

# Agent Mediated Retailing in the Connected Local Community

Mark Witkowski, Brendan Neville, Jeremy Pitt

Intelligent and Interactive Systems Group,  
Department of Electrical & Electronic Engineering,  
Imperial College of Science, Technology and Medicine,  
London, SW7 2BT.

{m.witkowski@ic.ac.uk; brendan.neville@ic.ac.uk; j.pitt@ic.ac.uk}

*Special Issue Editors: Kostas Stathis and Patrick Purcell*

## Abstract

This paper reports on and discusses recent research into software agent mediated retailing and considers how it may act as a catalyst to the formation of local electronic retail communities. Each member of the e-retail community has their own software agent that acts on their behalf. Their software agent acts as an *electronic Personal Sales Assistant* (ePSA), interacting with the person and assisting with the selection and purchase of products. We consider issues that arise when each person can express their likes and dislikes for various aspects of the product range to their software agent, both in terms of the software agent's ability to serve them better, and how such preferences may be shared with others in the larger community to help them – and others in the community - make better and more informed choices. The investigations reported use the MAPPA system, an experimental kiosk based e-retail system. MAPPA integrates a software agent based architecture, a novel form of product display and the use of an animated character agent to enhance the user's sense of personalisation. We report on user evaluation studies of the MAPPA system to investigate the effectiveness of the character agent as an interface to the e-retail system using both conventional usability evaluation techniques and eye-tracking technology. We also describe a simulation program that allows us to investigate the dynamics of e-retail community cohesion when a number of different parameters are considered, including the characteristics of the individual community members and the algorithms employed by the software agents to share information. Lastly, we develop and discuss notions of *loyalty*, *trust*, *reputation* and *preference*, four issues critical to the development and maintenance of the relationships between individual people, their personal software agents, businesses and the larger e-retail connected community. We propose these issues as “supra-functional” criteria for human-computer interaction design in this area.

**Keywords:** Connected Communities, Loyalty, Trust, Preference, Software Agents, Electronic Personal Sales Agents, Synthetic Personae.

# 1. Introduction

This paper considers several aspects of software agent mediated e-commerce within a localised on-line community comprised of both retail businesses and individuals who purchase the products or use the services of those businesses. The majority of business conducted on-line follows a mail-order catalogue approach. Customers view product descriptions using a conventional web-browser, make their selections and arrange to pay. Some products and services may be delivered electronically (such as software, music and banking) but most goods are dispatched via the mail service or by other traditional delivery methods. As pre-payment is the norm, problems encountered are often resolved by the customer returning the goods for a refund or by the vendor sending duplicate items in the case of non-delivery. While both purchaser and vendor behave in good faith the system works well. However, it is vulnerable to various abuses, primarily from rogue vendors, but also from ill-intentioned purchasers.

While the electronic mail-order approach has its place, and offers a convenient and well-tested mechanism to both parties, the question remains as to how extra advantage can be extracted from the use of the new technology and whether new and enhanced models of trade and community involvement can be implemented. Consider an alternative. You are shopping, possibly with your family, in your local shopping mall. Many of the purchases you intend to make are regular household items and on a list. There is little discussion about these, or where to purchase them. A smaller number of items, the wine to drink, a video to watch, some music to add to the collection or books to read, while not particularly expensive can be difficult to choose. This is not least because there is an extensive choice of quite similar products in each category. Each of these product ranges has been refined over the years to cater for a wide range of differing personal tastes and preferences. Some will be to your taste, some will not.

Placed about the mall, and in some of the stores, are clearly marked kiosks with computer screens. You insert your smart card and key in a PIN-number. This activates your personal software agent. On the screen an animated character, your *electronic Personal Sales Assistant* (ePSA) welcomes you, and you will be able to interact with it to view and select products from the stores. You can browse through the products, refining and making sub-selections, and confirm purchases if you wish. If you want, your ePSA will present details about individual products, indicating special features and in-store promotions or special deals. From time to time you can indicate to the ePSA various preferences you have for or against certain types of product. The ePSA notes these preferences and will change the way in which it interacts with you and makes suggestions in the future. Lastly your ePSA can act as a portal to the larger community of store users, enabling you to share (again, if you wish to) your preferences and experiences with the products – in exchange for ideas and suggestions derived from others in the local e-retail “community”. Many issues and questions arise from attempting to adopt this model of localised, computer-mediated retailing. This paper will report on our investigations into how software agents (semi-autonomous program modules that act for or on behalf of people or organisations, Jennings *et al.*, 1998; Maes, 1994), such as the ePSA, might be used to manage and enhance the experiences of the various participants in this form of commerce.

The MAPPA (“Multimedia Access through Personal Persistent Agents”) system was devised as an in-store/mall kiosk based retail information delivery and sales unit. The system was

developed to act as a test-bed to investigate a range of topics associated with the new opportunities that arise from the use of software agent technology within the context of the retail e-commerce environment. These topics of interest include:

- The development and application of software agent technologies and architectures suited to the retail domain.
- How software agent technology could be used to establish strong customer relationships and engender a sense of loyalty as might exist when business is done directly between people and organisations (business to consumer, B2C).
- How to present product information in a manner suited to in-store use, which is both clearly distinct from web-based presentation, but retains the advantages derived from user familiarity and experience.
- How to use character agent technology to enhance the sense of personalisation between the individual customer and host businesses.

This paper now extends the electronic metaphor to include communication between software agents working on behalf of individuals (consumer to consumer, C2C). This paper does not consider business to business (B2B) dynamics, although this is an integral factor in any overall e-commerce community.

We shall use the term “software agent” in two distinct ways. First, that of an on-screen animated character to present information to the user and, second, that of a software resource to manage aspects of the e-commerce activity such as product information management, database access or preference calculation.

The animated character agent (or *synthetic persona*) and associated screen display area in which it operates represents the *front-side* activities of the ePSA, presenting information to the individual user and embodying the personalisation of the system to that user. The MAPP system has been primarily tested in the guise of a wine store, and the majority of examples used in this paper will be drawn from that domain.

The *net-side* of the ePSA is concerned with the transfer and management of data between the user and businesses that have products on offer (B2C) and between individual users who have opinions and preferences to share (C2C) in the connected community. In both cases the relationship is potentially symbiotic. For the vendor business, it is able to better understand the specific requirements of its customers, and can therefore tailor the way information is presented and the products offered to the individual preferences of its customers. Equally, individuals, it is suggested, get increased levels of service from the business.

In the MAPP “kiosk” products on offer from the vendor business are presented on a “virtual shelf”, in which users may browse images of the products, moving backwards and forwards along the shelf and calling the animated character agent to provide additional information about individual products, as they might ask a sales assistant in the physical store. The contents of the shelf may be organised according to selections based on conventional product parameters, or according to the previously established user preferences. The user interface and character agent are considered in section three. We shall also briefly consider technical aspects of the software agent architecture (section four) as it relates to the MAPP system.

We have been concerned to evaluate several aspects of the interaction between the user and the software ePSA. Initially we adopted conventional user study techniques to investigate limited aspects of the effects of the ePSA and preference selections on user perception of loyalty. Recognising the limitations of these methods, we have also used eye-tracking

equipment to analyse user's visual attention while making use of the MAPPA system. Results from these studies are summarised in section five.

We also wish to investigate the potential dynamics of connected e-retail communities. In particular we wish to understand how specific on-line "relationships" can develop, forming sub-groups of individuals within the larger connected community who share similar preferences and whether these groupings can assist others in making better choices more quickly. In preparation for exploitation of these technologies in agent-mediated connected communities, we employ simulation techniques to evaluate the effects of communities with individuals with different preference and behavioural profiles and different strategies for information management. The simulation tool we have developed for this purpose is described in section six, and some results are presented.

In designing and evaluating the MAPPA ePSA, the relationship between individuals and businesses and the connected community in which it will operate we consider four closely interrelated issues: *loyalty*, *trust*, *reputation* and *preference*, which we believe can be used to guide the interaction between individuals and businesses (and their software agents). Each term has a strong connotation within normal, everyday usage. We adopt these terms to derive aspects of the mechanism by which the ePSA will operate, on its user's behalf, within the connected community of which it is a member. Taken together, we consider these to be the "supra-functional" specifications for this system and for e-retail systems in general. These topics are discussed further in section seven.

## 2. Connected Local Communities

With the current emphasis on worldwide Internet connectivity there is perhaps a concern that the traditional notion of "local community" might become compromised as individuals migrate to electronic communication with others regardless of locality. While the "global" approach offers many benefits, this would surely be a retrograde step if it were to completely replace those more traditional ideas of community centred on a particular place. The concept of "local community" itself needs some clarification, as its extent can be seen as being defined by a variety of geographical, social or political boundaries. Such communities can vary widely in size and encompassed population, perhaps starting at the street or neighbourhood level, moving through notions of village or town (or more likely suburban conurbation), to the borough or county level.

Much beyond this and the idea of "local" begins to seem untenable. In the context of this paper, with its emphasis on store-based retailing, a working definition will be adopted based on the idea of that business' *catchment area*, the population from the area that would normally use that business. In this respect we will use as examples businesses that typically have a significant passing trade, such as video, book, music, and (specifically) wine stores. Such businesses may seek to add value to the products they trade in by supplementing it with advice or other incentives with the anticipation of increasing customer loyalty and so enhancing future business. In any case the catchment area for an electronically based connected community must allow for a sufficiently large numbers of participants to form a viable on-line forum if it is to be sustainable.

Early attempts at on-line connected communities were based on the idea of a local dial-up service offering both access to the wider internet and providing local news and information (the Blacksburg Electronic Village; Noack, 1994, for example, or the Milan Civic Network, Casapulla *et al.*, 1995). Schuler (1994) discusses how such initiatives might be used to

strengthen various aspects of community, enhancing local cohesion, improving information for the citizen, give access to education and training and to strengthen democracy by providing information and interaction. In the event, dial-up access adopted a different business model. Internet Service Providers (ISPs) are not often locally based, but serve a wide physical area using non-geographical access charging schemes.

The World Wide Web (WWW) offers a partial solution to the notion of “Cybervillage”, with content and discussion areas specific to a locality. Such “villages” may be provided as a public service by community organisations, such as the “Lawrence Cybervillage” (Lawrence, 2002), a website developed by the local public library. It provides comprehensive access to a wide range of information about the region and services relevant to the local community. Such websites (it is representative of many similar) are largely passive, offering little opportunity for active participation through the browser by members of the local community (although there is a “local history” area for residents to provide local and family anecdotes relating to the area). Other initiatives can take as their starting point city development (e.g. Butte Development Corp., 2002) or various forms of local business promotion (e.g. Saint Paul, 2002).

Recent developments in connected communities (such as the Campiello project, Grasso *et al.*, 1998, and the Living Memory system, LiMe, Stathis *et al.*, 2002) have emphasised the use of large, public display screens to facilitate both general access to local information and to engender community spirit. Campiello is intended to serve areas with high tourist populations (including Venice in Italy), and provides information for both the transient and permanent populations. LiMe concentrates on notions of “community memory”, using information filtering techniques to ensure that documents available from the communal memory can be easily and appropriately accessed. The LiMe technical architecture builds on a variety of software agent types, designed to manage a range of social and technical issues (Mamdani *et al.*, 1999).

Several other approaches to electronic support for community or communal activity have been proposed. Erickson *et al.*, 2002 describe “social translucence”, which attempts to present representations of social cues when teams or work-groups interact remotely on-line. This is achieved through the “social proxy”, a minimalist visual representation of others participating (or not) in a session. Millen *et al.*, 2002 analysed the financial benefits of on-line “communities of practice”, groups defined by similarities in working methods or outlook, in a variety of industry types. Community Organiser (Kamei *et al.*, 2001) provides an intuitive visual representation of chat room participants so that they can easily locate others of similar interest and approach them. Within this view, clusters of individuals with similar interests appear to aggregate into a potential “community”. Sumi *et al.* (1998) generate a graphical view of shared interests by the various participants at exhibitions.

### 3. An e-Retail Interface in the Community

This section describes the appearance of the system to the individual community member user. In designing the MAPPA system we wished to produce an interface design well suited to presenting a large number of products in a highly intuitive and accessible way, but which gave users a graded path into the e-retail community as a whole. That first point of contact would, inevitably, be via a vendor business. This being so, the interface should in some sense directly reflect the shop experience. This experience can operate at three distinct levels, each aimed at drawing the individual further into the world centred about the vendor, but

increasingly acting as a portal to the greater e-retail user community. The casual visitor should be able to visit the shop anonymously, peruse the products on offer, obtain information about them and make purchases. The user should be able to sign-up and then be able to personalise the system to their interests and preferences and so obtain a better service. Lastly the user should be able to share information with others in the community and gain further advantage.

We wished to differentiate the interface from a conventional web-page format, yet not lose useful features common to browser technology that would both be familiar to (many) users and provide useful shortcuts through the system for those who wished to use them. Product presentation is based on the *Rapid Serial Visual Presentation (RSVP)* technique developed by Spence and others (deBruijn and Spence, 2000; Spence, 1999). The screen area also provides a canvas on which the animated character agent will operate.




Figure One: The MAPPA Display Area (left): Detail of Character Agent (right)

Figure one shows the layout of the main MAPPA interface. This interface has been through several design iterations and user evaluation stages, and represents the usual trade-off between providing clear and easy access to functions the user requires against excessive screen clutter. Two items of particular note are the RSVP display area (the area of the screen containing the bottle images, item (1) on figure one) and “James”, the character agent, who appears on the screen to present information to the user (2). In general, it is a design decision that the character agent does not appear unless the user takes some specific action, such as click on a bottle or select a “help” button.

The RSVP display is intended to emulate a “shelf” of products. The user may scan rapidly back-and-forth along the shelf using the tracker-slider control (3) or step left and right using the buttons (4) and (5). Having found a part of the “shelf” with products of interest, the user may “riffle” through those products by running the mouse pointer over them. As each product is touched by the pointer it positions slightly clear of the other images and a brief product description immediately appears above it, emulating a shelf product label (6). James, the character agent, may be called to provide further information by clicking on a product image, the purchase dialog by double clicking on it. The RSVP effect, which is both distinctive and visually effective, is not easily envisaged from the still image of figure one. Distinctive product images may be recognised quickly during the end-to-end tracking movement, and

details obtained immediately during the riffle process. The effect has garnered near unanimous interest whenever demonstrated informally or presented under controlled test conditions.

The character agent is animated over the top of the display. The character agent's utterances are generated in a variety of ways, some scripted within the control program, some derived from a product information database, and they can be generated by a JESS based expert system agent, which has access to both general product (wine) knowledge, and user preferences collected during the current and previous sessions. Information presented by the character agent is both vocalised using a male voiced electronic synthesiser, and appears in a scrolling "speech bubble" adjacent to the agent's head (item (7) and figure one – detail).

Initially there may be several hundred products on the display "shelf". The user may use the drop-down boxes, or call up product refinement dialogs. This is described in the next sections. The MAPPA display may be reconfigured to present alternative product ranges (8), by populating it with a different product database and (matching) product image directory. The display mechanism is best suited to product ranges with largely homogeneous packaging, such as wine bottles, video or DVD packages, CDs or books; or where the product images may be re-sized into regular image areas.

### 3.1 Refining the Selection

This section describes the two ways in the current display area by which the user can directly request sub-sets of the full product range by explicitly applying selection criteria. First the user may make selections from the series of "drop-down" selection boxes (9). Alternatively, the user may initiate detailed selection dialogs from button based selection panels (10). These dialogs appear at the top of the screen, and normal operation of the RSVP screen is suspended while they are displayed.



Figure Two: The "Goes-with" Selector Dialog

The first way is intended to reflect the method available from most websites, each drop-down allowing a number of broad categories. The second (dialog based) way allows for highly detailed or specialised options to be selected. Figure two shows the detailed "goes-with" selector. The user may refine a selection of wines on the basis of food it might best accompany. The main graphic buttons on this dialog access a broad category of dish, reflecting the categories available in the first method. The list box selector within this dialog (whose contents depend on the button selected) allows for a detailed choice to be made. Equally, the user might refine their choice on the basis of region, price, flavour or grape variety, each having its own dialog, each of broadly similar design.

Note, in particular the "Ask James" button on this selection dialog (an equivalent button is placed on the other selector dialogs and on the main screen, item (11), figure one). This

option invokes the appearance of James, who will explain the purpose and use of each facility associated with the dialog.

The ePSA could update the product selection at any time, but the “Find” button (12) is intended to allow the user to request the ePSA to present a new selection based on the criteria they have specified. Criteria may be selected sequentially, allowing the user to cumulatively refine the product search at successively greater degrees of detail.

### 3.2 Expressing Preference

Several possible strategies are available for determining the individual user’s preferences from his or her actions. The first, and most direct, is to present one or more “forms” for the user to complete at the beginning of the session. This may be an effective strategy for expert or highly opinionated users, who have already formed a strong viewpoint and will have confidence in it. However, we consider that this approach would be unnecessarily daunting for beginners and less confident users. This will vary between product domains. Most people will have reasonably strong opinions about the kind of films they like to watch or music they like to listen to, but may be less confident, for instance, about wines. Second, and apparently a direct measure of preference, on the assumption that people will buy things they like, is to review the product purchase history. Third, and least direct, is to monitor the activities of the user while they use the system. This has been proposed as a way of avoiding, or at least supplementing, the need to request explicit ratings while browsing web pages for information content (e.g. Claypool *et al.*, 2001).



Figure Three: Expressing Preferences

In general, we prefer to gather preference information incrementally by fully elective means. We consider that this will enhance the sense of interaction and ownership, and allows users to proceed at their preferred rate. Experts may wish to personalise the system immediately. All the product selection dialogs present two extra controls labelled “Like” and “Dislike”, shown circled in figure three. Pressing either of these buttons informs the ePSA that the individual user prefers (or does not) the current selection that is indicated. This information then forms part of the persistent knowledge the ePSA holds about the customer. The strength of a particular preference (or dislike) may be indicated by repeated application of a preference button and previously held preferences reversed using the opposite sense button.

Many of the product selector panels allow the user to select a broad category (i.e. “France”) with a pictorial button, or (for more knowledgeable users), to specify a detailed sub-category (such as “France-Bordeaux”) from the list-box. The preference always attaches to the specific option selected. In this manner the user may express a general dislike of a general category, such as sweet wines, but a liking for exceptions, such as the sweet Tokaji dessert style, or an individual product, in each instance using the preference buttons attached to the appropriate dialog. The influences on the expression of preference appear to be multidimensional and



there is no reason to suppose these will equate directly with “traditional” hierarchical product classification schemes (that is, general preferences may be overridden by the specific).

Past purchase data is important, but it is not a direct indicator of preference. In particular products purchased as gifts, ideally reflecting the taste of the recipient, represent a true counter-indicator. We would, however, endeavour to obtain preference feedback on a return visit to the system. If a small number of purchases had been made, this can be done directly by recalling the name of the product and asking for preference feedback. Clearly there must be a “no opinion” option. Where a large number of purchases had been made this activity is better interspersed into the remainder of the session. While the MAPPA system can accumulate detailed user activity data (the option was installed for monitoring during our user evaluation trials), we have not applied it in any of our selection methods.

### *3.3 Presenting Recommendations*

Where the user has requested a product selection based on the selection parameters, the RSVP “shelf” can be repopulated temporarily with the sub-set of products requested. Depending on the combination of criteria applied, this can range from about half the range (e.g. red/white for wines) to very small groups of specialised products. As criteria can be mutually exclusive, it is easily possible to leave the shelf empty. Where this occurs one of two strategies is adopted. First the character agent can apologise and suggest another try, second the ePSA can derive a selection that almost meets the requested criteria. Various conditions of this type are identified by the MAPPA wine rules and matching algorithms (such as when a red wine with a “white” grape type are specified), and reasonable substitutions made, allowing the session to continue smoothly and providing comments as how to better refine the selection. The main, business driven, criterion is to keep the user engaged, and to make purchases. Where the user has requested a selection based on recommendation the situation is more complex. Taking a pragmatic stance, “recommendations” would typically arise from one of the following classes:

- 1) The standard, popular or well-promoted items.
- 2) Recommendations from recognised “experts”, who have formulated product rankings (e.g. Gluck, 2002, for wines).
- 3) Special offers and discounted items offering particular value.
- 4) The user’s stated preference profile, as developed over previous sessions.
- 5) The user’s purchase history.
- 6) Collaborative, shared preference recommendations from the “community”.
- 7) General preferences “mined” from the community (product reputations).

Initially, without any prior knowledge of the customer, generic suggestions must be formulated and presented (these are items 1-3 above). As the user builds a personal profile, it becomes possible to introduce selections based on that information (items 4 and 5). Also, as the user builds a preferences profile, and as a “community” of users begins to form, items 6 and 7 become a viable option. Typically we would expect to present 10-15 products in this manner, sufficient to provide a reasonable chance that one or more will be acceptable, yet not so many as to overwhelm. To present a small number, and to have them all rejected, is to effectively invite the user to terminate the session.

One particular advantage of this strategy is that character agent may then verbalise the “reasons” for each choice, enhancing the sense of interaction and personalisation. A typical

commercial grade product database will include brief product descriptions, which for wine may include tasting notes, information about the vineyard or historical information. This data may not be extensive and we noticed that there is a particular need to ensure the character agent is not perceived as being repetitive, as this appears to be particularly damaging to the illusion of personalisation. We also considered it important to carefully check the database of “factual” utterances to ensure that the character agent is not seen to express a personal opinion based on taste or preference – which would clearly appear as ridiculous.

Note that we do not present here any definitive “best algorithm” for balancing these alternative sources for a recommendation. How the vendor balances this choice is very much a business issue. One vendor may be particularly keen to push particular products, another might wish to emphasise personalisation. A business must also consider how much it would like to encourage its customers to try new products and to explore amongst the “I’m not sure if I would like that” class. In the evaluations described in section five, we specifically used method four (“user’s stated preference”) to derive the selections presented.

## 4. An Agent System for e-Retail and e-Community

Figure four illustrates the MAPPa software agent architecture. Each of the major functionalities of the system is served by a corresponding software agent, and each is distinguished by the service type it provides. A combination of these software agent types are invoked and grouped together to perform the functions required for the user’s specific ePSA.

The “Personal Agent” (PA) is responsible for *presentation*, managing interactions with the user, recording information and calling on the “Database Agent” (DA) to provide product and preference data to provide *personalisation* to the current user. The PA generates all the actions and utterances by the character agent and coordinates these with the screen display (described previously in section three), ensuring that the two remain synchronised and that user input is reacted to in a timely manner. To achieve the necessary interactions within the connected e-community some part of the personal agent must be *persistent*, available to the community regardless of whether its user is present or not.

The display (“RSVP”) is coded as a C++ module, connected using a “Proxy Agent” (PrA) to the otherwise JAVA based system. The proxy agent acts as a software “wrapper”, interfacing between otherwise incompatible software modules. To achieve satisfactory visual performance the display code managed the Rapid Serial Visual Presentation product scanning and riffle functions, generating messages to the PA when intervention is required. In turn the PA issues instructions to reconfigure the display as required in coordination with the actions of the character agent.

The Personal Agent (PA) can call on the expert system shell “Jess Agent” (JA) to provide generic expertise on the product range, and to generate recommendations from known or assumed personalisation data held about the user.

In the full model each agent (and in the community, each ePSA grouping) communicates across a FIPA-OS backbone according to the FIPA-ACL (Agent Communication Language, FIPA, 2002) specification. A new implementation of the FIPA standard was created within the remit of the MAPPa project, and the MAPPa system was devised as a platform to evaluate that technology and provide a scenario within which to evaluate the effectiveness of the design issues raised.

The initial intention on this project was to store the user’s personal data on a smart-card (read by CR), as this could be retained by the user and would enhance the sense of privacy for

personal data. In practice, the quantity of data required to personalise the system effectively exceeded the available capacity of the card, and the data would not be available to the community under these conditions. Under these circumstances the card only acts as an access key, the data being retained in conventional persistent storage.

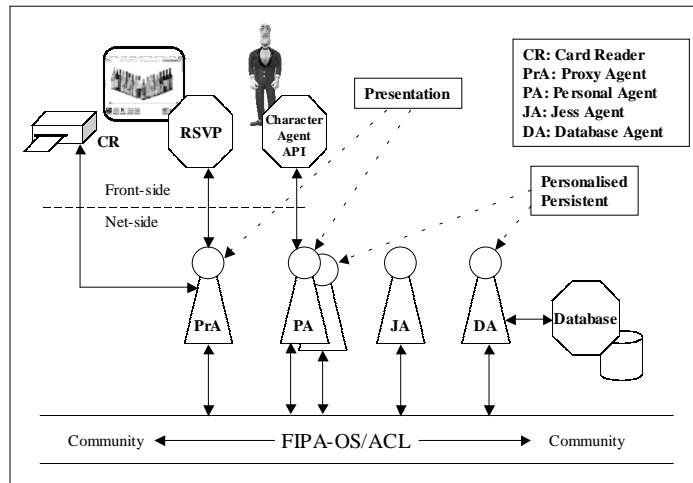


Figure Four: MAPPA ePSA Internal Architecture

#### 4.1 The Character Agent

The MAPPA synthetic persona, “James”, takes the form of a “butler”, which we consider offers an appropriate persona for an electronic Personal Sales Agent in this context. James has three primary roles in the system, (1) to act as a product advisor, (2) to provide help and to assist users whilst using the system and (3) to provide small amounts of “continuity” information about the progress of the session. A relatively large amount of work is required to program and coordinate the character’s agent’s movements with items on the screen area, to provide mannerisms appropriate to the persona and to ensure that the utterances required of the character are consistent both with the persona selected, and with the application and product range being presented. Although a number of alternative persona exteriors can be selected for testing, we have not prepared alternative speech patterns and mannerisms suited to these other characters. Synthetic speech vocalisation remains a problem, not only in the unnatural quality of the rendition, but also in the failure to pronounce various (typically non-English) words. This is particularly noticeable in the wine scenario, as many product names and other pertinent information is of an international nature.

Much work has been done to integrate synthetic personae into electronic display systems, notably in entertainment, CAE and e-retail contexts (see Dehn and van Mulken, 2000 for an overview). Cassell *et al.* (2000) describe Rea, a conversational synthetic persona which will accept speech input from users and generate task oriented dialogues mediated by detecting conversational cues (via video input) from the user. Rea has been demonstrated acting as a real estate sales advisor. Kießling *et al.* (2001) describe COSIMA, a sales support system presented by a speaking synthetic persona in the form of a young woman. The COSIMA system performs partial match searches from product databases, and generates conversational output based on the completeness of the match found. Kamyab *et al.*, (2001) are developing a software agent based system for virtual sales assistants that use high-level abstractions for representing, reasoning and planning human-like dialogs. Despite the number of systems that have been proposed and developed, the number of empirical evaluations of users’ reactions to

these characters is still relatively small and their interpretation can be problematic (McBreen *et al.*, 2000; van Mulken *et al.*, 1998).

## 5. Interface Evaluations

We have conducted two separate user evaluations with the MAPPA system. The first was aimed at understanding the perception by the user of the system, and specifically to determine whether the character agent and preference mechanism would have an effect on user loyalty in the e-retail context (Witkowski, Pitt, Fehin and Arafa, 2001). These investigations were conducted using conventional usability evaluation techniques. With a complex system such as MAPPA it is always difficult to extract reliable data and there are many sources of potential uncertainty. Nevertheless, we consider these investigations a worthwhile exercise (and much better than the alternative of not doing them) and summarise specific aspects of the procedure and results here.

The second was an investigation into the visual attention of users during a MAPPA session (Witkowski, Arafa and de Bruijn, 2001). This was performed using eye-tracking equipment (Eyelink, 2002), which gives a continuous indication of where on the screen the user's gaze is directed throughout the entire session. The purpose of this study was to determine the relative time spent by individual users attending to the different aspects of the interface, the products, the character agent, and searching for and using different interface features.

### 5.1 User Sessions

We asked 10 subjects to attend a MAPPA session lasting between 45 and 55 minutes. The subjects were obtained (and paid) through a local agency, and were self-selecting as being "somewhat" or "very" knowledgeable about wines. During this time the subjects were given an opportunity to familiarise themselves with the system, with the help of a trained facilitator if required.

The session was in two parts. In the first part the subjects were asked to enter various wine preferences into using the interface described. In the second part, the system presented wines on the basis of the preference options entered and the subjects were asked to inspect the choice and comment on its accuracy. The subjects were then asked to rank their preference for an e-store on three criteria: the ability to represent product preferences, to discount prices and a loyalty points scheme. On a scale of 1-3 (1 better) the subjects ranked preference first (average 1.4), discount second (1.7) and a loyalty points scheme a clear third (2.9). This was interesting as, when the subjects were initially asked to rank price, quality, shop advice, friend's advice and shop staff friendliness, they had consistently placed the price and quality criteria above advice.

At the end of the session the subjects were asked to rank the character agent on five characteristics (1) the extent to which it seemed "personal", (2) "pleasant", (3) "helpful", (4) "friendly" and (5) the extent to which they found the agent "engaging" on a seven point Lickert style scale. With the exception of one subject, who expressed a clear dislike of the character agent, the remaining subjects, on average, rated the agent above the median line. Additionally 8 of the 10 subjects preferred the character agent to alternate use of speech or text only output. These results are somewhat surprising as previous studies (Höök *et al.*, 1999; van Mulken *et al.*, 1998) have generally indicated a much stronger dichotomy between users who did and did not like animated character agents.

## 5.2 Eye-tracking

Visual attention, as recorded by eye-movements is partially volitional, under the “conscious” or active control of the participant, and partially autonomous, selected by an attention mechanism inherent in the brain’s visual system. Human eye movements are not smooth, but proceed as a series of *saccades*, rapid, ballistic, movements (usually taking between 30-120mS) interspersed with *fixations*, periods of low eyeball acceleration (typically lasting between 200 and 600mS). It is widely believed that visual processing only takes place during the fixation component (there are other eye-movement types, such as *tracking*, but we do not consider these here). By analysing the movements and fixations, it is possible to determine the relative attention being paid to different aspects of the interface.

Subjects were assisted during a MAPPA session (lasting between 5 and 10 minutes) and encouraged to use the various features of the system, calling up the character agent in both helper and advisor modes and using the RSVP product display. Considering first the analysis of fixations during an appearance of the character agent as product advisor (where a product is described) we recorded that users spent an average of 21.4% of the total gaze time looking at the character agent, 61.4% reading the text bubble, 11.8% looking at the product and the remaining 5.4% looking at other areas of the display (5 subjects, 34 character agent appearances). We conjecture that the time spent attending to the agent and text bubble can be interpreted as an indicator that the message being presented was being noted.

Considering all character agent appearances, the subjects spent more than 50% of total gaze time reading the text bubble and nearly 20% looking at the agent. There is little doubt that the appearance of the character agent attracts the user’s visual attention, the first fixation on the agent being less than 1 second from its appearance on the screen in 50 out of the 77 instances analysed. It is interesting to note that the users looked at the character agents “face” for over 76% of the time they attended to the agent, and 13% at its gestures (the remainder at the “body” area). The fact that these fixations are not randomly distributed across the character agent’s parts suggests that the subjects interact with it in a “social” manner.

## 6. Simulating Connected Community Interactions

This section describes our agent simulation package. Its purpose is to allow us to investigate the effects on an e-retail community in anticipation of an actual rollout to a larger community. In particular it allows us to consider a wide range of possible scenarios, including:

- The effects of different agent strategies for preference modelling and recommendation generation.
- Community growth and stability given different ratios of participant types, and different mixes of preference profile.
- The introduction and withdrawal of user participants and products.
- Effect of shifting preferences by individuals, simulating the effects of peer recommendation/pressure and other factors.

The observation that individuals share their many of their preferences with others (form into cliques) gives rise to notions of *recommender systems* based on *collaborative filtering*

methods (Alspector *et al.*, 1998; Herlocker *et al.*, 1999; Konstan *et al.*, 1997). Such systems identify individuals or clusters with similar tastes and preferences across the web or from (typically large) databases and suggest items drawn from the purchases or recommendations made by the clique. These techniques are now becoming widespread, increasingly forming the basis of choice recommendation strategies for agent systems and web delivery alike, and are likely to have significant economic impact (Fink and Kobsa, 2000). These applications have large catchment populations, our main concern is with smaller (e-retail) operations, and particularly the initial conditions under which they will operate.

The notion of a software agent in the simulation is distinct from that described in the MAPPA system. Here we are concerned only to extract out the fundamental algorithmic methods for choice and recommendation generation, and to only simulate those aspects of the (human) user actually necessary to investigate those strategies. Given the substantial investment required to implement and deploy a full e-retail system, it is critical to understand the effects of algorithmic choices prior to deployment. We assume that various individuals within the community will adopt different approaches to how they interact with the system (some adventurous, some cautious, for instance) and that the mix of different types will have a significant effect on how the on-line community develops. We also assume that the community will not be static, that individuals will join and leave, and that new products will be introduced and old ones withdrawn. We do not assume that the choice generating profile of the individuals is static, but that it may be influenced by peer recommendations and pressure, and in response to advertising and other marketing strategies.

## 6.1 Modelling the User

Individual users are modelled as having a range of different characteristics (related to, but not equivalent to their “personality”). These characteristics are intended to reflect how real users may react to various circumstances while using an e-retail system:

- Strength of their internal preferences (some have strong, some weak opinions).
- Willingness to provide recommendations.
- Willingness to be influenced by the recommendations of others.
- Willingness to be influenced by a product’s reputation.
- Willingness to experiment with new products (vs. product conservatism).

Each of these factors can be expressed as a probability that a choice will be made, or a threshold (as appropriate). The model assumes that the initial preferences of the users are pre-defined. This is the *user preference profile*, which is not known to the agent, but must be discovered or estimated by the agent during the simulation. The simulator allows the experimenter to create a population of users with individual (but differing) preference profiles and to define each individual’s characteristics as above. In a typical population to be simulated users will be given different roles. The “leader role” represents a user prepared to experiment with many wines, and provide their recommendations to the community. The “follower role” represents a user who mainly listens, and is influenced by those recommendations. The experimenter can define a range of types, modifying their predisposition to the characteristics listed above. Non-participants, those prepared to neither offer nor use recommendations, are not modelled.

## 6.2 Modelling the Agent

In the current simulator all agents are assumed to be the same at the beginning of a run, and must discover and adapt to the preference profile and other characteristics of its user. The internal characteristics of the agents can be adapted to modify the type and number of recommendations the agents will make to their respective users.

Each agent maintains four data-structures that represent the prevailing “beliefs” of the agent about its user and its peers in the community.

- The *user-preference structure*: records the user feedback about products that have been tried, and for which a preference has been expressed.
- The *peer-opinion structure*: records the recommendations made by the agent’s peers (reflecting their user’s stated preferences).
- The *shared-preference structure*: records the degree of commonality between its user’s preferences and those of its peers represented in the peer-opinion structure.
- The *product-reputation structure*: records an ordered list of the existing product range. Product reputation is calculated on the basis of the strength and number of peer testimonies, and the perceived credibility of the source of those assertions (“*The credibility of a belief depends on the credibility of its sources, evidences and supports*”, Castelfranchi and Falcone, 2001).

## 6.3 The Simulation Cycle

The simulator repeatedly iterates a cycle of events that would occur in the community. This trading cycle represents a self-contained time slice consisting of the following distinct stages:

1. The agent presents a set of products to its user. This step emulates the presentation of products on the MAPP terminal.
2. The simulated user chooses a product from those presented. The selection is made according to the individual user’s selection characteristics, and gives feedback on that product, based on the experimenter-defined user preference profile.
3. The agent updates its user preference structure based on this feedback.
4. The agent communicates with its peers, asserting its user’s belief about the product tried (if a knowledge sharing user), and receives testimony from its peers.
5. The agent updates its belief about each peer’s ability to make product recommendations (the shared preference structure) and updates its product reputation structure.

The experimenter defines the number of cycles to be simulated.

## 6.4 Experiments

The experiments reported here investigate the effects of different strategies for the generation of product reputation. In these simulations we define a “community” of 50 user/agent pairs. Four users are configured as “leaders” (always try new products, always share their opinions based on their internal preferences). The remainder are initialised as “followers”. These are not risk takers. In the absence of any recommendations from an agent/user whose recommendations they trust will purchase a wine they already know and like. The user

preference profiles are created such that each user liked and subsequently disliked half the product range and all the user preferences profiles are different, but overlap. This particular simulation set-up assumes a product type (wine) that will be purchased repeatedly, if liked.

#### 6.4.1 The agent shared preference strategies

We define *subjective product reputation* (SPRep) as the current estimate, by the specific user’s agent, of the potential desirability of a product to its user. Thus, SPRep(A,m) is the belief (of the Agent A) of the desirability of untried product (m) to its user (held in the product reputation structure). The agent uses its belief of the reputation of an untried product to rank that product against other untried items. This ranking would then be used to determine which products should be presented to the user (via the ePSA/MAPPA front-side, for instance). This reputation rating is intended to assist the user in their selection behaviour. In this simulation experiment the “user” chooses from the recommendations. In simulation, follower user types only select a new product if it has a high reputation rating (otherwise they opt for a safe, known, choice). Leaders take the untried product with the current best rating (until they have tried everything).

The community as a whole contributes to the reputation of a product. This occurs because members of the community are capable of communicating with each other their opinions (this creates the peer opinions structure of each agent). In the following example we examine the mechanism for determining the reputation of a single product, this is repeated for all the products in the community to create the agents product reputation structure.

$$\text{SPRep}(A,m) = \sum_{i=1}^n (\text{Rec}(i,m) \times \text{DoSp}(A,i)) \quad (1)$$

The function formulates the reputation of the product m according to agent A (SPRep(A,m)), from the influences of agent A’s peers. Equation 1 takes the form of a weighted sum in which assertions received from the agents peers regarding product m (Rec(B<sub>n</sub>,m), from the peer opinion structure), are weighted by the agent A’s belief in each of the n<sup>th</sup> peer agent’s (B<sub>n</sub>) ability to make recommendations (DoSp(A,B<sub>n</sub>), from A’s shared preference structure). This value is updated from the viewpoint of each agent (taking the role of A) whenever a new recommendation is received, such that each agent formulates a recommendation list of all untried products for its associated user.

*Pre-filtering* aims to remove from the weighted sum any recommendations from peers with a low degree of shared preference with the agent’s user. Recommendations from peers with less preference correlation are stopped from polluting the agent’s reputation evaluation mechanism. Two methods of pre-filtering are investigated. (1) The “Cut-off” Filter, taking into account recommendations only those peers whom the agent believes share a degree of preference greater than a threshold  $\theta$ , to contribute to the reputation of products. Any recommendations from agent with a shared preference value of less than the threshold are excluded. In method (2), “college size” filtering, the agent rejects all recommendations from agents other than the top n rated according to the current evaluated shared preference value.

#### 6.4.2 Simulation Results

Figure five shows the results from the cut-off filter and figure six the results from the college size filter. In each of the two figures the left hand graph indicates the cumulative number of wrong choices (i.e. where the agent recommends and the user accepts a product that is subsequently found not to be to the user’s taste). A lower number is better. The right hand



graph in each of the figures indicates the cumulative number of new products the user has tried during the simulation. Each figure compares the results from three separate simulation runs. In figure five the threshold,  $\theta$ , is set to 0.4, 0.6 and 0.8. Both leaders ('L0.4', 'L0.6' and 'L0.8', indicated by the plus, cross and diamond markers respectively) and followers ('F0.4', 'F0.6' and 'F0.8', indicated by dash, circle and triangle markers) are charted. In figure six the college size is set to 10, 15 and 20 peers for both leaders ('L10', 'L15' and 'L20', plus, cross and diamond markers) and followers ('F10', 'F15', 'F20', dash, circle and triangle).

The leaders always attempt to try a new product. Leaders are sensitive to product reputations, and this will influence the order in which they sample products. This is indicated by the slight bowing of the wrong count for L0.4 (and to a lesser extent for the other leaders) in the left hand graphs. Followers are heavily influenced, and the cumulative wrong choice count is substantially reduced in every case. The difference in final result, comparing the final wrong choice counts to those tried is negligible. Both result in a reduction to approximately 25-30% wrong in the followers. However it is the dynamic behaviour of the system that provides the determining factor. In the cut-off algorithm the agents F0.6 and F0.8 initially trust none of their peers to make recommendations and so rank all the products to have zero reputation. Consequently their users don't experiment and the system exhibits a period of stagnation. Figure five demonstrates this delay in product exploration. With the college size algorithm the users are encouraged to try products earlier, so helping to form relationships quickly within the community.




Figure Five: Cut-off Filter

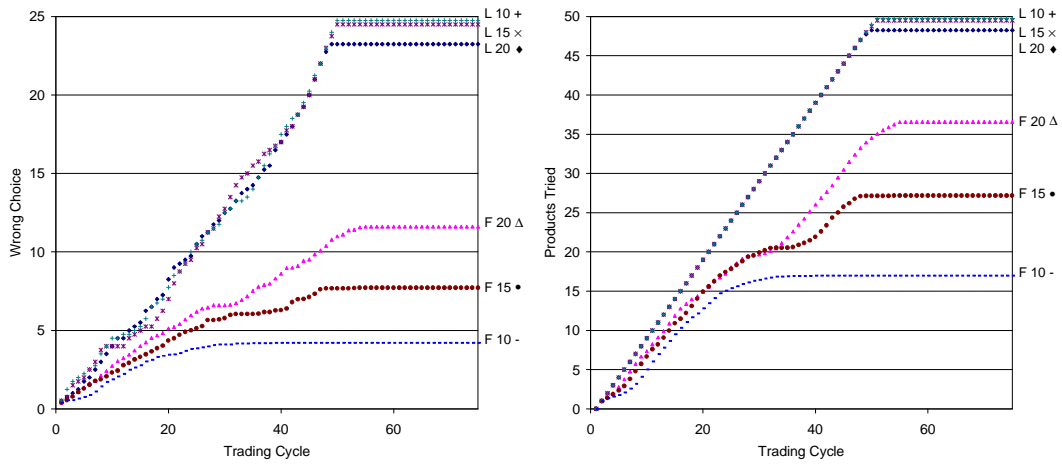


Figure Six: College Size Filter

Figure seven shows the development of the shared preference structure (from the perspective of agent/user zero). The left hand figure indicates the results from one of the four L10 agents from figure six. The right diagram is simulated under identical conditions, except that the users initial preferences are in five distinct groups, each with a leader. The chaotic behaviour at the start of the simulation indicates the importance of a good generic selection (section 3.3) during early visits. We assume in these simulations that the users are naïve, which will, of course not always be the case. Existing preferences, notified to the agent, can be used to pre-bias the user-preference structure (as though this had been entered via the MAPPA interface, for instance).




Figure Seven: 50 Agent Simulation Runs (from agent 0's perspective)

Figure eight (left) shows agent 48's evaluation of the shared preference value it holds relative to the other agents at cycle 9, the right hand diagram is the view from agent 20 at the same

time. These are stills from an animation of the data collected during the simulation run shown in figure seven (left).

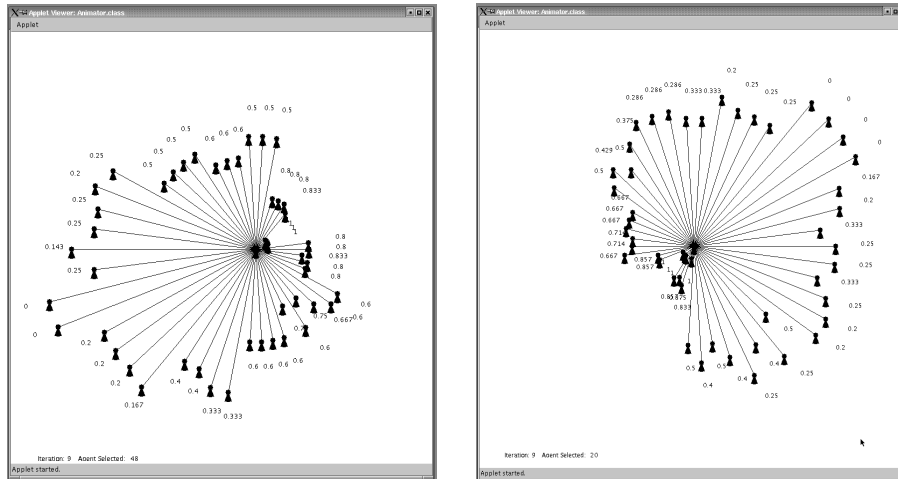


Figure Eight: Preference Neighbourhoods According to Agents 48 (left) and 20 (right)

## 7. Discussion

In this section we look at four closely related topics - loyalty, trust, reputation and preference - germane to the use of software agents in connected communities. These issues have in part motivated the design and implementation of the MAPPA system, and our view of their significance has been influenced in the subsequent development and investigation of the system. We identify these four topics as representing the “supra-functional” requirements for an e-retail business system. They stand beyond the more conventional notions of functional or technical specifications, such as the method of interconnection between software agents or the design of the user interface. Having identified them, each technical and functional specification and each business criteria can then be evaluated in the context of the anticipated impact it will have relative to each of the supra-functional requirements.

Loyalty refers to the willingness of entities to continue interacting, trust to the degree of reliance each may place upon the other, reputation to the shared opinion of entities and things and preference a personal assessment (by a person) of their disposition towards (almost) everything. In the current context we seek to find ways of exploiting these properties in the e-retail context and how to delegate these activities to software agents, so that they can better serve our needs. Loyalty and preference are specifically human qualities. Agents can be used to engender the first, and reflect the second. Software agents can directly evaluate measures for trustworthiness and reputation, and make “autonomous” decisions on behalf of their users based on these measures.

### 7.1 Loyalty, the Customer $I^2P^2R^2$

We believe that systems such as MAPPA can be used by retailers to create customer loyalty by adding value, if they are used and presented in appropriate ways. The problem is acute for most web based e-commerce operations, as each customer is only ever a few moves from being connected to a competitor, but a significant issue also for location based retailers. MAPPA adopts the  $I^2P^2R^2$  (Information, Interaction, Personalisation, Privacy, Respect and

Re-use) model of loyalty drivers (Arafa *et al.*, 2000). By understanding the value of information, interaction, personalisation, privacy, respect and re-use in creating customer loyalty (Vargas, 2000), we argue that this value can be transferred to the e-commerce domain. In this case we refer to loyalty as the “perception of loyalty” by which customers are willing to return to an electronic-based retail store. This is a critical issue of Customer Relationship Management (CRM) for the vendor, as the repeated return of loyal customers should lead directly to increased sales. More significantly, their non-return effectively ensures loss of potential sales.

The key issue here is personalisation, and the extent to which it can influence the user into a sense of “one to one relationship” with the retail operation. The principal conjecture is that by personalising an e-retail system to the individual customer, the customer will find it easier and more enjoyable to do business with the current retailer, and so prefer that business over other, competitor, retailers. Where personalisation is achieved through extra staff the concern has been that the additional costs involved have outweighed the advantage. Software agent systems offer the possibility that the advantages can be obtained, without incurring the majority of the costs.

Achieving personalisation clearly requires a degree of interaction and the accumulation of information. Interaction requires particular attention to the form of the interface. If the sense of personalisation is to be maintained the interface (in the current context the character agent) must appear to remember and use information that it has received. The accumulation of information must always raise privacy concerns. The MAPPA system advocates a smart-card access mechanism (like a bank cash-card), though if information about preferences is to be shared, some information must be available to other agents even when the individual user is not directly available. The appropriate balance, we suggest, can be achieved by the careful acquisition and management of permissions and the use of net-side aliases.

Mutual respect cements the relationship between two parties, and is a clear indicator when expressed between people, or between a person and an organisation. It remains to be seen if an equivalent relationship can be defined between a person and a software agent. The “re-use” of components: smart-card access, the same character ePSA across stores and the adoption of uniform presentation methods can give the user a sense of familiarity and reduce hurdles to participation.

## 7.2 Trust

Trust can be considered to be the establishment of a form of mutual dependency relationship between two entities, but also as the assessment by one or both parties that the other party will perform in the manner to which they have contracted. In this context then, we note that it appears as the reciprocal notion to loyalty. Where the vendor commits to offering a good product range, to offering it at a good price, to having it in stock, and providing reliable information about it and keeping private information secure, the customer can evaluate the vendor on those factors. If the vendor acts in a trustworthy manner, the user will have reason to show them loyalty – “*trust begets loyalty*” (Witkowski *et al.*, 2000). We note that the peer-to-peer relationship between members in the connected community should be viewed differently in the current context. The expression of a preference by an individual should not be seen as an obligation to any third party, that you do not share their opinion is not a reason to suppose them untrustworthy (though it may be a good reason to avoid their recommendations in the future).

### 7.3 Reputation

It is self-evident that trust is not a transitive property, that *A* trusts *B* and *B* trusts *C*, does not allow us to conclude that *A* will trust *C* (Jones and Firozabadi, 2002). Yet it is equally clear that *B*'s view of *C* will influence *A*, because of the trust *A* holds for *B*. In this context, reputation therefore becomes the adoption of some other entity's trust in a third party as a proxy or substitute for a directly observed trustworthiness measure. Distinction is made between reputation as applied to individuals or businesses in the community and product reputation. Reputations are by their very nature shared commodities, can be manipulated, are "fragile" and have true commercial value. Agent based systems are well placed to evaluate reputation based measures (e.g. Zacharia, 1999), as they can autonomously quiz other agents representing other members of the community for the "opinions" of their users and form group measures to formulate an overall assessment. The community as a whole contributes to the reputation of a product. This occurs because members of the society are capable of communicating with each other their opinions. Reputation can therefore be represented as a belief formed as a result of the testimonies of peers.

### 7.4 Preference

Last, we look at preference. The user's actual (stated) preference for, or against, something becomes the final arbiter of any selection or action made by a software agent on its user's behalf. An individual's preferences appear to be special to that individual and seem to pervade every aspect of that individual's choice making process. An individual may, of course, not always have the opportunity, or be willing, to actively express their true preference in a variety of circumstances. Users express their preferences in the MAPPA system through the user interface. By actively expressing a preference for some product, or class of products, and allowing their ePSA to either pass that information to a vendor or use the information locally the ePSA is able to present refined product selections, offering the user a more focussed choice. The ePSA can also use this preference profile to identify other members of the connected community who have similar preference profiles and obtain recommendations from them. Equally, the ePSA can, if permitted, make the product preferences of its owner available, and contribute to the community as a whole. Whether or not the user wishes to participate in this way is, naturally, a matter of personal preference.

Preference is at the heart of loyalty, trust and reputation. An individual may prefer to do business with an organisation he or she trusts, over one that offers lower cost. Similarly the individual may prefer to accept a shared reputation value as part of the purchase decision process, or not. The challenge, then, for the ePSA designer is to accurately capture the preferences of the associated user and apply this to the mutual benefit of that user and the retail business.

There is, of course, no right or wrong answer to the question "did you like *X*?" It is a subjective appraisal, by the individual, of their current disposition towards *X*. Alternatively, the individual may not feel particularly strongly one way or the other, or may present a false answer (for a variety of reasons). We are also not concerned here with how or why the individual came to like or dislike a particular product, but assume that the mechanism responsible for giving rise to preference expression is a product of natural inclination mediated by prior experiences. An individual's preferences can change with time, and in response to good and bad associations with the product. It is also clear that individual and group preferences can be influenced and manipulated in a variety of ways (marketing and

advertising) and that personal preferences can be influenced by peer group pressure and recommendations. Traditionally the analysis of preferences has based on multidimensional analysis of paired comparison data (Schonemann and Wang, 1972; Coombs, 1964). In the case of the MAPPA system we did not consider this an appropriate approach as users are only being asked to express an absolute like/dislike (or don't know) opinion about each product in isolation.

## *7.5 The e-Retail Conjectures*

This paper has, in part, been based on the conjecture that technology can help to restore a sense of local community and personal service that might otherwise be lost in an e-retail context. By physically locating a kiosk style terminal within a store or mall context there is a clear invitation to members of that community to engage with businesses, and by proxy other individuals, in that e-retail community. For this conjecture to succeed each party must accrue some significant benefit, over and above the available alternatives.

We also conjecture that businesses will benefit from an increase in perceived loyalty from their customers. By offering apparently personalised advice, by enabling access to a wider connected community of other individuals prepared to share their product preferences and by careful management of all the  $I^2P^2R^2$  issues relating to individual customer engagement with the business sales will be increased. Equally, by being presented with such advice, in being part of a wider connected community and by being treated in an appropriate manner the individual user must perceive sufficient advantage to overcome any initial inertia and the effort required to participate in a new technologically based approach.

The MAPPA system has also been based on a number of technical conjectures, relating both to front-side (user interface) design and net-side (community interface) design. In adopting these conjectures, as must be the case in any research funded activity of this nature, the work exposes a number of interesting conflicts and dilemmas for both the designers and adopters of such systems, and the users, at whom they are targeted. Clearly a project of this type must also make a large number of essential design decisions, about the nature of the interface, the use of front-side animated character agents, the net-side software agent technology, the selection of particular algorithmic methods for product recommendation (either in the individual or shared preference senses), and so on.

Although not directly related to either interaction or connected community issues, business and social considerations rather than technological ones will, we conjecture, determine the rate and form of uptake in terms of the connected e-retail community. The size of the user base that can be attained and the rate at which it can be built will surely be significant factors in this equation. Too small an active catchment population and the pressure from Return on Investment (RoI) requirements become excessive, too large and the nature of the community being formed changes (but this may not be all bad). Equally, the larger the e-community the more focussed the shared preference model becomes, and the more effective conventional collaborative filtering techniques become. What an individual store could not do, a national chain might justify on the grounds that it differentiates their business from that of their competitors at a significantly lower unit cost.

Cooperation between businesses within a physically co-located community, such as a shopping mall, will undoubtedly raise significant commercial challenges, even assuming they do not compete in the product space directly. However, with cooperation, unit costs could be managed, more terminals funded and greater publicity and awareness given to the scheme.

This form of cooperation would also provide added justification for adopting a software agent architecture in order to integrate different data sources and host systems.

There remain formidable and challenging technical issues relating to the use of net-side agent architectures relating to standards development and interoperability. Advantages clearly arise from the development of standards for platform independence and communications protocols, as well as essential support services, such as disparate database access, agent directory services and clear ownership of data (Bellifemine *et al.*, 1999; FIPA, 2002).

The use of character agents to present information also remains problematic. The majority of current research indicates a strong dichotomy of opinion across the broader potential user base. Some people clearly like the additional “personalisation” brought about by the appearance of an animated character agent, some, equally clearly, do not like them. To insist that all users are required to use the character agent would undoubtedly be counterproductive. The MAPPA system allows for the character agent utterances to be displayed in a text box (optionally with speech output generation). At present it remains largely a matter of faith, and user preference, as to whether the use of character agents adds to an e-retail experience.

Research continues into improving the interactive experience between user and character agent. André *et al.*, (2002) describe work to incorporate affect (the expression of emotion) into character agents and Kamyab *et al.*, (2001) advanced conversational abilities. The available evidence suggests that there is a strong correlation between the realism of the character rendition and perceived expectation of intelligence, putting extra effort and cost into the business cycle as ever greater attention to detail is required to maintain the persona’s credibility. In the final analysis, though, it is the preference of users and customers that must decide the fate of such systems – simple or complex.

## 8. Summary and Conclusion

We have considered the use of software agent technology as a supporting technology in the formation of e-retail based connected local communities. At the heart of this approach is access to retail outlets, and to a larger “community” of other users, through kiosk-style interface stations. These augment the physical presence of the retail store, and provide individual customers the opportunity to personalise their interaction with that store. We have used the MAPPA system as an example, which provides an interface incorporating an animated character agent (“James, the butler”) to interact with the customer. By expressing their product preferences to the system customers can achieve an improved level of service through personalisation. By allowing their software agent to share their preferences with others in the community, the user can both contribute to that community, and take advantage of recommendations from others.

There are many questions left unanswered about the role of these technologies in the localised retail situation, not least the effect of the World Wide Web. The main advantage it has, and that which the retail operations can best exploit, is the very fact of their locale. As it stands, the system is still largely impersonal – information is shared, but largely anonymously. The opportunity remains to coordinate these on-line activities with real events: wine tastings and wine clubs at the wine store, and readings and discussion groups at bookstores (Womak, 2002) - the challenge, how to do it in a way that enhances and retains the trust and loyalty of the community members.

# Acknowledgements

Development of the MAPPA system was financially supported by the European Union ESPRIT Programme (EP28831). Thanks are due to the many individuals who contributed to the work described here, including Alexander Artikis, Oscar de Bruijn, Simon Martin, Patrick Fehin and Gabrielle Dionisi. Particular thanks are due to Yasmine Arafa, who programmed the behaviour of the character agent “James”.

# References

- Alspector, J., Kołcz, A. and Karunanithi, N. (1998) “Comparing Feature-based and Clique-based User Models for Movie Selection”, in Proc 3<sup>rd</sup> Conf. on Digital Libraries, pp. 11-18.
- André, E., Arafa, Y., Costa, M., Figueiredo, P., Gebhard, P., Höök, K., Mamdani, A., Martinho, C., Mourão, D., Pavia, A., Petta, P., Sengars, P. and Vala, M. (2002) “SAFIRA – Supporting Affective Interactions in Real-time Applications”, in: Proc CAST-02, 5pp.
- Arafa, Y., Dionisi, G., Mamdani, A., Pitt, J., Martin, S. and Witkowski, M. (2000) “Towards Building Loyalty in e-Commerce Applications: Addressing Issues on Personalisation, Persistence & Presentation”, in Agents in Industry Workshop at Autonomous Agents 2000, 8 pp.
- Bellifemine, F., Poggi, A., and Rimassa, G. (1999) “JADE - A FIPA-compliant Agent Framework”. In Proc. 4<sup>th</sup> Int. Conf. on the Practical Application of Intelligent Agent and Multi Agent Technology (PAAM-99), pp. 97-108.
- Butte Development Corp. (2002) “Butte Cyber Village: Great Business Location”, hosted at: <http://www.cybervillagebutte.com/>.
- Casapulla, G., De Cindio, F. and Gentile, O. (1995) “The Milan Civic Network Experience and its Roots in the Town”, in: Proc. 2<sup>nd</sup> Int. Workshop on Community Networking, ACM SIGCOMM, June 1995. (<http://www.retecivica.milano.it/paper/ieee95.htm>).
- Cassell, J., Bickmore, T., Vilhjálmsón, H. and Yan, H. (2000) “More Than Just a Pretty Face: Affordances of Embodiment”, in Proc. 5<sup>th</sup> Int. Conf. on Intelligent User Interfaces (IUI-2000), pp. 52-59.
- Castelfranchi, C. and Falcone, R. (2001) “Socio-cognitive Dynamics of Trust”, in Technical Report, Alfebiite Deliverable D1.
- Claypool, M., Le, P., Wased, M. and Brown, D. (2001) “Implicit Interest Indicators”, in Proc. 6<sup>th</sup> Int. Conf. on Intelligent User Interfaces (IUI-01), pp. 33-40.
- Coombs, C.H. (1964). “*A Theory of Data*”, John Wiley & Sons.
- de Bruijn, O. and Spence, R. (2000) “Rapid Serial Visual Presentation: A Space-time Trade-off in Information Presentation”, in Proc. Int. Conf. on Advanced Visual Interfaces (AVI-2000), pp. 189-192.
- Dehn, D.M. and van Mulken, S. (2000) “The Impact of Animated Interface Agents: a Review of Empirical Research”, *Int. Journal of Human-Computer Studies*, **52**, pp. 1-22.
- Erickson, T., Halverson, C., Kellogg, W.A., Laff, M. and Wolf, T. (2002). “Social Translucence Designing Social Infrastructures that Make Collective Activity Visible”, *Communications of the ACM*, **45**(4), pp. 40-44.
- Eyelink (2002) “Eyelink Gaze Tracking”, hosted at: <http://www.smi.de/el/elsys.htm>.



- Fink, J. and Kobsa, A. (2000) "A Review and Analysis of Commercial User Modeling Servers for Personalization on the World Wide Web", *User Modeling and User-Adapted Interaction*, **10**, pp. 209-249.
- FIPA (2002). "The FIPA Agent Communication Language Specifications", hosted at: <http://www.fipa.org/repository/aclspecs.html>.
- Gluck, M. (2002) "*Superplonk*", London: Hodder and Stoughton.
- Grasso, A., Koch, M. and Snowdon, D. (1998) "Campiello – New User Interface Approaches for Community Networks", in Proc. Workshop Designing Across Borders: The Community Design of Community Networks (Schuler, D., ed.), Seattle, WA, Nov. 1998.
- Herlocker, J.L., Konstan, J.A., Borchers, A. and Riedl, J. (1999) "An Algorithmic Framework for Performing Collaborative Filtering", in Proc. 22<sup>nd</sup> Int. Conf. on Research and Development in Information Retrieval (SIGIR-99), pp. 230-237.
- Höök, K., Persson, P. and Sjölander, M. (1999) "Measuring Experience of Interactive Characters", in Workshop on Behavior Planning for Life-like Synthetic Characters, I3 Spring Days.
- Jennings, N. R., Sycara, K. P. and Wooldridge, M. (1998) "A Roadmap of Agent Research and Development", *Journal of Autonomous Agents and Multi-Agent Systems*, **1**(1), pp. 7-36.
- Jones, A.J.I. and Firozabadi, B.S. (2002) "On the Characterisation of a Trusting Agent - Aspects of a Formal Approach", in Castelfranchi, C. *et al.*, (eds.), Trust and Deception in Virtual Societies, Kluwer Academic Publishers, pp. 163-174.
- Kamei, K., Jettmar, E., Fujita, K., Yoshida, S. and Kuwabara, K. (2001) "Community Organiser: Supporting the Formation of Network Communities Through Spatial Representation", in IEEE 2001 Symposium on Applications and the Internet (SAINT 2001), pp. 207-214.
- Kamyab, K., Guerin, F., Goulev, P. and Mamdani, E. (2001) "Designing Agents for a Virtual Marketplace", *AISB Journal*, **1**(1), pp. 61-85.
- Kießling, W., Fischer, S., Holland, S. and Ehm, T. (2001) "Design and Implementation of COSIMA - A Smart and Speaking E-Sales Assistant", in 3<sup>rd</sup> Int. Workshop on Advanced Issues of E-Commerce and Web-Based Information Systems, Milpitas, California, 2001.
- Konstan, J.A., Miller, B.N., Maltz, D., Herlocker, J.L., Gordon, L.R. and Riedl, J. (1997) "Applying Collaborative Filtering to Usenet News", *Communications of the ACM*, **40**(3), pp. 77-87.
- Lawrence (2002) "The Lawrence (KS) Cybervillage", hosted at: <http://www.ci.lawrence.ks.us/>.
- Maes, P. (1994), "Agents that Reduce Work and Information Overload", *Communications of the ACM*, **37**(7), pp. 31-40.
- Mamdani, E., Pitt, J. V., and Stathis, K. (1999) "Connected Communities from the Standpoint of Multi-agent Systems", *Journal of New Generation Computing*, **14**(4), pp. 381-393.
- McBreen, H., Shade, P., Jack, M. and Wyard, P. (2000) "Experimental Assessment of the Effectiveness of Synthetic Personae for Multi-modal E-retail Applications", in Proc. 4<sup>th</sup> Int. Conf. on Autonomous Agents (Agents-2000), pp. 39-45.
- Millen, D.R., Fontaine, M.A. and Muller, M.J. (2002) "Understanding the Benefit and Costs of Communities of Practice", *Communications of the ACM*, **45**(4), pp. 69-73.
- Noack, D. (1994) "Blacksburg Electronic Village: Testbed for the Future", *Government Technology*, **7**(5), pp 20–76.
- Saint Paul (2002) "CyberVillage: Saint Paul", hosted at: <http://www.cybervillage.org/>

- Schonemann, P.H. and Wang M.M. (1972) "An Individual Difference Model for the Multidimensional Analysis of Preference Data", *Psychometrika*, **37**, pp. 275-309.
- Schuler, D. (1994) "Community Networks: Building a New Participatory Medium", *Communications of the ACM*, **37**(1), pp. 39-51.
- Spence, R. (1999) "A Framework for Navigation", *Int. Journal of Human Computer Studies*, **51**(5), pp. 919-945.
- Stathis, K., de Bruijn, O. and Macedo, S. (2002) "Living Memory: Agent-based Information Management for Connected Local Communities", *Interacting with Computers*, special issue on Intelligence & Interaction in Community Based Systems (pre-print).
- Sumi, Y., Etani, T., Fels, S., Simonet, N., Kobayashi, K, and Mase, K. (1998) "C-map: Building a Context-aware Mobile Assistant for Exhibition Tours", in *Community Computing and Support Systems*, LNCS-1519, Springer, pp. 137-154.
- van Mulken, S., André, E. and Müller, J. (1998) "The Persona Effect: How Substantial is It?", in *Proc. Int. Conf. on Human Computer Interaction (HCI-98)*, pp. 53-66.
- Vargas, M. (2000) "About Retail Loyalty and Customer Relationship Management", hosted at <http://retailindustry.about.com/industry/retailindustry/library/blloyalty.htm>.
- Witkowski, M., Artikis, A. and Pitt, J. (2000) "Trust and Cooperation in a Trading Society of Objective-Trust Based Agents", in *Workshop on Deception, Fraud and Trust in Agent Societies at Autonomous Agents 2000*, pp. 127-135.
- Witkowski, M., Arafa, Y. and de Bruijn, O. (2001) "Evaluating User Reaction to Character Agent Mediated Displays using Eye-tracking Technology", in *Proc. AISB'01 Symp. on Information Agents for Electronic Commerce*, pp. 79-87.
- Witkowski, M., Pitt, J., Fehin, P. and Arafa, Y. (2001) "Indicators to the Effects of Agent Technology on Consumer Loyalty", in *Stanford-Smith, B. and Chiozza, E. (eds.) E-work and E-Commerce: Novel Solutions and Practices for a Global Networked Economy*, IOS Press, pp. 1165-1171.
- Womak, S. (2002) "Love Blossoms at Book Clubs", *Daily Telegraph*, 29 April 2002, p.1.
- Zacharia, G. (1999) "Trust Management Through Reputation Mechanisms", in *Workshop on Deception, Fraud and Trust in Agent Societies at Autonomous Agents-99*, pp. 163-167.