HUMAN-ROBOT INTERACTION

When performing actions with robots, attribution of intentionality affects the sense of joint agency

Uma Prashant Navare[1,2]*†, **Francesca Ciardo**[1,3], **Kyveli Kompatsiari**[1,4], **Davide De Tommaso**[1], **Agnieszka Wykowska**[1]*

Sense of joint agency (SoJA) is the sense of control experienced by humans when acting with others to bring about changes in the shared environment. SoJA is proposed to arise from the sensorimotor predictive processes underlying action control and monitoring. Because SoJA is a ubiquitous phenomenon occurring when we perform actions with other humans, it is of great interest and importance to understand whether—and under what conditions—SoJA occurs in collaborative tasks with humanoid robots. In this study, using behavioral measures and neural responses measured by electroencephalography (EEG), we aimed to evaluate whether SoJA occurs in joint action with the humanoid robot iCub and whether its emergence is influenced by the perceived intentionality of the robot. Behavioral results show that participants experienced SoJA with the robot partner when it was presented as an intentional agent but not when it was presented as a mechanical artifact. EEG results show that the mechanism that influences the emergence of SoJA in the condition when the robot is presented as an intentional agent is the ability to form similarly accurate predictions about the sensory consequences of our own and others' actions, leading to similar modulatory activity over sensory processing. Together, our results shed light on the joint sensorimotor processing mechanisms underlying the emergence of SoJA in human-robot interaction and underscore the importance of attribution of intentionality to the robot in human-robot collaboration.

INTRODUCTION

Theoretical background

In recent years, humanoid robots, and their potential roles as assistants, educators, or social companions in elderly and health care, have been receiving increasing attention. It is predicted that robots will soon enter our lives in such roles (*1*), and research on human-robot interaction (HRI) and artificial intelligence has increasingly focused on developing solutions that enable humanoid robots to do so (*2*). However, humanoid robots have another, less widely recognized potential: as tools to study mechanisms of human cognition (*3*). On one hand, as embodied agents that can interact with the environment, humanoid robots can be used as the apparatus in sophisticated and, crucially, interactive experimental setups for cognitive neuroscience research, allowing for a higher degree of ecological validity than traditionally used screen-based stimuli. On the other hand, because they can be programmed to perform specific and repeatable behaviors, they offer higher experimental control than studying cognition in human-human interactions (*4*). Last, but certainly not least, it is possible to experimentally manipulate certain factors with the use of robots (for example, the degree to which a behavior or appearance is perceived as social or intentional), although this cannot be done with other humans, in whose case intentionality and socialness are attributed by default. Together, these features render humanoid robots and HRI paradigms powerful tools to examine the mechanisms underlying a variety of cognitive processes (*3*, *4*). Once those mechanisms are understood, one can use this knowledge to design robots that are better tailored to the way the human brain works for more efficient collaboration and perhaps more intuitive interaction.

As humans, we do not act in a social vacuum, and most of our actions require coordination with others in space and time to achieve a goal. A crucial aspect of interacting with others is the experience of a sense of joint agency (SoJA). When one performs an action together with someone—for example, a piano duet—one can form a SoJA with one's partner, meaning that one will perceive the sounds coming out of the instrument as a result of joint control over those sounds, regardless of who in the dyad actually pressed the specific key producing a specific sound. Various terms have been used in the literature to refer to this sense of agency (SoA) that one experiences in joint action, including "SoJA" and "joint sense of agency"—the former being the term we use throughout the present work because it reflects better the idea of experiencing (sensing) agency over action outcomes as joint (rather than separate and individual). Thus, SoJA refers to the feeling of control that we experience both for our own and our partner's actions and the subsequent outcomes when acting together with our partner (*5*).

In scientific literature, SoJA has been studied experimentally with the use of temporal interval estimation tasks (*6*, *7*). In such tasks, participants perform an action (commonly a button press), which elicits a sensory outcome (for example, a tone) and are then asked to judge the temporal duration between those two events. Typically, for individual actions, intervals are perceived as shorter for one's own voluntary, intentional actions as compared with passive or coerced actions (*8*–*9*). Thus, shorter intervals indicate a greater SoA. To study SoJA, the interval estimation paradigm is typically used while agents perform a task together. If SoJA is formed, then participants typically perceive the temporal interval similarly in terms of length for both their own and their partner's actions. If SoJA is not formed, then intervals between self-generated actions and the subsequent outcomes are perceived as shorter than intervals between partner-generated actions and outcomes (*6*, *7*). For example, Jenkins *et al.* (*6*) showed that when asked to

---

[1] Social Cognition in Human-Robot Interaction, Italian Institute of Technology, 16152 Genova, Italy. [2] Department of Computer Science, Faculty of Science and Engineering, University of Manchester, M13 9PL, Manchester, United Kingdom. [3] Department of Psychology, University of Milan-Bicocca, Milan, Italy. [4] Section for Cognitive Systems, DTU Compute, Kgs. Lyngby, Copenhagen, Denmark.
*Corresponding author. Email: unavare@gmail.com (U.P.N.); agnieszka.wykowska@iit.it (A.W.)
†Present address: ETIS UMR 8051, CY Cergy Paris Université, ENSEA, CNRS, F95000 Cergy, France.

complete a complementary joint task with a shared outcome (a tone), people exhibited identical temporal compression between actions and tones when they produced the tone as when their partner did, indicating emergence of SoJA (*6*).

Previous work suggests that SoJA seemingly arises from the motor control and monitoring processes that underlie joint action (*10*, *11*). This work is based on theories of individual action. According to prominent theories of motor control in individual action, two types of internal models are implemented in the central nervous system for the optimization and control of motor actions: the inverse model and the forward model. The inverse model is involved in calculating the (exact) motor command that will achieve a specific goal or desired change in state, whereas forward models compute the expected, or predicted, sensory outcomes and motor states that will be generated by a given motor command (*12*). Together, these computations are thought to underlie the precise execution and control of ongoing actions (*12*). Joint action requires predictions to be made both for one's own and the partner's actions and the outcomes. Thus, it has been proposed that, during joint action, both self- and partner-internal models are maintained (*10*, *11*, *13*, *14*) to support the joint action. According to several theories, during a joint action, both self- and partner-internal models are integrated into a joint internal model, forming a joint motor plan that takes into account the contributions of all partners toward the shared goal (*10*, *13*, *15*). Consequently, SoJA is thought to arise from comparing predicted outcomes and motor states, arising from both self- and partner-forward models, with actual sensory inputs and motor states (*10*, *16*).

It is intuitive then that SoJA does not seem to emerge when engaging in joint action with partners whose actions cannot be represented at the sensorimotor level or predicted using a human motor repertoire. For example, Sahaï *et al.* (*7*) showed that participants did not exhibit temporal compression between actions and outcomes when playing in joint action with a desktop machine (whose motor repertoire is different from that of humans), although they did when completing the task with a human partner.

However, the perceived intentionality of a partner may also be an important factor influencing SoJA in HRI. It has been suggested in previous work that attributing mental states or intentionality to a robot can increase the predictability of their behavior, can reduce stress or uncertainty, and can increase the sense of control experienced during interaction with a robot (*17*). Moreover, it is important to keep in mind that at the individual level, SoA occurs only for voluntary (intentional) actions. When sensory outcomes resulted from involuntary motor events [induced by transcranial magnetic stimulation (TMS)], the temporal compression effect was not observed (*8*). By extension, it is therefore plausible to assume that the intentionality of the partner's actions is a crucial factor for SoJA to emerge. Some evidence in this direction has been found in joint action literature. For example, in a joint Simon task, participants formed a shared task representation when interacting with a humanoid robot that was introduced as intentional and humanlike, but they did not do so when the robot was introduced as a passive and deterministic machine (in both conditions, unbeknownst to the participants, the robot was controlled by a human operator) (*18*). When people form a shared representation of a task with their partner, they include their partner's actions in their representation of the task and their action planning. Thus, the emergence of a shared task representation indicates the integration of self and partner models into a joint motor plan (*10*), a process that is also thought to underlie the emergence of SoJA in human-human interactions (*11*). Therefore, it is plausible that perceiving one's partner as intentional can affect the formation of a shared task representation and potentially also the emergence of SoJA with them. In the study of Sahaï *et al.* (*7*) described above, it is unlikely that participants perceived the desktop machine as a partner with intentionality, which may also have influenced the fact that SoJA did not emerge with the desktop machine. Furthermore, in another study, Sahaï *et al.* (*19*) found that participants did not experience SoJA when interacting with a servomotor device (presumably perceived as unintentional). Moreover, although participants in the study did seem to form a SoJA with a humanoid robot (which, according to the authors, may have arisen because of participants' attribution of a certain degree of intentional agency to the humanoid robot, although this was not explicitly measured in the study), overall, they experienced greater temporal compression (higher SoA) for self- and other-generated actions and outcomes when acting with a human partner as compared with the servomotor or humanoid robot partner (*19*).

Despite these hints in the literature suggesting that attribution of intentionality to an interaction partner might be crucial for the emergence of SoJA, especially when interacting with artificial agents, to our knowledge, no previous work has systematically investigated this issue. More specifically, to answer the question of whether the emergence of SoJA is related to intentionality attribution and not, for example, to sensorimotor repertoire of the partner, one needs to manipulate perceived intentionality while keeping the partner's sensorimotor repertoire constant and, crucially, similar to a human. The present study took exactly this approach. The aim was to evaluate whether SoJA is influenced by participants' perception of their partner as an intentional agent. To address this aim, we used the humanoid robot iCub (*20*) as an interaction partner. As a humanoid, iCub has a sensorimotor repertoire similar to that of a human, while at the same time, belief in iCub's intentionality can be manipulated experimentally.

Implementation of the SoJA paradigm with iCub

In our study, we asked participants to complete a task adapted from (*6*) in two conditions: alone and in joint action with iCub (Fig. 1). The task consisted of moving an onscreen cursor onto a target location and subsequently confirming the position of the cursor on the target, thereby triggering a tone (the so-called target-tracking-and-confirmation task). In the solo role, participants performed the task alone, bimanually. In the joint roles, participants either moved the cursor onto the target (mover role), whereas the robot subsequently confirmed and thus triggered the tone, or they confirmed the target and triggered the tone (confirmer role) once the robot had moved the cursor to the target location. In all three blocks (solo, mover, and confirmer), participants made an interval estimate of the time between the confirmation press and the tone, which is a behavioral measure of SoJA (*6*).

Across two experiments, we manipulated participants' attribution of intentionality toward iCub. In experiment 1, participants completed the task with iCub without any prior manipulation of the intentionality they attributed to it. The robot was programmed to simply execute the joint task. This behavior of the robot appeared quite mechanical, and thus, we anticipated that participants would perceive the robot as a mechanical artifact. We term this condition "seemingly mechanical robot." This expectation was confirmed by a manipulation check (Supplementary Materials). In experiment 2, before the joint task, participants engaged in an interaction with

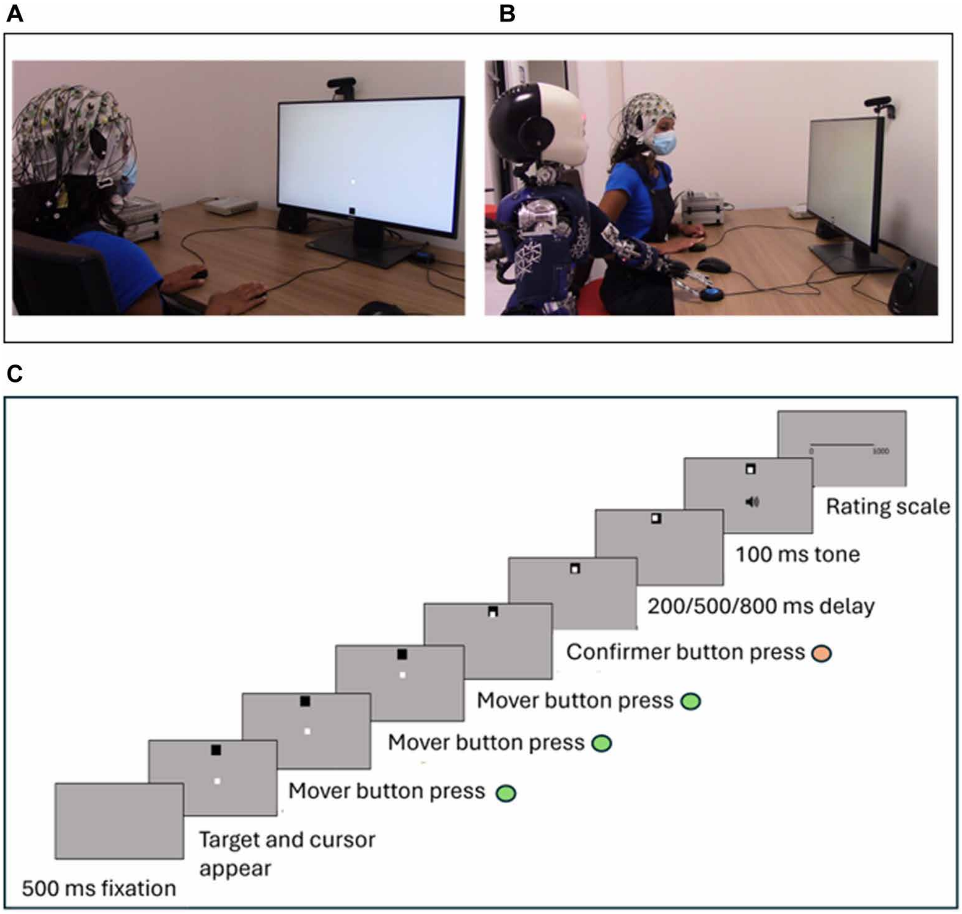

Fig. 1. Experimental setup. (**A**) Task setup in the solo block. (**B**) Task setup in the joint blocks. (**C**) Trial sequence of the target-tracking-and-confirmation task. In every trial, participants or iCub had to move the white cursor onto the black target with a series of button presses. Subsequently, the participant or iCub had to confirm the target. The confirmation press triggered a tone after a variable delay. Participants then had to rate the delay between the confirmation press and the tone on a scale provided on the screen.

iCub that has been demonstrated to increase the likelihood of attributing intentionality toward the robot (*21*). During this interaction, participants first engaged in a short dialogue with iCub, during which the robot asked the participants their names and then asked participants whether they wanted to watch some videos together. This was done through the Wizard-of-Oz technique (*22*), with the experimenter controlling iCub's verbal interaction. Participants were encouraged to respond to the robot's questions. After this, participants watched a series of three short National Geographic clips with the robot. The robot was programmed to respond to the videos at certain moments and exhibited humanlike behavior, expressing laughter, awe, or sadness depending on the video content. After the videos, iCub then spoke again and told participants they would see each other later. A detailed description of this interaction protocol is reported in (*21*). We expected that participants would perceive the robot in this condition as having some mental states [and thus having more intentionality attributed to it in the philosophical sense of the term (*21*, *23*)]. We term this condition "seemingly intentional robot." The two expectations regarding how the robot would be perceived were confirmed by a manipulation check (Supplementary Materials). A video of the experimental setup, as well as the theoretical background, hypotheses, and main results of the study, is presented in Movie 1.

Behavioral and neural markers of SoJA and hypotheses

We expected that increased attribution of intentionality to iCub would influence the emergence of SoJA. This expectation was based on findings from the literature on both individual SoA and joint action, previously summarized. Namely, individual SoA, in the form of temporal compression, emerges for voluntary and intentional actions (*8*, *9*). This suggests that SoJA too may be sensitive to the attributed intentionality with which actions are carried out, given the parallels between individual and joint SoA (*11*). Moreover, mechanisms in joint action, likely closely related to the emergence of SoJA, such as the formation of shared representations, appear to be sensitive to the intentionality attributed to one's joint action partner (*18*), further suggesting that mechanisms in joint action, such as SoJA, are sensitive to the perceived intentionality of the joint action partner. Accordingly, we hypothesized that, in experiment 1, participants would not form SoJA with the robot because their belief about the robot's intentionality was not deliberately induced. If SoJA is not formed, then there is a reduced SoA when the partner, rather than the participant themselves, produces the tone. In the current paradigm, when the participant was in the mover role, it was the robot who directly generated the tone outcome. Thus, if SoJA was not formed, this would be reflected in higher interval estimates (corresponding to a lower SoJA) for the participant's mover role relative to the solo and confirmer roles (note that the roles always refer to the role of the participant). In contrast, in experiment 2, when the likelihood of attribution of intentionality to the robot was increased, we expected to observe SoJA, meaning that the interval estimates were expected to be the same across the three roles, independent of who was producing the tone outcome, because joint SoA means that SoA is experienced regardless of who produces the action outcome. This was expected to replicate SoJA effects in human-human interaction [(*6*); see also the "Human-human control experiment" section in the Supplementary Materials for a behavioral control condition of our study].

In addition to behavioral measures, we also addressed neural correlates of SoJA. A neural marker frequently addressed in the context of SoA, and joint action, is the auditory N100, an event-related potential (ERP) in the electroencephalography (EEG) signal (*9*, *24*–*26*). The N100 is related to the auditory outcome of an action. The N100 is an ERP component of early auditory processing, sensitive to, among other things, the sensorimotor predictability of stimuli (*27*–*30*). For example, Poonian *et al.* (*25*) found greater temporal compression between actions and tones, and a reduced N100, in response

Movie 1. Theoretical background, experimental design, and summary of key findings.

to voluntarily produced tones relative to computer-generated tones. These results indicate that voluntary (self) produced tones, which are more predictable and associated with higher agency, result in smaller N100 amplitudes (*25*). However, previous work has also shown that the N100 may be smaller when participants experience lower SoA. Caspar *et al.* (*9*) found that coerced actions that produced a tone led to less temporal compression and a smaller N100 as compared with freely made actions.

Furthermore, the literature is also divided on whether predictability, in general, enhances neural responses to stimuli or reduces them. Previous work has found evidence for both reduction (*27*, *29*, *31*) and enhancement (*32*, *33*) of neural responses, such as the N100, to predictable sensory outcomes. The effect of predictability on neural responses may depend on the task and context, such that responses to task-relevant or attended stimuli might be enhanced when the stimuli are more predictable (*32*, *33*). Such enhancement of neural responses might be achieved through improved tuning of sensory responses to attended stimuli through top-down control (*34*). Putative mechanisms that play a role in modulating the processing of sensory stimuli are frontal-central theta activity and theta connectivity between frontal and sensory areas (*35*, *36*), which might then affect sensory responses indexed by the N100.

Fronto-central activity in the theta band (4 to 7 Hz) is thought to be a mechanism of cognitive control in the human brain. Moreover, theta activity has been proposed as a mechanism for communication between frontal and distal sites (for example, sensory sites) to facilitate various processes. For instance, increased fronto-central–to–sensory area connectivity in the theta range may be involved in the boosting of sensory gain in response to stimuli that may be relevant or attended (*35*). Although fronto-central theta activity has been widely studied in paradigms involving conflict or error monitoring, it has also been proposed that theta activity might represent a more general processing mechanism for the organization of neural activity during stimulus or response processing (*37*) and hence was a neural measure of interest in our joint task to understand how participants processed the joint (tone) outcomes produced in the task.

Thus, in the context of our study, we addressed both the N100 and theta activity as potential markers of sensory processing of the outcomes of joint actions (namely, tones). We expected the following: In experiment 1, if SoJA was not established with the seemingly mechanical partner, then the N100 amplitude would differ for self- and other-produced tones, paralleling the expected pattern of behavioral results. In experiment 2, under the assumption that attribution of intentionality to the partner matters for SoJA, we expected a similar N100 amplitude across both joint conditions (mover and confirmer roles), again in line with the expected pattern for the behavioral results. We did not hypothesize the directionality of this difference (enhanced or reduced N100) because the literature is not consistent on this matter, as described above. Last, we expected that theta activity would show a similar pattern as the N100 because theta connectivity between frontal and sensory areas has been related to sensory processing (*35*, *36*).

RESULTS

A manipulation check confirmed that before completing the joint task with the robot, participants were more likely to perceive the robot as intentional in experiment 2 as compared with experiment 1 (see the "Perceived intentionality manipulation and manipulation check" section in the Supplementary Materials).

Interval estimate results

Analysis was done on raw interval estimates. We also analyzed judgment errors (JEs). JEs are the over- or underestimates of temporal intervals in relation to actual intervals. The analysis of JEs, reported in the Supplementary Materials (the "Analysis on the JEs" section), confirmed the findings on the raw interval estimates. In addition, we also analyzed the *z*-normalized interval estimates. This analysis, reported in the Supplementary Materials (the "Analysis of *z*-scored interval estimates" section), also confirmed our findings from the analysis on the raw interval estimates.

Experiment 1

A repeated-measures analysis of variance (ANOVA) with role (solo, mover, or confirmer) as the factor revealed a significant main effect of role on the interval estimates [$F(2, 62) = 7.44$, $P_{\text{Greenhouse-Geisser}} = 0.001$, $\eta^2 = 0.19$] (Fig. 2). We then ran two-tailed pairwise *t* tests to further investigate the effect of role. The *t* tests showed that interval estimates in the solo role [$M = 347.51$ ms, confidence interval $(\text{CI})_{95\%} = 329.36$ to 365.66 ms] differed significantly from the mover role ($M = 381.49$ ms, $\text{CI}_{95\%} = 364.57$ to 398.41 ms) [$t(31) = -2.65$, $P_{\text{Bonferroni-corrected}} = 0.02$, Cohen's $d = -0.47$]. Interval estimates in the confirmer role ($M = 340.03$ ms, $\text{CI}_{95\%} = 325.81$ to 354.25 ms) also differed from the mover role [$t(31) = -4.07$, $P_{\text{Bonferroni-corrected}} = 0.001$, Cohen's $d = -0.72$]. The interval estimates in the solo role did not differ from the confirmer role [$t(31) = 0.67$, $P_{\text{Bonferroni-corrected}} = 0.51$, Cohen's $d = 0.12$]. These results suggest that in experiment 1, participants did not experience SoJA when interacting with a seemingly mechanical partner.

Experiment 2

A repeated-measures ANOVA with role as the factor revealed a nonsignificant effect of role on the interval estimates [$F(2, 52) = 1.56$, $P_{\text{Greenhouse-Geisser}} = 0.22$, $\eta^2 = 0.06$; solo: $M = 378.10$ ms, $\text{CI}_{95\%} = 361.29$ to 394.91 ms; mover: $M = 367.22$ ms, $\text{CI}_{95\%} = 346.15$ to 388.29 ms; confirmer: $M = 365.08$ ms, $\text{CI}_{95\%} = 340.08$ to 372.08 ms] (Fig. 2). These results indicate that participants formed a SoJA with the robot when it was introduced as an intentional partner. This pattern also parallels the results of SoJA in joint action between two humans [as in (*6*), and as in a human-human control experiment of our task, reported in the "Human-human control experiment" section in the Supplementary Materials].

ERP analysis results

Analysis of the N100 was done on the average N100 amplitude across fronto-central channels Fz, FCz, and Cz [in line with (*24*)] in the time window of ±20 ms surrounding the peak amplitude (for more details, see Materials and Methods).

Experiment 1

A repeated-measures ANOVA with role as the factor revealed a significant effect of role on the average N100 amplitude [$F(2, 62) = 6.74$, $P_{\text{Greenhouse-Geisser}} = 0.007$, $\eta^2 = 0.18$] (Figs. 3 and 4). We then ran two-tailed pairwise *t* tests to further check the effect of role. The *t* tests showed that the N100 amplitude differed significantly between the solo role ($M = -5.48$ μV, $\text{CI}_{95\%} = -5.93$ to -5.03 μV) and mover role ($M = -4.24$ μV, $\text{CI}_{95\%} = -4.87$ to -3.61 μV) [$t(31) = -3.01$, $P_{\text{FDR-corrected}} = 0.02$ (FDR, false discovery rate), Cohen's $d = -0.53$] and between the mover role and confirmer role ($M = -5.03$ μV, $\text{CI}_{95\%} = -5.40$ to -4.66 μV) [$t(31) = 2.15$, $P_{\text{FDR-corrected}} = 0.04$, Cohen's $d = 0.38$]. In addition, the solo and confirmer roles differed significantly [$t(31) = -2.12$, $P_{\text{FDR-corrected}} = 0.04$, Cohen's $d = -0.37$]. These results suggest that the tones in the solo role elicited the strongest neural responses, followed by the confirmer role, whereas the tones in the mover role evoked the weakest responses.

Experiment 2

A repeated-measures ANOVA with role as the factor revealed a significant effect of role on the N100 amplitude [$F(2, 52) = 22.44$, $P_{\text{Greenhouse-Geisser}} < 0.0001$, $\eta^2 = 0.46$] (Figs. 3 and 4). We ran two-tailed pairwise *t* tests to further check the effect of role. The *t* tests showed that the N100 amplitude in the solo role ($M = -7.44$ μV, $\text{CI}_{95\%} = -7.92$ to -6.96 μV) differed significantly from the confirmer role ($M = -5.79$ μV, $\text{CI}_{95\%} = -6.17$ to -5.41 μV) [$t(26) = -6.12$, $P_{\text{FDR-corrected}} < 0.0001$, Cohen's $d = -1.18$] and the mover role ($M = -5.42$ μV, $\text{CI}_{95\%} = -5.94$ to -4.88 μV) [$t(26) = -5.44$, $P_{\text{FDR-corrected}} < 0.0001$, Cohen's $d = -1.05$]. The N100 amplitude did not differ significantly between the confirmer and mover roles [$t(26) = -1.18$, $P_{\text{FDR-corrected}} = 0.25$, Cohen's $d = -0.23$]. These results indicate that although the tones in the solo role elicited the strongest neural responses, the tones in the two joint roles elicited similar responses regardless of who made the button press to trigger the tone, paralleling the behavioral results that indicated emergence of SoJA.

Spectral analysis results

The results of the event-related spectral power changes are reported in the Supplementary Materials (the "Event-related spectral perturbation analysis for experiment 1 and experiment 2" section). For the connectivity analysis, reported below, we applied a current source density (CSD) transform to the data to remove contributions of deep sources (*36*) before we extracted weighted phase lag index (wPLI) values averaged across trials, between the fronto-central (FCz/FC1/FC2) and temporal (T7/T8/TP7/TP8) electrodes, from epochs locked to the tone onset. We subsequently calculated the difference in connectivity between the tone stimulus period and a prestimulus period, the wPLI difference (for more details, see Materials and Methods).

Experiment 1

A repeated-measures ANOVA revealed a significant effect of role on frontotemporal wPLI difference [$F(2, 62) = 5.54$, $P_{\text{Greenhouse-Geisser}} = 0.006$, $\eta^2 = 0.15$] (Fig. 5). Two-tailed pairwise *t* tests showed that the wPLI difference in the mover role ($M = 0.02$, $\text{CI}_{95\%} = 0.0$ to 0.04) was significantly different from those of the solo role ($M = 0.06$, $\text{CI}_{95\%} = 0.04$ to 0.08) [$t(31) = -2.98$, $P_{\text{FDR-corrected}} = 0.02$, Cohen's $d = -0.53$] and the confirmer role ($M = 0.05$, $\text{CI}_{95\%} = 0.03$ to 0.07) [$t(31) = 2.68$, $P_{\text{FDR-corrected}} = 0.02$, Cohen's $d = -0.47$]. There was no significant difference between solo and confirmer roles [$t(31) = 0.44$, $P_{\text{FDR-corrected}} = 0.67$, Cohen's $d = 0.08$]. These results parallel the ERP analyses. They show that the roles in which the largest N100 amplitude was elicited, namely, solo and confirmer, were also the roles in which connectivity increased the most.

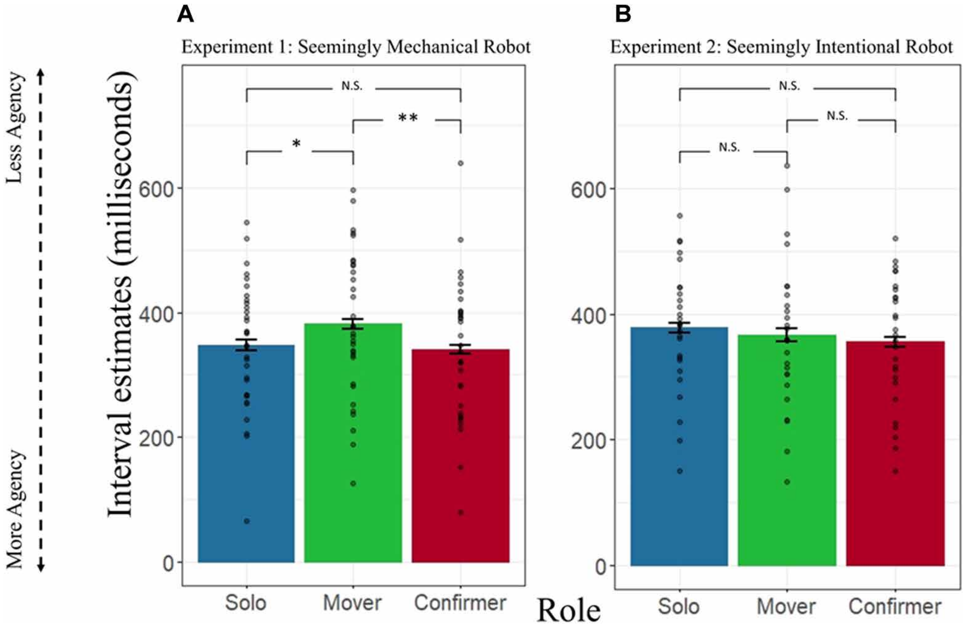

Fig. 2. Interval estimate results. Interval estimates, used as a behavioral measure of SoJA, across participants' roles in (**A**) experiment 1 with the seemingly mechanical robot ($n = 32$) and in (**B**) experiment 2 with the seemingly intentional robot ($n = 27$). Bars represent the mean interval estimates; error bars represent the mean ± SE. All pairwise *t* tests were two-tailed and corrected for multiple comparisons with the Bonferroni correction. * represents *P* between 0.05 and 0.01. ** represents *P* between 0.01 and 0.0001. N.S. represents a nonsignificant difference.

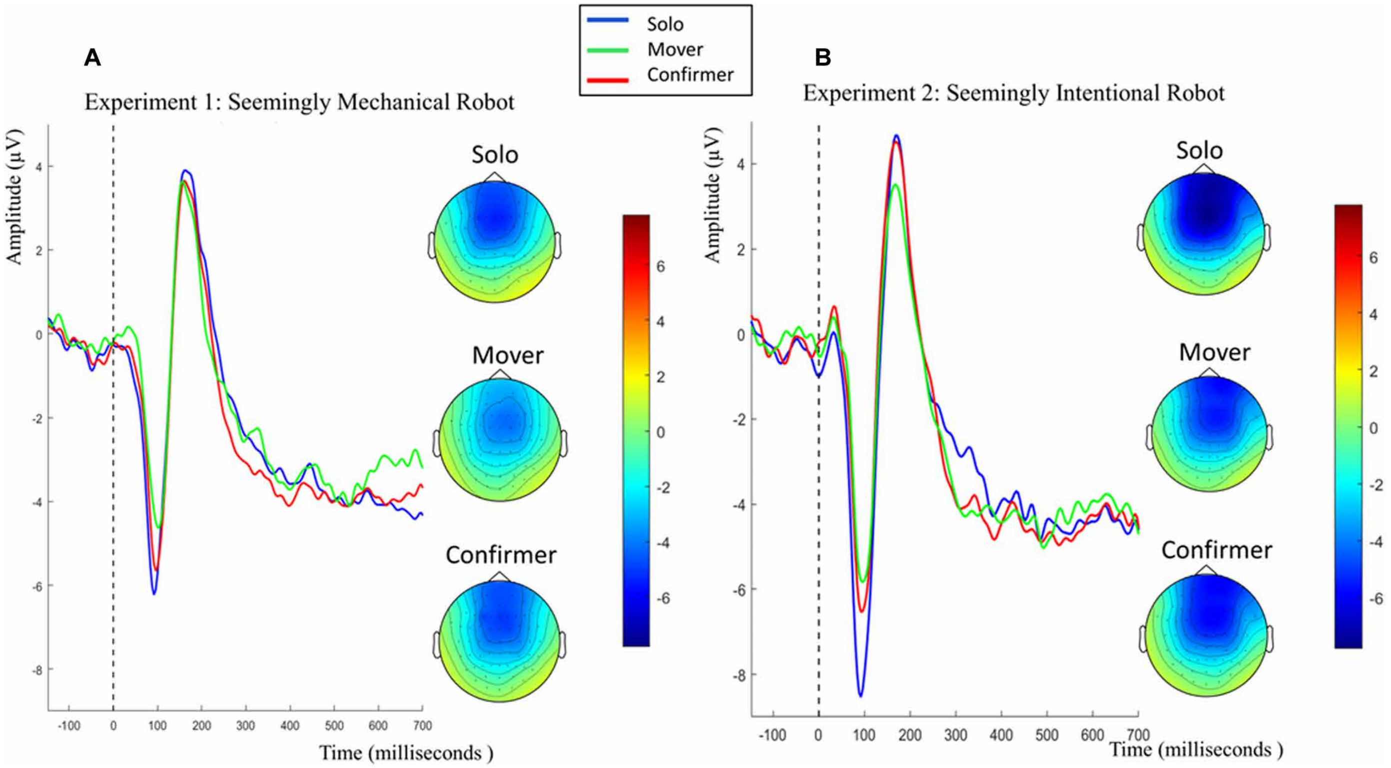

Fig. 3. N100 waveforms. N100 amplitude across electrodes Fz, FCz, and Cz and scalp topographies for the three roles in (**A**) experiment 1, with the seemingly mechanical robot ($n = 32$), and in (**B**) experiment 2, with the seemingly intentional robot ($n = 27$).

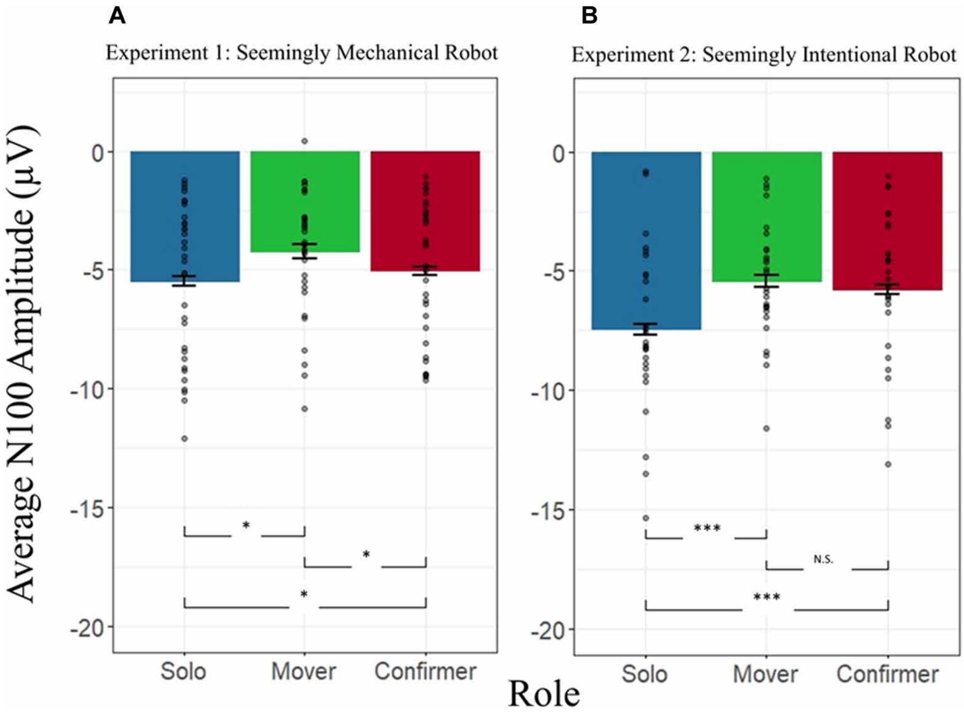

Fig. 4. N100 amplitude results. Average N100 amplitude across Fz, FCz, and Cz in (**A**) experiment 1 with the seemingly mechanical robot ($n = 32$) and in (**B**) experiment 2 with the seemingly intentional robot ($n = 27$). Bars represent the mean N100 amplitude; error bars represent the mean ± SE. All pairwise t tests were two-tailed and corrected for multiple comparisons using the FDR correction. * represents P between 0.05 and 0.01. *** represents P above 0.0001. N.S. represents a nonsignificant difference.

Experiment 2

A repeated-measures ANOVA for the effect of role revealed a significant effect of role on frontotemporal wPLI difference [$F(2, 52) = 5.78$, $P_{\text{Greenhouse-Geisser}} = 0.005$, $\eta^2 = 0.18$] (Fig. 5). Two-tailed pairwise t tests showed that the wPLI difference in the solo role ($M = 0.1$, $\text{CI}_{95\%} = 0.08$ to 0.12) was significantly different from those of the confirmer role ($M = 0.06$, $\text{CI}_{95\%} = 0.04$ to 0.08) [$t(26) = 2.97$, $P_{\text{FDR-corrected}} = 0.02$, Cohen's $d = 0.57$] and the mover role ($M = 0.06$, $\text{CI}_{95\%} = 0.04$ to 0.08) [$t(26) = 2.66$, $P_{\text{FDR-corrected}} = 0.02$, Cohen's $d = 0.51$]. There was no significant difference between confirmer and mover roles [$t(26) = 0.15$, $P_{\text{FDR-corrected}} = 0.89$, Cohen's $d = 0.03$]. These results once again parallel the results of the analysis on the N100 amplitude.

DISCUSSION

In the present study, we assessed—with behavioral and neural measures—whether SoJA is formed in joint action tasks with a robot, whether this is influenced by the perception of the robot as an intentional agent, and what the underlying neural mechanisms are. Participants played a target-tracking-and-confirmation task alone (solo role) and together with the humanoid robot iCub (in mover and confirmer roles). In experiment 1, participants played the task with the robot without any prior manipulation of their belief in iCub's intentionality. In experiment 2, we increased the likelihood of participants attributing intentionality to the robot. In all experiments, we collected interval estimates as a behavioral measure of SoJA and EEG signal to evaluate neural responses to action-generated sensory outcomes produced by the participants themselves (solo and confirmer roles) or by the robot (when participants were in the mover role). [Note that we also conducted a behavioral control experiment (Supplementary Materials) where participants performed the task with another human.]

We expected that if the attribution of intentionality [in the philosophical sense, namely, attribution of mental states (*23*)] toward the robot affects the emergence of SoJA, SoJA should not emerge with the robot in experiment 1, meaning that the interval estimates would be lower (indicating higher SoA) in the solo and confirmer roles (where the participant was directly producing the action outcome, namely, the tone) relative to the mover role (where the tone was produced by the robot). Along the same lines, if attribution of intentionality to the robot is a crucial factor for SoJA, then SoJA was predicted to emerge in experiment 2, meaning that the interval estimates would be the same across the three roles because a SoJA means that a sense of control is experienced regardless of who produces the action outcome in the joint action.

In experiment 1, in solo and confirmer roles, participants estimated the delays between their own actions and the subsequent outcomes as shorter than delays between the robot's actions and outcomes (mover role), indicating that they did not form a SoJA with the seemingly mechanical robot. Conversely, in experiment 2, participants' estimates of the delays were identical for their own and the robot's actions and outcomes, indicating that they experienced a SoJA with the seemingly intentional robot. This pattern of results paralleled the human-human control experiment (Supplementary Materials). Together, our behavioral results suggest that even if a robot has a human-like shape and sensorimotor repertoire, meaning that its actions can likely be represented by one's own sensorimotor mechanisms, it may not be enough to form a SoJA during HRI. It appears that attribution of intentionality to the partner is crucial. This is in line with evidence related to individual SoA: It is only in the case of voluntary/intentional actions (as opposed to passive or involuntary movements) that individual SoA emerges in the form of temporal compression (*8*, *38*).

Regarding neural correlates of action-generated outcomes, we focused first on the auditory N100. In experiment 1, we hypothesized that if participants did not form a SoJA with the robot, the N100 amplitude would differ for self-generated tones as compared with other-generated tones, mirroring the behavioral results. In experiment 2, we hypothesized that the N100 amplitude would be comparable across self-generated and other-generated tones in the joint roles, also mirroring the behavioral results.

Results showed that in experiment 1, the N100 amplitude was largest in the solo role, followed by the confirmer role, and was smallest in the mover role. Moreover, the N100 amplitudes in the two joint conditions differed significantly from each other, paralleling the behavioral results from these two roles. Conversely, in experiment 2, the N100 amplitude was largest in the solo role but did not differ significantly between the joint conditions, also in line with the behavioral data from the joint roles. Literature shows that the N100 is modulated by factors such as attention (*30*, *39*), as well as by the predictability of a sensory stimulus (*27*, *28*, *40*). As such, the N100 is related to perception and processing of stimuli (in the case of our study, the perception and processing of an auditory stimulus). Thus, a larger N100 amplitude might reflect enhanced sensory processing of stimuli as a consequence of a top-down modulatory activity (*34*). In the case of our study, this top-down modulation presumably had its source in more accurate sensorimotor predictions, which were more accurate for stimuli under one's own control relative to the predictions made for the joint action partner. Because, in the context of self-action, one has better access to one's own sensorimotor representations and more accurate internal models, sensorimotor prediction in the case of one's own actions might be more accurate than in the case of observed actions of others or joint actions (*10*). Thus, we posit that in our study, larger N100

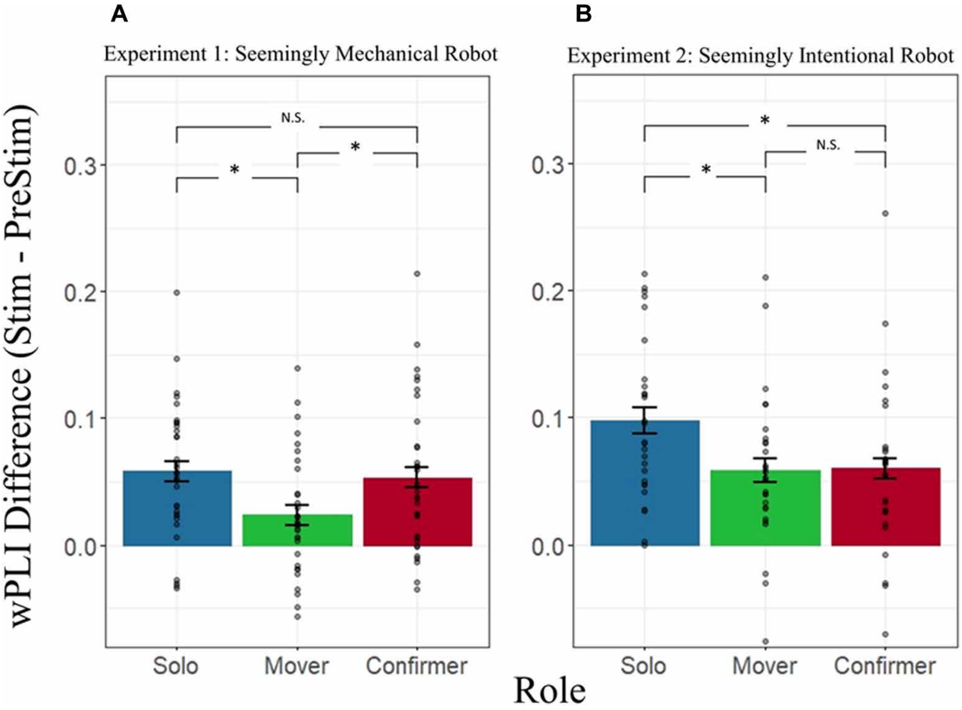

Fig. 5. Theta connectivity results. Stimulus period minus prestimulus period difference in the theta band (4 to 7 Hz) wPLI between fronto-central (FCz, FC1, and FC2) and temporal (T7, T8, TP7, and TP8) electrodes, across participants' roles (**A**) in experiment 1 with the seemingly mechanical robot ($n = 32$) and (**B**) in experiment 2 with the seemingly intentional robot ($n = 27$). Bars represent the mean wPLI difference; error bars represent the mean ± SE. All pairwise tests were two-tailed and corrected for multiple comparisons using the FDR correction. * represents P between 0.05 and 0.01. N.S. represents a nonsignificant difference.

amplitudes reflect enhanced sensory processing, which results from more accurate sensorimotor prediction produced by the forward model.

In the case of our study, in experiment 1, participants formed the best or most accurate sensorimotor prediction of the tone outcome in the solo condition because the tone was fully self-generated, meaning that, in this condition, participants were able to rely fully on their well-trained own (or "self") forward model. This prediction fed back to the areas of the brain where sensory stimuli are processed. Such feedback-based modulatory activity enhanced sensory processing of the tone to the greatest extent, as indexed by the largest N100 seen in the solo condition. In experiment 2, the most accurate sensorimotor prediction for the outcome, and thus the most enhanced processing of the tone, was again seen for the solo condition (again in which participants relied solely on their well-trained self forward model to make sensorimotor predictions), replicating the results of experiment 1 in this regard. The idea that sensorimotor predictions for individual actions are more accurate than for actions performed with others is in line with literature on motor control, showing that people are better at predicting their own actions as compared with those of others and are better at coordinating with themselves as compared with another person (*41*–*43*). Presumably, this is because self internal models are better trained (and thus more accurate) than internal models for partner-generated actions (hereafter denoted as "other" forward models) or joint models, making it easier to predict and coordinate with oneself than with a partner (*10*, *43*) [these sensorimotor prediction mechanisms most likely also underpinned participants' individual SoA in the solo role (*8*, *11*)]. In the case of the two joint action conditions, there was a difference in N100 amplitude between the joint roles in the experiment with the seemingly mechanical robot in experiment 1. More specifically, in experiment 1, the N100 amplitude was larger in the confirmer role relative to the mover role. This shows that the way the tone stimulus was processed at the sensory level was different (smaller amplitude) for the case in which the robot produced the outcome compared with when the participant produced the outcome directly. In line with our interpretation of the N100 effect, the predictive processes—being less accurate when observing the partner's actions relative to participants' own actions—lead to less enhanced sensory processing in the condition in which the participant did not produce the action outcome directly as compared with when they did. This aligns with the behavioral data indicating that the SoJA was not formed and points to an underlying neural mechanism that is a modulatory activity of sensory processing. The N100 amplitude in the confirmer condition, when participants produced the tone, was smaller than the solo condition. Presumably this is because in the solo condition, one relies solely on the self internal model, whereas in the confirmer role, although one relies largely on the self model, one still has to coordinate with one's partner, resulting in some reliance on the other model [as explained below through the predictive joint action model (PJAM) framework].

In experiment 2, although the N100 was largest in the solo condition, the N100 amplitude was equivalent in the two joint conditions, indicating that sensorimotor predictions (and thus processing) might have been equally accurate in the joint roles (and less accurate than the solo condition) in experiment 2. Presumably this is because the joint model was activated (as explained through the PJAM theoretical framework below), leading to similar sensory processing of the outcomes, and thus similar N100 amplitudes, in both joint roles.

Our data provide further evidence that the enhanced N100 reflects enhancement of sensory processing, resulting from top-down (feedback) modulatory activity having its source in sensorimotor-related predictions. Specifically, the results from our spectral connectivity analyses, which showed that the theta connectivity activity mirrors the pattern of results obtained from our analysis of the N100 component, support this interpretation. In experiment 1, the theta connectivity difference between the stimulus period and a prestimulus period was higher for the solo and confirmer roles compared with the mover role. In experiment 2, the theta connectivity difference was higher for the solo role relative to the mover and confirmer roles, which did not differ. Thus, in both experiments, the conditions in which we found larger N100 amplitudes corresponded to the conditions in which we observed a greater increase in fronto-central to temporal theta connectivity in response to the tone. Previous work suggests that long-range synchronized theta oscillations between fronto-central and sensory brain areas may play a role in the amplification of neural responses to task-relevant and attended sensory stimuli (*35*, *36*), perhaps in particular for more predictable stimuli (*33*, *35*). Thus, in our study, neural responses to tones that were more accurately predicted by the relevant forward model (for example, by the self-forward model in the solo role) might have been enhanced through top-down modulatory activity and indexed by increased theta synchrony between fronto-central and temporal areas, leading to larger N100 amplitudes for these stimuli. In addition, our analysis of the event-related power changes in the theta band also mirrors our N100 and connectivity results (see the "Event-related spectral perturbation analysis for experiment 1 and experiment 2" section in the Supplementary Materials).

Following this interpretation, the results from the joint blocks of our study can be interpreted in the framework of the PJAM (*10*). The PJAM is a hierarchical model of joint action with three levels. According to the PJAM, agents performing a joint action maintain internal models for both self and for actions generated by their partner, that is, other models. These internal models are integrated into a joint model by a joint goal representational level at the "top" and by incoming sensory information generated by oneself and by one's partner at the "bottom" (see Fig. 6 for a schematic representation and a more detailed explanation). In both experiments in our study, participants received sensory inputs regarding the consequences of their and their partner's actions in the joint conditions. These sensory inputs were fed into the self and other models, linking the models and integrating them both into the joint model to a certain extent. However, in experiment 1, participants completed the task with a seemingly mechanical robot, and thus, the joint goal representation level of the PJAM might not have been activated. Therefore, participants' self and other models were possibly linked only via the incoming sensory input but not by the joint goal level. Thus, participants might have maintained a distinction between self and other models. When in the confirmer role, participants were still able to rely largely on their own, more-accurate self model, whereas in the mover role, participants had to rely on the less accurate and less well-trained other model to a greater extent (Fig. 6). Because participants were relying more on their better trained self-model in the confirmer role, this allowed them to make more accurate sensorimotor predictions about the outcomes of their actions (namely, the tones) in the confirmer role as compared with the mover role. Through top-down modulatory activity, processing of these (more accurately) predicted stimuli may then have been enhanced. This

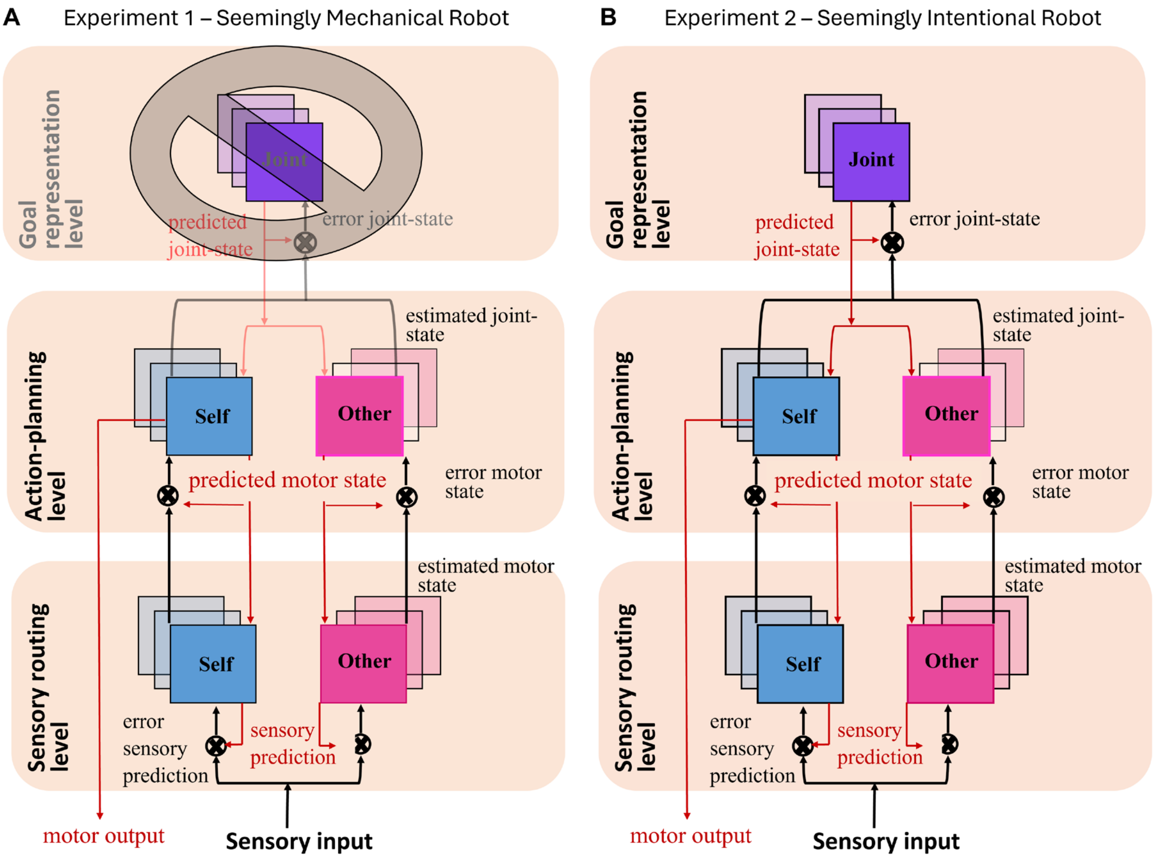

Fig. 6. Schematic representation of the PJAM, adapted from (*10*). The PJAM is a hierarchical model with three levels. According to the model, both self- and other-internal models are maintained during joint action. These models are organized into a hierarchical framework with three levels. At the top is the joint goal representation level, which consists of a set of modules that symbolically represent the agents' shared goals. On this level, partners are assumed to be mutually responsive to each other's actions, goals, and intentions. In the middle is the action planning level, which consists of both self- and other-forward models integrated into a dyadic motor plan that takes into account both agents' actions and their contribution toward the shared goal (*10*, *13*, *15*). Predictions about expected self and other motor states are made at this level and iteratively compared with actual motor states arising from the lowest level. The lowest level is the sensory routing level. On this level, information on actual motor states is processed, and predicted sensory input is compared with actual sensory input arising from the environment, both for self and other actions and sensory outcomes. Each of the three levels has connections to the levels above and below. Thus, according to the PJAM, the flow of information arising from both the shared goal level and the incoming sensory inputs inherently interlink self and other predictive models (*10*) into a joint model. The diagram on the left represents the situation that occurs in experiment 1, with the seemingly mechanical robot. In this case, the (joint) goal representation level is not activated. Thus, the self- and other-forward models are linked only by the incoming sensory input generated by each partner (human participant and robot) and, as a consequence, are only marginally linked together into the joint model. The diagram on the right represents the situation that occurs in our experiment 2, with the seemingly intentional robot [this diagram also represents scenarios of joint actions with human partners, initially proposed in (*10*)]. In this case, the joint goal representational level is activated. Therefore, the self and other forward models are linked both by the incoming sensory input and the flow of information from the (joint) goal representation level, resulting in self and other-forward models being better integrated into the joint model.

was indexed by a larger N100 amplitude in the confirmer role than in the mover role in experiment 1. Moreover, as a consequence of this lack of self-other integration, participants did not experience SoJA with the seemingly mechanical robot.

On the contrary, in experiment 2, because the robot was introduced as an intentional agent, the joint goal representation level was presumably activated. Thus, in experiment 2, both the incoming sensory inputs and the activated joint goal representation level linked participants' self and other models within the joint model. In consequence, participants' self and other models were presumably both integrated into the joint model to a greater extent in experiment 2 (as compared with experiment 1) (Fig. 6). Therefore, in experiment 2, sensorimotor predictions for the action outcomes (tones) in the confirmer role and mover role were both made by a joint model and thus were equally accurate across the two roles (and also less accurate than in the solo role). Hence, the modulatory activity over the sensory areas, and accordingly, the sensory processing of the tones, may have been similar in both conditions, as indicated by similar N100 amplitudes in both joint roles. Moreover, as a consequence of this self-other integration, participants experienced SoJA with the seemingly intentional robot. This suggests that in joint action, sensorimotor predictive mechanisms and implicit SoJA are sensitive to higher-level contextual cues, such as the intentionality with which actions are made or perceived. Moreover, it seems that for a SoJA to emerge in HRI, sensorimotor predictions, made by a joint forward model, do not necessarily have to be as "good" as in individual action but have to be equally accurate for self- and other-generated actions and outcomes.

Together, our results provide several theoretical contributions to the topic of SoJA. First, our results show that SoJA, measured using interval estimates, is sensitive to the intentionality attributed to the partner during a cooperative joint task. Second, our EEG results indicate mechanisms that are at stake in the emergence of SoJA: In joint action contexts, only when we attribute intentionality to our robotic partner, we form, or activate, a shared representation of the joint action goal. This leads to greater integration of self- and other-forward models into a joint model. This joint model consequently generates similarly accurate sensorimotor predictions for both self- and other-generated actions and outcomes. These predictive processes feed back to lower-level areas of the brain where sensory processing occurs and modulates processing of incoming stimuli, resulting in similar levels of activity in sensory areas of the brain (and similar processing of self- and other-produced sensory outcomes) and as a consequence the emergence of SoJA. For fully self-generated actions, that is, in individual action, our brains form more accurate sensorimotor predictions of the action outcomes because they are solely able to rely on our well-trained self forward model to make these sensorimotor predictions. As a result, more enhanced neural responses may occur for sensory stimuli that are outcomes of self-generated actions.

A potential limitation of the present study, and a topic for future research directions, is that although the robot was programmed to respond with a response profile similar to that of a human (see Materials and Methods), it was not identical to a human, which may have affected how self and other models were integrated into the joint model. A recently proposed model of joint agency suggests that in human-human joint action, implicit SoJA is exclusively shaped by sensorimotor cues (*11*). Thus, future studies might explore how SoJA emerges when interacting with a humanoid robot that is not perceived as intentional but does have a response profile identical to a human response profile, for example, through being remotely controlled by another human. Doing so would allow us to further unravel the mechanisms contributing to the emergence of SoJA.

In addition, we note that, in our study, the intentionality manipulation in experiment 2 was achieved through an interaction involving emotional expressiveness of the robot. Although we implemented this "interactive approach" because previous work has shown that it can effectively increase attribution of intentionality to the iCub robot (*21*), in the future, it would be interesting to investigate whether manipulation of intentionality attribution without interaction that involves emotional expressiveness of a robot would also be sufficient to influence SoJA. For example, future studies could investigate whether intentionality attribution can be sufficiently manipulated through verbal instruction and whether this also can influence the emergence of SoJA in HRI. Doing so may provide previously unidentified and perhaps even more efficient ways to induce the emergence of SoJA in instances of human-robot collaborations that could benefit from this.

Another avenue for future research could be to explore how the SoJA experienced during HRI is influenced by the type of task implemented and on the social context. Previous studies indicate that in certain HRI tasks, for example, when performing an action that is costly in terms of losing points, people experience a diffusion of responsibility when acting together with a robotic agent. That is, people exhibit a lower SoA over their own actions when acting together with a robot as compared with when acting alone (*44*). Moreover, when acting in a hierarchical social context, people may also shift the responsibility of their actions onto another, for example, onto someone giving them instructions, a displacement of responsibility (*45*). Similarly, the SoJA experienced may also be differently modulated in a hierarchal task, in which a human commands a robot to make an action or vice versa, or in a joint task in which actions may lead to negative outcomes, such as a loss of points. Therefore, future work could also investigate how the experience of SoJA, displacement, and diffusion of responsibility are affected by different task structures in HRI and different "roles" assigned to the robot.

In conclusion, the present study shows that SoJA with robots depends on attribution of intentionality to the latter. Attribution of intentionality has been suggested in previous literature to increase predictability and trust in HRIs, as well as the sense of control experienced (*17*). The current results further suggest that attribution of intentionality to the robot appears to be crucial for the emergence of SoJA in HRI. Furthermore, our results have elucidated the neural mechanisms underlying the emergence of SoJA in collaborative joint action, namely, enhancement of sensory processing of outcomes of actions through modulatory activity stemming from sensorimotor predictions.

Apart from the theoretical implications of our results, they also have a potential value for robotics. When performing collaborative tasks with robots, it may be important for certain tasks that the human-robot joint action is as smooth and effortless as possible, similar to how the task would be performed with another human. Our brain has developed several mechanisms that facilitate the smoothness and efficiency of joint actions with other humans, and one of those mechanisms is SoJA. In certain instances, such mechanisms may also facilitate human-robot joint actions, and therefore, when designing robots for human-robot collaboration, one needs to understand the conditions under which SoJA emerges, as well as the contexts and tasks in which SoJA can further facilitate human-robot collaborations. Conversely, in some other tasks, it might not be beneficial to form SoJA with a robot. For example, in cases of hierarchical relationship it might be crucial that the human retains their own individual SoA. In that case, it is better to not induce the impression of intentionality of the robot. Overall, understanding the factors that contribute to emergence of SoJA in HRI can provide practical insights to robotics and HRI, as well as further theoretical understanding of fundamental processes involved in human cognition.

MATERIALS AND METHODS

Participants

For experiment 1, sample size was estimated via an a priori power analysis based on effect sizes from similar behavioral (*7*) and neurophysiological (*24*) work. We estimated that a sample size of $n = 30$ would have sufficient power to reveal any effects [estimated behavioral effect size (η^2) ~ 0.44, estimated neurophysiological effect size (η^2) ~ 0.15, alpha error = 0.05, power = 0.95]. Forty-four right-handed adults took part in experiment 1 (mean age of 25.1 years, range = 20 to 44 years, 30 females).

For experiment 2, sample size was estimated via an a priori power analysis based on the effect sizes obtained in experiment 1 [estimated behavioral effect size (η^2) ~ 0.19, estimated N100 effect size (η^2) ~ 0.18, alpha error = 0.05, power = 0.95]. On the basis of the power analysis, we estimated that a sample size of $n = 25$ would have sufficient power to reveal any effects. Thirty right-handed adults took part in experiment 2 (mean age of 25.0 years, range = 20 to 44 years, 17 females).

All participants had normal or corrected-to-normal vision. The study was conducted according to the ethical standards specified in the 2013 Declaration of Helsinki and was approved by the local ethical committee (Comitato Etico Regione Liguria). All participants gave written consent before participating in the experiment. Participants received monetary compensation of 14 euros/hour for their participation.

Apparatus and stimuli

The experimental task was presented on a 27-inch screen, with a refresh rate of 140 Hz, placed in front of the participant and iCub. Participants and iCub were placed side by side (Fig. 1). To collect responses from both the participant and the robot, we used a custom response box, based on a MicroPython controller, with buttons provided in the Logitech Gaming Kit. The experiment was implemented in PsychoPy v2021.2.0 (*46*). The robot was controlled using Yet Another Robot Platform (YARP) (*47*) and its Python wrappers, which allowed integration with PsychoPy. During the task, tones were played via standard external speakers placed to the right and left of the screen.

During the joint blocks, the robot was programmed to execute responses by pressing the Logitech button placed in front of it. This behavior was achieved by controlling iCub's arm with the YARP IPositionControl controller (*47*), which ensured control of the joints in position with good repeatability of the movements and trajectories that follow a minimum jerk profile. We based the timings of the robot's button presses on reaction time (RT) values collected during a similar human-human joint action pilot of the task. On the basis of these data, we constructed discrete uniform distributions of RTs for each button press, with 30 values per distribution. The human-human pilot revealed that each of the "mover" button presses (from 1 to 7) and the "confirmation" press had different RTs. Therefore, we constructed separate distributions for each mover press (from 1 to 7) and the final confirmation press. For each key press, each trial in a given block was assigned a value selected from the respective distribution. Because we had 90 trials per block, the values for each button press were repeated three times per block. The trial order was then shuffled, thereby shuffling the values in the distribution, to randomize their presentation. The values were independently assigned to trials for each of the button presses. (This was especially relevant for the confirmer role of the participant, in which the robot played as the mover. In these trials, the RTs for each key press were independently assigned to each trial. Therefore, although a specific discrete value was repeated three times over the entire block, the exact combination of mover RTs was not necessarily repeated.) The mean and SDs for the distributions of each button press for both experiments are provided in Table 1.

The distributions from which the timing of the robot's button presses were taken during the target tracking-and-confirmation task remained identical across the two experiments, and no extra behaviors were added to the robot during trials. Between the two experiments, the robot underwent hardware maintenance. This included adjustments to the robot's arm joints, which subsequently necessitated adjusting the parameters of the starting and target positions (pressing the button being the "target") of the robot's arm, between experiments. However, these adjustments did not significantly affect the joint task, as detailed in the dedicated section in the Supplementary Materials [the "Comparison of reaction times (RTs) of the iCub across experiments" section].

In experiment 2, to maintain the appearance of the robot as an intentional agent during the task, we added two behaviors. The first was that after trials in the joint task, the robot was programmed to look toward the participant, with an independent probability of 70% on each trial. The range of motion of the robot's head movements was drawn from discrete uniform distributions (elevation: −5° to +5°; azimuth: −5° to −50°) (*48*). The second was that the robot was programmed to look toward the participants and speak to them before starting the joint task blocks and between the blocks. The robot always repeated the same phrase at each time point. (Before starting the joint task blocks, the robot was programmed to say, "I can't wait to play with you. Let's start when you want!" Between the blocks, it was programmed to say, "This game is fun! Let's keep going." All sentences were presented in Italian.) The experimenter controlled this behavior via a button press. Questionnaires were presented on a 27-inch screen, in PsychoPy v2020.1.3 (*46*).

Procedure and experimental task

In experiment 1, participants first completed the Intentional Stance Test (IST-PRE) (*49*) to assess the degree to which they attributed intentionality to the robot. Then, they performed the entire target-tracking-and-confirmation task, and lastly, they completed the second half of the Intentional Stance Test (IST-POST). In experiment 2, participants first completed the target-tracking and confirmation task in solo action, then they completed the IST-PRE, and lastly, they engaged in an interaction with iCub during which the robot's behaviors were programmed to appear humanlike (intentionality attribution manipulation). The purpose of this interaction was to increase participants' likelihood of perceiving iCub as an intentional agent (*21*). After the interaction with iCub, participants completed the IST-POST. Last, participants performed the joint blocks of the target-tracking-and-confirmation task. A more detailed description of the IST, the interaction protocol, and iCub's behaviors can be found in the Supplementary Materials (the "Perceived intentionality manipulation and manipulation check" section) and in the paper by Marchesi *et al.* (*21*).

The target-tracking-and-confirmation task was adapted from (*6*) to make it suitable to play with the robot [the changes we made are described in the Supplementary Materials (the "Human-human control experiment" section)]. In each trial, a cursor and a target were presented on a central screen. The starting position of the cursor was always in the center of the screen, whereas the target was presented randomly above or below the cursor at a fixed position. The spatial units for the size and location of the cursor and target were set to "height" in PsychoPy. The size of the cursor was set to [width, height] = [0.03, 0.03], and the starting position was set to $[x, y] = [0, 0]$. The size of the target was set to [width, height] = [0.06, 0.06], and the position was set to $[x, y] = [0, \pm 0.46]$ (depending on whether the target was presented above or below the cursor). Participants were seated approximately 70 cm from the screen. In the solo role, they sat straight in front of the screen, whereas in the joint roles, they sat at an approximately 45° angle, with the robot positioned at a similar angle (see Fig. 1, A and B). Participants were instructed to make a series of button presses to move the cursor to the target location. In each trial, the cursor only moved in one direction, toward the target (1 degree of freedom). The number of button presses needed to reach the target location was randomly determined between three and seven. Once the cursor was in the center of the target, participants were instructed to press another button to confirm that the cursor had reached the target. The confirmation button press triggered a tone of 440 Hz and 100-ms duration. The delay between the confirmation

Table 1. Mean values and SD for the distributions of the robot's button press timings (RTs). A separate distribution was constructed for each button press on the basis of RTs taken from a human-human pilot of the task. When acting as the mover [i.e., the participants' role was confirmer (confirmer condition)], the robot made between three and seven button presses, whereas when acting as the confirmer [i.e., the participants' role was mover (mover condition)], the robot made only the confirmation press.

Button press	Mean RT (ms)	SD RT (ms)
Mover press 1	454.7	248.7
Mover press 2	174.4	56.3
Mover press 3	139.5	35.4
Mover press 4	167.6	65.6
Mover press 5	212.3	75.5
Mover press 6	225.9	68.9
Mover press 7	323.3	126.3
Confirmation press	490.7	178.2

press and the tone was fixed to be 200, 500, or 800 ms (*6*, *9*, *24*). The delays were pseudorandomized such that each delay was presented an equal number of times within each block. Participants were then asked to estimate the time interval between the confirmation press and the tone, on a rating scale provided on the screen. The scale had a range of 0 to 1000 ms, with a mark every 100 ms (see Fig. 1C). Participants performed the task alone by responding bimanually (solo role), or they shared the task with iCub. When the task was shared, participants performed either the mover or the confirmer role, and the robot completed the complementary role (joint roles). The roles were completed in separate blocks. Participants always completed the solo role first, followed by the joint roles. The presentation order of the two joint roles was counterbalanced across participants. In the solo role, participants moved the cursor with their left hand and confirmed the target with their right hand and then made the delay rating. Participants first completed a short practice round in the solo role to get acquainted with the trial sequence, after which they completed the solo block. In the two subsequent joint roles, when assigned to the mover role, participants moved the cursor onto the target location with their right hand, while iCub made the confirmation press to trigger the tone with its left hand. Conversely, when participants were assigned the confirmer role, iCub moved the cursor onto the target location with its left hand, while the participant made the confirmation press with their right hand.

In the joint roles, participants sat next to iCub, and both faced the screen. Both the participant and iCub were assigned their unique buttons (Fig. 1B).

After the three task blocks, participants completed two baseline action blocks with the robot in the two joint roles. In these blocks, the confirmation press did not trigger a tone, and participants did not make a delay rating. The baseline action blocks were used to correct the task block ERPs for any overlapping components, as described in the Supplementary Materials (the "Overlapping component correction for the tone locked data from experiment 1 and experiment 2" section). In total, participants completed 330 trials, 90 trials for each of the task blocks and 30 trials for each of the baseline action blocks.

General data processing and analysis

In experiment 1, we removed the data of 12 participants in total from the main analysis. Three participants were removed because of poor performance on the delay rating task (analysis showed that they were unable to distinguish the three different delays). Two participants were removed because of technical issues with the recording of the behavioral data (the data were not saved properly). Four participants were removed because of technical issues with the EEG equipment during the task (we were unable to reduce the impedances of the electrodes below 10 kohm because of issues with either the ground or the online reference electrodes). Three participants were removed because of poor EEG data quality (the data of one participant had a lot of low frequency drifts in the data; the data of the other two participants contained several epochs with large ocular and muscle artifacts), and we were unable to clean the noise from the data with our preprocessing pipeline [the independent components analysis (ICA) decompositions resulted in more than 75% of nonbrain-related ICs that had to be rejected].

In experiment 2, we removed the data of three participants in total from the final analysis: One participant was removed because they completed the joint task with the wrong hand, and two participants were removed because of poor EEG data quality (the ICA decompositions resulted in more than 75% of nonbrain-related ICs that had to be rejected). Therefore, data from 32 participants were further analyzed in experiment 1 and data from 27 participants in experiment 2.

EEG data acquisition, processing, and analysis

EEG acquisition

In both experiments, EEG data were recorded using a 64-channel Ag-AgCl electrode system (ActiCap, Brain Products GmbH, Munich, Germany), arranged according to the international 10-20 layout. The data were referenced online to FCz. The EEG signal was amplified with a BrainAmp amplifier (Brain Products GmbH), digitized at a 1000-Hz sampling rate, and recorded. No filters were applied during EEG signal recording. Electrode impedances were kept below 10 kohm. No external electrodes were used to record eye movements because participants had to wear FFP2 masks for the duration of the entire experiment. However, horizontal eye movements were captured by displacing the FT9 and FT10 electrodes to F9 and F10 electrode positions, making them closer to the outer canthi of the eyes. Vertical eye movements were captured with the Fp1 and Fp2 electrodes on the cap.

General EEG preprocessing

EEG data were preprocessed in BrainVision Analyzer and MATLAB (R2020b), together with the EEGLAB and FieldTrip toolboxes and custom scripts in MATLAB. As a first step, data were re-referenced to the average of all electrodes in BrainVision Analyzer, and FCz was interpolated back into the data. Data were then imported into EEGLAB, down-sampled to 250 Hz, high-pass filtered with a cutoff of 0.1 Hz (−6-dB cutoff: 0.05 Hz), and low-pass filtered with a cutoff of 80 Hz (−6-dB cutoff: 90 Hz). A 50-Hz notch filter was then applied to the data to filter out line noise (−6-dB cutoffs: [49.5 50.5] Hz). We then epoched the data from the task blocks to create epochs related to the onset of the tone (task epochs). The task epochs were locked to the onset of the tone, with a length of −1200 to +3500 ms relative to the tone onset. A baseline correction of −200 to 0 ms relative to tone onset was used. We visually inspected the data, and noisy channels and epochs containing large artifacts (i.e., large muscle or ocular movements and electric artifacts) were rejected (on average, we removed 1.3 ± 1.30 channels per participant). After the visual inspection, the data were once again re-referenced to the average of all (clean) channels. We then applied an ICA on the epoched data to isolate sources of noise in the data, and we rejected ICs related to eye blinks, saccades, muscular movements, and other noise from the data. In experiment 1, on average, we removed 1.3 ± 1.30 channels per participant and 38.41 ± 4.52 ICs per participant. In experiment 2, we removed an average of 2.62 ± 1.80 channels per participant and 33.48 ± 4.13 ICs per participant. After cleaning the data with the ICA, we once again inspected the data for noise and removed epochs that still contained large artifacts. In experiment 1, on average, 85.64 ± 4.02 epochs remained in the solo role, 85.48 ± 3.16 epochs in the mover role, and 85.06 ± 3.79 epochs in the confirmer role. In experiment 2, on average, 85.56 ± 4.76 epochs remained in the solo role, 82.48 ± 5.71 epochs in the mover role, and 81.44 ± 6.07 epochs in the confirmer role. We then performed a spherical interpolation to add back rejected channels. We subsequently re-referenced the data using two different references to create two datasets. One dataset was used in the analysis of the N100 and the other in the spectral analyses on the theta band. Data used in the N100 analysis were re-referenced to the average of the mastoids (TP9 and TP10). Data used in the spectral analyses were re-referenced to the average of all electrodes to include all electrodes in the re-reference because some were removed before the ICA. The experimental epochs were then sorted by tone delay and corrected for overlapping components. The procedures we used to correct the data for overlapping components are described in detail in a dedicated section of the Supplementary Materials (the "Overlapping component correction for the tone locked data from experiment 1 and experiment 2" section).

N100 specific processing and analysis

Once the overlapping component correction was performed, we applied a low-pass filter to the data at 40 Hz before extracting the N100 amplitude for further analysis. For each participant and condition, the N100 amplitude was calculated as the average amplitude in the time window "time of peak amplitude ± 20 ms." The peak amplitude was defined for each participant and condition as the most negative amplitude between 50 and 200 ms. We applied this individualized peak detection strategy to account for any variability between participants and conditions in terms of peak amplitude timing.

Spectral data specific processing and analysis

After the overlapping component correction, we performed two analyses: an event-related spectral perturbation analysis and a phase-based connectivity analysis. Further details related to the event-related spectral perturbation analysis, as well as the results of this analysis, are provided in the Supplementary Materials (the "Event-related spectral perturbation analysis for experiment 1 and experiment 2" section).

For the phase-based connectivity analysis on the tone-locked data, given that we were analyzing data in a time period with very prominent ERPs, we wanted to make sure any differences we found in connectivity were not actually due to contributions from a common deep source. Therefore, before performing the connectivity analysis, we applied a CSD transform to the data using the default parameters specified in the Fieldtrip toolbox. The CSD transform acts to effectively subtract the volume conducted activity arising from common deep sources (*36*). We then performed a Morlet wavelet analysis to extract the Fourier spectrum of the data. The parameters of this wavelet analysis were nearly identical to the event-related spectral analysis. However, because we were specifically interested in the 200 ms before and after tone onset, we performed the analysis on a shorter time window, from −800 to +1000 ms relative to tone onset, to save computational time. We then computed the wPLI between temporal channels (T7, T8, TP7, and TP8) and all other channels (phase-based connectivity measures also reduce the effects of volume conduction). We did this analysis separately for each combination of role and tone delay. Temporal channels were used because they are the closest to the temporal lobe (involved in auditory processing) (*50*).

We then extracted the wPLI values in the theta band (4 to 7 Hz) between temporal channels (T7, T8, TP7, TP8) and fronto-central channels (FCz, FC1, and FC2) (*35*). We extracted wPLI values averaged over time for two time windows: −200 to 0 ms relative to tone onset (prestimulus period) and 0 to +200 ms relative to tone onset (stimulus period).

Because we were interested in increases in theta connectivity in response to task-relevant stimuli, we calculated a difference measure to represent the increase in theta connectivity between the prestimulus period and stimulus period, the wPLI difference. We calculated this measure by subtracting the wPLI values in the prestimulus period from the wPLI values in the stimulus period (wPLI difference = $\text{wPLI}_{\text{Stimulus period}} - \text{wPLI}_{\text{Prestimulus period}}$).

Statistical analysis

In the analysis of the behavioral data, when correcting for multiple comparisons, we used the Bonferroni method to limit the type 1 error rate as far as possible. However, for the neurophysiological data, in which the effects may be less robust (*51*) (but nevertheless require corrections for multiple comparisons), to avoid an excessive loss of power, we used the FDR method when correcting for multiple comparisons. We performed all statistical analysis in RStudio (2021.09.0, Build 351) with the packages rstatix, car, and ez.

Behavioral data

We first analyzed the data from each participant separately, with a between-trials ANOVA, to make sure that participants were able to accurately distinguish the three delay intervals, and participants who were not able to do so were removed from further analysis (see above). Subsequent analyses, the results of which are reported in the Results section, were performed on raw interval estimates, pooled across all tone delays (*6*, *9*, *24*). Namely, we conducted a repeated-measures ANOVA on the data with role as the main factor, as described in Results. Where necessary, we then performed follow-up paired pairwise tests on the data, as described in Results. For the sake of completeness, we also performed an identical analysis on the participants' JEs, reported in the Supplementary Materials (the "Analysis on the JEs" section).

N100 data

N100 amplitude data were combined across tone delays and averaged across the channels Fz, FCz, and Cz (*9*, *24*). The data were analyzed using the same model as the behavioral data, namely, a repeated measures ANOVA with role as the main factor. When necessary, we ran follow-up paired pairwise *t* tests on the data.

Spectral data

For the statistical analysis of the spectral data, we combined the data across tone delays, as we did for the behavioral and ERP data. We analyzed the phased-based connectivity data with the same models described above for the behavioral and N100 data. Namely, we first conducted a repeated-measures ANOVA on the data with role as the main factor, followed by paired pairwise *t* test where necessary. We performed the event-related spectral perturbation analysis by means of cluster-based permutation analyses, described in greater detail in the Supplementary Materials.

Supplementary Materials

This PDF file includes:

Supplementary Text
Figs. S1 to S5
Table S1
References (*52*–*55*)

Other Supplementary Material for this manuscript includes the following:

MDAR Reproducibility Checklist

REFERENCES AND NOTES

1. L. L. Royakkers, Q. R. Van Est, A literature review on new robotics: Automation from love to war. *Int. J. Soc. Robot.* **7**, 549–570 (2015).
2. A. Tapus, M. Mataric, B. Scassellati, Socially assistive robotics [Grand Challenges of Robotics]. *IEEE Robot. Automation Mag.* **14**, 35–42 (2007).
3. A. Wykowska, Social robots to test flexibility of human social cognition. *Int. J. Soc. Robot.* **12**, 1203–1211 (2020).
4. A. Wykowska, Robots as mirrors of the human mind. *Curr. Dir. Psychol. Sci.* **30**, 34–40 (2021).
5. J. D. Loehr, The sense of agency in joint action: An integrative review. *Psychon. Bull. Rev.* **29**, 1089–1117 (2022).
6. M. Jenkins, O. Esemezie, V. Lee, M. Mensingh, K. Nagales, S. S. Obhi, An investigation of "We" agency in co-operative joint actions. *Psychol. Res.* **85**, 3167–3181 (2021).
7. A. Sahaï, A. Desantis, O. Grynszpan, E. Pacherie, B. Berberian, Action co-representation and the sense of agency during a joint Simon task: Comparing human and machine co-agents. *Conscious. Cogn.* **67**, 44–55 (2019).
8. P. Haggard, S. Clark, J. Kalogeras, Voluntary action and conscious awareness. *Nat. Neurosci.* **5**, 382–385 (2002).
9. E. A. Caspar, J. F. Christensen, A. Cleeremans, P. Haggard, Coercion changes the sense of agency in the human brain. *Curr. Biol.* **26**, 585–592 (2016).
10. A. Pesquita, R. L. Whitwell, J. T. Enns, Predictive joint-action model: A hierarchical predictive approach to human cooperation. *Psychon. Bull. Rev.* **25**, 1751–1769 (2018).
11. L. Zapparoli, E. Paulesu, M. Mariano, A. Ravani, L. M. Sacheli, The sense of agency in joint actions: A theory-driven meta-analysis. *Cortex* **148**, 99–120 (2022).
12. D. M. Wolpert, R. C. Miall, M. Kawato, Internal models in the cerebellum. *Trends Cogn. Sci.* **2**, 338–347 (1998).
13. P. E. Keller, G. Novembre, J. Loehr, "Musical ensemble performance: Representing self, other and joint action outcomes" in *Shared Representations: Sensorimotor Foundations of Social Life*, S. S. Obhi, E. S. Cross, Eds. (Cambridge Univ. Press, 2016), pp. 280–310.
14. D. M. Wolpert, K. Doya, M. Kawato, A unifying computational framework for motor control and social interaction. *Philos. Trans. R Soc. Lond. B Biol. Sci.* **358**, 593–602 (2003).
15. L. M. Sacheli, E. Arcangeli, E. Paulesu, Evidence for a dyadic motor plan in joint action. *Sci. Rep.* **8**, 5027 (2018).
16. A. Sahaï, E. Pacherie, O. Grynszpan, B. Berberian, Predictive mechanisms are not involved the same way during human-human vs. human-machine interactions: A review. *Front. Neurorobot.* **11**, 52 (2017).
17. S. Thellman, M. M. De Graaf, T. Ziemke, Mental state attribution to robots: A systematic review of conceptions, methods, and findings. *ACM Trans. Hum. Robot Interact.* **11**, 1–51 (2022).
18. A. Stenzel, E. Chinellato, M. A. T. Bou, A. P. Del Pobil, M. Lappe, R. Liepelt, When humanoid robots become human-like interaction partners: Corepresentation of robotic actions. *J. Exp. Psychol. Hum. Percept. Perform.* **38**, 1073–1077 (2012).
19. A. Sahaï, E. Caspar, A. De Beir, O. Grynszpan, E. Pacherie, B. Berberian, Modulations of one's sense of agency during human-machine interactions: A behavioural study using a full humanoid robot. *Q. J. Exp. Psychol.* **76**, 606–620 (2023).
20. G. Metta, L. Natale, F. Nori, G. Sandini, D. Vernon, L. Fadiga, C. von Hofsten, K. Rosander, M. Lopes, J. Santos-Victor, A. Bernardino, L. Montesano, The iCub humanoid robot: An open-systems platform for research in cognitive development. *Neural Netw.* **23**, 1125–1134 (2010).
21. S. Marchesi, D. D. Tommaso, J. Perez-Osorio, A. Wykowska, Belief in sharing the same phenomenological experience increases the likelihood of adopting the intentional stance toward a humanoid robot. *Technol. Mind Behav.* **3**, 10.1037/tmb0000072 (2022).
22. J. F. Kelley, An empirical methodology for writing user-friendly natural language computer applications, in *Proceedings of the SIGCHI Conference on Human Factors in Computing Systems - CHI '83* (ACM, 1983), pp. 193–196.
23. D. Dennett, "Intentional systems theory" in *The Oxford Handbook of Philosophy of Mind*, A. Beckermann, B. P. McLaughlin, S. Walter, Eds. (Oxford Academic, 2009), pp. 339–350.
24. E. A. Caspar, S. Lo Bue, P. A. Magalhães De Saldanha da Gama, P. Haggard, A. Cleeremans, The effect of military training on the sense of agency and outcome processing. *Nat. Commun.* **11**, 4366 (2020).
25. S. K. Poonian, J. McFadyen, J. Ogden, R. Cunnington, Implicit agency in observed actions: Evidence for N1 suppression of tones caused by self-made and observed actions. *J. Cogn. Neurosci.* **27**, 752–764 (2015).
26. N. K. Bolt, J. D. Loehr, Sensory attenuation of the auditory P2 differentiates self- from partner-produced sounds during joint action. *J. Cogn. Neurosci.* **33**, 2297–2310 (2021).
27. A. Bendixen, I. SanMiguel, E. Schröger, Early electrophysiological indicators for predictive processing in audition: A review. *Psychophysiology* **83**, 120–131 (2012).
28. I. SanMiguel, J. Todd, E. Schröger, Sensory suppression effects to self-initiated sounds reflect the attenuation of the unspecific N1 component of the auditory ERP. *Psychophysiology* **50**, 334–343 (2013).
29. K. Lange, Brain correlates of early auditory processing are attenuated by expectations for time and pitch. *Brain Cogn.* **69**, 127–137 (2009).
30. E. Schröger, A. Marzecová, I. SanMiguel, Attention and prediction in human audition: A lesson from cognitive psychophysiology. *Eur. J. Neurosci.* **41**, 641–664 (2015).
31. P. Bäss, T. Jacobsen, E. Schröger, Suppression of the auditory N1 event-related potential component with unpredictable self-initiated tones: Evidence for internal forward models with dynamic stimulation. *Int. J. Psychophysiol.* **70**, 137–143 (2008).
32. Y. F. Hsu, J. A. Hämäläinen, F. Waszak, Both attention and prediction are necessary for adaptive neuronal tuning in sensory processing. *Front. Hum. Neurosci.* **8**, 152 (2014).
33. P. Kok, D. Rahnev, J. F. Jehee, H. C. Lau, F. P. de Lange, Attention reverses the effect of prediction in silencing sensory signals. *Cereb. Cortex* **22**, 2197–2206 (2012).
34. S. A. Hillyard, E. K. Vogel, S. J. Luck, Sensory gain control (amplification) as a mechanism of selective attention: Electrophysiological and neuroimaging evidence. *Philos. Trans. R Soc. Lond. B Biol. Sci.* **353**, 1257–1270 (1998).
35. J. F. Cavanagh, M. J. Frank, Frontal theta as a mechanism for cognitive control. *Trends Cogn. Sci.* **18**, 414–421 (2014).
36. M. X. Cohen, S. van Gaal, K. R. Ridderinkhof, V. A. Lamme, Unconscious errors enhance prefrontal-occipital oscillatory synchrony. *Front. Hum. Neurosci.* **3**, 54 (2009).
37. J. F. Cavanagh, L. Zambrano-Vazquez, J. J. Allen, Theta lingua franca: A common mid-frontal substrate for action monitoring processes. *Psychophysiology* **49**, 220–238 (2012).
38. A. M. Cravo, P. M. Claessens, M. V. Baldo, Voluntary action and causality in temporal binding. *Exp. Brain Res.* **199**, 95–99 (2009).
39. S. J. Luck, S. Fan, S. A. Hillyard, Attention-related modulation of sensory-evoked brain activity in a visual search task. *J. Cogn. Neurosci.* **5**, 188–195 (1993).
40. P. Baess, J. Horváth, T. Jacobsen, E. Schröger, Selective suppression of self-initiated sounds in an auditory stream: An ERP study. *Psychophysiology* **48**, 1276–1283 (2011).
41. R. Flach, G. Knoblich, W. Prinz, Off-line authorship effects in action perception. *Brain Cogn.* **53**, 503–513 (2003).
42. G. Knoblich, R. Flach, Action identity: Evidence from self-recognition, prediction, and coordination. *Conscious. Cogn.* **12**, 620–632 (2003).
43. P. E. Keller, G. Knoblich, B. H. Repp, Pianists duet better when they play with themselves: On the possible role of action simulation in synchronization. *Conscious. Cogn.* **16**, 102–111 (2007).
44. F. Ciardo, F. Beyer, D. Tommaso, A. Wykowska, Attribution of intentional agency towards robots reduces one's own sense of agency. *Cognition* **194**, 104109 (2020).
45. A. Bandura, "Social cognitive theory of personality" in *Handbook of Personality: Theory and Research*, L. A. Pervin, O. P. John, Eds. (Guilford Press, 1999), pp. 154–196.

46. J. Peirce, J. R. Gray, S. Simpson, M. MacAskill, R. Höchenberger, H. Sogo, E. Kastman, J. K. Lindeløv, PsychoPy2: Experiments in behavior made easy. *Behav. Res. Methods* **51**, 195–203 (2019).
47. G. Metta, P. Fitzpatrick, L. Natale, Yarp: Yet another robot platform. *Int. J. Adv. Robot. Syst.* **3**, 8 (2006).
48. A. Roncone, U. Pattacini, G. Metta, L. Natale, "A cartesian 6-DoF gaze controller for humanoid robots" in *Proceedings of Robotics: Science and Systems* (Robotics: Science and Systems, 2016).
49. S. Marchesi, D. Ghiglino, F. Ciardo, J. Perez-Osorio, E. Baykara, A. Wykowska, Do we adopt the intentional stance toward humanoid robots? *Front. Psychol.* **10**, 450 (2019).
50. R. K. Chikara, W. C. Lo, L. W. Ko, Exploration of brain connectivity during human inhibitory control using inter-trial coherence. *Sensors* **20**, 1722 (2020).
51. S. J. Luck, N. Gaspelin, How to get statistically significant effects in any ERP experiment (and why you shouldn't). *Psychophysiology* **54**, 146–157 (2017).
52. M. Siebertz, P. Jansen, Diverging implicit measurement of sense of agency using interval estimation and Libet clock. *Conscious. Cogn.* **99**, 103287 (2022).
53. M. X. Cohen, *Analyzing Neural Time Series Data: Theory and Practice* (MIT Press, 2014).
54. E. Maris, R. Oostenveld, Nonparametric statistical testing of EEG- and MEG-data. *J. Neurosci. Methods* **164**, 177–190 (2007).
55. S. Crottaz-Herbette, V. Menon, Where and when the anterior cingulate cortex modulates attentional response: Combined fMRI and ERP evidence. *J. Cogn. Neurosci.* **18**, 766–780 (2006).

Acknowledgments: We would like to thank G. Siri, A. Bazzuro, and F. Rospo for assistance with the data collection. We would also like to thank S. Marchesi for comments and contributions regarding the intentional interaction paradigm. **Funding:** This work has received support from the following funding schemes: European Research Council under the European Union's Horizon 2020 research and innovation program; ERC Starting Grant, G.A. number: ERC-2016-StG-715058, awarded to AW; F.C. was supported by H2020–Marie Skłodowska-Curie Action agreement no. 893960. The content of this article is the sole responsibility of the authors. The European Commission or its services cannot be held responsible for any use that may be made of the information it contains. **Author contributions:** U.P.N., F.C., K.K., and A.W. conceptualized and designed the study. U.P.N. and D.D.T. programmed the experiments. U.P.N. collected the data and analyzed the data. U.P.N., F.C., K.K., and A.W. discussed and interpreted the results. U.P.N., F.C., and K.K. wrote the original draft of the manuscript. U.P.N., F.C., K.K., and A.W. wrote and edited subsequent versions of the manuscript. All authors reviewed the manuscript and approved it. **Competing interests:** The authors declare that they have no competing interests. **Data and materials availability:** The datasets analyzed for the current study, as well as the analysis scripts, are available at the following link: https://osf.io/rtb8f/ (OSF project: Neural markers of self-other integration in joint action—why attribution of intentionality matters).

Submitted 22 June 2023
Accepted 30 May 2024
Published 26 June 2024
10.1126/scirobotics.adj3665

Science Robotics

When performing actions with robots, attribution of intentionality affects the sense of joint agency

Uma Prashant Navare, Francesca Ciardo, Kyveli Kompatsiari, Davide De Tommaso, and Agnieszka Wykowska

Sci. Robot. **9** (91), eadj3665. DOI: 10.1126/scirobotics.adj3665

Editor's summary

The sense of joint agency can emerge when humans collaborate to complete tasks. They experience a perception or feeling of shared control over joint actions and their outcomes. Navare *et al*. investigated whether the sense of joint agency can occur between humans and humanoid robots during collaborative tasks. They used neural and behavioral measures to show that human participants experienced this sense of joint agency, underpinned by mechanisms of joint sensorimotor processing, when a robot was presented as an intentional agent. However, a sense of joint agency was not felt when the robot was presented as a mechanical artifact. —Amos Matsiko

---

Science Robotics (ISSN 2470-9476) is published by the American Association for the Advancement of Science. 1200 New York Avenue NW, Washington, DC 20005. The title *Science Robotics* is a registered trademark of AAAS.