

I did that! Measuring Users' Experience of Agency in their own Actions

David Coyle^{1,5,6} James Moore^{2,3} Per Ola Kristensson⁴ Paul C. Fletcher² Alan F. Blackwell¹

¹ Computer Laboratory and ² BCNI ³ Dept. of Psychology, Goldsmiths ⁴ School of Computer Science
University of Cambridge, UK University of London, UK University of St Andrews, UK

⁵ Trinity Centre for Bioengineering
Trinity College Dublin, Ireland

⁶ Department of Computer Science
University of Bristol, UK

csydc@bristol.ac.uk, jm697@cam.ac.uk, pok@st-andrews.ac.uk, {pcf22, afb21}@cam.ac.uk

ABSTRACT

Cognitive neuroscience defines the sense of agency as the experience of controlling one's own actions and, through this control, affecting the external world. We believe that the sense of personal agency is a key factor in how people experience interactions with technology. This paper draws on theoretical perspectives in cognitive neuroscience and describes two implicit methods through which personal agency can be empirically investigated. We report two experiments applying these methods to HCI problems. One shows that a new input modality - skin-based interaction - can substantially increase users' sense of agency. The second demonstrates that variations in the parameters of assistance techniques such as predictive mouse acceleration can have a significant impact on users' sense of agency. The methods presented provide designers with new ways of evaluating and refining empowering interaction techniques and interfaces, in which users experience an instinctive sense of control and ownership over their actions.

Author Keywords

Personal agency; experience of agency; implicit measures; cognitive science; evaluation methods;

ACM Classification Keywords

H.5.2 [Information Interfaces and Presentation]: User Interfaces - Evaluation/methodology;

INTRODUCTION

The seventh of Shneiderman's Eight Golden Rules of Interface Design states that designers should strive to create interfaces that "*support an internal locus of control*" [33]. This is based on the observation that users "*strongly desire the sense that they are in charge of the system and that the system responds to their actions*". The presence of this rule in Shneiderman's list reflects the critical importance of the sense of control in human computer interactions. Early

research by Rotter laid the foundations for the concept of locus of control [32]. In more recent years the sense of control - or "Experience of Agency" - has become the focus for a significant body of research in cognitive neuroscience [12, 13, 16-19, 26, 27, 35-38]. Within this literature the experience of agency is defined as a person's innate sense of being in control of their actions and through this control of being responsible for, or having ownership of, the consequences of those actions. Researchers in this field have developed techniques through which the experience of agency can be empirically investigated.

We believe that the experience of agency, as defined in cognitive neuroscience literature, is key to how people experience interactions with technology. In this paper we consider how implicit measures for the experience of agency can be applied in HCI research and practice. We provide a brief review of literature on agency and introduce two implicit methods through which it can be empirically investigated. We then describe two experiments that investigate the sense of agency in simple human computer interactions. The first experiment investigates the effect of a change in the input modality of an interaction on user's sense of agency. It compares traditional keyboard input with skin or body-based input and shows that changes in input modality can have a substantial effect on the sense of agency. The second experiment investigates what happens when a computer assists users in achieving their goals. The results of this experiment suggest that up to a certain point, the computer can assist users whilst they still retained a strong sense of agency. Beyond this point users experienced a significant loss in their sense of agency.

This paper makes several contributions to the field of HCI:

1. It introduces two methods through which the sense of agency can be implicitly measured.
2. It demonstrates that changes in input modalities can have a substantial effect on the user's sense of agency.
3. It demonstrates that assistance from a computer can have a significant effect on the user's sense of agency.
4. It demonstrates an experimental means of mapping the personal agency characteristics of an assisted interaction technique.

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 '12, May 5–10, 2012, Austin, Texas, USA.

Copyright 2012 ACM 978-1-4503-1015-4/12/05...\$10.00.

5. It discusses the implications of these findings and provides an agenda for future research in this area.

From a more general perspective this paper demonstrates the benefits of applying methods derived from cognitive neuroscience to provide a more in depth understanding of interactions with technology [25]. We therefore start with a brief comparative review of diverse perspectives on agency in cognitive neuroscience and HCI literature.

THE SENSE OF AGENCY

Humans are agents. That is, we have the capacity to *intentionally* bring about change in the world through our actions. Humans also have a corresponding conscious experience of this capacity. This is the *sense* of agency. There is a distinction, then, between the fact of controlling an action and the immediate sense or experience of having done so. It is the experience of ourselves as agents that allows us to say, “*I did that*”.

Neurocognitive research on the sense of agency aims to shed light on the processes supporting this experience. Broadly speaking there are two theoretical positions. One view focuses on the motor system and suggests that the sense of agency is generated by neural process dedicated to the initiation and control of voluntary movement [16]. In this view, the sense of agency is based on the result of a comparison between the predicted and actual sensory consequences of movement [16]. If this comparison reveals no discrepancy then ‘I’ am the agent responsible for the particular event. The alternative theoretical position downplays the specific contribution of the motor system [36-38]. In this view, the sense of agency is produced by a more general-purpose cognitive system that monitors the relationship between thoughts (i.e. intentions), actions and their outcomes, with the mind inferring and reconstructing a path between conscious intention and effect [37]. Here the sense of agency is taken to be a reconstructive inference that one’s intention has caused an external event.

Within cognitive neuroscience much of the interest in the sense of agency stems from the fact that degradation of this experience is characteristic of certain psychiatric and neurological disorders. One of the most striking examples is the passivity phenomena in schizophrenia. Here, people feel that their actions (and sometimes their thoughts and emotions) are not under their own control, and are instead under the control of some external force or agent. For example, a patient described by Mellor [23] reports, “It is my hand and arm that move, and my fingers pick up the pen, but I don’t control them.” Such reports highlight a profound disruption in the mechanisms supporting the sense of agency, the consequences of which can be immensely disabling. However, the sense of agency is a remarkably fragile and malleable experience in all of us. There are instances in which we feel an exaggerated sense of agency. For example, people have been shown to place higher bets when they, rather than someone else, initiate a gamble, despite the odds being the same. On the flip side there are

instances, such as facilitated communication, table turning and Ouija boards, in which we feel a reduced sense of agency [36]. This malleability in the human experience of agency offers both opportunities and pitfalls for HCI.

AGENCY IN HCI

In subsequent sections we will introduce an implicit means of measuring the experience of personal agency. Before this, however, it is useful to briefly consider some of the diverse perspectives on agency in previous HCI research:

Media agency: together with collaborators, Reeves and Nass have undertaken an extensive body of research on the ‘Media Equation’ [31]. This work suggests that people treat computers as social *actors* [29] and respond to computers and other media in ways derived from their response to other people. In HCI, the creation of user interfaces which include embodied agents [5, 10] - photographs or animated characters - is one example of an explicit attempt to co-opt this response as a design resource that engages users and improves the interaction experience.

Intelligent agents/interfaces: in HCI contexts, components of intelligent user interfaces often demonstrate autonomous behaviours [21]. These systems use machine learning and other AI techniques to either aid users’ intended actions (e.g. [2]) or to act on users’ behalf (e.g. [4]). Such systems have the potential to blur the boundaries of user and machine agency and are the subject of the second experiment reported in this paper.

Design agency: semiotic analyses of user interfaces understand the designed product as a text, ‘read’ by users as a message from the designer. This understanding of the interface as the ‘designer’s deputy’ [9] relies on the recognition of a design stance, in which the designer as an agent is recognised to have had a particular intent in the presentation of the product [8].

Agency in the laboratory: in science and technology studies, the status of machines in laboratories as either neutral instruments, or objects of human observation, draws attention to the way in which they have characteristics attributed to them through their participation in social processes of knowledge production and interpretation [6]. Suchman in particular has drawn rich connections between agency of robots and software agents in scientific contexts, and the political and philosophical implications of artificial intelligence, robotics and HCI [34].

It is important at this point to highlight a significant difference between the concepts examined in this paper and those examined in prior HCI research. The Media Equation, for example, examines how users respond, or attribute characteristics, to computer and media agents. In this paper we do not focus on attributions of agency to other people or things. Instead we offer a new perspective, complementary to those described above, which focuses on how people experience their own sense of personal agency when interacting with technology.

INTENTIONAL BINDING: AN IMPLICIT METRIC

Building on the work of Libet et al. [22], Haggard et al. introduced an implicit metric through which the sense of agency can be empirically assessed [18]. Repeated experiments have shown that voluntary human actions are associated with systematic changes in our perception of time. The interval between a voluntary or intentional action and the action outcome is typically perceived as shorter than the actual interval. This is illustrated in figure 1(a). If a person takes a simple action - e.g. they press a button - and this action causes an outcome - e.g. a beep - it is highly likely that the person will perceive their action as having happened later than they it actually did (action binding). They are also likely to perceive the outcome of their action as having happened earlier than it actually did (outcome binding). Haggard coined the term ‘*Intentional Binding*’ to describe this phenomenon as it is contingent on several factors. In the absence of outcomes people are found to more accurately report the timing of actions. For the temporal binding effect to occur, an action must be *intentional* and it must lead to an *outcome*. Under these conditions our perception of the timings of our own actions and of their outcomes become bound more closely together.

Haggard also demonstrated that unintentional actions - e.g. an involuntary button press - cause changes in our perception of time. In this case however, the perceptual change tends to occur in the opposite direction (figure 1(b)). Now the interval between actions and outcomes is perceived as being longer than the actual interval. In experimental settings involuntary actions can be induced in a number of ways, ranging from transcranial magnetic stimulation to moving a person’s finger with a mechanical arm or even with a piece of string tied to their finger.

In the years since his first experiments [18], a large number of studies have been undertaken which validate and build on Haggard’s initial observations [19, 26-28]. The scientific consensus is now supports the conclusion that time perception in voluntary actions - and the binding effects associated with such actions - provides a robust implicit metric for the sense of agency. Higher intentional binding values correlate to a greater sense of personal agency.

Method 1 – the Libet Clock

In order to investigate intentional binding, Haggard developed an empirical method through which the binding effect can be implicitly assessed. He adapted Libet’s clock method [22] for measuring peoples’ perceived timing of events. The Libet clock (figure 2) consists of an analogue clock face with an arm that rotates clockwise through a




Figure 2: a Libet Clock

full cycle once every 2560ms. The clock is typically displayed at the centre of a screen and is quite small relative to the overall screen size. In our experiments for example the clock is 100 pixels in diameter and is displayed on a

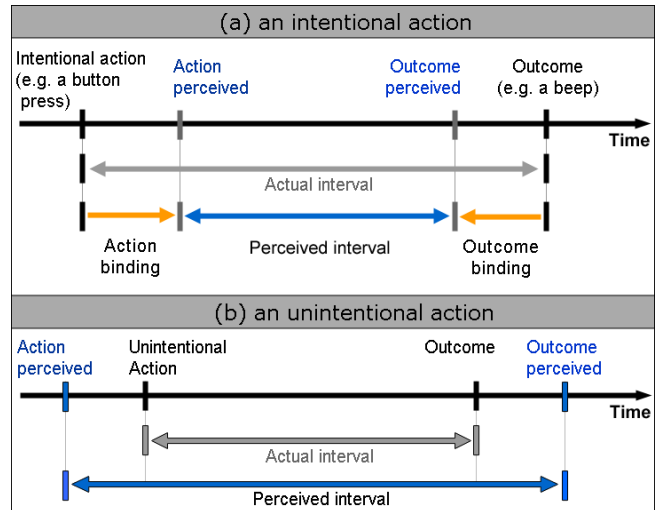


Figure 1: time perception in an intentional action (a), and an unintentional action (b).

Action measurements	
<p>Baseline error</p>	<p>The participant is asked to attend to the Libet clock and then press the button whenever they want to. In this case the button press does not cause a beep. The person reports the time on the clock (perceived time) when they pressed the button. The actual time of the button press is also recorded.</p> <p><i>Baseline error = actual time - perceived time</i></p>
<p>Active error</p>	<p>The participant attends to the clock face and presses the button whenever they want to. In this case the button press causes a beep. The person is asked to report the time on the clock (perceived time) when they pressed the button. The actual time of the button press is also recorded.</p> <p><i>Active error = actual time - perceived time</i></p>
Outcome measurements	
<p>Baseline error</p>	<p>The participant attends to the clock face. On this occasion the person takes no action and the computer randomly generates a beep. The participant reports the time on the clock (perceived time) when they heard the beep. The actual time of the beep is also recorded.</p> <p><i>Baseline error = actual time - perceived time</i></p>
<p>Active error</p>	<p>The participant attends to the clock face and presses the button whenever they want to. In this case the button press causes a beep. The participant reports the time on the clock (perceived time) when they heard the beep. The actual time of the beep is also recorded.</p> <p><i>Active error = actual time - perceived time</i></p>
<p>Each measure outlined above is repeated multiple times and the mean value for each is then used to calculate Intentional Binding.</p>	
<p>Intentional Binding calculations</p> <p><i>Action binding</i> = action error[active] - action error[baseline]</p> <p><i>Outcome binding</i> = outcome error[baseline] – outcome error[active]</p> <p><i>Total binding</i> = action binding + outcome binding</p>	

Table 1: the measurements and calculations used to assess Intentional Binding for an action/outcome condition, using the Libet Clock method. For illustrative purposes we use the example of a button press that causes a beep. A fixed action/outcome interval is used for all active trials.

screen with a resolution of 1920 x 1080 pixels. A person sits in front of the screen and is asked to report where the hand on the clock is pointing when certain events occur. The clock face is deliberately small so that the person can attend to the clock without making significant head or eye movements. Using the example of a button press, table 1 illustrates the four measurements and the calculations used to measure the intentional binding effect. In the active conditions there is a fixed action/outcome interval for all trials. Each measurement is repeated multiple times (typically 40) and mean time perceptions errors are then calculated. These mean values are then used to calculate the action, outcome and total Intentional Binding values.

The Libet clock method has a number of strengths and limitations. The most attractive feature is that it provides a robust measure of the total intentional binding effect. It can also distinguish and quantify both action and outcome binding. However, as it requires people to visually attend to a clock, this method is not suitable if the task being studied includes significant visual elements. It also requires four blocks of measures for each action/outcome condition and is time consuming to administer if the experimental design includes a large number of action/outcome conditions.

Method 2 – Interval Estimation

Subsequent to Haggard's initial work, Engbert et al introduced interval estimation [12], a second method through which the sense of agency can be implicitly assessed. As with the Libet clock method, this method involves measurements of the perception of time for actions that cause an outcome. In this case however people are simply asked to give an estimate of the time interval between the action and outcome. Trials are again repeated multiple times, but here the interval between the action and outcome is randomly varied. After each trial the participant gives their estimate of the time between the action and outcome. This estimate is then compared with the actual interval and the interval estimation error is calculated.

This method is considered less robust than the Libet clock method. Another limitation is that it cannot distinguish where a binding effect may be occurring; i.e. is it an action or outcome binding effect? However, it has the benefit of requiring just one measurement for each action/outcome condition under observation. It is also suitable for use with a wider range of experimental tasks, including designs that involve significant visual elements. Overall this method is best suited to studies comparing binding effects across experimental conditions, rather than studies aimed at robustly measuring binding for a given condition. It is also well suited to studies comparing a larger number of experimental conditions - as is the case in experiment 2 below. By comparing people's mean estimation errors in different interaction conditions it enables empirical comparisons of the sense of agency in different conditions.

We have now introduced two implicit methods through which people's sense of agency - defined as the experience

of control of actions and ownership of the outcomes of those actions - can be empirically investigated. Both methods make use of intentional binding, a phenomenon extensively examined in cognitive neuroscience literature, but not yet investigated or exploited in HCI research. In the following sections we present two experiments in which these methods have been used to empirically investigate peoples' experiences of interactions with technology.

EXPERIMENT 1 – WHAT HAPPENS WHEN THE INPUT MODALITY CHANGES?

At CHI 2010 Harrison et al. published an award winning paper on Skinput, a radical new input device that allows users to interact with technology by 'appropriating' their bodies, and more specifically their skin, as an input surface [20]. Once the device is attached to a person's arm they can control a computer by tapping different places on their hand or arm. For the appreciative conference audience this paper gave rise to many questions. In particular, one audience member asked: "*So, what's it like to be a button?*"

Whilst this question may appear frivolous at first sight, we believe it taps into an important issue for the HCI community. In recent years researchers have developed a wide range of new interaction techniques and devices, investigating concepts such as embodied [11], affective [30] and brain-computer [14, 39] interaction. These new approaches offer the potential to dramatically reshape the experience of interacting with technology. The question "what's it like to be a button?" challenges HCI researchers to think more deeply and provide a more rigorous understanding of these new experiences.

Inspired by the work of Harrison et al., our first experiment provides an empirical investigation of the question "what's it like to be a button". More generally this experiment asks if changes in the input modality of an interaction can have an effect on the user's sense of agency. As such, the techniques described in this experiment provide an empirical means of assessing and comparing the experiences provided by new interaction techniques. We compare a traditional input device - a keypad - with a skin-based input device.

Capturing skin-based interactions

For the purposes of this experiment we did not require a skin-based device as complex as that described by Harrison et al. We did however require that our device be able to capture the user's actions reliably, with a high degree of temporal accuracy and without requiring a training period. Our device, figure 3, bears a stronger resemblance to an earlier device described by Amento et al. [1]. A piezo-electric contact microphone is placed on the user's arm and connected to a computer through a high-fidelity external soundcard (Edirol UA-25EX). Vibrations in the user's arm are monitored using this microphone and analysed to detect when the user taps on their arm. Following a process of refinement and testing we were able to detect gentle arm taps with a reliability of greater than 95% and with no




Figure 3: a piezo-electric contact microphone is used to detect when the user taps on their arm.

training required for new users. The overall capture device and software is sufficiently robust that users can move their arms and hands without triggering false arm taps. It has a temporal accuracy of less than 3ms. This is equivalent to the temporal accuracy with which the MAC used in our experiment reports button presses from our USB keypad.

Procedure

Figure 4 illustrates the overall procedure for trials in this experiment. As the task undertaken by participants does not involve a significant visual component, we make use of the Libet clock method. One independent variable is manipulated: the input modality. In condition 1 - the button press condition - participants press a button on a keypad to cause a beep. In condition 2 - the skin-based condition - participants cause a beep by tapping on their arm. The skin-based capture device is attached to the participant’s left arm and they tap the inside of this arm with their right hand, (avoiding any sensitive carpal tunnel region), as in figure 3. There is a fixed interval of 250ms between the user’s action and the beep in both input conditions.

Participants sit at a desk, with both arms on the desk in each condition. They face a computer screen showing a Libet clock. To initiate a trial they press a footswitch. The onscreen clock begins to rotate. Participants are instructed to allow the clock to rotate for a short while and then make the action (button press or arm tap according to condition) whenever they want to. They then report the time on the clock face for one of the four measures outlined in Table 1.

The trial was repeated 40 times for each measure, with specific measures blocked together. A within-subjects design was used. Across the two input conditions this gave 320 trials per participant (160 per input condition). The measurement blocks for each condition were grouped together and the order of the input conditions was alternated for odd and even numbered participants. The order of the measurement blocks was randomised within the input conditions, but was balanced for odd and even participants.

This study was approved by the University of Cambridge, Computer Laboratory Ethics Committee. Participation was on the basis of written informed consent. The experiment lasted approximately 70 minutes for each participant.




Figure 4: the protocol for the active conditions in our first experiment. Baseline conditions omit either the action or beep.

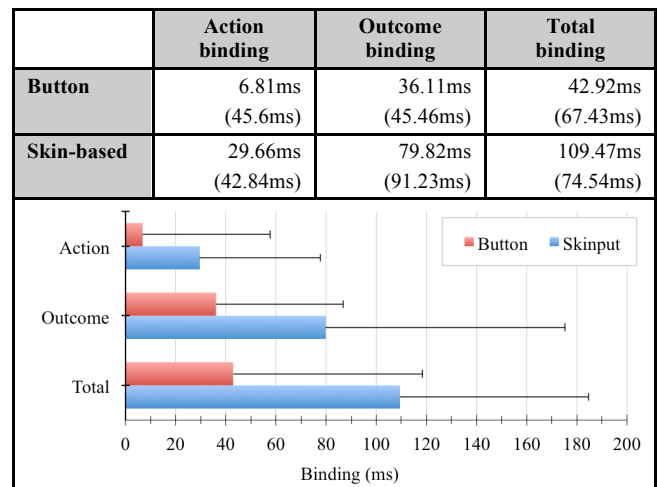


Table 2: the mean action, outcome and total binding times for each input condition. Higher binding values indicate a greater sense of agency. Standard deviations in brackets.

Participants

Participants for this experiment were recruited through an email to a University bulletin board. 21 participants - all right handed and aged 20 to 40 - completed the experiment. Three participants were excluded from our analysis. In two cases our capture device did not work effectively. A third participant incorrectly interpreted the experimental instructions. Each participant received a £15 gift certificate.

Results

Table 2 shows the action, outcome and total binding effects for the button press and skin-based conditions. These values represent the overall means across 18 participants. (A more detailed breakdown of specific measurements is in an Appendix at the end of the paper.) To test the statistical significance of our results we first conducted a 2x2 repeated measures analysis of variance comparing participants’ action and outcome binding values across the two input conditions. The results $F(1,17)=16,397, p<.001$, indicated a significant difference within the data. We then conducted a Bonferroni corrected paired sample t-test, with a null hypothesis of no difference, to compare the total binding values. The result shows a significant difference in binding between the two conditions, $t(18)=4.05, p<0.01$. Overall 15 of the 18 participants experienced greater binding in the skin-based condition than in the button press condition.

Two further paired sample t-tests were conducted to compare the baseline errors for the two input conditions. The comparison of baseline action errors gives $t(18)=0.753$, $p=0.461$, indicating no significant difference. The comparison of baseline outcome errors also indicates no significant difference, $t(18)=0.477$, $p=0.320$.

Analysis and Discussion

As the comparisons of baseline errors for the two input conditions indicate no significant differences, we can conclude that any binding differences between button and skin-based input are caused by differences in the active condition, i.e. the condition in which an action causes an outcome. With a fixed delay of 250ms between actions and outcomes participants experienced higher binding on both actions and outcomes for skin-based interactions. A total binding of 42.92ms was observed in the button input condition. It is worth noting that this value is consistent with the results of prior binding experiments that have used button inputs [19, 26]. In the skin-based condition participants experienced a total binding of 109.47ms. Given the correlation between intentional binding and the sense of agency, this experiment allows us to conclude that people experience a greater sense of agency when they interact with technology via skin-based input, as compared with traditional keypad input.

We believe that the finding in this experiment has a number of important implications for HCI research. It provides empirical evidence that new interaction techniques can provide different experiences of control and ownership to those offered by traditional mouse or keyboard interactions. The methods described also provide an implicit means of quantifying this difference. We chose to assess two dramatically different input conditions. In future the methods described here can be used to assess and quantify the differences for a larger range of interaction techniques, including changes more subtle than those assessed here. We could for example assess the difference in interactions with a touch screen via a stylus versus direct finger interaction, or differences in interactions that incorporate techniques such as haptic, embodied or physiological input.

In addition to comparing input techniques the method described here can be used to compare the different interaction experiences for a specific input technique when other conditions of the interactions change: e.g. when the user is experiencing different cognitive loads, when the environmental conditions of the interaction change or when different user groups interact with the same piece of technology. Greater consideration is given to these possibilities in our final discussion and conclusions section, but at this point we will move on to our second experiment.

EXPERIMENT 2 – WHAT HAPPENS WHEN A COMPUTER HELPS OUT?

Cognitive neuroscience experiments on intentional binding have typically examined voluntary and involuntary actions. From an HCI perspective, this might be considered an

unnecessarily black or white disjunction. Many user interactions with technology are more intermediate. In particular ‘intelligent’ user interfaces often seek to interpret and act on the intentions of the user. The user’s actions are voluntary, but the outcomes may be assisted. Our second experiment is designed to investigate what happens to the user’s sense of agency when a computer interprets their intention and helps them to achieve their goal.

Assistance from a computer

Our experiment investigates agency in a machine-assisted point-and-click task. We chose this task because pointing is a basic and ubiquitous action in user interfaces. As a consequence, several pointing enhancement techniques have been designed and evaluated [2]. One of our primary concerns in designing this experiment was model simplicity. As this is the first study of machine-assisted binding we chose a machine-assisting pointing model with as few free parameters as possible. This allows us to manipulate the assistance given to participants, while at the same time having to change a minimum number of system parameters. In our experiment participants use a mouse and are asked to hit targets on the screen as quickly and as accurately as possible. Participants are assisted by an algorithm that has the effect of adding “gravity” to targets, causing the mouse pointer to be attracted to the spatially closest target. Our algorithm models horizontal and vertical movement but updates them independently of each other. Let x be a horizontal or vertical component of the user’s mouse position and let x_t be the corresponding horizontal or vertical component of the target position with the shortest Euclidean distance to the user’s mouse position. At each mouse position update from the user we compute a new corresponding horizontal or vertical component x' according to the following model:

$$x' = x + \text{sgn}(\Delta x_t) \alpha \left(1 + \exp \left[-\frac{|\Delta x_t|}{\beta} \right] \right),$$

where $\Delta x_t = x_t - x$. Since the algorithm is one-dimensional it is applied twice: once for the horizontal and once for the vertical update of the mouse position. This model has α and β as two free parameters that can be manipulated to adjust the level of assistance participants receive. The parameter α controls the extent of pixel movement towards the target that can be applied at each update. The parameter β controls the degree of change of gain as a function of the distance to the target and typically models a range corresponding to expected differentials between initial mouse positions and target positions. In our experiment (see Figure 5) the distance from an initial start area to a target does not change and we use a fixed β -value of 800. This gives us α as a single parameter that we can then manipulate to change the level of assistance participants receive in our task. Higher values of α correspond to higher levels of assistance.

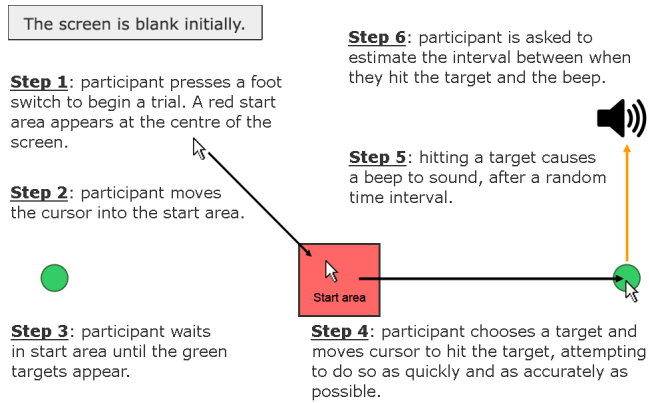


Figure 5: the overall procedure for the point and click task.

Procedure

Figure 5 describes the overall procedure for trials in this experiment. As this experiment involves significant visual elements we make use of the interval estimation method to implicitly assess people's sense of agency. Participants sit at a desk facing a computer screen that is initially blank. They press a footswitch to start a trial. This causes a red start area to appear at the centre of the screen. Participants use a mouse to move the cursor into the start area. After waiting in the start area for a short period two green targets appear, one to the left and one to the right of the start area. The targets are equidistant from the centre of the start area. The participant is asked to hit one of the targets (an action) as quickly and as accurately as they can. Hitting a target causes a beep (an outcome), with a randomised time interval between hitting the target and the beep. The participant is then prompted to estimate the time between hitting the target and hearing the beep.

Two targets are provided (left and right) as this allows the participant to make a voluntary choice in each trial. The computer has to interpret this choice before giving assistance. Participants are given a number of important instructions about the targeting task prior to the experiment:

- Participants are told that there is no difference between the two targets. They can choose the target they want to hit in each trial. They are asked not to hit the same target each time or simply alternate between the targets.
- Participants are specifically told that the experiment is not designed to test their reflexes. Participants are told not to move toward the green targets as soon as they appear. Rather they are told they can wait in the start area for as long as they want, moving toward their chosen target whenever they are ready.
- Having begun the targeting action, participants are told to hit the target a quickly and as accurately as possible.

It is important to note that we are not primarily interested in the speed with which participants can hit targets. Rather we are interested in differences in participants' interval estimates when our independent variable - the assistance level - is changed.

We tested four distinct levels of computer assistance: *no assistance*, *mild assistance*, *medium assistance* and *high assistance*. In the no assistance condition the gravity algorithm is not used and the computer gives the participant no assistance. In the mild assistance condition the α parameter in the gravity algorithm is set to 3. In the medium assistance condition α is set to 6. In pilot studies we found that many people were unaware of the different assistance levels given in these conditions. In both the mild and medium conditions the computer only assists the user if they are actually moving the mouse. In the high assistance condition the cursor moves quickly to a target once the user leaves the start area, even if the user stops moving the mouse. The α parameter is set to 9. Our aim at this level was to make the assistance obvious to participants. Given the potential for ordering effects in this experiment, we chose to have 24 participants, thus allowing full counterbalancing of the order of the four assistance level blocks.

In designing this experiment we following the procedure for interval estimation experiments described by Ebert and Wegner [12]. The experiment has a within-subjects design, with each participant completing one block of trials for each of the four assistance levels. Rather than using an entirely random interval between actions and outcomes they recommend the use of three fixed intervals: 150ms, 400ms and 700ms. Participants are told that the beep interval is fully random and that it ranges from 50 to 950ms in steps of 50ms. Each of three intervals - 150ms, 400ms and 700ms - was randomly repeated 12 times per assistance level, giving totals of 36 trials in each block and 144 trials per participant. Participants use an input box with radio buttons to give their estimates of time intervals.

Prior to beginning the experimental trials each participant undertook a short practice period involving 12 trials. Participants pressed a button on a computer screen. This triggered a beep and they estimated the interval between pressing the button and beep. They were then told the actual interval. In the practice trials the interval did vary randomly between 50ms and 950ms. These trials served multiple purposes. Firstly, they were intended to improve participants' accuracy by giving them a better sense of the brief intervals used in the experiment. Second, by exposing participants to a full range of random delays, we aimed to reinforce their expectation that the experimental trials would also involve randomised delays.

This experiment was run on a desktop PC running Windows 7. Participants used a Microsoft optical USB mouse. The pointer sensitivity was set to the default level for Windows 7, with enhanced pointer precision switched off throughout all trials. The red start area was 100 pixels high and wide. The green targets each had a diameter of 35 pixels. We used a 24-inch widescreen monitor with a resolution of 1920x1080 pixels. The centre of each target was 810 pixels from the centre of the start area. Participants

were excluded if they repeatedly missed targets or if they hovered over targets for long periods before clicking them.

Whilst this experiment was primarily focused on an empirical assessment of the sense of agency, participants were given an opportunity to comment on their experiences. Upon completion of each of the four blocks of trials an experimenter asked participants: “Did you notice any problems with that block of trials?” This question was deliberately open-ended and did not make any reference to agency or assistance from the computer.

This study was approved by the University of Cambridge, Computer Laboratory Ethics Committee. Participation was on the basis of written informed consent. The experiment lasted approximately 50 minutes for each participant.

Participants

Participants were recruited through an email to a University bulletin board. In all 27 participants - all right handed and aged 20 to 40 - completed the experiment. Three were excluded from our analysis. Two participants hovered over the targets for long periods (>500ms) and one missed the target 18 times in one experimental condition. This left us with 24 participants, with the order of the experimental blocks fully counter-balancing for the four assistance levels. Each participant received a £15 gift certificate.

Results

Table 3 shows the mean interval estimation errors and standard deviations for each assistance level. These values represent the overall means across 24 counter balanced participants. Lower estimation errors indicated a higher sense of agency. To test the statistical significance of our results we first conducted a repeated measures analysis of variance that compared the interval estimation errors of participants across the four assistance levels. The result, $F(3,69)=2.74, p=0.05$, indicates a significant difference within the results. Bonferroni corrected paired sample t-tests were then conducted between the successive assistance levels, with a null hypothesis in each case of no difference in the estimation errors. We do not find a significant difference between the no assistance and mild assistance conditions, $t(24)=0.036, p=0.97$. Similarly we do not find a significant difference between the medium and high assistance levels, $t(24)=0.419, p=0.679$. There is however a significant difference in estimation errors between the mild and medium assistance levels, $t(24)=3.08, p < .01$.

Analysis and Discussion

Prior to undertaking this experiment our initial expectation was that – if assistance from a computer did indeed have a measureable effect on peoples’ sense of agency – we would find a gradual effect, with successively higher levels of assistance leading to a successively lower sense of agency. The results paint a more complex and compelling picture.

This experiment used interval estimation as an implicit metric for participants’ sense of agency. The results of our experiment indicate a high sense of personal agency

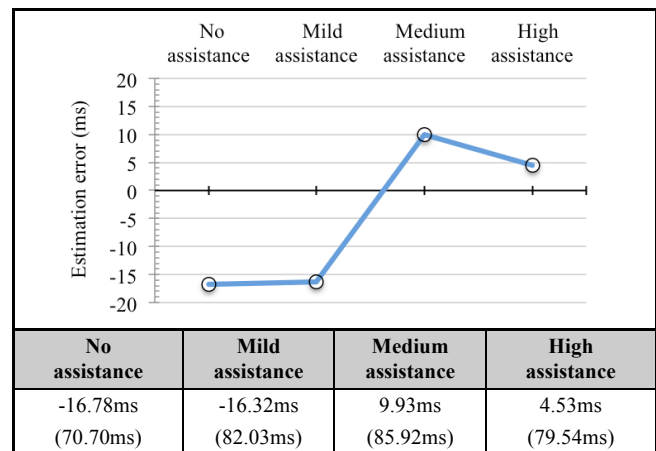


Table 3: the mean interval estimation errors in milliseconds for each assistance level. Lower estimation errors indicate a greater sense of agency. Standard deviations in brackets.

(negative estimation errors) in the no assistance condition, where participants completed the task on their own with no help from the computer. They also indicate a similarly high sense of agency when the computer gave them a mild level of assistance, with no significant difference in participants’ interval estimation errors at the no assistance and mild assistance levels. When the assistance level was taken one step further, to the level we have termed medium assistance, results indicate a significant loss in the sense of agency, as indicated by positive interval estimation errors. There was no significant difference between participants’ estimation errors at the medium and high assistance levels.

Participants’ comments in response to our queries about problems after each block of trials strongly indicate that they were aware of differences between the no assistance level and each of the three assisted levels (mild, medium and high). This suggests that while participants were aware of the assistance from the computer in the mild assistance condition, they still experienced a high sense of agency.

This overall result has a number of important implications. Firstly it shows empirically that, up to a certain point, the computer could assist users whilst also allowing them to maintain a sense of control and ownership of their actions and the outcomes of those actions. Secondly it suggests that beyond a certain level of assistance users experienced a detectable loss in their sense of agency. This experiment suggests that for this particular assisted input algorithm - and possibly for assisted input systems more generally - there may exist a tipping point or sweet spot. This is the point at which a computer can help people and potentially maximise task performance - e.g. speed or accuracy - without significant detriment to the experience of agency. A third important implication of this work is that the methods applied here can be used to map the personal agency characteristic of computer assisted interactions techniques. In this experiment we investigated four distinct assistance levels - with α values of 0, 3, 6 and 9 in our gravity algorithm. In future experiments a more fine-grained

analysis - testing further intermediate α values e.g. 1, 2, 4, and 5 - would provide a more detailed mapping of users' sense of agency when interacting with the system.

Our primary aim in this experiment was to investigate if changes in assistance from a computer led to experimentally measurable changes in peoples' sense of agency. However, as was the case with the first experiment in this paper, we believe the methods described here have a broad range of applications. They could be used to compare different machine learning/user assistance algorithms. They could also be used to refine the parameters of a particular assistance technique or algorithm. As stated previously, we focused on simplicity in designing this experiment. One of our choices was to have an experimental design in which it was extremely unlikely that the computer would incorrectly interpret the user's intentions and therefore hinder, rather than assist, them. In future we believe it will be interesting to use the techniques described here to investigate the effects on people's sense of agency when 'intelligent' systems make mistakes and incorrectly interprets users' intentions. For example, what quantifiable impact does a violation of the user's intention have on their experience of control and, once lost, how long does it take before users recover their sense of control? It may also be possible to implicitly investigate techniques that mitigate losses in the sense of control. The methods described in this experiment can provide a valuable means of empirically addressing such questions and thereby help developers to make informed design trade-offs.

DISCUSSION, CONCLUSIONS AND FUTURE WORK

Personal agency is a crosscutting experience essential to many aspects of our day-to-day lives. It links to concepts such as control, ownership, responsibility and causality, and has a significant impact on how we perceive and interact with our world and with the people and things in it. In this paper we have sought to provide a new insight into the experience of personal agency as it relates to interactions with technology. In doing so we have drawn upon prior research in cognitive neuroscience and have introduced two implicit measures through which the experience of agency can be empirically investigated. These methods have been extensively investigated in cognitive neuroscience literature, but, prior to this paper, they have not been exploited in HCI research.

Measurement and evaluation are key components of HCI research and practice. Measurement enables comparison, which enables us to make systematic improvements and consider trade-offs in the design of systems. Prior HCI research has provided a rich set of techniques through which many task performance aspects of system usability - e.g. efficiency, error rate, memorability - can be evaluated. In this paper we have demonstrated a means through which an underlying experience - the sense of agency - can be empirically investigated. Interfaces and interaction techniques that facilitate a high sense of personal agency

are likely to have a strongly empowering effect for users. They will allow people interacting with technology to instinctively sense that "*I did that*".

Cognitive neuroscience research has demonstrated that the experience of agency is highly malleable. In this paper we have demonstrated two specific ways in which changes in interactions can significantly affect the user's experience of agency. We believe this work has tapped into a research area rich in possibilities. We chose to focus on input modality and assistance from a computer, and have offered specific examples of the future research questions in each area. However, a wide range of factors relevant to HCI research have been shown to have an effect on the experience of personal agency. For example, distortions in feedback or inconsistencies across different feedback modalities can be detrimental [15, 28] and it will be interesting to investigate and compare the impact of different feedback techniques on user agency. Uncertainty and contingency [26-28] also have an impact and are worthy of investigation in the context of HCI research on uncertain control [39].

Implicit measures of the experience of agency will also be relevant to cross-population investigations of interaction experiences. For example, studies suggest that this experience changes across the lifespan. In particular older groups may experience substantial differences on their sense of agency [24]. It will be interesting to examine the design factors that maximise the sense of personal agency in such groups. Giving the growth in HCI research on health and wellbeing, and in research that addresses mental health [7] and seeks to support people experiencing psychotic difficulties [3], we believe the methods described in this paper can also play a valuable role in understanding the unique requirements of these user groups.

Finally, this paper has focused on implicit measures for the sense of agency. Research in cognitive neuroscience has provided robust evidence for the validity of these methods and they have the benefit of enabling researchers to empirically assess the experience of agency as it happens. In future we also believe it will be interesting to combine these methods with subjective methods, in which users are explicitly asked to report on their experiences of control and ownership. This will allow us to provide a deeper understanding of users' preferences regarding the experience of agency in human computer interactions.

ACKNOWLEDGMENTS

David Coyle's work was supported by an IRCSET-Marie Curie International Mobility Fellowship in Science, Engineering and Technology. Per Ola Kristensson was supported by the Engineering and Physical Sciences Research Council (grant number EP/H027408/1). Paul Fletcher is funded by the Wellcome Trust and the Bernard Wolfe Health Neuroscience fund and is part of the Behavioural and Clinical Neuroscience Institute.

APPENDIX: MEASUREMENTS IN EXPERIMENT 1

	Action errors		Outcome errors	
	Baseline	Active	Baseline	Active
Button	-4.10ms	2.71ms	16.46ms	-19.64ms
Skin-based	3.43ms	33.08ms	14.28ms	-65.54ms

Action binding button = 2.71 - (-4.71) = 6.81ms
 Outcome binding button = 16.46 - (-19.64) = 36.11ms
 Action binding skin-based = 33.08 - 3.43 = 29.66ms
 Outcome binding skin-based = 14.28 - (-65.54) = 79.82ms

Table 4: baseline and active errors in the perceived timing of events in experiment 1.

REFERENCES

- Amento, B., Hill, W., & Terveen, L. *The sound of one hand: a wrist-mounted bio-acoustic fingertip gesture interface*. In Proc ACM CHI Extended Abstracts 2002, p.724-725.
- Balakrishnan, R., "Beating" Fitts' law: virtual enhancements for pointing facilitation. *Int Journal of Human-Computer Studies*, 2004. 66(6): p.857-874.
- Bickmore, T. & Gruber, A., *Relational Agents in Clinical Psychiatry*. Harvard Review of Psychiatry, 2010. 18(2).
- Boone, G. *Concept features in Re:Agent, an intelligent Email agent*. In Proc 2nd Intl. Conf. on Autonomous Agents 1998.
- Cassell, J., Sullivan, J., Prevost, S., & Churchill, E.F., eds. *Embodied Conversational Agents*. 2000, MIT Press.
- Collins, H.M. & Kusch, M., *The Shape of Actions: What Humans and Machines Can Do*. 1998, MIT Press.
- Coyle, D., Doherty, G., Matthews, M., & Sharry, J., *Computers in Talk-Based Mental Health Interventions*. *Interacting with Computers*, 2007. 19(4): p.545-62.
- Crilly, N., Good, D., Matravers, D., & Clarkson, P., *Design as communication: exploring the validity and utility of relating intention to interpretation*. *Design Studies*, 2008. 29: p.425-57.
- de Souza, C.S., *The Semiotic Engineering of Human-Computer Interaction*. 2005, Boston, MA: MIT Press.
- Dehn, D.M., *The impact of animated interface agents: a review of empirical research*. *International Journal of Human-Computer Studies*, 2000. 52: p.1-22.
- Dourish, P., *Where The Action Is: The Foundations of Embodied Interaction*. 2001: MIT Press.
- Ebert, J.P. & Wegner, D.M., *Time Warp: authorship shapes the perceived timing of actions and events*. *Consciousness and Cognition*, 2010. 19: p.481-89.
- Engbert, K., Wohlschläger, A., Thomas, R., & Haggard, P., *Agency, subjective time, and other minds*. *J Exp Psychol Hum Percept Perform*, 2007. 33(6): p.1261-8.
- Fairclough, S.H., *Fundamentals of physiological computing*. *Interacting with Computers*, 2009. 21(1): p.133-145.
- Farrer, C., Bouchereau, M., Jeannerod, M., & Franck, N., *Effect of distorted visual feedback on the sense of agency*. *Behavioural Neurology*, 2008. 19: p.53-57.
- Frith, C.D., Blakemore, S.J., & Wolpert, D.M., *Abnormalities in the awareness and control of action*. *Philosophical Transactions of the Royal Society of London – B: Biological Sciences*, 2000. 355: p.1771-1788.
- Gallagher, S., *Multiple aspects of agency*. *New Ideas in Psychology*, 2010. In press(Available online, March 2010).
- Haggard, P., Clark, S., & Kalogeras, J., *Voluntary action and conscious awareness*. *Nature Neuroscience* 2002. 5(4).
- Haggard, P. & Tsakiris, M., *The Experience of Agency: Feelings, Judgments, and Responsibility*. *Curr Dir Psychol Sci*, 2009. 18(4): p.242-46.
- Harrison, C., Tan, D., & Morris, D. *Skinput: appropriating the body as an input surface*. In Proc ACM CHI 2010, p.453-62.
- Lieberman, H. *Autonomous Interface Agents*. In Proc ACM CHI 1997, p.67-74.
- Libet, B., Gleason, C.A., Wright, E.W., & Pearl, D.K., *Time of conscious intention to act in relation to onset of cerebral activity (readiness potential): The unconscious initiation of a freely voluntary act*. *Brain*, 1983. 106: p.623-642.
- Mellor, C.S., *First rank symptoms of schizophrenia*. *Br J Psychiatry*, 1970. 117: p.15-23.
- Metcalfe, J., Eich, T.S., & Castel, A.D., *Metacognition of agency across the lifespan*. *Cognition*, 2010. 116: p.267-82.
- Minnery, B.S. & Fine, M.S., *Neuroscience and the Future of Human-Computer Interaction*. *Interactions*, 2009. 16: p.70-75.
- Moore, J. & Haggard, P., *Awareness of action: inference and prediction*. *Consciousness and Cognition*, 2008. p.136-144.
- Moore, J.W., Lagnado, D., Deal, D.C., & Haggard, P., *Feelings of Control: contingency determines experience of action*. *Cognition*, 2009. 110: p.279-83.
- Moore, J.W., Wegner, D.M., & Haggard, P., *Modulating the sense of agency with external cues*. *Consciousness and Cognition*, 2009. 18: p.1056-64.
- Naas, C., Steuer, J., & Tauber, E. *Computers are Social Actors*. In Proc ACM CHI 1994, p.72-78.
- Picard, R.W., *Affective Computing*. 1997: MIT Press.
- Reeves, B. & Nass, C., *The Media Equation: How People Treat Computers, Television, and New Media like Real People and Places*. 1998: Cambridge University Press.
- Rotter, J.B., *Generalized expectancies of internal versus external control of reinforcements*. *Psychological Monographs*, 1966. 80(609).
- Shneiderman, B. & Plaisant, C., *Designing the User Interface: Strategies for Effective Human-Computer Interaction*. 4th ed. 2004: Pearson Education.
- Suchman, L., *Human-Machine Reconfigurations*. 2007: Cambridge University Press.
- Synofzik, M., Vosgerau, G., & Newen, A., *Beyond the comparator model: a multifactorial two-step account of agency*. *Consciousness and Cognition*, 2008. 17: p.219-39.
- Wegner, D., *The Illusion of Conscious Will*. 2002 MIT Press.
- Wegner, D., *The mind's best trick: how we experience conscious will*. *Trends in Cog Science*, 2003. 7(2): p.65-69.
- Wegner, D. & Wheatley, T., *Apparent Mental Causation: Sources of the Experience of Will*. *American Psychologist*, 1999. 54(7): p.480-92.
- Williamson, J., Murray-Smith, R., Blankertz, B., Krauledat, M., & Muller, K.R., *Designing for uncertain, asymmetric control: Interaction design for brain-computer interfaces*. *Int. J. Human-Computer Studies*, 2009. 67: p.827-41.