Action enhances predicted touch

Emily R. Thomas[1*], Daniel Yon[1,2], Floris P. de Lange[3] & Clare Press[1]

1. Department of Psychological Sciences, Birkbeck, University of London, UK.

2. Department of Psychology, Goldsmiths, University of London, UK.

3. Donders Institute for Brain, Cognition and Behaviour, Radboud University, NL.

* Email: ethoma09@mail.bbk.ac.uk

https://orcid.org/0000-0001-8875-6678
https://orcid.org/0000-0002-5511-3884
https://orcid.org/0000-0002-6730-1452
https://orcid.org/0000-0002-4777-2249

Keywords

Prediction, Perception, Attenuation, Motor Processes

Author Contributions

This PDF file includes:

> Main Text
> Figures 1 to 3

Abstract

It is widely believed that predicted action outcomes are perceptually attenuated. Such mechanisms would prioritise processing of unexpected events, enabling us to perform corrective actions and update models of the world. However, evidence that action predictions attenuate perception is sparse. The present experiments determined whether predictive mechanisms in fact generate attenuation, or instead enhance perception – as assumed in sensory cognition domains outside of action. Under these accounts, it is thought that we combine prior expectations and sensory input to generate percepts, resulting in upweighted perception of the expected. Such mechanisms would enable rapid generation of veridical percepts in the face of sensory noise. We therefore manipulated probabilistic expectations in a paradigm typically used to demonstrate sensory attenuation during action. Participants produced actions and rated the intensity of concurrent passive tactile forces. Experiment 1 first confirmed previous findings that action outcomes are perceived less intensely than similar passive stimulation, but demonstrated *more* intense perception when reducing the contribution of non-predictive gating processes. Experiments 2 and 3 manipulated prediction explicitly and found that expected outcomes are perceived more, not less, intensely than unexpected outcomes. Computational modelling suggested that expectations alter the way sensory evidence is integrated – increasing the gain afforded to expected tactile signals. These findings challenge a central tenet of prominent motor control theories and demonstrate that sensorimotor predictions do not exhibit a qualitatively distinct influence on perception.

Significance Statement

Perception of expected action outcomes is thought to be attenuated. Such a mechanism is reasoned to be adaptive because surprising inputs are more useful, e.g., signalling the need to take new courses of action. This mechanism is thought to explain why we cannot tickle ourselves and explain unusual aspects of action and awareness in clinical populations. However, we do not know whether action predictions really attenuate perception. We demonstrate here that in fact action predictions enhance, rather than attenuate, touch. These findings are consistent with prediction theories outside of action, thought to explain how we generate accurate representations of our noisy sensory world. Our data importantly thereby demonstrate that action predictions do not have a qualitatively distinct influence on perception.

Main Text

Introduction

Prominent motor theories (1–4) propose that we attenuate – or downweight – perceptual processing of expected action outcomes. Such downweighting mechanisms are thought to finesse the limited capacity of our sensory systems, prioritising processing of more informative unexpected sensory events that signal the need to perform new actions or update of our models of the world (5, 6). For example, if we lift a cup of coffee that is lighter than expected, attenuated processing of expected signals (e.g., touch on our fingertips) will allow dedicated processing of unexpected events (e.g., accelerating motion of the cup) allowing swift updating of our beliefs about the environment (e.g., the weight of the cup) and support corrective action to avoid spillage. These downweighting mechanisms are invoked to explain that self-produced tactile sensations generate lower activity in bilateral secondary somatosensory cortex (1, 7–9), and are perceived to be less intense (3, 10–15), than externally-produced forces. This theory also provides an explanation for why it is difficult to tickle oneself (1).

However, while these accounts have been influential for the last two decades, evidence that attenuation results from predictive mechanisms is sparse – especially considering the number of non-predictive mechanisms known to influence perception during action (6, 16–18, see *SI*

Discussion). The defining feature of prediction mechanisms is that they operate according to stimulus probabilities (19). As such, effects of prediction mechanisms are typically measured by presenting events with high and low conditional probabilities, and comparing perception of the 'expected' with the 'unexpected'. In contrast, typical experiments demonstrating attenuation during action compare the perception of events in the presence or absence of action, or when events are coincident versus delayed with respect to action (1, 3, 7, 10, 11, 13–15). In these experiments it is assumed that the sensory events which coincide with action are the (more) predicted consequences, which explains why perception of them is attenuated. It is however unclear whether these effects are indeed reflective of predictive mechanisms operating according to conditional probabilities.

This question is especially crucial to answer because the neural mechanisms that have been characterised to underlie predictive influences on perception in domains outside of action are thought to generate a qualitatively opposite influence on perception. In these theories – typically couched in Bayesian frameworks – it is proposed that we combine our expectations (prior) with the input (likelihood) to determine what we perceive (posterior; 19, 20). In other words, we *upweight*, rather than downweight, perception of expected events, increasing the detectability and apparent intensity of events (21–23) and enabling rapid generation of largely veridical experiences in the face of sensory noise (19, 24, 25). For example, some such theories propose that when we predict a sensory event we increase the gain on sensory units tuned to that event and via competitive local interactions relatively inhibit sensory populations tuned to unpredicted events (6, 19, 26–30). It is perhaps unclear why the adaptive arguments presented for downweighting (informativeness) and upweighting (veridicality) predicted percepts should apply differentially in the domain of action, and some preliminary evidence from the visual domain suggests that they may not (6, 30–34).

A stark difference between studies purporting to demonstrate upweighting and downweighting is that the former study visual perception whereas the latter study tactile perception. It is widely believed that action predictions shape tactile perception in a qualitatively distinct way, including proposals that differences relate to tactile events being body-related (2) and tightly coupled with the motor system (35), in a way that many predicted visual or auditory events are not. Similarly, differences may also relate to assumptions that tactile attenuation during action is dependent upon somatosensory-cerebellar connectivity (1, 8), in contrast with hippocampal mediation of prediction in visual processing (36, 37). However, given that studies examining touch perception during action have not explicitly manipulated predictability, such conclusions would be premature. The present studies therefore examine whether action predictions attenuate touch perception, or in fact enhance perception similarly to other predictions. We adopted the force judgement paradigm used widely in the action domain to examine tactile attenuation, and employed predictive manipulations. Such force judgement paradigms (e.g. 10, 11, 15) require participants to move an active right index finger towards a passive left finger, and the left finger receives stimulation mechanically. In these paradigms, touch on the left finger is reported as less forceful when accompanied by action, relative to similar judgements when passive.

Experiment 1 was designed to replicate the typical set-up and contrast it with one that removes potentially confounding influences of non-predictive sensory gating (see 38, 39) – thought to occur at the earliest relay in the spinal cord (17) and therefore influence perception regardless of conditional probabilities. Upon finding that tactile events presented alongside action can in fact be perceived *more* forcefully when reducing the contribution of non-predictive gating processes, Experiments 2 and 3 explicitly examined the contribution of predictive mechanisms. Thus, participants always judged the force of stimulation while concurrently performing an action, but where the tactile events were either predicted or not according to learned action-outcome probabilities. We thereby tested whether tactile events predicted by action are indeed perceived less intensely than unexpected events – as downweighting accounts would propose – or instead

are perceived as more intense – consistent with upweighting theories of prediction outside of action domains.

Results

Experiment 1: Replication of typical attenuation finding and contribution of non-predictive sensory gating processes. Participants moved an active right index finger towards a passive left index finger. The 'Contact' condition closely resembled previous studies where participants pressed a device with the active finger to generate a mechanical force on the passive finger. We were also interested in whether similar effects would be found in a 'No Contact' condition, where a similar downward motion of the active finger triggered the same stimulation but without making contact with any device. This condition was included to examine whether attenuation previously thought to be determined by predictive mechanisms could in principle be generated by generalised sensory gating mechanisms (see *SI Discussion*). Generalised gating is thought to reduce perception of any events delivered to moving effectors (38, 39), regardless of whether they are predicted outcomes of action, which is why studies seeking to establish predictive sensory attenuation have examined perception on passive effectors. However, perception of concurrent gated stimulation on an active finger – as generated when pressing a button – could still bias responses about passive stimulation (40). We predicted that we would replicate the typical attenuation finding (e.g. 10, 11) in the Contact condition, such that tactile events would be perceived as less forceful when accompanied by action. If this attenuation reflects the operation of predictive downweighting mechanisms, the reduction in intensity ratings during action would likely be equivalent in Contact and No Contact conditions. However, if instead attenuation is modulated by generalised sensory gating, the pattern should be different in the No Contact condition – to the extent that events may even be perceived with greater intensity when accompanied by action, as predicted by upweighting theories.

At the start of each trial, participants were cued onscreen to move their right index finger ('move'; Active trials – 50%) or remain stationary ('do not move'; Passive trials – 50%). On Active trials, they moved their finger downwards to a button (Contact blocks) or performed a similar movement without making button contact (No Contact blocks – motion detected via infrared tracking, generating no tactile feedback). The target stimulus was delivered to the left index finger for 30 ms, resulting in apparent synchrony of stimulation with movement. Participants judged the force of this stimulus against a subsequently presented reference. In Passive trials, the target stimulus was delivered 500 ms after the cue to remain still. Modelled Points of Subjective Equivalence (PSEs) reflect the point at which participants judge the target and reference events to have equal force. Lower values indicate more forceful target percepts.

PSE values were analysed in a 2x2 within-participants ANOVA, revealing no main effect of Contact ($F(1, 29) = 3.11$, $p = .089$ $\eta p^2 = .10$) or Movement ($F(1, 29) = 1.24$, $p = .274$, $\eta p^2 = .04$). However, there was a significant interaction between Contact and Movement ($F(1, 29) = 15.39$, $p < .001$, $\eta p^2 = .35$), driven by lower force judgements (higher PSEs) in Active ($M = 4.73$, $SD = 1.22$) compared to Passive trials ($M = 4.35$, $SD = .80$) in the Contact condition ($t(29) = 2.07$, $p = .047$, $d = .38$), but higher force judgements (lower PSEs) in Active ($M = 3.82$, $SD = 1.11$) compared to Passive trials ($M = 4.45$, $SD = .97$) in the No Contact condition ($t(29) = -3.80$, $p = .001$, $d = .69$, see Fig. 1C). Thus, when a moving effector receives cutaneous stimulation simultaneously with passive effector stimulation, tactile events are perceived less intensely during movement. Conversely, when the moving effector does not receive such stimulation, tactile events are perceived more intensely during movement. These results therefore replicate previous findings (10, 11, 15) that tactile events on a passive finger are perceived less intensely during

active movement, but only when the specifics of the paradigm are replicated such that the active finger is also stimulated.

Experiment 2: Probabilistic manipulation of tactile events expected on the basis of action. If the mechanism generating effects in the No Contact condition of Experiment 1 operate according to predictive upweighting mechanisms – given that the tactile event can be anticipated on the basis of action – this mechanism should be sensitive to conditional probabilities between actions and outcomes (19). Therefore, in Experiment 2 we compared perception of tactile events when they were expected or unexpected based on learned action-outcome probabilities established in a preceding training session. Downweighting accounts predict that participants will rate expected events as less intense (forceful), whereas upweighting theories predict that expected tactile stimulation will be rated more intensely, than unexpected stimulation.

Participants performed one of two movements that predicted one of two tactile effects (Fig. 2A). Participants now moved their right index finger upwards or downwards in response to a shape (square or circle) imperative stimulus. Such movement triggered the delivery of a target tactile stimulus to the left index finger or left middle finger. Presenting two action types and two stimulation types allowed us to compare perception of expected and unexpected events, while controlling for repetition effects (see *SI Discussion*). During the training session, each action (e.g., upward movement) was 100% predictive of a specific tactile event (e.g., stimulation to the index finger). In a test session 24 hours later, the action-outcome relationship was degraded to measure perception of unexpected events – the expected finger was stimulated on 66.6% of trials, and the unexpected was stimulated on the remaining 33.3% of trials.

PSE values were lower on expected trials ($M = 3.72$, $SD = .96$) than unexpected trials ($M = 3.93$, $SD = .80$; $t(29) = -2.13$, $p = .041$, $d = .39$, see Fig. 2C), demonstrating that – as predicted under upweighting accounts – expected target events were perceived to be more forceful than unexpected events. This finding is more difficult to reconcile with claims in the action literature that predicted percepts should be attenuated.

Experiment 3: Probabilistic manipulation of expected tactile events when movements are always made towards another effector. Experiment 3 was designed to provide a conceptual replication of Experiment 2, using a set-up more closely aligned with typical action paradigms – whereby one always makes a movement towards another effector. Similarly to Experiment 2, this movement could predict the stimulation or not, on the basis of conditional probabilities established in the training phase. Participants' right hand was vertically aligned with their left hand, such that movements of their right index finger were directed towards their left index finger (Fig. 2B). We hypothesised that – like in Experiment 2 – tactile events predicted on the basis of the preceding movement would be perceived as more intense than unexpected events. Like in Experiment 2, PSE values were lower in expected ($M = 4.08$, $SD = 1.00$) than unexpected ($M = 4.28$, $SD = .99$) trials ($t(29) = -2.56$, $p = .016$, $d = .47$, see Fig. 2C). These results show that tactile events expected on the basis of action were perceived as more forceful than unexpected events, consistent with the results of Experiment 2 and upweighting theories of perception.

Computational modelling: The contribution of sensory (drift) biasing to upweighting. These findings are consistent with predictive upweighting theories of perception, which propose that it is adaptive for observers to combine sampled sensory evidence with prior knowledge – biasing perception towards what we expect (25). This may be achieved mechanistically by altering the weights on sensory channels, increasing the gain of expected relative to unexpected signals (19, 26). However, an alternative explanation is that expectation effects reflect biasing in response-generation circuits – such that action biases people to *respond* that events are more intense when they are expected, rather than altering perception itself (see 40, 41). These different kinds of bias can be dissociated in computational models that conceptualise perceptual decisions as a process of evidence accumulation. Perceptual biases are thought to grow across time –

every time response units sample from perceptual units they will be sampling from a biased representation, therefore increasing the magnitude of biasing effects across a larger number of samples (34, 42). In contrast, response biases are thought to operate regardless of current incoming evidence and to be present from the outset of a trial (43). According to this logic, we can model the decision process with drift diffusion modelling (DDM; 44) to identify the nature of the biasing process. We can thus establish whether action expectations shift the starting point of evidence accumulation towards a response boundary ('start biasing', z parameter; Fig. 3A), or instead bias the rate of evidence accumulation (db parameter, 'drift biasing', Fig. 3B).

We fit hierarchical DDMs to participant choice and reaction time data from Experiment 3 (NB: reaction times were not collected in Experiment 2). We specified four different models: 1) a *null* model where no parameters were permitted to vary between expected and unexpected trials; 2) a *start bias* model where the start point of evidence accumulation (z) could vary between expectation conditions; 3) a *drift bias* model where a constant added to evidence accumulation (db) could vary according to expectation; 4) a *start + drift bias* model where both parameters could vary according to expectation. Overall model fit values are presented in *SI Results*. We conducted a posterior predictive check to evaluate how well simulated data from each of the models could reproduce key patterns in our data (see Materials and Methods for details). Correlations were calculated to quantify how well simulated expectation effects reproduced empirical expectation effects, which revealed significant relationships for all three models (Start bias model: $r_{30}= .39$, $p = .034$; Drift bias model: $r_{30}= .43$, $p = .017$; Start + Drift bias model: $r_{30}= .53$, $p = .003$, see Fig. 3C).

Given that we were interested in whether any of the PSE expectation effect is generated by sensory biasing – rather than possible additional contributions of response biasing – we examined whether drift biasing accounted for any further variance in expectation effects than start biasing alone, by conducting a stepwise linear regression to predict the empirical expectation effect (Unexpected PSE – Expected PSE). In the first step, we included the simulated expectation effect from the start bias model to predict the empirical expectation effect. The simulated start bias data was able to predict the empirical expectation effect ($R^2 = .15$, $F(1,28) = 4.96$, $p = .034$). In the second step, we included the simulated expectation effect from the drift bias model as an additional predictor of the empirical expectation effect, importantly providing a significant improvement to the model fit ($Fchange(1,27) = 6.72$, $p = .015$; $R^2 = .32$, $F(2,27) = 6.34$, $p = .006$). This analysis reveals that a model implementing a drift biasing mechanism better predicts empirical effects of expectation on perceptual decisions, by explaining unique variance in participant decisions that cannot be explained by response biasing.

Discussion

Extant models disagree about how predictions should shape perception of action outcomes. We examined whether sensorimotor prediction attenuates perception of tactile events, as is widely assumed in the action literature, or instead whether predicted events may be perceptually upweighted – in line with theories in the wider sensory cognition literature. We adapted a force judgement paradigm used widely in the action literature, and applied predictive manipulations. Experiment 1 replicated typical findings that self-produced forces were rated as less intense than similar passive forces presented in the absence of action. However, this effect was reversed when cutaneous stimulation was removed from the active finger, consistent with the suggestion that general sensory gating is involved in generating such attenuation, rather than specific prediction mechanisms. Experiments 2 and 3 adapted the paradigm to examine the relative intensity of tactile events expected or unexpected on the basis of conditional probabilities, always in the presence of action. Both of these experiments found that expected tactile action outcomes were perceived *more*, not less, intensely than those that were unexpected. Computational

modelling suggested that expectations alter the way sensory evidence is integrated – increasing the gain afforded to expected tactile signals.

These findings are consistent with upweighting perceptual accounts from outside of action domains, proposing that we use knowledge about statistical likelihoods to bias our percepts towards those that are more probable (19, 20, 24, 25). Such accounts propose that prior expectations – which can be learned through context or statistical regularities – are combined with sensory evidence to generate veridical perceptual interpretations in our noisy environment. The present findings indicate that these mechanisms operate similarly in touch. The findings are harder to reconcile with prominent downweighting theories from action (1–4), which propose that expected action outcomes are perceptually attenuated. It is therefore essential to consider how the present findings can be resolved with the multitude of data cited in support of downweighting theories. A large body of work has attributed attenuation to predictive processes in humans as well as in a variety of other species and sensory systems. For example, attenuating internally-generated electric fields in Mormyrid fish improves detection of prey-like stimuli, a finding thought to be a consequence of predicting self-generated sensory input (45). Similarly, virtual reality trained mice show suppressed auditory responses to self-produced tones generated by treadmill running (46) or licking behaviours (47), compared to no movement. In humans, studies measuring the perceived force of a tactile stimulus during movement show similar attenuation effects during movement relative to no movement, and commonly attribute these effects to predictive mechanisms (e.g. 10, 11).

However, none of the above studies have demonstrated whether underlying mechanisms operate according to stimulus probabilities (19). There are a number of non-predictive mechanisms which could instead explain attenuation, and on the basis of the current findings we propose that some effects are instead generated by identity-general gating mechanisms, and others (e.g. 15) by mechanisms shaping perception according to event repetition (see 6, 18 and *SI Discussion*). Furthermore, it is essential to note that findings of reduced somatosensory (1, 7, 8) or visual (48, 49) neural response when perceiving expected action outcomes – frequently taken to support attenuation theories – may not be reflective of predictive downweighting processes. It has been demonstrated in visual cognition that a global signal reduction can emerge via neural upweighting (e.g., 'sharpening') mechanisms rather than processes proposed in attenuation accounts (19, 29, 30). Future work must unpack precisely what determines these attenuating influences in action, but importantly the present study indicates that present theories are incomplete.

Resolution of these conflicts will be crucial for determining the typical influence of prediction on perception in healthy young populations, as well as older and clinical populations. Tactile attenuation during action – and even more specifically, this particular force judgement paradigm – has been used to demonstrate sensory differences associated with healthy ageing (13), motor severity in Parkinson's disease (14) and hallucinatory severity in schizophrenia (9). Such differences are typically attributed to aberrant prediction of action consequences, but a closer inspection is necessary if the underlying mechanisms are not predictive. It will also be essential for future work to determine the level of overlap and interaction between expectation-based and attention-based processes when examining these mechanisms (26). While expectation influences perception according to statistical likelihood of event occurrence, attention prioritises perceptual information according to task relevance. However, in our natural environment as well as the majority of experimental paradigms, more probable events are also more relevant for task performance. For example, many classic attentional paradigms in fact manipulate stimulus probabilities (50), while expectation manipulations render the relatively more expected also more task relevant (e.g. 51). Future work must disentangle the relative influences of probabilities and task relevance to both upweighting and downweighting mechanisms to understand where these mechanisms can and cannot be dissociated.

To conclude, these findings suggest that sensorimotor prediction may increase, rather than decrease, the perceived intensity of tactile events, likely reflecting the operation of a mechanism that aids veridical perception of action outcomes in a noisy sensory environment. These findings challenge a central tenet of prominent motor control theories and suggest that sensorimotor prediction operates via qualitatively similar mechanisms to other prediction and regardless of the sensory domain.

Materials and Methods

Participants. Thirty participants participated in Experiment 1 (16 female, mean age = 25.53 years [SD = 5.25]), Experiment 2 (20 female, mean age = 22.80 years [SD = 3.18]) and Experiment 3 (22 female, mean age = 24.3 years [SD = 4.34]). Participants were recruited from Birkbeck, University of London and paid a small honorarium for their participation. The sample size was determined a priori on the basis of pilot testing to estimate effect size. The experiment was performed with local ethical committee approval and in accordance with the ethical standards laid down in the 1964 Declaration of Helsinki for all experiments.

Experiment 1 procedure. Participants held their left hand palm upwards (Fig. 1A) with their index fingertip positioned against a solenoid, and their right hand resting on a shelf above. At the start of each trial, participants were cued onscreen to move their right index finger or remain stationary. On Active trials, they rotated their index finger downwards at the metacarpophalangeal joint. When motion was detected – by a button press in Contact blocks and infrared tracking in No Contact blocks – the target stimulus was delivered to the left index finger for 30 ms, resulting in apparent synchrony of stimulation with movement. After 1000 ms, a reference stimulus was presented for 30 ms. The target stimulus presented one of seven logarithmically-spaced forces, and the reference stimulus always presented the fourth (middle) force. After a 300 – 500 ms delay, participants were asked which tap was more forceful, responding with a left foot pedal for the first stimulus and a right foot pedal for the second stimulus. The next trial started after 1000 ms. In Passive trials, the target stimulus was delivered 500 ms after the cue to remain still. Participant responses were modelled by cumulative Gaussians to estimate psychometric functions (see *SI Materials & Methods*).

There were 560 trials in total; 140 for each of the Active and Passive conditions, in both the Contact and No Contact blocks. The order of blocks was counterbalanced across participants and trial type order was randomized. Participants completed eight practice trials before the main test blocks.

Experiment 2 procedure changes relative to Experiment 1. Participants were positioned with their left index and middle finger making contact with independent solenoids. At the beginning of each trial an arbitrary cue (either a square or circle) instructed participants to move their right index finger either upwards or downwards from the metacarpophalangeal joint, tracked by an infrared motion sensor. This action triggered delivery of the target stimulus to either solenoid. During training, participants' action (e.g. downwards movement) was 100% predictive of the location of tactile event (e.g. left index finger). In the test session this relationship was reduced to 66.6% predictive of the tactile event. The training and test sessions were carried out at the same time on consecutive days. There were 420 trials in each session. Trial order was randomised and the action-stimulus mapping was counterbalanced across participants. Participants completed 12 practice trials before the main session trials.

Experiment 3 procedure changes relative to Experiment 2. At the start of each trial, participants selected to make a downwards movement with their right index or middle finger. Participants' hands, and therefore index and middle fingers, were spatially aligned with each other. Participants were asked whether they perceived the test force to be more or less forceful

than the average force intensity. An example of the average force was presented to each finger once every 21 trials. Training and test sessions now occurred in the same session on the same day. There were 210 trials in each session. During the training sessions, movements were still perfectly predictive of tactile events, reduced to 66.6% predictive of the tactile event in test sessions.

Computational modelling. We fit DDMs to participant choice and reaction time data from Experiment 3 using the hDDM package implemented in Python (52). The posterior model parameters for the *start bias*, *drift bias*, and *start + drift bias* models were used to simulate a distribution of 500 reaction times and choices for each trial for each participant. From this simulated data we calculated the probability that a 'stronger than average' response was given at each intensity level, separately for expected and unexpected trials. This allowed us to model simulated psychometric functions for expected and unexpected trials, exactly as we had done for empirical decisions. Performing this procedure for each model yielded separate simulated expectation effects (Unexpected PSE – Expected PSE) for each participant under the *start bias*, *drift bias* and *start + drift bias* models.

Further information is provided in *SI Materials and Methods*. All data and documentation will be deposited in the Birkbeck Research Data Depository (BiRD - https://researchdata.bbk.ac.uk/) and be openly accessible by an associated DOI.

Acknowledgments
The work was supported by a Leverhulme Trust (RPG-2016-105) and Wellcome Trust (204770/Z/16/Z) grant awarded to C.Press. F.deLange was supported by a Vidi Grant (Nederlandse Organisatie voor Wetenschappelijk Onderzoek, 452-13-016) and ERC Starting Grant (Horizon 2020 Framework Programme, 678286).

References
1.   S. J. Blakemore, D. M. Wolpert, C. D. Frith, Central cancellation of self-produced tickle sensation. *Nat. Neurosci.* **1**, 635–640 (1998).
2.   M. Dogge, R. Custers, H. Aarts, Moving Forward: On the Limits of Motor-Based Forward Models. *Trends Cogn. Sci.* **23**, 743–753 (2019).
3.   K. Kilteni, H. H. Ehrsson, Body ownership determines the attenuation of self-generated tactile sensations. *Proc. Natl. Acad. Sci.* **114**, 8426–8431 (2017).
4.   K. Fiehler, E. Brenner, M. Spering, Prediction in goal-directed action. *J. Vis.* **19**, 1–21 (2019).
5.   D. M. Wolpert, J. R. Flanagan, Motor prediction. *Curr. Biol.* **11**, R729–R732 (2001).
6.   C. Press, P. Kok, D. Yon, The Perceptual Prediction Paradox. *Trends Cogn. Sci.* **24**, 13–24 (2020).
7.   S. S. Shergill, *et al.*, Modulation of somatosensory processing by action. *Neuroimage* **70**, 356–362 (2013).
8.   K. Kilteni, H. Ehrsson, Functional connectivity between the cerebellum and somatosensory areas implements the attenuation of self-generated touch. *J. Neurosci.* **40**, 894–906 (2020).
9.   S. S. Shergill, *et al.*, Functional magnetic resonance imaging of impaired sensory prediction in schizophrenia. *JAMA Psychiatry* **71**, 28–35 (2014).
10.   P. M. Bays, J. R. Flanagan, D. M. Wolpert, Attenuation of self-generated tactile sensations is predictive, not postdictive. *PLoS Biol.* **4**, 281–284 (2006).
11.   P. M. Bays, D. M. Wolpert, J. R. Flanagan, Perception of the consequences of self-action is temporally tuned and event driven. *Curr. Biol.* **15**, 1125–1128 (2005).
12.   S. S. Shergill, P. M. Bays, C. D. Frith, D. M. Wolpert, Two Eyes for an Eye: The Neuroscience of Force Escalation. *Science (80-.).* **301**, 187 (2003).

13. N. Wolpe, *et al.*, Ageing increases reliance on sensorimotor prediction through structural and functional differences in frontostriatal circuits. *Nat. Commun.* **7**, 13034 (2016).

14. N. Wolpe, *et al.*, Sensory attenuation in Parkinson's disease is related to disease severity and dopamine dose. *Sci. Rep.* **8**, 15643 (2018).

15. K. Kilteni, C. Houborg, H. H. Ehrsson, Rapid learning and unlearning of predicted sensory delays in self-generated touch. *Elife* **8** (2019).

16. C. Press, R. Cook, Beyond action-specific simulation: domain-general motor contributions to perception. *Trends Cogn. Sci.* **19**, 176–178 (2015).

17. K. Seki, E. E. Fetz, Gating of Sensory Input at Spinal and Cortical Levels during Preparation and Execution of Voluntary Movement. *J. Neurosci.* **32**, 890–902 (2012).

18. C. Press, P. Kok, D. Yon, Learning to Perceive and Perceiving to Learn. *Trends Cogn. Sci.* **24**, 260–261 (2020).

19. F. P. de Lange, M. Heilbron, P. Kok, How Do Expectations Shape Perception? *Trends Cogn. Sci.* **22**, 764–779 (2018).

20. D. J. Heeger, Theory of cortical function. *PNAS* **114**, 1773–1782 (2017).

21. H. Brown, R. A. Adams, I. Parees, M. Edwards, K. Friston, Active inference, sensory attenuation and illusions. *Cogn. Process.* **14**, 411–427 (2013).

22. M. Carrasco, S. Ling, S. Read, Attention alters appearance. *Nat. Neurosci.* **7**, 308–313 (2004).

23. V. Wyart, A. C. Nobre, C. Summerfield, Dissociable prior influences of signal probability and relevance on visual contrast sensitivity. *Proc. Natl. Acad. Sci.* **109**, 3593–3598 (2012).

24. D. Kersten, P. Mamassian, A. Yuille, Object Perception as Bayesian Inference. *Annu. Rev. Psychol.* **55**, 271–304 (2004).

25. A. Yuille, D. Kersten, Vision as Bayesian inference: analysis by synthesis? *Trends Cogn. Sci.* **10**, 301–308 (2006).

26. C. Summerfield, F. P. de Lange, Expectation in perceptual decision making: neural and computational mechanisms. *Nat. Rev. Neurosci.* **15**, 745–756 (2014).

27. C. Press, D. Yon, Perceptual Prediction: Rapidly Making Sense of a Noisy World. *Curr. Biol.* **29**, R751–R753 (2019).

28. M. Gandolfo, P. E. Downing, Causal Evidence for Expression of Perceptual Expectations in Category-Selective Extrastriate Regions. *Curr. Biol.* **29**, 2496–2500.e3 (2019).

29. P. Kok, J. F. M. Jehee, F. P. de Lange, Less Is More: Expectation Sharpens Representations in the Primary Visual Cortex. *Neuron* **75**, 265–270 (2012).

30. D. Yon, S. J. Gilbert, F. P. De Lange, C. Press, Action sharpens sensory representations of expected outcomes. *Nat. Commun.* **9**, 4288 (2018).

31. A. Christensen, W. Ilg, M. A. Giese, Spatiotemporal tuning of the facilitation of biological motion perception by concurrent motor execution. *J. Neurosci.* **31**, 3493–9 (2011).

32. M. Dogge, R. Custers, S. Gayet, H. Hoijtink, H. Aarts, Perception of action-outcomes is shaped by life-long and contextual expectations. *Sci. Rep.* **9**, 5225 (2019).

33. D. Yon, C. Press, Predicted action consequences are perceptually facilitated before cancellation. *J. Exp. Psychol. Hum. Percept. Perform.* **43**, 1073–1083 (2017).

34. D. Yon, V. Zainzinger, F. de Lange, M. Eimer, C. Press, Action biases perceptual decisions toward expected outcomes. *PsyArXiv* (2019) https://doi.org/10.31234/OSF.IO/3ZP8N.

35. F. Kusnir, S. Pesin, G. Moscona, A. N. Landau, When temporal certainty doesn't help. *J. Cogn. Neurosci.* **32**, 315–325 (2019).

36. P. Kok, L. I. Rait, N. B. Turk-Browne, Content-based Dissociation of Hippocampal Involvement in Prediction. *J. Cogn. Neurosci.* **32**, 527–545 (2020).

37. P. Kok, N. B. Turk-Browne, Associative prediction of visual shape in the hippocampus. *J. Neurosci.* **38**, 6888–6899 (2018).

38. S. R. Williams, J. Shenasa, E. C. Chapman, Time course and magnitude of movement-related gating of tactile detection in humans. I. Importance of stimulus location. *J. Neurophysiol.* **79**, 947–963 (1998).

39. S. R. Williams, E. C. Chapman, Time-course and magnitude of movement-related gating

of tactile detection in humans. II. Effect of motor tasks. *J. Neurophysiol.* **84**, 863–875 (2000).

40. C. Firestone, B. Scholl, Cognition does not affect perception: Evaluating the evidence for "top-down" effects. *Behav. Brain Sci.* **39**, e229 (2016).

41. F. P. de Lange, D. A. Rahnev, T. H. Donner, H. Lau, Prestimulus Oscillatory Activity over Motor Cortex Reflects Perceptual Expectations. *J. Neurosci.* **34**, 1400–1410 (2013).

42. A. E. Urai, J. W. de Gee, K. Tsetsos, T. H. Donner, Choice history biases subsequent evidence accumulation. *Elife* **8** (2019).

43. F. P. Leite, R. Ratcliff, What cognitive processes drive response biases? A diffusion model analysis. *Judgm. Decis. Mak.* **6**, 651–687 (2011).

44. R. Ratcliff, G. McKoon, The Diffusion Decision Model: Theory and Data for Two-Choice Decision Tasks. *Neural Comput.* **20**, 873–922 (2008).

45. A. G. Enikolopov, L. F. Abbott, N. B. Sawtell, Internally Generated Predictions Enhance Neural and Behavioral Detection of Sensory Stimuli in an Electric Fish. *Neuron* **99**, 135–146 (2018).

46. D. M. Schneider, J. Sundararajan, R. Mooney, A cortical filter that learns to suppress the acoustic consequences of movement. *Nature* **561**, 391–395 (2018).

47. S. Singla, C. Dempsey, R. Warren, A. G. Enikolopov, N. B. Sawtell, A cerebellum-like circuit in the auditory system cancels responses to self-generated sounds. *Nat. Neurosci.* **20**, 943–950 (2017).

48. J. Stanley, R. C. Miall, Functional activation in parieto-premotor and visual areas dependent on congruency between hand movement and visual stimuli during motor-visual priming. *Neuroimage* **34**, 290–299 (2007).

49. I. Kontaris, A. J. Wiggett, P. E. Downing, Dissociation of extrastriate body and biological-motion selective areas by manipulation of visual-motor congruency. *Neuropsychologia* **47**, 3118–3124 (2009).

50. M. Posner, C. Snyder, B. Davidson, Attention and the detection of signals. *J. Exp. Psychol. Gen.* **109**, 160–174 (1980).

51. G. Rohenkohl, I. C. Gould, J. Pessoa, A. C. Nobre, Combining spatial and temporal expectations to improve visual perception. *J. Vis.* **14**, 8–8 (2014).

52. T. V. Wiecki, I. Sofer, M. J. Frank, HDDM: Hierarchical Bayesian estimation of the Drift-Diffusion Model in Python. *Front. Neuroinform.* **7**, 14 (2013).

53. M. Allen, D. Poggiali, K. Whitaker, T. R. Marshall, R. A. Kievit, Raincloud plots: A multi-platform tool for robust data visualization [version 1; peer review: 2 approved]. *Wellcome Open Res.* **4**, 63 (2019).

Figure 1: Experiment 1. (A) On each trial, participants made downward movements with their right index finger either over a motion tracker (No Contact condition), or towards a button with which they made contact (Contact condition). Each movement elicited a tactile punctate event to the left index finger positioned directly below. (B) PSEs were calculated for each participant (data represents an example participant in the No Contact condition for Active [dark blue] and Passive [light blue] trials). (C) Mean PSEs were higher in Active than Passive trials in the Contact conditions, but lower in Active than Passive trials in the No Contact condition. Larger PSEs indicate less intense target percepts (* $p < .05$, ** $p = .001$). (D) PSE effect of movement (Passive – Active) for the Contact (top) and No Contact (bottom) condition, plotted with raincloud plots (53) displaying probability density estimates (upper) and box and scatter plots (lower). Boxes denote lower, middle and upper quartiles, whiskers denote 1.5 interquartile range, and dots denote difference scores for each participant (N=30). Positive effects of movement indicate more intensely perceived active events relative to passive events, but negative values indicate the reverse – less intensely perceived active events.

Figure 2. (A) Experiment 2. On each trial, participants made a downwards or upwards movement with their right index finger over a motion tracker, which elicited tactile punctate events to the left index or middle finger. Movements were perfectly predictive of tactile events during the training session, and 66.6% predictive during the test session. (B) The design of Experiment 3 was similar to Experiment 2, but participants instead made only downwards movements, now with either their right index or middle finger. (C) Mean PSEs were lower for expected than unexpected trials in both Experiment 2 and Experiment 3. Larger PSEs indicate a less intensely perceived target stimulus (* p < .05). (D) PSE expectation effect (Unexpected – Expected) plotted with raincloud plots (53) displaying probability density estimates (upper) and box and scatter plots (lower), for Experiment 2 (top) and Experiment 3 (bottom). Boxes denote lower, middle and upper quartiles, whiskers denote 1.5 interquartile range, and dots denote difference scores for each participant (N=30). Positive expectation effect values indicate more intensely perceived expected events relative to unexpected events.

Figure 3. Illustration of how the DDM could explain expectation biases, and results of computational modelling. (A) For an unbiased decision process (black lines), sensory evidence integrates towards the upper response boundary when stimuli are stronger than average (solid lines) and towards the lower response boundary when weaker than average (dotted lines). Baseline shifts in decision circuits could shift the start point of the accumulation process nearer to the upper boundary for expected events (influencing the parameter *z*; blue lines - Start bias model). (B) Alternatively, selectively altering the weights on sensory channels could bias evidence accumulation in line with expectations (influencing parameter *db*; red lines – Drift bias model). (C) Simulated Start + Drift bias (winning DIC model) expectation effect plotted against the empirical expectation effect. (D) Simulated Drift bias expectation effect plotted against the empirical expectation effect accounting for simulated Start bias effects (plotted as the residuals from a model where the simulated Start bias effect predicts the empirical effect). Importantly, our regression analysis revealed that drift biases accounted for significant additional variance once accounting for start biases. All expectation effects were calculated by subtracting Expected PSEs from the Unexpected PSEs.

Supplementary Information

SI Materials & Methods

Participants. Eight participants in Experiment 1, six participants in Experiment 2, and nine participants in Experiment 3, were replacements for those where acceptable psychometric functions could not be modelled to their responses (flat functions), where they were unable to follow instructions concerning movement performance (>20% recorded movement errors), or where there was technical malfunction. One participant's PSE scores from Experiment 2 were winsorized to meet the normality assumptions of parametric tests (from z = 3.34 to z = 3; (from z = 3.34 to z = 3; 1).

Experiment 1 procedure. The experiment was conducted in MATLAB using the Cogent toolbox. Participants positioned their left hand palm upwards and held their index fingertip against a downward-facing solenoid (diameter of metal rod = 4 mm; diameter of solenoid = 15 mm; TACT-CONTR-TR2, Heijo Research Electronics), positioned so that the metal rod of the solenoid sat on the apex of the fingertip. Their right hand rested on a shelf, positioned such that the index finger distal phalange was directly above the left hand distal phalange, but rotated 90 degrees anticlockwise relative to their left hand (Fig 1A, Main Text). An infrared motion tracker (Leap Motion Controller using the Matleap MATLAB interface) was placed on the shelf supporting the solenoids at the midpoint between them. Participants' hands were visually occluded during the experiment and white noise was played through headphones (53 Db; piloting confirmed that this level resulted in inaudible solenoid movement) throughout testing.

In Contact blocks, participants' right hand was positioned 5 cm above their left hand, and in No Contact blocks it was moved to 12 cm above to allow movements to be made without touching the shelf. Termination points of movements were approximately the same in Contact and No Contact conditions. It is worth noting that this palm separation generated a difference between our Contact and No Contact conditions additionally to contact. This allowed us to replicate the typical setup in the Contact condition while allowing movement to be registered with the infrared tracking in the No Contact condition. Importantly, all conclusions relate primarily to the simple effects within the Contact and No Contact conditions, so this additional difference should not alter the conclusions.

Experiment 3 procedure changes relative to Experiment 2. We made two further changes in Experiment 3 to remove confounds from Experiment 2. In Experiment 2, a cue instructed participants of actions to perform. In principle, the expectation effects observed in Experiment 2 may have resulted from cue-outcome learning rather than action-outcome learning. Experiment 3 thus removed these cues and required free selection of action. The explicit reference stimulus was also removed from Experiment 3 and comparisons were made against an implicit reference, eliminating the possibility that effects in Experiment 2 were determined by influences of expectation on perception of the reference stimulus.

Independent solenoids were now attached to both fingers via adhesive tape (diameter of metal rod = 4 mm; diameter of solenoid = 18 mm; TactAmp 4.2 Dancer Design). Participants' hands, and therefore index and middle fingers, were spatially aligned with each other. The foot pedals were positioned at either 45 (for stronger) or 90 (for weaker) degree angles relative to their right foot to record responses. The experiment consisted of two training blocks followed by a test block, all occurring in the same session of testing. In the first training block participants responded yes/no to the question 'Tap on index or middle finger?'. In the second training block they were asked about the force, similarly to in Experiment 2 and in subsequent test blocks. Half of the participants experienced a mapping whereby moving the right hand index finger resulted in left hand index stimulation and middle finger movement resulted in middle stimulation. The other half experienced a mapping whereby index finger movement resulted in middle stimulation, and middle finger movement in index stimulation.

Modelling Psychometric Functions. Participant responses were modelled by cumulative Gaussians to estimate psychometric functions, using the Palamedes Toolbox (2) in Matlab. This procedure was performed separately for expected and unexpected trials during the test phase. The mean of the modelled Gaussian was taken as the PSE, describing the point at which participants judge the target and reference events to have equal force. Lower values are indicative of more intense target percepts.

Computational modelling. In the hDDM, model parameters for each participant are treated as random effects drawn from group-level distributions, and Bayesian Markov Chain Monte Carlo (MCMC) sampling is used to estimate group and participant level parameters simultaneously. All models were estimated with MCMC sampling, and parameters were estimated with 30,000 samples ('burn-in'=7,500). Model convergence was assessed by inspecting chain posteriors and simulating reaction time distributions for each participant.

SI Results
Computational modelling. DDM models were compared using deviance information criteria (DIC) as an approximation of Bayesian model evidence. Lower DIC values indicate better model fit. Fitting the DDM to the behavioural data found that the model allowing both start and drift biases to vary according to expectation provided the best fit (DIC relative to null = -234.8) relative to both the start bias (DIC relative to null = -191.06) and drift bias (DIC relative to null = -8.62) models. This finding may suggest that observed biases are a product of both start and drift rate biasing. However, although the DIC measure does include a penalty for model complexity, it is thought to be biased towards models with higher complexity (3) and it indeed favoured the most complex model here. The analyses presented in the Main Text consider whether drift biasing contributes to the effects, which was the question of primary interest.

SI Discussion
Non-predictive mechanisms that may influence perception during action. While downweighting predictive accounts have been influential for the last two decades, evidence that attenuation results from predictive mechanisms is sparse – especially considering the number of non-predictive mechanisms known to influence perception during action (4–7). For this reason, a recent study (8) aimed to determine whether the attenuating influence of action on tactile perception in fact reflected the operation of predictive mechanisms. As outlined in the Main Text, the defining feature of prediction mechanisms is that they operate according to stimulus probabilities (9). This important extension therefore trained participants that sensory events would be delayed by 100 ms, and inverted the typical attenuation effect such that delayed (but expected) events were now perceived less intensely than simultaneous (but unexpected) events. They therefore concluded that we do indeed downweight perception of tactile effects predicted by action – as has been assumed for the last 20 years but never explicitly demonstrated.

However, while the environmental statistics were altered to make delayed events more probable (8), this training protocol confounds prediction with repetition. Participants in the test group were trained in only one repeated mapping – consisting of a 100 ms action-outcome delay – and subsequently presented with events at that delay or at no delay. Given that repeated presentation of stimulus features can attenuate perception of those features (10) and neural sensory processing (11), downweighting effects observed in this study could be determined by repetition of the 100 ms delay rather than prediction of it. The distinction between mechanisms underlying repetition and prediction effects has been scrutinised widely for the last decade and a number of differences emerge. For instance, when repetition and prediction are orthogonalised by sometimes rendering stimulus alternations predicted, repetition and prediction neural effects exhibit distinct time-courses (12). More recently, evidence has emerged that repetition neural effects may be mediated by mechanisms functionally distinct from prediction effects – with

prediction effects perhaps emerging via neural sharpening (13, 14) and repetition effects via local neural scaling (15).

Furthermore, as discussed briefly in the Main Text, another common non-predictive mechanism known to have an influence on perception during action is generalised sensory gating. Previously, the force judgement paradigm has been employed by action researchers partly to allow for dissociation of identity-specific prediction effects from those of identity-general gating effects (16) – where action generates generalised gating of all sensation on a moving effector (17, 18), perhaps mediated by neural mechanisms observed at the earliest relay in the spinal cord (6). These mechanisms are not considered predictive, because they influence perception of events regardless of whether they were predicted outcomes of action, e.g. gaps in ongoing wrist vibrations are not a typical outcome predicted by juggling but are perceptually attenuated during this task (19). Force judgement paradigms are therefore considered optimal for examining the influence of predictive mechanisms because they probe perception of events presented to passive effectors. However, given that the active finger is concurrently stimulated by pressing the trigger device, generalised gating mechanisms may still influence responses – perception of this low intensity stimulation on the active finger could bias responses about the passive stimulation (20). The results of Experiment 1, where the effect is reversed without active finger stimulation, are consistent with this possibility (see also 16).

SI References

1. J. W. Tukey, The Future of Data Analysis. *Ann. Math. Stat.* **33**, 1–67 (1962).
2. N. Prins, F. A. A. Kingdom, Applying the model-comparison approach to test specific research hypotheses in psychophysical research using the Palamedes toolbox. *Front. Psychol.* **9** (2018).
3. T. V. Wiecki, I. Sofer, M. J. Frank, HDDM: Hierarchical Bayesian estimation of the Drift-Diffusion Model in Python. *Front. Neuroinform.* **7**, 14 (2013).
4. C. Press, P. Kok, D. Yon, The Perceptual Prediction Paradox. *Trends Cogn. Sci.* **24**, 13–24 (2020).
5. C. Press, R. Cook, Beyond action-specific simulation: domain-general motor contributions to perception. *Trends Cogn. Sci.* **19**, 176–178 (2015).
6. K. Seki, E. E. Fetz, Gating of Sensory Input at Spinal and Cortical Levels during Preparation and Execution of Voluntary Movement. *J. Neurosci.* **32**, 890–902 (2012).
7. C. Press, P. Kok, D. Yon, Learning to Perceive and Perceiving to Learn. *Trends Cogn. Sci.* **24**, 260–261 (2020).
8. K. Kilteni, C. Houborg, H. H. Ehrsson, Rapid learning and unlearning of predicted sensory delays in self-generated touch. *Elife* **8** (2019).
9. F. P. de Lange, M. Heilbron, P. Kok, How Do Expectations Shape Perception? *Trends Cogn. Sci.* **22**, 764–779 (2018).
10. N. Ofen, A. Moran, D. Sagi, Effects of trial repetition in texture discrimination. *Vision Res.* **47**, 1094–1102 (2007).
11. K. Grill-Spector, R. Henson, A. Martin, Repetition and the brain: neural models of stimulus-specific effects. *Trends Cogn. Sci.* **10**, 14–23 (2006).
12. A. Todorovic, F. P. De Lange, Repetition Suppression and Expectation Suppression Are Dissociable in Time in Early Auditory Evoked Fields. *J. Neurosci.* **32**, 13389 –13395 (2012).
13. P. Kok, J. F. M. Jehee, F. P. de Lange, Less Is More: Expectation Sharpens Representations in the Primary Visual Cortex. *Neuron* **75**, 265–270 (2012).
14. D. Yon, S. J. Gilbert, F. P. De Lange, C. Press, Action sharpens sensory representations of expected outcomes. *Nat. Commun.* **9**, 4288 (2018).
15. A. Alink, H. Abdulrahman, R. N. Henson, Forward models demonstrate that repetition suppression is best modelled by local neural scaling. *Nat. Commun.* **9**, 1–10 (2018).
16. P. M. Bays, D. M. Wolpert, J. R. Flanagan, Perception of the consequences of self-action is temporally tuned and event driven. *Curr. Biol.* **15**, 1125–1128 (2005).

17. S. R. Williams, E. C. Chapman, Time-course and magnitude of movement-related gating of tactile detection in humans. II. Effect of motor tasks. *J. Neurophysiol.* **84**, 863–875 (2000).
18. S. R. Williams, J. Shenasa, E. C. Chapman, Time course and magnitude of movement-related gating of tactile detection in humans. I. Importance of stimulus location. *J. Neurophysiol.* **79**, 947–963 (1998).
19. G. Juravle, C. Spence, Juggling reveals a decisional component to tactile suppression. *Exp. Brain Res.* **213**, 87–97 (2011).
20. C. Firestone, B. Scholl, Cognition does not affect perception: Evaluating the evidence for "top-down" effects. *Behav. Brain Sci.* **39**, e229 (2016).