

Alice: Assisting Online Shoppers through Ontologies and Novel Interface Metaphors

John Domingue¹, Maria Martins², Jaicheng Tan¹, Arthur Stutt¹ and Helgi Pertusson^{3,4}

¹The Knowledge Media Institute
The Open University, Milton Keynes, UK
Tel: +44 1908 655014, Fax: -3169
{J.B.Domingue, J.Tan, A.Stutt}@open.ac.uk
<http://kmi.open.ac.uk/>
³Innn Inc., Laugavegur 26, 101 Reykjavik, Iceland
<http://www.innn.com/>

Abstract. In this paper we describe some results of the Alice project. Alice is an ontology based e-commerce project which aims to support online users in the task of shopping. Ontologies describing customers, products, typical shopping tasks and the external context form the basis for the Alice architecture. We also exploit two novel interface metaphors originally developed for navigating databases: the Guides metaphor and Dynamic Queries. The Guides metaphor was developed at Apple to reduce the cognitive load on learners navigating a large hypermedia database. Within Alice we use the Guides metaphor to allow online shoppers to classify themselves. We discuss the link between Alice Guides and Kozinet's notion of e-tribes or Virtual Communities of Consumption. Our second interface metaphor Dynamic Queries (coupled with Starfield displays) allow users to very quickly find relevant items by displaying the results of queries, posed via specialised slider widgets, within 100 milliseconds. We have constructed a tool, Quiver, which constructs Dynamic Query interfaces on-the-fly as the result of queries to knowledge models stored on the Alice server.

Introduction

Currently shopping on the internet is not always a pleasant experience. Navigating websites with thousands of products by browsing virtual aisles or by keyword search is time consuming and often frustrating. Each aisle will typically contain hundreds of items that are hard to differentiate. The shopper has to rely on a product's name and sometimes on a small accompanying picture. Neither of which are particularly descriptive. Keyword searches over generic product types (e.g. flour) will often return hundreds of irrelevant items (e.g. wholemeal flour bread). Online shopping websites also contain a lot of irrelevant information related to new types of products or reduced items.

² Current contact details: Business School, The University of Gloucestershire, Cheltenham, UK
Email: mmartins@glos.ac.uk

⁴ Current contact details: cTarget Inc., Brautarholt 1, 105 Reykjavik, Iceland.
Email: hp@ctarget.com

Contrast the above with the local ‘corner shop’ which was prevalent in villages in England in the 1950s. Of course there are a variety of differences between a corner shop and an online shopping site. These include the fact that one is physical and the other virtual. Also corner shops sometimes have a better layout. Neither the less we believe that one of the key differences to the customer’s shopping experience was due to the fact that the shop had a human agent, the shopkeeper, who used his or her knowledge to *personalise* the interaction. Typically, the shopkeeper would know which products were currently in stock and products that could easily be obtained. Additionally, the shopkeeper understood the relationships between the products, for example, when one product could be substituted for another (out of stock product), or how one product complemented another (e.g. a particular cheese and wine combination). Regular customers would also benefit from the fact that their personal tastes and preferences, their current situation (e.g. number of dependents), and their previous purchases were known to the shopkeeper. The shopkeeper was also able to relate desired products to the local context including the surrounding geography and community and the resident culture and events.

The overall goal of the Alice project is to make the experience of online shopping seem more like visiting a local corner shop than browsing or searching long lists. In the rest of this paper we will describe some of the results of the project structured in the following fashion. In the next section of the paper we describe the Alice approach. We then illustrate the approach through a short scenario. The subsequent two sections describe the architecture of the system and an interface for detecting the patterns of behaviour of online customers. Finally, a discussion section, linking the Alice approach to a notion of e-tribes, is followed by some conclusions.

Approach

The Alice framework is based on the use of ontologies for representing knowledge related to online shopping. An ontology [13] is an explicit representation of a view of a domain of discourse (a conceptualisation) usually composed of a set of concepts and relationships. Over the last few years the use of ontologies has become relatively popular, for example a web search for ontology will now return more than 64,000 web pages [14]. Moreover, ontologies are widely deployed within the knowledge acquisition and modelling communities, have been successfully used in a variety of web based applications (e.g. [7]) and form one of the cornerstones of the semantic web [3]. Within Alice we use five ontologies to create a personalised online shopping experience. The five ontologies are:

- *Products* – this ontology describes the main attributes of products, for example, how a product is used, its components, complementary products and a product’s geographic origin.
- *Shopping Tasks* – this ontology represents typical shopping tasks, for example a monthly shop for household essentials and shopping for an evening meal.
- *External Context* – appropriate items from the local context are described within this ontology. For example, relevant local social events, groups, and small businesses.
- *Customer* – this ontology represents the main attributes of a customer including his or her shopping and browsing histories.
- *Alice Media* – this ontology maps between the other four ontologies and relevant web resources.

The first four ontologies to a greater or lesser degree reflect the four categories of knowledge used by a local corner shopkeeper. Our definition of a product borrows from the product ontology available on the ontolingua server. Specifically, we use the relations list-price, has-model-number and has-special-discount.

In contrast with the other four ontologies, modelling customers within a formal representation is particularly problematic. This is for a number of reasons. Firstly, there is no definitive knowledge source for classifying customers' according to their shopping behaviour. There exist many competing marketing and economic models of consumer behaviour, but none of these give a 'foolproof' account of why and when humans purchase goods. Second, a customers' behaviour will vary depending on his or her current situation, for example, buying a single urgently needed item vs. buying goods for a week, or having a personal cash flow problem until the next salary payment. Major life events, such as having a baby, will also dramatically affect behaviour.

Determining a customer's current state is also non trivial. Two options are to infer the customer's state from their interactions or to explicitly ask the customer. The former option is prone to error because the individual user interactions - selecting hyperlinks or options from menus - contains little informational content. Whilst accurate the latter option has to be carefully applied because customers are, in general, unwilling to spend significant amounts of time on any task that does not have an immediate benefit.

An additional factor linked to the above is that it is imperative that any online system only offers pertinent advice. Offering a customer advice based on an incorrect model would result in the system being quickly discarded.

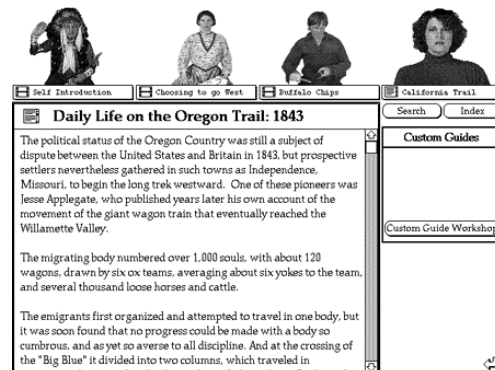


Fig. 1. A screen snapshot of the Guides system (taken from <http://www.abbedon.com/project/guides.html> with permission (copyright 1990 Apple Computer)). The first three Guides above explain early American history from the perspective of: a native American, a female settler, and a frontiersman. The last icon represents the system Guide who gave overview information.

Our approach within the Alice project has been to use the Guides metaphor [8, 18, 26] as a mechanism to allow customers to classifying themselves. Guides were produced at Apple in the late 1980s as an interface for an educational hypermedia database depicting early American history. The Guides who were characters drawn from this period, delivered stories from specific viewpoints. Each story consisted of a series of video clips. A screen snapshot of the Guides system can be seen in Fig. 1. Four Guides are shown at

the top of the screen. The first three deliver stories on early American history from the viewpoint of a native American, a female settler and a frontiersman. The icon on the far right represents the system Guide who delivered overview information.

Within Alice we decided to use the Guides metaphor to enable customers to state their shopping preferences. We shall give an overview of our implementation using a short scenario.

A Scenario

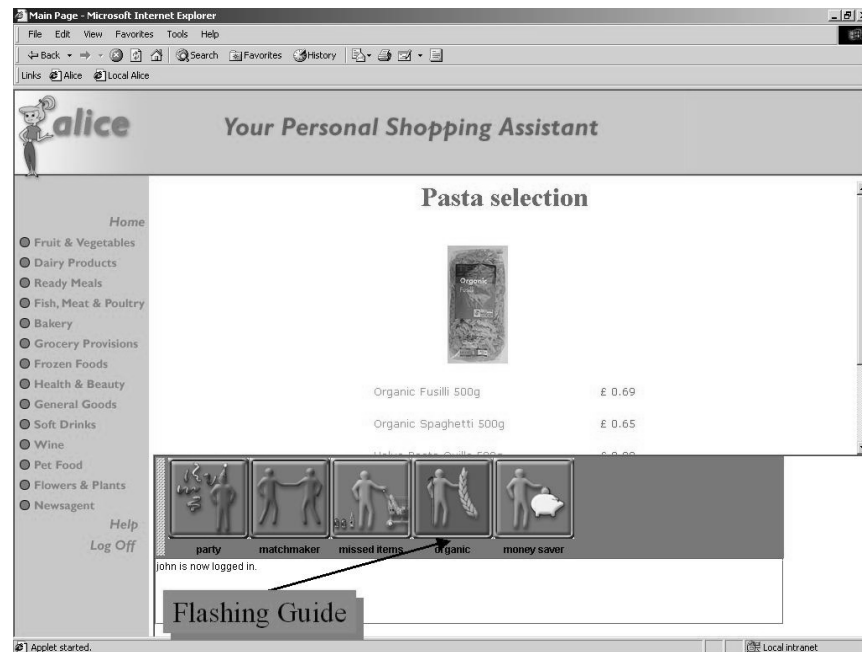


Fig. 2. A screen snapshot of the Alice Guides interface. The left panel contains a standard navigation bar as found in most online supermarkets. The products are shown in the large pane in the centre of the browser. The Guides interface is shown in the bottom panel. The customer is browsing a selection of pastas and the *Organic Guide* is indicating that it would like to start a dialogue by blinking red.

In the following scenario an online shopper is looking to buy some pasta within a fictional 'Alice Supermarket'. A screen snapshot from the shopper's web browser is shown in Fig. 2. The navigation menu on the left and the product display area in the centre of the window are similar to those found in most online supermarkets. The Alice Guides interface is contained in the panel at the bottom. The shopper has selected five Guides to go shopping with:



Party Guide – this Guide assumes that the task for the session is to buy products for a party. The Guide prompts with related offers (e.g. the free loan of wine glasses), recipes and local services (e.g. marquee hire).



Matchmaker Guide – this Guide matches products that are purchased to similar or complementary products. For example, it would match pasta with bottled pesto sauce.



Missing Items Guide – we have found that shoppers will sometimes forget to select the ‘Add to Basket’ button and consequently fail to purchase a desired product. When the customer goes to the checkout this Guide collects a list of items that the customer browsed in detail but did not add to his or her basket.



Organic Guide – when appropriate this Guide recommends organic versions of goods that are being viewed.



Money Saver Guide – this Guide informs the shopper of any offers or promotional items which are related to the currently viewed item.

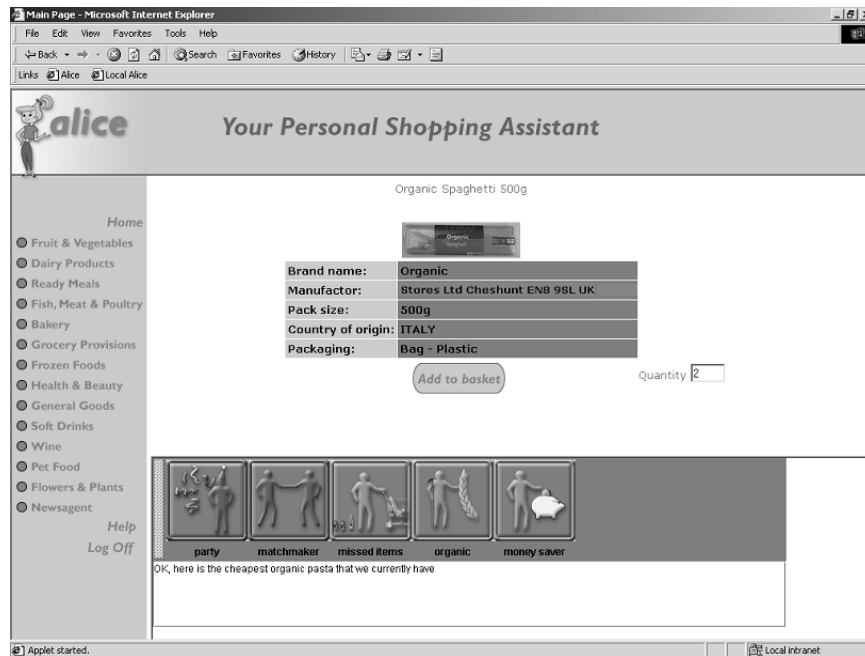


Fig. 3. A screen snapshot just after the *Organic Guide* has displayed the cheapest organic pasta.

In Fig. 2 the shopper is browsing the pasta section of the online store. The *Organic Guide* indicates that it has something to contribute by blinking red a number of times (the Guide’s normal colour is blue). The shopper is free to ignore the Guide and to carry on browsing but she elects to see what the Guide has to say and selects the *Organic Guide* icon. The *Organic Guide* offers the cheapest organic pasta. The shopper agrees and the display changes to Fig. 3. Note that hundreds of potential items (a well known online store

methods, and roles. Export mechanisms exist from OCML to Ontolingua [13], XML [34], RDF [33], and GXL [15].

The ontologies are split into three levels: retail, retail sector and client. At the retail level there are the five ontologies described earlier. Each of these five ontologies contain definitions which are applicable across the whole retail industry. Below is the retail sector level where definitions applicable to specific retail sectors are stored. For example, for childrens' toys important attributes would be age-range and educational value.

The client level knowledge models represent specific companies. These models would contain mappings from the generic Alice models to the existing corporate resources. For example, to the company's database schemas. The client specific models are also used to link to the client web server via the *Alice Client Plugin* module. This module communicates with the *Alice Plugin Interface* via a set of HTTP like messages. The server responds with a message that is either plain text, HTML, a list or a set of attribute value pairs. The number and type of messages is set by an XML based configuration file. An implementation of the plugin module exists in PHP and future implementations are planned for Java, Perl and active server pages.

The customer's interactions with the web based interface are sent via the *Alice Plugin Interface* to a customer history. The customer history is used in two ways. Firstly, the products browsed and purchased are asserted as facts within the company's specific knowledge model and are used to trigger a customer's Guides. The customer history is also used by a module which clusters the history according to customer and product attributes. The results are fed through to a manager's interface allowing a company's sales and marketing departments to discover new relationships between products and new clusters of customer behaviour. When appropriate new knowledge gleaned from the results of clustering will be used to create new Guides. This module is under construction and will be based on an unsupervised clustering technique [4].

The knowledge models are created and maintained using WebOnto [6]. In addition to its use in over a dozen projects within our lab WebOnto has been available as a public service since autumn 1999. The public library contains over a hundred models and has just over 150 registered users. In a comparative evaluation of several knowledge modelling tools WebOnto was evaluated very favourably, in particular being judged as the most user-friendly and as the one requiring the shortest learning curve [11].

```
(def-guide-trigger (amount-trigger party-guide)
  "Check that the customer buys the required amount of food
  for a party"
  (current-customer ?customer)
  (user-name ?customer ?user-name)
  (has-profile ?customer ?profile)
  (has-history ?customer ?customer-history)
  (party-number-of-guests ?profile ?n)
  (last-item-bought ?customer-history ?item)
  (item-name ?item ?item-name)
  (amount-too-low ?item ?n)
  :action
  (low-amount-for-party ?user-name ?item-name ?n))
```

Fig. 5. The definition of an amount trigger for the party Guide.

The Guide module contains the server part of the Guide system which is implemented on top of the OCML forward chaining system. Guides have a set of associated triggers and actions. Triggers enable a Guide to be activated when certain conditions occur. The definition of a trigger contains a set of clauses and an action. When the clauses match the

contents of the current knowledge base the action is invoked. An example trigger, amount-trigger, for the party Guide is shown in Fig. 5. Amount-trigger invokes the low-amount-for-party action if the customer buys insufficient quantity of a product to satisfy the specified number of guests at a party. Actions provide a high level mechanism for defining how a Guide will interact with a customer.

The Guide Applet sits within the supermarket's existing online shopping interface. Communication between the Guide server and the Guide applet is via a CORBA interface. The main types of messages defined including logging in, and adding, removing and alerting Guides. The underlying infrastructures for the server and client, Xanalys LispWorks™ and Java v.1.4, and have inbuilt CORBA interfaces.

The Manager's Interface

Once the Alice system has been installed within a company managers will need to analyse customers' browsing and shopping behaviours in order to identify new types of products and customer characteristics. These characteristics may then lead to changes to the website including the design and creation of new Guides. From a marketing perspective, following the Pareto rule of 80-20, the overall goal of any analysis is to determine the significant attributes with respect to the 20% of customers who purchase the 80% of products. For example, 16% of US beer drinkers account for 88% of annual consumption [17].

Within Alice we have created a visual query tool, Quiver, to support the analysis of shopping behaviour through the Alice ontologies. Quiver couples Dynamic Query and Starfield like interfaces [1, 30] to our ontology server. Dynamic queries and starfield displays were developed within the Human Computer Interaction Lab at the University of Maryland in the early 1990s. A number of control widgets – sliders, checkboxes and buttons – generate queries in real time to a database. The results of each query is presented in a specialised two dimensional graphical display, termed a starfield display, within 100 milliseconds. The tight coupling of the widgets generating database queries to the graphical display enables users to quickly navigate large data stores. Quiver, creates a two dimensional graphical display and coupled sliders from a query to a knowledge model held on the Alice server.

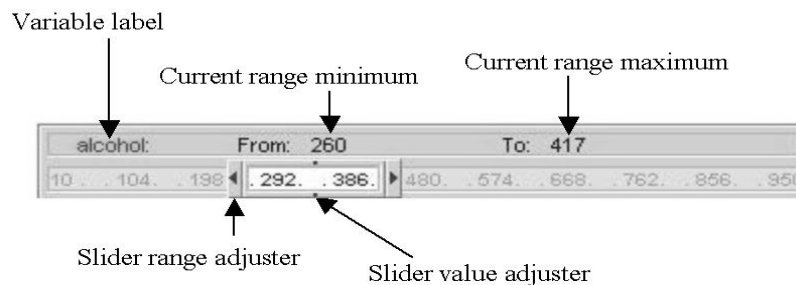


Fig. 6. An annotated screen snapshot of the Quiver slider for the variable 'alcohol'. The slider filters the graphical display area to only show those items whose alcohol rating is between 260 and 417.

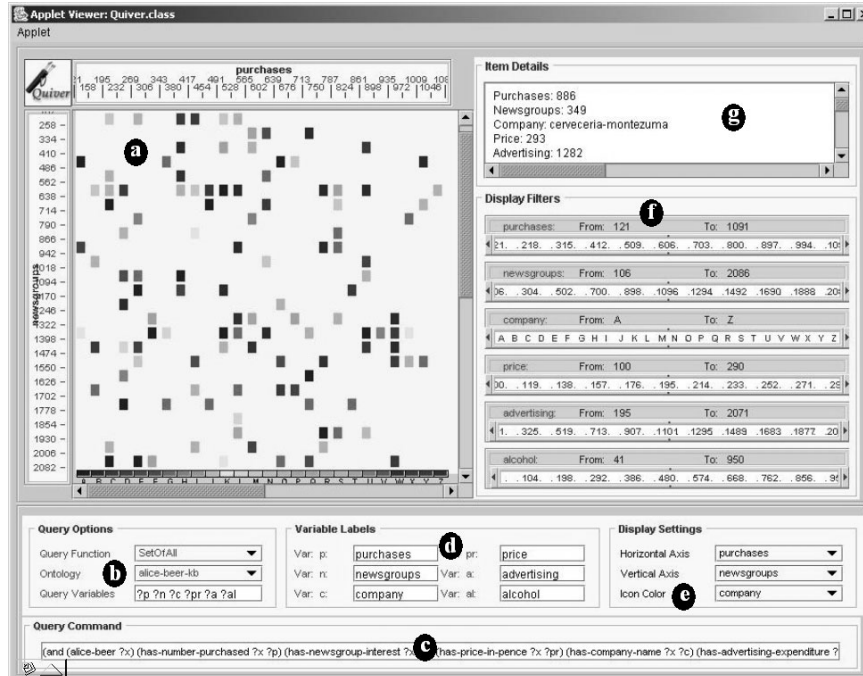


Fig. 7. A screen snapshot of Quiver displaying items from a beer knowledge base. The full query which generated this snapshot is shown in Fig. 8.

We shall now describe Quiver through a mini scenario. Imagine that a manager responsible for the marketing of wines and spirits has decided to investigate the possible influences on the sales of bottled beer over the last week. In particular, she wants to explore the relationships between:

- The number of bottles of beer sold over the last week,
- The number of mentions for a beer within related newsgroups,
- The beer producer,
- The retail price for the beer,
- The amount of advertising expenditure, and
- The beer's alcohol rating.

Assuming that the above is stored within a knowledge model on the Alice server then the manager can use Quiver to explore the above relationships. A screen snapshot of the Quiver interface shown is in Fig. 7. This was created in the following way. First the manager selected the query function *SetOfAll*, the ontology *alice-beer-kb* and specified the query variables (*?p ?n ?c ?pr ?a ?al*) using the *Query Options* panel (b). Then using the *Variable Labels* panel (d) new names (*purchases*, *newsgroups*, *company*, *price*, *advertising* and *alcohol*) were given to the variables. The query was then entered into *Query Command* panel (c). The display then took on the appearance of the screen snapshot in Fig. 7. The manager wanted to see if the number of purchases of

medium to low alcohol beer was linked to high advertising expenditure. She set the sliders in the *Display Filters* panel (f) to reflect her interests.

Let us now examine the Quiver interface shown in Fig. 7 in more detail. The large pane on the left, (a), contains a graphical display. Each icon in this display represents an item returned from a query. The other panels have the following functionalities:

- (b) *Query Options* – this panel enables the manager to set the query function, the target ontology and the query variables to display.
- (c) *Query Command* – the command is entered in OCML syntax in this panel.
- (d) *Variable Labels* – using this panel the manager can attach arbitrary labels to the variables specified in the *Query Options* panel.
- (e) *Display Settings* – this panel is used to specify which variables correspond to the horizontal and vertical axes and to the icon colour.
- (f) *Display Filters* – the sliders generated for each variable specified in the *Query Options* panel are shown in this panel.
- (g) *Item Details* – when an icon within the graphical display is selected detailed information about the item is displayed in this window.

```
(setofall ' (?p ?n ?c ?pr ?a ?al)
  '(and (alice-beer ?x) (has-number-purchased ?x ?p)
    (has-newsgroup-interest ?x ?n)
    (has-price-in-pence ?x ?pr)
    (has-company-name ?x ?c)
    (has-advertising-expenditure ?x ?a)
    (has-alcoholic-rating ?x ?al)))
```

Fig. 8. The OCML query which generated the display shown in Fig. 7.

The display in Fig. 7 is a result of the query shown in Fig. 8 posed directly in OCML. Additionally, queries can also be created using our visual query tool Lois [7].

Specific sets of items can be quickly homed in on using Quiver's set of sliders. A detailed view of a slider is shown in Fig. 6. As a slider is dragged horizontally the items within the graphical display are instantaneously updated. Changing the width of a slider - using the slider's range adjusters - changes the range of data covered.

Discussion

The Guides metaphor was created to resolve the tension between providing flexible routes through a hypermedia database and not placing an undue cognitive load on the user by offering a large selection of links. Although each Guide provided a fixed path through the database, flexibility was still supported through the choice of Guides.

Within the Alice framework Guides provide a mechanism for customers to classify themselves. The first time a customer logs onto an Alice enhanced shopping site she will be assigned a number of Guides by default. When the shopper feels the need she will be able to add or delete her current Guides. Because it is the customer who says 'this is who I am' they will feel an ownership of their characterisation and consequently be more tolerant of any mismatches between their preferences and the recommended products.

We envisage that Guides will be designed by a company's sales department and broadly fall into the categories of knowledge that we stated our corner shopkeeper would use, namely, customer, products, shopping tasks and external context.

Guides in Alice form a bridge between the formal knowledge models and the individual user. We also want to argue that Alice Guides can help in the formation and support of online communities.

According to Rheingold, the Web encourages the growth of virtual communities of various kinds [28]. Rheingold's perspective has been the subject of a great deal of criticism on philosophical and political grounds. Although Rheingold attempts to answer these criticisms in a new chapter in the latest edition of his book, there are still critics such as Dreyfus [10]. Dreyfus suggests that Rheingold's *electronic agora* is in fact 'dangerously distopian' since its participants can remain anonymous and are not exposed to the real-world risk associated with the vulnerabilities of embodiment. Despite these criticisms there is no doubt that such communities do exist. Furthermore, the communities of interest to e-commerce are less open to the sorts of philosophical criticisms deployed by Dreyfus and others—real-world risk and the inability to make public commitments are largely irrelevant to a community centred around an interest in communicating, say, about the music of Bob Dylan.

While many e-commerce sites attempt to foster virtual communities through, for example, their facilities for the publication of consumer reviews of products, their approach is half-hearted at best since they remain attached to a form of marketing which is directed at the individual consumer. Kozinets [17] calls this *database marketing*, and suggests that the marketer (or e-store) bases efforts to influence consumer behaviour on the incorrect assumption of a one-way relationship between active seller and passive buyer. Of the two main forms of personalization (see below) *contented-based filtering* is the most individualistic with information technology being used essentially to track and make inferences about what consumers have purchased in the past. Although *collaborative filtering* tries to make inferences about what is relevant for a particular consumer based on some measure of similarity with other consumers it remains essentially oriented to individuals. Indeed, the choice of the term *personalization* suggests an individualistic approach to the relationship between seller and buyer.

Another approach is possible and may be more relevant to virtual communities. Kozinets [17] defines *Virtual Communities of Consumption* as "affiliative groups whose online interactions are based upon shared enthusiasm for, and knowledge of, a specific consumption activity or related group of activities" (p. 254). He mentions Barbie doll collectors, X-Files fans and wine lovers. Kozinets stresses that the consumption of a particular product is only part of what is important to members of these groups. Of equal importance is knowledge of various kinds, for example, knowledge about a product and its context, knowledge of a community's cultural norms and knowledge of its specialized language. In addition the identity of the community member may be more or less defined in terms of the consumption of the particular cultural or commercial product. According to Kozinets types of group members can be defined in terms of two axes: the degree of self-centrality of the consumption activity and the degree of social ties to a community. From this perspective *devotees* (who identify closely with the product but less so with the group) and *insiders* (who identify strongly with both) are most significant to the marketer. Thus it is not only important to determine the community that a consumer belongs to, it is also important to determine the correct type of community member. Kozinets mentions three characteristics of community-based as opposed to individualistic consumers: (1) they are more proactive; (2) they are more influenced by the community they belong to; and (3) they can provide valuable, multi-faceted information to marketers. He concludes that

marketers “must provide community members with the raw materials they need to construct a meaningful community” (p. 264).

While in the long run, some hybrid of individualistic and community based marketing will prove to be more attractive to many e-commerce sites, we have emphasized the latter in the Alice project. If we take each of Kozinets’ three points in turn we can indicate to what extent Alice can (or could in future) comply with them:

Consumers are proactive. In Alice we provide a default set of Guides but allow the consumer to select their own. Their selection reflects their self-assessment of themselves as consumers or in Kozinets’ terms as members of particular e-tribes. While many of the Guides act as critics of or assistants with the consumer’s interaction with Alice (e.g., Matchmaker, Missing Item) others can be seen as system components knowledgeable about the characteristics of particular communities (e.g., Organic, Money Saver). For instance, the Organic Guides allow consumers to express the activist tendencies associated with these communities. In future versions of the system we might include more awareness of meta-categories of community member such as devotee and insider. We might also provide the means for communication between consumers and stores, for example, in the form of a complaints procedure. More also needs to be done to allow community members to creatively review and criticize products, policies and marketing strategies perhaps through some form of Web log. In addition, since Alice is part of the Semantic Web, intelligent, ontologically-guided searches could be instigated for additional consumer reviews, buying Guides and so on.

Consumers are communal. As well as the community-oriented Guides discussed above Guides could also be used as an interface for building communities. A future version of our Guides will be linked to an instant messaging server such as Jabber [20]. A ‘Seek Soulmates’ Guide would attempt to establish chat or email sessions with other online customers who employ a similar set of Guides. In addition, the Quiver tool could be adapted to identify, visualize and provide the means of contacting individuals both for managers and other consumers. While many consumers would find such a tool overly sophisticated, the kind of technology-savvy community member identified by Rheingold would have little trouble. As we have said we also intend to use clustering to “notice” new communities and create new (community) Guides based on these. By so doing we would go a long way towards the provision of raw materials for constructing communities demanded by Kozinets.

Community-based consumers provide valuable information to marketers. In addition to the usual information about products bought, items viewed and ratings given to products, the e-tribe aware e-store should be able to derive what Kozinets calls a “cultural profile” of its consumers. According to Kozinets this would lead to a more detailed picture of the interests of the community (or, indeed, communities) the consumer belongs to, which could be used to assess “interconnections between seemingly disparate forms of consumption” and to see where “consumers are focussing their attention” (p. 260). He singles out insiders and devotees as important here. While Alice does not currently provide such a facility for marketers, the Quiver tool coupled with customer histories and clustering might form the basis for a more multi-faceted approach to understanding consumers in the future. We might also extend the learning abilities of Guides (so that they could track the individuals they interact with and learn from them) and create a new tier of system component which can notice significant patterns in the combination of Guides employed.

Related work

The Alice Guides are a particular approach to personalization. Jakob Nielsen defines personalization as:

“...to serve up individualized pages to the user based on some form of model of that user's needs.” [24]

Other approaches to personalization include content based and collaborative filtering. Content-based filtering recommends items based on their similarity to what the customer has bought in the past. An example of this approach is the Intelligent Personalised TV Guides [5].

Collaborative filtering makes recommendations based on the preferences of customers from the same group. Users are compared based on how similar their ratings are, and they are recommended items favoured by other people with similar interests. A well known example of collaborative filtering is www.amazon.com. ALEXA (<http://www.alexa.com>) is a web browser that recommends related links based in part on other people's web surfing habits.

The main problem in some types of business is the lack of information about customers' habits. Customers do not want to fill forms about themselves, unless they can clearly see the advantage of doing so (for instance, credit card companies often offer a prize draw for filling in a survey). Thus, it is difficult to fully understand their shopping behaviour. An alternative approach is to use Knowledge Discovery and Data Mining techniques on retailer's databases [1, 16]. In Alice, we intend to use an unsupervised clustering technique based on [4], to cluster customers according to their buying patterns (i.e. their shopping baskets). Alice will then extract rules encoding the consumption patterns. A similar approach was adopted by Lawrence, et. al [19] to identify groups of shoppers with similar spending histories.

Stereotypes [29] assume that facts about people are not statistically independent. This suggests that facts can be clustered into groups that frequently co-occur. Thus, a user model built with stereotypes adds a whole cluster of user facts at once, as soon as some evidence that is known to be a predictor of the cluster is observed. Therefore, it might be possible to make predictions about the behaviour of users on the basis of an amount of evidence – which can be acquired before an action is performed. The role of these predictions is to provide a basis for an action until specific knowledge becomes available.

Let us now examine initiatives related to the development of shareable product data. ISO 10303 (STEP) is an International Standard for product data representation and exchange which has existed since 1994. The development of STEP was initiated and is still driven by industry's need for technologies that enable application systems to exchange and share data about technical products. A STEP model is not however the same as an ontology - STEP definitions tend to be semantically weak. An overview of the problems in precisely capturing semantics within STEP models are discussed in [21].

A five level hierarchical categorisation of products is contained in the United Nation Standard Products and Services Codes (UN/SPSC) taxonomy. This structure however has no attributes. The Universal Content Extended Classification (UCEC) is an extension of UN/SPSC, developed by ContentEurope.com S.A. and now managed by The Electronic Commerce Code Management Association (ECCMA), has over 12,000 categories of products but again there are no attributes. A classification of a number of content standards including a number of product and e-business standards can be found at [9]

Although, visualization is considered relatively important by the ontology and e-commerce community (visualization is stated as a key issue for a proposed Internet Services Operating System [27] from the OntoWeb industrial applications special interest group), the number of visualization systems targeted at ontology based tools are relatively few. The most popular visualization used for browsing ontologies is the folder based tree views supplied as part of Java. A notable exception is Jambalaya [32], an application of the SHRIMP visualization framework [31] to view knowledge models in Protégé [25].

Conclusions

Alice is an example of a semantic web [3] e-commerce application. One of the main problems that the semantic web aims to solve is that of information overload. By indexing web resources with a formal representation items of interest can be found from their semantics. Although a lot of work is under way in creating the semantic web (see for example, [12]) most of this work focuses on infrastructure issues. Within the Alice project we have focused on how interface metaphors can augment semantic web technology to aid in user interaction within an e-commerce context. In this paper we have described two metaphors that we currently use: Dynamic Queries and Guides.

Quiver couples knowledge modelling technology to highly interactive navigation mechanisms through its on-the-fly dynamic query interface generation. Combining the strengths of ontology based queries and dynamic queries will benefit both knowledge engineers developing knowledge systems and end users looking for relationships in large volumes of data. Moreover, we expect that as the semantic web grows tools like Quiver that can present semantic data in a form that non computer specialists can understand will become ubiquitous.

The Alice Guides form a bridge between online communities of users and semantically enriched web resources. We believe that a community-based approach to marketing coupled with tribalized Guides begin to provide an online approximation of the old style community shop which we mentioned in our introduction. It is paradoxical that this particular type of store may be disappearing from the real world just as it is beginning to materialize in cyberspace.

Acknowledgements

This work was sponsored by INNN (www.innn.com) and by the Advanced Knowledge Technologies (AKT) project. AKT is an Interdisciplinary Research Collaboration (IRC), which is sponsored by the UK Engineering and Physical Sciences Research Council under grant number GR/N15764/01. The AKT IRC comprises the Universities of Aberdeen, Edinburgh, Sheffield, Southampton and The Open University.

The Alice graphic design work was carried out by Harriett Cornish.

The authors are grateful to feedback received on versions of this paper from Simon Buckingham Shum, Marc Eisenstadt and Paul Mulholland.

Reference

1. Agrawal, R., Mannila, H., Srikant, R., Toivonen, H. and Verkamo, A. I. (1996). Fast discovery of association rules. In Fayyad, U. et al. Eds. *Advances in Knowledge Discovery and Data Mining*. MIT Press, Cambridge.
2. Ahlberg, C., Williamson, C., and Shneiderman, B., (1992) Dynamic queries for information exploration: An implementation and evaluation, *Proceedings of ACM CHI'92: Human Factors in Computing Systems*, pp. 619-626.
3. Berners-Lee T., Hendler J., and Lassila O. (2001) *The Semantic Web*, Scientific American, May, 2001.
4. Cheeseman, P. and Stutz, J. (1996) Bayesian Classification (AutoClass): Theory and Results. In Fayyad, M. et al. (editors), *Advances in knowledge discovery and data mining*. MIT Press.
5. Cotter, P. and Smyth, B. (2001) PTV: Intelligent Personalised TV Guides. Smart Media Institute, Dep. of Computer Science, University College Dublin (available at <http://www.ptv.ie>).
6. Domingue, J. (1998) Tadzebao and WebOnto: Discussing, Browsing, and Editing Ontologies on the Web. In B. Gaines and M. Musen (editors), *Proceedings of the 11th Knowledge Acquisition for Knowledge-Based Systems Workshop*, April 18th-23th, Banff, Canada, (available at <http://kmi.open.ac.uk/people/domingue/banff98-paper/domingue.html>).
7. Domingue, J. and Motta, E. (2000) Planet-Onto: From News Publishing to Integrated Knowledge Management Support. *IEEE Intelligent Systems Special Issue on "Knowledge Management and Knowledge Distribution over the Internet"*, pp. 26-32, May/June, 2000.
8. Don, A., Oren, T. and Laurel, B. (1991) Guides 3.0, *Proceedings of ACM CHI'91 Conference*, New Orleans, LA, April 27th-May 2nd, pp. 44-448.
9. Martin Dörr, Nicola Guarino, Mariano Fernández López, Ellen Schulten Milena Stefanova, Austin Tate. State of the Art in Content Standards. *OntoWeb* (available at <http://www.ontoweb.org/download/deliverables/D3.1.pdf>.)
10. Dreyfus, H. L. (2001) *On the Internet*. Routledge: London.
11. Duineveld, A., Stoter, R., Weiden, M., Kenepa, B., and Benjamins, V. R. (200) Wondertools? A comparative study of ontological engineering tools. *International Journal of Human Computer Studies*, 52(5), pp. 1111-1133.
12. Fensel D. and Musen, M. (2001) Special Issue on The Semantic Web. *IEEE Intelligent Systems*, March/April, 16(2).
13. Gruber, T. R. (1993) A Translation Approach to Portable Ontology Specifications. *Knowledge Acquisition*, 5(2).
14. Gruninger, M. and Lee, J. (2002) Ontology Applications and Design. *Communications of the ACM* Vol. 45, No. 2, pp. 39-41 February, 2002.
15. Holt, R., Schürr, A., Sim, S. E., and Winter, A. (2001) GXL (1.0) Document Type Definition, Dagstuhl Edition, February 14, 2001, (available at <http://www.gupro.de/GXL/>).
16. Jain, A.K.; M.N.Murty; P.J.Flynn. Data Clustering: A Review. *ACM Computing Surveys*, Vol. 31, N.3, 1999.
17. Kozinets, R. V. (1999) E-Tribalized marketing?: The Strategic Implications of Virtual Communities of Consumption. *European Management Journal*, 17(3), pp. 252-264.
18. Laurel, B., Oren, T. and Don, A. (1990) Issues in Multimedia Interface Design: Media Integration and Interface Agents, *CHI '90 Conference Proceedings*, Seattle, WA, pp. 133-139, April 1-5, 1990.
19. Lawrence, R. D., Almasi, G. S., Kotlyar, V., Viveros, M. S., and Duri, S. S. (2001) Personalization of Supermarket product recommendations. *Data Mining and Knowledge Discovery*, 5, pp. 11-32.
20. Lee, S. and Smelser, T. (2002) *Jabber Programming*, Hungry Minds.

21. Metzger, F. (1996) The Challenge of Capturing the Semantics of STEP Data Models Precisely, Workshop on Product Knowledge Sharing for Integrated Enterprises, held in conjunction with the First International Conference on Practical Aspects of Knowledge Management, October 30-31, Basel, Switzerland, (available at <http://www.ladseb.pd.cnr.it/infor/Ontology/BaselPapers/Metzger.pdf>).
22. Motta, E. (1999) Reusable Components for Knowledge Models. IOS Press, Amsterdam.
23. Motta, E., Buckingham Shum, S. and Domingue, J. (2001) Ontology-Driven Document Enrichment: Principles, Tools and Applications. *International Journal of Human Computer Studies*, 52(5), pp. 1071-1109, 2000.
24. Nielsen, J. (1998) Personalization is Over-rated. Alertbox, October 4, 1998, (available at www.useit.com/alertbox/981004.html).
25. Noy, N. F., Sintek, M., Decker, S., Crubezy, M., Fergerson, R. W. and Musen, M. A. (2001) Creating Semantic Web Contents with Protege-2000. *IEEE Intelligent Systems* 16(2) pp. 60-71, 2001.
26. Oren, T., Salomon, G., Kreitman, K., and Don, A. (1990) Guides: Characterizing the Interface. In Laurel, B. cd., *The Art of Human-Computer Interface Design*, Reading, MA Addison-Wesley, pp. 367-381.
27. Persidis, A. (2001) The Working Group on Web Services Integration White Paper, ONT-SIG4-WG2-P1, June 20, 2001 (available at http://www.cs.vu.nl/~maksym/sig4/papers/SIG4_wp.doc).
28. Rheingold, H. (2000) *The Virtual Community: Homesteading on the Electronic Frontier*. Revised edition. MIT Press: Cambridge, MA.
29. Rich, E. (1985) Stereotypes and User Modeling. In *User's model in Dialog Systems* Kobsa, A. and W. Wahlster (eds.). Springer-Verlag.
30. Shneiderman, B. (1994) Dynamic queries for visual information seeking, *IEEE Software* 11, 6 pp. 70-77.
31. Storey, M.-A. D., Fracchia, F.D. and Müller, H. A. (1999) Customizing A Fisheye View Algorithm to Preserve the Mental Map. *Journal of Visual Languages and Computing*, 10, pp. 245-267, 1999.
32. Storey, M.-A. D., Musen, M. A., Silva, J., Best, C., Ernst, N., Fergerson, R. W. and Noy, F. (2001) Jambalaya: Interactive visualization to enhance ontology authoring and knowledge acquisition in Protégé Workshop on Interactive Tools for Knowledge Capture, held in conjunction with KCAP'01, The First International Conference on Knowledge Capture. Victoria, British Columbia, Canada, October 21-23, 2001.
33. W3C. (1999) Resource Description Framework, (RDF) Model and Syntax Specification, W3C Recommendation 22 February 1999 (available at <http://www.w3.org/TR/REC-rdf-syntax/>).
34. W3C (2000) Extensible Markup Language (XML) 1.0 (Second Edition), October, 2000, (available at <http://www.w3.org/TR/2000/REC-xml-20001006>).