Agency Attribution in Human-Computer Interaction

John E. McEneaney

Department of Reading and Language Arts
Oakland University, Rochester, MI USA
mceneane@oakland.edu

Abstract. Social psychologists have documented that people attribute a human-like agency to computers. Work in human motor cognition has identified a related effect known as "intentional binding" that may help explain this phenomenon. Briefly, intentional binding refers to an unconscious attribution of agency to sufficiently complex entities in our environments that influences how we perceive and interact with those entities. Two studies are presented that examine whether intentional binding, an agency effect observed when people interact with physical objects, also applies in virtual environments typical of human-computer interaction (HCI). Results of the studies indicate that agency effects are observed in human-computer interaction but these effects differ from those reported in physical environments. Results of the studies suggest that human perception and action may operate differently in virtual environments than in physical interactions.

Keywords: social interface theory, intentional binding, cognition, perception, agency attribution.

1 Introduction

A substantial body of work in social interface theory [1] has demonstrated that people often attribute a human-like agency to computers. Gender bias, for example, has been observed depending on whether the voice used by the computer is male or female [2]. Computer users are more likely to disclose personal information to a computer if that computer has disclosed "personal" information about itself [3]. Participants in HCI studies have also been observed to adopt a bias attributing credit or blame to a computer depending on whether the computer is perceived to share qualities with the experiment participants [4] and people using an online learning system that fails by design during an experiment tend to soften criticism of the failure when they report the problem directly to the computer rather than an independent human experimenter [5]. Furthermore, results from work in affective computing suggest that the perception of agency can be influenced by interface factors [6]. Our understanding of agency attribution by computer users, however, is limited. This work examines the attribution of agency in tasks typical of HCI.

Recent neurocognitive work exploring human agency has, perhaps not surprisingly, focused on motor behaviors. Motor cognition is better understood than higher-order cognitive phenomena because the behaviors studied are defined in more precise

D. Harris (Ed.): Engin. Psychol. and Cog. Ergonomics, HCII 2009, LNAI 5639, pp. 81–90, 2009.

ways and the mapping of brain and motor behavior is clearer and more direct. Correlations between specific motor behaviors and activation in the motor cortex are reliable enough for meaningful generalization and even simple motor behaviors such as moving a single finger still incorporate an essential feature of intentional behavior – the subjective experience of free action. Moreover, the history of work exploring connections between motor behavior and agency dates back to the early 1980s and during this time researchers have defined a variety of techniques to assess the subjective experience of agency and its neural correlates. Among the earliest of these studies was work by Libet and coauthors [7], who sought to establish a timeline relating central neural activity in the brain, peripheral neural activity of muscles, and the conscious subjective experience of action.

Libet used scalp electrodes to chart the broadly distributed electrical activity of the brain known as the readiness potential or RP [8], a finger electrode to measure muscle-specific electrical activity, and a clock with a single rapidly moving hand so that subjects could report when they became aware of the decision to move their finger. Libet instructed his subjects to move a finger of their own free will, using the position of the clock hand to indicate when they had made the decision to move. Libet's goal was to determine the chronology of these three events and, from a neurological perspective, his results were unsurprising. Briefly, the RP associated with a finger movement usually emerged from a subject's neural baseline about 800 milliseconds before the finger movement, the subject became aware of the decision to move about 200 ms before the movement, and the increase of electrical activity in the finger muscle began about 50 ms before the subject's finger movement triggered an electrical switch. Libet's study, however, stirred a debate among philosophers of mind as it provided rather compelling evidence that the subjective experience of free will (at least with respect to motor volition) followed, rather than preceded, the broader brain activity (i.e., the RP) responsible for the motor action. Apparently, a brain has a mind of its own.

More recent work, relying on Libet's clock method, explores social aspects of the experience of agency and has its origins in research with non-human primates. Recent neurocognitive work [9] suggests biologically plausible mechanisms to support agency-specialized neural structures (mirror neurons) that might account for cognitive correlates of agency attribution. In the mid-1990s researchers at the University of Parma were engaged in mapping brain function in macaque monkeys [10]. As a part of this research a specific neuron had been mapped that always fired when a monkey picked up a peanut that had been placed on a surface before it. Every time the monkey acted in this manner, the neuron fired, suggesting that the brain activity corresponded to "I (the monkey) am picking up this peanut." One day a monkey, fully wired and waiting to begin a study, watched as a researcher came and picked up a peanut from the table and the same "I am picking up the peanut" neuron fired. Apparently, the same neural circuit could mean "*I* am picking up the peanut," as well as "*You* are picking up the peanut". This neural circuit did not distinguish between these two rather different situations although the same neuron did *not* fire when the monkey watched a mechanical arm pick up a peanut [11]. It appears that this neural activity does not simply represent perception or motor response; it seems to represent agent-initiated action. Research since Rizzolatti's serendipitous discovery has confirmed both the presence of these so-called mirror neurons in people and Rizzolatti's initial hunch that mirror neurons seem to be related to a wide range of social behaviors [12].

Other work in motor cognition [13] provides independent support for the idea that the brain has specialized processes for perceiving the actions of other agents. Study participants were instructed to perform simple arm movements (rhythmic side-to-side or up-and-down arm motions) while watching another person or a robot perform similar or different movements. The hypothesis driving the study was that “actions are intrinsically linked to perception (p. 522)” and that, if the mirror system is activated, this should increase the likelihood of imitative behavior. Subjects who watched another person make incongruent movements (e.g., up-and-down when the subject was moving side-to-side) had significantly more variability in their movements than when the other person was moving congruently. Furthermore, arm movements by people in the study were *not* influenced by watching a robot make incongruent movements. This interference effect brought on by observation of another agent demonstrates that perception and action are linked and, more importantly, that perceived agency can influence action.

Another extensive line of work exploring agency effects has shown that subjective timing of events depends on whether or not people attribute agency to the source of an observed action. Briefly, our perceptual experiences are subtly different when we watch ourselves or other people than when we watch non-intentional events. One expression of this agency effect is a perceptual shift that delays perception of an event relative to its actual time when the event is perceived as an intentional action. A second manifestation of this agency effect is a perceptual shift that anticipates perception of an event that is perceived to be a consequence of intentional action. When, for example, a buzzer sounds as a result of pressing a button, the perceived timing of the button press is delayed while the timing of the buzzer is reported as occurring earlier (see Fig. 1). These same effects are also observed when subjects watch other people press the button. Taken together, these two effects have been referred to as “intentional binding” [14].

Prior work on intentional binding, however, has focused on *physical* environments where people observe the movements of machines, physical objects, or human hands. In these studies, the subjective experience of actions and effects are reliably influenced by the agency effect illustrated in Figure 1. The subjective timing between intentional actions and their effects is consistently reduced and this effect is not observed in circumstances where subjects watch objects or even their own fingers when their movements are under external control. Although work by social psychologists indicates people attribute a kind of agency to computers, it is unclear whether we can expect to see the subtle perceptual effects of intentional binding when events and consequences are represented in *virtual* environments typical of HCI. The following studies were designed to address this question: Is intentional binding observed in the virtual environments typical of HCI?

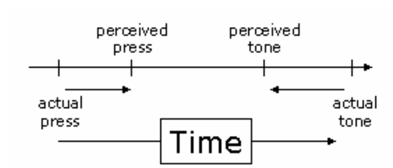

Fig. 1. Intentional binding of an action and its effect

2 Experimental Procedures

Both studies relied on the same experimental procedures and stimuli (see Figure 2). The only difference between the two studies was the target of the timing task. In study 1, participants were asked to report the timing of a red flash that coincided (within an average of 8 milliseconds) with a self-, other-, or machine-initiated mouse click. In study 2 participants reported the timing of an auditory tone that followed a mouse click by 250 milliseconds (ms). Subjective timing relied on Libet's clock methodology [7] [15]. Experimental materials were developed in Python using PsychoPy [16], a collection of Python modules that provide access to hardware-level system functions for precision stimulus display and response timing. Participants reported the observed position of a rotating spot at the time of a mouse click, a visual stimulus, or an audible tone.

Participants indicated the position of the spot by clicking the clock where the spot was at the time of the target event. For example, a participant who saw the rotating spot in the position shown in Fig. 2B when the clock flashed red will move the cursor and click the "28-minute" mark on the clock (Fig. 2C). In the other-click condition, the participant watches another person (a researcher) interact with the computer, again reporting the timing of a target event. Finally, in the two computer-click conditions the participant started the trial by clicking the "Next Trial" button but, once initiated, the computer controlled cursor movement and clicks. In the machine-simulated-click condition the movement of the cursor and timing of the click replicated the movement and timing of one of the participant's own trials randomly selected from data collected during training trials. In the machine-random-click condition the cursor was invisible so there was no cursor movement and timing of the target event was random. The first set of trials in both studies was a spontaneous self-click condition requiring participants to report when they had clicked the mouse. This set of trials was designed to familiarize participants with the experimental protocol and assess whether the mouse click and visual signal were subjectively perceived as simultaneous. Presentation of the four subsequent blocks of experimental trials employed a Latin square with conditions presented in random order, with each condition occurring once in each position.

A. When a trial begins, a rotating spot appears at a random position. The task is to report the position of the spot when a target event occurs.

B. The target in study 1 was a red flash when the mouse was clicked. Study 2 subjects reported the timing of a tone following the flash by 250 ms.

C. After the tone, the spot continues for a random interval then disappears. Subjects click to show where the spot was at the time of the target event.

Fig. 2. The Libet clock used in the studies at three points in a trial

All trials in both studies presented the same sequence of events. The only differences between blocks of trials and studies were the timing of events (which could be determined by the participant, the experimenter, or the computer) and the target of the timing task (i.e., a click, red flash, or audible tone). Each experimental trial was initiated by the participant or researcher clicking a "Next Trial" button, starting the clock rotating from a random starting position. Participants were asked to let the spot make one complete revolution, move the cursor to the center of the clock, and click the mouse at a time of their own choosing. When the clock was clicked, its center flashed red for 100 milliseconds, providing an immediate virtual "action signal" when the mouse was clicked. In addition, 250 ms after the click there was a 1000 Hz auditory tone that sounded for 100 milliseconds, providing an action effect that always followed the click and visual flash. After these events occurred, the spot continued to rotate for a random period of time that ranged from 1.5 to 2.5 seconds so that participants would not be aided by after-image effects.

Both studies adopted fully-factorial one-way repeated measures designs with four levels of an agency factor. Study 1 focused on a visual cue associated with the mouse click. In the self-click (SC) condition, the participant reported the timing of the visual cue associated with a self-initiated mouse click. In the other-click (OC) condition, the participant observed an experimenter click the mouse and reported the visual cue associated with that action. In addition to these two "human" conditions, there were two "machine" conditions. In the machine-simulated-click (MSC) condition the participant observed a computer-controlled movement of the cursor and a mouse click that matched the timing and motion of trials randomly selected from the participant's data in the training block. Lastly, in the machine-random-click (MRC) machine condition, the timing of the mouse click was random and the cursor was not visible; the participant simply reported the visual cue when it appeared. Materials and conditions in Study 2 were identical to those used in Study 1. The only difference in Study 2 was that participants reported the timing of the auditory tone that followed the mouse click rather than the red flash.

3 Study 1 Results and Discussion

Results of a repeated-measures GLM ANOVA indicated an agency effect, with significant differences in both the general test (Hotellings Trace $F(3,32) = 5.62$, $p = .003$, $\eta^2 = .345$) and in the within-participants test corrected for possible violation of the sphericity assumption (Huyhn-Feldt $F(2.124,72.221) = 4.527$, $p = .013$, $\eta^2 = .118$) with observed power of .915 and .773 respectively. Mean error scores across the four conditions are illustrated in Figure 3.

Follow-up analysis confirmed that the observed agency effect could be attributed to differences between the human and machine conditions. A paired t-test examined the data with self- and other-click conditions collapsed into a single "human" condition and the computer-controlled conditions treated as a single "machine" condition. Results of the paired t-test revealed a significant difference $t(34) = -2.511$, $p = .017$. Further paired t-tests showed no significant differences within the two human and machine conditions. Results of Study 1 replicate prior findings that perceptual judgments of observed actions are influenced by whether or not participants attribute agency to the source of those actions. Unlike prior work assessing intentional binding in physical environments, however, the effects observed showed a greater delay in perceiving *machine* initiated action than that of human action, the opposite of findings in physical environments. Prior work has consistently documented relative delays in perception of self- and other-initiated (i.e., human) actions compared to machine generated events. In study 1, however, this pattern of timing errors is *reversed*.

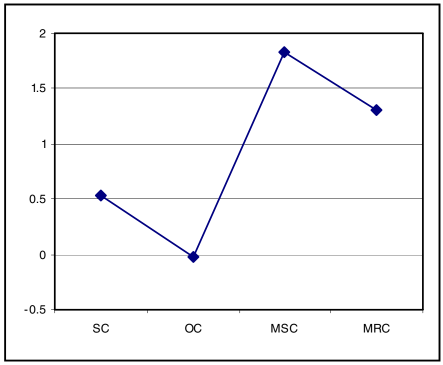

Fig 3. Mean error scores from Study 1 for self-click (SC), other-click (OC), machine-simulated-click (MSC), and machine-random-click (MRC) conditions

Table 1. Study 1 error scores for data collapsed across Human (SC & OC) and Machine (MSC & MRC) conditions

Collapsed Conditions	Mean	Std. Dev.	N
Human	.2557	3.20063	35
Machine	1.5657	2.43723	35

One possible explanation for this reversal of the agency effect draws on the use of a virtual target (a red flash) that is *associated* with a mouse click rather than the direct observation of a physical movement. Perhaps a time interval as short as 8 ms is enough for participants to perceive the red flash as an *effect* of the mouse click rather than a simultaneous associated event. This is a perception that may be reinforced in everyday computer use, where mouse clicks are understood to cause computer events. In order to test this possibility, a paired t-test compared the timing of physical mouse clicks in the training blocks that preceded experimental trials with the self-click experimental trials in which participants were prompted to use the on-screen flash as the target in the timing task. Results of this analysis showed no significant difference ($t(34) = 1.052$, $p = .300$), suggesting that physical and virtual markers of action are functionally equivalent, at least with respect to the timing tasks used in study 1.

4 Study 2 Results and Discussion

Analysis in study 2 also began with a repeated-measure GLM ANOVA. This analysis revealed an agency effect for the auditory tone following observed actions, with significant differences in both the general test (Hotellings Trace $F(3,33) = 3.117$, $p = .039$, $\eta^2 = .221$) and in the within-subjects test corrected for possible violation of the sphericity assumption (Huyhn-Feldt $F(2.304,80.653) = 4.461$, $p = .011$, $\eta^2 = .113$) with observed power of .673 and .792, respectively. Mean error scores across the four conditions are illustrated in Figure 4.

Follow-up analysis showed that the observed agency effect could be attributed to differences between the human and machine conditions. A paired t-test examined the data with conditions collapsed into human and machine conditions as in study 1. Results of a paired t-test revealed a significant difference, t(35) = 2.681, p = .011. Further paired t-tests showed no significant differences within the two human and machine conditions. Results of study 2 corroborate the findings of study 1, including the anomalous reversal of agency effects in a virtual environment. As before, there was a statistically significant agency effect and this effect could be attributed to differences between the human and machine conditions.

Results show that participants' perceptions of the timing of audible tones differed depending on whether those tones followed a human- or machine-initiated action, although the sequence and relative timing of the target stimuli were similar in all conditions. These results confirm the findings of study 1 indicating that intentional binding is observed in virtual environments typical of HCI. As before, however, results indicate that intentional binding operates differently in a virtual environment, where the observed effects are reversed compared to human-machine differences noted in physical environments

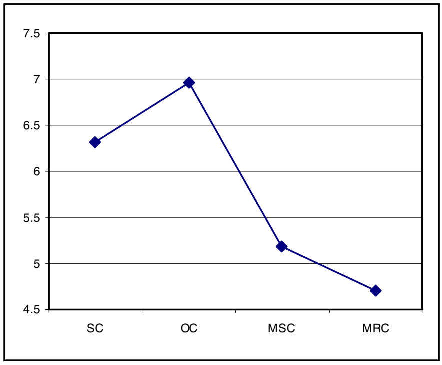

Fig 4. Mean error scores from Study 2 for self-click (SC), other-click (OC), machine-simulated-click (MSC), and machine-random-click (MRC) conditions

Table 2. Study 1 error scores for data collapsed across Human (SC & OC) and Machine (MSC & MRC) conditions

Collapsed Conditions	Mean	Std. Dev.	N
Human	6.6442	3.20063	36
Machine	4.9450	2.43723	36

5 General Discussion

Results of studies 1 and 2 indicate that agency effects influence users' subjective experiences of action and response in even very simple HCI tasks involving mouse manipulations with auditory and visual feedback. The results that have been observed, however, differ from results of prior work. Macro-interactive studies of people interacting with computers on social time scales (on the order of minutes) show that people tend to attribute agency to computers with which they interact. In the present studies, however, subjects attributed agency only when the source of an action or effect was human. Subjects did *not* attribute agency in machine-action conditions suggesting that there may be an interaction threshold for social agency effects that was not attained by either machine condition in the present studies but was attained in the macro-interactive tasks used in prior work.

There are also differences in the way agency effects are expressed in virtual environments. Micro-interactive tasks in previous work document an attraction between action and effect when agency is attributed to the source of action. The present studies, however, show greater attraction between action and effect under non-agency *machine* conditions. Data across all four conditions and both studies are depicted in Figure 5. In the two human agency conditions (SC & OC) the subject or an experimenter manipulated the mouse. In the two machine conditions (MSC & MRC) the computer simulated agency or triggered events randomly without any semblance of agency. The relative attraction between click and tone is apparent: click and tone are more attracted in the machine conditions than in human conditions.

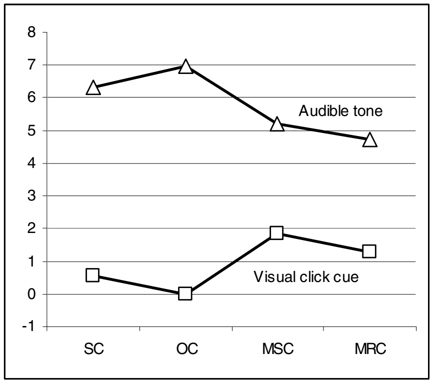

Fig. 5. Mean perceptual errors scores from study 1 (squares) and study 2 (triangles)

6 Conclusions and Implications

Results of these studies support the claim that intentional binding operates in virtual environments typical of HCI. Human users perceive even simple actions and effects differently depending on whether agency is attributed to the source of the action. Given the significance and ubiquity of simple actions and effects in HCI, it may be important to examine the influence of intentional binding and other agency effects on software use and the user experience. Furthermore, the fact that unconscious agency effects also influence macro-interactive social behaviors suggests these effects operate across a broad spectrum of contexts. It cannot, however, be assumed that the agency effects expressed in micro-interactive contexts like those in the two studies described here are related to macro-interactive counterparts in simple ways. As demonstrated by the discrepancies between physical and virtual effects, even within micro-interactive contexts there are differences in the way agency effects are expressed.

The significance of agency effects however is broader than suggested by the technical focus adopted in these studies. Personal agency is one of the most fundamental experiences we have in our interactions with others and the physical world. It is, however, becoming increasingly clear that the folk psychology of agency and the conscious experience on which it is based are not consistent with neurobiology and cognition. Our experience of intention follows rather than precedes the brain activity that drives it and there are circumstances where we may be mistaken about even the most basic aspects of personal agency. As we learn more about the neurocognition of agency it may be possible for complex interactive systems to unconsciously influence our experience of agency [17], particularly since there are well documented examples of similar types of unconscious influence [18]. Semantic priming studies have consistently shown that subliminally displayed words influence the perception of subsequent semantically associated words [19]. Studies of blindsight [20] have shown that human behavior may rely on visual information even when there is no conscious experience of vision. Finally, a series of related studies that involve HCI tasks similar to those in the present work [21] have shown that people can be induced to *incorrectly* attribute actions to themselves when those actions are actually a result of someone else's action. Finally, numerous recent studies suggest that our experience of personal agency relies on many of the same neural circuits that support the attribution of agency to others [22]. Learning about how we attribute agency to computers, therefore, may also help us better understand ourselves and our interactions with other people, as well as contribute to the neurocognitive foundations of HCI.

References

1. Reeves, B., Nass, C.: The media equation: How people treat computers, television, and new media like real people and places. Cambridge University Press, New York (1996)
2. Nass, C., Moon, Y., Green, M.: Are computers gender-neutral? Gender stereotypic responses to computers. Journal of Applied Social Psychology 27(10), 864–876 (1997)
3. Nass, C., Moon, Y.: Machines and mindlessness: Social responses to computers. Journal of Social Issues 56(1), 81–103 (2000)
4. Moon, Y., Nass, C.: Are computers scapegoats? Attributions of responsibility in human-computer interaction. International Journal of Human-Computer Interaction 49(1), 79–94 (1998)
5. Sundar, S.S., Nass, C.: Source orientation in human-computer interaction. Communication Research 27(6), 683–703 (2000)
6. Sengers, P., Liesendahi, R., Magar, W., Seibert, C., Muller, B., Joachims, T., Geng, W., Martensson, P., Hook, K.: The Enigmatics of Affect. Anonymous. In: Proceedings of the Conference on Designing Interactive Systems: Processes, practices, methods, and techniques, London, pp. 87–98 (2002)
7. 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 106, 623–642 (1983)
8. Coles, M.G.H., Rugg, M.D.: Event-related brain potentials. In: Rugg, M.D., Coles, M.G.H. (eds.) Electrophysiology of Mind. Oxford University Press, Oxford (1996)
9. Society for Neuroscience: Mirror, mirror in the brain: Mirror neurons, self-understanding, and autism research. Science daily (2007)
10. Gallese, V., Fadiga, L., Fogassi, L., Rizolatti, G.: Action recognition in the premotor cortex. Brain 119, 593–609 (1996)
11. Rizolatti, G., Fogassi, L., Gallese, V.: Neurophysiological mechanisms underlying the understanding and imitation of action. Nat. Rev. Neurosci. 2, 661–670 (2001)
12. Gallese, V., Keysers, C., Rizzolatti, G.: A unifying view of the basis of social cognition. Trends in Cognitive Sciences 8, 396–403 (2004)
13. Kilner, J.M., Paulignan, Y., Blakemore, S.J.: An interference effect of observed biological movement on action. Current Biology 13(6), 522–525 (2003)
14. Haggard, P., Clark, S., Kalogeras, J.: Voluntary action and conscious awareness. Nat. Neurosci. 5(4), 382–385 (2002)
15. Pockett, S., Miller, A.: The rotating spot method of timing participantive events. Consciousness and Cognition 16(2), 241–254 (2006)
16. Peirce, J.W.: Psychopy - Psychophysics software in Python. Journal of Neuroscientific Methods 162(1-2), 8–13 (2006)

17. Stetson, C., Cui, X., Montague, P.R., Eagleman, D.M.: Motor-sensory recalibration leads to an illusory reversal of action and sensation. Neuron 51, 651–659 (2006)
18. Dixon, N.F.: Subliminal perception: The nature of a controversy. McGraw Hill, New York (1971)
19. Marcel, A.J.: Conscious and unconscious perception: Experiments on visual masking and word recognition. Cognit. Psychol. 15, 197–237 (1983)
20. Weiskrantz, L.: Prime-sight and blindsight. Conscious. Cogn. 11, 568–581 (2002)
21. Wegner, D.M.: The illusion of conscious will. MIT Press, Cambridge (2002)
22. Jackson, P.L., Decety, J.: Motor cognition: A new paradigm to study self-other interactions. Current Opinion in Neurobiology 14, 259–263 (2004)