Marketing Letters 9:1 (1998): 79–91

Using a Community of Knowledge to Build Intelligent Agents

ANDREW D. GERSHOFF
Doctoral Student, University of Texas at Austin, CBA 7.202, Austin, TX 78712; e-mail: a.gershoff@mail.utexas.edu.

PATRICIA M. WEST
Assistant Professor of Marketing, University of Texas at Austin, CBA 7.202, Austin, TX 78712; e-mail: pwest@mail.utexas.edu.

Abstract

The modeling of individual consumer preference can be aided by incorporating others' opinions which contain information above and beyond identified product attributes. The value of others' opinions is tested using two empirical data sets. The results indicate that incorporating others' opinions into an attribute-based model can reduce systematic error and increase predictive accuracy by serving as a proxy for missing information (e.g., undiscovered attributes or attribute interactions, sensory or experiential aspects of the product, as well as advertising or word of mouth effects). Additionally, modeling individual preference based on others' opinions alone is shown to predict as well or better than traditional multiattribute models thus bypassing the need for defining a product attribute space.

Key words: Consumer preferences, prediction

Recent advances in technology are reshaping the way in which businesses service their customers. Marketing communication is increasingly becoming "interactive," engaging customers in a dialogue that enables companies to provide tailored services, manage customer relationships, and generate increased loyalty (Blattberg and Deighton 1991; Blattberg and Glazer 1994; Deighton, Peppers, and Rogers 1994; McKenna 1995; Pine, Peppers, and Rogers 1995). The benefits of establishing a learning relationship with customers are evident in the success of new businesses such as Peapod, an on-line grocery delivery service that refines its knowledge of individual customers with each transaction. By tracking customer purchase histories, Peapod is able to create customized shopping lists based on household buying patterns.

Additionally, a company such as Peapod can apply the wealth of "community knowledge" gained from comparing individual customer buying patterns with others who share similar tastes to provide useful recommendations of other products in which a customer might be interested (Peppers and Rogers 1995). The popularity of "Firefly," a web site (Firefly Network, Inc. 1997) offering movie and music recommendations based on individual profiles, is a testament to the perceived value of the service to its 150,000 members (Judge 1996). Firefly builds individual profiles based on a consumer's direct experience in a product category (i.e., having individuals rate various movies they have seen or recording artists they listen to) rather than standard demographic or psychographic information. This artificial intelligence system bases its recommendations on the likes and dislikes of other members identified as "nearest neighbors" (as defined by their similarity in tastes rather than demographic, geodemographic, or psychographic characteristics).

Intelligent agent tools such as Firefly depend on accuracy of preference prediction to provide a number of benefits to the consumer. In product categories such as movies, books, restaurants, homes, and home furnishings, which all offer a virtually limitless set of available alternatives, accurate intelligent agents can reduce the search time and effort required to find the "right" product. Second, when consumers are unaware of the determinants of their preferences (West, Brown, and Hoch 1996; Loewenstein and Schkade 1997) intelligent agents can deduce a preference structure based on an individual's overall evaluation of previously experienced products and recommend alternatives that the consumer would otherwise not identify. Third, just as consumers have been shown to respond positively to customized products (Pine 1993), tailored recommendations may increase consumer satisfaction with the choice process (Huffman and Kahn 1996).

Marketers also stand to gain from the prediction capabilities of intelligent agents. The sharing of information required to build an effective customer profile may result in switching costs. Thus to the extent that an individual is satisfied with the recommendations offered by an intelligent agent, he or she will loyally patronize the firm (Pine et al. 1995). Beyond the information that individual customers share about themselves, the firm maintains a wealth of information about others' preferences. This community knowledge is relevant for, but unavailable to, the consumer. When properly utilized, this information may allow the firm to provide a value-added service to its customers. Finally, accurate individual level preference prediction can increase the efficiency of marketing communications by reducing the costs associated with misdirected marketing efforts.

1. The value of others' opinions

The recommendations offered by today's intelligent agents have the capability to go beyond models traditionally used by marketing researchers that require a clearly specified attribute space for the product category (Griffin and Hauser 1993; Wilkie and Pessemier 1973). As marketers, we can learn from artificial intelligence systems that incorporate community knowledge. In particular, information available from others' opinions can and should be used to refine the predictive ability of individual level models above and beyond what can be achieved by incorporating product attribute information alone. Others' opinions can serve as a proxy for missing information such as undiscovered attributes or important attribute interactions, sensory or experiential aspects not conveyed by attribute information alone, and advertising or word of mouth influence.

The purpose of this paper is to present a simple methodology for optimally combining the information available from a database of others to build individual level preference prediction models. This method steps beyond the traditional multiattribute approach (Wilkie and Pessemier 1973) or conjoint analysis (Green and Srinivasan 1990) by incorporating others' opinions as an additional source of information. We present results from two experiments that demonstrate that the predictive accuracy of individual level attribute-based models can be improved by including a uniquely specified and optimally weighted linear combination of others' opinions. In addition, we illustrate how others' opinions act as a proxy for missing information and can be used to help identify undiscovered attributes. Finally, we discuss the method's applications to marketing problems, as well as its limitations and directions for future research.

2. Using others' opinions to predict individual preference

In order to present a more formal discussion the following notation which is extensively used in the judgment bootstrapping literature is introduced (see Blattberg and Hoch 1990; Camerer 1981; Einhorn 1974). We begin with a vector of preference ratings from a target individual (Y_t), and a matrix of i others' preference ratings for a set of alternatives (Y_{oi}). All ratings are a probabilistic function of a vector of known product attributes (X). Therefore, a multiattribute model of each individual can be built using multiple regression,

$$Y_t = X\beta_t + \varepsilon_t \text{ and } Y_{oi} = X\beta_{oi} + \varepsilon_{oi} \tag{1}$$

Let $\hat{Y}_t = X\beta_t$ be a vector of the target's model predictions and $\hat{Y}_{oi} = X\beta_{oi}$ be a matrix of others' model predictions. The resulting model predictions ($\hat{Y}_t$ and $\hat{Y}_{oi}$) represent the upper bound on the linear information extractable from the identified product attributes. Let $Z_t = \varepsilon_t$ and $Z_{oi} = \varepsilon_{oi}$ represent the residual portion of each individual's judgments, not explained by a linear combination of identified attribute information. Einhorn (1974) pointed out that the variance in the residual portion of judgment can include both systematic and random components and can be expressed in three parts:

$$\sigma_{z_t}{}^2 = \underbrace{\sigma_{\text{attributes}}{}^2 + \sigma_{\text{model}}{}^2}_{\text{systematic}} \underbrace{+\sigma_{\text{error}}{}^2}_{\text{Random}} \tag{2}$$

The first term on the right refers to the error associated with unknown or omitted product attributes, nonlinearities, and configural attribute relationships, including non-product attributes that systematically influence consumer preference (e.g., sensory or experiential aspects, advertising, or word of mouth). The second term refers to error associated with misspecification of the integration function. The final error term represents random error due to judgmental inconsistency, measurement, and sampling error.

We propose that others' opinions can act as a proxy for missing information and effectively be used to capture systematic error from the first two terms in equation (2) resulting in improved model predictions. The size of the random error component can be estimated by having the target re-rate a subsample of alternatives. The correlation between the two ratings, Y_{t1} and Y_{t2}, represents a measure of judgmental consistency, while $1 - r(Y_{t1}, Y_{t2})$ represents judgmental inconsistency or random error in judgment.[1]

Assuming that a target and one or more other individuals share a preference or distaste for an undiscovered attribute (or for a combination of existing attributes) then the preference model of the target and the preference models of these other n individuals (see equation 1) will have correlated errors ($r(Z_i, Z_{oi}) \neq 0$). The residuals from these other individuals can be used to explain the systematic error in the target's preferences in the following manner:

$$Z_t = \sum_{i=1}^{n} \alpha_i Z_{oi} + \varepsilon \tag{3}$$

Let $\hat{Z}_t = \Sigma a_i Z_{oi}$ be a vector of model predictions that represent the upper bound on the information extractable from others' opinions, controlling for all linear information associated with identified product attributes. To the extent that others' residuals are correlated with the target's residuals we can combine the attribute model predictions ($\hat{Y}_t$) and an optimally weighted linear combination of others' opinions ($\hat{Z}_t$) to produce improved predictions. Consider the following combined model that incorporates both identified product attribute information and additional information extracted from others' residuals:

$$\hat{Y}_t' = b_1\hat{Y}_t + b_2\hat{Z}_t. \tag{4}$$

Whenever the target and at least one other individual's judgments include a shared preference or distaste for an unidentified attribute, nonlinearities or configural attribute relationships ($r[Z_t, Z_o] \neq 0$), the combined model will generate more accurate predictions than a model built on identified attribute information alone. In essence, we are proposing that some portion of an individual's preference judgment not directly explainable as a linear function of known product attributes can be modeled as a unique combination of optimally weighted others' residuals which serve as a proxy for missing information.

In summary, we propose that others' opinions can be used to supplement product attribute information and improve individual preference predictions. Additionally, the proposed method of residual analysis can help in identifying individuals with similar (or opposing) tastes, and provide insight for identifying undiscovered attributes.

3. Empirical testing

3.1. Stimuli

Two stimulus sets were used to test the value of using others' opinions to improve the predictive accuracy of a multiattribute model. Aesthetic product categories were purposefully selected because they are known to have numerous attributes, configural and nonlinear relationships between attributes (Holbrook and Moore 1981) as well as sensory and experiential aspects not conveyed by tangible attributes (West and Broniarczyk 1998). Others' opinions are expected to be particularly valuable in improving predictive accuracy in these heterogeneous preference categories because non-product attributes, including experiential and sensory aspects, as well as symbolic and hedonic benefits are difficult to uncover and reliably evaluate.

3.1.1. House images and floor plans. A set of 63 realistic drawings of exterior images of homes was compiled along with an accompanying set of 63 floor plans. Both the house images and floor plans were reproduced and scaled to cover the same amount of space on an individual page by digitally scanning the images commonly found in magazines specializing in home building plans. The set of house images and floor plans was judged by two professional realtors to be representative of the stylistic elements and configurations typically considered by first time home buyers.[2]

A set of seventeen house attributes and a set of twenty floor plan attributes were constructed by examining descriptors found in magazine advertisements, a focus group with seven individuals, and consultation with two realtors. Two judges identified the key attributes of both the house images and floor plans; differences in attribute coding were resolved by discussion.

3.1.2. Quilt patterns. A collection of one hundred quilt images developed in previous research were used for the task (West 1996; West, Brown and Hoch 1996). The set of quilt patterns was chosen with the assistance of a quilting expert to represent a sample of patterns which would be available to a typical quilt buyer. A set of thirteen quilt attributes was constructed based on consultation with the same quilting expert.

3.2. Study one

3.2.1. Subjects and procedure. Thirty-five undergraduate students rated the house images and floor plans. These subjects were recruited from an introductory marketing class and given extra credit in their course in exchange for participation.

Subjects were asked to provide a preference rating for each of 63 house images and 63 floor plans reproduced on individual pages in separate booklets. Each booklet consisted of a different random ordering of house images or floor plans. Half of the subjects were asked to rate the 63 house images first and then the 63 floor plans, while the other half rated the floor plans first and then the house images. An eleven-point rating scale (anchored at one end with 'dislike like very much' and the other end with 'like very much') was located at the bottom of each page. Subjects were instructed to circle their evaluation of each of the house images and each of the floor plans on the scale provided. After completing both the house and floor plan booklets each subject re-rated a random subset of 10 house images and 10 floor plans and answered a brief demographics questionnaire.[3]

3.2.2. Analyses and results. Individual preference models were estimated for the house images and the floor plans. Two-stage hierarchical regression analyses were performed to test for the incremental improvement in explained variance resulting from incorporating other's opinions. For each subject, we first regressed their preference ratings on the appropriate set of attributes using ordinary least squares regression, and then regressed the residuals from the attribute-only model on all other subjects' residuals. In the second-stage of the procedure (i.e., when incorporating others' opinions) a stepwise regression procedure was used to limit overfitting.[4] Both the mean adjusted R^2 values and average mean squared error (MSE) between the subject ratings and the predicted values of the model are reported in Table 1.

Adding any new variables to the attribute-only model will increase the overall fit, however this may not hold true when testing out-of-sample predictive accuracy (i.e., forecasting). In order to determine whether the predictive accuracy of the models improved by using others' opinions as predictors a successive hold-one-out cross-validation procedure was used (Efron and Gong 1983).

3.2.3. Variance explained. Adding others' opinions to the model significantly increased the amount of explained variance for house images and the floor plans. For the house images, the mean adjusted R^2 value associated with the attributes alone was .34, which increased to a mean of .89 with the inclusion of others' opinions ($.89 - .34 = .55$; $t[34] = 20.65$, $p < .001$). A similar decrease in the MSE was indicated by the inclusion of others' opinions ($.40 - 2.28 = -1.89$; $t[34] = -12.08$, $p < .001$). For the floor plans, the mean adjusted R^2 value went from .22 to .74, ($.74 - .22 = .52$; $t[34] = 20.85$, $p < .001$), and the MSE decreased from a mean of 2.71 to .99 with the inclusion of others' opinions ($.99 - 2.71 = -1.72$; $t[34] = 9.548$, $p < .001$).

3.2.4. Predictive accuracy. Because of the potential problem of overfitting the data, an out-of-sample estimate of predictive accuracy was measured. Specifically, for each subject, the same two-stage hierarchical analysis was performed 63 times, successively holding out one of the house image or floor plan ratings. The remaining 62 ratings were used to predict the held out observation using the product attributes alone, then using attributes plus others' opinions. This method produced a complete set of 63 predicted scores for each subject allowing for an unbiased estimate of the accuracy of models outside of the model building sample. These results are also presented in Table 1.

For the house images the average correlation between the actual ratings and predicted values associated with the attributes alone was .34, which increased to .86 with the addition of others' residuals ($.86 - .34 = .52$; $t[34] = 20.33$, $p < .001$). Similarly, for the floor plans the average correlation between the actual ratings and predicted values started at .23 for models containing attributes alone and increased to .67 with the addition of other's opinions ($.67 - .23 = .44$; $t[34] = 16.01$, $p < .001$).

The MSE results from the cross-validation procedure also show improvement in predictive accuracy with the inclusion of others' residuals in the model. For the house plans, the MSE decreased from a 4.48 average for an attribute only model to a .97 average with the inclusion of others' residuals ($.97 - 4.48 = -3.51$; $t[34] = 14.85$, $p < .001$). Likewise,

Table 1. Summary of Results

Within Sample:		Study 1				Study 2	
		House Images		Floor Plans		Quilts	
Model	Predictors	Adjusted R^2	Mean Squared Error	Adjusted R^2	Mean Squared Error	Adjusted R^2	Mean Squared Error
Stage 1:	Attributes Only	0.336	2.283	0.218	2.708	0.318	1.299
Stage 2:	Attributes and Others' Residuals	0.886***	0.398***	0.734***	0.987***	0.945***	0.101***
"Attribute-less":	Others' Only	0.441**	1.901**	0.299**	2.430*	0.660***	0.628***
Out-of-Sample:		**Study 1**				**Study 2**	
		House Images		Floor Plans		Quilts	
Model	Predictors	r (actual, predicted)	Mean Squared Error	r (actual, predicted)	Mean Squared Error	r (actual, predicted)	Mean Squared Error
Stage 1:	Attributes Only	0.336	4.478	0.229	4.889	0.372	2.106
Stage 2:	Attributes and Others' Residuals	0.856***	0.971***	0.669***	1.965***	0.861***	0.576***
"Attribute-less":	Others' Only	0.434*	3.010***	0.195	3.470***	0.413*	2.104

***$p < .001$
**$p < .01$
*$p < .05$

for the floor plans the average MSE decreased from 4.89 for models containing attributes alone to 1.97 with the inclusion of others' residuals ($1.97 - 4.97 = -2.92$; $t[34] = 7.604$, $p < .001$).

3.2.5. Recovery of withheld information. We have argued that residual analyses of others' opinions improves the predictive accuracy of an attribute-based model because the residuals act as a proxy for missing information associated with undiscovered attributes, or configural attribute relationships, and sensory or experiential aspects of the product not conveyed by attributes. Our results indicate that inclusion of others' opinions decreased the systematic variance and thus improved the predictive accuracy of the individual models. If inclusion of others' residuals are, in fact, reducing systematic error associated with missing information, then we should be able to observe this effect by examining what happens when a known attribute is purposefully withheld from the model.

In order to test whether we could successfully recover missing information, we purposefully withheld an important floor plan attribute "separated master bedroom." This attribute was found to be significant in the models of fourteen of the thirty-five subjects. After regressing each of these subjects' ratings on the reduced set of attributes, the resulting residuals were compared to the value of the withheld attribute. Not surprisingly, each of the fourteen subjects for whom that attribute was known to be important had residuals that correlated significantly with the missing attribute, the average absolute value of the correlation being .42. Hence, each of these fourteen subjects' residuals included systematic variance associated with the withheld attribute. After others' residuals were included in the model, the absolute value of the correlation between subjects' residuals and the withheld attribute decreased by 71.4% to .12 ($.12 - .42 = -.30$; $t[13] = 7.13$ $p < .001$). This lends support to the argument that others' opinions can act as a proxy for missing attribute information.

3.2.6. Using others' opinions in place of product attributes. In order to test the boundary conditions of this approach, we were interested in determining the degree to which others' opinions alone could explain individual ratings (i.e., how well a model with only others' opinions and no attribute information can predict individual preference). Blattberg and Hoch (1990) found that an equally weighted combined model containing the forecast of an expert's opinion derived from a mechanical model (similar to our attribute-only models) and that same individual's subjective judgment outperformed the model alone and the expert alone. Similarly, we expected to find that a combined model containing both attribute information and the subjective evaluation of others' will outperform a model based purely on others' opinions. The combined model offers an optimal weighting of the identified product attributes and then attempts to explain the residual systematic variance using others' opinions. A model based solely on others' opinions is expected to be less precise due to differences in identified product attributes.

However, it is unclear how well a model built using others' opinions only would compare to a model built using attribute information alone. Obviously, the answer to this question depends on whether a well-defined attribute space has been identified, the degree of attribute configurality and nonlinearity in individual judgments, as well as the nature of the integration rule used. It also depends on both the number of other individuals whose judgments are available for modeling and the diversity in tastes between these individuals and the target of interest.

To examine the extent to which preference for aesthetic objects can be modeled using others' opinions alone, each subject's preference ratings were regressed on all others' ratings of the same alternative, the results are reported in Table 1. The average adjusted R^2 value for the house images was .44 which is a significant improvement over the average R^2 from the models using the attributes alone as explanatory variables ($.44-.34 = .10$; $t[34] = 3.047, p < .005$). For the floor plans the average adjusted R^2 was .30 which is also significantly better than the average R^2 of the attribute only models ($.30-.22 = .08$; $t[34] = 2.75$, $p < .01$). Similarly, the average MSE decreased using the others-only models compared to "attribute-only" models (house images: others-only = 1.90, attribute-only =2.28, $t[34] = 2.89, p < .01$; floor plans: others-only = 2.43, attribute-only = 2.71, $t[34] = 2.03$, $p = .05$).

The predictive accuracy results, based on a similar cross-validation procedure, indicate that using others' opinions instead of attribute information produce models that predict as well as or better than traditional attribute based models. For the house images, the average correlation between the target ratings and predicted values based on others' opinions alone was .43 which is significantly greater than the average correlation between the target ratings and predicted values based on attribute information alone, .34 ($.43 - .34 = .09$; $t[34] = 2.605$, $p < .05$). For the floor plans, there was not a significant difference in the average correlations between the target ratings and the predicted values of the models based on others' opinions only, .20, and the models based on attribute information alone, .23 ($.20 - .23 = -.03$, $t[34] = -.574$, $p = .57$). Significantly reduced MSE results produced in the cross-validation procedure indicated superior predictive accuracy using others' opinions alone compared to attributes alone for both the house images and the floor plans (house images: others-only=3.01, attribute-only = 4.48, $t[34] = 5.94$, $p < .001$; floor plans: others-only = 3.47, attribute-only = 4.89, $t[34] = 4.88$, $p < .001$).

These results indicate that using only others' opinions to predict individual preference may result in levels of predictive accuracy superior to those of multiattribute models, with the added benefit of not having to specify an attribute space. Considering the small sample size of available others' judgments (only 34 other opinions), the opinion-based model performed surprisingly well relative to the attribute-based model. With a larger sample of others' opinions we expect that the opinion-based model will perform even better than we have observed here. We recognize that improvement in predictive accuracy over an attribute-only model will depend, in part, on the quality of the definition of the attribute space. However, the painstaking efforts made to develop the attribute space together with assurances from experts that the set of attributes was representative of the level and quality available to, and used by, marketers suggests that our results are not likely to be due to insufficiency in modeling the subjects' responses to product attributes.

3.3. Study two

The purpose of Study 2 was twofold. First, we wanted to replicate the results found in Study One in another product category. Second, we wanted to obtain a larger sample size of others' opinions in order to see how much additional improvement in predictive accuracy can be gained by expanding the set of "potential predictors."

3.3.1. Subjects and procedure. Ninety-two graduate and undergraduate students rated the quilt patterns. These subjects participated in groups of two to eight, and were recruited either by posting signs on campus, or from undergraduate introductory marketing courses. Subjects were paid an average of \$8 for their participation.

Subjects were asked to provide their preference ratings for each of 100 quilt patterns as part of another task (West 1996). The quilt images were presented on a color computer screen, where the subject both viewed and rated each of the alternatives. The 100 quilts were presented in a different random order for each of the subjects, and no time limit was imposed. The task began with a practice session in which the subjects were trained how to input their ratings into the computer. A 7-point scale ranging from "dislike very much" to "like very much" appeared next to the quilt image on the screen. Subjects controlled the scale by manipulating the up and down arrow keys on a computer keyboard.

3.3.2. Analyses and results. Data were analyzed in the same manner as in Study 1 (see Table 1). The average adjusted R^2 value associated with the attribute-only model was .32, which increased to .95 by including others' residuals ($.95 - .32 = .63$; $t[91] = 41.20$, $p < .001$). Similarly, the average MSE between subject ratings and the predicted values went from 1.30 based on attribute information alone to .10 when others' residuals were added ($.10 - 1.30 = -1.20$; $t[91] = 17.63$, $p < .001$).

The cross-validation results indicate that the average correlation between the subject ratings and the predicted values of the combined model, .86, substantially exceeded the attribute only model, .37, ($.86 - .37 = .49$; $t[91] = 30.11$, $p < .001$). Likewise, the MSE of the attribute-only model, 2.11, was reduced to .56 with the inclusion of others' residuals ($.56 - 2.11 = -1.55$; $t[91] = -14.83$, $p < .001$).

Once again the performance of an opinion-only model was examined relative to an attribute-only model. The average adjusted R^2 for the opinion-only model was .66 compared to .32 for the attribute-only model ($.66 - .32 = .34$; $t[91] = 16.53$, $p < .001$). The average MSE associated with the opinion-only model was .63, which was significantly less than the attribute-only model, 1.30 ($.63 - 1.30 = -.67$; $t[91] = -8.23$ $p < .001$). The cross-validation results were similar to those observed in Study 1. The average correlation between subject ratings and predicted values for the opinion-only model was .41, which was significantly better than the attribute-only models ($.41 - .37 = .04$; $t[91] = 2.45$, $p < .05$). However, the average MSE for the opinion-only model was 2.10, which does not represent a significant difference from the attribute-only model ($2.10 - 2.10 = .00$; $t\ [91] = .014$, $p = .989$). Thus the others-only models performed better on one measure of predictive accuracy and equally as well on another. Similar to the results found in Study 1, others' opinions can be at least as accurate as multiattribute models but can do so without the requirement of a well defined attribute space.

4. Discussion and implications

Our results indicate that incorporating the residuals of others' opinions as a secondary source of information in addition to product attributes can significantly improve the accuracy of individual-level preference prediction. The residuals from others' preference models serve as a proxy for missing information, such as undiscovered attributes or important attribute interactions, sensory and experiential aspects of the product, and symbolic or hedonic benefits not conveyed by attribute information alone, as well as advertising or word of mouth influence which are difficult to uncover and reliably evaluate. Additionally, we found that others' opinions alone can serve as an effective proxy for attribute information providing similar or superior predictive accuracy without the identification of an attribute space.

The degree of improvement in predictive accuracy that results from incorporating the residuals of others' opinions in models with product attributes depends on a number of factors: First is the degree of preference heterogeneity across consumers. Heterogeneity in preferences among consumers can come about either through variance in the weight placed on a given set of attributes, or through variance in the set of attributes considered (Feick and Higie 1992). The more diversity in tastes due to variance in attribute importance across consumers, the greater the need for having a large sample of others' opinions to draw from when building preference models. In general, when there is substantial systematic variance not accounted for by known product attributes, others' opinions have greater potential for providing valuable information. Aesthetic and/or experiential product categories (movies, books, music, homes, and home furnishings) are more likely than functional products (computers, appliances, and cameras) to inspire perceptual processing that involves complex and configural attribute relationships (see Holbrook and Moore 1981; West and Broniarczyk 1997) as well as heterogeneity in preference (Feick and Higie 1992).

Second, the degree to which a consumer has well-defined and consistent preferences within the product category will impact predictive accuracy as well as the value of others' opinions. For consumers with well-defined preferences, others' opinions have greater potential for improving predictive accuracy because these individuals are providing consistent information. These same individuals also serve as better information sources (i.e., their residuals are likely contain more systematic variance than consumers with ill-defined preferences) for building community knowledge.

Our findings suggest a promising avenue for companies to pursue. Existing technologies can be used to build intelligent agents that can assist customers in finding the optimal alternative to satisfy their personal tastes. Given the vast amounts of information about individual customers that many companies are gathering today, the proposed method of using others' opinions to increase predictive accuracy over and above the traditional attribute-based preference models allows for efficient use of existing information. Predicting preference in this method allows marketers to tailor their recommendations and potentially their product offerings to the individual consumer. In practice, these models can be refined by updating the models as new information becomes available. A more precise preference model of a target consumer can be built as that target provides new information about his or her preferences (e.g., additional product evaluations or feedback regarding model recommendations).

Finally, with these new capabilities marketers will also face new challenges of how to best utilize and enhance the effectiveness of these intelligent agents that rely on community knowledge. For example, increases in precision of predicting the target's tastes will come about as other consumers provide additional information, thus expanding the community knowledge base. Product categories where consumers exhibit a high degree of preference heterogeneity require the marketer to either actively seek out a diverse set of consumers to build a community knowledge base, or to cater to a select segment of the population with shared tastes (e.g. cult video stores).

Acknowledgments

The authors wish to thank Mark Alpert, Susan Broniarczyk, Bart Bronnenberg, Leigh McAlister, and Elizabeth Thompson for their helpful comments.

Notes

1. This estimate of random error assumes that the target's preferences are stable (i.e., do not change as a result of exposure to additional alternatives). If the target's preferences are not well-defined to begin with but become crystallized in the process of rating alternatives than this estimate of random error may be upwardly biased.
2. Differences in individual taste were expected to vary more for the house images than the floor plans due to the more prominent aesthetic elements.
3. Subjects were asked to re-rate a subset of the alternatives in order to estimate their "judgmental consistency" (Camerer 1981).
4. The models were also estimated using stepwise regression in stage one. All results were consistent with the ordinary least squares results reported.

References

Blattberg, Robert C. and Stephen J. Hoch. (1990). "Database Models and Managerial Intuition: 50% Model +50% Manager," *Management Science* 36 (8), 887–900.

Blattberg, Robert C. and John Deighton. (1991). "Interactive Marketing: Exploiting the Age of Addressability," *Sloan Management Review* Fall, 5–14.

Blattberg, Robert C. and Rashi Glazer. (1994). "Marketing in the Information Revolution." In Robert C. Blattberg, Rashi Glazer, and John D. C. Little (eds.), *The Marketing Information Revolution*. Boston: Harvard Business School Press.

Camerer, Colin. (1981). "General Conditions for the Success of Bootstrapping Models," *Organizational Behavior and Human Performance* 27, 411–422.

Deighton, John, Don Peppers, and Martha Rogers. (1994). "Consumer Transaction Data Bases: Present Status and Prospects." In Robert C. Blattberg, Rashi Glazer, and John D. C. Little (eds.), *The Marketing Information Revolution*. Boston: Harvard Business School Press.

Efron, Bradley and Gail Gong. (1983). "A Leisurely Look at the Bootstrap, the Jackknife, and Cross-Validation," *The American Statistician*, February 37 (1), 36–48.

Einhorn, Hillel J. (1974). "Cue Definition and Residual Judgment," *Organizational Behavior and Human Performance* 12, 30–49.

Feick, Lawrence and Robin Higie. (1992)., "The Effects of Preference Heterogeneity and Source Characteristics on Ad Processing and Judgments about Endorsers," *Journal of Advertising* 21 (June), 9–24.

Firefly Network, Inc. (1997). "Firefly Network Inc.," ⟨http://www.firefly.net⟩ (Accessed 25 February 1997).

Green, Paul E. and V. Srinivasan. (1990). "Conjoint Analysis in Marketing: New Developments With Implications for Research and Practice," *Journal of Marketing* 54(4), 3–19.

Griffin, Abbie and John R. Hauser. (1993). "The Voice of the Customer," *Marketing Science* 12 (1), 1–27.

Holbrook, Morris B. and William L. Moore. (1981). "Feature Interactions in Consumer Judgments of Verbal Versus Pictorial Presentations," *Journal of Consumer Research* 8 (June), 103–113.

Huffman, Cynthia, and Barbara Kahn. (1996). "Variety for Sale: Mass Customization or Mass Confusion," working paper, The Wharton School, University of Pennsylvania.

Judge, Paul C. (1996). "Why Firefly Has Mad Ave. Buzzing," *Business Week* (October 7), 100–101.

Loewenstein, George and David Schkade. (1997). "Wouldn't it be Nice? Predicting Future Feelings." In E. Diener, N. Schwartz, and D. Kahneman (eds.), *Hedonic Psychology: Scientific Perspectives on Enjoyment and Suffering*. Sage Foundation.

McKenna, Regis. (1995). "Real Time Marketing," *Harvard Business Review* (July–August), 87–95.

Peppers, Don and Martha Rogers. (1995). "The End of Mass Marketing," *American Demographics (Marketing Tools Supplement)* (March–April), 42–52.

Pine, B. Joseph. (1993). *Mass Customization: The New Frontier in Business Competition*. Boston: Harvard Business School Press.

Pine, B. Joseph, Don Peppers, and Martha Rogers. (1995). "Do You Want to Keep Your Customers Forever?," *Harvard Business Review* (March–April), 103–114.

West, Patricia M. (1996). "Predicting Preferences: An Examination of Agent Learning," *Journal of Consumer Research* 23 (June), 68–80.

West, Patricia M., Christina L. Brown and Stephen J. Hoch. (1996). "Consumption Vocabulary and Preference Formation," *Journal of Consumer Research* 23 (September), 120–135.

West, Patricia M., and Susan M. Broniarczyk. (1998). "Integrating Multiple Critic Opinions: The Role of Aspiration Level on Consumer Response to Critic Consensus," *Journal of Consumer Research,* forthcoming.

Wilkie, William L., and Edgar A. Pessemier. (1973). "Issues In Marketing's Use of Multi-Attribute Models," *Journal of Marketing Research* 10 (November), 428–441.