

Recommendations for Long-Term Profit Optimization

Patrick Hosein, Inzamam Rahaman, Keanu Nichols, Kiran Maharaj

patrick.hosein@sta.uwi.edu,inzamam@lab.tt,nichols.keanu9@gmail.com,kiranma@hotmail.com

The University of the West Indies, St. Augustine, Trinidad and Tobago

ABSTRACT

Recommender Systems have traditionally sought to identify the most relevant products for a customer with the aim of maximizing expected purchases. While this may be an appropriate objective for services such as Netflix and Spotify, it may not necessarily be so for others. For example, an insurance firm may want to recommend the most suitable insurance plan to a customer but it may also want to take into account the profitability of the product being recommended. This, however, is a delicate trade-off since if the offered product is not sufficiently suitable then the customer may switch to a competitor. Therefore, the product should be sufficiently likable to keep the customer while also being sufficiently profitable. We consider this problem and introduce a recommender system that picks the products that maximize the long-term profit for the company. In this way the company benefits, by having acceptable profits from a long-term customer, while the customer benefits by receiving satisfactory recommendations. The long-term profit is the sum of the immediate profit (i.e., for the immediately recommended product) and the expected profit of future product purchases made by the customer.

KEYWORDS

Recommender Systems, Profit Optimization, Long-Term Profit, Naive Bayes

1 INTRODUCTION

In a Recommender System, one uses historical customer information to determine a suitable product to recommend to a present or new customer. Such determinations are typically made based on the users' past interactions and the underlying characteristics associated with the user and the domain's items. Recommender systems are primarily modelled on user preferences and are developed as a personalized service made available on major e-commerce and leisure websites such as Amazon, Netflix, Spotify, Pandora and YouTube.

Research into Recommender Systems has traditionally been user-focused with the emphasis being on the maximization of the predictive accuracy of the recommender system. While it is important to maximize the utility of the system

for the user, perhaps of equal importance should be the utility or business value to the company providing the product. The development of a recommender system is typically motivated by the assumption that assisting users in finding more relevant products would lead to future purchases from the vendor. However, there is merit in including the profitability of items within the recommendation process itself [2, 3, 7, 13, 19].

In traditional Recommender Systems, a company recommends a product or service with the highest probability of being accepted by the customer. However, if the chosen product is not profitable then it may not be in the company's best interest to offer it. On the other hand one can recommend the product with the highest profitability but such an item may be unlikely to be purchased by the customer and hence the sale is lost. A third option is to choose the product with the highest expected profitability, i.e., the product for which the probability that the customer chooses the product times the profit associated with the product is the largest. Although this approach maximizes the immediate expected profit, this choice may not maximize the long term profits possible from the customer [1, 13]. In this paper we instead maximize the long-term profit which takes into account the short-term profit as well as the long term retention (and hence continued profits) of the customer. Such a model has not been widely studied in the literature but we believe that it most appropriately represents the objective of the company while at the same time providing benefits to the customer.

2 RELATED WORK AND CONTRIBUTIONS

The work in [4] addresses the profit maximization problem but uses a different approach. They attempt to maximize the sum of the prices of the recommended products. They then define a trust factor which represents the difference between the profit-based recommendations and those that would have been made if chosen independently of profit. The optimization problem includes a constraint with a lower bound on the trust. In other words, only products that are relatively close to those that are highly probable to be chosen by the customer (and hence maintains the customer's trust) are chosen. In our case we take into account the trust factor by including it directly into the objective function of the proposed model.

Similarly, the paper [12] considers other objectives in addition to profit maximization. In addition to the most likely

item to recommend, one may want to consider whether the item is in stock. Also, the user may already be familiar with popular items and hence it might be more beneficial to recommend a less popular item that they may like. The latter case is addressed by assigning weights to rankings and re-ordering. The former is addressed by using thresholds as was done in [4]. However these require the specification of weights which may have to be constantly adjusted.

The paper by Jannach et. al. [13] summarizes the various approaches described above, but it also brings up the issue of long-term revenue, although it is not directly addressed. Their conclusion was that the incorporation of the long term perspective has been largely under-explored and hence we believe that our proposed model is new. The paper by Hosanagar et. al. [10] also addresses revenue optimization but only the short-term revenue optimization case. Additional research, but similar to the above, can be found in [2, 5, 7, 8, 14, 15, 17, 18, 20–22].

Our contributions include, (a) a model for optimizing long-term revenue when making recommendations, (b) a simple illustrative example to show the benefits that can be gained, (c) a more realistic example to illustrate the potential gains of looking at the long term rather than short term profits. Note that the underlying premise is that, typically, high demand items (e.g. those that are inexpensive or placed on sale) tend to have low profitability.

3 PROBLEM FORMULATION

We consider recommendations of a set of products to a set of customers using collaborative filtering. In other words, the recommendation made to a customer is based on the products chosen by other customers with similar features or purchasing habits. Prior work considered either maximizing the probability that the recommended product is chosen (i.e., accuracy) or maximizing the profit achieved by the company. Note that these objectives can conflict.

In this paper we provide a different formulation that captures both aspects of accuracy and profit by considering the long-term profit that can be achieved from a customer. Let us assume that a customer is to be provided with a sequence of product recommendations over time. We assume that the set of available products is continuously updated and hence a suitable product is always available for each customer. At each offering of a product, the customer can either purchase the product (and the company achieves the associated profit or reward) or ignore the product. Furthermore, if the product is not a suitable recommendation for the customer then the customer may decide to stop making purchases from this company because of a loss of trust in their recommendations, i.e., the customer believes that the company is not making its best offers but rather trying to push its profitable products.

We assume that K products are available and we index these by k . Let R_k denote the reward or profit that the company receives if the customer is offered the product and it is purchased. Let p_k denote the probability that the customer purchases the product if recommended. For example, one can use techniques such as Naive Bayes to determine this probability. Therefore, the expected profit (reward) for product k is $p_k R_k$. One can then find the product for which this is maximum if the intent is to maximize revenue only for this purchase. Let q_k denote the probability that the customer rejects the product if offered and also stops using the company (loss of trust). We assume that a product is offered at most once to a customer. Hence, when a product is recommended, it is removed from the set of available products and a new product (taken from the same distributions for p_k , R_k and q_k) is added to the list of products. We assume that we are in steady state so that the products before and after a recommendation have the same statistical properties.

Let \bar{R} denote the expected long term reward for the concerned customer and let R denote the maximum expected total reward given the present recommendation (i.e, the maximum over products of the expected value of the sum of the present purchase reward plus the expected future rewards). Note that future rewards are zero if the customer is not retained. We can therefore write:

$$R = \max_k \{p_k R_k + (1 - q_k)\bar{R}\} \quad (1)$$

Note that we do not know the value of \bar{R} . However, assuming stationarity, in steady state the expected value of R will be \bar{R} . Assuming the variation of R from one recommendation to the next is small we make the approximation $R \approx \bar{R}$. Using this approximation we can now solve for \bar{R}

$$\bar{R} = \max_k \{p_k R_k + (1 - q_k)\bar{R}\} \quad (2)$$

Subtracting \bar{R} from both sides we have

$$\max_k \{p_k R_k - q_k \bar{R}\} = 0 \quad (3)$$

which can be re-written as

$$\max_k \left\{ q_k \left\{ \frac{p_k R_k}{q_k} - \bar{R} \right\} \right\} = 0. \quad (4)$$

Consider the following equation

$$\max_k \left\{ \frac{p_k R_k}{q_k} - \bar{R} \right\} = 0 \quad (5)$$

Note that this can be solved to obtain

$$\bar{R}^* = \max_k \left\{ \frac{p_k R_k}{q_k} \right\} = \frac{p_{k^*} R_{k^*}}{q_{k^*}} \quad (6)$$

LEMMA 3.1. *The values \bar{R}^* and k^* defined in (6) are also solutions to the original optimization problem stated in (4).*

PROOF. We prove by contradiction. Note that

$$\max_k \left\{ q_k \left\{ \frac{p_k R_k}{q_k} - \bar{R}^* \right\} \right\} \geq \left\{ q_{k^*} \left\{ \frac{p_{k^*} R_{k^*}}{q_{k^*}} - \bar{R}^* \right\} \right\} = 0 \quad (7)$$

If the first two expressions are equal then k^* and \bar{R}^* are also optimal for (4) and we are done. Suppose that we have strict inequality then this means

$$\max_k \left\{ q_k \left\{ \frac{p_k R_k}{q_k} - \bar{R}^* \right\} \right\} > 0 \quad (8)$$

If we denote the optimal solution values by k' then for (8) to be true we must have

$$\frac{p_{k'} R_{k'}}{q_{k'}} > \bar{R}^* = \max_k \left\{ \frac{p_k R_k}{q_k} \right\} \quad (9)$$

where the equality is a result of (6). This is a contradiction and hence strict inequality cannot hold. \square

4 AN ILLUSTRATIVE EXAMPLE

Let us illustrate the approach via a simple example. Note that the profit of a product will typically be related to its desirability. If this relationship is positive and profit increases with demand (or desirability) then the optimal strategy is to recommend the most desirable product since that strategy will optimize accuracy as well as profit and, in turn, long-term profit. However, in practice, the relationship between profit and desirability tends to be a negative one. Items which have low profitability (e.g. they are on sale) typically are highly desired (and hence if recommended the product will be purchased with high probability). Those products with high profitability are less desirable since consumers may decide to wait for a price reduction.

For illustration purposes, let us assume that the profit of a product decreases with increasing probability of purchase. This is an agreement with the traditional linear demand curve used in economics whereby price (revenue) decreases linearly with demand (acceptance probability) [6]. Hence

$$R = 1 - p - \theta \quad (10)$$

where $\theta > 0$, to account for the case of negative revenue and $\theta < 1$ since the company is not expected to support large losses. This includes the case whereby a company provides the product at a loss (e.g., to gain more customers) since we can have $R < 0$. This is sometimes termed a "loss leader" and such products will be in very high demand. Vendors may sell a small number of items at a loss in order to attract customers in the hope that the attracted customers would also purchase items with higher profitability [9]. The relationship in equation 10, which is based on the assumption that low profit items tend to be in higher demand, is our initial guess and more work is needed for a more precise model.

Next let us consider q which is the probability that a customer leaves because of inappropriate recommendations.

This probability will have a positive relationship with $1 - p$ so we assume a linear relationship such that

$$q = \varepsilon(1 - p) \quad (11)$$

where $\varepsilon < 1$ determines the degree by which the customer is affected by the recommendation. We are presently investigating other models for this relationship, such as proportional hazard models [11], on which we will report at a later date. We next compute the optimal values for various optimization objectives.

Maximum Long-Term Profit

Substituting for R_k and q_k in 6 we have

$$\bar{R} = \max_k \left\{ \frac{p_k(1 - p_k - \theta)}{\varepsilon(1 - p_k)} \right\} \quad (12)$$

For this illustrative example let us assume a continuous function of p_k (i.e., for any p_k , we can find a corresponding product) we can then obtain the maximum by taking derivatives. Let us define

$$F(p_k) = \left\{ \frac{p_k(1 - p_k - \theta)}{(1 - p_k)} \right\} \quad (13)$$

where F is the function giving the long term profit from selecting a product with a particular purchase probability. Note that

$$F''(p_k) = -2\theta(1 - p_k)^{-3} < 0 \quad (14)$$

and hence a single maximum exists.

$$F'(p_k) = 1 - \theta(1 - p_k)^{-2} \quad (15)$$

and setting to zero and solving for p_{k^*} we obtain

$$p_{k^*} = 1 - \sqrt{\theta} \quad (16)$$

where k^* is the product that achieves this optimal point. Finally we obtain

$$\bar{R}_{LT} = \frac{1 + \theta - 2\sqrt{\theta}}{\varepsilon} \quad (17)$$

where \bar{R}_{LT} is the expected long term profit obtained by explicitly maximizing the long term profit.

Maximizing Short-Term Profit

In the case of maximizing short term profit we instead have

$$F(p_k) = p_k(1 - p_k - \theta) \quad (18)$$

where F is the function giving the short term profit from selecting a product with a particular purchase probability. We can again show a single maximum and that this occurs when

$$p_{k^*} = \frac{1 - \theta}{2}. \quad (19)$$

We are interested in the long term profit (for this short term optimization problem) and this is given by

$$\bar{R}_{ST} = \frac{p_{k^*}(1 - p_{k^*} - \theta)}{\varepsilon(1 - p_{k^*})} = \frac{(1 - \theta)}{2\varepsilon} \quad (20)$$

Maximizing Accuracy

In this case we have

$$F(p_k) = p_k \quad (21)$$

where F is the accuracy of selecting a product with a particular purchase probability. The optimal solution is

$$p_{k^*} = 1 \quad (22)$$

and so the long term profit in this case is given by

$$\bar{R}_{AC} = \lim_{p_{k^*} \rightarrow 1} \frac{p_{k^*}(1 - p_{k^*} - \theta)}{\varepsilon(1 - p_{k^*})} = -\infty \quad (23)$$

Comparison

Clearly maximizing accuracy is not appropriate since the customer is retained forever but for each product sold the company loses money (i.e., the loss leader approach is not valid in this case). The ratio of the long term profit to the short-term profit can be written as

$$\rho \equiv \frac{\bar{R}_{LT}}{\bar{R}_{ST}} = \frac{2(1 - \sqrt{\theta})}{(1 + \sqrt{\theta})}. \quad (24)$$

In Figure 1 we plot this ratio as a function of the parameter θ for small values. Note that θ will typically be small since this represents the loss the company is willing to incur for the "loss leader" item. For such values optimizing for long term instead of short term profit results in an almost doubling of the expected profit. Note that this ratio is independent of the factor ε which determines how likely a customer will leave the vendor if recommended products are not suitable.




Figure 1: Ratio of Long-Term to Short-Term Expected Profits

5 NUMERICAL RESULTS

We used the MovieLens dataset that comprises 100,000 ratings (1-5) (bad-excellent) from 943 users on 1682 movies. From this dataset 563 users were used for training for a Naive Bayes model which was imported from *sklearn*. If a user assigned a rating of 4 or 5 to a movie, we took this to indicate that a user liked a movie. Using sex, age, and occupation as the user attributes, we then used Naive Bayes to learn the probability of a user liking a particular movie.

According to a CNN article [16], "The percentage of ticket sales that the studio takes decreases on each week that a movie is in the theater" and so we assume a similar profit model (for the vendor) with profit linearly increasing with age. More precisely, if t_m is the age of movie m in years, r_m , the profit of movie m , was taken to be $r_m = 2 + 0.75t_m + \eta$, where $\eta \sim \mathcal{N}(0, 1)$.

The results obtained are provided in Table 1 where the objectives (maximizing accuracy, short-term profit and long-term profit) are listed in the columns and the evaluation criteria are listed in the rows. As expected, the optimal metric value corresponds to the associated objective function. However the distinction is not as clear as expected. We believe the reason to be the fact that the collaborative filtering of the MovieLens dataset results in older rather than newer movies being highly recommended.

Table 1: Revenue and Accuracy Metrics

	$\max\{p_k\}$	$\max\{p_k R_k\}$	$\max\left\{\frac{p_k R_k}{q_k}\right\}$
Accuracy	0.127	3.655	63665.336
Short-Term	0.121	3.976	63667.110
Long-Term	0.121	3.975	63667.130

6 CONCLUSIONS AND FUTURE WORK

We developed a model for optimizing long-term revenue in a recommender system. We then formulated and solved the associated optimization problem and illustrated the approach with the MovieLens dataset. Although we did find a distinction, we believe that a more appropriate dataset that includes pricing information and is less skewed would provide more insightful results. We are presently in the process of identifying such a dataset and will provide results in a future paper. We also plan to investigate customer retention models which is the probability of customer loss as a function of the suitability of the product offered to them.

REFERENCES

- [1] Gediminas Adomavicius, Jesse C. Bockstedt, Shawn P. Curley, and Jingjing Zhang. 2018. Effects of Online Recommendations on Consumers' Willingness to Pay. *Info. Sys. Research* 29, 1 (March 2018). <https://doi.org/10.1287/isre.2017.0703>
- [2] Amos Azaria, Avinatan Hassidim, Sarit Kraus, Adi Eshkol, Ofer Weintraub, and Irit Netanel. 2013. Movie Recommender System for Profit Maximization. In *Proceedings of the 7th ACM Conference on Recommender Systems (RecSys '13)*. ACM, New York, NY, USA, 121–128. <https://doi.org/10.1145/2507157.2507162>
- [3] Long-Sheng Chen, Fei-Hao Hsu, Mu-Chen Chen, and Yuan-Chia Hsu. 2008. Developing recommender systems with the consideration of product profitability for sellers. *Information Sciences* 178, 4 (2008), 1032–1048.
- [4] Aparna Das, Claire Mathieu, and Daniel Ricketts. 2009. Maximizing profit using recommender systems. *CoRR* abs/0908.3633 (2009). [arXiv:0908.3633](http://arxiv.org/abs/0908.3633) <http://arxiv.org/abs/0908.3633>
- [5] Michael Fellmann, Novica Zarvič, and Oliver Thomas. 2018. Business processes modeling recommender systems: User expectations and empirical evidence. *Complex Systems Informatics and Modeling Quarterly* 14 (2018), 64–79.
- [6] M. Friedman. 2017. *Price Theory*. Taylor & Francis. <https://books.google.tt/books?id=SyEuDwAAQBAJ>
- [7] Michele Gorgoglione, Umberto Panniello, and Alexander Tuzhilin. 2019. Recommendation strategies in personalization applications. *Information & Management* 56, 6 (2019), 103143. <https://doi.org/10.1016/j.im.2019.01.005>
- [8] Mikael Hammar, Robin Karlsson, and Bengt J. Nilsson. 2013. Using Maximum Coverage to Optimize Recommendation Systems in e-Commerce. In *Proceedings of the 7th ACM Conference on Recommender Systems (RecSys '13)*. ACM, New York, NY, USA, 265–272. <https://doi.org/10.1145/2507157.2507169>
- [9] James D. Hess and Eitan Gerstner. 1987. Loss Leader Pricing and Rain Check Policy. *Marketing Science* 6, 4 (Nov. 1987), 358–374. <https://doi.org/10.1287/mksc.6.4.358>
- [10] Kartik Hosanagar, Ramayya Krishnan, and Liye Ma. 2008. Recommended for you: The impact of profit incentives on the relevance of online recommendations. *ICIS 2008 Proceedings* (2008), 31.
- [11] Tomoharu Iwata, Kazumi Saito, and Takeshi Yamada. 2006. Recommendation Method for Extending Subscription Periods. In *Proceedings of the 12th ACM SIGKDD International Conference on Knowledge Discovery and Data Mining (KDD '06)*. ACM, New York, NY, USA, 574–579. <https://doi.org/10.1145/1150402.1150470>
- [12] Tamas Jambor and Jun Wang. 2010. Optimizing multiple objectives in collaborative filtering. In *Proceedings of the fourth ACM conference on Recommender systems*. ACM, ACM, New York, NY, USA, 55–62.
- [13] Dietmar Jannach and Gediminas Adomavicius. 2017. Price and profit awareness in recommender systems. *arXiv preprint arXiv:1707.08029* (2017).
- [14] Georgia Koutrika. 2018. Modern Recommender Systems: From Computing Matrices to Thinking with Neurons. In *Proceedings of the 2018 International Conference on Management of Data (SIGMOD '18)*. ACM, New York, NY, USA, 1651–1654. <https://doi.org/10.1145/3183713.3197389>
- [15] Shanyun Liu, Yunquan Dong, Pingyi Fan, Rui She, and Shuo Wan. 2019. Matching Users' Preference Under Target Revenue Constraints in Optimal Data Recommendation Systems. *ArXiv* abs/1901.10721 (2019).
- [16] Annelena Lobb. 2002. How a movie theater makes a buck. (2002). https://money.cnn.com/2002/03/08/smbusiness/q_movies/
- [17] Cheng Long, Raymond Chi-Wing Wong, and Victor Junqiu Wei. 2018. Profit Maximization with Sufficient Customer Satisfaction. *ACM Trans. Knowl. Discov. Data* 12, 2, Article 19 (Jan. 2018), 34 pages. <https://doi.org/10.1145/3110216>
- [18] Umberto Panniello, Michele Gorgoglione, Shawndra Hill, and Kartik Hosanagar. 2014. Incorporating profit margins into recommender systems: A randomized field experiment of purchasing behavior and consumer trust. *University of Pennsylvania, ScholarlyCommons* (2014).
- [19] D. Paraschakis. 2017. Towards an ethical recommendation framework. In *2017 11th International Conference on Research Challenges in Information Science (RCIS)*. IEEE, US, 211–220. <https://doi.org/10.1109/RCIS.2017.7956539>
- [20] Sandeep K. Raghuvanshi and R. K. Pateriya. 2019. Recommendation Systems: Techniques, Challenges, Application, and Evaluation. In *Soft Computing for Problem Solving*, Jagdish Chand Bansal, Kedar Nath Das, Atulya Nagar, Kusum Deep, and Akshay Kumar Ojha (Eds.). Springer Singapore, Singapore, 151–164.
- [21] Jing Ren, Robert Kauffman, and Dave King. 2019. Two-Sided Value-Based Music Artist Recommendation in Streaming Music Services. In *Proceedings of the 52nd Hawaii International Conference on System Sciences*.
- [22] Jian Wang and Yi Zhang. 2011. Utilizing marginal net utility for recommendation in e-commerce. In *Proceedings of the 34th international ACM SIGIR conference on Research and development in Information Retrieval*. ACM, ACM, New York, NY, USA, 1003–1012.