

Supporting Online Shopping through a Combination of Ontologies and Interface Metaphors

John Domingue¹, Arthur Stutt¹, Maria Martins^{1,3} Jiacheng Tan^{1,4} Helgi Petursson^{2,5} and Enrico Motta¹

¹The Knowledge Media Institute
The Open University, Milton Keynes, UK
Tel: +44 1908 655014, Fax: -3169
{J.B.Domingue, A.Stutt}@open.ac.uk
<http://kmi.open.ac.uk/>

²Innn 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 build upon 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 interact with the system in a variety of ways. In effect, by choosing these options they are classifying themselves for the purpose of customizing system responses. We discuss the link between Alice Guides and Kozinets' notion of e-tribes or Virtual Communities of Consumption. Our second interface metaphor Dynamic Queries (coupled with Starfield displays) allows 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.

³ Current contact details: Dept. of Marketing and Strategy, The Business School, University of Gloucestershire, Pallas, 208, The Park, Cheltenham, Gloucestershire. GL50 2QF, UK
Email: mmartins@glos.ac.uk

⁴ Current contact details: Dept. of Comp. Science and Soft. Eng., University of Portsmouth, Buckingham Building, Lion Terrace, Portsmouth, Hampshire, PO1 3HE, UK
Email: jiacheng.tan@port.ac.uk

⁵ Current contact details: cTarget Inc., Brautarholt 1, 105 Reykjavik, Iceland.
Email: hp@ctarget.com

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.

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. None 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. Empirical findings (Lohse and Spiller, 1998; Koivumaki, 2001) suggest that improving the browsing and navigation capabilities of stores and especially product information (such as price, image, better descriptive name) can generate significantly higher traffic and sales per store. However, while the friendly shop-keeper is the ideal it will not be viable in the online situation. We may trust the shopkeeper who knew our mother when she was a child but we do not (and should not) trust some remote corporate entity in the same way. In practice we have found that the best way to limit the knowledge held about us as individuals while still providing customization facilities is to model individuals as members of particular communities. Interestingly, we can view community facilitation as one of the roles of our shopkeeper. The shopkeeper would often make associations between customers, such as *Mr. X always buys this product and might happen to know how to solve the problem Mrs Y is having*. In other words, she mediates communication between users, either introducing one to each other, or gathering and exchanging information that could be useful to users. This suggests the beginning of *communities of consumption* (Kozinets, 1999) which nowadays form a significant part of cyberspace.

The paper will describe some of the results of the project structured in the following fashion. In the next section we describe the Alice approach followed by a short

scenario. Section 3 describes the architecture of the system and Section 4 discusses the implementation of an interface for detecting the patterns of behaviour of online customers. Finally, a discussion section, linking the Alice approach to the notion of e-tribes (a kind of on-line community), is followed by some conclusions. The reader should note at the outset that while we in general espouse a user centred approach to system design, we have not been able to use this in Alice. This is because: (a) there are potentially many million users of a system like Alice; (b) it is precisely because it is difficult to identify typical users that we have adopted a community-centred approach; and (c) Alice is a generic module which can be further customized by a vendor who can take fuller account of their particular customers.

Approach

The Alice framework is based on the use of ontologies for representing knowledge related to online shopping. An ontology (Gruber, 1993) is an explicit representation of a domain of discourse (a conceptualisation) usually composed of a set of concepts and relationships. In one view ontologies contain data about data (or metadata), therefore making it possible for applications and people to ‘talk’ and ‘share data’ based on a common ground⁶. This feature enables computers and people to work in cooperation, which is the main principle of the next web generation, the Semantic Web⁷, being

⁶ Ontologies go beyond simple metadata schemes in that they incorporate features derived from first order logics which allow them to model subtle aspects of a domain and to be used in making inferences.

⁷ <http://www.w3.org/DesignIssues/Semantic.html>

developed since 1998 at W3C⁸ (Berners-Lee et al., 2001). As a result 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 441,000 web pages⁹. Moreover, ontologies are widely deployed within the knowledge acquisition and modelling communities, and have been successfully used in a variety of web based applications (e.g. Planet-Onto, a news server employing semantic search, Domingue and Motta, 2000). 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. This ontology also provides the structure for users to interact with the system in subjects they are interested in. It is common nowadays for supermarkets to have a message board for clients to publish messages (small ads, invitations, etc.)

⁸ www.w3c.org

⁹ We can ensure that most of the web pages devoted to ontology as a branch of philosophy are filtered out by using the query “ontology –philosophy”. This returns 175,000 hits but it is important to remember that many of these will be relevant to the computer science sense of the term. If nothing else this indicates why semantic search is essential.

- *Customer* – this ontology represents the main attributes of a customer including his or her shopping and browsing histories. Note that a customer’s interactions can also be recorded anonymously.
- *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 server¹⁰. Specifically, we use the relations list-price (indicating the listed price of the product), has-model-number (indicating the product type) and has-special-discount (indicating any current discount offers)¹¹.

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 customer’s 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

¹⁰ Available at <http://ontolingua.Stanford.edu/>

¹¹ We have in mind Prolog-like relations here with any arity. For example, the binary relation “Product1 list-price £200”.

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.

Our approach within the Alice project has been to use the Guides metaphor (Don et al., 1991; Laurel et al., 1990; Oren et al., 1990) as a mechanism to allow customers to implicitly classify 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 figure 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.

Guides in Alice embody different approaches to consumption (organic, luxurious, etc.), which are familiar to shoppers. Different guides (or combination of guides) would suggest different products to customers. Guides do not replace the well-known online supermarket store interface. They are designed so as to interfere as little as possible with the interface the user is familiar with. Therefore, the design of the guides includes ways of catching the user's attention when something relevant appears. In our prototype, the guide catches the user's attention by changing its colour to red so as to differentiate from other guides already selected but that have nothing to say¹². Guides' names and pictures were chosen so as to give an instantaneous understanding of their purpose, either by association (*productmatchmaker* for complementary products, shortened to *matchmaker* in the icon), or by incorporating symbols which are well known to the user because they represent a given aspect of a virtual community (e.g., the symbol for organic produce).

Although Alice uses Guides as a means of accessing users' self-classifications and linking these to different communities, from the user point of view, the selection merely represents a set of (possibly momentary) preferences which is entered as a means of accessing relevant information. This accords with Schoeder's comments (Schoeder, 1998), "users would like to use the website as in a reference book: open it, find what they want, and close it". We have adopted this approach to avoid any user

¹² While red was selected for its contrast with the normal Guide colour of blue, we can further justify our use of a colour often associated with emergencies or danger. According to Madden et al. (2000), in most countries red is perceived as a unique colour, consistently associated with 'hot', 'active' and 'vibrant'. Blinking is used to catch the user's attention specially for those shoppers that cannot distinguish colours properly.

resistance to self-classification. In the following section we shall give an overview of our implementation of Guides using a short scenario.

Insert Figure 1 Here

Figure 1. A screen snapshot of the Guides system (taken from <http://www.abbendon.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.

A Scenario

Insert Figure 2 Here

Figure 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 figure 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).



(Product)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.

Insert Figure 3 Here

Figure 3. A screen snapshot where the *Organic Guide* offers to show the cheapest organic pasta.

Insert Figure 4 Here

Figure 4. A screen snapshot just after the Organic Guide has displayed the cheapest organic pasta.

In figure 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 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 display changes to figure 3 – the *Organic Guide* offers to show the cheapest organic pasta. The shopper agrees and the display changes to figure 4. Note that hundreds of potential items (a well known online store we checked has over a hundred different pastas) have been narrowed to one in precisely two mouse clicks. The shopper decides to buy 2 packets and the *Party Guide* steps in to say that for a party of 15 people (previously specified) at least 1500g of pasta are required. This is shown in figure 5.

The key design feature of the Alice Guides is that the customer selects them. This means that they reflect the customer's own perspective of themselves (e.g. rich, poor or ethical) and therefore there is a greater likelihood that appropriate suggestions will be made¹³. Also, depending on the current situation the customer can chose to de-select or ignore certain Guides, for example, the *Money Saver Guide* when shopping for a specific luxurious item.

Insert Figure 5 Here

¹³ Our plan is to carry out an empirical evaluation of this around a deployed system.

Figure 5. A screen snapshot showing where the *Party Guide* is warning the customer that she may need more pasta for her party.

The Alice Architecture

The Alice architecture is composed of a server and several specific clients. One of the overriding goals when conceiving the system was that it should be easy to incorporate into an existing online infrastructure. Consequently, the Alice server was designed to sit alongside an existing web server. The architecture is shown in figure 6. The Alice server contains a library of ontologies implemented in OCML (Motta, 1999). OCML, which can be conceived as an ‘Operational Ontolingua’ has been used in over a dozen knowledge management and knowledge modelling projects within our lab (e.g. Motta et al., 2001). Knowledge modelling in OCML is supported by a library of reusable definitions, which is structured according to the basic categories of the OCML modelling framework: task, method, domain and application (Motta, 1999). The library also relies on a number of base ontologies, which provide definitions for basic modelling concepts, such as numbers, sets, relations, tasks, methods, roles, etc. Export mechanisms exist from OCML to Ontolingua (Gruber, 1993), XML (W3C, 2000), RDF (W3C, 1999), and GXL (Holt et al., 2001).

Insert Figure 6 Here

Figure 6. The Alice Architecture.

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 fed through to a manager's interface allowing a company's sales and marketing department to discover new relationships between products and clusters of customer behaviour. When appropriate new knowledge gleaned from the analysis of customer behaviour will be used to create new Guides.

The knowledge models are created and maintained using WebOnto (Domingue, 1998). 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 (Duineveld, 2000).

```
(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))
```

Figure 7. 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 figure 7.

`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.¹⁴

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

Insert Figure 8 Here

Figure 8. A screen snapshot of Quiver displaying items from a beer knowledge base. The full query which generated this snapshot is shown in figure 9.

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 classic Pareto rule of 80-20, the overall goal of any analysis is to determine the significant attributes with respect to the 20% of

¹⁴ This view of Guides as interface components which depend on simple trigger mechanisms explains why we are reluctant to describe them as fully-fledged agents. They are more like the critiquers described in Fischer et al. (1993).

customers who purchase the 80% of products. For example, 16% of US beer drinkers account for 88% of annual consumption (Kozinets, 1999).

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 (Ahlberg et al., 1992; Shneiderman, 1994) 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.

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 figure 8. 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 resulting display is shown in screen snapshot in figure 8. 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 and noted from the changes to the graphical display area, shown in figure 12 that a) relatively few medium to low alcohol beer beers had been advertised extensively and b) there didn't seem to be any obvious relationship.

Let us now examine the Quiver interface shown in figure 8 in more detail. The large pane on the left (labelled 'a') of the window 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)))
```

Figure 9. The OCML query which generated the display shown in figure 8.

The display in figure 8 is a result of the query shown in figure 9 posed directly in OCML. Additionally, queries can also be created using our visual query tool Lois (Domingue and Motta, 2000). Lois enables queries to be created by direct manipulation. A Lois user navigates through a knowledge model selecting classes and relevant binary relations. These selections result in a set of rows from which a query is created. A query equivalent to the one shown in figure 9 (used to create the quiver display in figure 8) is shown in figure 10.

Insert Figure 10 Here

Figure 10. A screen snapshot showing the same query used to create the Quiver display in figure 8 within our visual query tool Lois (Domingue and Motta, 2000). Queries are created through the addition of rows. Each row is built by selecting on the 'Beer' button, selecting one of the displayed attributes, an attribute type, optionally a value, and finally the 'Add Row' button.

Insert Figure 11 Here

Figure 11. 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.

Insert Figure 12 Here

Figure 12. A screen snapshot showing how sliders can greatly reduce the number of visible items. The hundred plus items shown in figure 8 have been reduced to precisely 6.

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 figure 11. 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. In

figure 11 (taken from figure 12) the slider restricts the graphical display to only show items whose alcohol rating is between 260 and 417.

It is also possible to map sliders onto class hierarchies using a quiver-mapping relation. This relation assigns the range 0-100 to the root of a hierarchy. The children are recursively assigned contiguous sub-ranges of their parents. For example, if a root class had four children they would be assigned the ranges 0-24, 25-49, 50-74 and 76-100. The resulting slider is then attached to a graph view which highlights the currently active portion of the hierarchy.

More on Guides and communities

The Guides metaphor was created to resolve the tension between providing flexible routes through a hypermedia database and 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.

We envisage that Guides will be designed by a company's sales department (perhaps after an analysis using Quiver) and broadly fall into the categories of knowledge that we stated our corner shopkeeper would use, namely, customer, products, shopping tasks and external context. In a future version of Alice we might offer users the chance to design their own Guides. In Table 1 we set out the typical customer oriented Guides one could find in an online supermarket. The table includes a column for Guide orientation which is explained below.

Table 1. A range of customer oriented Guides that one might expect to find in an online supermarket

| Guide Name | Guide Orientation | Description |
|---------------------|--------------------------|---------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| Family | Community | Alerts the customer if any versions of products suitable for families are available (e.g. family size). |
| Single | Community | Notifies the customer if single portion versions of products are available. |
| Large Income | Community | Directs the customer to high value non-essential products. |
| Low Income | Community | Directs the customer to essential, low-cost products. |
| Allergy | Community | The Guide warns the customer if they attempt to buy a product which contains an allergen. |
| Health | Community | Directs the customer towards products that are healthy, for example, low in unsaturated fats and high in fibre. |
| Ecological | Community | Notifies the customer when goods have been produced locally or when they are part of a 'fair trade' scheme. |
| Party | Individual | 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). |
| (Product)Matchmaker | Individual | Matches products that are purchased to similar or complementary products. For example, it would match pasta with bottled pesto sauce. |
| Missing Items | Individual | 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 | Community | When appropriate this Guide recommends organic versions of goods that are being viewed. |
| Money Saver | Individual | Informs the shopper of any offers or promotional items which are related to the |

| | | |
|----------------|----------------|-------------------------------------------------------------------------------------------------------------------|
| | | currently viewed item. |
| Seek Soulmates | Meta-Community | Would attempt to establish chat or email sessions with other online customers who employ a similar set of Guides. |

Within the Alice framework Guides primarily act to provide a range of services to the customer. However, in selecting Guides customers also classify themselves. As we said above, this is important since we do not want the user to resist self-classification either because the categories are unintelligible or unpalatable thereby failing to use one of Alice's most important features. 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 implicitly says 'this is who I am' there will be less likelihood that they will feel the need to question system choices since there will be fewer mismatches between their preferences and the recommended products.

Thus, unlike other approaches to tailoring the e-commerce environment Alice does not need to accumulate large quantities of information about a particular customer or to keep track of all their interactions (although an analysis of the overall patterns of customer interaction allows the creation of new Guides). Recommendations are made not on the basis of a store of information but on the characteristics indicated by the choice of Guides. In fact Alice will operate well even if nothing is known about a customer other than her choice of Guides. Of course some information will be required if the customer wishes information to be stored over time and/or to make a purchase. Apart from this minimal amount of monitoring, Alice has no need to build a complex user profile. Alice only needs to acquire knowledge of the mappings

between the generic characteristics embodied in its Guides and particular products instead of reasoning from characteristics of individuals such as age, social group, spending power and previous shopping history. For example, the combination of the Organic and Low Income Guides would determine recommendations for a limited, cheaper range of organic produce.

A significant difference between the self categorization in Alice and similar experiments (Schoeder, 1998) is that the Alice interface does not constrain the user to only one type of characterization. Users are free to choose any number of guides. Moreover, the selected guides may represent conflicting categories (e.g. Large Income and Money Saver). In Alice, each guide represents the essence of a virtual community by representing mappings between its characteristics and products, allowing it to recommend suitable products.

Finally, Alice can create new Guides by analyzing a semantic¹⁵ log of mostly anonymous interactions with the system as a means of discovering patterns. These patterns can be presented to managers and new Guides created. For instance a pattern might indicate that there is a substantial group of customers who identify themselves by selecting the Organic Guide and who are only interested in fresh produce. Another group might only be interested in fruit and vegetables. Alice can thus learn from customer interactions without learning anything private about them. Its new knowledge would be embodied in Guides for, say, Fresh-Organics and Fruit&Vegetable-Organics. This evolutionary approach to Guides fits well with Hagel

and Armstrong's ideas on the formation of *niche communities* as aggregations of members of *core communities* with similar interests. By core community Hagel and Armstrong mean a large community interested, for example, in travel. By niche community they mean a sub-community of those interested in, in their example, travel to the Amalfi coast of Italy. As they say, "there is in principle no reason why a well-developed core travel community could not identify and aggregate its members who share a passion and interest in travel to the Amalfi coast. It might then create a distinct set of offerings and environments for these members to enjoy within its broader community offering" (Hagel and Armstrong, 1997, p. 93).

The knowledge of community characteristics embodied in a Guide is not simply a product list. For example the Organic Guide addresses the community of consumers who are interested in the benefits of organic food. This means that as well as including a list of products which can be described as organic (i.e., those with an organic certification), the Organic Guide will know that organic consumers are interested in educational material about how organic food is grown, the implications of genetic modification and so on. On the other hand, for example, the community self-identified as Low Income will be less interested in such material if only because the significant difference in cost will prevent them from buying organic items. Of course there is nothing to stop the lower income purchaser from switching on the Organic Guide as a form of window-shopping and for the information and advice provided.

¹⁵ By 'semantic' here we mean that the log is composed of assertions to Alice's knowledge base. This can be searched using the various product and community ontologies.

Our notion of community-oriented Guides is based on the work of Kozinets and others on *e-tribes* or *virtual communities of consumption*. 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 (1999) 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.

Kozinets favours another approach: *virtual communities of consumption*. He defines these 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) and mentions, as examples, Barbie doll collectors and wine lovers. In fact, there are already virtual communities of users that care for some sort of user characteristics, such as those that are allergic to gluten (<http://www.glutino.com>), or diabetics (<http://www.diabeticsdelightonline.com/>).

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 (p. 254). 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 extent to which they are committed (a) to a product or consumption activity and

(b) to the community based around it. Kozinets identifies four types of community member: *devotees*, *insiders*, *minglers* and *tourists* depending on their location in relation to these two axes. For example, the mingler has strong ties to the community but is less committed to the product, while the devotee is strongly committed to the product but less so to the community.

From this perspective devotees and insiders (who identify strongly with both product and community) 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.

Since we have adopted a minimal monitoring approach (where, although we collect information about user interactions we do not attempt to acquire other than the basic information about individuals) we cannot use Alice to track members of communities and make inferences about their changing roles. We can however add detail to our Guides so that when a user chooses a community-oriented Guide, he or she can, in addition, choose to identify themselves as devotee, insider, mingler or tourist. This could be done, for example, using Quiver's slider mechanism which would allow the user to give a self-description in terms of Kozinets' two dimensional space. This latter approach would avoid any problems of resistance to self-categorization in terms of the four categories. Indeed this approach allows us to go beyond Kozinets' simple set of categories to a space of possible community member types.

Kozinets mentions three characteristics of community-based as opposed to individualistic consumption: "Virtual communal marketing centers on consumers as (1) more proactive and (2) more communally influenced, and (3) the information that

they provide online as more multifaceted than more passive, one-to-one, and constrained database marketing practices” (p. 261).

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., (Product)Matchmaker, Missing Items) others can be seen as system components knowledgeable about the characteristics of particular communities (e.g., Organic, Allergy, Health, Ecological). For instance, the Ecological and Organic Guides should allow consumers to express the activist tendencies associated with these communities. In future versions of the system we might 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 made for additional consumer reviews, information on product defects 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. Kozinets

suggests that marketers “must provide community members with the raw materials they need to construct a meaningful community” (p. 264). To do this a future version of our Guides will be linked to an instant messaging server such as Jabber (Lee and Smelser, 2002). The 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.

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 involves 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”, the “reasons for customer (dis)satisfaction” and to see where “consumers are focussing their attention” (p. 260). While Alice does not currently provide such a facility for marketers, the Quiver tool coupled with semantic logs might form the basis for a more multi-faceted approach to understanding consumers in the future.

Kozinets thus provides the means for giving some structure to communities of consumption and more detail to the knowledge contained in particular Guides. Using the ontology derived from Kozinets we can create a generic *community profile* (as opposed to a user profile) which can be augmented as we learn more about a

community through the many interactions with the system and with each other performed by community members. See figure 13 for an example.

| |
|---------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <p><i>Community Profile CP123</i> Community Name: organic-vegan Community Insiders: user1, user22 Community Devotees: user2, user33 Community Minglers: unknown Community Tourists: unknown Community Events: Christmas, Easter, New-Product-Release1, ... Community Meeting Places: Alice-community-chat-room-COV Community Norms: strictly-no-meat-or-animal-derived-products Community Goals: learn-more-about-organics, learn-more-about-vegan-lifestyle, force-better-product-labelling Community Popular Products: product-ov1, product2-ov2...</p> |
|---------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|

Figure 13. An example community profile.

In short, community based Guides have the following advantages over alternative means of tailoring e-commerce sites to customers:

- a. They are less invasive. They operate using knowledge about communities rather than individuals.
- b. It is easier to form a community profile for an e-tribe than it is to know an individual if only because there are more interactions from a group of individuals. Instead of acquiring and storing a large amount of difficult to obtain information about individuals, all that is needed is a basic ontological framework for tribal structures, some knowledge of the characteristics of specific communities, a knowledge base derived from these, and the capacity to learn about communities from interactions with tribal members. Mining the company database of anonymous interactions will give clues about trends in customer behaviour. Guides can be validated by being frequently selected.

The need for new Guides is shown by, for instance, frequent combinations of Guides or combinations of Guides and particular types of product choices.

The really successful Guides will be those which closely match the needs of a community and which reinforce the self-identification and community feeling of users who select the Guide.

- c. It is easier to successfully track a community than an individual since there are more interactions with a community.
- d. Single interactions are unlikely to unbalance the system's suggestions. One problem with personalization on e-commerce sites is that a one-off purchase (perhaps for a friend) may figure too highly in the recommendations a site makes.

Related Work

The Alice Guides are a particular approach to personalization. Jakob Nielsen defines personalization, in the context of the web, as:

“...to serve up individualized pages to the user based on some form of model of that user's needs.” (Nielsen, 1998).

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 (Cotter and Smyth, 2001) which relates the user's profile to content of television programs to make viewing recommendations.

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 to solve this problem is to cluster customers according to their buying pattern (the shopping basket), and to extract the rules from each customer cluster. In this way, the rules encode consumption patterns which are used for personalization. A similar approach was adopted by Lawrence, et al. (2001) to identify groups of shoppers with similar spending histories. In contrast Alice situates clustering within an ontology-based approach which uses its understanding of community needs to infer mappings between products and users without invading the privacy of the latter.

There are a number of initiatives related to the development of shareable product data. Even though they delineate broad taxonomies, the semantic information about

products they contain is weak. To properly exchange information over the Semantic Web, those taxonomies need to be enriched with product related information suitable for a semantic web application. In Alice, we have used the UCEC classification as the basis for our product ontology. Below we list some of the existing standards:

- ISO 10303 (STEP)¹⁶ 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 (Metzger, 1996).
- UN/SPSC (United Nation Standard Products and Services Codes)¹⁷ is a five level hierarchical categorisation of products. This structure however has no attributes other than the taxonomy.

Conclusions

Alice is an example of a semantic web (Berners-Lee et al., 2001) 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

¹⁶ <http://www.nist.gov/sc4/www/stepdocs.htm>

creating the semantic web (see, for example, Fensel and Musen, 2001) 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 begins to provide an online approximation of the old style community shop which we mentioned in our introduction. Even if Alice does not have access to the store of personal information available to the shopkeeper, it is adept at differentiating between different customers by identifying their membership of one or more communities. This enables Alice to target information and advice in a more focussed manner. It is paradoxical that this particular type of interaction may be disappearing from the real world with its bland,

¹⁷ <http://www.un-spsc.net/>

homogenized and impersonal shopping centres 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.

This work has benefited from the numerous conversations we have had with Annie Brooking. The authors are grateful to feedback received on versions of this paper from Simon Buckingham Shum, Marc Eisenstadt and Paul Mulholland. We are also grateful for the insightful comments from an anonymous IJHCS reviewer.

Reference

- 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.
- Berners-Lee T., Hendler J., and Lassila O. (2001) The Semantic Web, Scientific American, May, 2001.
- 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>).
- 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>).
- 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.
- 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.

- Duineveld, A., Stoter, R., Weiden, M., Kenepa, B., and Benjamins, V. R. (2000) Wondertools? A comparative study of ontological engineering tools. *International Journal of Human computer Studies*, 52(5), pp. 1111-1133.
- Fensel D. and Musen, M. (2001) Special Issue on The Semantic Web. *IEEE Intelligent Systems*, March/April, 16(2).
- Fischer, G., Nakakoji, K., Ostwald, J., Stahl, G. and Sumner, T. (1993) Embedding Critics in Design Environments, *The Knowledge Engineering Review*, 8(4).
- Gruber, T. R. (1993) A Translation Approach to Portable Ontology Specifications. *Knowledge Acquisition*, 5(2).
- Hagel, J. and Armstrong, A. G. (1997) *Net Gain - Expanding Markets through Virtual Communities*. Harvard Business School Press.
- 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/>).
- Koivumaki, T. (2001) Customer Satisfaction and Purchasing Behaviour in a Web-based Shopping Environment. *Electronic Markets*, 11(3), pp. 186-192.
- Kozinets, R. V. (1999) E-Tribalized marketing?: The Strategic Implications of Virtual Communities of Consumption. *European Management Journal*, 17(3), pp. 252-264.
- 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.

- 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.
- Lee, S. and Smelser, T. (2002) *Jabber Programming*, Hungry Minds.
- Lohse, G. L. and Spiller, P. (1998) Quantifying the Effect of User Interface Design Features in Cyberstore Traffic and Sales. CHI, Los Angeles, USA.
- Madden, T. J., Hewett, K., Roth, M. S. (2000). Managing Images in Different Cultures: A Cross-National Study of Color Meanings and Preferences. *Journal of International Marketing*, 2000, Vol. 8 Issue 4.
- 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>)
- Motta, E. (1999) *Reusable Components for Knowledge Models*. IOS Press, Amsterdam.
- 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.
- Nielsen, J. (1998) Personalization is Over-rated. Alertbox, October 4, 1998 (available at www.useit.com/alertbox/981004.html).
- 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.

Schoeder, W. (1998) Steering Users isn't Easy. Eye for Design, 5(2), pp. 8-10.

Shneiderman, B. (1994) Dynamic queries for visual information seeking, IEEE Software 11, 6 pp. 70-77.

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/>).

W3C (2000) Extensible Markup Language (XML) 1.0 (Second Edition), October, 2000. (available at <http://www.w3.org/TR/2000/REC-xml-20001006>).

Figures

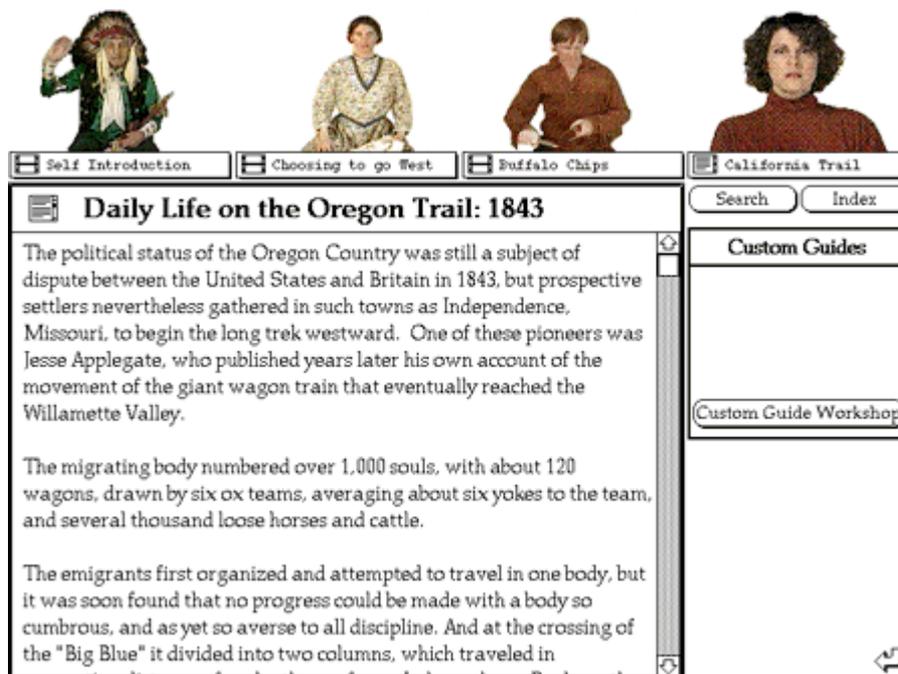


Figure 1. A screen snapshot of the Guides system (taken from <http://www.abbodon.com/project/guides.html> with permission (copyright 1990

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.

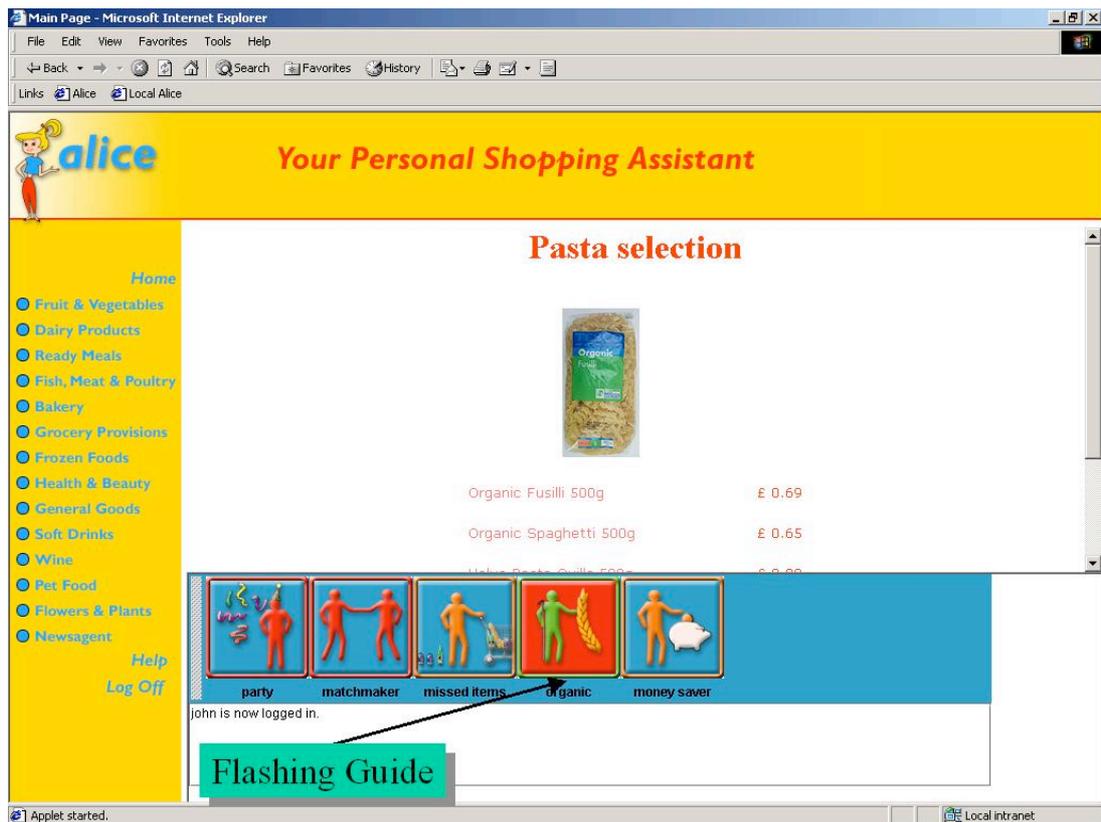


Figure 2. A screen snapshot of the Alice Guides interface. The left panel contains a standard navigation bar as found in most online supermarkets. The products are shown in the large pane in the centre of the browser. The Guides interface is shown in the bottom panel. The customer is browsing a selection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

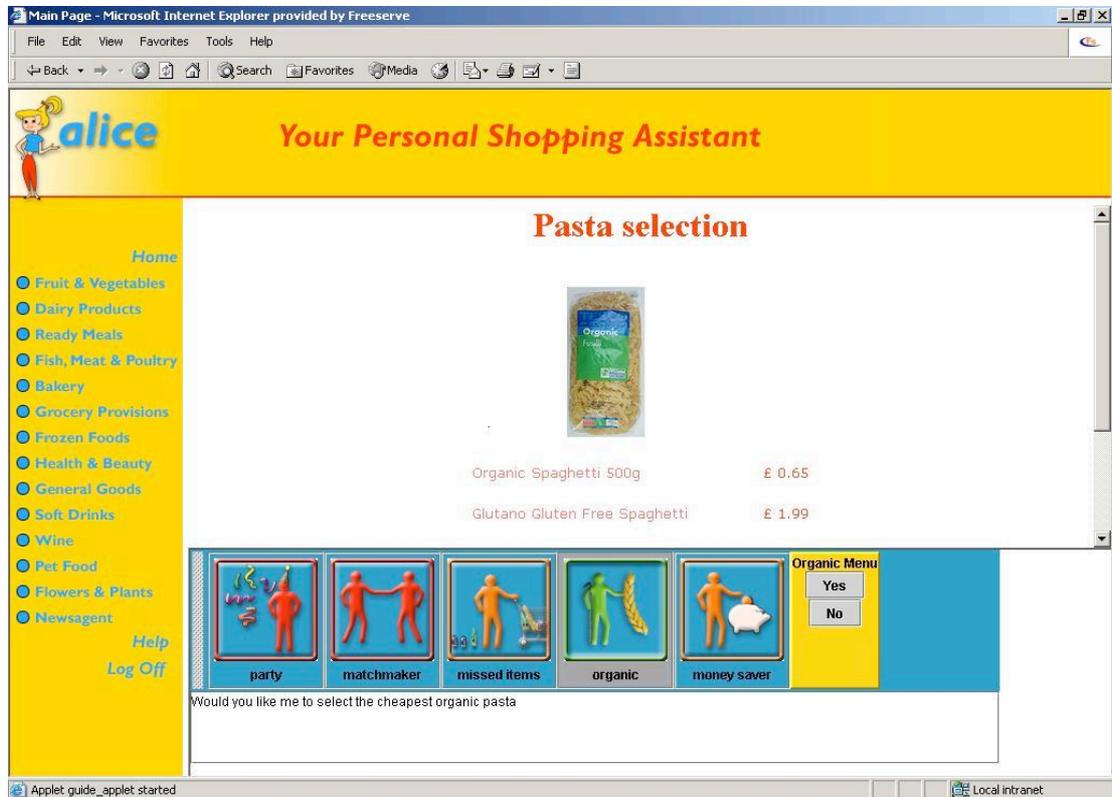


Figure 3. A screen snapshot where the *Organic Guide* offers to show the cheapest organic pasta.

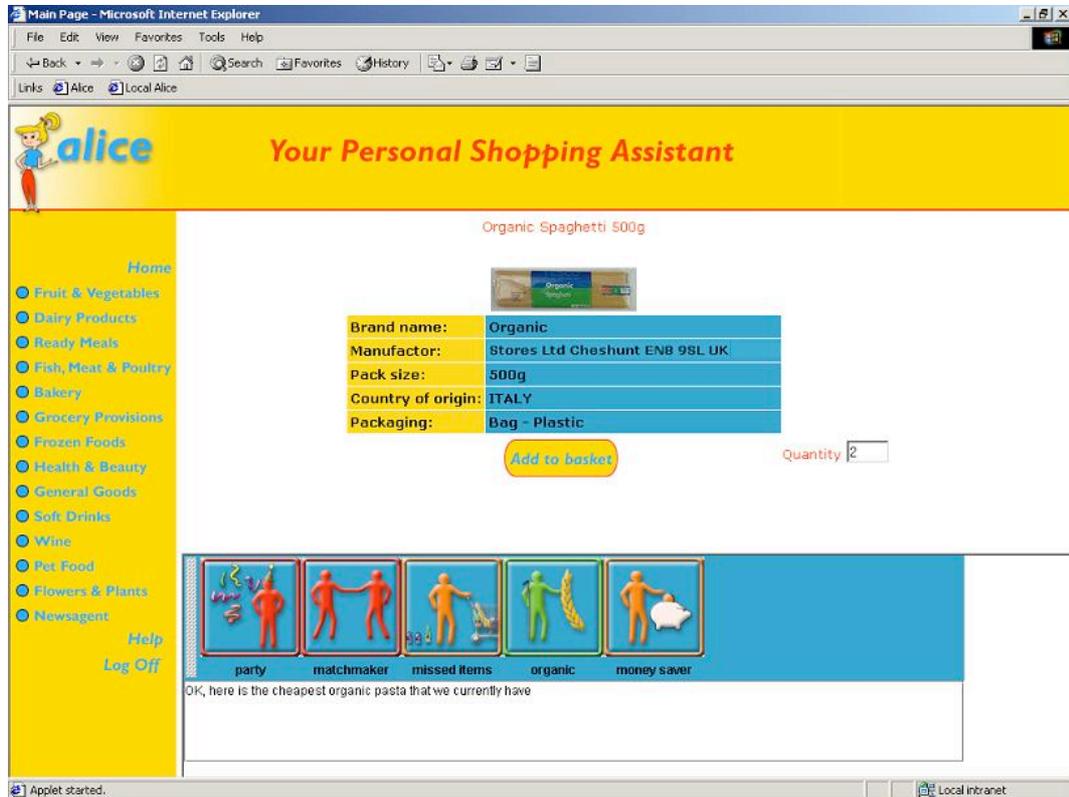


Figure 4. A screen snapshot just after the Organic Guide has displayed the cheapest organic pasta.

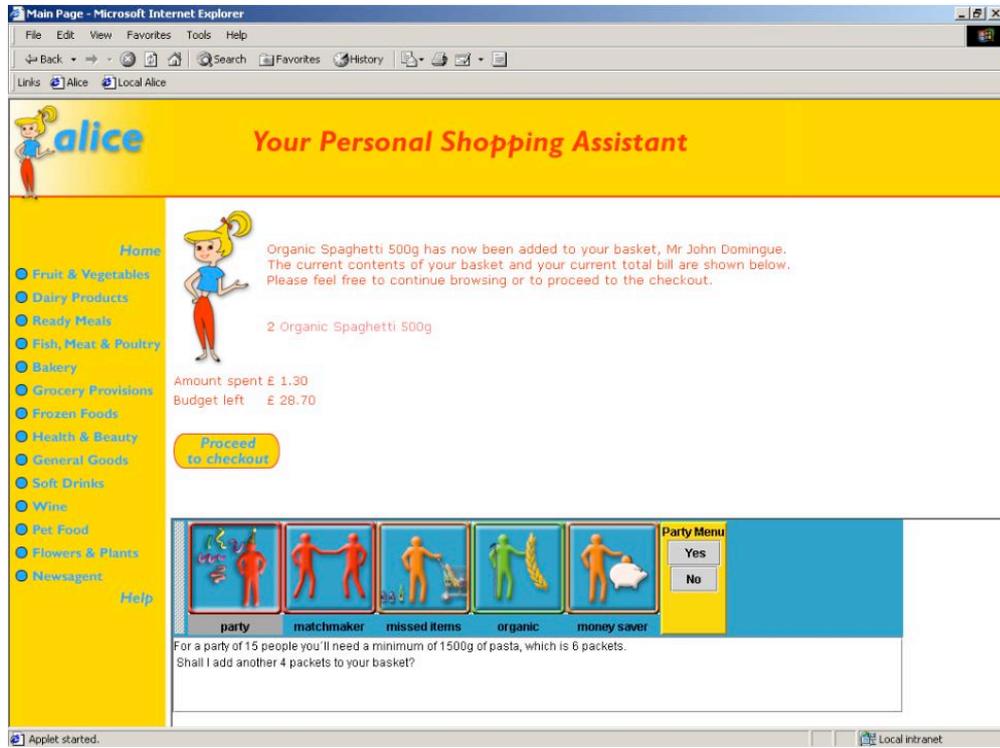


Figure 5. A screen snapshot showing where the *Party Guide* is warning the customer that she may need more pasta for her party.

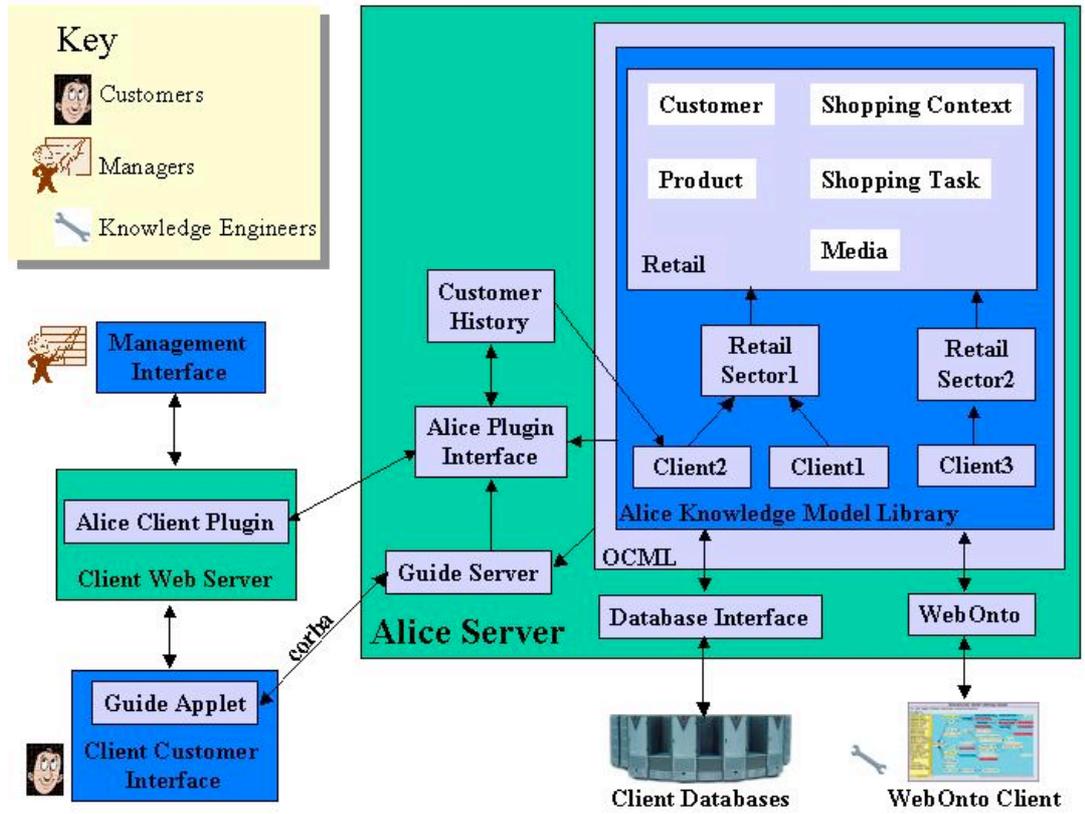


Figure 6. The Alice Architecture.

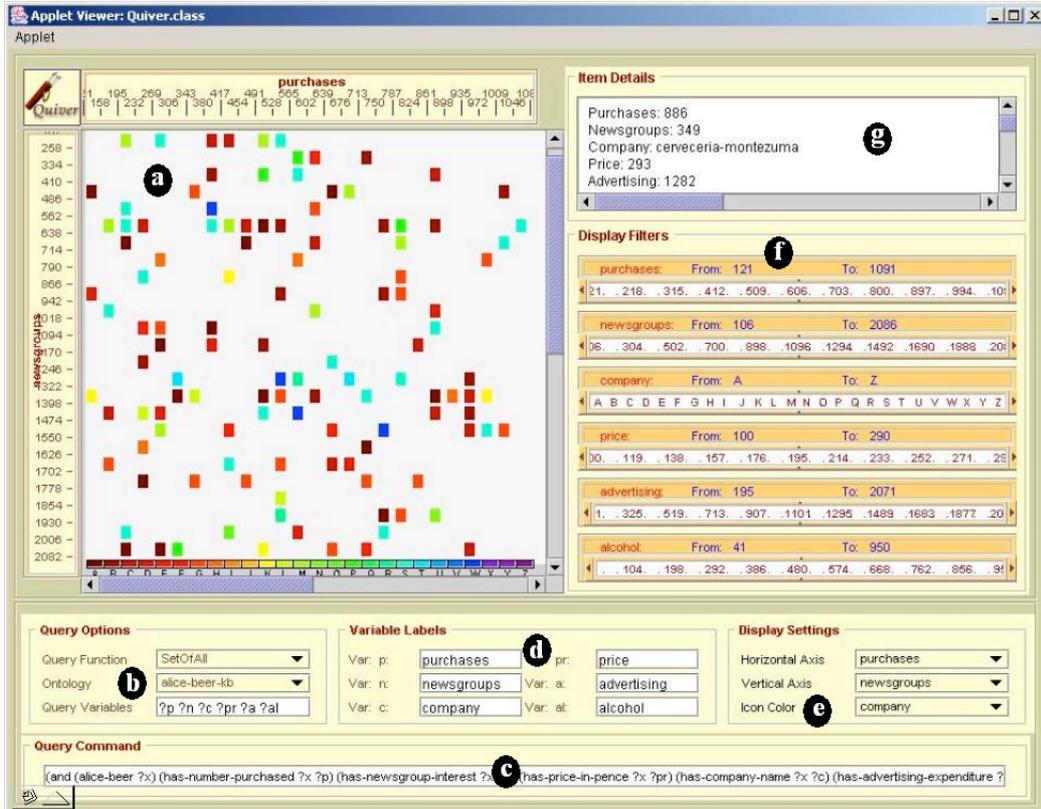


Figure 8. A screen snapshot of Quiver displaying items from a beer knowledge base. The full query which generated this snapshot is shown in figure 9.

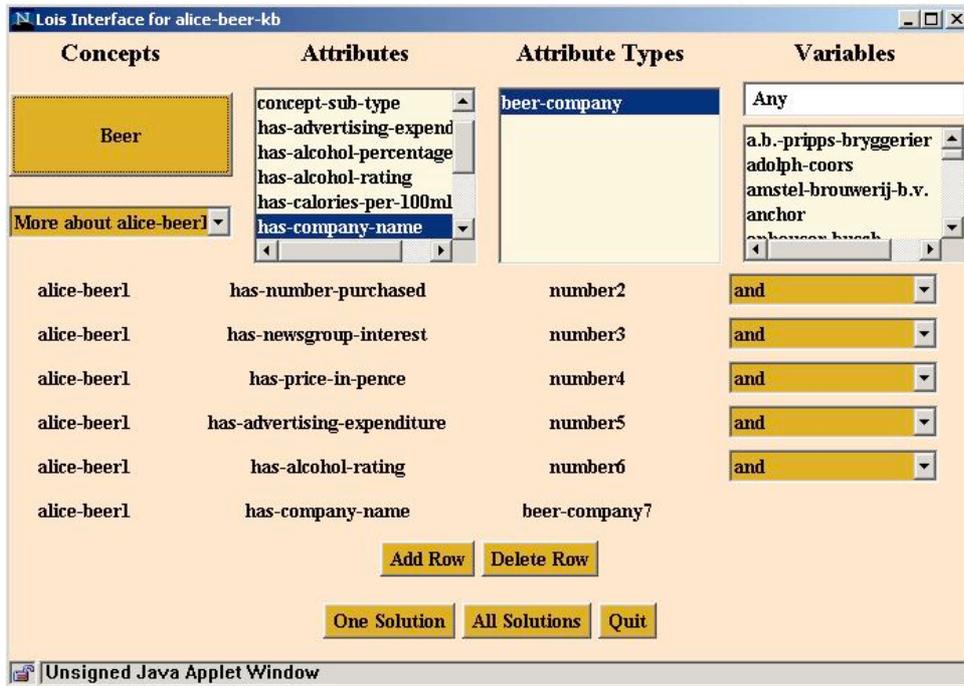


Figure 10. A screen snapshot showing the same query used to create the Quiver display in figure 8 within our visual query tool Lois (Domingue and Motta, 2000). Queries are created through the addition of rows. Each row is built by selecting on the 'Beer' button, selecting one of the displayed attributes, an attribute type, optionally a value, and finally the 'Add Row' button.

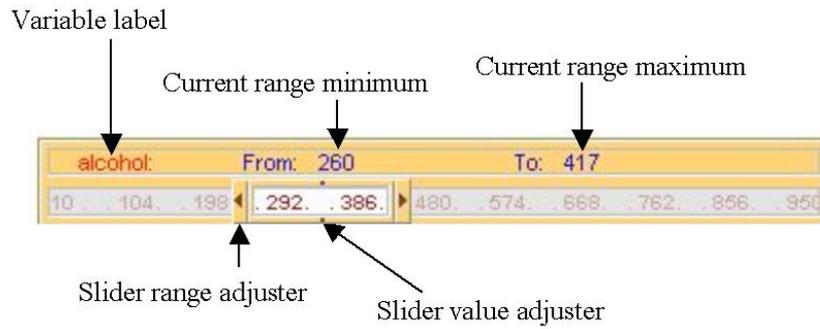


Figure 11. An annotated screen snapshot of the Quiver slider for the variable 'alcohol'. The slider filters the graphical display area to only show those items whose alcohol rating is between 260 and 417.

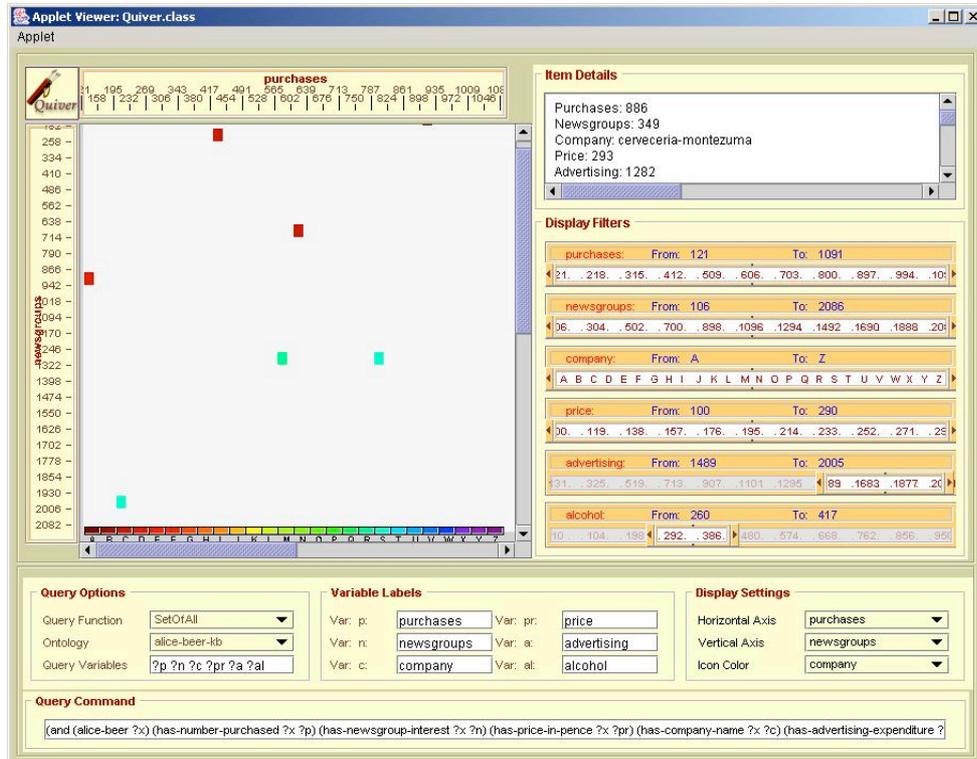


Figure 12. A screen snapshot showing how sliders can greatly reduce the number of visible items. The hundred plus items shown in figure 8 have been reduced to precisely 6.