Stay Elsewhere? Improving Local Search for Hotels Using Econometric Modeling and Image Classification*

Beibei Li
bli@stern.nyu.edu

Anindya Ghose
aghose@stern.nyu.edu

Panagiotis G. Ipeirotis
panos@stern.nyu.edu

Department of Information, Operations, and Management Sciences
Leonard N. Stern School of Business, New York University

ABSTRACT

One of the common Web searches that have a strong local component is the search for hotel accommodation. Customers try to identify hotels that satisfy particular criteria, such as service, food quality, and so on. Unfortunately, today, the travel search engines provide only rudimentary ranking facilities, typically using a single ranking criterion such as distance from city center, number of stars, price per night, or, more recently, customer reviews. This approach has obvious shortcomings. First, it ignores the multidimensional preferences of the consumer and, second, it largely ignores characteristics related to the location of the hotel, for instance, proximity to the beach or proximity to a downtown shopping area. These location-based features represent important characteristics that influence the desirability of a particular hotel. However, currently there are no established metrics that can isolate the importance of the location characteristics of hotels. In our work, we use the fact that the overall desirability of the hotel is reflected in the price of the rooms; therefore, using hedonic regressions, an established technique from econometrics, we estimate the weight that consumers place on different hotel characteristics. Furthermore, since some location-based characteristics, such as proximity to the beach, are not directly measurable, we use image classification techniques to infer such features from the satellite images of the area. Our technique is validated on a unique panel dataset consisting of 9463 different hotels located in the United States, observed over a period of 5 months. The final outcome of our analysis allows us to compute the "residual value" of a hotel, which roughly corresponds to the "value for the money" of a particular hotel. By ranking the hotels as using our "value for the money" approach we generate rankings that are significantly superior to existing techniques. Our preliminary user studies show that users overwhelmingly favor the hotel rankings generated by our system.

*This work was partially supported by a Microsoft Virtual Earth Award, and by NSF grants IIS-0643847 and IIS-0643846. Any opinions, findings, and conclusions expressed in this material are those of the authors and do not necessarily reflect the views of the Microsoft Corporation or of the National Science Foundation.

Proceedings of the 11th International Workshop on Web and Databases (WebDB 2008), June 13, 2008, Vancouver, Canada

Keywords

local search, hedonic pricing, econometrics, remote sensing image classification, texture feature extraction, ranking services

1. INTRODUCTION

Local search for hotel accommodations is a component of general Web searches that is increasing in popularity as more users book and arrange their trips online. Customers try to identify hotels that satisfy particular criteria, such as food quality, service, and so on. Unfortunately, today, travel search engines provide only rudimentary ranking facilities, typically using a single ranking criterion such as distance from the city center, number of stars, price per night, or, more recently, customer reviews. This approach has obvious shortcomings. First, it ignores the multidimensional preferences of the consumer and, second, it largely ignores characteristics related to the location of the hotel, for instance, in terms of proximity to the beach or proximity to a downtown shopping area. These location-based features represent important characteristics that influence the desirability of a particular hotel.

However, currently there are no established metrics that can isolate the importance of the different characteristics of the hotels. Existing empirical work only focused on 1-2 location-based characteristics for merely 10-20 hotels within small geographical areas [4, 20]. Besides, most of these studies used rather rudimentary and non-scalable data collection approaches such as interviews of hotel managers or personal observation [4]. Furthermore, due to the small sample size, such approaches have been criticized for potential selection bias, due to the fact that the data sets consisted of "convenience samples" and the precision of the empirical studies has been discounted.

In our work, we try to generate a general hotel ranking that can serve the general customer population and which we can later extend to include a personalization component. (Personalized ranking is out of the scope of this paper.) For this, we start by using user surveys, to see what characteristics are important for consumers. Then, for a large data set of almost 10 thousand hotels, we identify these features, using a variety of techniques, including customer review analysis and on-demand annotations using Amazon Mechanical Turk. Furthermore, since some location-based characteristics, such as proximity to the beach, are not directly measurable, we use image classification techniques to infer such features from the satellite images of the area. Once we have collected and measured all the important hotel characteristics, we estimate their importance by using the fact that the overall desirability of the hotel is reflected on the price of rooms. Using hedonic regressions, an established technique from econometrics, we estimate the weight that consumers place on different hotel characteristics. The final outcome of our analysis allows us to compute the "residual value" of a hotel, which roughly corresponds to the "value for the money" of a particular hotel. By ranking the hotels as using our "value for the money" approach we generate rankings that are significantly superior to existing techniques. Our preliminary user studies show that users overwhelmingly favor the hotel rankings generated by our system.

The rest of the paper is organized as follows. Section 2 gives the overview of our work. Section 3 describes how we conducted our survey using Amazon Mechanical Turk to identify the important hotel characteristics. Section 4 discusses the data collection procedures that we employed. Section 5 discusses the details of our econometric model and Section 6 describes how we use the econometric model to derive a new ranking approach, which our preliminary results indicate to be better than existing baselines. Finally, Section 7 gives the conclusion as well as future directions.

2. OVERVIEW

In general, our goal is to empirically estimate the economic value of different hotel characteristics, especially the location-based characteristics given the associated local infrastructure. Then, using this economic analysis, we can locate the hotels with specific criteria that provide "the best value for the money." We achieve this by combining state-of-the-art econometric modeling with user-generated content data and image classification methods. Our work involves three stages:

1. Identify the important hotel characteristics that influence the hotel prices and measure them efficiently and effectively.
2. Estimate how these hotel characteristics influence the hotel prices.
3. Improve local search for hotels by incorporating the economic impact of the hotel characteristics.

Specifically, in the first stage, we want to find out particular hotel characteristics that are most highly valued by customers and hence, contribute to the aggregate prices of the hotels. Users, beyond the directly measurable characteristics, such as "number of stars," tend to value location characteristics such as proximity to the beach, or proximity of downtown and shopping areas. In our work, we incorporate the satellite image classification and use both human and computer intelligence, which in the end contributes to a more comprehensive dataset.

In the second stage, we use *hedonic regressions* [17] and estimate the economic value of each hotel characteristic; in this way, we quantitatively analyze how each feature influences the price for a hotel and hence, we estimate its importance.

In the third stage, after inferring the economic significance of the location- and service-based hotel characteristics, we incorporate them into a local ranking function. By doing so, we can provide customers with the "best-value" hotels, hence improving the quality of local search for such hotels.

3. IDENTIFICATION OF HOTEL CHARACTERISTICS

In this section, we discuss the procedure for identifying the important hotel characteristics. Our analysis is based on the idea that characteristics mentioned frequently by consumers are the ones that ultimately determine the aggregate prices of the hotels. To perform the survey, we used the *Amazon Mechanical Turk*[1] (MTurk) service, which is an online tool for distributing small tasks to a large number of users; each user receives a small monetary compensation for completing the task. For our survey, we asked 100 anonymous MTurk users for hotel characteristics that would influence their choice.[2] Our analysis identified two broad categories of hotel characteristics:

1. Location-based hotel characteristics
2. Service-based hotel characteristics

Location-Based Characteristics: Location-based characteristics refer to features that describe the geographic environment information of a hotel accommodation. There are 7 characteristics in this category:

- Near the beach
- Near the waterfront (sea, lake, river), not necessarily with a beach
- Near public transportation
- Near downtown
- Close to the interstate highway
- External amenities (i.e., near restaurants and shops)
- Safe neighborhood

We describe in the next section how we automate the collection of such features that are not always trivial to compute.

Service-Based Characteristics: Service-based characteristics are used for specifying the performance of a hotel accommodation, including hotel amenities, appearance, service, and so on. There are 4 broad characteristics in this category:

- Hotel class
- Customer reviews
- Total number of rooms
- Internal amenities

Here, "Hotel class" is an international standard ranging from 1-5 stars representing low to high hotel grades. "Consumer reviews" is a broad option covering the word-of-mouth that the hotel has received on the Internet, such as on the popular TripAdvisor site; we measured word-of-mouth by using the "Popularity rank" of a hotel as a proxy of its popularity, together with the number of reviews and the reviews rating. "Internal amenities" is the aggregation of hotel internal amenities, including "24 hour front desk," "ice machine," "beautiful furnishings," "credit card payment," "cable TV," "pets allowed," "size of the room," "wheelchair accessible," "friendly staff," "free breakfast," "cleanliness," "wakeup call service," "nonsmoking," "gym," "iron," "internet reservation available," "high speed internet," "kids friendly service," "laundry services," "swimming pool," "parking," "kitchenette," and "spa."

4. DATA

After identifying the important hotel characteristics for our analysis, we now shift our discussion on how we effectively collect the corresponding data.

4.1 Methodology Overview

We collected our data in different ways from September 2007. We monitored a total of almost 10 thousand hotels in the US, gathering prices and service-based hotel characteristics from the website of TripAdvisor.[3]

For the location-based characteristics, which contribute to the larger component of our work, we collected the data utilizing the Microsoft Virtual Earth Interactive SDK, which we will discuss further in Section 4.2.

[1] http://www.mturk.com

[2] We also conducted an survey to determine the demographics of MTurk users. Most of the users are based in the US, their income distribution is similar to the income distribution in the US, and there is a slight bias towards younger users.

[3] http://www.tripadvisor.com

Method	Characteristic
Image Classification	Near the beach Near downtown
Virtual Earth Interactive SDK	Number of restaurants Shopping destinations
MTurk	Close to the interstate highway Near the waterfront Near public transportation
FBI online Statistics	City annual crime rate City population
TripAdvisor	Hotel class Customers' review count Total number of rooms Hotel internal amenities Popularity rank

Table 1: Methods for Measuring Hotel Characteristics

The difficulty of extracting location-based characteristics varies for each characteristic. A characteristic like "Near restaurants and shops" can be computed using an online API that allows such "local search" queries. However, a characteristic like "Near the beach" cannot be answered by existing mapping services. To measure such characteristics, we used automatic image classification of satellite images.

In general, we collected the location-based hotel characteristics primarily through four different methods:

1. Commercial characteristics were computed via local search queries, through the Virtual Earth Interactive SDK.
2. Geographical characteristics, such as proximity to a waterfront, were derived by image classification.
3. Geographical characteristics too difficult even for image classification algorithms were classified using on-demand human annotation through the Amazon MTurk service.
4. Characteristics related to neighborhood safety, containing two sub-characteristics, were acquired from the FBI online statistics:
    - City annual crime rate (mean 2000-2006)
    - City population (mean 2000-2006)

The hotel characteristics and their corresponding extracting methods are listed in Table 1. Except for the image classification part, the other location characteristics did not pose any significant challenges, therefore in the rest of this section we focus and describe the image classification part.

4.2 Image Data Retrieving

As mentioned in the previous sections, our research is based on data from Microsoft Virtual Earth, a service for interactive mapping applications. Virtual Earth provides both a main mapping site[4] and a JavaScript API[5] for developers to embed the maps in their own site. It supports three different types of imageries: road, aerial and hybrid, down to foot-per-pixel quality in urban areas. The hybrid pictures are satellite images with mapping information (e.g., roads) superimposed over the satellite image.

The goal of our research was to automatically identify whether a hotel is located in a downtown area, or next to a beach. (As a reminder, it is impossible to get this information from a mapping service, or from the TripAdvisor website.) For this, we extracted hybrid satellite images (sized 256×256 pixels) using the Visual Earth Tile System[6], for each of the 9463 hotel venues located in the United States, with 4 different zoom levels for each. These 9463×4 images were then used to extract information about the surroundings of the hotel, through image classification and through human inspection using MTurk.

4.3 Texture Feature Extraction

Remote sensing image classification requires consideration of many factors [6, 14, 15, 18] such as: determination of classification system, selection of training samples, image preprocessing, feature extraction, post-classification processing, and accuracy assessment [12]. Among all of these, selection of effective features and determination of a suitable classification method are especially significant for improving classification performance. For land-cover classification [12], the most important aspect is textural and contextual information, which captures the spectral response as a unique feature of the images.

In our work, this information is extracted using Gabor wavelets. Gabor texture feature extraction and Gabor wavelets provide the best overall performance compared to other multi-resolution texture features using the Brodatz texture database [11–13]. It shows stable performances on capturing the repeating patterns of local variation of pixel intensities. The experimental results on large aerial photographs also indicate that Gabor wavelets give good pattern retrieval accuracy [13].

Specifically, we started by transforming the 256×256 color images into greyscale. Then, before conducting the texture features extraction, we performed Principal Component Analysis (PCA) on each 256×256 pixel-based image. By doing so, we decomposed each image into $7 \times 7 = 49$ overlapping regions of 64×64 pixels. Then, we computed a texture feature vector represented by Gabor wavelets for each region. A smaller region size will not cover sufficient spatial or texture information to effectively characterize land cover types, whereas a larger region size may involve too much information from other neighboring land cover types. In our special case, a region size of 64×64 gives us the best accuracy, compared with other region sizes, such as 128×128 and 32×32. A feature vector is constructed using 4 scales and 6 orientations, resulting in a 1×48 feature vector for each region:

$$\vec{R} = [\mu_{1,1}, \sigma_{1,1}, \ldots, \mu_{4,6}, \sigma_{4,6}]. \tag{1}$$

Then, we represent each of our original 256×256 satellite images as a $1 \times (48 \times 49) = 1 \times 2352$ vector:

$$\vec{I} = [\mu_{1,1,1,1}, \sigma_{1,1,1,1}, \ldots, \mu_{7,7,4,6},\ \sigma_{7,7,4,6}], \tag{2}$$

where the first two indices (from 1 to 7) represent the row and column positions of a certain region in the original image; the third subscript (from 1 to 4) represents the scale, and the last index (from 1 to 6) represents the orientation.

4.4 Image Classification

While effective feature representation has significant impact on the outcome, a suitable classification method also plays an important role in the results. Non-parametric classifiers, such as Neural Network, Decision Tree and Support Vector Machines(SVM) have been proved for better results than parametric classifiers in complex landscapes [12]. Images that reveal the "Beach" characteristic and those that contain the "Downtown" characteristic provide typically good, highly discriminating features; however, the type of features varies (see

[4] http://maps.live.com

[5] http://dev.live.com/virtualearth/sdk

[6] http://msdn2.microsoft.com/en-us/library/bb259689.aspx

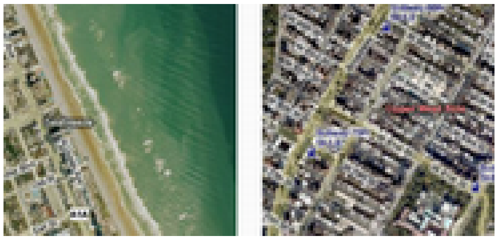

Figure 1: Beach Image and Downtown Image

		SVM	Decision Tree
Beach	Precision	0.908	**0.919**
	Recall	0.852	**0.911**
	Accuracy	0.912	**0.944**
	F-Measure	0.879	**0.915**
Downtown	Precision	**0.808**	0.778
	Recall	**0.915**	0.802
	Accuracy	**0.807**	0.720
	F-Measure	**0.858**	0.790

Table 2: Image Classification Results

Figure 1 for an example) so a classification method that works well for one may not work well for the other. For example, a beach is typically revealed through a straight, thick line of light color that can be nicely represented in a linear classifier. On the other hand, the downtown image is revealed by dense intersection of streets, together with dense "landmark" pointers for the different services offered in the area.

Therefore, we tested various non-parametric classification techniques: (1) Decision Trees, which are widely used for training and classification of remotely sensed image data, (due to its capability to generate human interpretable decision rules and its relatively fast speed in training and classification), and (2) SVM, which are highly accurate and perform well for a wide variety of classification tasks [2, 7].

We built the image classifiers as follows: First, we selected a set of 121 hotels and we requested 5 MTurk users to label each example according to its corresponding satellite images from 4 different zoom levels. The labelers answered whether there is a beach in the image, or whether the area is a downtown area. We applied a simple majority voting method to make the final decision from the multi-labels of the example. Second, we trained an SVM classifier on this dataset and used the trained SVM classifier to classify the images that corresponded to the remaining 9,342 hotels. To evaluate the performance of the classifier on truly unseen data, we also classified these images using Mechanical Turk; our results show that our SVM classifier has an accuracy of 91.2% for "Beach" image classification and 80.7% for "Downtown" image classification. We also used the C4.5 algorithm for classification, and noticed an accuracy increase for "Beach" and a decrease for "Downtown." The classification results are shown in Table 2.

According to the results presented in Table 2, the overall classification performance for the "Beach" images is much higher than the "Downtown" images. The reason of this may be that the "Beach" images often present a "sand strip," together with an "ocean margin" well distributed in density. This may provide more stable and distinct textural information for the "Beach" images, thus making them much easier to distinguish. From the results above, in order to acquire more accurate results, we choose Decision Tree C4.5 for the "Beach" images classification and choose SVM for "Downtown."

5. ECONOMETRIC ANALYSIS

In this section, we discuss how we estimate the economic value of different hotel characteristics. We first give a brief introduction of the hedonic regressions and then we describe how we used them for our setting.

5.1 Hedonic Regressions

The hedonic regression model is derived from characteristics theory, based on work by Lancaster [9, 10] and dates largely from Rosen's model [17]. It assumes that differentiated goods can be described by vectors of objectively and deterministically measured features and the consumer's valuation of a good can be decomposed into an implicit value of each product feature [17]. Implicit in the hedonic price framework is the assumption that a particular product can be viewed as consisting of various bundles of a small number of characteristics or basic attributes [3, 5]. In other words, hedonic models can be used to estimate the value that different product aspects contribute to a consumer's utility. For instance, a hotel can be decomposed to characteristics such as class(c), size(s), service(r), and the utility of the hotel for a consumer can be represented as a function $u(c, s, r, \dots)$ [1]. In order to estimate parameters of the hedonic models, it is common to reduce the problem into a linear instantiation, and regressions are used to estimate the parameters of the model. Hedonic regressions are commonly used in real estate economics to identify implicit market price for characteristics such as location variables and amenities with location-specific commodities, for example, housing and recreation services [4], as well as hotel accommodations which will be discussed in the following sections.

5.2 Definition of Variables

We use the hedonic pricing model to estimate the relationship between the equilibrium price of a hotel as a heterogeneous product and the contribution to that price from each hotel characteristic. In this way, each hotel k at time t can be represented as a bundle of s differentiated characteristics:

$$Z_{kt} = (Z_{1kt}, \dots, Z_{skt}) \tag{3}$$

According to the two different categories of hotel characteristics, this bundle is then divided into two corresponding types, such that there are l location-based characteristics $\bar{Z}_{kt}$ and $(s - l)$ service-based characteristics Z_{kt}. Thus, the new characteristics bundle is described as :

$$\begin{aligned} Z_{kt} &= (\bar{Z}_{kt(location-based)}, Z_{kt(service-based)}) \\ &= (\bar{Z}_{1kt}, \dots, \bar{Z}_{lkt}, Z_{(l+1)kt}, \dots, Z_{skt}) \end{aligned}$$

We added two additional dummy variables to capture the effects of a "*large city*", C_{kt}, and by the *holiday season*, S_{kt}. Let $p(Z_{kt})$ denote the market price for this bundle of characteristics, hence we have

$$p(Z_{kt}) = p(\bar{Z}_{1kt}, \dots, \bar{Z}_{lkt}, Z_{(l+1)kt}, \dots, Z_{skt}, C_{kt}, S_{kt}). \tag{4}$$

In our following empirical longitudinal study, price is the dependent variable, and all the location-and service-based characteristics, as well as city and season dummy variables are independent variables. We use a linear function of location-based and service-based characteristics in the following form:

$$\ln(PRICE_{kt}) = \beta_0 + \sum_{i=1}^{l} \beta_i \bar{Z}_{ikt} + \sum_{i=l+1}^{s} \beta_i Z_{ikt} + \delta_1 C_{kt} + \delta_2 S_{kt} + \mu_{kt}, \tag{5}$$

where $\bar{Z}_{ikt}$ are location-base variables, Z_{ikt} are service-based variables, C_{kt} is dummy variable for large cities, S_{kt} is dummy

Variable	Coef.	Variable	Coef.
$CLASS_{kt}$	0.6885	$ROOM_{kt}$	0.0292
C_{kt}	0.3505	S_{kt}	0.0247
$BEACH_{kt}$	0.2942	DT_{kt}	0.0227
$PUBTRAN_{kt}$	0.0818	$REVIEW_{kt}$	0.0079
EXT_{kt}	0.0754	POP_{kt}	-0.0003
$LAKE_{kt}$	0.0735	HW_{kt}	-0.0866
INT_{kt}	0.0573	$CRIME_{kt}$	-0.1693

Table 3: Estimation Results

variable for holiday season, $\beta_0 \ldots \beta_s$, and δ_1, δ_2 are parameters, and μ_{kt} are random error terms, i.i.d. as $N(0, \sigma)$.

5.3 Economic Value of Hotel Characteristics

Since our dataset is a *panel* dataset, i.e., it contains observations of the same hotel over a period of time, we need to use panel models that can capture the non-independence of the different observations. We estimate our model using Fixed Effect Vector Decomposition(FEVD) [16]. This is a recent technique that allows estimating time-invariant variables in panel data models with unit effects in an augmented fixed effects approach. First, we run a fixed-effects model to obtain the unit effects; then, we break down the unit effects into two parts, a part explained by the time-invariant and/or rarely changing variables and an error term; finally, we re-estimate the first stage by pooled OLS. In this way, we are able to capture more precise estimation by taking into consideration the time-invariant variables correlated with the unit effects. Meanwhile, we also look into the robust Random-effects GLS. Directionally, the results from both approaches are similar. Table 3 shows the estimates of coefficients by using FEVD.

Our results have shown that at the significance level of 1% with R^2 fit equal to 0.9763, all of the 14 independent variables are useful for the model. According to the estimated signs of the coefficients, we can qualitatively analyze the trend for the economic impacts of hotel characteristics.

There are two location-based characteristics which bring negative influences on the hotel price. Not surprisingly, one is the "Average Crime Rate." The higher the average crime rate reported in a local area, the lower price their hotel rooms can make a sale at. This indicates that neighborhood safety usually plays a vital role in hotel industry. Another negative factor is the "Highway." Hotels located close to the interstate highway exits (within 0.6 miles) are not as preferable as those which are located a little farther away. This is because a location adjacent to highway often leads to noise and insecurity, which may prevent customers from choosing to stay, thus causing a compromise in the hotel price. Notice that "Popularity rank" also has a negative sign, but it has a positive impact on the hotel price. (Because the smaller the hotel rank index is, the more popular the hotel.)

Meanwhile, there are 11 other characteristics which have positive coefficients: "Beach," "Lake/River," "Public transportation," "Downtown," "External amenities," "Hotel class," "Customers' review count," "Total number of rooms," "Internal amenities," "Large cities," and "Holiday season." Their economic impacts on the hotel price are positive. Among all the location-based characteristics, "Beach" has the greatest significance. Hotels that offer a "walkable beachfront"- a beachfront within 0.6 miles, attract travelers the most. Hence, it will lead to a dramatic raise in their room rates. Besides, a location in the downtown area, or near a lake or river can also increase the hotel price. Hotels providing easy access (within 0.6 miles) to public transportation, such as subway stations, airport shuttles, or on a bus line, can charge higher prices as well.

6. LOCAL RANKING FUNCTION

After estimating the economic impact for each hotel characteristic, we propose a local ranking function as follows:

$$Value_k = avg_k(Pred.Price) - avg_k(RealPrice). \qquad (6)$$

We define "Value" of a hotel as the difference between the average predicted price and the average real price for that particular hotel. In other words, given the hotel characteristics and their average price, does an *average*[7] consumer overpay for these characteristics or not? For example, if the average predicted price for a hotel is \$734, whereas its average real price is \$583, then its "value" will be \$734 - \$583= \$151. Here, the predicted prices are obtained by incorporating all the economic value of hotel characteristics. Therefore the \$151 is the "residual value" that remains after accounting for all the features of the hotel.

Then, we rank all the hotels according to their "value for the money" in a descending order, which gives a best evaluation on the hotel cost performance and provides customers with the best valued hotels consequently. Our preliminary ranking result for top 10 hotels with "best value for the money" in New York City, during the 2007 holiday season is shown in Table 4.

6.1 User Study

To evaluate the quality of our ranking technique, we conducted user study, using Amazon Mechanical Turk.

Comparison using titled rankings and correct hotel lists: First, we generated 7 different rankings for the top-10 New York City hotels: "price low to high," "price high to low," "maximum review count," "hotel class," "number of internal amenities," "total number of rooms," and "popularity rank" (generated by TripAdvisor). Then, we presented our ranking together with the 7 alternative rankings and asked 100 anonymous customers to choose their favorite. At the beginning, in order to help customers distinguish, we provided a short title for each ranking. When the rankings were titled, more than 50% of the customers chose our ranking, over the other 7 rankings listed. We also presented the lists in different order (e.g., "hotel class" first, followed by "popularity rank" and so on) and our observation remained the same, indicating that presentation order did not matter in this experiment. When we asked the users to justify their choice, most people said that they would prefer better hotel experiences if the price is right.

Comparison using titled rankings and *incorrect* hotel lists (Robustness Test): While the experiment with the titled rankings indicated that consumers like our approach, we wanted to estimate the reliability of the result. For this, we decided to keep the titles intact and swap the underlying hotel lists, putting for example the list produced by "hotel class" under our "value for the money" title. We noticed that consumers again voted overwhelmingly for the "value for the money"-titled ranking, even if the hotels were different. Although this indicated customers' strong preference for hotels with the "best value," this highlighted the need for a truly blind test since it implied that customers do not look at the actual hotels but are simply influenced by the title of the ranking.

Comparison using blinded lists: In order to obtain more objective evaluation, we decided to conduct a blind test to further eliminate the bias caused by the titles. Asking consumers to choose among 8 different, untitled rankings was proven to be

[7] Under a personalized ranking approach the value assigned to each characteristic may be different, but we leave personalization as a topic for interesting future research.

Top 10 New York City Hotels with "Best Value for the Money"	
1.	The Roosevelt Hotel
2.	Gramercy Park Hotel
3.	Hotel Gansevoort
4.	Four Season Hotel New York
5.	Crowne Plaza Manhattan
6.	Lowell Hotel
7.	Soho Grand Hotel
8.	Warwick New York Hotel
9.	The Sherry-Netherland Hotel
10.	Intercontinental The Barclay New York

Table 4: Local Ranking Results

difficult for users, as we received complaints about the difficulty of understanding the difference between the lists. Therefore, we compared now our technique in a *pairwise* fashion with the competing alternatives. In pairwise comparisons, the responses have shown that more than 80% of customers prefer our ranking ($p = 0.001$, sign test). The results of the blind test indicate that consumers indeed prefer our ranking approach, even when they are not aware of the ranking title.

Reasoning: We also asked consumers why they choose a particular ranking. The majority of the users indicated that they liked the *diversity* of the returned results, which were also priced competitively, while the other ranking approaches tend to list hotels of only one type (e.g., luxury hotels). Our "value-for-the-money" ranking show a variety of 30% 5-star, 40% 4-star, and 30% 3-star (or lower) hotels in the city, presenting information in a way that helps customers in their decision making process. Based on the qualitative opinions of the users, it appears that diversity is indeed an important factor that improves consumers' satisfaction. Our economic approach for ranking seems to introduce diversity naturally, without using any diversity-aware ranking techniques [19].

7. CONCLUSION AND FUTURE WORK

In this paper, we empirically estimate the economic value of different hotel characteristics, especially the location-based ones. We combine the state-of-the-art econometric modeling with user-generated content and image classification methods, on a unique dataset consisting of totally 157,414 observations based on 9,463 hotels located in the United States, over 5 months. Our research quantifies the economic impact of hotel characteristics and identifies the most crucial characteristics that influence the desirability of a particular hotel. After inferring the economic significance of each characteristic, we incorporate the economic value of hotels characteristics into a local ranking function and we improve the quality of local search for hotels, as indicated by our preliminary user studies.

In the future, we plan to analyze more hotel characteristics based on datasets observed over longer periods of time. Also, instead of price we plan to use demand and revenue as "objective variables", which will allow us to perform a more robust economic analysis. (Our analysis implicitly assumes that consumers *buy* at the posted prices, something that is not always the case.) We also aim to incorporate more advanced geomapping techniques to achieve a more comprehensive dataset. For example, we will use the multimap.com API to get precise information regarding public transportation. We also plan to use geonames.org, a site where users geotag locations with tags such as "beach", "hospital", "park" and so on; such tags can be used to train multiple image classifiers and expand our analysis. Finally, we want to use reviews and reviewer profiles [8] and examine potential personalization techniques for ranking.

References

[1] Archak, N., Ghose, A., and Ipeirotis, P. G. Show me the money! Deriving the pricing power of product features by mining consumer reviews. In *Proceedings of the Twelveth ACM SIGKDD International Conference on Knowledge Discovery and Data Mining (KDD-2007)* (2007), pp. 56–65.

[2] Azimi-Sadjadi, M. R., and Zekavat, S. A. Cloud classification using support vector machines. In *Proceedings of the IEEE 2000 International Geoscience and Remote Sensing Symposium (IGARSS 2000)* (2000), pp. 669–671.

[3] Berndt, E. R. *The Practice of Econometrics: Classic and Contemporary*. Addison-Wesley, 1996.

[4] Bull, A. O. *The effects of location and other attributes on the price of products which are place-sensitive in demand*. PhD thesis, Griffith University, 1998.

[5] Feenstra, R. C. Exact hedonic price indexes. *The Review of Economics and Statistics 77*, 4 (Nov. 1995), 634–653.

[6] Franklin, S. E., and Wulder, M. A. Remote sensing methods in medium spatial resolution satellite data land cover classification of large areas. *Progress in Physical Geography 26* (Mar. 2002), 173–205.

[7] Fukuda, S., and Hirosawa, H. Support vector machine classification of land cover: Application topolarimetric SAR data. In *Proceedings of the IEEE 2001 International Geoscience and Remote Sensing Symposium (IGARSS 2001)* (2001), pp. 187–189.

[8] Ghose, A., and Ipeirotis, P. G. Designing novel review ranking systems: Predicting the usefulness and impact of reviews. In *Proceedings of the 9th International Conference on Electronic Commerce (ICEC 2007)* (2007), pp. 303–310.

[9] Lancaster, K. J. A new approach to consumer theory. *The Journal of Political Economy 74*, 2 (Apr. 1966), 132–157.

[10] Lancaster, K. J. *Consumer Demand: A New Approach*. Columbia University Press, 1971.

[11] Li, J., and Narayanan, R. M. Integrated spectral and spatial information mining in remote sensing imagery. *IEEE Transactions on Geoscience and Remote Sensing 42*, 3 (Mar. 2004), 673–685.

[12] Lu, D., and Weng, Q. A survey of image classification methods and techniques for improving classification performance. *International Journal of Remote Sensing 28*, 5 (Jan. 2007), 832–870.

[13] Manjunath, B. S., and Ma, W.-Y. Texture features for browsing and retrieval of image data. *IEEE Transactions on Pattern Analysis and Machine Intelligence 18*, 8 (Aug. 1996), 837–842.

[14] Pal, M., and Mather, P. M. An assessment of the effectiveness of decision tree methods for land cover classification. *Remote Sensing of Environment 86*, 4 (Aug. 2003), 554–565.

[15] Pal, M., and Mather, P. M. Assessment of the effectiveness of support vector machines for hyperspectral data. *Future Generation Computer Systems 20*, 7 (Oct. 2004), 1215–1225.

[16] Plümper, T., and Troeger, V. E. Efficient estimation of time-invariant and rarely changing variables in finite sample panel analyses with unit fixed effects. *Political Analysis 15*, 2 (2007), 124–139.

[17] Rosen, S. Hedonic prices and implicit markets: Product differentiation in pure competition. *The Journal of Political Economy 82*, 1 (Jan.-Feb. 1974), 34–55.

[18] Stuckens, J., Coppin, P. R., and Bauer, M. E. Integrating contextual information with per-pixel classification for improved land cover classification. *Remote Sensing of Environment 71* (Mar. 2000), 282–296.

[19] Vee, E., Srivastava, U., Shanmugasundaram, J., Bhat, P., and Amer-Yahia, S. Efficient computation of diverse query results. In *Proceedings of the 24th IEEE International Conference on Data Engineering (ICDE 2008)* (2008), pp. 228–236.

[20] White, P. J. Site and situation determinants of hotel room rates. Master's thesis, University of Arizona, 2000.