Human-Computer Interaction -- INTERACT'03
M. Rauterberg et al. (Eds.)
Published by IOS Press, (c) IFIP, 2003, pp. 876-879

The Intelligent E-Sales Clerk: the Basic Ideas

Giovanni M. Sacco

Dipartimento di Informatica, Università di Torino,
Corso Svizzera 185, 10149 Torino, Italy

sacco@di.unito.it

Abstract: E-commerce is one of the most active and important Internet application areas, yet selecting a product to buy is normally quite a frustrating experience. We identify the principal user tasks: the thinning-game and the end game. The thinning-game is used to find a suitably small set of candidate items on the basis of personal requirements. The end game is used to compare the features of a set of candidate items in order to find a single "right" item to purchase. We propose effective solutions for both tasks. The thinning game is solved by dynamic taxonomies, a recently introduced knowledge management model. For the end game, which is inherently an information presentation problem, a color-coding scheme is used.

Keywords: e-commerce, dynamic taxonomies, color-coding

1 Introduction

Current e-stores use traditional search tools such as database queries, text retrieval queries (Blair and Maron, 1985), hypermedia (Groenbaek and Trigg, 1994), etc., to help users locate products. However, these tools are useful only if the user knows exactly what he wants. We contend that the typical user does not know exactly the specific item he wants, but is rather looking for the item that best fits his individual requirements. This situation requires browsing and exploration as opposed to retrieval on the basis of a precise specification.

We propose a new access model for this typical user. His interaction is seen as iteratively discarding irrelevant items until he arrives at a set of candidate items sufficiently small to be manually inspected (the thinning-game). At this point, an exhaustive comparison of the features of the candidate set is to be performed (the end-game) in order to find a single "right" item to purchase. Solutions for both games are discussed. Online demos are available at www.knowledgeprocessors.com.

2 The user access model

The typical user accessing an electronic store has basically the task of finding the "right" product in a possibly quite large set of alternative products. Which product is the right one really depends on how competing features rate according to user perceptions, interests and financial capabilities: these really depend on the specific user and it is quite likely that different users or even the same user at different times will use different weights for each feature. Note that while it is unlikely that a patron is able to associate a precise numeric weight to each feature, it is usually very easy for him to identify the most important features and to rank them in decreasing order of importance.

In addition to a primary interest focus, the user will have a secondary, tertiary, etc. focus: e.g. a fledging pro might be interested in budget cameras with the highest resolution, while a cost-conscious amateur will possibly prefer the lightest budget cameras available. Thus, a secondary focus depends not only on the user preferences but also on the features that items in the primary focus exhibit, and so on.

We split the interaction into two stages in cascade: the thinning-game and the end game. Although the primary action (discarding irrelevant items) is the same, the main difference between the two stages is the size of their input: in the thinning-game, the initial number of items is potentially large, while in the end-game a limited number of items is used. In turn, this has the implication that the thinning game is closely related to data management, while the end game is related to information presentation.

In the ***thinning game,*** the user is confronted with a large number of items and has to derive a relatively small set of candidate items to be further exhaustively inspected. In order to be able to efficiently thin the

number of alternatives the user has to find all the available features, focus on the most relevant one for him (the primary focus), discard all the items without that feature, find all the features for the items retained, select the next focus among them, and iterate the process until the number of candidates is sufficiently small.

There are a number of critical points in this interaction (Sacco, 2002a), the main one being the display of all the features correlated with the selected ones. What are the features (e.g. resolution, zoom, etc.) for cameras under $200? If the user is not able to find them out easily, the next focus cannot be set and the thinning game is already over: the user has to inspect all the cheap cameras and find their features by enumeration. On the other hand, if related features are available, he can add to the current focus the next feature in the order of perceived importance and focus on it, thereby discarding other documents that do not have that feature and consequently further thinning the number of candidate items.

The most important points for the thinning game are the ability to operate on items at a set-at-a-time rather than at an instance-at-a-time level (the primary focus defines a set of items, a secondary focus intersects the primary focus set with the set defined by the secondary focus, etc.), and to have systematic summaries of sets (the current focus) in real time. Presentation tends to be a second-order concern.

The second stage, the ***end game***, takes over when a suitably small set of candidate items has been located. At this point, the user must compare all the features for all the candidate items. The usual organization is a table with features on the rows and items on the columns. In real applications, the number of features for each item may be relatively high (tens of features) and the number of candidate items often larger than 10-20. While from the data management point of view the quantity of data to be managed is trivial, even the simplest situation may well result in hundreds of comparisons, which lead to total user disorientation.

The primary concern here is to assist the user to select the single "right" item by minimizing the number of comparisons to be performed. In addition, the screen real estate represents a critical resource that is easily and quickly exhausted when more than a few items and features are to be compared: appropriate commands to selectively hide or show features and items are required. The end game is therefore primarily concerned with presentation and human factors rather than data management.

In order to minimize the number of comparisons, the user should be assisted in quickly finding "discriminants" among different items, i.e. features with different values that can guide the selection. This implies that, at the very minimum, features whose values are the same over all the items, and are therefore useless as discriminants, should be quickly perceived as such, and discarded on demand.

The user discriminates among different items on the basis of the values of some features: some values of these features are more desirable than others and the user informally "weighs" a combination of features of interest in order to find the best match. In many practical cases, the values of a specific feature can be ranked a priori from the less desirable to the most desirable value. For instance, a smaller price tag is always better than a higher one, even if the user is not cost-conscious. These rankings should be used in such a way that the user quickly perceives the desirableness of feature values in a row, instead of comparing them exhaustively.

3 Dynamic Taxonomies and the Thinning Game

A dynamic taxonomy (Sacco, 2000) is a taxonomy with a multidimensional classification: a document D can be classified under several topics at any level of abstraction as required. No other relationships in addition to subsumptions (e.g. IS-A, PART-OF) are needed at the intensional level. Other relationships among concepts can be dynamically inferred through the following **extensional inference rule**: two concepts A and B are related if there exists a document D such that D is classified both under A (or one of its descendants) and under B (or one of its descendants). This simple rule has a dramatic importance because relationships between concepts can be inferred on the basis of the actual classification only. This means that we can easily account for dynamic, unanticipated relationships (information discovery), and at the same time keep the taxonomy simple since compound concepts, which are responsible for the combinatorial growth of classification taxonomies, are not really needed because they can be synthesized on the fly.

The browse and retrieval system is based on the extensional inference rule. The user is initially presented with the complete taxonomy. He selects one or more topics of interest: these define a set of documents, which is the current focus. Then, he hits Zoom. The system presents a ***reduced taxonomy***, which describes the focus by retaining all and only the concepts that are related with the current focus (that is, a taxonomic summary of the current focus). Any concepts in the reduced taxonomy can be used to set an additional focus and zoom, so that the user can

iterate this process until the focus is sufficiently small for manual inspection.
The user is guided to reach his goal and analysis and simulation (Sacco, 2002b) show that the reducing power of zooms is extremely effective and that 2 zoom operations are sufficient to reduce a 100,000 item information base to 10 items on the average.
Dynamic taxonomies offer three additional benefits. First, the extensional inference rule can be extended to produce conceptual summaries for set of documents retrieved through any retrieval method (database queries, text retrieval, etc.), so that dynamic taxonomies can be effectively integrated with other search methods. Second, multilingual access is very easily supported, since it only requires the translation of concept labels. Finally, they allow the transparent and unobtrusive gathering of user preferences, by simply monitoring the concepts used for zooming, which are the real user interests.
An example of interaction is shown in figures 1 to 3.

4 The end game

The normal display for product comparison is shown in figure 4. All the features for three Nikon cameras are shown. Alternating gray and white backgrounds are used to increase row readability: however, the user has no orientation in the comparison of features, so that selecting the "right" camera requires the comparison of all the features. Even the comparison of two cameras is difficult: not only all the features must be compared, but all the different features must also be remembered. Since different features will be stored in the user short-term memory (Miller, 1956), which holds 7±2 items, a comparison between items with more than nine different features becomes quite complex and usually either disorients the user or requires that he use additional tools (e.g. pencil and paper).
Figure 5 shows the enhanced display (Sacco, 2001), according to the requirements discussed in section 2. Here a gray background identifies features with the same value that can be hidden or shown through a button. For most of the features (e.g. price, resolution, etc.), values have been ranked according to their desirableness: the color of the background goes from a bright green (for the best values) to a bright red (for the worst values); a white background is used for mean values. Color ranking immediately provides a guidance for the user not only on what's different and what's not, but also on where items differ more dramatically: a preliminary assessment can be usually done at a glance. In figure 5, the comparison of the three Nikon cameras is immediate and simply based on the number of red features of each camera, whereas it requires a careful inspection of all the features in figure 4.

5 Conclusions

We have proposed a new approach to product selection in large e-stores. Dynamic taxonomies greatly simplify the thinning game because they supply a focus mechanism and continuing conceptual summaries of features. The current version performs zooms and conceptual summaries in real time even for very large infobases. For the end-game, a color-coding scheme based on feature value desirableness make qualitative assessments quick and easy, and appropriate coding allows focusing the attention to those values that are useful for discrimination.
Users have found the current solution easily understood and useful. Formal tests are planned in order to validate the model and to explore alternate solutions. The inclusion of recommendations (Reinick and Varian, 1997), popularity ranking, similarity searches and adaptive hypermedia (Brusilovsky and Maybury, 2002) is also being investigated.

References

Blair, David C., Maron, M. E. (1985), An evaluation of retrieval effectiveness for a full-text document-retrieval system**,** *Comm. of the ACM*, 28:3, 289-299

Brusilovsky, P., Maybury, M. T. eds. (2002), The Adaptive Web, *Comm. of the ACM 45*, 5, 30-67

Groenbaek, K, Trigg, R. eds. (1994), Hypermedia, *Comm. of the ACM 37*, 2, 26-96

Knowledge Processors, The Universal Knowledge Processor (patents pending) www.knowledgeprocessors.com

Miller, G. A. (1956), The magic number seven plus or minus two: Some limits of our capacity for information processing, *Psychological Review*, 63:2, 81-87

Reinick, P., Varian, H. R. eds. (1997), Recommender Systems, *Comm. of the ACM 40*, 3, 56-89

Sacco, G. M. (2000), Dynamic Taxonomies: A Model for Large Information Bases, *IEEE Transactions on Knowledge and Data Engineering*, 12:3, 468-479

Sacco, G. M. (2001), Procedimento per il confronto assistito di oggetti descritti da un insieme di caratteristiche, Patent application

Sacco, G.M (2002a), The Intelligent E-Sales Clerk, Univ. of Torino, Tech Rep Jan 2002 (revised: May 2002)

Sacco, G. M. (2002b), Conventional vs. Dynamic Taxonomies, Univ. of Torino, Tech Rep June 2002

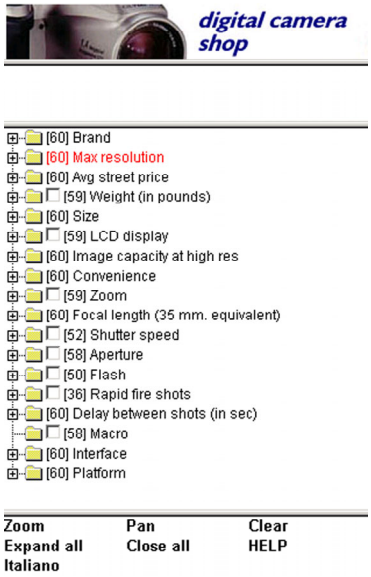

Fig. 1 – Initial taxonomic summary of the features of 60 digital cameras

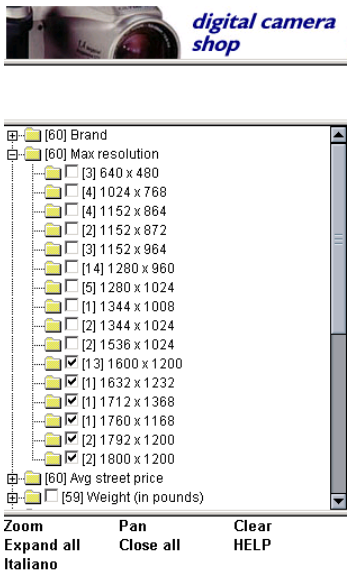

Fig. 2 – Preparing to zoom (focus) on hi-res (1600x1200+) cameras. 20 such cameras exist.

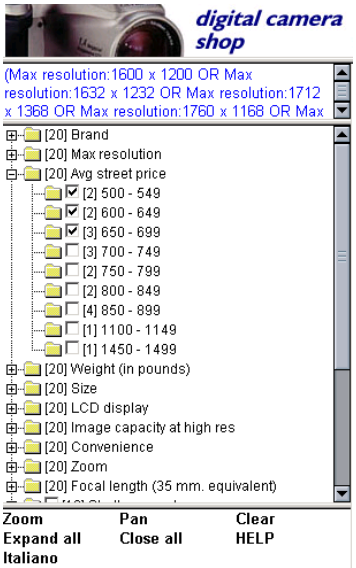

Fig. 3 – Taxonomic summary for the features of hi-res cameras. Features and values that do not apply to hi-res cameras were automatically filtered out. Here, we are preparing to zoom on budget (<$700) cameras.

1 to 3 out of 3

Click here to enlarge images

Click here for smart compare

	Nikon Coolpix 700	Nikon Coolpix 800	Nikon Coolpix 950
Brand	Nikon	Nikon	Nikon
Max resolution	1600 x 1200	1600 x 1200	1600 x 1200
Avg street price	500 - 549	650 - 699	850 - 899
Weight (in pounds)	0.6 - 0.69	0.5 - 0.59	0.7 - 0.79
Size	Medium	Medium	Medium
LCD display	1.8 inch	1.8 inch	2.0 inch
Image capacity at high res	1	8	8
Zoom			
Max optical zoom	no zoom	2 X	3 X
Max digital zoom (optical x digital)	2.5 X	5 X	7.5 X
Focal length (35 mm. equivalent)			
Min	normal (36-60)	normal (36-60)	normal (36-60)
Max	normal (36-60)	low tele (61-100)	tele (101-200)
Shutter speed			
Max	1/750 sec	1/750 sec	1/750 sec
Min	1 sec	8 sec	8 sec
Aperture			
Max	F 2.6	F 3.5	F 2.6
Min	F 8	F 4.7	F 11
Flash	synch red eye reduction	synch red eye reduction	synch red eye reduction
Rapid fire shots	10	10	10
Delay between shots (in sec)	1	1	1
Macro	Macro	Macro	Macro
Interface	serial parallel	serial parallel	serial parallel
Platform	PC Mac	PC Mac	PC Mac

Fig. 4 – Features for Nikon cameras

1 to 3 out of 3

Click here to enlarge images

Click here for dumb compare

	Nikon Coolpix 700	Nikon Coolpix 800	Nikon Coolpix 950
Brand	Nikon	Nikon	Nikon
Max resolution	1600 x 1200	1600 x 1200	1600 x 1200
Avg street price	500 - 549	650 - 699	850 - 899
Weight (in pounds)	0.6 - 0.69	0.5 - 0.59	0.7 - 0.79
Size	Medium	Medium	Medium
LCD display	1.8 inch	1.8 inch	2.0 inch
Image capacity at high res	1	8	8
Zoom			
Max optical zoom	no zoom	2 X	3 X
Max digital zoom (optical x digital)	2.5 X	5 X	7.5 X
Focal length (35 mm. equivalent)			
Min	normal (36-60)	normal (36-60)	normal (36-60)
Max	normal (36-60)	low tele (61-100)	tele (101-200)
Shutter speed			
Max	1/750 sec	1/750 sec	1/750 sec
Min	1 sec	8 sec	8 sec
Aperture			
Max	F 2.6	F 3.5	F 2.6
Min	F 8	F 4.7	F 11
Flash	synch red eye reduction	synch red eye reduction	synch red eye reduction
Rapid fire shots	10	10	10
Delay between shots (in sec)	1	1	1
Macro	Macro	Macro	Macro
Interface	serial parallel	serial parallel	serial parallel
Platform	PC Mac	PC Mac	PC Mac

Fig. 5 – Features for Nikon cameras: color-coded enhanced display