Supporting Subtlety with Deceptive Devices and Illusory Interactions

Fraser Anderson[1], Tovi Grossman[1], Daniel Wigdor[2], George Fitzmaurice[1]

[1]Autodesk Research
{firstname.lastname}@autodesk.com

[2]University of Toronto
daniel@dgp.toronto.edu

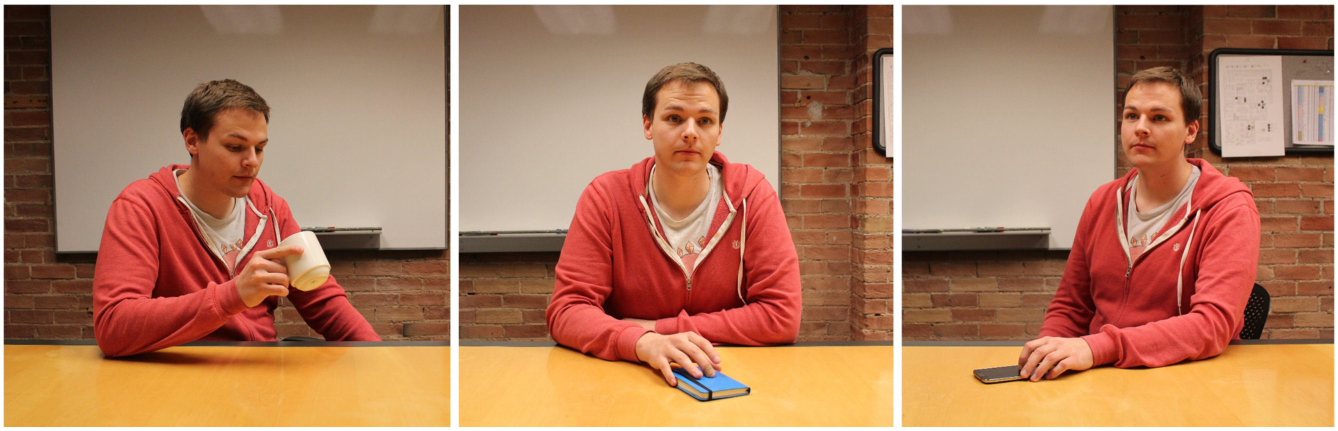

Figure 1: Examples of subtle interactions achieved through the use of the presented gimmicks. From left to right: reading text messages, composing an email, and checking the time.

ABSTRACT

Mobile devices offer constant connectivity to the world, which can negatively affect in-person interaction. Current approaches to minimizing the social disruption and improving the subtlety of interactions tend to focus on the development of inconspicuous devices that provide basic input or output. This paper presents a more general approach to subtle interaction and demonstrates how a number of principles from magic can be leveraged to improve subtlety. It also presents a framework that can be used to classify subtle interfaces along with a modular set of novel interfaces that fit within this framework. Lastly, the paper presents a new evaluation paradigm specifically designed to assess the subtlety of interactions. This paradigm is used to compare traditional approaches to our new subtle approaches. We find our new approaches are over five times more subtle than traditional interactions, even when participants are aware of the technologies being used.

In submission to ACM CHI 2015. Submission # 1086

INTRODUCTION

Mobile devices enable always-available access to information, supplying a perpetual connection to the world outside our immediate surroundings. While compelling, it is often inappropriate to use one's phone or other electronic device, be it for privacy or social reasons [22, 27]. In many situations, users wishing to access their mobile device resort to simplistic methods, such as hiding their phone under a desk or turning down the screen brightness, to avoid detection. These methods have limited effectiveness as observers often notice such behaviors.

Our work focuses on the development of technologies that enable *subtle interaction* (Figure 1). We define subtle interaction as *providing input to, or receiving output from, systems without being observed*. The primary goal is the development of a suite of technologies that enable users to leverage always-available computing without compromising privacy or social interaction.

Previous commercial and research efforts have investigated approaches to subtle interaction. These efforts typically focus on miniaturization, or decreasing the cognitive load of interactions [3, 27]. Further, these efforts have largely produced technologies, which once known to an observer, are no longer subtle – that is, they can be readily noticed and observed. In contrast, our work focuses on developing technologies and techniques, which even if known, would not telegraph their use to others.

Efforts to advance subtle interfaces have also been hampered somewhat by a lack of systematic identification of the space of subtle interaction or guiding principles for the design of the devices. To help us in developing these, we focused on the related domain of magic and illusion. Magicians have been developing and perfecting means of deceiving observers for centuries, amassing a vast quantity of knowledge. Their effects rely on established methods and principles that can be combined in various ways to conceal their true intent. While specific magic tricks are of little value to designers of subtle interfaces, the principles upon which they are built can be leveraged for subtle interactions.

This paper makes several contributions to the area of subtle interfaces. First, we present a survey of the techniques that users currently employ in the wild to achieve subtle interaction. Next, we provide design guidelines for subtle interfaces that are derived from established magic principles. We also provide a framework for the classification of subtle interfaces. Leveraging the guidelines and framework, we developed a system of subtle interfaces comprised of input and output devices. Lastly, we present an evaluation method that can be used to quantify the subtlety of interactions.

RELATED WORK

The present work is inspired by research from intersection of magic and human-computer interaction (HCI) as well as within wearable and ubiquitous computing.

Illusion and Deception in Human-Computer Interaction

Stage magic and HCI have many parallels [25], such as a need for consistency, the use of metaphors, and smoothness throughout the interactions. Busk et al. separated magical knowledge into three layers of *effects*, *methods*, and *principles* that can be used as a framework for magic-inspired interfaces [7]. Marshall, Benford, and Pridmore described a technology-based magic trick that applied the principles of misdirecting attention and creating false expectations to deceive the user [19]. This prior work focuses on high-level similarities between HCI and magic, rather than adapting magical knowledge for subtle interactions.

Adar, Tan, and Teevan discuss the uses and morality of deception in computer interfaces [1]. They argue that deception can be a useful feature of design, as in placebo buttons, which give the operator a sense of control and decrease user frustration [28]. In contrast, our work focuses on systems that are designed to aid users in deceiving others, rather than systems that are designed to deceive the user.

Wearable and Ubiquitous Computing

Subtle interactions for wearable interfaces typically focus on actions with low cognitive load, so they can be performed without a perceived loss of attention, or small and fast movements, so they can be performed before observers notice them [4]. Ashbrook introduced microinteractions for mobile interfaces [5], which can be initiated and completed in under four seconds. Holleis et al. evaluated the placement of fabric-based capacitive sensors, identifying locations that are quick and easy to access, but not necessarily subtle [14].

A number of approaches provide subtle, low-bandwidth input through novel hardware designs. ShoeSense made use of a depth camera mounted on a shoe to sense gestures made by users [6], while ShoeSoleSense made use of pressure sensors and vibrotactile feedback to provide input and output [20]. An approach by Sumitomo et al. used a strain sensor placed around the waist to allow users to input sequences by contracting their abdomen [24]. Electromyogram sensors have been used to enable users to interact hands-free, or without visibly moving [10, 23]. Nenya was a magnetic ring that could be tracked by a wrist-worn sensor to allow users to subtly access menu items [3]. In other work, covert 'duress' signals were built into authentication schemes, allowing users to signal that the authentication is not genuine using a small modification of their input sequence [18]. Clawson et al. described a scenario where devices could be used to subtly send a signal from one party to the other [8].

With respect to subtle output, several approaches have made use of small LEDs to provide low-bandwidth information channels without being visible to observers [11, 13]. Toney et al. described a suit that integrated LEDs, vibrotactile feedback and an interactive watch [26]. Similarly, Lee developed wearable tactile displays that relayed alerts to the user [15]. To disguise output, Worddit and Timesify provide tools that modify how Reddit and gossip websites were rendered so they appeared to be Microsoft Word or the New York Times webpage, respectively [29, 30].

While these approaches contribute to the development of novel devices, the use of the individual devices is not sufficient for truly subtle interaction. Just as there is more to a great magician than their tricks, there is also more to subtle interaction than the devices. In contrast to prior work, we address the greater context of subtle interaction, and provide guidelines derived from magic principles. We also explicitly evaluate the subtlety of our approaches. Lastly, we describe how a system of devices can work together to improve the overall subtlety of the interactions.

SURVEY OF EXISTING PRACTICES

To increase our understanding of how technology is used subtly today, a crowd-sourced, online survey was conducted. The survey was intended to gauge interest in subtle interaction, as well as understand how users currently utilize their technology when they are attempting to be subtle. It also provided information on common use cases, enabling the design of an evaluation that reflects realistic settings.

Survey Instrument

The survey consisted of 19 questions with Likert-scale, binary, and free form responses regarding aspects of subtle interaction. The tasks people currently attempt to perform subtly were probed, for example, as were the steps they take to conceal their interactions, and their emotions towards people who use technology subtly. The questions focused on both social and professional context and was posted on classified sites (e.g., Craigslist), as well as Mechanical Turk.

Responses

In total, 227 responses were received. The age of respondents ranged from 14-69 years (M = 30; 124 male), with respondents located in North America (143), Asia (61), Europe (20), and South America (3).

When asked which tasks were frequently performed, checking the time, texting, and taking photos were the most common (Figure 2). Most respondents, i.e., 86%, reported that they attempt to use their mobile devices undetected during meetings for tasks unrelated to the meeting, while 94% reported doing the same during social events. A less frequent, but commonly reported behavior was pretending to receive a call or text message to leave a situation (73%).

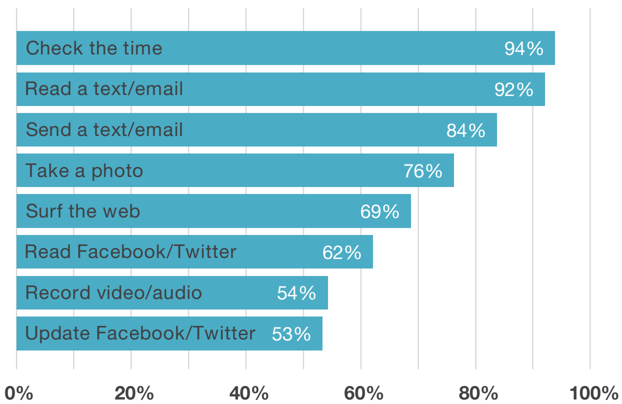

Figure 2: Proportion of respondents who reported performing each task subtly.

Nearly all respondents (94%) reported being ‘caught’ using technology where they were not supposed to (e.g., where interaction was prohibited by laws, employers, or social norms). When asked about the primary reason for attempting to using technology discreetly, the most frequent response was to avoid violating social norms (35%) followed by prohibition by their employers (25%) and prohibition by location (e.g., at a museum, 18%). Almost all respondents (94%) responded that they would use technology to help them interact with their devices more subtly.

When asked about the most common situations where subtle interaction was used, three general themes were evident. Respondents reported using technology subtly to *distract themselves* from undesirable situations (e.g., “during moments when things get boring or awkward”). Another common purpose was to *record events or information* (e.g., “record a video from a meeting … I needed evidence”). Finally, several participants used subtle interactions to *keep others informed* of their current situation (e.g., “I used [my] mobile phone to message my wife about the situation”).

Respondents also reported a variety of techniques they adopted to facilitate subtle interaction. Respondents reported *hiding the device* out of sight (e.g., “check the phone under the desk”) or *leave the area* where they may be caught (*“*At work, I go to the washroom”). Others *disguise their interactions* (e.g., “I pretend I'm trying to get reception … in actuality I am snapping a photo”).

The responses to the survey, indicated that users attempt to use technology subtly in many scenarios, and that they desire technology that provides them with more discretion. Additionally, the responses demonstrate that users are willing to put conscious effort into concealing their interactions from others to prevent others from being disturbed or upset that they are using technology.

DESIGNING FOR SUBTLETY

As users desire technology that can disguise their actions, we looked to magic to inform the design of these technologies. As magic requires expertise in deception and illusion, we consulted with professional magicians, as well as the magic literature, to explore how magic could be leveraged for the design of subtle interactions. As there are many aspects of magic (i.e., methods, principles and effects) that could be utilized, we initally present guidelines that are derived from a subset of magic principles: *user customization, modularity, simulation and dissimulation*, *separating cause and effect* and *user training*. These principles were chosen because of their broad applicability to many subtle interface designs.

P1) User Customization

> *we don’t need any more new tricks … we do need more tricks fitting the specific personalities.*
>
> D. Fitzkee [12]

For a convincing performance, magicians customize and tailor their tricks to their routine and character. Similarly, users that desire subtle interaction should be able to customize their devices. A construction worker may look out of place wearing a large ring that conceals a vibrotactile motor. However, if the vibrotactile motor was contained within a more generic form-factor, it could be integrated into a pair of safety glasses or other object that is consistent with the user’s character. More importantly, devices that are not customizable will be easily recognized once they are widely available. If a manufacturer develops a smart-ring with a distinctive look, its functionality will be known and observers will be weary of the users’ interactions with it.

P2) Modularity

> *If the magician performs the same trick twice for the same audience, there is an increased chance that the audience will identify the method.*
>
> S. Macknik et al. [16]

Using the same method to misdirect the audience’s attention can lead to revelation of the magician’s methods. Likewise, repeated use of a subtle device is likely to raise suspicion from observers. By varying the location, modality, and method of input, users of subtle interfaces can improve their ability to interact undetected. For instance, if the user has a joystick that can be concealed in a pocket, up a sleeve, or within a lapel, they can move the joystick between locations to prevent observers from becoming suspicious of the continued manipulation of their pocket, for example. Thus, a modular system that provides redundancies allows users to vary their behavior and select the interaction technique most subtle for a given transaction. This, however, adds complexity to interactions with the system

P3) Simulation and Dissimulation

> *Where simulation is disguising a thing to make it appear similar to something else, dissimulation is the act of making something appear to be dissimilar to what it truly is.*
>
> D. Fitzkee [12]

If a magician is attempting to convince the audience that a box is empty, he must interact with the box and behave as if it were empty and weightless. Similarly, users who wish to interact subtly must conceal the true nature of their devices and integrate their interactions with their natural behaviours. If a sensor is concealed within a baseball cap, then the interactions should be extensions of natural manipulations of the cap, such as adjusting the angle of the brim. Additionally, if users leverage existing technology to interact subtly, then efforts should be made to interact overtly with the device when it is appropriate to do so. This will lead observers to believe their device is normal and unremarkable which will aid in concealing its subtle functionality.

P4) Separating Cause and Effect

> *... introduce delays between the method behind a trick and its effect, preventing you from causally linking the two ... the magician must separate the method from the magical effect.*
>
> S. Macknik & S. Martinez-Conde [17]

Just as magicians introduce delays to misdirect the audience, designers of subtle interfaces can introduce temporal and spatial separations between an action and resultant effect to misdirect observers. If a user adjusts his watchband and the screen immediately changes, observers may correlate the watchband manipulation with the change in output. However, if the user adjusts his watchband and thirty seconds later his phone displays a notification, observers would have difficulty associating the input, which was unremarkable at the time, with the output.

This separation is counter to traditional user-interface guidelines, which promote a tight coupling between cause and effect to minimize the gulf of evaluation [21]. By implementing this separation, interactions can become more subtle, but the user interface may become more cumbersome and usability may suffer as a result.

P5) User Training

> *I promise never to perform any illusion for any non-magician without first practicing the effect until I can perform it well enough to maintain the illusion of magic.*
>
> Magician's oath [17]

Just as practice is a critical component in the performance of magic, it is important with subtle interaction as well. Clumsy, untrained movements can cause erroneous input and be spotted by observers quickly. To circumvent this, users should be aware of the necessity of training and be provided with tools to support that training. Guides can help train users on how to efficiently provide input, such that behaviours become automatic and require less of the users' attention.

DEVICE FRAMEWORK

While there are several devices that enable subtle interaction, there is no framework that can describe or categorize them in a useful manner that will facilitate future developments. We have identified five factors that describe the functionality of an interaction technique within the context of subtle interfaces: the *type of device*, the *customizability*, the *device visibility*, the *interaction observability*, and the *task transparency*.

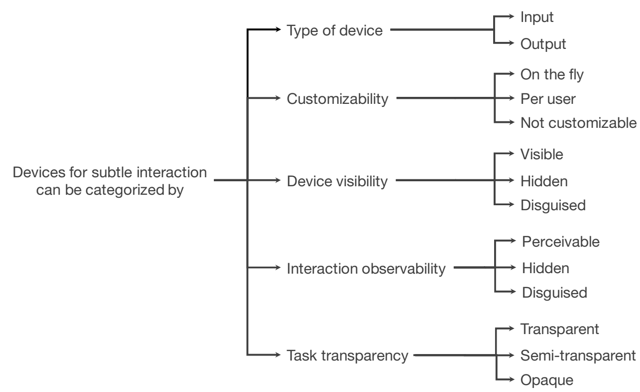

Figure 3: Overview of the framework for subtle devices.

The *type of device* references whether the interface provides *input* or *output* capabilities.

The *customizability* refers to the degree that the interface can be tailored by the end user. Techniques may be re-configured *on-the-fly* (e.g., modifiable during a meeting), for each *user* (e.g., to match their personality or existing technology), or *not customizable*. Increased customizability increases the potential for more subtle interaction.

Device visibility represents whether or not the hardware associated with the device is *visible*, *hidden*, or *disguised.* Visible devices, e.g., a traditional smartphone, are the least subtle and must be used sparingly. Hidden devices are those that can be used while out of sight of the observers, e.g., within a pocket or a shoe. These are the most subtle but also have limited bandwidth and can be difficult for the user to operate. Disguised devices are those that are visible to the observer, but provide functionality beyond their appearance.

Interaction observability refers to how the observers may perceive the user's interaction. This could be *perceivable, hidden*, or *disguised.* Perceivable interactions are those that are noticeable by observers, but may enable subtle interaction if performed quickly enough. Hidden interactions are not able to be detected by observers, and as with hidden devices, often provide maximal subtlety with minimal bandwidth. Disguised interactions may be observable but could be identified as a normal behaviour.

Task transparency refers to how readily observers can detect the task that is being performed by the user. This is important, as some tasks may be acceptable for a given scenario and others may not. Thus, by obscuring the task, the user may be afforded additional subtlety. The task may be *transparent,* in the case of a traditional watch with a single function, or it may be *opaque,* in the case of a smart phone that possesses a wide range of functionality. The task may also be *semi-transparent* if observers can identify a subset of potential tasks (e.g., they can observe that the user is receiving a notification, but cannot determine the type).

SAMPLE DECEPTIVE DEVICES

To explore the framework, several representative input and output devices were constructed (Figure 4). Used collectively, these devices demonstrate the utility of the design principle of Modularity (**P2**).

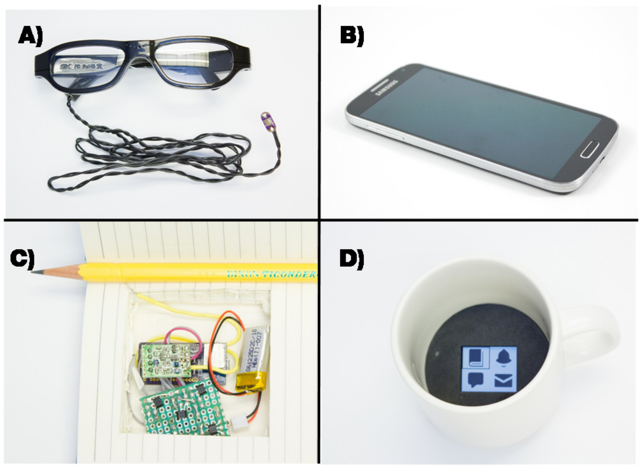

Figure 4: Subtle devices. A) Super Sneaky Spyglasses, B) Samsung Galaxy S5 running Phony Phone, C) Magput in a book, and D) Lil LCD in a mug

Input Devices

Super Sneaky Spyglasses

A small camera mounted inside a pair of glasses (Figure 4A) can allow the user to take photographs or record video. This *disguised device* can be triggered through *hidden interactions*, i.e., orienting their head then pressing a button that is tethered to the frame of the glasses and concealed within a sleeve or a pocket of the user. The task is *transparent* and not obfuscated at all, as knowledgeable observers will be aware of the sole function of the device. The current form factor is *not customizable*. The success of this device relies on the user's use of *dissimulation* (**P3**) when using the device, as they must treat the glasses as they would unmodified glasses - they must not touch them excessively or prolong eye contact while taking photos.

To implement the Super Sneaky Spyglasses commercial spyglasses[1] were modified to include a tethered button rather than the more obvious button mounted on their frame.

Phony Phone

Phony Phone is a mobile application that renders a black screen on the cell phone (Figure 4B), providing the appearance of a sleeping phone while still allowing input to be recorded and processed. Users can interact with the device by swiping and tapping on the black screen, *disguising* their interactions as idle fidgeting. Phony Phone is a *disguised* device, the task is *semi-transparent*, as a knowing observer could determine that the user was providing input but could not identify the specific task. Phony Phone can provide a full range of input supported by modern touch screen phones (touches, swipes, gestural input) and can be *customizable* by the end user (i.e., the application can be deployed on their phone and does not require a specific, identifiable piece of hardware). This device enables subtle interaction through the spatial *separation of input and output* (**P4**); the touches on the screen have no visible effect on the device itself.

To implement Phony Phone, an Android application was developed and was run on a Samsung Galaxy S5.

Magput

Five linear Hall Effect sensors mounted in an 'x' shape can be used to sense the location of a magnet over the grid enabling continuous 2D or gestural input in a small space (Figure 4C). Magput can be placed within many common objects to provide an input device that is easily disguisable. For example, the device can be placed within a book with a magnet inside of a pen, or it could be mounted inside a jacket pocket with a magnet placed beneath an adhesive bandage. In both cases, the sensors and magnet are *disguised*, and the users' interaction is *disguised* as idle doodling or casual fidgeting with the device. As with Phony Phone, the task is *semi-transparent* as expert observers could determine that the user is providing input to the system but could not determine the nature of the input. Magput can be *customized on-the-fly* (**P1**) by moving it between various locations on the user's body, or amongst their accessories. To subtly interact using EdgeWrite gestures, users must *practice* (**P5**) prior to use. To aid in this, a small, physical training guide was developed that can be placed over top of the sensor to provide the user with tactile cues to guide them to the correct locations. While this is a simple device, training systems that are more complex could be implemented (e.g., [2]).

In our implementation, the sensors are connected to a Femtoduino BLE[2], which relays sensor data via an RF link to an Arduino Uno, which relays the data to a PC which is tethered to the mobile phone via Bluetooth.

Output Devices

Lil LCD

Small, modular displays can be leveraged to provide repositionable, high-bandwidth output that could be concealed within innocuous objects. These displays can be hidden inside a coffee cup form factor (Figure 4D), for example, and allow for the subtle viewing of information during a meeting. The device, as well as the interactions with it, are *disguised*, as the user can place it out of sight and casually glance at it when appropriate. If noticed, the task is *semi-transparent*, observers are aware that there is some information being obtained, but are not able to determine what that information is from a distance. Lil LCD is

[1] http://www.newegg.com/Product/Product.aspx?Item=9SIA34P0ZN7835

[2] www.femtoduino.com

customizable *on-the-fly*, as the display can be relocated with relative ease. By *separating the output* (**P4**) from the interaction, observers will have difficulty identifying when the user is interacting.

In our implementation, the module is comprised of a waterproof Sony Smartwatch 2 display that is tethered via Bluetooth to a mobile phone. To aid in concealment, custom 3D printed 'shells' were printed and affixed to the module.

Numerical Sonification

To disguise numerical data, an audio alert can be used to encode numbers in a manner that could be mistaken for a message or calendar notification. Four notes of increasing pitch can be used to signify the numbers 30, 15, 5, and 1 respectively and can be summed together to derive the encoded number. For instance, the note sequence 'A C E E G G' could inform the user that it is 57 minutes past the hour (30 + 15 + 5 + 5 + 1 + 1 = 57). The device used in this approach is *customizable per-user* (**P1**) and *disguised,* as it is deployed on the user's own mobile phone. Users can further customize it by changing the timbre of the audio tones to be more consistent with their existing notifications. The interactions with the device are *disguised* and the task is *semi-transparent* as observers may be aware that the tones convey information, but would not know if the information was the time, a sports score, or the number of unread emails in the user's inbox. To provide misdirection, the audio is *separated temporally* (**P4**) from the input that triggers it by adding a delay of five seconds.

To develop and test Numerical Sonification a number of audio files were created and the playback functionality was integrated into the Phony Phone application.

MANAGING ILLUSIONS

As the context of use rapidly changes with wearable and mobile interactions, it is imperative that users have their technology tailored to their immediate situation. The use of a large number of highly configurable modules and devices has the potential to overwhelm new users with the volume of potential combinations of devices. To address this, a configuration system (Figure 5) was developed to allow users to assign actions and data to their various devices.

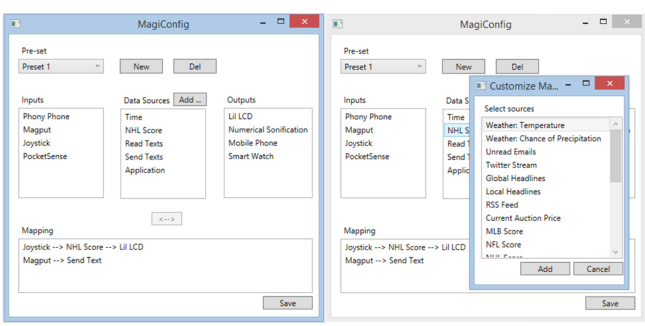

Figure 5: Configuration interface showing mappings between input devices, data, and output devices.

The user begins the configuration process by creating a new preset. A preset represents a set of mappings that are suitable for the given context. For instance, the user may have one preset for casual meetings or social events where some use of technology is appropriate, and another for formal meetings where interaction is restricted. In the current system, there is no automatic detection of context, so users must manually select the active configuration.

With each preset, users can create mappings between input triggers, a list of pre-defined data streams (e.g., incoming text messages, sports scores, or the weather), and output devices (e.g., the Lil LCD or Numerical Sonification). The functional capabilities of each device (e.g., what type of data it is capable of generating or displaying) are encoded within the program logic and simplify the user experience. As a user selects a particular data stream, only those devices capable of triggering or displaying the data stream become active. While this interface does not expose the complete functionality of the modular system, it provides enough flexibility to satisfy many common use cases.

EVALUATION

While previous efforts have been made to develop subtle interactions, there is no accepted methodology for evaluating the subtlety of an interaction. In this section, we propose a novel evaluation methodology centered on the real-world context of a round-table meeting (Figure 6).

Figure 6: Sample frame from the staged meeting videos.

We recorded video of a staged five-person meeting from the perspective of one of the meeting attendees (Figure 6). During the meeting, which was divided into a series of two-minute clips, the four attendees (actors) within the view of the camera would subtly interact with technology in some way. After all recordings had been completed, the crowd-sourced participants viewed the resultant videos and assessed when the attendees were interacting with technology. In contrast to the prior evaluation methods [9], this method better reflects a real-world use case of subtle interaction and allows for a wider range of hypotheses to be tested. Additionally, our survey identified workplace meetings as one a common scenario where subtle interaction is desired.

Recording Procedure

Two blocks of ten two minute clips were recorded, resulting in twenty total video clips. The content of each of the two-minute clips was a game of 'twenty questions' played by all five attendees. During eight of the ten clips in each block, one of the on-camera attendees would interact with technology, while the remaining two rounds did not contain any interaction. The actors were instructed to behave naturally and to perform the tasks as they would in a regular meeting (e.g., relatively discreetly). For each block each actor was assigned two interaction techniques and was given time to familiarize themselves with their operation and to practice using them subtly. Between blocks, actors were re-assigned interaction techniques to ensure that the measured subtlety was not due to an actor's performance. The meetings were recorded at 1080p, via a camera positioned at eye-level, which captured all attendees within the frame (Figure 6).

Subtle Tasks and Devices

Four tasks were analyzed in the evaluation: checking the time (*check time*), reading a text (*read text*), sending a text (*send text*), and taking a picture (*picture*). These tasks are representative of the range of input and output tasks most commonly performed subtly today (Figure 2). Each of the four tasks was completed using currently available technology (*traditional*) as well as using interactions designed to support subtle interaction (*subtle*).

Using *traditional* interaction, checking the time was accomplished using a wristwatch, while reading and sending texts, as well as taking photos was accomplished using mobile phones (i.e., Samsung Galaxy Nexus, iPhone 4, and Nokia Lumia 1020). Unlock passcodes were removed from the phones to minimize the amount of interaction needed.

To evaluate the *subtle* interactions, checking the time was accomplished using numerical sonification in combination with Phony Phone – the attendee tapped the black screen and five seconds later, a tone encoding the time played. Reading a text was accomplished using a modular display affixed to the inside of a coffee cup. Sending a text was done using EdgeWrite-based input with Magput. The sensors were placed inside a hollowed-out notebook and activated by a magnet that was embedded in the back of a pencil or underneath an adhesive bandage on the finger. Taking a picture was accomplished using the Super Sneaky Spyglasses with the tethered button resting in the lap.

Participants

Sixty participants, age 21-55 years (M = 34 years; 27 male) were recruited via crowd-sourcing to watch the videos and determine when an actor was interacting with a device.

Design

A three-factor, mixed study design was used with a between-subjects factor of Condition (i.e., *no knowledge*, *module knowledge*, and *customization knowledge*) and within-subject factors of Task (i.e., *check time, read time, send text, picture*) and InteractionType (i.e., *subtle*, *traditional*).

For each condition, a different introductory video was shown. For participants in the *no knowledge* condition, an introductory video explained their task and informed them that they should be monitoring for technology use in addition to paying attention to the meeting content. Participants in the *module knowledge* condition were also shown the various subtle modules that would be used (e.g., the Magput, the Lil LCD, etc.). Participants in the *customization knowledge* condition were also shown how those modules would be used within the experiment (e.g., embedded in a book, placed at the bottom of a cup, etc).

Each participant watched a single block of ten videos, which included two instances of each task (*subtle* and *traditional*), along with two videos with no interactions. The presentation order of the videos was counterbalanced across participants using a Latin-square design.

Procedure

After watching each video, participants indicated which meeting attendant, if any, interacted with technology during the clip. Participants were also asked to specify the device, the interaction, as well as the task that was being completed by the identified attendant. Participants were asked to guess the item being described by the 20 questions. This ensured that they were not focusing their attention solely on detecting technology use, which better reflects a real-world scenario.

RESULTS

Responses to the post-video questionnaire were collected, and where appropriate, free-form text was manually coded.

Measures

Three measures were captured to analyze the subtlety of each interaction technique: *device subtlety*, *interaction subtlety*, and *task subtlety*. Task subtlety is defined as the proportion of trials in which the participant correctly identifies the task that was being performed subtly. Device subtlety is defined as the proportion of trials in which the device that was being used was correctly identified. Interaction subtlety is defined as the proportion of trials in which the observer correctly identified the attendee performing the interaction technique.

The results aggregated across task and condition (Figure 7) demonstrated that participants had trouble correctly identifying the task in comparison to identification of the device or the interactions. This is because it is difficult to observe the content of the phone's screen from the viewpoint of another meeting attendee so participants would have to guess if the attendee was reading texts, sending texts, or just surfing the web. In contrast, the device and the interaction were identified more readily, as they are recognizable from the vantage point of the recorded video. In addition, the identification rates for both device and interaction were nearly identical, as observers have to see both the device and interaction to realize that the attendee is using technology.

While participants had difficulty identifying the task being performed across all interaction techniques, they were more accurate in determining the direction of the interaction (i.e., input or output) with the subtle devices. As input and output were separated for many of the techniques, they could, for example, identify movements of the pencil on the book as being used to input data, but could not identify whether the attendee was texting or triggering another device.

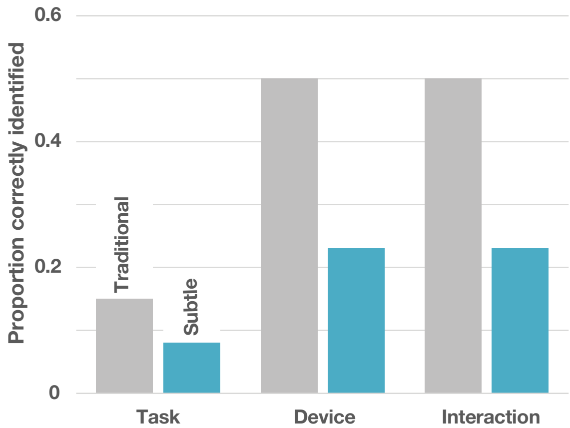

Figure 7: Correct identifications of task, device, and interactions by device type, aggregated over all knowledge conditions.

In the analyses that follow, we consider an interaction to be correctly identified if the participant correctly identified either the device or the interaction.

Analysis

A loglinear analysis was conducted in SPSS on the cell frequencies for the correct identification of the interaction using the factors of Condition (i.e., *no knowledge, module knowledge, customization knowledge*), Task (i.e., *check time, read text*, *send text*, *picture*), and InteractionType (i.e., *traditional, subtle*). The analysis demonstrated significant main effects of Condition ($\chi^2(2) = 53.99$; $p < 0.001$), Task ($\chi^2(3) = 22.63$; $p < 0.001$) and InteractionType ($\chi^2(1) = 43.85$; $p < 0.001$), as well as significant interactions between InteractionType and Condition ($\chi^2(2) = 14.72$; $p < 0.001$) and InteractionType and Task ($\chi^2(3) = 34.17$; $p < 0.001$). Due to the interaction effects, the main effects of Condition, Task and InteractionType were not analyzed further.

To examine the interaction between InteractionType and Condition (Figure 8), three separate Chi-square tests were performed for each of the three conditions. The analyses found significant differences between the noticeability of traditional and subtle interaction for the no knowledge and module knowledge conditions ($\chi^2(1) = 25.31$; $p < 0.001$; $\chi^2(1) = 25.58$; $p < 0.001$ respectively). There was, however, no significant difference between the traditional and subtle interaction types within the customization knowledge condition ($\chi^2(1) = 1.26$; $p = 0.26$). It thus appears that the amount of knowledge a participant was given influenced how well they noticed the interactions. Those participants with no knowledge, or knowledge of only the modules, were not able to reliably identify the subtle interactions, with participants only noticing 8-9% of interactions. In contrast, those same participants were able to notice 44% of traditional interactions, representing a five-fold increase in subtlety with the new approaches. It is worth noting that the correct identifications of the subtle interactions in the no knowledge condition stem from Numerical Sonification.

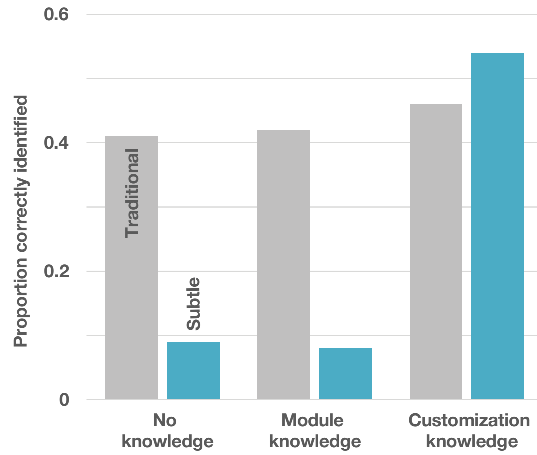

Figure 8: Correct identifications by knowledge condition and interaction type, aggregated over task.

To examine the interaction between InteractionType and Task (Figure 9) separate Chi-square tests were performed. Task was found to significantly affect the identification of traditional interactions, with checking the time being identified significantly less often than the other tasks ($\chi^2(3) = 25.73$; $p < 0.001$). Task also affected noticeability of subtle interactions, with taking a photo being significantly less noticeable than the other tasks ($\chi^2(3) = 19.19$; $p < 0.001$).

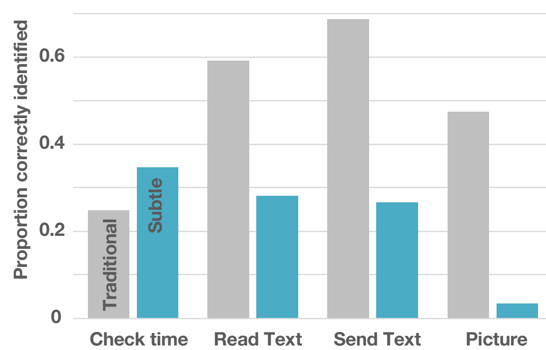

Figure 9: Correct identifications by interaction type and task, aggregated over all knowledge conditions.

The interactions with the watch to check the time were rarely noticed, as the meeting attendees could glance down at the watch briefly. Similarly, taking photos with the spyglasses required only a quick press of a button. In contrast, checking text messages required prolonged manipulation of the device in a relatively overt manner. With the exception of checking the time, the subtle techniques were less conspicuous than the traditional techniques. This is because the numerical sonification technique required touching a mobile device followed by an audible tone. In our evaluation paradigm, this stood out to observers, as none of the other technology generated audio notifications. However, in real-world use we expect this approach to be less noticeable.

DISCUSSION

The results of our study demonstrate that our proposed methodology can quantify the subtlety of various forms of interaction, and distinguish these techniques with regards to their level of subtlety. The results also show that the majority of the new approaches to subtle interaction are successful in deceiving observers.

Knowledge of the Observers

The results for *module knowledge* are of particular interest as they demonstrate that end-user customization is a powerful tool for subtle interactions. Despite being shown all of the devices, participants were not able to imagine using them in the manner that they were used in the video, allowing the meeting attendees to interact without being noticed. In contrast, once participants were aware of the customization, they were able to identify the subtle interactions at nearly the same rate as the traditional interactions, with the exception of the spyglasses, which remained nearly imperceptible.

Participants in the *customization knowledge* condition identified more of the traditional interactions then participants in the *module* or *no knowledge* conditions. This could be the result of the participants being more attentive as they had a greater proportion of trials where they could successfully identify an interaction. It is possible that if modular subtle devices become prevalent, that the set of possible customizations may also become well known, and therefore easier for observers to spot. However, if this should happen, then the set of possible customizations will be so great as to make observers suspicious of everyday actions.

Supporting Devices

Situational awareness and the knowledge of observer's attention is essential for successful subtle interaction, thus we explored support devices for subtle interaction. While the role of input and output devices is clear, these support devices are used in conjunction with other devices to increase the resulting subtlety of interactions. Support devices are unexplored in the literature, and here we provide two examples which may be used to drive future research. The first alerts the users when observers may be perceiving their interaction, the second provides smoother transitions into and out of social situations.

If a user is engaged with their device, their attention may be diverted from their environment. If an observer approaches the user from behind, the user may be unaware and may continue to interact with their device while the observer monitors their actions. To prevent this, a proximity sensor is placed on the user's shoulder facing away from the user and provides pulsing vibrotactile feedback if an observer approaches the user. As the approaching observer moves closer the frequency of pulsing increases, giving the user an awareness of the speed and distance of the observer.

As the use of technology diverts the user's attention from the current situation, it may be difficult to follow and rejoin conversations when the interaction is complete. To enable a more fluid transition in and out of social contexts, our buffered audio technique records audio in a short buffer, and replays it at 1.5x speed through a bone-conduction speaker on demand. This technique provides contextual awareness of the situation, and helps conceal the user's diverted attention using a deceptive device by bringing them back "up-to-speed" with the current conversation.

Ethics of Deception

Adar et al. [1] discuss well-meaning uses of deception for interaction. What is being proposed with subtle interaction, however, could be perceived as ill meaning. Deceiving observers into believing you are cognitively present while you attend to remote information can have consequences. We take the stance that users will interact with their devices regardless of the subtlety of the interaction. By supporting interactions that are more secretive, we increase the probability that such interaction goes unnoticed and the observers remain un-offended. In this light, advances in subtle interaction could be beneficial as users could check the time, for example, without seeming as if they are bored with the current situation.

Evaluation Limitations

While the study captures the visual and audio aspects of interactions as well as changes in the user's attention, it is not able to capture whether or not the behavior is consistent with the user's personality. To capture this, a study would have to intentionally deceive participants that were known to the experimenter. This type of study is logistically difficult and the results would be subjective. Additionally, the study did not assess the utility of using a number of devices rather than a single device. With a number of devices, we suspect the overall subtlety of the interaction will increase as users can interact with the most subtle interface given their situation.

Future work

This work should serve as a basis for the development of many future subtle interaction techniques that leverage the presented magical principles to improve the subtlety of the interfaces. We hope that researchers and practitioners seek out magical principles and methods beyond what is presented here to develop novel interactions. Truly subtle interaction relies on the existence of a multitude of customizable modules to be successful, so the development of new approaches is critical to its success.

There is also great potential for subtle interaction in non-mobile scenarios. For instance, encoding dynamic information within presentations (e.g., augmenting on-screen information such as slide numbers, or by encoding information within artificial errors) could provide valuable data to the presenter or knowing members of the audience.

CONCLUSION

The design of subtle interfaces extends beyond the design of imperceptible input and output devices. Driven by a desire from users for more subtle interactions, we have leveraged knowledge and techniques from the domain of magic to inform the design of subtle interfaces. We developed a framework to classify the various devices and techniques that support subtle interaction, and have built representative interfaces that sample areas of the framework. Lastly, we provided an evaluation paradigm by which the subtlety of interactions can be explored and understood.

REFERENCES

1. Adar, E., Tan, D.S., & Teevan, J. (2013). Benevolent Deception in Human Computer Interaction. *ACM CHI*, 1863-1872.
2. Anderson, F., Grossman, T., Matejka, J., & Fitzmaurice, G. (2013). YouMove: Enhancing Movement Training with an Augmented Reality Mirror. *ACM UIST*, 311-320.
3. Ashbrook, D., Baudisch, P., & White, S. (2011). Nenya: Subtle and Eyes-free Mobile Input with a Magnetically-tracked Finger Ring. *ACM CHI*, 2043-2046.
4. Ashbrook, D.L., Clawson, J.R., Lyons, K., Starner, T.E., & Patel, N. (2008). Quickdraw: The Impact of Mobility and On-body Placement on Device Access Time. *ACM CHI*, 219-222.
5. Ashbrook, D.L. (2010). Enabling mobile microinteractions. https://smartech.gatech.edu/ handle/1853/33986.
6. Bailly, G., Müller, J., Rohs, M., Wigdor, D., & Kratz, S. (2012) ShoeSense: A New Perspective on Gestural Interaction and Wearable Applications. *ACM CHI*, 1239-1248.
7. Busk, N.K., Forstholm, J.S., Lessél, L.S., & Sørensen, S.G. (2012). Introducing Magical Experiences in UX. *SIDeR*, 21-24.
8. Clawson, J., Patel, N., & Starner, T. (2010). Digital kick in the shin: on-body communication tools for couples trapped in face-to-face group conversations. *Workshop on Ensembles of On-Body Devices, ACM MobileHCI*.
9. Costanza, E., Inverso, S.A., Allen, R., & Maes, P. (2007). Intimate Interfaces in Action: Assessing the Usability and Subtlety of Emg-based Motionless Gestures. *ACM CHI,* 819-828.
10. Costanza, E., Inverso, S.A., & Allen, R. (2005). Toward Subtle Intimate Interfaces for Mobile Devices Using an EMG Controller. *ACM CHI*, 481-489.
11. Costanza, E., Inverso, S.A., Pavlov, E., Allen, R., & Maes, P. (2006). Eye-q: Eyeglass Peripheral Display for Subtle Intimate Notifications. *ACM MobileHCI*, 211-218.
12. Fitzkee, D. (1945). *Magic by Misdirection*. Saint Raphael House.
13. Hansson, R. & Ljungstrand, P. (2000). The Reminder Bracelet: Subtle Notification Cues for Mobile Devices. *ACM CHI EA*, 323-324.
14. Holleis, P., Otto, F., Hussmann, H., & Schmidt, A. (2007). Keystroke-level model for advanced mobile phone interaction. *ACM CHI*, 1505-1514.
15. Lee, S. & Starner, T. (2010). BuzzWear: Alert Perception in Wearable Tactile Displays on the Wrist. *ACM CHI*, 433-442.
16. Macknik, S.L., King, M., Randi, J., Robbins, A., Teller, Thompson, J & Martinez-Conde S. (2008). Attention and awareness in stage magic: turning tricks into research. *Nature Reviews Neuroscience 9*, 11, 871-879.
17. Macknik, S.L., Martinez-Conde, S., & Blakeslee, S. (2010). *Sleights of mind: What the neuroscience of magic reveals about our everyday deceptions*. Macmillan.
18. Marques, D., Carrico, L., & Guerreiro, T. (2014). Duress Unlock: Using Covert Signals to Restrict Access to Mobile Devices. *ACM CHI EA.*
19. Marshall, J., Benford, S., & Pridmore, T. (2010). Deception and Magic in Collaborative Interaction. *ACM CHI*, 567-576.
20. Matthies, D.J.C., Müller, F., Anthes, C., & Kranzlmüller, D. (2013). ShoeSoleSense: Proof of Concept for a Wearable Foot Interface for Virtual and Real Environments. *ACM VRST*, 93-96.
21. Norman, D.A. (2002). *The design of everyday things*. Basic books.
22. Profita, H.P. (2011). Social acceptability of wearable technology use in public: an exploration of the societal perceptions of a gesture-based mobile textile interface. Thesis, *Georgia Institute of Technology*.
23. Saponas, T.S., Tan, D.S., Morris, D., Balakrishnan, R., Turner, J., & Landay, J.A. (2009). Enabling Always-available Input with Muscle-computer Interfaces. *ACM UIST*, 167-176.
24. Sumitomo, H., Katayama, T., Terada, T., & Tsukamoto, M. (2014). Implementation and Evaluation on a Concealed Interface Using Abdominal Circumference. *ACM AH*, 51:1-51:8.
25. Tognazzini, B. (1993). Principles, Techniques, and Ethics of Stage Magic and Their Application to Human Interface Design. *ACM CHI*, 355-362.
26. Toney, A., Mulley, B., Thomas, B.H., & Piekarski, W. (2002). Minimal social weight user interactions for wearable computers in business suits. *IEEE ISWC*, 57-64.
27. Yi, B., Cao, X., Fjeld, M., & Zhao, S. (2012). Exploring User Motivations for Eyes-free Interaction on Mobile Devices. *ACM CHI,* 2789-2792.
28. "Dummy" Thermostats Cool Down Tempers, Not Temperatures. http://news.thomasnet.com/IMT/2003/ 04/11/dummy_thermostat
29. Timesify. http://www.timesify.com/.
30. Worddit Document. http://pcottle.github.io/MSWorddit.