

Challenges and opportunities in high-dimensional choice data analyses

Prasad Naik · Michel Wedel · Lynd Bacon ·
Anand Bodapati · Eric Bradlow ·
Wagner Kamakura · Jeffrey Kreulen · Peter Lenk ·
David M. Madigan · Alan Montgomery

Published online: 31 May 2008
© Springer Science + Business Media, LLC 2008

Abstract Modern businesses routinely capture data on millions of observations across subjects, brand SKUs, time periods, predictor variables, and store locations, thereby generating massive high-dimensional datasets. For example, Netflix has

Prasad Naik and Michel Wedel are co-chairs.

P. Naik
University of California, Davis, Davis, CA, USA
e-mail: panaik@ucdavis.edu

M. Wedel (✉)
University of Maryland, College Park, MD, USA
e-mail: mwedel@rhsmith.umd.edu

L. Bacon
Polimetrix Inc., Palo Alto, CA, USA

A. Bodapati
University of California, Los Angeles, Los Angeles, CA, USA

E. Bradlow
University of Pennsylvania, Philadelphia, PA, USA

W. Kamakura
Duke University, Durham, NC, USA

J. Kreulen
IBM Almaden Research Center, San Jose, CA, USA

P. Lenk
University of Michigan, Ann Arbor, MI, USA

D. M. Madigan
Columbia University, New York, NY, USA

A. Montgomery
Carnegie Mellon University, Pittsburgh, PA, USA

choice data on billions of movies selected, user ratings, and geodemographic characteristics. Similar datasets emerge in retailing with potential use of RFIDs, online auctions (e.g., eBay), social networking sites (e.g., mySpace), product reviews (e.g., ePinion), customer relationship marketing, internet commerce, and mobile marketing. We envision massive databases as four-way VAST matrix arrays of Variables \times Alternatives \times Subjects \times Time where at least one dimension is very large. Predictive choice modeling of such massive databases poses novel computational and modeling issues, and the negligence of academic research to address them will result in a disconnect from the marketing practice and an impoverishment of marketing theory. To address these issues, we discuss and identify the challenges and opportunities for both practicing and academic marketers. Thus, we offer an impetus for advancing research in this nascent area and fostering collaboration across scientific disciplines to improve the practice of marketing in information-rich environment.

Keywords Challenges · Opportunities · High-dimensional choice data analyses · Modern businesses · Four-way VAST matrix arrays

1 Introduction

Advances in computing power have revolutionized marketing practice. Companies not only possess massive database on billions of consumer choices, user ratings, and geodemographics, but also recognize that the predictive modeling of massive choice data offers potentially high return on investment. As firms become customer-focused, they amass customer databases by recording various aspects resulting from interactions with customers. A fundamental consequence of building massive databases to marketing is the shift in focus from reaching anonymous consumers (e.g., in mass marketing via television) to targeting identifiable customers (e.g., in direct marketing via online stores or services). In mass marketing, researchers have traditionally relied on data that came from a representative random sample of the population so that inferences could be drawn from that small sample to the population at large. In direct marketing, on the other hand, firms want specific insights at the customer level to enable them to develop marketing plans for each of its customers. To this end, we need models and methods that are scalable up to the entire customer base (e.g., Balasubramanian et al. 1998).

The size of choice datasets has been increasing due to the scanning devices in supermarkets since the 1970s. In the late 1980s, with the advent of loyalty cards, marketing companies expanded the details of customer-level information, which enabled a better targeting of marketing activities. Since the late 1990s, due to Internet-based technologies, marketers in diverse sectors such as retailing, direct marketing, banking, and telecom began to compile large databases on consumers' past transactions. More recently, due to the Internet becoming a marketing channel, companies augmented their databases with clickstream data and web-based transaction data. The basic nature of all these data sets, however, is numerical.

Besides numerical data, a qualitatively different kind of data known as “unstructured data” is increasingly compiled and used. It includes textual information from customer complaints, product reviews by other customers or expert opinions (ePinion, Netflix), social networking sites (e.g., mySpace), blogs, or patents (Google Patents). In the twenty-first century, due to advances in software engineering, marketers can use this unstructured information to “listen” to what customers have to say about one’s own existing products (or competitors’ products) and to identify gaps in the product space, trends in innovation, or emerging technologies. In *Mining the Talk*, Spangler and Kreulen (2007) describe how to cull insights from textual data using illustrative business cases. In addition, new high-dimensional data sources continue to emerge due to RFID devices in stores and on products; eye-tracking data (e.g., Tobii) from natural settings such as PCs, TV, and billboards; gaming, video, and metadata from semantic Web 2.0.

To summarize, marketing data are expanding in several modes; first, the number of Variables to explain choice behavior has greatly increased (large V); second, the number of choice Alternatives in packaged goods exploded to hundreds per product category (large A); third, automated data collection led to the recording of choice decisions from large samples of Subjects (large S); finally, the frequency of data collection on weekly, daily, or even hourly basis led to long Time series (large T). These VAST data matrices offer opportunities for improving the marketing analysis, planning, and implementation. But, they raise novel computational and statistical challenges. These challenges include the need to incorporate substantive issues such as consumer heterogeneity, cross-category preferences, spatial covariances, and dynamic and forward-looking aspects of choice behavior. To manage this enhanced complexity, practitioners resort to simpler models that are scalable to real-life size of the databases. On the other hand, to provide greater rigor and better explanation, academic researchers tend to enrich the models by focusing on a subset of variables, alternatives, subjects, or times (e.g., by sampling from or aggregating across one or more of the VAST modes). But the enriched models have limited applicability because they lack scalability required by marketing practitioners. This dilemma becomes even more salient in (1) real-time applications that technology-intense marketing environment facilitates (e.g., Google or Amazon) and (2) decision-theoretic applications that require not only model estimation but also decision optimization. It points to the urgent need for faster numerical methods, efficient statistical estimators, and simpler models.

Some solutions to the challenges in massive data analyses are offered, singularly or in combination, by (1) increasing computational power, for example, through parallelization and grid computing; (2) developing scalable estimation methods via closed-form expressions for posterior distributions of parameters of marketing models or one-pass particle filters and other sequential MCMC methods and variational inference; and (3) creating alternative approaches for data analyses, for example, factor models, regularization methods, and inverse regression methods. The rest of the paper describes the sources of emerging data, then reviews the current computational solutions for analyzing massive data, and finally presents the recent advances in model estimation.

2 Sources of high-dimensional data

The main phenomena that drive the emergence of novel massive databases are Web 2.0, virtual social networks, predictive markets, recommendation systems, and unstructured data, which we discuss below.

2.1 Web 2.0

Recent considerations for web standards focus on creating the “Semantic Web,” an environment that supports reasoning and enables nontechnical users to create their own content. As a result, web sites are moving away from a publishing (i.e., pushing of content) paradigm to a participatory functionality where large numbers of consumers can post product evaluations and comments. To extract insights from such consumer-generated data, companies like Umbria and Wize attempt to solve the challenges of locating the data and getting it into analyzable form. Alternatively, when consumers participate in web-based activities, they can produce more usable data and metadata by structuring differently what they do and earning rewards appropriately for their efforts. An example is in the generation of “usefulness” ratings for product reviews on Amazon.com. Such user- and site-related interactions create large databases that contain potentially valuable behavioral information which as yet not much research in marketing has addressed.

2.2 Virtual social networks

Social networks are a centerpiece phenomenon in Web 2.0. Online users generate content and create online activity for other users. Hence, the audience base in social networks tends to be self-propelling or self-deflating. As for the social networking web site, it generates revenues via advertising by exposing advertisements to those who visit the site. From a managerial perspective, understanding who the users are and who keeps the social network site running is important for marketing strategies, in particular viral marketing. Who are the key players? What (or who) drives site visitation and, therefore, advertising revenues? How should the social network site segment its customers? How to advertise on social networking sites? To manage the site, to make it a better place for consumers, and to obtain better information for advertisers, the site’s management needs to know exactly what role is being played by each individual user within the network and whose actions have an impact on whom. To shed light on these questions, massive databases on user activity, connectivity, and referrals are being compiled. Managers need modeling approaches to cull insights and identify site users who are influential in affecting a firm’s primary revenue driver: other members’ site usage. Recent research in marketing has begun to address these issues (e.g., Trusov et al. 2007a, b).

2.3 Predictive markets

Predictive markets, sometimes called virtual stock markets, are gaining popularity (e.g., Iowa Electronic Markets or Hollywood Stock Exchange). Carnegie Mellon University offers an ESP game that mobilizes collective efforts to label images

available on the Web. Such applications are “games” in the sense that there are rules and desired outcomes. Other game-like methods have been used for ideation and product development. Prelec (2001) describes a guessing game played over a network that can be used to develop customer descriptions of products and product concepts. Toubia (2006) describes a system for generating ideas in which incentives are aligned with performance. Bacon and Sridhar (2006) present online “innovation tools” that can be used to identify preferences and solve hard problems, including a game-oriented discussion system that aligns rewards with incentives. Similarly, Ding et al. (2007) propose a bartering mechanism to collect data in conjoint settings. Foutz and Jank (2007) show how to use virtual stock market to make early forecasts of new product sales. The European Union plans to build an online environment for collaborative innovation (<http://www.laboranova.com>). Such purposeful and structured activities will produce rich data on preferences that can be modeled effectively if participant contributions are constrained appropriately, metadata are defined adequately, and scalable methods are developed rapidly.

2.4 Recommendation systems

Product recommendation system based on collaborative filtering is the backbone of major Internet-based companies. For example, Netflix and Blockbuster ask their customers to rate the movies they have viewed as a basis for generating recommendations; Amazon and Barnes & Noble implement the recommendation systems to suggest books to read; CDNow and iTunes use it to recommend music recordings; Findory and Tivo recommend news items and TV shows. A recent trend, however, is to shift this collaborative filtering from the companies who do not utilize all the product ratings they collect to generate the recommendations to stand-alone services. Such open source collaborative filtering systems enable consumers to update their own information about actual listening, reading, or viewing behaviors.

Following this industry practice, marketing literature has refined the methods for making recommendations; for example, Bayesian network models (Breese et al. 1998), mixture models (Bodapati 2008), hierarchical Bayes’ models (Ansari et al. 2000), and hierarchical Bayes selection models (Ying et al. 2006). These models show substantial improvements in the quality of recommendations on test datasets. However, these methods are not yet applicable to the entire customer database available with the companies because the estimation methods are computationally intensive. Besides scalability, recommendation systems need to incorporate VAST choice alternatives, large-scale missing data, scale usage heterogeneity, endogeneity of recommendations, and sensitivity to shilling attacks; these topics will be a fruitful area of future research (e.g., Chung et al. 2007).

2.5 Unstructured data

Using database management systems, companies build numerical databases and apply data mining techniques to extract customer insights and gain competitive advantage. However, extracting information from text is underutilized in corporate marketing and underresearched in academic marketing. For example, companies can learn how consumers view brands in a certain product category (i.e., brand

positioning) by analyzing textual conversations obtained from chat rooms, blogs, or social networking sites.

To analyze textual information, a first step is to create structure from this inherently unstructured textual data by classifying words into meaningful categories or taxonomies, i.e., natural hierarchy of information consistent with the business goals and processes (see, e.g., Kreulen et al. 2003; Kreulen and Spangler 2005). Then, the resulting datasets, which are much larger than traditional (numerical) databases, can be analyzed by adapting statistical techniques for classification, prediction, and visualization to extract business intelligence and knowledge management (e.g., Kreulen et al. 2002). IBM has explored the application of these technologies in several domains including customer relationship management, enterprise content analysis, and World Wide Web analytics. For further details, see Spangler and Kreulen (2007).

3 Computational issues

High-dimensional data leads to computational challenges because many algorithms do not scale linearly, they scale exponentially. A few select examples of computationally intensive research include the use of store transaction data to set optimal pricing strategies (Montgomery 1997), clickstream data for adaptive web design (Montgomery et al. 2004), or consumer choices for shopbot design (Brynjolfson et al. 2007). Although MCMC-based procedures facilitated the estimation of intractable models, these techniques are time-consuming: the estimation of a hierarchical choice model could take hours to converge (Allenby et al. 1996). In practice, companies such as Amazon or Google typically require its system to respond in 2 s or less. Hence, fast computational methods are necessary. We next discuss two such methods: grid computing and specialized electronic devices.

Grid computing environments harness the power of thousands of low-cost personal computers by splitting up the overall computational task into many smaller tasks that can be executed in parallel. Even techniques like MCMC that are inherently sequential can be parallelized. Specifically, Brockwell and Kadane (2005) propose a method of parallelizing these algorithms that relies on regeneration, while Brockwell (2006) suggest a prefetching scheme for large-scale problems.

Application-specific integrated circuit (ASIC) and field-programmable gate array (FPGA) are specialized electronic devices that can increase computational speed (Brown and Rose 1996). Traditional microprocessors execute a set of instructions to perform a computation. By changing the software instructions, the functionality of the system is altered without changing the hardware. The downside of this flexibility is that the performance can suffer. The processor must read each instruction from memory, decode its meaning, and only then execute it. This results in a high execution overhead for each individual operation. The second method of algorithm execution is to use hardwired technology like ASICs, which are designed specifically to perform a given computation and so they are fast and efficient when executing the exact computation for which they were designed. However, the circuit cannot be altered after fabrication. FPGAs lie between these two extremes, being

cheaper and more flexible than ASICs, while also being much faster than software-programmable microprocessors. FPGAs contain an array of computational elements whose functionality is determined through multiple programmable configuration bits. For integer or fixed-point arithmetic, FPGAs can improve speed by two orders of magnitude compared to a conventional cluster CPU. Future research needs to apply these computational methods to not only perform statistical tasks such as estimation, inference, and prediction, but also for decision-theoretic optimization.

4 Models and estimation methods

Although massive databases pose computational challenges, is there a conceptual difference between small versus large datasets? The answer is affirmative. In a multivariate probability density function, as the number of variables increases, any local neighborhood is almost surely empty (i.e., sparseness of data cloud), whereas a nonempty neighborhood is almost surely not local (Simonoff 1996, p. 101). To understand this so-called empty space phenomenon, consider the standard normal density whose 68% of its total mass lies within ± 1 standard deviation from the origin. In a bivariate normal, the mass within a unit square centered at the origin is about $(0.68)^2$, which is less than one half of its total mass. As the number of variables p increases, the mass within the unit hypercube, which equals $(0.68)^p$, tends to zero rapidly. With just ten variables (i.e., $p=10$), the center of the density function where the mode with the largest frequency is located contains merely 2% of the sample.

Consequently, as Silverman (1986, p. 92) cautions, “large regions of high [dimensional] density may be completely devoid of observations in a sample of moderate size.” In contrast to small datasets, most observations in large databases are outliers rather than representative. Thus, the empty space phenomenon induces different conceptual issues when analyzing massive data; nonparametric models cannot be estimated accurately for ten or more regressors and dimension-reduction methods are necessary to project the data into a low-dimensional subspace.

In many application areas, the move from small datasets to large datasets requires researchers to rethink, in a qualitatively different manner, how one approaches the problem of data analysis. In several cases, the data analysis is initially tackled using combinatorial approaches because that is how the problem is first conceptualized; with larger datasets, the combinatorial approach often becomes computationally impracticable and one is forced to adopt model-based approaches where the researcher needs to postulate, in the abstract, the stochastic processes that could generate the dataset. As an example, for many decades, cluster analysis was approached as a combinatorial NP-hard problem, but researchers now approach the problem from the perspective of statistical mixture distributions (Wedel and Kamakura 2000). While the problem continues to be hard, the computational burden increases much more slowly with the size of the dataset. A similar story has played out with social network analysis: the early approaches in Wasserman and Faust (1994) are combinatorial in nature, whereas Handcock et al. (2007) offer model-based approaches that are computationally feasible for gigantic networks. Below we

describe methods for analyzing high-dimensional data using factor analysis, inverse regression, Bayesian methods, and regularization methods.

4.1 Factor analytic approach

Consider a retailer with data on sales and pricing for each of 30 brands across 300 customers for 36 months who wants to understand the (cross) effects of price on brand sales, while allowing the price elasticities to vary across stores and over time. An aggregate model would require the estimation of 900 price coefficients, which is feasible given the available data. But a random-coefficients model that accounts for unobserved heterogeneity would additionally require the estimation of 405,450 covariance terms. To transform this random-coefficients model into a feasible one while allowing for the possibility that the (cross)elasticities might be correlated across stores and over time, Kamakura and Kang (2007) propose a factor regression model where the covariance of the random regression coefficients is decomposed into its principal components across stores and over time. With a two-factor solution, the model would require 1,800 estimates for the covariance of the price coefficients, still large but more manageable than a full-covariance model.

A more traditional approach for reducing a high-dimensional problem to a more manageable one is by factor-analyzing the dependent variables (Wedel and Kamakura 2001). Du and Kamakura (2007), for example, propose a stochastic frontier factor model that takes advantage of the covariance structure among the dependent variables to define efficiency frontiers for each of multiple product categories based on the observed sales in all categories, as well as exogenous variables. But even if the “number of variables problem” is solved with some form of dimension-reduction model, it is still critical that the calibrated model can be implemented in the entire database. Simulation-based methods requiring many iterations per case are not as scalable as the simpler data mining models based on singular-value decomposition. To this end, a pragmatic solution that has been employed is to calibrate the model on a random sample from the database and then implement the estimated model on all the other cases. This approach, however, does not efficiently use all of the available data.

4.2 Inverse regression methods

A broad class of methods for dimension reduction (see Naik et al. 2000) is based on the work by Li (1991) on sliced inverse regression. Recent advances allow, for example, the extraction of confirmatory factors via constrained inverse regression (Naik and Tsai 2005) or the estimation of models even if the number of regressors exceeds the number of observations via partial inverse regression (Li et al. 2007). Next, we describe an application of inverse regression for nonparametric estimation of binary choice models for data with many variables.

For parametric (choice) models, the sensitivity to the specification of the link function may be limited for small- to moderate-sized data, but it is likely much larger for massive data where if the prespecified link departs from the true unknown link, the model cannot accurately predict rare events with low probability of occurrence. Nonparametric binary models overcome this drawback but break down when

datasets contain ten or more regressors. Hence, semiparametric models have been proposed that relate an index (or a factor) a linear combination of regressors to the expected response via a nonparametric link function (Naik and Tsai 2004). The multiindex binary response (MBR) model extend this idea to include more than one index and can be estimated using a noniterative two-step method, thus allowing for flexibility of the link function and scalability to the large datasets (Naik et al. 2008). The two-step estimation method involves projection and calibration. In the projection step, information available in high-dimensional predictor space is combined and projected into a low-dimensional index space. Dimension reduction is achieved by estimating an index structure (i.e., the composition of indexes in terms of original regressors and the number of indexes to retain) without specifying the link function. In the calibration step, the unknown link function is estimated via local polynomial regression (or a parametric model). One of the main advantages of the MBR model is that, by allowing for multiple indexes, it facilitates a more refined understanding of consumer behavior.

In an application to predict customer defection for a telecom company, Naik et al. (2008) formulated the multifactor binary choice model where they reduced dimensionality from 124 predictors to four factors using a new method for high-dimensional data analysis known as sliced average variance estimation (Cook and Weisberg 1991). Having reduced dimensionality markedly, in the resulting low-dimensional subspace, they deploy local polynomial regression, which would be practically impossible in the original predictor space to visualize the regions where customers least or most likely to defect reside and to discover that attractive regions of prospective customers need not be contiguous or connected.

4.3 Bayesian methods

As previously mentioned, VAST datasets pose computational problems in Bayesian inference due to its computationally intensive procedures. Specifically, MCMC algorithms recursively generate the model's parameters from the full conditional distributions given the data. Standard implementation of MCMC requires holding all the data in memory or repeatedly reading from the data file in every iteration of the MCMC chain. Because these approaches are not feasible with massive datasets, DuMouchel (1999) uses empirical Bayes (EB) to analyze very large frequency tables. EB can be viewed as a large sample approximation to hierarchical Bayes models where the parameters for the population-level distributions are estimated by non-Bayes methods, such as maximum likelihood, often after marginalizing over individual-level parameters. Ridgeway and Madigan (2002) proposed a fully Bayesian method that first performs traditional MCMC on a subset of the data, then applies importance sampling and resampling to the remainder of the data. Balakrishnan and Madigan (2006) modified this method by applying particle filters (Liu and Chen 1998) to obtain a method that requires only a single pass through the data.

One approach to reduce computational costs relies on the property of conditional independence of observations. One may divide the data into manageable blocks and recursively analyze each block by using the posterior distribution from the previous block as the “updated prior” for the next one. The recursion continues until all the

data has been processed. Whereas this idea can be directly implemented if the all prior and posterior distributions are conjugate, for many models, this is not the case and approximations need to be used. This idea needs further work, but holds promise for estimating mixtures of Dirichlet processes (Escobar and West 1996) to describe heterogeneity distribution, which is of interest in marketing (Wedel and Zhang 2004). Huang and Gelman (2006) have also pursued related ideas.

Another approach to reduce computational cost is to derive analytical approximations. Recently, Bradlow et al. (2002), Everson and Bradlow (2002), and Miller et al. (2006) offer closed-form Bayesian inference to models previously thought to be nonconjugate problems. Bradlow et al. (2002) demonstrate this approach for commonly used models in marketing such as the negative binomial, beta-binomial, and heterogeneous logit model. The general approach is to (1) approximate the likelihood function by a polynomial with a user-specified number of terms in the summation (which could be a very large number, even thousands to get the desired accuracy), (2) find a conjugate prior to the approximated likelihood function, and (3) integrate term-by-term the “posterior polynomial” since the approximated likelihood and prior are now conjugate.

Finally, as an alternative to MCMC methods, variational approximations have recently attracted considerable attention, for example, see Wainwright and Jordan (2003) and the references therein. Like MCMC, variational inference methods have their roots in statistical physics; unlike MCMC, they are deterministic. The basic idea of variational inference is to formulate the computation of marginal or conditional probability as an optimization problem, perturb the optimization problems (i.e., induce “variations”) and find solutions to the (simpler) perturbed problems (see Jordan et al. 1998). The application of variational inference to VAST marketing databases is as yet unexplored.

4.4 Regularization methods

Regularization methods offer a promising route to tackle the problem of many variables. For example, Tibshirani (1996) developed the least absolute shrinkage and selection operator (LASSO) regression, which estimates the coefficients and selects the variables simultaneously. It minimizes the usual sum of squared errors in the linear regression model subject to constraining the sum of the absolute values of the coefficients. Fan and Li (2001) developed the smoothly clipped absolute deviation (SCAD) regression, which correctly eliminates with one probability all the variables whose true effects are negligible (i.e., the so-called oracle property). The Bayesian variant of the LASSO regression places a Laplace prior on the coefficients of the regression model (Genkin et al. 2007; Balakrishnan and Madigan 2007).

As an example of “many” variables, consider the monitoring of the safety of licensed drugs (Hauben et al. 2005). To detect novel drug–adverse event associations, patients, physicians, and drug companies submit reports containing drug(s), adverse event(s), and basic demographic information. Because most reports contain multiple drugs and multiple adverse events, complex multidrug interactions can exist. About 15,000 drugs are available in the US market; coincidentally, the dictionary of adverse events employed by the FDA also contains about 15,000 unique adverse events. Thus, the challenge reduces to performing a multivariate

regression with a 15,000-dimensional response variable and 15,000 input variables. By including just two-way interactions, one obtains over 100,000,000 input variables, which is well beyond the capabilities of current algorithms (for more details, see Hauben et al. 2005). Regularization methods, dimension-reduction methods (i.e., factor analysis, inverse regression) and their Bayesian variants hold the promise to provide effective solutions for VAST data.

5 Conclusion

In the mid-1960s, diagonalizing a 7×7 matrix represented a significant computational challenge. Progress during the past decades has led some, for example, Benzécri (2005), to claim that in data analysis “there is no longer any problem of computation.” But two interrelated events have changed the landscape: ubiquitous computing and ubiquitous databases. Massive databases will alter the computational scale of things to come, requiring the overall task to be divided into thousands (even millions) of smaller tasks and to be distributed to legions of computers, which solve the problem posed by each task and combine the resulting nuggets of information to arrive at the overall answer in a fraction of a second. Google’s MapReduce technology applies this grid computing idea when providing an output to a user’s query by searching massive databases on the Internet (Baker 2007). (Yahoo! and IBM also apply similar computing technology.) But failure to address the problems posed by VAST data in the near future may potentially lead to an impoverished marketing theory and can disconnect us from the marketing practice. Also, because a single solution to all computational challenges does not exist, a combination of computational, statistical, and algorithmic approaches is needed to solve specific problems. The VAST dimensions of massive data will require different approaches as well: solutions to handling massive samples (vaSt) will differ from those dealing with high-frequency data (vasT), which will differ from those addressing massive numbers of variables (Vast). To address these challenges expeditiously, marketing scholars may benefit from collaborations with different disciplines such as statistics and computer science. They can utilize Google’s grid computing technology when made available for scientific purposes (e.g., at Carnegie Mellon University or University of Maryland) to analyze massive data using the models and techniques discussed above.

We hope this article sets an agenda and stimulates progress in marketing for modeling VAST choice databases.

References

- Allenby, G. M., McCulloch, R., & Rossi, P. E. (1996). The value of purchase history data in target marketing. *Marketing Science*, 15, 321–340.
- Ansari, A., Essegai, S., & Kohli, R. (2000). Internet recommendation systems. *Journal of Marketing Research*, 37, 363–375.
- Bacon, L., & Sridhar, A. (2006). *Interactive innovation tools and methods*. Annual Convention of the Marketing Research Association, Washington DC, June.
- Baker, S. (2007). Google and the wisdom of clouds. *Business Week*, December 13th issue.

- Balakrishnan, S., & Madigan, D. (2006). A one-pass sequential Monte Carlo method for Bayesian analysis of massive datasets. *Bayesian Analysis*, *1*(2), 345–362.
- Balakrishnan, S., & Madigan, D. (2007). LAPS: LASSO with partition search. Manuscript.
- Balasubramanian, S., Gupta, S., Kamakura, W. A., & Wedel, M. (1998). Modeling large datasets in marketing. *Statistica Neerlandica*, *52*(3), 303–324.
- Benzécri, J.-P. (2005). Foreword. In F. Murtaugh (Ed.), *Correspondence analysis and data coding with JAVA and R*. London, UK: Chapman and Hall.
- Bodapati, A. (2008). Recommendation systems with purchase data. *Journal of Marketing Research*, *45*, 77–93.
- Bradlow, E. T., Hardie, B. G. S., & Fader, P. S. (2002). Bayesian inference for the negative binomial distribution via polynomial expansions. *Journal of Computational and Graphical Statistics*, *11*(1), 189–201.
- Breese, J., Heckerman, D., & Kadie, C. (1998). *Empirical analysis of predictive algorithms for collaborative filtering*. Madison, WI: Morgan Kaufmann.
- Brockwell, A. E. (2006). Parallel Markov chain Monte Carlo simulation by pre-fetching. *Journal of Computational and Graphical Statistics*, *15*(1), 246–261.
- Brockwell, A. E., & Kadane, J. B. (2005). Identification of regeneration times in MCMC simulation, with application to adaptive schemes. *Journal of Computational and Graphical Statistics*, *14*(2), 436–458.
- Brown, S., & Rose, J. (1996). Architecture of FPGAs and CPLDs: A tutorial. *IEEE Design and Test of Computers*, *13*(2), 42–57.
- Brynjolfson, E., Smith, M., & Montgomery, A. (2007). *The great equalizer: An empirical study of choice in shopbots*. Working Paper, Carnegie Mellon University, Tepper School of Business.
- Chung, T., Siong, R. R., & Wedel, M. (2007). My mobile music: Automatic adaptive play-list personalization. *Marketing Science*, in press.
- Cook, R. D., & Weisberg, S. (1991). Discussion of Li (1991). *Journal of the American Statistical Association*, *86*, 328–332.
- Ding, M., Park, Y.-H., & Bradlow, E. (2007). *Barter markets*. Working Paper, The Wharton School.
- Du, R., & Kamakura, W. A. (2007). *How efficient is your category management? A stochastic-frontier factor model for internal benchmarking*. Working Paper.
- DuMouchel, W. (1999). Bayesian data mining in large frequency tables, with an application to the FDA spontaneous reporting system. *The American Statistician*, *53*(3), 177–190.
- Escobar, M. D., & West, M. (1996). Bayesian density estimation and inference using mixtures. *Journal of the American Statistical Association*, *90*, 577–588.
- Everson, P. J., & Bradlow, E. T. (2002). Bayesian inference for the beta-binomial distribution via polynomial expansions. *Journal of Computational and Graphical Statistics*, *11*(1), 202–207.
- Fan, J., & Li, R. (2001). Variable selection via nonconcave penalized likelihood and its oracle properties. *Journal of the American Statistical Association*, *96*, 1348–1360.
- Foutz, N. Z., & Jank, W. (2007). *Forecasting new product revenues via online virtual stock market*. MSI Report.
- Genkin, A., Lewis, D. D., & Madigan, D. (2007). Large-scale Bayesian logistic regression for text categorization. *Technometrics*, *49*, 291–304.
- Handcock, M. S., Raftery, A. E., & Tantrum, J. M. (2007). Model-based clustering for social networks. *Journal of the Royal Statistical Society. Series A*, *170*(2), 301–352.
- Hauben, M., Madigan, D., Gerrits, C., & Meyboom, R. (2005). The role of data mining in pharmacovigilance. *Expert Opinion in Drug Safety*, *4*(5), 929–948.
- Huang, Z., & Gelman, A. (2006). *Sampling for Bayesian computation with large datasets*. Retrieved from <http://www.stat.columbia.edu/~gelman/research/unpublished/comp7.pdf>.
- Jordan, M. I., Ghahramani, Z., Jaakkola, T. S., & Saul, L. K. (1998). An introduction to variational methods for graphical models. In M. I. Jordan (Ed.), *Learning graphical models, vol 89 of series D: Behavioural and social sciences* (pp. 105–162). Dordrecht, The Netherlands: Kluwer.
- Kamakura, W. A., & Kang, W. (2007). Chain-wide and store-level analysis for cross-category management. *Journal of Retailing*, *83*(2), 159–170.
- Kreulen, J., Cody, W., Spangler, W., & Krishna, V. (2002). The integration of business intelligence and knowledge management. *IBM Systems Journal*, *41*(4), 2002.
- Kreulen, J., & Spangler, W. (2005). *Interactive methods for taxonomy editing and validation. Next generation of data-mining applications, chapter 20* pp. 495–522. New York: Wiley.
- Kreulen, J., Spangler, W., & Lessler, J. (2003). Generating and browsing multiple taxonomies over a document collection. *Journal of Management Information Systems*, *19*(4), 191–212.

- Li, K.-C. (1991). Sliced inverse regression for dimension reduction. *Journal of the American Statistical Association*, 86, 316–342.
- Li, L., Cook, R. D., & Tsai, C.-L. (2007). Partial inverse regression. *Biometrika*, 94, 615–625.
- Liu, J. S., & Chen, R. R. (1998). Sequential Monte Carlo methods for dynamical systems. *Journal of the American Statistical Association*, 93, 1032–1044.
- Miller, S. J., Bradlow, E. T., & Dayartna, K. (2006). Closed-form Bayesian inferences for the logit model via polynomial expansions. *Quantitative Marketing and Economics*, 4(2), 173–206.
- Montgomery, A. L. (1997). Creating micro-marketing pricing strategies using supermarket scanner data. *Marketing Science*, 16(4), 315–337.
- Montgomery, A. L., Li, S., Srinivasan, K., & Liechty, J. (2004). Modeling online browsing and path analysis using clickstream data. *Marketing Science*, 23(4), 579–595.
- Naik, P. A., Hagerty, M., & Tsai, C.-L. (2000). A new dimension reduction approach for data-rich marketing environments: Sliced inverse regression. *Journal of Marketing Research*, 37(1), 88–101.
- Naik, P. A., & Tsai, C.-L. (2004). Isotonic single-index model for high-dimensional database marketing. *Computational Statistics and Data Analysis*, 47(4), 775–790.
- Naik, P. A., & Tsai, C.-L. (2005). Constrained inverse regression for incorporating prior information. *Journal of the American Statistical Association*, 100(469), 204–211.
- Naik, P. A., Wedel, M., & Kamakura, W. (2008). Multi-index binary response model for analysis of large datasets. *Journal of Business and Economic Statistics*, in press.
- Prelec, D. (2001). *Readings packet on the information pump*. Boston, MA: MIT Sloan School of Management.
- Ridgeway, G., & Madigan, D. (2002). A sequential Monte Carlo method for Bayesian analysis of massive datasets. *Journal of Knowledge Discovery and Data Mining*, 7, 301–319.
- Silverman, B. W. (1986). *Density estimation*. London, UK: Chapman and Hall.
- Simonoff, J. S. (1996). *Smoothing methods in statistics*. New York, NY: Springer.
- Spangler, S., & Kreulen, J. (2007). *Mining the talk: Unlocking the business value in unstructured information*. Indianapolis, IN: IBM.
- Tibshirani, R. (1996). Regression shrinkage and selection via the lasso. *Journal of the Royal Statistical Society. Series B*, 58(1), 267–288.
- Toubia, O. (2006). Idea generation, creativity, and incentives. *Marketing Science*, 25(5), 411–425.
- Trusov, M., Bodapati, A., & Bucklin, R. E. (2007a). *Determining influential users in internet social networks*. Working Paper, Robert H. Smith School of Business, University of Maryland.
- Trusov, M., Bucklin, R. E., & Pauwels, K. (2007b). *Estimating the dynamic effects of online word-of-mouth on member growth of a social network site*. Working Paper, Robert H. Smith School of Business, University of Maryland.
- Wainwright, M., & Jordan, M. (2003). *Graphical models, exponential families, and variational inference*. Technical Report 649, Department of Statistics, UC Berkeley.
- Wasserman, S., & Faust, K. (1994). *Social network analysis*. Cambridge: Cambridge University Press.
- Wedel, M., & Kamakura, W. (2000). *Market segmentation: Conceptual and methodological foundations* (2nd edn.). Dordrecht: Kluwer.
- Wedel, M., & Kamakura, W. A. (2001). Factor analysis with mixed observed and latent variables in the exponential family. *Psychometrika*, 66(4), 515–530.
- Wedel, M., & Zhang, J. (2004). Analyzing brand competition across subcategories. *Journal of Marketing Research*, 41(4), 448–456.
- Ying, Y., Feinberg, F., & Wedel, M. (2006). Improving online product recommendations by including nonrated items. *Journal of Marketing Research*, 43, 355–365.