2012 2nd International Conference on Management and Artificial Intelligence
IPEDR Vol.35 (2012) © (2012) IACSIT Press, Singapore

Data-driven vs. Theory-driven Approaches to E-Selling

Petri Parvinen+, Maurits Kaptein and Essi Pöyry

Aalto University School of Economics 00076 Aalto, Finland

Abstract. This research explores the ways in which computer learning is advancing the development of a new area of electronic business: e-selling. With the increase of computing power, online businesses are beginning to face similar challenges to those experienced by salespersons in face-to-face settings; there is a demand for higher degree of personalization, adaptation and persuasion. The ability to use e.g. Bayesian learning algorithms for real-time online persuasion has challenged the dominance of *ex ante* knowledge of customers. This paper is intended as a conceptual landmark and a discussion-opener for linking dynamic real-time data-based learning and selling by presenting two examples of studies testing data- vs. theory-driven approaches in e-selling. The results indicate that data-driven learning can be used to increase online sales yet the interpretative skills possessed by humans are still needed.

Keywords: E-selling, adaptation, Bayesian learning, algorithms

1. Introduction

Digital business exchange has developed in the form of e-marketing, e-(re)tailing and e-support/e-automation, but little in the form of e-selling. This article investigates how computer based learning can advance the ways in which customer interactions can be adapted in the same way salespeople in traditional offline settings have adapted to each customer at hand. Adaptation is seen as the key competence in contemporary sales work [21, 23].

Although many online vendors have personalized their product offerings by employing recommender systems, the ways in which these products are typically presented to customers are static. For example, Amazon.com's product recommendations are still based on an analysis of what similar customers have done in similar situations, and not on a constant analysis of how the current customer responds to different stimuli.

Amazon.com's practice represents the high-end of large-scale personalized e-commerce that is based on *ex ante* understanding of customers. There are two types of conventional approaches to designing artificial intelligence that attempt to tailor individual offerings. Firstly, there are theory-driven approaches, in which a model is developed from theory and the practice bases on the model. Customer categorizations, segmentations or behavioral stereotypes represent typical practical examples. As an example, theory states that hedonic and utilitarian consumers behave differently in their web-consumption behavior [5, 12]. Based on this, e-commerce systems can be designed to differentiate between the offerings that are primarily offered for hedonic or utilitarian consumers.

Data-driven artificial intelligence approaches, on the other hand, work with empirical data or with input-output behavior of real systems. Typically, this approach has been based on gathering data into data warehouses. The behavior or decision-making situation of an individual customer is then associated with earlier similar customer cases and the tailoring of approaches and offerings are done accordingly.

These practices are far from what face-to-face salespersons are able to do with customers. Most importantly, sales persons are able to adapt in real-time to customer reactions and the history of that

+ Corresponding author, Tel.: +358 50 312 0905
Email address: petri.parvinen@aalto.fi

particular customer.

This article reviews and discusses the different ways in which data- and theory driven learning and optimization take place within e-commerce. Particular attention is paid to how data-driven learning advances e-commerce towards e-selling. The primary research question of this paper is: *What are the key differences in data- and theory-driven approaches to e-commerce?* The secondary research question is: *How does data-driven learning help e-commerce applications resemble human-like e-selling?* We approach these questions by reviewing the relevant literature and two empirical studies conducted in e-commerce settings.

2. Theory-driven vs. Data-driven Approaches to Learning and Adaptation in E-Commerce

In social sciences, there are two alternative ways of using computational power in the form adaptive, dynamic models. Artificial intelligence in this form can be used for both theory-driven and data-driven model building [17]. Likewise, the optimization of variables and parameters in e-commerce applications can be based on a) existing theory about customer behavior, b) existing data about their behavior or c) their combination.

In e-commerce applications, existing data about customers' behavior (b) has traditionally been based on aggregate level data, i.e. data on the behavior and responses of a large group of individual customers. Recently, the increase in computational power and the development and application of dynamic learning have transformed the scene. E-commerce applications use e.g. sequential Bayesian learning [11, 22] to optimize offerings.

2.1. Theory-driven e-commerce adaptation

Traditional theory-building and optimization is based on the assumption that theoretical thinking is powerful in framing an understanding about customers' behavioral responses to given commercial offerings. Theories are built by testing hypotheses about the relationships between e.g. customer behavior, offerings and business results. The resulting theoretical understanding is then used to categorize customers. Based on the theory, metrics for identifying each type of customer behavior can be developed. These metrics, often in the form of behavioral scales, can thus serve both as theoretical tools and metrics for practical applications.

According to this logic, optimization takes place by designing offerings that are theoretically fit for each customer category. The success of the fit between the theory-based customer behaviors and the created offerings is measured, and, afterwards, the customer categories, metrics and offerings are fine-tuned for better results. The results can also fine-tune the theories. As new interdependencies emerge, either in empirical tests or as observed in real-life, new outlooks to existing theories about consumer behavior are created.

2.2. Data-driven e-commerce optimization

Conventionally, e-commerce applications start using data-driven approaches by gathering data about large amounts of customers, and using their behavioral data on an aggregate level to predict customer responses in a given situation. In e-commerce, this is analogous to the classical Amazon.com algorithm "other people, who bought this, also bought". By assimilating a given decision-making situation with the ones experienced by other customers earlier on, a probabilistic calculation guides the offering. This use of customer data for e-commerce optimization is based on a theoretical premise that customers behave categorically, and that assimilating a customers' situation to earlier 'batches' of customers delivers the best result.

Recently, however, the operation of algorithms has started to resemble data-driven theory building. Sequential Bayesian learning algorithms [11, 22] are able to optimize offerings on e-commerce platforms without existing theory. By reading the behavior of individuals and building the e-commerce offerings based on real-time adaptation to their individual-level responses, e-commerce applications are able to adapt dynamically to individual situations without any relationship to any understanding about how customers are theoretically supposed to behave, or how other customers have behaved on aggregate.

This represents a major departure in the philosophy of building e-commerce applications, optimizing their performance and ultimately also theorizing about customer behavior. The following examples highlight how this works in practice, why it resembles human behavior and how it represents a move towards e-selling.

3. Towards E-selling: Data-driven Learning

To illustrate how data-driven optimization can take place in e-commerce, we present two categorical examples of its application. The first example shows how simple parameters like price can be based on data-driven real-time learning at the level of an individual. The second example describes how different theory-based assumptions work in the data-driven learning practice.

3.1. Example 1 – Data-driven learning in online price setting

The first example regards setting prices to online services. Traditionally, prices are determined based on value/cost theories that provide models to optimize one or few fixed prices for a given product. In offline settings these approaches are usually supported and used, as fully flexible pricing is hardly cost-effective. However, when operating online, companies can offer different prices to different customers in (near) real-time [15]. Thus, sellers can run 'field experiments' on the prices of their services and sequentially find out the optimal price for their product.

The sequential learning of the profit maximizing sales price of a product can be treated as a variation of the multi-armed bandit problem [22]: a seller wants to learn her payoffs for each possible price (probability of a sale) so that she is able to determine the sales maximizing price. While the multi-armed bandit is notoriously resistant to analysis [8], optimal solutions are available in certain cases [8, 26]. The problem of finding the profit maximizing sales price to a product sold online compares to the multi-armed bandit problem in the sense that the seller wants to learn the respective probabilities of a sale (success) given different prices (arms). Contrary to the standard formulation of the multi-armed bandit problem, however, is the fact that the pay-offs for each successful sale given different prices are not equal but rather depend on the price offered.

Methods that have been put forward to solve the multi-armed bandit problem range from the use of the Gittins Index [8] to multiple heuristic strategies. One heuristic that corresponds well to the tradition of conducting market research to set optimal sales prices would be to use a batch learning approach: The seller randomly allocates the first N customers to each of the possible prices p of the product and observes customers' willingness to buy with that certain price. After this learning phase the seller has sufficient information to compute the optimal price. In the batch learning method, the size of N is critical but cumbersome to determine: large N will lead to a correct estimation of willingness to buy given a price, but will also lead to incurred losses by selling the product too cheap to a number of customers or by losing deals by asking for a too high price.

Another more recent heuristic to approach the multi-armed bandit problem is the use of randomized probability matching [22]. Here, each of the 'arms' of the bandit is treated as independent and their respective success probabilities are modeled using the Bayesian approach. Next, one performs a random draw of the posterior probability distribution of each arm, and selects the arm with the highest draw. Finally, the probability distribution describing the estimated success of that arm is updated based on the observed success or failure.

To compare the performance of the proposed weighted randomized probability matching algorithm with a batch estimation algorithm, a range of simulations were performed. Based on the results, we can see that an algorithm that corresponds to the current practice of A-B testing in online marketing outperforms the batch estimation. These algorithms rely on conducting an experiment to estimate the values of interest, and next exploiting the estimated value. This setup is problematic because the learning costs can be ever increasing if the wrong decision has been made. Also, the A-B testing algorithms implicitly assume that the first buyers of a product have similar price sensitivities to later buyers. Weighted randomized probability matching guards against the problems described above. Furthermore, the algorithm is computationally lightweight and can thus be implemented in real time.

3.2. Example 2 – Dynamic adaptation of sales influence tactics to individual online customers

Salespeople have a set of sales influence tactics (SITs) at their disposal in face-to-face settings, and they dynamically adapt their behavior to what they experience in individual customer interactions [18, 20]. As an example, in some situations customers might be made a special offer and in others the sales pitch can stress popularity of the product [3, 4]. According to the literature, there is a considerable heterogeneity in the effects of these strategies [2, 9, 16].

However, despite the advantages in using SITs, such adaptation in online e-commerce settings has been rare. By dynamically adapting SITs to individual customers rather than the average customer, e-commerce could move towards e-selling. In this example we show, via a field experiment using Bayesian learning algorithms, that the business performance of an e-commerce platform can be increased by dynamically adapting SITs to individual customer responses. It was expected that using SITs on an e-commerce platform would help achieve better sales results, individual customers would respond differently but consistently to SITs and that it would be more effective to adapt SITs on an individual level rather than statically using them on an aggregate-level.

The usage of dynamic SIT adaptation was tested on a Dutch affiliate site selling children's clothing. During a four-month period, half of the website visitors were randomly assigned to the version that implemented the dynamic updating of SITs. The other half were assigned to browse the original website (without any usage of SITs). Two SITs were used for the trial. Firstly, a 'Special offer' SIT combining the power of scarcity and threats [3, 18] was employed. Secondly, a 'Bestseller' SIT relying on the informational power of social proof [3, 20] was employed. Given that the goal of the affiliate store is to obtain high click-through rates to the vending website(s), the click-through rate was used as a measure of success: any click on the image or link of the product implementing an SIT was regarded as a success for that specific SIT. A failure to get the customer to click the product was regarded as an unsuccessful influence attempt.

Based on the experiment, the results give support to the previous expectations. By using SITs, the website got better click-through rates (yet the result was statistically insignificant), individuals responded differently to the tested SITs, and, instead of using the best SIT for all customers, better results were obtained by adapting the SITs to each customer.

The results are of primary importance because they show the ease with which adaptive selling can be implemented in e-commerce settings. The identification of different SITs and a relatively simple scheme of tracking their performance for individual users were sufficient to significantly improve the performance of an e-commerce website. The learning algorithms used for this application were lightweight and are easily implemented on existing infrastructures to provide real-time adaptation of SITs to individual customers.

The results are also significant for the further development of practices in e-selling. Many theories explaining the effectiveness of SITs currently build on their *average* effectiveness among groups of people. However, as the effectiveness of face-to-face selling depends on the ability to adapt to reactions of individuals, the use of SITs should focus on the consistency and dynamics of individual responses rather than on the responses of groups. Systems like the one used in the study can monitor individuals' responses to SITs over time and the use of different SITs can be iteratively developed.

Both of the described examples illustrate a finding according to which a real-time, iterative and data-driven approach to e-commerce is extremely powerful both in single business-decisions (setting prices) as well as in more general operating processes (adaptive sales). From a methodological point of view, the shift in the level of analysis implies that the key procedures required in e-commerce research involve dynamic modeling and algorithms. The ability to hypothesize, measure and ultimately predict the behavior of individuals is paramount.

4. Challenges in Extending E-selling

Even though operating in a data-driven mode seems desirable, there are challenges in its utilization. The challenges highlight especially the difficulties in the computation of data but also the need for human-produced theories.

Firstly, the more data a website produces, the slower the iterative process becomes. Furthermore, in situations with fewer variables, real-time dynamic learning algorithms outperform other ways of optimizing e-commerce offerings but when the complexity increases, the accuracy and the power of mathematical estimation decreases. Hence, a solid theory helps the analyzer to structure the variables in the data-driven mode, as it gives a starting point to the data-driven process and a way to categorize and simplify the variables. Therefore, it is important to understand the relative strengths and the weaknesses of humans' and computers' analytical power. Humans' analytical capacities should be used for structuring problems, creating contents and interpreting results. Computers' processing power should be harnessed for, preferably multivariate, optimization and tireless dynamic learning.

Secondly, from the perspective of recreating the interactions of face-to-face salespeople, two limitations to relying on data-driven approaches emerge. In e-commerce, the embedded nature of a transaction in the digital locus needs to be understood. The concept of the digital spatio-temporal locus captures the boundaries and outcomes of human-like activities in interactions in the digital media. Actions are governed by the digital interaction medium. Here, human immersion [1, 19] is elevated as a key theoretical tenet particular due to its ability to complement the vagueness of the concept of flow [13] with an understanding of the digitalization of the spatio-temporal locus of interaction and the psychology of persuasion. Human immersion is argued to both be a key premise for conceptualizing e-selling and moderating e-selling success, much due to the same reason it was originally found relevant in the acceptance of human-computer interaction (HCI) [7, 24, 25], and later e-commerce technology [6, 10, 14]. The impact of a feeling of presence, personality and body language fall under the same category. The space in which human interactions occur and the immediacy of the ability to react unconsciously are essential parts of professional face-to-face saleswork. For e-selling to advance, the particularities of the digital spatio-temporal locus need to be included in the way systems learn and adapt to situations. This is analogous to cultural contexts or meeting spaces, both of which are always considered when face-to-face sales people decide how to behave.

Furthermore, while theory-led approaches are not necessarily most useful for optimizing offerings, persuasive cues or prices, they can be invaluable in designing contents. A deeper theoretical understanding, or the process of achieving it, can produce insights that represent major departures from existing ways of working. In practice, this would entail systems and algorithms that are creative and experimental in their learning. The pharmaceutical industry is a good example of a context where data-driven innovation (based on databases of molecular interactions) produces numerous incremental advances but disruptive innovations emerge from theoretically fundamental research initiatives about e.g. depression, ageing or sex life. There, data miners are working hard to extend the power of data-led approaches to more radical innovations.

5. References

[1] J. Bjork and J. Holopainen. *Patterns in Game Design*. Charles River Media, 2005.

[2] J.T. Cacioppo, R.E. Petty, and C.F. Kao. The efficient assessment of need for cognition. *Journal of Personality Assessment* 1984, **48** (3): 306-307.

[3] R.B. Cialdini. Systematic opportunism: An approach to the study of tactical social influence. In J.P. Forgas and K.D. Williams (Eds.). *Social Influence: Direct and Indirect Processes*. Psychology Press, 2001.

[4] R.B. Cialdini and N.J. Goldstein. Social influence: Compliance and conformity. *Annual Review of Psychology* 2004, **55**: 591-621.

[5] J. Cotte, T.G. Chowdhury, S. Ratneshwar and L.M. Ricci. Pleasure or utility? Time planning style and web usage behaviors. Journal of Interactive Marketing 2006, 20 (1): 45–57.

[6] M.K. Chang, W. Cheung, and V.S. Lai. Literature derived reference models for the adoption of online shopping. *Information and Management* 2005 **42** (4): 543-559.

[7] J.A. Ghani and S.P. Deshpande. Task characteristics and the experience of optimal flow in human-computer interaction. *Journal of Psychology: Interdisciplinary and Applied* 1994, **128** (4): 381-391.

[8] J.C. Gittins. Bandit processes and dynamic allocation indices. *Journal of the Royal Statistical Society Series B Methodological* 1979, **41** (2): 148–177.

[9] R. Guadagno and R.B. Cialdini. Online persuasion and compliance: Social influence on the Internet and beyond. In Y. Amichai-Hamburger (Ed.). *The social net: The social psychology of the Internet*. Oxford University, 2005.

[10] Y. Ha, M. Yoon and M. Choi. Determinants of adoption of mobile games under mobile broadband wireless access environment. *Information and Management* 2007, **44** (3): 276–286.

[11] J.R. Hauser, G.L. Urban, G. Liberali and M. Braun. Website morphing. *Marketing Science* 2009, **28** (2): 202-223.

[12] J.B. Hartman, S. Shim, B. Barber and M. O'Brien. Adolescents' utilitarian and hedonic web-consumption behavior: Hierarchical influence of personal values and innovativeness. *Psychology and Marketing* 2006, **23** (10): 813–839.

[13] D.L. Hoffman and T.P. Novak. Flow online: Lessons learned and future prospects. *Journal of Interactive Marketing* 2009, **23** (1): 23-34.

[14] C-L. Hsu and H-P. Lu. Why do people play on-line games? An extended TAM with social influences and flow experience. *Information and Management* 2004, **41** (7): 853–68.

[15] P.K. Kannan and P.K. Kopalle. Dynamic pricing on the Internet: Importance and implications for consumer behavior. *International Journal of Electronic Commerce* 2001, **5** (3): 63–83.

[16] C.T.F. Klein and D.M. Webster. Individual differences in argument scrutiny as motivated by need for cognitive closure. *Basic and Applied Social Psychology* 2000, **22** (2): 119–129.

[17] K. Manhart. Artificial intelligence modelling: Data driven and theory driven approaches, 2007. In K. Troitzsch, U. Müller, G. Nigel and J.E. Doran (Eds.). *Social science micro simulation*. Springer, 1996.

[18] R.G. McFarland, G.N. Challagalla and T.A. Shervani. Influence tactics for effective adaptive selling. *Journal of Marketing* 2006, **70** (4): 103-117.

[19] J. Nechvatal. *Immersive ideals / Critical distances*. LAP Lambert Academic Publishing, 2009.

[20] J.M. Payan and R.G. McFarland. Decomposing influence strategies: Argument structure and dependence as determinants of the effectiveness of influence strategies in gaining channel member compliance. *Journal of Marketing* 2005, **69** (3): 66-79.

[21] C.R. Plouffe, J. Hulland and T. Wachner. Customer-directed selling behaviors and performance: A comparison of existing perspectives. *Journal of the Academy of Marketing Science* 2009, **37** (4): 422-439.

[22] S.L. Scott. A modern Bayesian look at the multi-armed bandit. *Applied Stochastic Models in Business and Industry* 2010, **26** (6): 639-658.

[23] R.L. Spiro and B.A. Weitz. Adaptive selling: Conceptualization, measurement, and nomological validity. *Journal of Marketing Research* 1990, **27** (1): 61-69.

[24] L.K. Trevino and J. Webster. Flow in computer-mediated communication. *Communication Research* 1992, **19** (5): 539-573.

[25] L.K. Webster and L.R. Trevino. The dimensionality and correlates of flow in human computer interactions. *Computer-Human Behavior* 1993, **9** (4): 411-426.

[26] P. Whittle. Multi-armed bandits and the Gittins index. *Journal of the Royal Statistical Society Series B Methodological* 1980, **42** (2): 143-149.