Deception and Magic in Collaborative Interaction

Joe Marshall, Steve Benford, Tony Pridmore
Department of Computer Science
University of Nottingham, Nottingham UK
{jqm,sdb,tpp}@cs.nott.ac.uk

ABSTRACT
We explore the ways in which interfaces can be designed to deceive users so as to create the illusion of magic. We present a study of an experimental performance in which a magician used a computer vision system to conduct a series of illusions based on the well-known 'three cups' magic trick. We explain our findings in terms of the two broad strategies of misdirecting attention and setting false expectations, articulating specific tactics that were employed in each case. We draw on existing theories of collaborative and spectator interfaces, ambiguity and interpretation, and trajectories through experiences to explain our findings in broader HCI terms. We also extend and integrate current theory to provide refined sensitising concepts for analysing deceptive interactions.

Author Keywords
Deception, interaction, performance, magic, ambiguity, trajectories, feedthrough, misdirection, spectator interface.

ACM Classification Keywords
H5.m. Information interfaces and presentation (e.g., HCI)

General Terms
Design

INTRODUCTION
The idea of deliberately using computers to deceive people would appear to run counter to the traditional tenets of HCI that emphasise legibility, predictability and consistency as being essential goals of usability. Indeed, in many circumstances it would be at best unethical, and quite possibly illegal, to use a computer to deceive a user. However, there is one growing area of HCI in which deception may in fact be a valid and indeed powerful design strategy. This is in interactive performances, installations, games, rides and other 'cultural applications' of computing that need to create a sense of magical illusion as part of an entertaining and engaging user experience. These kinds of leisure and entertainment applications are a growing market for computing and correspondingly are of increasing interest to those who design and study interaction. This paper therefore sets out to explore how, in practice, we can design deceptive interactions to create magical experiences.

For this, we naturally turn to stage magic as the domain in which deception is most routinely and professionally employed to create magical illusions. We present a study of an experimental performance in which a magician used a computer vision interface to conduct a series of illusions based on the well-known 'three cups' trick. By analysing video recordings and system logs from 17 performances we are able to draw out the detailed ways in which the interface was designed and used to create various deceptions, and also how these were experienced by the audience.

Our study makes two contributions to HCI. The first is a pragmatic one. Drawing on previous studies of (non-computer-based) stage magic we show how our findings can be understood in terms of the two overarching strategies of misdirecting attention and setting false expectations. We then articulate specific tactics for using computers to create deceptions as part of these. Our aim here is to demonstrate that deception is indeed a valid strategy for creating magical interfaces and to explain to others how to set about employing this approach in practice.

Our second contribution is theoretical. Here we wish to provide a more general explanation in HCI's own terms of how these deceptive strategies and tactics work. While Tognazzini has previously adopted stage magic as a metaphor for single-user interaction [17], our aim is instead to account for how magical deceptions are actually carried out in practice in collaborative settings. There is a growing body of theory to account for collaboration through and around shared interfaces including Dix's framework for relating direct and indirect channels of communication [6] and Reeves' taxonomy of design strategies for spectator interfaces, one of which is 'magical' [12]. Others have argued that, in contrast to traditional design goals of clarity and ease of use [11], it may sometimes be useful to design ambiguous interfaces [7] that provoke users into their own interpretations [15]. Finally, the idea of trajectories has been proposed as a way of describing the overall design of complex cultural experiences [2]. This body of theory provides us with concepts to explain and generalise our findings. In return, we extend and integrate these theories to yield refined concepts to guide future studies that wish to focus on deceptive interaction, certainly when studying cultural experiences, but potentially also in relation to using computers as part of everyday social interaction, or perhaps even to help understand issues of computer security.

Permission to make digital or hard copies of all or part of this work for personal or classroom use is granted without fee provided that copies are not made or distributed for profit or commercial advantage and that copies bear this notice and the full citation on the first page. To copy otherwise, or republish, to post on servers or to redistribute to lists, requires prior specific permission and/or a fee.
CHI 2010, April 10–15, 2010, Atlanta, Georgia, USA.
Copyright 2010 ACM 978-1-60558-929-9/10/04....$10.00.

THE CUP GAME PERFORMANCE

We begin by describing a study of an experimental magical illusion called 'The Cup Game', run as a performance for members of the public. We follow a growing tradition in HCI of naturalistic studies of installations and artistic performances from museum 'interactives' [19], to pervasive games and fairground rides [14], to designer furniture in the home [8]. In our case, this approach makes it possible for us to understand the practical detail of how performers create, and audiences actually experience, deceptions and illusions 'in the wild' of a live performance. Magicians explicitly instruct and direct participants in a way somewhat similar to someone running a user study experiment. This meant that we were able to integrate experimental data capture into the study whilst maintaining a sense of performance throughout. The (experienced) performer aimed to make the performance as close as possible to a normal performance.

The performance is an intimate one-to-one experience in which an audience member (the 'punter' in traditional parlance) sits down in front of the performer (the 'magician') and plays the cup game. It is structured as an unfolding series of games, each of which involves a different way of tricking the punter as described below. The workings of each trick are revealed to the punter at its end before passing to the next, creating a structure in which they are led through unfolding layers of illusion and trickery as part of an overall experience. Conveniently, this structure also enables us to study several different approaches to using an interface to deceive within a single performance.

Given this structure, our approach to capturing data involved the magician interviewing the punter about their experience at the end of each trick as part of the explanation of its mechanics. These interviews sought to answer two types of question. First, we wished to discover whether the combination of techniques and technology we employed enabled a successful performance. This required us to determine whether each trick did or did not work to fool the punter and why this was so, something that could be ascertained by noting which cup the punter chose in each trick, observing what the magician and punter did during the trick, and asking the punter for an explanation of the trick. Secondly, data was collected on people's explanations for the effects that occurred during the tricks and how they felt about them. Some demographic data also allowed us to consider possible effects of prior technical or magic experience on susceptibility to the various deceptions.

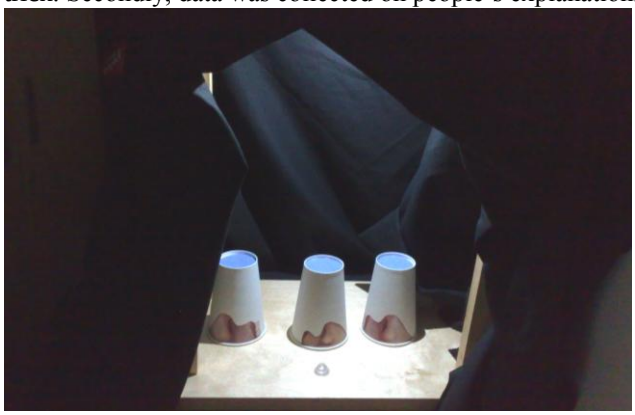

Figure 1 The Cup Table initial set-up

The overall performance therefore unfolded as a series of tricks, interspersed by explanations and interviews so that data capture was integrated into the performance itself. Each entire performance was captured by two video cameras, one looking over the shoulder of the magician and one over the shoulder of the punter. In addition, we also recorded video from the 'tracking camera' that formed part of the computer-vision used during the trick along with system logs describing the system's view of where the cups were located. These four data streams – three video views and one data log file – were then synchronised so that they could be replayed in step, enabling our analysis to piece together the fine details of how the performance unfolded. A final interview at the end allowed the capture of demographic data and some feedback on the general experience of the performance.

The performance was piloted with a group of 10 people known to the magician, before running it for real with 17 unknown volunteers, recruited via posters around our university and paid a small amount for participation. Piloting allowed the magician to get the moves of the tricks smooth and to fine-tune the spoken script before going live. The following sections summarise the design of the final performance as experienced by the 17 volunteers and analysed and discussed in the remainder of the paper.

Stage 1: Tracking the Cups by Sound

On arrival each punter is ushered into a small, dark room, and asked to sit down on a chair in front of a special cabinet style table. This is covered in black cloth, and has a light embedded in the top, making it look like a small table-top theatre surrounded by curtains. Three cups are placed on this 'stage' along with a glass bead (Figure 1). The magician stands on the other side of the table, so the cups are not visible to him behind the cloth and begins his script: "*In this trick, the bead is placed under one of the three cups, and shuffled by one person, and a guess is made by the other person as to which cup the bead is under.*"

In the first trick, the magician puts on a blindfold, and turns so that they are facing away from the table so as to emphasise that they cannot see what is happening. Then he tells the punter that he will detect which cup the bead is under from only the sound of the cups being moved. The punter is asked to put the bead under the middle cup, and then shuffle the cups. Once the cups have been well shuffled, the magician says '*I know which one it is under*', turns round, sits down and takes off the blindfold, and without looking at the cups, identifies the correct cup. At this point, the performance is paused while the magician asks some questions of the punter in order to find out whether they believe he detected the cup by sound, and if not, what alternative explanation they are able to offer.

Stage 2: In Front of Your Eyes

This section begins with the magician revealing that the punter was tricked in the previous stage because the magician used a hidden computer system that tracked the movement of the cups on their behalf. First, he takes the black cloth off the table, to reveal a camera in the top of the table, looking down on the cups (Figure 2(a)), and a laptop underneath the table, showing the view from the camera (Figure 2(b)). The laptop is placed on the floor so that its screen can be seen by the magician (and now by the punter too). The camera image can be seen, overlaid with a circle marking the top of the central cup. This circle is generated by a computer vision system that tracks the movements of the three cups, the underlying interactive technology that is used to support the tick. This system uses a simple computer vision tracking system to track the bright blue coloured ends of the cups, using the view from the camera directly. This is relatively simple, as the tapered shape of the cups stops the tops of them from ever touching, so there are always three separate areas of colour in the camera's view. The light built into the table is set close to the camera, to ensure consistent lighting of the cups.

The magician now begins the second trick, bending down to press a key on the laptop, causing an instruction to appear on the screen saying to put the bead under the central cup. The video image fades out, showing only solid circles representing where the three cups are according to the vision system. He shuffles the cups, saying '*if you look at the screen now, as I move a cup, the circle on the screen moves. Watch carefully and pay attention to this one [shakes centre cup] which is the one with the bead under it.*' The cups are then shuffled for ten seconds, at which point the computer displays the message '*I know which cup it is under*' and the punter is asked to say which cup they think the bead is under. If the trick works, they think that it is under the right hand cup from the magician's viewpoint. They are asked if they are sure, and a key is pressed to reveal that the computer has guessed the left hand cup. The two cups are lifted, and the computer is shown to be correct. The punter is asked what they think happened, whether they think it was a trick, or if the computer was just better at tracking the cups than they were.

Stage 3: Get Your Own Back

The magician says that having been 'tricked by the computer' (or more accurately by the magician using a computer), the punter now has a chance to '*get their own back*' by trying to trick the computer themselves. However, the hard thing here is that the computer has been designed to catch people trying to trick it. The punter must make the computer think that the bead is under a different cup, without setting off a 'cheating alarm' in the tracking system.

The computer is started, and gives instructions to put the bead under the central cup, and shuffle, as before, and after a certain amount of shuffling, guesses which cup it is under. If it detects obvious cheating (such as taking a cup off the table), it stops, with a loud beeping noise and '*cheating detected*' on the screen. The cheating detector basically detects when the computer cannot see a cup for a certain amount of time, which catches covering the tops of cups, taking cups off the table or turning off the light, but does not catch anything being done to the bead. The punter must try and make the computer guess wrong. If they have trouble the first time, the magician emphasises that they can do '*absolutely anything*' to trick the computer, as long as it doesn't notice that it has been tricked. This aims to avoid the situation where a person repeatedly tries to do the trick in the same way and fails. They are allowed to try to trick the computer up to 3 times. Once they have managed to trick the computer or failed after three attempts, the punter is again asked some questions, this time about the type of strategy they were trying to use, and especially whether they were trying to beat the computer in the same way as they would beat a person, or if they were trying to do something different because they knew it was a computer.

(a)

(b)

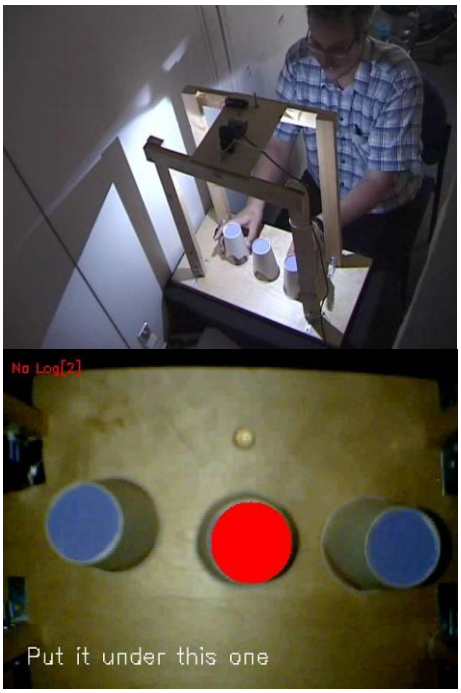

Figure 2 The Cup Table apparatus revealed to a punter

Stage 4: The Big Reveal

At this point, the magician demonstrates how the punter was led to pick the wrong cup in Stage 2. This was by lifting the centre cup slightly, and flicking the bead to be under the left hand cup, which could be done while they were attending to the computer display rather than to the cups. In other words the magician employed traditional sleight of hand to quickly move the bead to be under the left hand cup without being seen while the punter was otherwise distracted. He demonstrates how this is a fundamental weakness in the tracking system, that it is actually tracking the cup, and not the bead itself, and how this can easily be used to trick the computer. He also demonstrates the

secondary part of the trick, how the computer appears to guess the correct cup, which is actually done by using a secret key press on the laptop to tell the computer to display a prescripted message confirming the outcome. Finally, he shows how the cheating alarm works including detecting covering of cups with hands, moving cups out of view, and other methods of cheating the computer.

Stage 5: Experience and Demographic Information

Finally the magician asks the punter a set of questions about their background knowledge of the type of technology used, and also of magic tricks. They are then asked about their experience of the tricks, how they felt after each, how it felt to trick the computer themselves, and generally whether they enjoyed the performance.

FINDINGS FROM THE STUDY

There were 17 performances of the Cup Game, 9 to males and 8 to females. 14 of these were undergraduate students, 2 were postgraduate students, and 1 was a lecturer. Of these, all professed to use computers regularly and 10 had programming experience. Only 2 claimed to have used computer vision technology, and a further 2 had heard of this kind of technology before, meaning that for 13 of the punters, this was the first they had heard of or used computer vision interfaces. As for experience of magic, all had seen magic performed on television, and 9 had seen magic performed live. 11 had performed a magic trick themselves, mostly as children, although 1 person still performed magic tricks as a hobby, and 1 other had performed stage illusions at a semi-professional level.

Overall, the magician was very successful at tricking the punters and the supporting video tracking technology proved reliable. There was only one failure in the first Cups by Sound trick, which analysis of the log files and video showed to be due to the punter failing to follow instructions and holding the cups from above with their hands, which unknown to them caused the tracking to fail. The second, In Front of Your Eyes trick was equally successful, failing in just one case, where a (different) punter spotted the bead being switched. The Get Your Own Back stage also proved successful, with beating the computer being a difficult, but not impossible, challenge for most punters.

When asked whether they enjoyed the tricks, punters were unanimously positive in their responses, suggesting that the performance was an overall success: *"Yeah… very interesting"*, *"A bit sneaky, but yeah it was fun"*, "*Yeah, definitely enjoyed them"*. However, in the final interview, 3 people expressed disappointment at the use of the computers in the tricks. One said that on being told shown that the computer was doing the first trick, their initial reaction was "*I don't trust magic - [I felt] slightly cross*". Another expressed disappointment that magic tricks are not real magic - "*I'm a bit disappointed that there's no such thing as magic ... I realised that all those magic tricks have nothing behind them*", which was somewhat unexpected. One more technically-minded person expressed disappointment with the way that the computer guessed which cup the bead is under correctly during the In Front of Your Eyes trick, "*I thought that ball must be special, RFID, something, I don't know [but] you just pushed a button!*". The other participants were not disappointed, because they knew it was a trick already. One participant actually thought that the use of technology enhanced the tricks – "*amazing technology ... it was quite clever, a lot of tricks, when you find them out they're quite lame really*".

Given these overall impressions, we now progress to a more detailed discussion of two specific aspects of the Cup Game: an analysis of how and why the two tricks by the magician appear to have worked, followed by an account of what happened when people who did not know the system tried to trick the computer in return.

Tricking People with Computers

One obvious factor in the success of the Cups by Sound trick was that the mechanism for the trick was hidden from the punters. The initial cover story did not mention computers and the technology was not directly visible so why would they guess that computer vision was being used? Moreover, not only did the cover story fail to mention the true mechanism, but it also gave a misleading explanation as to how the trick was done. Somewhat unexpectedly, 10 of the 17 punters reported that they believed that the magician could track the cups by sound alone. Although the technology was largely hidden, the camera was visible to an inquisitive punter and here the cover story played a second important role of distraction. Inspection of the video recordings revealed that the magician's constant patter was successful at distracting the punters, keeping their eyes off the camera and preventing them from examining the table in detail. Only one punter spotted a camera; this person was one of 2 people who had previous experience using video tracking systems, and had considered that the trick could be performed with it.

Revealing the secret of the Tracking Cups by Sound trick then set up an anticipation that the computer would be used in a similar way in the second trick; as one person said "*you tell people how you play first time, you use the computer ...I don't realise you actually play the trick [yourself]*". Using computers in a trick also made use of a longer term expectation that people have, that computers are trustworthy. For example, during the In Front of Your Eyes trick, people were asked whether they had followed the cups and were certain that they had the right cup. However, after the computer had shown them the correct cup, 10 people believed that they must just not have been as good at following the cup movements as the computer. This trick revealed a high level of trust in the computer: roughly 60% of the punters were willing to change their mind about their ability to visually track objects when the computer beat them. The following transcript demonstrates the way in which one participant lost confidence in their own judgement as the computer's answer was revealed.

Magician: *Which one do you think it's under?*
Punter: *I think it's under this one.* (Cup 3)
M: *Let's see what the computer thinks.*
P: *I've done it wrong I know*
M *The computer thinks it's under that one.* (Cup 1)
P: *The computer's gonna be right*

This trust in computers is also demonstrated by the exclamation of a punter in the final reveal section - "*I didn't know computers could cheat.*"

The In Front of Your Eyes trick also involved distraction. By blurring the boundary between the reveal of the first trick and the performance of the second trick, the magician exploited a moment when the punter was off their guard to switch the bead. The following transcript shows how the magician deliberately shifted the punter's attention to the computer screen at the vital moment:

M: *So, it tells me to pop it under this one*
(puts bead under cup)
M: *In a second it will tell me to shuffle.*
M: *OK so it's telling me to shuffle, so it's under this one*
(moves the centre cup so that the link between the movement and screen is clear. At this point the punter is looking between both screen & cups)
M: *Pay attention to which one it's under*
(Magician looks down at screen, cueing the punter to look down, and does the switch when they do)

At the critical moment, the punter has a choice of watching the real cups or the visual representation on the computer screen – the magician looks down, a strong cue to encourage them to watch the computer screen (as demonstrated by previous experiments [16]), then switches once he is sure they are watching the screen. This moment was key to the performance of the second trick – the manual switch of the bead from the centre to the left had to be performed late enough so that punter's attention was directed to the screen, but before they realised this was also a trick. This was successful in all but one case. Analysis of this single failure showed that this was a mistake by the magician, who performed the switch too early, before he finished the line of patter directing the punter's eyes to the laptop screen and away from the real cups.

Although this was not a controlled experiment, it is interesting to briefly consider how the demographics of the audience might relate to their susceptibility to these tricks. Knowledge of magic did not appear to have any great effect on the suggestibility of the participants; those who had seen magic live showed no differences in levels of belief compared to those who had only seen magic on television. Similarly, there were no great differences between those who had performed tricks themselves and those who hadn't. There did however appear to a more interesting relationship to the level of prior technical knowledge. 10 of the participants had some software development experience and, of these, 8 believed that the Cups by Sound trick was done by sound alone, as described in the cover story. In contrast, of 7 users with no development experience, only 2 believed the cover story. It may appear that having technical knowledge affords no great protection against being tricked using a computer in this way. This may perhaps be because these participants were more able to come up with explanations for how it might be possible, such as "*You counted the number of times I moved it*" and "*You could hear the bead moving in the cup*", whereas those who did not believe it, thought that the trick was simply not possible or that other explanations were more likely: "*it is feasible that somebody might have a really good sense of hearing but quite improbable*", *"maybe there is a mirror somewhere", "something in the table", "I have no idea...".* Only 2 of the 17 punters suggested that any electronic technology might be used, this is perhaps due to the deliberately low tech wooden and cloth construction of the visible parts of the cup table. Finally, in much of the earlier magic literature (e.g. [18] p492), it is stated that women are more suggestible than men. This does not seem to have been the case in this experiment – of men doing the experiment 7 of the 9 believed the cover story in the first trick, compared to only 3 of 8 women. However, 7 software developers were male, so it is hard to separate gender and technical knowledge without a controlled study.

People Tricking Computers

As well as being tricked with the help of the computer system, in stage 3 of the performance punters were asked to try and trick the computer themselves. This gives us an interesting insight into how non-expert users might fare when trying to deceive a computer.

Tricking the computer turned out to be challenging, with only 2 people able to trick the computer first time, and 5 completely failing to trick it. This appeared to be partly because they had built up expectations of the correct way to move the cups from the previous two tricks, which meant that even though they were prompted that they could do 'anything' to beat the computer, they tended to at least start by trying the same moves that they had seen demonstrated. Tricking the computer was also hard because people had no idea of how to fool a computer, perhaps showing again a level of trust in the reliability of computers: "*it's gotta be ... how can you trick a computer? ... I've no idea*". Those who beat the computer, expressed satisfaction in managing to beat it: "*it felt GOOOD [to beat the computer]*."

The strategies people adopted could be divided into two types, firstly simply trying to fool the computer by moving the cups quickly, or moving them erratically as one might do if trying to trick a person and secondly by specifically trying to trick the computer, for example by covering the top of the cups, piling the cups on top of each other, or simply taking the bead out from under the cup. The quick movement strategy did not fool the computer. When people tried to fool the computer, the most common thing to try was covering the cups, which set off the 'cheating alarm',

as it was easy for the system to detect. Most people only finally succeeded by accidentally covering a cup too quickly to set off the cheating alarm, or other things that caused tracker errors. Interestingly, 2 punters who did notice that the bead was not tracked managed to trick both magician and computer, by performing fast movements along with switching or removing the bead from the cups.

The effect of the magician's prompt on punters' expectations of how to interact with the cups was very clear – on the first try at beating the computer, only 3 people performed strategies designed to beat the computer, with the rest simply moving the cups fast as if trying to fool a person. On their second try, 9 of the 15 punters who didn't beat the computer first time used a strategy directed at the computer. This indicates the power of the expectations built up in the first two tricks, and how the magician's patter was able to alter the built up expectations. At this stage, punters with software development experience did have an advantage, with a larger proportion of non-developers (3/7) compared to developers (2/10) failing to trick the computer at all. Those who tricked the computer first time were also developers. In practice it seemed that developers were more likely to work out that the computer was just tracking the cups and not the bead.

In summary, the performance appears have been a success as the magician was able to trick nearly all of the punters who found the experience to be enjoyable. This success arose from a complex performance structure that involved a constantly shifting set of relationships between the magician, punter and computer, both in terms of their momentary division of attention and distraction, and also in continually confounding their understanding as the performance passed through its various stages. Finally, the idea of getting the punter to trick the computer introduced a further novel and enjoyable twist to the whole experience.

MISDIRECTION AND FALSE EXPECTATIONS

We now generalise our findings from studying the Cup Game, articulating the practical strategies and tactics that a designer/performer might use to deceive a spectator as part of a magical experience and relating these to current theories of interaction within HCI. Many magic performers use technology, from traditional props such as altered playing cards to modern digital technologies, although it should be noted that there is resistance amongst magicians to the use of certain technologies such as video special effects (it is seen as cheating to produce an effect on TV that cannot be produced live). We turn to previous accounts of conventional stage magic to provide us with two initial sensitizing concepts [5] around which we can frame and structure our discussion. Specifically, we propose that the success of the various tricks in the Cup Game relied on the combination of two overarching strategies: *misdirecting attention* and *setting false expectations*; both of which have previously been described by psychologists and magicians wishing to understand the mechanisms behind magical deception. We consider each of these strategies in turn, briefly reviewing the relevant literature on stage magic, drawing on our findings to articulate specific tactics for deception that employ computer interfaces, and then relating these to HCI theory.

Misdirecting attention

Key to the success of magic tricks is the misdirection of the audience – directing them away from the real explanation of a trick towards a false interpretation of events. This is not just about hiding actions, but requires active direction of the audience's attention. Experimental work investigating magical performance has demonstrated that magicians lead people to not see things that are happening directly in their view [16], and even to believe 'positive illusions', where people believe they see something purely due to suggestion by the magician [20]. Magicians perform misdirection by keeping track of multiple narratives, both the real one and the one that they are trying to convince their audience of. Magicians place particular emphasis on the importance of the spoken element of the performance or patter, in making clear the external story of the trick, and hiding the secret: "*he says what he does not do, he does not do what he says, and what he actually does he takes particular care not to say anything about.*" [18]. Multiple narratives may not even occur at the same time, magicians often reframe tricks in time [9], for example when the external story is that they are shuffling cards between tricks, they may set up the deck for the next trick. In these various ways, a magician misdirects a spectator's attention away from deceptive actions and towards those that reinforce a cover story.

Computers introduce new possibilities for misdirection. Drawing on the findings from our study, we suggest that a performer can draw on several tactics to (mis)direct a spectator's attention towards and away from different elements of a computer interface at different times.

1) Hiding and revealing aspects of interaction. Perhaps the most basic tactic is to simply hide the computer interface so that the spectator does not even know that it exists as seen in the initial Cups by Sound trick. However, other tricks involved deliberately revealing some aspects of interaction (e.g., the laptop interface) so as to make it easier to hide others (the moving of the bead).

2) Split attention between elements of the interface. A more subtle tactic is to reveal the interface, but to split the spectator's attention between its different elements. Thus, the interface is clearly visible in the In Front of Your Eyes trick, but the physical separation of the cups from the laptop screen makes it difficult for the spectator to attend to both, while enabling the performer to steer their attention.

3) Use patter to distract and direct attention. The performers' patter is essential throughout, both to distract the spectator from inspecting the set-up in detail, but also to deliberately steer their attention, either through words ('look over here') or through gaze and gesture.

These various tactics can be related to existing HCI theory. Reeves and colleagues' taxonomy classifies spectator interfaces according to the extent to which they hide, partially-reveal, transform, reveal or even amplify a performer's manipulations of an interface, versus the extent to which they do the same for the visible effects of these manipulations [12]. Several aspects of their taxonomy are relevant here. First, they explain that manipulations include those that directly control the interface but also 'performative gestures' that occur around it, which we see when the performer uses gesture to direct a spectator's attention. Secondly, their observation that manipulations may be hidden accounts for our first tactic. Reeves and colleagues also argue that the category of interfaces that hide a performer's manipulations while revealing their effects should be labeled as 'magical'. While in our study we do see moments when the magician hides manipulations and reveals their effects, our findings suggest that this characterization of magical interfaces is too simple. Rather than a blanket revealing and hiding of everything, the performer often reveals only some aspects of manipulations and effects while hiding others, simultaneously occupying several areas of Reeves' taxonomy. Moreover, they may subsequently readjust what is hidden and revealed as a series of deceptions unfolds, thereby moving around the taxonomy. Finally, the spectator as well as the performer may manipulate the interface (e.g., move the cups).

We therefore offer the interpretation of Reeves' taxonomy in Figure 3 in which we classify the computer-vision interface at three successive moments during the cup game: (1) in the middle of the 'Cups by Sound Trick'; (2) at the moment of switching the bead; and (3) when pressing the button to reveal the computer's choice. Unlike The original taxonomy which only presented the spectator's view of interaction, this figure shows both the performer's (P) and spectator's (S) views overlaid on the same space. It also reveals how these migrate around this space throughout the performance, including how the spectator's attention is dynamically shifted between different manipulations and effects and away from others as the performance unfolds.

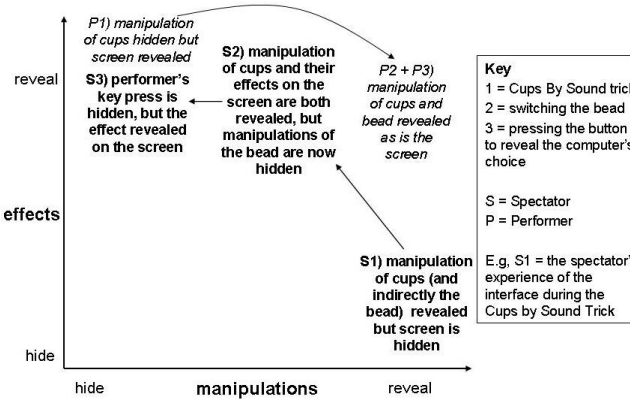

Figure 3: Charting deceptive manipulations and effects

While Reeves' taxonomy accounts for some aspects of misdirection, it does not naturally express the manner in which the performer splits and directs attention between the cup, beads and screen, or the role of their patter in this. To explain these, we instead turn to the cooperative work diagrams introduced by Dix to represent situations where two people are interacting both with each other and with an 'artefact' (typically a computer) that mediates this interaction [6]. These diagrams highlight two important concepts for our discussion. The first (figure 4a) is the presence of multiple channels of interaction between the computer (C) and the two participants (P). The second is the concept of 'feedthrough' (figure 4b) whereby one participant's interactions are processed by the computer before being passed onto the other.

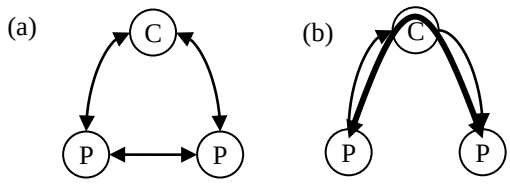

Figure 4 Dix's CSCW Diagrams

Figure 5 now explains the tactics of misdirection by extending Dix's diagrams and combining them with elements of Reeves' taxonomy. We distinguish the two participants as performer (P) and spectator (S) and also separate the computer (C) from other more conventional artefacts such as cups and beads (A), giving us three channels of communication, one direct and two feedthrough (more complex tricks involving many computers or artefacts would introduce further channels). In Reeves' terms, feedthrough in one direction arises from one participant manipulating the computer or artefact which then generates effects that may be passed onto the other.

The presence of two (or more) feedthrough channels enables the performer to divide the spectator's attention between them (tactic 2), while the direct channel carries the patter that steers their attention towards one and away from the other (tactic 3). Moreover, the performer designs and arranges the computer and other artefacts to hide and reveal different combinations of manipulations and effects creating what we shall call *deceptive feedthrough*. Take the Cups by Sound Trick as an example (Figure 6). The spectator's manipulations of the cups (A) and the subsequent effects of

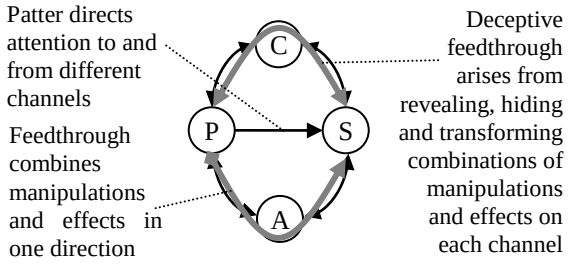

Figure 5 Extending Dix's CSCW diagrams

the cups moving are visible to themselves but not to the performer. However, the performer can hear some noise and there is an *implied effect* through the cover story that this enables them to follow the cups. This is deceptive feedthrough. In the other channel, the spectator unwittingly manipulates the computer vision interface, but does not see the effects of this. The performer however does see these effects, which is how they know the location of the bead, although the spectator does not know that this is what they are doing. This is another kind of deceptive feedthrough.

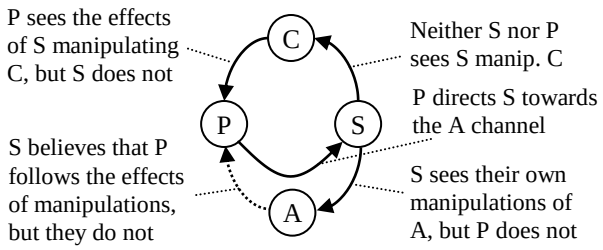

Figure 6 Extending Dix's CSCW diagrams

In the In Front of Your Eyes trick, the manipulations of cups and vision interface and the effects of these are visible to both spectator and performer. However, the performer's manipulation of the button to reveal the computers 'guess' is invisible to the spectator, although they do see its effects. The performer's manipulation of the bead, while potentially visible to the spectator, is not usually seen due to the distraction of looking at the screen.

There is clearly more to these deceptions than what is directly revealed or hidden. They also rely on what is known and believed. Thus, in Cups by Sound the spectator believes that the performer can follow the sound effects of the cups being moved, while in In Front of Your Eyes they believe that the computer's guess arose from tracking the cups, not from the unseen press of a hidden key. This brings us to our second strategy.

Setting false expectations

While the strategy of misdirecting attention accounts for the moment-by-moment details of executing deceptions within specific tricks, this only becomes possible because the performer has carefully set the spectator's expectations in advance. More specifically, they have invariably set *false expectations* as we now discuss.

Manipulating the audience's expectations is another well-known aspect of conventional magical technique. One of the most basic 'rules' of magic, never to perform a trick twice in a row, is a classic example of this, as many misdirections will no longer work once audience members know what effect is going to be produced [20]. The flip side to this is that in many tricks a performer performs actions once and then the second time only pretends to repeat them, exploiting expectations that were build up the first time around [9]. An audience's lack of expectations may also be exploited, for example if a trick fails in some way, such as the wrong card is picked, the magician may use the fact that the audience do not know the expected outcome to recover by seamlessly moving onto a different trick, in effect changing the hidden narrative, whilst keeping the visible narrative moving smoothly [9]. Expectations are also fundamental to another key aspect of magical performance, that of suspension of disbelief. Binet [3] argues that a key part of the success of illusions in the hands of magicians is that the audience expects and enjoys trickery. The expectation of the audience that they will be tricked is also an important part of magicians' ethics, with it being seen as acceptable to perform stage magic, but not to use the same methods to trick people out of money. Real deception may occur within a trick itself, but it is always within an act that is known by the audience to be trickery.

As with misdirection, the introduction of computers creates new opportunities for setting false expectations, as in the following tactics that were seen in the Cup Game.

4) Build on expectations of computers. People may have a natural tendency to trust computers, in which case exposing them as part of the trick may positively build false expectations. Even if not, people may not fully comprehend their capabilities, especially where invisible sensing systems (visual tracking in our study) are employed.

5) Build on expectations of stage magic. It is both ethically and practically important to make spectators aware that the deception is in fact staged magic. While on the one hand this may encourage them to look for the deception, it also enables them to suspend disbelief and play along.

6) Lead by demonstration. People naturally mimic the interactions that you show them, reducing the likelihood of exploring other interactions that might reveal a deception.

7) Establish a cover story. It is important to provide at least one credible alternative explanation to the truth and to actively guide people towards this.

8) Use the reveal of one deception to set up the next. Appearing to reveal a deception can be an entertaining payoff to a trick, but can also be the ideal way of actually setting up new expectations to enable the next trick.

Once again, we can draw on existing HCI theory to explain these tactics. The need to understand a user's prior knowledge and expectations is widely recognised as being essential to good interface design [11], though this is usually to create interfaces that can be understood rather than ones that cannot as we see in tactics 4 and 5. Tactic 6 on the other hand, resonates with previous studies of museum interactives in which one visitor's interactions have been observed to configure those of the next visitor who learns what to do by watching them [19].

Various authors have discussed the particular challenges of interacting with invisible sensing systems, including Bellotti and colleagues who argue that, in contrast to graphical user interfaces, sensor-based interfaces challenge users in terms

of knowing how to address the interface, understanding whether the system is attending to actions, controlling the interface, confirming correct responses, and avoiding mistakes [1]. While these may present problems in many applications, they may actually create opportunities for deception in magical performance. In other words, the very invisibility of sensor-based systems makes them an ideal candidate for magical interfaces. Moreover, when confronted with an unknown, inexplicable interface, users may tend to draw on their everyday knowledge of human perception to explain how it works. Thus, asked to trick the computer, participants in our study would try to move the cups very quickly (a good way to fool a human, but not to fool our vision system) while not trying to move the bead (which a human can easily see but our system cannot).

Turning to the overall structure of the experience, Benford et al. [2] propose that complex cultural user experiences can be described in terms of trajectories through hybrid structures of space, time, roles and interfaces. Canonical trajectories express the author/performer's ideal journey through the experience, whereas participant trajectories express the routes that individual participants actually follow. These trajectories must carefully negotiate various transitions along the way if they are to maintain an overall coherence, including beginnings, endings, role and interface handovers, disengagements and seams. Using these concepts, tactics 7 and 8 enable a performer to establish an overall trajectory through the performance that passes through several distinct stages. Here, the performer actually designs two canonical trajectories, one being the narrative that the spectator is intended to believe (the cover story), and the other the sequence of actions that the performer will actually carry out.

The cover story is important because it enables an alternative interpretation of the events that are unfolding. Use of a cover story relies to some extent on the audience's cooperation to flesh it out and make it believable, with the spectator being steered away from the true interpretation towards one or more false interpretations. This idea directly reflects the arguments of Sengers and Gaver, who suggest that in contrast to conventional task-oriented interfaces, artistic or cultural interfaces may deliberately create space for multiple interpretations, in part by discouraging obvious real interpretations of what the interface is doing [15] and also by exploiting ambiguity in interface design [7].

Returning to trajectories, there are also important transitions between the stages in which one trick is set-up even as another appears to be revealed. Indeed, 'set-ups' and 'reveals' are two new examples of transitions. We would also extend the argument of [2] to say that our trajectories pass through a layered structure of knowledge as much as they do through space, time, roles and interfaces. Thus, in the fictional canonical trajectory, the spectator is first intended to believe the initial cover story of tracking 'cups by sound'; then they are intended to believe that the computer reliably tracks cups; next they learn that the computer is not always used honestly; finally they learn the limitations of the computer by tricking it themselves. By way of summary, Figure 7 offers a schematic overview of the structure of our performance, showing two parallel trajectories running through a series of tricks, each involving various misdirections, and interspersed by transitions in which expectations are reset.

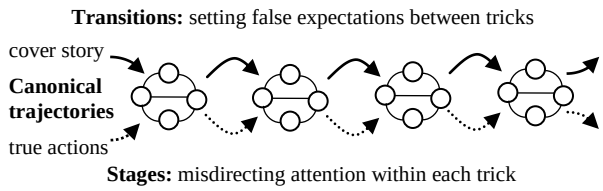

Figure 7 Trajectories through unfolding knowledge

CONCLUSIONS

By studying an experimental magic performance we have seen how performers can employ computers to deceive spectators and thus create magical illusion. This study has revealed how creating and sustaining such deceptions is a complex business involving two broad strategies:

- Misdirecting attention through the tactics of: hiding some aspects of interaction while revealing others; dividing attention between different elements of the interface; and using patter to distract and direct attention.
- Setting false expectations through the tactics of: exploiting general expectations of computers and stage magic; providing an alternative cover story; demonstrating intended interactions; and also resetting expectations by embedding the set-up of the next trick into the reveal of the previous one.

We have also seen that concepts from current HCI theory can help explain how these various strategies and tactics work from an HCI point of view, relating them to contemporary discussions in interface design. Overall, we see that the performer establishes competing canonical trajectories, one illusory and one actual. We also see how these trajectories chart out a journey through progressively unfolding knowledge and confounded expectations as a series of tricks unfolds. Key transitions along the way involve resetting expectations, carefully framing the next interaction while employing ambiguity to create a plausible cover story. The detailed mechanics of each trick rely on the misdirection of attention and can be explained in terms of switching attention between multiple channels of interaction, and by variously hiding, revealing or transforming manipulations and effects in each channel to create various kinds of deceptive feedthrough.

In turn, our study suggests various extensions to current HCI theory. We have seen that it is necessary to combine Reeves' and Dix's existing frameworks in order to explain the subtleties of misdirection and deception. We have also

seen that the conceptual framework of trajectories needs to adopt further kinds of transition, consider relations between multiple canonical trajectories, and also address structures of expectations and knowledge.

These findings and subsequent theoretical discussions are intended to be of use to HCI researchers in several ways. In the short term, just as recent work on ambiguity has challenged traditional views of interaction, so we seek to broaden HCI's agenda to consider the currently unfamiliar idea that the active deception of one user by another can be a valid approach to interaction design. In the medium term, we aim to provide sensitizing concepts to guide further studies of deception and interaction. Ultimately, we wish to support the designers of future installation, performances, games and rides with practical strategies and tactics for creating increasingly magical interactions.

While our focus here has been restricted to stage magic, we anticipate that our findings may be of value in other areas of HCI. For example, in the same way that [4] has used ideas of ambiguity to explain aspects of everyday social interaction, so our concepts might potentially help explain the role of deceptions as part of the 'performance' of everyday social life. Finally, there is the issue of the deceptive use of computers as part of computer crime, pointing us towards the field of computer security. While this has not been the focus of our work here, computer security has long involved aspects of 'social engineering', where people are tricked in order to gain access to computer systems. More recent developments have clear parallels with HCI work on multiple interpretations and ambiguity, including the use of deliberately inconsistent responses in computer firewalls, designed to remove an obvious interpretation from the computer's response [10] and the creation of computer systems that study their attackers and create deceptions that are customised to trick a particular person [13]. It is an intriguing topic for future work to explore how the strategies and tactics used to deceive as part of the entertainment of stage magic might ultimately inform our understanding of malicious deceptions using computers.

ACKNOWLEDGEMENTS

This work was funded by the Engineering and Physical Sciences Research Council under the Horizon (EP/G065802/1) and Challenge of Widespread Ubiquitous Computing (EP/F03038X/1) projects.

REFERENCES

1. Bellotti, V., Back, M., Edwards, W.K., Grinter, R.E., Henderson, A., and Lopes, C. Making sense of sensing systems: five questions for designers and researchers. In *CHI*, 2002, 415–422.
2. Benford, S., Giannachi, G., Koleva, B., and Rodden, T. From interaction to trajectories: Designing coherent journeys through user experiences. In *CHI*, 2009, 709–718.
3. Binet, A. Psychology of prestidigitation. *Annual Report of the Smithsonian Institution, 1894* 555–571.
4. Boehner, K. and Hancock, J.T. Advancing ambiguity. In *CHI*, 2006, 103–106.
5. Bowen, G. Grounded theory and sensitizing concepts. *Int. Jnl. Qualitative Methods 5*, 3 (2006), 1–9.
6. Dix, A. Challenges for cooperative work on the web: An analytical approach. *Computer Supported Cooperative Work 6*, 2-3 (1997), 135–156.
7. Gaver, W., Beaver, J., and Benford, S. Ambiguity as a resource for design. In *CHI*, 2003, 233–240.
8. Gaver, W., Bowers, J., Boucher, A., Law, A., Pennington, S., and Villar, N. The history tablecloth: illuminating domestic activity. In *DIS*, 2006, 199–208.
9. Nardi, P. Toward a social psychology of entertainment magic. *Symbolic Interaction 7*, 1 (1984), 25–42.
10. Neagoe, V. and Bishop, M. Inconsistency in deception for defense. In *NSPW '06: Proceedings of the 2006 workshop on New security paradigms*, 2006, 31–38.
11. Norman, D.A. *The Design of Everyday Things*. MIT Press, Cambridge, MA, USA, 1990.
12. Reeves, S., Benford, S., O'Malley, C., and Fraser, M. Designing the spectator experience. In *CHI*, 2005, 741–750.
13. Rowe, N.C. Designing good deceptions in defense of information systems. In *20th Annual Computer Security Applications Conference*, 2004, 418–427.
14. Schnädelbach, H., Rennick Egglestone, S., Reeves, S., Benford, S., Walker, B., and Wright, M. Performing thrill: designing telemetry systems and spectator interfaces for amusement rides. In *CHI*, 2008, 1167–76.
15. Sengers, P. and Gaver, B. Staying open to interpretation: engaging multiple meanings in design and evaluation. In *DIS*, 2006, 99–108.
16. Tatler, B.W. and Kuhn, G. Don't look now: The magic of misdirection. In R. van Gompel, M. Fischer, W. Murray, and R. Hill, eds., *Eye Movement Research: Insights into mind and brain*. Elsevier, Amsterdam, Netherlands, 2007, 697–714.
17. Tognazzini, B. Principles, techniques, and ethics of stage magic and their application to human interface design. In *CHI*, 1993, 355–362.
18. Triplett, N. The psychology of conjuring deceptions. *American Journal of Psychology 11*, 4 (1900), 493–510.
19. vom Lehn, D., Hindmarsh, J., Luff, P., and Heath, C. Engaging constable: revealing art with new technology. In *CHI*, 2007, 1485–1494.
20. Wiseman, R. and Greening, E. 'it's still bending': Verbal suggestion and alleged psychokinetic ability. *British Journal of Psychology 96*, 1 (2005), 115–127.