I Need Your Help... or Do I? Maintaining Situation Awareness Through Performative Autonomy

Sayanti Roy
sayantiroy@pnw.edu
Purdue University Northwest

Trey Smith
trey.smith@nasa.gov
NASA Ames Research Center

Brian Coltin
brian.coltin@nasa.gov
NASA Ames Research Center

Tom Williams
twilliams@mines.edu
Colorado School of Mines

ABSTRACT

Interactive intelligent systems are increasingly being deployed in safety critical contexts like Space Exploration. For humans to safely and successfully complete collaborative tasks with robots in these contexts, they must maintain Situational Awareness of their task context without being cognitively overloaded – regardless of whether they are co-located with robots or interacting with them from a distance of thousands or millions of miles. In this paper, we present a novel autonomy design strategy we term *Performative Autonomy*, in which robots behave as if they have a lower level of autonomy than they are truly capable of (i.e., asking for advice they do not believe they truly need), for the sole purpose of maintaining interactants' Situational Awareness. In our first experiment (n=264), we begin by demonstrating that Performative Autonomy can increase Situational Awareness (SA) without overly increasing workload, and that this is true across tasks with different baseline levels of Mental Workload. In our second experiment (n=318), we consider cases where robots do not *believe* they need advice, but in fact have faulty perception or decision making capabilities. In this experiment, we only observed benefits to Performative Autonomy for specific types of questions, and only when there was significant cognitive load imposed by a secondary task; yet we observed uniform benefit on task performance for asking these types of questions regardless of task-imposed Mental workload. Our results from these two studies (total n=582) thus provide strong support for using this autonomy design strategy in future safety-critical missions as humanity explores the Moon, Mars, and beyond.

CCS CONCEPTS

• **Human-centered computing** → **Empirical studies in HCI**; • **Computing methodologies** → Natural language generation; • **Computer systems organization** → Robotic autonomy.

HRI '23, March 13–16, 2023, Stockholm, Sweden

ACM ISBN 978-1-4503-9964-7/23/03.
https://doi.org/10.1145/3568162.3576954

KEYWORDS

Situation Awareness, Workload, Levels of Autonomy, Human Robot-Teaming, Space Robotics

ACM Reference Format:
Sayanti Roy, Trey Smith, Brian Coltin, and Tom Williams. 2023. I Need Your Help... or Do I? Maintaining Situation Awareness Through Performative Autonomy. In *Proceedings of the 2023 ACM/IEEE International Conference on Human-Robot Interaction (HRI '23), March 13–16, 2023, Stockholm, Sweden.* ACM, New York, NY, USA, 10 pages. https://doi.org/10.1145/3568162.3576954

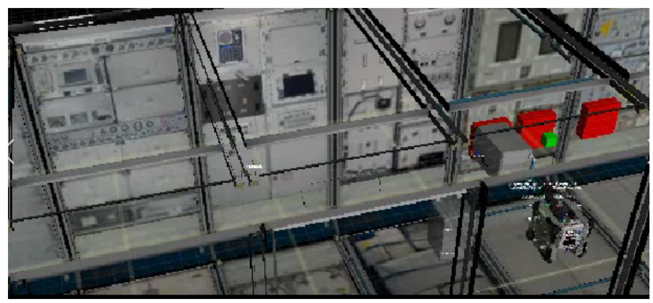

Figure 1: Simulation of an Astrobee Robot aboard the International Space Station, as shown to experimental participants.

1 INTRODUCTION

Many human-robot collaboration domains are safety-critical, characterized by high cognitive load and by high costs that can be incurred by teammate errors or from teammates becoming out-of-the-loop. Subsets of these characteristics hold, for example, in search and rescue [27, 30, 64], autonomous driving [45], industrial robotics [9, 24, 25, 40], robot-assisted surgery [55, 66], conversational agents [4, 5, 46], and, a focus of this work, space robotics.

A key task of NASA mission control is the supervision of missions involving autonomous or semi-autonomous agents, including lunar and planetary rovers, on-station robots like the Robonaut, Valkyrie, and Astrobee, Cimon, and the ISS itself. With the construction of the Lunar Orbital Platform-Gateway (LOP-G) and NASA missions increasingly more to MARS and beyond, stations and habitats will increasingly be uncrewed for long periods of time, making remote agent supervision an increasingly common task.

While these robots promise substantial benefit, their supervision imposes challenges for both robot-colocated astronauts and robot-remote mission control workers. Supervision can impose serious challenges from a human factors perspective, especially in terms of cognitive load and situation awareness. NASA investigations have shown that during complex robot operations like the deployment of the Curiosity rover, the cognitive load of the mission control room workers can reach especially high levels, leading to loss of situational awareness and potentially hazardous situations. These challenges must be addressed before autonomous solutions can be confidently pursued in space exploration contexts [33].

Several approaches have been proposed in the HRI literature to help human teammates maintain Situational Awareness. One approach for *language-capable* robots in safety-critical domains (like the free flying Astrobee and Cimon robots deployed on the ISS to perform Inter-Vehicular Activities (IVAS)) is to communicate task-relevant information to human teammates. However, simply stating information in high-stress contexts in which teammates are already cognitively overloaded may not sufficiently capture teammate attention in order to promote the deep (Level 2 or Level 3) Situation Awareness (SA) needed to truly head off catastrophe.

We propose a novel autonomy design strategy that we term *Performative Autonomy*, which could genuinely increase teammate Situation Awareness in safety-critical domains without overly burdening them in terms of Mental Workload. To briefly summarize this strategy, our key insight is to allow robots to "perform" lower levels of autonomy than they truly need to, by asking questions that they do not (so far as they are aware) actually need answers to, solely to encourage teammates to reflect on the robot's situation and stay "in the loop". In this paper we present this novel autonomy design strategy, and then evaluate it across two human-subjects experiments, each of which tests two key scientific hypotheses.

In our first experiment (n=264) crowdworkers performed a cargo transport task with Astrobee [54] robots aboard a simulated ISS, as shown in Fig. 1. By systematically varying robot communication strategy and baseline levels of Imposed Mental Workload, we assess two key research questions:

RQ1 In contexts where robots truly do not need assistance with their tasks, can Performative Autonomy increase interactant Situation Awareness without overly increasing Mental Workload?

RQ2 If so, can these benefits be gleaned across tasks with different baseline levels of Imposed Mental Workload?

While our first experiment assumes robots know the optimal action to take, this is of course not always the case. Robot sensing and decision making capabilities are, despite the wishes of roboticists worldwide, still unfortunately prone to significant sources of error. Yet we argue this should only reinforce the benefits of our proposed autonomy design strategy. That is, in contexts where robot autonomy inevitably falls short, the increased SA provided by our approach should serve the exact intended purpose of enabling human teammates to detect and correct suboptimal robot decisions stemming from faulty robot perception.

In our second experiment (n=318), participants thus engaged in a similar cargo transport scenario with a similar experimental design, but with a robot with faulty sensor capabilities, leading to situations in which robots, believing that they know precisely what to do and only stating or asking about their intended decisions for human benefit, nevertheless occasionally reveal their intent to make a suboptimal decision. Through this experiment, we assess two further research questions:

RQ3 In contexts where robots in fact do (unknowingly) need assistance, assuming Performative Autonomy increase interactant Situation Awareness without overly increasing perceived Mental Workload, does this benefit lead to increased task performance as interactants catch and correct robot error at increased rates?

RQ4 If so, as before, can these benefits be gleaned across tasks with different baseline levels of Imposed Mental Workload?

2 RELATED WORK

2.1 Human Factors Concerns in Space Exploration

In this work we are interested in two key Human Factors constructs of relevance in Space Exploration contexts: Mental Workload and Situation Awareness. Mental Workload (MW) is a Human Factors construct that captures the demand that a task places on a human's limited cognitive resources [36, 60]. Of particular importance for MW is the different ways that performance is impacted on either side of the divide that defines the limits of one's cognitive resources [62]. Specifically, when MW exceeds these limits and passes the "red line of workload" [22], performance quickly degrades and breaks down.

Situation Awareness (SA) is a Human Factors construct that captures human awareness of environmental stimuli and the ability to comprehend and make predictions based on these stimuli [15, 18]. SA is correspondingly divided into three distinct levels: Level 1 SA simply reflects a human's awareness of important stimuli within their environment; Level 2 reflects the ability to comprehend the significance of these stimuli; Level 3 reflects the ability to predict the future state of the environment based on this comprehension. SA is typically difficult to develop due to human cognitive limitations. Attentional limits prevent humans from juggling their attention in a way that allows them to achieve their goals while maintaining SA [2, 62]; Working memory limitations produce bottlenecks for the information processing needed to maintain SA [19]; MW and other stressors further constrain working memory and attentional resources [18]; and human reliance on autonomy exacerbates these problems through Out-of-the-Loop Syndrome [20, 62]

Both MW and SA are of significant importance to Space Exploration contexts. MW is especially critical to manage in hazardous situations like on-station medical events [65]. Moreover, task-imposed MW may magnify the existing physical and psychological toll taken on astronauts by long duration missions [6]. As such, addressing MW concerns have long been a priority for Space Exploration research [59]. MW is also of import due to its aforementioned downstream effects on SA. Space exploration missions often have significant SA demand both during mission planning and execution [32], yet human participants in such missions often have decreased SA [7], due to high MW and ongoing communication challenges [7]. In the next section, we will discuss one category of approach taken in previous work to try to address these challenges.

2.2 Levels of Autonomy & Adaptive Automation

One way that robotics and automation researchers has traditionally sought to manage the types of human factors concerns we have discussed is through adjustment of a robot's level of autonomy [3, 52] to meet the needs of the current context an approach known as *Adaptive Automation*. In this paradigm, a robot may increase or decrease its level of autonomy based on a number of contextual factors, such as changing goals [56] or changes in teammate workload or performance [12, 13].

Adaptive autonomy could alleviate MW and SA concerns by helping users avoid complacency. Previous research on the relationship between autonomy, workload, and SA suggests that high-autonomy robots may increase overall task performance and decrease teammate workload [26] by reducing response demands [42, 43]; but high autonomy without interaction can lead to problematic misuse and disuse [39], loss of SA [37], and Out of the Loop Unfamiliarity [17, 20]. Researchers have thus investigated how robots can strategically bringing humans back into the robot's loop [38].

But in this paradigm, robots typically bolster interactant Situation Awareness through periodic status updates, or by explaining their actions [14, 29, 31, 57]. We argue that this approach may run the risk of promoting only shallow engagement, increasing Level 1 SA alone. In this work we thus consider how robots might proactively attempt to encourage deeper reflection in order to facilitate more advanced levels of SA.

Our approach towards this goal is grounded in a key inversion of traditional approaches to Adaptive Automation. While adaptive automation might typically seek to *increase* autonomy to help manage interactant MW, we instead seek to *decrease* autonomy to help manage SA, at sufficiently sparse rate that MW is unaffected. We term this approach to autonomy design *Performative Autonomy*.

3 PERFORMATIVE AUTONOMY

If a negative side effect of increasing autonomy (acting on one's own rather than asking for assistance, even if assistance would be helpful) has the negative side effect of decreasing SA, perhaps we can increase user SA by decreasing autonomy (engaging users for assistance rather than acting on one's own, even if assistance is not clearly helpful). If a robot were to do so, it would act "as if" it had a lower level of autonomy than it was truly capable of. We thus term this strategic lowering of autonomy *Performative Autonomy*.

We specifically consider performance of lower levels of autonomy through *human-robot dialogue*. Dialogue-theoretic approaches to autonomy design have a long history, including approaches like Mixed-Initiative Interaction [1], *Collaborative Control* [21, 24, 44]. We propose six levels of dialogue autonomy, inspired by [41] and [61]. These levels (shown in Tab. 1, Col. 2) are ordered from demonstrating most autonomy to demonstrating least autonomy.

These six levels of dialogue autonomy can be understood through a Speech Act theoretic perspective [51], as the six levels align with six categories of Illocutionary Acts [50] (shown in Tab. 1, Col. 3). From this perspective, we observe that (1) an agent's choice of illocutionary act has an associated level of autonomy, (2) its level of autonomy changes from utterance to utterance, and (3) an agent can use lower levels of autonomy to achieve various social goals, e.g., using indirect speech acts (ISAs) [49] whose literal and intended meanings differ. For example, a speaker desiring a listener to bring them a wrench will often, for reasons of politeness [8], communicate indirectly, using an utterance like "Could you bring me a wrench?", which is literally a YN-Question (Yes/No Question) but is understood by sociocultural convention as a Request/Command.

Level	Strategy	Speech Act
6	Selecting option without proposal	(None)
5	Proposing and selecting a single option without opportunity for veto	Requests/ Commands
4	Proposing and selecting a single option with opportunity for veto	Statements/ Assertions
3	Proposing a single option	Suggestions
2	Requesting confirmation of a single option	YN-Questions
1	Requesting selection between multiple options	WH-Questions

Table 1: Dialogue Autonomy Levels & associated Speech Acts

In this work, we are interested in dialogue-based *Performative Autonomy* to achieve the social goal of facilitating interactant SA. When viewed through this lens, we can see previous work in which robots have sought to promote SA by proactively stating or explaining their actions [10, 11, 35, 48] as special cases of this approach that have restricted themselves to the top few levels of dialogue autonomy. In contrast, we believe that *Performative Autonomy* stands to be uniquely successful in achieving this goal when it extends beyond these initial levels. In other areas of HRI, researchers have considered the way that different types of dialogue strategies (including different types of Speech Acts) could help interactants to engage in deeper reflection [28, 58, 63, 67]. While much of that work is focused on moral robotic communication, we believe the same lessons could hold true. That is, by keeping interactants "in the loop" by asking *questions*, robots may be able to promote deeper reflection, and thus promote higher levels of SA.

One potential challenge for this approach, of course, is the workload cost imposed by this approach. As described above, higher levels of autonomy are often taken to explicitly alleviate teammates' MW. As such, it is possible that *Performative Autonomy* could increase SA *at the cost* of increasing perceived MW; or it is possible that increases in MW caused by *Performative Autonomy* could wash out any effects on SA. Moreover, MW costs have varying impacts based on the level of MW imposed by an underlying task; as such, it's possible that *Performative Autonomy* could be effective only in contexts that are not already imposing a high level of MW.

Our intuition is that this should not be the case. A robot using *Performative Autonomy* would only need to employ this strategy when making certain decisions of high importance. As such, *Performative Autonomy* should not *consistently* add to users' MW. If this is true, then *Performative Autonomy* should not impose significant MW over the course of a task. In other words, benefits to SA should outlast momentary costs to MW.

To verify these intuitions, we performed a human-subject experiment studying the differences in induced SA and MW between Dialogue Autonomy Levels 1, 2, and 6 for a robot capable of Level 6, across tasks imposing different baseline levels of MW.

4 EXPERIMENT ONE

To understand the potential effectiveness of *Performative Autonomy* as an Autonomy Design Strategy, we specifically begin by considering three key research hypotheses:

H1 In contexts where robots truly do not need assistance with their tasks, *Performative Autonomy* will increase interactant Situation Awareness.

H2 These benefits will be gleaned across tasks with different baseline levels of *Imposed Mental Workload.*

H3 These benefits will be gleaned without overly increasing perceived Mental Workload.

To test these hypotheses, we conducted a human-subject experiment in which we systematically varied two key independent variables (*Performative Autonomy* strategy and baseline level of *Imposed Mental Workload*); and systematically measured two key dependent variables (SA and perceived MW). We will now describe the design and results of this experiment.

4.1 Task Design

Prolific crowdworkers engaged in an eight-minute Cargo Transport task in an online simulation of the ISS. Participants viewed the robot's traversal of the ISS through a video window, and were tasked with providing assistance when (ostensibly) required (if an obstacle prevented successful placement of cargo into a cargo berth). The video feed watched by participants was composed at a series of clips pre-recorded in RViz, which created the illusion of a continuous live stream. While participants monitored the robot and provided input, they participated in a secondary *n-back*: a classic psychological method for imposing MW. In a separate window, participants were presented with a sequence of numbers, after each of which they are asked to enter the number that had appeared *n items back* in the sequence. This was a demanding task that made it difficult for participants to consistently monitor the robot.

Performative Autonomy was manipulated in this experiment by changing the communication the robot would use when it encountered a blocked berth (every 80 seconds). In the *Performed High Autonomy* condition, the robot did not communicate, and simply took the optimal corrective action. This was equivalent to not using *Performative Autonomy* at all. In the *Performed Medium Autonomy* condition, the robot used a YN-Question to ask for confirmation on its decision before acting on it. When using a YN-Question, if the user rejected the robot's suggestion, the robot followed up by asking for clarification between which option they *would* prefer. In the *Performed Low Autonomy* condition, the robot directly used a WH-Question to ask for arbitration between multiple possible solutions without first making its own proposal.

Imposed Mental Workload was manipulated by varying the *n* in the secondary *n-back* task. In the *Low Mental Workload* condition, n was set to zero. That is, the user merely needed to repeat back whatever number was shown. In the *Medium Mental Workload* condition, n was set to one. That is, the user needed to report the previously shown number in the sequence. In the *High Mental Workload* condition, n was set to two. That is, the user needed to report the number that had been shown *before* the previously shown number in the sequence.

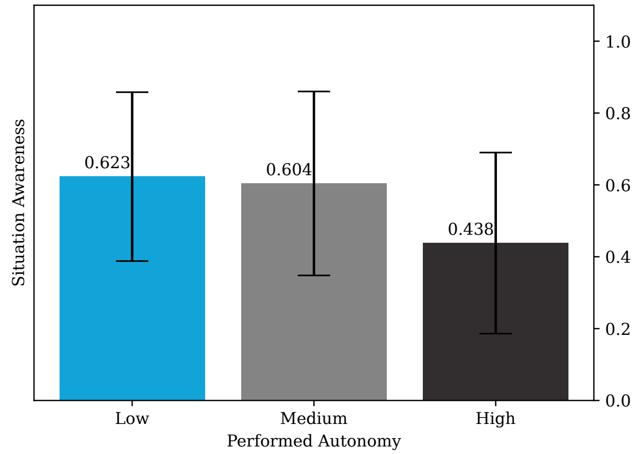

Figure 2: Study 1: Effects of Performative Autonomy on SA. In this and all future charts, error bars represent Standard Deviation.

Situation Awareness was measured by periodically (every 100 seconds) asking participants to answer a question about the task environment. Specifically, participants were asked which Berth was currently blocked, with all three listed as options. During these questions, the screen was blocked out, preventing visual inspection. Correct answers were used as evidence for higher SA.

Perceived Mental Workload was measured by periodically (every 50 seconds) asking participants to self report their level of perceived Mental Workload on a 1-5 Likert Item.

4.2 Participants

263 American participants were recruited through Prolific (M=111, F =145, Other=8). Participant ages ranged from 18 to 69 (M=33.83, SD=11.50). Each was randomly assigned to one of our three *Performative Autonomy* conditions and one of our three *Imposed Mental Workload* conditions. All experimental data and analysis scripts can be found at https://osf.io/xajpc/.

4.3 Results

4.3.1 Situation Awareness. To assess H1 and H2, we performed a Bayesian Analysis of Variance of the effect of the *Performative Autonomy* and *Imposed Mental Workload* conditions on SA. Our results provided extreme evidence for a main effect of *Performative Autonomy* (*BF* =30833.841). Post-Hoc Bayesian t-tests revealed extreme evidence that *Performed Low Autonomy* promoted more SA (M=0.62,SD=0.24) than did *Performed High Autonomy* (M=0.44,SD=0.26) (BF=12016.97), as did *Performed Medium Autonomy* (M=0.60, SD=0.25) (BF=772.28) These results (Fig. 2) thus supported both H1 and H2.

4.3.2 Perceived Mental Workload. To assess H1 and H2, we performed a Bayesian Analysis of Variance of the effect of the *Performative Autonomy* and *Imposed Mental Workload* conditions on perceived Mental Workload.

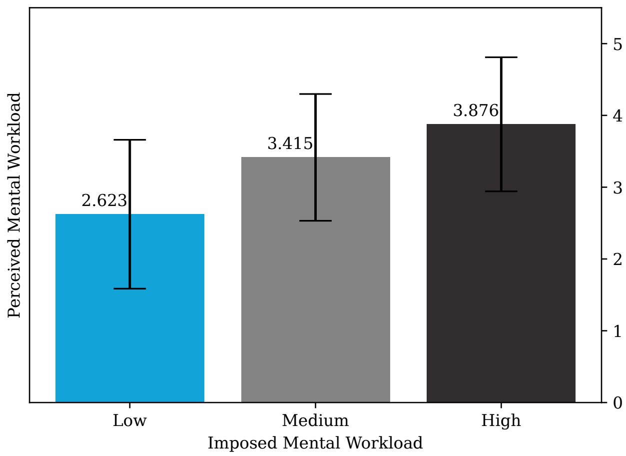

Figure 3: Study 1: Effects of Imposed Mental Workload on perceived Mental Workload.

Our results provided extreme evidence for a main effect of *Imposed Mental Workload* condition 5.109×10^{12}. Post-Hoc Bayesian t-tests revealed extreme evidence that the *Low Mental Workload* condition led to less perceived Mental Workload (M=2.623,SD=1.037) than did the *Medium Mental Workload* condition (M=3.42,SD=0.88) (BF=63550.83) or the *High Mental Workload* condition (M=3.88, SD=0.93) (BF=5.645×10^{11}); and strong evidence that the *Medium Mental Workload* condition led to less perceived Mental Workload than did the *High Mental Workload* condition (BF=26.26). These results (Fig. 2) serve as evidence that our *Imposed Mental Workload* conditions successfully manipulated perceived Mental Workload.

Evidence was found *against* a main effect of *Performative Autonomy* (BF=0.04) or an interaction between *Performative Autonomy* and *Imposed Mental Workload* (BF= 0.06), thus supporting H3.

4.4 Discussion

Our results support H1, H2 and H3. That is, (1) *Performative Autonomy* was a successful autonomy design strategy for increasing interactant SA, even in contexts with high baseline levels of cognitive load, without meaningfully increasing cognitive load.

This confirms our key intuition: that by asking meaningful questions when making important decisions, robots can encourage engagement with their actions. And because YN-Questions and Wh-Questions were equally effective, yet neither resulted in any meaningful increase in perceived Mental Workload over silence, our results suggest that either level of performed autonomy represented a reasonable and effective strategy for the robot to take.

One explicit assumption made by this experiment, however, was that *Performative Autonomy* was a strategy to be used when robots knew the optimal action to take, and thus did not actually require the advice they requested from human teammates. However, it may be valuable to consider cases where this assumption fails, as robots' perception and decision making capabilities are obviously not flawless in practice. Yet we believe that in such situations, *Performative Autonomy* as an autonomy design strategy should be no less effective, and in fact it is in just such situations that this strategy may truly shine. Specifically, contexts where robot autonomy is imperfect are those in which interactant vigilance and high SA are most critical to maintain. In these situations, the SA gains promoted by *Performative Autonomy* may be just what is needed to maintain high task performance and avoid catastrophic error. In our second experiment, we aimed to test this intuition.

5 EXPERIMENT TWO

Our second experiment considered four research hypotheses:

H4 In contexts where robots need assistance with their tasks, Performative Autonomy will, again, increase interactant Situation Awareness.

H5 These benefits will be gleaned across tasks with different baseline levels of Imposed Mental Workload.

H6 These benefits will be gleaned without overly increasing perceived Mental Workload.

H7 Performative Autonomy will thus lead to increased task performance.

5.1 Task Design

Our second experiment used an identical task design to Experiment One with one small change: the robot's decisions as to what to do were occasionally faulty. That is, when detecting a blocked berth, the robot misinterpreted which berth was blocked with some probability, and erred in its judgment of where the cargo should instead be placed in at least three out of five placement tasks. The transparency of these errors thus varied by experimental condition.

Otherwise, this experiment used the same experimental manipulations as our first experiment, but also included several key task performance measures.

Task Performance was measured in terms of Accuracy and Reaction Time. Accuracy was measured based on user responses to robot questions, with responses deemed accurate if they led to the robot performing the task-optimal action for the choice asked about by the robot. Meanwhile, Reaction Time was measured by the time taken by the user to respond to the robot's queries. Since under *Performed High Autonomy* the robot did not ask for user input, task success was only measured under *Performed Medium Autonomy* and *Performed Low Autonomy*.

5.2 Participants

318 American participants were recruited through Prolific (M=124, F=168, Other=26).Participant ages ranged from 18 to 69 (M=35.25, SD=12.24). Each was randomly assigned to one of our three *Performative Autonomy* conditions and one of our three *Imposed Mental Workload* conditions. All experimental data and analysis scripts can be found at https://osf.io/xajpc/.

5.3 Results

5.3.1 Situation Awareness. To assess H4 and H5, we performed a Bayesian Analysis of Variance of the effect of the *Performative Autonomy* and *Imposed Mental Workload* conditions on SA. Our results provided extreme evidence for an interaction effect between *Performative Autonomy* and *Imposed Mental Workload* (BF=298.08), as shown in Fig. 4, and strong evidence against a main effect of

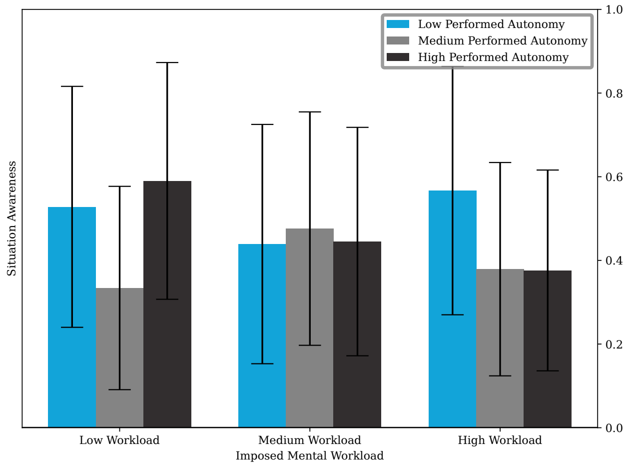

Figure 4: Study 2: Effects of Performative Autonomy and Imposed Mental Workload on SA.

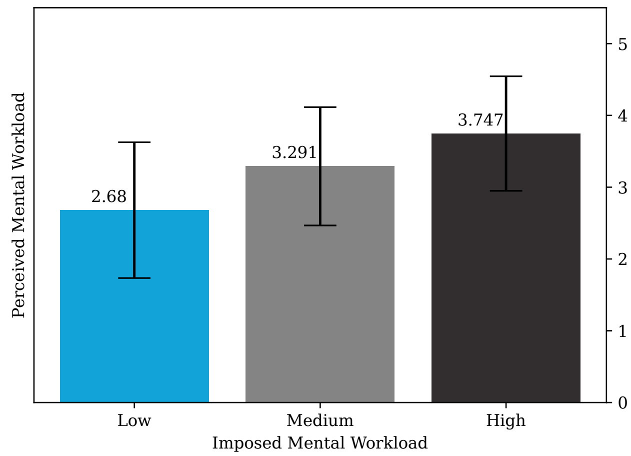

Figure 5: Study 2: Effects of Imposed Mental Workload on perceived Mental Workload.

Performative Autonomy (BF=0.09) or a main effect of *Imposed Mental Workload* (BF=0.05). These results thus failed to support H4 or H5.

In the *Low Mental Workload* condition, Post-Hoc Bayesian t-tests revealed extreme evidence (BF=236.58) that *Performed High Autonomy* (M=0.59, SD=0.28) promoted more SA than *Performed Medium Autonomy* (M=0.34,SD=0.24), and strong evidence (BF 11.87) that *Performed Low Autonomy* (M=0.53, SD=0.29) promoted more SA than *Performed Medium Autonomy*. In the *High Mental Workload* condition, Post-Hoc Bayesian t-tests revealed moderate evidence (BF=6.76) that *Performed Low Autonomy* (M=0.57, SD=0.30) promoted more SA than *Performed Medium Autonomy* (M=0.38,SD=0.26), and moderate evidence (BF=5.10) that *Performed Low Autonomy* promoted more SA than *Performed High Autonomy* (M=0.38,SD=0.24). At least moderate evidence was found against all other differences within each *Imposed Mental Workload* condition.

5.3.2 *Perceived Mental Workload.*

To assess Hypothesis H6, we performed a Bayesian Analysis of Variance of the effect of the *Performative Autonomy* and *Imposed Mental Workload* conditions on perceived Mental Workload.

Our results provided extreme evidence in favor of a main effect of *Imposed Mental Workload* (BF=4.168×10^{13}), as shown in Fig. 5. Strong evidence was found against a main effect of *Performative Autonomy* (BF 0.07) and against an interaction effect (BF 0.07), thus partially (due to the lack of support for H4) supporting H6.

Post-Hoc Bayesian t-tests revealed extreme evidence (15351.89) that the *High Mental Workload* condition led to higher perceived Mental Workload (M=3.75, SD=0.80) than did the *Medium Mental Workload* condition (M=3.29, SD=0.82), extreme evidence (BF 9.550×10^{12}) that the *High Mental Workload* condition led to higher perceived Mental Workload than did the *Low Mental Workload* condition (M=2.68, SD=0.95), and extreme evidence (BF 201.04) that the textitMedium Mental Workload condition led to higher perceived Mental Workload than did the *Low Mental Workload* condition. These results demonstrated that our *Imposed Mental Workload* conditions successfully manipulated perceived Mental Workload.

5.3.3 *Task Performance.*

To assess H7, we performed a Bayesian Analysis of Variance of the effect of the *Performative Autonomy* and *Imposed Mental Workload* conditions on Task Performance, in terms of Accuracy and Reaction Time. Because task performance was measured in this work in terms of how participants responded to robots' questions, task performance was only analyzed in the *Performed Medium Autonomy* and *Performed Low Autonomy* conditions.

Accuracy of Responses. Our results provided extreme evidence for a main effect of *Performative Autonomy* (BF=8.635×10^7), as shown in Fig. 6, suggesting that *Performed Low Autonomy* led to a higher proportion of correct responses to the robot's questions (M=77.9%, SD=23%) than did *Performed Medium Autonomy* (M=53.7%, SD=21.7%. Moderate evidence was found against a main effect of *Imposed Mental Workload* (BF 0.15). Anecdotal evidence was found against an interaction effect (BF 0.07), as shown in Fig. 7. Because an interaction effect could not be ruled out, we performed a post-hoc analysis comprised of pairwise Bayesian t-tests.

In the *Low Mental Workload* condition, Post-Hoc Bayesian t-tests revealed extreme evidence (BF=130323) that *Performed Low Autonomy* led to a higher proportion of correct responses to the robot's questions (M=.846, SD=.221) than *Performed Medium Autonomy* (M=.559,SD=.202). In the *Medium Mental Workload* condition, Post-Hoc Bayesian t-tests revealed anecdotal evidence (BF=0.92) against a difference between *Performed Low Autonomy* (M=.71, SD=.22) and *Performed Medium Autonomy* (M=.611,SD=.231). In the *High Mental Workload* condition, Post-Hoc Bayesian t-tests revealed very strong evidence (BF=63.15) that *Performed Low Autonomy* led to a higher proportion of correct responses to the robot's questions (M=.78, SD=.23) than *Performed Medium Autonomy* (M=.57,SD=.22).

Reaction Time of Responses. Our results provided extreme evidence in favor of a main effect of *Performative Autonomy* (BF=4.433×10^{34}), as shown in Fig. 8, suggesting that the *Performed Low Autonomy* condition led to faster reaction times (M=8.74s, SD=2.85s) than did the *Performed Medium Autonomy* condition (M=15.90s,

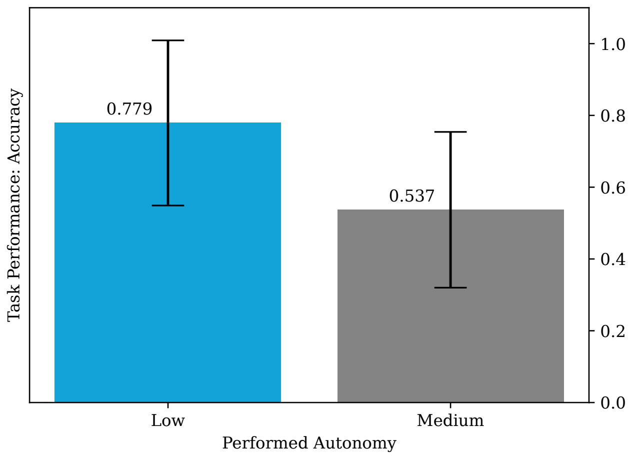

Figure 6: Study 2: Effects of Performative Autonomy on Task Performance: Accuracy.

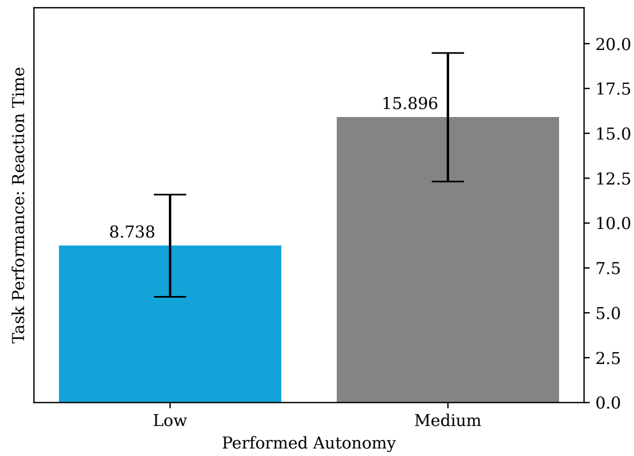

Figure 8: Study 2: Effects of Performative Autonomy on Task Performance: Reaction Time.

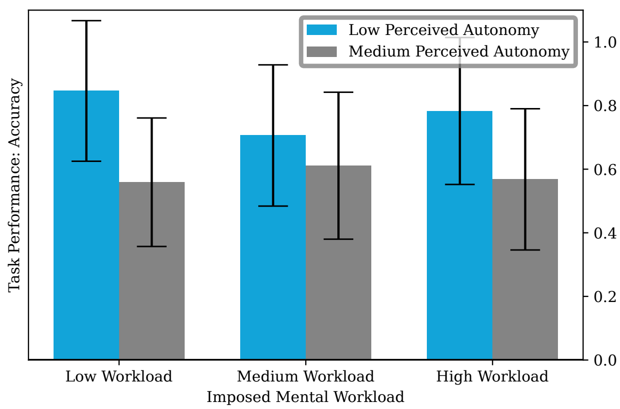

Figure 7: Study 2: Effects of Performative Autonomy and Imposed Mental Workload on Task Performance: Accuracy.

SD=3.58s). Strong evidence was found against a main effect of *Imposed Mental Workload* (BF 0.18) or an interaction (BF 0.27).

5.4 Discussion

The results of our second experiment failed to support H4 or H5. That is, while we had expected to see the same benefits of *Performative Autonomy* for SA as we had seen in Experiment One, this was not the case. Specifically, we observed that *Performative Autonomy* as an autonomy design strategy only promoted benefits to SA in this experiment when participants were under high workload, and thus more likely to be suffering from decreased awareness and most in need of being brought back into the loop. But moreover, we observed that when workload was especially low, *Performed Medium Autonomy* did more harm than good. These results carried over into our task performance metrics, supporting H7. That is, the places where we saw benefits of *Performed Low Autonomy* over *Performed Medium Autonomy* in terms of SA, we also saw benefits in terms of task accuracy; and overall, *Performed Low Autonomy* required a shorter reaction time than did *Performed Medium Autonomy*. These results raise two key questions.

The first question raised by our results is why, in this experiment, *Performed Medium Autonomy* (asking Yes/No Questions) did more harm than good? These results may be primarily due to the extra step needed by interactants under this autonomy strategy. While under the *Low Performed Autonomy* strategy, the robot merely asks for selection between multiple options. In contrast, under the *Medium Performed Autonomy* strategy, the robot first asked for approval or disapproval of its suggestion. While in the first experiment, the robot's suggestions were guaranteed optimal, and thus participants were unlikely to reject these requests. In contrast, in the second experiment, the robot's suggestions were occasionally incorrect. In such cases, the participant needed to reject the robot's suggestion before selecting a more appropriate one.

But moreover, the added time needed in these cases may have gone beyond the need to simply make an additional selection, as reflected in the seven extra seconds needed on average in the *Medium Performed Autonomy* condition. First, it is possible that processing YN-Questions themselves required an extra cognitive step for human interactants, as they needed to (1) determine the best option, and then (2) determine whether this aligned with the robot's choice, whereas WH-Questions only required the user to perform that first step. That is, perhaps the YN-Question, although requiring additional decision making by the robot, did not actually save the human any cognitive effort. Second, it is possible that YN-Questions in fact required a significantly deeper reasoning process involving social or moral cognition. Following Malle et al. [34]'s Path Model of Blame, while WH-Questions merely asked participants to make a decision, YN-Questions may have further prompted participants to consider whether an adverse or blameworthy event occurred, and if so, decide how to address it. That is, there may have been an extra cognitive cost to deciding whether an error had occurred, *and* an extra cognitive cost to deciding how to respond to the event, and whether and how to communicate blame and issue corrections.

Overall, these results suggest a complicated nonlinear relationship between *Performative Autonomy* and cognitive processing. Differentiating between these possibilities will be important in future work, both through refined experimental design and task design, as well as through providing opportunities for free-response through surveys or interviews post-experiment.

The second question raised by our results is why, in Experiment Two, our results were moderated by *Imposed Mental Workload* when they were not in Experiment One? It could be that robot error led to increased cognitive load overall, changing the overall level of MW imposed by each of our three conditions. Yet the perceived Mental Workload levels in Experiment Two were nearly identical to those observed in Experiment One, and certainly not higher. Moreover, as in Experiment One, we saw no effect of *Performative Autonomy* on perceived Mental Workload, supporting H6. However, given the differences in reaction time observed in this experiment, it appears that *Performative Autonomy did* affect participants' cognitive processes, even if they were not aware of these differences, and even if these results were not felt as imposing on MW itself. This question might be further investigated in future work through the use of other measures of MW and related constructs.

6 GENERAL DISCUSSION

Through our two experiments, we demonstrated the benefits and pitfalls of *Performative Autonomy*, a novel autonomy design strategy. Our results suggest that for robots with *deserved* confidence in their own decisions, *Performative Autonomy* may be a valuable strategy for promoting SA without decreasing Cognitive Load. Yet our results also suggest that for robots with *undeserved* confidence in their own decisions, *Performative Autonomy* may only be effective when underlying task-imposed MW is sufficiently high, and may itself impose additional cognitive and temporal demands in cases where the robot's faults become observable.

Overall, our results should encourage robot designers to employ *Performative Autonomy when deserved*. This autonomy design strategy may not be effective or worthwhile in domains that are not safety-critical: if errors or out-of-loop unfamiliarity do not come at high costs, the SA benefits from this strategy may not be worth the cost of interruption. Similarly, in domains where users do not struggle to maintain SA due to overly high (or overly low) levels of cognitive load, there may be little need to consider this type of strategy. But we argue, based on our results, that this approach could be a useful strategy across domains where these considerations do hold, such as search and rescue, autonomous driving, industrial robotics, robot-assisted surgery, and of course, space robotics.

These results should also encourage Human-Robot Interaction scientists to further explore this autonomy design strategy in future empirical work. A number of additional variables may be manipulated and measured in future work to develop a deeper understanding of the novel ideas presented in this work.

First, a wider range of strategies (corresponding to additional levels of Tab. 1) can be explored as independent variables in future work. In particular, we suggested that asking questions should encourage deeper engagement with a robot's situation, but we did not test this as a formal hypothesis. Comparison of the *Performative Autonomy* strategies explored in this work to mere Statements or Explanations would allow us to test this hypothesis. Future work could also consider the frequency of Performative Autonomy strategies. It would be natural to expect that as this strategy is used more frequently, SA benefits may increase, but so too may MW consequences. Similarly, while in this work, we emphasized imposed *Mental Workload*, future work could also manipulate other closely related dimensions of context, such as potential for harm and time pressure, due to the role these play in mediating how people choose whether and how to be polite [53].

Second, a wider array of measures may now be explored as dependent variables in future work. Most directly, a key limitation of this work is the abbreviated nature of our SA and Workload measurements. In this work we used simple Likert items and awareness questions so that we could repeatedly survey participants without overly distracting them. However, it may have been valuable to supplement these with infrequent yet more robust procedures, such as the NASA TLX [23], Bedford scale [47], or SAGAT procedures [16]. Future work should also examine the effects of Performative Autonomy on human-robot trust, and the way that trust might moderate or explain the findings observed in this paper. Of particular importance, we note that *Performed Low Autonomy* (WH-Questions) achieved the desired goal of providing an opportunity for interactants to "check in" on their robot teammates, thus promoting enhanced SA of the robot and the problems it may be encountering. Yet this dialogue strategy would *not* demonstrate the flaws in the robot's perception or reasoning when these occurred. In contrast, these flaws would be clear in the case of *Performed Medium Autonomy* (YN-Questions). While YN-Questions did "more harm than good" in our second experiment in terms of SA and Task Performance, it is possible that this strategy would have been beneficial in terms of promoting transparency and calibrated trust. This presents an intriguing design dichotomy worthy of future exploration.

7 CONCLUSION

We have presented a novel autonomy design strategy we term *Performative Autonomy*, in which robots behave *as if* they have a lower level of autonomy than they are truly capable of (i.e., asking for advice they do not believe they truly need), for the sole purpose of maintaining interactants' Situational Awareness. Our experimental assessment of this strategy suggest that for robots with *deserved* confidence in their own decisions, *Performative Autonomy* may be a valuable strategy for promoting Situation Awareness without decreasing Cognitive Load, but that for robots with *undeserved* confidence in their own decisions, *Performative Autonomy* may only be effective when underlying task-imposed Mental Workload is sufficiently high, and may itself impose additional cognitive and temporal demands when the robot's faults become observable. As we have discussed, our results thus provide strong support for the use of this autonomy design strategy in future safety-critical missions as humanity explores the Moon, Mars, and beyond, and open a host of new directions for future research in service of this goal.

ACKNOWLEDGMENTS

This work was supported by NASA Early Career Faculty award 80NSSC20K0070.

REFERENCES

[1] James E Allen, Curry I Guinn, and Eric Horvtz. 1999. Mixed-initiative interaction. *IEEE Intelligent Systems and their Applications* 14, 5 (1999), 14–23.

[2] Alan D Baddeley. 1972. Selective attention and performance in dangerous environments. *British journal of psychology* 63, 4 (1972), 537–546.

[3] Michael Baker and Holly A Yanco. 2004. Autonomy mode suggestions for improving human-robot interaction. In *2004 IEEE International Conference on Systems, Man and Cybernetics (IEEE Cat. No. 04CH37583),* Vol. 3. IEEE, 2948–2953.

[4] Adam Bogg, Stewart Birrell, Michael A Bromfield, and Andrew M Parkes. 2021. Can we talk? How a talking agent can improve human autonomy team performance. *Theoretical Issues in Ergonomics Science* 22, 4 (2021), 488–509.

[5] Adam Bogg, Andrew Parkes, and Mike Bromfield. 2020. Can we talk?–the impact of conversational interfaces on human autonomy teaming perception, performance and situation awareness. In *International Conference on Intelligent Human Systems Integration.* Springer, 938–944.

[6] Robert L. Bond. 1976. *Skylab experience bulletin No.26: The methods and importance of man-machine engineering evaluations in zero-g.* Technical Report. NASA-JSC, Houston, TX.

[7] Guy Andre Boy and Donald Platt. 2013. A situation awareness assistant for human deep space exploration. In *International Conference on Human-Computer Interaction.* Springer, 629–636.

[8] Penelope Brown, Stephen C Levinson, and Stephen C Levinson. 1987. *Politeness: Some universals in language usage.* Vol. 4. Cambridge university press.

[9] Benjamin Camblor, Jean-Marc Salotti, Charles Fage, and David Daney. 2022. Degraded situation awareness in a robotic workspace: accident report analysis. *Theoretical Issues in Ergonomics Science* 23, 1 (2022), 60–79.

[10] Erin K Chiou, Mustafa Demir, Verica Buchanan, Christopher C Corral, Mica R Endsley, Glenn J Lematta, Nancy J Cooke, and Nathan J McNeese. 2021. Towards human–robot teaming: tradeoffs of explanation-based communication strategies in a virtual search and rescue task. *International Journal of Social Robotics* (2021), 1–20.

[11] Mustafa Demir, Nathan J McNeese, and Nancy J Cooke. 2016. Team communication behaviors of the human-automation teaming. In *2016 IEEE international multi-disciplinary conference on cognitive methods in situation awareness and decision support (CogSIMA).* IEEE, 28–34.

[12] Michael C Dorneich, William Rogers, Stephen D Whitlow, and Robert DeMers. 2016. Human performance risks and benefits of adaptive systems on the flight deck. *The International Journal of Aviation Psychology* 26, 1-2 (2016), 15–35.

[13] Michael C Dorneich, Stephen D Whitlow, Santosh Mathan, Patricia May Ververs, Deniz Erdogmus, Andre Adami, Misha Pavel, and Tian Lan. 2007. Supporting real-time cognitive state classification on a mobile individual. *Journal of Cognitive Engineering and Decision Making* 1, 3 (2007), 240–270.

[14] Mary T Dzindolet, Scott A Peterson, Regina A Pomranky, Linda G Pierce, and Hall P Beck. 2003. The role of trust in automation reliance. *International journal of human-computer studies* 58, 6 (2003), 697–718.

[15] Mica R Endsley. 1988. Design and evaluation for situation awareness enhancement. In *Proceedings of the Human Factors Society annual meeting*, Vol. 32. Sage Publications Sage CA: Los Angeles, CA, 97–101.

[16] Mica R Endsley. 1988. Situation awareness global assessment technique (SAGAT). In *Proceedings of the IEEE 1988 national aerospace and electronics conference.* IEEE, 789–795.

[17] Mica R Endsley. 2017. From here to autonomy: lessons learned from human–automation research. *Human factors* 59, 1 (2017), 5–27.

[18] Mica R Endsley. 2021. Situation awareness. *Handbook of human factors and ergonomics* (2021), 434–455.

[19] Mica R Endsley, Betty Bolté, and Debra G Jones. 2003. *Designing for situation awareness: An approach to user-centered design.* CRC press.

[20] Mica R Endsley and Esin O Kiris. 1995. The out-of-the-loop performance problem and level of control in automation. *Human factors* 37, 2 (1995), 381–394.

[21] Terrence Fong, Charles Thorpe, and Charles Baur. 2001. *Collaborative control: A robot-centric model for vehicle teleoperation.* Vol. 1. Carnegie Mellon University, The Robotics Institute Pittsburgh.

[22] Rebecca Grier, Christopher Wickens, David Kaber, David Strayer, Deborah Boehm-Davis, J Gregory Trafton, and Mark St. John. 2008. The red-line of workload: Theory, research, and design. In *Proceedings of the human factors and ergonomics society annual meeting*, Vol. 52. Sage Publications Sage CA: Los Angeles, CA, 1204–1208.

[23] Sandra G Hart and Lowell E Staveland. 1988. Development of NASA-TLX (Task Load Index): Results of empirical and theoretical research. In *Advances in psychology.* Vol. 52. Elsevier, 139–183.

[24] Sarah K Hopko, Riya Khurana, Ranjana K Mehta, and Prabhakar R Pagilla. 2021. Effect of cognitive fatigue, operator sex, and robot assistance on task performance metrics, workload, and situation awareness in human-robot collaboration. *IEEE Robotics and Automation Letters* 6, 2 (2021), 3049–3056.

[25] David B Kaber and Mica R Endsley. 2004. The effects of level of automation and adaptive automation on human performance, situation awareness and workload in a dynamic control task. *Theoretical issues in ergonomics science* 5, 2 (2004), 113–153.

[26] David B Kaber, Emrah Onal, and Mica R Endsley. 2000. Design of automation for telerobots and the effect on performance, operator situation awareness, and subjective workload. *Human factors and ergonomics in manufacturing & service industries* 10, 4 (2000), 409–430.

[27] Amro Khasawneh, Hunter Rogers, Jeffery Bertrand, Kapil Chalil Madathil, and Anand Gramopadhye. 2019. Human adaptation to latency in teleoperated multi-robot human-agent search and rescue teams. *Automation in Construction* 99 (2019), 265–277.

[28] Boyoung Kim, Ruchen Wen, Qin Zhu, Tom Williams, and Elizabeth Phillips. 2021. Robots as moral advisors: The effects of deontological, virtue, and confucian role ethics on encouraging honest behavior. In *Companion of the 2021 ACM/IEEE International Conference on Human-Robot Interaction.* 10–18.

[29] Jeamin Koo, Jungsuk Kwac, Wendy Ju, Martin Steinert, Larry Leifer, and Clifford Nass. 2015. Why did my car just do that? Explaining semi-autonomous driving actions to improve driver understanding, trust, and performance. *International Journal on Interactive Design and Manufacturing (IJIDeM)* 9, 4 (2015), 269–275.

[30] Benoit Larochelle, Geert-Jan M Kruijff, Nanja Smets, Tina Mioch, and Peter Groenewegen. 2011. Establishing human situation awareness using a multi-modal operator control unit in an urban search & rescue human-robot team. In *IEEE Symposium on Human-Robot Interactive Communication (RO-MAN).*

[31] John Lee and Neville Moray. 1992. Trust, control strategies and allocation of function in human-machine systems. *Ergonomics* 35, 10 (1992), 1243–1270.

[32] Dan Lester. 2014. Putting Human Cognition and Awareness on Other Worlds: A Challenge for Human and Robotic Space Exploration. *Space Operations Communicator* (2014).

[33] Matthew Luebbers, Christine Chang, Aaqib Tabrez, Jordan Dixon, and Bradley Hayes. 2021. Emerging Autonomy Solutions for Human and Robotic Deep Space Exploration. (2021).

[34] Bertram F Malle, Steve Guglielmo, and Andrew E Monroe. 2014. A theory of blame. *Psychological Inquiry* 25, 2 (2014), 147–186.

[35] Nathan J McNeese, Mustafa Demir, Nancy J Cooke, and Christopher Myers. 2018. Teaming with a synthetic teammate: Insights into human-autonomy teaming. *Human factors* 60, 2 (2018), 262–273.

[36] Neville Moray. 1979. Models and measures of mental workload. In *Mental workload.* Springer, 13–21.

[37] Linda Onnasch, Christopher D Wickens, Huiyang Li, and Dietrich Manzey. 2014. Human performance consequences of stages and levels of automation: An integrated meta-analysis. *Human factors* 56, 3 (2014), 476–488.

[38] Raja Parasuraman, Mustapha Mouloua, and Robert Molloy. 1996. Effects of adaptive task allocation on monitoring of automated systems. *Human factors* 38, 4 (1996), 665–679.

[39] Raja Parasuraman and Victor Riley. 1997. Humans and automation: Use, misuse, disuse, abuse. *Human factors* 39, 2 (1997), 230–253.

[40] Jayam Patel and Carlo Pinciroli. 2020. Improving Human Performance Using Mixed Granularity of Control in Multi-Human Multi-Robot Interaction. In *2020 29th IEEE International Conference on Robot and Human Interactive Communication (RO-MAN).* 1135–1142. https://doi.org/10.1109/RO-MAN47096.2020.9223553

[41] Karen Petersen and Oskar Von Stryk. 2011. Towards a general communication concept for human supervision of autonomous robot teams. In *Proceedings of the fourth international conference on advances in computer-human interactions (ACHI).* Citeseer, 228–235.

[42] Matthew S Prewett, Ryan C Johnson, Kristin N Saboe, Linda R Elliott, and Michael D Coovert. 2010. Managing workload in human–robot interaction: A review of empirical studies. *Computers in Human Behavior* 26, 5 (2010), 840–856.

[43] Lawrence J Prinzel III, Alan T Pope, and Frederick G Freeman. 2001. *Application of physiological self-regulation and adaptive task allocation techniques for controlling operator hazardous states of awareness.* Technical Report.

[44] S. Rajendran, M. H. Nordin, S. Sharma, A. Khan, M. Gianni, and R. Sutton. 2022. Extended Abstract: Supervisory Intelligent Operator/Scheme for optimal shared control authority between human-vessel cooperation for increased autonomy. In *2022 UKACC 13th International Conference on Control (CONTROL).* 154–155. https://doi.org/10.1109/Control55989.2022.9781464

[45] Dale Richards and Alex Stedmon. 2016. To delegate or not to delegate: A review of control frameworks for autonomous cars. *Applied ergonomics* 53 (2016), 383–388.

[46] David A Robb, Francisco J Chiyah Garcia, Atanas Laskov, Xingkun Liu, Pedro Patron, and Helen Hastie. 2018. Keep me in the loop: Increasing operator situation awareness through a conversational multimodal interface. In *Proceedings of the 20th ACM International Conference on Multimodal Interaction.* 384–392.

[47] Alan H Roscoe and Georges A Ellis. 1990. *A subjective rating scale for assessing pilot workload in flight: A decade of practical use.* Technical Report. Royal Aerospace Establishment Farnborough (United Kingdom).

[48] Lindsay Sanneman and Julie A Shah. 2022. The Situation Awareness Framework for Explainable AI (SAFE-AI) and Human Factors Considerations for XAI Systems. *International Journal of Human–Computer Interaction* (2022), 1–17.

[49] John R Searle. 1975. Indirect speech acts. In *Speech acts.* Brill, 59–82.

[50] John R Searle. 1976. A classification of illocutionary acts1. *Language in society* 5, 1 (1976), 1–23.

[51] John R Searle and John Rogers Searle. 1969. *Speech acts: An essay in the philosophy of language*. Vol. 626. Cambridge university press.

[52] Thomas B Sheridan and William L Verplank. 1978. *Human and computer control of undersea teleoperators*. Technical Report. Massachusetts Inst of Tech Cambridge Man-Machine Systems Lab.

[53] Cailyn Smith, Charlotte Gorgemans, Ruchen Wen, Saad Elbeleidy, Sayanti Roy, and Tom Williams. 2022. Leveraging Intentional Factors and Task Context to Predict Linguistic Norm Adherence. In *Proceedings of the Annual Meeting of the Cognitive Science Society*, Vol. 44.

[54] Trey Smith, Jonathan Barlow, Maria Bualat, Terrence Fong, Christopher Provencher, Hugo Sanchez, and Ernest Smith. 2016. Astrobee: A new platform for free-flying robotics on the international space station. In *International Symposium on Artificial Intelligence, Robotics, and Automation in Space (i-SAIRAS)*.

[55] Dimitrios Stefanidis, Fikre Wang, James R Korndorffer, J Bruce Dunne, and Daniel J Scott. 2010. Robotic assistance improves intracorporeal suturing performance and safety in the operating room while decreasing operator workload. *Surgical endoscopy* 24, 2 (2010), 377–382.

[56] K Suzanne Barber, Anuj Goel, and Cheryl E Martin. 2000. Dynamic adaptive autonomy in multi-agent systems. *Journal of Experimental & Theoretical Artificial Intelligence* 12, 2 (2000), 129–147.

[57] Ning Wang, David V Pynadath, and Susan G Hill. 2016. Trust calibration within a human-robot team: Comparing automatically generated explanations. In *2016 11th ACM/IEEE International Conference on Human-Robot Interaction (HRI)*. IEEE, 109–116.

[58] Ruchen Wen, Boyoung Kim, Elizabeth Phillips, Qin Zhu, and Tom Williams. 2022. Comparing Norm-Based and Role-Based Strategies for Robot Communication of Role-Grounded Moral Norms. *ACM Transactions on Human-Robot Interaction (T-HRI)* (2022).

[59] Mihriban Whitmore, Jurine A Adolf, and Barbara J Woolford. 1999. Research Priorities for the International Space Station and Beyond. *Aviat Space Environ Med.* (1999).

[60] Christopher D Wickens. 2008. Multiple resources and mental workload. *Human factors* 50, 3 (2008), 449–455.

[61] Christopher D Wickens. 2020. Automation Lessons from Other Domains. *Handbook of Human Factors for Automated, Connected, and Intelligent Vehicles* (2020).

[62] Christopher D Wickens, William S Helton, Justin G Hollands, and Simon Banbury. 2021. *Engineering psychology and human performance*. Routledge.

[63] Tom Williams, Qin Zhu, Ruchen Wen, and Ewart J de Visser. 2020. The confucian matador: three defenses against the mechanical bull. In *Companion of the 2020 ACM/IEEE International Conference on Human-Robot Interaction*. 25–33.

[64] Holly A Yanco and Jill Drury. 2004. "Where Am I?" Acquiring situation awareness using a remote robot platform. In *2004 IEEE International Conference on Systems, Man and Cybernetics*, Vol. 3. IEEE, 2835–2840.

[65] Steven Yule, Jamie M Robertson, Benjamin Mormann, Douglas S Smink, Stuart Lipsitz, Egide Abahuje, Lauren Kennedy-Metz, Sandra Park, Christian Miccile, Charles N Pozner, et al. 2022. Crew Autonomy During Simulated Medical Event Management on Long Duration Space Exploration Missions. *Human Factors* (2022), 00187208211067575.

[66] Marco A Zenati, Lauren Kennedy-Metz, and Roger D Dias. 2020. Cognitive engineering to improve patient safety and outcomes in cardiothoracic surgery. In *Seminars in thoracic and cardiovascular surgery*, Vol. 32. Elsevier, 1–7.

[67] Qin Zhu, Tom Williams, Blake Jackson, and Ruchen Wen. 2020. Blame-laden moral rebukes and the morally competent robot: A Confucian ethical perspective. *Science and Engineering Ethics* 26, 5 (2020), 2511–2526.