

Automated Learning of Pricing and Bundling Strategies in Information Economies

Christopher H. Brooks and Edmund H. Durfee *

Artificial Intelligence Laboratory
University of Michigan
Ann Arbor, MI 48109
{chbrooks,durfee}@umich.edu
<http://www.eecs.umich.edu/~chbrooks>

The advent of automated commerce and the ability to electronically distribute information goods, or articles, have created a new set of problems for producers of information content. Since information goods have negligible marginal cost, producers are presented with a huge space of potential bundling and pricing strategies to choose from. Producers potentially have two separate decision problems to solve: what to offer and how to price it. For example, in deciding what to offer, producers may choose to offer articles from a wide variety of categories, or just offer in-depth content in a few selected categories. In deciding what to charge for their goods, a producer might choose from a number of schedules such as a flat fee, a per-article price, or a schedule which changes nonlinearly based on the number of articles purchased. The choice of schedule and goods offered will impact both the producer's per-iteration profit and also what it is able to learn about consumer preferences. In an online environment with a large number of goods and rapid transaction times, it will be advantageous for a producer to automate this decision-making. In (Brooks *et al.* 1999), we describe this problem and show how a monopolist producer can acquire higher aggregate profit by learning a simpler pricing schedule. Our current work considers two other factors that complicate the producer's learning and decision-making problems: the presence of other producers and differences in article content.

In a world in which a producer is competing with other producers for a market of consumers, the producer must consider both the strategy of other producers and the preferences of the consumers. Typically, neither of these are known, but instead must be learned over time. What makes this difficult is that the other producers are also trying to learn about consumer preferences, so producer strategies are not necessarily stationary. One question we are concerned with is the extent to which reasoning about opponent strategies is actually useful; there will be a tradeoff between an increase in the degree of modeling detail and the marginal gain in the value of the information acquired.

One factor that makes this problem different from a traditional game theory problem is that we are explicitly interested in the nonequilibrium rewards that a producer receives. That is, it's not enough for a producer to eventually learn an

optimal pricing schedule; the producer also wants to maximize its interim profits while learning this price schedule. This leads to a tradeoff between the quality of the solution that is learned and the time needed to perform this learning. The reason for this is that, in real-world e-commerce problems, the consumer population is typically changing, either due to entry and exit of consumers or shifting consumer tastes. Therefore, a producer may not have enough time to find an optimal strategy.

The decision as to how to price a set of articles is only half of the problem. A producer must also decide which articles to offer. Typically, consumers will have heterogeneous preferences over articles. The producer's decision problem is then to find a set of articles which can be priced so as to maximize its profits.

One advantage of information goods is that it is their content which is valued by consumers. This content can often be categorized and placed within a taxonomy. Under some assumptions, similarity within a taxonomy can be used as a predictor of consumer utility for an article. For example, if a producer sees that consumers highly value an article on Michael Jordan, it might conclude that consumers would also value articles on Magic Johnson, since both men are former NBA players. Of course, this is only a bias; it may be that consumers use a different taxonomy than the producer assumes. Determining the taxonomy used by the consumer then becomes a learning problem in itself.

Again, we would like for the producer to not only find a good set of articles to offer, but also to perform well during learning, since the consumer population is likely to be changing.

All of this work takes place under a decision-theoretic framework, in which a producer estimates the value of its actions with respect to a possibly unknown problem horizon and takes a series of actions which will maximize its profits. By viewing the choice of articles to sell and the selection of prices as a set of decisions which yield both information and profit and then estimating the value of each decision, we will gain a greater understanding of how to explore complex decision spaces.

References

Brooks, C. H.; Fay, S.; Das, R.; MacKie-Mason, J. K.; Kephart, J. O.; and Durfee, E. H. 1999. Automated strategy searches in an electronic goods market: Learning and complex price schedules. In *Proceedings of ACM EC-99*.

*This work was supported in part by an IBM University Partnership Grant and by the National Science Foundation under grant IIS-9872057
Copyright © 2000, American Association for Artificial Intelligence (www.aaai.org). All rights reserved.