Spatial Analytic Interfaces: Spatial User Interfaces for In-Situ Visual Analytics

Barrett Ens, Pourang Irani

{bens, irani}@cs.umanitoba.ca

Department of Computer Science, University of Manitoba, Canada

Abstract

As wearable devices gain acceptance, we ask "What do user interfaces look like in a post-smartphone world?" and "Can these future interfaces support sophisticated interactions in a mobile context?" In stark contrast to the micro-interactions of current wearable interfaces lies visual analytics. A hallmark of such platforms is the ability to simultaneously view multiple linked visualizations of diverse datasets. We draw from visual analytic concepts to address the growing need of individuals to manage information on personal devices. We propose Spatial Analytic Interfaces to leverage the benefits of spatial interaction to enable everyday visual analytic tasks to be performed in-situ, at the most beneficial place and time. We explore the possibilities for such interfaces using head-worn display technology, to integrate multiple information views into the user's physical environment. We discuss current developments and propose research goals for the successful development of SUI for in-situ visual analytics.

Keywords: spatial user interface; visual analytics; augmented reality; head-worn displays

Introduction

Personal computing devices are becoming smaller yet more powerful, allowing greater user mobility, increased quantities of personal data, and enhanced control in manipulating these data to benefit our everyday lives. A catalyst in this shift in computer usage is increased access to sensors and interfaces, which are becoming integrated with what we normally wear. As we have already witnessed computers moving from entire rooms to desktops to pocket-sized devices, we are now experiencing a continuing shift to wearable form factors such as smart watches and digital eyewear. This new generation of interactive information displays has great potential to enrich our lives. Unlike with current mobile technology, information from these devices can be ingested from a glance at the wrist or even a slight eye movement. Such always-available information access allows *in-situ* computing: access to situationally appropriate data at an ideal time and place. By providing wearable technology with suitable information-seeking interfaces we can make computing a natural and 'invisible' part of our daily activities.

The complexity of mobile computing interfaces has so far been limited by the small available space for input and display. For example, some common tasks performed on mobile devices include consumption tasks such as reading or viewing videos, mobile communication tasks such as sending or receiving short messages, and organizational tasks such as keeping a list of contacts or setting reminders. As wearable device interfaces continue to shrink, current design solutions are trending further toward simplicity; new interface paradigms (e.g. Google Glass[1], Android Wear[2]) are designed to support *micro-interactions*, short bursts of activity that avoid impinging on one's day-to-day activities by minimizing task duration.

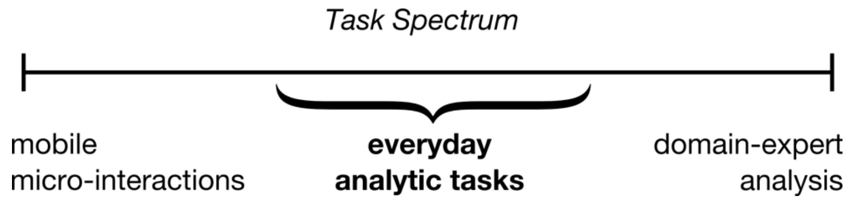

Figure 1. The nature of computer-assisted tasks varies widely from short-term, simple tasks carried out by mobile or wearable device users to intensive, analytic tasks carried out by teams of domain experts. We propose Spatial Analytic Interfaces for supporting everyday analytic tasks, which reside between the extremes on this spectrum.

In contrast to these current trends, we are interested in designing interfaces that support *analytic tasks* (Figure 1). Properties of such tasks include the requirement for concerted thought, the integration of information from multiple sources, and the application of human sensemaking abilities (Figure 2). Typical examples of everyday analytic tasks include balancing a cheque book, planning a vacation itinerary, or conducting a price search for the best available deal on a particular item. Such tasks are commonly carried out with the assistance of computers, yet are not necessarily well supported by today's mobile device interfaces.

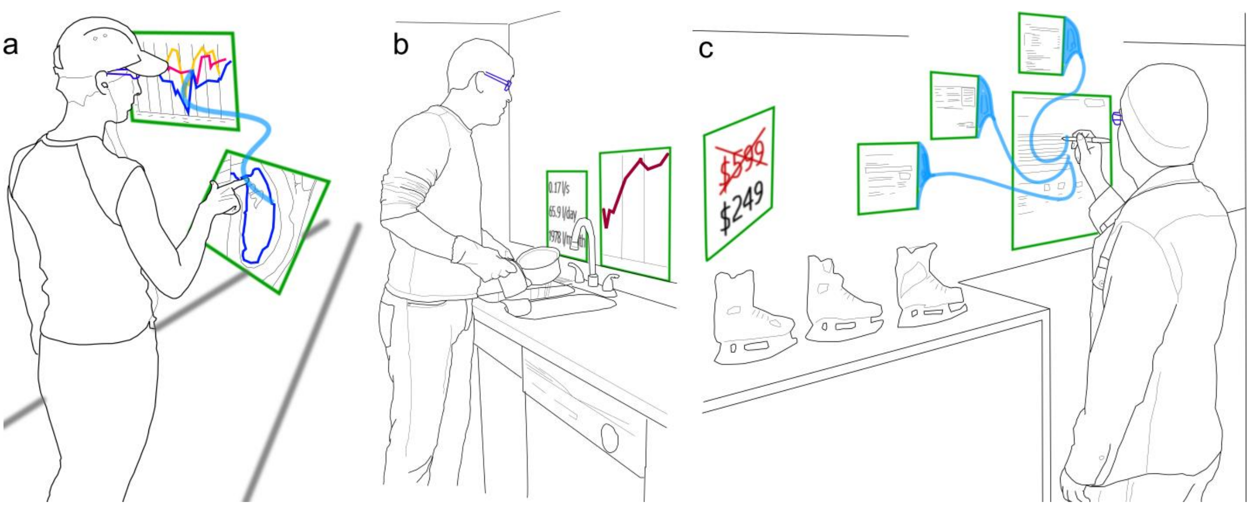

Figure 2. Three scenarios depicting beneficial uses of Spatial Analytic Interfaces: comparing heart rate and route records during exercise with visualizations that hover in surrounding space (a); monitoring home water consumption using virtual information panels that appear on the kitchen backsplash (b); and completing a quick budget before making a purchase, using virtual documents overlaid on a nearby surface and a spatially-tracked stylus (c).

To design interfaces that support analytic tasks, we can draw from the field of *visual analytics*. Visual analytics is devoted to developing tools that help users gain insights through deep exploration of multiple interlinked visualizations of diverse data sets. Although originally aimed at supporting domain experts with intensive analysis, for instance of biomedical data [19] or military intelligence reports [18], visual analytic methods have been recently adopted for analysis of an

[1] https://www.google.ca/glass/start/
[2] https://www.android.com/wear/

increasing wealth of everyday personal information [14, 17]. For example, sensors in people's homes track energy consumption and resource usage patterns; mobile computers such as smartphones and embedded automobile software continuously track their owners' everyday movements; and wearable accessories are popular for tracking personal health and fitness data. This trend of ubiquitous data collection presents a growing need for tools to comprehend and digest the patterns of importance and to provide actionable results [5].

The benefits to be realized from an increasing prevalence of mobile and wearable technology are then twofold: While these devices allow the routine collection of useful activity data, they also provide an opportunity to facilitate in-situ data analysis. Homeowners concerned with minimizing their energy consumption, for instance, are better able to make informed choices if appropriate information is available at the time when they are choosing how to consume resources or energy (i.e. turning up a thermostat). Similarly, if people are able to consult their banking history through a mobile app, they can make use of this information directly before making significant purchases. The mobile component is essential to in-situ computing, since waiting to view data at home on a personal computer results in the situational context becoming lost. However, viewing data on the small screen of a personal mobile device may be prohibitively cumbersome, and lacks the potential for gaining insight by coordinating multiple, coordinated views of the data [18].

One promising approach to provide mobile interfaces for in-situ use, with advanced features to support analysis and sensemaking is the application of *spatial user interfaces*. Spatial user interfaces leverage benefits such as spatial memory and proprioception to map information to physical space and have been shown to improve performance on some analytic tasks [7]. For instance, arranging multiple visualizations side-by-side can allow for faster and easier comparison than navigating between multiple components on a single abstract interface; the user can easily switch views using physical head or body motion and apply spatial memory to recall the location of important items, making for an efficient and intuitive experience.

We propose the concept of *Spatial Analytic Interfaces* (SAIs) as a solution for everyday data-monitoring and decision-making based on in-situ analysis. SAIs leverage the benefits of *spatial user interfaces* for completing in-situ, analytic tasks (Figure 2). The concept of SAIs is platform-agnostic, however we focus on *head-worn display* (HWD) technology as a particularly appropriate platform for meeting the requirements for supporting in-situ analytic taskwork. Such digital eyewear is recently available in lightweight form factors at an affordable cost for general consumers and the technology is rapidly advancing (see sidebar on page 7). HWDs are becoming equipped with depth cameras and inertial sensors that allow tracking of hand, fingertip and body motion (e.g Meta[3], Microsoft Hololens[4]). These features will facilitate intuitive spatial interaction, for instance the ability to switch between spatially situated displays by turning one's head [19]. With robust spatial tracking, these devices essentially provide unlimited 'display' space; multiple information visualizations can be integrated directly into the appropriate home, work or mobile environment. Furthermore, virtual displays rendered by these wearable systems can be situated where they are most convenient for a given context, for instance on the kitchen counter or

[3] https://www.getameta.com/
[4] https://www.microsoft.com/microsoft-hololens/

backsplash for monitoring home energy consumption, or in a hemispherical formation around the user's body in mobile situations such as shopping or jogging. This spatial paradigm can also support advanced techniques not possible with standard desktop displays; for example, visual links can span physical space to connect data across multiple displays or guide users to information that is not currently in their focus of attention [20].

The goal of this paper is to introduce the concept of Spatial Analytic Interfaces, discuss their benefits over current mobile interfaces, and to define the challenges in implementing them. We first describe several example scenarios in which this concept can be applied. Next, we propose a number of design requirements for novel wearable platforms to facilitate in-situ analytic tasks. We then describe the characteristics of HWDs and focus our discussion on how this particular platform can satisfy many of the requirements for in-situ analytics. Finally, we highlight several open research areas in which work is needed to make practicable implementations of Spatial Analytic Interfaces.

Scenarios

To demonstrate the breadth of potential of the opportunities for SAI, we give three scenarios where data visualizations presented on HWDs can be of potential value for in-situ analytic tasks. These scenarios include tracking personal health information during a morning run, monitoring home water consumption, and managing a quick overview of finances while shopping.

A Morning Run – First we visit Elie on her morning run (Figure 2a). Following along is a pair virtual display windows. The display to her left shows her step count, heart rate and estimated calories burned. The other, on her right, contains a map showing her current location and her predicted route, based on logs from previous runs. Neither of these windows occludes her forward view and she periodically consults them by turning her head slightly to either side. Elie pauses for a short break on a hilltop to drink some water and look at her progress. With a hand gesture, she makes the map window larger and places it at a sloped angle at about waist level. At eye level, Elie opens a new window showing a visualization of her heart rate, a line graph with several different coloured lines representing the pulse readings from her wrist band, with one line for each of the past few days. Sure enough, each of the lines has a peak at approximately the same time. Elie 'taps' one of these peaks on the floating virtual display and then makes a gesture toward the map. A virtual link appears, connecting the high point on the graph to a spot on the marked path on the map. As suspected, the peak in the heart rate coincides with the location of a hill on Elie's route.

Doing the Dishes – Next we visit Zak, who has just eaten breakfast with his family (Figure 2b). His modern home looks like many others, however it is equipped with sensors that can record the usage of resources such as electricity, gas and water. After breakfast, he heads toward the kitchen with a pile of dishes. He loads many of them into the dishwasher and begins to wash the remaining items in the kitchen sink. As he turns on the tap, he sees some information appear behind the sink; Zak is barely aware that the visualization is actually produced by his HWD as it appears to be on the surface of the sink's backsplash. Figures show the rate at which the water is flowing and the cost per unit. On the adjacent panel, a chart shows the amount of water that was used each day for the past month along with the total cost of the water. Seeing that the amount has been steadily increasing over several days, he turns off the tap with just enough water to wash the remaining dishes.

A Shopping Excursion – One afternoon Marcus is at the local sporting goods store looking for a new pair of ice skates (Figure 2c), as his old ones have finally worn out. As he walks down the display aisle, a small virtual tag appears above each pair of skates he looks at, showing the cost. Each number appears in either red or black, depending on whether it is higher or lower than the amount he entered in his budget on his desktop computer before leaving home. Suddenly, he comes across a comfortable looking pair of skates that is on sale. The regular price is far higher than what he planned on spending, however the sale price is very tempting. Marcus decides to quickly re-examine his budget. He walks over to a nearby wall pulls a stylus from his pocket and begins making some 'virtual' strokes on the wall. This opens a spreadsheet containing his budget. A few more strokes bring up a pile of virtual bills from last month. He spreads the bills around the budget on the surrounding wall space. Marcus is not heavily worried about the privacy of his information, as the items are visible to only him through his HWD. Using the stylus, he copies the amount due from each bill and pastes the amount on a line of the current month's budget, after which a virtual link connects each amount to the corresponding bill. After entering a few calculations, he comes up with an estimate of his expenses that will soon be due. Marcus makes a few changes in the numbers he previously entered and decides that he can afford a few more dollars to purchase the ice skates.

These examples explore many of the possible features that can be enabled with the opportunities of a HWD discussed above. All of them illustrate typical everyday activities that rely on analytic processes. These are the type of activities that many people already do on a daily basis, however we show how our lives can be enriched by increasing the availability of information and the convenience of access using the in-situ visual analytic tools of an HWD.

Requirements for supporting In-Situ, Visual Analytic Tasks

To begin our discussion on what SAIs have to offer to the everyday user, we pose the following question: *as mobile and wearable technologies become an integral part of our everyday lives, what are the design requirements for an ideal platform to facilitate in-situ data analysis?* In response to this question, we propose a list of requirements, which we derive from several sources. We draw from our own experience designing interactive systems, from inspirations given by the above scenarios, and from existing literature surveys on visual analytics. Among the latter seminal works are an exploration of interaction in visual analytic systems from Yi et al., [21] and an early look at adapting information visualization for everyday use by Pousman et al. [17]. More recently, a survey by Huang et al. [14] distills a general taxonomy for the design space of Personal Visual Analytics. From these and other relevant works we define a set of requirements specific to *is-situ* visual analytic tasks. This list contains five primary categories: *Mobility*, *Integration*, *Interpretation*, *Multiple Views* and *Interactivity*. In the following descriptions of these items, we demonstrate how each builds upon the previous core concept.

Mobility – One implication of mobile devices is their ability to implicitly collect sensor data and infer activities of the user. This opportunity has been recently exploited by industry with the introduction of numerous tracking devices and has resulted in the recent 'Quantified Self' movement aimed at making use of this data [14], for example to benefit users' health. However, the activities of data collection and analysis are primarily conducted separately, for instance by periodic recommendations (e.g. a reminder stand up every 30 minutes) or by more intensive analysis supported by desktop tools. In contrast, we believe that supporting in-situ analysis, allowing users to analyze data directly in the situations where they are applied, will help users gain the most benefit from their data. Based on their in-depth survey, Huang et al. [14] suggest that incorporating analysis tasks into users' daily activities can help encourage adoption of analysis tools. For instance, presenting data about commuting habits at the time of the activity [11] can help users make informed choices. Likewise, if a jogger (Figure 2b) wishes to track her heart rate and estimated calories burned for training purposes, she may benefit from the ability to monitor these data during a run. This would allow her to alter her physical activity levels immediately, in contrast to comparing daily records at home on a desktop computer. In many instances such access requires the analysis tools to be mobile and usable is a range of potential situations.

Integration – In addition to being embedded in mobile or wearable devices, sensors that collect data about our daily activities can be embedded in places frequented by users, such as homes and offices. Likewise, many potential scenarios for using analytic tools can be done in-situ in these environments. Similarly, another method proposed by Huang et al. [14] for encouraging user adoption of analytic tools is to integrate visualizations into the environment. By integrating visualizations into the surrounding environment, the visualizations become readily available to the user while interfering minimally with their task. For instance, a reminder about the costs of excess water consumption (Figure 2a) is most actionable if available when and where the water is being used, say on a vanity mirror to inform a homeowner about the cost of leaving the water running while shaving.

Interpretation – Whether in-situ analysis is conducted at home or work, or while on-the-go, the adoption of analytic tools will depend on their ease of use. Pousman et al. [17] made several recommendations toward adopting visual analytic techniques to everyday situations. They suggest that visualizations should provide the most immediately relevant information, should present data in a form that is intuitive or easy to learn, and should be aesthetically pleasing to encourage contemplation. In relation to the above criteria, i.e. in the case of mobile scenarios or in those were visualizations are integrated into the environment, we add that the format of a particular visualization should be adapted to the given context; for example, information consumed in a mobile context should be highly simplified, while that integrated into a home appliance should fit both the physical form and use case of the appliance.

Multiple Views – In some contexts, sensemaking can be assisted by distributing data into multiple visualizations. For example, multiple data views are useful for making side-by-side comparisons, or for viewing an overview and a detailed view simultaneously [19]. Baldonado et al. [1] propose that the cognitive overhead of interpreting a single complex visualization can be reduced by dividing the same information into multiple simpler views that can be viewed in parallel. Each set of multiple views may contain only a subset of components from the full data set, however analysts can form mental links by switching their attention between them. As a caveat, browsing information across multiple views may incur additional costs such as additional required display space, increased memory load and effort for context switching [1]. However, visual analytic research indicates that there are cases when the benefits outweigh the costs [18]. Challenges we outline in this article (see Challenges, below) focus on how to incorporate multiple views in combination with the other requirements, such as mobility.

Head-Worn Display Technology

The concept of a display worn on the user's head originated in the late 1960s[†] and a wide variety of realizations have undergone development since. Many advances in 3D interface design have occurred as a result of Virtual Reality (VR) research since the early '90s. VR has seen a recent resurgence in popular culture as advances in hardware have progressed to the stage where relatively light-weight, low-latency devices such as Oculus Rift[††] and HTC Vive[†††] are entering the market. Optical see-through HWDs are most widely known through the introduction of Google Glass, whose introduction revealed user concerns about privacy and social acceptability. In contrast to Glass, which was designed for micro-interactions on a small, peripheral display, another class of see-through HWDs place binocular displays in the user's line of sight. These stereoscopic devices, which allow objects to be superimposed in 3D space, are ideally suited for the development of SAIs. Robust sensing technologies are also being incorporated into such devices, for tracking the user's hands or the external environment. Microsoft's Hololens, for example, can construct a model of the user's surroundings in real time and use this information to integrate virtual displays on nearby walls. Meanwhile, hardware is becoming miniaturized so that we can soon expect devices that look similar to typical eyewear in common use today – as a result, the current social acceptance barriers will be reduced to the point where such devices may be commonly worn in a wide variety of daily activities.

[†] Sutherland, I.E., A head-mounted three dimensional display. In *Proc. AFIPS '68*, 757-764.

[††] https://www.oculus.com/

[†††] http://www.htcvive.com/

Interactivity – Although actionable choices can sometimes be presented with a well-timed summary (e.g. the efficiency of a particular thermostat setting), many analytic tasks require a human decision-making component. The importance of interaction has been strongly highlighted in the visual analytics community. For example, two extensive surveys on interactive information visualization, one by Heer and Shneiderman [12] and another by Yi et al. [21] describe how interactions such as item selection, exploration of different representations, data filtering, and navigating through various levels of abstraction are essential to sensemaking in visual analytics. Although personal information visualization occupies a smaller scale, Huang et al. [14] note that human input can help to overcome the limitations of using automated data-mining techniques to identify patterns. Furthermore, these operations should be coordinated across multiple views. For instance, using a technique known as *brushing* [18] causes a selection made in one view to be reflected through visual feedback (i.e. highlights) across related items on all views. Likewise, navigation such as zooming or filtering that selects a subset of data in one view can be made to concurrently filter the subsets of other views. For example, an examination of personal finances (Figure 2c), can be assisted by several automated processes (e.g. sorting, filtering, finding sums), but ultimately requires a user 'in the loop' to understand the data and make decisions such as shopping purchases. Such a process may entail several component tasks such as navigating through multiple bills and receipts, identifying items of interest and making calculations.

Opportunities presented by HWD interfaces

Upcoming see-through HWD technologies provide many opportunities for meeting the above set of requirements for SAIs. These opportunities result from the mobile nature of HWDs, their spatial presence, and their ability to augment the real world with digital information. Because several aspects of the requirements are drawn from previous display and interaction technologies, some implementation details must be updated for HWD applications, however the primary requirements likely remain valid. For instance, viewing multiple simple visualizations side-by-side may be more efficient than viewing a single, complex visualization, however the design of these visualizations must take into account the benefits and limitations of the novel HWD platform. We elaborate on several of these opportunities below. A summary is available in Table 2.

Wearable – Being wearable devices, HWDs are inherently mobile and the interface is always available. This property makes them ideal devices for in-situ visual analytics. HWDs can be worn in virtually any situation – at home, during work, or while on the go – making them more versatile than projection-based approaches that require equipment to be installed. Also, unlike current mobile devices they can provide information with hands-free access, making use practical in situations when the user's hands are occupied, such as carrying groceries or holding on to a subway handrail.

Spatial User Interfaces – HWDs are capable of providing a far richer experience than is available with current mobile technology (see sidebar). Embedded sensors and stereoscopic viewing capabilities can provide an 'immersive' experience, where virtual objects can be made to appear in physical space, or integrated with surrounding real-world objects. Whereas the small display of a smartphone requires its user to divert their attention from the outside world to a handheld object, HWD content can instead be integrated with our surroundings. Thus HWDs have potential to attract our attention toward, rather than away from, objects in the real world. This level of integration provides a range of display possibilities, from ambient displays that require little attention, to a set of multiple display panels laid out in space.

Furthermore, any region of real world space can be used to host a virtual display, thus the amount of ‘display space’ available for use by HWDs is limited only by the ergonomic viewing constraints of the user. Multiple displays can be situated in space, for example in a ring or sphere that follows the user as she walks, or arranged to coincide with nearby surfaces such as walls or desktops (Figure 2). Switching between different views laid out in space provides a more natural and efficient experience than navigating between application views on a display that is fixed in the user’s line-of sight [7] or on a handheld mobile device. Since rearranging view does not require moving physical objects, displays can easily be placed on any existing surface or even in mid-air. Adding additional displays for multiple views does not require the expense of additional monitors, and can be used anywhere for in-situ analytic tasks.

Augmentation – Virtual displays can produce some effects that are not easily obtainable with conventional display technologies. For example, a display can easily change size, say to shrink out of the way when someone enters the room and interrupts the analytic task. Visual links can connect related items across different visualizations such as items that are jointly highlighted in a coordinated selection [4]. Such links have been shown to help users find related entities more quickly than highlights alone in a desktop environment [19]. On physical displays, visual links can only connect items across views within the same display space, whereas with virtual ‘floating’ displays, such links can connect views across interstitial space (Figure 9). While a similar effect is possible by rendering a 3D environment on a flat display [4], the spatial user interface of a HWD allows the user to actually move among and between the visualizations and links to gain the best perspective without the need for abstract virtual navigation. In this spatial environment, such links can serve the dual purpose of guiding users’ attention to related items, while also guiding users to the physical locations of other displays distributed in the physical surroundings. Other possibilities with greater leverage on Augmented Reality (AR) techniques can be imagined to integrate information more directly with the surroundings. For example, a building that contains a hotel or restaurant can be overlaid with information such as reviews, menus or room availability. Aggregated location tracking data can be overlaid on the floor of a plaza to show the paths of various visitors. Or in a home environment, different rooms can be overlaid with visualizations showing trends about temperature, overall energy consumption and human traffic flow, which could provide a useful context for programming a thermostat and control of air vents. Virtual displays can also be used in conjunction with physical displays, for example to provide peripheral display space for sorting bills around the screen of a home desktop computer, or to provide a large overview map which can be viewed alongside a detailed view on a smartphone.

Embedded Sensors – While the ideal method for controlling content on a HWD remains an open problem (see Challenges, below), many interesting possibilities are presented through the availability of embedded sensors. One such possibility is speech recognition, used by Google Glass to present and respond to a menu of available voice commands. In cases where interactivity is minimal, speech or context-based interaction can allow hands-free operation. For instance, water usage can be displayed beside a sink when it is used, or a jogger's heart rate can be continuously displayed while she is running. One potentially useful interaction mode that remains to be thoroughly explored is the use of head-tracking. By combining gyroscopic

Table 2. This table summarizes our proposed design requirements for in-situ, everyday analytics, as well as how the opportunities afforded by upcoming Head Worn Displays (HWDs) support each of these requirements.

Requirement	Description	HWD Opportunities
Mobility	Supports analysis in the environment or situation where the data are collected or applied	HWDs and wearable input devices and can be used while at home or work, or while on the go (see Scenarios, below). HWDs can support hands-free use
Integration	Information should be integrated into the user's environment through ambient displays or overlaid onto objects in use	Spatial interfaces place content in surrounding space and embedded sensors allow precise alignment for augmenting real-world objects with information displays (see sidebar, above and Content Organization, below)
Interpretation	Information should be easy to interpret for non-experts and presented in an engaging fashion	Interpretation can be simplified by augmenting objects with information in the correct context (see Scenarios and Visual Design, below). Flexibility of HWD interfaces allows 2D or 3D objects to be placed anywhere to provide imaginative and fun experiences
Multiple Views	Multiple views allow introduction of additional information for overview or comparison. Interpretation can be simplified by distributing multiple simple views instead of a single complex representation	Augmentation allows an unlimited number of displays to be placed anywhere without extra cost. Spatial interfaces spread multiple views in space for fast, intuitive switching (see Content Organization, below)
Interactivity	Gaining insights requires exploration of the data through interactive visualizations. Selection and navigation operations should be coordinated across views	Embedded sensors can track gaze, hands and other objects to provide many possible interaction methods. HWDs can work in conjunction with other devices (see Interaction Methods, below) to enable interaction both for manipulating display views and interacting with their contents (Table 3). Augmentation allows views to be coordinated with interspatial links, while a spatial interface allows users to find the best physical viewpoint

readings with the forward camera view, the device can sense where a person is directing their attention, be it toward virtual content or toward people and objects in the real world. This can be used to facilitate context-oriented interactions, such as presenting a virtual business card alongside a colleague's face or activating visualizations related to particular objects. It is also possible to embed devices with eye-tracking sensors to enable more precise gaze-based interactions. Sensors that track hands can enable ordinary surfaces to become interactive touch surfaces, allowing the use of standard gestures such as tapping for selection, flicking for scrolling or pinch-to-zoom. In-air gestures are also possible with floating displays while on the go or when touch interaction is impractical, for instance while following a messy kitchen recipe. Proxy objects can potentially be tracked using computer vision or network-connected inertial sensors to allow other forms of input such as raycasting with a stylus (Figure 3) or virtual cursor manipulation using a mouse.

Ethereal Planes Metaphor

In our current work we root our interface designs in a metaphor we call Ethereal Planes [9], in which content is placed within a set of 2D virtual windows situated in 3D physical space. In this metaphor windows act as 'containers' in a similar vein to traditional desktop interfaces, however the windows in Ethereal Planes are not constrained to the boundaries of a physical display. SAIs leverage several benefits from the situation of these 2D windows in 3D space. For example, spatial memory and proprioception can be utilized to store and retrieve information components. The virtual windows can be manipulated and organized in such a fashion to benefit interpretation, for instance by placing related information sources side-by-side for cross-referencing. Physical space can also be leveraged by placing windows in the vicinity of appropriate objects or by drawing meaningful visual links across intervening space between windows or to connect data points to physical locations.

We contrast Ethereal Planes from the concept introduced by ElSayed et al. of *situated analytics* [6], where information is rendered directly onto related objects in the environment. Both SAIs and situated analytics are similar in their use of AR display technology to support in-situ, analytic tasks. However, situated analytics assumes an explicit spatial relationship between the data and the outside world, making it particularly appropriate for particular datasets, for example geographical data. With SAIs, in-situ opportunities may be found without such an explicit spatial relationship, for instance determined by temporal or opportunistic nature of a given task. The SAI concept also places a greater emphasis on spatial interaction, which leverages body motion, whereas a situated analytic AR interface might be viewed and controlled through the screen of a smart phone or tablet. Furthermore, simple versions of SAIs (e.g. a body-centric array as in Figure 1) do not require the degree of sensing and tracking precision to overlay content directly on real-world locations as is required for situated analytics, and can be implemented using today's technology.

Despite some apparent limitations of a window-based interface, there are several practical reasons why we choose the Ethereal Planes metaphor. First, even in a spatial visualization of geographic or other spatially-related data, it is easy to imagine cases where additional window interfaces would be useful. For example, if one is viewing the paths of people's movements projected onto the floor of an environment, the analytic task may benefit from a map showing the same paths in a top-down overview of the entire location. Many applications could further benefit from flat panels containing abstract controls or text. Also, windows act as containers for organizing and compartmentalizing information, preventing it from unnecessarily cluttering or obscuring important information in the real world, leading to unwanted or even dangerous distraction. Furthermore, there is evidence, in many cases, that 2D visualizations are more easily interpreted and can be more easily manipulated than 3D visualizations. Finally, 2D interfaces are familiar to users and can in some cases incorporate existing applications or familiar elements.

In the following section, we present several scenarios that demonstrate how many of the principles described here can be applied to support analytic tasks using SAIs.

Challenges

Toward the realization of our vision for SAIs, we define a roadmap consisting of several challenges we have identified through our research to date. Several of the requirements outlined above have been partially satisfied by years of research invested in hardware and low-level software. For instance, there is high potential for *Mobility* due to an impressive variety of lightweight yet powerful devices currently available or under development by device manufacturers, together with advances in network communication and widespread distribution of networks that allow routine access to network services. Likewise, many of the spatial components necessary for distribution of *Multiple Views* and *Integration* of these in the surrounding environment are made possible by advances in sensor quality and compactness along with robust algorithms for interpreting data in real time (e.g. Microsoft Hololens). Dedicated chipsets and software libraries (e.g. Qualcomm Vuforia[5], Metaio[6]) now make it possible to robustly detect and track nearby surfaces or other objects. Likewise, robust hand tracking for natural *Interactivity* is currently offered by several low-cost devices (e.g. Microsoft Kinect[7], Leap Motion[8]).

[5] https://www.qualcomm.com/products/vuforia
[6] https://www.metaio.com/
[7] https://dev.windows.com/en-us/kinect
[8] https://www.leapmotion.com/

Building on these many promising advances, our work turns the focus toward important user interface issues. We move beyond the 'viewing' experience of most existing AR implementations toward interfaces that allow users to drill down into the data set. Furthermore, we leave behind touchscreen interfaces, currently the dominant platform, spreading usage to alternate devices such as HWDs. These goals require a fresh perspective on the look and feel of interface design, which we provide though our requirements-based approach to SAIs. We frame our past, present and future work through a set of three primary challenges: interaction methods, content organization and visual design.

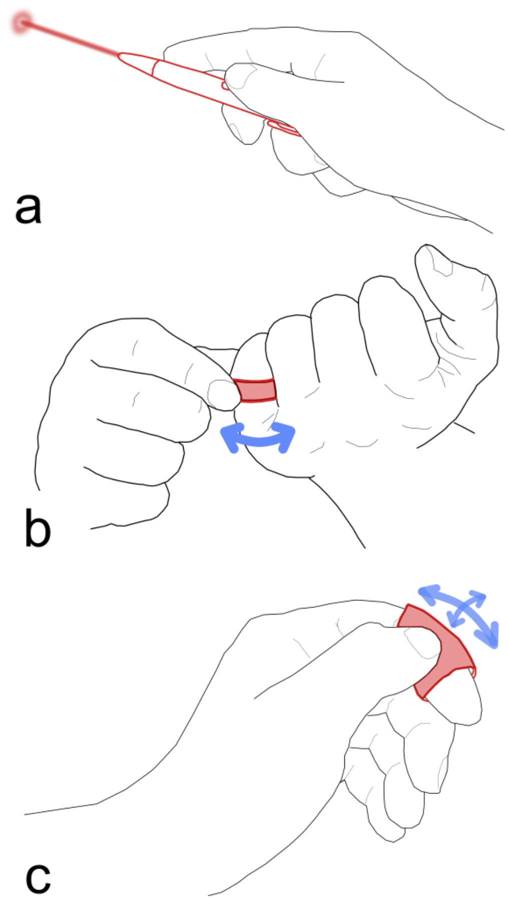

Figure 3. A variety of handheld or wearable devices can potentially be developed to provide interaction with SAIs. Possible form factors include a stylus for pointing (a), a ring for scrolling (b) or a finger pad for 2D input (c).

We single out these research areas for three reasons. First, to a certain extent, these challenges are interdependent: progress in anyone of these areas would also spur advances in another. For instance, interaction methods are closely tied to the visual design of widgets suited for a task. Take scrolling as an example: the design of a scrollbar on current WIMP interfaces is closely tied to the manner in which a pointer operates using a mouse. Conversely, the design of mice has undergone numerous iterations, including the addition of a mouse-wheel to accommodate the pervasiveness of scrolling tasks. The second reason is the fundamental nature of these areas in interface design and their necessity for the support of many practical tasks. The novelty of wearable devices prompts many fundamental questions about how to accommodate the yet unknown tasks that will become commonplace with such technology. Finally, we foresee in-situ analytic tasks to a large extent driving innovation in these specific areas.

Interaction Methods – There are many potential design options for providing interactivity with HWD content, however there is as yet no common 'standard' method that satisfies requirements such as user efficiency and social acceptability. Current market-ready solutions use voice commands (e.g. Google Glass[1]) or are equipped with trackpads (e.g. Epson Moverio[9], Optinvent ORA-1[10]). It is an open question whether these methods will gain wide user acceptance. Guided by existing research, including a substantial amount of work on interaction techniques for immersive VR environments, we can predict several properties of a successful SAI interface.

[9] http://www.epson.com/moverio
[10] http://optinvent.com/see-through-glasses-ORA

User interactions should not require large gestures, both to avoid drawing unwanted attention in public spaces [15], and to prevent user fatigue from large arm, neck, or eye motions. Conversely, interaction methods should also provide cues to make others aware when the user is engaged with the interactive system [8, 16]. Overly subtle interactions that lack such social cues can be problematic, for instance when interaction interrupts a conversation, the other person would likely prefer to know whether the user's attention is directed at them or at the computer.

Interaction with SAIs must allow a number of basic operations, such as selecting, moving and filtering items. Our work has explored several options for manipulating window layouts, for instance selecting, resizing and grouping windows. In one implementation [7], the user can shrink the window array into a palm-sized sphere to provide an overview of the current views (Figure 4b). Currently, we are exploring how to integrate two co-existing 'tiers' of operations within the Ethereal Planes metaphor; one higher tier for managing the layout of 2D views in the surrounding 3D space, and a lower tier for interacting with content within those views.

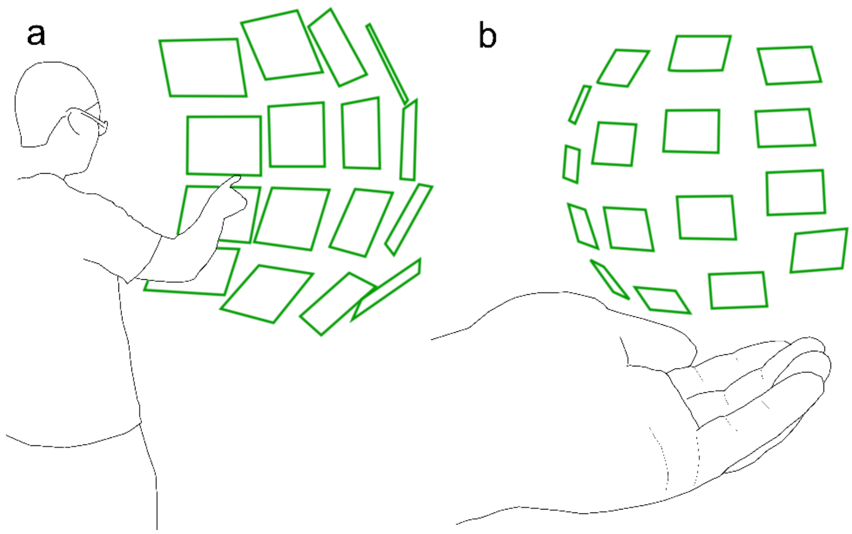

Figure 4. An important question for SAI on HWDs is how to lay out multiple views in a spatial interface. Our initial work explored parameters for arranging a body-centric array of applications (a) and of various interactions such as shrinking the array into a palm-sized overview (b).

Table 3. Operations for interacting with virtual 2D views must consider interaction at two different levels. One set of interactions is required for fine-grained interaction with visualization content while others are needed to manipulate the layout of multiple views in the surrounding 3D space.

Operation	Layout	Content
select	choose window in focus	highlight one or more items
move	translate or rotate windows in 3D space	pan content to bring items into view
resize	make a window larger or smaller	zoom in or out to change scale of items
change	open or close a visualization	change the representation of a chosen view
filter	choose which views are relevant	reduce the amount of content shown in a view
symbolic input	invoke system or menu commands	text entry, numeric input, sketching

Table 3 lists several common operations, with examples of how these can be interpreted within each tier.

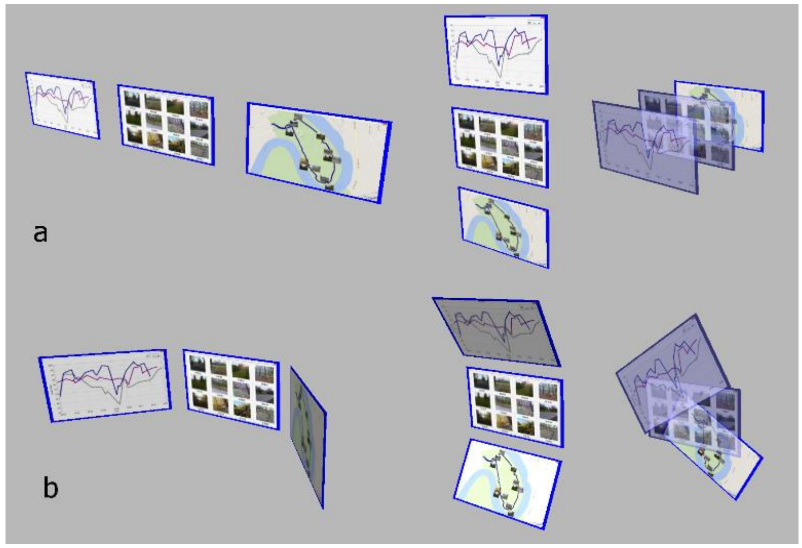

Figure 5. Expanded from [14], layout configurations can be classified as combinations of translations (a) and rotations (b).

There are a wide number of possible devices and methods for providing user interaction. One method is the direct 'touch' input explored in the scenarios above. Because the content is visible only to the user, reaching with hands may be deemed socially awkward, particularly without a visible support surface. However, direct input is highly intuitive and straightforward, thus may nonetheless be favourable in some circumstances. Handheld input devices are another option; a familiar object such as a stylus may be attractive to users, and affords many types of use, such as writing, pointing (Figure 3a), tapping and rolling. Other physical objects such as disks or cubes could be used as proxies for interacting with visualizations such pie or bar charts. Fully wearable interfaces are an attractive option for mobile situations where the user's hands may be occupied. Aside from the HWD's temple region, which is used for selection on Google Glass, possible wearable formats include: watches, which feature a flat surface for tapping, dragging or flicking; rings, which could be rotated to provide scrolling or discrete item selection (Figure 3b); gloves or other hand-worn objects [15] that provide small surfaces for 2D touch input (Figure 3c).

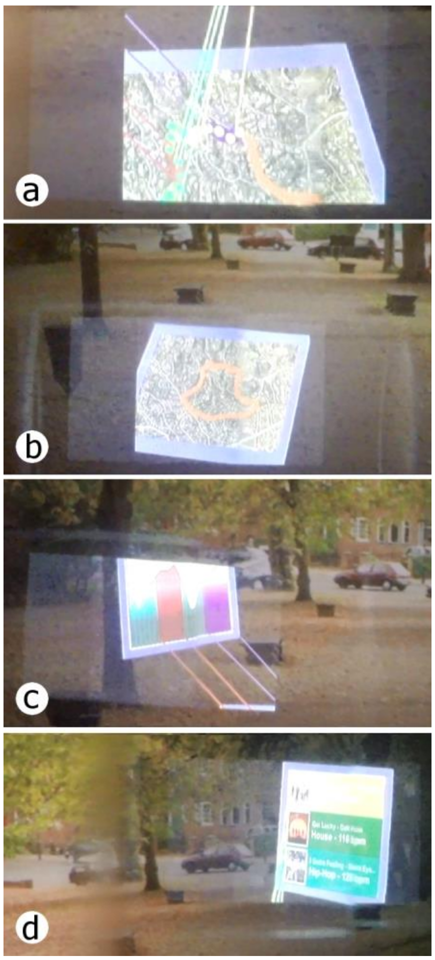

Figure 6. Images of our implementation, through a Moverio BT-100 HWD. The spatial UI assists the user in gaining insight from interspatial links (a) between a map (b), a heart rate chart (c), and a song playlist (d).

Content Organization – A vital question concerning in-situ visual analytics is how to organize a set of multiple views. How does the layout differ in a mobile context with 'floating windows', versus in a home or office with windows mapped to the surfaces of appliances or furnishings? Should the view arrangement be primarily automated or should the layout be managed manually by the user? We have begun to answer some of these questions in our research and in our current implementations of SAIs. One goal is to categorize different window layout configurations [14] (Figure 5) and to determine which layouts work best in different situations. We ran a series of studies [7] to determine parameters for the size, distance and separation distance of multiple displays in a spherical, body-centric configuration (Figure 4a) when given a restricted viewing field (40° width). One outcome of this work is evidence that situating views in the world-coordinates of a spatial interface allows users to complete a multi-view analytic tasks faster than with a ‘baseline’ interface that requires users to navigate views that are fixed to the user’s forward view. To explore such benefits, we are currently appying SAIs to a set of data collected over several days from a user traversing a route in a park, to determine how layouts can assist the interpretation of data such as heart rate, GPS and the location of events such as where a particular song was playing (Figure 6).

In other work we have explored how such spatial layouts can be integrated into surrounding surfaces in the environment. There are many existing algorithms for arranging items on a see-through display, for example to keep labels close to their objects of origin, however there has been little comparable exploration of display placement on surfaces in the surrounding environment. We developed a window manager [10] that transitions body-centric layouts to world-fixed form, with data view embedded in the user’s current environment. In addition to constraints such as surface fit, avoidance of scene objects and relative window order, this window manager applies a constraint of spatial constancy to keep layouts consistent between different environments (Figure 7). For example, if a user keeps their calendar application to the lower-right in the body-centric window layout, they can expect to always find the calendar to the lower-right in the corresponding room-fixed layout, regardless of the particular configuration of the current environment (Figure 8). Further work is required to evaluate the benefits of spatial memory using this layout manager and to measure consistency across a variety of diverse locations. There are further deeper questions to explore such as when body-centric spatial memory is preferred to contextual memory (e.g. a calendar application always near a physical clock) and how to manipulate layouts dynamically in environments with frequently moving objects.

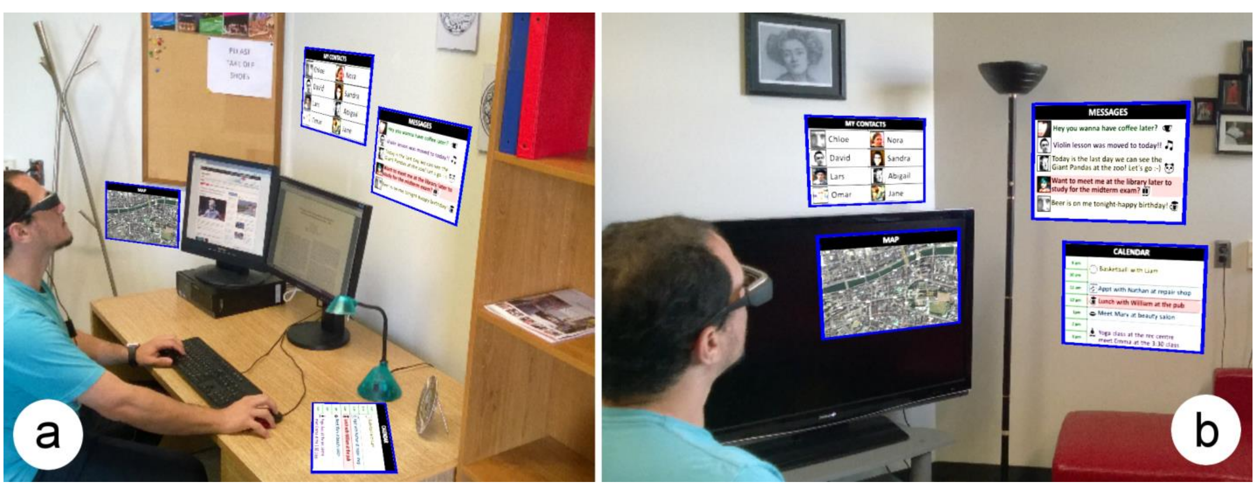

Figure 7. Another important question is how to arrange window layouts when integrating them into the user’s surroundings. One approach we have explored is to apply spatial constancy to keep windows in predictable locations in different environments (a and b), and consistent with their locations in a user’s preconfigured body-centric layout (Figure

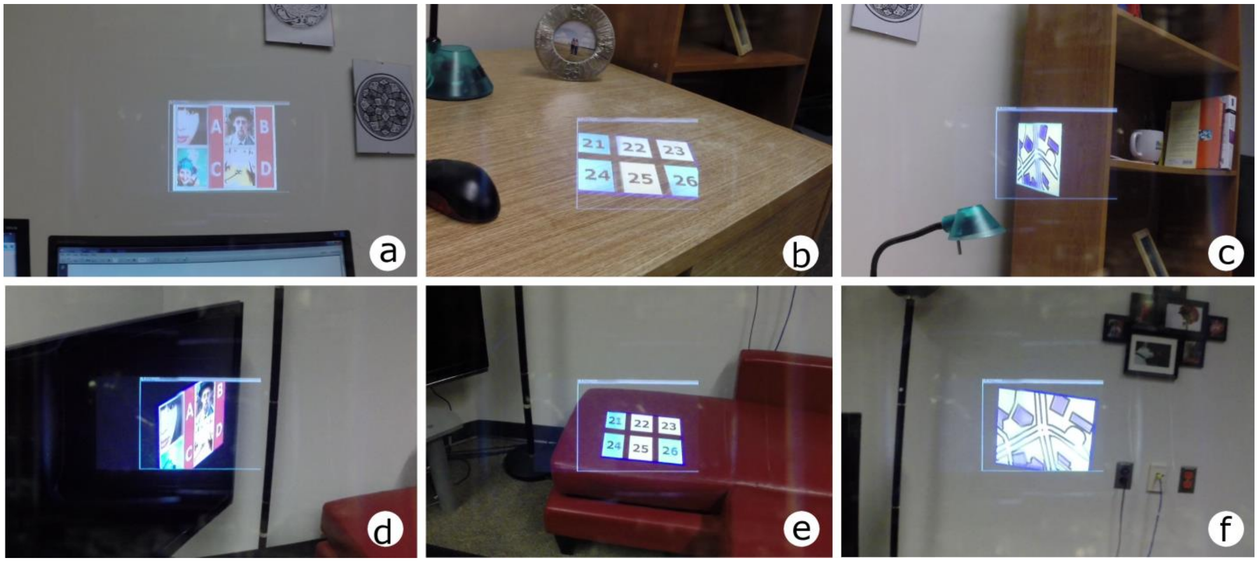

Figure 8. Images of our implementation through a Moverio BT-100 HWD showing spatially consistent layouts across office (a-c) and living room (d-f) environments. Regardless of the current location the user's applications can be found in similar relative locations. For example in this layout, the contact list (a, d) is in or above the user's forward view, the calendar (b, e) is on the lower right and a map application is on the upper right (c, f).

Visual Design – HWDs have some unique properties that set them apart from thouchscreen devices in popular use. For instance, the display of opaque objects is not possible with current transparent displays, necessitating solutions for colour blending and contrast with changing background textures. Current HWDs have a limited viewing field, causing virtual content to be cropped to a relatively small region of the human visual field. Also, the initial generations of HWDs will have limitations in display resolution and brightness, with consequences for designing applications for outdoor use. All of these inherent properties and limitations must be taken into account when designing visual content for HWDs to ensure that information can be easily interpreted using a given device, and that sufficient insights can be obtained by a casual audience of everyday visual analysts.

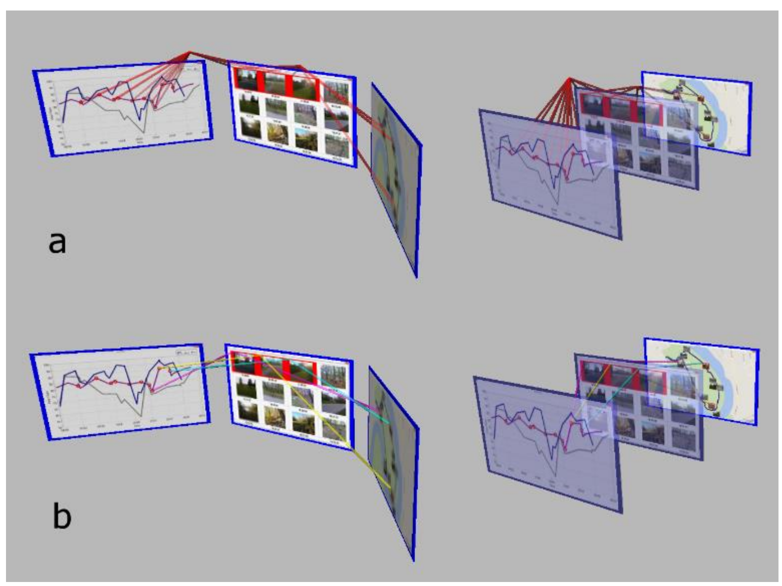

Figure 9. Interspatial visual links can join a set of data points belonging to a group selection (a). Different coloured links can show relations across views while contrasting between different selections (b).

Conversely, visual designers may also take advantage of several opportunities provided by the spatial context of data exploration in SAIs that do not exist with current mobile applications. The ability for 3D spatial view layouts in the Ethereal Planes metaphor is such an advantage. Also, these 2D views can selectively be embellished with 3D content. For example, list of items on an online shopping page can be accompanied by stereoscopic views of the corresponding products in place of 2D images. Similarly, a relief map that projects terrain or a cityscape outward from the window frame may be preferable in some instances to a flat version of the same map. Another possibility in SAIs is to display visual links that extend across physical space to reveal relationships between data points in separate visualizations. For example, one use of visual links is to tie together a number of data points that belong to a group selection (Figure 9a). Alternatively, different colours can be used to join related data points across views while contrasting a set of individual selections (Figure 9b). While evidence toward the benefits of visual links has been empirically shown for desktop interfaces [19], links between views have yet to be explored within spatial interfaces. Initial pilot studies we conducted show promise for these links to benefit both drawing attention to related content across views, and for guiding users to the spatial location of the views.

Conclusion

While wearable devices have become an integral component in personal visual analytics, much work to date has focused on using such devices for collecting contextual and biometric data. Few systems exist to support broader and advanced analytic exploration of personal data, particularly in mobile contexts. As a solution, we propose Spatial Analytic Interfaces, which combine the advantages of spatial user interfaces with principles derived from the field of visual analytics. We outline a roadmap toward the design and development of SAIs by laying out a set of design requirements and challenges. The requirements focus on applying principles from visual analytics to a new breed of spatial interfaces. The challenges discuss the need for advanced natural interaction methods to explore large personal datasets, numerous options for organizing information to support analytic tasks and the nuances of depicting visual information through HWD technologies in the SAI paradigm. Our work has taken several steps in addressing these challenges, however there remain many possibilities to explore in meeting the requirements for SAIs. However, as wearable technology and HWDs gain prominence in the general consumer market, we are hopeful that SAIs will bring powerful visual analytic capabilities to these mobile devices of the future.

References

1. Baldonado, M.Q.W., Woodruff, A. and Kuchinsky, A. Guidelines for using multiple views in information visualization. *Proc. AVI '00*, ACM (2000), 110-119.
2. Bellcore, A.B., McDonald, J.A., Michalak, J. and Stuetzle, W. Interactive data visualization using focusing and linking. *Proc. Visualization '91*, IEEE (1991), 156-163.
3. Billinghurst, M., Bowskill, J., Dyer, N. and Morphett, J. An evaluation of wearable information spaces. *Proc. VRAIS '98*, IEEE (1998), 20-27.

4. Collins, C. and Carpendale, S. VisLink: Revealing relationships amongst visualizations. *IEEE Transactions on Visualization and Computer Graphics* (2007), 1192-1199.
5. Elmqvist, N. and Irani, P. Ubiquitous Analytics: Interacting with big data anywhere, anytime. IEEE Computer 46, 4 (2013), 86-89.
6. ElSayed, N.A.M., Thomas, B.H., Smith, R.T. and Marriott, K. Using augmented reality to support situated analytics. *Proc. VR '15*, IEEE (2015), 175-176.
7. Ens, B., Finnegan, R and Irani, P. The Personal Cockpit: A spatial interface for effective task switching on head-worn displays. *Proc. CHI '14*, ACM (2014), 3171-3180.
8. Ens, B., Grossman, T., Anderson, F., Matejka, J. and Fitzmaurice, G. Candid Interaction: Revealing hidden mobile and wearable computing activities. *Proc. UIST '15*, ACM (2015), in press.
9. Ens, B., Hincapié-Ramos, J.D. and Irani, P. Ethereal Planes: A design framework for 2D information spaces in 3D mixed reality environments. *Proc. SUI '14* (2014), 2-12.
10. Ens, B., Ofek, E., Bruce, N. and Irani, P. Spatial constancy of surface-embedded layouts across multiple environments. *Proc. SUI '15*, ACM (2015), 65-68.
11. Froehlich, J., Dillahunt, T., Klasnja, P., Mankoff, J., Consolvo, S., Harrison, B. and Landay, J.A. UbiGreen: Investigating a mobile tool for tracking and supporting green transportation habits. *Proc. CHI '09,* ACM (2009), 1043-1052.
12. Heer, J. and Shneiderman, B. Interactive dynamics for visual analytics: A taxonomy of tools that support the fluent and flexible use of visualizations. ACM Queue 10, 2 (2012).
13. Hinckley, K., Pausch, R., Goble, J.C. and Kassell, N.F. A survey of design issues in spatial input. *Proc. UIST '94*, ACM (1994), 213-222.
14. Huang, D., Tory, M., Aseniero, B., Bartram, L., Bateman, S., Carpendale, S., Tang, A. and Woodbury, R. Personal Visualization and Personal Visual Analytics. *Visualization and Computer Graphics, IEEE Transactions on* , vol.PP, no.99, pp.1,1.
15. Lucero, A., Järvenpää, T., Lyons, K., White, S., Vetek, A. and Salmimaa, M. Exploring the interaction design space for interactive glasses. *CHI'13 Extended Abstracts*, ACM (2013), 1341-1346.
16. Lyons, K., Kim, S.W., Seko, S. Nguyen, D., Desjardins, A., Vidal, M., Dobbelstein, D. and Rubin, J. Loupe: a handheld near-eye display. *Proc. UIST'14*, ACM (2014), 351-354.
17. Pousman, Z., Stasko, J. and Mateas, M. Casual information visualization: Depictions of data in everyday life. *Proc. TVCG '07*. IEEE (2007), 1145-1152.
18. Stasko, John, Görg, Carsten, and Liu, Zhicheng. Jigsaw: supporting investigative analysis through interactive visualization. *Proc. VAST '07*. IEEE (2007), 131-138.
19. Steinberger, M., Waldner, M., Streit, M., Lex, A. and Schmalstieg, D. Context-preserving visual links. *Transactions on Visualization and Computer Graphics*, 17 (12). IEEE (2011), 2249-2258.
20. Waldner, M., Puff, W., Lex, A., Streit, M. and Schmalstieg, D. Visual Links Across Applications. *Proc. GI '10*, ACM (2010), 129-136.
21. Yi, J.S., ah Kang, Y. and Stasko, J.T. Toward a deeper understanding of the role of interaction in information visualization. *Visualization and Computer Graphics*, 13(6), IEEE (2007), 12224-1231.

Barrett Ens is a PhD Candidate in the Department of Computer Science at the University of Manitoba. His research interests include human-computer interaction, augmented reality and wearable computing. Ens has a BSc in Computer Science from the University of Manitoba and a BMus in Music Theory from the University of Calgary. Contact him at bens@cs.umanitoba.ca.

Pourang Irani is a Professor in the Department of Computer Science at the University of Manitoba and Canada Research Chair in Ubiquitous Analytics. His research interests lie in the areas of human-computer interaction and information visualization, with an emphasis on the emerging interdisciplinary field of ubiquitous analytics. Irani has a PhD in computer science from the University of New Brunswick. Contact him at irani@cs.umanitoba.ca.