Направление «Менеджмент»

Профиль «Маркетинг» **Код - 050**

Время выполнения задания – 180 мин.

Прочитайте статью[1] и сделайте ее критический анализ на русском языке.

Marketing research, the set of methods to collect and draw inferences from market-level customer and business information, has been the lifeblood of the field of marketing practice and the focus of much academic research in the past 30-plus years. Simply put, marketing research, the methods that surround it, and the inferences derived from it have put marketing as an academic discipline and as a functional area within the firm "on the map." Although these methods are here to stay, the radical changes resulting from the Internet and user-generated media promise to fundamentally alter the data and collection methods used to perform these methods.

In this study, we propose to harness the growing body of free, unsolicited, user-generated online content for automated market research. Specifically, we describe a novel text-mining algorithm for analyzing online customer reviews to facilitate the analysis of market structure in two ways. First, the VOC, as presented in user-generated comments, provides a simple approach to generating and selecting product attributes for market structure analysis. In contrast, traditional methods rely on a predefined set of product attributes (external analysis) or *ex post* interpretation of derived dimensions from consumer surveys (internal analysis). Second, the preponderance of opinion, as rep- resented in the continuous stream of reviews over time, provides practical input to augment traditional approaches (e.g., surveys, focus groups) for conducting brand sentiment analysis and can be done (unlike traditional methods) continuously, automatically, inexpensively, and in real time.

Our focus on market structure analysis is not by chance but rather due to its centrality in marketing practice and its fit with text mining of user-generated content. Analysis of market structure is a key step in the design and development of new products as well as the repositioning of existing products (Urban and Hauser 1993). Market structure analysis describes the substitution and complementary relationships between the brands (alternatives) that define the market (Elrod et al. 2002). Thus, if automated in a fast, inexpensive way (as described here), it can have a significant impact on marketing research and the decisions that emanate from it.

In market structure analysis approaches vary by the type of data analyzed (e.g., panel-level scanner data, aggregate sales, consumer survey response) and by the analytic approach (Elrod et al. 2002). Regardless of an internal or external approach, with few exceptions, market structure analysis begins with a pre-determined set of attributes and their underlying dimensions, the same set of survey or transaction sales data, and assumes that "all customers perceive all products the same way and differ only in their evaluation of product attributes" (Elrod et al. 2002, p. 229). In this study, we visualize market structure by

---

[1] На основе статьи Thomas Y. Lee and Eric T. BradLow Automated marketing research using online customer reviews Journal of Marketing Research Vol. XLVIII (October 2011), 881–894

applying correspondence analysis (CA) to product attributes mined from the VOC.

METHODOLOGY

Figure 1 summarizes our approach. We review the entire process here in a high-level way because many of these techniques are new to the marketing audience. The process begins with a set of online reviews in a product category over a specified time frame. For example, here, we consider the reviews for all digital cameras available at Epinions.com between July 5, 2004, and January 1, 2008. Figure 1, Step 1, shows three reviews: one each for cameras manufactured by Olympus, Hewlett-Packard (HP), and Fuji. In Step 2, screen scraping software automatically extracts details from each review, including the brand and explicitly labeled lists of pros and cons. Our goal is to group all phrases discussing a common attribute into one or more clusters to reveal what customers are saying about the product space, the level of detail used to describe an attribute, and the specific word choices that customers make. While some review sites do not provide user-authored pro and con summaries (e.g. Amazon.com), many others, including Epinions.com, BizRate, and CNET, do (Hu and Liu 2004). Exploiting the structure provided by pro and con lists enables us to avoid numerous complexities and limitations of automated language processing of prose-like text. This enables us to “automate” our process relative to most extant research.

Then, we separated all the pros and cons into individual phrases, as depicted in Column 1 of the table in Figure 1, Step 3. Preprocessing transforms the phrases from Column 1 into a normalized form. Column 2 depicts one step of pre- processing, namely, the elimination of uninformative stop- words such as articles (*the*) and prepositions (*of*). Next, we rendered each phrase as a word vector. The matrix of word vectors is depicted in the remaining columns of the table in Step 3. Each row is the vector for one phrase. Each vector index represents a word, and the vector value records a weighted, normalized count of the number of times the indexed word appears in the corresponding phrase.

Phrases are automatically grouped together according to their similarity. Similarity is measured as the cosine angular distance between word vectors and is identical to the Hamming distance used in computer science–based research. Step 4 depicts the K-means clustering of the phrases from Step 1. Conceptually, we can think of each cluster as representing a different product attribute. The example shows clusters for “zoom,” “battery life,” “picture quality,” and “memory.”

While any number of clustering algorithms is acceptable, we selected K-means for its simplicity and its familiarity to both the text-mining and marketing communities. As we noted previously, the principal contribution of our work is the elicitation of product attributes directly from the customer.

In product design, a product architecture defines the hierarchical decomposition of a product into components (Ulrich and Eppinger 2003). In the same way, Step 5 depicts the hierarchical decomposition of a product attribute into its constituent dimensions (i.e., attribute levels, similar to con- joint analysis). In Step 5a, we show a conceptual decomposition of the digital camera product attribute “memory.” In Step 5b, we show an actual decomposition using only the phrases from Step 1. The decomposition is treated as a linear programming assignment problem (Hillier and Lieberman 2001). The objective is to assign each word in the attribute cluster to an attribute dimension. Each word phrase defines a constraint on the assignment: Any two words that co-occur in the same phrase cannot be assigned to the same attribute dimension. Thus, we know that “smart” and “media” cannot appear as a value for the attribute dimension quantity (4, 8, or 16) or for the attribute dimension of memory unit (MB). Note that not all phrases include a word for every dimension. Intuitively, this is both reasonable and an important aspect of capturing the VOC. We want to know not only what customers say but also the level of detail with which they say it.

For the algorithm, phrases that do not include a word for each attribute simply represent a smaller set of co-occurrence constraints than a phrase containing more words.

ANALYSIS AND EVALUATION

We evaluate the quality and efficacy of analyzing market structure using online product reviews in three ways: First, we ask whether any new information can be discovered from user-generated product reviews not otherwise obtained from existing methods. Second, we measure the importance of attributes discovered from within product reviews. Finally, we use CA to analyze the VOC and ask whether any meaningful managerial insight is gained by visualizing market structure using product reviews. The first two measures support the use of reviews as a complementary method for generating possible attributes for use in market structure analysis. The third measure supports reviews and their corresponding word counts as a complementary method for selecting attributes and deriving pairwise brand distances for visualizing market structure. In this section, we evaluate our approach using all three measures on an actual set of digital camera product reviews.

Digital Cameras

Our initial data set consists of 8226 online digital camera reviews downloaded from Epinions.com on July 5, 2004. The reviews span 575 products and product bundles that range in price from $45 to more than $1,000. Parsing the pro and con lists produces the aforementioned phrase × word matrix that is 14,081 phrases × 3364 words.

Generating Attributes from Product Reviews

The purpose of mining the text of user-generated online reviews is to complement rather than replace existing methods for analyzing market structure. As a measure of the value within the text, we ask whether reviews reveal product attributes not found using traditional measures, such as those used to create expert buying guides. As a metric, we compute precision (P) (Salton and McGill 1983) as the number of automatically generated attributes and dimensions also used by experts in published buying guides. Conversely, recall (R) counts the number of attributes and levels named in professional buying guides that are automatically discovered in the VOC.

The columns of Table 2 correspond to each of ten expert guides aimed at the same consumer audience and available during the period covered by our reviews. For this study, the guides represent traditional methods for identifying meaningful product attributes. Epinions (A) comprises the attributes and levels by which customers can browse the digital camera product mix at Epinions (in 2005), and Epinions (B) represents a buying guide available on the Epinions website (in 2005). CR02–CR05 represent print buying guides for digital cameras from *Consumer Reports* for 2002–2005. The “Auto” column represents the attributes and dimensions derived automatically from online reviews. Row 1 indicates the average precision, and Row 2 indicates the average recall between the column source and all other reference sources.

Comparing the expert guides with one another reveals that there is no consensus among the experts on what the “important” attributes are. The absence of consistency indicates that analyzing online product reviews is not necessarily a redundant exercise; the VOC reveals interest in specific product attributes that are not identified by experts. If we observed high precision at any level of recall that would indicate polling the VOC is redundant; that is, consumers do not reveal any information not already captured by existing methods. Low precision and high recall would indicate that, in their reviews, customers mention every possible attribute, which offers no discriminatory power. The data indicate that analyzing reviews yields high recall compared with expert attributes. This means that users report nearly all the attributes that experts do. In pairwise comparisons, the VOC has better recall than any other single expert. In addition, median precision suggests that consumers mention what the experts do but that consumers also mention some additional attributes.

Importance of Attributes from Product Reviews

Through comparison with expert guides, we find that customers mention unseen attributes. To determine the significance, we conducted a laboratory survey in which participants evaluated the importance of different attributes for the purpose of purchasing a new digital camera. We find that automated analysis of online product reviews can support managerial decision making in at least two ways. First, our approach can identify significant attributes that experts otherwise overlook. Second, the reviews can serve as a filter for other attribute elicitation methods; attributes that experts and also customers identify may have more salience for purposes of product marketing and design.

We constructed a set of product attributes for testing by reconciling the 39 attributes shown in Table 1 with all the attributes identified in the ten reference buying guides listed in Table 2. After we eliminated duplicate attributes, we divided the resulting set of 55 attributes into overlapping thirds to reduce any individual respondent's burden. A few attributes were repeated in each third as an additional validity check. Each participant viewed between 20 and 21 attributes rated their "familiarity with" and the "importance of" each attribute using a 1–7 scale.

Participants were prompted to answer all questions. In particular, they were reminded to answer the second question for each attribute even if they answered [1] for the first question. In Part 2 of the survey, to understand the role that expertise might play, participants provided their self-assessed expertise on digital cameras (1 = "novice," and 7 = "expert"). Finally, participants provided a standard set of demographic variables such as age, gender, and education level. We used these variables as covariates to verify our main findings.

Figure 2 summarizes our main results, plotting the mean familiarity versus mean importance for each of the 55 attributes. We labeled attributes that appeared only in one or more buying guides "Expert Only," symbolized by circles. We labeled attributes that appear in at least one buying guide and also in our automated results "Expert + VOC," identified by squares. Finally, we labeled attributes that emerged only from our automated analysis of reviews "VOC Only," plotted as "X" symbols. As we would expect, the graph indicates a general trend upward and to the right. Users are more familiar with those product attributes that they tend to consider important.

Selecting Attributes and Visualizing Market Structure Using Product Reviews

By comparing expert guides and consumer surveys, we demonstrate that customer reviews can complement existing methods for generating attributes used in market structure analysis. As a third measure of efficacy, we use attribute counts within customer reviews to select attributes, calculate brand distances, and visualize market structure. A product brand is associated with each online review. Using the automatically generated attribute clusters, we generate a brand by attribute matrix, counting the number of brand occurrences (number of phrases) for each attribute. We row (brand) normalize the matrix by the total number of phrases for that brand (Nagpaul 1999). Then, to turn this into a visual "map," we use CA, a technique for analyzing two- way, two-mode frequency data (Everitt and Dunn 2001), making it more appropriate for this task than the continuously scaled multidimensional scaling procedures commonly used for market structure maps. The CA approach is designed to help generate hypotheses by representing the data in a reduced space as determined by consulting the eigenvalues and the corresponding scree plot (Greenacre 1992).

To help interpret the dimensions in the reduced space, we use (stepwise) regression of each brand's (x, y) coordinates on the derived attributes. The probability for entry is .05 and probability for removal is .1, though other values were tested and yielded high robustness. To ensure stability of the regression results, we use a rotationally invariant, asymmetric CA (Bond and Michailidis 1997).

To make the dimensions both interpretable and actionable, we follow the "house of quality" (Hauser and Clausing 1988) in relating customer needs, as represented by user-generated reviews, to actionable manufacturer specifications. Specifically, product attributes elicited from online reviews include different granularities ("zoom" vs. "10xoptical zoom") or reflect different levels of technical sophistication ("low-light settings" vs. "ISO settings"). Consulting the set of professional buying guides, we manually mapped our automatically generated attributes onto a coarser but actionable categorization of specifications.

For attribute selection, Figure 3 depicts the brand map in two dimensions based on the initial set of digital camera reviews analyzed for this study alongside the scree plot of the eigenvalues and cumulative percentage of inertia. In this instance, the average percentage explained is 20%, indicating a clear separation between two and three dimensions. More generally, we recognize that the reliability of any such model is dependent on fit as represented by the percentage of inertia captured by those dimensions. At the limit, we imagine that marketing managers might use the two-dimensional figures as a point of departure for interpretation with respect to decision making. For example, we generated the additional F1–F3 and F2–F3 plots and found that the results are consistent with our two-dimensional interpretation.

When interpreting CA maps, it is important to note that the dimensions only reflect relative relationships between brands and the underlying attributes. Being further from the origin does not necessarily imply higher term counts along one dimension or another. Rather, distance from the origin indicates greater deviation from the "average" brand.

We next assess the degree to which the brand map we derived from the VOC, as derived in our research, is consistent with empirical market share and documented brand strategies. Based on 2004 market share numbers from IDC, the top three manufacturers, in descending order, were Kodak, Sony, and Canon. All other brands had share percentages of 10% or less. The corresponding brand map reveals three distinct clusters, with Nikon as a notable outlier. In the mind of the customer, the secondary manufacturers all cluster by one of the three market leaders in a classic imitation strategy.

The positioning is also consistent with documented brand strategy, based on proprietary market research reports from the period. HP and Kodak explicitly attempted to differentiate themselves by emphasizing the camera-to-personal computer and camera-to-printer connection, linking them in the mind of the consumer (Worthington 2003). Among the best-known vendors, Nikon was the brand of professional photographers rather than lay consumers, late in its commitment to digital, and focused on the advanced amateur (Lyra Research 2003). Therefore, it is unsurprising that Nikon would be furthest from the origin (brand average).

At the same time, the brand map derived from the VOC reveals information that complements market share numbers and the proprietary market research reports. Traditional market research might cluster Fuji and Olympus on the basis of their shared commitment to and introduction of the xD-Picture Card memory format (Lyra Research 2003). Likewise, the “Four-Thirds” partnership between Olympus and Kodak to standardize image-sensor size and a removable lens mount might lead us to cluster those two brands (Lyra Research 2003). Instead, we observe that in the mind of the consumer, Olympus is more closely associated with market leader Sony and fellow follower Panasonic. Despite introducing a unique memory format with Olympus at the time, Fuji is paired with market leader Canon and Casio in the mind of the consumer.

One advantage of analyzing market structure on the basis of user-generated product reviews is the ability to quickly repeat the process and analyze changes in market structure over time. To this end, we collected a parallel set of 5567 digital camera reviews from Epinions.com dated between January 1, 2005, and January 28, 2008. The new set of reviews revealed five new attributes replacing five previously formed clusters (see Table 3). Although it is difficult to discern the underlying causes, the changes have face validity. As the customer base becomes increasingly sophisticated and increasingly connected online, the need for instructions and support has shifted toward online self-service. Like-wise, the ubiquity of personal computers and online photo management software may have shifted such functions away from the camera. In keeping with the theme of a more technically sophisticated audience, functions such as ISO settings, multiple-shot modes, and white/color balance are more significant or reflect active campaigns.

Although many attributes remain the same across time periods, changes in both customers and the marketplace may drive (or reflect) brand repositioning over time. To visualize these changes, we use the supplemental points method of CA (Nagpaul 1999). We construct a brand x attribute count matrix for Period 2; applying the transformation weights from the Period 1 decomposition, we map the Period 2 brand positions onto the Period 1 factors. Using supplemental points, we can trace market evolution between two (or more) periods on a common set of factors (axes). Figure 4 overlays Period 2 (year ≥ 2005) on the two-dimensional space defined for Period 1 (year < 2005).

Market share numbers for 2006 and 2007 published by the NPD Group and Hon Hai Worldwide show that Kodak, Sony, and Canon remained the three major manufacturers, though by 2007, Canon had emerged as the overall market leader, with Kodak falling to third. Consistent with market share, the Period 2 market structure continues to echo the overall positioning of a market leader with a cluster of imitators around each. However, in Period 2, the clusters exhibit a convergence in which even Nikon has now drawn closer to the three market leaders.

To visualize market structure, the VOC is mapped onto a set of product specifications (Hauser and Clausing 1988). Different users may refer to the same product attribute using different terminology such as “ISO settings” versus “photos in dim lighting.” General terms such as “easy to use” need to be translated into actionable features such as “menus,” “navigation,” “auto set- tings,” and so on. As noted previously, in this research, we followed the “house of quality” (Hauser and Clausing 1988) in manually rolling up different user terms into a common, actionable specification.

However, the manager may discover knowledge within specific consumer word choices. Focusing on F1 and F2 from our original market structure analysis, we analyze the degree of association between the three market leaders (Kodak, Canon, and Sony) and customer word choice on F1 and customer word choice on F2. Word choice in F1 is separated into two classes: the generic terms “ease of use” and “easy to use” and the specifications “navigation” “menus,” and so on. Likewise, word choice in F2 is separated into two classes: lay terminology “low-light” photography, “dim-light,” and so on, and advanced terminology “ISO settings,” “ISO adjustment,” and so on. Because we manually map different word choices onto a common factor, it is entirely possible that we are missing additional terminology. For example consumers citing “ease of use” could also be implicitly referring to ergonomic factors such as “grip” or “lightweight.” Our analysis here is intended to convey the potential of drilling down on word choice rather than provide an authoritative accounting for a specific product attribute.

A chi-square of 9.087 (p = .011) for F1and a chi-square of 21.409 (p < .0001) for F2 suggest a clear relationship between brand and word choice. We constructed the corresponding 2 by 2 contingency tables and report both the chi- square and the z-statistic from the log (odds ratio) in Tables 4 and 5.

We observe that of the big three players, Kodak is clearly the brand most associated with lay terminology such as “low-light photography” (F2) rather than “ISO” and generic words such as “ease of use” (F1). In contrast, Canon users clearly tend toward more sophisticated terms such as “ISO settings.” These findings reinforce market research from the time period, which reports that gender, age, and education level all exhibit strong brand preferences for Kodak and Canon (Lyra Research 2005). For example, older consumers and those most likely to have stopped schooling after high school or a few years of college prefer Kodak. Canon users are typically younger and have at least a college degree.

The visualization characterizes the “must-have” product attributes as defined by Kano’s needs hierarchy (Ulrich and Eppinger 2003). Must-haves define basic product requirements. Customers criticize products for their failure on must-haves; they do not praise products for must-haves because these attributes are expected. Consistent with this definition, the figure exhibits the expected brand differentiation based on cons. Positive comments mapped onto the same (F1, F2) dimensions reveal a single cluster of all brands, which share a limited number of pro remarks. Although attributes are not strictly “needs” in Kano’s terms, we can think in terms of those attributes that support specific needs.

Using stepwise regression, F1 defines the minimum that *any* digital camera is expected to provide: picture quality, start-up time, lens cover, and shutter lag. These attributes certainly match our intuitive sense of must-haves for digital cameras. In contrast, F2 defines those must-have attributes that distinguish specific product segments within the market space. Specifically, F2 (size, interchangeability of lenses, and the degree of in-camera programmability [e.g., menus, options]) separates the compact, “prosumer,” and digital SLR product categories.

The structural map by polarity is most revealing when interpreted in the context of our two-period market structure map (Figure 4). Situated at the origin on both pro and con, Canon essentially defines the space. This is consistent with the Canon’s position as a leader whose market share grew from Period 1 to Period 2 and is further evidenced by the Period 2 convergence of all groups toward the Canon brand cluster.

Market structure suggests three distinct clusters with leaders and imitators (Figure 4); focusing on polarity helps explains differences in market share. We know that Kodak and HP share a similar market strategy. However, HP clearly lags behind in sales; performance in the must-haves and the relative importance of the particular attribute dimensions help explain this difference. Likewise, imitators Fuji and Casio are closer to leader Canon than Olympus and Panasonic are to leader Sony. This again makes sense considering the overall market convergence toward the Canon, Fuji, and Casio cluster in Period 2.

Nikon, like Canon, is effectively neutral with respect to pros and cons. This is consistent with Nikon’s strategy during the study period to produce excellent technology but to focus on a different customer base than the rest. Casio, whose market share numbers were too low to break out by IDC, NPD, or Hon Hai during the study period, is an outlier in the polarity maps.

Questions for your consideration

1. What is the main research question of the article?
2. Which research approaches and (or) methods you consider to be controversial or not quite suitable for this research? Why?
3. What are the limitations of the current research and what recommendations could be given for further research?
4. How the research method described can be used by managers in Russian firms and multinational corporations?
5. What methods of data analysis do you consider suitable for this research?

Figure 1

HEURISTIC DESCRIPTION OF TEXT PROCESSING OF USER-GENERATED PRO–CON LISTS

Step 1. Reviews

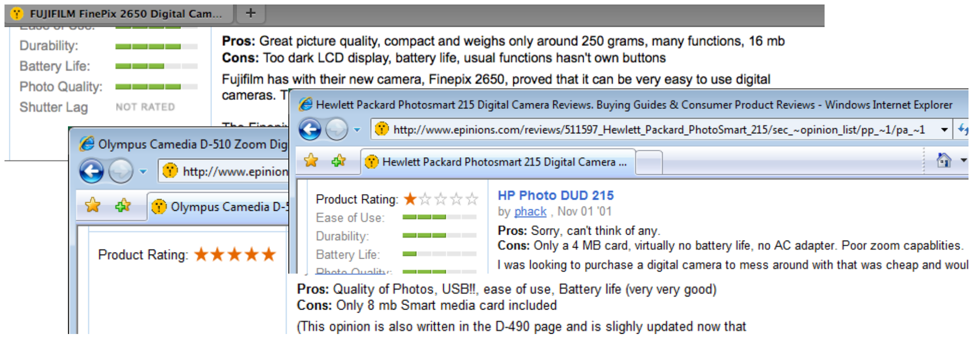

Step 2. Screen Scraping

Olympus: Quality of Photos, …, Battery life (very very good), Only 8 mb Smart media card included
HP: …, Only a 4 mb card, virtually no battery life, no AC adapter, Poor zoom capabilities
Fuji: Great picture quality, 16 mb, battery life, …

Step 3. Matrix of Word Vectors

Brand	Original phrase	Stop-words removed	only	life	mb	card	zoom	battery	…
Olympus	Quality of Photos	quality photos							
Olympus	Battery life (very very good)	battery life good		1				1	
Olympus	Only 8 mb Smart media card …	8 mb smart media card inc	1		1	1			
HP	Only a 4MB card	4 mb card			1	1			
HP	Virtually no battery life	virtually battery life		1				1	
HP	No AC adapter	no ac adapter							
HP	Poor zoom capabilities	zoom capabilities					1		
Fuji	Great picture quality	great picture quality							
Fuji	16 mb	16 mb			1				
Fuji	battery life	battery life		1				1	

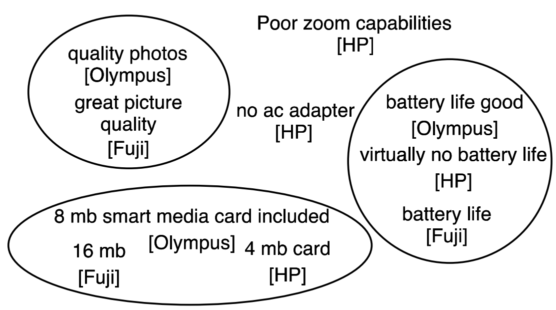

Step 4. Clustering Phrases

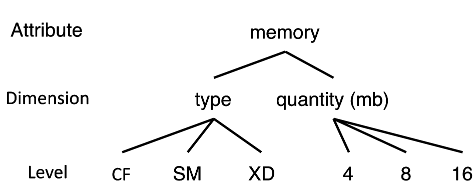

Step 5a. Attributes and Their Dimensions

D1	D2	D3	D4	D5	D6
8	mb	smart	media	card	included
4	mb			card	
16	mb				
8	mb			card	

Step 5b. Attributes and Their Dimensions in the Text

Table 1

AUTOMATICALLY GENERATED PRODUCT ATTRIBUTES

LCD	Memory/Screen	Shutter (delay, lag)	Optical (zoom, viewfinder)	Floppy (storage media)	Support (service)	Shoot
Red eye	Lens (cap, quality, manufacturer)	Print (size, quality, output)	Slow (start-up, turn on, recovery)	Flash (memory card, photo)	Body (design, construction)	USB
Price	Picture (what, where)	Cover (lens, LCD, battery)	Feel (manufacturer, construction)	Battery (life, use, type)	Movie (audio, visual)	Size
Focus	Edit (in camera)	Disk	Instruction	Photo quality	Menu	Control
Features	Adapter (AC)	MB (memory)	Low light	Picture quality	Resolution	Zoom
Software	Image (quality)	Macro (lens)	Megapixel			

Table 2

INTERNAL CONSISTENCY: AVERAGE PRECISION AND RECALL BETWEEN ONE SOURCE AND ALL OTHERS

	Auto	*E(A)*	*DP*	*Mega*	*Biz*	*CNET*	*E(B)*	*CR02*	*CR03*	*CR04*	*CR05*	*Mean*
Precision	.37	.69	.33	.57	.41	.27	.48	.24	.26	.27	.37	.42
Recall	.72	.23	.55	.23	.48	.33	.46	.38	.37	.37	.50	.39

Figure 2

USER ASSESSMENTS OF FAMILIARITY VERSUS IMPORTANCE

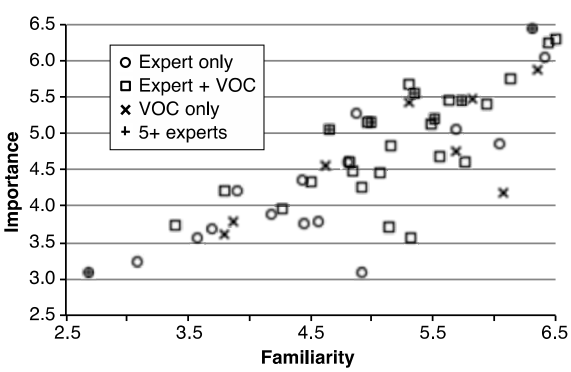

Figure 3

MAPPING THE MARKET USING CUSTOMER REVIEWS

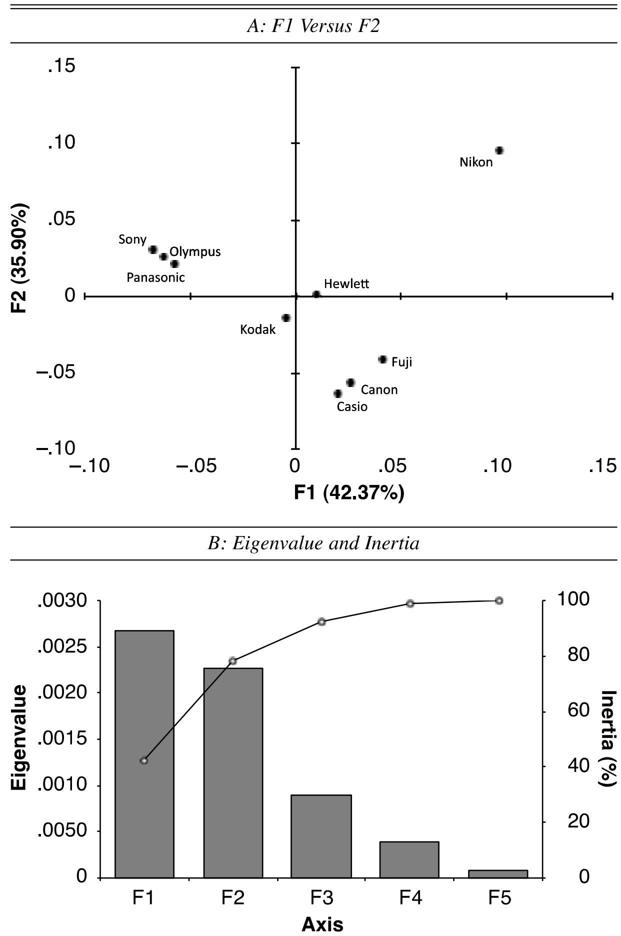

Figure 4

EVOLUTION OF MARKET STRUCTURE WITH PERIOD 2 AS SUPPLEMENTARY POINTS

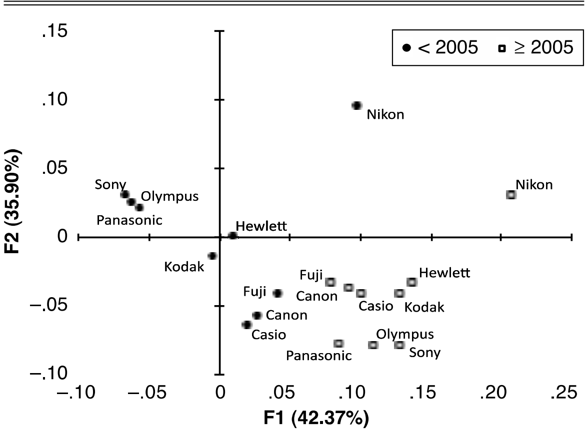

Table 3

CHANGES IN AUTOMATICALLY GENERATED ATTRIBUTES BEFORE AND AFTER 2005

Before 2005	Size	Support (service)	Feel (manufacturer)	Instruction	Edit (in camera)
After 2005	ISO	Modes	Accessories	Easy to use	White/color balance