Smart Survey Tool: A Multi Device Platform for Museum Visitor Tracking and Tracking Data Visualization

Author 1106[1]

1 1106 Institution

ABSTRACT

This paper describes the Smart Survey Tool, a novel multi-device application for museum visitor tracking and tracking data visualization. The application allows museum staff to capture detailed information describing how visitors move around an exhibition and interact with individual exhibits. They can then visualize the results of tracking either on a single mobile device or with multiple mobile devices connected to a large display. The platform uses orthogonal views of the exhibition space for tracking and visualization, with a 'chess-piece' icon to represent visitors during tracking, and curved semi-transparent lines with animated semi-circles to communicate the path and direction of visitor movement. Our visualization is novel in its use of an orthogonal projection for pedestrian tracking and animation to communicate the flow of visitors around the exhibition space, as well as allowing users to dynamically switch between views representing different groups of visitors. The design of our application was informed through an extensive requirements analysis study conducted with Nanjing Museum and evaluated by conducting expert interviews with museum managers who considered that the application allowed for more effective and efficient recording and analysis of visitor tracking data.

Keywords: Information Visualisation, Pedestrian Tracking.

1 INTRODUCTION

There is currently a massive effort within the academic and commercial research communities to develop indoor tracking and indoor GPS type technologies that can locate pedestrians inside a building and track their movement [1-5]. One of the primary motivations behind this research is to give building designers, and the people who manage public spaces, a better understanding of how people move around a building so they can design the building or modify the design (by placing attractions, moving shops, or renovating) to better control the movement of people in order to improve their experience and maximize commercial outcomes. However, this outcome depends not only on the effectiveness of technologies to collect tracking data, but also the availability of suitable analysis software that can help users to find important patterns in the data.

The people in charge of public spaces need to be able to know about the density and flow of different types of visitor, at different times and under different conditions, in order to be able to make informed decisions about building space design and management. The focus of this paper is to develop a mobile software platform that helps with the collection and proper analysis of visitor tracking data to support this requirement. We also look to learn from the design decisions taken during the development of our system and would hope that anyone working on a similar problem could learn from our findings as presented in this paper.

As a case study for the general problem of indoor pedestrian tracking and information visualization, we have been working with staff at Nanjing Museum develop new techniques to help us understand how visitors move around different museum exhibitions. Our initial requirements analysis with Nanjing Museum revealed that that tracking data could be extremely valuable toward understanding how visitors moved around museum exhibitions. It also revealed that the analysis of this data was expensive and limited with regard to what they could find out from larger scale data collected over a longer period of time.

The results of tracking-data analysis were found to be very useful, but the museum staff had the feeling that they could get a lot more from their data if they could ask more questions of the data and had access to a more powerful interface that was capable of answering those questions. This called for us to support them by developing a dedicated application to collect and visualize tracking data.

Prior to the development of our tracking and visualization application, the main procedure for tracking visitors involved either pen-and-paper manual tracking (see figure 1) or automatic tracking using RFID tagging technology. The manual pen-and-paper approach involved using a paper copy of the exhibition floorplan to record; the time a visitor arrives at an exhibit, when they leave the exhibit, and what they do at the exhibit (with activities such as taking a photo, reading text, concentrating etc.). This method was found to be quite flexible, and allowed trackers to capture a good range of different types of data. The problem with this technique is that writing down times to capture timing information gets to be tedious and different trackers tend to drift towards using different notation. The data also needs to be entered into a computer in order to generate any sort of meaningful summary of the data or perform anything more than the most basic analysis. This data-entry could be somewhat of an arduous task when data is collected for larger numbers of visitors.

The alternative to pen-and-paper tracking is to use RFID tagging technology to automatically record the time spent by each visitor at each exhibit. This allows staff to record the movement of more visitors (i.e. all the visitors who use an audio guide with an RFID tag reader fitted) but has the disadvantage of not being able to record information about visitor demographics or activity at different exhibits.

After tracking data was collected, museum staff would commission a report to summarize the data. This would include some statistics and images to describe different aspects of the data such as the density of visitors or the general movement of visitors around the exhibition space. Statistics would include mostly demographic data, such as the proportion of visitors from different economic groups, but also things like the most popular exhibit by total visits or time spent. The images included in the report would be heat-maps to show the density of visitors in different areas of the exhibition and directed graph view with arrows showing the path of each user around the exhibition (see figure 2). These reports were found to work well as a summary of the data but were less effective for more in-depth analysis of larger scale data.

The primary limitation of relying on a printed report for the analysis of tracking data was described to us as being a lack of flexibility to explore less obvious patterns contained in the data. For example, a printed report might have a heat-map that shows the most popular areas of the exhibition for all visitors. It may even show separate heat-maps showing the most popular exhibits for different genders or different hours of the week. It would, however, be impractical to have separate heat-maps for every possible type of visitor (by gender, age group, type of visitor) or ever different time (different hours, days of the week, weeks of the year etc.). The report would end up being overloaded with too many images and many of the images would be redundant as the patterns of visitor activity for groups of visitors would be similar to the pattern for all visitors. It would also be difficult to detect patterns without comparing different groups. Here, it makes more sense to allow museum staff to interact with the data in a way that allows then to quickly switch between different groups of visitors to observe how patterns of movement are different.

Another significant limitation of the graphics used in the printed report was that the graph view, showing the movement of visitors, suffered from the problem of clutter and crossing lines [6] and was ineffective at showing the direction of visitor movement for larger groups of visitors. Anything more than around ten or twenty visitors (depending on the layout of the exhibition) would render the graph view pretty much useless as it would just appear as a grey mass of crossing lines (see Figure 2 for an example of this).

The focus of the project described in this paper was to develop an interactive visual interface to overcome the limitations of the existing method for visitor tracking and tracking data visualization. This led us to develop a new mobile software platform including a tracking application to allow museum staff to track a sample of visitors to a museum exhibition, and a visualization application to support more effective analysis of the results.

2 Related Work

There is a long history of tracking visitors to museum and galleries exhibitions [7-9] with most systematic methods developed over the past 20 years. The main challenges of visitor tracking are managing the cost associated with gathering data, and performing effective analysis of the data once data is gathered.

There are a number of different methods that can be can be used to track visitors to museums or other public spaces. The most common method is for people to be tracked manually with their time at each exhibit recorded using pen and paper [9]. Tracking can also be automated using pedestrian tracking systems that employ video camera footage, RFID tags [3, 5, 10, 11], or applications that monitor WIFI signal strength on the visitor's smartphone [1, 2]. While automated tracking has the advantage of potentially higher throughput with less investment in manpower, manual tracking is generally thought of as being more accurate and can record more detailed information such as the type of activity that each visitor is engaged in at each exhibit.

Once visitor tracking data is collected it can be visualized in order to understand patterns of visitor movement. The most common visualizations for this use heat-map type displays that show the density of visitors in different areas and graph views that show the most common paths for visitors around the exhibits [12, 13]. Additional insights are often gleaned from the

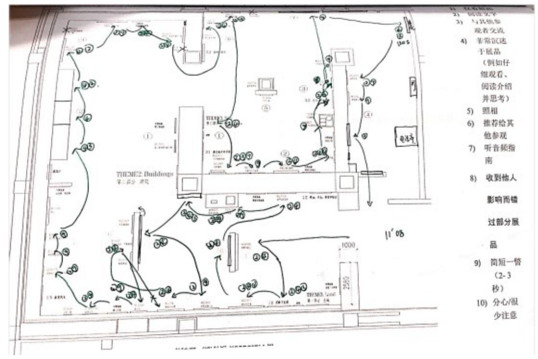

Figure 1: An annotated print-out of the museum floor-plan used for pen-and-paper tracking.

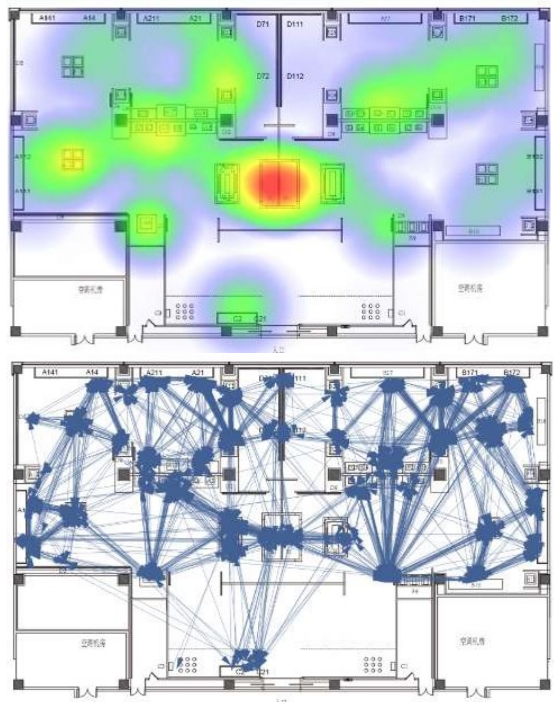

Figure 2: Examples of the type of visual representation used in printed reports for tracking data. These are heat-maps to show the density of visitors in different areas of the exhibition, and directed graph views with arrows showing the path of visitors around the exhibition.

data when it is filtered according to sample attributes such as age or gender [13] or attributes of the data such as the door by which visitors enter the exhibition or the first turn they take when they enter [14].

The most challenging aspect of visualizing visitor tracking data is showing the path and movement of visitors around the exhibition space [12, 13]. While existing graphical representations worked well for smaller scale data [12, 14, 15], when the data was scaled up these techniques are found to suffer from the problems of clutter and crossing lines [6]. Certainly, more powerful analysis techniques are needed to properly exploit the potential of large-scale tracking data [6, 16].

For the purposes of information visualization, this sort of data can be considered as a specialized type of large-scale directed graph data where the positions of nodes are fixed. Display

Table 1. Different types of data the users wanted to record.

Type of data	Description	Characterization
Visitor	Including data about the visitor, such as their gender and age, as well as general information about their visit to the exhibition such as the type of group they belonged to or how busy the exhibition was when they arrived.	We initially identified a total of nine classifications (gender, age group, etc.) and around one hundred individual classes that a person could belong to (male, female, child, teenager etc.).
Behaviour	What the visitor does at each exhibit, for example, taking a photo, reading text, or thinking about the exhibit.	We initially identified a total of nine actions. The visitor could be recorded with multiple actions at each exhibit. The timing of these actions is not important and the exhibition managers were happy to rely on observer to judge if an action was taking place.
Timing	The duration of the visitors stop at each exhibit and the order in which they visit exhibits. The exact path of visitors was not considered important.	The time when a visitor arrives at an exhibit and when they leave the exhibit.

techniques that can be applied to this sort of data are curved lines to distinguish between paths [17, 18], edge bundling to simplify the view by gather similar paths together [19-21], and different methods (such as arrow heads, tapering, curves, color, and animation) for showing the direction of edges [22]. The curved line techniques applied to parallel plots [17] can be thought of as particularly applicable to the problem of visualizing visitor paths as different the edges in parallel plots also form distinct paths and the curves should be capable of helping us follow visitor paths in a similar way to how they help us follow parallel-plot paths.

3 Requirements Analysis

Our initial requirements analysis for this project took place over approximately one month with different groups of urban planners, museum exhibition planners and museum staff based in the United Kingdom and China. The number of personnel involved in our requirements analysis study was relatively small, including just eleven key project stakeholders, but we felt it was important to involve people who had a genuine interest in the projected functionality of the software. These were the staff involved in the United Kingdom Arts and Humanities Research Council funded 'Romantic Scotland' exhibition to be held at Nanjing Museum, and included a core team of three urban planners, five permanent Chinese museum staff, and three of the United Kingdom staff responsible for the design and management of the exhibition. The project team also included visualization experts, mobile device developers and personnel familiar with a variety of indoor pedestrian tracking technologies. Over the course of four months we held a series of eight meetings with various members of the project team to arrive at our final application design.

This initial requirements analysis was supplemented with small-scale evaluations of early prototypes (involving twenty or so students at the XJTLU university museum) for tracking and visualization. This allowed us to iron out usability issues and refine the interface design before a final trial run at the Nanjing Museum, and the actual tracking and visualization of the

Table 2. Patterns the users wanted to be able to see in the data.

Type of pattern	Description	Rationale	Existing analysis and limitations
Density of visitors	Density of visitors in different areas of the exhibition.	To identify potential bottlenecks, where to space is too busy at times, or parts of the exhibition space that are underutilized.	Heat-map representation is already quite effective for this type of pattern (see figure 2).
Flow of visitors	how visitors generally move around the exhibition including their general around the exhibits and the direction of movement.	This could tell them why parts of the exhibition are not being visited enough, or how their strategy to guide visitors towards certain exhibits is working.	A line graph representation works well for smaller numbers of visitors but for larger numbers it becomes unusable due to the density of lines (see figure 2).
Statistics for exhibits	Total and Average number of visitors at an exhibit, time spent, number of photos taken, etc.	To tell us how different exhibits appeal to visitors in different ways.	Tables are used to present data for individual exhibits, but users wanted a more intuitive graphical representation (using the exhibition floorplan) that also allowed them to explore the data to view data for different demographic groups.
Patterns for different demographics	Above types of pattern for different demographic groups including different genders, age groups, and visitor groups (tour groups, couples families etc.).	To see how different groups behaved (e.g. young people versus elderly people) in order to ensure that the exhibition caters to a wider demographic, or is able to target a particular demographic such as those who are more likely to become tourists in the location that the exhibition is promoting.	Bar charts and pie-charts are used for broad overviews (such as male visitors versus female visitors) of the data but users wanted to be able to explore the data to see more detailed data for smaller groups (such as the poplar exhibits for couples or young men).

Table 3. Different ways the users wanted to analyse the data.

Analysis	Description
Mobile analysis	The users wanted to be able to analyze the data in different locations such as meeting rooms or the exhibition space itself. Using a desktop PC would not be possible in most of these locations and a laptop PC would be also impractical.
Collaborative analysis	Often the users would want to work together to analyze the data so they could combine their knowledge of the exhibition and discuss findings.
Publication of results.	The users also wanted to be able to share their findings in printed form, so the visualization needed to be aesthetically appealing.

Romantic Scotland exhibition.

The first stage of requirements analysis was to gather requirements for visitor tracking. This included site trips to the museum and Skype meetings with the exhibition planners in the UK to determine what technology would be most appropriate for tracking the museum visitors. The main outcome of this stage of requirements analysis was the information that the personnel responsible for managing the exhibition required three types of data. These were data related to the visitor, the visitors behaviour in the exhibition and the timing of their movement through the exhibition. These data are described in table 1.

The requirement of the exhibition managers to collect behavioural data in addition to timing data made it necessary for us to use human observers rather than an automated technique such as RFID tagging [3, 5, 11] or Wi-Fi tracking [1, 2]. However, the range of behaviours that they wanted to record made pen-and-paper tracking impractical. This made it necessary to develop some way to make the process of manual tracking more manageable. As the layout of the exhibition made it necessary for trackers to actually follow visitors in order to observe their behaviour, the development of a mobile tracking application seemed the most plausible solution.

The second stage of requirements analysis focused on the potential users' objective for tracking data analysis. Here we held a number of meetings, initially using the reports commissioned by the museum to help guide our dialogue to determine how the data should be visualized and later using our own sketches of how the application should look. These meetings also allowed us to identify the main types of pattern that the users wanted to be able to find in the data (summarised in table 2), and the different ways in which they wanted to be able to analyses the data (table 3).

A requirement that was stressed as particularly important for our users was to be able to view and compare patterns for different demographic groups. As well as providing clues as to why particular parts of the exhibition may be visited more or less for different groups, data related to the behaviour of different demographics could also allow the analysts to test theories related to how different groups are attracted to different types of exhibit (looking at different genders [23, 24], types of visitor group [24], and age groups [7, 25] etc.).

The requirement for some analysis to be done in locations where it would not be practical to use a laptop or desktop PC was also significant in that it made it necessary for us to design our visualization application to run on mobile devices. These were chosen as six eight inch tablets, to be used primarily for tracking, and a couple of larger ten inch tablets to be used mostly for visualization and analysis. It would also be necessary to connect mobile devices to synchronize the data collected on different devices, and these devices could use the same connection for collaboration on a large display.

4 The Smart Survey Tool

The process of requirements analysis described above gave rise to our initial platform design. This was for a mobile application that allows us to track visitors (running on multiple devices so different tracking staff can work at the same time) and a mobile application that connects to a large display screen for visualization. We used an Android socket connection running over a shared Wi-Fi connection to synchronize data across devices and coordinate mobile devices with the large display for collaborative visualization. This connection could be initiated by using one device to scan a QR code encoding the IP address of the second device. Figure 3 shows the device configuration and outlines the basic functionality of our platform.

4.1 Visitor Tracking

Figure 4 shows a screenshot of the interface of our tracking application. The interface uses a military projection of the exhibition space using an image of the exhibition floor plan with the visitor currently being tracked represented as an icon resembling a pawn in chess. The military projection (see figure 8) is an oblique projection where the angles of the *x*- and *z*-axes are at 45°, meaning that the angle between the *x*-axis and the *z*-axis is 90°. This allows vertical elements such as walls and doors to have a kind of 3D representation without having to skew the xz-plane. The icon representing visitors is designed to fit the overall aesthetic of the interface without specifically representing any particular type of visitor. The interface is designed to run on a tablet to be held in landscape orientation and operated using the thumbs.

To operate the interface and track visitors moving around the exhibition space the tracking staff click on doors and exhibits to record visitor movement and uses the thumbwheel dials at the bottom corners of the interface to record additional details about the visit. While a human tracker is certainly less efficient that automated technologies such as RFID for tracking motion, the advantage of employing a person to track visitors is that they can also record the activity of visitors at different exhibits. This indeed was one of the requirements of our project stakeholders who were interested not only in visitor movement but also how people behaved in front to exhibits by doing things like taking photos or reading exhibit text.

When a visitor enters the exhibition space, tracking begins by the tracker clicking on the door that the visitor entered by. After clicking on the door an icon representing the visitor appears as if they are entering through the door. The tracker can then record how the visitor moves around the exhibition by clicking on each exhibit the person stops at. This is done in real time, and as the tracker clicks on exhibits the icon representing the visitor is animated to follow an approximation the path of the actual visitor. If the tracker makes a mistake they can hit an undo button to reverse their previous action. Clicking on a space away from an exhibit moves the visitor icon away from any current

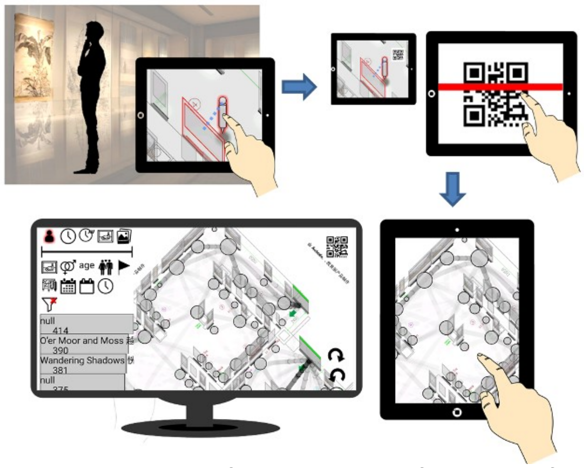

Figure 3: The device configuration and basic functionality of our platform. Tracking staff use a mobile device to track visitors moving around the exhibition. Devices used to track visitors and display the data are connected on the local Wi-Fi network by scanning a QR code. Data is automatically synchronized between connected devices and the data can be visualized on a mobile device or a mobile device connected to a large display.

Figure 4: The tracking application interface. The user presses on different exhibits to indicate the movement of a visitor and uses the thumbwheel dials to specify visitor characteristics or actions.

Figure 5: Using the tracking app to record the movement of a museum visitor around the Romantic Scotland exhibition in Nanjing Museum.

exhibit but not on to another exhibit. This whole process records the start and end time of the museum visitor's attention being given to individual exhibits.

Thumbwheel dials at the bottom corners of the interface can be used to indicate details such as the type of visitor (approximate age group, gender, type of group etc.) and details about the visitor's activity at each exhibit. So, if the visitor performs actions such as taking a photo, reading the exhibit text, or conversing with a friend, all this can be recorded. Figure 5 shows the tracking application running on a mobile device.

4.2 Tracking Data Visualization

The interface for our visualization application (see figure 6) uses different projected views of the exhibition space to show the density or flow of visitors for the tracking data. The different types of projections used are an isometric projection, a military projection, and top-down plan view (see figure 8). The interface aggregates data for all visitors or a selected group of visitors. The interface can be set to either show the density of visitors in different areas using a heat-map, the stats for different exhibits using a bubble-map, or the main path of visitors using a graph view with semi-transparent curved lines and animated arrowheads. Figure 6 shows the interface with the bubble map selected, figure 8 shows the path view for different projections and figure 9 shows the heat-map showing areas of the exhibition where visitors spent the most time.

A bar-chart menu with icons at the left hand side of the screen (seen in figure 6) allows the user to select from a total of thirteen different attributes to display (number of visits, average visit time, number of photos taken etc.) and nine different data classifications (age group, gender, day-of-week, hour etc.) with a total of around one hundred different classes. The different classes for the selected classification are shown in a horizontal bar chart together with the values for the selected attribute. This allows the user to look at things like the most photographed exhibits, or the number of males and females or different types of group visiting the exhibition. The bar chart can also be used to select different groups in the main map view to look at popular exhibits or patterns of movement within these groups. The selection in figure 6 shows the number of visitors by group type (individual, friends, couple, family, school group or guided tour). Figure 7 shows a part of the heatmap for different genders in the sixteen to thirty year age bracket. Here males are found to spend more time than their female counterparts sitting on the bench (on average 1 minute 12 seconds versus 1 minute 6) but a lot less time in the interactive booth (33 seconds versus 59 second for females). This suggested to the museum staff that younger female visitors might be more outgoing and that more could be done to engage young male visitors in interactive exhibits.

In addition to being able to select different types of visitor the user can also interact directly with the map to select different exhibits or paths and use the bar chart and icon to view statistics for that part of the data. Pinch, splay and drag actions, that were already familiar to our sample-users, can also be used to navigate around the map and the buttons at the right hand side of the screen allow the user to rotate the view 45 degrees in either direction or connect to another device to synchronize data or control a large display.

The graph view of visualization interface uses curved lines to show the path of visitors between exhibits. Animated arrow heads communicate the direction of visitor movement, and color value is used to communicate the number of users moving along any given path. The curve used is a simple quadratic curve starting mid-way between exhibits with a control point at the exhibit. Transparency is used so that only the most popular paths are fully opaque and less popular paths are more transparent. This has the effect of reducing the problems of clutter and crossing lines, as less popular paths have less impact and less weight to occlude the more travelled paths that the analyst is likely to be interested in.

The approach of making less common paths less visible made sense to our users as they only wanted to be able to see the main patterns of movement for each selected group. Outlying patterns of visitor movement would only be considered as significant if they involved a significant number of visitors or a smaller numbers of users in a smaller portion of the data. In either of these cases the pattern would be evident as the opacity of a path is proportional to the total visitors taking that path and inversely proportional to the total number of selected visitors. Smaller numbers of visitors who took an unusual route within a larger group were not considered to be significant for the purpose of this analysis.

Users can connect the application running on their own device to any other device running the application by scanning a QR code displayed on the top right of the screen. This connects the devices together using an Android socket connection running over a shared Wi-Fi connection. Typically, a mobile device will be used to connect to an Android box powering a large display so that the smaller mobile device can be used to control the large display (see figure 10). When multiple users connect to the same large display, they can take turns controlling the large display and collaborate to analyze the data.

4.3 Design Rationale and Methodology

The design of our platform, and our platform components, were developed and refined through a series of meetings and small-scale prototype evaluations with potential users and project stakeholders. This proceeded over a period of roughly four months. The initial design was inspired by our initial requirements analysis (described in section 1) where staff at Nanjing were able tell us not just about what they needed, but also what they liked or didn't like about the existing method of analysis and what they wanted from the new system. This was followed by the development and evaluation of several small scale paper prototypes and consultation sessions with urban planning experts at XJTLU. Finally, we refined our design by testing specific functional components attached to early functional prototypes.

4.3.1 Bar chart widget

The first thing we gathered from the requirements analysis was the importance of ranked lists for the visualization and analysis of the tracking data. The museum staff wanted to be able to view things like the most popular exhibits by number of visitors, average visit time, and things like the number of photos taken at different exhibits and by different types of visitor. They also wanted to be able to look at things like what different types of visitor they had, how much time different age groups spent in the exhibition and what where the most popular exhibits for different age groups. Written reports were good at presenting this sort of data in lists, but they were limited in the number of lists that could be included. Displaying each of the thirteen different types of attribute for each of the nine classifications would give us a total of one hundred and seventeen different lists to display. This would take a lot of paper, even before we consider that we might want to display a filtered version of any of these lists with any combination of the hundred or so distinct classes (showing, for example, the most photographed exhibits for a particular gender or age group). This made it important to have some easy way for the users to be able to select the

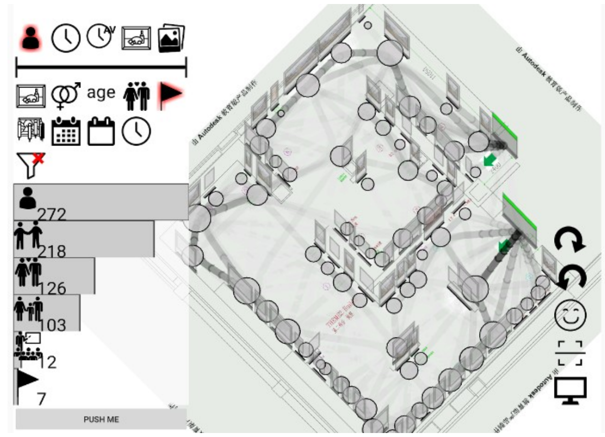

Figure 6: The main information visualization display with the graph view selected to show how visitors move around the exhibition.

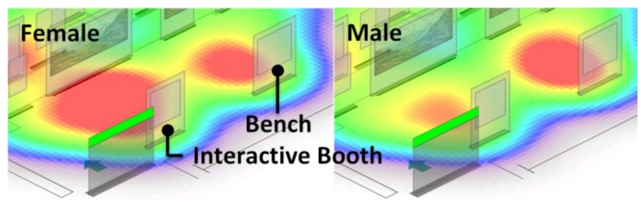

Figure 7: A heatmap with the 16-30 age group and *average visit duration* selected shows us that males (left) tend to spend more time at the bench, while females (right) spend more time in the interactive booth.

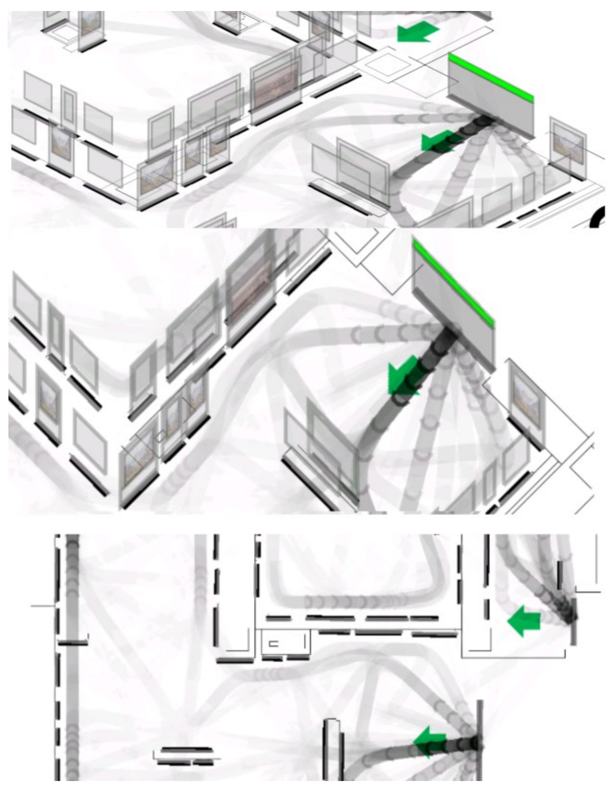

Figure 8: The different types of projections used in the Smart Survey Tool visualization interface. These are an isometric projection, a military projection, and a top-down plan view.

parameters for different lists and select different settings for the lists.

Our interactive bar chart widget for looking at lists of tracking stats can be seen in figure 6. The motivation behind this design is to use icons to make a more efficient use of screen space but otherwise make the interface as clear as possible by having three distinct lists for views, attributes, classifications, and class filters. Attributes include number of visits, average visit time and behavioural data like number of photos taken. Classifications include timing data (day of the week, hour or day etc.) as well as information about the visitor (age-group, gender etc.) and the conditions (quiet, bust etc.). Selected attributes and classifications are highlighted using a red glow, applied filters are shown using icons, and a help mode is available to provide tool tips describing the function of the icons.

4.3.2 Heat-map and bubbles-map

One thing that was reported to work well for our sample users using the existing method of analysis was the heat-map view of the data. This was found to be effective at communicating the density of different visitors in different areas of the exhibition space. The only major limitation of this type of representation was that it could not show the density of different types of visitors or for different times, and it could be difficult to see which individual exhibits were popular if two exhibits were close together.

To allow the user to be able to look at different heat-maps for different types of visitor of different times, it was easy to link the heat-map view in our visualization interface to the bar chart widget (described in section 4.3.1). Here the user could use the bar chart to select a class of visitor or a time (a week of the year, weekday or hour of the day) and show a heat-map for just that data. The user could also look at a different heat map for different data attributes such as average visit time or photos taken. This ability to view different aspects of the data in the heat-map view made it possible to replicate all the existing heat-map images included in the commissioned report and also generate any of the other types of heat-map that the museum staff told us they might want to see.

To make it easier to distinguish between statistics for exhibits close to each other, we decided to supplement the heat-map with a bubble-map type display showing statistic values for individual exhibits. If this setting is selected a circle is shown next to each exhibit with the size of the circle proportionate to some statistic such as the average visit time or the total number of visitors.

4.3.3 Curved line graph view

The museum staff felt that the graph views included in their printed reports, showing the direction and flow of visitors around the exhibition space, did not work well (see figure 1). These suffered from the problem of clutter and crossing lines [6] and were ineffective at showing the direction of visitor movement for larger groups of visitors. In order to resolve this issue we developed a new visualisation method using semi-transparent curved lines which could be combined to show the main patterns in each view allowing the user to view more detail by interacting with the interface to select different filters. Without filtering, this representation gives a simple overview of the data. If the user interacts by selecting exhibits in the map or classes of visitor in the bar-chart widget, they can drill down into the data to reveal more detail. This type of representation worked well in a small scale evaluation of an earlier prototype and was refined to be used in the latest prototype as shown in figures 6 and 7.

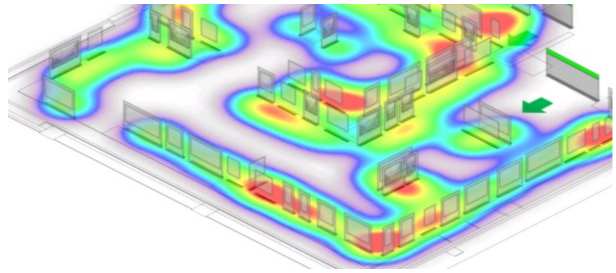

Figure 9: The heat-map view with an isometric projection.

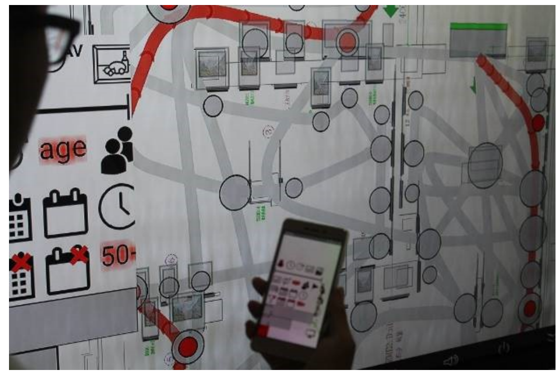

Figure 10: Using the visualization interface with a mobile device controlling a large display.

We also tried various encodings to communicate the direction of movement (including arrow-heads, tapering, curves and color-coding, from Holten et al. [22]). The most effective representation appeared to be animated semi-circle *waves* moving along the path. This had the advantages of not requiring us to change the shape of the path (which could result in a path that is confusing) and of allowing us to superimpose two paths in different directions without losing clarity. Animation and movement also has the advantage of being a pre attentive visual variable that grabs the users' attention [26]. It can also be thought of as being quite a literal sort of representation as the graphic moves in the direction of movement of the actual visitors. The potential disadvantage of animation is that it can attract too much of the user's attention and distract them from other important parts of the interface. In our case we did not find this to be a problem. Our small scale evaluation with our users told us that this representation gave them a better sense of how visitors would move around the exhibition and this did not detract in a negative way from any other aspects of the interface. The semi-circles showing the direction of motion can be seen in figure 8.

4.3.4 Military, top-down and isometric projections

The military, top-down and isometric projections use in our tracking and visualization interfaces where sketched out and discussed at the very start of the project with a small group of urban planners based in the university. These are based on a style used for architectural drawings and it was felt that the 3D effect of the isometric and military projections could give the user a better feel for the exhibition space.

In general 3D representations tend not to work well for information visualization due to the problem of occlusion where placing objects in a 3D space can mean that objects of interest get hidden behind other things [27]. Three dimensional visual cues are however proven to help users remember objects and can help with spatial orientation [28, 29]. As the task associated with

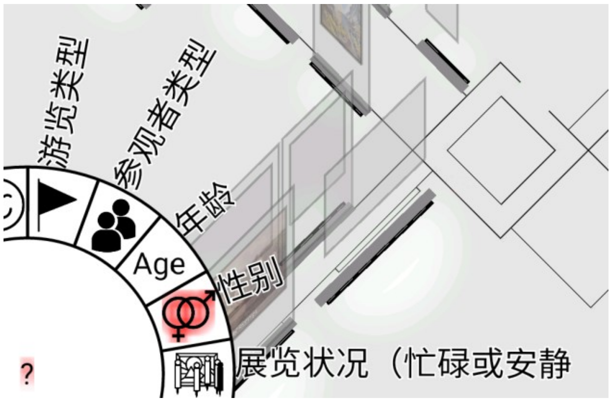

Figure 11: Thumbwheel widget with help text (in Chinese) activated by pressing the question mark icon

this visualization was associated with a real physical space, we felt that the use of three dimensional visual cues to represent exhibits and doors was warranted. To reduce the problem of occlusion we made exhibits semi-transparent and allowed the users to rotate the display in increments of 45 degrees.

Some initial testing of an early prototype of the tracking app indicated that the military projection would be preferred for tracking. Out of sixteen trackers, twelve expressed a preference for the military projection. The average rating for the military and isometric projections was significantly higher that the rating for the top-down view which scored badly.

The users told us they preferred the military projection because exhibits where more easily recognizable with a more natural side-on perspective. They also considered it easier to use than the isometric projection because the spacing for exhibits is more regular on both x and z axes. This made it more convenient for the user to select different exhibits quickly and accurately regardless of whether the exhibits were spaced on the x or z axes. This finding led us to make the military projection the default projection for tracking.

For the visualization interface users expressed a preference for isometric and top-down views. Out of the nine users involved in the evaluation seven preferred the isometric view, but they also scored the top-down view well and explained that they would like to be able to use both views. This was irrespective of whether analysis was done on a mobile device or mobile devices connected to the large display.

The top-down view was felt to give a better overview of the data without the problem of occlusion and the isometric view gave a more natural perspective. There was less need to select items quickly during visualization, so this advantage of the military perspective to better support this was not important. As no single view was significantly preferred over the others, we gave the user the option of being able to switch between military, top-down and isometric projections while using the visualization interface.

4.3.5 Tracking application interface design

For the tracking application interface design our principle concerns were ease of use and efficiency. The thumbwheel dials of the tracking interface are designed so that trackers can specify this information as efficiently as possible with the minimum required effort while holding the tablet in both hands and maintaining a reasonable level of accuracy. The view of the interface also moves automatically to keep the tracked visitor in the centre of the screen so that the tracker does not need to adjust the view themselves. Both these features are aimed at reducing the effort needed from the tracker who may need to operate the app over extended periods of time.

In a prototype evaluation with ten trackers we tested the tracking interface. During these tests a thumbwheel menu with icons was compared with a more traditional list type menu. Our thumbwheel menu (see figure 11) could be rotated using the users thumb with icons being tapped to specify functions. Icons where either reused from established iconic representation (for example the representations of gender), designed using symbols associated with the concept (for example a triangular flag for a tourist group, where in China it is common for the tour guide to carry a such flag), adapted from a physical representation (one person beside a painting for a quiet exhibition space), or designed using text when a small enough text label can fit into the button space (for 'age' and different age ranges). The list type menu is simply stacked textual descriptions of each concept.

Comparing the thumbwheel and the text menu, the thumbwheel was found to be more efficient as it required less movement and the user did not need to adjust their grip on the tablet. The icons also took up less space than text in the interface and could be recognised and pressed quickly if the user understood their meaning. The users told us that they wanted to be able to operate the interface without reading and that they found the thumbwheels be more efficient than the text based menu. Unfortunately, some of the interface functions related to concepts that were not so easy to represent using an image that would be easily recognisable for the user. These included things like how busy the exhibition space was, or if the visitor was deep in thought. To help the user learn the icons for these actions we supplied an instruction manual describing the meaning of each icon and incorporated a help-text into the interface (see figure 11).

5 Expert Interviews and Evaluation

Our platform was used to track just over 750 individual visitors to Nanjing museum over a period of two months with each visitor viewing an average of seventeen individual exhibits each. The data successfully captured and visualized on mobile devices by individual museum staff and on mobile devices linked to a large display by multiple staff. Overall this was a success as the tracking and visualization proceeded smoothly with a small number of technical glitches that were easily resolved.

Figures 12, 13 and 14 show a selection of the insights into the data learned from our visualization app that the users considered they would not be able to find using any other methods they had already tried. Our users were able to find a variety of patterns like these and it was evident through further investigation that they would not have been able, or would have found it significantly more difficult, to find most of these patterns using the previous analysis methods.

Figure 12 shows how visitors tracked the left-hand-side wall or took a shortcut past an enclave to view the striking portrait of Sir Joshua Renolds in a kilt (or exhibition organizers considered that this type of image would be particularly attractive to Chinese tourists and it featured heavily in marketing material). This pattern is clear using the new interface, but less evident using the traditional technique where the volume of crossing lines and the abrupt change in direction of straight lines makes it difficult to tell where the visitors come from before taking the short-cut or that most of the visitors who move along the wall move onto the portrait. Figure 13 shows how visitors tend to be attracted to the Wandering Shadows painting when they enter the exhibition then move to the left-hand-side wall. This is very clear from the new visualisation with curved lines and transparency, but not clear from the traditional representation where the volume of data makes it difficult to see any trend of this sort.

Both figures 12 and 13 illustrate behaviour which is contrary to the analysts' expectations, where visitor are known to have a tendency to follow the wall on the right hand side. While neither of these findings led the exhibition designers to change the layout of the exhibition (which would be very unusual for a temporary exhibition), they felt that it might change how they approach the layout of future exhibitions.

Figure 14 shows how patterns of behaviour change when the exhibition is more or less busy with visitors tending to spend more time at interactive exhibitions near the exit when the museum is quieter. This pattern was revealed by alternating between different items in the bar-chart menu. This pattern would not be seen using the traditional method as a limited number of images need to be pre-rendered without giving the user the opportunity to make selections to filter and explore the data.

Overall, the final evaluation was considered a success and expert interviews with museum managers revealed that they were satisfied that the software was more efficient than the existing technique employed and effective at supporting all their requirements for analysis. The main benefit of the tool was found to be that it allowed the users to interactively explore the data and allowed them quickly analyse the data without having to rely on data analysis experts. The users described the software as 'flexible', 'interactive', and said that it provided 'good access to the data'. The users were also impressed with the use of animation and semi-transparent lines to show visitor flow as this was felt to be clearer and give a better sense of how visitors moved around the exhibition. They also felt that the different projections each had advantages for different aspects of the data. The top down projection was felt to give a better overview of the data, while the isometric view gave a better feel for the exhibition space when zoomed-in. The limitation of the software was that museum staff could not change the mapping of the exhibition inside the software, and this is something they would like to add in future versions. Overall the user feedback was largely positive and as a result of their experience the museum have committed to using the software again for other temporary exhibitions.

We also plan to adapt the platform for use with the museum's permanent exhibitions. This will involve working to adapting the tracking software to work with RFID technology so that tracking can be automated. The system will be integrated with audio-guide equipment so we can incorporate information about the type of person being tracked. The disadvantage of using RFID tracking, rather than manual tracking, will be that it will not be able to record the users behaviour at individual exhibits unless we can use some other type of technology such as live-video object recognition and develop an additional system component to do this. This is being considered as a potential direction for further work.

6 Conclusion

We have developed a software platform to support museum visitor tracking and data visualization. The software allows museum visitors to have their behaviour and movement recorded

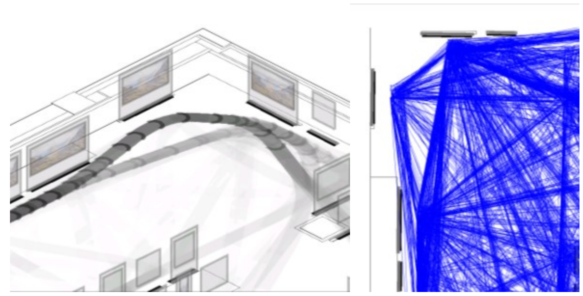

Figure 12: Visitors take a shortcut past an enclave to look at the portrait of Sir Joshua Renolds.

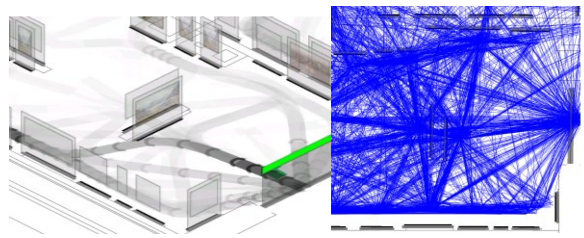

Figure 13: Visitors tend to be attracted to the Wandering Shadows painting when they enter the exhibition, then move to the left-hand-side wall.

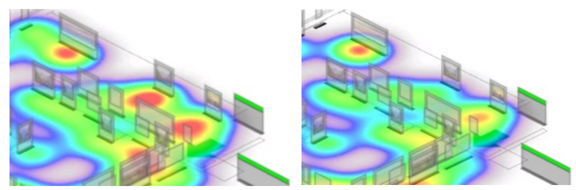

Figure 14: The heat-map views show the average visit times when the exhibition is quiet (top) and when it is busy (bottom). Visitors tend to spend less time at the more interactive exhibits (at the top left of the view) when the exhibition is busy.

using a specialized tracking app running on a mobile device with the data visualized on a mobile device or mobile devices connected to a large screen display. The visualization uses a custom bar chart widget, and a novel combination of semi-transparent curved lines and animated semi-circle waves to show the flow of visitors around the exhibition. Our visualization is novel in its use of an orthogonal projection for pedestrian tracking and animation to communicate the flow of visitors around a physical space. Another significant advantage of our technique of established techniques was that it allowed users to dynamically switch between views representing different groups of visitors.

Our study began by recording and gathering together user requirements (described in section 3) which informed the design of our software (section 4). The software was tested to refine the design and evaluated to gauge a positive response from sample users. It was also used to track just over 750 visitors to Nanjing museum over a period of two months and visualize the results. This led to positive results and our users who were able to find a significant number of patterns in the data that could not be found using the previous technique. A series of expert interviews revealed that our users felt the flexibility of the interface to allow them to be able to explore the data, rather than rely on data analysis experts, to be the main advantage.

The design and successful evaluation of our platform demonstrates the power of a more interactive information visualization display to reveal patterns in large scale data, and also the potential of mobile devices for this sort of visualization task [30]. Our interface worked well by relying on the user to interact with the display to reveal patterns in the data, rather than us trying to show too much of the data in any single frame of the visualization. This was particularly important for our graph view, where less popular visitor paths would stay in the background until the user selected a smaller portion of the data containing these paths so they became significant.

It is also worth noting that the limitations of mobile devices (with less accurate interaction, less screen space, and reduced processing power) did not cause any significant issue for our users in the final evaluation. This was partially down to advancements in mobile device capabilities but also down to design decisions such as using a thumbwheel menu and military projection for less accurate interaction in the tracking app, or the design of the bar-chart widget for more predictable reversible single touch actions. It is also very likely that allowing the mobile devices to control a large display, when the app was used in a meeting setting with multiple users, also helped us to overcome some of the limitations of mobile devices for information visualization.

Reflecting on the complete design process the most important lessons for us, of potential interest to designers of similar applications, are as follows.

- Having an interactive interface that allows the user to explore the data by selecting different facets of the data (e.g. different genders, age groups etc.) and different views (e.g. heat-map or curves for paths) can be very powerful. This was seen as the main advantage of this technique over existing methods.
- Animation can be used to effectively to show the flow of visitors around an exhibition.
- Different projections, including 2.5D projections such as isometric projections, can have value in the presentation of visitor movement data and heatmaps.

Due on the success of this operation Nanjing Museum have committed to using the Smart Survey software for future tracking exercises and we will look to improve the software and adapt it for more challenging data (e.g. in multiple levels of a building or over longer periods of time) and data collected using different technologies such as RFID. More work could also be done looking at how to apply or adapt this sort of technique for other environments such as shopping centres and other public spaces where there may be a commercial advantage to tracking visitor movement but patterns of movement may be less predictable or more difficult to characterize.

ACKNOWLEDGMENTS

The authors wish to thank the students who helped with the evaluations and development of prototype applications. These are Ke Wang, Chao Xiang, Kun Wang, Wenjun Chen, Yuxuan Zhao, Zifan Wang, Yuexuan Li, Zhejian Zhang, Xiaoyu Wu, Weiyi Zhang and Yuhao Wu. We would also like to thank the staff at Nanjing museum and our collaborators on the UK AHRC project 'Producing/Consuming Romantic Scotland'.

REFERENCES

[1] Dyuldina, K., Snopko, S., Shakhgeldyan, K. and Kryukov, V. *Indoor navigation service based on Wi-Fi positioning*. IEEE, City, 2017.

[2] Hill, M. R. Stalking the Urban Pedestrian:" A Comparison of Questionnaire and Tracking Methodologies for Behavioral Mapping in Large-Scale Environments". *Environment and Behavior*, 16, 5 1984), 539.

[3] Choi, J. S. *Accurate and cost efficient object localization using passive UHF RFID*. The University of Texas at Arlington, 2011.

[4] Zhang, D., Xia, F., Yang, Z., Yao, L. and Zhao, W. *Localization technologies for indoor human tracking*. IEEE, City, 2010.

[5] Want, R. An introduction to RFID technology. *IEEE pervasive computing*, 5, 1 2006), 25-33.

[6] Yabushita, H. and Itoh, T. *Summarization and visualization of pedestrian tracking data*. IEEE, City, 2011.

[7] Robinson, E. S. The behavior of the museum visitor1928).

[8] Melton, A. W. Problems of installation in museums of art1935).

[9] Yalowitz, S. S. and Bronnenkant, K. Timing and tracking: Unlocking visitor behavior. *Visitor Studies*, 12, 1 2009), 47-64.

[10] Hiyama, A., Yamashita, J., Kuzuoka, H., Hirota, K. and Hirose, M. *Position tracking using infra-red signals for museum guiding system*. Springer, City, 2004.

[11] Montaser, A. and Moselhi, O. RFID indoor location identification for construction projects. *Automation in Construction*, 392014), 167-179.

[12] Lanir, J., Kuflik, T., Sheidin, J., Yavin, N., Leiderman, K. and Segal, M. Visualizing museum visitors' behavior: Where do they go and what do they do there? *Personal and Ubiquitous Computing*, 21, 2 2017), 313-326.

[13] Strohmaier, R., Sprung, G., Nischelwitzer, A. and Schadenbauer, S. Using Visitor Flow Visualization to Improve Visitor Experience in Museums and Exhibitions.

[14] Corini, A. *Evaluating Visitor Experience in the Citi Money Gallery at the British Museum*. Worcester Polytechnic Institute, 2013.

[15] Ma, J. *New tools for tracking and timing*. City, 2007.

[16] Alavi, H. S., Lalanne, D., Nembrini, J., Churchill, E., Kirk, D. and Moncur, W. *Future of human-building interaction*. ACM, City, 2016.

[17] Graham, M. and Kennedy, J. *Using curves to enhance parallel coordinate visualisations*. IEEE, City, 2003.

[18] Craig, P. and Roa-Seïler, N. A Vertical Timeline Visualization for the Exploratory Analysis of Dialogue Data. In *Proceedings of the Information Visualisation* (Montpellier, France, 10 - 13 July 2012, 2012), [insert City of Publication],[insert 2012 of Publication].

[19] Holten, D. and Van Wijk, J. J. *Force-Directed Edge Bundling for Graph Visualization*. Wiley Online Library, City, 2009.

[20] Gansner, E. R., Hu, Y., North, S. and Scheidegger, C. *Multilevel agglomerative edge bundling for visualizing large graphs*. IEEE, City, 2011.

[21] Bach, B., Riche, N. H., Hurter, C., Marriott, K. and Dwyer, T. Towards unambiguous edge bundling: Investigating confluent drawings for network visualization. *IEEE transactions on visualization and computer graphics*, 23, 1 2017), 541-550.

[22] Holten, D., Isenberg, P., Van Wijk, J. J. and Fekete, J.-D. *An extended evaluation of the readability of tapered, animated, and textured directed-edge representations in node-link graphs*. IEEE, City, 2011.

[23] Borun, M. and Chambers, M. Gender roles in science museum learning. *Visitor Studies Today*, 3, 3 1999), 11-14.

[24] Black, G. *The engaging museum: Developing museums for visitor involvement*. Routledge, 2012.

[25] Dean, D. *Museum exhibition: Theory and practice*. Routledge, 2002.

[26] Treisman, A. Preattentive processing in vision. *Comput. Vision Graph. Image Process.*, 31, 2 1985), 156-177.

[27] Yi, J. S., ah Kang, Y. and Stasko, J. Toward a deeper understanding of the role of interaction in information visualization. *IEEE transactions on visualization and computer graphics*, 13, 6 2007), 1224-1231.

[28] Robertson, G. G., Mackinlay, J. D. and Card, S. K. *Cone trees: animated 3D visualizations of hierarchical information*. ACM, City, 1991.

[29] Alper, B., Hollerer, T., Kuchera-Morin, J. and Forbes, A. Stereoscopic highlighting: 2d graph visualization on stereo displays. *IEEE Transactions on Visualization and Computer Graphics*, 17, 12 2011), 2325-2333.

[30] Craig, P. Interactive animated mobile information visualisation. In *Proceedings of the SIGGRAPH Asia 2015 Mobile Graphics and Interactive Applications* (Kobe, Japan, 2015). ACM