
Using quality criteria to assist in information searching

Mikhaila S.E. Burgess*, W. Alex Gray
and Nick J. Fiddian

Cardiff School of Computer Science,
Cardiff University, Wales, UK
E-mail: M.Burgess@cs.cardiff.ac.uk
E-mail: W.A.Gray@cs.cardiff.ac.uk
E-mail: N.J.Fiddian@cs.cardiff.ac.uk

Abstract: One of the challenges facing today's information consumer is how to find information that meets their personal needs, within an acceptable time frame, and at an appropriate level of quality. One potential method for assisting these consumers is to employ a personalisable, explicit definition of quality to focus information search results. In this paper we discuss the feasibility of this approach by demonstrating how a consumer-refined definition of quality can be used to drive an information search, initially within a closed-world environment. This paves the way for further research, transferring lessons learned and techniques developed to an open, heterogeneous environment.

Keywords: information quality; information overload; information search; ranking algorithm; quality framework; information consumer.

Reference to this paper should be made as follows: Burgess, M.S.E., Gray, W.A. and Fiddian, N.J. (2007) 'Using quality criteria to assist in information searching', *Int. J. Information Quality*, Vol. 1, No. 1, pp.83–99.

Biographical notes: Mikhaila S.E. Burgess is a Lecturer in Information Systems in the School of Computer Science at Cardiff University. Having received a BSc in Computer Science from Cardiff University in 1998, she worked for a year as a Software Developer, returning to Cardiff in 1999 as a PhD Researcher. Since receiving her PhD in 2004, she has worked as a Researcher on a number of national and international projects, primarily in the fields of ethics and biodiversity. Her current research interests include the use of quality measures and ontologies in information retrieval.

W. Alex Gray is a Professor of Advanced Information Systems in the School of Computer Science at Cardiff University, where he leads the School's Knowledge and Information Systems research group. He is also Deputy Director of WeSC – the Welsh e-Science Centre, which is one of eight Regional Grid facilities in UK academia. He was educated at Edinburgh and Newcastle upon Tyne Universities. His current research is in interoperability of distributed data and knowledge-based systems with particular interest in how knowledge and semantics can be exploited to enrich the interoperation process and allow users to have different perspectives of data linkage.

Nick J. Fiddian is Head of School and Professor of Computer Science at Cardiff University. He was educated at London (LSE and ULICS) and Southampton Universities. His current research interests lie in the field of meta-translation – meta-programmed automatic translation between

syntactically diverse but semantically similar programming languages in specific areas of computer application, particularly data/knowledge-base system interoperation.

1 Introduction

A number of problems face the present day information consumer. When searching for information amongst the exponentially increasing amount now freely available, for example, via the internet, it is becoming extremely difficult to find the best information that meets the consumers' requirements within an acceptable amount of time.

One potential method for assisting the consumer in searching for required information is to use a personal, explicit, definition of quality to focus information search results. In this paper we discuss the feasibility of this approach by showing how a consumer-created definition of quality can be used to drive an information search by ranking search results based on these quality preferences.

In this paper we also present our prototype ISE. This software enables the information consumer to conduct searches within a number of experimental subject domains, using previously developed frameworks of quality to explicitly define their quality requirements, which are then exploited to focus the search results based on their personal preferences.

The first section of this paper introduces the background to our work by discussing the two main challenges facing information consumers today, that of 'information overload' and 'information quality'. It goes on to discuss the problems associated with defining quality, followed by an overview of our previously published work on the creation of hierarchical, flexible and extensible frameworks of quality. Our prototype ISE is then introduced with an explanation of how the consumer can use this tool to create a dynamically adjustable personal definition of quality and how this definition is employed in the search process. The paper concludes with a discussion on the results of our experimentation and areas requiring further work.

2 The problem of quality

In this section we discuss the two main issues affecting the consumer's search for information: information overload and information quality. It concludes with a discussion on Google's approach to tackling these issues via their PageRank algorithm, followed by a description of the approach taken in our work.

2.1 Information overload

The internet is currently growing at an exponential rate. According to the bi-annually conducted Domain Name Survey the number of internet domains has increased from under 30 million in January 1998 to over 400 million in July 2006 (ISC, 2006). With so much information now readily available online and being comparatively easy to access, an assumption could be made that finding information on a desired topic should be a straightforward task. However, as the amount of information increases, the task of the

individual to analyse that which is available in order to decide which to accept or reject is becoming increasingly difficult, or even impossible. We, as information searchers, are increasingly suffering from ‘information overload’ (also commonly referred to as ‘cognitive overload’ or ‘information fatigue syndrome’).

“The term information overload describes situations in which the individual is no longer able to integrate new information for decision-making, due to the great amount of information he or she is exposed to. He or she can no longer productively use the quantity of information in the available time scale. In consequence, decision quality, efficiency, and even well-being may be reduced.” (Eppler, 2001)

Research conducted at the University of York found that the relentless flow of information that modern workers have to handle can seriously reduce an individual’s ability to focus on his/her work. During experimentation they discovered that information overload resulted in an average ten point reduction in IQ score, double the amount seen in studies involving cannabis users (Knight, 2005).

Internet search engines such as Google (2006) assist users in finding the information that they require and go some way in alleviating the information overload problems associated with internet searching. However, the results of the WebTop Search Rage Study (Sullivan, 2001) conducted in 2001 showed that one third of search engine users find searching ‘very frustrating’, with three quarters of respondents reporting frustration to some significant degree. It also reported that only 60% of the respondents regularly find the information they require.

The potential for further increasing the amount of information returned from an internet search exists in utilising metasearch engines such as Dogpile (Infospace, 2006). In this instance, a single search query is sent to multiple search engines simultaneously, the results of which are then combined (with duplicate results often removed) and presented to the user. This often provides a greater internet search coverage, but also results in a further increase in the amount of potentially irrelevant information returned, through which the user must seek out that information which is relevant to their needs.

As information overload affects increasing numbers of people, more work is being conducted into the causes, effects and potential solutions. The result is that towards the end of 2006, there exist three million internet sites addressing this issue, and thus themselves contributing to the problem!

2.2 Information quality

Alongside the sheer amount of information now available, there also exists the problem that information obtained from the increasing number of available sources will typically be of differing levels of quality. Due to lack of a controlling body and data constantly changing in such a fluid environment, when searching for information on the internet, sites can contain out of date or unavailable information (Spinellis, 2003), incorrect (BBC, 2003a) or potentially dangerous information – such as providing access to people selling substances not intended for general use; for instance, prescription drugs with little or no safety instructions (BBC, 2003b). This lack of quality becomes a concern when people believe what they see, with little or no regard as to its accuracy (Graham and Metaxas, 2003) and results in the following statement from Naumann: “Information quality is the main discriminator of data and data sources on the web” (Naumann, 2001).

2.3 *Google and the PageRank algorithm*

Google's PageRank algorithm (Page et al., 1998) goes some way towards improving the quality of internet search results, by ranking highly those websites that are referenced by many other sites, with more emphasis on links from sites which are themselves considered as being of high quality. This, in effect, incorporates peer reviews of websites from other site developers. Alongside this, Google's ranking method also takes into account user feedback for each site and favours sites that are frequently updated.

Google's method for ranking the relevance of web pages retrieved from an information search has made it one of the most popular internet search engines currently available. However, the main problem with this, and other, online search services is the assumption that all individuals have the same information requirements. Two users querying an online search system at the same time, entering the same keywords, will receive identical search results, yet they are likely to have different views regarding the quality of that information due to their different opinions, needs, prior knowledge and experience, etc. A one-method-suits-all approach can not meet the needs of all individuals.

2.4 *Our approach*

Our concern is how the information consumer copes with this information quality problem and how, when searching for information, they can be assisted to find information of the best possible quality, from the available sources.

Our proposed approach involves the creation of a flexible model of quality that can be used by the information consumers to assist them in explicitly stating their quality requirements. This personalised definition of quality can then be exploited to focus information search results by ranking highly those items that closely match the consumer's quality preferences. This enhances current search technology by enabling search personalisation, rather than relying on a one-method-suits-all approach.

3 **Defining quality**

Quality is not an easily definable term, as it is not absolute. It has many different facets, and its meaning varies across different situations, users (Firquin, 1992) and queries (Bouch, 2000). A formal method is therefore needed to describe this term, eliminating the need to produce a detailed quality definition each time information is requested.

The paradox of quality is that, as Robert Pirsig notes, "even though quality can not be defined, you know what quality is" (Pirsig, 1974). With a fair amount of thought, the typical consumer may be able to produce some qualitative example of what quality means to them when discussing a specific product. Unfortunately, this definition will be non-transferable to other products or even other consumers when discussing the same product, as each individual's perception of quality will vary depending upon their current circumstances and quality requirements.

As stated by Dromey,

"people make a judgement, depending on their particular needs or perspective, that something they use, encounter, or examine is 'good' or has 'quality'."
(Dromey, 1996)

Current definitions of quality vary from the classic, yet vague, ‘fitness for purpose’ (Juran and Godfrey, 1999) and ‘conformance to requirements’ (Crosby, 1979), to those which have broken quality into a number of attributes, varying dependent on the domain in which they are applicable. Examples of the latter include the ISO standard for software quality, comprising six dimensions and 34 metrics, and Wang and Strong’s definition of data quality containing four categories and 15 dimensions (Wang and Strong, 1996).

3.1 The information consumer

The problem with most current definitions of quality is that they focus on the information provider rather than the consumer. The providers typically have control over their information, data, product or service, and can therefore control its quality, for example by working towards some quality standard or conducting data cleansing. The consumers, however, have no direct control of the quality of the information, product, etc. They have to select from that which is available, without the ability to directly affect quality. The challenge is therefore to assist the consumer in finding the information they require, of the best currently available level of quality, from that which is currently available, based upon their perception of what constitutes quality at the present time.

3.2 Flexible, extensible quality frameworks

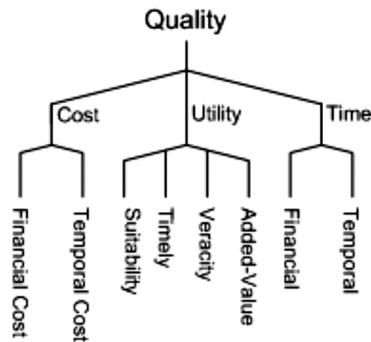
As stated earlier, our approach involves the creation of a flexible model of quality that can be used by the information consumers to explicitly state their quality requirements. The initial stage in the development of this model was the creation of a generic framework of quality, based on the results of extensive research conducted into previous attempts at defining quality.

The principal features of this initial framework are as follows:

- *Generic.* The framework contains a set of quality criteria applicable across a range of subject domains, based on the most frequently identified characteristics of quality stated within the published literature.
- *Hierarchical.* This allows criteria to be grouped into related categories and sub-categories.
- *Intuitive.* Quality criteria are locatable within the hierarchy by following an intuitive path through the structure, to ensure its ease of use. This has resulted in some criteria being visible in more than one location in the hierarchy, to improve navigability.
- *Flexible.* The development of an accompanying software tool, the Quality Toolkit, provides facilities for the relocation of criteria within the hierarchy. The structure can also be used to enable the selection of a subset of criteria to create a personalised definition of quality.
- *Extensible.* The aforementioned accompanying tool provides facilities for the addition of new criteria and the deletion of un-required criteria, meaning the framework is not static and can be updated over time as the result of feedback.

An illustration representing the highest level of this hierarchical model can be seen in Figure 1. This figure illustrates the structure of the generic quality framework by presenting the high level quality dimensions. Each of these dimensions contains further sub-dimensions of quality, each containing either further dimensions or a set of individual quality criteria. For more information about this generic framework, including a complete description plus details of its development, the reader is referred to Burgess et al. (2002, 2006b).

Figure 1 Upper level of the generic quality framework



A generic model is not always appropriate, as some quality criteria are only applicable within specific domains. For example, ‘research quality’ is an important criterion when rating the quality of research institutions such as universities, but is of little value elsewhere. Therefore, the next stage of this project involved the development of a set of domain-specific frameworks of quality, each containing both relevant generic quality criteria and those that are specifically relevant to that domain. Although personal quality definitions can be created using the generic model, when a domain-specific model also exists this provides the consumer with a richer set of quality criteria from which to select when creating their personalised definition.

An example domain-specific quality framework for the domain of UK universities can be viewed in Burgess et al. (2003). Within the remainder of this paper, when demonstrating how these quality frameworks can be used to assist in the information searching process, we will primarily use the European new car domain, the complete framework for which can be viewed in Burgess (2003).

4 Experimental Information Search Environment (ISE)

To demonstrate how quality criteria can be used to assist in information searching, we developed the prototypical ISE. This proof-of-concept application illustrates how the information consumers can search for information within their current domain of interest by enabling them to create a personalised definition of quality, which can then be used to focus information search results.

In the following section we describe ISE, its user interface, the current level of information searching supported and conclude with planned future developments. We start, however, by discussing the data used for initial experimentation during the early stages of our research.

4.1 Experimental data

The ultimate aim of our research is to provide a facility to support information searching within a large, heterogeneous environment such as the internet, where information of various levels of quality is readily available and the potential exists for cognitive overload. However, to demonstrate the feasibility of our ideas initial experimentation was conducted in a closed-world environment. Data were initially obtained on a number of topics from a variety of online sources and stored locally. This method was followed for the following reasons:

- to guarantee availability of data for experimentation by eliminating the potential for network connection problems
- to ensure experiment consistency by using a constant, static data set.

A number of domains were selected for inclusion in the early experimentation stage of this research, including the two principal domains of UK universities and European new cars. The reasons for the selection of these two fields in particular include:

- data are freely available online for both domains
- data are provided by a number of information providers, thus resulting in richer and potentially conflicting, datasets.

One future aim of our research is to expand ISE further to use live online data across a larger, potentially unlimited, number of domains, once the feasibility of our approach has been demonstrated in a controlled environment.

4.2 System overview

The prototype ISE developed during our research currently supports the user when searching for tangible products or services, based on personal quality preferences. The following sections describe the current functionality of this system, how the consumer can create an explicit, personalised definition of quality which can then be used, and dynamically adjusted, to focus information searches within the experimental domains.

5 Conducting an information search

The first stage in conducting an information search within the prototype ISE involves the selection of the experimental subject domain within which the search is to be conducted. The domains currently available for selection are those for which data have been previously downloaded from online sources and for which a domain-specific quality framework has been created. Once a domain has been selected the user is presented with the appropriate quality framework from which they can select those criteria appropriate to their current requirements. This is achieved by navigating through the hierarchical structure and selecting desired quality criteria by using the appropriate interface options. The interface for this stage can be seen in Figure 2.

Figure 2 Selecting quality criteria in the car domain

Criteria Selection Process

Select desired criteria from within this framework

Highlighted criterion details

125 drivability

rating as to the drivability of a vehicle

Type of criterion: Positive

Mapping Information

TableName	FieldName	Rank
ParkersCarRatings	Driving	1
WhatCarRatings	RideAndHandling	2
WhatCarRatings	BehindTheWheel	1

Select this criterion

Selected Criteria

Name	Definition	State	Scale
drivability	rating as to the drivability of a vehicle	0	numeric
image	rating as to the image of a vehicle	0	numeric
build quality	evaluation of the quality of the manufacturing process	0	numeric
safety	abiding by industry safety standards	0	numeric

Get default criteria selection

Remove chosen criterion

Remove all criteria

Select a Different Domain

Continue

5.1 Selecting dimensions of quality

The current version of ISE allows the selection of individual quality criteria, but has not yet been extended to incorporate a facility for selecting quality dimensions (groups of related criteria). In a fully developed search system the user would also be able to select a dimension of quality and then state values for that dimension alone (rather than the lower-level individual quality criteria). Although the user then only sees that selected dimension, the system would be developed to allow the selection of one of three following options:

- using a small set of general quality criteria values related to that dimension to produce quick-and-dirty search results, providing a basis for the user to refine their search preferences and conduct further detailed searches
- the automatic selection of a set of predetermined default quality criteria relating to that dimension, rather than all sub-criteria, thus resulting in a more in-depth search, which can also be used as a basis for further user investigation
- the automatic selection of all criteria within the selected dimension, but hiding the details of those criteria from the user to ensure a simple search from the user's perspective.

In the current prototype system the only way to select a dimension of quality is to select all individual criteria within that chosen dimension. As the aim of this stage of our research is to demonstrate the feasibility of using quality criteria in an information search this limitation does not adversely affect the results of this work. Further development of our prototype application to allow quality dimension selections, and thus experimentation as to the feasibility of the three potential domain selection methods, is, therefore, an area pending further research.

5.2 Quality criterion importance weightings


Once desired criteria have been selected, each criterion must be rated according to its importance in the information search. In the current version of ISE the importance of each criterion is rated on a scale of 0...100, as illustrated in Figure 3.

Figure 3 Weighting IQ criterion importance



The use of sliders as a method for inputting the importance weightings of each criterion provides the users with a graphical representation of the values they select, resulting in easy comparison between multiple criteria preferences. This provides a clearer method for the user to perceive the comparative values of selected criteria weightings than if values are displayed textually, as illustrated in Figure 4.

Figure 4 Comparison between graphical and textual weighting value input methods



This method also gives a user the ability to change criterion weightings quickly and easily, thereby allowing them to experiment with different importance weightings to see how their chosen values affect the ordering of the returned results. The effect of using 'dynamic querying' sliders has been investigated by Ahlberg et al. (1992) and later by Schneiderman (1994) alone. Using dynamic querying in the three exploratory domains of the chemical table of elements, computer directories, and a real estate database, Schneiderman's experiments resulted in significantly improved performance and increased user enthusiasm.


5.3 Stating quality criterion preference values

The usual position when searching for information based on quality is to find the best available item that meets your needs. It may, however, be the case that the best item for the user is one that closely matches a predefined value, such as good, poor or average. If conducting research into a subject the user may be interested in finding information on

the item that performs worst on a particular criterion; for example, they may want to find the university with the lowest teaching quality. When wanting to purchase a tangible item, such as a car, the consumer may also have preferences regarding price and delivery time.

The use of sliders was also chosen for allowing the user to state desired values for selected quality criteria. For example, if setting a preference for the price to pay for an item the user can select a point on a scale, between the minimum and maximum values within a price range, as illustrated in Figure 5. By stating these preferences the set of results can be ranked appropriately by using these preferences, plus importance weightings, as a basis for the ranking. Future developments will incorporate a facility to input user-stated ranges within which values will be acceptable, plus the automatic relaxation of values based upon whether values are constraints or preferences.

Figure 5 Selecting a value between available minimum and maximum values



Occasionally, a preference scale from 0 to 100 is not appropriate. In these instances other methods of stating preferences may be used, such as the excellence scale illustrated in Figure 6.

Figure 6 Example of a non-numeric ‘excellence’ scale

This non-numeric scale option is not yet implemented within ISE as the current sliding scales are adequate for the purpose of experimentation.

5.4 *Quality criterion options*

In some circumstances it is not possible to state an importance rating for a chosen quality criterion. For example, if looking for information on cars, one of the qualities the user may consider as being important is colour. Simply placing an importance weighting on such a quality would be meaningless, without stating which colour is actually preferred. In this instance instead of presenting the user with a weighting facility, a set of available options also needs to be presented to the user, allowing them to state which option is the most desirable. By entering importance values for the available options the user can state which are of most importance to them and therefore, should be given preference when ranking the search results. They may wish to state that two or more options are of equal

importance by giving them the same value, state a dislike for an option by entering zero, or leave blank to indicate no preference. An example of this type of option is shown in Figure 7.

Figure 7 Example options for the 'colour' criterion

Available	<input checked="" type="checkbox"/> 3	Red	<input checked="" type="checkbox"/> 2	Black
Colours	<input type="checkbox"/>	Blue	<input type="checkbox"/> 0	White
	<input type="checkbox"/>	Green	<input type="checkbox"/> 1	Silver

Stating preference values for available options does not however mean that the most highly ranked results will exactly match this requirement. All user preferences are taken into consideration when calculating overall personalised quality values for the available data, but undesired values, such as 'white' cars in the above example, may still be ranked relatively highly if there are high quality values for the other criteria the user has chosen. This could occur, for example, if the user states low importance weightings for criteria with undesired values.

The facility for selecting quality criterion options, where appropriate, is planned to be incorporated into the next version of ISE, once our initial research premise, the feasibility of using quality criteria to focus information searches, has been fully explored.

5.5 Source quality ratings

When data are obtained from multiple sources the quality of these sources should also be considered, as this will affect the perception of the quality of the information they provide. For example, information from a source considered to be of high quality is likely to be preferable to information obtained from a poor quality source.

Previous projects have investigated methods for ascertaining source quality, such as that conducted by Naumann (2002). This subject has, therefore, not been selected as a focus for our research, as the methods used for ranking information based on quality criteria can be combined with other research into quality assessment of information sources. As the data selected for use in initial ISE experimentation were obtained from a number of third-party organisations, are stored locally, and were cleaned before being used (e.g., by the elimination of synonyms), the quality of the various data sources is not an issue in this project, and therefore does not affect the validity of the results of this stage of our research.

However, within ISE a facility does exist for explicitly rating the quality of the information provider. Instead of being an accurate representation of the quality rating of the available suppliers, this facility allows users to state their supplier preferences by making these preferences explicit. For example, a user may have a preference for one information supplier over another, and therefore, would rate his/her preferred supplier higher than the others. This also allows for the exclusion of suppliers from an information search, as setting a value of zero for any source results in no data from that source being incorporated into the quality score calculations. The interface for setting these supplier preferences is shown in Figure 8.

Figure 8 User stated source quality ratings for the car domain

The screenshot shows a window titled "Information Supplier Preference Ratings". Inside, there is a table with the following data:

SupplierID	Supplier Name	Current Rating
4	WhatCar	1
5	Perkers	1
6	EuroNCAP	1

Below the table, the "Selected Supplier" is "EuroNCAP". The "Current preference rating" is shown as "1" in a text box. There is an "Enter new preference rating:" text box. An "Update Rating" button is located to the right of the text boxes. A "Done" button is at the bottom center of the dialog.

5.6 *Quality-driven information searching*

Based on the selection of quality criteria, stated importance weightings and preference values for each criterion and information about supplier preference ratings, quality scores are calculated for each of the available data items in the chosen domain. Currently, all data stored for experimentation are ranked according to quality preferences. However, future developments in this area are planned, and include the following:

- Taking the results obtained from an information search as the base data (such as the results obtained from an internet search or metasearch engine), which are then ranked by the most appropriate ranking algorithm.
- Using the quality preferences to develop an advanced information search statement to augment the typical keyword-based search statement, the results of which are then ranked, again by means of the most appropriate ranking algorithm.
- Using quality preferences to filter out irrelevant search results, producing a smaller more focused data set. These results can then be ranked according to user preferences as before. Using quality preferences during these two stages of an information search will result in both a ranked and smaller data set, thus further reducing information overload.

Once quality values have been calculated by the chosen ranking algorithm for each data item or search result, the results are then presented to the user, positively ranked with those most closely meeting the user's requirements coming higher in the ranking order. This is achieved by applying one of the ranking algorithms incorporated within ISE.

5.7 *Ranking algorithms*

Three ranking algorithms are currently implemented in ISE: Simple Additive Weighting (SAW), Technique for Order Preference by Similarity to Ideal Solution (TOPSIS) and TOPSIS with Given Preferences (TOPSIS-GP). These are described briefly below.

SAW is the simplest and best known of the three algorithms. It orders results by multiplying each criterion value by the corresponding weighting value, then summing all weighted criterion values for each data item. The items are then ranked according to this final value.

TOPSIS and TOPSIS-GP are more sophisticated in their approach by ranking items according to how close they are to some ideal solution: ‘best’ value for positive criteria (where a high value is preferable over a low value, such as for ‘reliability’), and ‘worst’ value for negative criteria (where a low value is preferable over a high value, such as for ‘access time’). The original TOPSIS algorithm assumes the best and worst values for each criterion. However, in TOPSIS-GP the users can themselves state their preferred values, meaning that results are ranked according to their personal preferences rather than an assumed preference.

For further information about these algorithms the reader is referred to other literature published on this topic. A detailed overview of SAW and TOPSIS is available in Hwang and Youn (1981). Full details on TOPSIS-GP can be found in Burgess et al. (2006b).

When using ISE it is possible to state which ranking algorithm should be employed, thus enabling comparisons between the different methods. However, if no specific selection is made the ranking algorithm automatically selected will depend on how the user has decided to specify their quality requirements. If the user has opted to state preferred values for their selected quality criteria then TOPSIS-GP will be used. Otherwise, the traditional version of TOPSIS is automatically invoked.

5.8 Search results

Figure 9 shows the current ISE search results interface, displaying the results of the selected ranking method. All available information items are presented to the user via this screen. For large data sets the majority of the items will not be immediately visible, but can be accessed via the scroll bars. However, due to users typically only wishing to know about those data items that closely meet their quality requirements only those present at the top of the result list are likely to be of interest.

Figure 9 Search results presented in ISE in the example domain of European new cars

The screenshot shows a software window titled "Search Process Form" with two main panes: "Search Type" and "Search Results".

Search Type:

- Buttons for "SAW" and "TOPSIS" are present, with "TOPSIS" selected.
- Under "Scaling Method", there are radio buttons for "Vector scaling" (selected) and "Linear scaling".
- A "Run TOPSIS Search" button is located below the scaling options.
- The "TOPSIS Notes" section contains a scrollable list of text:
 - Negative Ideal for metric 1 = 0
 - Negative Ideal for metric 2 = 0
 - Negative Ideal for metric 3 = 0
 - Negative Ideal for metric 4 = 0
 - Calculating euclidean distances
 - Calculating closeness to ideal solution
 - Search complete
- At the bottom of the "Search Type" pane, there are radio buttons for "Simple" (selected) and "Advanced", and a "Display TOPSIS Results" button.

Search Results:

A table with two columns: "Name" and "Result". The results are as follows:

Name	Result
BMW: 318i SE 4 door	000.86449
BMW: 520i SE 4 door	000.86187
Saab: 9-3 2.0i	000.85737
Audi: A4 2.0 SE 4 door	000.85277
BMW: X5	000.81168
Volvo: S60 2.0 T	000.80250
Audi: A3 1.6i 3 door	000.79702
Ford: Focus 1.8 LX 5 door	000.79659
Volvo: S80 2.4s	000.79029
Audi: A2 1.4	000.78518
Volkswagen: Polo 1.2 S 3 door	000.78497
Lexus: LS430 4.3	000.78366
Vauxhall: Vectra 1.8	000.78197
Mercedes-Benz: E240 Elegance 4 door	000.77794
Renault: Laguna 1.6 16v Expression 5 door	000.77639
Skoda: Fabia 1.4 16v	000.77633
Ford: Ka	000.77462
Jaguar: X-Type 3.0 SE	000.76801
Audi: A6 2.8 v6 Quattro 4 door	000.76507
Renault: Megane 1.6e sport 5 door	000.75023
Fiat: Multipla 100 ELX	000.74950
Mercedes-Benz: C200 Kompressor Elegance 4 door	000.74913
Peugeot: 307 1.6	000.74866
Toyota: Corolla 1.6 T3 5 door	000.74461
Seat: Leon 1.6s 16v	000.73737
Honda: Accord 2.0 SE 5 door	000.73666
Peugeot: 307 SW 1.6	000.73319
Toyota: Prius	000.72640

At the bottom of the window, there are four buttons: "Refine This Search", "Start New Search", "Compare Search Results", and "Close".

With the inclusion of a filtering stage as part of the information searching process a large number of data items can be removed from the final set of results. These eliminated data items would include those with low quality scores, with undesirable quality criterion values, or information from suppliers undesired by the user. The challenge of finding the most appropriate method for information filtering is an area requiring further work.

The results presented in ISE currently contain all available data items for the selected domain, plus more information than is needed by the average information searcher. This extra information, including values for each quality criterion at each stage of the TOPSIS ranking method, is presented for the purpose of insight into the prototype system. The information presented to the user of the final system would therefore be clearer, and of a less detailed nature.

5.9 Dynamic searching

Schneiderman (1994), a leader in the field of Human-Computer Interaction, conducted research into the area of dynamic querying and concluded that when users are able to dynamically explore an information space their experience of the process is significantly improved. This idea has therefore been incorporated into ISE, via the inclusion of sliders for stating criteria importance weightings and preference values. The inclusion of these sliders allows users to visually state their quality preferences. After viewing the results of the search based on their initial quality settings they can return to the criteria settings and experiment with changing quality ratings. This dynamic approach to querying a set of data items allows for search criteria to be adjusted until the user is satisfied with the final result set.

Alongside dynamic refining of current search criteria, all previous information searches can be revisited and refined. A facility is provided whereby the user can view search history data and select a specific previous information search. They are then presented with the quality-driven search parameters which can be adjusted as required and further searches conducted.

6 Evaluation and results

The development of ISE has shown that quality criteria can be used to drive an information search in a variety of subject domains. Quality scores for individual data items are calculated according to the set of quality criteria selected for use during the current search, plus the setting of quality criteria importance weightings and preference values. ISE has also shown that it is possible to state preference values for information suppliers, when data have been obtained from more than one source. This results in a favouring of data from preferred sources, a reduction in the importance of data from others, and the potential for elimination of data sources that are not desired.

The principal aim of creating such a model of quality was to assist users in searching for information by focusing search results based upon their personal quality preferences. Experimentation was therefore conducted whereby different quality requirements were defined and used within ISE to determine whether the ranked result sets significantly differed. If a statistically significant difference can be shown between the ordering of results based on two different user-defined quality search preferences it will confirm that

changing the personalised quality definition can result in a changed focus in the search results.

Search simulations were therefore conducted to discover whether the ranking order of search results can significantly vary in the following situations:

- changing the selected quality criteria
- changing the importance weighting for a constant set of quality criteria
- changing preference values for a constant set of quality criteria
- changing both importance weightings and preference values for a constant set of quality criteria.

The results of these simulations showed that by altering the personalised definition of quality, very high statistical significance can be seen in the ranking order of the search results, thus demonstrating that quality settings can alter the focus of an information search. This is not always the case, as minor changes in quality preferences will not have as great an impact on altering the focus as a large change. Neither will some larger changes, such as when all selected criterion weighting values are changed, but still remain identical when compared with the other criteria. For example, when searching for information in the university domain the user may select to search based upon the two criteria of 'Research Quality' and 'Teaching Quality'. Placing importance weighting of *high* and *very high* respectively will result in an identical result ranking order as rating them as *low* and *average*. In this example, although the weighting values have changed they remain in the same preference order with respect to each other – 'Teaching Quality' being rated more highly than 'Research Quality'.

A fuller description of the experiments conducted to validate this work can be found in Burgess et al. (2006b), where there are further details of these experiments, the method used to measure statistical significance, and sample experiment settings with statistical evaluations.

7 Conclusion and future developments

As stated previously, the purpose of this research was to ascertain the feasibility of using a personalised, explicit, definition of quality to assist in information searching. By using this definition to calculate quality values for each item in the dataset, or set of search results, the chosen ranking algorithm can order the results accordingly and thus provide the searcher with an ordered set of results focused on their personal quality requirements.

Having demonstrated that quality criteria can be used to focus information search results within a closed-world environment using our prototype ISE we can now proceed to investigate a number of further issues. Some of the areas planned for further exploration are listed below:

- Continued experimentation using the generic framework to search for information when domain-specific frameworks are not available. The development and use of domain-specific quality frameworks is useful and feasible when searching within a specific subject area. However, this would not be appropriate when conducting a search in a heterogeneous environment, such as the internet.

- Removal of the closed-world restriction to support online information searches, initially concentrating on how user-defined quality preferences can be used to focus searches conducted via current search services such as Google and Dogpile.
- Using quality to assist in information searching within the healthcare domain. A great deal of information is available online for patients, their families and their caretakers, but this is also at various levels of quality – such as incorrect, unavailable, or presented at the wrong level, thus not meeting the needs of the user. Further details on our proposal for research in this domain can be found in Burgess et al. (2006a).
- Experimentation into combining filtering and ranking. At present, all data items or search results are rated and ranked according to the full set of user quality preferences. However, with large result sets this would not be feasible. Initial filtering of data to remove those items which do not meet the user's quality requirements would increase the speed of ranking and produce a smaller, more focused set of results. However, this requires further examination to ascertain the most appropriate approach to filtering, as if the user has stated a set of conflicting preferences this initial filtering could potentially lead to an empty result set.
- Automatic relaxation of quality preferences. The previous point raises the issue that users may specify conflicting values for quality criteria; for example, it would not usually be possible to obtain a car from a highly prestigious manufacturer for less than \$1000. Therefore, a method for automatically relaxing user quality preferences is required to increase their chance of finding information that meets their needs as closely as possible.
- The incorporation of previously developed methods for obtaining user feedback (or attention metadata), either explicitly or implicitly, in order to learn quality preferences and assist future users by suggesting potential quality 'profiles'.

As can be seen in this non-exhaustive section, a number of areas still require further investigation and experimentation. By continuing on this path we aim to help the information consumer combat the two principal online search hindrances of 'information overload' and 'information quality', in order for them to find the information that best meets their needs in our ever growing information society.

References

- Ahlberg, C., Williamson, C. and Schneiderman, B. (1992) 'Dynamic queries for information exploration: an implementation and evaluation', *Proc. ACM Conference on Human Factors in Computer Systems*, Monterey, California, USA, pp.619–626.
- BBC (2003a) *Skin Cancer Web Sites 'Misleading'*, BBC News, <http://news.bbc.co.uk/1/hi/health/1730436.stm>.
- BBC (2003b) *Student Died after Buying Web Drugs*, BBC News, <http://news.bbc.co.uk/1/hi/england/3130187.stm>.
- Bouch, A. (2000) 'Quality is in the eye of the beholder: meeting users' requirements for internet quality of service', *Proc. CHI 2000 Conference on Human Factors in Computing Systems (CHI 2000)*, The Hague, Netherlands, pp.297–304.

- Burgess, M.S.E., Gray, W.A. and Fiddian, N.J. (2002) 'Establishing a taxonomy of quality for use in information filtering', *Proc. 19th British National Conference on Databases (BNCOD 19)*, Sheffield, UK, pp.103–113.
- Burgess, M.S.E. (2003) *Using Multiple Quality Criteria to Focus Information Search Results*, PhD Thesis, School of Computer Science, Cardiff University, Wales, UK.
- Burgess, M.S.E., Gray, W.A. and Fiddian, N.J. (2003) 'A flexible quality framework for use in information retrieval', *Proc. 8th Int. Conf. on Information Quality (ICIQ-2003)*, Cambridge, MA, USA, pp.297–313.
- Burgess, M.S.E., Allam, O., Gray, W.A. and Al-Busaidi, A. (2006a) 'Patient-centred approach to focusing online health information search results', *Proc. Int. Symposium on Electronic Personal Health Records 2006 (ISePHR06)*, Trondheim, Norway, pp.5–16.
- Burgess, M.S.E., Gray, W.A. and Fiddian, N.J. (2006b) 'Quality measures and the information consumer', in Al-Hakim, L. (Ed.): *Challenges of Managing Information Quality in Service Organizations*, Idea Press Group, Hershey, USA, pp.213–242.
- Crosby, P.B. (1979) *Quality is Free: The Art of Making Quality Certain*, McGraw-Hill, New York, USA.
- Dromey, R.G. (1996) 'Concerning the chimera', *IEEE Software*, Vol. 13, No. 1, pp.33–43.
- Eppler, M.J. (2001) 'A generic framework for information quality in knowledge-intensive processes', *Proc. Int. Conf. on Information Quality (ICIQ 2001)*, Boston, MA, USA, pp.329–346.
- Firquin, B.L. (1992) 'Quality is free ... but how do you implement total quality and restructure information technology at the same time?', *Proc. 20th (SIGUCCS) User Services Conference*, New York, NY, USA, pp.93–96.
- Google (2006) *Google*, <http://www.google.com/>.
- Graham, L. and Metaxas, P.T. (2003) 'Of course it's true; I saw it on the internet!': critical thinking in the internet era', *Communications of the ACM*, Vol. 46, No. 5, pp.71–75.
- Hwang, C.L. and Youn, K. (1981) *Multiple Attribute Decision Making: Methods and Applications*, Springer-Verlag, New York, USA.
- Infospace (2006) *Dogpile Web Search*, <http://www.dogpile.com/>.
- ISC (2006) *Internet Systems Consortium, Inc.*, <http://www.isc.org/>.
- Juran, J.M. and Godfrey, A.B. (1999) *Juran's Quality Handbook*, 5th ed., McGraw-Hill, New York, USA.
- Knight, W. (2005) 'Info-mania' dents IQ more than marijuana', *New Scientist*, *NewScientist.com News Service*, <http://www.newscientisttech.com>.
- Naumann, F. (2001) 'From databases to information systems information quality makes the difference', *Proc. Int. Conf. on Information Quality (ICIQ 2001)*, Boston, MA, USA, pp.244–260.
- Naumann, F. (2002) *Quality-Driven Query Answering for Integrated Information Systems*, Springer-Verlag, Heidelberg, Germany.
- Page, L., Brin, S., Motwani, R. and Winograd, T. (1998) 'The PageRank citation ranking: bringing order to the web', *Stanford Digital Library Technologies Project*, Vol. 5, p.26.
- Pirsig, R.M. (1974) *Zen and the Art of Motorcycle Maintenance*, Bantam Books, New York, USA.
- Schneiderman, B. (1994) 'Dynamic queries for visual information seeking', *IEEE Software*, Vol. 11, No. 6, pp.70–77.
- Spinellis, D. (2003) 'The decay and failures of web references', *Communications of the ACM*, Vol. 46, No. 1, pp.71–77.
- Sullivan, D. (2001) *WebTop Search Rage Study*, <http://searchenginewatch.com/showPage.html?page=2163451>.
- Wang, R.Y. and Strong, D. (1996) 'Beyond accuracy: what data quality means to data consumers', *Journal of Management Information Systems*, Vol. 12, No. 4, pp.5–34.