

The Design and Evaluation of a Virtual Trusted Advisor

by

Stanley W. Cheung


Submitted to the Department of Electrical Engineering and Computer Science
in Partial Fulfillment of the Requirements for the Degrees of
Bachelor of Science in Computer Science and Engineering
and Master of Engineering in Electrical Engineering and Computer Science

at

Massachusetts Institute of Technology

May 24, 2002

Copyright 2002 Stanley W. Cheung. All rights reserved.



The author hereby grants to M.I.T. permission to reproduce and distribute publicly paper
and electronic copies of this thesis and to grant others the right to do so.

Author _____

Stanley W. Cheung

Department of Electrical Engineering and Computer Science

May 24, 2002

Certified by _____

Glen L. Urban

Thesis Supervisor

David Austin Professor of Management

Dean Emeritus, Sloan School of Management

Accepted by _____

Arthur C. Smith

Chairman, Department Committee on Graduate Theses

The Design and Evaluation of a Virtual Trusted Advisor

by

Stanley W. Cheung

Submitted to the
Department of Electrical Engineering and Computer Science

May 24, 2002

In Partial Fulfillment of the Requirements for the Degree of
Bachelor of Science in Computer Science and Engineering
and Master of Engineering in Electrical Engineering and Computer Science

ABSTRACT

Internet users may be overwhelmed by the vast amount of information available on numerous shopping sites. To help them select the best product that suits their needs, we propose to use a virtual trusted advisor to assist online users in making shopping decisions. This document details the design and evaluation of a software system, called Trucktown, which implements this idea using pickup trucks as an application. After analyzing a large number of truck owners' experiences on the system, major unmet needs in the US pickup truck market are identified. The clustering procedure which enables this analysis is also validated by simulation testing. Furthermore, a prototype of a "talking-head" version of the trusted advisor is built to investigate how different modes of user interactions affect trust. This system represents a revolutionary method for new product development on the Internet.

Thesis Supervisor: Glen L. Urban

Title: David Austin Professor of Management, Sloan School of Management

Acknowledgement

First of all, I would like to thank Professor Glen Urban for his valuable guidance, his advice and insights into this interesting project over the past year. I learned a lot, especially the managerial aspects of the project, from him. His confidence and patience in me is greatly appreciated.

I would like to extend my gratitude to all the team members – Iakov Bart, Thomas Cheng and Chris McCall for giving me such an enjoyable experience here. Special thanks to Iakov for his work in some of the statistical analysis work for this project.

I am grateful to the generous support of General Motors and the staff members during the visit in the MIT E-Business Conference on April 2002.

Finally, I would like to thank my Mom and Alice for giving me the support, the care and the love over the years. Today, I can proudly tell them, “I made it!”

Contents

Chapter 1: Introduction	6
1.1 Background	6
1.2 Project Plan	7
1.3 Organization	7
Chapter 2: Previous Works	8
2.1 Contributions of team members	8
2.2 Trusted Advisor	9
2.3 Apriori probability	9
2.4 Bayesian Attribute Effects	10
2.5 Recommendations	11
2.6 Updating truck database	11
2.7 AutoChoiceAdvisor.com	12
Chapter 3: Identifying Customers' Unmet Needs: Empirical Results	13
3.1 Introduction	13
3.2 Research Design	13
3.2.1 Virtual Advisor	13
3.2.2 Virtual Engineer	14
3.2.3 Design Pallet	16
3.2.4 Need Opportunity Criteria	16
3.2.5 Identifying Unmet Need	18
3.2.6 Virtual Engineer Questions	20
3.2.7 Determining Question Order – Two-Step Bayesian Look Ahead	21
3.2.8 Field Test	22
3.3 Results and Analysis	22
3.3.1 Users Profile	22
3.3.2 Clustering Analysis	23
3.3.2.1 Procedure	23
3.3.2.2 Results	24
3.3.3 Virtual Engineer	25
3.3.4 Design Pallet	28
3.3.5 Simulated Market Potential	33

Chapter 4: Simulation Testing	35
4.1 Motivation	35
4.2 Methodology	35
4.2.1 Profiles	35
4.2.2 Errors	36
4.2.3 Simulation Plan	37
4.2.4 Clustering Procedure	38
4.3 Results and Discussion	38
4.3.1 Sample Run	38
4.3.2 Correct Classification Rate	40
4.3.3 Existence of Small or Uninterpretable Clusters	41
4.3.4 Relations with Real User Data	42
Chapter 5: Future User Interface – Talking-Head Trusted Advisor	43
5.1 Motivation	43
5.2 Evolving Trucktown	43
5.2.1 Static Picture Persona	44
5.2.2 Talking-Head Trusted Advisor	45
5.2.3 Free-Format Text Input	46
5.2.4 Future Possibilities	47
5.3 Technology	47
Chapter 6: Discussion	49
6.1 Using Slide Bar as Input Method	49
6.2 LifeFX Player	50
6.3 Quality of Text-to-Speech Engine	51
Chapter 7: Conclusion	52
Reference	53
Appendix 1: Profile Descriptions for Simulation Testing	54
Appendix 2: Results for Simulation Testing	57

Chapter 1: Introduction

1.1 Background

As technology improves in a much faster rate than in the past, many new products are making their way into the daily lives of many people. Unfortunately, most people do not have time to understand all the features and benefits of these new products. To assist them in making purchasing decisions, many people will choose to consult salespersons of that product or to simply go for what they see in the mass media. However, salespersons may be driven by commission and be biased in giving advice. Going for what most other people buy generally does not satisfy a person's particular need.

Many people thus turn to online materials to help them make better decisions. Lots of shopping sites emerge to help users narrow down their searches. However, many of these sites assume that users understand the technical details of the product, which often lead to frustration. For example, a user who does not know what the word "resolution" means probably will not have much success finding the most suitable digital camera in normal shopping sites.

Professor Urban proposes that shopping sites can incorporate a virtual trusted advisor to assist users in making shopping decisions [1]. A virtual advisor is a piece of software which asks users simple questions to reveal their true preferences and uses utility-based algorithms to provide the most suitable recommendations. The advisor has visualizations and other trust cues to establish trust with the users. A prototype of this software, called Trucktown, has been developed and improved over the past five years to implement this idea [3][4][5].

1.2 Project Plan

The research team's objective is to extend the work by previous team members. Specifically, a field test on the Trucktown program will be carried out to gather real user data. Statistical analysis including a clustering procedure will be applied to the data to discover any major unmet needs in the US pickup truck market. A simulation testing will be performed in order to validate our clustering procedure. We would also investigate whether users prefer this idea of having a virtual advisor to give them shopping advice. As a new initiative of the project, a talking-head 3D-animation model will be used to represent the trusted advisor. Our goal is to find out how different kind of user-advisor interactions will affect trust.

1.3 Organization

This document is organized into seven chapters. Chapter 2 describes the past work by previous team members. Chapter 3 details the design of the Trucktown system which includes three main components: the virtual advisor, the virtual engineer and the design pallet. This chapter also presents the statistical analysis results of the user data collected on July 2001. The simulation testing to validate our clustering procedure is presented in Chapter 4. Chapter 5 describes the implementation of a talking-head 3D-model of a human face to replace the static picture representation of the trusted advisor. Practical issues encountered in the project and other discussion will be presented in Chapter 6. Lastly, the conclusion in Chapter 7 summarizes the implication of the project.

Chapter 2: Previous Works

2.1 Contributions of team members

Professor Urban et al. [2] [3] suggested a novel approach of using a virtual advisor to assist online customers to make purchasing decisions in 1996. The theories were implemented in a software prototype called “Trucktown” [3] [4] [5]. Trucktown assists potential pickup truck buyers to identify the most suitable truck for them. Since General Motors (GM) is the main sponsor of the research project, pickup trucks are chosen as the application, but the theories are general and applicable to other products. Since the beginning of the project, the methodology has evolved from our team’s weekly meetings. In addition, each member contributed in different ways. Table 1 is the summary of the members’ main contributions.

Team member	Main contributions	Participation period
Glen Urban	Started and supervised the project	Since 96
Stanley Cheung	Collected and analyzed field test data. Updated two-step bayesian look-ahead algorithm. Performed simulation testing on clustering procedure. Developed talking-head prototype.	Feb 01 – Jul 02
Iakov Bart	Performed statistical analysis on field test data.	Jul 01 – Aug 02
Thomas Cheng	Programmed and integrated the virtual engineer and the design pallet. Designed the system architecture and the clustering process.	Feb 99 - May 01
Chris Mann	Modeled the virtual engineer and dialog	Sep 99 - Jan 00
Hunter Chen	Made cost models of the design pallet	Feb 00 - May 01
James Ryan	Modeled the Bayesian probabilities of 2001 trucks model	Sep 00 - May 01
Brian Chan	Made 3D truck models for the design pallet	Oct 00 - May 01

Table 1: Team members and summary of their main contributions

2.2 Trusted Advisor

The recommendations of the virtual advisor are useful only if the customers trust the advisor as a dependable and knowledgeable expert. Hence, the first task for the virtual advisor is to gain the customers' trust, which can be accomplished by providing various "trust cues" [2][3]. One of the examples of trust cues is to declare that the virtual advisor does not receive any money from manufacturers, so that it gives unbiased recommendations.

2.3 Apriori probability

The apriori probability of a product is the initial probability of recommending the product to the customer over others. The calculation of the apriori probabilities is based on knowledge of the truck utility rankings from previous third party marketing research, plus customer's preferences of a set of pre-defined perceptual dimensions of the product. Customers indicate the preferences by answering a constant-sum preference question, in which they assign the relative importance of the perceptual dimensions.

The a-priori probabilities are calculated using the following formula [1]:

$$P(A_a) = \frac{\exp(\sum_d \alpha_d D_{d,a})}{\sum_a \exp(\sum_d \alpha_d D_{d,a})} \quad \text{Eq. 1}$$

$P(A_a)$ = Probability of purchase of product alternative a ,

i.e., the a-priori probability

α_d = constant sum importance of dimension d given by customer, and

$D_{d,a}$ = standardized database value of dimension d for alternative a

i.e., the values stored in the utility knowledge database.

It has been acknowledged that questions about constant-sum preferences are difficult to answer [6]. Default values for the dimensions can be used until the answers to the questions about constant-sum preferences are collected. Hence, these questions can be asked at any time during the conversation. The trusted advisor could establish a rapport with the customers before asking these difficult questions. However, in Trucktown, the question about constant-sum preferences is always asked first, so that Bayesian ordering of questions is more effective. There are five dimensions in Trucktown: price, performance, fuel economy, reliability, and safety.

2.4 Bayesian Attribute Effects

After the customer has answered a question, the Bayesian probabilities of every product in the database will be updated using the following equation [1]:

$$P(A_a) = P(A_a | R_{r,q}) = \frac{P(A_a)P(R_{r,q} | A_a)}{P(R_{r,q})} \quad \text{Eq. 2}$$

- $P(A_a | R_{r,q})$ = the conditional probability that the customer will purchase the product a given that he answered with response r to question q ;
- $P(R_{r,q} | A_a)$ = the conditional probability of answering question q with response r given that the purchased alternative is a (these data are stored in the knowledge database of Bayesian Probabilities);
- $P(A_a)$ = the prior probability of purchase of product alternative a before the question q is answered;
- $P(R_{r,q})$ = the marginal probability that a customer answers question q with the response r ,

$$P(R_{r,q}) = \sum_a P(A_a)P(R_{r,q} | A_a) \quad \text{Eq. 2a; and}$$

$P'(A_a)$ = the new probability of purchasing the product alternative a .

The probability can be iteratively updated across all responses to questions. The probabilities are calculated whenever the customer answers a question. The list of products is sorted according to the updated probabilities. Products with the highest probabilities will be recommended to the customers.

2.5 Recommendations

The virtual advisor decides the recommendations in two ways. First, three products with the top three Bayesian probabilities are recommended. If the customer's market segment is identified, one truck based on the segment analysis are recommended. Otherwise, the virtual advisor would recommend the next product with the highest Bayesian probability.

2.6 Updating truck database

To prepare for the field test in July 2001, the trucks in the database have been updated to the 2001 models. We decided to follow the definitions of *Autosite.com*, an authoritative third-party website for autos. Most of the trucks' specifications are from the manufacturers' websites, with some customers' ratings from *Edmunds.com* and the Consumer Reports, *ConsumerReports.org*.

2.7 AutoChoiceAdvisor.com

GM has developed and launched their own site www.AutoChoiceAdvisor.com on summer 2001 to give advice to people on auto purchasing. The site takes a similar approach as Trucktown as it emphasizes the unbiasedness of the information and recommendation. However, it is different from Trucktown because it does not use any form of persona to represent the advisor. The site also covers all types of autos sold in US, in addition to only pickup trucks. The site obtains auto specifications from JDPowers.com

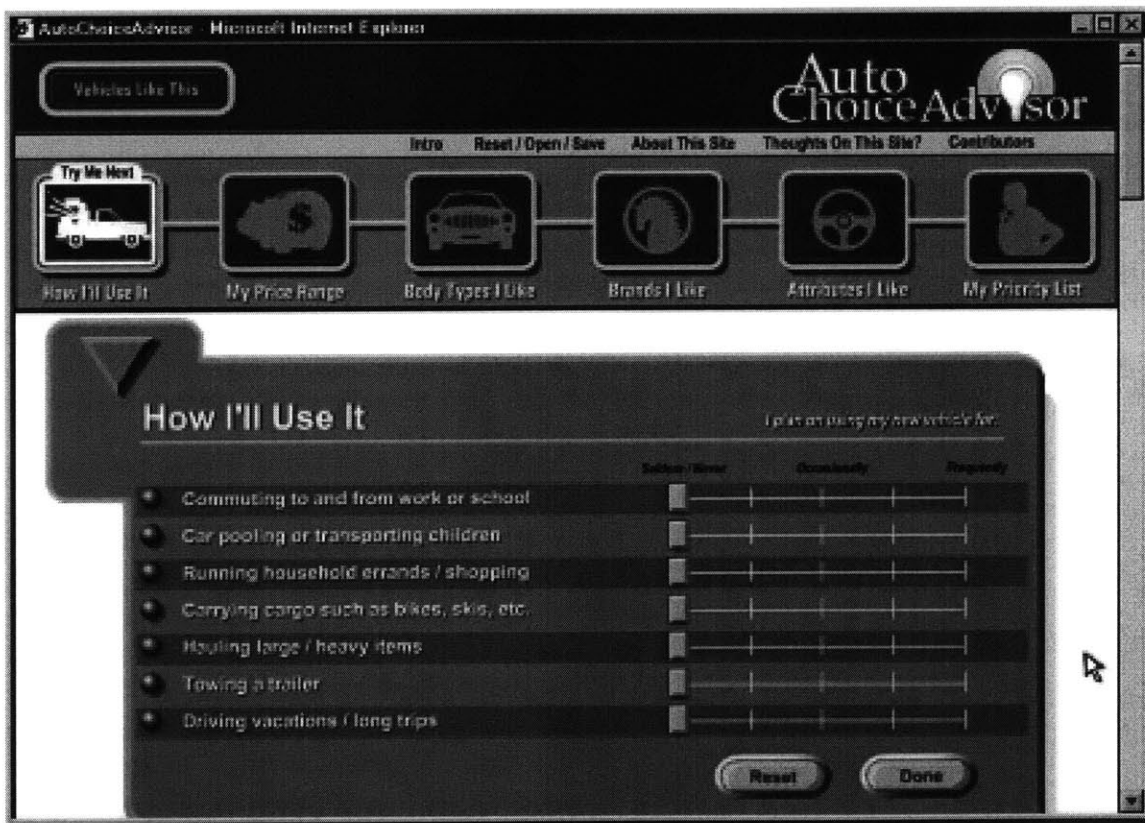


Figure 1: AutoChoiceAdvisor.com

Chapter 3: Identifying Customers' Unmet Needs: Empirical Results

3.1 Introduction

Identifying customers' unmet needs are very important to companies. By designing new products to explore those opportunities, company can generate new sources of revenue. Therefore, we need measures to identify such unmet needs. This chapter details the design and evaluation of the Trucktown system, which enables companies to discover such market opportunities.

Trucktown consists of three main components: the Virtual Advisor inquires about a user's preference and makes appropriate recommendations, the Virtual Engineer identifies inconsistent responses and tries to resolve the conflict, and the Design Pallet allows the users to design their own dream trucks.

An extensive market research analysis has been carried out based on 1092 users' experiences on the Trucktown system, and the major results are included in this chapter.

3.2 Research Design

3.2.1 Virtual Advisor

The virtual advisor inquires about a user's preference on certain attributes of a pickup truck, such as the engine type, drivetrain and bed length, etc. The system contains a database of priori probabilities for each truck and a constant-sum preference question is used to determine the initial utility of each truck. The dialog between the advisor and the user is conducted in a Q&A fashion, during which each truck's utility is updated after each question. The question order is determined by a 2-step bayesian look-ahead algorithm which finds the question that carries the most amount of information. After all

fifteen questions have been answered, the top three trucks with the highest utilities are recommended to the users. Here is a sample screen shot in the virtual advisor session:

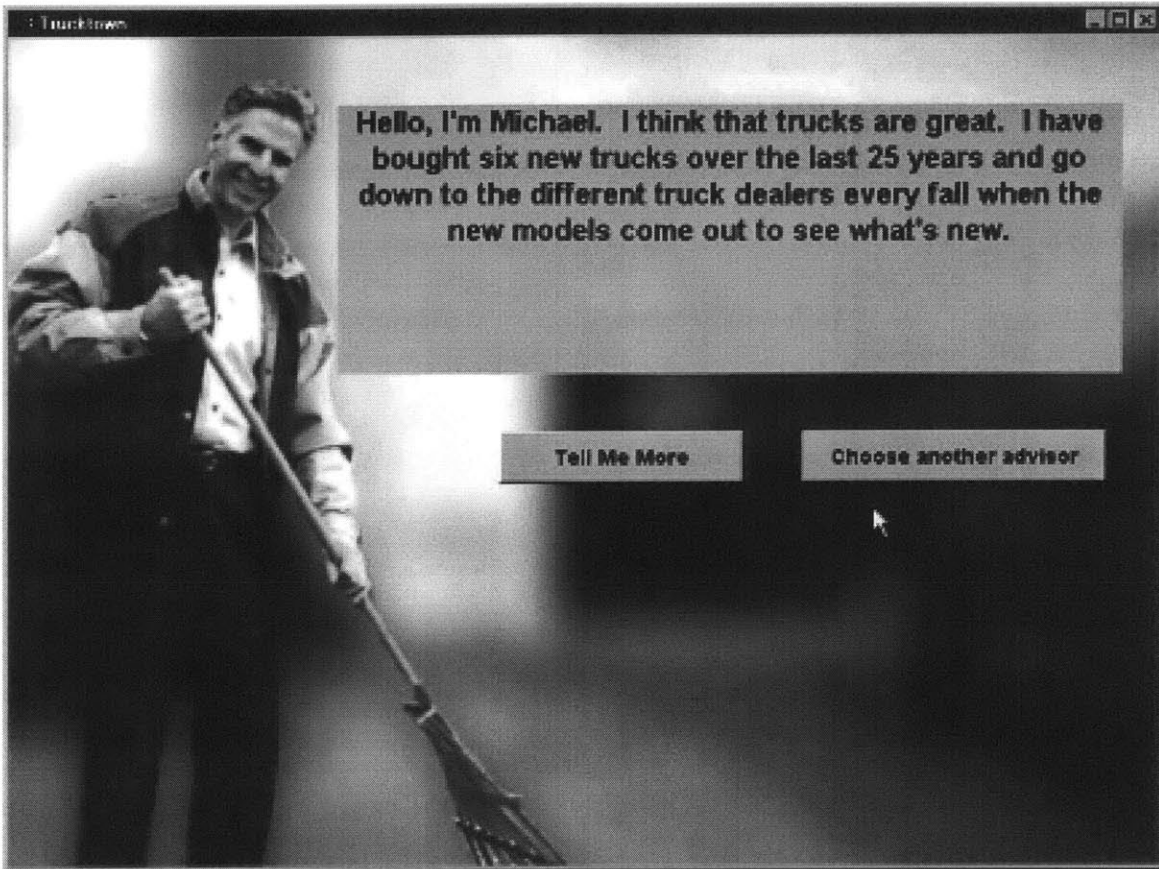


Figure 2: The Virtual Advisor

3.2.2 Virtual Engineer

During the virtual advisor session, a user may provide inconsistent answers in the sense that no trucks in the current market satisfies all the user's needs. An engineer, having "listened in" to the conversation between the user and the advisor, comes in and tries to resolve the conflict. The engineer determines a conflicting pair of attributes which causes an utility drop in the most preferred truck during the virtual advisor session.

The algorithm for determining the conflict pair will be explained in more details in Section 3.2.4.

For example, a user indicates that he would like a compact truck which can tow and haul. No such truck in the market can satisfy the user's need and therefore the engineer comes in and asks why the user wants a compact truck and what exactly the user is planning to tow and haul. This information will be helpful to auto manufacturers to design the next generation of trucks which can cater to this segment of the market.




Figure 3: The Virtual Engineer

3.2.3 Design Pallet

The user can freely design his/her own dream truck in the design pallet, varying attributes such as engines, cab dimension and wheel drive, etc. The image and the parameters of the truck such as price, fuel economy and towing capacities, etc will be updated whenever corresponding attributes are modified.

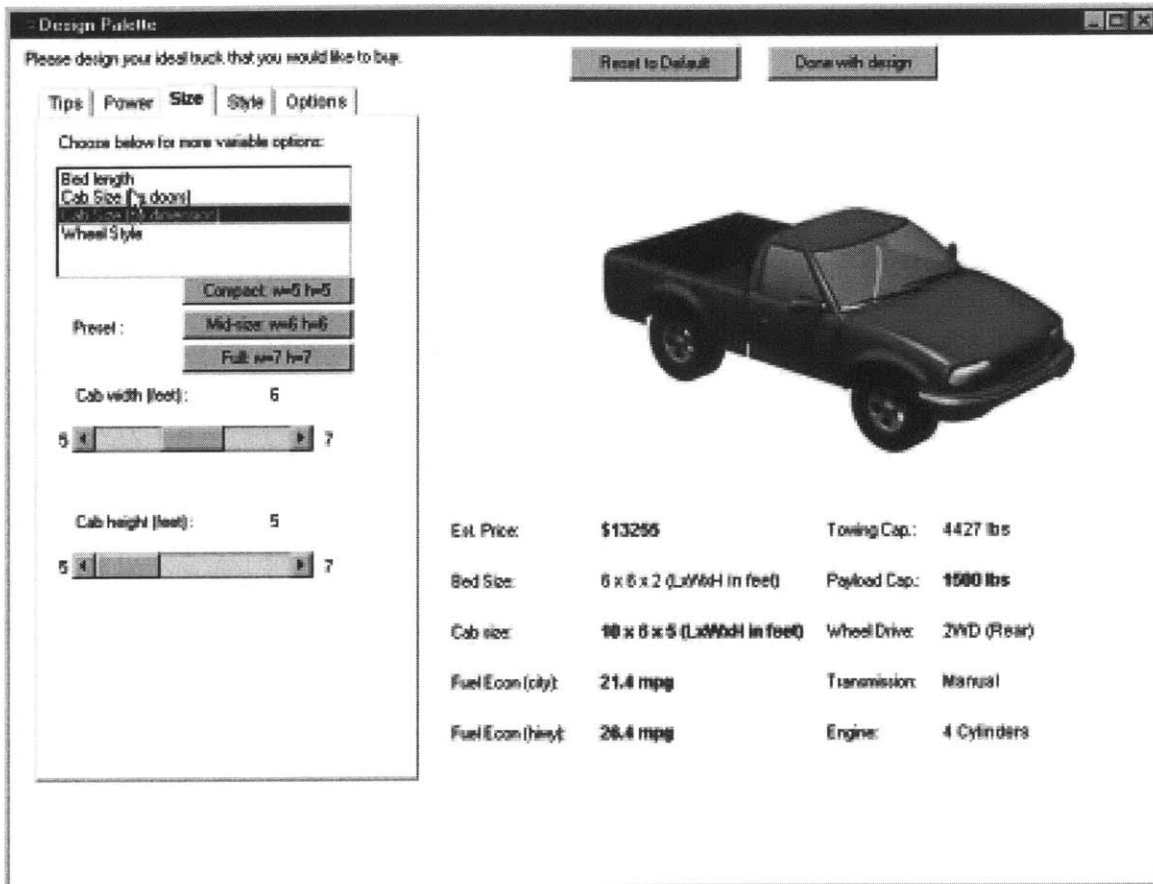


Figure 4: The Design Pallet

3.2.4 Need Opportunity Criteria

After each question in the virtual advisor dialog, each truck in the database will be updated with an utility which describes the likelihood that the user will prefer that truck. Therefore, the current best truck and the current highest utility can be determined after

each question by sorting the utilities. In normal cases, we would expect that the current highest utility increases monotonically and the best truck converges to one which satisfies all of the user's preferences, like the following example:

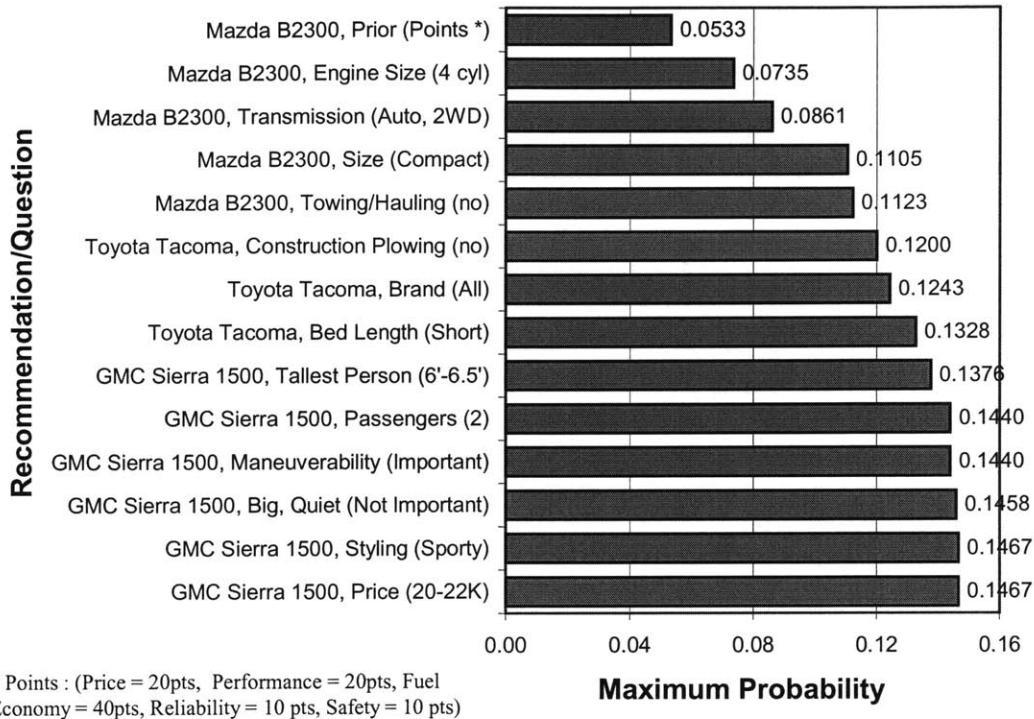


Figure 5: An Utility Profile where highest utility is monotonically increasing

However, sometimes the answers provided by the user may reveal inconsistencies. For example, a user may have already indicated that he would like a compact truck which can drive on icy road. The current best truck is Mazda B2300. Now suppose that the advisor asks the user about the normal usage of the truck and the user replies that he would like to tow a boat, the highest utility will drop because no existing truck satisfies all the needs. Note that the highest utility truck may not be the same before and after the question. Here is a sample utility profile to illustrate such utility drop:

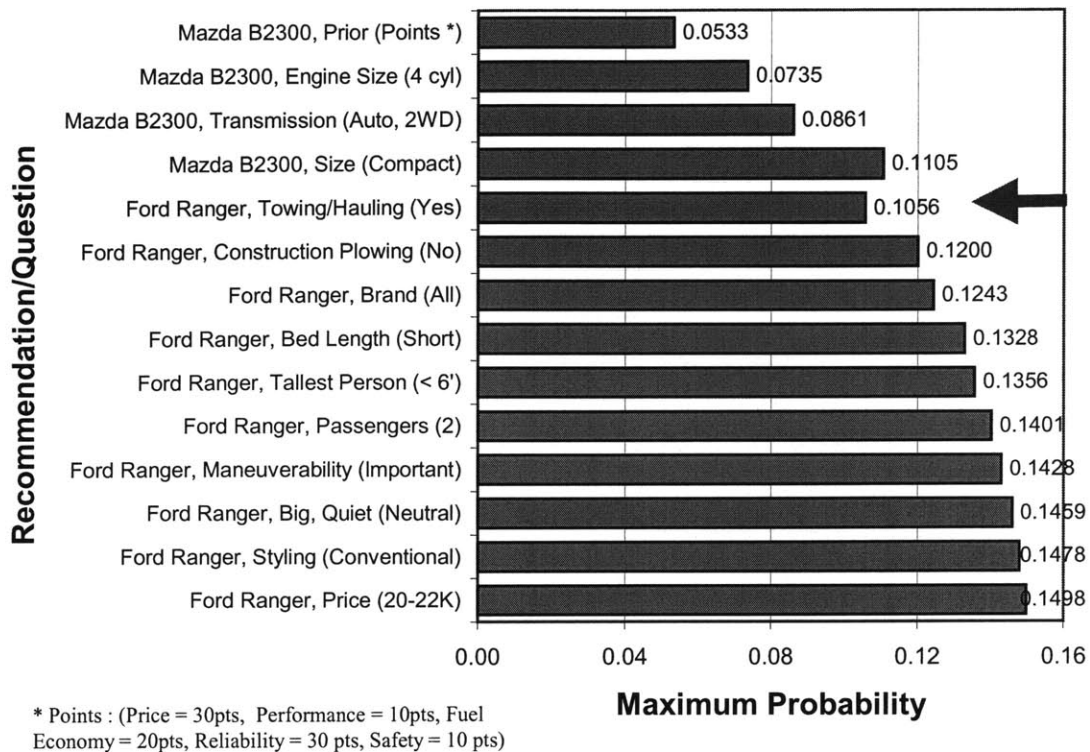


Figure 6: A Utility Profile with a Utility Drop

Such utility drop is worth further investigating and therefore, we define a drop in the highest utility to be the criteria of identifying an unmet need.

3.2.5 Identifying Unmet Need

After we have detected a drop in the highest utility, we determine the user's actual unmet needs by looking at every pair of answers the user provided to see which of those pairs causes the drop. We look at every pairing of the answers because the utility drop may not be necessarily caused by the last question asked. In the priori probability database, each truck is associated with a conditional probability for each possible answers in the virtual advisor dialog. The correlation among each column (each answer) therefore

provides a good indication of which pair of answers corresponds to a conflict. For example, the correlation between the answers “driving offroad” and “two wheel-drive” is -1.0 , indicating that no existing two wheel drive pickup truck is good at driving offroad and this defines a particular conflict.

	A	B	C	D	E	F	G	
1		compact	full	Manual	Auto	2WD_yes	4WD_yes	C
2	compact	1.00						
3	full	-1.00	1.00					
4	Manual	0.01	-0.01	1.00				
5	Auto	-0.01	0.01	-1.00	1.00			
6	2WD_yes	0.05	-0.05	0.14	-0.14	1.00		
7	4WD_yes	-0.05	0.05	-0.14	0.14	-1.00	1.00	
8	Off_road_yes	-0.05	0.05	-0.14	0.14	-1.00	1.00	
9	Towing_yes	-0.66	0.66	-0.01	0.01	-0.09	0.09	
10	Hauling_yes	-0.54	0.54	0.14	-0.14	0.05	-0.05	
11	Construction_yes	-0.71	0.71	-0.14	0.14	-0.07	0.07	
12	Plowing_Snow_ye	-0.72	0.72	-0.10	0.10	-0.73	0.73	
13	body_entry_pickup	0.17	0.17	0.33	0.33	0.08	0.08	

Figure 7: Correlation Table

We define a pair of answers to be a conflict if their correlation is less than -0.3 , which means that it is unlikely that a truck can satisfy both answers. Note that one particular conflicting pair of answers may not cause an overall utility drop. However, in the other way round, an overall utility drop must be caused by some number of conflicts.

In the current truck database, there are a total of 208 conflict pairs. For each user after the virtual advisor session, a vector of 208 entries is logged, with each entry corresponding to a particular conflict. Most entries in this vector will be zero. If a user has a particular conflict, the corresponding entry will contain the negative correlation of

the conflict. If a user does not incur an utility drop in the whole advisor session, all entries in the vector will be zero.

3.2.6 Virtual Engineer Questions

Having determined the conflicts that cause an utility drop, the virtual engineer comes in and tries to clarify why the user has such a conflict. In this session, the engineer will ask the user for more information about why he wants those features. The virtual engineer will pick the 3 conflicts (if there are more than 3 conflicts) with the largest negative correlation (most negative) and ask a maximum of 6 questions. Here is a sample question the virtual engineer may ask:




Figure 8: A Sample Virtual Engineer Question

3.2.7 Determining Question Order – Two-Step Bayesian Look Ahead

To select the next question, the virtual advisor attempts to gain as much information as possible from the customer. For example, if, after looking ahead at all the responses, the advisor decides that a question on engine type is likely to make one truck highly preferable and all other trucks less preferable, then that question would be a good candidate to ask next. In order to compute the expected information, we need to take the expectation over all possible responses to each question and over all possible trucks. To improve the question ranking, we can look-ahead two questions that have not been asked instead of one. Here is the algorithm we have implemented:

For each question q , which has not been asked yet

For each response $z \in Q_q$

For each question $p \neq q$ AND p has not been asked yet

$$f(p, q, z) = \sum_a \sum_{s \in Q_p} P(A_a | R_{sp}, R_{zq}) * \log \frac{P(A_a | R_{sp}, R_{zq})}{P(A_a)} \quad \text{Eq. 3a}$$

$$p'(z) = \text{argmax } f(p, q, z)$$

$$f(q) = \sum_a \sum_{z \in Q_q} \left(\sum_a \sum_{s \in Q_{p'(z)}} P(A_a | R_{sp'(z)}, R_{zq}) * \log \frac{P(A_a | R_{sp'(z)}, R_{zq})}{P(A_a)} \right) \quad \text{Eq. 3b}$$

$$q' = \text{argmax } f(q)$$

where Q_q = set of responses for question q

R_{sp} = the response of s to question p

A_a = the probability that the user will prefer truck a

The question, q' , with the highest ranking from all the unanswered question will be picked as the next question in the dialog.

3.2.8 Field Test

The Trucktown program has been hosted in the Harris Interactive website, www.HarrisInteractive.com on July 2001. The truck database is updated with all the current models in the market. 10,000 emails have been sent out to invite pickup-truck owners across the US to go through the Trucktown exercise including the Virtual Advisor, Virtual Engineer and Design Pallet session for 30 minutes. Users are also required to fill out an exit survey to provide feedback on their demographics and specific opinions on the Trucktown program. Each user who successfully completed the whole process is awarded with \$20 incentive.

1092 users have completed the exercise and all their responses including the conflict vector are logged. Extensive statistical analysis has been conducted on the data and the results.

3.3 Results and Analysis

3.3.1 Users Profile

Here are some of the top-line characteristics of the sample pool who have successfully completed the Trucktown program and the exit survey.

- 73% of the respondents are male.
- 44% have better internet connection (cable modem, T1, DSL) than a dial-up.
- Users come from 46 different states while 20% of them come from Texas.
- 40% consider himself an expert with respect to pickup trucks.
- 55% keep up-to-date to the newest model and trend in the pickup market.

- 66% trust the Trucktown recommendation process, compared to 8% preferring traditional dealership.
- 82% would utilize an online program to advise them on buying vehicles.
- 56% would consider buying a vehicle online.
- 44% chooses the garage mechanic as their advisor while 30% chooses the editor for consumer reports magazine and 26% chooses the neighbor.

3.3.2 Clustering Analysis

3.3.2.1 Procedure

As mentioned in Section 3.2.5, all the possible conflicts of a user will be logged if there is an utility drop in the virtual advisor session. Therefore, after the field test, we have gathered a database consisting of 1092 rows (one per user) and 208 columns (one per conflict). To identify what the major unmet needs in the market are, intuitively, we should try to look for large groups of people who have similar conflicts. To achieve this goal, we cluster the users into clusters, such that users within a cluster share similar conflicts.

We use the FASTCLUS procedure in the SAS software to do the clustering. The variables used in the clustering procedure are all the 208 conflict variables. Initially, we force a small number of clusters, i.e. setting the parameter `maxclusters` to 3. Recall that an entry in the conflict vector has the value of the negative correlation of the two attributes involved if the user has that conflict and has value zero otherwise. For each conflict, we decide whether the majority of the people who have that conflict fall into the same cluster. If it is the case, that conflict will be classified as one of the representative

conflicts of that particular cluster. Otherwise, if it is ambiguous, we will just discard it. To decide whether a conflict should be classified as representative, we observe which cluster the centroid of that conflict falls into.

After all the representative conflicts of each cluster have been determined, we decide whether those conflicts give interpretable meaning to that cluster, such as “compact, towing and hauling”. If those conflicts do not provide interpretable meaning to most of the clusters, we increment the `maxclusters` parameter by 1 and repeat the experiment. We stop until we get clusters with relatively large number of users and a coherent interpretation, and if further fine-graining of the clusters does not give significant results. Using too few clusters may collapse different clusters together, while using too many will simply subdivide already meaningful clusters. The optimal number of clusters should be somewhere in between.

3.3.2.2 Results

Here are the results from our clustering procedure:

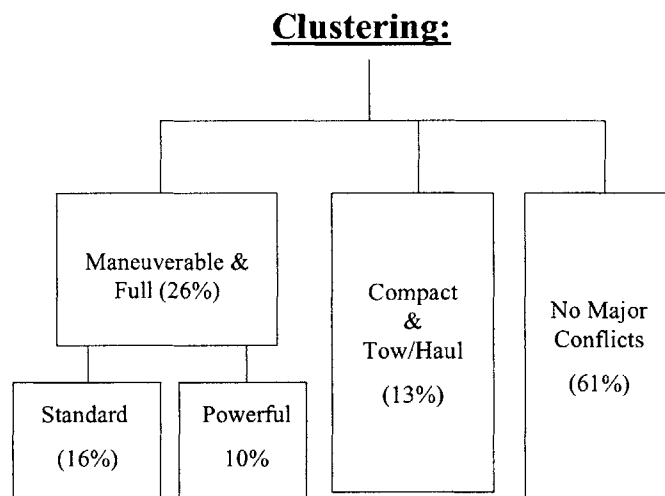


Figure 9: Cluster Result

26% of the 1092 users fall into Cluster 1. Users in Cluster 1 want to have a full-size truck which is highly maneuverable. This is an unmet need because in the current market, all full-size trucks have relatively large turning radius which make them difficult to turn tight corners and maneuver in small parking lots. This cluster can be further broken down into 2 smaller cluster. People in Cluster 1A (16%) want to have a V8 engine to go with a conventional style outlook, while people in Cluster 1B (10%) prefer a rugged and sporty styled truck with a long bed and a diesel engine.

13% of the users belong to Cluster 2. They want to have a compact truck which can tow and haul. Many of them also want a V6 engine and a short bed. In the current market, only mid-size and full-size trucks have higher towing and hauling capacities.

The remaining 61% of the user fall into a big cluster which has no easily interpretable meaning. Individuals may have small number of conflicts but no groups among them are large enough and still share common conflicts. Basically those users have random conflicts which are not representative.

3.3.3 Virtual Engineer

We first start with some statistics of the general opinion on the Virtual Engineer session:

- 88% finds the engineer's questions easy to answer.
- 75% feels that the questions relate to their needs.
- 60% feels that the engineer has asked all the necessary questions about their needs.
- 62% believes that their input will help design future vehicles.

In the previous section, we have identified major clusters of unmet needs. Why exactly do the user provide such responses which lead to conflicts? Let us look at the results from the Virtual Engineer dialog to find out why. One important point to note is that a user may not be asked the corresponding questions in the dialog even though they have certain conflicts, because the virtual engineer only chooses the 3 most important conflicts to investigate.

Cluster 1: Full-size vs. Maneuverability




Figure 10: Reasons why Users want Full-size Truck and Maneuverability

Cluster 2: Compact vs. Towing

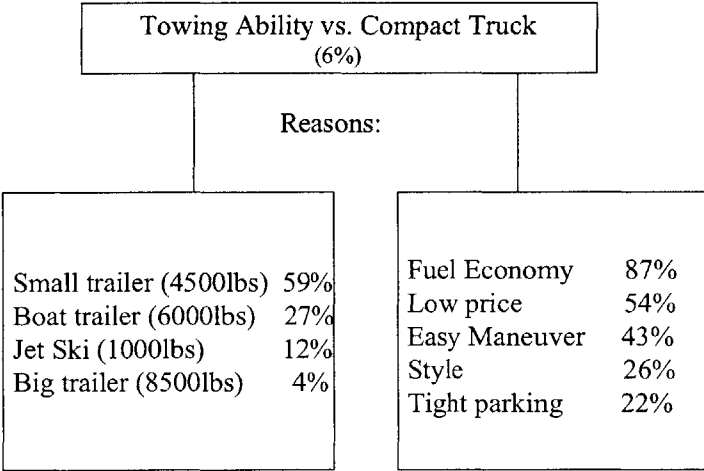


Figure 11: Reasons why Users want Compact Truck and High Towing Capacities

Cluster 2: Compact vs. Hauling

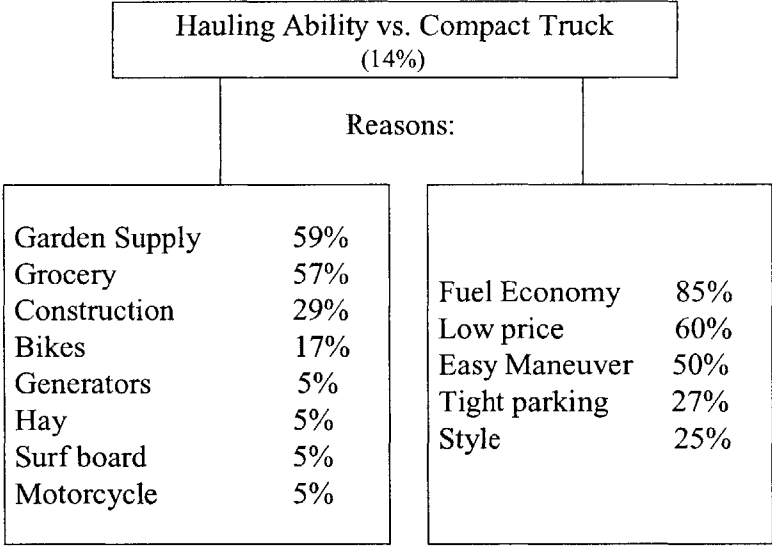


Figure 12: Reasons why Users want Compact Truck and High Hauling Capacities

From these results, we can understand better why users want certain features in their trucks. After all, compact trucks are certainly capable of doing light duty towing and hauling. By knowing what exactly and how often users need to tow and haul stuff, engineers may be able to design new generation of trucks which satisfy both needs.

3.3.4 Design Pallet

The design pallet allows the users to design their own dream trucks, aided by a visualization of the truck they have built and an estimate to various parameters of the truck. Here are some basic statistics of the general opinion on the design pallet session:

- 79% find this an enjoyable experience.
- 82% regard it as a serious exercise.
- 53% believe that their design will help design future vehicles
- 73% would buy the truck they designed.

We apply a similar clustering procedure on the user final design, using all the attributes as variables, hoping to find out the common designs. Again, we are only interested in relatively large clusters with coherent interpretation. Here are the results:

	Size	Estimated Price	Cab	Wheel Drive	Transmission
Design A	Full-Size	\$30,000	4-Door	4WD	Auto
Design B	Mid-Size	\$27,000	4-Door	4WD	Auto
Design C	Mid-Size	\$22,000	4-Door	2WD	Auto
Design D	Mid-Size	\$21,000	4-Door	2WD	Manual
Design E	Compact	\$13,000	2-Door	2WD	Manual

Table 2: Summary of User Design in Design Pallet

We can see that some users are interested in large vehicles with loaded functionality, while some others just want more basic and economical models. Here are the pictures and details of these five final designs.

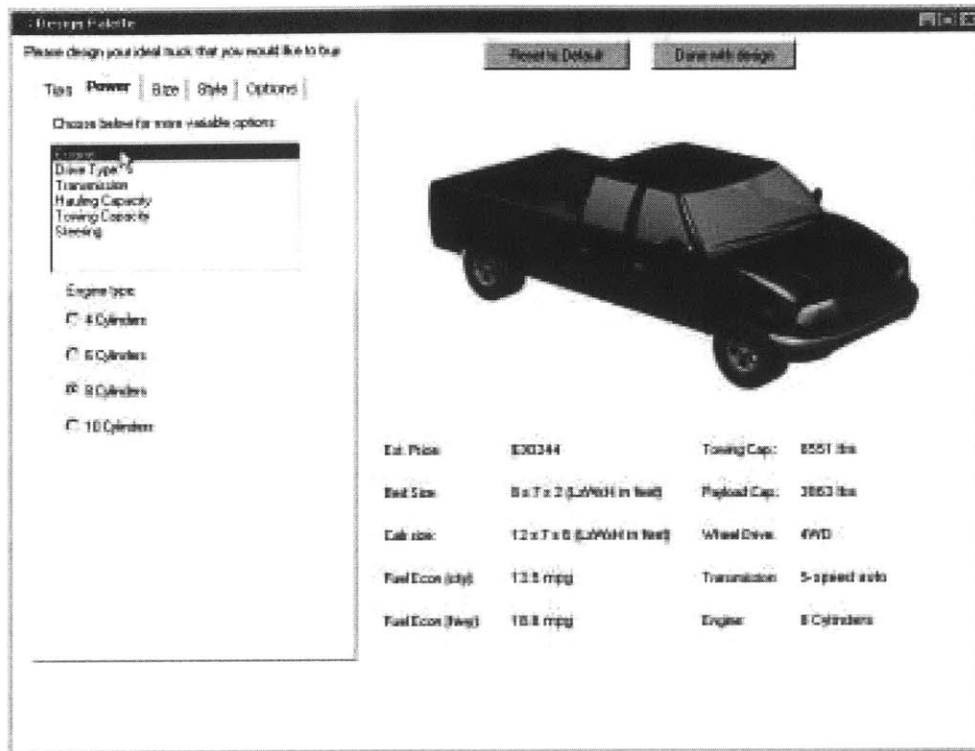


Figure 13: User Design A – Full-size and Powerful

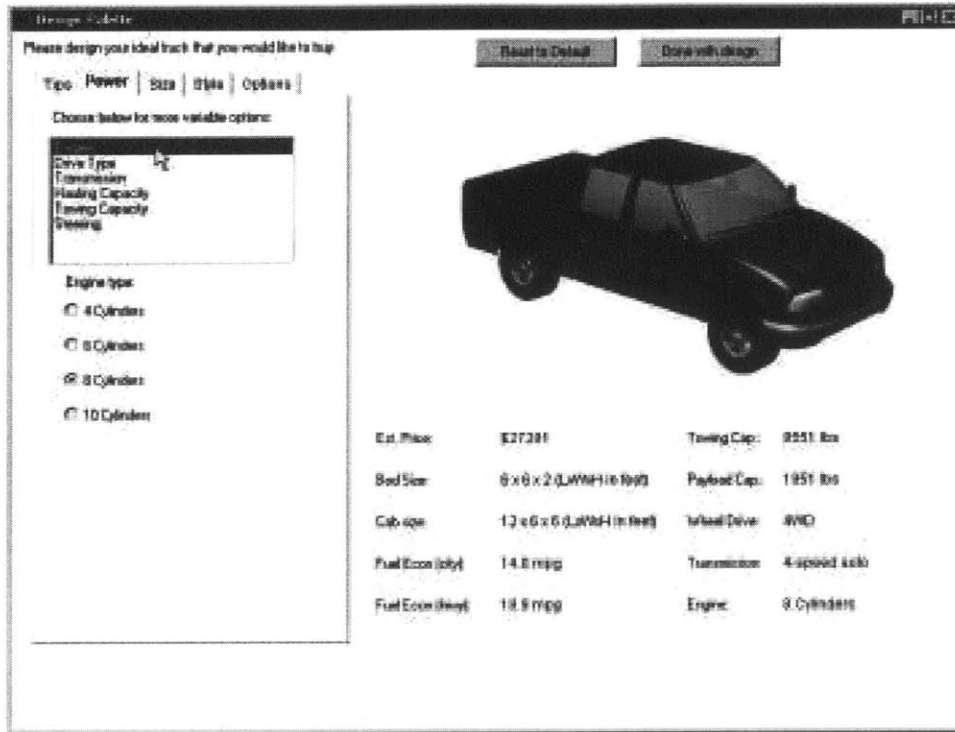


Figure 14: User Design B – Mid-Size, 4WD, loaded with options

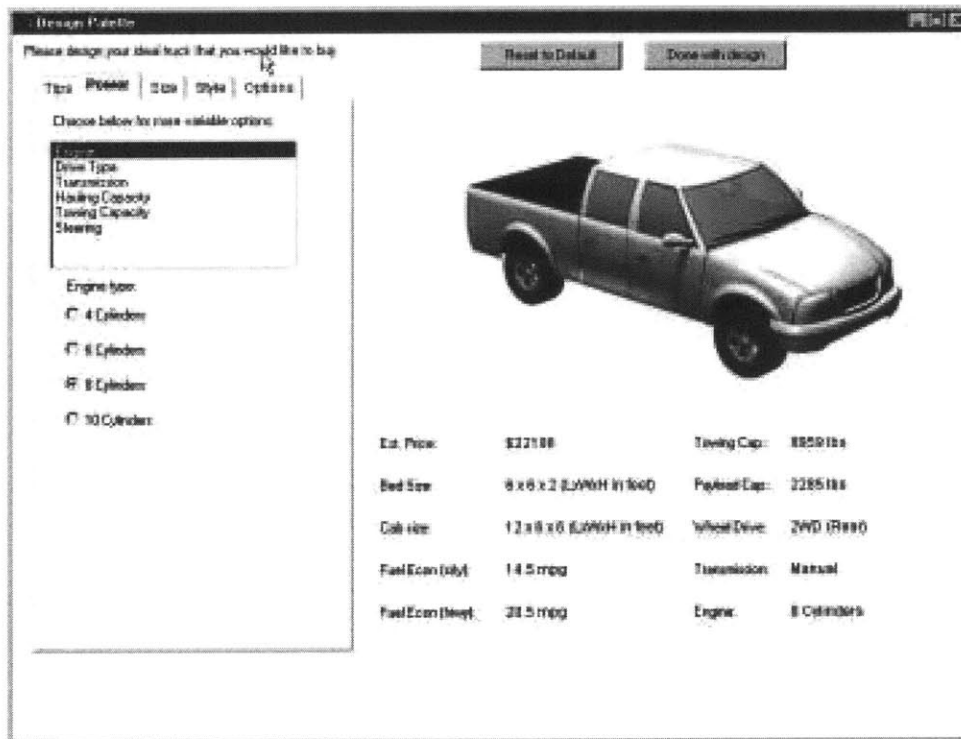


Figure 15: User Design C – Mid-Size and 2WD

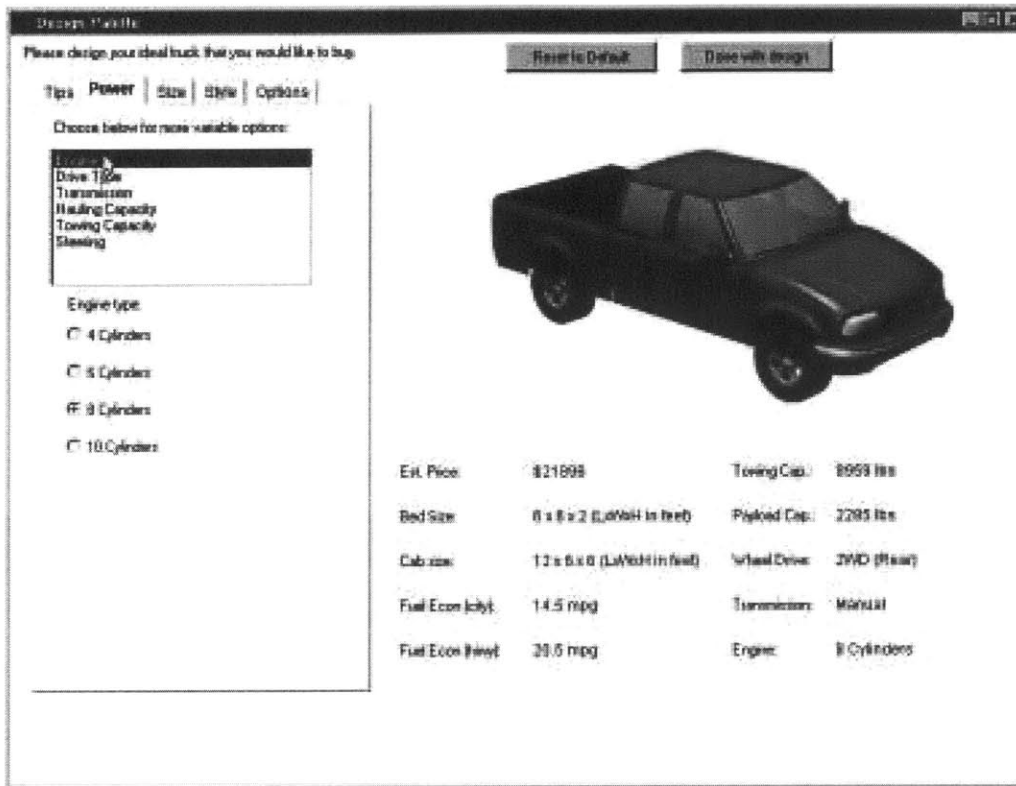


Figure 16: User Design D – Mid-Size with basic package

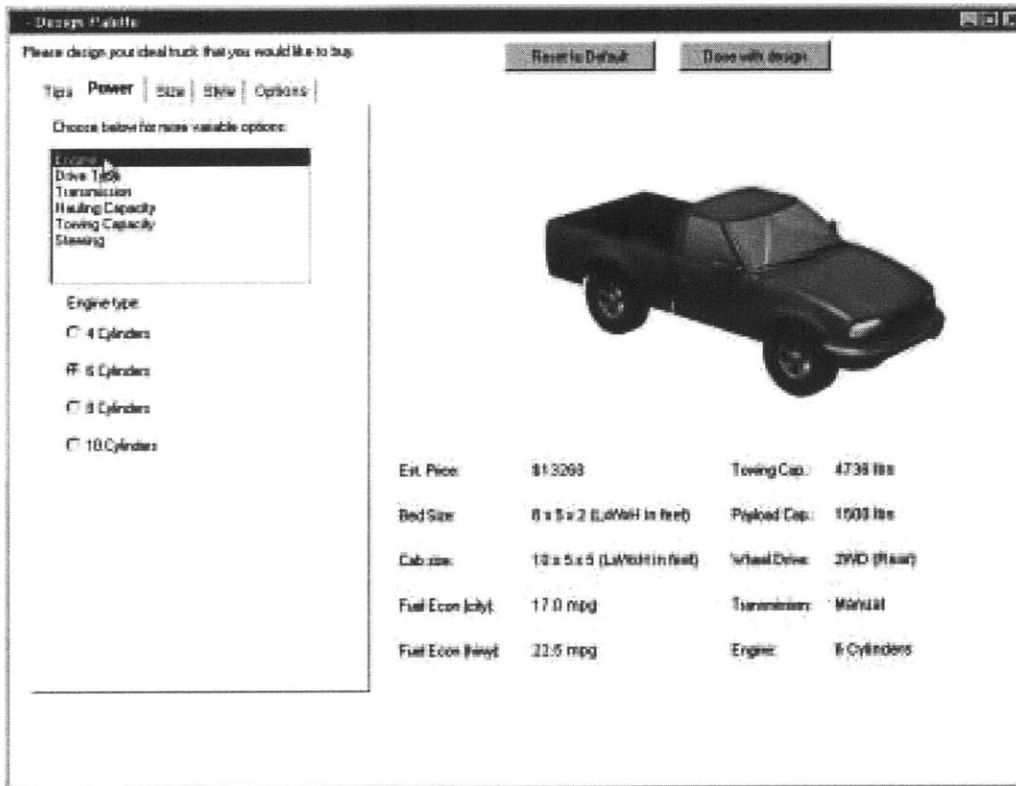


Figure 17: User Design E – Compact and Economical

In the design pallet, the user is initially given a truck which is closest to the best truck the virtual advisor has recommended. This truck may not satisfy all the user's needs. Therefore, from the final user designs, we can also investigate how the users themselves try to resolve the conflicts. To achieve this, we look at the designs by users from each of the major clusters identified in Section 3.3.2 and see what they change in their vehicles. Here are the results:

Cluster 1: Full-size and Maneuverability

- 6% change two-wheel steering to four-wheel steering
- 15% change the truck width from 6 feet to 7 feet
- 16% change the truck height from 6 feet to 7 feet

As we can see, some users are aware that adding four-wheel steering can lower the turning radius of the vehicle and thus increase maneuverability. On the other hand, some users would very much prefer to have a full size truck that they actually increase the cab dimensions.

Cluster 2: Compact, Towing and Hauling

- 21% change from hauling grocery (500 lbs) to hauling lumber (3,000 lbs)
- 2% change from hauling grocery (500 lbs) to hauling generator (4,500 lbs)
- 10% change from towing jet ski (2000 lbs) to towing a horse (4,500 lbs)
- 8% change from towing jet ski (2000 lbs) to towing a boat (6,000 lbs)

- 10% change the bed length from 6 feet to 7 feet
- 7% change the bed length from 6 feet to 8 feet

Again, we can see how users resolve their own conflicts by making their own design decisions. Many users increase the towing and hauling capacities while some lengthen the truck bed. Technically, with the added metals, the final design may not be classified as a compact truck anymore, but the idea to illustrate in this exercise is how users resolve their own conflicts.

3.3.5 Simulated Market Potential

So far, we have discovered what the major unmet need groups there are in the market. How much market potential are there for future truck designs which specifically cater for these needs?

We have designed an imaginary truck for each of the major unmet needs in Section 3.3.2 and added these trucks to our truck database. Users' responses are fed back into the system and the new recommendations are logged. We count the number of times these imaginary trucks make it to the top of the recommendation list (with highest utility). If we assume that the users would buy the best truck the virtual advisor recommends, we can estimate the percentage of the market which will buy this new imaginary truck. These are the 2 trucks which we have designed and added to the database:

Cluster 1: Full-size and Maneuverability

Adapted from a Sierra 3500 CrewCab but is as maneuverable as a compact truck. Price is \$3000 more expensive. Potential Market Share: 6.77%

Cluster 2: Compact, Towing and Hauling

Adapted from a Chevy S-10 but have higher towing and hauling capacities. Price is \$2000 more expensive. Potential Market Share: 3.89%

As a rough estimate, one percentage point of the current pickup truck market corresponds to approximately 800 million of revenue in 2001.

Chapter 4: Simulation Testing

4.1 Motivation

In the previous chapter, we have showed how major unmet need groups can be identified by clustering each user's conflict row. If there is a large group of users who have similar conflicts, we report that the most common conflicts within the group represent a major unmet need. In retrospect, how robust is this procedure? How accurate are we able to identify unmet need groups? In addition, it is inevitable that real users may give incorrect responses accidentally or not be serious in the survey. This chapter describes how we validate our clustering procedure by running the procedure on simulated users with given unmet needs in the presence of errors. We show that our clustering procedure achieves high recovery rate in classifying users into the original unmet need groups.

4.2 Methodology

4.2.1 Profiles

In our simulation, each user's response is generated from a designated "profile". Each profile represents a true unmet need. For example, Profile 1 in our simulation represents a group of users who primarily want to have a compact truck which can tow and haul. There are six of these conflict profiles – each of them represent an essentially different unmet need. Each of these profiles is associated with a set of standard answers to each of the virtual advisor questions. See Appendix 1 for the details for each profile. Here are the general descriptions for each profile:

Profile 1 – Compact truck which can tow and haul

Profile 2 – Full-size 8-cylinder truck with short bed

Profile 3 – Conventional-styled small truck with diesel engine

Profile 4 – Compact truck with extra short bed and a tall driver

Profile 5 – Small short-bed truck with a 10-cylinder engine

Profile 6 – Long-bed truck with high maneuverability

We have also included 3 additional profiles which do not generate any conflicts to show that users with no conflicts should not be confused with users with conflicts. Here are the descriptions of those profiles:

Profile 7 – A compact truck with no additional features

Profile 8 – A full-size truck with no additional features

Profile 9 – A full-size truck which can tow and haul

4.2.2 Errors

It is almost impossible that real users always give the same answers even though they have similar needs. People may not be serious in survey. They may interpret the questions differently or they simply mistype the answers. This motivates us to add variance to the simulated users' responses and to see how well our clustering procedure recovers the true clusters.

To generate responses for each user in a particular profile:

a) For the constant-sum question which sets up the prior probability, we use the individual value as mean and add x points standard deviation under normal distribution, where x is 0, 5 or 10. If the generated value is less than zero, we truncate it and set it to zero.

b) For all other choice-based questions, most of the time a user will be assigned the standard answer in the profile. However, with probability y , the answer will be one of the remaining choices with equal probability. In this case, y can be 0%, 10% or 20%.

For example, with $y = 10\%$, if the answer to the question “Compact or Full” is “compact” in a particular profile, each user will have a 0.9 probability to answer “compact” and a 0.05 probability each to answer “full” or “not sure”.

4.2.3 Simulation Plan

The above sets up a simulation plan which nine different simulations are run:

	$y = 0\%$	$y = 10\%$	$y = 20\%$
$x = 0$ -Point	Run A	Run B	Run C
$x = 5$ -Point	Run D	Run E	Run F
$x = 10$ -Point	Run G	Run H	Run I

Table 3: Simulation Plan

In each of the nine runs, 500 users will be generated from each of the nine profiles for a total of 4500 users per run. Each user then goes through the virtual advisor dialog and the conflict row is logged.

4.2.4 Clustering Procedure

As a result, for each run, a conflict matrix with 208 columns (each column represents a particular conflict) and 4500 rows (one user per row) is generated. We run the FASTCLUS procedure of the SAS software twice on this matrix:

Case 1: on only the first 3000 users, corresponding to the first 6 conflict profiles, hoping to recover 6 major clusters.

Case 2: on all 4500 users, corresponding to all 9 profiles, hoping to recover 7 major clusters (because the last 3 profiles should generate no conflict, so all those lines should be zero).

As in the clustering procedure we performed on real user data, we start with a small, also minimum, number of clusters (for Case 1, 6 clusters and for Case 2, 7 clusters). After getting the clustering result, we look at each cluster to see if we can provide an interpretable meaning to that cluster. In this case in the simulation, we decide whether each resulting cluster corresponds to one of the original profiles. If some of the profiles are collapsed into one big cluster, we repeat the clustering procedure by incrementing the parameter `maxclusters` by 1. We stop until each of the major resulting clusters correspond to their original profiles.

4.3 Results and Discussion

4.3.1 Sample Run

Here is a sample result after running the clustering procedure on run F ($x = 5$ -Point, $y = 20\%$) with all 4500 users. See Appendix 2 for a complete list of results.

9F	Clusters									
Profiles	1	2	3	4	5	6	7	8	9	Total
1	441	0	0	11	0	3	31	9	5	500
2	9	377	0	6	0	52	25	0	31	500
3	12	0	289	5	9	6	169	5	5	500
4	29	0	0	405	0	0	66	0	0	500
5	0	2	0	0	384	0	100	0	14	500
6	0	0	0	12	0	288	146	6	48	500
7	88	0	0	92	0	1	1302	5	12	1500
Rate:	77.5%									

Table 4: Sample Simulation Result

As we can see, most of the users are successfully classified into their original profiles. Some of the users are misclassified into other clusters because after we have introduced errors in the answers, those particular users' answers may deviate from the original set. Therefore, those users may produce drastically different conflicts. Also note that many users from Profile 1 – 6 are misclassified into Cluster 7 which corresponds to the non-conflict cluster. This happens because those users' responses may not generate conflicts anymore after errors are introduced.

The correct classification rate is calculated by summing up the numbers along the diagonal and dividing it by 4500. Note that Profile 7, 8 and 9 are treated as the same profile (Profile 7) because there should be no way to distinguish among them as they don't generate any conflict. Therefore, there should be 1500 users in that profile.

4.3.2 Correct Classification Rate

The following tables summarize the result of the simulation testing:

6 Profiles	y = 0%	y = 10%	y = 20%
x = 0-Point	100.0% (6)	86.8% (8)	66.6% (7)
x = 5-Point	97.6% (7)	77.3% (11)	65.7% (9)
x = 10-Point	98.3% (7)	78.7% (11)	66.8% (10)

9 Profiles	y = 0%	y = 10%	y = 20%
x = 0-Point	100.0% (7)	86.0% (8)	79.4% (9)
x = 5-Point	98.5% (8)	83.2% (12)	77.5% (9)
x = 10-Point	99.2% (8)	83.4% (12)	71.1% (11)

Table 5: Correct Classification Rate

Note: The number in bracket is the total number of clusters used in the clustering procedure.

As expected, our clustering procedure identifies the unmet needs perfectly when there are no response errors. In the case where we only cluster the first 3,000 users, only 6 clusters are needed to classify them perfectly. Similarly for all 4,500 users, only a minimum of 7 clusters is needed. For other cases, since errors and variations may introduce new and random groups of conflicts, a small number of users will be classified into small and separate clusters and therefore we need more clusters to identify the 6 or 7 major profiles.

As an example, let us take a look at the result of Run C (x = 0-Point, y = 20%) in the 9 Profile case – the correct classification rate is 79.4%. If we consider Profile 1 alone, which corresponds to “Compact, Towing and Hauling”, introducing 20% variations to the responses means that 20% of the users will not answer “Compact” for the question

“Compact/Full Size”, but will answer “Full-size” or “Not sure”. That will automatically bring them to different clusters because the majority of the users in the profile answers “Compact”. These 20% of users will contribute to the misclassifications. However, not all of the advisor questions are definite “yes or no” type question so not all variations will bring users to different clusters. Therefore, we would consider a 79.4% recovery rate a good enough indicator that our clustering procedure is able to discover major unmet need groups.

Furthermore, if we look across different rows of Table 5, which corresponds to varying the constant-sum responses, we see that the correct classification rate does not drop by too much. In the $y=20\%$ column, the rate only drops from 79.4% to 71.1% after we have introduced 10-point standard deviation to the constant-sum responses. This shows that our procedure is robust enough under high variations.

4.3.3 Existence of Small or Uninterpretable Clusters

As we can see in the result of run F in Table 4, 2 clusters (cluster 8 and 9) does not contribute to the correct classification. Both of the clusters contain too few samples to be representative and Cluster 9 does not have any interpretable meaning. These kinds of clusters arise mainly due to the errors introduced. Some of the users across different original profiles may so happened to end up answering the questions in a very similar way and thus be clustered into a separate cluster. We think that it is justified to discard these small clusters and treat them as experimental noises.

4.3.4 Relations with Real User Data

The clustering procedure used in this simulation is the same as what we use to apply on the real user data mentioned in Chapter 3. In both cases, we start with using a small number of clusters. Then we look at each resulting cluster to see if the cluster contains a large number of users and has a coherent interpretation. If not, we increment the number of clusters until large and interpretable clusters are discovered, while small clusters are discarded. Given the high recovery rate we get in the simulation testing, we show that our clustering procedure can successfully discover major unmet need groups from user data.

Chapter 5: Future User Interface – Talking-Head Trusted Advisor

5.1 Motivation

So far, our efforts have been focused on the algorithm behind making the recommendations in the Trucktown program and on the validation of the clustering procedure. In the presentation side, the virtual advisor and the virtual engineer in Trucktown are represented by a static picture of a human person. Are there better ways to represent such a persona to enhance trust? What capabilities can we add to the user-advisor interaction? In this new extension of the research, we propose to use a talking head 3D-model of a human face to represent the trusted advisor. The advisor will behave the same way as in the original Trucktown, asking simple questions and making recommendations, but with added capabilities of speaking with voice and communicating with facial expressions. We intend to use this talking-head as a platform to test out other variations of user-advisor interactions. This new system represents an opportunity for embedding the next generation of online trusted advisor on corporation websites, and can possibly bring more satisfaction to online shoppers.

5.2 Evolving Trucktown

To achieve our goal to investigate how different mode of user-advisor interactions affect trust, we will build several versions of our virtual advisor with different capabilities, and then survey users for feedback. This section will detail the design of these different virtual advisors.

5.2.1 Static Picture Persona

We will start with the original Trucktown, in which the virtual advisor is represented by a static picture of a human person.

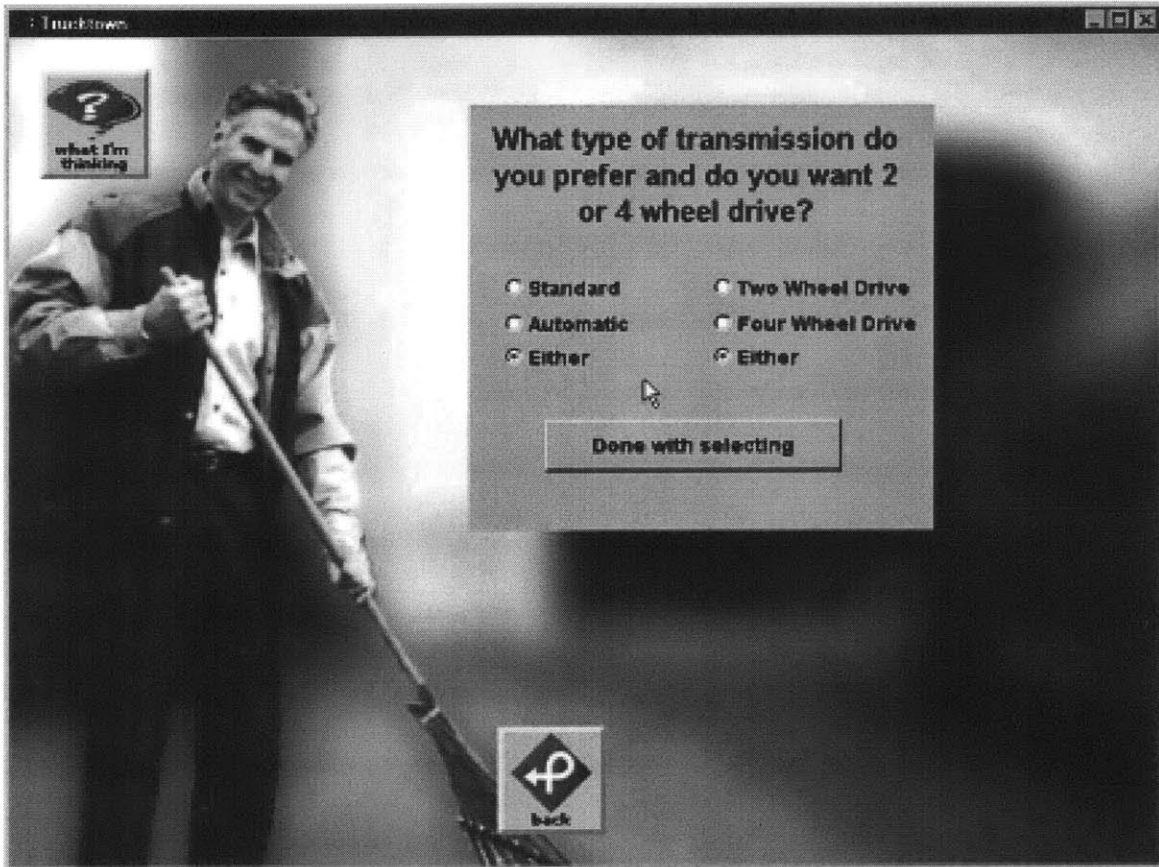


Figure 18: Original Trucktown

No sound is used in this presentation. The picture remains the same in the course of the dialog session. User input is a basic point-and-click approach in which the user is given choice buttons. We have found in the previously mentioned market research that this is an effective way to establish trust with the user.

5.2.2 Talking-Head Trusted Advisor

In the next step, we will add animation to the presentation. The virtual advisor will be a “talking-head”, which is implemented by a 3D-model of a realistic human face. Pre-recorded human voice or synthesized voice are combined with the model to create the “talking-head” effect. The model will also be able to communicate with his facial expressions and emotions. Here are some sample screen shots of the virtual talking-head advisor in action:

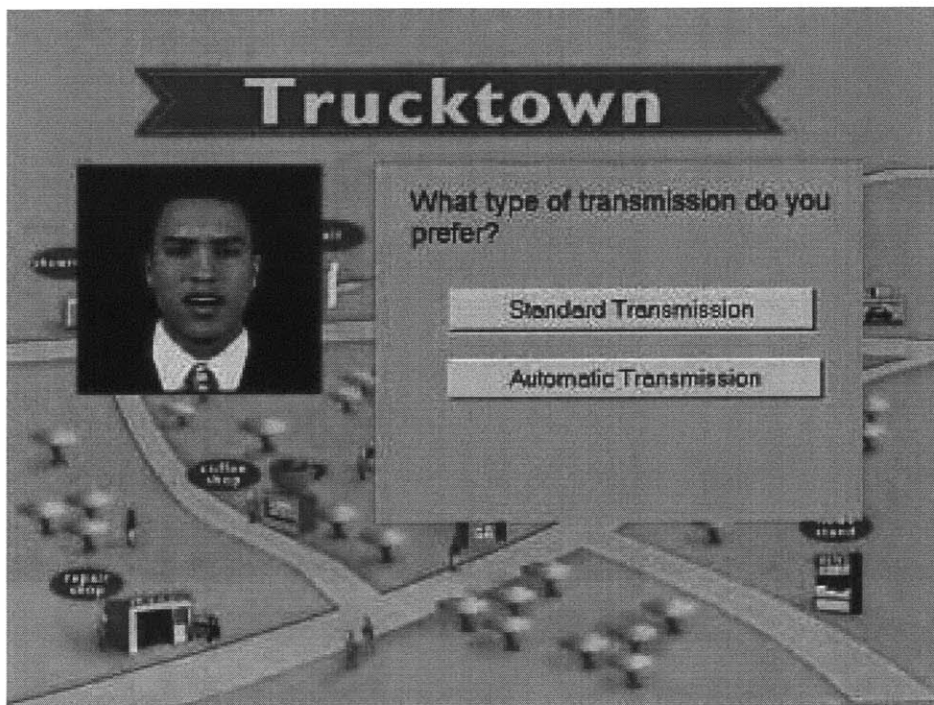


Figure 19: Talking-Head Animation with Voice-Out

The talking-head will read all the original questions to the user. In the animation, the lip movement of the talking-head is synchronized with the pace of the speech. Towards the end of a question, the talking-head may show an inquiring look, and he may nod his head to confirm a user’s response. User input will remain a basic point-and-click

style, and all the other aspects of Trucktown including the truck recommendation algorithm will remain the same.

The talking-head advisor engages in a more lively interaction with the user by showing different kinds of facial expressions while it speaks. We believe that this is a more effective way to establish trust with the user.

5.2.3 Free-Format Text Input

Taking the talking-head advisor idea further, we will try to allow users to type in free-format text input instead of to click on pre-defined choices. For example, instead of clicking on the button “Four-wheel Drive”, the user may type in “I’d like to have four wheel drive in my truck.” Here are some screen shots from a prototype of the system.

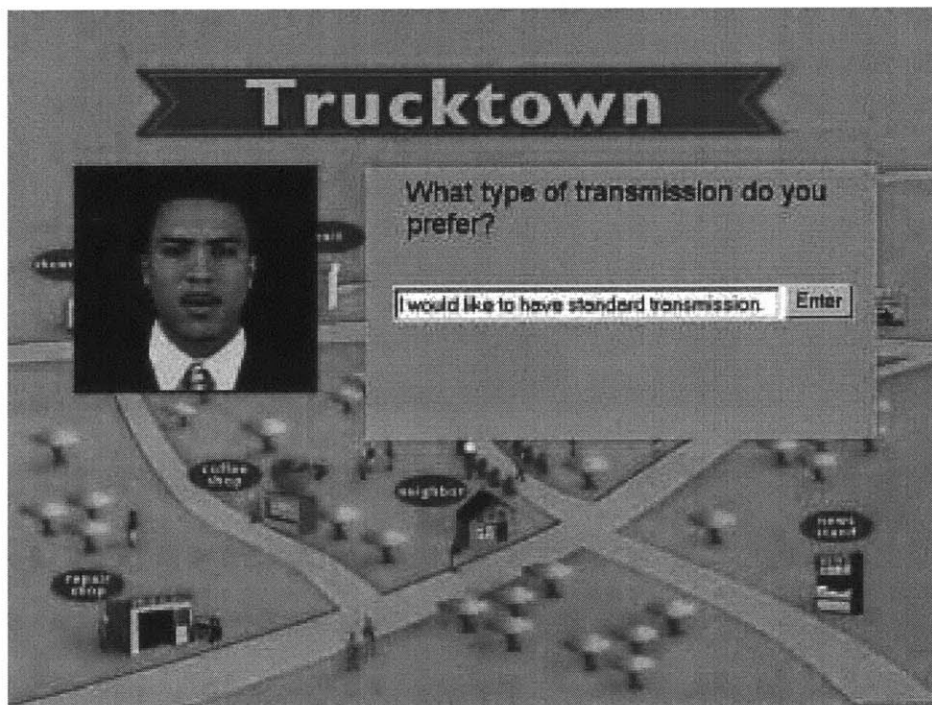


Figure 20: Free-Format Text Input

The same talking-head will be used as voice and animation output. A natural language understanding engine will be implemented to generate appropriate responses for the advisor dynamically. In the two systems mentioned above, the dialog is basically a question-and-answer session with the advisor asking the questions and the user providing the answer. With the new capability, the user will not be limited to only providing the answers, but can raise questions as well. This can potentially increase user satisfaction.

5.2.4 Future Possibilities

In the future, virtual trusted advisor may take the form of a virtual reality avatar, sensing input signals from user's speech, gestures or body movements. Users may stand in front of different sensors which detect their movements. The program can then try to "understand" the user's intention and react appropriately.

5.3 Technology

The software that enables the talking-head concept is provided by LifeFX.com. It provides proprietary development tool to generate the content file from a pre-recorded voice file (.wav) and a text script. The facial movement, emotions and lip synchronizations are generated dynamically to match the voice. The development tool also allows the developer to adjust the magnitude and timing of some gestures such as eye-blinking, nodding and smiling.

The content file, including the voice and animation, is incredibly small (approximately 10Kb per file) due to the proprietary compression scheme. The small file size allows the contents to be downloaded easily even in low-bandwidth scenarios, such

as a dial-up connection. The LifeFX player, which is freely downloadable, is required in the user's browser to play the content file and a license is required to publish such contents on the server side.

The LifeFX player is an ActiveX control that can be embedded in a web page. Users can use a standard Internet Explorer browser to play the content. Here is a snippet of code used to embed the player into a web page.

```
<OBJECT
  ID="LifeFXPlayer"
  HEIGHT="235"
  WIDTH="235"
  CLASSID="CLSID:32634F75-03FF-11D4-B346-00C04FA06E32"

CODEBASE="http://betamirror2.lifefx.com/FaceOfTheInternet/lfxplr.cab#version=2,50,0,0">
<PARAM NAME="StandIn" VALUE="{00000000-0017-11D4-80A4-010000000000}">
<PARAM NAME="PlayStream" VALUE="sound01.lfxi">
<PARAM Name="RemoteStandinDirectory"
VALUE="http://underthedome.mit.edu/Lifefx_work">
</OBJECT>
```

Pre-recorded human voices are used to create content file in the current project. Alternatively, synthesized voice generated from a text-to-speech engine, the IBM ViaVoice system, can also be used. The advantage of using such a text-to-speech engine is the flexibility. The scripts or sentences spoken by the talking-head can be synthesized dynamically and be turned into voice in real-time without having to pre-record every possible sentences in advance.

Chapter 6: Discussion

6.1 Using Slide Bar as Input Method

In the final screen of the design pallet, the user is presented with the original truck recommended by the advisor, and the final truck designed by the user himself side by side. The user is then asked to use a slide bar to indicate which truck they prefer. The value of the slide bar is by default set in the middle, to indicate that the user would be neutral toward the two trucks. From the user data we collected, we found that 37% of the users leave the value in the middle. We are not sure if those users really mean that they are neutral towards both trucks, or they actually miss the question. Here is a screen shot of the final screen of the design pallet:

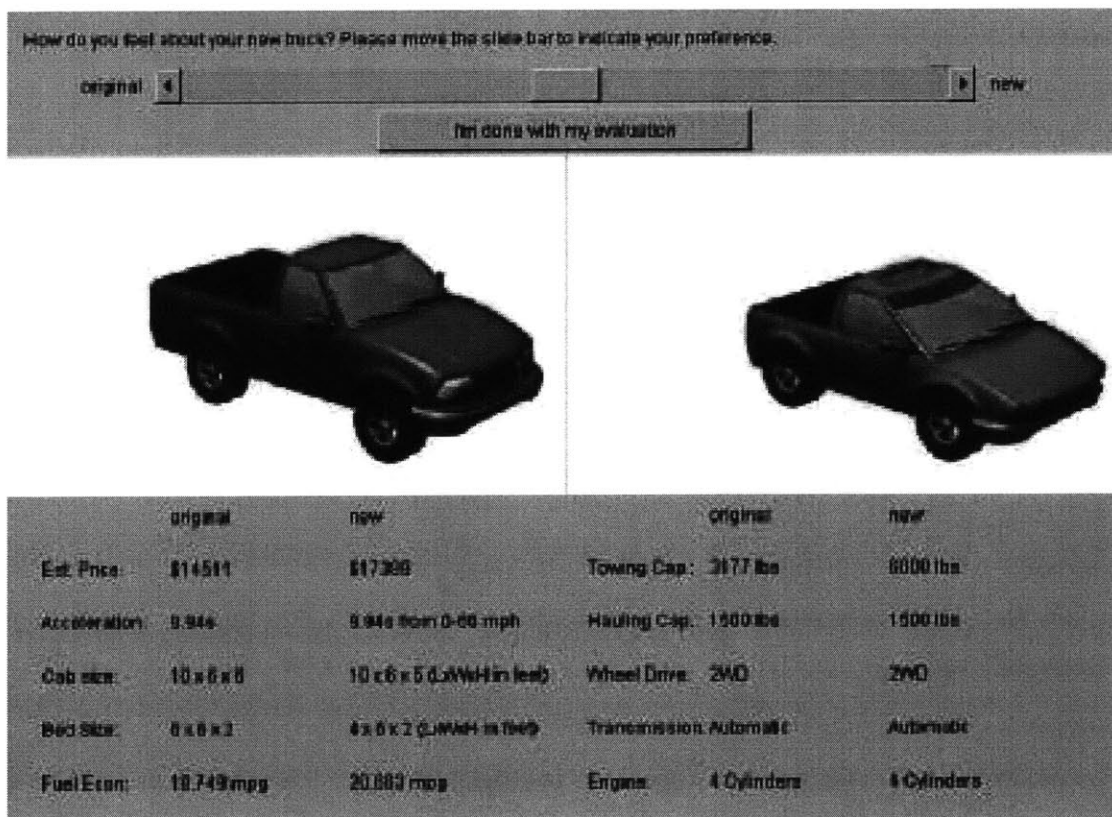


Figure 21: Original Design Pallet Final Screen

In order to make sure users actually respond to that question, we change the slide bar to a set of radio buttons with no default value. Users must provide a choice to indicate their preferences before leaving the design pallet.

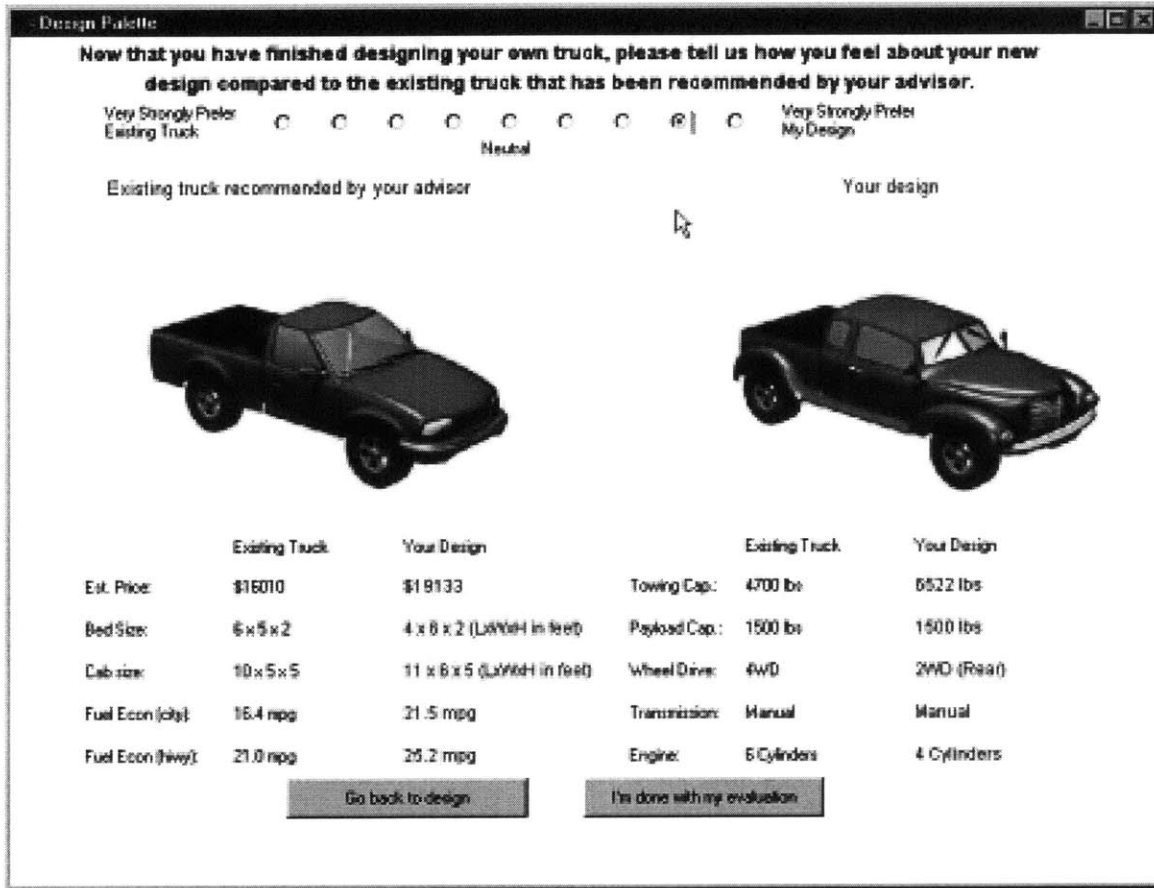


Figure 22: Modified Design Pallet Final Screen

6.2 LifeFX Player

The talking-head mentioned in Chapter 5 needs to be played on the proprietary LifeFX player. Although the content file holding the voice and animation is optimized for real-time online transfer, each user still needs to download and install the LifeFX player to their browser once before the first time they can see the animation content. The player itself is about 300KB in size and may take 5 minutes or more to download in a 56K dial-

up connection. This may create problems to users who have slow and unreliable connections. For now, we have to assume that participating users should have a fast and reliable connection.

6.3 Quality of Text-to-Speech Engine

To create the sound for the talking-head animation, we can either pre-record human voice or use the a text-to-speech engine to dynamically synthesize the voice. To gain flexibility during development of the system, it would be more preferable to use the text-to-speech engine so as to avoid the effort of pre-recording all possible sentences the talking-head will be speaking. On the other hand, we regard the quality of the voice an important trust cue which can affect user experiences. Therefore, it would be most preferable to use a text-to-speech synthesizer which can produce highly realistic human voice. However, our current experience with the system is not satisfactory. Sometimes the synthesizer is not “smart” enough to raise the tone toward the end of a question. More research needs to be done in this area.

Chapter 7: Conclusion

This document details the design and evaluation of the Trucktown system. The statistical analysis on real user data discovers major unmet needs in the current US pickup truck market. The results also reveal that users generally find their experiences enjoyable and believe that their input will help design the next generation of vehicles. A simulation testing has been performed on our clustering procedure to show the robustness of the procedure. As an extension of the research, we have built a prototype of a talking-head representation of the trusted advisor. This 5-year research provides a methodology that represents a revolutionary method for new product development on the Internet.

Reference

- [1] Glen L. Urban, Farreena Sultan, William Qualls, "Design and evaluation of a trust based advisor on the Internet," July 1999, MIT-SSM working paper, MIT 1999.
- [2] Glen L. Urban, Farreena Sultan, William Qualls, "Placing Trust at the center of your Internet strategy," MIT Sloan Management Review Fall 2000 Vol 42 Number 1.
- [3] Ken Lynch, "AI Needs-based configurator implemented in Java," MIT M.Eng Thesis, MIT 1998.
- [4] Andy Tian, "Nonlinear navigation and context-sensitive help in the Sloan electronic commerce project," MIT M.Eng Thesis, MIT 1998.
- [5] Xingheng Wang, "Extending the Sloan E-commerce project with intelligent customer interactions," MIT M.Eng Thesis, MIT 1999.
- [6] Glen L. Urban, John R. Hauser, "Design and Marketing of New Products," Prentice-Hall 1993.
- [7] Christopher Mann, "Listening In – developing a virtual engineer for the online identification of unmet customer needs," MIT Master Thesis, MIT 2000.

Appendix 1

Profile Descriptions for Simulation Testing

Each profile is associated with a standard set of answers. The answers to all virtual advisor questions will be, by default, the same as the set listed below. Each profile will have different answers for a subset of questions.

Standard Answers (if not specified within profile)

1.	Compact/Full	= not sure
2a.	Number of Passengers	= 3
2b.	Number of Front Passengers	= 1
2c.	Easy Rear Entry	= 3
3,4.	Construction	= no
	Towing	= no
	Hauling	= no
	Offroad	= no
	Snow Plowing	= no
6.	Budget	= >32K
7.	Brand	= All no
8a.	Drivetrain	= either
8b.	WheelDrive	= either
9.	Style	= not sure
10.	Height	= under 6 feet
11.	Bed	= not sure
12.	Cylinders	= not sure (neither 4,6,8,10,diesel)
13.	Big & Comfort	= 3
14.	Maneuverability	= 3

Profile 1

Cylinders	= 4
Compact / Full	= Compact
Towing	= Yes
Hauling	= Yes

Profile 2

Style	= Sporty
Compact / Full	= Full
Bed	= Short
Cylinders	= 8
Big & Comfort	= 5

Profile 3

Style	= Conventional
Construction	= Yes
Cylinders	= Diesel
Big & Comfort	= 1

Profile 4

Style	= Rugged & Sporty
Compact / Full	= Compact
Bed	= Extra Short
Height	= 6-6'5
Cylinders	= 4

Profile 5

Style	= Conventional & Sporty
Towing	= Yes
Bed	= Short
Cylinders	= 10
Big & Comfort	= 1

Profile 6

Style	= Conventional
Maneuverability	= 5
Bed	= Long
Cylinders	= 6
Big & Comfort	= 4

Profile 7

Compact / Full = Compact

Profile 8

Compact / Full = Full
Construction = Yes
Towing = Yes
Hauling = Yes

Profile 9

Compact / Full = Full

Appendix 2

Results for Simulation Testing

The simulation is run 9 times based of the errors defined below. For each run, 500 users are generated from each of the 9 profiles for a total of 4500 users.

	y = 0%	y = 10%	y = 20%
x = 0-Point	Run A	Run B	Run C
x = 5-Point	Run D	Run E	Run F
x = 10-Point	Run G	Run H	Run I

For the clustering procedure, the resulting conflict matrix is clustered twice: the first time on only the first 3,000 users (6 profiles) and the second time on all 4,500 users (9 profiles). Here are all the results:

6A	Clusters										
Profiles	1	2	3	4	5	6	Total				
1	500	0	0	0	0	0	500				
2	0	500	0	0	0	0	500				
3	0	0	500	0	0	0	500				
4	0	0	0	500	0	0	500				
5	0	0	0	0	500	0	500				
6	0	0	0	0	0	500	500				
Rate:	100.0%										

6B	Clusters										
Profiles	1	2	3	4	5	6	7	8	Total		
1	492	0	8	0	0	0	0	0	500		
2	4	443	53	0	0	0	0	0	500		
3	0	0	500	0	0	0	0	0	500		
4	28	0	3	469	0	0	0	0	500		
5	7	5	28	0	369	0	0	91	500		
6	5	0	28	0	0	332	135	0	500		
Rate:	86.8%										

6C	Clusters												
Profiles	1	2	3	4	5	6	7	Total					
1	368	122	0	0	0	3	7	500					
2	6	392	0	5	0	69	28	500					
3	0	77	299	0	0	12	112	500					
4	115	65	0	320	0	0	0	500					
5	3	241	0	0	229	0	27	500					
6	0	110	1	0	0	389	0	500					
Rate:	66.6%												

6D	Clusters												
Profiles	1	2	3	4	5	6	7	Total					
1	500	0	0	0	0	0	0	500					
2	0	427	0	0	0	0	73	500					
3	0	0	500	0	0	0	0	500					
4	0	0	0	500	0	0	0	500					
5	0	0	0	0	500	0	0	500					
6	0	0	0	0	0	500	0	500					
Rate:	97.6%												

6E	Clusters																				
Profiles	1	2	3	4	5	6	7	8	9	10	11	Total									
1	305	0	0	5	0	3	0	128	59	0	0	500									
2	15	319	0	0	0	0	0	166	0	0	0	500									
3	0	0	441	0	0	4	0	55	0	0	0	500									
4	0	0	0	453	0	0	0	28	6	0	13	500									
5	0	0	4	0	418	0	0	20	0	58	0	500									
6	0	0	0	0	0	382	30	88	0	0	0	500									
Rate:	77.3%																				

6F	Clusters																			
Profiles	1	2	3	4	5	6	7	8	9	Total										
1	240	153	5	76	19	0	0	0	7	500										
2	0	383	0	21	46	0	0	50	0	500										
3	0	186	287	2	5	0	17	0	3	500										
4	15	62	0	423	0	0	0	0	0	500										
5	0	26	0	0	357	0	115	2	0	500										
6	0	121	2	7	82	281	0	0	7	500										
Rate:	65.7%																			

6G	Clusters								
Profiles	1	2	3	4	5	6	7	Total	
1	500	0	0	0	0	0	0	500	
2	0	450	0	0	0	0	50	500	
3	0	0	500	0	0	0	0	500	
4	0	0	0	500	0	0	0	500	
5	0	0	0	0	500	0	0	500	
6	0	0	0	0	0	500	0	500	
Rate:	98.3%								

6H	Clusters											
Profiles	1	2	3	4	5	6	7	8	9	10	11	Total
1	447	0	0	1	0	0	52	0	0	0	0	500
2	36	459	0	0	0	0	0	0	5	0	0	500
3	53	5	383	0	0	0	0	0	0	59	0	500
4	53	0	0	401	0	0	46	0	0	0	0	500
5	157	0	0	0	318	0	0	0	0	0	25	500
6	54	7	0	0	0	352	0	84	3	0	0	500
Rate:	78.7%											

6I	Clusters										
Profiles	1	2	3	4	5	6	7	8	9	10	Total
1	366	0	3	5	0	0	0	3	79	44	500
2	78	366	10	0	0	0	0	46	0	0	500
3	48	0	364	0	0	0	7	78	0	3	500
4	65	0	0	279	0	0	0	0	144	12	500
5	204	0	12	0	272	0	12	0	0	0	500
6	23	21	9	0	0	356	0	91	0	0	500
Rate:	66.8%										

9A	Clusters								
Profiles	1	2	3	4	5	6	7	Total	
1	500	0	0	0	0	0	0	500	
2	0	500	0	0	0	0	0	500	
3	0	0	500	0	0	0	0	500	
4	0	0	0	500	0	0	0	500	
5	0	0	0	0	500	0	0	500	
6	0	0	0	0	0	500	0	500	
7	0	0	0	0	0	0	1500	1500	
Rate:	100.0%								

9B	Clusters								
Profiles	1	2	3	4	5	6	7	8	Total
1	447	12	0	0	0	0	0	41	500
2	0	498	2	0	0	0	0	0	500
3	0	29	469	0	0	0	0	2	500
4	3	6	0	432	0	0	0	59	500
5	0	168	0	0	332	0	0	0	500
6	2	25	2	0	0	468	3	0	500
7	9	248	0	0	2	0	1225	16	1500
Rate:	86.0%								

9C	Clusters									
Profiles	1	2	3	4	5	6	7	8	9	Total
1	426	5	3	10	1	9	46	0	0	500
2	0	306	21	2	4	28	133	0	6	500
3	4	10	408	0	0	6	45	27	0	500
4	54	19	0	357	0	20	50	0	0	500
5	15	29	10	0	401	0	45	0	0	500
6	5	0	0	0	0	352	83	0	60	500
7	110	5	39	0	3	18	1325	0	0	1500
Rate:	79.4%									

9D	Clusters								
Profiles	1	2	3	4	5	6	7	8	Total
1	500	0	0	0	0	0	0	0	500
2	0	431	0	0	0	0	0	69	500
3	0	0	500	0	0	0	0	0	500
4	0	0	0	500	0	0	0	0	500
5	0	0	0	0	500	0	0	0	500
6	0	0	0	0	0	500	0	0	500
7	0	0	0	0	0	0	1500	0	1500
Rate:	98.5%								

9E	Clusters												
Profiles	1	2	3	4	5	6	7	8	9	10	11	12	Total
1	399	46	0	0	0	3	0	0	51	0	1	0	500
2	4	476	0	0	0	0	0	0	0	0	20	0	500
3	0	80	416	0	0	4	0	0	0	0	0	0	500
4	0	20	0	454	0	0	0	0	0	24	2	0	500
5	0	28	1	0	416	0	0	55	0	0	0	0	500
6	0	95	0	0	0	369	0	0	0	0	0	36	500
7	11	261	7	0	0	0	1216	0	0	0	5	0	1500
Rate:	83.2%												

9F	Clusters									
Profiles	1	2	3	4	5	6	7	8	9	Total
1	441	0	0	11	0	3	31	9	5	500
2	9	377	0	6	0	52	25	0	31	500
3	12	0	289	5	9	6	169	5	5	500
4	29	0	0	405	0	0	66	0	0	500
5	0	2	0	0	384	0	100	0	14	500
6	0	0	0	12	0	288	146	6	48	500
7	88	0	0	92	0	1	1302	5	12	1500
Rate:	77.5%									

9G	Clusters								
Profiles	1	2	3	4	5	6	7	8	Total
1	500	0	0	0	0	0	0	0	500
2	0	462	0	0	0	0	0	38	500
3	0	0	500	0	0	0	0	0	500
4	0	0	0	500	0	0	0	0	500
5	0	0	0	0	500	0	0	0	500
6	0	0	0	0	0	500	0	0	500
7	0	0	0	0	0	0	1500	0	1500
Rate:	99.2%								

9H	Clusters												
Profiles	1	2	3	4	5	6	7	8	9	10	11	12	Total
1	389	0	0	44	0	0	67	0	0	0	0	0	500
2	0	459	0	0	0	0	41	0	0	0	0	0	500
3	0	0	408	0	29	0	58	0	5	0	0	0	500
4	0	0	0	381	0	0	25	94	0	0	0	0	500
5	0	0	0	0	327	0	160	0	0	0	13	0	500
6	0	0	0	0	0	350	69	0	6	75	0	0	500
7	20	0	0	1	13	0	1437	0	0	0	0	29	1500
Rate:	83.4%												

9I	Clusters											
Profiles	1	2	3	4	5	6	7	8	9	10	11	Total
1	244	0	1	88	0	2	76	51	34	4	0	500
2	63	249	0	0	0	62	88	0	0	38	0	500
3	2	0	373	0	0	7	115	1	0	0	2	500
4	3	0	4	455	0	0	32	3	3	0	0	500
5	121	0	4	0	291	0	54	23	0	0	7	500
6	2	42	0	7	0	370	67	5	0	5	2	500
7	57	0	52	14	0	3	1218	99	10	36	11	1500
Rate:	71.1%											