

Expected, Sensed, and Desired: A Framework for Designing Sensing-Based Interaction

STEVE BENFORD, HOLGER SCHNÄDELBACH, BORIANA KOLEVA,
ROB ANASTASI, CHRIS GREENHALGH, TOM RODDEN,
JONATHAN GREEN, AHMED GHALI, and TONY PRIDMORE

The University of Nottingham

BILL GAVER, ANDY BOUCHER, BRENDAN WALKER,
and SARAH PENNINGTON

The Royal College of Art

ALBRECHT SCHMIDT, and HANS GELLERSEN

Lancaster University

and

ANTHONY STEED

University College London

Movements of interfaces can be analyzed in terms of whether they are expected, sensed, and desired. Expected movements are those that users naturally perform; sensed are those that can be measured by a computer; and desired movements are those that are required by a given application. We show how a systematic comparison of expected, sensed, and desired movements, especially with regard to how they do not precisely overlap, can reveal potential problems with an interface and also inspire new features. We describe how this approach has been applied to the design of three interfaces: pointing flashlights at walls and posters in order to play sounds; the Augurscope II, a mobile augmented reality interface for outdoors; and the Drift Table, an item of furniture that uses load sensing to control the display of aerial photographs. We propose that this approach can help to build a bridge between the analytic and inspirational approaches to design and can help designers meet

This work has been supported by the EPSRC funded Equator Interdisciplinary Research Collaboration [Gr-N-15986]. We also gratefully acknowledge the support of Getmapping.com and Nottingham Castle Museum.

Authors' addresses: S. Benford, H. Schnädelbach, B. Koleva, R. Anastasi, C. Greenhalgh, T. Rodden, J. Green, A. Ghali and T. Pridmore, School of Computer Science, Jubilee Campus, The University of Nottingham, Nottingham, NG8 1BB, UK; email: sdb@cs.nott.ac.uk; B. Gaver, A. Boucher, B. Walker and S. Pennington Department of Interaction Design, The Royal College of Art, Kensington Gore, London, SW7 2EU, UK; A. Schmidt, and H. Gellersen, Computing Department, Lancaster University, Lancaster, LA1 4WA, UK; A. Steed, Department of Computer Science, University College, London WC1E 6BT, UK.

Permission to make digital or hard copies of part or all of this work for personal or classroom use is granted without fee provided that copies are not made or distributed for profit or direct commercial advantage and that copies show this notice on the first page or initial screen of a display along with the full citation. Copyrights for components of this work owned by others than ACM must be honored. Abstracting with credit is permitted. To copy otherwise, to republish, to post on servers, to redistribute to lists, or to use any component of this work in other works requires prior specific permission and/or a fee. Permissions may be requested from Publications Dept., ACM, Inc., 1515 Broadway, New York, NY 10036 USA, fax: +1 (212) 869-0481, or permissions@acm.org.

© 2005 ACM 1073-0616/05/0300-0003 \$5.00

the challenges raised by a diversification of sensing technologies and interface forms, increased mobility, and an emerging focus on technologies for everyday life.

Categories and Subject Descriptors: H.5.2 **[Information Interfaces and Presentation]**: User Interfaces—*User-centered design; interaction styles; theory and methods*

General Terms: Human Factors, Design, Theory

Additional Key Words and Phrases: Sensing, mixed reality, augmented reality, mobile and wireless applications, interactive furniture

1. INTRODUCTION

This article introduces a design framework for sensing-based interfaces in which designers are encouraged to compare expected physical movements with those that can be sensed by a computer system and those that are desired by a particular application. They are asked to treat the boundaries between these as interesting areas of the design space, both in terms of problems to be solved and also opportunities to be exploited. Our framework is motivated by four recent trends in human computer interaction (HCI).

First is the growth of interest in sensing technologies that enable interfaces to actively respond to a wide variety of user behaviors. Video and audio tracking, electronic tagging, load sensing, light sensing, physiological sensing, and other kinds of sensing underpin proposals for new styles of interface such as location-based and context-aware interfaces, smart environments, tangible interfaces, ambient interfaces, and affective interfaces. Instead of users directly manipulating the interface with their hands on the controls, these interfaces often autonomously react to users who are neither directly tethered to them or necessarily always in control or even actively engaged. However, such interfaces still need to be interpretable and to some extent predictable, raising new challenges for interface designers [Bellotti 2002].

Second, the physical forms of interfaces are diversifying. They are simultaneously getting smaller (e.g., wearable, portable, and embedded displays) and also larger (e.g., immersive displays such as CAVE-like systems [Cruz-Neira 1992]). There is also a trend towards more purpose-designed and specialized one-off appliances in contrast to the general purpose “one size fits all” PC. Consequently, designers increasingly mix and match technologies to create new interfaces, requiring them to be aware of the boundary conditions that result from attaching different sensors to different physical forms.

Third is an increase in mobility. The increasing power of handheld computers and mobile phones, coupled with the spread of wide-area positioning technologies such as GPS and cellular positioning, means that the nature of our physical interaction with computers is changing. We now see interfaces that require users to walk over large areas [Cheverst 2000], carry objects [Rekimoto 2000; Ullmer 1998], run [Flintham 2003], kick footballs [Mueller 2003], play table tennis [Ishii 1999], dance [Griffith 1998], and otherwise engage in physical movements that could be considered extreme when compared to using a keyboard and mouse. This requires designers to take a broader view of how an interface might potentially be used, more carefully considering the possibilities of surprising or even bizarre physical interactions.

Fourth, the nature of applications is changing. Looking beyond traditional productivity oriented workplace technologies where performance is a key objective, HCI is increasingly considering applications for everyday life. Interface design now encompasses leisure, play, culture, and art, and in some cases the design of computer interfaces is merging with the design of everyday appliances. Consequently, there is a shift in emphasis towards interfaces that are pleasurable, aesthetic, expressive, creative, culturally relevant, and even provocative. This trend requires designers to take a fresh perspective on application “requirements”, adopting new approaches to design, including those that stimulate imaginative thinking and even exploit ambiguity.

These four trends raise significant new challenges for interface designers. Some of these are concerned with how users interact with sensing systems. Others are more concerned with envisioning new kinds of interaction, opening up new design possibilities and considering how interfaces might potentially be (mis)treated and (ab)used in future situations. Together, they suggest designing systems in which physical input-output (I/O) devices are no longer treated as specialized and separate components, but rather are seen as an integral part of what the thing “is”. In turn, this requires a holistic approach to design in which the mechanics of interaction and new design possibilities are combined.

2. BUILDING ON PREVIOUS TAXONOMIES AND METHODS

There is already a wide variety of taxonomies, methods, and guidelines available to interface designers to support different aspects of the design process.

2.1 Taxonomies for Input Devices

Various taxonomies have been proposed to help designers reason about the detailed mechanics of how users interact with different input devices, several of which have considered how physical movements map onto the sensing capabilities of the interface.

Buxton [1983] reviews some early user interface taxonomies and concludes that there is not a sufficient focus on the pragmatic aspects of interaction, for example, on how the choice of particular input devices affects interaction. He introduces a taxonomy that classifies input devices (mostly for desktop direct-manipulation interfaces) according to the input property sensed (position, motion, or pressure) and the number of dimensions sensed (one, two, or three), enabling designers to reason about pragmatic choices with regard to input devices. Foley et al. [1984] focus on the range of tasks that are undertaken in an application—selection, positioning, orienting, path specification, quantification, and text input—and how these can be supported with different input devices. Bleser and Sibert [1990] introduce a tool that uses heuristics to suggest interaction methods given a task description. Card et al. [1991] produce a very wide-ranging review of input devices used in desktop interfaces, characterizing individual one-dimensional sensors in terms of force versus position, linear versus rotary, and absolute versus relative. Finally, Jacob et al. [1994] argue that such taxonomies should also consider which actions can be expressed simultaneously with a given device.

Taxonomies such as these offer designers detailed insights into the relationship between users' physical manipulations of input devices and the capabilities of sensors. However, they are limited with regard to addressing the four trends noted in our introduction. They assume as a starting point that the user wants to interact, knows how to interact, and has a goal in mind, and then helps the designer make this interaction more efficient. This may not be a good assumption for smart environments, location-based services, and other interfaces that actively engage passing users and push information at them, and where users' intentions may be less clear. Furthermore, although such taxonomies can inspire new classes of device as both Buxton [1983] and Card et al. [1991] demonstrate, they are not primarily focused on generating new design ideas. Rather, they are analytic tools for refining an interface once its functionality has been nailed down, typically by matching the right input device to each interaction task.

These taxonomies also tend to focus on relatively direct and precise sensors. In this article, we are interested in extending their analysis to less precise sensors such as video tracking and GPS which involve a much higher level of uncertainty. With less precise sensors, the areas of potential mismatch between actions and sensing become broader and, we argue, more interesting as design spaces.

Finally, previous taxonomies tend to assume that the user's focus is on the computer interface, and that physical I/O devices are peripherals, that is, they are tools to get at what you're interested in, and not the focus of interest in and of themselves. As such, their form can justifiably be determined almost entirely by their function as I/O devices. In contrast, a focus on individually designed appliances and augmented physical artifacts brings the design of the artifact itself more into focus. The forms of "designer" devices will be strongly influenced by preexisting functionality and cultural connotations and the fact that they are I/O devices is only a part of their meaning to users. This naturally leads us to our second thread of related research, inspirational design methods.

2.2 Inspirational Design Methods

There is a long and extensive history of user-centered design methods in HCI, including task-analysis techniques that draw on cognitive psychology in order to understand how individuals plan and carry out detailed interactions with particular interfaces, for example, GOMs [John 1996], the use of ethnography to inform system design with an understanding of the social and situated use of technologies in particular environments [Hughes 1992] and participatory design methods that directly involve users as partners in the design process, sometimes through working with low-tech physical prototypes (e.g., Ehn [1991]).

Of particular relevance to this article are inspirational methods such as cultural probes [Gaver 1999] whose primary aim is to inspire new design ideas and that are targeted at designing products for everyday life rather than the work place, focusing on creative, engaging, and playful applications of computer technologies. In one example, a community of seniors in Amsterdam was

mailed activity packs that included digital cameras, maps, and postcards which they could use to record snapshots of their lives and feelings about their environment. The completed cultural probes were then used by the design team to create fictional scenarios as the basis for new inspirations. An underlying idea here is that working with physical artifacts allows people to tap their embodied understanding of things, from their affordances [Gaver 1992; Norman 1999] to their cultural connotations, in reasoning about how new designs might work. A related approach that also draws upon the disciplines of art and design is to recognize the potentially positive role of ambiguity in creating interfaces that stimulate engagement and provoke reflection [Gaver 2003].

These inspirational design methods support our framework in two ways. First is the importance of deliberately undermining assumptions. In our case, we wish designers to explicitly consider the unexpected—unlikely patterns of use that might lead to extreme movements or might result in unlikely sensor data. Second, is the idea of looking at boundary conditions, the ambiguous area where physical movement may not precisely match the capabilities of sensors, as a new source of design opportunities.

However, these methods suffer from their own limitations. In particular, they do not support the kinds of detailed analysis of design trade-offs that were the focus of the interface taxonomies that we reviewed previously. In response to these observations and the four trends noted in our introduction, we now introduce a design framework that aims to build a bridge between the idea-generation phase of design (supported by ethnography, participatory design, and inspirational design methods) and the refinement phase where detailed trade-offs are explored (supported by analytic frameworks and taxonomies), and that encourages designers to focus on extraordinary or quirky behaviors and boundary conditions for sensor technologies. We now introduce our framework, beginning with definitions of expected, sensed, and desired movements.

3. EXPECTED MOVEMENTS

The physical form of an interface fundamentally shapes the kinds of interactions that users can and will perform. We define expected movements as being those movements that users might be expected to carry out; they are natural movements for a given combination of user, interface, and environment. Besides expected movements, there are less expected movements. These are unusual, although certainly possible movements, and when they occur, they indicate that the interface is being used in an atypical manner or context. Outside the realm of these movements are nonsensical movements that are impossible without breaking the interface or the laws of physics. We are interested in identifying unexpected and nonsensical movements as well as expected ones. We briefly illustrate this idea in relation to existing interfaces.

Handheld Computer (PDA). Expected movements include holding the PDA in one or two hands while standing still and looking at the screen. Movements of the interface can then be expected to follow the principal axes of rotation of the human body (about wrists, elbows, shoulders, spine, etc.). Examples of less expected (but possible) movements might be carrying the device above

your head, interacting while running, throwing the device from one person to another, or attaching it to a balloon, although note that Paulos [1999] presents a system in which a small computer, camera, and microphone are attached to a remote controlled blimp in order to create a Personal ROving Proxy (PROP). An example of a nonsensical movement is moving through a solid wall.

Tangible Interface Object Moved Across a Surface. (e.g., a block on a table [Underkoffler 1999] or a post-it note on a drawing board [Klemmer 2002]. It is expected movement to place the device on the surface in an orientation suggested by its shape. It is also expected that users will carry objects between different surfaces, a possibility exploited in mediaBlocks [Ullmer 1998] and the work of Rekimoto et al. [2000] that treat physical objects as containers for digital information. It may be less expected to stack objects, turn them upside down, raise them into the air, or change their shapes and colors.

Laser Pointer Used to Interact With a Screen [Olsen 2001; Myers 2002]. Here it is expected to hold the laser pointer in one hand and point it at targets for short periods. Again, movements will typically follow natural arm movements. It is less expected to wave it about wildly, or to hold the beam perfectly still and point at an object for many minutes (this could be achieved by resting the pointer on a surface with the switch taped down). It is impossible to move the beam instantly from one surface to another.

A Virtual Reality Head-Mounted Display (HMD). Normal head movements will be slow and will not feature extreme pitch and roll rotations and the hands will stay within arm-extension distance of the head/body. The user will not move far in physical space due to connecting cables and the possibility of colliding with objects that they cannot see. Less expected movements are rapid and extreme head movements (but perhaps the head-mount is being held in the hands) or a large separation between head and hands (but perhaps several people are holding the equipment).

The Common Mouse. It is expected to move a mouse horizontally on its mouse mat. It is common to lift it off the surface, move it through the air, and then place it down on the surface again, and also to rotate it. It may be less expected to move the mouse entirely away from its mouse mat or surface, turn it over, and use it as a trackball, or carry it away altogether (though rollerball and wireless mice offer different possibilities here).

We offer some general observations on these examples. First, they concern different properties of movement:

- degrees of freedom: which combinations of translations and rotations are expected?
- range: how far is the device expected to move in each degree of freedom?
- speed: how quickly is it expected to move in each degree of freedom?
- accuracy: how precisely is it expected to move in each degree of freedom?
- stability and maintainability: how stable will expected movement be over time?

Second, distinctions between expected and less expected movements emerge from a combination of other factors. The physical form of the interface (its

size, shape, texture, weight, joints, supports, handles, etc.) suggests particular movements. Drawing on the ideas of Norman [1999] and also Gaver [1992], both based upon Gibson [1977], the relationship between the user and the interface affords certain movements. The form of the interface also constrains possible movements [Norman 1988], for example, through size, weight, shape, tethering by physical cables, and physical joints that constrain rotation and extension. Furthermore, the human body imposes constraints on movement in terms of reach, natural postures, and rotations.

Third, the surrounding environment implies and constrains expected movements through its size and shape and through the presence and absence of obstacles, including boundaries. This relationship between the environment and the movement of an interface has not featured strongly in previous taxonomies and frameworks but is one that takes on an increased significance as interfaces become mobile.

Finally, the designer will hold assumptions about how the interface will be moved based upon their own experience and vision of how the interface is intended to be used. A key feature of our approach is encouraging designers to deliberately question these assumptions by imagining extreme and bizarre scenarios in which unexpected movements could occur.

4. SENSED MOVEMENTS

Next, we turn our attention to an interface's sensed movements, defined as those that can actually be measured by a computer. These are determined by the particular combination of sensing technologies that are used with the interface. There is an increasingly wide range of such technologies to choose from, each with its own capabilities and limitations. The following list considers a few representative examples.

Global Positioning System (GPS). This is a versatile technology for sensing position on and above the Earth's surface that can be integrated into PDAs and wearables. However, a GPS does not generally work indoors or underground, at extreme northerly or southerly latitudes, and can be problematic in built-up urban environments or in poor weather. Compared to the transducers used in devices for traditional direct manipulation interfaces, GPSs can suffer from considerable inaccuracy which varies over space and time.

Video Tracking. Video tracking can be readily combined with interfaces such as laser pointers, flashlights, and tangible objects. This technology can track the presence, identities, number, position, orientation, and movement of known objects, including people. However, the number of cameras and their fields of view limit the extent of the surfaces that can be tracked. Stereo or mono deployment determines the ability to track depth. Camera resolution and the rate at which frames can be processed limit accuracy. Systems are also usually tailored to follow specific objects in particular environments and may be unable to cope with different objects, multiple objects, occlusion, and changes in lighting.

Electro-Magnetic Tracking. This technology is widely used with immersive virtual reality (VR) and usually trades off tracking range for accuracy, with typical examples providing roughly a centimeter of accuracy over a couple of

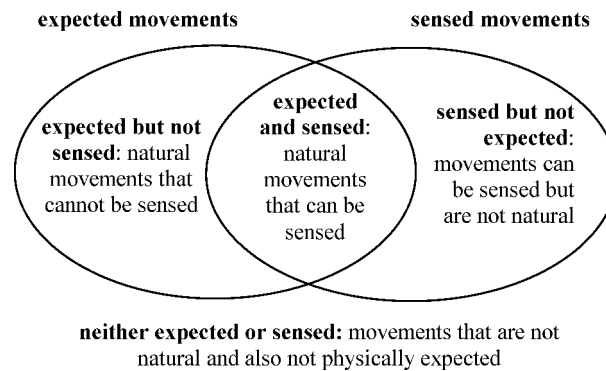


Fig. 1. Expected versus sensed movements.

meters range, or only several centimeters accuracy over several meters. Systems also suffer from interference and accuracy and stability decays towards the edge of the tracking volume.

Radio Frequency Identification (RFID). Widely used to recognize the identities of objects that are placed on surfaces *RFID* is characterized by different reading ranges, responsiveness (limiting how quickly an object can pass by), and by the number of tags that may be read simultaneously.

Mechanical Tracking. Mechanical tracking involves instrumenting the moving parts of an interface such as the joints of a moving arm or the rollerball of a mouse. This typically provides accurate and stable readings, but for limited degrees of freedom. For example, the rotation of a normal mouse is not sensed, although there are designs such as the two-ball [MacKenzie 1997] and Rockin'Mouse [Balakrishnan 1997] that overcome this limitation.

As with expected movement, we can consider many different properties of sensed movement including degrees of freedom, range, speed, accuracy, and stability. It is also worth drawing attention to the wide variety of factors that lead to their limitations including, inherent limitations in the technologies, manufacturing cost (budget models may be less instrumented or accurate), environmental conditions (weather, lighting, and interference), computing power (requiring trade-offs between accuracy and responsiveness), and political control (e.g., the accuracy of civilian versus military GPS).

5. EXPECTED VERSUS SENSED

A key point of this article is that the expected and sensed movements of a given interface may only partially overlap with interesting consequences emerging at the boundaries. Figure 1 shows the four possible relationships between expected and sensed movements. Designers should consider what could happen in each of the four areas.

Expected and Sensed. These are natural movements that can be sensed and define the 'normal' operation of the interface. This area has been the dominant focus for previous frameworks and taxonomies that have been mainly oriented towards achieving the best possible match between expected and sensed movements.

Expected But Not Sensed. These are natural physical movements that cannot be sensed by the computer. Consider as examples: taking a PDA equipped with a GPS indoors, tilting it, or moving it more precisely than the GPS can follow; using a laser pointer to point at an object that is outside video tracking range; stepping outside tracking range in an HMD; and, rotating a conventional mouse. The potential problem with such movements is that they may confuse users. For example, an interface may appear to stop working as it moves out of sensing range. The user is performing natural movements but suddenly is getting no response. Several options are open to the designer at this point.

- Improve the sensing by adding additional sensors or sensor technologies so that sensed movement matches the expected movement. This adds additional cost and may not be possible.
- Constrain expected movements so that they match sensed movements, for example, by adding a physical constraint or tether to prevent such movement. This may be appropriate for already jointed or tethered displays where, for example, rotations can be limited, but is less so for wireless interfaces that can be moved freely.
- Change the application to work in a more static or less spatially precise mode when out of sensor range. When no sensor information is available, the display can present static information that is clearly not expected to respond to movement. When less precise sensor data is available, the display can present information in a way that is less suggestive of a precise location or orientation, an approach demonstrated by the Guide tourist information system [Cheverst 2000].
- Communicate the limits of sensed movement to the user, either in software (e.g., messages on the interface to indicate that they are now out of tracking range), or in the physical design of the environment (e.g., clearly delineating the extent of video tracking in a room through visible markers, barriers, and furniture).
- Ignore the issue and assume that users will adapt to the interface (e.g., we soon learn that rotating a mouse has no effect on the cursor).

However, we further suggest that movements that are expected but not sensed can present designers with opportunities as well as problems. They enable the user to reposition the interface without making input to the computer. Perhaps the most familiar example here is lifting a mouse off the edge of a mouse mat so as to reposition it back to the center without affecting the cursor. Moving out of sensing range might be used as a way of deliberately pausing an application.

Related to this, deliberate pauses allow the user to take a rest by disengaging from the interface and entering a state where their physical actions no longer trigger (now unwanted) effects. This is typically not a major concern for the kinds of direct manipulation interface that have been the focus of previous frameworks since with these the user can often simply take their hands off the controls. It is much more of a concern with ubiquitous sensing-based interfaces where users may not be able to disengage, or where it may not be clear how

to disengage. We suggest that rather than aiming for full sensor coverage, designers should consider deliberately building rest spaces into experiences and should make it clear to users how to enter them, especially where those experiences involve prolonged or extreme physical activity such as virtual sports, dance, or other performance.

Such movements also allow physical preparation for, and follow through after, the point of interaction. The principle that a moment of interaction is actually embedded in an entire gesture that determines its timing and feel is familiar from sports (e.g. a golfer's swing). Physical movement around an interface also facilitates expressive interaction during public performance as seen with traditional musical instruments such as pianos [Sudnow 1978]. In discussing electronic instruments, Bowers and Hellstrom [2000] refer to “expressive latitude”—designing interfaces to not sense every movement so as to leave space for physical performance. It seems that, far from being a “dead zone”, expected but not sensed movements may actually provide an important space of opportunities for readjustment, rest, preparation, follow-through, and performance—important features of physical movement.

Sensed and Not Expected. These are movements that can be sensed but not naturally or easily physically carried out. Perhaps the interface cannot easily be moved through all of the available sensing range, or it is being used in a bizarre way or an unanticipated context. For example, GPS can sense when our example PDA is raised several hundreds of meters above the ground (perhaps the user is hang-gliding) or is moving faster than walking speed (perhaps they are running or are in a vehicle). Video tracking can detect a laser pointer that is being held perfectly still for many minutes (perhaps it has been left switched on, resting on a table and the battery is in danger of running out). Electromagnetic trackers can sense full 360-degree rotations of an HMD (perhaps it is in the user's hands instead of on their head). Again, these may be problems or opportunities.

Treating these movements as problems, the designer can extend the expected range of movement, although this may involve a radical physical redesign, or applications can monitor and react to sensor data that is outside the expected norm. Such data could indicate that the sensors are erroneous, that the device is physically broken (e.g., a part has become detached), or that someone is behaving inappropriately with it (e.g., moving more quickly with it than they should) and, as a result, the application might raise a warning or alarm.

Treating the movements as opportunities, especially where the physical movement involved is safe for both the user and the technology, but just not normally expected, the designer might trigger special application functionality that is rarely used or not generally available, for example, resetting or reconfiguring a system or swapping into another mode of operation. Designers might reward users with an equally “odd” experience, for example, revealing mysterious information, or offering strange perspectives. In this way, sensed but nonexpected movements create a space for playful or mysterious uses of interfaces that otherwise behave conventionally, potentially a useful strategy for entertainment, performance, and artistic applications.

Neither Expected or Sensed. These movements cannot be achieved easily and cannot be sensed anyway. In practice, the interface cannot distinguish them




Fig. 2. Expected, sensed, and desired movements.

from movements that are expected but not sensed, and so designers may best treat them in the same way.

As a final note, although our discussion has focused primarily on the range of movement, previous taxonomies of input devices show that designers can also compare expected and sensed movement across other properties of movement including speed, accuracy, and stability. Even if the same basic physical movements are possible, there may be mismatches in other properties (e.g., the user may move an interface more quickly than its sensors can follow, or sensors may be less accurate than physical movements).

6. DESIRED MOVEMENTS

So far, we have discussed the design of interfaces independently of particular applications. For interfaces that are used with a variety of different applications there is a further issue to be considered—how does the range of available expected and sensed movements relate to those that are needed for the application? What is it that we want the application to do? This leads us to the third component of our framework, movements that are desired, or conversely that are possible but not desired. Understanding of desired movements emerges from the kinds of participatory, observational, and inspirational design methods that we reviewed earlier (more than they emerge from analytic frameworks for analyzing input devices).

Desired movements may only partially overlap with expected and/or sensed movements. In other words, there may be movements that are desired for the application but that are not expected and/or sensed (it might be very desirable for users to be able to fly in some 3D entertainment applications), and other movements that are expected and/or sensed but that are not desired. Figure 2, therefore, extends Figure 1 to include desired movements.

Once again, the designer can consider each of the outlying regions as a space of design problems or opportunities. We would like to raise one particularly interesting design strategy, the idea of compensating between movements that are desired, but that are not expected and/or sensed, and those that are expected

and sensed, but not desired. For example, in the Go Go immersive VR interface, the physical movement of extending one's arm right has the effect of extending one's reach in the virtual world beyond its normal range [Poupyrev 1996]. This nonlinear mapping is taking one action that is expected and sensed, but not particularly desired, and making it more desired.

7. APPLYING THE FRAMEWORK

We propose that our framework can help with refining an outlined design concept or sketch towards a more detailed design specification, evaluating how different sensing technologies match proposed application requirements, or identifying detailed potential problems or opportunities with a prototype. It might also support the repurposing of an existing artifact by suggesting unusual or playful ways in which it might be used. Applying the framework of Figure 2 to the design of an interface involves the following steps.

- (1) Analyze expected movements, exploring the impacts of physical form, envisaged users, and environments. Consider for each degree of freedom, the range, speed, accuracy, and stability of expected movements. Spend time imagining scenarios that could result in less-expected movements. Consider which movements are genuinely impossible (rather than just unlikely). This step can draw on existing analytic frameworks but we would also encourage designers to deliberately imagine and discuss the extreme boundaries of unexpected physical movement, envisaging situations in which the interface might be accidentally or even willfully misused.
- (2) Analyze sensed movements by identifying all of the known limitations of the sensing technologies. Again consider the range, speed, accuracy, and stability of sensing for each degree of freedom. Deliberately try to imagine how you could fool sensing systems. This step can also utilize existing analytic frameworks but we emphasize the importance of explicitly identifying the extremes of what can be sensed.
- (3) Analyze desired movements for your application. Apply inspirational design methods to determine how your ideal interface would move if unconstrained by the limitations of the physical world and available sensing technologies.
- (4) Consider each of the different regions of Figure 2, trying to find possible movements to fit each. Consider whether each issue raised represents a problem to be solved or an opportunity to be exploited. In each case, consider the design options outlined previously and whether users will require rest or will perform with an interface.

We now describe how our framework has been applied to the design of three contrasting interfaces: the use of flashlights as interaction devices in underground caves and with public wall displays; the design of a wheeled, mobile 3D display for use outdoors at a museum; and the design of a piece of interactive domestic furniture. Between them, these examples span a variety of physical forms (handheld flashlight, display attached to wheeled base, and a table); include both purpose-designed physical objects (the 3D display) and the augmentation of everyday objects (flashlights and the table); use different

sensing technologies with different degrees of precision (video tracking, a combination of GPS and other sensors, and load sensing); are intended for use in different environments (museums and the home); and are more or less “task oriented”, ranging from the defined task of exploring a historical 3D recreation to a much more open-ended style of engagement with an item of domestic furniture.

It should be noted from the outset that development of the framework and development of the three examples have occurred in parallel in such a way that there has been a flow of ideas from the framework to the designs and back again. This is particularly true of the first two examples that we present which both informed the framework, and were informed by it, across several iterations.

8. EXAMPLE 1: INTERACTIVE FLASHLIGHTS

Our first example focuses on the use of flashlights for interacting with surfaces such as walls and posters. In this case, the sensing technology is based on visual tracking; a video camera captures an image of the surface onto which a user directs the flashlight beam. Image processing software extracts key features of the beam including its position, shape, and extent in real time and uses these to trigger events, for example playing a sound whenever the beam illuminates a designated target area. This is similar in principle to using visually tracked laser pointers to interact with large displays [Olsen 2000; Myers 2002; Davis 2002], although there are significant differences, too; most notably that a flashlight casts a pool rather than a point of light whose size and shape varies according to the user’s position relative to the surface and the kind of flashlight being used. Consequently, flashlight beams can select areas of a surface and can overlap, potentially enabling different kinds of collaboration where several beams are brought together.

Our visually-tracked flashlights technology is targeted at museums, exhibitions, tradeshow, and even classrooms since it involves the use of everyday devices that are familiar, cheap, fun, and safe. We have explored three applications to date. In the first, children used flashlights to control objects in a virtual environment that was projected onto the tent-like immersive interface shown in Figure 3 [Green 2002]. In the second, visitors to Nottingham Castle Museum used flashlights to trigger ghostly voices when exploring a series of underground caves [Ghali 2003]. In the third, flashlights were used to create interactive posters such as the solar system poster, also shown in Figure 3 (right), that replays audio descriptions of each planet as it is illuminated [Ghali 2003].

Experience with these applications was one of the motivating factors for developing our framework, and in the following discussion, we distinguish between those design issues that inspired the framework (i.e., cases where we first encountered an issue which led us to make a generalization) and those that were directly inspired by the framework (i.e., where we then applied the generalization to redesign the technology).

8.1 Summary of Expected, Sensed, and Desired Movements for Flashlights

Expected. Considering expected movements, flashlights can be carried or worn; they can be handheld, head-mounted (e.g., when caving), stand-mounted or




Fig. 3. Using flashlights with the Storytent (left) and with an interactive poster (right).

even vehicle-mounted and vary greatly in size and scale, ranging from small handheld flashlights to large directional spotlights. We expect visitors to point flashlights at a sequence of target features on a surface in a systematic way and then listen to the resulting sounds and sweep across a surface in order to find targets. Flashlights also serve their traditional purpose of illuminating dark spaces so that we expect visitors to point them into dark areas in order to find out what is there or see where they are going. We have observed that it is difficult to hold larger flashlights perfectly still and that visitors will sometimes shine multiple flashlights onto a single surface causing their beams to overlap. Battery life is an issue, with the intensity of a beam becoming noticeably weaker as the batteries begin to run out, and it is expected that a person would switch a flashlight on and off, if only to conserve battery life. We can also expect ambient lighting levels to change in some environments, for example, as people switch lights on and off, or open and close doors.

Less expected, but still possible, movements include not looking where you are pointing a flashlight (but people may do this when talking to one another), or waving a flashlight around very rapidly (children have been observed doing this). We consider it to be a less expected action to defocus a flashlight (where it has a variable focus) or to obscure the beam in some way although this may be unavoidable and people may even deliberately cast shadows. Repeatedly switching the flashlight on and off may be less expected, unless perhaps signaling to others. Also, leaving the flashlight switched on and pointing at one spot for a long time is less expected, although perhaps the user has put it down so that they can work with both hands. It is also less expected to shine the flashlight directly into the video camera (although we observed some people doing this when they first noticed the camera), or to shine it into people's eyes (although this happens when wearing a head-mounted flashlight as people face one another to talk).

Sensed. Turning to sensed movements, our tracking software extracts the position of the centroid and extent of the area of the image of the flashlight

beam on the surface. However, this is only possible when the beam image falls within the camera's field of view which may only cover a part of the total environment, especially when deployed in a large area. This was apparent in our cave experience where, even with three cameras, we could only cover a fraction of the total surface area of the cave and where the boundaries of the interactive parts of the surface were not clearly visible as they would be if interacting with a projection screen. Second, the sensing technology can be fooled by changes in ambient lighting conditions which effectively change the background image (a well known problem for visual tracking technologies in general). Objects that obscure the camera's view (e.g., people walking in front of it) can also confuse the tracking software.

Desired. Finally, we consider general characteristics of desired movements. It should be easy to reach the targets with the flashlight beam and to hold the beam on a target once found. It will be necessary to support a variety of target sizes, shapes, and placements. Groups of visitors may wish to share an experience, and this has implications for how they can position themselves to view the surface and also how they can share audio output. It may also be necessary to deal with potential interference between groups of visitors, for example, shielding them from the sound that is triggered by other groups, avoiding conflicting use of flashlights on a single surface, and generally managing visitor flow.

8.2 Comparison of Expected, Sensed, and Desired Movements for Flashlights

We now compare these expected, sensed, and desired movements. The first two issues below arose from our initial experience and inspired us to develop the framework in the first place. The framework is therefore playing an explanatory role in these cases.

Shining the beam outside tracking range (expected, and maybe desired, but not sensed). Users can be expected to point the beam outside of the camera view. Indeed, this may be desired if they are finding their way around a dark environment. This observation inspired an option in the system to play a background sound whenever the flashlight beam is recognized as being in tracking range but is not currently on a target in order to confirm to the user that they are in the right area. This technique also allows the user to understand when they can safely use a flashlight for other purposes such as general illumination, or signaling without accidentally triggering targets. Early experience with this technique has suggested a further refinement. It is often the case that the extent of the camera viewpoint does not precisely match the relevant visual features of the surface. For example, it was not feasible to position the camera so that its field of view exactly matched the edge of the poster shown in Figure 3 (in this case, the camera could see an area of the wall on either side of the poster). One implication is to constrain the active tracking range to be the subset of the field of view that matches the desired tracking range (e.g., specifying the edge of the poster as being the active region).

Wobbly flashlights (expected, sensed, but not desired). The observed wobble of a flashlight beam as a user tries to hold it on a target may be both expected and sensed, but it is not desired if it produces an annoying effect where the

associated action is repeatedly stopped and retriggered. Our solution has been to refine the mechanism for triggering a target. Early implementations triggered a target whenever the centroid of the beam entered the defined target area based on the idea that the user is pointing with the beam, rather like they would with a laser pointer. Later implementations use a revised mechanism that measures the proportion of the target that is illuminated by the area of the beam. The target is triggered when this exceeds a critical threshold. This approach assumes that the user is illuminating the target rather than pointing at it which we suggest is more in line with the expected use of flashlights, and has proved to be more accommodating to wobble.

Our next two issues arose from reapplying the framework back to the design of the technology and provide examples of how it can help to generate new design possibilities.

Detecting a very static beam (sensed, but maybe not expected or desired). The tracking system can potentially detect when a beam is held precisely still in one spot for an extended period of time (say several minutes). Given the tendency to wobble when handheld, this would suggest that the flashlight has been put down on a surface while still switched on and might indicate a potential problem (perhaps it has been left behind and forgotten and perhaps the batteries will run down) which might, in turn, generate a warning. Again, this might also be seen as an opportunity. Perhaps users could leave flashlights in position for a while in order to achieve special effects, for example, metaphorically “drilling” into a surface to reveal new content.

Using a defocused flashlight (not expected, potentially not sensed, and not desired). Initially the idea of defocusing the beam seems problematic as it can no longer be tracked. This problem can be solved by physical constraint, using flashlights that can't be refocused, or by jamming the refocusing mechanism. However, this might also be an opportunity. It might sometimes be desired to be able to use a flashlight for general illumination without triggering any targets. This could be achieved by deliberately defocusing the flashlight so that the beam is no longer visible to the sensing system. This demonstrates an interesting approach to disengaging from a sensing system without leaving the physical sensing area.

9. EXAMPLE 2: THE AUGURSCOPE II

Our second example, the Augurscope II, is a portable mixed-reality display for viewing 3D historical reconstructions when outdoors. Users wheel it around a physical site and rotate and tilt its display in order to view a 3D model as it would appear from different physical vantage points. Our design, shown in Figure 4, responds to issues that were raised by public trials of an earlier prototype, particularly limited mobility [Koleva 2002]. It supports two modes of use: stand-mounted in which the top is attached to the wheeled-base by a mounting that allows rotation, tilting, and various other adjustments; and hand-held in which the top can be detached and moved more freely. The design of the stand features two handles, one attached to the base and one to the rotating mount for the top, so that users can rotate the top while pushing the display along.




Fig. 4. The Augurscope II, stand-mounted and hand-held.

9.1 Summary of Expected, Sensed, and Desired Movements for the Augurscope

Expected. As designers, we expected individuals or small groups to wheel our interface across relatively flat terrain at slow walking pace, occasionally stopping to rotate, and tilt it in order to explore a particular viewpoint. More experienced users may be able to study the display as they wheel it, combining pushing, rotating, and tilting movements. Some may detach the top and use it in hand-held mode in which case we would expect relatively cautious movements and probably no long, sustained poses due to its weight. They might also lay the top flat on the ground, especially if tired, or try to take it indoors. In terms of less-expected movements, we would not expect users to run quickly with the device, to spin the top around rapidly, to turn it over and over in their hands, to move through solid walls, to take it underground, or to fly above the ground.

As well as clarifying our general expectations of use and potential misuse, we carried out a more systematic analysis of how the proposed physical form of our design would afford and constrain expected movements. Table I summarizes expected movements for the six degrees of freedom of possible movement: translate sideways ($\uparrow X$), raise and lower ($\uparrow Y$), push and pull forwards and backwards ($\uparrow Z$), tilt forwards and backwards (θX), rotate around vertical axis (θY), and tilt sideways (θZ).

As noted by Foley et al. [1984], it is possible to combine two or more degrees of freedom into a single movement. Rotating (θY) while tilting (θX) the top is expected as this can easily be done with one handle. Using two hands to rotate and tilt the top while pushing the base ($\uparrow Z$) might be expected. Rotating and tilting the top while raising and lowering it (in unlocked mode) is possible, but is much harder and hence less expected. Raising and lowering while pushing is possible, but only with great difficulty.

Table I. Details of Expected Movements of the Augurscope II

DoF	Range of Movement	Accuracy, Speed, Stability
$\uparrow X$	Not expected (the user would have to lift it off of the ground and carry sideways)	
$\uparrow Y$	Expected height adjustment involves unlocking adjusting screw, raising or lowering, and locking again (range of 70cm). More extreme movement not expected. Less expected to raise and lower dynamically while unlocked.	Millimeter accuracy. Takes seconds in locking mode, but is stable. Quicker but less stable when unlocked
$\uparrow Z$	Expected to push forwards and pull backwards. Unconstrained in range (unless by obstacles).	cm accuracy. Expected at walking pace.
θX	Freely and indefinitely tilt the top forwards and backwards. Not expected to rotate through many loops.	Better than 1° accuracy. Takes seconds. Stable.
θY	Can freely and indefinitely rotate the top unit on the stand. May not be expected to rotate through many loops.	Better than 1° accuracy. Takes seconds. Stable.
θZ	Not physically possible on when stand mounted	

Using the top in hand-held mode affects this analysis in several ways. $\uparrow X$ and θZ are now possible. The range of $\uparrow Y$ is extended down to ground level and up to the maximum height to which a user can lift the screen and still see it. The speed of movement can also be increased. However, stability will be reduced because the user has to hold the display in position rather than it resting on a supporting stand. It is also easier to combine degrees of freedom in hand-held mode.

Sensed. Turning now to sensed movements, the Augurscope II uses two sensing technologies. A Trimble GPS receiver provides global position and a Honeywell HMR3000 digital compass measures rotation and tilt. Both are integrated into the top unit which communicates wirelessly with the base using 802.11b networking. Table II summarizes the range, accuracy, and delay associated with sensing each degree of freedom of movement. Several aspects need emphasizing. First, the range of GPS sensing extends a long way above ground, but not below it. Second, the digital compass can only sense up to 45° of tilt downwards or upwards. Third, the top unit can stray out of communications range of the base unit when in hand-held mode.

Desired. When considering desired movements, we focus on the example application of viewing a 3D recreation of a historic castle when exploring its present day site, the same application that was used for testing the first prototype [Koleva 2002]. The interesting detail of the 3D model is in buildings as well as in an underground cave section. Conversely, there is no significant detail on the surface of the ground or high in the sky. The model also covers a restricted geographical area. It is therefore desired to move around this area and to look at objects at building height. It is less desired to look at the ground or high into the sky or to move outside of the geographical area that is modeled. At some points it would be desired to be able to fly under the ground. Our experience shows that it is also desired for users to be able to see a bird's eye view of such models, both to be able to orientate themselves and for the novel perspective that this brings.

Table II. Details of Sensed Movements for the Augurscope II

DoF	Sensed by	Sensed range	Accuracy	Delay
$\uparrow X, \uparrow Y, \uparrow Z$	GPS receiver	not indoors, in black spots, underground or extreme latitudes	cm to meters varies	1 hz.
θX θZ	compass	+45° to -45°	< 1°	20 hz
θY	compass	360°	< 1°	20 hz

9.2 Comparison of Expected, Sensed, and Desired Movements for the Augurscope

Extreme tilting (expected, but not sensed or desired). The User may tilt the interface beyond its sensed range of 45° up and down, especially in hand-held mode, in which case it will appear to suddenly stop responding. However, given that there is no interesting detail on the floor or sky of the 3D model, this is not an especially desired movement.

Flying (sensed and desired, but not expected). The GPS can sense the desired movement of flying into the air (although looking down for a bird's eye view is not sensed as noted above), but the user cannot lift the device off the ground.

These two issues have been addressed together through a single extension to the design. We have altered the mapping between the tilt of the top and the movement of the virtual camera in the 3D model. The tilt has been exaggerated so that for every sensed degree of tilt, two degrees are rendered. Additionally between 20° to 45°, the camera pulls upwards. At 45°, the limit of sensed movement, the virtual camera has tilted to 90° (i.e., is looking straight down) and has raised several tens of meters into the air to give a bird's eye view as shown in Figure 5. The view remains static beyond 45°. This provides an example of the compensation strategy described earlier in which a sensed, expected, but not especially desired movement (20° to 45° tilt) is remapped to support desired, but not expected and/or sensed movements.

Exploring caves (desired, but not expected or sensed). The desired act of exploring the caves in the 3D model requires taking the interface underground. However, there is no suitable physical access and GPS will not work there. Our framework suggests that we might exploit a similar compensation strategy as previously mentioned, using extreme upward tilting to drop the viewpoint below ground level into a virtual cave. However, two further issues have to be addressed. First, the caves only exist at limited locations under the castle grounds, and so it is appropriate to trigger this mechanism only when the Augurscope is above a cave, requiring an additional indication that there are caves below. Second, it would not be appropriate for the viewpoint to remain fixed in the cave ceiling once underground. A better solution might be for an extreme upward tilt to trigger a downward navigation of the viewpoint, after which the Augurscope could be rotated as normal to explore a panoramic view of a cave (but not translated as there is no GPS) and for a subsequent extreme downward tilt to take the viewpoint back up above ground.

Running with the Augurscope II (sensed, but not expected or desired). Running while looking at the screen is not expected and not desired, and is debatably




Fig. 5. Comparing normal with exaggerated tilt.

sensed (the GPS can follow, although with some latency). One possibility here is to replace the view of the 3D model with a warning message. In terms of the framework, this deliberately makes the device less desired when it is not being used sensibly.

Moving outside the model (expected and sensed, but not desired). The user might wheel the augurscope outside the castle grounds and hence outside of the scope of the virtual model. There is no benefit associated with this since there is no virtual model to explore, but there is certainly a risk (the augurscope might be stolen!). Our proposal is to raise an alarm and to encourage the user to take the device back into range.

Areas of poor GPS (desired and expected, but not sensed). We are concerned about the effects of variable GPS (inside buildings, undercover, or by a wall). One proposed solution is to switch to a mode in which there is a dialogue with the user to confirm their location, perhaps using a touch-screen. Another is to present information in a less precise way, for example, turning off the exact mapping between GPS and the 3D viewpoint and instead exploring a panorama from a predetermined viewpoint.

10. EXAMPLE 3: THE DRIFT TABLE

Our third example is the Drift Table, an interactive coffee table (Figure 6) that allows its owners to take an aerial trip over Great Britain. The Drift Table is an example of a load-sensing interactive surface [Schmidt 2002]. It uses four load sensors to determine the total weight and center of gravity of the objects on its surface. These two measures control a viewpoint that looks down onto a series of aerial photographs that are drawn from a database that covers the whole of



Fig. 6. The Drift Table.

England and Wales at 25cm resolution (kindly provided by Getmapping.com). These are stitched together so that the viewer appears to be smoothly and seamlessly traveling across the entire country and are viewed on a small circular display that is sunk into the center of the table. The direction of movement is given by the direction of the center of gravity relative to the center of the table, and the speed of movement is a function of the total weight on the table multiplied by the distance of the center of gravity from the center of the table. Finally, the more weight there is on the table, the lower its apparent altitude.

The Drift Table is primarily a coffee table, but one that is also intended to provide its users with an engaging, stimulating, and provoking experience that is deliberately designed to be open-ended rather than focused on achieving a particular task. The Drift Table experience is also open-ended in terms of its duration. At one extreme, users might move objects on its surface to see immediate changes in the viewpoint. At the other, they might leave objects and return to the table after hours or maybe even days to see changes.

10.1 Summary of Expected, Sensed, and Desired Movements for the Drift Table

An initial design was in place and the construction of a prototype underway when our framework was introduced to the project through a series of brainstorming sessions in which designers tried to envisage a wide variety of scenarios of use. The Drift Table has since been completed and tested.

Expected. It is expected to place everyday objects such as books, magazines, and drinks on the Drift Table, to take them off again, and to move them around on its surface, including translating, rotating, and stacking them. It is expected movement to clean the table and to occasionally move it to clean around or underneath it. Moving the table might involve rolling (it has castors) or lifting it. Objects with changing or shifting weight such as candles or plants might also find their way onto the table. People can be expected to lean on the table (e.g., to look into the display) or even write or sit on it.

Less expected, but certainly possible, are standing and jumping on the table. Pets such as cats might walk over the table from time to time. People might also lift one or two legs (perhaps for a quick clean underneath) or jog the table when passing by or using a vacuum cleaner. Less expected also is using the table as a support when hammering nails or putting a plank or trestle between two tables when decorating (although this might be more expected with less “designer” tables). It is also not expected to turn the table over or to stand it on its side. Another unlikely possibility is placing the table in a moving environment where it might be subject to external forces or might change its orientation with respect to the outside world, on a yacht for example.

Sensed. There are two key factors to be considered with regard to the Drift Table’s sensed movements: the characteristics of each load sensor, and the way in which the four sensors combine to measure overall weight and center of gravity.

The load sensors are industrial precision load cells based on resistive strain gauges. The physical characteristics of the load cells used in the design for the Drift Table provide a response for a force from 0 to 500N. If the load cells are used in a horizontal scale, this is the gravitational force created by a weight of about 50Kg. Between 500N and 1000N, the response is not linear anymore but there is no damage to the load cell. If a force greater than 1000N is applied, it might damage the load cell permanently. It takes on the order of 500ms for load measurements to reach a stable reading because the surface of the table wobbles and then settles down due to the materials used, its construction, and the force with which load is applied (e.g., throwing a book onto the table causes a greater and longer period of instability than gently placing it).

Each sensor measures one degree of freedom—the load acting on it from above. However, the use of four sensors allows two degrees of freedom to be derived—the overall weight and the eccentricity of the center of gravity. If only one object is moved at a time, it is possible to infer its movement across the surface, for example a finger can be used as a drawing tool [Schmidt 2002] although moving multiple objects at a time would fool such an interpretation. It might be possible to identify which objects have been added to, or removed from, the table if each has a unique weight, although again, a person could easily be fooled if different objects turn up with the same weight or objects change their weight over time (e.g., a burning candle).

The surface of the table has a weight, and so the whole sensing system has to be calibrated to output an effective weight of zero whenever just the surface is present. This raises the possibility of sensing “negative weight”, for example, if the surface is somehow lifted. Calibration is done automatically when switching on the table. For a period of 5 seconds, the average of each sensor is taken and stored as a base weight for this sensor. All further calculations are then relative to this base weight. If during this initialization process, there is additional weight on the table that is then taken off, it would lead to the empty table registering negative weight.

Desired. Our initial design was guided by a general principle that using the Drift Table should be analogous to riding in a hotair balloon. Movements that

break this metaphor (e.g., zooming and displaying extraneous information) are generally undesired.

The range of movement needs to cover the whole area of Britain and altitude also needs to vary, but within a fixed range, so that the image remains interestingly visible at the highest altitude and yet readable without being grainy at the lowest.

In terms of speed, movement should be responsive enough to quickly and visibly react to a sudden change in load and yet should move slowly enough to make a gradual journey across the country over the course of several hours or maybe even days. The maximum speed of movement will be about 150 km/hr so that an appreciable amount of time will be required to traverse Britain even at its narrowest point.

Although one of the primary functions of the Drift Table is to get lost over the British landscape, it should be possible to “find” oneself from time to time or frustration might ensue. Thus it is desired to know the current (virtual) location of the table, checked against a map if necessary. In addition, it is desired that the orientation of the image corresponds to the orientation of the user’s environment so that heading north from a starting point over one’s home will cause the image to move in the appropriate direction. Finally, in order to avoid being in less desired parts of the country for days at a time, it is desired to reset the table’s location to be over the user’s home on rare occasions—the only discontinuous movement deemed to be acceptable.

10.2 Comparison of Expected, Sensed, and Desired Movements for the Drift Table

We now describe how a comparison of these expected, sensed, and desired movements through a series of face to face meetings and email exchanges among the design team inspired new design ideas. The overall goal of keeping the table open-ended rather than too task specific, combined with the desire to achieve a strong and quite minimal aesthetic, has meant that many of these new ideas will not be included in the final prototype. However, we discuss them here as they demonstrate the use of the framework to raise new possibilities and might also be relevant to other load-sensing surfaces.

Keeping the display oriented to North when the table is turned (desired and expected, but not sensed). A discussion of the possibility of the table being moved (expected) raised the question of what should happen to the orientation of the maps being displayed. The design team felt that they should always orient correctly to the North whatever the orientation of the table and hence its display. However, this would require being able to sense the rotation of the table (currently not sensed). As a result, the team extended the sensing capability of the table by adding a digital compass to its hardware.

Resetting the viewpoint (desired, but not sensed or expected). The initial design required a way of resetting the viewpoint to a default position. The application of our framework revealed the possibility of generating negative weight by pulling upwards on the load sensors, and it was felt that this might be an appropriate way of triggering a viewpoint reset. Given the boxed-in design

of the nearly completed prototype, there was no easy way to achieve this (i.e., it is not currently expected) and a reset button, hidden away near the floor, was provided. However, an option for the future is to build a new housing so that the surface stretches a few centimeters away from the base whenever the table is lifted, allowing negative weight to be generated.

Journey objects (sensed, expected, and desired). Our discussions of objects that change weight or position over time led to a new idea: journey objects that take the viewpoint on a predictable journey. For example, a burning candle loses weight in a predictable way. This appears to the Drift Table as a gradual reduction in overall weight and a shift in the center of weight away from the candle back towards the center of the table. Users could “program” the table to undertake a journey of a given distance in a given direction by placing candles of specific weights at specific locations on the table. A “rotating compass candelabra” might even allow the user to place a candle at a set distance from the center in a particular compass direction. Different candles could then be manufactured to travel set distances.

Limit speed in order to limit weight (expected, sensed, and not desired). It was decided to introduce a maximum speed limit for the movement of the viewpoint, beyond which adding more weight to the table wouldn’t make any difference. The aim here was to deliberately discourage less expected behaviors such as standing or jumping on the table, or loading it heavily in order to see how fast it might go. This is an example of deliberately constraining sensed movement in order to avoid less expected physical behavior. If this strategy fails, then more drastic measures might be called for such as removing the images altogether (making the table useless during nonexpected use).

Reducing sensitivity to sudden weight change (expected, sensed, but not desired). A related issue concerns managing acceleration. While it is desired for the table to be visibly responsive to the movement of objects, sudden and large changes in velocity (due to heavy weight being placed on the table) can cause problems. In particular, a software caching mechanism needs to predict which aerial photographs to preload from the database into the rendering software in order to ensure a smooth viewing experience. This prediction becomes difficult when velocity changes rapidly. Our solution here is to treat any sensed change in the total weight and center of gravity as defining a target velocity to which the table gradually accelerates over a configurable period of time so that some change is immediately noticeable, but that the caching mechanism has time to adapt to it. This effectively reduces the (apparent) sensitivity of the table (i.e., increases its sensing delay) in order to achieve a tradeoff between two potentially conflicting desired movements.

11. SUMMARY AND REFLECTION

Driven by four trends in interfaces—the growth of sensor-based interaction, the diversification of physical forms, increasing mobility, and a focus on playful, engaging and creative applications—we believe that interface designers will increasingly have to wrestle with the complex problem of matching physical form to the capabilities of sensors and the shifting requirements of applications.

There is already a range of existing frameworks and methods available to support them. Existing taxonomies of input devices support detailed analysis of how the physical form of a device matches its sensing capabilities. Existing design methods such as participatory design, ethnography, and emerging inspirational methods can be used to generate new design ideas. However, we have argued that none of these existing frameworks and methods is in itself sufficient to address the entire problem. First, they have not been focused on the specific challenges raised by working with imprecise sensor technologies and augmented everyday artifacts. Second, we believe that successful design needs to combine both analytic and inspirational perspectives.

We have introduced a new framework that encourages designers to tackle this problem head-on by analyzing and comparing expected, sensed, and desired movements. Our framework focuses on the boundaries between these, drawing on analytic and inspirational approaches, and treating mismatches as opportunities as well as problems. We have applied our framework to three example interfaces.

For the interactive flashlights, the framework was used to help understand initial experiences. Indeed, development of the framework was in part inspired by these experiences. We then reapplied the framework back to the technology, which generated further insights in areas such as defocusing a flashlight in order to disengage from the sensing system; the potential for setting up a very static beam to trigger special effects; and a reconsideration of the relationship between the expected extent of users' movements with a flashlight, the extent of sensed movement in terms of the video camera's view, and the extent of desirable movement in terms of the interactive content (e.g., a poster).

For the Augurscope II, the framework was introduced relatively late in the design process to refine and extend the design of a second generation prototype. In this example, which is closest to designing a new physical input device, we combined an analytic approach in reasoning about expected movements with a deliberate attempt to question assumptions about users' likely actions. Applying the framework led to several new design ideas including our extended tilt mechanism.

For the Drift Table, the framework was employed earlier in the design process to help inspire design ideas in moving from a general concept to a first concrete prototype. Our analysis here focused on brainstorming a wide variety of potential uses of a table and considering how these matched the capabilities of the load-sensing technology as well as the more open-ended goals of the "application". Several new insights emerged from this discussion including the use of journey objects, the need to maintain orientation to the North, and introducing constraints on the speed and acceleration of apparent movement.

Across all three examples, the framework has helped us clarify design trade-offs, identify and explain likely problems with interaction, and has sometimes helped inspire new interaction possibilities. We believe that these examples show that the framework has the potential to support both the analysis of designs and the generation of new ideas. However, it is not a silver bullet. There is no guarantee that the resulting ideas are good ones. Applying the framework to the Drift Table generated many new possibilities, most of which were rejected

by the design team as over-complicating the design. Indeed, one of the main uses of the framework was to strongly focus discussion on the question of what was desired.

Future work will focus on applying the framework to a broader set of designs in order to better understand its role within the design process. This will involve packaging the framework in a way that makes it easier for others to apply. We are intrigued by the possibility that focusing on mismatches between expected, sensed, and desired movements can lead to new design opportunities and believe that this will be a useful design tactic as interfaces become more graceful, sporting, artistic, and playful. With this in mind, we are particularly interested in applying the framework to emerging applications that involve extreme physical interaction such as sports and performance.

REFERENCES

- BALAKRISHNAN, R., BAUDEL, T., KURTENBACH, G., AND FITZMAURICE, G. 1997. The Rockin'Mouse: Integral 3D manipulation on a plane. In *Proceedings of the Conference on Human Factors in Computing Systems*. Atlanta, GA. 311–318.
- BELLOTTI, V., BACK, M., EDWARDS, W. K., GRINTER, R. E., HENDERSON, A., AND LOPES, C. 2002. Making sense of sensing systems: Five questions for designers and researchers. In *Conference on Human Factors in Computing Systems*. Minneapolis, MN. 415–422.
- BLESER, T. W. AND SIBERT, J. L. 1990. Toto: A tool for selecting interaction techniques. In *Proceedings of the 3rd annual ACM SIGGRAPH Symposium on User interface Software and Technology*. Snowbird, UT. 135–142.
- BOUCHER, A., GAVER, W., PENNINGTON, S., AND WALKER, B. 2003. Workbook one: Ideas, scenarios and proposals for the home, available at www.interaction.rca.ac.uk/equator/papers/workbook1.pdf.
- BOWERS, J. AND HELSTROM, S. 2000. Simple interfaces to complex sound in improvised electronic music. In *Proceedings of the Conference on Human Factors in Computing Systems, Supplementary Proceedings*. The Hague, Netherlands. 125–126.
- BOWERS, J. AND PYCOCK, J. 1994. Talking through design: Requirements and resistance in cooperative prototyping. In *Proceedings of the Conference on Human Factors in Computing Systems, Supplementary Proceedings*. New York, NY. 299–305.
- BUXTON, W. 1983. Lexical and pragmatic considerations of input structures. *Comput. Graph.* 17, 1, 31–37.
- CARD, S. K., MACKINLAY, J. D., AND ROBERTSON, G. G. 1991. A morphological analysis of the design space of input devices. *ACM Trans. Inform. Sys.* 9, 2, 99–122.
- CHEVERST, K., DAVIES, N., MITCHELL, K., FRIDAY, A., AND EFSTRATIOU C. 2000. Developing a context-aware electronic tourist guide: Some issues and experiences. In *Proceedings of the Conference on Human Factors in Computing Systems*. The Hague, Netherlands. 17–24.
- CRUZ-NEIRA, C., SANDIN, D. J., DEFANT, T. A., KENYON, R. V., AND HART, J. C. 1992. The cave-audio visual experience virtual environment. *Comm. ACM* 35, 6, 65–72.
- DAHLEY, A., WISNESKI, C., AND ISHII, H. 1998. Water lamp and pinwheels: Ambient projection of digital information into architectural space. In *Proceedings of the Conference on Human Factors in Computing Systems*. Los Angeles, CA. 269–270.
- DAVIS, J. AND CHEN, X. 2002. LumiPoint: Multi-user location-based interaction on large tiled displays. *Elsevier Science* 23, 5.
- EHN, P. AND KYNG, M. Cardboard computers: Mocking-it-up or hands-on the future. In *Design at Work: Cooperative Design of Computer Systems*. J. Greenbaum and M. Kyng, Eds. Lawrence Erlbaum Ass., Hillsdale, CA, 169–197.
- FOLEY, J. D., WALLACE, V. L., AND CHAN, P. 1984. The human factors of computer graphics interaction techniques. *IEEE Comput. Graph. Appl.* 4, 11, 13–48.
- GAVER, W. AND MARTIN, H. 2000. Alternatives: Exploring information appliances through conceptual design proposals. In *Proceedings of the Conference on Human Factors in Computing Systems*. The Hague, Netherlands. 209–216.

- GAVER, W., BEAVER, J., AND BENFORD S. 2003. Ambiguity as a resource for design. In *Proceedings of the Conference on Human Factors in Computing Systems*. For Lauderdale, FL. 233–240.
- GAVER, W. 1992. The affordances of media spaces. In *Proceedings of the 1992 ACM Conference on Computer-Supported Cooperative Work*. Toronto, Canada. 1, 7–24.
- GAVER, W. AND DUNNE, A. 1999. Projected realities: Projected design for cultural effect. In *Proceedings of the Conference on Human Factors in Computing Systems*. Pittsburgh, PA. 600–607.
- GHALI, A., BOUMI, S., BENFORD, S., GREEN, J., AND PRIDMORE, T. 2003. Visually tracked flashlights as interaction devices. In *Proceedings of Interact*. Zurich, Switzerland. 487–494.
- GIBSON, J. J. 1977. The Theory of Affordances. In *Perceiving, Acting and Knowing*. R. E. Sahw & J. Bransford, Eds. Lawrence Erlbaum Ass. Hillsdale, CA. 67–82.
- GREEN, J., SCHNÄDELBACH, H., KOLEVA, B., BENFORD S., PRIDMORE, T., MEDINA, K., HARRIS, E., AND SMITH, H. 2002. Camping in the digital wilderness: Tents and flashlights as interfaces to virtual worlds. In *Proceedings of the Conference on Human Factors in Computing Systems Extended Abstracts*. Minneapolis, MN. 780–781.
- GRIFFITH, N. AND FERNSTROM, M. 1998. Litefoot—a floor space for recording dance and controlling media. In *Proceedings of the International Computer Music Conference*. Ann Arbor, MI. 475–481.
- HUGHES, J. A., RANDALL, D., AND SHAPIRO, D. 1992. Faltering from ethnography to design. In *Proceedings of the 1992 ACM Conference on Computer-Supported Cooperative Work*. Toronto, Canada. 115–122.
- ISHII, H. AND ULLMER, B. 1997. Tangible bits: Towards seamless interfaces between people, bits and atoms. In *Proceedings of the Conference on Human Factors in Computing Systems*. Atlanta, GA. 234–241.
- ISHII, H., ORBANES, J., CHUN, B., AND PARADISO, J. 1999. PingPongPlus: Design of an athletic-tangible interface for computer-supported cooperative play. In *Proceedings of the Conference on Human Factors in Computing Systems*. Pittsburgh, PA. 394–401.
- JACOB, R. J. K., SIBERT, L. E., MCFARLANE, D. C., AND MULLEN, M. P. 1994. Integrality and separability of input devices. *ACM Trans. Comput.-Hum. Interact.* 1, 1, 3–26.
- JOHN, B. AND KIERAS, D. 1996. The GOMS family of user interface analysis techniques: Compare and contrast. *ACM Trans. Comput.-Hum. Interact.* 3, 4, 320–351.
- KLEMMER, S. R., THOMSEN, M., PHELPS-GOODMAN, E., LEE, R., AND LANDAY, J. A. 2002. Where do web sites come from?: Capturing and interacting with design theory. In *Proceedings of the Conference on Human Factors in Computing Systems*. Minneapolis, MN. 1–8.
- KOLEVA, B., SCHNÄDELBACH, H., FLINTHAM, FRASER M., IZADI S., CHANDLER P., FOSTER M., BENFORD S., GREENHALGH C. AND RODDEN T. 2002. The Augurscope: A mixed reality interface for outdoors. In *Proceedings of the Conference on Human Factors in Computing Systems*. Minneapolis, MN. 9–16.
- MACKENZIE, I. S., SOUKOREFF, R. W., AND PAL, C. 1997. A two-ball mouse affords three degrees of freedom. In *Proceedings of the Conference on Human Factors in Computing Systems Conference Companion*. New York, NY. 303–304.
- MACKINLAY, J. D., CARD, S. K., AND ROBERTSON, G. G. 1990. A semantic analysis of the design space of input devices. *Hum.-Comput. Interact.* 5, 23, 145–190.
- MUELLER, F., AGAMANOLIS, S., AND PICARD, R. 2003. Exertion interfaces: Sports over a distance for social bonding and fun, In *Proceedings of the Conference on Human Factors in Computing Systems*. For Lauderdale, FL. 651–668.
- MYERS, B., BHATNAGER, R., NICHOLS, J., PECK, C., KING D., MILLER, R., AND LONG, C. 2002. Interacting at a distance: Measuring the performance of laser pointers and other devices. In *Proceedings of the Conference on Human Factors in Computing Systems*. Minneapolis, MN. 33–40.
- NORMAN, D. A. 1988. *The Psychology of Everyday Things*. Basic Books, New York, NY.
- NORMAN, D. A. 1999. Affordances, conventions and design. *Interact.* 6, 3, 38–43.
- OLSEN, D. AND NIELSEN, T. 2001. Laser pointer interaction. In *Proceedings of the Conference on Human Factors in Computing Systems*. Seattle, WA. 17–22.
- PAULOS, E. AND CANNY, J. 1998. PProp: Personal roving presence. In *Proceedings of the Conference on Human Factors in Computing Systems*. Los Angeles, CA. 296–303.
- POUPYREV, I., BILLINGHURST, M., WEGHORST, S., AND ICHIKAWA, T. 1996. Go go interaction technique: Non-linear mapping for direct manipulation in VR. In *Proceedings of the 10th Annual ACM Symposium on User Interface Software and Technology*. Banff, Canada. 79–80.

- REKIMOTO, J. AND SCIAMMARELLA, E. 2000. ToolStone: Effective use of the physical manipulation vocabularies of input devices. In *Proceedings of the 13th Annual ACM Symposium on User Interface Software and Technology*. San Diego, CA. 109–117.
- SCHMIDT, A., STROHBACH, M., VAN LEERHOVEN, K., FRIDAY, A., AND GELLERSEN, H. 2002. Context acquisition based on load sensing. In *Proceedings of the International Conference on Ubiquitous Computing*. Goteborg, Sweden. 333–350.
- SUDNOW, D. 1978. *Ways of the Hand: The Organization of Improvised Conduct*. Harvard University Press, Cambridge, MA.
- ULLMER, B., ISHII, H., AND GLAS, D. 1998. mediaBlocks: Physical containers, transports, and controls for online media. In *Proceedings of the 25th Annual Conference on Computer Graphics and Interactive Techniques*. Orlando, FL. 379–386.
- UNDERKOFFLER, J. AND ISHII, H. 1999. Urp: A luminous-tangible workbench for urban planning and design. In *Proceedings of the Conference on Human Factors in Computing Systems*. Pittsburgh, PA. 386–393.

Received February 2003; revised August 2003, April 2004; accepted February 2004 by Shumin Zhai and Victoria Bellotti