2008

User Context Analysis from Spatial Interface Interactions

E Mac Aoidh
University College Dublin, eoin.macaoidh@ucd.ie

M Bertolotto
University College Dublin, michela.bertolotto@ucd.ie

D Wilson
University College Dublin, davils@uncc.edu

USER CONTEXT ANALYSIS FROM SPATIAL INTERFACE INTERACTIONS

Mac Aoidh, Eoin, School of Computer Science and Informatics, University College Dublin, Belfield, Co. Dublin, Ireland, eoin.macaoidh@ucd.ie

Bertolotto, Michela, School of Computer Science and Informatics, University College Dublin, Belfield, Co. Dublin, Ireland, michela.bertolotto@ucd.ie

Wilson, David C, Department of Software and Information Systems, University of North Carolina at Charlotte, 9201 University City Boulevard, Charlotte, NC 28223-0001, USA, davils@uncc.edu

Abstract

Our work addresses the problem of information overload in the spatial domain. Information overload can confuse and overwhelm users by providing too much detail, often complicating simple interaction tasks. Knowing users' interests allows the system to alleviate information overload, and therefore facilitate interaction, by tailoring the information displayed at the interface to individual user preferences. Spatial data (e.g. topographic maps) are rich in content and typically require explicit actions such as zooming and panning to view their extent. This paper presents preliminary results which show that user's spatial interests can viably be ascertained by monitoring mouse movements and clicks as spatial implicit interest indicators. An interest modelling algorithm analyses this interaction information to produce user interest models, reflecting each user's latest interests as they change over time. These models can be used for the personalisation of user datasets and interfaces, balancing the display content with the relevance of the available information in the system to the user and his context.

1   INTRODUCTION

Information overload occurs when there is too much data available. For example, non-spatial Web search engines return millions of documents for common search terms such as 'house' or 'car'. It would be impossible for any one user to view all of these documents. If the documents were un-ranked, the user would have too much information to sift through. As a result, search engines usually display the document titles ranked in order of relevancy as a filtering method. In the spatial domain, the same problem arises, viewing a map with an application such as Google Earth (2007) with the available "dining and lodging" facilities displayed in a region such as New York provides too much information for one user. Figure 1 shows a map representation with icons showing all dining and lodging facilities in a small area of New York, and right, the same area, personalised to show dining and lodging icons suitable for the user's tastes, dramatically reducing information overload. Without employing a means of ranking the data and prioritising the display for the most relevant information, users must sift through a large amount of data in the hope of finding the information they require. As with many domains, GIS (Geographic Information Systems) is subject to information overload.
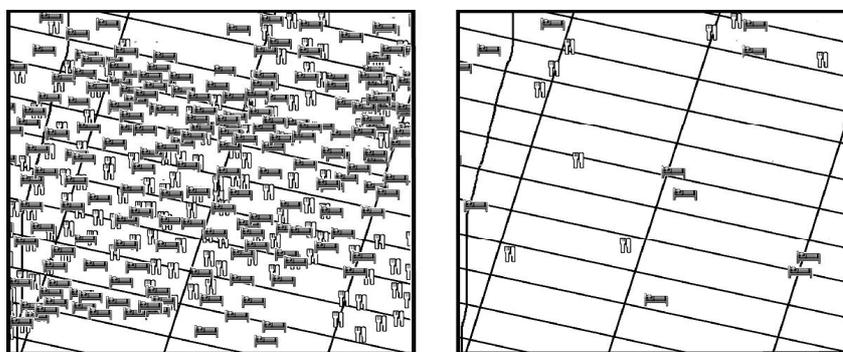

Figure 1.   Information overload display and personalised information display, spatial simulations

Increasingly, spatial metrics are being incorporated into everyday use. For example, straightforward distance and compass directions are offered in local searches by applications such as Google Maps (2007). A search for "Restaurants near Central Park" will return a series of results with the distance and direction from Central Park included, such as "Uncle Joe's, 0.3mi, E". As spatial information becomes more pervasive, both its availability and quantity increase dramatically, making information overload more prevalent. Our work tackles this problem in the spatial domain. We have developed a means of implicitly ascertaining users' spatial interests by analysing their map browsing behaviour. Knowing users' interests will allow us to filter the information returned to them, giving users personalised content preference in the interface to suit their spatial information requirements. This reduces the amount of time and effort required by the user to find what they are looking for, thus improving their interaction experience.

Our approach logs the user's actions as they interact with the system. In particular we focus on mouse movement[1] and click data. The movement and click information is processed by an interest determining algorithm. The algorithm computes a ranked list of every object in the map, based on its relevance to the user for a given session. For example, hotels might be more relevant than swimming pools in a given session according to the user's actions. By maintaining an average of the lists from each session for a particular user, we produce an up-to-date interest model that reflects his changing interests over time. This interest model can be used to produce personalised content for the user. By

---

[1] Mouse movement data refers to the time each mouse movement occurs at, the duration of each movement, the location (Latitude, Longitude) of the mouse over the map, and the direction of the movement

personalising content, it is possible to eliminate extraneous detail and recommend new, relevant detail to the user, reducing information overload and enhancing the user's contact with relevant information.

Mouse movements are regarded as relatively weak indicators of interest for *non-spatial* data, in comparison to book-marking a web page for instance, which is regarded as a strong indicator of interest (Lieberman 1997). The work outlined in this paper examines the validity of assumptions of interest made using mouse movements and mouse clicks with *spatial data*, as stronger operations such as book-marking a Webpage are not always relevant to spatial data. The implicit approach is adopted, as collecting explicit information can be time consuming and distracting for the user, often prompting users to ignore requests for such explicitly collected information. We also examine the strengthening in the accuracy of interest assumptions made by combining these two implicit methods. The results of a series of initial experiments are presented and analysed in this paper. The results show that users' interests can accurately be inferred based on their interactions.

The remainder of the paper is structured as follows: Section 2 discusses related research, section 3 introduces our method for calculating users' interests. Section 4 details the experiments carried out, validating our approach. The results are analysed in section 5. Section 6 provides a discussion on future considerations. The paper concludes with ideas for further development in section 7.

2   RELATED WORK

Contextual information about users' interests can be attained either implicitly or explicitly. To procure such information explicitly necessitates the interruption of users' work flows, as their contexts and preferences must be explicitly stated by filling out a form or answering a set of questions. This can be irritating and time consuming for users. As a result many users may skip this step and proceed directly with their task. In such an instance, no information can be obtained about the user's context. Implicit profiling avoids such a scenario by making implicit inferences based on users' interactions. It is a non-intrusive method of gaining information about the user. Claypool et al. (2001) consider a number of implicit interest indicator techniques in detail, including keyword extraction from documents and event logging (such as book-marking web pages, and the recording of mouse movements) in relation to non-spatial data. Our work examines the importance of users' mouse movements and clicks as implicit interest indicators in the context of spatial information.

A three point, or triangular relationship exists between mouse movements, thought processing and eye movements (Crazyegg 2007, Chen et al. 2001, Cox & Silva 2006). All three elements of the triangular relationship have been shown to be interrelated. A significant body of research, conducted within the field of implicit interest indicators including Arroyo et al. (2006), Atterer et al. (2006), Hijikata (2004), Mueller & Lockerd (2001) and commercial endeavours such as Crazyegg (2007), give evidence of a correlation between one's mouse movements and thought processing. It is reasonable to assume a correlation between eye movements and thought processes as demonstrated by experiments conducted with eye tracking technology such as (Cox & Silva 2006) and (Pan et al. 2004). Research projects including work by Cox & Silva (2006) and Chen et al. (2001) have completed the triangle by showing that there is relationship between a user's eye movements and mouse movements. All of the aforementioned research has been conducted with non-spatial data. To the best of our knowledge no such research has been conducted pertaining to implicit interest indicators with data of a spatial nature.

Systems such as the Curious Browser (Claypool et al. 2001) determine users' interests based on their mouse input while browsing standard web pages. The system records the amount of time spent moving the mouse, and also counts clicks. MouseTrack (Arroyo et al. 2006) and Cheese (Mueller & Lockerd 2001) (both monitoring web page interaction), consider the location of the mouse and mouse clicks, and the duration spent in each location to produce a heat map visualisations of user interaction. We consider the elements monitored by all three of these systems to be of equal importance. Our system incorporates a combination of the indicators used by MouseTrack, Cheese and the Curious Browser, and makes adaptations for considering spatial data.

According to results documented by Arroyo et al. (2006), Atterer et al. (2006), Chen et al. (2001), Cox & Silva (2006), Mueller & Lockerd (2001) and Pan et al.(2004) when considering the links forming a triangle between eye movements, thought processing and mouse movements, the link between a user's thoughts and his mouse movements is difficult to validate without expensive eye tracking hardware which is complicated to use and out of reach to many. We wish to show that a strong link can be forged without the use of eye tracking, by thoroughly examining mouse movement data, furthermore we wish to make our techniques readily available for use on any workstation. We seek to strengthen the link between thought processing and mouse movements by considering additional browsing interactions such as mouse clicks and map navigational actions.

Our research involves data with a spatial dimension. Viewing spatial data such as a map is different to viewing non-spatial data (e.g. a webpage). Tasks involving maps generally require the basic ability to zoom in and out, pan, and re-centre the map. The space used to represent the objects in a map is the key to the interpretation of spatial data, in contrast to web pages and other documents containing text and graphics, where white space is generally used for layout purposes to improve readability in non-spatial circumstances. White space, such as the empty horizontal space between paragraphs on this page for instance contains no 'meaning' in itself. It separates each paragraph for aesthetic reasons. In a map however, the space between two points represents an inherent object in itself.

Techniques used by Arroyo et al. (2006), Atterer et al. (2006), Hijikata (2004), Mueller & Lockerd (2001) focus around mouse movements with non-spatial data. Some of these measures, such as mouse position relative to hyperlinks, as used by Cheese (Mueller & Lockerd 2001) and MouseTrack (Arroyo et al. 2006) need to be adjusted accordingly for spatial data, while others, such as page visits, hits and views as used by USAPROXY (Atterer et al. 2006) and TF-IDF key word extraction from web pages by Hijikata (2004) are not applicable to spatial data. New techniques must be developed in their place. By implicitly taking spatial navigational patterns into account in addition to considering users' mouse movements we strengthen the link between a user's thoughts and mouse movements.

Spatial systems are lagging behind other systems when it comes to inferring implicit interests. To date, many spatial applications such as CRUMPET (Zipf 2002) rely on explicit user profiling. More recent research in the domain (e.g. Weakliam et al. 2005) sees the introduction of implicit techniques by monitoring user's map browsing behaviour including pan and zoom actions and actions such as adding or removing map content signifying interest and disinterest respectively. These applications provide an insight into the applicability and benefits of user profiling (Fischer 2001). By knowing the user's likes and dislikes, map content can be tailored to suit the user's needs. Experiments conducted with the CoMPASS system (Weakliam et al. 2005) revealed that information overload was reduced, and user's tasks made easier through the use of dataset personalisation.

3   APPROACH

The application developed in the course of our research is based on a case study in the cultural domain, specifically we deal with ancient burial tombs in Italy belonging to excavations in the area of Tarquinia. The data covers an area of approximately 8km^2 and roughly 80 tombs are represented as objects of interest on the map. The information is displayed in an interface consisting of two browsers, a map browser and a HTML browser presented side by side (a full system description is available in Mac Aoidh et al. (2006)). Standard pan, zoom and search functionality is provided for the map interface. Objects in the map are identified by a mouse-over function, basic information such as object name and type are revealed by mouse-overs. Clicking on a particular map object shows its associated, detailed information in the HTML browser. A screen shot of the interface is shown in figure 2.

As users complete their tasks, their interactions with the system interface are logged, specifically we capture mouse movement data. Upon session termination, the logged information is sent to an Oracle 9i Spatial database. We have developed an interest determining algorithm which acts on the logged information to produce a ranked list of a particular user's interests for a given session. Initially, we

wish to establish a basic hypothesis that users' interactions disclose significant information about their preferences. We begin initially with straight-forward implicit indicators, thus mouse position and dwell time in each location in relation to the underlying map objects are logged. The contents of the map browser at any point during a session is referred to as a frame. Each time the user completes a navigational operation such as a pan or zoom, which alters the frame, a new frame is logged.
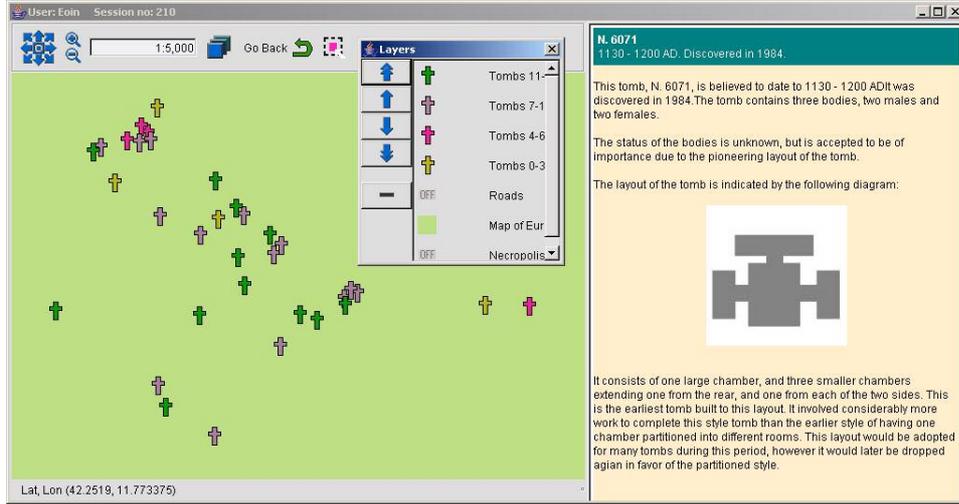

Figure 2.   User interface screen shot as displayed during experiments

Our basic assumption is that a longer dwell time indicates greater significance, and that the proximity of the mouse to an object gives an indication as to the importance of the object. In addition, the scale of the map is also of importance. We deem the few objects present in the frame when the map is zoomed in tightly to be of greater importance than the greater number of objects in the frame when the map is zoomed out further. With these basic assumptions in mind, we calculate a user's interest in a particular object as follows.

$$(\sum_{0}^{n} ObjectScore = \frac{Dpt(0)}{dist(Obj,Pt(0))} + \frac{Dpt(1)}{dist(Obj,Pt(1))} ... + \frac{Dpt(n)}{dist(Obj,Pt(n))}) * F$$

To calculate the score of an object of interest, *Obj*, in a specific frame of a specific session, the distance between each *Obj* and each mouse dwelling point, *Pt*, within the frame is calculated as *dist(Obj,Pt)*. The dwelling time of the mouse at *Pt*, *Dpt*, is divided by *dist(Obj,Pt)*. Each of these fractions are added together for a particular object to give it a score based on all of the mouse movements within the frame. The score for *Obj* is then weighted by the scale of the frame, *F*.

By dividing the dwell time of the mouse at *Pt(i)* by *dist(Obj,Pt(i))*, a measure of *Pt(i)*'s importance in the context of the session is obtained. Using this method, dwelling points furthest from *Obj* will receive the lowest scores. Additionally, a longer dwell time will give a greater value for the numerator, *Dpt(i)*, allocating higher scores to points where the mouse rested for a longer duration.

Most objects appear in more than one frame, thus their score is computed for each frame within which they appear, and summed together to give a final score over the entire session. These scores are used to rank the map objects according to the user's level of interest. It is from this list that the user interest model is computed. It is adjusted after each session in order to keep up with the user's interests as they change over time.

In practice, our algorithm also makes two additional simple assumptions. Firstly, suppose a user leaves his computer alone, and goes for a cup of coffee, it would be incorrect of our algorithm to assume the mouse is intentionally, and meaningfully focused in the position in which the user left it, thus we

disregard dwell times of longer than λ seconds (λ was set to 15 for these experiments). We have found series of slow, small movements to be much more indicative of users' interests than considering extended periods of dwelling in a single position. Secondly, while we use proximity as an interest measure, we also do our best to distinguish the semantic object type. For instance, consider a hotel, a restaurant and a hospital all located near each other on a map. If the user dwells near the hotel, the algorithm, taken at face value, would assume the user to be equally interested in all three features. However, the maps used in our system are vector maps, which allows the map to be composed of various transparent layers of features which can be turned on and off by the user. Thus the user could turn off the restaurant and hospital layers on the map. With both of these layers invisible to the user at the time the mouse dwelling near the hotel took place; the algorithm could conclude that the user is only interested in the hotel near where the mouse rested. For the sake of simplicity we test the Boolean value of a feature, i.e. restaurants and hospitals are either visible or invisible to the user. Weakliam et al. (2005) carried out a series of experiments using similar vector-based maps which show that it is possible to discern which semantic types of feature the user is most interested in. As this research has yielded successful results, we focus on determining which specific feature of a given semantic type the user is most interested in, with a view to integrating both methods in the future.

4   EXPERIMENT DESCRIPTION

As part of a system evaluation we conducted some preliminary experiments. Their function is to assess the validity of our approach and to ascertain the correct procedures and comparisons that need to be followed in order to produce more significant results in later experiments. The data for the experiments was acquired from human subject volunteers, each of whom completed a number of tasks designed to encourage them to interact with our system. A total of 20 tasks were defined in order to provide a scenario-based usability testing environment. The spatial data used in the case study is real data, however the volunteers were not 'real' users of the system, thus we designed the tasks to stimulate the volunteers into using the system in the most natural way possible. Each task required them to identify certain information about objects on the map (available in the information browser by clicking on the object in the map), and to write the name(s) of the object(s) they had selected to answer their task on an answer sheet. For example, one of the tasks asked users to "Identify four objects on the map with the following features in common..." Another required them to "give the name and location of three objects on the map which you would like to visit."

A total of 12 volunteers were used for the experiments. They were carefully selected to include a variety of males (8) and females (4), of varying familiarity with spatial applications. The volunteers ranged in age from 24 to 33. Data for 74 sessions was obtained and used as the experiment data. Each task was repeated by at least three different users to allow for comparison between different users completing the same task. All sessions were carried out on the same machine, thus each user had an identical environment within which to experience the system.

Our initial evaluation is based on the data obtained from these 74 sessions. As the volunteers conducted the tasks, every aspect of their interaction was logged. Upon completion of a session (each session representing a task), the interest determining algorithm computes an object score for each object in the map based on how the user interacted with it. An object's score is indicative of the user's level of interest in the object during his task. A user interest model is produced by ranking the objects according to their scores.

In assessing the validity of our approach, we examine the experiment results to show that analysing mouse movements is an effective means of profiling users as they interact with spatial data, and to show that this method can be used to strengthen spatial data recommendations. To this effect we examine the accuracy of inferred interests when interacting with spatial data, which will subsequently be used to create a user interest model. In our system, interests are inferred based on: 1) mouse click data, 2) mouse movement data, and 3) mouse movement combined with mouse click data. In addition to this analysis, we also identify two distinct categories of user, based on their mouse usage.
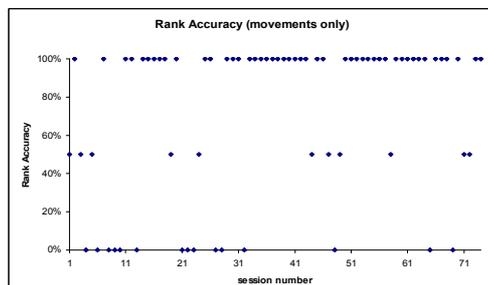
4.1 Metrics

In order to judge the performance of our system's algorithm, we examine its accuracy in three areas: Rank Accuracy (RA), Relative Preference (RP) and Absolute Preference (AP). Rank Accuracy is defined as the accuracy of the rank assigned to the objects in a user's answer. For instance, if a user specifies three objects on the answer sheet, we would expect the algorithm to rank the same three objects given in the user's answer as #1, #2 and #3. If only two of the user's objects are returned in the top three, then the RA is 66.66%. RA is not concerned with the order of the objects, merely the percentage of them present in the top ranking objects.

Relative Preference, RP, is a metric to judge a user's preference for one object in the ranked list over his preference for the following object, based on the object score, as computed by the algorithm. It is a measure of the degree of change in preference. For example if object X has a score of 60 and the next highest ranked object, object Y, has a score of 40, RP is computed as (20/60)*100. The degree of change is 33.3%. In other words we can say that the user had a Relative Preference of 33.3% for object X over object Y. The closer the object scores are together, the smaller the RP, thus the less confident we can be that the user had a preference for one object over another.

Absolute Preference, AP, is a means of determining the level of interaction associated with any object in a ranked list. It constitutes a mapping of all object scores in the list to a scale of 0 to 100. The object at 0 is ranked lowest, and received the least amount of attention, while the object at 100 is ranked highest, and received the most attention from the user. While RP tells us the degree of preference for one object over another, AP gives an insight into why one object was considered to be preferable to another. It gives an indication as to the quantity of interaction (clicks and/or movements) that was focused around a particular object in the context of all the interactions for all the objects in a session. We use these three metrics to examine the results of our algorithm with some initial test data. The metrics are applied to results of the algorithm with mouse movement data alone, and are then applied to results obtained by mouse movement combined with mouse click data.

5 RESULTS

Initially we examine Rank Accuracy, an aggregate indication of the effectiveness of our algorithm. Figure 3A shows the RA across all 74 recorded tasks when mouse movements alone are considered. The graph shows that in the majority of cases, the objects identified on the user's answer sheet were 100% the same as the objects identified by the algorithm. There were however, 15 cases (i.e. 20.3%) where the Rank Accuracy was at 0%.

Figure 3A: Rank Accuracy for movements only

Figure 3B: Rank Accuracy for movements and clicks

Figure 3B gives the rank accuracy of the algorithm for the same tasks while incorporating mouse click data. The number of cases where the algorithm was incorrect was reduced to 2 (i.e. 2.7%). The RA of tasks which previously held a RA of 50% was largely accentuated to 100% due to the reinforcement of assumptions made by clicks on objects of interest.

Tables 1A and 1B give a break down of results for all of the tasks conducted in our experiment. Table 1A pertains to interests calculated using mouse movements only, while table 1B considers both mouse movements and clicks. Each table initially gives results for all 74 tasks, and subsequently gives greater detail. 26 of the tasks did not specify objects to interact with in the task description. The objects interacted with were at the users' discretion. The remaining 48 tasks specified objects in the task description. The user was specifically instructed to find a named object. E.g. "Find Tomb X, dating to 10 B.C. and compare its contents to Tomb Y's contents, also dating to the same period". The 26 tasks with unspecified objects could be considered to hold a greater weight for evaluation purposes, as users had a choice of objects to interact with. In contrast, in the remaining 48 tasks, interactions were slightly 'constrained' or 'forced' as users were told what objects to interact with.

	RA	RP	AP
All 74 sessions	72.30%	48.97%	55.89
Unspecified objects (26 sessions)	73.08%	33.87%	36.55
Specified objects (48 sessions)	71.88%	57.15%	66.36

Table 1A: Average results for mouse movements with no clicks.

	RA	RP	AP
All 74 sessions	94.26%	85.60%	57.65
Unspecified objects (26 sessions)	87.50%	70.34%	24.42
Specified objects (48 sessions)	97.92%	93.87%	75.61

Table 1B: average results for mouse movements and clicks combined

The first column of each table shows the average RA percentage, i.e., on average how correctly each user's interests were ranked (the RA for all 74 tasks is the data depicted in figures 3A and B). The second column gives the average RP percentage. The average given is the RP of the last of the objects listed on the user's answer sheet over the next highest ranked object. This illustrates the division between the objects included in a user's answer from the remaining objects in the map. The third column gives the AP of the same object. The AP is an indication of the amount of interaction that focused around the object in the context of all interactions in the session.

On average, over all 74 tasks, the system was able to correctly infer the users' interests in 72.30% of cases based on mouse movements alone. Results of a Paired T-test reveal that 72.30% correctly inferred interests is significant for $p<0.00$ with 73 degrees of freedom. By including information from mouse clicks, the accuracy of determining users' interests rose to 94.26%. The results shown for the 26 tasks where objects were unspecified allowed users to interact in a freer manner, they are more prudent, and closer to reality. For these tasks, users' interests were inferred correctly in 73.08% of cases based on movements alone, however by combining our two implicit indicators an average of 87.5% was achieved. By including clicks in this case, the average Relative Preference was accentuated from just 33.9% to 70.3%, giving a clear distinction between the objects of importance to the user, and the next highest ranked objects. Absolute Preference is unaffected, remaining at similar levels in both figures due to the lack of distinction between clicks and movements. Both are considered equally important in calculating Absolute Preference.

Our evaluation results are encouraging. They show that while interests can be inferred effectively for spatial information using mouse movement data, they might not be sufficiently accurate as a standalone interest indicator. By combining information gained from other implicit interest indicators, such as mouse clicks, user profiling techniques for spatial data can be strengthened significantly.

5.1 Mouse Usage Patterns

While conducting the experiments, we carefully observed users' behaviour as they interacted with the system. We noticed two extremes in terms of mouse usage behavioural patterns. Some users make very little use of their mouse, reading the map almost exclusively by sight. They only move the mouse when absolutely necessary. We named exhibitors of this behaviour "Lazy Mouse Users" (LMU), by contrast at the other extreme we identified "Fast & frequent Mouse Users" (FMU). The latter tend to make exhaustive use of the mouse, clicking on everything, and constantly moving the mouse.

We carried out a post-experiment visualisation of movement patterns for users exhibiting these traits. The visualisation interface provided by the system, shown below, allows for the creation of a transparent layer showing a user's logged interactions overlaid on the map. Figures 4A and B show mouse dwelling positions represented by circles. The smallest circles correspond to dwelling points of 30ms. The largest circles represent points greater than 4sec. Tomb objects are represented by crosses.
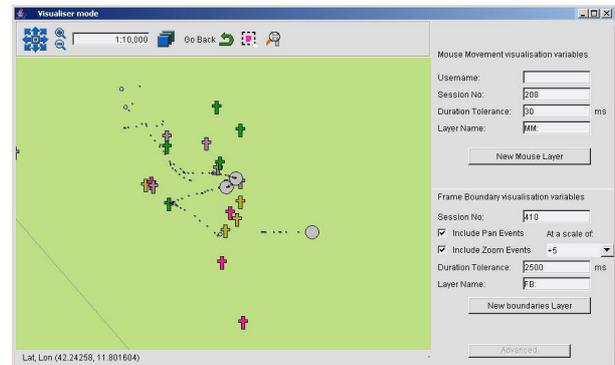

Figure 4A: FMU visualised   *Figure 4B: LMU visualised:*
Circle sizes correspond to duration of mouse dwell in the location at the centre of the circle.

The task visualisations showed that LMUs rested their mouse in few locations, but for significant time periods (between 500ms and 3 seconds) and moved slowly when they did move, while FMUs rarely rested their mouse. Our system captures a movement on average every 15ms, Movements by FMUs would be captured up to 5cm apart every 15ms, while LMU's movements would be captured closer together, less than 2cm apart.

The movement categories we identified bear similarities to those identified and analysed in detail by Cox & Silva (2006) during mouse movement and eye tracking studies using non-spatial data. Their study concerned item selection in file menu systems. They identified "Mouse On Side", where the user hovers the mouse to the side while scanning the screen with their eyes until the target is located, and "Mouse Hovering Target" where the user rests on the target while scanning with their eyes. These two groups bear similarities to our LMU group. The final group identified by Cox and Silva, "Mouse With Eyes", where the user's mouse followed their eye movements, bears similarities to our FMU category.

It could be argued that it is more difficult to establish a LMU's interests using our interest determining method, as LMUs make fewer movements, and therefore disclose less data. On the contrary, our experiments showed that, due to fewer movements, there is also less noise. In fact, we found it quite reasonable to calculate user's interests (for FMUs and LMUs) while ignoring movements with a dwell time of less than 200ms. This eliminates most of the quick movements collected for FMUs, and puts both categories of user on level pegging. Decreasing the threshold to 30ms means including most of the movements captured for FMUs. This results in a substantially larger number of calculations being carried out; however, the increase in accuracy of the algorithm, with the current approach is negligible.

Not all users fit into these two categories, as they are defined as extreme cases. Some users exhibit these characteristics more so than others. Our study has shown that it is usually subjects with less spatial data usage experience who show signs of FMU characteristics, while more accomplished users of spatial data tend to fit into the LMU category. The identification of these mouse movement categories could be useful in further strengthening our user profiles by determining user context. They could be used as a kind of implicit interest indicator in conjunction with case based reasoning. For instance, users exhibiting FMU characteristics could be compared to one another, if they show an interest in the same semantic type of objects it would be reasonable to assume they are contextually similar. By associating a user with a particular movement category, and considering the types of objects they are interested in, it is possible to establish a clearer contextual picture of the user.

There is a fuzzy discrepancy between users' contextual interest (the objects they are actually interested in) and users' behavioural intent (how their behaviour at the interface mirrors their intent in exploring the objects of interest to them). It is a significant step to pick out with certainty what the user's contextual interests are. By considering movement categories such as FMU and LMU, and performing comparisons with other users exhibiting similar behavioural characteristics, a clearer picture of the user's contextual interests could be constructed. With this information it is possible to construct a more concrete user profile, in turn allowing for more accurate personalisation of the system for an individual, leading to a reduction in the user's exposure to information overload.

6   DISCUSSION

In addition to establishing context based on the user's mouse movements, it is also possible to establish context based on a user's click stream pattern. Click patterns contain much less data than movement patterns, however they are reasonably concrete indicators of interest and as such, are a favoured means of interest indication by companies such as Digimaker (2007) providing Website traffic analysis tools, and dealing exclusively with non-spatial data. In a Web environment, there are inherent links connecting relevant information. Thus analysis of clicks can disclose the information relevant to the user. In an exclusively spatial environment however (unlike the environment used in our case study), there are no links to click on, thus clicks become redundant. Without any clickable information in a map the user has no links to click on, thus click stream information becomes comparatively useless and mouse movements become the better interest indicator. There may still be a small value in clicks, as users may still subconsciously click on map objects of interest out of habit, or to make sure there is no additional information obtainable from an object by clicking on it.

Many interfaces for systems whose primary concern is not necessarily spatial also contain spatial information, similarly, as is the case with our case study, interfaces for spatial information systems can also contain areas used for displaying non-spatial information (as shown in figure 2). Given the mix of spatial and non-spatial information, both clicks and movements can be important indicators of users' interests. Due to the exploitation of click stream data by Web traffic analysts, we initially focused our work on the lesser utilised area of mouse movements. We simply counted the number of clicks on each object as an indicator of interest, without exploring the click data any further. It could be expanded however to take in the amount of time elapsed between each click, and also the sequence of clicks. For instance a short elapsed time between clicks means only a short time was spent looking at the information on the object clicked on, while a longer time lapse usually means a longer time was spent looking at the information. The number of total clicks, and duration between clicks could also reinforce the categorisation of users according to FMU and LMU.

Click sequences could expose user intentions in relation to an object clicked on. For instance, a user clicking on Object *A*, followed by *B*, followed by *A* again would likely indicate a comparison being made between Object *A* and Object *B*. Further analysis of commonalities shared by Objects *A* and *B*, the user's category of mouse user (FMU or LMU), and elements deemed relevant by the interest determiner algorithm based on mouse movements can all be combined to produce as clear a picture as possible of a user's interests and context to produce an accurate user profile. This profile is based entirely on implicit information derived from the user's explicit actions. When the system is subjected to personalisation for an individual's user profile, information overload for that user can be reduced.

6.1   Limitations

Our system was initially developed for a specific case study, thus the variety of data types, and data distribution and spread were restricted initially to the data in the case study. As we are expanding the system's appeal, we are developing a means to deal with a larger variety of data types and varying spreads of data. The quantity of data available to the system is not really an issue, as algorithm processing is carried out when the user has finished interacting with the system. While larger

quantities of data will require more processing, this will go unnoticed by the user. Spatial data is composed of points (points of interest), lines (such as roads and rivers) and polygons (such as lakes and parks). The case study was only concerned with the user's level of interest in point objects, i.e. burial tombs. Variations in the algorithm would be required for line and polygon objects.

Our work is now being expanded to also determine users' levels of interest in line and polygon objects based on their mouse movements. The previous method of calculating distance between mouse dwelling point and point of interest is being experimented with to find the best alternative. Distances between the dwell point and the end points of the line, and the dwell point and the nearest point in the line are being tested in the case of line objects. Techniques for polygons include calculating the distance between the dwell point and the corners of the polygon's minimum bounding rectangle, and the distance between the dwell point and the polygons centroid.

The distribution of the map data is certainly a factor in the accuracy of interests determined by our approach. Data used in the case study was relatively evenly distributed over a small area. This even spread of data was an ideal testing ground for our approach, however the distribution of the data is not likely to be as ideal in all situations. Taking a map of a rural town as an example dataset, it would have a hinterland, usually comprised of large spaces with the road network converging on a central point, and the density of data increasing towards the centre of the town. Alternatively the data could be clustered around a number of points interspersed with sparse areas. For example, a small country is comprised of many towns. The towns account for clusters of data, with a relatively sparse amount of data between towns. We intend to develop a pre-processing algorithm to examine the spread and distribution of datasets provided, while also determining the best approach to adopt for inferring users' interests based on their spatial interactions for the particular dataset in question. This will also give a clearer indication of the conditions in which our methods are most applicable, and at what point they become stretched to their limits.

7   CONCLUSIONS AND FUTURE WORK

In this paper we have shown how analysis of a user's spatial interactions can be used to infer their spatial interests. Mouse movements have proven to be a strong indicator of interest with the dataset in our case study. The inclusion of subsequent interest indicators serves to improve the accuracy of the system even further. The proficiency of this method of interest indication remains to be tested for varying spatial distributions and other spatial data types. We have also discussed the possible limitations of our approach.

While it is evident from our results that a user's interests can be inferred in this manner in order to build a user interest model, we have not yet progressed to employing the user model to improve users' experiences with the system. It could be used for instance, to highlight features relevant to the user, or eliminate irrelevant features according to the user's individual interest model. This constitutes a personalisation of map content according to the user's context, reducing information overload and drawing the user's attention to features considered relevant to their profile. Content recommendations could be made to the user based on their context. Van Setten et al. (2004) discuss the recommendation of content specifically in relation to mobile devices. Interface tools could also be personalised to suit the user's interaction style, easing access and usability of the most used tools according to the user's individual profile.

Future work includes extending our interest determining algorithm for mouse movements to incorporate line and polygon objects. Modifications to our method will be required to this effect, indeed there may be more than one valid method for determining interests when considering lines and polygons. In such a case we will also need to develop a means of examining and classifying the dataset according to the most suitable interest indicator.

In addition to exploring the viability of our methods with various data types, we wish to revisit the analysis of click stream logs. In our initial experiments we took a simplistic view of mouse click logs,

considering only the number of clicks on a particular object. A further examination of patterns produced within this data would be worth exploring. These patterns include patterns which might indicate comparison between two objects, and clicking on an object in error. By taking such detailed information into account, we would expect the accuracy of interests inferred for users by our algorithm to improve considerably, in turn allowing for improved accuracy of information overload reduction and content recommendation.

References

Arroyo, E., Selker, T., and Wei. W. (2006) Usability Tool for Analysis of Web Designs Using Mouse Tracks. *In Proceedings of CHI'06 Extended Abstracts*, pages 484 - 489, Quebec, Canada, ACM.

Atterer, R., Wnuk, M., and Schmidt, A. (2006) Knowing the Users Every Move - User Activity Tracking for Website Usability Evaluation and Implicit Interaction In *Proceedings of the 15th international conference on World Wide Web*, pages 203- 212, Edinburgh, Scotland.

Chen, M., Anderson, J., and Sohn, M. (2001) What Can a Mouse Cursor Tell Us More? Correlation of Eye/Mouse Movements on Web Browsing. In *CHI '01*, pp. 281 - 282, Washington, USA, ACM.

Claypool, M. Le, P., Waseda, M., and Brown, D. Implicit Interest Indicators. (2001) In *Proceedings of IUI'01*, pages 33 - 40, Santa Fe, New Mexico, USA, January 14-17. ACM.

Cox, A. and Silva, M. The Role of Mouse Movements in Interactive Search. (2006) In *Proceedings of the 28th Annual CogSci Conference*, pages 1156 - 1162, Vancouver, Canada, July 26-29.

CrazyEgg: http://www.crazyegg.com (2007)

Digimaker: http://www.digimaker.com (2007)

Fischer, G. (2001) User Modeling in Human–Computer Interaction, User Modeling and User-Adapted Interaction, v.11 n.1-2, pp.65-86

Google Earth: http://www.earth.google.com (2007)

Google Maps: http://www.maps.google.com (2007)

Hijikata, Y. (2004) Implicit User Profiling for on Demand Relevance Feedback. In *Proceedings IUI'04*, pages 198 - 205, Madeira, Portugal. ACM.

Lieberman, H. (1997) Autonomous Interface Agents. In Proceedings of Conference on Human Factors in Computing Systems, CHI-97, Atlanta, Georgia, March 1997, pages 67 - 74.

MacAoidh, E., Koinis, A., and Bertolotto, M. (2006) Improving Archaeological Heritage Information Access Through a Personalised GIS Interface. In *Proceedings W2GIS'06*, pp 135 - 145, Springer.

Mueller F. and Lockerd, (2001) A. Cheese: Tracking Mouse Movement Activity on Websites a Tool for User Modeling. In *CHI'01*, *Extended Abstracts* pp. 279-280, Seattle, Washington, USA. ACM.

Nivala, A. and Sarjakoski, L.T. (2003) Need for Context-Aware Topographic Maps in Mobile Devices, in Proceedings of the 9th Scandinavian Research Conference on Geographic Information Science ScanGIS'2003, pages 15 - 29, Espoo, Finland, June 4-6.

Pan, B., Hembrooke, H., Gay, G., Granka, L., Feusner, M., and Newman, J. (2004) The Determinants of Web Page Viewing Behavior: An Eye Tracking Study. In *2004 symposium on Eye tracking research and applications. ETRA*, pages 147 - 154, San Antonio, Texas, USA, 2004. ACM Press.

Van Setten, M., Pokraev, S. and Koolwaai, J. (2004) Context-Aware Recommendation in the Mobile Tourist Application Compass in Proceedings of Adaptive Hypermedia 2004, pages 235-244, Eindhoven, August 23-26, Springer Verlag.

Weakliam, J., Bertolotto, M. and Wilson, D. (2005) Implicit Interaction Profiling for Recommending Spatial Content. In *Proceedings of the 13th annual ACM international workshop on Geographic information systems*, pages 285 - 294, Bremen, Germany. ACM Press.

Zipf, A. (2002) User-adaptive Maps for Location Based Services (LBS) for Tourism. In *Proceedings of the 9th nternational Conference for Information and Communication Technologies in Tourism*, pages 329 - 338, Innsbruck, Austria.