Cognition 259 (2025) 106112

Contents lists available at ScienceDirect

Cognition

journal homepage: www.elsevier.com/locate/cognit

Observers translate information about other agents' higher-order goals into expectations about their forthcoming action kinematics

Katrina L. McDonough[a,c,1], Eleonora Parrotta[a,b,1,*], Camilla Ucheoma Enwereuzor[a], Patric Bach[a]

[a] *School of Psychology, University of Aberdeen, Scotland, UK*
[b] *Department of Psychology, Sapienza University of Rome, Rome, Italy*
[c] *School of Psychology, University of East Anglia, UK*

ARTICLE INFO

Keywords:
Predictive processing
Social perception
Action prediction
Action observation
Theory of mind
Action goals

ABSTRACT

Social perception relies on the ability to understand the higher-order goals that drive other people's behaviour. Under predictive coding views, this ability relies on a Bayesian-like hypothesis-testing mechanism, which translates prior higher-order information about another agent's goals into perceptual predictions of the actions with which these goals can be realised and tests these predictions against the actual behaviour. We tested this hypothesis in three preregistered experiments. Participants viewed an agent's hand next to two possible target objects (e.g., donut, hammer) and heard the agent state a higher-order goal, which could be fulfilled by one of the two objects (e.g., "I'm really hungry!"). The hand then reached towards the objects and disappeared at an unpredictable point mid-motion, and participants reported its last seen location. The results revealed the hypothesized integration of prior goals and observed hand trajectories. Reported hand disappearance points were predictively shifted towards the object with which the goal could be best realised. These biases were stronger when goal statements were explicitly processed (Experiment 1) than when passively heard (Experiment 2), more robust for more ambiguous reaches, and they could not be explained by attentional shifts towards the objects or participants' awareness of the experimental hypotheses. Moreover, similar biases were not elicited (Experiment 3) when the agent's statements referred to the same objects but did not specify them as action goals (e.g., "I'm really not hungry!"). These findings link action understanding to predictive/Bayesian mechanisms of social perception and Theory of Mind and provide the first evidence that prior knowledge about others' higher-level goals cascades to lower-level action expectations, which ultimately influence the visuospatial representation of others' behaviour.

People effortlessly interpret the meaning behind other people's actions. We only need to see someone reach for a pen to bring to mind their goal of jotting something down, or we see them favour a sandwich over a notebook and perceive it as an expression of their hunger. Social inferences like these demonstrate a key aspect of human social cognition: the ability to "read" the higher-order goals that drive others' behaviours. Making sense of others' actions is essential for navigating the ever-changing world we live in and this ability underpins all our interactions with conspecifics (Sebanz et al., 2006). Yet, the effortlessness with which people derive the hidden mental states driving others' overt behaviour belies the complexity of the underlying process. Understanding others' behaviour requires observers to go beyond the sensory input to solve an inductive problem with multiple solutions (Bach et al., 2014; Bach & Schenke, 2017; Csibra & Gergely, 2007; Jacob & Jeannerod, 2005): there are usually multiple possible combinations of goal states that could explain any single behaviour (e.g., people can smile for various reasons), and the same goal can often be achieved by multiple behaviours (e.g., a mother may tell off her child by speaking up or by making eye contact).

The lack of a direct mapping between overt behaviours and underlying mental states poses difficulties for conventional "direct matching" accounts of social perception, in which observed behaviours are analysed from the bottom-up, and lower-level sensory action features directly activate higher-order motor or semantic knowledge the

---

* Corresponding author at: Department of Psychology, Sapienza University of Rome, Via dei Marsi 78, I-00185 Rome, Italy.
E-mail address: eleonora.parrotta@uniroma1.it (E. Parrotta).

[1] joint first authorship

https://doi.org/10.1016/j.cognition.2025.106112
Received 24 August 2023; Received in revised form 3 March 2025; Accepted 4 March 2025

observer has about these actions, without “inferential processing” or “backwards projections” (e.g., Rizzolatti et al., 2014; Rizzolatti & Sinigaglia, 2010). More recently it has therefore been argued that, to solve such inverse problems, the brain instead adopts a probabilistic Bayesian-like hypothesis testing strategy, as proposed by predictive processing frameworks of perception (Bach & Schenke, 2017; Bubic et al., 2010; Clark, 2013; Den Ouden et al., 2012; Friston & Kiebel, 2009; Markov & Kennedy, 2013; Otten et al., 2017; Yuille & Kersten, 2006) and their extensions to social perception (Bach et al., 2014; Bach & Schenke, 2017; Otten et al., 2017) and Theory of Mind (Baker et al., 2011; Baker & Tenenbaum, 2014; de Bruin & Strijbos, 2015; Koster-Hale & Saxe, 2013). Under these views, action understanding is a hierarchical process of top-down hypothesis testing, in which prior assumptions about others’ goals are translated into specific action expectations through which this mental state would manifest in the given situation, depending on the means of goal pursuit it affords (for explicit proposals for such a translation, see Bach et al., 2014; Bach & Schenke, 2017). By projecting these expectations onto the sensory input, even ambiguous behaviours of others can be imbued with the (inferred) goals. Moreover, action planning can be based on the predicted behaviour of others rather than sensory information available right now, enabling anticipative control of one’s behaviour within shared task spaces (Nijhawan, 1994; Nijhawan, 2002; Sebanz & Knoblich, 2009). In contrast, behaviours that do not match the prior assumptions stand out, prompting revisions of one’s original assumption until a different hypothesis of the agent’s goals provides a better fit (for more details, see Summerfield & Egner, 2009).

We have developed a task that can probe this hypothesis-testing process and reveal the expectations people project onto other people’s actions, in terms of the difference between where people perceive a briefly seen movement to terminate and where it really did. A series of studies from our and other labs have shown that, when people view brief glimpses of other people’s actions, their perceptual reports of what they saw diverge subtly from what was really presented, and that these distortions reflect the prior expectation about how this action will develop (e.g., McDonough & Bach, 2023, McDonough et al., 2019; McDonough et al., 2020; Hudson, Bach, & Nicholson, 2018; Hudson, Nicholson, Ellis, & Bach, 2016; Hudson, Nicholson, Simpson, et al., 2016; for similar findings, see Han, Gandolfo, & Peelen, 2024; Hudson & Jellema, 2011; Ianì et al., 2023; Vandenberghe & Vannuscorps, 2023). For example, when viewing brief glimpses of a hand reaching for an object, the hand is mis-perceived to have moved further when the actor has just said that he would “take” the object than when he would “leave” it, but vice versa for withdrawals (Hudson, Nicholson, Ellis, & Bach, 2016), and these biases increase the more reliably the agent’s statements forecast their actual behaviour (Hudson, McDonough, et al., 2018). Other studies showed that action expectations are spontaneously derived during action observation, correcting the perception of moving hands away from obstacles and towards target objects for example (Hudson, McDonough, et al., 2018; McDonough et al., 2019)McDonough & Bach, 2023)

Importantly, studies to date have always manipulated lower-level action expectations directly (e.g., that the agent would “take” or “leave” a glass of water in Hudson, Nicholson, Simpson, et al., 2016), but not the higher-order information about the agent’s goals from which these actions are hypothesized to emerge (i.e., whether the agent is thirsty). The findings therefore only provide evidence that observers test expectations about forthcoming actions against the observed behaviour, which induces subtle perceptual biases towards the expectations. They say nothing about whether these biases reflect the proposed social hypothesis-testing mechanism that tests observed behaviour against the higher-order goals it is hypothesized to serve. As noted above, a central proposal of predictive processing accounts of action understanding (Bach & Schenke, 2017; Bach et al., 2014; Csibra and Gergely, 2007; Kilner et al., 2007) is that the meaning attributed to other’s behaviour emerges precisely from such a hierarchical translation of higher-order goal information into specific behaviour expectations, and projecting them onto what is indeed observed.

The current study was designed to explicitly test whether predictive biases in action observation can emerge from higher-order goal information. In three pre-registered experiments, we presented participants with a bird’s eye view of an agent reaching from the bottom of the screen to the top, where two objects were located. The objects differed in the higher-order goals they were able to support (see Table 1 for the full set of objects). For example, a cup of tea supports the goal of quenching thirst, while a hammer supports the goal of driving in nails. In Experiments 1 and 2, to induce prior hypotheses about the action’s goal, participants heard the actor make a verbal statement reflecting their current higher-order goal state. These statements were framed in terms of needs (e.g., “I need something refreshing”), intentions (e.g., “Let me jot that down”), or desires (e.g., “I’m craving something salty”), and always matched one of the two objects but not the other (while never explicitly identifying it). For instance, the actor might say, “I’m starving,” and the scene would include a donut and a screwdriver, ensuring that only one object could fulfil the implied goal of satisfying hunger. The actor’s hand would then reach towards the objects but disappeared at an unpredictable point mid-motion. Participants were asked to indicate its last seen location as accurately as possible with the mouse cursor.

This task allows us to reveal participants’ subjective visuospatial representation of the observed actions and how it differs from what was really observed. According to Bayesian/predictive processing accounts of social perception, observers test assumptions about others’ higher-order goals against their observed behaviour. If so, then the agent’s goal statements should affect how objectively identical reach trajectories are represented, biasing them towards the object with which the stated goal can be best achieved. Thus, hearing the actor saying “I am thirsty” would subtly bias the hand’s last seen location towards a drink rather than a pencil, and vice versa if people heard “I need to jot that down”.

We varied, between experiments, whether participants were asked to explicitly evaluate which of the available objects best supports the actor’s stated goal (Experiment 1, Explicit goal processing) or whether they just passively listened to the goal statements (Experiment 2,

Table 1
Audio statement and goal object pairings in each experiment. In Experiment 1 and 2, each of the twelve audio statements conveyed a higher-order goal towards one of the twelve goal objects. In Experiment 3, the statements referred to the same goal states but did *not* convey an intention to interact with the object. These pairings were categorised into one of four object types: tool items (wrench, hammer, screwdriver), stationery items (notepad, pen, scissors), food items (donut, crisps, banana) and drink items (tea, cola, coffee).

Category	Goal Object	Goal Statement (Experiment 1 & 2)	Non-goal Statement (Experiment 3)
Tool	Wrench	That really needs tightening	It really doesn’t need any more tightening
Tool	Screwdriver	I’ll fix that creaky chair	I don’t need to fix that creaky chair anymore
Tool	Hammer	I’ll hang up the painting	I don’t need to hang up the painting anymore
Stationery	Notepad	I’ll jot that down	I don’t need to jot anything down
Stationery	Scissors	I need to cut it in half	I don’t need to cut anything anymore
Stationery	Pen	I’ll check that off my list	I don’t need to check anything off my list
Food	Donut	I’m starving	I really don’t want more food
Food	Crisps	I’m craving something salty	I really don’t want anything salty
Food	Banana	I’m really hungry	I’m really not hungry
Drink	Tea	I’m really thirsty	I’m really not thirsty
Drink	Coffee	My mouth is really dry	I really don’t want any caffeine
Drink	Cola	I need something refreshing	I really don’t want something fizzy

Spontaneous goal processing). This allows us to test to what extent goals are spontaneously translated into precise action expectations, or whether action goal expectations are weighted more when observers explicitly assess the object's suitability for the actor's goal pursuit, as some accounts of Theory of Mind propose (e.g., Schneider et al., 2012; Schneider et al., 2014). Prior work with similar experimental paradigms has revealed an increased weighting of certain action expectations with explicit processing (e.g., expectations of obstacle avoidance, Hudson, McDonough, et al., 2018; McDonough & Bach, 2023), but not for others (e.g., approach expectations towards spatially matching target objects, McDonough, Costantini, Hudson, Ward and Bach, 2020).

Between trials within each experiment, we varied whether the reach was clearly directed leftwards or rightwards towards one of the objects, or whether the hand took a middle path between both objects so that its target remained ambiguous. In Bayesian/predictive processing frameworks, the contribution of priors and sensory signals to perceptual judgments is weighted based on their precision. More reliable (or precise) signals are assigned greater weight, meaning they have a stronger influence on perception. In contrast, less reliable signals are down-weighted and have a reduced impact on the final perceptual outcome (i.e., precision-weighting, e.g., Yon & Frith, 2021). The magnitude of any induced perceptual biases should therefore be present for all reach trajectories (right, left and centre) but be more pronounced for more ambiguous centre reaches than for outer ones, as the latter provide more precise sensory information about the actual action targets.

One concern is that any observed bias could reflect attention directed towards the objects implied by the sentences, rather than an expectation for action towards it. This may be particularly true for Experiment 1, where the instruction to identify which object best supports the agent's goal may be likely to induce such an attentional shift. Experiment 3 therefore replicates Experiment 1 and uses statements that draw attention to the same objects but now express them as non-goals of the agent (e.g., "I am really not hungry anymore!", "I don't need to jot anything down!"). Moreover, like in Experiment 1, participants were explicitly asked to identify, before action onset, which object the agent did *not* want. Thus, if the biases observed in Experiment 1 and 2 reflect the mere shifting of attention to the relevant objects, then they should also be present here, as the statements refer to the same action goals and the instruction to identify the target object itself should lead to a similar shift of attention towards the implied objects. In contrast, if the biases do reflect the hypothesized goal attribution, then they should not be elicited here, where the statement refers to the same objects, but make clear that they do *not* support the actor's current goals. If anything, one may hypothesize that these "non-goals" may induce subtle shifts away from the objects; note however that prior work has consistently suggested that negation induces a *reduced* activation of the negated content, not its reversal to the opposite (Foroni & Semin, 2013; Giora et al., 2005; Liuzza et al., 2011; Tettamanti et al., 2008; Tomasino et al., 2010; Vitale et al., 2022).

To the best of our knowledge, this is the first study to test whether prior knowledge of others' higher-order mental states cascades to lower levels to induce action expectations and is tested against the actual behaviour that is observed. The evidence would therefore show for the first time that the ability to understand people's actions in terms of their thoughts, desires, and inner states (Astington et al., 1988; Gopnik & Meltzoff, 1997; Wellman, 1990) is perceptually instantiated, in terms of the expected action kinematics that these mental states would entail.

1. Methods

1.1. Participants

A total of 162 participants took part across all three experiments (Experiment 1: $n = 54$, 22 females, 31 males, 1 preferred not to say, mean age = 31 years, SD = 6.0; Experiment 2: n = 54, 33 females, 21 males, mean age = 30 years, SD = 6.5; Experiment 3: n = 54, 20 females, 34 males, mean age = 31 years, SD = 6.6). Additional participants (11 in Experiment 1; 9 in Experiment 2; 12 in Experiment 3) were excluded following our pre-registered criteria (see Data pre-processing). All participants gave informed consent, had normal or corrected-to-normal vision and hearing, and were recruited from the Prolific participation pool in exchange for payment. The study was approved by the University of Aberdeen Ethics Committee. Sample size was determined by an a priori power analysis conducted on pilot data. The analysis revealed that a sample size of 54 participants in each experiment provides 90 % power to detect medium size effects (Cohen's $d = 0.45$) with an alpha level of 0.05.

1.2. Apparatus and stimuli

Inquisit (Millisecond) software was used to program and host each experiment online via the Prolific participant recruitment platform.

Stimuli were derived from videos of a male actor's right hand reaching from a rest position at the bottom centre of the screen towards an object that was placed either in the top left, top centre, or top right of the screen. Each reach trajectory was filmed three times, creating a set of nine different reach trajectories. Each video was converted into eight still frames where frame 1 depicted the hand in the same rest position, and frame 8 depicted the hand half-way towards the target object. A final frame was created in which the hand was digitally removed, serving as the response stimulus.

The original objects were removed from all frames. Twelve new object stimuli were created that matched the size of the original objects. The new object stimuli depicted tool items (wrench, hammer, screwdriver), stationery items (notepad, pen, scissors), food items (donut, crisps, banana) and drink items (tea, cola, coffee). Each object stimulus could appear either in the top left or in the top right of the screen (but never in the top centre). We mirrored each object stimulus along the horizontal axis so that their affordance was equal when placed on the top left or the top right. Example stimuli are depicted in Fig. 1A.

For Experiment 1 and 2, twelve audio stimuli were created from voice recordings of a male actor announcing a higher-order goal that could be fulfilled with one of the objects. For example, when the goal was to reach the banana, the actor said "I'm really hungry!" and when the goal was to reach the notepad, he said, "I'll jot that down.". For Experiment 3, twelve corresponding audio stimuli were created by the same voice actor that also identified one of the objects available but did so by announcing the higher-order goal to *not* interact with it. For example, when the (non-)goal was to *not* reach for the banana, the actor said "I'm really not hungry!" and when the (non-)goal was to *not* reach for the notepad, he said "I don't need to jot anything down". See Table 1 for a full list of audio stimuli and associated objects.

1.3. Procedure

For each experiment, participants completed two blocks of 48 trials in which each condition (3 Reach trajectories x 2 Goal object locations x 4 Action sequence lengths) was repeated 2 times, creating a total of 96 trials. Participants were first presented with a fixation cross which they were required to click with their mouse. Clicking the fixation cross made sure that the mouse cursor was centralised at the start of each trial. The mouse cursor was then hidden from the screen for the rest of the trial, to prevent participants from tracing the movements with the cursor. After a randomised interval of 500–1000 ms the first frame of the action sequence appeared. It depicted the actor's hand at rest at the bottom centre of the screen, one object in the top left of the screen and one object in the top right of the screen. The objects were always from different object categories and were chosen at random. After a randomised interval of 1000–2000 ms, the actor announced their higher-order goal, chosen so that it could either be fulfilled by either the object on the left or the object on the right (counterbalanced). In Experiment 1 and 3, participants were instructed to listen carefully to the statement and

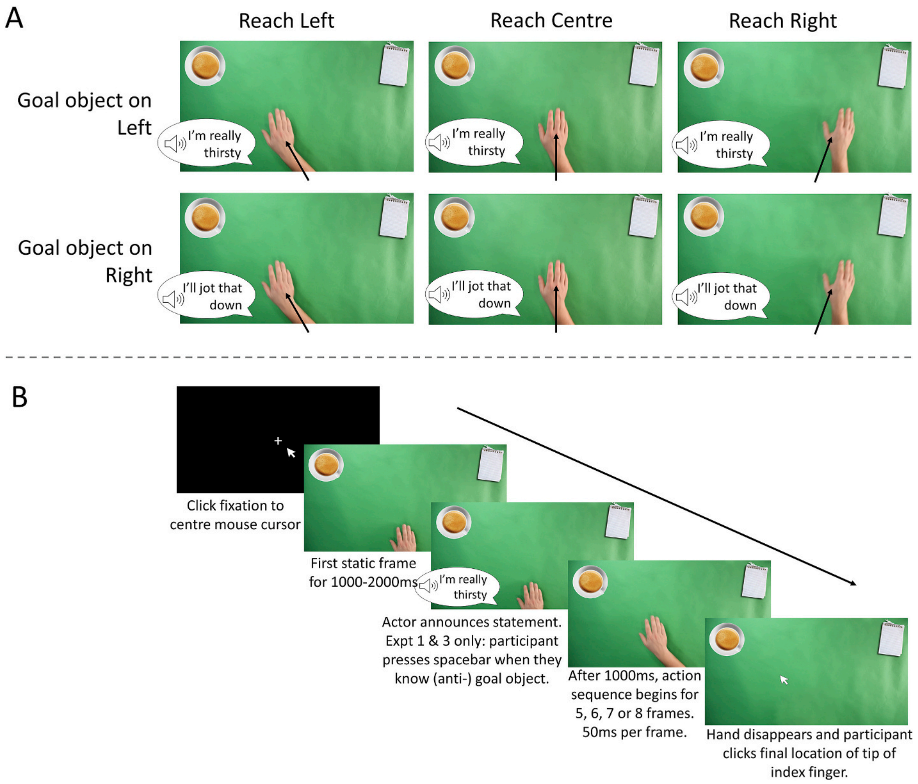

Fig. 1. Stimulus conditions and trial sequence. Panel A depicts each of the Reach trajectory and Goal object location conditions. The reach trajectory was either to the left, to the centre, or to the right. The actor's statement determined which object was the goal object (Experiment 1 & 2) or non-goal object (Experiment 3). The goal/non-goal object was either the left object (top row) or the right object (bottom row). Panel B depicts an example trial sequence. Participants clicked the fixation cross to centre their mouse and then saw the first frame of the action. They then heard the actor's statement followed by the action sequence. Once the hand disappeared, participants clicked the final location of the tip of the actor's index finger using their computer mouse.

think about which object the person might want (Experiment 1) or not want (Experiment 3) and press the space bar once they had identified the (non-)goal object. In Experiment 2, participants just passively listened to the intention statements. After 1000 ms from the spacebar response (Experiment 1 & 3) or audio offset (Experiment 2), the action sequence began. The action sequence was depicted by presenting each frame of the action sequence for 50 ms for a total sequence length of either 5, 6, 7 or 8 frames. It either showed a reach towards the top left object, the top centre or towards the top right object. The final frame of the action sequence was then immediately replaced by the response stimulus, which gave the impression of the hand disappearing from the scene. The mouse cursor was made visible, and participants were asked to click where the last seen position of the tip of the actor's finger was just before it disappeared. The next trial began 1000 ms after this response. An example trial is depicted in Fig. 1B.

At the end of the experiment, participants were asked to provide feedback about the experiment, and to indicate if they had experienced any technical problems, were disrupted during the task, or if they had trouble completing the task (i.e., did not understand the instructions). Participants were also asked to report what they thought the hypothesis of the study was.

1.4. Data pre-processing

Data was pre-processed in line with our pre-registered criteria (Experiment 1 & 2: https://aspredicted.org/BT4_FQS; Experiment 3: https://aspredicted.org/JZ9_LMX). For each participant, individual trials were excluded if the explicit response times when listening to the goal statements (Experiment 1 & 3 only) were shorter than 200 ms (Experiment 1, 17.7 % trials; Experiment 3, 12.1 % trials), and if the mouse click response time were shorter than 200 ms or longer than 3000 ms (Experiment 1, 1.35 % trials; Experiment 2, 2.44 % trials; Experiment 3, 1.08 % trials). Individual trials for which the Pythagorean distance between the real final coordinates and participant's mouse response exceeded 100 pixels were also excluded (Experiment 1, 2.87 % trials; Experiment 2, 11.3 % trials; Experiment 3, 3.26 % trials). Participants who had too few trials remaining (<50 %) were removed (Experiment 1, 2 participants; Experiment 2, 5 participants; Experiment 3, 4 participants). We further excluded participants if they showed an unreliable relationship between their mouse localisation responses and the true hand disappearance point. To this end, like in our prior research (e.g., Hudson, McDonough, et al., 2018), we computed, for each participant separately, the across-trial correlation between the real final coordinates of the hand disappearance point on X and Y-axes in a given trial and their mouse localisation response in the same trial. To the extent that a participant is able to accurately resolve the motions they observe, follow the task instruction, and accurately direct their mouse cursor to the disappearance locations, then their localisations should closely track actual disappearance points. As preregistered, participants were excluded when the resulting correlation coefficient was smaller than 0.70 (Experiment 1, 9 participants; Experiment 2, 4 participants; Experiment 3, 8 participants). Finally, participants were removed if they indicated any reasons in their feedback to suggest that they did not conduct the experiment appropriately, for example, if they were distracted/interrupted during the task (no participants).

2. Results – Experiments 1 and 2

Primary data analysis of Experiment 1 and 2 was conducted on participants' localisation errors following our preregistered analysis plan. All data required to replicate these results is available in Parrotta (2025) ([dataset] Parrotta, E., *High level intentions*, Zenodo, version 1.0, 2025, DOI: https://doi.org/10.5281/zenodo.14917302). Localisation error in each trial was calculated for the X axis and Y axis separately by subtracting the real final coordinates from participants' mouse response coordinates on each trial. As the experiment was conducted online, the screen resolution differed between participants. Therefore, the real final coordinates and participants' mouse response coordinates were standardised onto a 1920 × 1080 resolution. For localisation errors on the X axis, positive scores denote rightwards displacement of participants' reports relative to real disappearance points and negative scores denote leftward displacement. For localisation errors on the Y axis, positive scores denote upwards displacement, and negative scores denote downward displacements. Scoring 0 on both the X and Y axis denotes that the participant selected the real final position exactly.

Participants' mean localisation errors were entered into a 3 × 2 × 2 mixed ANOVA for the X and Y axes separately, with Reach Trajectory (left, centre, right) and Goal object location (left, right) as repeated measures factors and Experiment (1: Explicit goal processing, 2: Spontaneous goal processing) as between-subjects factors. Also following our pre-registered analysis plan, we analysed the 3 × 2 effects separately for each experiment.

In the main ANOVA, the hypothesized effect will be revealed as a main effect of Goal object location in the analysis of localisation errors along the X axis, such that reaches are reported further rightwards when the actor's statement indicates that the right object is the Goal object, and further leftwards when the left object is the Goal. This main effect of Goal object location should be present in both experiments, but should interact with Experiment, being larger when participants were asked to explicitly identify the Goal object (Experiment 1) than when they passively heard the goal statements (Experiment 2). As we had no further predictions, all other effects should be evaluated against a Bonferroni-corrected alpha of 0.004 to account for multiple comparisons in multifactor ANOVAs (Cramer et al., 2016). We did not have any predictions for the Y axis.

2.1. *X-axis*

Primary analyses. As can be seen in Fig. 2a, the localisation errors on the X axis showed an overall bias to the right across all conditions (M = 6.03px), $t(107) = 6.15$, $p < .001$, $d = 0.59$), reflecting that participants generally localised the hand disappearance point more rightwards than it really was. This reflects a known bias for spatial localisation responses away from the tip of the index finger towards the (right) hand's centre of mass (Coren & Hoenig, 1972; Hudson, Bach, & Nicholson, 2018; Hudson, McDonough, et al., 2018; McDonough et al., 2020), which is independent of our hypotheses.

Importantly, the ANOVA revealed the predicted main effect of Goal object location, $F(1,106) = 22.80$, $p < .001$, $\eta_p^2 = 0.177$. Localisation errors were biased more strongly rightwards when the Goal object was on the right (7.28px) than when it was on the left (4.79px), see Fig. 2A. As expected, this effect of Goal object location interacted with Experiment, $F(1,106) = 4.09$, $p = .046$, $\eta_p^2 = 0.037$, with a larger bias towards the goal object in Experiment 1 (Explicit goal processing; 3.56px) than in Experiment 2 (Spontaneous goal processing; 1.44px). Separate 3 × 2 ANOVAs confirmed that the main effect of Goal object location was present in both experiments (Experiment 1: $F(1,53) = 21.9$, $p < .001$, $\eta_p^2 = 0.292$; Experiment 2: $F(1,53) = 4.01$, $p = .050$, $\eta_p^2 = 0.070$, see Fig. 2B & 2C). There were no further main effects or interactions.

Secondary analyses. Our pre-registered secondary analyses investigated how the main effect of Goal object location varied across the three reach trajectories (left, centre, right). We predicted that the bias towards the goal would be present for all Reach Trajectories, but that it would be largest for centre reaches compared to outer reaches (the average of left reaches and right reaches), because left and right reaches provide additional kinematic information of the reach goal while centre reaches remain ambiguous. Pairwise comparisons of the data across both experiments revealed that the main effect of Goal object location was present for all reach trajectories (Left Reach: $t(107) = 3.64$, $p < .001$, $d = 0.35$; Centre Reach: $t(107) = 4.86$, $p < .001$, $d = 0.47$; Right Reach: $t(107) = 2.06$, $p = .042$, $d = 0.20$). Although the biasing effect of goal statements was numerically larger for centre reaches (3.29px) than for outer reaches (2.10px), this difference was not significant, $t(107) = 1.68$, $p = .096$, $d = 0.16$.

In exploratory analyses, we evaluated these findings separately for each experiment. These analyses revealed that the main effect of Goal object location was present in all reach conditions for Experiment 1 (Left Reach: $t(53) = 3.24$, $p = .002$, $d = 0.44$; Centre Reach: $t(53) = 4.71$, $p < .001$, $d = 0.64$; Right Reach: $t(53) = 2.54$, $p = .014$, $d = 0.35$), but only for Centre reaches in Experiment 2 (Left Reach: $t(53) = 1.89$, $p = .065$, $d = 0.26$; Centre Reach: $t(53) = 2.45$, $p = .018$, $d = 0.33$; Right Reach: $t(53) = 0.039$, $p = .969$, $d = 0.01$, see Fig. 2B). Furthermore, the biasing effect of goal statements was numerically larger for centre reaches than for outer reaches in both experiments, although this difference was significant neither in Experiment 1, $t(53) = 0.645$, $p = .522$, $d = 0.09$, nor in Experiment 2, $t(53) = 1.76$, $p = .084$, $d = 0.24$.

Finally, we were interested in whether we would find effects of reach trajectory itself. While we predicted (see Preregistration) that left reaches would be reported as being more leftwards and right reaches reported as more rightwards (i.e., the classic representational momentum effect, Freyd & Finke, 1984; Hudson, Bach, & Nicholson, 2018), we suspected that this effect could be counteracted by a tendency to shift responses closer to the centre (bias towards the mean, Manassi et al., 2024). A pairwise t-test on left reaches vs right reaches revealed no differences, $t(107) = 0.482$, $p = .631$, d=. 05.

2.2. *Y-axis*

Primary analyses. Participants' localisation errors on the Y axis were analysed with the same ANOVA model as the X axis data. We had no predictions for Y-axis data (see Preregistration), as the possible coordinate range of hand disappearance points was less than 30 % than on the X-axis, and particularly at early stages of motions similar upwards motion is expected for all reaches. Even when it becomes clearer later in the trajectory that the hand approaches an unexpected object, the hand is still expected to move further upwards, even if its direction on the X-axis is expected to change. Due to this lack of differential predictions, all effects should be evaluated against a Bonferroni-corrected alpha of 0.004 to account for multiple comparisons in multifactor ANOVAs (Cramer et al., 2016). The results revealed no main effects or interactions that would pass this corrected threshold (all $Fs < 2.3$, all $ps > 0.096$).

In exploratory analyses, we evaluated overall perceptual biases on the Y-axis, as visual inspection of the data suggested a general downward bias, which increased for longer action lengths. A one-sample t-test revealed a general downward bias ($t(107) = 11.57$, $p < .001$, $d = 1.1$), revealing that hands were generally perceived to have reached less far than they really did. A one-factor ANOVA with the factor action length (5, 6, 7 or 8 frames) revealed that this downward bias increased with the length of the action sequence, $F(1,107) = 982.7$, $p < .001$, $\eta_p^2 = 0.902$. This most likely again reflects the known tendency to shift responses closer to the mean of responses (i.e., serial bias, for review, see Manassi et al., 2024).

Secondary analyses. Our pre-registered secondary analyses investigated whether congruent reaches (i.e., a left reach for Goal objects on the left and right reaches for Goal objects on the right) would be reported to have reached further towards the objects (i.e., upwards) than incongruent reaches. This tests for the possibility that participants could expect a hand moving towards an unexpected object to slow down compared to an expected one. Please note that this was unlikely to be observed due to: (a) the more limited range of movement on the Y-axis, (b) that a general upwards motion is still expected even if the action is directed towards an unexpected object, and (c) the fact that any differential expectations could only emerge later in the trajectories when it

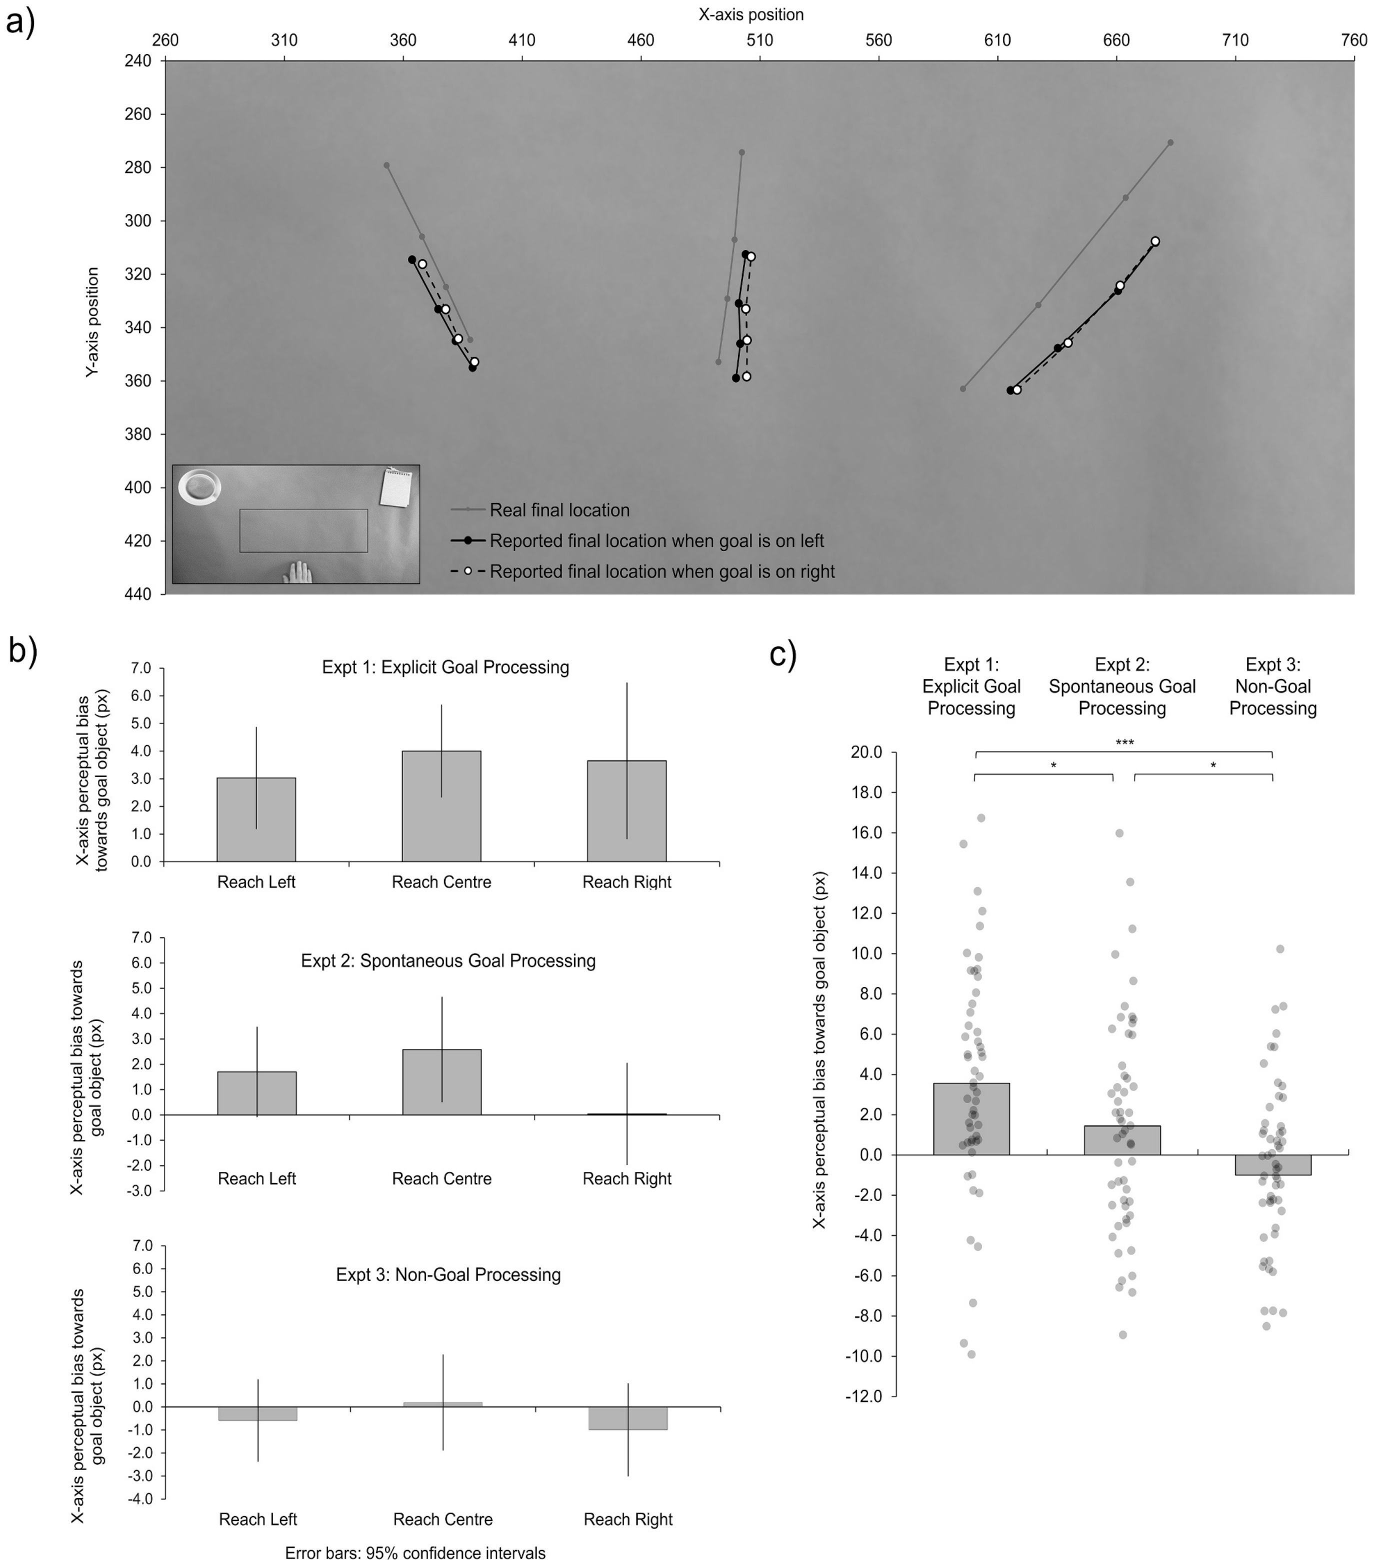

Fig. 2. Panel a) depicts the hand's disappearance points as reported by participants in Experiment 1 and 2, depending on whether a goal object was on the left (solid lines) or on the right (dotted lines), for reaches towards the left, the centre, and the right. The real disappearance points are plotted in grey. The inset shows the area plotted relative to the full stimulus display. Panel b) shows the amount of perceptual bias towards the goal object (Experiment 1 & 2) or non-goal object (Experiment 3), for each experiment separately, depending on whether participants were watching a hand movement to the left, to the centre, or to the right. Error bars represent 95 % confidence intervals. Panel c) depicts participant's individual location errors towards the goal object (or non-goal object) for each Experiment separately. Each datapoint reports the bias towards the (non-)goal object for each participant in each experiment, averaged across all three reach directions.

is clear that the hand was directed towards an unexpected object (note that from the first frame, where an incongruent action may become apparent, only 200 to 350 ms. are available for such an incongruent action direction to be detected and then to influence perceptual judgments, before the last frame is presented, given the possible sequence lengths of 5 to 8 frames of 50 ms. each). Indeed, while this predicted difference was present numerically, it was not significant, $t(107) = 0.905$, $p = .367$, $d =. 09$.

3. Results – Experiment 3

Experiment 3 was identical to Experiment 1, with the exception that the agent's statements were now phrased as non-goals (e.g., "I'm really not hungry!", "I don't need to jot anything down!"). As in Experiment 1, however, participants were asked to identify which of the two objects is implied by the sentence, but now with the requirement to identify "the object the person does not want". This tests for the possibility that (part of) the effects in Experiment 1 and 2 is not due to induced action expectations, but due to simple attentional biases towards the objects implied by the sentences (see also: Additional Analyses). This may appear particularly likely as effect sizes were quadrupled when the task instruction required explicitly identifying the objects implied by the sentences in Experiment 1, which is liable to induce such an attentional shift. If this were the case, however, then any attentional effects should be observed in Experiment 3 as well, as the sentences imply the same objects, and the task instruction similarly requires participants to explicitly identify the (non-wanted) objects before action onset. This instruction should therefore induce strong attentional shifts towards these objects, just like in Experiment 1. In contrast, if our effects reflect the coding of action goals, as we hypothesize, the effects should be substantially reduced or eliminated.[2]

As in Experiment 1 and 2, participant's mean localisation errors were entered into a 3 × 2 repeated measures ANOVA for the X and Y axes separately, with Reach Trajectory (left, centre, right) and Non-Goal object location (left, right) as repeated measures factors. We hypothesized that, in contrast to the main effect of goal object location found for the x-axis in Experiments 1 and 2, the corresponding main effect of Non-goal object location here would be absent, or if anything reversed. As preregistered, equivalence testing was planned to confirm the absence of a main effect of Non-Goal object location, as well as a between-experiment analysis to compare any effect of Non-Goal object location here with the corresponding Goal object location effects in Experiment 1 and 2. As we had no further predictions, all other effects should be evaluated against a Bonferroni-corrected alpha of 0.004 to account for multiple comparisons in multifactor ANOVAs (Cramer et al., 2016). We did not have any predictions for the Y axis.

3.1. X-axis

Primary analyses. The localisation errors on the X axis showed the known (see Experiment 1 and 2) overall bias to the right across all conditions ($M = 7.36$px), $t(53) = 4.64$, $p < .001$, $d = 0.63$), reflecting that participants generally localised the hand disappearance point more rightwards than it really was. Importantly, the ANOVA showed that, as predicted and in contrast to Experiment 1 and 2, no main effect of Non-Goal object location, $F(1,53) = 0.694$, $p = .409$, $\eta_p^2 = 0.013$. There were no further main effects or interactions.

As preregistered, the absence of an effect of Non-Goal object location was confirmed with equivalence testing using a Two One-Sided Tests (TOST) procedure. Since we hypothesized either a null result or an effect in the opposite direction, an inferiority test (Lakens et al., 2018) was performed with no lower bound and an upper bound of $\Delta U = 0.27$ (Cohen's d), which reflects the critical test value based on Experiment 1 and 2. This test confirmed that the observed effect size for the main effect of Non-Goal object location ($d = -0.11$) was significantly within the equivalence bounds, $t(53) = 2.83$, $p = .003$, providing evidence for the absence of the effect. Furthermore, between-experiment *t*-tests revealed that the absent effect of Non-Goal object location here (−.46px) was significantly different to the corresponding main effects of Goal object location in Experiment 1 (Explicit goal processing; 3.56px), $t(106) = 4.28$, $p < .001$, $d = 0.82$, and in Experiment 2 (Spontaneous goal processing; 1.44px), $t(106) = 2.10$, $p = .039$, $d = 0.40$, see Fig. 2C. Thus, while goal statements in Experiment 1 and Experiment 2 induced robust biases in participants' localisation errors, they were absent when the same objects were identified as non-targets of the action in Experiment 3.

Secondary analyses. As with Experiments 1 and 2, pre-registered secondary analyses tested how the main effect of Non-Goal object location varied across the three reach trajectories (left, centre, right). Pairwise comparisons revealed that the main effect of Non-Goal object location was absent for all reach trajectories (Left Reach: $t(53) = -0.511$, $p = .611$, $d = 0.07$; Centre Reach: $t(53) = 0.234$, $p = .816$, $d = 0.03$; Right Reach: $t(53) = -1.03$, $p = .310$, $d = 0.14$, see Fig. 2C), and there was no significant difference between centre reaches (.20px) and outer reaches (−.79px), $t(53) = 0.875$, $p = .385$, $d = 0.12$.

3.2. Y-axis

Primary analyses. Participants' localisation errors on the Y axis were analysed with the same ANOVA model as the X axis data. As for Experiment 1 and 2, we had no predictions for Y-axis data, as (a) the possible coordinate range of hand disappearance points was less than 30 % than on the X-axis, and any differential upwards biases (e.g., an increased downwards bias) could only emerge later in the trajectory when it becomes clear whether the hand approaches an unexpected object. Due to this lack of predictions, all effects on the Y-axis should be evaluated against a Bonferroni-corrected alpha of 0.004 to account for multiple comparisons in multifactor ANOVAs (Cramer et al., 2016). The results revealed no main effects or interactions that would pass this corrected threshold (all *F*s < 1.98, all *p*s > 0.144).

As before, we evaluated overall perceptual biases on the Y-axis, as visual inspection of the data appeared to show a general downward bias, which increased for longer action lengths. A one-sample t-test revealed a general downward bias ($t(53) = 6.50$, $p < .001$, $d = 0.88$), revealing that hands were generally perceived to have reached less far than they really did. A one-factor ANOVA with the factor action length (5, 6, 7 or 8 frames) revealed that this downward bias increased with the length of the action sequence, $F(1, 53) = 20.0$, $p < .001$, $\eta_p^2 = 0.221$.

Secondary analyses. We investigated whether reaches towards the non-goal object mentioned in the sentence (i.e., a left reach for Non-goals on the left and right reaches for Non-goals on the right) would be reported to have reached less closer to the objects than reaches directed towards the other object. In principle, observers could have expected a hand moving towards the non-goal mentioned in the sentence to slow down compared to one directed to the other object. However, like for Experiment 1 and 2, this was unlikely to be observed, and therefore not preregistered. The reason is that any differential expectation could only emerge after motion onset when it has become clear that the hand was directed towards an unexpected object, but maximally 200 to 350 ms are available for such an influence to manifest, given the sequence lengths of 5 to 8 frames with 50 ms duration.

---

[2] Prior research on negation consistently shows that negation reduces or eliminates the activation that a positive sentence otherwise induces but does not invert it (Foroni & Semin, 2013; Giora et al., 2005; Liuzza et al., 2011; Tettamanti et al., 2008; Tomasino et al., 2010; Vitale et al., 2022; Zuanazzi et al., 2024) As such, while we accepted the possibility that an inverted effect (e.g., a bias away from the implied object) might be observed here, it was not expected. Therefore, our preregistration focused on an absent or reduced effect.

4. Additional analyses

4.1. Testing for demand effects

We evaluated, for Experiment 1 and 2, whether the crucial main effect of Goal Object location on the X-axis could have been influenced by participant's ability to guess the hypothesis of the experiment, which may have led them to bias their responses in the desired direction (e.g., Firestone & Scholl, 2016). To do so, each of the three experimenters blindly rated each participant's free text guesses of the experimental hypothesis that was given at the end of the experiment. Responses were scored on a scale of 0 to 2 where 0 reflected completely incorrect guesses of the hypothesis (e.g. measuring reaction times), 1 depicted that some of their guess aligned with the experimental hypothesis (e.g. that the actor's statement was related to one of the objects), and 2 depicted that the response aligned well with the experimental hypothesis (e.g. that responses will be biased towards the object that the actor said they wanted). Scores given from the three individual experimenters for all 108 participants across both experiments were highly correlated, ranging from $r = 0.588$ to $r = 0.794$ (all $p < .001$), suggesting a good level of agreement. Scores from the three experimenters were therefore averaged for each participant to give one hypothesis score per participant. The median score of successful hypotheses guesses was $M = 0.33$, on the scale of 0 to 2, suggesting that most participants had no or little insight into the experimental hypotheses. Correlating this hypothesis score with the magnitude of their perceptual bias towards the goal object location (the crucial contrast value for the main effect of Goal object location across reach trajectories) revealed no relationship across all participants, $r = 0.053$, $t = 0.546$, $p = .586$, nor in either experiment individually (Experiment 1: $r = 0.091$, $t = 0.657$, $p = .514$; Experiment 2: $r = 0.142$, $t = 1.04$, $p = .305$). Participant's ability to identify the hypothesis tested by Experiments 1 and 2 was therefore unrelated to the size of the perceptual biases they showed, providing no indication that the results were due, in part, to demand effects (for a similar result, see Parrotta et al., 2023).

Hypothesis guessing scores in Experiment 3 were analysed analogously, to confirm whether the absence of any effect may be due to participants understanding that no effect was expected here. However, as these were new participants who were not aware of the previous studies, none reported this hypothesis and, as before, most participants had no insight into the experimental hypothesis more generally (median hypothesis guessing score of $M = 0$). Interestingly, those participants who came closer to a correct hypothesis guess (i.e., scores of 1 and 2) either hypothesized an advantage (e.g., in response times or accuracy) towards the (non-goal) object referred to in the sentence (despite the negation), or an advantage for the other object. To test whether these guesses predict their actual biases in location error, we scored those predicting a bias towards the non-goal object positively (+1, +2), those predicting a bias away from it negatively (−1,−2) and omitted those for whom the hypothesized direction of the influence could not be identified. However, like in Experiment 1 and 2, no correlation emerged ($r = 0.023$, $p = .892$) between their guess of the experimental hypothesis and the measured biases in localisation errors towards the relevant object.

4.2. Do the statements induce attentional biases towards the target objects?

In exploratory analyses, we evaluated whether the effects in Experiment 1 and 2 could reflect attentional biases towards the object mentioned in the actor's statements. Simply identifying the object referred to in the actor's statement could elicit a shift of attention towards that object, which may then induce a perceptual bias towards it if it persists over the 1000 ms until the hand starts moving. While Experiment 3 already shows that such unspecific influences are highly unlikely to be responsible for our effects, here we provide two converging measures that the actor's statement did not induce such a shift of spatial attention. If spatial attention was shifted towards the inferred object, then actions spatially closer to this object (e.g., leftwards hand trajectories when the inferred object is on the left) should be represented more precisely – and reported more accurately – than actions away from this object (e.g., leftward trajectories when the inferred object is on the right) (e.g., Fernández et al., 2019, for a review see Anton-Erxleben & Carrasco, 2013.

To assess this possibility, we first compared participants' variability in location responses when reporting the hand disappearance points. If the goal statements in Experiment 1 and 2 drive attention towards the region of space of the inferred object, then participants should be more precise in reporting hands disappearing in this region of space compared to the space around the other object. We therefore calculated, for each participant, each axis, and each combination of reach trajectory and target object location, the standard deviation in localisation error for hands reaching towards the relevant goal object compared to hands reaching towards the other object and compared the results with paired sample *t*-tests. The results revealed no difference in response precision, for neither Experiment 1 (hands reaching towards target object, mean $SD = 28.4$; hands reaching towards alternative object, mean $SD = 29.0$; $t = 0.79$, $p = .433$) nor Experiment 2 (hands reaching towards target object, mean $SD = 30.5$; hands reaching towards alternative object, mean $SD = 30.4$; $t = 0.02$, $p = .850$). This analysis therefore reveals no evidence for a reduction in response uncertainty in the space around the goal object, which would have been expected if attention was shifted towards it.

Second, a converging measure of any attentional bias is provided by participants' hand tracking performance for trajectories towards the inferred target object compared with those towards the alternative object. We calculated, again for each participant, each of the two axes, and each combination of reach trajectory and target object location separately, the trial-by-trial correlations between the hand's real and reported disappearance points. We then averaged the (Fisher-transformed) correlation coefficients across left and right trajectories and X and Y axes and compared, with paired sample t-tests, whether tracking performance was higher when the hands reached towards the goal object's location compared to away from it. The results revealed no difference in tracking performance in either Experiment 1 (hands reaching towards target object, mean $r = 0.84$; hands reaching towards other object, mean $r = 0.83$; $t = 1.12$; $p = .270$) or Experiment 2 (hands reaching towards target object, mean $r = 0.82$; hands reaching towards other object, mean $r = 0.81$; $t = 0.42$; $p = .680$). Like the standard deviation measure, analysis of tracking performance therefore provides no evidence that spatial attention was shifted towards the inferred goal object.

4.3. Item-based analyses

It is important to verify that our results in Experiments 1 and 2 do not only generalize from the tested participants to the population they were sampled from (as confirmed by the inferential statistics reported for each experiment above), but from the sampled stimulus items (the goal statements and objects) to the wider population of similar goal statements people make in everyday life. In an exploratory analysis, we therefore re-ran the main analysis of localisation errors on the X axis but now treated the 12 different goal statements as independent measurements (instead of the 54 participants for each experiment).

The main ANOVA across Experiments 1 and 2 fully replicated the main analysis. It revealed a main effect of Goal object location, $F(1,11) = 82.62$, $p < .001$, $\eta_p^2 = 0.883$, confirming the biasing of localisation errors towards the object referred to in the sentence also in this item-based analysis. Moreover, it replicated the interaction of Goal object location and Experiment, $F(1,106) = 5.31$, $p = .042$, $\eta_p^2 = 0.334$, confirming the larger bias towards the goal object in Experiment 1 (Explicit goal processing) than in Experiment 2 (Spontaneous goal processing). Separate 3 × 2 ANOVAs for each experiment confirmed the main effect of Goal object location in Experiment 1, $F(1,11) = 73.25$, $p < .001$, $\eta_p^2 = 0.869$, but it was not as robust in Experiment 2, $F(1,11) = 3.843$, $p = .076$, $\eta_p^2 = 0.259$. These analyses therefore confirm, at least for Experiment 1, that our results are not driven by individual items but that they generalize from the sampled stimulus items to the large variety of goal statements and objects in other people's everyday interactions with objects.

In an exploratory analysis suggested by a reviewer, we also investigated a possible factor that should determine how strongly each goal statement evoked action expectations. If, as we hypothesize, our effects reflect action expectations derived from higher-order goal information, then such biases should be larger for statements that suggest a more immediate action (i.e., "I'll jot this down", "I'll fix that creaky chair", "I'll hang up the painting", "I need to cut it in half") than for the eight other items for which the evoked action goal is perhaps more remote (e.g., "I am starving", "That really needs tightening"). Indeed, when re-running the main ANOVA across participants of Experiment 1 and 2, split between both item types, it replicated all effects and additionally provided suggestive evidence for larger biases towards the inferred goal object for items that suggest more immediate action goals compared to more remote ones, $F(1,106) = 2.98$, $p = .087$. Moreover, in Experiment 1 the induced biases were robust for both types of goal statement (immediate goal statements, 4.3 pixels, $t = 4.053$, $p < .001$; remote goal statements, 3.1 pixels, $t = 3.632$, $p = .005$). For the spontaneously evoked action expectations in Experiment 2, they could only be robustly demonstrated for the immediate action items (2.0 pixels, $t = 2.514$, $p = .015$) but not for those evoking more remote goals (1.0 pixels, $t = 1.174$, $p = .128$).

The findings of this exploratory analysis therefore add suggestive evidence that action expectations are more readily derived from higher-order goal statements that convey more immediate action goals. Please note however that robust behavioural designs like ours minimize, per definition, differences between participants or items (e.g., Hedge et al., 2018; Parsons et al., 2019), so that any observed variability reflects to a large extent measurement noise instead of true between-item differences. Indeed, in split-half analyses with repeated random allocation of items to each half, average across-item reliability was low in both experiments (average $r \sim 0.30$). Thus, while in line the idea that spontaneous action expectations are derived particularly for more immediate action goals, the results of this analysis must be taken with caution, before being replicated in a study that is designed to test for such between-item differences.

5. Discussion

In three preregistered experiments, we tested a central proposal of Bayesian/predictive processing accounts of action understanding and social perception (Bach & Schenke, 2017; Bach et al., 2014; Csibra and Gergely, 2007; Kilner et al., 2007): that observers test their prior hypotheses about an action's higher-order goal by projecting it onto the behaviour that is actually observed. We showed participants videos of hands reaching towards two objects, each supporting different higher-order goals (e.g., a pen, a glass of coke). To probe their visuospatial representation of these actions, we asked them to report the hand's last seen location after it had disappeared at an unexpected point along its way. We tested whether prior information about the agent's higher-order goals – them saying that they would like to jot something down, drink something refreshing, for example – induces subtle predictive misperceptions of the action kinematics towards the object that matches the goal.

Experiments 1 and 2 revealed, as hypothesized, that participants' visuospatial representation of the action kinematics was biased by higher-order goal information inferred from the actor's stated mental state. Participants (mis-)reported the same hand disappearance points more rightwards when an object with which the agent's higher-order goal could be achieved was on the right and more leftwards when it was on the left. The bias towards the expected action kinematics was present when participants passively heard the goal statements before the actions started (Experiment 2) but was substantially increased when they were explicitly instructed to identify, before action onset, which of the two objects would help the person achieve their goal (Experiment 1).

Importantly, the same biases were not induced (Experiment 3) when the statements referred to the same objects in the scene but now expressed higher-order goal states that the actor currently did *not have* (e.g., "I don't need to jot this down."), and participants had to identify, before the action began, which of the objects the agent did *not* want. One concern was that the observed biases in Experiment 1 and 2 could reflect the allocation of attention towards the objects mentioned in the sentences, rather than their coding as action goals. In particular, Experiment 1's instruction to explicitly identify the object mentioned in the sentence could be expected to induce such an attentional shift. If so, then the same biases towards the objects should be observed in Experiment 3 as well, where the statements referred to the same objects, and participants were similarly asked to identify the relevant (non-goal) objects before action onset. However, no such biases were observed, with their absence being confirmed by both equivalence testing (TOST) procedures and comparisons to Experiments 1 and 2, in line with prior work from psycholinguistics that negation should eliminate or reduce the higher-order goal representations that the positive sentences had otherwise evoked (Foroni & Semin, 2013; Tettamanti et al., 2008; Tomasino et al., 2010; Zuanazzi et al., 2024). Moreover, additional analyses confirmed that key signatures of attentional processing, such as a reduced localisation variability and a heightened ability to distinguish hand disappearance points closer to the relevant object, were absent in Experiments 1 and 2. Together, these findings therefore strongly tie the evoked biases specifically to a coding of the implied objects as (non-) action goals, instead of more unspecific (e.g., attentional) effects that emerge simply from identifying the objects implied by the actor's statements.

The errors in participants' perceptual reports we observe here provide the first evidence for a central tenet of Bayesian/predictive processing accounts of action understanding (Bach & Schenke, 2017; Bach et al., 2014; Csibra and Gergely, 2007; Kilner et al., 2007). They support the proposal for a mechanism that tests higher-order inferences about an agent's goals against their observed behaviour, thereby biasing it in the predicted direction. Prior work had shown that action observation is shaped by the observer's expectations about how the actions will likely develop in the immediate future (Han, Gandolfo, & Peelen, 2024; Hudson, Bach, & Nicholson, 2018; Hudson, Nicholson, Ellis, & Bach, 2016; McDonough et al., 2019; McDonough et al., 2020; Vandenberghe & Vannuscorps, 2023). However, in these studies, action expectations were always induced directly, by giving explicit cues about the forthcoming action (e.g., that someone would "take" or "leave" an object, Hudson, Bach, & Nicholson, 2018) or by constraining which trajectories were more or less likely (McDonough et al., 2019; Vandenberghe & Vannuscorps, 2023). In contrast, ours is the first study to demonstrate that similar perceptual errors can be induced by higher-order information about an actor's goals. In Experiment 1 and 2, the actor's statements never mentioned the relevant goal object directly, but only implicitly referred to it in terms of the actor's distal action intentions, needs, or desires. The statements were therefore ambiguous about which action was likely to follow, unless these goals were integrated with the observer's knowledge of the available objects and goals they supported. The findings are therefore fully in line not only with the notion of a general shaping of perception through expectation, but with the more specific proposal that, during social perception, these expectations can emerge from a hierarchical translation of others' inferred higher-order mental states into the specific lower-order behaviours through which they would be expressed in the current situation (Bach et al., 2014; Bach & Schenke, 2017; Csibra, 2008; Kilner et al., 2007). Indeed, when in Experiment 3 the actor's statements mentioned the same intentions, needs, or desires towards one of the objects, but made clear that they were *not* current mental states of the actor, then no such localisation errors were induced. This puts the results of Experiment 3 in line with prior research from psycholinguistics that negation reduces activation otherwise induced by a (goal) concept (Foroni & Semin, 2013; Giora et al., 2005; Liuzza et al., 2011; Tettamanti et al., 2008; Tomasino et al., 2010; Vitale et al., 2022; Zuanazzi et al., 2024), and with prior research on similar action observation tasks as here that show that even attentionally salient objects do not induce attractive biases if they are not action goals (Hudson, Nicholson, Simpson, et al., 2016; Hudson et al., 2018; McDonough et al., 2020).

The current study provides new insights into core mechanisms that allow our knowledge of other people's mental states – their thoughts, beliefs, and goals – to become perceptually instantiated, so that we can represent others' behaviour in terms of the mental states that drive it (i.e., the intentional stance, Dennett, 1987). Bayesian/predictive processing accounts of social perception (Bubic et al., 2010; Clark, 2013; Den Ouden et al., 2012; Friston & Kiebel, 2009; Markov & Kennedy, 2013) argue that this attribution of meaning emerges from such a hierarchical translation of (inferred) mental states into expected behaviours. By projecting these behaviour expectations onto what was actually observed, people can test whether the inferred mental states can explain what they see other people do. In such a way, even ambiguous behaviour of others could be abstracted away from its visually apparent kinematic features, and be represented in terms of the hidden mental states that explain it. Mismatching behaviours, in contrast, stand out and can prompt revisions of one's mental state attributions until they provide a better fit (Bach & Schenke, 2017; Summerfield & Egner, 2009). Moreover, a predictively enriched representation of others' actions can help resolve uncertainty during motion perception (e.g., during obstructions, etc.) and allows one's own actions to be planned based on the other person's future state rather than on the visuospatial information that is available right now (Nijhawan, 1994, 2002). It has been argued that such an anticipative planning of own behaviour in response to others' predictive behaviour is what makes human social interactions so smooth, dynamic and effective (Sebanz et al., 2006; Sebanz & Knoblich, 2009). Our findings support these accounts and show that the translation of mental states into action expectations can, at least for the higher-order intention statements studied here, occur spontaneously but is largely under cognitive control.

The increase in observed biases when goal statements were explicitly processed in Experiment 1, compared to Experiment 2, contrasts previous studies in which expectation-induced biases did not differ when the relevant cues were task irrelevant or had to be explicitly evaluated. Crucially, as mentioned above, these prior studies required no translation of higher-order goals into behaviours because the predictive cues always directly specified forthcoming actions, for example, when the agent explicitly stated that they would "take" or "leave" an object (Hudson, Bach, & Nicholson, 2018; Hudson, Nicholson, Ellis, & Bach, 2016) or when they formed a hand shape that spatially matched one object but not the other (McDonough et al., 2020). This suggests that what limits such influences in Experiment 2 is not the perceptual testing of expectations against observed behaviours, but the prior *translation* of higher-order goals into lower-order action expectations. This conclusion dovetails with research on Theory of Mind, which also finds that lower-order associations between mental states and behaviour can drive judgments spontaneously (or even automatically), but that deriving precise action expectations from higher-order mental states requires cognitive resources (e.g., Schneider et al., 2012; Schneider et al., 2014). For this reason, research now proposes two separate systems for Theory of Mind, one that operates largely automatically based on simple heuristics and associations, and one for more sophisticated mental state reasoning that requires cognitive resources (e.g., Apperly & Butterfill, 2009; Butterfill & Apperly, 2013). In light of these considerations, it is interesting that a robust (but weaker) biasing effect of intention statement was still observed in Experiment 2 even though action goals and resulting behaviours were task irrelevant. It suggests that the translation of goals into expected actions is efficient enough to occur spontaneously, even if not explicitly prompted, as long as cognitive resources are not otherwise taxed, and perhaps only when the statements suggest immediate action goals (see item-based additional analyses). However, our results also suggest that such spontaneously formed action expectations are less precise and may therefore only affect perceptual judgments of more ambiguous actions.

In this respect, an interesting observation was that the spontaneously emerging biases in Experiment 2 specifically affected the ambiguous centre reaches, but not the left- and rightwards reaches that were clearly directed towards one of the two objects. Such a differential impact of expectation is consistent with the proposal that the integration of different sources of information when predicting an action's next steps depends on the relative precision of the information source (Yon & Frith, 2021). When one source provides highly reliable evidence – such as the observed kinematics clearly favouring one object – it allows prior expectations derived from another source – like the agent's intention statements – to be challenged and revised, reducing their influence (e.g., Ainley et al., 2016; Summerfield & Egner, 2009; see Ambrosini et al., 2015, for a similar finding in anticipative gaze behaviour). A clearly evident action target (e.g., seeing the hand reaching towards the hammer) may therefore have triggered a revision of the prior belief that the person seeks "something refreshing", reducing its influence on perceptual judgments. To our knowledge, this finding would be the first demonstration of a dynamic revision of prior expectation during action observation. However, please note that while such an influence on ambiguous reaches was a priori hypothesized (and preregistered), it was not the core focus of our study, and our methodology was not optimised to draw robust conclusions from such a difference; the finding should therefore be considered tentative until confirmed by further work.

Future work now needs to determine at what level prior expectations are integrated with the visuospatial representation of the observed action. For example, recent studies have shown that expectations about physical state changes (e.g., ice melting, Hafri et al., 2022) are subject to similar biases as the social behaviours here, with an assumed post-perceptual origin in short term memory or decision making. In contrast, research with the current paradigm has pointed to a more low-level perceptual locus. As we found here, the induced biases are unrelated to participants' understanding of the hypotheses (Parrotta et al., 2023) or indeed to their awareness of the manipulation itself (Currie et al., 2023). Moreover, they replicate with purely perceptual response modes without visuomotor biases and very short gaps (e.g., 250 ms) between stimulus disappearance and localisation (Hudson, Bach, & Nicholson, 2018; Hudson, McDonough, et al., 2018; Hudson, Nicholson, Ellis, & Bach, 2016; Hudson, Nicholson, Simpson, et al., 2016), and they can be disrupted by transcranial magnetic stimulation (TMS) of higher-order visual body representation cortex (e.g., Gandolfo & Downing, 2019) or dynamic visual noise masks presented briefly after visual offset (Hudson, McDonough, et al., 2018). These findings link the biases to processes in iconic visual memory that underpins the conscious representation of stimuli (Becker et al., 2000; for a full discussion, Öğmen & Herzog, 2016), instead of later post-perceptual influences. It should be noted of course that, in these prior studies, the action expectations were elicited by lower-level perceptual cues so that any biases could be induced without any top-down translation. It remains to be seen whether a similar penetration of perception may occur for predictive information that is not already perceptually represented.

A second important step for future research is to resolve how observers adjust their inferences when the observed behaviour mismatches expectations. In accounts of predictive social perception (Bach & Schenke, 2017; Kilner et al., 2007; Otten et al., 2017), observers project the inferred mental states onto others' behaviour and revise these inferences in case of a prediction error (Bubic et al., 2010; Clark, 2013; Den Ouden et al., 2012; Friston & Kiebel, 2009; Markov & Kennedy, 2013). So far, research has mainly focussed on demonstrating the proposed projection mechanism – in terms of predictive biases in perceptual judgments like here, or changes to eye movements in other work (e.g., Ambrosini et al., 2015). Few studies have tested whether mismatching behaviour indeed leads to a revision of previously attributed action goals. As noted above, our finding that in Experiment 2 prior goal expectations did not reliably bias perceptual judgments when the action kinematics itself provided clear goal information may provide evidence for such a revision. Other evidence comes from studies testing the attribution of goals to actions directly. These studies have also shown that reliable kinematic information reduces the influence of prior expectations (Betti et al., 2022; Koul et al., 2019). Combining these measures of intention attribution with the current perceptual measures may therefore provide a promising testing bed to trace how intentions are first attributed to observed behaviour and then revised by conflicting evidence.

6. Conclusions

Our study supports the proposal that social perception reflects a top-down hypothesis testing process, in which people project their prior assumptions about other people's goals against their actual behaviour. It provides evidence that people translate information about an agent's higher-order goals into specific expectations for forthcoming body movements they will carry out to achieve these goals, and project them onto the sensory input they receive. The findings may have important implications for the study of social perception and action planning in social interactions and suggest further investigation into how misperceptions in action observation may intersect with everyday experiences. Moreover, they raise intriguing questions about whether these measurable perceptual biases may be altered in clinical populations (Hudson et al., 2012) or whether they change over a lifespan. Furthermore, identifying a perceptual bias associated to the ability to attribute mental states as causes of others' actions provides a fruitful area for further work. The development of this task can then constitute a platform to uncover the neuroscientific basis of perceptual representations of others' mental states as drivers of their actions, as well as to explore other aspects of mentalizing and their impact on action perception (e.g., false beliefs, Wellman et al., 2001). Moreover, here we limited the study to actions involving hands reaching for specific objects, opening new avenues to explore whether the findings extend to other types of actions and human behaviour, such as facial expressions or body postures.

CRediT authorship contribution statement

Katrina L. McDonough: Writing – review & editing, Writing – original draft, Visualization, Validation, Project administration, Methodology, Investigation, Formal analysis, Data curation, Conceptualization. **Eleonora Parrotta:** Writing – review & editing, Writing – original draft, Visualization, Validation, Project administration, Methodology, Investigation, Formal analysis, Data curation, Conceptualization. **Camilla Ucheoma Enwereuzor:** Methodology. **Patric Bach:** Writing – review & editing, Writing – original draft, Visualization, Validation, Supervision, Software, Resources, Project administration, Methodology, Investigation, Funding acquisition, Formal analysis, Data curation, Conceptualization.

Acknowledgements

The work was funded by Leverhulme Trust research grant RPG-2019-248 to PB, and a PhD studentship awarded to Eleonora Parrotta from the Universities of Plymouth and Aberdeen. We thank the members of the Action Prediction Lab (https://www.actionprediction.org), School of Psychology, University of Aberdeen, for insightful discussion, Dr. Gareth Edwards, University of East Anglia, for recording the goal statements, Dr. Lenard Dome for providing us with the original hand motion recordings, and Dr. Joost Rommers for psycholinguistics guidance for experiment 3.

Data availability

The full dataset used for the analyses is available in an open repository: [dataset] Parrotta, E., High level intentions, Zenodo, version 1.0, 2025, DOI: https://doi.org/10.5281/zenodo.14917302

The data used for the analyses is available upon publication in an open repository here: https://zenodo.org/records/14917302

References

Ainley, V., Apps, M. A., Fotopoulou, A., & Tsakiris, M. (2016). 'Bodily precision': A predictive coding account of individual differences in interoceptive accuracy. *Philosophical Transactions of the Royal Society, B: Biological Sciences, 371*(1708), Article 20160003.

Ambrosini, E., Pezzulo, G., & Costantini, M. (2015). The eye in hand: Predicting others' behavior by integrating multiple sources of information. *Journal of Neurophysiology, 113*(7), 2271–2279. https://doi.org/10.1152/jn.00464.2014

Anton-Erxleben, K., & Carrasco, M. (2013). Attentional enhancement of spatial resolution: Linking behavioural and neurophysiological evidence. *Nature Reviews Neuroscience, 14*(3), 188–200.

Apperly, I. A., & Butterfill, S. A. (2009). Do humans have two systems to track beliefs and belief-like states? *Psychological Review, 116*(4), 953.

Astington, J. W., Harris, P. L., & Olson, D. R. (1988). *Developing theories of mind.* CUP Archive.

Bach, P., Nicholson, T., & Hudson, M. (2014). The affordance-matching hypothesis: How objects guide action understanding and prediction. *Frontiers in Human Neuroscience, 8*, 254.

Bach, P., & Schenke, K. C. (2017). Predictive social perception: Towards a unifying framework from action observation to person knowledge. *Social and Personality Psychology Compass, 11*(7), Article e12312.

Baker, C., Saxe, R., & Tenenbaum, J. (2011). Bayesian theory of mind: Modeling joint belief-desire attribution. *Proceedings of the Annual Meeting of the Cognitive Science Society, 33*(33).

Baker, C. L., & Tenenbaum, J. B. (2014). Modeling human plan recognition using Bayesian theory of mind. In *Plan, activity, and intent recognition* (pp. 177–204). Elsevier. https://doi.org/10.1016/B978-0-12-398532-3.00007-5.

Becker, M. W., Pashler, H., & Anstis, S. M. (2000). The role of iconic memory in change-detection tasks. *Perception, 29*(3), 273–286.

Betti, S., Finisguerra, A., Amoruso, L., & Urgesi, C. (2022). Contextual priors guide perception and motor responses to observed actions. *Cerebral Cortex, 32*(3), 608–625.

de Bruin, L., & Strijbos, D. (2015). Direct social perception, mindreading and Bayesian predictive coding. *Consciousness and Cognition, 36*, 565–570.

Bubic, A., Von Cramon, D. Y., & Schubotz, R. I. (2010). Prediction, cognition and the brain. *Frontiers in Human Neuroscience, 25*.

Butterfill, S. A., & Apperly, I. A. (2013). How to construct a minimal theory of mind. *Mind & Language, 28*(5), 606–637.

Clark, A. (2013). Whatever next? Predictive brains, situated agents, and the future of cognitive science. *Behavioral and Brain Sciences, 36*(3), 181–204.

Coren, S., & Hoenig, P. (1972). Effect of non-target stimuli upon length of voluntary saccades. *Perceptual and Motor Skills, 34*(2), 499–508.

Cramer, A. O., van Ravenzwaaij, D., Matzke, D., Steingroever, H., Wetzels, R., Grasman, R. P., … Wagenmakers, E.-J. (2016). Hidden multiplicity in exploratory multiway ANOVA: Prevalence and remedies. *Psychonomic Bulletin & Review, 23*(2), 640–647.

Csibra, G. (2008). Action mirroring and action understanding: An alternative account. *Sensorymotor Foundations of Higher Cognition. Attention and performance, XXII*, 435–459.

Csibra, G., & Gergely, G. (2007). 'Obsessed with goals': Functions and mechanisms of teleological interpretation of actions in humans. *Acta Psychologica, 124*(1), 60–78.

Currie, J., Giannaccini, M. E., & Bach, P. (2023). *Consequential sound induces illusory distortions in the perception and prediction of robot motion*.

Den Ouden, H. E., Kok, P., & De Lange, F. P. (2012). How prediction errors shape perception, attention, and motivation. *Frontiers in Psychology, 3*, 548.

Dennett, D. C. (1987). *The intentional stance*. MIT press.

Fernández, A., Li, H. H., & Carrasco, M. (2019). How exogenous spatial attention affects visual representation. *Journal of Vision, 19*(11), 4-4.

Firestone, C., & Scholl, B. J. (2016). Cognition does not affect perception: Evaluating the evidence for "top-down" effects. *Behavioral and Brain Sciences, 39*.

Foroni, F., & Semin, G. R. (2013). Comprehension of action negation involves inhibitory simulation. *Frontiers in Human Neuroscience, 7*. https://doi.org/10.3389/fnhum.2013.00209

Freyd, J. J., & Finke, R. A. (1984). Representational momentum. *Journal of Experimental Psychology: Learning, Memory, and Cognition, 10*(1), 126.

Friston, K., & Kiebel, S. (2009). Cortical circuits for perceptual inference. *Neural Networks, 22*(8), 1093–1104.

Gandolfo, M., & Downing, P. E. (2019). Causal evidence for expression of perceptual expectations in category-selective extrastriate regions. *Current Biology, 29*(15), 2496–2500.e3. https://doi.org/10.1016/j.cub.2019.06.024

Giora, R., Fein, O., Ganzi, J., Levi, N. A., & Sabah, H. (2005). On negation as mitigation: The case of negative irony. *Discourse Processes, 39*(1), 81–100. https://doi.org/10.1207/s15326950dp3901_3

Gopnik, A., & Meltzoff, A. N. (1997). *Words, thoughts, and theories.* Mit Press.

Hafri, A., Boger, T., & Firestone, C. (2022). Melting ice with your mind: Representational momentum for physical states. *Psychological Science, 33*(5), 725–735.

Han, Q., Gandolfo, M., & Peelen, M. V. (2024). Prior knowledge biases the visual memory of body postures. *Iscience, 27*(4).

Hedge, C., Powell, G., & Sumner, P. (2018). The reliability paradox: Why robust cognitive tasks do not produce reliable individual differences. *Behavior Research Methods, 50*(3), 1166–1186. https://doi.org/10.3758/s13428-017-0935-1

Hudson, M., Bach, P., & Nicholson, T. (2018). You said you would! The predictability of other's behavior from their intentions determines predictive biases in action perception. *Journal of Experimental Psychology: Human Perception and Performance, 44*(2), 320.

Hudson, M., Burnett, H. G., & Jellema, T. (2012). Anticipation of action intentions in autism spectrum disorder. *Journal of Autism and Developmental Disorders, 42*, 1684–1693.

Hudson, M., & Jellema, T. (2011). Resolving ambiguous behavioral intentions by means of involuntary prioritization of gaze processing. *Emotion, 11*(3), 681.

Hudson, M., McDonough, K. L., Edwards, R., & Bach, P. (2018). Perceptual teleology: Expectations of action efficiency bias social perception. *Proceedings of the Royal Society B, 285*(1884), Article 20180638.

Hudson, M., Nicholson, T., Ellis, R., & Bach, P. (2016). I see what you say: Prior knowledge of other's goals automatically biases the perception of their actions. *Cognition, 146*, 245–250.

Hudson, M., Nicholson, T., Simpson, W. A., Ellis, R., & Bach, P. (2016). One step ahead: The perceived kinematics of others' actions are biased toward expected goals. *Journal of Experimental Psychology: General, 145*(1), 1.

Ianì, F., Limata, T., Bucciarelli, M., & Mazzoni, G. (2023). The implicit effect of action mental simulation on action evaluation. *Quarterly Journal of Experimental Psychology, 76*(2), 257–270.

Jacob, P., & Jeannerod, M. (2005). The motor theory of social cognition: A critique. *Trends in Cognitive Sciences, 9*(1), 21–25.

Kilner, J. M., Friston, K. J., & Frith, C. D. (2007). Predictive coding: An account of the mirror neuron system. *Cognitive Processing, 8*(3), 159–166.

Koster-Hale, J., & Saxe, R. (2013). Theory of mind: A neural prediction problem. *Neuron, 79*(5), 836–848. https://doi.org/10.1016/j.neuron.2013.08.020

Koul, A., Soriano, M., Tversky, B., Becchio, C., & Cavallo, A. (2019). The kinematics that you do not expect: Integrating prior information and kinematics to understand intentions. *Cognition, 182*, 213–219.

Lakens, D., Scheel, A. M., & Isager, P. M. (2018). Equivalence testing for psychological research: A tutorial. *Advances in methods and practices in psychological science, 1*(2), 259–269.

Liuzza, M. T., Candidi, M., & Aglioti, S. M. (2011). Do not resonate with actions: Sentence polarity modulates Cortico-spinal excitability during action-related sentence reading. *PLoS ONE, 6*(2), Article e16855. https://doi.org/10.1371/journal.pone.0016855

Manassi, M., & Whitney, D. (2024). Continuity fields enhance visual perception through positive serial dependence. *Nature Reviews Psychology, 3*(5), 352–366.

Markov, N. T., & Kennedy, H. (2013). The importance of being hierarchical. *Current Opinion in Neurobiology, 23*(2), 187–194.

McDonough, K. L., & Bach, P. (2023). Expectations of efficient actions bias social perception: a pre-registered online replication. *Royal Society Open Science, 10*(2), 220889.

McDonough, K. L., Costantini, M., Hudson, M., Ward, E., & Bach, P. (2020). Affordance matching predictively shapes the perceptual representation of others' ongoing actions. *Journal of Experimental Psychology: Human Perception and Performance, 46*(8), 847.

McDonough, K. L., Hudson, M., & Bach, P. (2019). Cues to intention bias action perception toward the most efficient trajectory. *Scientific Reports, 9*(1), 1–10.

Nijhawan, R. (1994). Motion extrapolation in catching. *Nature, 370*, 256–257.

Nijhawan, R. (2002). Neural delays, visual motion and the flash-lag effect. *Trends in Cognitive Sciences, 6*(9), 387–393.

Öğmen, H., & Herzog, M. H. (2016). A new conceptualization of human visual sensory-memory. *Frontiers in Psychology, 7*. https://doi.org/10.3389/fpsyg.2016.00830

Otten, M., Seth, A. K., & Pinto, Y. (2017). A social Bayesian brain: How social knowledge can shape visual perception. *Brain and Cognition, 112*, 69–77.

Parrotta, E., McDonough, K., & Bach, P. (2023). Imagination as predictive perception: Mental imagery predictively biases perceptual judgments of observed action kinematics. *bioRxiv*, 2023–06.

Parsons, S., Kruijt, A.-W., & Fox, E. (2019). Psychological science needs a standard practice of reporting the reliability of cognitive-behavioral measurements. *Advances in Methods and Practices in Psychological Science, 2*(4), 378–395. https://doi.org/10.1177/2515245919879695

Rizzolatti, G., Cattaneo, L., Fabbri-Destro, M., & Rozzi, S. (2014). Cortical mechanisms underlying the organization of goal-directed actions and mirror neuron-based action understanding. *Physiological Reviews, 94*(2), 655–706.

Rizzolatti, G., & Sinigaglia, C. (2010). The functional role of the parieto-frontal mirror circuit: Interpretations and misinterpretations. *Nature Reviews Neuroscience, 11*(4), 264–274.

Schneider, D., Lam, R., Bayliss, A. P., & Dux, P. E. (2012). Cognitive load disrupts implicit theory-of-mind processing. *Psychological Science, 23*(8), 842–847.

Schneider, D., Nott, Z. E., & Dux, P. E. (2014). Task instructions and implicit theory of mind. *Cognition, 133*(1), 43–47.

Sebanz, N., Bekkering, H., & Knoblich, G. (2006). Joint action: Bodies and minds moving together. *Trends in Cognitive Sciences, 10*(2), 70–76.

Sebanz, N., & Knoblich, G. (2009). Prediction in joint action: What, when, and where. *Topics in Cognitive Science, 1*(2), 353–367.

Summerfield, C., & Egner, T. (2009). Expectation (and attention) in visual cognition. *Trends in Cognitive Sciences, 13*(9), 403–409.

Tettamanti, M., Manenti, R., Della Rosa, P. A., Falini, A., Perani, D., Cappa, S. F., & Moro, A. (2008). Negation in the brain: Modulating action representations. *NeuroImage, 43*(2), 358–367. https://doi.org/10.1016/j.neuroimage.2008.08.004

Tomasino, B., Weiss, P. H., & Fink, G. R. (2010). To move or not to move: Imperatives modulate action-related verb processing in the motor system. *Neuroscience, 169*(1), 246–258. https://doi.org/10.1016/j.neuroscience.2010.04.039

Vandenberghe, A., & Vannuscorps, G. (2023). Predictive extrapolation of observed body movements is tuned by knowledge of the body biomechanics. *Journal of Experimental Psychology: Human Perception and Performance, 49*(2), 188.

Vitale, F., Monti, I., Padrón, I., Avenanti, A., & de Vega, M. (2022). The neural inhibition network is causally involved in the disembodiment effect of linguistic negation. *Cortex, 147*, 72–82. https://doi.org/10.1016/j.cortex.2021.11.015

Wellman, H. M. (1990). *The child's theory of mind*. Cambridge, MA: Bradford. MIT press.

Wellman, H. M., Cross, D., & Watson, J. (2001). Meta-analysis of theory-of-mind development: The truth about false belief. *Child Development, 72*(3), 655–684.

Yon, D., & Frith, C. D. (2021). Precision and the Bayesian brain. *Current Biology, 31*(17), R1026–R1032.

Yuille, A., & Kersten, D. (2006). Vision as Bayesian inference: Analysis by synthesis? *Trends in Cognitive Sciences, 10*(7), 301–308. https://doi.org/10.1016/j.tics.2006.05.002

Zuanazzi, A., Ripollés, P., Lin, W. M., Gwilliams, L., King, J.-R., & Poeppel, D. (2024). Negation mitigates rather than inverts the neural representations of adjectives. *PLoS Biology, 22*(5), Article e3002622. https://doi.org/10.1371/journal.pbio.3002622