COGNITIVE SCIENCE
A Multidisciplinary Journal

Cognitive Science 47 (2023) e13316

ISSN: 1551-6709 online
DOI: 10.1111/cogs.13316

Strategic Task Decomposition in Joint Action

Jeremy Gordon,[a] Guenther Knoblich,[b] Giovanni Pezzulo[c]

[a]*School of Information, University of California, Berkeley*
[b]*Social Mind and Body Group, Department of Cognitive Science, Central European University*
[c]*Institute of Cognitive Sciences and Technologies, National Research Council*

Received 19 September 2022; received in revised form 15 June 2023; accepted 23 June 2023

Abstract

The core of human cooperation is people's ability to perform joint actions. Frequently, this requires effectively decomposing a joint task into individual subtasks, for example, when jointly shopping at the market to buy food. Surprisingly, little is known about how collaborators balance the costs of establishing a joint strategy for such decompositions and its expected benefits for a joint goal. We created a new online task that required pairs of randomly matched participants to jointly collect colored items. We then systematically varied the cognitive costs and benefits of applying a color-splitting strategy. The results showed that pairs adopted a color-splitting strategy more often when necessary to lower cognitive costs. However, once the strategy was jointly adopted, it continued to be used even when the cost–benefits changed. Our results provide first insights on how people decompose joint tasks into individual components and how decomposition strategies may evolve into conventions.

Keywords: Joint action; Joint planning; Cooperation strategies; Task decomposition

1. Introduction

We often engage in cooperative tasks, such as collecting a list of items at the market (Pacherie, 2008; Rand & Nowak, 2013; Sebanz & Knoblich, 2009; Sebanz, Bekkering, & Knoblich, 2006; van der Wel, Becchio, Curioni, & Wolf, 2021). Such tasks require the coordination of actions and plans of two (or more) people for prolonged periods of time and can

Correspondence should be sent to Giovanni Pezzulo, Institute of Cognitive Sciences and Technologies, National Research Council, Via S. Martino della Battaglia 44, 00185 Rome, Italy. E-mail: giovanni.pezzulo@istc.cnr.it

be *decomposed* into multiple subtasks (such as grouping vegetables and cheeses found in different parts of the market) that can be executed—at least in part—in parallel by different people, resulting in improved efficiency through cooperation (Clark, 1996; Henrich, 2017; Tomasello, 2014; Tomasello, Carpenter, Call, Behne, & Moll, 2005; Wu et al., 2021). Such a decomposition requires decisions about how, specifically, to divide the work, that is, *who does what and when*. Collaborators who devise and coordinate on good strategies to split their work can increase the benefits of cooperation (e.g., not buying duplicate items).

Prior work on collective decision-making and action coordination has shown that features of the information channels available between collaborators (and their environment) are critical to any cooperative pursuit. In particular, the kind and frequency of feedback provided about the performance of others influence a dyad's ability to effectively divide task demands. As one example, in a cooperative visuospatial tracking task, dyads received different kinds of information about the other's performance following each trial (Wahn, Kingstone, & König, 2017). Dyads receiving both performance information (partner efficacy), as well as target selection information (partner's task choices), performed better than dyads receiving only one type of feedback. Teams can also learn to decompose work even in the absence of objective feedback (e.g., the performance of other team members), through verbal or non-verbal social interaction (Bahrami, 2012; Bahrami et al., 2010; Pezzulo, Roche, & Saint-Bauzel, 2021). An additional example of this is a study of motor coordination in the absence of online perceptual feedback (Vesper, van der Wel, Knoblich, & Sebanz, 2013). Dyads who could not see or hear one another were tasked with performing jumps, given independent target distances, that landed at the same time. Participants were able to coordinate their jumps successfully, adjusting their motion onset timing and trajectories to land at a similar time, even when target distances were different. Of particular interest to the present work, this study also showed that participants converged upon an asymmetric (and therefore perhaps "fairer") convention to solve the task, where the jumper with the shorter distance adjusted their action more than their partner with the larger distance (and therefore the harder task). Motor simulation is putatively the driving mechanism behind this result, as the shorter jumper must evaluate the time dynamics of their partner's task even as they plan their own movement.

In parallel, work in computational cognitive science has attempted to explore whether agents endowed with a human-inspired theory of mind are better able to effectively share work in complex tasks containing a hierarchy of subtasks to complete successfully. One recent example is a model called Bayesian Delegation, trained to solve a joint cooking task, which infers which subtasks collaborators are working on, and then chooses whether to help (cooperate) or divide and conquer (work independently; Wu et al., 2021). This line of work highlights the fact that improving our understanding of human collaboration techniques is not only valuable in its own right but potentially critical to developing safe and effective Artificial Intelligence and robotics applications intended to work in conjunction with human collaborators (a key objective in the fields of Human Robot Interaction and more broadly Human Computer Interaction).

Despite this progress, and the fundamental importance of division of labor strategies to human cooperation, we lack a good understanding of how individuals split work during cooperative planning (Braun, Ortega, & Wolpert, 2009; Candidi, Curioni, Donnarumma, Sacheli,

& Pezzulo, 2015; Chackochan & Sanguineti, 2019; Chouhan & Niyogi, 2017; Clodic et al., 2017; Curioni et al., 2022; Kourtis, Sebanz, & Knoblich, 2013; Pezzulo & Dindo, 2011; Pezzulo, Donnarumma, & Dindo, 2013; Shum, Kleiman-Weiner, Littman, & Tenenbaum, 2019; Wu et al., 2021). To successfully decompose a problem, group members could benefit from agreeing upon a single shared strategy and implementing it in a consistent way. Importantly, there are costs and benefits associated with strategy selection and maintenance. Accordingly, the *establishment cost* of selecting an appropriate strategy—which we define as the cumulative cognitive resources required to jointly discover, communicate, and enact a strategy—can be significant. At the same time, successful strategies alleviate the *cognitive cost of monitoring* the others' actions (the cognitive resources unavailable to perform a subtask due to attention on a partner's actions), by enabling collaborators to independently complete their own subtasks. Notably for the present work, these monitoring costs are directly related to the amount of information required about the activity of a collaborator in order to perform one's own subtask successfully—that is, the extent to which participants' subtasks are coupled. Hence, in planning a joint strategy, collaborators must weigh the expected establishment cost of forming a strategy against the cost of performing the task without it (i.e., "cognitive monitoring" cost). Which parameters affect this balance in a joint decision-making process is not fully understood (Ang, 2021; Meyer, van der Wel, & Hunnius, 2013; Sacheli, Arcangeli, & Paulesu, 2018; Török, Pomiechowska, Csibra, & Sebanz, 2019).

To test what strategies people use to divide work we investigated (1) how collaborators balance the establishment and monitoring costs of decomposition strategies, (2) under which conditions they maintain existing strategies when cost–benefits change, (3) whether they adapt their strategies when they lead to an unfair allocation of effort, and (4) how strategy selection influences overall task performance. To address these questions, we created a new joint planning task in which two players navigated in a shared grid world to collect "gems" of two different colors in a limited amount of time. For each trial, the dyad was assigned a joint goal (e.g., collect exactly six red gems and two blue gems) and was free to decide how to accomplish the task, without being able to communicate verbally. Players could only move to connected nodes, and they had to press a button to collect a gem at their current location.

In a basic (non-strategic) approach, this task can be completed by monitoring another player's movements and tracking the number of gems of each color collected so far, to meet the joint goal of collecting an exact number of gems of each color. However, we hypothesized that dyads would develop a "color polarization" strategy to decompose the whole set of gems into subsets of gems of the same color. Crucially, forming such a strategy substantially reduces cognitive monitoring costs but incurs some establishment costs. To study the trade-offs implicit in strategy formation and maintenance, we manipulated three experimental variables (see Fig. 1).

First, we varied the monitoring costs of joint performance by introducing two different color counters. In the *low monitoring cost* (LMC) condition, participants could monitor progress toward the joint goal with the aid of a color-specific progress counter that showed the number and color of gems collected so far. In the *high monitoring cost* (HMC) condition, a similar progress counter (shown in gray) tracked only the number of gems collected but not their colors. The presence of the LMC counter should significantly decrease cognitive

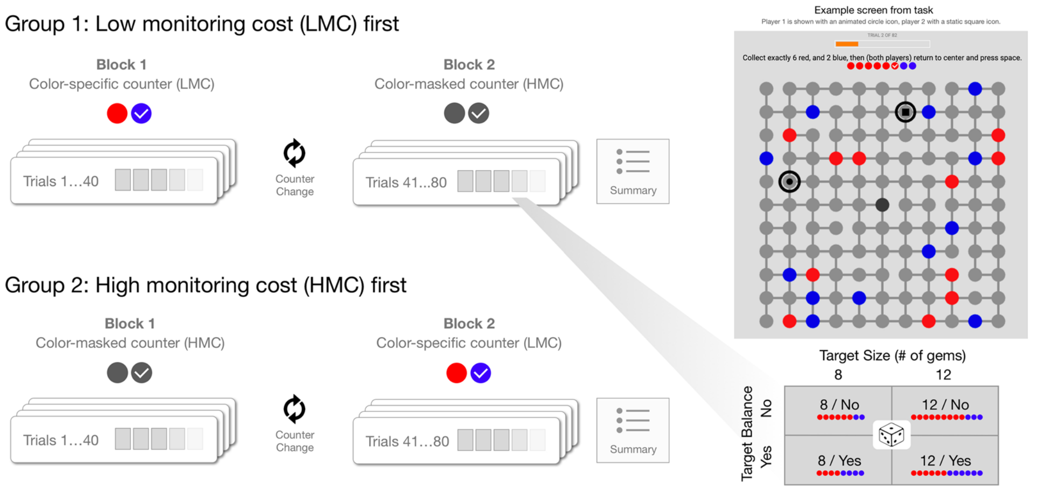

Fig 1. Experimental design. A between-subjects factor manipulated which counter type each dyad saw first (the gray "high monitoring cost" [HMC] counter, or the colored "low monitoring cost" [LMC] counter). Additionally, a 2 × 2 within-subjects manipulation randomized both the target size (eight or 12 gems total), and the target's color balance (either balanced with an even number or red and blue gems or not balanced with 75% of gems red).

monitoring costs due to the shared, color-specific informational aid displaying current task needs (e.g., "I can pick up this blue gem since we still need two more"). With the non-color specific HMC counter, players' subtasks were tightly coupled, and, therefore, required costly continuous monitoring ("My teammate is about to pick up the last blue gem, so I will look for a red"). We randomized dyads into two groups: Group 1 saw the LMC block first before switching to HMC trials, and Group 2 saw the HMC block first before switching to LMC trials. Comparing the behavior of each group during Block 1, we expected fewer dyads in the LMC block to invest in establishing a task decomposition strategy since the task was feasible without it. Due to higher monitoring costs, HMC dyads should be more motivated to establish such a strategy, for example, to split the task by color.

Second, we showed each group of dyads a second block, following the first, using the alternate counter type. We hypothesized that if color polarization had already been established in the first block, and it remained effective at reducing cognitive costs under the new dynamics of the second block's task, then it would remain in place, effectively becoming a "convention" (Centola & Baronchelli, 2015; Galantucci & Garrod, 2011; Hechter & Opp, 2001; Young, 1993). Therefore, HMC-first dyads should conventionalize their previously developed strategy into the second block, despite the inherent reduction of monitoring costs due to the LMC counter. LMC-first dyads would need to develop a color polarization strategy once confronted with the more cognitively demanding HMC block.

Third, we introduced a potential imbalance of effort by creating trials in which the joint target was to collect either the same number (*balanced*, e.g., four red, four blue) or an imbalanced number of gems of each color (*imbalanced*, e.g., six red, two blue). Note that color-imbalanced trials were always imbalanced in the same direction, that is, there were no trials with a goal to collect more blue than red gems. We hypothesized that in trials with an unequal allocation, dyads who had developed color polarization would loosen this strategy to increase fairness, especially in the LMC condition where monitoring costs are low.

2. Methods

2.1. Map generation

Each of the 80 unique maps was generated by the following process. An 11 by 11 grid of nodes was created. Nodes could be connected to adjacent neighbors in each of the four cardinal directions. To produce more interesting map connectivity, we flipped a p-coin (with $p_{connected} = 0.9$) for each pair of potentially connected nodes. We confirmed that all nodes could be visited by confirming that the resultant graph was composed of a single connected component.

To ensure gems were well distributed across each map, and not too close to the central starting location, we divided the grid into four 4×4 quadrants (containing 25 nodes each), with centers at coordinates (0-indexed x & y): [2, 2], [2, 8], [8, 2] and [8, 8]. Finally, 12 gems of each color were placed in random order, such that exactly six gems would appear in each quadrant, and both the color distribution and location of gems within each quadrant would be fully random. A 3×3 middle region and each of the four corner nodes were also disallowed for gem placement. Finally, we assigned each generated, valid, map a task condition, based on our two by two by two design: number of gems to collect (eight or 12), color balance of target gems (even or 75% red), and counter type (easy monitoring, colored, or difficult monitoring, gray). While this task condition assignment influenced the dyad's goal for a trial, as communicated by the in-game task message, as well as the appearance of the counter, the map generation process was consistent across all conditions (i.e., all shared the same size, connectivity, gem color balance, and randomized placement process for gems). Note that the map design was fixed, such that all dyads saw the same set of trials in pseudo-randomized order based on which counter group they were assigned to. See Fig. 1 for a schematic of the experimental design.

2.2. Recruitment and study procedure

Participants were recruited using Prolific (www.prolific.co) between August and October of 2021. The only exclusion criteria were that participants be of at least 18 years of age, and list English as a primary language. A sample of 24 dyads (48 participants) completed the study during the data collection period, with an equal number randomly assigned to each between-subjects condition (12 in the easy monitoring condition and 12 in the difficult monitoring condition). This sample size is consistent with prior work on joint action coordination, which typically looks for large effect sizes (Curioni, Vesper, Knoblich, & Sebanz, 2019; Konvalinka, Sebanz, & Knoblich, 2023; Vesper et al., 2013). No dyads completing the experiment were excluded. The mean age was 26.2 ± 6.1. 60.4% of participants identified as male, and 37.5% identified as female.

A study coordinator scheduled each dyad for a particular time window allowing us to match each dyad through the web interface and initiate the two-player experiment.

In preparation for the experiment, both members of each dyad were instructed to visit the experiment website independently to complete the online consent form, respond to basic demographic questions, and read detailed instructions introducing the game dynamics (including navigating their player around the map, collecting gems, monitoring progress using the counter, and then finishing each trial). A unique dyad code was provided to participants, identifying each participant's dyad and randomized study condition. As a result, the instructions could be customized based on which block the dyad would see first. Instructions included only information about the type of color that would appear in the first block, to avoid biasing strategic choices based on knowledge of the eventual counter-based manipulation. Following the instructions was an invitation to try out several independent (one player) practice trials, to ensure participants entered the main experiment with sufficient familiarity with the game controls and dynamics. The practice trials, like the instructions, were customized based on the first counter block to be seen in the main experiment.

At the scheduled time, both members of each dyad signed into the study website from independent locations. When both players were present, they were automatically matched and the main two-player study began. On each trial, each participant used the keyboard to navigate a single, shared map for up to 30 s. The positions of both players were visible on the map at all times (see Figs. 1 and S4). Each player was assigned a unique visual marker, and the player's own position was animated to clearly differentiate it from their collaborator's marker. A trial was completed when either (a) both players returned to the starting location (map center) and pressed space, (b) the trial's goal was exceeded (too many red or blue gems were collected), or (c) the 30-s trial timer expired. If successful, the dyad received points for each trial equal to the number of seconds remaining in the timer. If unsuccessful (i.e., either color target was not exactly met), no points were awarded. Upon completion of the first block of 40 trials, an interstitial was presented to both members of the dyad explaining that in the next block, the counter's presentation would be changing but that the task would otherwise remain the same. This interstitial was customized to describe and show the changed appearance of the counter, based on the dyad's block order. After acknowledging the information in the interstitial, and completing a single practice round with the new counter, the experiment resumed through the remaining 40 trials in the second block.

At the end of the experiment, a summary page was presented to the dyad, reporting the dyad's score on each trial and cumulative final score. A unique confirmation code was also generated for each participant, allowing him or her to confirm the completion of the study to the coordinator. All participants who completed the study were compensated 10 euros for their participation.

2.3. Data collected

The web-based experiment allowed telemetry to be collected and stored in an online database, which constituted the primary dataset for analysis. Telemetry included: the timestamp and parameters of each player's movements, gem collection actions, and the results of each trial (number of each gem color collected by each participant, trial success, and points received). These data allowed us to reconstruct a precise trajectory of each player's movements, and therefore absolute and relative spatial locations, which was necessary for various analyses.

2.4. Dependent variables

Our key dependent variables for time series and mixed ANOVA analyses were defined as follows.

Color polarization: $\frac{1}{2}[\max(G_{R,1}, G_{B,1})/G_1 + \max(G_{R,2}, G_{B,2})/G_2]$, where $G_{C,P}$ is the single-trial count of all collected gems with color C collected by player P. This trial-level metric is most easily interpreted as the average percent of gems collected by each player in their preferred color.

Range: 0.5 when each player collects 50% of the gems of each color, 1.0 when each player collects 100% of the gems of one color and 0% of the gems of the other color.

Work balance ("fairness"): $\frac{1}{2}[1 - |\gamma_1 - \gamma_2|]$, where γ_P is the decimal percentage of all collected gems, on a given trial, collected by player P.

Range: 0 when a single player collects 100% of gems and 0.5 when each player collects 50% of gems.

Success rate: This is the percentage of trials completed successfully (in which dyads collected the exact target number of gems of each color).

Range: 0 when no trials are completed successfully, and 1.0 when all trials are completed successfully.

We also analyzed the trial score. Results were in line with trends and group differences and are reported in the Supplementary Materials (see Table S2 and Fig. S5).

3. Results

For our main analysis, we devised two high-level variables that index key aspects of the behavioral patterns and strategies expected: color strategy ("color polarization") and fairness ("work balance"), see Fig. 2. For definitions of each of these dependent variables, see the Methods section. We additionally considered the "success rate" (i.e., percent of successful trials) as a dependent variable. Fig. 3 visualizes the dynamics of these three variables as they evolve throughout the course of the experiment.

Our analysis used a mixed ANOVA with group (i.e., block order: high monitoring first "HMC-first" vs. low monitoring first "LMC-first") as between-subjects factor and task color balance (balanced "yes" vs. not balanced, i.e., more red "no") and counter type ("HMC" or "LMC") as within-subjects factors. The ANOVA has model DF of 1 and error DF of 22 for all results. We followed up the ANOVA with pairwise *t* test post hoc tests when ANOVA showed significant interactions at $p < .05$.

Table 1 shows the results of the analysis of the "color polarization" variable, which indexes the choice of a color strategy to split the task. This variable takes a minimum at 0.5 when both

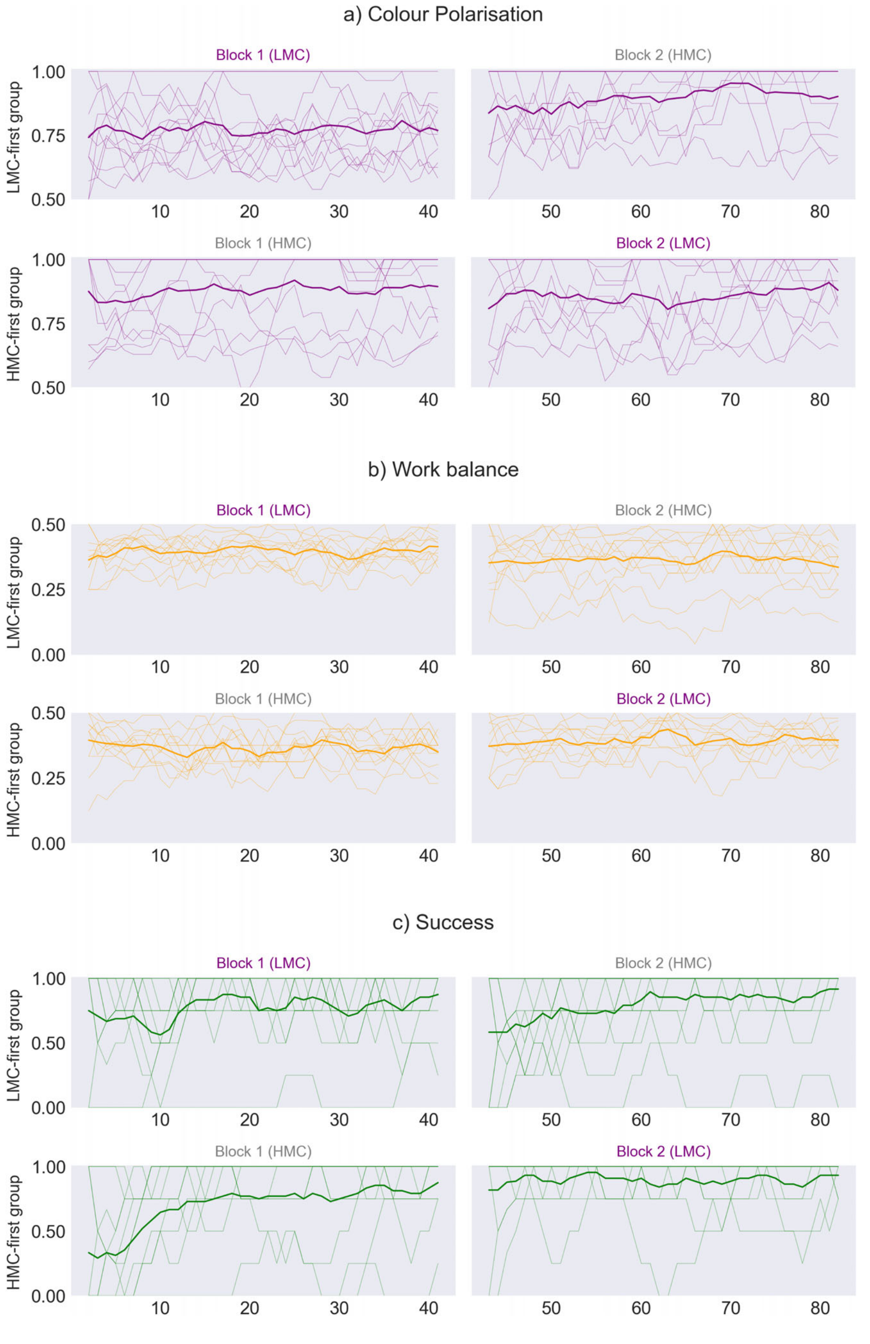

Fig 2. Time series plots of three key dependent variables, split into blocks. Each 2 × 2 panel shows a complete trial progression (1–80) for LMC-first group (top row) and HMC-first group (bottom row). Variables in each panel are as follows: (a) color polarization, (b) work balance or "fairness," and (c) success rate. The plots show both the average time series (in bold) and the time series trace for each dyad. For simplicity, these figures are not further segmented by the color-balance condition; these results are reported in Tables 1–3.

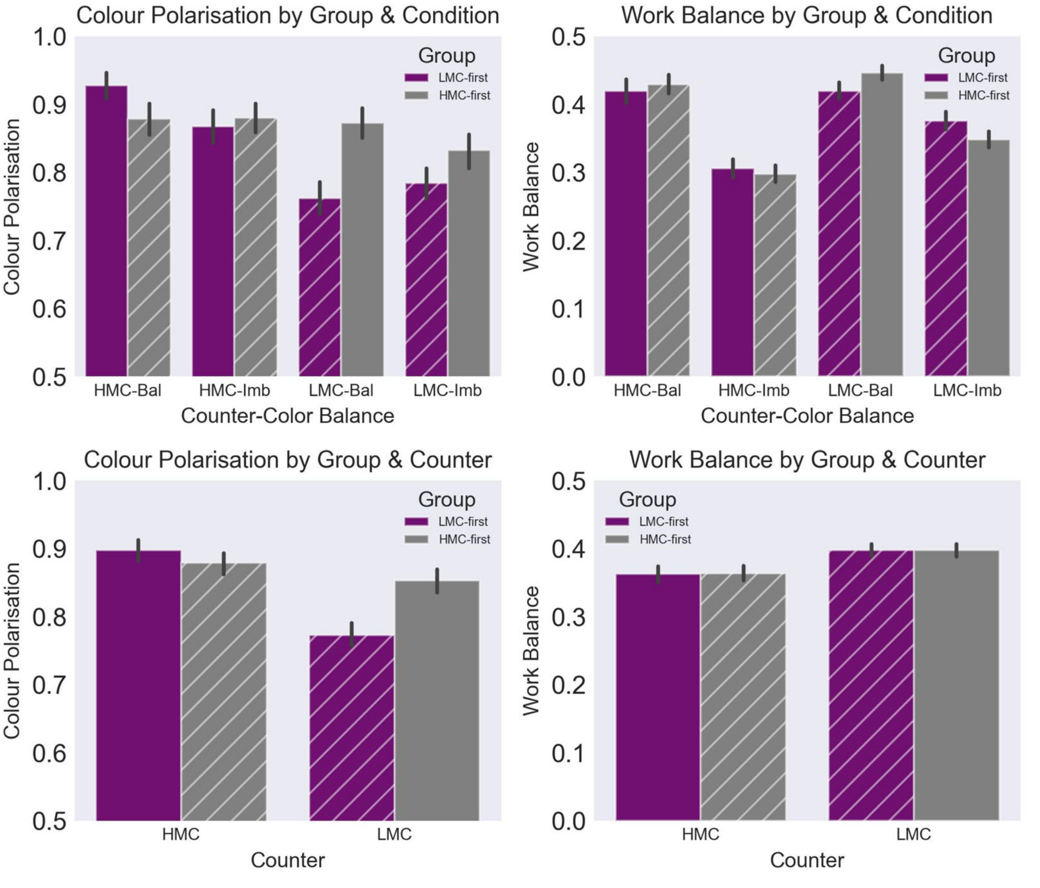

Fig 3. Plots of selected ANOVA results. Upper-left: color polarization by group and condition. Upper-right: work balance by group and condition. Lower-left: color polarization by counter type only. Lower-right: work balance by counter type only. The LMC-first group is plotted in purple, and the HMC-first group in gray. Bars are grouped as per condition (trial counter, and trial color balance for the top row, and counter type only for the bottom row). Angled hash marks indicate trials within the first block. Note that axis limits differ and are based on the range of each dependent variable.

Table 1
Color polarization: An index of use of a color-based strategy for splitting work. Means are reported, with standard deviation shown in parentheses. These results are plotted in Fig. 1

Block Order (Group)	Counter: Low Monitoring Cost (LMC) and Balance: Yes	Counter: LMC and Balance: No	Counter: High Monitoring Cost (HMC) and Balance: Yes	Counter: HMC and Balance: No
LMC-first	0.762 (0.120)	0.784 (0.114)	0.928 (0.087)	0.867 (0.143)
HMC-first	0.873 (0.158)	0.833 (0.143)	0.878 (0.135)	0.880 (0.129)

Table 2
Work balance: An index of fairness in the split of the work at every trial. These results are plotted in Fig. 1

Block Order (Group)	Counter: LMC and Balance: Yes	Counter: LMC and Balance: No	Counter: HMC and Balance: Yes	Counter: HMC and Balance: No
LMC-first	0.420 (0.049)	0.376 (0.051)	0.419 (0.121)	0.306 (0.074)
HMC-first	0.446 (0.050)	0.348 (0.075)	0.429 (0.070)	0.297 (0.048)

Table 3
Success rate: An index of task performance calculated as the percent of trials for which the exact target number of blue and red gems was collected

Block Order (Group)	Counter: LMC and Balance: Yes	Counter: LMC and Balance: No	Counter: HMC and Balance: Yes	Counter: HMC and Balance: No
LMC first	0.771 (0.227)	0.779 (0.243)	0.838 (0.209)	0.771 (0.262)
HMC first	0.888 (0.115)	0.899 (0.165)	0.679 (0.270)	0.758 (0.300)

players take half of each color and a maximum of 1.0 when each player collects only one color. We found a main effect of counter type ($F(1, 22) = 13.4, p < .001, \eta^2{}_p = .378$) showing increased split by color in the HMC trials, which was expected given the challenges of monitoring without a convention. This main effect can be seen especially clearly in comparison of Block 1 alone (prior to any effects of task switching) as shown in the bottom-left panel in Fig. 3 ($T = -6.1$, p-uncorrected $< .001$, $\eta^2{}_p = .163$, Bayes Factor >100). Comparison within the second block alone shows a weaker effect in the same direction ($T = -2.3$, p-unc $= .021$, $\eta^2{}_p = .027$, BF $= 1.9$). Additionally, we find an interaction between group and counter type ($F(1, 22) = 5.7$, $p = 0.026$, $\eta^2{}_p = .206$) showing that participants seeing the LMC counter first were significantly less likely to exhibit color polarization during the first block of LMC trials but increased their color polarization during the second block of HMC trials ($T = -3.6$, p-uncorrected $= .004$, $\eta^2{}_p = .323$, BF $= 11.86$).

Table 2 shows the results of the analysis of the "work balance" variable, which indexes fairness. This variable takes a minimum of 0.0 when one player collects all gems and a maximum of 0.5 when each player collects exactly half the gems. We found a main effect of counter ($F(1, 22) = 7.2, p = .014, \eta^2{}_p = .246$), with a greater work balance for LMC; a main effect of target balance ($F(1, 22) = 33.2$, $p < .001$, $\eta^2{}_p = .602$), with greater fairness when target colors are balanced; and an interaction ($F(1, 22) = 7.3$, $p = .013$, $\eta^2{}_p = .249$), that is, effect of target imbalance is larger for HMC trials where splitting colors is the main strategy ($T = 3.9$, p-unc $< .001$, BF $= 50.35$).

Table 3 shows the results of the analysis of the "success rate" variable. We found two 2-way interaction effects, and one 3-way effect. An interaction between counter type and group ($F(1, 22) = 7.7$, $p = .011$, $\eta^2{}_p = .260$) indicates that participants seeing the HMC block first were more successful during the LMC block ($T = 2.97$, p-unc $= .013$, Bayes BF $= 4.87$). An interaction between target color balance and group ($F(1, 22) = 5.4$, $p = .029$, $\eta^2{}_p = .198$) shows that the HMC-first group succeeded more often in imbalanced trials, while the

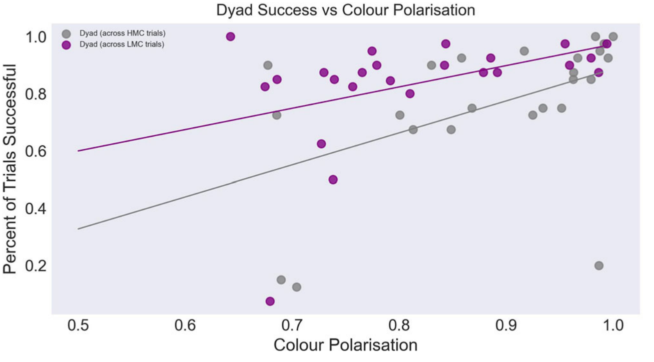

Fig 4. Dyad-level mean color polarization was predictive of success across all trials. The figure shows a dyad-level scatter plot of mean trial success against mean color polarization segmented by trial type (dyad mean within HMC trials in gray and within LMC trials in purple). We find a stronger positive relationship in HMC trials, where the color polarization strategy was especially effective. Linear regression shows a slope of 0.75 (Pearson $r = .412$, $p = .045$) for LMC trials, and a slope of 1.12 (Pearson $r = .481$, $p = .017$) for HMC trials.

LMC-first group succeeded more often in balanced trials. Finally, a three-way interaction of group, counter type, and color balance ($F(1, 22) = 8.4$, $p = .008$, $\eta^2{}_p = .276$) shows that the HMC-first group struggled most with balanced trials in the HMC block.

Finally, as an additional way to evaluate the efficacy of developing a color polarization strategy, we analyzed the relationship between dyad-level success rate and mean color polarization using linear regression (see Fig. 4). When segmented by counter type, we find a positive relationship in both trial types (Pearson $r = .41$, $p = .045$, LMC trials, $r = .48$, $p = .017$, HMC trials) but a steeper regression coefficient among HMC trials (plotted in gray).

4. Discussion

We investigated how collaborators, when decomposing a joint task, balance the costs of establishing a joint strategy of "color polarization" and the expected benefits of its execution. To understand this balance, we manipulated task parameters in such a way as to "tilt" this trade-off and then measured the effect on dyads' behavioral and strategic responses.

As pointed out by Wu et al. (2021), successful strategy selection in such a joint task requires an understanding of the future actions and intentions of a partner. It can be seen, then, as a joint planning exercise and, as such, may entail the prospective mental simulation (akin to the motor simulation posited in Vesper et al., 2013) of various counterfactual strategy choices (Klein & Crandall, 1995). As an example, let us consider a dyad having established a strong color polarization strategy (participant *A* takes red, and participant *B* takes blue) during the first, HMC, block. Confronted with a balanced trial (collect four of each color), participant *A* notices that three of the blue gems their partner will likely pursue are clustered, but the others are positioned far away, and several are in regions easily collected by *A*'s likely path. *A* may well evaluate a decision to break from their color-based work division by collecting a blue gem, thus earning additional points for the dyad through time savings (and thereby establish a new, more lenient joint strategy). Evaluation of this decision may be supported by mental simulation (prior to commitment) of plausible trajectories for each through the map. Perhaps *B* will notice their (strategy breaking) collection of the final blue gem and continue home with the dyad's target achieved. Alternatively, *B* may be distracted by their own navigation and gem collection, and confident in the efficacy of the color strategy used to date, and move to collect an additional blue gem, thus overshooting the target and failing the trial. The probability of these two outcomes (and perhaps others) may be approximated during such counterfactual simulation (Kahneman & Tversky, 1981), as well as the anticipated cognitive costs (monitoring *B*'s movements) and value of each outcome. In our example, value may be assessed as the additional points earned for time savings in the first outcome and the zero value of a failure in the second. Of course, due to the repeated nature of our task, such simulation-based reasoning might well be replaced or complemented by a more online approach (try a policy change, and check on the resultant joint performance), enabling strategy adaptation through experience across trials.

In the present work, we report four main novel findings, consistent with the evaluation and balancing of expected costs and benefits when changing strategy. First, task demands can alter the trade-off between the execution costs of a candidate strategy and the expected costs of its establishment. Second, these establishment costs create an asymmetry between groups at the transition between blocks: We observe persistence (of the established strategy) for one group and adaptation (to develop a more effective strategy) for the other. Third, considerations of equal allocation of work were present but difficult to prioritize when in conflict with overall task performance. Fourth, we found that "color polarization" is beneficial for task performance. We discuss these four findings in order.

First, despite the task being simple and potentially solvable using a strategy that is insensitive to our manipulations (e.g., to collect every gem in the surrounding until the counter is filled in), the unique establishment-monitoring cost trade-off confronted by each group in the first block drove differences in strategy selection. Specifically, in the first block, the LMC group primarily satisfied the task's demands with a simple "collect-while-monitoring" behavior and only a loose tendency toward color preferences. The LMC counter made it possible for this group to avoid the costs of establishing a strict convention. In contrast, the HMC group, contending with the large cognitive costs of monitoring and remembering task progress with the less specific counter, exhibited stronger use of a color polarization strategy to render the task feasible. Though establishment costs were significant (e.g., see the delayed rise of mean task success at the start of the first block, as well as several individual dyad examples of early failures prior to convention establishment in Fig. S3), once adopted, this strategy allowed dyads to "decouple" the task in a way that significantly lowers monitoring costs and increases success rate (see Fig. 4). Overall, these findings illustrate that task demands significantly affect strategy selection. It is worth noting, however, that within these trends, we found substantial variation both in dyad-level strategies, as well as individual behaviors (see the Supplementary Material for analysis and a brief discussion).

Second, we find a history dependence of strategy selection when tasks dynamics changed. This implies an asymmetry in either the expected establishment cost, cognitive monitoring cost, or both. Indeed, participants in the HMC-first condition largely reused the color polarization strategy that they had developed in the first block through the second, despite the changed demands. Their behavior was markedly different from participants who encountered the LMC block first, who did not develop strong color polarization. The color polarization of the HMC-first condition could therefore be considered a "convention" that was maintained even if no longer strictly necessary under the new LMC task dynamics (i.e., akin to a carryover effect of strategy but which influences joint rather than individual behavior in our cooperative task paradigm). In principle, these dyads could have benefitted from switching to a more liberal strategy that did not include a strict color assignment, or a strategy permitting a fairer allocation of work, but such considerations did not justify the costs of switching to a new convention.

Participants in the LMC-first group, who initially avoided the establishment of a strict color polarization strategy, faced a different trade-off. Satisfactory performance could no longer be achieved within a tightly coupled teammate-monitoring strategy, and thus a strategy shift, despite the costs of establishing it, was necessary. The significant drop in success following the transition for this group illustrates the cost of a weak strategy (resulting in highly coupled actions), under the HMC task dynamics. Like their counterparts encountering the HMC counter in Block 1, the majority of dyads in this group established a strict color polarization strategy, despite their success without it in the prior block. Having said this, a strict color polarization strategy was not truly mandatory to solve the task. Indeed, several dyads used spatial preferences as an alternative to (or augmentation of) the more common color polarization (for an analysis of this spatial polarization, see the Supplementary Materials).

Notably, by the second half of the HMC block, the LMC-first group exhibits the highest rate of color polarization seen. These results imply that the cost of establishing a strict color polarization strategy was significant and that dyads were only willing to make this investment after seeing they could not be successful without it. This result may also indicate that establishing an initial strategy is less costly than moving away from an already established convention. A carryover of old strategies to novel situations could be common to both individual and joint action setups. However, while in individual settings, the carryover could be linked to the cost of devising a novel strategy (Lancia et al., 2022; Todorov, 2009), in joint action settings, it could be additionally due to (or amplified by) the communication costs inherent to strategy change. Distinguishing between these possibilities is an intriguing question for future experimental study.

Third, work balance (fairness) was lower when target colors were imbalanced and during the most difficult (HMC) blocks. Furthermore, fairness was lowest in color-imbalanced HMC trials, in which strict color polarization does not allow an equal split. In these cases, it seems that considerations about the effectiveness of the strategy prevail over fair work distribution. Previous studies of fairness during joint action allocation reported inconsistent findings about the prioritization of efficiency and fairness. One study reported findings consistent with ours, suggesting that efficiency is prioritized over fairness during joint action allocation (Strachan & Török, 2020). However, another study reported that planning during joint action was more difficult and took longer time if the selected strategy implied an unequal allocation of effort between partners (Ang, 2021), suggesting that people's perception of fairness of work allocation influences decisions about joint plans. Furthermore, fairness could motivate the selection of apparently counterproductive actions, such as those that alter others' incomes at a personal cost (Dawes, Fowler, Johnson, McElreath, & Smirnov, 2007), suggesting that people care about fairness despite its potential costs (Fehr & Schmidt, 1999). Finally, considerations of fairness could help strategy selection; for example, they could help group members coordinating their decisions tacitly in conditions of uncertainty (Schelling, 1980; van Dijk & Wilke, 1996). A deeper exploration is still needed of how collaborators handle trade-offs between the fairness of individual contributions to a jointly performed task and the cognitive limitations related to its planning and performance.

Finally, the costs and benefits of selecting "color polarization" are well reflected in performance measures (see Fig. 3c). While this result might seem trivial, it is important to remark that participants following a color polarization strategy must sometimes skip gems that are close to them but have the wrong color—which might potentially slow down task completion. Despite this, in our task, the benefits of the strategy (in terms of lowering monitoring costs and the uncertainty about which gems to collect) surpass the costs of skipping gems. In the first block, participants who saw the LMC counter exhibited improved performance, whereas participants who saw the HMC counter performed poorly in the first half of the block, prior to developing a strong color polarization strategy. Across blocks, the performance of the HMC-first group remained stable, as the color convention established by this group remained effective. In contrast, the performance of the LMC-first group dropped significantly at the beginning of the second block; this drop in performance was likely a key driver of the strategy shift across blocks. Overall, the most successful block was the HMC-first group during their second block, which is the only block that benefited from an already learned (color polarization) strategy. Interestingly, color polarization was predictive of success in both HMC and LMC trials, with a stronger relationship seen among HMC trials (see Fig. 4).

In sum, our findings suggest that subtask allocation strategies are modulated by costs and benefits of joint actions (Curioni et al., 2022; Pezzulo et al., 2018; Sebanz et al., 2006) and that strict decomposition strategies emerge more often when simpler (though cognitively costly) behaviors fail to satisfy task demands. However, once a certain decomposition strategy has been adopted, it continues to be used even in cases where simpler behaviors could satisfy task demands. This raises the interesting possibility that repeated interactions result in an increased use of decomposition strategies across a range of task and cost–benefit parameters that is stable enough to be culturally transmitted (Sperber & Hirschfeld, 2004).

An open question is how such a stability trades off with the necessity to change individual and joint behavior in a continuously changing social environment (Becchio, Sartori, & Castiello, 2010; Pezzulo et al., 2018; Tomasello, 2014). Understanding how the human mind deals with such trade-offs might be a critical ingredient in future multi-agent AI systems that have to face complex problems of joint planning and coordination (Belhassein et al., 2022; Castelfranchi & Falcone, 2010; Cohen & Levesque, 1991; Donnarumma, Dindo, & Pezzulo, 2017; Pezzulo et al., 2013; Ramirez & Geffner, 2010; Tambe, 1997). Another open question regards the generalization of our findings about decomposition strategies to situations in which participants can communicate explicitly (e.g., linguistically). Converging on a specific decomposition strategy (e.g., color polarization) would be plausibly much simpler if the agreement can be reached verbally, but explicit communication could be redundant when the strategies to be followed are already established or obvious. Future studies are needed to assess under which conditions people prefer using implicit or explicit forms of communication to allocate subtasks.

Acknowledgments

This research received funding from the European Union's Horizon 2020 Framework Programme for Research and Innovation under the Specific Grant Agreement No. 945539 (Human Brain Project SGA3) and 952215(TAILOR) to Giovanni Pezzulo, the PNRR MUR project PE0000013-FAIR to Giovanni Pezzulo, and the European Research Council under the Grant Agreement No. 820213 (ThinkAhead) to Giovanni Pezzulo. Jeremy Gordon received support from the National Science Foundation Graduate Research Fellowship Program under Grant No. 2019236659.

Data availability statement

The data that support the findings of this study are openly available in Open Science Framework (OSF) at https://osf.io/db32m/?view_only=1d4312d1f9ea4cc7aea82ffdc401d2cb

References

Ang, J. (2021). *Planning for a joint task distribution: To be fair and timely* (Doctoral dissertation, Central European University), https://cognitivescience.ceu.edu/sites/cognitivescience.ceu.hu/files/attachment/basicpage/3173/angjennifer.pdf.

Bahrami, B. (2012). Together, slowly but surely: The role of social interaction and feedback on the build-up of benefit in collective decision-making. *Journal of Experimental Psychology: Human Perception and Performance*, *38*(1), 3–8. https://doi.org/10.1037/a0025708

Bahrami, B., Olsen, K., Latham, P. E., Roepstorff, A., Rees, G., & Frith, C. D. (2010). Optimally interacting minds. *Science (New York, N.Y.)*, *329*(5995), 1081–1085. https://doi.org/10.1126/science.1185718

Becchio, C., Sartori, L., & Castiello, U. (2010). Toward you: The social side of actions. *Current Directions in Psychological Science*, *19*(3), 183–188. https://doi.org/10.1177/0963721410370131

Belhassein, K., Fernández-Castro, V., Mayima, A., Clodic, A., Pacherie, E., Guidetti, M., Alami, R., & Cochet, H. (2022). Addressing joint action challenges in HRI: Insights from psychology and philosophy. *Acta Psychologica*, *222*, 103476. https://doi.org/10.1016/j.actpsy.2021.103476

Braun, D. A., Ortega, P. A., & Wolpert, D. M. (2009). Nash equilibria in multi-agent motor interactions. *PLOS Computational Biology*, *5*(8), e1000468. https://doi.org/10.1371/journal.pcbi.1000468

Candidi, M., Curioni, A., Donnarumma, F., Sacheli, L. M., & Pezzulo, G. (2015). Interactional leader–follower sensorimotor communication strategies during repetitive joint actions. *Journal of The Royal Society Interface*, *12*(110), 20150644. https://doi.org/10.1098/rsif.2015.0644

Castelfranchi, C., & Falcone, R. (2010). *Trust theory: A socio-cognitive and computational model*. Chichester: Wiley.

Centola, D., & Baronchelli, A. (2015). The spontaneous emergence of conventions: An experimental study of cultural evolution. *Proceedings of the National Academy of Sciences*, *112*(7), 1989–1994. https://doi.org/10.1073/pnas.1418838112

Chackochan, V. T., & Sanguineti, V. (2019). Incomplete information about the partner affects the development of collaborative strategies in joint action. *PLOS Computational Biology*, *15*(12), e1006385. https://doi.org/10.1371/journal.pcbi.1006385

Chouhan, S. S., & Niyogi, R. (2017). MAPJA: Multi-agent planning with joint actions. *Applied Intelligence*, *47*(4), 1044–1058. https://doi.org/10.1007/s10489-017-0938-8

Clark, H. H. (1996). *Using language*. Cambridge, UK: Cambridge University Press.

Clodic, A., Pacherie, E., Alami, R., & Chatila, R. (2017). Key elements for human-robot joint action. In R. Hakli & J. Seibt (Eds.), *Sociality and normativity for robots. Studies in the philosophy of sociality* (pp. 159–177). Cham: Springer. https://doi.org/10.1007/978-3-319-53133-5_8

Cohen, P. R., & Levesque, H. J. (1991). Teamwork. *Nous*, *25*(4), 487–512.

Curioni, A., Vesper, C., Knoblich, G., & Sebanz, N. (2019). Reciprocal information flow and role distribution support joint action coordination. *Cognition*, *187*, 21–31. https://doi.org/10.1016/j.cognition.2019.02.006

Curioni, A., Voinov, P., Allritz, M., Wolf, T., Call, J., & Knoblich, G. (2022). Human adults prefer to cooperate even when it is costly. *Proceedings of the Royal Society B: Biological Sciences*, *289*(1973), 20220128. https://doi.org/10.1098/rspb.2022.0128

Dawes, C. T., Fowler, J. H., Johnson, T., McElreath, R., & Smirnov, O. (2007). Egalitarian motives in humans. *Nature*, *446*(7137), Article, 7137. https://doi.org/10.1038/nature05651

Donnarumma, F., Dindo, H., & Pezzulo, G. (2017). Sensorimotor communication for humans and robots: Improving interactive skills by sending coordination signals. *IEEE Transactions on Cognitive and Developmental Systems*, *PP*, (99), 1–1. https://doi.org/10.1109/TCDS.2017.2756107

Fehr, E., & Schmidt, K. M. (1999). A theory of fairness, competition, and cooperation. *The Quarterly Journal of Economics*, *114*(3), 817–868.

Galantucci, B., & Garrod, S. (2011). Experimental semiotics: A review. *Frontiers in Human Neuroscience*, *5*, 11. https://doi.org/10.3389/fnhum.2011.00011

Hechter, M., & Opp, K.-D. (Eds.). (2001). *Social norms*. New York: Russell Sage Foundation.

Henrich, J. (2017). *The secret of our success. How our collective intelligence has helped us to evolve and prosper*. Princeton, NJ: Princeton University Press. https://press.princeton.edu/books/paperback/9780691178431/the-secret-of-our-success

Kahneman, D., & Tversky, A. (1981). The simulation heuristic. Technical Report 5, Department of Psychology, Stanford University, CA. https://www.dtic.mil/docs/citations/ADA099504

Klein, G., & Crandall, B. (1995). The role of mental simulation in problem solving and decision making. In P. A. Hancock (Ed.), *Local applications of the ecological approach to human-machine systems*. Hillsdale, NJ: L. Erlbaum Associates, 324–358.

Konvalinka, I., Sebanz, N., & Knoblich, G. (2023). The role of reciprocity in dynamic interpersonal coordination of physiological rhythms. *Cognition*, *230*, 105307. https://doi.org/10.1016/j.cognition.2022.105307

Kourtis, D., Sebanz, N., & Knoblich, G. (2013). Predictive representation of other people's actions in joint action planning: An EEG study. *Social Neuroscience*, *8*(1), 31–42. https://doi.org/10.1080/17470919.2012.694823

Lancia, G. L., Eluchans, M., D'Alessandro, M., Spiers, H., & Pezzulo, G. (2022). Humans account for cognitive costs when finding shortcuts: An information-theoretic analysis of navigation. *PLoS Computational Biology*, *19*(1), e1010829. https://doi.org/10.1101/2022.08.06.503020

Meyer, M., van der Wel, R. P. R. D., & Hunnius, S. (2013). Higher-order action planning for individual and joint object manipulations. *Experimental Brain Research*, *225*, (4), 579–588. https://doi.org/10.1007/s00221-012-3398-8

Pacherie, E. (2008). The phenomenology of action: A conceptual framework. *Cognition*, *107*, 179–217.

Pezzulo, G., & Dindo, H. (2011). What should I do next? Using shared representations to solve interaction problems. *Experimental Brain Research*, *211*(3), 613–630.

Pezzulo, G., Donnarumma, F., & Dindo, H. (2013). Human sensorimotor communication: A theory of signaling in online social interactions. *PLoS One*, *8*(11), e79876. https://doi.org/10.1371/journal.pone.0079876

Pezzulo, G., Donnarumma, F., Dindo, H., D'Ausilio, A., Konvalinka, I., & Castelfranchi, C. (2018). The body talks: Sensorimotor communication and its brain and kinematic signatures. *Physics of Life Reviews*, *28*, 1–21. https://doi.org/10.1016/j.plrev.2018.06.014

Pezzulo, G., Roche, L., & Saint-Bauzel, L. (2021). Haptic communication optimises joint decisions and affords implicit confidence sharing. *Scientific Reports*, *11*(1), 1051. https://doi.org/10.1038/s41598-020-80041-6

Ramirez, M., & Geffner, H. (2010). Probabilistic Plan recognition using off-the-shelf classical planners. In *Proceedings of the AAAI conference on artificial intelligence* (Vol. 24, No. 1, pp. 1121–1126).

Rand, D. G., & Nowak, M. A. (2013). Human cooperation. *Trends in Cognitive Sciences*, *17*(8), 413–425. https://doi.org/10.1016/j.tics.2013.06.003

Sacheli, L. M., Arcangeli, E., & Paulesu, E. (2018). Evidence for a dyadic motor plan in joint action. *Scientific Reports*, *8*(1), Article 5027. https://doi.org/10.1038/s41598-018-23275-9

Schelling, T. C. (1980). *The strategy of conflict*. Cambridge, MA: Harvard University Press.

Sebanz, N., Bekkering, H., & Knoblich, G. (2006). Joint action: Bodies and minds moving together. *Trends in Cognitive Scienc*, *10*(2), 70–76. https://doi.org/10.1016/j.tics.2005.12.009

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

Shum, M., Kleiman-Weiner, M., Littman, M. L., & Tenenbaum, J. B. (2019). Theory of minds: Understanding behavior in groups through inverse planning. *Proceedings of the AAAI Conference on Artificial Intelligence*, *33*(1), 6163–6370. https://doi.org/10.1609/aaai.v33i01.33016163

Sperber, D., & Hirschfeld, L. A. (2004). The cognitive foundations of cultural stability and diversity. *Trends in Cognitive Sciences*, *8*(1), 40–46. https://doi.org/10.1016/j.tics.2003.11.002

Strachan, J. W., & Török, G. (2020). Efficiency is prioritised over fairness when distributing joint actions. *Acta Psychologica*, *210*, 103158.

Tambe, M. (1997). Towards flexible teamwork. *Journal of Artificial Intelligence Research*, *7*, 83–124.

Todorov, E. (2009). Efficient computation of optimal actions. *Proceedings of the National Academy of Sciences of U S A*, *106*(28), 11478–11483. https://doi.org/10.1073/pnas.0710743106

Tomasello, M. (2014). *A natural history of human thinking*. Cambridge, MA: Harvard University Press.

Tomasello, M., Carpenter, M., Call, J., Behne, T., & Moll, H. (2005). Understanding and sharing intentions: The origins of cultural cognition. *Behavioral and Brain Sciences*, *28*(5), 675–691; discussion 691–735. https://doi.org/10.1017/S0140525x05000129

Török, G., Pomiechowska, B., Csibra, G., & Sebanz, N. (2019). Rationality in joint action: Maximizing coefficiency in coordination. *Psychological Science*, *30*(6), 930–941. https://doi.org/10.1177/0956797619842550

van der Wel, R. P. R. D., Becchio, C., Curioni, A., & Wolf, T. (2021). Understanding joint action: Current theoretical and empirical approaches. *Acta Psychologica*, *215*, 103285. https://doi.org/10.1016/j.actpsy.2021.103285

van Dijk, E., & Wilke, H. (1996). Tacit coordination and fairness judgments in social dilemmas. In W. B. G. Liebrand & D. M. Messick (Eds.), *Frontiers in social dilemmas research* (pp. 117–134). Cham: Springer.

Vesper, C., van der Wel, R. P. R. D., Knoblich, G., & Sebanz, N. (2013). Are you ready to jump? Predictive mechanisms in interpersonal coordination. *Journal of Experimental Psychology: Human Perception and Performance*, *39*, 48–61. https://doi.org/10.1037/a0028066

Wahn, B., Kingstone, A., & König, P. (2017). Two trackers are better than one: Information about the co-actor's actions and performance scores contribute to the collective benefit in a joint visuospatial task. *Frontiers in Psychology*, *8*, 669. https://www.frontiersin.org/articles/10.3389/fpsyg.2017.00669

Wu, S. A., Wang, R. E., Evans, J. A., Tenenbaum, J. B., Parkes, D. C., & Kleiman-Weiner, M. (2021). Too many cooks: Bayesian Inference for coordinating multi-agent collaboration. *Topics in Cognitive Science*, *13*(2), 414–432. https://doi.org/10.1111/tops.12525

Young, H. P. (1993). The evolution of conventions. *Econometrica*, *61*(1), 57–84. https://doi.org/10.2307/2951778

Supporting Information

Additional supporting information may be found online in the Supporting Information section at the end of the article.

Fig. S1.
Fig. S2.
Fig. S3.
Fig. S4.
Fig. S5.
Supplementary Material
Supplementary Material