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

Barrett Ens and Pourang Irani ▪ *University of Manitoba*

What will user interfaces look like in a post-smartphone world, and will these future interfaces support sophisticated interactions in a mobile context? Spatial analytic interfaces can leverage the benefits of spatial interaction to enable everyday visual analytic tasks to be performed in situ, at the most beneficial place and time.

Personal computing devices are becoming smaller, yet more powerful, allowing greater user mobility, increased personal data collection, and greater ability to manipulate these data to benefit our everyday lives. A catalyst in this shift is the increased access to and use of sensors and interfaces that are being integrated into what we wear. We are 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 current mobile technology, information from these devices can be ingested with 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 available space for input and display. For example, some common tasks performed on mobile devices include consumption (such as reading and watching videos), mobile communication (such as sending and receiving short messages), and organization (such as keeping a list of contacts and setting reminders). As wearable device interfaces continue to shrink, current design solutions are trending further toward simplicity; new interface paradigms (such as Google Glass and Android Wear) are designed to support *micro-interactions*, short bursts of activity that avoid impinging on the user's daily activities by minimizing task duration.

In contrast to these trends, analytic tasks require unique types of interfaces. Such tasks often require concerted thought, integrating information from multiple sources, and applying human sensemaking abilities. Typical examples of *everyday analytic tasks* include balancing a checkbook, planning a vacation, and doing a price search for the best available deal on a particular item. Although we are well accustomed to computer support for these tasks, they are not well supported by today's mobile interfaces.

To design interfaces that support analytic tasks, we can draw from the visual analytics field, which is devoted to developing tools that help users gain insights through deep exploration of multiple interlinked visualizations of diverse datasets. Although originally aimed at supporting domain experts with intensive analysis, for instance with biomedical data[1] or military intelligence reports,[2] visual analytic methods have recently been adopted for analysis of everyday personal information.[3,4] 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' movements, and wearable accessories track personal health and fitness data. This ubiquitous

 Published by the IEEE Computer Society

data-collection trend presents a growing need for tools to comprehend and digest the important patterns and to provide actionable results.[5]

The benefits to be realized from an increasing prevalence of mobile and wearable technology are 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 (such as when adjusting a thermostat). Similarly, if people are able to consult their banking history through a mobile app, they can use this information directly before making significant purchases. The mobile component is essential to in situ computing because the situational context is lost if the user has to wait to view data at home on a personal computer. Nevertheless, viewing data on the small screen of a personal mobile device is often still prohibitively cumbersome, and it lacks the potential to provide insight using multiple, coordinated views of the data.[2]

One promising approach to developing 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 a physical space, and research has shown they can improve performance on some analytic tasks.[6] 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 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) 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. Although in this article we focus on head-worn display (HWD) technology as the particularly appropriate platform for in situ analytic tasks, the concept of SAIs is platform-agnostic. We chose to work with HWDs because the technology is advancing rapidly and they are available in lightweight form factors at an affordable cost for general consumers. (See the "Head-Worn Display Technology" sidebar for more details.) HWDs such as Meta and Microsoft Hololens can come equipped with depth cameras and inertial sensors that allow us to track the user's hand, fingertip, and body motion. These features facilitate intuitive spatial interaction, such as the ability to switch between spatially situated displays by turning one's head.[7] 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, such as on a kitchen counter or backsplash for monitoring home energy consumption or in a hemispherical formation around the body in mobile situations when a user is shopping or jogging. This spatial paradigm can also support advanced techniques not possible with standard desktop displays; for example, visual links can span a physical space to connect data across multiple displays or guide users to information that is not currently in their focus of attention.[8]

Head-Worn Display Technology

The concept of a display worn on the user's head originated in the late 1960s,[1] and a wide variety of realizations have undergone development since. Many advances in 3D interface design have occurred as a result of VR research since the early 1990s. VR has seen a recent resurgence in popular culture as advances in hardware have progressed to the stage where relatively lightweight, low-latency devices such as the Oculus Rift and HTC Vive are entering the market.

Optical see-through HWDs are most widely known through the introduction of Google Glass, which revealed user concerns about privacy and social acceptability. Unlike Google 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 superimpose objects in 3D space, are ideally suited for the development of spatial analytic interfaces (SAIs). Robust sensing technologies are also being incorporated into such devices to track 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.

At the same, hardware is being miniaturized so we can soon expect devices that look similar to typical eyewear in common use today. As a result, social acceptance will likely increase to the point where such devices may be commonly worn in a variety of daily activities.

Reference

1. I.E. Sutherland, "A Head-Mounted Three Dimensional Display," *Proc. Fall Joint Computer Conf., Part I* (AFIPS), 1968, pp. 757–764.

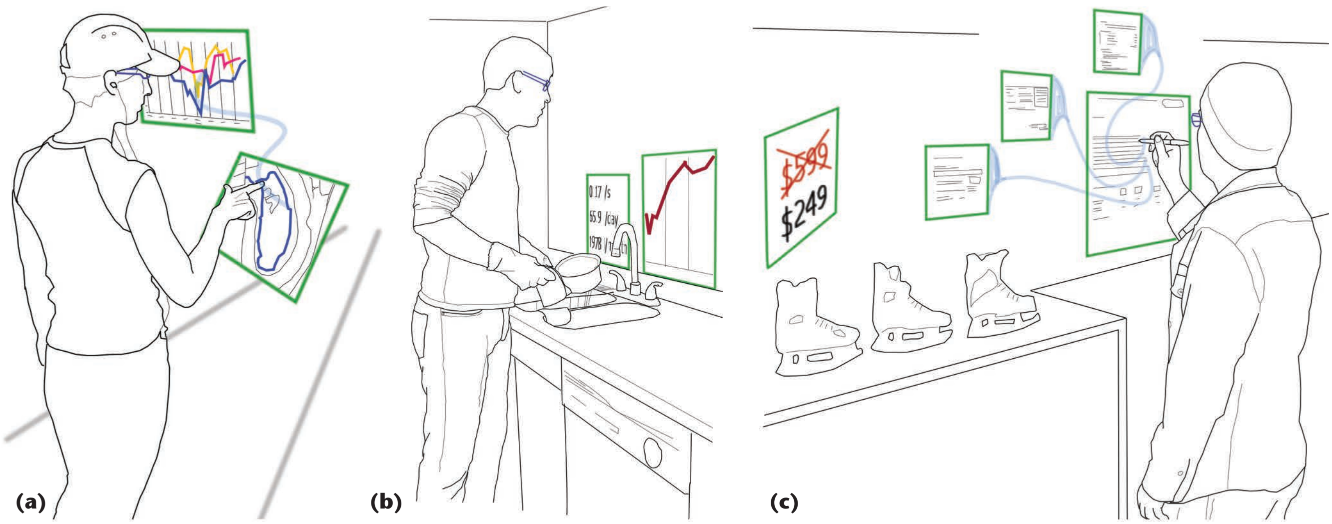

Figure 1. Three scenarios depicting beneficial uses of spatial analytic interfaces: (a) comparing heart rate and route records during exercise with visualizations that hover in space, (b) monitoring home water consumption using virtual information panels that appear on a kitchen backsplash, and (a) completing a quick budget before making a purchase using a spatially tracked stylus and virtual documents overlaid on a nearby surface.

The goal of this article is to introduce the SAI concept, discuss its benefits over current mobile interfaces, and define the implementation challenges. After describing several example scenarios, we propose a number of design requirements for novel wearable platforms to facilitate in situ analytic tasks. Finally, we discuss how this particular platform can satisfy many of the requirements for in situ analytics and highlight several open research areas where work is needed to enable practical SAI implementations.

Example Scenarios

To demonstrate the breadth of potential opportunities for SAI, we give three scenarios where data visualizations presented on HWDs can be of 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 (see Figure 1).

The following examples explore the possible features that can be enabled with a HWD. All of them illustrate typical everyday activities that rely on analytic processes. Many people already do these types of activities on a daily basis. Here, we show how our lives can be enriched by increasing the availability of information and the convenience of access using the in situ, visual analytics tools of an HWD.

A Morning Run

First we visit Elie on her morning run (see Figure 1a), during which she is accompanied by a pair of 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 window occludes her forward view, and she periodically consults them by turning her head slightly to either side.

During the run, 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 colored 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 gestures 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 (see Figure 1b). His home is equipped with sensors that record its electricity, gas, and water usage. After breakfast, he takes a pile of dishes to the kitchen. He loads many of them into the dishwasher and begins to wash the remaining items in the kitchen sink.

As Zak turns on the tap, he sees some information appear behind the sink. Although the visualization is actually produced by his HWD, 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 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 a local sporting goods store looking for a new pair of ice skates (see Figure 1c) because his old ones have 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. He then comes across a comfortable looking pair of skates that is on sale. The regular price is far higher than what he planned to spend, but the sale price is tempting.

Marcus decides to quickly reexamine his budget. He walks over to a nearby wall, pulls a stylus from his pocket, and begins making 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 worried about the privacy of his information because the items are visible only to him through his HWD. Using the stylus, he copies the amount due from each bill and pastes the amount on a line in 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 to spend few extra dollars to purchase the ice skates.

Requirements for In Situ Analysis

To begin our discussion on what SAIs have to offer 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 derived from several sources. We draw from our own experience designing interactive systems, from inspirations given by the example scenarios, and from existing literature surveys on visual analytics. Among the latter seminal works are an exploration of interaction in visual analytics systems[9] and an early look at adapting information visualization for everyday use.[4] More recently, a survey distilled a general taxonomy for the personal visual analytics design space.[3]

From these and other relevant works, we defined a set of requirements specific to in situ, visual analytics tasks. This list contains five primary categories: mobility, integration, interpretation, multiple views, and interactivity. In the following discussion, we demonstrate how each builds upon the previous core concept.

Mobility

Mobile devices can implicitly collect sensor data and infer the user's activities. In an effort to exploit such information, industry has introduced numerous tracking devices. The Quantified Self movement also aims at making use of this data,[3] for example, to benefit users' health. However, data collection and analysis activities are primarily conducted separately, for instance, with periodic recommendations (such as a reminder to 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 them gain the most benefit from their data. Based on their in-depth survey, Dandan Huang and her colleagues suggested that incorporating analysis tasks into users' daily activities can help encourage the adoption of analysis tools.[3] For instance, presenting data about commuting habits at the time of the activity can help users make informed choices.[10] Likewise, if a jogger 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 (see Figure 1a). This would help her to alter her physical activity levels immediately, as opposed to comparing daily records later on a desktop computer. In many instances, such access requires the analysis tools to be mobile and usable in a range of potential situations.

Integration

In addition to being embedded in mobile or wearable devices, sensors can be embedded in places frequented by users, such as their homes and offices. Likewise, many potential scenarios for using analytics tools can be done in situ in these environments.

Another method proposed by Huang and her colleagues for encouraging user adoption of analytic tools is to integrate visualizations into the environment.[3] By doing so, the visualizations become readily available to the user while interfering minimally with the task at hand. For instance, a reminder about the cost of excess water consumption (see Figure 1b) is most actionable if it is available when and where the water is being used—for example, 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 analytics tools will depend on their ease of use. Zachary Pousman and his colleagues made several recommendations to encourage the adoption of visual analytics techniques in everyday situations.[4] They suggested 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 addition to these criteria, for mobile scenarios or in those where visualizations are integrated into the environment, we add that a particular visualization format should be adapted to the given context. For example, information consumed in a mobile context should be highly simplified, whereas that integrated into a home appliance should fit the appliance's physical form and use case.

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 when using an overview and a detailed view simultaneously.[1] Michelle Baldonado and her colleagues argued 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.[11] Each set of multiple views may contain only a subset of components from the full dataset, but 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.[11] However, visual analytics research indicates that there are cases when the benefits of multiple views may outweigh the costs.[2] Later on in this article, we outline related challenges, focusing on how to incorporate multiple views in combination with the other requirements, such as mobility.

Interactivity

Although actionable choices can be presented with a well-timed summary (such as the efficiency of a particular thermostat setting), many analytic tasks require a human decision-making component. The visual analytics community has strongly highlighted the importance of interaction. For example, two extensive surveys on interactive information visualization described 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.[9,12]

Although personal information visualization occupies a smaller scale, Huang and her colleagues noted that human input can help to overcome the limitations of using automated data-mining techniques to identify patterns.[3] Furthermore, these operations should be coordinated across multiple views. For instance, using a technique known as brushing[13] causes a selection made in one view to be reflected through visual feedback (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 (see Figure 1c) can be assisted by several automated processes (sorting, filtering, and finding sums), but ultimately a user needs to understand the data and make decisions. 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 (see the sidebar for more details) provide many opportunities for meeting our prescribed set of SAI requirements. 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, but 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, but the design of these visualizations

Table 1. Design requirements for in situ, everyday analytics.

Requirement	Description	HWD opportunities
Mobility	Mobility allows in situ analysis to be performed in the environment or situation where the data are collected or applied.	HWDs and wearable input devices support hands-free use and can be used while at home/work or on the go.
Integration	Information should be integrated into the user's environment via 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.
Interpretation	Information should be engaging and easy for nonexperts to interpret.	Interpretation can be simplified by augmenting objects with information in the correct context. HWD interfaces allow 2D or 3D objects to be placed anywhere to provide imaginative and fun experiences.
Multiple views	Multiple views allow additional information for overview or comparison. Multiple simple views may be simpler than 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.
Interactivity	Gaining insights requires data exploration via interactive visualizations. Selection and navigation operations should be coordinated across views.	Embedded sensors can track gaze, hands, and other objects to provide possible interaction methods. HWDs can work in conjunction with other devices to enable interaction both for manipulating display views and interacting with their contents (see Table 2). Augmentation allows views to be coordinated with interspatial links, and a spatial interface allows users to find the best physical viewpoint.

must take into account the benefits and limitations of the novel HWD platform. Table 1 summarizes several of these opportunities.

Wearable

Because they are 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, which makes 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 while carrying groceries or holding on to a subway handrail.

Spatial User Interfaces

HWDs are capable of providing a far richer experience than current mobile technologies (see the sidebar for more details). 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 a smartphone's small display requires that users divert their attention from the outside world, HWD content can be integrated with our surroundings. Thus, HWDs have potential to attract 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 real-world space can be used to host a virtual display. Thus, the amount of "display space" available to HWDs is limited only by the user's ergonomic viewing constraints. 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 (see Figure 1b). 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[6] or on a handheld mobile device. Since rearranging a view with a HWD does not require moving physical objects, displays can easily be placed on any existing surface or even in mid-air. Adding displays for multiple views does not require additional monitors, so HWDs 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, shrink out of the way when someone enters the room or interrupts the analytic task. Visual links can connect related items across different visualizations, such as items that are jointly highlighted in a coordinated selection.[14] Such links have been shown to help users find related entities more quickly than highlights alone in a desktop environment.[1]

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. Although a similar effect is possible by rendering a 3D environment on a flat display,[14] a HWD's spatial user interface lets the user 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 that leverage augmented-reality (AR) techniques can be imagined to integrate information more directly within 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 controlling air vents.

Virtual displays can also be used in conjunction with physical displays, for example, to provide a peripheral display space for sorting bills around a home desktop computer screen or to provide a large overview map that can be viewed alongside a detailed view on a smartphone.

Embedded Sensors

Although the ideal method for controlling content on a HWD remains an open problem, the availability of embedded sensors offers many interesting possibilities. One such possibility is speech recognition, which Google Glass uses 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 has yet to be thoroughly explored is the use of head tracking. By combining gyroscopic readings with the forward camera view, the device can sense where a person is directing his/her 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, and pinch-to-zoom. In-air gestures are also possible with floating displays and when touch interaction is impractical—for instance, while following a kitchen recipe with messy hands. 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 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*,[15] in which content is placed within a set of 2D virtual windows situated in 3D physical space. In this metaphor, windows act as containers much like traditional desktop interfaces, but ethereal plane windows are not constrained by the boundaries of a physical display. SAIs leverage several benefits 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, such as 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 the intervening space between windows or to connect data points to physical locations.

Ethereal planes differ from the situated analytics concept introduced by Neven ElSayed and her colleagues, where information is rendered directly onto related objects in the environment.[16] 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, such as geographical data. With SAIs, in situ opportunities may be found without such an explicit spatial relationship; for instance, it may be determined by the 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 a smartphone or tablet screen. Furthermore, simple versions of SAIs (such as the body-centric array in Figure 2a) do not require the degree of sensing and tracking precision to overlay content directly on real-world locations that situated analytics requires. In fact, SAIs 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 with 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 we are 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, which could lead to unwanted or even dangerous distractions. Furthermore, there is evidence 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 we describe here can be applied to support analytic tasks using SAIs.

Challenges

Toward the realization of our vision for SAIs, we defined a roadmap consisting of several challenges we have identified through our research to date. Several of the requirements outlined previously 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 the widespread distribution of networks that allow routine access to network services. Likewise, many of the spatial components necessary for the distribution of multiple views and the 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 (such as Microsoft Hololens). Dedicated chipsets and software libraries (such as Qualcomm Vuforia and Metaio) now make it possible to robustly detect and track nearby surfaces or other objects. Likewise, several low-cost devices (Microsoft Kinect and Leap Motion) currently offer robust hand tracking for natural interactivity.

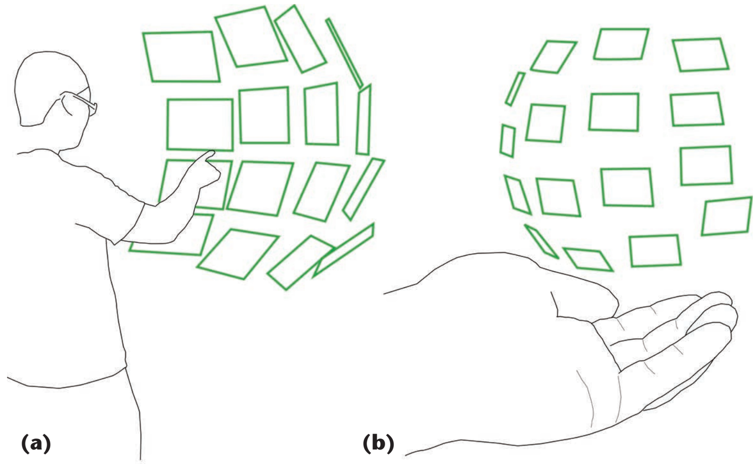

Figure 2. An important question for SAIs on HWDs is how to lay out multiple views in a spatial interface. Our initial work explores parameters for arranging (a) a body-centric array of applications and (b) of various interactions such as shrinking the array into a palm-sized overview.

Building on these many promising advances, our work focuses on 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 dataset. 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.

We single out these research areas for three reasons. First, to a certain extent, these challenges are interdependent: progress in any one area 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 for 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.

Table 2. Operations for interacting with virtual 2D views must consider interaction at two different levels.*

Operation	High-level interation tier (layout)	Low-level interaction tier (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 the item scale
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
Symbolic input	Invoke system or menu commands	Text entry, numeric input, sketching

**One set of interactions is required for fine-grained interaction with visualization content, whereas others are needed to manipulate the layout of multiple views in the surrounding 3D space.*

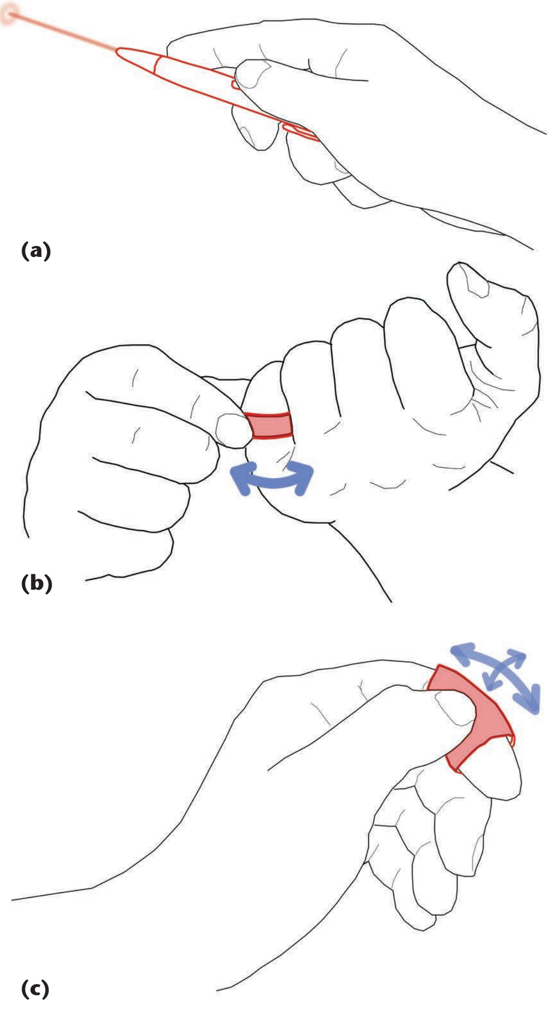

Figure 3. Example handheld or wearable devices that could be developed to provide interaction with SAIs. Possible form factors include (a) a stylus for pointing, (b) a ring for scrolling, and (c) a finger pad for 2D input.

Interaction Methods

There are many potential design options for providing interactivity with HWD content, but as yet no common standard method that satisfies requirements such as user efficiency and social acceptability. Current market-ready solutions use voice commands (such as Google Glass) or are equipped with track pads (such as Epson Moverio and Optinvent ORA-1). 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.

User interactions should not require large gestures, both to avoid drawing unwanted attention in public spaces[17,18] 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.[19,20] Overly subtle interactions that lack such social cues can be problematic, for instance, when interaction interrupts a conversation. As a rule, people 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, such as selecting, resizing, and grouping windows. In one implementation,[6] the user can shrink the window array into a palm-sized sphere to provide an overview of the current views (see Figure 2b). Currently, we are exploring how to integrate two coexisting 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.[21] Table 2 lists several common operations, with examples of how these can be interpreted within each tier.

There are numerous possible devices and methods for providing user interaction. One method is the direct "touch" input we explored in the scenarios earlier. 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 and may nonetheless be favorable in some circumstances.

Handheld input devices are another option; a familiar object such as a stylus may be attractive to users, and it 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 visu-

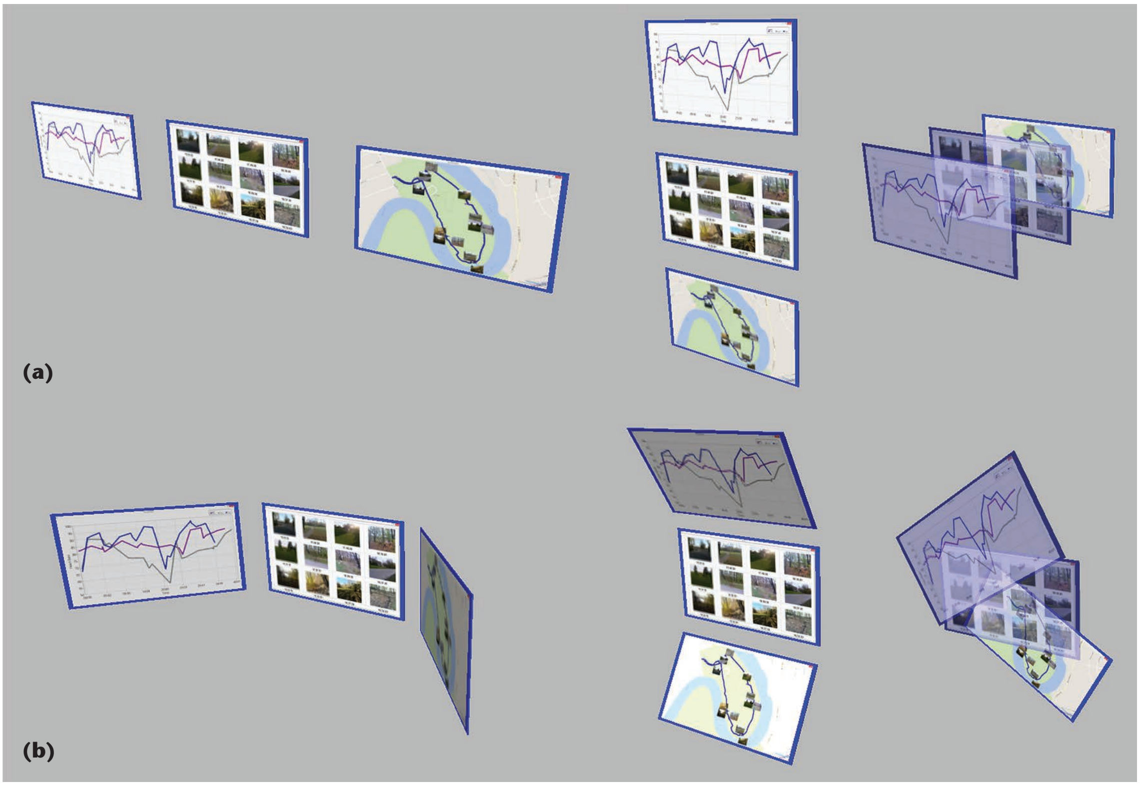

Figure 4. Window layout configurations. Layout configurations can be combinations of (a) translations and (b) rotations.

alizations such as 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, and flicking), rings (which could be rotated to provide scrolling and discrete item selection, see Figure 3b), and gloves or other hand-worn objects[17,21] (which provide small surfaces for 2D touch input, see Figure 3c).

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 SAIs implementations. One goal is to categorize different window layout configurations (see Figure 4) and to determine which layouts work best in different situations.[14] We ran a series of studies to determine parameters for the size, distance, and separation distance of multiple displays in a spherical, body-centric configuration (see Figure 2a) when given a restricted viewing field (40° width).[6] One outcome of this work is evidence that situating views in the world coordinates of a spatial interface lets users complete a multiview 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 applying SAIs to a set of data collected over several days from a user traversing a park. Our goal is to determine how layouts can assist the interpretation of data such as heart rate, GPS, and event location, such as where a particular song was playing (see Figure 5).[22]

In other work, we explored how such spatial layouts can be integrated into 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), but there has been little comparable exploration of display placement on surfaces in the surrounding environment. We developed a window manager that transitions body-centric layouts to a world-fixed form, with data view embedded in the user's current environment.[23] In addition to constraints such as surface fit, avoidance of scene objects, and relative window order, this window manager applies a spatial constancy constraint to keep layouts consistent between different environments (see Figure 6). For example, if a user keeps a calendar application on the lower right in the body-centric window layout, he or she can expect to always find the calendar in the lower right in the corresponding room-fixed layout, regardless of the particular configuration

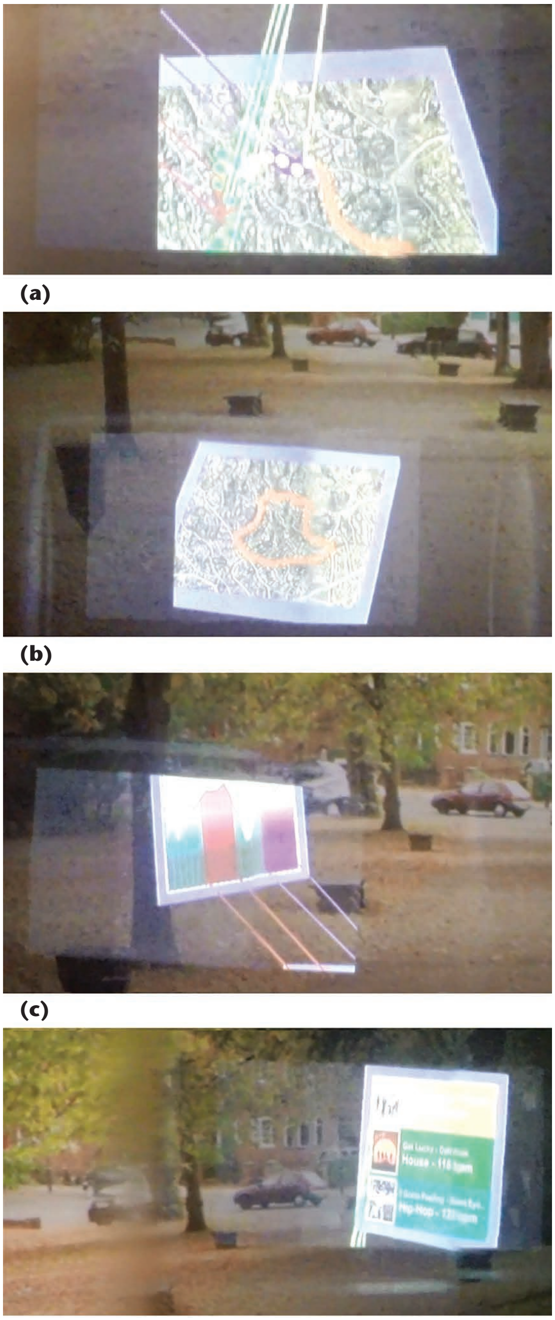

(a) (b) (c) (d)

Figure 5. Image captures from an implementation using a Moverio BT-100 HWD. The spatial user interface helps the user gain insight from (a) interspatial links between (b) a map, (c) a heart rate chart, and (d) a song playlist.

of the current environment (see Figure 7). Further work is required to evaluate the benefits of spatial memory using this layout manager and to measure consistency across various locations. There are further deeper questions to explore such as when body-centric spatial memory is preferred to contextual memory (for example, a calendar application always near a physical clock) and how to manipulate layouts dynamically in environments with frequently moving objects.

Visual Design

HWDs have some unique properties that set them apart from the touchscreen devices in popular use. For instance, the display of opaque objects is not possible with current transparent displays, necessitating solutions for color 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, which will affect the design of applications for outdoor use. All 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.

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, an item list 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 (see Figure 8a). Alternatively, different colors can be used to join related data points across views while contrasting a set of individual selections (see Figure 8b). Although researchers have shown evidence of the benefits of visual links for desktop interfaces,[1] links between views have yet to be explored within spatial interfaces. Initial pilot studies we conducted show promise for these links, both in their ability to draw attention toward related content across views and to guide users to the spatial location of the views.

Although 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,

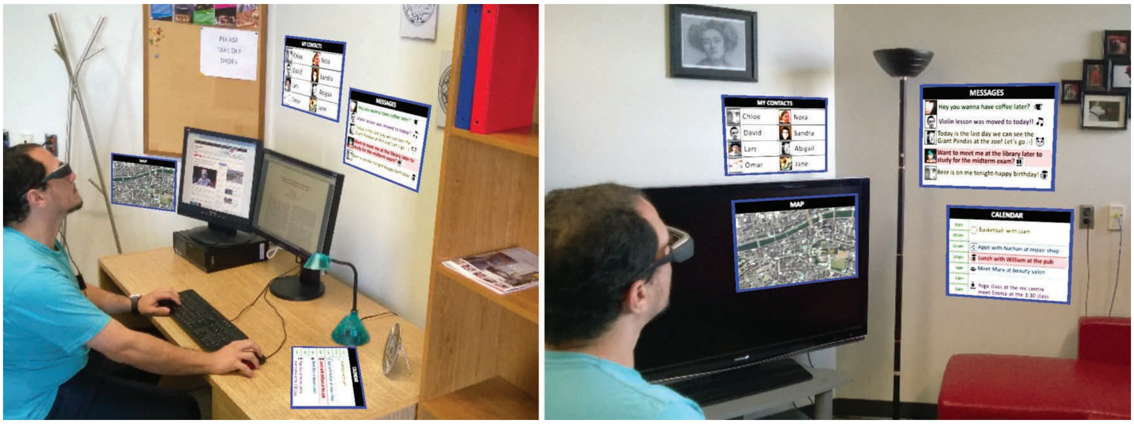

(a) **(b)**

Figure 6. Another important question is how to arrange window layouts when integrating them into the user's surroundings. We have explored applying spatial constancy (a and b) to keep windows in predictable locations in different environments and consistent with their locations in a user's preconfigured body-centric layout (see Figure 2).

particularly in mobile contexts. As a solution, we propose SAIs, which combine the advantages of spatial user interfaces with principles derived from the visual analytics field. Here, we outlined a roadmap toward the design and development of SAIs by laying out a set of design requirements and challenges. Our work has taken several steps toward addressing these challenges, but many possibilities remain to be explored in meeting the requirements for SAIs. 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. M. Steinberger et al., "Context-Preserving Visual Links," *IEEE Trans. Visualization and Computer Graphics*, vol. 17, no. 12, 2011, pp. 2249–2258.
2. J. Stasko, C. Görg, and X. Liu, "Jigsaw: Supporting Investigative Analysis through Interactive Visualization," *Proc. IEEE Symp. Visual Analytics Science and Technology* (VAST), 2007, pp. 131–138.

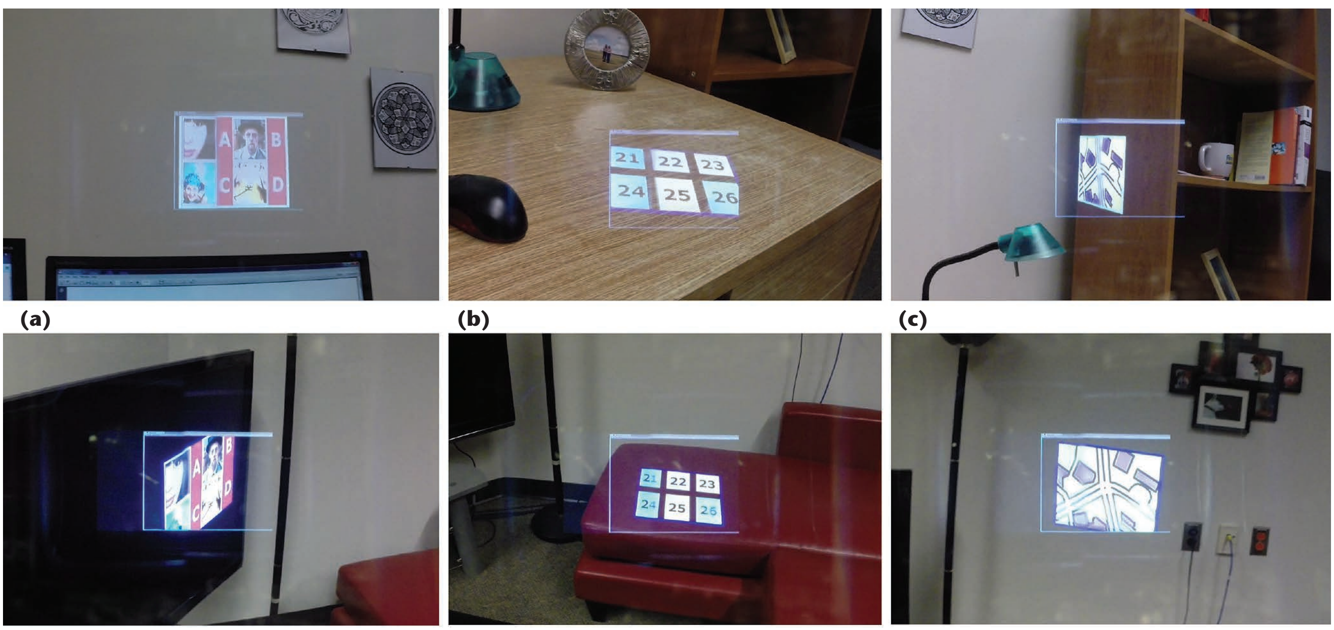

(a) **(b)** **(c)** **(d)** **(e)** **(f)**

Figure 7. Image captures from an implementation using a Moverio BT-100 HWD showing spatially consistent layouts in (a–c) office and (d–f) living room environments. Regardless of the current location, the user's applications can be found in similar relative locations. In this layout, (a, d) the contact list is in or above the user's forward view, (b, e) the calendar is on the lower right, and (c, f) a map application is on the upper right.

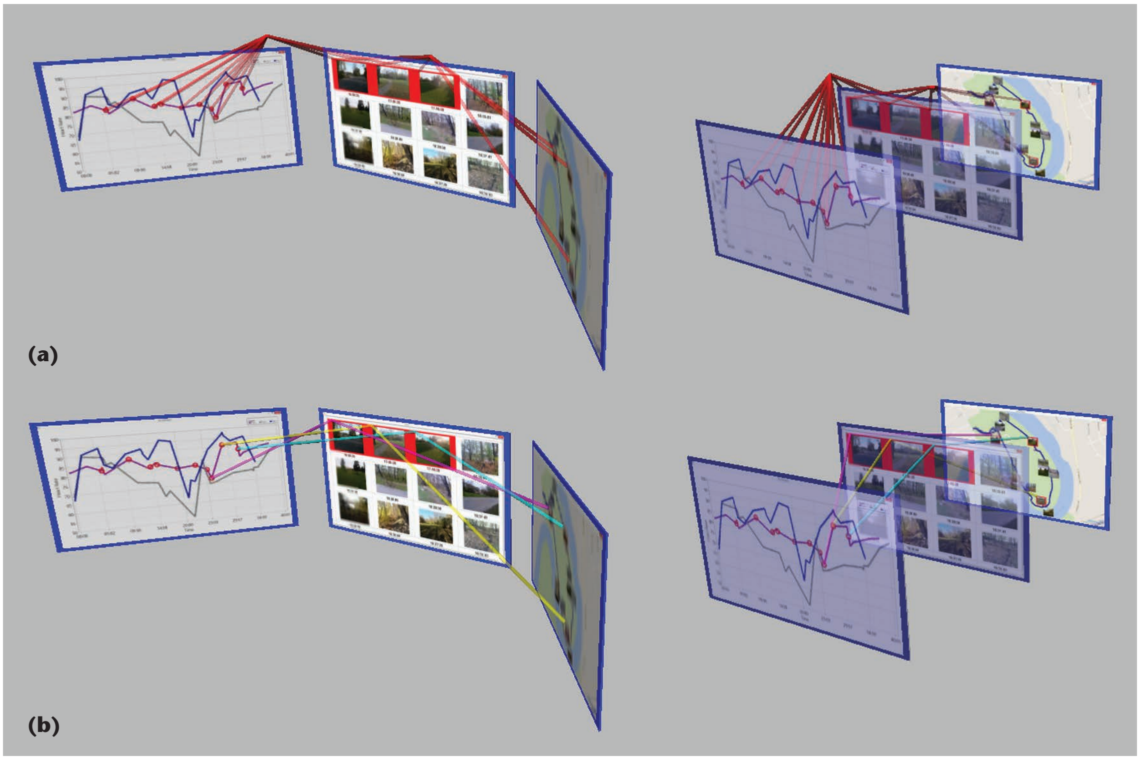

Figure 8. Displaying visual links. (a) Interspatial visual links can join a set of data points belonging to a group selection. (b) Different colored links can show relations across views while contrasting between different selections.

3. D. Huang et al., "Personal Visualization and Personal Visual Analytics," *IEEE Trans. Visualization and Computer Graphics*, vol. 21, no. 3, 2015, pp. 420–433.
4. Z. Pousman, J. Stasko, and M. Mateas, "Casual Information Visualization: Depictions of Data in Everyday Life," *IEEE Trans. Visualization and Computer Graphics*, vol. 13, no. 6, 2007, pp. 1145–1152.
5. N. Elmqvist and P. Irani, "Ubiquitous Analytics: Interacting with Big Data Anywhere, Anytime," *Computer*, vol. 46, no. 4, 2013, pp. 86–89.
6. B. Ens, R. Finnegan, and P. Irani, "The Personal Cockpit: A Spatial Interface for Effective Task Switching on Head-Worn Displays," *Proc. SIGCHI Conf. Human Factors in Computing Systems* (CHI), 2014, pp. 3171–3180.
7. M. Billinghurst et al., "An Evaluation of Wearable Information Spaces," *Proc. IEEE Virtual Reality Ann. Int'l Symp.* (VRAIS), 1998, pp. 20–27.
8. M. Waldner et al., "Visual Links across Applications," *Proc. Graphics Interface* (GI), 2010, pp. 129–136.
9. J. Heer and B. Shneiderman, "Interactive Dynamics for Visual Analytics: A Taxonomy of Tools that Support the Fluent and Flexible Use of Visualizations," *ACM Queue*, vol. 10, no. 2, 2012; queue.acm.org/detail.cfm?id=2146416.
10. J. Froehlich et al., "UbiGreen: Investigating a Mobile Tool for Tracking and Supporting Green Transportation Habits," *Proc. SIGCHI Conf. Human Factors in Computing Systems* (CHI), 2009, pp. 1043–1052.
11. M.Q.W. Baldonado, A. Woodruff, and A. Kuchinsky, "Guidelines for Using Multiple Views in Information Visualization," *Proc. Working Conf. Advanced Visual Interfaces* (AVI), 2000, pp. 110–119.
12. J.S. Yi, Y. Kang, and J.T. Stasko, "Toward a Deeper Understanding of the Role of Interaction in Information Visualization," *IEEE Trans. Visualization and Computer Graphics*, vol. 13, no. 6, 2007, pp. 1224–1231.
13. A.B. Bellcore et al., "Interactive Data Visualization Using Focusing and Linking," *Proc. 2nd IEEE Conf. Visualization*, 1991, pp. 156–163.
14. C. Collins and S. Carpendale, "VisLink: Revealing Relationships Amongst Visualizations," *IEEE Trans. Visualization and Computer Graphics*, vol. 13, no. 5, 2007, pp. 1192–1199.
15. B. Ens et al., "Ethereal Planes: A Design Framework for 2D Information Spaces in 3D Mixed Reality Environments," *Proc. 2nd ACM Symp. Spatial User Interaction* (SUI), 2014, pp. 2–12.
16. N.A.M. ElSayed et al., "Using Augmented Reality to Support Situated Analytics," *Proc. IEEE Virtual Reality* (VR), 2015, pp. 175–176.
17. A. Lucero et al., "Exploring the Interaction Design Space for Interactive Glasses," *CHI'13 Extended Abstracts Human Factors in Computing Systems* (CHI EA), 2013, pp. 1341–1346.
18. M. Serrano, B. Ens, and P. Irani, "Exploring the Use of Hand-To-Face Input for Interacting with Head-Worn Displays," *Proc. SIGCHI Conf. Human Factors in Computing Systems* (CHI), 2014, pp. 3181–3190.
19. B. Ens et al., "Candid Interaction: Revealing Hidden Mobile and Wearable Computing Activities," *Proc.*

28th Ann. ACM Symp. User Interface Software & Technology (UIST), 2015, pp. 467–476.
20. K. Lyons et al., "Loupe: A Handheld Near-Eye Display," *Proc. 27th Ann. ACM Symp. User Interface Software & Technology* (UIST), 2014, pp. 351–354.
21. B. Ens et al., "Combining Ring Input with Hand Tracking for Precise, Natural Interaction with SAIs," *Proc. 4th ACM Symp Spatial User Interaction* (SUI), 2016, pp. 99–102.
22. B. Ens and P. Irani, "Personal Command and Control: A Spatial Interface for Head-Worn Displays as a Platform for Everyday Visual Analytics," *Proc. Workshop: A Personal Perspective on Visualization and Visual Analytics at DIS 2014, Conf. Designing Interactive Systems*, 2014.
23. B. Ens et al., "Spatial Constancy of Surface-Embedded Layouts across Multiple Environments," *Proc 3rd ACM Symp. Spatial User Interaction* (SUI), 2015, pp. 65–68.

Barrett Ens *received his PhD from 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 BS 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 a Canada Research Chair in Ubiquitous Analytics. His research interests include 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.*