Extending User Vision in the iCARE System

John Doody, Edwin Costello and Lorraine McGinty

School of Computer Science and Informatics,
UCD Dublin, Belfield, Dublin 4, Ireland.
Email: {firstname.surname}@ucd.ie

Abstract. iCARE is a web-based application that aims to provide intelligent customer assistance for recommending eyeglasses. In this domain *case* descriptions take the form of highly technical feature-value pairs that tend to be difficult for users to relate to. In our previous work we have described how we can map (limited) feedback from users, provided over visual examples, to these highly complex case descriptions. In this way we have demonstrated that users can be effectively guided to products of choice, and reduced session lengths can be also be realised by managing in-session feedback provided by the user. Following a series of feedback trials on an early version of the iCARE system we have revised and extended our application in line with the user responses collected. In this paper we describe and illustrate our extensions to the system, and comment on how these revisions were received in our latest live user trials with the iCARE application.

Keywords: Query Revision, User Feedback, Recommender Systems, User Interfaces, E-Commerce Applications;

1 Introduction

Product domains that lend themselves well to case-based recommendation [1, 9] include examples such as *books*, *movies*, *cameras*, *mobile phones*, etc where product descriptions (i.e., product cases) are available in the form of feature-value pairs. For example, in the literature domain, *cases* are often described in terms of the author, publisher, genre, price, etc. and users are usually well able to express their preferences in relation to these features. Taking another example, in the movie domain, users do not usually find it difficult to provide feedback because they have a deep understanding of the domain and can easily express their preferences as a consequence. However, it is important to realise that domains do exist whereby the technical feature descriptions are available for product cases, but users are usually unable to relate their preferences to these (e.g., jewellery, apparel, cosmetics, etc.). Given that they do not understand these domains, users are challenged by the task of providing feedback for recommendation alternatives. This problem is called the *Vocabulary Gap* [2, 3]. In our research we are interested in looking at ways of dealing with this challenging problem by gaining feedback from the user through alternative means (e.g., preference-based feedback over visual examples).

At last year's UK-CBR conference, we described one such domain (i.e., the domain of eyeglass recommendation), and introduced the *i*CARE system (our early-stage demonstrative prototype). This is a domain where users are unlikely to describe their "ideal"' frames in terms of their precise technical features. However, it is also a domain that lends itself well to a case-based representation because such technical feature case descriptions do exist. In addition, it is a domain where alternative recommendation techniques, such as collaborative filtering [7, 8], tend to be unsuitable because irrespective of (dis)similarities that might exist between user characteristics and previous purchase decisions, ultimate decisions are usually very subjective and non-transferable. Following a series of feedback trials on our early version of the *i*CARE system we have revised and extended our application in line with the user responses collected. Our studies have revealed, for instance, that users were often unhappy to end their shopping session when first presented with an option (i.e., a pair of eyeglasses frames) they would be happy to buy; a critical assumption made by

the original *i*CARE system. A further limitation of the original version that was highlighted by our early trialists was that it did not cater for the fact that the opinion of other trusted friends plays a significant role in their final purchase decision. In this paper we describe these extensions to our system, and report on user responses to this latest version of the *i*CARE application.

1.1 Providing a Broader Choice of Vision

The everyday scenario of selecting suitable eyeglass frames presents a number of problems for the customer. These problems include: 1) there is a limited choice, restricted to only those brands and models stocked by a particular store; 2) the selection process is normally oriented by the salesclerk, who makes suggestions, therefore, it is further limited by the patience, expertise and preferences of the salesclerk; 3) choosing the perfect set of frames is often not easy, requiring a detailed analysis by the customer, considering particularly its impact on his cosmetic appearance; 4) the customer usually presents a low uncorrected visual acuity; 5) oftentimes the customer values the options of trusted friends when choosing between a small set of what they feel to be suitable alternatives; 6) the customer may need to visit several optical stores in order to evaluate a sufficient number of alternatives and to compare prices, exhausting a great deal of time, energy and money.

Motivated by the above observations, an increasing number of eye-wear specialists are now establishing an online presence. The vast majority of sites here (e.g., Specsavers.com, EyeGlasses.com) simply allow users to browse through pictures of their product catalogue by presenting category options that satisfy their preferred feature constraints. A recent and growing trend is that eye-wear specialists have started to provide their customers with a better quality of service. Examples include FramesDirect.com and FrameFinder.com whereby the user is afforded the opportunity to virtually "try-on" product options as they browse the product catalogue for each store. Specifically, side-by-side displays of the customer virtually fitted with aesthetically realistic images of eyeglass frames are shown to the customer in order to help them make a purchase decision.

*i*CARE is a web-based application that aims to provide **I**ntelligent **C**ustomer **A**ssistance for **R**ecommending **E**yewear. Akin to the services provided by FramesDirect.com, and other similar online specialists, the *i*CARE system enables a side-by-side comparison of product cases. In addition to this, *i*CARE aims to provide more intelligent assistance (e.g., personalized recommendation generation, opinion gathering and summarisation) through its user interface. The *i*CARE system has the following capabilities: 1) the ability to upload one or more digital images of a customer to be fitted with frame alternatives; 2) enables a customer to select eye-wear from a variety of frame choices for virtual try-on; 3) facilitates the provision of suitable recommendations to the customer based on their limited preference interaction with the system; 4) allows the customer to seek the opinion of trusted friends/advisors on frame alternatives and summarises any feedback gathered; 5) the ability to incrementally reduce the relevant search space (and search time expended) on the basis of incremental feedback provided by the customer (and/or advisors) at each recommendation cycle.

2 Architectural Overview

Figure 1 shows an overview of the *i*CARE system architecture. Central to the functioning of the *i*Care system is the application layer. There are now three crucial components here: the case-based recommendation engine, the product visualisation engine, and the collaborative manager. Particulars relating to these central components will be discussed later in the section. Recent extensions to the system did not affect the Data Representation Layer. Here case descriptions for 3061 pairs of frames are stored. The (partial) case description, shown in Figure 1, is representative of how these are represented by the data-set layer. Each description is represented by a list of 13 *feature-value* instances (e.g., Price, Temple Size, Material etc.).

The primary role of the user interface layer is to handle message passing between the user and the application layer. It is responsible for displaying the recommendations made, and allows the user to *Try-On* a pair of frames, and accepts user feedback to relay back to the recommendation engine. The basic look and feel of the *i*CARE interface remains unchanged. As in the previous version of the system the user has the ability to

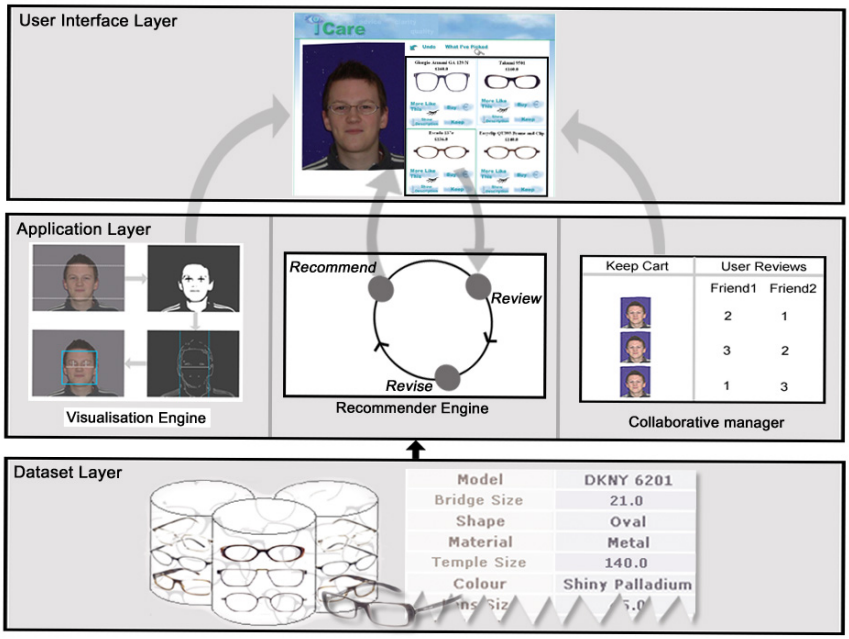

Fig. 1. *i*CARE System Architecture

view technical feature descriptions for each recommendation. In addition, to this we have also included a "keep-cart" facility whereby a user can choose to set aside those recommendations that they like (and may be happy to buy) but would like to come back to at a later stage. The user can also seek collective feedback from a group of trusted advisors/friends through the *i*CARE interface in order to make a more informed purchasing decision.

Importantly, our trialists indicated that options in the keep-cart should remain visible as the session continues. The vast majority of feedback we received here shows that users value being able to compare subsequent recommendation alternatives against those in the session keep-cart, and in fact users demonstrate a greater effort to update their cart (e.g., by deleting options) once they were presented with more suitable choices.

2.1 Recommendation, Feedback and Retrieval

The core algorithm behind the *i*CARE system is based on the comparison-based conversational model proposed by McGinty and Smyth [5]. This conversational recommender engine supports an iterative interaction with the user, providing them with cyclic feedback opportunities to influence retrieval. There are three stages to each cycle: 1) k cases are recommended to the user; 2) the user provides preference-based feedback over these visual examples by indicating a preference; 3) the system revises the query which represents the user's current needs, using only the limited feedback they provide. *i*CARE manages two main strands of feedback communication; individual feedback from a target user collected in recommendation cycles, and collective feedback from the advisors to the target user. In the case of the latter (i.e., *collective feedback*), *ratings-based* information is collected from a set of advisors to the target user. A user profile is stored which contains the entire session information (i.e. the current query, the preference case and the rejected cases for each cycle), the current items in the keep-cart, and all reviews associated with these items. When the target user requests information from advisors, the advisors are required to rate all of the times in the keep-cart.

The review page shown in Figure 4 is generated from the data stored in the user profile. For the individual reviews, we average and normalise all the advisor's ratings generate a *Hot-or-Not* score (see Equation 1).

$$hot\_or\_not(C) = \left(\frac{\sum rating(C)}{n}\right)\left(\frac{10}{Max(rating(P))}\right) \quad (1)$$

Where rating(C) retrieves a rating for a product case C, and is normalised based on the maximum rating over all product cases. The *Hot-or-Not* rating is a value between 1 and 10, where 10 is the best rating that can be applied to a case. The User-Interface uses these values to display the appropriate information in the review screens.

Utilising the example-based [6] feedback collected from individual recommendation cycles involves the system updating its understanding of a user's personal requirements (i.e., the evolving query) at the revise stage of each recommendation cycle. Our early version of *i*CARE took a very simple approach here by using the technical features of the user's preferred case c_p as the query for the subsequent cycle. Updating the query in this way provides the system with crucial fine-grained feature information that it needs to guide its search, without having to burden the user to provide (or understand) these technical features; instead they can simply base their feedback on whether or not they think the presented frame options suit them or not. Importantly, there is a direct mapping between these technical features and the frame options they describe. In our previous work we referred to this (the simplest form of query revision for preference-based feedback) as the *More-Like-This* approach.

Finally, before the recommender can present the user with k recommendations the remaining product cases are ranked in decreasing order of their similarity to the current query, Q, according to the Equation 2. Accordingly, the final score is always a number between 0 and 1.

$$sim(Q,C) = \frac{\sum featureSim(F_{Qi}, F_{Ci})}{n} \quad (2)$$

When calculating similarity at the feature level nominal and numeric values need to be handled differently. For nominal features an exact match comparison is carried out, returning the value 1 when the values match and 0 otherwise. Numeric values, on the other hand, use their relative difference as a basis for similarity calculation. The equation for this is shown in Equation 3 where F_Q and F_C are the values for the numeric features being compared.

$$featureSim(F_Q, F_C) = 1 - \frac{|F_Q - F_C|}{\max(F_Q, F_C)} \quad (3)$$

3 Session Walkthrough

On the outset of a recommendation session with the *i*CARE System, a user is required to upload their digital picture to the system, and provide some basic feature preferences (i.e., *price, shape, gender*) if they wish. The resulting query is made up of the users facial dimensions, (e.g., the location of the eyes), and a list of these feature-value pairs. Once the user has stated their initial preferences regarding a pair of glasses, they are presented with the main recommendation screen. Here, 4 pairs of eyeglasses are presented to the user (see Figure 3). The user can do any of the following: 1) select their preference to generate a new set of recommendations; 2) view the technical features of a product; 3) Add an item to the keep-cart; 4) buy a pair of eyeglasses; 5) *try-on* a pair of eyeglasses. The session summary screens (as shown in Figure 4) displays the results of the advisors ratings and allows the customer to review these reviews and has the option to continue shopping from where they left off, start a new shopping session or to purchase an item currently in their keep-cart.

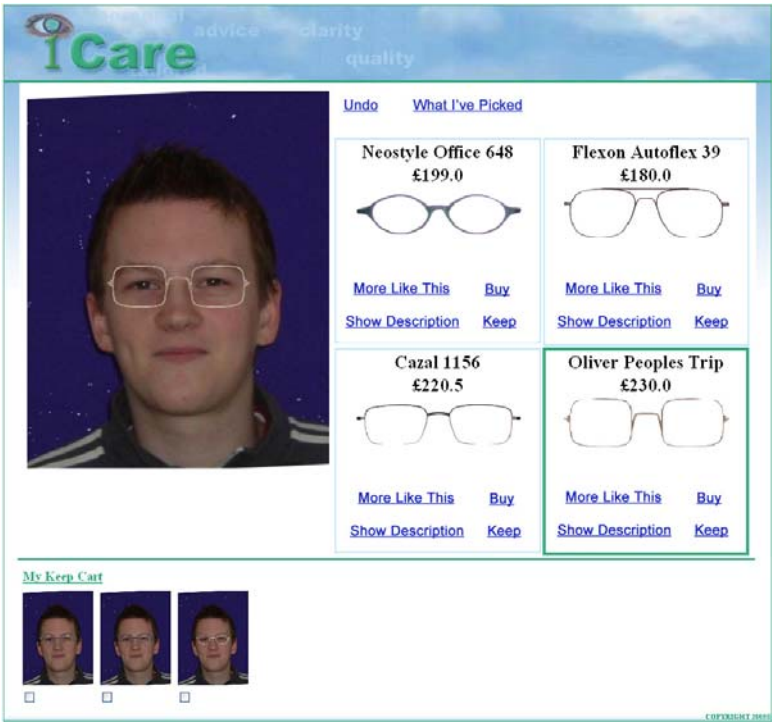

Fig. 2. Illustration of cycle *k* with items in the keep-cart

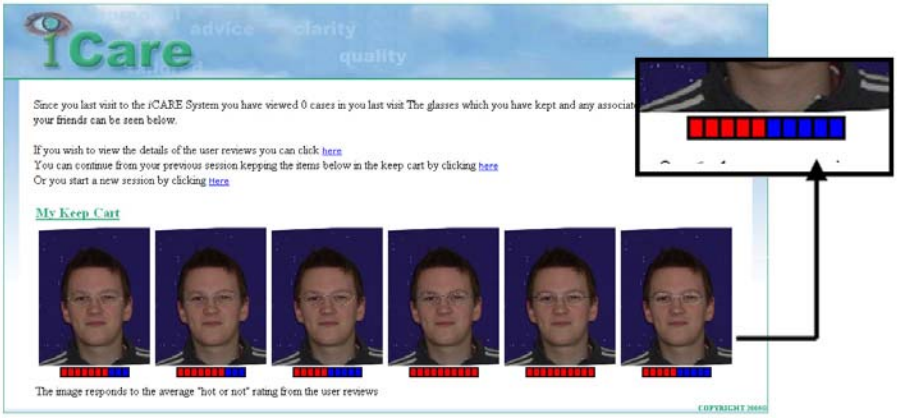

Fig. 3. Summary Screen, showing the items in the keep-cart and the averaged rating

4 Real User Evaluation

The focus of this trial differs from any previous study we have carried out using the *i*Care system [3, 4]. Previous evaluations have mainly concentrated on evaluating the efficiency of the recommendation engine, and investigating ways of further reducing the number of recommendation cycles required to bring users to their preferred target case. In this trial we were more interested in gathering feedback from users in relation to their interaction experiences with the extended system, and seeing what lessons could be learned in relation to implementing further extensions.

4.1 Trial Set-Up

In total 150 users participated in the trial; 40 of these interacted with *i*CARE to find suitable eyeglasses (henceforth referred to as the *target users*), and these target users sought the opinions of a total of 110 *advisors* (i.e., their trusted friends and family). This trial took place over a two week period in August 2006. For the purposes of the trial the *i*CARE application drew on a catalogue of 3061 cases available from a reputable online eyewear specialist. For each of these 3061 cases a technical description and an image of the frames they describe was stored by the system.

Although in the current version of *i*Care the number of recommendation options that can be presented is parameterised and changeable, for the purposes of this trial we preset the recommendation window-size for each session to be 4. The typical user-system interaction that took place in each session was as illustrated in the previously described session walk-through. Target users were simply asked to use the system to find a pair of eyeglasses they would be happy to buy in reality. Importantly, it was not a requirement that they use the *keep-cart* or *ask a friend* facility. Finally, all session interaction information was recorded for each user; that is, what cases they tried on and preferred, which cases were added or removed from the keep-cart and when these additions or removals took place, whether or not they consulted the technical feature descriptions associated with each case, the collective feedback they received back from their advisors etc.

Figure 5 shows some general information about the participants of the user trial. It details their previous experience with recommender systems and online purchasing. Of the 40 target user that participated, 33 had prior experience of using recommender systems, and of those 90% considered them useful in helping to find a desired case.

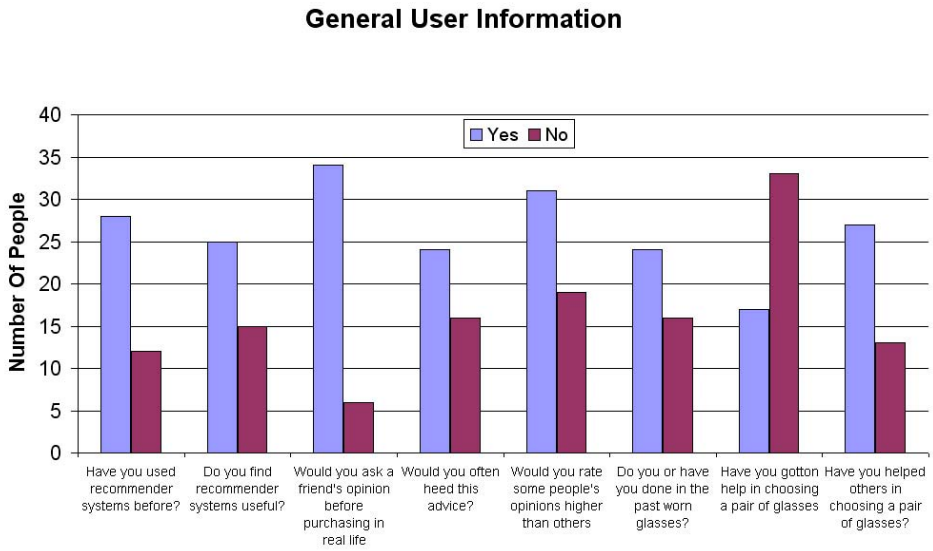

Fig. 4. General User Information

3.3 Key Findings

All trial participants agreed that the try-on facility was crucial and none preferred to interact with the system without this. From the 40 target users 18 consulted the technical features during their interaction sessions with *i*CARE. However, all of these users agreed that their feedback was solely on the basis of their subjective response to seeing the product options on themselves, and so they did not find these features particularly useful. In total 70% of the users liked the simplistic *preference-based feedback* mode of interaction, as it does not require a lot of effort or expertise in the domain from the user. However, in our findings we note that approximately 30% of users did report that they would like to see alternative forms of feedback supported.

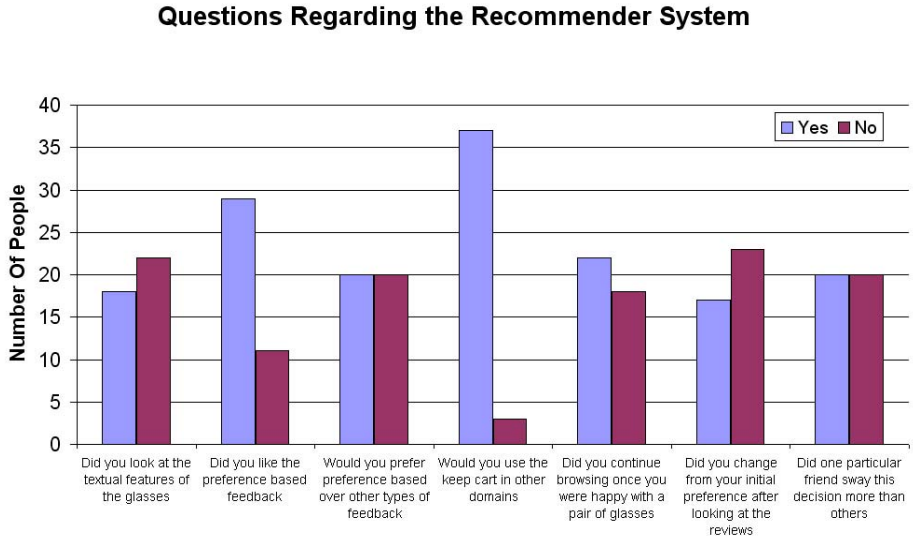

Fig. 5. Questions regarding recommendation

User response to the integration of the *keep-cart* facility was very positive. 95% of users found the keep-cart very useful, with the main reason being that it allowed them to easily compare subsequent cases with those they already liked, and allowed them to continue their session despite having located a case they would be happy to buy. Interestingly, we found that 22 of all target users, having located a satisfactory case, continued to browse in the hope that they might find a "better" alternative.

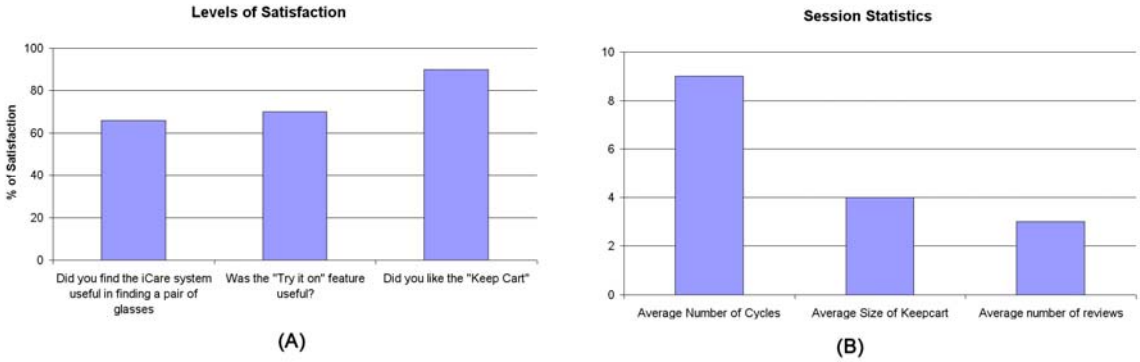

Fig. 6. (A) Degree of satisfaction. (B) Average session statistics

Although only less than 50% of the target users revised their final preference based on their advisor reviews, those that did not change their mind stated that this was because the opinion of the advisors was in agreement with them. All agreed that this was a domain where they would seek the option of others and be influenced heavily by this feedback. Moreover, 50% of users admitted that they would value the opinion of some advisors more than others.

Figure 7(A) gives a brief summary of the recommendation session particulars we recorded averaged over all users. It shows that the average user interacted with the *i*CARE application for 9 cycles before finishing their session or seeking advice from friends. While the number of options that made it into the keep-cart varied widely from user-to-user, the average number works out to be 4 per session. Similarly, target users differed in terms of the number of advisors they consulted (ranging from 2 to 10) the average being 3 advisors per session.

5 Conclusions

The task of building a case-based recommender that concentrates on making relevant suggestions for eyeglasses is a challenging one. Asking individuals to describe their preferences in this domain is unreasonable; that is, users are usually unable to provide information about the type of frames they are looking for using the vocabulary the recommender system understands (e.g., papillary distance, metal type, price). For this reason, to build an effective case-based recommender in this domain, feedback needs to be elicited from a user in a much simpler way. The provision of a virtual *try-on* facility allows customers to better evaluate and compare cases in this domain.

*i*CARE is a web-based application that aims to provide intelligent customer assistance for recommending eyeglasses using this example-based feedback from users over visual examples. In previous related work we have concentrated on how to improve the quality and efficiency of the recommendation engine behind its operation [4] .In this paper we have described how we have extended the *i*CARE application. Feedback gathered in previous user trials highlighted a number of interfacing limitations of the original application that have now been addressed. For instance, users are often reluctant to end their shopping sessions when first presented with an option (i.e., a pair of eyeglasses frames) they would be happy to buy, and the opinion of other trusted friends plays a significant role in their final purchase decision.

In this paper we describe and illustrate how we have incorporated this `ask-a-friend' facility as well as discuss other extensions. User responses gathered in our recent feedback trial with the latest version of the *i*CARE system have been positive. However, some users have highlighted that the system could/should be further extended to support different feedback modes. We think this would be worthwhile, and so as part of our future work agenda we will investigate how best to handle query revision using alternative feedback modes such as critiquing and ratings based feedback.

6 Acknowledgements

This publication has emanated from research conducted with the financial support of Science Foundation Ireland. Sincere thanks are extended to the trailists who volunteered to participate in our recent user study. This feedback is very valuable to us. Your time spent on this task and interacting with the system is most appreciated.

References

[1]R. Burke, K. Hammond, and B. Young. The findme approach to assisted browsing. *Journal of IEEE Expert*, 12(4):32–40, 1997.

[2]E. Costello, J. Doody, L. McGinty and B. Smyth. Using product visualization for preference elicitation in case-based recommenders. In M. Petridis, editor, *Proceedings of the 10th UK Workshop on Case-Based Reasoning*, pages 49-58, Cambridge, UK, 2005.

[3]E. Costello, J. Doody, L. McGinty and B. Smyth. iCARE: intelligent customer assistance for recommending eyewear. In *IUI '06: Proceedings of the 11th international conference on Intelligent user interfaces*, pages 282-284, New York, NY, USA, 2006. ACM Press

[4]J. Doody, E. Costello, L. McGinty and B. Smyth. Combining visualization and feedback for eyewear recommendation. *Proceeding of the Nineteenth International FLAIRS Conference (FLAIRS06).*, 2006.

[5]L. McGinty and B. Smyth. Comparison-based recommendation. In S. Ed. Craw, editor, *Proceedings of the Sixth European Conference on Case-Based Reasoning (ECCBR-02)*, Aberdeen, Scotland, 2002. Springer.

[6]P. Pu and P. Kumar. Evaluating example-based search tools. In *Proceedings of the ACM Conference on Electronic Commerce (EC'04)*, pages 208-217, 2004. New York, USA.

[7]P. Resnick, N. Iacovou, M. Suchak, P. Bergstrom and J. Reidl. Grouplens: An open architecture for collaborative filtering of netnews. In *Proceedings of ACM 1994 Conference on Computer Supported Cooperative Work*, pages 175-186, Chapel Hill, North Carolina, 1994. ACM.

[8]B. M. Sarwar, G. Karypis, J.A. Konstan nad J. Reidl. Item-Based collaborative filtering recommendation algorithms. In *World Wide Web*, pages 285-295, 2001.

[9]H. Shimazu. Expertclerk: Navigating shoppers buying process with the combination of asking and proposing. In *IJCAI*, pages 1443-1450, 2001.