Awareness of action: inference and prediction

Authors: James Moore & Patrick Haggard

Correspondence address:

James Moore

Institute of Cognitive Neuroscience

17 Queen Square

London WC1N 3AR

Tel: 020 7679 1128

Email: j.w.moore@ucl.ac.uk

Running head: Action awareness

Acknowledgements: JM was supported by an ESRC/MRC studentship grant. This work was further supported by ESRC grant RES00231571 to PH.

Submitted to Consciousness and Cognition (accepted, in press article)

Abstract

This study investigates whether the conscious awareness of action is based on predictive motor control processes, or on inferential "sense-making" process that occur after the action itself. We investigated whether the temporal binding between perceptual estimates of operant actions and their effects depends on the occurrence of the effect (inferential processes) or on the prediction that the effect will occur (predictive processes). By varying the probability with which a simple manual action produced an auditory effect, we showed that both the actual and the predicted occurrence of the effect played a role. When predictability of the effect of action was low, temporal binding was found only on those trials where the auditory effect occurred. In contrast, when predictability of the effect of action was high, temporal binding occurred even on trials where the action produced no effect. Further analysis showed that the predictive process is modulated by recent experience of the action-effect relation. We conclude that the experience of action depends on a dynamic combination of predictive and inferential processes.

Keywords: Action, Intention, Action effects, Time judgment, Awareness, Agency

Introduction

Operant or instrumental actions are those that are causally related to certain outcomes (Thorndike, 1911). The ability to make operant actions which manipulate and control the external environment is of great significance in the evolution of the mind: operance has made us masters of the environment, not enslaved by it.

Operant actions are frequently explained in terms of an agent's intentions, carried out in order to satisfy particular desires or needs (Dickinson, 1994). In humans, they often involve a conscious representation of both the desired or goal state, and a conscious intention to achieve it. The explanatory framework for operant action implies that such actions are internally generated, rather than responses to external stimuli. They thus have intimate relationship with consciousness and the Self.

Recent research has focussed on the epistemic content and neural origins of the experience of action. One recent account proposes that the awareness of action arises from predictive motor control processes (Blakemore, Wolpert, & Frith, 2002). On this position, we may be aware of motor predictions made by the 'forward model' component of the motor control system (for review see Wolpert & Ghahramani, 2000).

Support for this position comes from a number of sources. For example, Fried et al (1991) stimulated the supplementary motor area in humans as part of an exploratory procedure prior to neurosurgery. At low levels of electrical stimulation, patients reported having an urge to move. When stimulation was increased, they actually performed the urged movement. More direct evidence comes from amputee patients experiencing phantom limb phenomena. Such patients will feel that they can voluntarily move and control the amputated limb (Ramachandran & Hirstein, 1998). Blakemore, Wolpert and

Frith (2002) suggest that this phenomenon arises because conscious awareness of action is based not on actual limb position, but on the *predicted* limb position, derived from efferent commands. On this view, we have internal motoric information about our actions, which is, at least in part, available to conscious awareness.

An opposing view holds that our conscious awareness of action is subserved by an inferential process. On this view, we use sensory evidence to ‘make sense’ of our actions and their antecedent/ subsequent events (e.g. Wegner & Wheatley, 1999; Wegner, 2002).

Wegner and Wheatley (1999) fooled participants into perceiving movements performed by a confederate as self-generated. They achieved this by priming participants with thoughts relevant to a subsequent movement made by the confederate. This finding suggests that we are not intrinsically informed of our actions, and instead rely on an inferential process to make sense of them. The inferential process would generate the experience of action by accumulating sensory evidence about actions in the same way that other perceptual/inferential processes rely on sensory evidence about external events. Crucially, the inferential view does not give any special role to efferent information internal to the motor system in generating the experience of action. Wegner and Wheatley identify three principles of *priority*, *consistency*, and *exclusivity* that underlie the inferential processes of action experience. They argue that when a thought occurs prior to an action, is consistent with the action, and the action has no plausible alternative cause, then we experience the feeling of consciously willing the action.

Further support for the inferential account of action awareness comes from a study by Johansson, Hall, Sikstrom, and Olsson (2005). They presented participants with photographs of two faces, and the participants had to choose which of the two was more attractive. The ‘chosen’ photograph was then re-presented to the participant, and they were asked to justify their decision. Johansson et al covertly manipulated the relationship between the choice and the photograph subsequently presented by swapping the chosen and non-chosen cards. They found that when there was a mismatch between initial choice and actual outcome (i.e., when the cards were swapped) participants would nevertheless offer a justification for a choice they had not in fact made. The participants appeared to infer that they had chosen a card from the fact that it was subsequently re-presented, and proceeded to retrospectively invent the reasons for their choice. Again, this suggests that we are not intrinsically informed of our actions and their consequences, and that instead we build up an account of our actions based on current evidence.

There are two principle ways to distinguish between the predictive and inferential accounts of action awareness. First, the two accounts differ *epistemologically*, because predictive efferent information is necessarily intrinsic to the agent. The inferential account suggests that “we are not intrinsically informed of our own authorship” (Wegner, 2002, p. 218), and instead, we represent our actions using general cognitive mechanisms for inferring information about events. Second, the two accounts differ in their *chronometric* implications. Predictive processes contribute to action awareness only before action. Inferential processes will typically contribute after action, when external information about the action becomes available. However, Wegner has suggested that the inferential system may also contribute prior to movement onset (see Wegner, 2002, footnote 3, p. 68).

We have used the intentional binding effect (Haggard, Aschersleben, Gehrke, & Prinz, 2002; Haggard, Clark, & Kalogeras, 2002) to measure the extent to which experiences of action are influenced by subsequent effects. In the intentional binding paradigm, participants perform a keypress at a time of their choosing, which in turn causes a tone to sound 250msec later. This creates an operant context because the participant causes the tone. Participants judge either the time of their keypress or the time at which they heard the tone, in separate blocks of trials. In previous studies, participants judged the time of actions followed by tones to be later than the time of actions in a control block which were not followed by tones (Haggard, Clark, & Kalogeras 2002; Haggard & Cole, in press). They also judged the time of tones produced by their own actions as earlier than the time of tones in a control block which occur at similar random latencies in the absence of any voluntary action. In this way, there is a perceived temporal attraction between operant actions and their effects. Since this attraction was found for voluntary but not involuntary actions, it was called ‘intentional binding’ (Haggard, Clark, & Kalogeras, 2002). The intentional binding paradigm is thus an implicit measure of the sense of agency. In particular, the intentional binding effect suggests that the experience of a voluntary action reflects wider, contextual factors, viz the effects that the action has in the external environment. Awareness of an individual action depends on a more general sense of agency.

Here we have investigated the link between action awareness and agency in more detail, by manipulating the relation between actions and their external effects, as part of the experimental design. This manipulation serves to clarify the contributions of predictive and inferential processes to action awareness and to the more general sense of agency. Specifically, we asked whether the shift of the action towards the tone, as seen in the intentional binding effect, occurred only on those trials where the effect actually occurred, or whether it also occurred on trials where a tone was merely predicted to occur but did not. The former result would imply that the actual occurrence of the tone lead to retrospective postdiction of action awareness. The latter result would imply that the prediction of the tone was sufficient to influence conscious awareness of operant action. We also investigated the origin of the putative predictive processes by investigating whether the perceived time of actions depended on recent experience of the operant relation on the immediately preceding trials.

Method

Participants

10 right-handed individuals (9 female, 1 male; aged 18-31 years; mean 23 years) took part in the study. They were tested individually in a designated testing room over two one-and-a-half hour sessions, each on separate days.

Procedure

The experiment was based on that of Haggard, Clark, and Kalogeras (2002), and is shown in figure 1. On each trial participants were instructed to press a key with their right index finger at a time of their choosing which caused an effect (a tone) 250ms later. At the same time, a clock hand was rotating about a clock face at a rate of one revolution every 2560ms. After the tone's offset the clock hand continued rotating for a random time, and then stopped. The participant was then prompted to judge the position of the clock hand at which they pressed the key, or the time at which they heard the tone. The participant entered the judgements using a keypad with their left hand. The judgements were blocked, so that in each condition the participant always judged just one of the two events.

In the current study we focus on awareness of action, rather than awareness of the subsequent tone. However, we reasoned that obtaining tone timing judgements on separate blocks was necessary to prevent subjects from ignoring the tones completely, and to encourage them to represent the relation between action and effect.

[Figure 1 about here]

In the current study, the probability of the action causing the tone was manipulated. In some conditions, 75% of actions were followed by tones; in other conditions, 50% of actions were followed by tones. For those blocks of trials in which subjects were asked to judge the time of tones, they were instructed to report a dummy value on trials where no tone in fact occurred. For blocks of trials in which subjects were asked to judge the time of actions, they did so on every trial irrespective of whether their action produced a tone or not.

The experimental procedure consisted of two sessions (lasting 1.5 hours each), completed on separate days. The blocks of trials each participant completed were spread equally over the two sessions. For each effect probability (50% and 75%), participants completed six blocks of action timing judgements, and six blocks of tone timing judgements in total. Further to this there were two baseline conditions. In the *baseline* *action condition* participants made a key press at a time of their choosing. This key press did not cause a tone, and participants simply had to judge the time of their key press in the absence of the effect. In the *baseline tone condition*, participants were instructed not to press the key and instead wait for a tone to sound at a random latency generated by the computer. Participants had to then judge the time of this tone in the absence of any action. Participants completed two blocks of baseline action trials, and two blocks of baseline effect trials in total.

Each block consisted of 32 trials. The effect probabilities to which participants were exposed were counterbalanced. Within this, the order of the action/ tone judgement blocks was randomised anew for each subject. Finally, the baseline conditions were randomly inserted into the running order between effect-probability conditions.

Data analysis

Our analysis focussed on the judgements of action, and estimates of the tone onset time are not considered here. Mean baseline-corrected estimates were obtained for action judgements for each participant. These were generated by subtracting judgements of actions in baseline conditions from the mean judgments of the same event in operant conditions. This procedure gives a measure of binding for actions. By convention a positive shift indicates a delay in the awareness of action in the direction of the effect. This shift is taken as an implicit measure of the extent to which the subject had the experience of agency.

Results

The mean time estimates in each condition are shown in table 1.

Binding for actions

Figure 2 shows the mean baseline-corrected estimates of the time of action. We first investigated whether the intentional binding effect was present for each trial type at each level of effect probability. Because the direction of the binding effect was predicted, one-tailed tests were used. We performed 4 separate tests and therefore the one-tailed significance level was Bonferroni-corrected to 0.0125. In the 75% probability condition, binding was found for 'action only' estimates ($t(9) = 4.81$, $p=.0005$, one-tailed), and 'action + tone' estimates ($t(9) = 2.99$, $p=.008$). In the 50% probability condition, there was a trend towards binding for 'action + tone' trials in the 50% probability condition which came close to significance ($t(9)$ 2.42, $p=0.019$). However, no binding was found for 'action only' trials in the 50% probability condition ($t(9) = .861$, $p = 0.206$). These preliminary analyses partly replicate previous reports of binding of actions towards subsequent effects, while suggesting possible differences between the conditions of our experimental design.

[Figure 2 about here]

Predictive and inferential influences on the conscious awareness of operant actions

We next compared the binding effects between conditions, to investigate the contribution of predictive and inferential processes to action awareness. A 2x2 ANOVA was conducted with factors of effect-probability (50%, 75%) and trial type (‘action only’, ‘action + tone’). There was no main effect of effect-probability ($F(1,9) = 1.84$, $MSE = 170.17$, $p = .208$), or of trial type ($F(1,9) = 1.38$, $MSE = 53.51$, $p = .271$). However, the analysis revealed a significant interaction ($F(1,9) = 6.06$, $MSE = 47.60$, $p = .036$).

Simple effects t tests were used to investigate this interaction further. First, we found a significant difference between ‘action only’ estimates in the 50% and 75% conditions ($t(9) = -2.33$, $SD = 14.89$, $p = .045$). That is, actions which did not cause a tone in conditions where the tone was of high predictability (75% probability) were shifted towards the expected time of the tone’s occurrence, whereas actions not causing a tone where the tone was of low predictability (50% probability) showed only a minimal shift (see Figure 2). This confirms the predictive account.

Second, in the 50% effect-probability blocks where predictability of tones was low, there was little shift towards the tone for ‘action only’ trials, yet there was a greater shift for ‘action + tone’ trials (see Figure 2). A planned paired sample t-test showed this difference to be significant ($t(9) = 2.70$, $SD = 9.47$, $p = .024$). This result is consistent with the inferential account.

No other pairwise differences between conditions were significant (75% probability ‘action + tone’ vs ‘action only’: $t(9)$ -.792, $SD = 10.24$, $p = .449$, & ‘action + tone’ 50% vs. 75% probability: $t(9) = -.049$, $SD = 14.62$, $p = .962$).

The influence of recent experience on conscious awareness of operant action

We performed an additional analysis to investigate the role of recent experience in binding of actions towards tones. The method for this analysis was taken from Whitney, Vetter and Wolpert (2001). First, we classified each trial according to how recently an ‘action + tone’ trial had previously been experienced. We restricted our analysis to trials for which just one of the previous three trials had involved ‘action + tone’. Thus, the amount of recent information about action-effect relations was held constant, but the recency of this information varied. Accordingly, each ‘action only’ trial was classified according to the nature of the 3 trials that preceded it (1 indicates an ‘action + tone’ trial, and 0 indicates an ‘action only’ trial):

1 0 0 – An ‘action + tone’ trial occurred three trials previous to the current trial.

0 1 0 – The ‘action + tone’ trial occurs in the penultimate trial.

0 0 1 – The ‘action + tone’ trial occurs in the previous trial.

The data from both probability blocks was combined to increase the number of trials in each of these classes for this analysis. Our focus was thus on the influence of *when* a single *recent* operant event occurred, rather than on the frequency of these operant events over the *long run*, since the latter factor was already addressed by comparing 50% and 75% probability blocks.

Figure 3 shows action shifts for ‘action only’ trials and ‘action + tone’ trials, as a function of the previous experience triplets. A 2 x 3 ANOVA with factors of trial type (‘action only’, ‘action + tone’) and recent experience triplet (‘1 0 0’, ‘0 1 0’, ‘0 0 1’) was performed on action shift data. This analysis demonstrated no main effect of trial type ($F(1, 9) = 3.72$, MSE = 339.39, $p = .086$), and no main effect of recent experience triplet ($F(2, 18) = .85$, MSE = 193.68 , $p = .443$), but a significant interaction between recent experience triplet and action trial type (F(2, 19) = 3.56, MSE = 492.83, p = .05). This interaction was investigated further by simple effects t-tests. These showed that the interaction arose because 'action only' trials which involved more recent 'action + tone' experience ('0 0 1') showed stronger binding effects than trials with penultimate experience triplet ('0 1 0'; t(9) = - 2.64, SD = 24.95, p = .027). In contrast, the corresponding comparison for 'action + tone' trials was not significant (t(9) = 2.01, SD = 22.32, p= .075). In addition, the '1 0 0' vs. '0 0 1' comparison was not significant for both 'action only' trials (t(9) = - 1.81, SD = 35.25 , p = .104) and 'action + tone' trials (t(9) = 1.28, SD = 21.50, p = .232). To summarise, recent experience generates a predictive shift in the awareness of action on 'action only' trials, but not a significant change in inferential action awareness on 'action + tone' trials.

[Figure 3 about here]

Discussion

We have investigated shifts in the perceived time of voluntary actions towards a subsequent effect, under a number of different action-effect relations. First, we replicated the shift reported previously in studies of intentional binding (Haggard, Clark, & Kalogeras, 2002). Second, we showed that the magnitude of this shift depends both on the long-run probability of actions being followed by the effect, and also on the actual occurrence of the effect on the current trial. That is, actions which were not followed by effects nevertheless showed binding in blocks where an effect was highly probable, but not in blocks where an effect was less probable. Third, in addition to these relatively long-term effects operating over an entire block, we showed an effect of recent operant experience. Very recent reinforcement of the action-effect relation increased binding on action only trials, compared to identical experience in the more distant past.

To investigate influences on action awareness, we used an intentional binding paradigm developed in our previous work (Haggard, Aschersleben, Gehrke, & Prinz, 2002; Haggard, Clark, & Kalogeras, 2002), and based on Libet et al's event timing method. This method has proved to be one of the few viable experimental approaches to conscious experience of voluntary action. In spite of this, the timing method has attracted considerable criticism (see Banks & Pockett, 2006, for a review). Factors such as attention and 'prior entry' (Titchener, 1908) mean that any single time estimate needs to be treated with caution. However, differences between estimates of the same physical event in different contexts can avoid some of these criticisms. In the present study, all inferences are based on shifts in the perceived time of action relative to a baseline condition. We thus focus on differences in action awareness between conditions, rather than the time of action awareness on individual trials.

Current theories suggest two possible sources of the conscious awareness of operant action. Action awareness might depend on predictive motor control processes, such as efferent motor command signals (e.g. Blakemore, Wolpert, & Frith, 2002). Alternatively, awareness of action may arise from inferential 'sense-making' processes, which use sensory evidence about physical events (e.g. Wegner & Wheatley, 1999).

Our results suggest that *both* predictive and inferential processes contribute to the conscious awareness of operant action. The relative contribution of each of these processes seems to be context dependent. When we can predict the consequences of our actions, as in a high action-effect contingency block, the awareness of action reflects these predictions. This would provide us with a *predictive* sense of our own agency. In addition, our results show clear evidence that inferential processes also influence the conscious awareness of operant action; the occurrence of a tone *after* the action influenced the perceived time of the action. These processes could therefore be called reconstructive or postdictive (Haggard, 2005). Our experiment used an indirect measure of sense of agency, based on the perceived time of individual actions. We show that the experience of agency resulting from prediction, and that arising from inference, have similar effects on time perception. However, we cannot rule out the possibility that the two processes differ in other phenomenal aspects. This would be a fruitful point for further study using phenomenologically richer measures.

The interaction between predictive and inferential processes is of particular interest. Figure 2 suggest that this interaction is less than additive. An estimate of the predictive process comes from the binding of ‘action only’ trials in the 75% block, while an estimate of the inferential process comes from the binding of ‘action + tone’ in the 50% block. Interestingly, the amount of binding in the ‘action + tone’ trials of the 75% probability blocks was not a simple linear sum of predictive shifts and inferential shifts, but was similar in size to each shift individually. Estimates of the time of action seem to involve an underadditive combination of two quite different kinds of information. Moreover, these two kinds of information become available at different times. We now speculate on the principles governing this cognitive combination of information, resulting in a sense of agency.

The best developed framework for explaining the combination of different sources of information in such settings is the Bayesian approach. Specifically, this describes a method for the combining prior knowledge or beliefs, and new sensory data (Chater, Tenenbaum, & Yuille, 2006). A Bayesian framework weights each of these sources according to its reliability. For example, given a strong prior and noisy sensory feedback, a good estimate should rely heavily on prior information. Given a weak prior and reliable sensory feedback, a good estimate should rely heavily on sensory feedback. We propose that a form of Bayesian integration underpins 'sense of agency' as measured by intentional binding. In this context, we consider the prior to be the knowledge about the action-effect relation built up over the course of the block. The various sources of information about the time of action on the current trial are regarded as sensory data, including efferent and somatosensory information, and the auditory information about the beep following the action.

Where effect probability is high (as in the 75% probability condition) there should be a strong prior representation of the action-effect relation. This strong prior representation should make clear predictions about the occurrence of an effect. Consequently, a predictive shift in the awareness of action towards an expected tone should occur, even in the absence of the tone.

Alternatively, where effect probability is low (as in the 50% probability condition) there should be a weak representation of the action-effect relation. As a result, clear predictions about the occurrence of the effect after any particular action are not possible. In this case, the only evidence about the action comes from current sensory data. Since there is not enough evidence to predict the occurrence of the tone, the modulation of action awareness due to *predicting* the tone is absent. If, however, the tone is actually experienced, it may still modulate the experience of action, but will now do so postdictively rather than predictively.

Thus, two dissociable representations of the tone contribute to action awareness, a predictive one and a sensory one. Lau, Rogers, and Passingham (in press) used a Bayesian framework to model the perceived onset of actions and intentions as a weighted mixture of information about preparation and about actual motor activity. This framework fits well with our interpretation, which extends the application of the Bayesian approach to agency and the effects of action. Furthermore, we suggest that the awareness of action involves integration not only of different signals from a single trial, but also the integration of predictions based on previous trials with information from sensory events on the current trial.

The time course over which information about action is built up may be an important clue to this interaction. Specifically, predictive information about actions and effects may operate only over the brief timescale of motor preparation, and may be discarded when sensory evidence makes inference possible. Predictions may be based on an internal forward model (Blakemore, Wolpert and Frith, 2002). This representation is available for the control of action and on-line guidance of behaviour, but does not outlast the current action. Sensory feedback provides more precise evidence about actions and their effects. This evidence becomes available only after a short sensory delay, but can then be transferred to memory. Thus, reliable and enduring sensory evidence replaces short-lived predictive estimates. We suggest that awareness of action therefore switches from a predictive to an inferential source as the action itself occurs, and as sensory information becomes available. This time-varying mixture of predictive and inferential information may ensure that our experience of our own action is an optimal reflection of the actual relation between our voluntary motor system and the outside world.

The present study also investigated the modulation of predictive and inferential influences by recent prior experience. We believe this is the first report of experience-dependent effects on action awareness. We show that very recent prior experience influences the predictive contribution to action awareness, but not the inferential contribution. In particular, ‘action only’ trials showed significant binding only when they immediately followed an ‘action + tone’ trial. Even a single intervening trial which did not involve any experience of agency (e.g., ‘0 1 0’ trials) was sufficient to abolish the predictive binding effect. A similar effect of recent prior experience on predictive motor control of grasping was reported by Whitney, Vetter and Wolpert (2001). In that study, participants made upward movements of a lever with their left hand which in turn either did or did not cause a lever held in their right hand to move. Participants were instructed to keep the right hand as still as possible. This design enabled the experimenters to obtain a measure of anticipatory grip force modulation in the right hand on trials where the left hand movement did not cause a movement of the right hand. Interestingly, the more recently a perturbation of the right hand had occurred within the previous three trials, the greater the anticipatory grip force modulation. This shows that the motor system uses short term predictions based on recent sensorimotor experience. Our result is based on a similar sequential analysis of trials. We show that action awareness reflects short-term learning in the same way as Whitney et al.’s measure of motor performance. This supports the account that conscious awareness of action is, in part, subserved by predictions made by the motor control system.

To conclude, this study distinguishes between two processes that contribute to the conscious awareness of intentional actions. Inferential processes enable an evidence-based experience of action, akin to general ‘sense-making’ mechanisms of cognition. In addition, predictions, perhaps based on copies of efferent commands within the motor system, and built up over the course of previous actions, also influence the conscious awareness of intentional action, and contribute to the sense of agency. The consistency of the action-effect relation influences action awareness over long time scales, reflecting the overall probability of an action causing an effect, *and* over the shorter timescale of recent experience of actions causing effects. Finally, the predictive component operates in advance of the action. After the action itself, predictive estimates appear to be replaced by an inference-based representation. Thus, action awareness represents a dynamic fusion of predictive and inferential processes in the brain.

References

Banks, W. P., & Pockett, S. (2006). Libet's work on the neuroscience of free will. In M. Velmans (Ed.) *The Blackwell Companion to Consciousness*. Cambridge, England: Blackwell.

Blakemore, S-J, Wolpert, D. M., & Frith, C. D. (2002). Abnormalities in the awareness of action. *Trends in Cognitive Sciences, 6(6)*, 237-242.

Chater, N., Tenenbaum, J. B., & Yuille, A. (2006). Probabilistic models of cognition: Conceptual foundations. *Trends in Cognitive Sciences, 10(7)*, 287-291.

Churchland, P. (1984). *Matter and Consciousness*. Cambridge: MIT Press.

Dickinson, A. (1994). Instrumental conditioning. In N. J. (Ed.) *Animal learning and conditioning* (pp. 45-79). New York: Academic Press.

Fried, I. Katz, A., McCarthy, G., Sass, K. J., Williamson, P., Spencer, S. S., & Spencer, D. D. (1991). Functional organisation of human supplementary motor cortex studied by electrical stimulation. *Journal of Neuroscience, 11,* 3656-3666.

Haggard, P. (2005). Conscious intention and motor cognition. *Trends in Cognitive Sciences, 9,* 290-295.

Haggard, P., Aschersleben, G., Gehrke, J., & Prinz, W. (2002). Action, binding and awareness. In W. Prinz & B Hommel (Eds.) *Common Mechanisms in Perception and Action* (Attention and Performance, Vol. XIX, pp266-285). Oxford, UK: Oxford University Press.

Haggard, P., Clark, S., & Kalogeras, J. (2002). Voluntary action and conscious awareness. *Nature Neuroscience, 5(4)*, 382-385.

Haggard, P. & Cole, J. (in press). Intention, attention, and the temporal experience of action. *Consciousness and Cognition*.

Johansson, P., Hall, L., Sikstrom, S., & Olsson, A. (2005). Failure to detect mismatches between intention and outcome in a simple decision task. *Science, 310*, 116-119.

Kording, K. P. & Wolpert, D. M. (2004). Bayesian integration in sensorimotor learning. *Nature, 427*, 244-247.

Libet, B., Gleason, C. A., Wright, E. W. & Pearl, D. K. (1983). 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.

Lau, H. C., Rogers, R. D., & Passingham, R. E. (in press). Manipulating the experienced onset of intention after action execution. *Journal of Cognitive Neuroscience.*

Ramachandran, V. S. & Hirstein, W. (1998). The perception of phantom limbs. *Brain, 121*, 1603-1630.

Thorndike, E. L. (1911). *Animal Intelligence*. New York: Macmillan.

Titchener, E. B. (1908). *Lectures on the elementary psychology of feeling and attention*. New York: MacMillan.

Wegner, D. M. (2002). *The Illusion of Conscious Will*. Cambridge, MA: MIT Press.

Wegner, D. M. & Wheatley, T. (1999). Apparent mental causation: sources of the experience of will. *American Psychologist, 54(7),* 480-492.

Whitney, A. G., Vetter, P., & Wolpert D. M. (2001). The influence of previous experience on predictive motor control. *NeuroReport, 12(4),* 649-653

Wolpert, D. M. & Ghahramani, Z. (2000). Computational principles of movement neuroscience. *Nature Neuroscience, 3,* 1212-1217.

Table 1. Mean judgment errors in ms (SD across subjects) and shifts relative to baseline conditions in ms.

	Judged Event	Mean (SD) judgement error (ms)	Mean shift from baseline (ms) (sd)
Baseline conditions			
	Action	-73 (95)	
Operant actions			
50 % probability	Action only	-68 (99)	5 (85)
	Action (followed by tone)	-60 (104)	13 (90)
75% probability	Action only	-57 (98)	16 (84)
	Action (followed by tone)	-60 (92)	13 (80)

Figure Caption

Figure 1. Schematic of the experimental set-up used in the study.

Figure 2. Mean baseline-corrected action estimates (ms) for each probability block, and for ‘action only’ and ‘action + tone’ trials. Bars show standard error across subjects.

Figure 3. Mean baseline-corrected action estimates (ms) as a function of experience of action-effect relation on the immediately preceding three trials, for ‘action only’ and ‘action + tone’ trials. See text for further explanation.

Figure 1

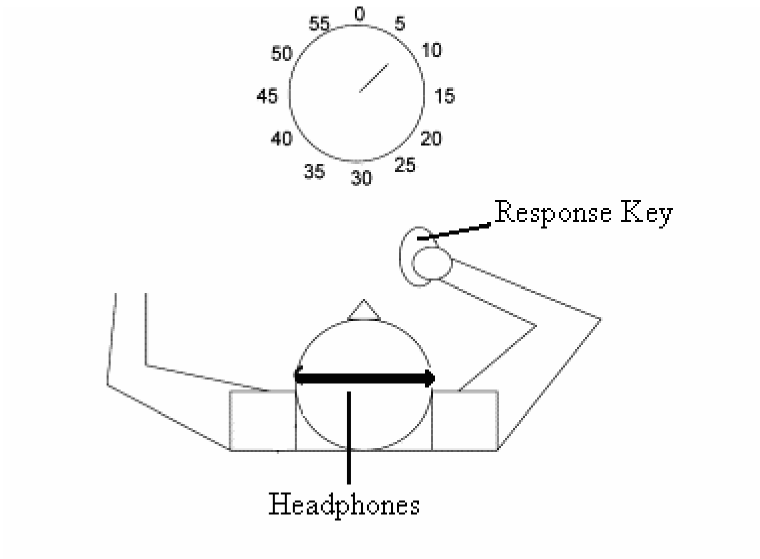

Figure 2

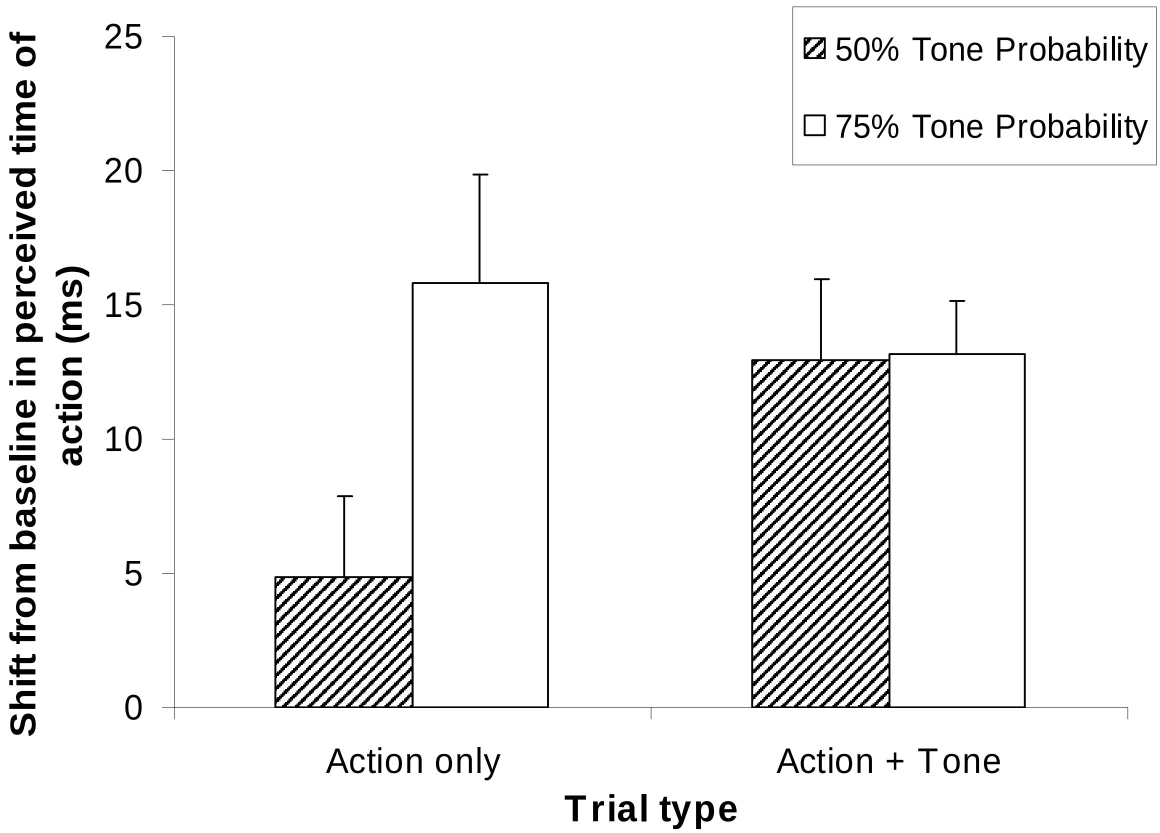

Shift from baseline in perceived time of action (ms)
25
20
15
10
5
0
50% Tone Probability
75% Tone Probability
Action only
Action + Tone
Trial type

Figure 3

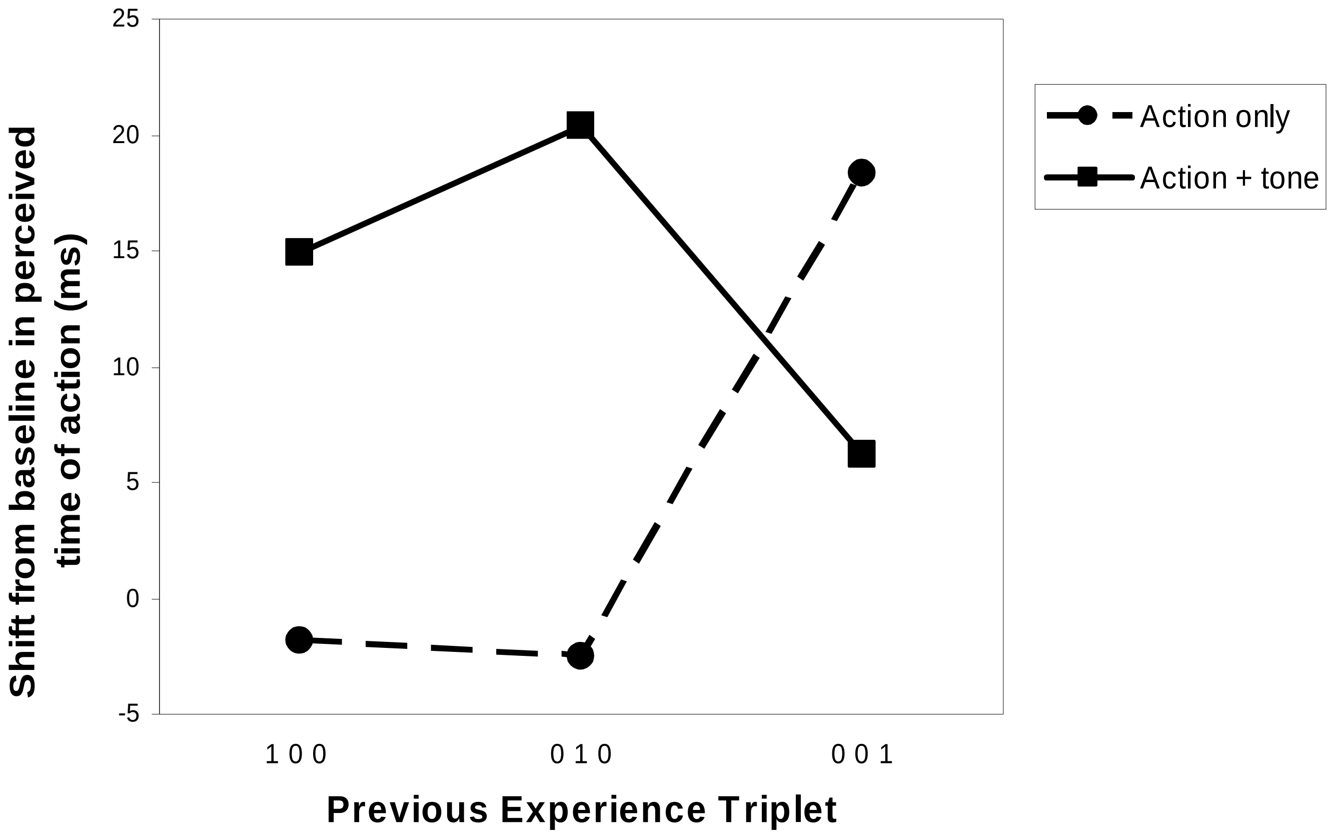
25
20
15
10
5
0
-5
Shift from baseline in perceived time of action (ms)
1 0 0
0 1 0
0 0 1
Previous Experience Triplet
Action only
Action + tone