them when deciding which item to buy. In that case, her purchasing behavior is unclear.

The vendor's main goal is to maximize profit. However, to maintain customer trust he is required to present vector \vec{r} such that $T(\vec{r}) \geq \tau$ for some constant τ . Denote $\vec{\varphi}(\vec{r})$ be a vector valued function whose i th component gives the probability that the customer will purchase item i at a given time step. The customer, at any step, can purchase zero or more items, so the components of $\vec{\varphi}(\vec{r})$ need not sum to one. Let \vec{p} be the profit vector whose i th component, p_i , gives the profit received when item i is purchased. The vendor's expected profit is given by $E_p = \vec{p} \cdot \vec{\varphi}(\vec{r})$. Formally our problem is to maximize the vendor's expected profit while maintaining a level τ of trust with the customer:

$$\text{Max } \vec{p} \cdot \vec{\varphi}(\vec{r}) \quad \text{s.t. } T(\vec{r}) \geq \tau \quad (1)$$

3. SIMILARITY MEASURES FOR TRUST

In this section we argue that the *Dice coefficient*, which measures similarity between two vectors, is an appropriate measure of consumer trust. We briefly discuss why some of the other common singularity measures are lacking.

Dice coefficient. We adopt the *Dice coefficient* given in equation 2, to measure trust $T(\vec{r})$. The Dice coefficient is a popular measure which has previously been used to measure recommendation accuracy [5, 7, 14].

$$\text{Dice}(\vec{r}) = \frac{2 \sum_i c_i r_i}{\sum_i c_i^2 + \sum_i r_i^2} = \frac{2\vec{c} \cdot \vec{r}}{\|\vec{c}\|^2 + \|\vec{r}\|^2} \quad (2)$$

Above $\|\vec{x}\| = \sqrt{\sum_i x_i^2}$ denotes the length of vector x . Normally the Dice coefficient is denoted as a function of the two vectors whose similarity is being measured but here we denote it as only a function of \vec{r} to emphasize the fact that \vec{c} is constant known to the vendor. Let θ denote the angle between \vec{c} and \vec{r} . An equivalent definition of the Dice coefficient is,

$$\text{Dice}(\vec{r}) = \cos(\theta) \cdot \frac{2\|\vec{c}\|\|\vec{r}\|}{\|\vec{c}\|^2 + \|\vec{r}\|^2} \quad (3)$$

We now list some properties of the Dice coefficient that makes it a reasonable function to measure trust.

Property 3.1. *Dice*(\vec{r}) is always between zero and one. *Dice*(\vec{r}) = 1, if and only if $\vec{c} = \vec{r}$ for every item i . *Dice*(\vec{r}) = 0, if and only if $r_i = 0$ on all items i such that $c_i > 0$.

Thus the *Dice*(\vec{r}) is one only when \vec{r} is in complete agreement with \vec{c} , and it is zero only when \vec{r} disagrees with \vec{c} on all relevant items.

PROOF. (Proof of Property 3.1)

$Dice(\vec{r}) \geq 0$ because no item is rated less than zero. Thus the numerator of Equation 2 is always positive. Now suppose that $Dice(\vec{r}) > 1$. By Equation 2, this implies that $0 > \sum_i c_i^2 - 2r_i c_i + r_i^2 = \sum_i (r_i - c_i)^2$ which is a contradiction. Thus $Dice(\vec{r}) \leq 1$.

Suppose that $r_i = c_i$ for every i . Then using Equation 2 $Dice(\vec{r}) = 1$. To prove the other direction suppose that $Dice(\vec{r}) = 1$. The formulation of the Dice coefficient given in Equation 3 can be used to show that $r_i = c_i$ for all i . Note that $0 \leq \cos(\theta) \leq 1$ as no item is rated less than zero. The second term of Equation 3 is non-negative by definition and the following proof by contradiction shows it is at most 1: suppose $2\|\vec{c}\|\|\vec{r}\|/(\|\vec{c}\|^2 + \|\vec{r}\|^2) > 1$ then $2\|\vec{c}\|\|\vec{r}\| > \|\vec{c}\|^2 + \|\vec{r}\|^2$ implying that $0 > (\|\vec{c}\| - \|\vec{r}\|)^2$ which is false. Thus if $Dice(\vec{r}) = 1$ both terms of Equation 3 must be 1. As $\cos(\theta) = 1$ the angle between \vec{c} and \vec{r} is zero and the second term being 1 implies that $(\|\vec{c}\| - \|\vec{r}\|) = 0$ i.e that \vec{r} and \vec{c} have equal length. Together this implies that $r_i = c_i$ for all i .

Finally note that when $Dice(\vec{r}) = 0$, the numerator of Equation 2 is zero which implies that r_i must be zero for each i such that $c_i > 0$. \square

Jaccard measure. The Jaccard similarity measure, given below, behaves similarly to the Dice coefficient and is used widely in information retrieval and data mining [17, 10, 3].

$$Jac(\vec{r}) = (\vec{c} \cdot \vec{r}) / (\|\vec{c}\|^2 + \|\vec{r}\|^2 - \vec{c} \cdot \vec{r})$$

It is another appropriate measure of consumer trust and all results extend to the setting where $T(\vec{r}) = Jac(\vec{r})$.

Cosine measure. The cosine similarity measure, given in equation 4, is equal to the cosine of θ . It is always between zero and one as no item is rated less than zero.

$$Cos(\vec{r}) = \cos(\theta) = \frac{\sum_i c_i r_i}{\sqrt{\sum_i r_i^2} \sqrt{\sum_i c_i^2}} \quad (4)$$

The cosine measure is not influenced by difference in the lengths of \vec{r} and \vec{c} and this is the main reason it seems unsuitable for measuring trust as demonstrated in the following example.

Example. Suppose the items are rated between 0 and 5. Consider a picky customer who rates all items somewhat low and a vendor that recommends all items high, i.e $c_i = 1$ and $r_i = 5$ for all i . The cosine measure will be to one because $\theta = 0$ indicating that the picky customer has high trust for the vendor's recommendation, which is surely not the case.

Mean squared error and distance measures. The mean squared error (MSE) measures the average dissimilarity between \vec{r} and \vec{c} . Let m denote the maximum possible rating. To measure similarity we could use 1-MSE, Equation 5, which is always between zero and one.

$$MSE(\vec{r}) = \frac{\sum_i ((c_i/m - r_i/m))^2}{n} \quad (5)$$

The 1-MSE measure gives equal credit for agreements on items the customer dislikes as on agreements on items she prefers and this makes it unsuitable for measuring the trust as demonstrated by the example below. This example also applies to other distance based similarity measures such as Euclidean distance and Manhattan distance.

Example. Suppose the items are rated between 0 and 5. Let $\vec{c} = [5, 5, 5, 1, 1, 1, \dots, 1]$ which represents a customer who rates a few items very highly but dislikes most items. Consider the following recommendation $\vec{r} = [1, 1, 1, \dots, 1, 5, 5, 5]$ where the vendor gives the highest ratings to a few items that the customer dislikes and gives the low ratings to all other items including the items the customer prefers. MSE is $\Theta(1/n)$ so that 1-MSE approaches 1 as the number of items n gets large indicating that the customer would have high trust such recommendations which is not the case.

4. PROFIT MAXIMIZATION

Using the Dice similarity coefficient defined in Section 3 as the trust function the vendor's optimization problem is,

$$\text{Max } E_p(\vec{r}) = \vec{p} \cdot \vec{\varphi}(\vec{r}) \text{ s.t. } Dice(\vec{r}) = \frac{2\sum_i c_i \cdot r_i}{\sum_i c_i^2 + \sum_i r_i^2} \geq \tau \quad (6)$$

We outline a general approach for solving Equation 6 and then apply our technique on two different objective functions obtained by alternative definitions of $\vec{\varphi}(\vec{r})$. We analyze how much profit the vendor gains by presenting the customer with recommendation \vec{r} rather than \vec{c} .

4.1 General Approach

Calculus gives us a general approach for solving the vendor's maximization problem stated in Equation 1. Adding $1/\tau$ to both sides of the Dice constraint above and simplifying reveals that it is equivalent to,

$$\sum_i \left(r_i - \frac{c_i}{\tau} \right)^2 \leq \left(\frac{1}{\tau^2} - 1 \right) \sum_i c_i^2 \quad (7)$$

As \vec{c} is a constant for our setting, the feasible \vec{r} for our problem lie in a region enclosed by an n-sphere with radius $\sqrt{(\frac{1}{\tau^2} - 1) \sum_i c_i^2}$ which we refer to as the *Dice sphere*.

The general approach to solving maximization problem of Equation 6 involves two parts. The first is to determine if there are any local maxima that lie strictly inside the Dice sphere i.e. which satisfy $Dice(\vec{r}, \vec{c}) > \tau$. The second part is to find the vector \vec{r} that maximizes expected profit over all vectors on the surface of the Dice sphere, i.e. which satisfy $Dice(\vec{r}, \vec{c}) = \tau$. The largest of the local maxima in the sphere and the maximum on the surface is the global maximum.

The gradient of the objective function is zero at each local maximum inside the Dice sphere. Thus all solutions to $\nabla E_p = 0$ are candidate vectors. Let $\vec{r}_1^*, \vec{r}_2^*, \dots, \vec{r}_k^*$ be the list of all vectors that satisfy this property. While these vectors could be local minima or saddle points instead of local maxima, it is not necessary to distinguish between them. It is only necessary to find the greatest value of $E_p(\vec{r}_i^*)$ for all i and compare it to the maximum of all vectors on the surface of the sphere. Let \vec{r}_{in}^* denote the maxima vector inside the Dice sphere.

The maximum vector on the surface of the Dice sphere can be found using the method of Lagrange multipliers. The maximum valued vector \vec{r}_s^* will satisfy $\nabla Dice(\vec{r}_s^*, \vec{c}) = \lambda \nabla E_p(\vec{r}_s^*)$ and $Dice(\vec{r}_s^*, \vec{c}) = \tau$. The maximum of \vec{r}_s^* and \vec{r}_{in}^* is the solution to the optimization problem [4, 16].

4.2 Simple probability function

Consider a scenario where the customer can purchase zero or more items at each time step where the probability that

the customer purchases item i is independent of the vendor's ratings for other items. Here is a simple way to define the probability function which satisfies this assumption,

$$\vec{\varphi}(\vec{r}) = \left(\frac{r_1}{m}, \frac{r_2}{m}, \dots, \frac{r_n}{m} \right) \quad (8)$$

Recall that m is the highest possible rating for an item. With this definition of $\vec{\varphi}$, the probability that a customer purchases an item is linearly proportional to the vendor's rating of that item and the vendor's expected profit is

$$E_p(\vec{r}) = \vec{p} \cdot \vec{\varphi}(\vec{r}) = \frac{1}{m} \sum_i p_i \cdot r_i \quad (9)$$

As prices are all greater than zero, $\nabla E_p = (\frac{p_1}{m}, \dots, \frac{p_n}{m}) \neq 0$ so there are no local maxima inside the sphere. Thus we can proceed to finding the maximum be on the surface of the Dice sphere. Using Equation 7 we have $\nabla Dice(\vec{r}, \vec{c}) = (2(r_1 - \frac{c_1}{\tau}), \dots, 2(r_n - \frac{c_n}{\tau}))$, so the maxima on the Dice sphere surface must satisfy the following system of equations,

$$\begin{aligned} \frac{p_i}{m} &= 2\lambda(r_i - c_i/\tau) \quad \text{for all } i \\ \text{and } \sum_i (r_i - c_i/\tau)^2 &= (\frac{1}{\tau^2} - 1) \sum_i c_i^2 \end{aligned}$$

where the last equation requires \vec{r} to lie within the Dice sphere. Solving the first set of equations we obtain that

$$r_i = \frac{p_i}{2m\lambda} + \frac{c_i}{\tau} \quad (10)$$

Substituting in the value of r_i into Equation 7 we get $\lambda = \frac{1}{2m} \sqrt{\frac{\sum_i p_i^2}{(\frac{1}{\tau^2} - 1) \sum_i c_i^2}}$. The final solution is obtained by plugging λ into Equation 10.

$$r_i = p_i \sqrt{\frac{(\frac{1}{\tau^2} - 1) \sum_j c_j^2}{\sum_j p_j^2}} + \frac{c_i}{\tau} \quad (11)$$

Profit Gains. By presenting the customer with recommendation \vec{r} derived in Equation 11 the vendor earns expected profit $E_p(\vec{r}) = \frac{1}{m} \left(\sum_i p_i^2 \sqrt{(1/\tau^2 - 1) \sum_j c_j^2 / \sum_j p_j^2} + p_i c_i / \tau \right)$, which is simplified via the Cauchy-Schwarz inequality to ²,

$$E_p(\vec{r}) \geq \left(\sqrt{\frac{1}{\tau^2} - 1} + \frac{1}{\tau} \right) \frac{\sum_i p_i c_i}{m}. \quad (12)$$

The expected profit from \vec{c} is $E_p(\vec{c}) = (\sum_i p_i c_i) / m$. The vendor's profit gain from presenting \vec{r} rather than \vec{c} is,

$$\frac{E_p(\vec{r}) - E_p(\vec{c})}{E_p(\vec{c})} = \sqrt{\frac{1}{\tau^2} - 1} + \frac{1}{\tau} - 1 \geq 2(1/\tau - 1).$$

Example. If the vendor presents recommendation vectors that are within similarity threshold $\tau = .9$, allowing a 10% deviation to \vec{c} , then in expectation he earns at least $2(10/9 - 1) > 22\%$ more profit by presenting \vec{r} rather than \vec{c} .

4.3 Simple distribution

Now we consider the scenario where at each time step the customer purchases only one item but chooses which item to purchase based on how its rating compares to ratings of other items. A simple way to model this is to set the

²The Cauchy-Schwarz inequality is $\sum_i p_i^2 \sum_i c_i^2 \geq (\sum_i p_i c_i)^2$

purchase probability of item i to be $r_i / \sum_j r_j$.

$$\vec{\varphi}(\vec{r}) = \left(\frac{r_1}{\sum_i r_i}, \frac{r_2}{\sum_i r_i}, \dots, \frac{r_n}{\sum_i r_i} \right)$$

Thus the customer scales each item by the sum of the vendor's ratings, and then chooses uniformly among all offered items. With this definition of $\vec{\varphi}$ the expected profit is,

$$E_p(\vec{r}) = \vec{p} \cdot \vec{\varphi}(\vec{r}) = \sum_i \frac{p_i \cdot r_i}{\sum_j r_j} \quad (13)$$

The local maxima inside the Dice sphere occur where the gradient of the expected profit is zero. The gradient is

$$\nabla E_p = \left\langle \frac{p_1 \sum_j r_j - r_1 p_1}{(\sum_j r_j)^2}, \dots, \frac{p_n \sum_j r_j - r_n \cdot p_n}{(\sum_j r_j)^2} \right\rangle$$

For $\nabla E_p = 0$, it must be the case that $\sum_j r_j = r_i$ for all i . Thus if there are at least two items with different ratings then there are no local maxima inside the Dice sphere.

We proceed to finding local maxima on the surface of the Dice sphere. Applying the method of Lagrange Multipliers as before we end up having to solve the following system of equations for variables λ and r_i for all i :

$$\frac{p_i \sum_j r_j - r_i \cdot p_i}{(\sum_j r_j)^2} = 2\lambda(r_i - \frac{c_i}{\tau}) \quad \text{for all } i \quad \text{and } Dice(\vec{r}) \geq \tau \quad (14)$$

Unfortunately we do not know how to find a solution for this as the first set of equations involves r_i for all i . Recall that in section 4.2 the corresponding equations for finding the maxima on the Dice surface were each a function of only one r_i . This lead us to seek an alternative approach. We will reduce solving the optimization problem under the definition of E_p given in Equation 13 to solving a series of simpler optimization problems on which we can effortlessly apply the method of Lagrange Multiplier. To do so, first consider the decision version of our problem: *Does there exist a \vec{r} such that $E_p(\vec{r}) \geq V$ and $Dice(\vec{r}) \geq \tau$?*

Under Equation 13, having $E_p(\vec{r}) \geq V$ is equivalent to having $\sum_i (p_i - V)r_i \geq 0$. Thus the decision version of our problem is equivalent to solving the following maximization problem and checking that its solution has $r_i \geq 0$ for all i ,

$$\mathbf{Max} \ E_p(\vec{r}) = \sum_i (p_i - V)r_i \quad \mathbf{s.t.} \quad Dice(\vec{r}) \geq \tau \quad (15)$$

Equation 15 can be solved using the general approach outlined in Section 4.1. The gradient $\nabla E_p(\vec{r}) = 0$ iff $p_i - V = 0$ for all i . Thus as long some item is not priced V , there are no local maxima for Equation 15 inside the Dice sphere ³.

To find the maxima on the surface of the sphere, we solve the following system of equations,

$$p_i - V = 2\lambda(r_i - c_i/\tau) \quad \text{for all } i \quad \text{and } Dice(\vec{r}) \geq \tau \quad (16)$$

Note that Equation 16 differ from Equation 10 only by constants so the solution derived in Section 4.2 shifted by constants is a solution for Equation 16. We get that,

$$r_i = \frac{p_i - V}{2\lambda} + \frac{c_i}{\tau} \quad \text{where } \lambda = \frac{1}{2} \sqrt{\frac{\sum_i (p_i - V)^2}{(1/\tau^2 - 1) \sum_i c_i^2}}$$

³If all item are priced V , all \vec{r} yield expected profit V and we could pick any \vec{r} which lies inside the Dice sphere.

If for all i , $r_i \geq 0$, we return a “yes” for the solution of the decision problem and otherwise we return “no”.

To find a solution for the original optimization problem which is arbitrarily close to optimal, we can do binary search along a bounded interval of possible values of E_p , checking the existences of solutions using decision version algorithm described above. Let $V_{\max} = \max_i p_i$. The initial binary search interval can be set to $[0, V_{\max}]$ since $E_p(\vec{r}) \leq V_{\max}$ as the customer purchases one item per time step. Each binary search step reduces the search interval by half and doing more and more binary search steps brings us closer and closer to the optimal solution for the optimization problem. Let $\delta < 1$. A solution which is within distance $V_{\max}\delta$ of the optimal can be found by performing $\log(\frac{V_{\max}}{\delta})$ binary search steps. For example, let \vec{r}_* denote the optimal solution. With $\delta = 1/V_{\max}$, we can obtain an approximate solution \vec{r}_a such that $E_p(\vec{r}_a) + 1 \geq E_p(\vec{r}_*)$ in $\log(\frac{V_{\max}}{\delta}) = \log(V_{\max}^2) = O(\log V_{\max})$ binary search steps.

Profit Gains. Thus we are able to find a near optimal solution to the optimization problem with $E_p(\vec{r})$ as defined in Equation 13 by solving a series of optimization problems of the kind solved in Section 4.2. The profit gains analysis from section 4.2 extends to the last “yes” solution obtained for a decision problem. However as this “yes” solution is near optimal for the original optimization problem, the profit gains will be close to that from section 4.2.

5. CONCLUSIONS AND FUTURE WORK

Traditional recommendation systems do not directly incorporate the profitability of items into its recommendations and in this work we propose one of the first direct methods to do so. Our approach maximizes vendor profit while providing trustworthy recommendations for the customer. While our approach is simple, its simplicity allows it to be used in conjunction with any traditional recommendation system. Our method is also tunable and allows the vendor to control how much the profit based recommendation should deviate from the traditional recommendation. Our work is a starting point and we hope it will simulate new research on incorporating profits into recommendation systems. The main future direction for this work is conduct user studies to verify our assumption that the Dice coefficient is a suitable measure of customer trust. It would also be useful to analyze how the customer’s trust is affected small deviations from traditional recommendations systems.

6. REFERENCES

- [1] ADOMAVICIUS, G., AND TUZHILIN, A. Toward the next generation of recommender systems: a survey of the state-of-the-art and possible extensions. In *IEEE Trans. Know. Data. Eng.* (2005), IEEE, pp. 734–749.
- [2] AMATRIAIN, X., PUJOL, J., AND OLIVER, N. I like it... i like it not: Evaluating user ratings noise in recommender systems. In *Proceedings of the Conference on User Modeling, Adaptation, and Personalization* (2009).
- [3] BERRY, M., AND BROWNE, M. *Lecture Notes in Data Mining*. World Scientific, 2006.
- [4] CAIN, G., AND HEROD, J. *Multivariable Calculus*. www.math.gatech.edu/~cain/notes/calculus.html, 1997.
- [5] CHEN, L., HSU, F., CHEN, M., AND Y., H. Developing recommender systems with the consideration of product profitability for sellers. In *Information Sciences* (2008), Elsevier, pp. 1032–1048.
- [6] CHEN, P., AND WU, S. How does collaborative filtering technology impact sales? empirical evidence from amazon.com. working paper. Available at SSRN: <http://ssrn.com/abstract=1002698>, 2007.
- [7] CHO, Y., AND KIM, J. Applications of web usage and product taxonomy to collaborative recommendations in e-commerce. In *Expert systems with Applications* (2004), Elsevier, pp. 233–246.
- [8] HILL, W., STEAD, L., ROSENSTEIN, M., AND FURNAS, G. Recommending and evaluating choices in a virtual community of use. In *Proceedings of ACM CHI’95 Conference on Human Factors in Computing Systems* (1995), ACM Press, pp. 194–201.
- [9] J.S., B., HECKERMAN, D., AND C, K. Empirical analysis of predictive algorithms for collaborative filtering. In *Proceedings of the 14th Conference of Uncertainty in AI* (1998), Morgan Kaufmann Publishers, pp. 42–52.
- [10] KAUFMAN, L., AND ROUSSEEUW, P. *Finding Groups in Data, An Introduction to Cluster Analysis*. Wiley, New York, 1990.
- [11] LINDEN, G., B., S., AND YORK, J. Amazon.com recommendations:item-to-item collaborative filtering. In *IEEE Internet Computing* (2003), IEEE, pp. 76–80.
- [12] NETFLIX, INC. Netflix prize., Dec 2008.
- [13] RESNICK, P., AND VARIAN, H. R. Recommender systems. In *Communication of ACM* (1997), ACM Press, pp. 56–58.
- [14] SARWAR, B., KARYPIS, G., KONSTAN, J., AND RIEDL, J. Analysis of recommendation algorithms for e-commerce. In *Proceedings of the 3rd ACM Conference on Electronic Commerce* (2000), ACM Press, pp. 158–167.
- [15] SCHAFER, J., KONSTAN, J., AND RIEDL, J. Recommender systems in e-commerce. In *Proceedings of the 2nd ACM Conference on Electronic Commerce* (1999), ACM Press, pp. 158–166.
- [16] STEWART, J. *Multivariable Calculus 5e (Fifth Edition)*. Thompson-Brooks/Cole, Canada, 2003.
- [17] STREHL, A., AND GHOSH, J. Value-based customer grouping from large retail data-sets. In *Proceedings of SPIE Conference on Data Mining and Knowledge Discovery* (2000), pp. 33–42.
- [18] WHARTON. Reinforcing the blockbuster nature of media: The impact of online recommender, Dec 2008.