Topic: Cognition

Know your cognitive environment!
Mental models as crucial determinant of offloading preferences

Patrick P. Weis & Eva Wiese
George Mason University, Fairfax, VA, USA

Précis: Here, we examine the impact of a human problem solver's mental models regarding fellow humans, robots, and smartphone apps on offloading cognitive processing onto these "helpers". We found that refining mental models is an easy and crucial approach to adjust offloading preferences and thus improve human problem solvers' interactions in cognitive environments.

Running head: Mental Modals and Offloading Preferences

Manuscript type: Research Article

Word count: 5956

Corresponding author:
Patrick Weis
Email: pweis@gmu.edu

The final version of this manuscript has been published in Human Factors at https://doi.org/10.1177/0018720820956861

ABSTRACT

Objective: Human problem solvers possess the ability to outsource parts of their mental processing onto cognitive “helpers” (*cognitive offloading*). However, suboptimal decisions regarding which helper to recruit for which task occur frequently. Here, we investigate if understanding and adjusting a specific subcomponent of mental models –beliefs about task-specific expertise – regarding these helpers could provide a comparatively easy way to improve offloading decisions. **Background:** Mental models afford storage of beliefs about a helper that can be retrieved when needed. **Methods:** Arithmetic and social problems were solved by 192 participants. Participants could – in addition to solving a task on their own – offload cognitive processing onto a human, a robot, or one of two smartphone apps. These helpers were introduced with either task-specific (e.g., stating that an app would use machine learning to “recognize faces” and “read emotions”) or task-unspecific (e.g., stating that an app was built for solving “complex cognitive tasks”) descriptions of their expertise. **Results:** Providing task-specific expertise information heavily altered offloading behavior for apps but much less so for humans or robots. This suggests 1) strong pre-existing mental models of human and robot helpers and 2) a strong impact of mental model adjustment for novel helpers like unfamiliar smartphone apps. **Conclusion**: Creating and refining mental models is an easy approach to adjust offloading preferences and thus improve interactions with cognitive environments. **Application:** To efficiently work in environments in which problem solving includes consulting other people or cognitive tools (“helpers”), accurate mental models –especially regarding task-relevant expertise– are a crucial prerequisite.

Keywords: Cognitive Offloading; Mental Models; Distributed Cognition; Extended Cognition; Metacognition; Strategy Selection

INTRODUCTION

Primer: cognitive environments

Technological advances related to computer hardware (e.g., the steady increase in processing power; Schaller, 1997), software and algorithms (e.g., modeling uncertainty in probabilistic programming; Ghahramani, 2015), and embodiment (e.g., the creation of intelligent virtual agents; Cassell et al., 2000; or improving the social component of robot agents; Wiese et al., 2017) contribute to a world with a plethora of opportunities to support our brain's limited abilities (i.e., *cognitive offloading*; Risko & Gilbert, 2016). These advances have the potential to "supersize our minds" (Clark, 2011). However, the continuously changing landscape of these opportunities also comes with a challenge: how do we decide which of the opportunities to take? When leaving for dinner with a friend, would we (1) navigate on our own or seek support by relying on (2) our friend's navigational ability, (3) a smartphone app, or (4) a robot companion? Current evidence suggests that we frequently make biased and suboptimal choices when seeking to support our brain (Gilbert et al., 2019; Jérémy Virgo et al., 2017; Risko & Dunn, 2015; Weis & Wiese, 2019a). Consequently, we are in need of interventions that inform unbiased choices (*cf*. Risko & Gilbert, 2016), which requires that we improve our understanding of the underlying decision mechanisms. The current manuscript caters to these needs by exploring how mental models about our fellow humans, smartphone apps, and embodied robots (i.e., *cognitive helpers*) influence offloading choice and how these models can be updated so as to readjust suboptimal choice behavior.

Primer: mental models of cognitive environments

A problem solver's mental model of a cognitive helper reflects "his or her beliefs about the [...] system, acquired either through observation, instruction or inference" (Norman, 2014, p. 12).

Mental models enable problem solvers to retrieve these beliefs from their memory such that the beliefs can subsequently guide interaction behavior like cognitive offloading (also known as *information-based metacognition*; for a review, see Koriat & Levy-Sadot, 2000). It should be noted that mental models have been crucial for seeking cognitive support long before the advent of human-computer-interaction: when interacting with fellow humans. For example, when asked to remember topic-specific information in concert with another person, social problem solvers will remember less information when they believe that the other person is an expert in the respective topic (Wegner, 1987). The dynamic way humans use mental models to distribute information across the minds of other group members (*transactive memory*; Wegner, 1987) has long been at the core of human society.

What has changed in recent decades, however, is the variety of non-human entities that can be accessed for such cognitive support. For example, humans can nowadays access internet-based rather than fellow-human-based information (Clowes, 2013; Wegner & Ward, 2013). In general, humans are increasingly inter-connected with computers that can enhance their cognitive abilities way beyond information-seeking (Clark, 2004, 2011) and consequently are in need of mechanisms to decide when to rely on computer-based processing. The straight-forward assumption that we adopt in the present paper is that this decision process can be informed by the same mental-model-based mechanism that holds when interacting with humans rather than computers. That beliefs are relevant for a human's decision to seek cognitive support is highly likely. For example, when a user's mental model of a calculator's CLEAR button includes beliefs that suggest low reliability, the user will press the button multiple times rather than only once (Norman, 2014). Similarly, beliefs about an input device's reliability have been shown to alter use frequency independently of actual reliability (Weis & Wiese, 2019a).

Current study: do mental models shape how cognitive environments are used?

In the present study, we therefore argue that, and investigate if, understanding and adjusting mental models of cognitive environments could provide a comparatively easy way to guide and improve cognitive support seeking (i.e., *cognitive offloading*) behavior [1]. What is known is that if helpful information regarding the cognitive environment is missing, it is likely that pre-existing mental models are accessed to guide offloading choice. For example, when asked to solve arithmetic and social problems, humans preferred to seek advice from computers and robots when solving arithmetic and advice from humans when solving social problems (Hertz & Wiese, 2019). Although not explicitly investigated in that study, we assume these task-specific preferences to have emerged due to stereotypical beliefs about the expertise of specific human and robotic entities that are part of an individual's mental model of the generic entity (e.g., "all humans are social beings", "all robots can rely on precise computers to calculate", etc.). In a similar vein, the way humans cognitively interact with other agents has been shown to depend on whether they believe that the agent possessed a mind (Wiese et al., 2012; Wykowska et al., 2014), which likely has extensive consequences for how humans structure their mental model of that agent.

To put the importance of mental models for cognitive support seeking to a test, we used a novel computer-based paradigm in which participants can either solve arithmetic or social problems on their own or offload it onto a human, a robot, or one of two smartphone applications. Note that novel smartphone applications are, just like robots, created by humans and, also just like robots, likely perceived superior to humans in analytical tasks (compare to Hertz & Wiese, 2019). However, they are not embodied, less present in the news, and usually more specialized in

[1] Please note that other factors like performance (Risko et al., 2014; Walsh & Anderson, 2009; Weis & Wiese, 2019b), effort (Ballard et al., 1997; Kool et al., 2010), or trust (de Visser et al., 2012, 2016) likely also influence cognitive interactions with humans, computers, and robots, but are addressed in the current paper only insofar as they might be mediated by an associated belief system (i.e., a mental model).

a specific domain (e.g., entertainment or finance) than their embodied counterparts. We therefore assume that our participants have little to none pre-existing mental models regarding novel smartphone apps.

Before engaging in the tasks, participants were to read short texts that were supposed to alter the participants' mental models of these cognitive helpers. The texts were – inspired by our interpretation of the data provided by Hertz and Wiese (2019) – supposed to specifically alter beliefs about task-specific expertise. In other words, we assume expertise beliefs to be a subcomponent of a mental model (of a cognitive helper) that has particularly high relevance for cognitive offloading choice. Therefore, we designed texts that could either provide task-unspecific (e.g.., the human is called "Michael" and studies English) or provide tasks-specific (e.g., the human is called "Michael" and is a social worker who is used to read emotions in people's faces on an everyday basis) information about the helpers' cognitive expertise.

Current study: hypotheses

H1-1. Based on the human advice seeking behavior reported by Hertz and Wiese (2019) and in absence of information about a cognitive helper's task-specific cognitive expertise, we assume that our participants' offloading preferences are based on pre-existing generic mental models of the cognitive helpers available in a particular situation. Thus, when familiar cognitive helper types like a human or an embodied robot are available, we assume our participants to make use of these generic mental models. Expertise beliefs stored in the generic model are then accessed and participants consequently prefer offloading arithmetic tasks to the robot and social tasks to the human even when no information about the cognitive helpers' expertise is provided.

H1-2. If that mechanism was true, providing specific expertise information that is consistent with pre-existing beliefs (i.e., that suggest arithmetic expertise for the robot and social expertise for the human) should hardly change these offloading preferences.

H2-1. Analogously, if pre-existing generic mental models do not differ, no differences in offloading preference should be exhibited. To test this hypothesis, we introduced two novel smartphone apps in a task-unspecific manner, observed offloading patterns for both arithmetic and social task, and expected no offloading preferences for any of the apps in either of the tasks.

H2-2. However, when presenting information that suggests differential task-specific expertise of both apps, clear offloading preferences should emerge again. In other words, we hypothesize that offloading preferences similar to the ones existing for humans and robots can be established for novel cognitive environments solely by adjusting the environment's mental model. Such a finding would suggest that human problem solvers use the same principles for deciding whether to offload cognition onto embodied agents like humans or robots, or onto non-embodied entities like smartphone apps.

Hypotheses have been preregistered. The preregistration can be accessed using the OSF repository associated with this manuscript (osf.io/s93tv)[2].

[2] Note that factor names and hypotheses are phrased differently in the present manuscript to improve readability. The factor "External Helpers" is now called "Environment", the factor "Metacognitive Priors" is now called "Mental Model". H1 has been split into H1-1 and H1-2 in the present manuscript, H2 has been split into H2-1 and H2-2.

METHODS & MATERIALS

Participants

In total, 323 participants were recruited via Amazon Mechanical Turk (www.mturk.com). Six participants were excluded because they took less than 10 or more than 45 minutes for a study that was designed to take 20 minutes. Additionally, 121 participants were excluded because they failed the manipulation check (for details on the manipulation check, see last paragraph of section *Procedure*) at the end of the study. We acknowledge that the exclusion rate is substantial but retained the manipulation check as exclusion criterion because it (1) was determined a priori and (2) is crucial that our participants did attend to and remembered the information given to them as this information constitutes our main manipulation (i.e., the Mental Model factor, see **Figure 1**) and we assume that some online participants do read texts only casually. Each participant that spent on average less than one second for each of the Perceived Competence ratings was also excluded (e.g., answered the question "How proficient do you think Michael is in solving the Dot task?" on a 21-point sliding scale in less than one second). This led to an additional exclusion of four participants, resulting in a final sample size of 192 participants (121 females, mean age: 40.1[3], age range: 21 – 75). The rigorous and extensive exclusion of participants was necessary to avoid biased results that underestimate the actual effects due to inattentiveness. All participants gave informed consent prior to participating. The study took on average about 20 minutes to complete and participants received $ 0.50 for their participation. This research complied with the tenets of the Declaration of Helsinki was approved by the Institutional Review Board at George Mason University. Informed consent was obtained from each participant prior to participation.

[3] Age was comparable between groups: Mean age was 39.1 years for the *Task-unspecific Agent and Task-specific App Expertise Beliefs* Mental Model group (for details on the factor, see section *Design*) and 41.3 years for the *Task-specific Agent and Task-unspecific App Expertise Beliefs* Mental Model group.

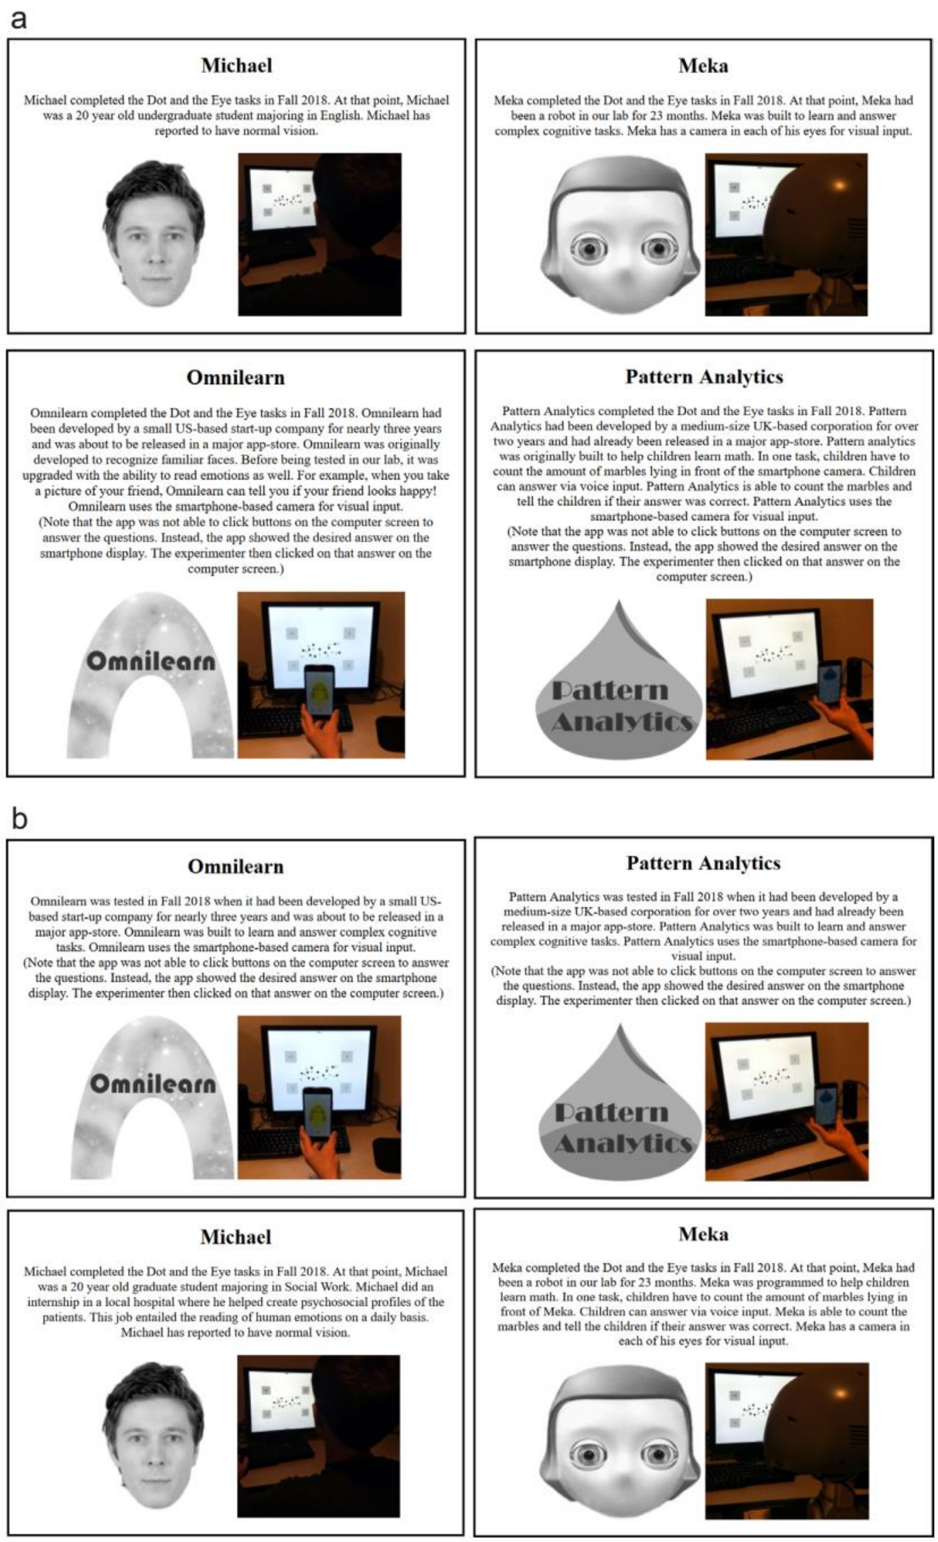

Figure 1: Instructions for the Mental Model Factor. Instructions as shown in the *Task-unspecific Agent Expertise Beliefs and Task-specific App Expertise Beliefs* (a) and the *Task-specific Agent Expertise Beliefs and Task-unspecific App Expertise Beliefs* (b) Mental Model conditions. Instructions are either suggesting task-unspecific cognitive skill or suggesting expertise specific to either arithmetic or social tasks; see *Design* for details.

Apparatus

Participants took the survey online on their own devices. The experiment was presented using the psychological testing software Inquisit (version 5; Millisecond Software, www.millisecond.com). Stimulus presentation scaled with the size of the participant's screen.

Stimuli

In total, 72 stimuli were used, 36 for the "eye task" and 36 for the "dot task" (see section *Tasks*). For the eye task, stimuli were extracted from the Reading the Mind in the Eyes test (Baron-Cohen et al., 2001). For the dot task, nine stimuli were custom-made using a common image editing software. All nine stimuli consisted of either nineteen or twenty dots and the following numeric difference between black and gray dots: -4, -3, -2, -1, 0, 1, 2, 3, 4. The remaining 27 dot task stimuli were created by mirroring the existing stimuli on the horizontal axis and then further mirroring both mirrored and original stimuli on the vertical axis. In addition, one unique practice stimulus was used for both eye and dot task that was not used in the main experiment. All stimuli can be accessed using the linked OSF repository.

Tasks

Similar to the paradigm used by Hertz and Wiese (2019), participants engaged in two tasks: an arithmetic ("dot task") and a social ("eye task") one. In the social task, participants were seeing pictures of human eyes and asked to "select which word best describes what the person in the picture is thinking or feeling" *(Reading the Mind in the Eyes Task*; Baron-Cohen et al., 2001). In the arithmetic task, participants were seeing black and gray dots and were to count and report the difference between the count of black and gray dots (for details on the dot stimuli, see *Stimuli*). Participants were asked to solve the tasks as accurately as possible. In both tasks, participants had six answer options. Participants could either choose to answer the question on their own (four options) or they could offload the cognitive task to one of two apps or agents (two options). Participants were instructed that all apps and agents that they can choose from had already been completing the eye and the dot task in our lab and that by clicking an app or an agent they would thereby chose the answer that the app or agent had given when solving the task in our lab.

For example, participants might see a stimulus with nine black dots and ten gray dots and could select to solve the task on their own by clicking one of the four numeric answer options (e.g., +1, 0, +2, and -1; see **Figure 2**, top row). Participants could also choose to offload the task to one of two agents instead of clicking one of the numeric answer options. For example, in the top row of **Figure 2**, participants were able to offload the task to the robot Meka (center top) or the human Michael (center bottom). In the figure, the participant chose to offload the task to Meka and is provided with the answer that Meka selected ("Meka chose for you: -1").

Design

Three factors were employed in the current study. First, participants engaged in two different tasks (Factor: Task Type, Levels: *Arithmetic*, *Social*). Second, participants were able to offload the task-related cognitive processing onto different entities. In one of two experimental blocks, participants were able to offload cognitive processing (Factor: Cognitive Environment, Levels: *Agents, Apps*) onto the human Michael and the robot Meka (i.e., level *Agents*). In the other experimental block, participants were able to offload cognitive processing onto a smartphone app called Omnilearn and another smartphone app called Pattern Analytics (i.e., level *Apps*). Third and lastly, participants had to read through different texts introducing the human, the robot, and the smartphone apps (Factor: Mental Model[4], Levels: *Task-unspecific Agent and Task-specific App Expertise Beliefs*, *Task-specific Agent and Task-unspecific App Expertise Beliefs;* **Figure 1**).

To establish or update mental models of apps and agents, each participant was provided with text-based information describing their cognitive abilities (i.e., Factor: Mental Model). Providing text-based information should be sufficient to establish or update mental models given that mental models represent "beliefs about [a] [...] system [that are] acquired either through observation, instruction or inference" (Norman, 2014, p. 12). Subsequently, participants are able to access the established mental models and recall the associated beliefs to guide their interactive behavior (also known as information-based metacognition; for a review, see Koriat & Levy-Sadot, 2000). Specifically, the provided information could either describe the helper as having task-unspecific or task-specific cognitive abilities in either the arithmetic or the social domain.

Whether the provided information described the helpers as having task-specific or more general (task-unspecific) cognitive abilities differed between blocks, and participants always engaged in one block with helpers that were described as having task-specific and one block with helpers that were described as having task-unspecific cognitive abilities. Which description type (task-specific or task-unspecific) was paired with which helpers (i.e., with "Agents" or "Apps") was randomly assigned and which helpers were available differed between blocks. Participants thus belonged to one of two Mental Model groups:

(1) *Task-unspecific Agent and Task-specific App Expertise Beliefs*: In the *Agents* Cognitive Environment block, the human was introduced as an undergrad majoring in English and the robot was introduced as being built for learning and answering complex cognitive

[4] We want to acknowledge a Reviewer's suggestion to name the factor "Information provided about the Helper" because it would more closely describe what we manipulated. In other words, a mental model cannot be *directly* manipulated but is only manipulated *via* the provided information. Although we appreciate the suggestion and think the proposed name would be more precise we decided to keep the current factor name because of its relative brevity and the theoretical framework associated with it.

tasks (task-unspecific cognitive abilities). In the *Apps* Cognitive Environment block, Omnilearn was introduced as an app built for recognizing familiar faces and reading emotions, and Pattern Analytics was introduced as an app built for helping children learn math in real-life surroundings by being able to count and provide feedback about the amount of marbles lying in front of the child (task-specific cognitive abilities). The exact wording can be inspected in **Figure 1a**.

(2) *Task-specific Agent and Task-unspecific App Expertise Beliefs*. In the *Agents* Cognitive Environment block, the human was introduced as an undergrad majoring in Social Work and is proficient in reading human emotions and the robot was introduced as being built for helping children learn math in real-life surroundings by being able to count and provide feedback about the amount of marbles lying in front of the child (task-specific cognitive abilities). In the *Apps* Cognitive Environment block, both Omnilearn and Pattern Analytics were introduced as apps built to learn and answer complex cognitive tasks. The exact wording can be inspected in **Figure 1b**.

Block order (i.e., whether the *Agents* or *Apps* Cognitive Environment was encountered first) was randomized.

Procedure

After clicking a link provided on MTurk, participants were to read a consent form. If a participant gave consent, general instructions concerning the two task types were given. One task required the participant to answer arithmetic questions; the other one required to answer social questions (for details, see section *Tasks*). Importantly, participants could either choose to answer the question on their own (four options) or they could offload the cognitive task to one of two apps or agents (two options). Participants then completed one practice trial for each task with four answer options, i.e. without the possibility to get cognitive support from a human, robot, or an app. Only then participants were introduced to the possibility to offload their cognitive processing to their *cognitive environment*, i.e. onto a human, a robot, or one of two apps. Participants then completed one trial for each task with only two answer options, a human and an app. A unique human and app that did not appear in the main experiment were used for that purpose. Right before the beginning of the main experiment, participants were explicitly instructed to "*Remember: Whenever you like, you can click on some of the humans, robots, or apps to choose the answer that they gave last fall! However, keep in mind that their answer is not necessarily correct and that your task is to score as many correct answers as possible*." Prototypical trials as well as timing details are provided in **Figure 2**. For details on the different apps and agents, see *Design*.

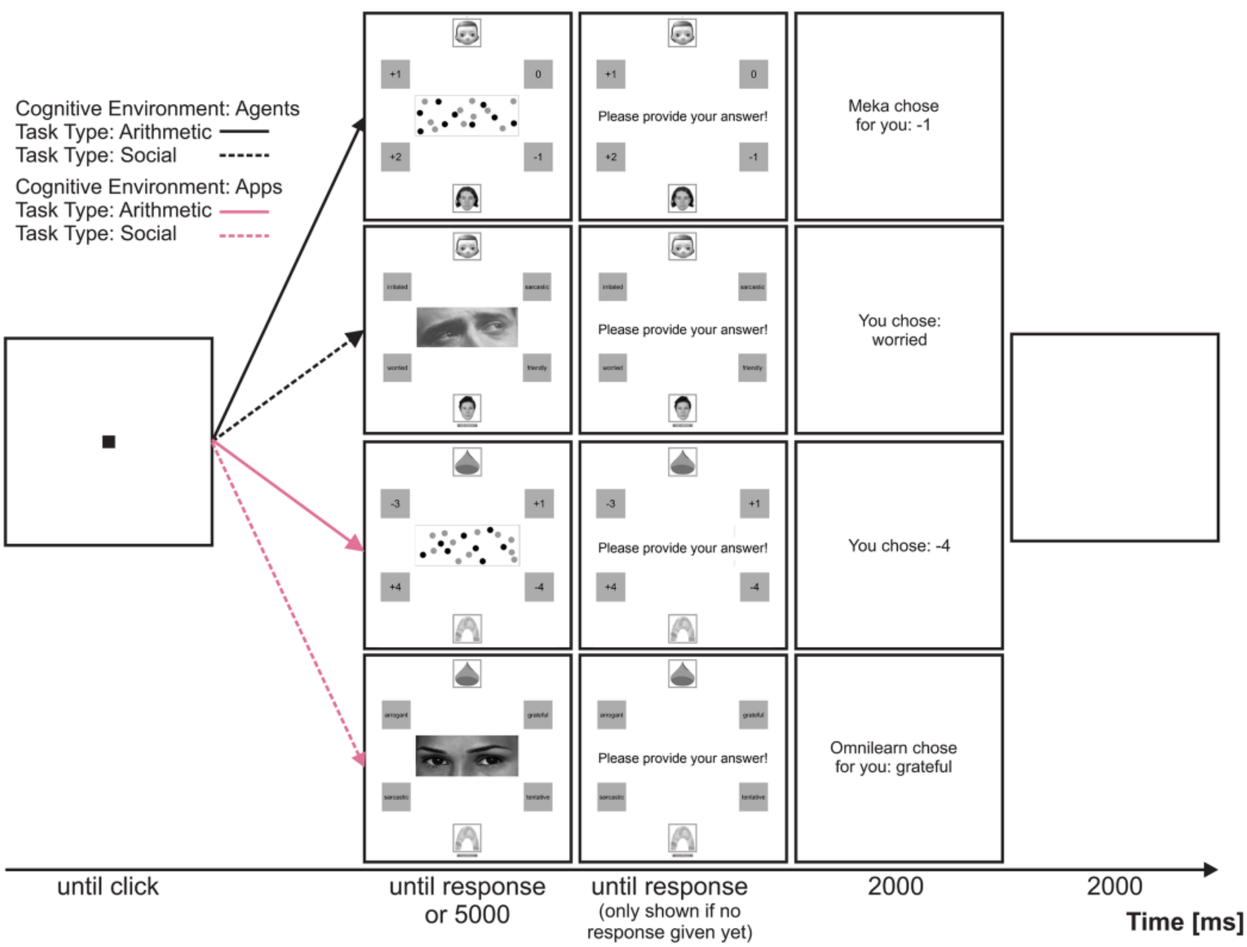

Figure 2: Trial Sequence. At the beginning of a trial, participants had to click a square to center the mouse cursor. After clicking the square, the task-related stimulus and the answer options were shown. If a participant took longer than five seconds to pick a response, the task-related stimulus disappeared. A five second window was chosen to keep response times roughly comparable between tasks and to provide a challenging experience that encourages the use of cognitive helpers. After choosing a response, feedback was provided for two seconds. Between trials, a blank screen was shown for two seconds. Stimuli and answer options are drawn to scale; other text is not drawn to scale. Note that answer options are provided in squares of equal size and that the centers of the squares are presented at equal distance to the center of the screen for all six answer options.

Participants then started one out of two experimental blocks. Both blocks consisted of the following. First, participants read a brief description of the two agents or apps that they could offload their cognitive processing to in the respective block (Mental Model manipulation; see **Figure 1**). Second, participants had to answer one question about each agent that ensured that they read and understood the instructions. For example, when asked "What is Michael trained in?", out of four answer options (Answering Complex Cognitive Tasks, English Language, Reading Emotions, Counting Objects), participants would have to select “English Language” if they read the instruction for Michael provided in **Figure 1a** and “Social Work” if they read the instruction for Michael provided in **Figure 1b**. If they answered at least one of both questions incorrectly, participants had to read the descriptions once more until they could provide correct answers to both questions. Third, participants were to rate the two apps’ or agents’ as well as their own abilities to perform the arithmetic and the social task on a 21-point sliding scale that closely resembled a visual analogue scale. Questions followed the following format: “How proficient do you think ‘Meka‘/‘Michael‘/‘Omnilearn‘/‘Pattern Analytics‘ is in solving the ‘Dot‘/‘Social‘ task?” The scale ranged from “Very Unproficient” on the left side to “Very Proficient” on the right side. Fourth, participants engaged in a total of 36 trials consisting of 18 arithmetic and 18 social trials (compare **Figure 2**). Trial order was randomized within the block and in the first block, problems were chosen randomly from the pool of 36 arithmetic and 36 social problems. At the end of the second block, each problem had been shown exactly once. At the end of the first block, participants were allowed to take a self-paced break.

After completing both experimental blocks, participants completed a brief demographic survey, rated all four agents and themselves once more in their abilities to complete the arithmetic and the social tasks, and completed a final manipulation check. For the manipulation check, participants were once more asked to select, out of four options, what each of the four agents and apps were trained in. This manipulation check allowed us to test whether participants retained the information provided in the agent and app descriptions (i.e., Mental Model manipulation; see **Figure 1**). Participants then were thanked for participating in the study and a unique code that participants were to enter on MTurk to receive payment was presented.

Measure: offloading preference

For the main analysis, offloading preference was used as dependent variable. Offloading preference is defined as the difference between how frequently a participant offloaded cognitive processing onto the human as compared to the robot in the Cognitive Environment *Agents* condition and onto Omnilearn as compared to Pattern Analytics in the Cognitive Environment *Apps* condition. Within each block, offloading preference can therefore range between -18 and 18. A value of -18 means that participants offloaded the task exclusively onto the robot (in the Cognitive Environment *Agents)* or the Pattern Analytics app (in the Cognitive Environment *Apps*) condition.

Analyses

As an omnibus test, we employed a 2 x 2 x 2 ANOVA with the within-participants factors Cognitive Environment and Task Type and the between-participants factor Mental Model. To test our specific hypotheses (see *Introduction: Hypotheses*), *t*-tests were employed. For details about the *t*-tests, see *Results: Hypotheses-driven analyses*.

RESULTS[5]

To provide an overview over our participants' problem solving behaviors, full data on how frequently participants chose to rely on their own cognitive processing and how frequently they chose to rely on the human, the robot, or on one of the smartphone apps, is depicted in **Figure 3**. In the following, hypotheses-driven and explorative statistical analyses are reported.

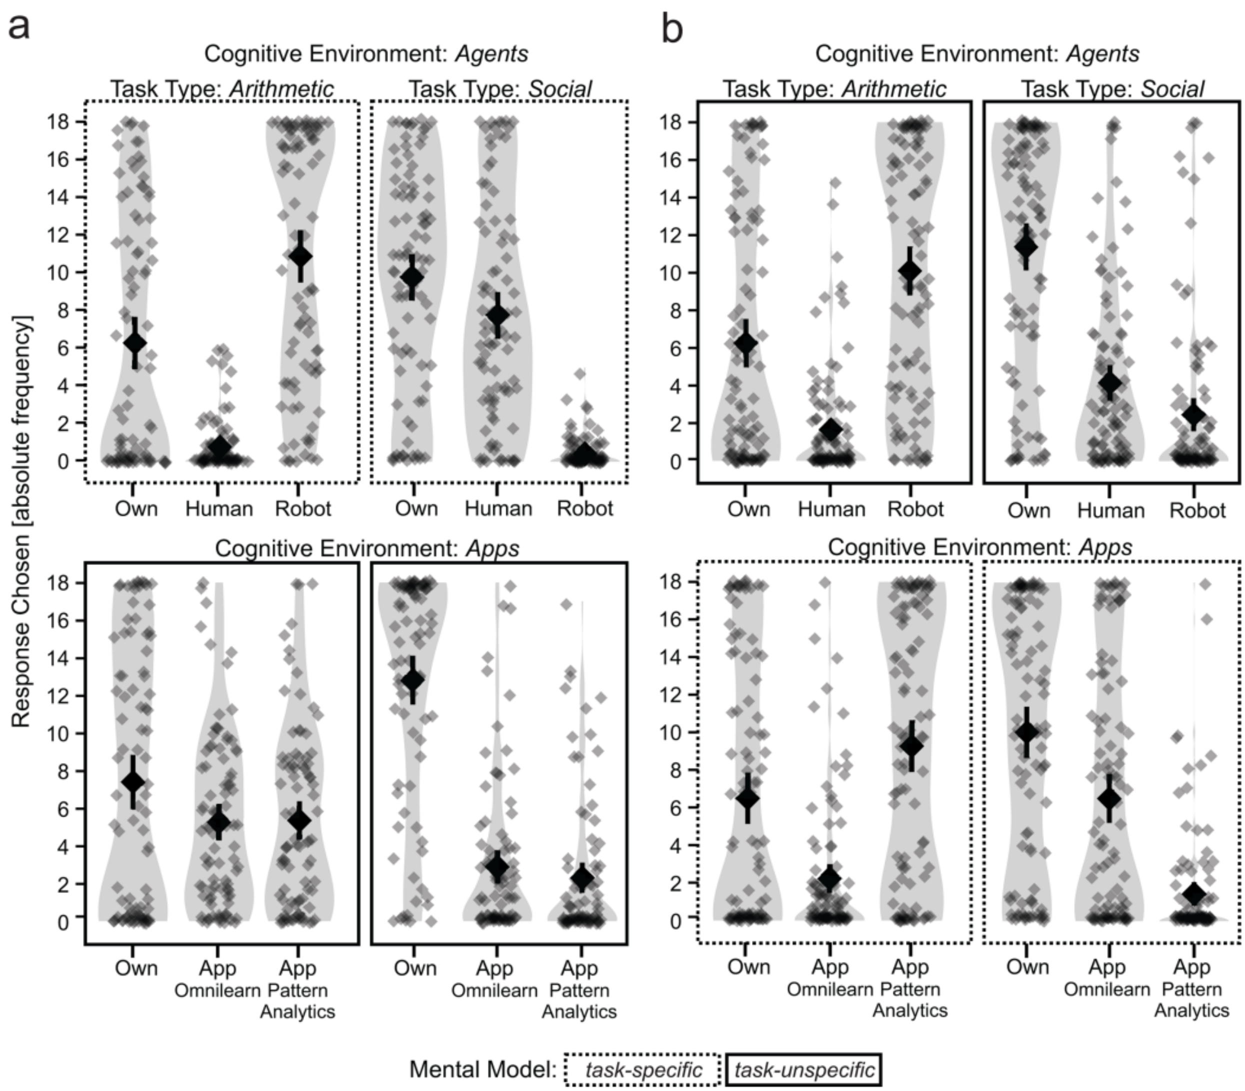

[5] The associated R analysis script and data files can be freely accessed online through the Open Science Framework at osf.io/s93tv.

Figure 3: Responses. Response counts for the *Task-specific Agent and Task-unspecific App Expertise Beliefs* (a) and the *Task-unspecific Agent and Task-specific App Expertise Beliefs* (b) Mental Modal conditions. Each box summarizes data of the 18 trials per participant in the respective condition. The x-axis specifies whether participants solved the task on their own or chose to offload to the available apps or agents. Response counts can thus range from 0 (response chosen in 0% of trials) to 18 (chosen in 100% of trials) for each answer option and sum up to 18 within each box. Black diamonds represent means. Error bars represent 95% confidence intervals. Gray diamonds represent raw data points. Gray shapes represent violin plots as implemented by ggplot2 (Wickham, 2016). The numeric values depicted in this plot can be inspected in Table S1 in the Supplemental Material.

Hypotheses-driven analyses

In the following, results of the omnibus ANOVA are reported to allow the reader to inspect participants' response patterns and to deduce whether the hypothesis-driven *t*-tests are backed by significant interactions in the data set as a whole. Subsequently, hypotheses-driven analyses are reported.

(1) *Omnibus ANOVA*. Confirming our expectations, the omnibus test indicated that offloading preference was altered as a function of the three-way-interaction between Mental Model, Cognitive Environment, and Task Type ($F(1, 190) = 69.4$, $p < .0001$, $\eta_G^2 = .10$; **Figure 4**). Recall that offloading preference is defined as the difference between how frequently the human agent was chosen in comparison to the robot agent (in the Cognitive Environment *Agents* condition) or how frequently Omnilearn was chosen in comparison to Pattern Analytics (in the Cognitive Environment *Apps* condition). In addition, Task Type and Cognitive Environment interacted in their influence on offloading difference ($F(1, 190) = 43.4$, $p < .0001$, $\eta_G^2 = .06$) whereas the interaction effects of Mental Model and Cognitive Environment ($F(1, 190) = 1.0$, $p = .3077$, $\eta_G^2 < .01$) and Mental Model and Task Type ($F(1, 190) = 2.1$, $p = .1482$, $\eta_G^2 < .01$) were not significant at a .05 alpha level. All three main effects, i.e. Cognitive Environment ($F(1, 190) = 30.0$, $p < .0001$, $\eta_G^2 = .02$), Mental Model ($F(1, 190) = 16.0$, $p < .0001$, $\eta_G^2 = .01$), and Task Type ($F(1, 190) =$

211.2, p < .0001, $\eta_G^2 = .33$) were significant. The ANOVA results suggest that human problem solvers prefer specific environments (i.e., specific apps, humans, robots) for solving specific tasks (i.e., arithmetic or social tasks) and that updating a mental model with task-specific information has a different effect for different environments.

(2) *Hypotheses H1-1 and H1-2*. Specifically, in the *Agents* Cognitive Environment, participants changed their offloading preferences based on the Task Type for both Mental Model conditions: Participants showed a higher preference for the human agent for the *Social* in comparison to the *Arithmetic* Task Type for both the *Task-specific Agent and Task-unspecific App Expertise Beliefs* ($t(88) = 13.9$, $p < .0001$; $M_{Social} - M_{Arithmetic} = 17.5$; in line with *H1-1*) and the *Task-unspecific Agent and Task-specific App Expertise Beliefs* ($t(102) = 9.00$, $p < .0001$; $M_{Social} - M_{Arithmetic} = 10.1$; in line with *H1-1*) Mental Model conditions. Mental Model did not alter offloading preferences for the *Arithmetic* Task Type ($t(190) = 1.56$, $p = .1202$; $M_{Task\text{-}specific\ Agent\ and\ Task\text{-}unspecific\ App\ Expertise\ Beliefs} = -10.2$, $M_{Task\text{-}unspecific\ Agent\ and\ Task\text{-}specific\ App\ Expertise\ Beliefs} = -8.4$; in line with *H1-2*) in the *Agents* Cognitive Environment condition. Mental Model however did alter offloading preferences for the *Social* Task Type ($t(190) = 6.01$, $p < .0001$; $M_{Task\text{-}specific\ Agent\ and\ Task\text{-}unspecific\ App\ Expertise\ Beliefs} = 7.4$, $M_{Task\text{-}unspecific\ Agent\ and\ Task\text{-}specific\ App\ Expertise\ Beliefs} = 1.7$; contradicting *H1-2*). In sum, in the *Agents* Cognitive Environment, our human problem solvers showed task-specific offloading preferences for different agents; **Figure 4a**. In alignment with H1-1, these offloading preferences existed even when only task-unspecific metacognitive information was provided. In alignment with H1-2, providing information describing the human as highly capable of reading emotions and the robot as highly capable of object recognition and object counting was not able to alter our problem solvers' offloading preferences in the arithmetic task. Unexpectedly and not aligned with H1-2 however, the ascription of social ability

to the human was able to change offloading preferences. H1-2 is therefore only partially confirmed.

(3) *Hypotheses H2-1 and H2-2*. In the *Apps* Cognitive Environment on the other hand, participants changed their offloading preferences based on the Task Type only in the *Task-unspecific Agent and Task-specific App Expertise Beliefs* Mental Model condition: Participants showed a higher preference for Omnilearn for the *Social* in comparison to the *Arithmetic* Task Type for the *Task-unspecific Agent and Task-specific App Expertise Beliefs* ($t(102) = 7.74$, $p < .0001$; $M_{Social} - M_{Arithmetic} = 12.1$; in line with *H2-2*) but not for the *Task-specific Agent and Task-unspecific App Expertise Beliefs* ($t(88) = .73$, $p = .47$; $M_{Social} - M_{Arithmetic} = .7$; in line with *H2-1*) Mental Model condition. Mental Model in the *Apps* Cognitive Environment altered offloading preferences for both the *Arithmetic* ($t(190) = 5.97$, $p < .0001$; $M_{Task\text{-}specific\ Agent\ and\ Task\text{-}unspecific\ App\ Expertise\ Beliefs} = -.1$, $M_{Task\text{-}unspecific\ Agent\ and\ Task\text{-}specific\ App\ Expertise\ Beliefs} = -7.0$; in line with *H2-2*) and the *Social* ($t(190) = 4.60$, $p < .0001$; $M_{Task\text{-}specific\ Agent\ and\ Task\text{-}unspecific\ App\ Expertise\ Beliefs} = .6$, $M_{Task\text{-}unspecific\ Agent\ and\ Task\text{-}specific\ App\ Expertise\ Beliefs} = 5.1$; in line with *H2-2*) Tasks. In sum, results for the *Apps* Cognitive Environment condition show that our participants had no prior task-related offloading preferences for the Omnilearn or the Pattern Analytics app, thus confirming H2-1. Results also show that updating a mental model with task-specific information is sufficient to establish strong offloading preferences, thus confirming H2-2. Results for the *Apps* Cognitive Environment condition are depicted in **Figure 4b**. Providing task-specific metacognitive information about a cognitive environment can thus outmatch the relevance of pre-existing mental models. In particular, participants in the *Task-unspecific Agent and Task-specific App Expertise Beliefs* Mental Model condition showed more extreme offloading preferences for apps than for agents in the *Social* ($t(102) = 3.64$, $p = .0004$; $M_{\text{Apps-Agents}} = 3.38$) and similar

offloading preferences in the *Arithmetic* ($t(102) = 1.50$, $p = .1366$; $M_{\text{Apps-Agents}} = 1.41$) Task Type.

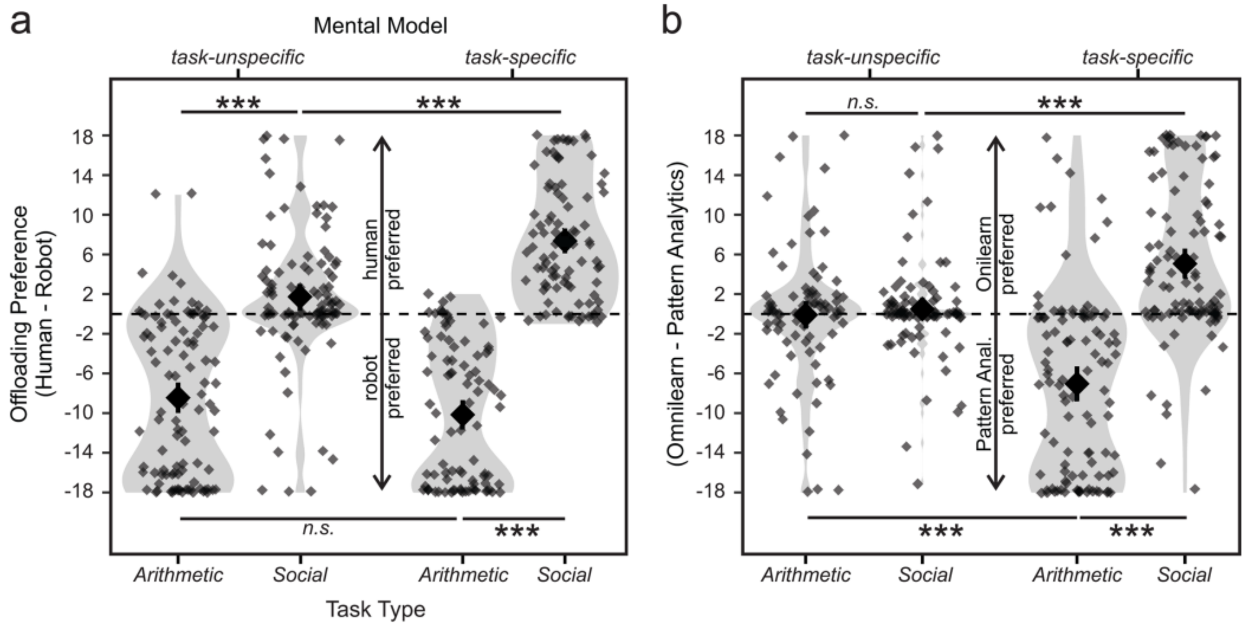

Figure 4: Offloading Preferences. Offloading preferences, as measured in absolute frequencies, for the *Agents* (a) and the *Apps* (b) Cognitive Environment conditions. Note that the *Task-unspecific Agent and Task-specific App Expertise Beliefs* Mental Model condition comprises the left half of (a) and the right half of (b) while the *Task-specific Agent and Task-unspecific App Expertise Beliefs* Mental Model condition comprises the right half of (a) and the left half of (b). An individual's preference scores can range from -18 to +18 for each permutation of Task Type and Cognitive Environment. Black diamonds represent means. Error bars represent 95% confidence intervals. Gray diamonds represent raw individual data points. Gray shape represents the distribution of the raw data as implemented by ggplot2's geom_violin function (Wickham, 2016). *** $p < .0001$; *n.s.* $p > .1$

Exploratory analyses

In addition to the hypothesis-driven analyses, we explored whether the hypothesized effect of Cognitive Environment and Task Type on offloading preference is mediated by task-specific perceived competence of the respectively available humans, robots, or smartphone apps. Mediation would suggest perceived competence to be a crucial property of a cognitive environment. It would furthermore suggest that the Mental Model manipulation induced consciously accessible competence beliefs that are a source of the offloading preference. Specifically, we ran two multilevel mediation models – one for each Mental Model level. For details on the Bayesian parameter estimation, consult Vuorre and Bolger (2018). As we expected similar and substantial mediation for both Mental Model levels, running two models allowed for cross-validation of the parameters.

Model results showed that for *Task-specific Agent and Task-unspecific App Expertise Beliefs*, none of the bootstrapped 95% confidence intervals of path a, b, c, or c' included 0, which sets the stage for mediation tests. Mediation tests revealed that both the indirect effect (M = .28, 95% CI = [.10 .48]) as well as the percentage mediated (M = .37, 95% CI = [.13 .62]) were significantly greater than zero. Analogously, for *Task-unspecific Agent and Task-specific App Expertise Beliefs*, none of the bootstrapped 95% confidence intervals of path a, b, c, or c' included 0. Both the indirect effect (M = .16, 95% CI = [.05 .28]) as well as the percentage mediated (M = .24, 95% CI = [.07 .42]) were significantly greater than zero. Results of both mediation models suggest partial mediation. Task-specific competence ratings therefore seem to be a relevant part of a human problem solver's mental model of an agent or an app. All mediation model parameter estimates are depicted in **Figure 5**. More details regarding the statistical procedure as well as parameter estimates are provided in the Supplemental Material. Mean rating data is provided in **Figure S1**[6].

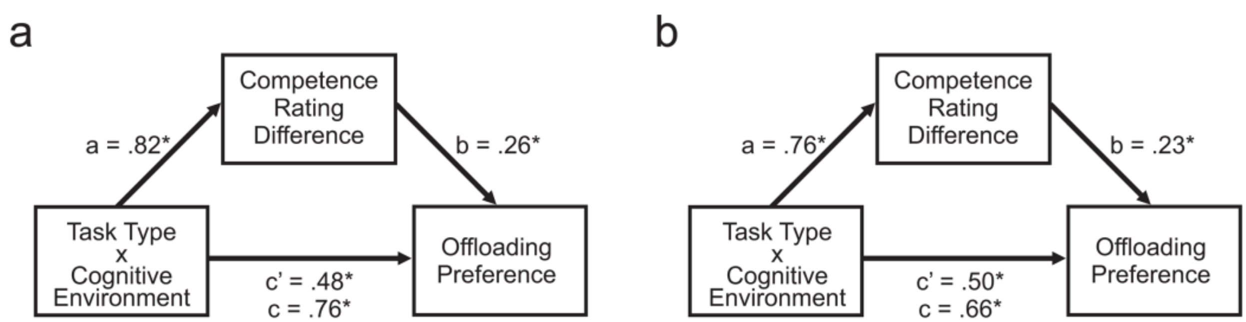

Figure 5: Mediation Models. Standardized Bayesian multilevel mediation model estimates for the *Task-specific Agent and Task-unspecific App Expertise Beliefs* (a) and the *Task-unspecific Agent and Task-specific App Expertise Beliefs* (b) Mental Model conditions. Both models suggest partial mediation (see text). *: 95% CI does not include 0

[6] Note that, for exploratory purposes, we obtained perceived competence ratings before and after participants engaged in the task. However, also note that, as indicated by the means, both ratings seem to be highly correlated.

DISCUSSION

The current paper investigated the four hypotheses H1-1, H1-2, H2-1, and H2-2 (1-4) regarding the influence of mental models – specifically, beliefs about task-specific expertise – on cognitive offloading. Hypothesis testing was complemented by exploratory mediation analyses (5).

(1) *H1-1*. We confirmed in the offloading domain what previous research has already shown in the advice seeking domain (Hertz & Wiese, 2019): Human problem solvers seem to have pronounced pre-existing beliefs regarding human and robotic agents that inform their decision whether to offload cognitive tasks to human or robotic agents (confirmation of H1-1).

(2) *H1-2*. When adding to these pre-existing beliefs by providing information about task-specific competencies, we found that introducing a robot as proficient in arithmetic and a human as proficient in social tasks did not alter offloading preferences in the arithmetic task (partial confirmation of H1-2). We argue that offloading preferences did not change because our participants' pre-existing generic beliefs have already been in congruence with the description of task-specific arithmetic expertise before the description was presented (compare to **Figure S1ab**, first graph from the left, for associated perceived competence ratings). However, offloading preferences after providing task-specific information did change for the social task (partial rejection of H1-2). The description's impact on offloading preferences was likely due to the fact that – contradicting our expectations – pre-existing competence beliefs have not been in congruence with the task-specific social expertise suggested in the description (compare **Figure S1ab**, third graph from the left, for associated perceived competence ratings). Thus, our participants' pre-existing mental models contained beliefs ascribing high arithmetic proficiency to the robot but surprisingly only suboptimal social proficiency to the human used in the present study. It should be

noted that these results might not generalize to all human stimuli. For example, we only used a male human stimulus image and males are known to score lower on social skill measures than females (Petrides & Furnham, 2000), which makes it questionable whether initial social proficiency ratings would have been as low as in the present study if a female was used as the human agent instead[7].

(3) *H2-1*. We unsurprisingly found no task-specific offloading preferences for novel smartphone apps when introducing both apps in a task-unspecific manner (confirmation of H2-1). We argue that is because our participants' mental models regarding the smartphone apps did not contain differential task-relevant beliefs. The finding thus supports the relevance of mental models for cognitive offloading and sets the stage for H2-2.

(4) *H2-2*. We found that providing task-relevant information about the smartphone apps was sufficient to induce substantial offloading preferences (confirmation of H2-2). These preferences were of comparable magnitude to the preferences for humans and robots. Thus, providing task-relevant information about novel cognitive tools like smartphone apps can be sufficient to induce offloading preferences that are as strong as prior beliefs humans have about embodied agents like humans and robots.

(5) *Exploratory analyses*. Lastly, when conducting follow-up explorative analyses, we found that offloading preference was partially mediated by competence ratings, suggesting an at least partially information-based (Koriat & Levy-Sadot, 2000) decision process that further highlights the importance of mental models for cognitive offloading. In other words, providing information about a cognitive helper's task-specific expertise can update our

[7] We thank the anonymous reviewer who made us aware of this issue.

mental model of this helper. The updated model will subsequently provide consciously accessible competence beliefs that can inform offloading preferences.

The present results provide evidence for how substantially mental models regarding fellow humans but also evolutionary novel cognitive partners like robots or smartphone apps can influence cognitive offloading preferences. We argue that refining mental models is an easy and crucial approach to adjust offloading preferences and thus to improve our cognitive interactions with our social or tech-infused environments. To realize the potential benefit of such refinement, it is crucial to note that establishing valid and accurate mental models does not necessarily occur automatically. For example, it is known that the elderly frequently underrate their mnemonic abilities which leads to an overreliance on external memory aids (Touron, 2015). Similarly, it has been shown that false beliefs about the reliability of a specific human-computer-interface can have prolonged maladaptive effects on offloading preferences (Weis & Wiese, 2019a).

The present results suggest a general mechanism for learning how to cognitively interact with our environment that holds for embodied (e.g., human, robot) and non-embodied (e.g., smartphone apps) helpers with varying degrees of social features alike: establishing and refining mental models. This establishing-mental-models-mechanism is well compatible with a view that emphasizes human technical reasoning skills when engaging in cognitive interactions (Osiurak & Reynaud, 2019). Such technical reasoning (here: inferring a cognitive helper's task expertise from an introductory text and pre-existing beliefs) is largely independent of social components of the interaction (*social learning*; e.g., Laland, 2004) or whether the cognitive interaction "partner" is assumed to possess a mind (top-down social cognition can heavily impact cognitive interactions with the environment; Wiese et al., 2012). Thus, while social learning (e.g., copying others) and social cognition (e.g., gaze following) can provide feasible means for human problem solvers to establish novel tool use behavior, asocial mechanisms based on technical reasoning seem to be equally feasible.

Several issues should be kept in mind when interpreting the present results. First, we want to emphasize that mental models can only partially explain how human problem solvers establish offloading preferences. For example, it has been shown that one's beliefs about the own prospective memory ability and actual ability are distinct from each other and have separable effects on offloading preferences (Gilbert, 2015). Accordingly, the moderate relationship between perceived competence and offloading preferences found in this study does leave room for additional explanations. In principle, the moderate relationship could also be due to methodological issues like a poor validity of our perceived competence measure. We however deem this possibility unlikely given the strong correlation with the "Task Type x Cognitive Environment" manipulation. Second, it should be noted that the mediation analysis only captured one aspect of the mental models: beliefs about the cognitive helpers' competence/expertise. It might well be that the metacognitive information we provided led to beliefs that are not directly related to competence and still affected offloading preferences. For example, we might have unwillingly established beliefs about how trustworthy or likable an entity is. In the case of trust, it has been shown that humans, robots, and non-embodied computers can receive similar pre-task trust ratings (de Visser et al., 2012). However, trust has been shown to be more stable for human than non-human cognitive helpers (de Visser et al., 2012), which might have in turn affected offloading preference over the course of the present study. Further complexity is added by the fact that individual differences regarding trust towards machines (e.g., Merritt & Ilgen, 2008) and towards own cognitive functioning (e.g., Touron, 2015) are likely to factor in as well. Note that task-specific trust towards own cognitive functioning can possibly be inferred from the perceived competence ratings shown in the Supplemental Material (**Figure S1**) but that domain-general cognitive functioning has not been measured in the present study and is known to influence offloading preference as well (Gilbert, 2015). Third, in the present paradigm, participants were continuously confronted with two helpers, a situation that might deviate from everyday problem solving and obscure absolute offloading rates. Relatedly, the discrete depictions as well as the novelty of the human helper, the robot helper, and the smartphone applications, might have further influenced absolute offloading rates, which should be considered when interpreting absolute offloading rates. However, note that the present analyses were focused on relative offloading differences between helpers, a measure that should not substantially be influenced by helper availability or novelty. Fourth, agent and app description (as provided in **Figure 1**) lengths differed between Mental Model conditions. Although description length was comparable within Cognitive Environment conditions and our main DV (i.e., relative offloading preference) is thus not impacted, comparisons of absolute offloading preference between Mental Model conditions as depicted in **Figure 2** might be confounded. However, we are not aware of any theory that would suggest this potential confound to be substantial.

One other highly interesting potential predictor of offloading preference that was not captured in the present study is experience-based (i.e., gut-feeling-based) rather than only information-based (i.e., based on memory retrieval; compare Koriat & Levy-Sadot, 2000) processing[8]. Consequently, our participants' decisions to offload to a specific agent or app could have been due a gut feeling response that was developed when reading the agent and app descriptions rather than due to recalling the respective description (i.e., **Figure 1**). The unexplained variance in the present mediation results would provide enough room for such a possibility. In general, it is well estab-

[8] The difference is nicely illustrated by Koriat and Levy-Sadot (2000) on p. 194: “A person who does not like tuna fish may feel some repulsion toward a salad offered in a buffet when she learns that it contains tuna fish. Her choice to avoid the salad may then be based on the explicit information gained (information-based action) or on the immediate repulsive feeling (experience-based action)”.

lished that some characteristics that inform strategy selection processes might not be consciously accessible (Cary & Reder, 2002). Such unconscious processes are also in line with the finding that belief manipulations can influence offloading preferences without changing subjective ratings of the cognitive environment’s usefulness (Weis & Wiese, 2019a).

KEY POINTS

- Naive human problem solvers possess mental models that encompass beliefs about task-specific expertise of human and robot agents
- These pre-existing mental models are reflected by how willing human problem solvers are to make use of such agents to help them solve specific cognitive tasks
- Accordingly, when confronted with two similar and novel cognitive tools like smartphone apps, humans are indifferent about which one to use
- However, providing a paragraph describing each app’s task-specific capabilities is enough to update the mental model and create as much behavioral relevance as the strong pre-existing mental models that are in place for human and robotic agents do
- We argue that creating or refining mental models (specifically, beliefs about expertise) is an easy and crucial approach to adjust offloading preferences and thus improve human problem solvers’ interactions in cognitive environments

REFERENCES

Ballard, D. H., Hayhoe, M. M., Pook, P. K., & Rao, R. P. (1997). Deictic codes for the embodiment of cognition. *Behavioral and Brain Sciences*, *20*(4), 723–742.

Baron-Cohen, S., Wheelwright, S., Hill, J., Raste, Y., & Plumb, I. (2001). The “Reading the Mind in the Eyes” Test revised version: A study with normal adults, and adults with Asperger syndrome or high-functioning autism. *The Journal of Child Psychology and Psychiatry and Allied Disciplines*, *42*(2), 241–251.

Cary, M., & Reder, L. M. (2002). Metacognition in strategy selection. In *Metacognition* (pp. 63–77). Springer.

Cassell, J., Sullivan, J., Churchill, E., & Prevost, S. (2000). *Embodied conversational agents*. MIT press.

Clark, A. (2004). *Natural-born cyborgs: Minds, technologies, and the future of human intelligence*. Oxford University Press.

Clark, A. (2011). *Supersizing the mind: Embodiment, action, and cognitive extension: Embodiment, action, and cognitive extension*. Oxford University Press.

Clowes, R. W. (2013). The Cognitive Integration of E-Memory. *Review of Philosophy and Psychology*, *4*(1), 107–133. https://doi.org/10.1007/s13164-013-0130-y

de Visser, E. J., Krueger, F., McKnight, P., Scheid, S., Smith, M., Chalk, S., & Parasuraman, R. (2012). The World is not Enough: Trust in *Cognitive* Agents. *Proceedings of the Human Factors and Ergonomics Society Annual Meeting*, *56*(1), 263–267. https://doi.org/10.1177/1071181312561062

de Visser, E. J., Monfort, S. S., McKendrick, R., Smith, M. A. B., McKnight, P. E., Krueger, F., & Parasuraman, R. (2016). Almost human: Anthropomorphism increases trust resilience in cognitive agents. *Journal of Experimental Psychology: Applied*, *22*(3), 331–349. https://doi.org/10.1037/xap0000092

Ghahramani, Z. (2015). Probabilistic machine learning and artificial intelligence. *Nature*, *521*(7553), 452–459. https://doi.org/10.1038/nature14541

Gilbert, S. J. (2015). Strategic use of reminders: Influence of both domain-general and task-specific metacognitive confidence, independent of objective memory ability. *Consciousness and Cognition*, *33*, 245–260. https://doi.org/10.1016/j.concog.2015.01.006

Gilbert, S. J., Bird, A., Carpenter, J. M., Fleming, S. M., Sachdeva, C., & Tsai, P.-C. (2019). Optimal use of reminders: Metacognition, effort, and cognitive offloading. *Journal of Experimental Psychology: General*.

Hertz, N., & Wiese, E. (2019). Good advice is beyond all price, but what if it comes from a machine? *Journal of Experimental Psychology. Applied*.

Jérémy Virgo, Jonathan Pillon, Jordan Navarro, Emanuelle Reynaud, & François Osiurak. (2017). Are You Sure You're Faster When Using a Cognitive Tool? *The American Journal of Psychology*, *130*(4), 493. https://doi.org/10.5406/amerjpsyc.130.4.0493

Kool, W., McGuire, J. T., Rosen, Z. B., & Botvinick, M. M. (2010). Decision making and the avoidance of cognitive demand. *Journal of Experimental Psychology: General*, *139*(4), 665–682. https://doi.org/10.1037/a0020198

Koriat, A., & Levy-Sadot, R. (2000). Conscious and Unconscious Metacognition: A Rejoinder. *Consciousness and Cognition*, *9*(2), 193–202. https://doi.org/10.1006/ccog.2000.0436

Laland, K. N. (2004). Social learning strategies. *Animal Learning & Behavior*, *32*(1), 4–14.

Merritt, S. M., & Ilgen, D. R. (2008). Not All Trust Is Created Equal: Dispositional and History-Based Trust in Human-Automation Interactions. *Human Factors: The Journal of the Hu-*

man Factors and Ergonomics Society, *50*(2), 194–210. https://doi.org/10.1518/001872008X288574

Norman, D. A. (2014). Some observations on mental models. In *Mental models* (pp. 15–22). Psychology Press.

Osiurak, F., & Reynaud, E. (2019). The Elephant in the Room: What Matters Cognitively in Cumulative Technological Culture. *Behavioral and Brain Sciences*, 1–57. https://doi.org/10.1017/S0140525X19003236

Petrides, K. V., & Furnham, A. (2000). Gender differences in measured and self-estimated trait emotional intelligence. *Sex Roles*, *42*(5–6), 449–461.

Risko, E. F., & Dunn, T. L. (2015). Storing information in-the-world: Metacognition and cognitive offloading in a short-term memory task. *Consciousness and Cognition*, *36*, 61–74.

Risko, E. F., & Gilbert, S. J. (2016). Cognitive Offloading. *Trends in Cognitive Sciences*, *20*(9), 676–688. https://doi.org/10.1016/j.tics.2016.07.002

Risko, E. F., Medimorec, S., Chisholm, J., & Kingstone, A. (2014). Rotating with rotated text: A natural behavior approach to investigating cognitive offloading. *Cognitive Science*, *38*(3), 537–564.

Schaller, R. R. (1997). Moore’s law: Past, present and future. *IEEE Spectrum*, *34*(6), 52–59. https://doi.org/10.1109/6.591665

Touron, D. R. (2015). Memory avoidance by older adults: When “old dogs” won’t perform their “new tricks.” *Current Directions in Psychological Science*, *24*(3), 170–176.

Vuorre, M. (2017). *Bmlm: Bayesian Multilevel Mediation. R package version 1.3.4.* https://cran.r-project.org/package=bmlm

Vuorre, Matti, & Bolger, N. (2018). Within-subject mediation analysis for experimental data in cognitive psychology and neuroscience. *Behavior Research Methods*, *50*(5), 2125–2143.

Walsh, M. M., & Anderson, J. R. (2009). The strategic nature of changing your mind. *Cognitive Psychology*, *58*(3), 416–440.

Wegner, D. M. (1987). Transactive memory: A contemporary analysis of the group mind. In B. Mullen & G. R. Goethals (Eds.), *Theories of group behavior* (pp. 185–208). Springer.

Wegner, D. M., & Ward, A. F. (2013). How Google Is Changing Your Brain. *Scientific American*, *309*(6), 58–61. https://doi.org/10.1038/scientificamerican1213-58

Weis, P. P., & Wiese, E. (2019a). Using Tools to Help Us Think: Actual But also Believed Reliability Modulates Cognitive Offloading. *Human Factors*, *61*(2), 243–254. https://doi.org/10.1177/0018720818797553

Weis, P. P., & Wiese, E. (2019b). Problem Solvers Adjust Cognitive Offloading Based on Performance Goals. *Cognitive Science*, *43*(12). https://doi.org/10.1111/cogs.12802

Wickham, H. (2016). *ggplot2: Elegant graphics for data analysis*. Springer.

Wiese, E., Metta, G., & Wykowska, A. (2017). Robots as intentional agents: Using neuroscientific methods to make robots appear more social. *Frontiers in Psychology*, *8*, 1663.

Wiese, E., Wykowska, A., Zwickel, J., & Müller, H. J. (2012). I See What You Mean: How Attentional Selection Is Shaped by Ascribing Intentions to Others. *PLoS ONE*, *7*(9), e45391. https://doi.org/10.1371/journal.pone.0045391

Wykowska, A., Wiese, E., Prosser, A., & Müller, H. J. (2014). Beliefs about the Minds of Others Influence How We Process Sensory Information. *PLoS ONE*, *9*(4), e94339. https://doi.org/10.1371/journal.pone.0094339

BIOGRAPHIES

Patrick Weis is a postdoc in the Psychology Departments at Nicolaus Copernicus University in Torun, Poland, and at Ulm University in Ulm, Germany. He received his PhD in Human Factors and Applied Cognition from George Mason University in 2019. He also received an MS in Neuroscience from the University of Tuebingen in 2014.

Eva Wiese is an associate professor in the Human Factors and Applied Cognition Program at George Mason University. She received her PhD in Neuroscience from the Ludwig-Maximilian University of Munich, Germany in 2013.

SUPPLEMENTAL MATERIALS

Perceived competence ratings: Mediation models

In both mediation models, offloading preference was estimated using a multiple regression procedure with Task Type, Cognitive Environment, and participants as predictors. Scores for the predictor variable were computed with participants as a categorical variable in a regular multiple regression rather than as a random factor in a multilevel model because adding participants as a random factor would have turned the multilevel model singular. The difference between the task-specific perceived competence rating scores[9] of the available cognitive environments (e.g., human - robot) was used as mediator. In other words, the mediator "competence rating difference" is defined as the task-specific difference in perceived competence between the available offloading options (i.e., between both agents for trials in the *Agents* Cognitive Environment and between both apps the *Apps* Cognitive Environment). For example, if a participant in the *Agents* Cognitive Environment rated the human's competence in the arithmetic task as 80 and the robot's competence as 95, the competence rating difference score would be "-15". Offloading preference was used as predicted variable. To estimate mediation parameters, R's bmlm package (M. Vuorre, 2017) was used. Variables were centered and standardized within each subject. We thus report standardized weights.

Results of the mediation model in the *Task-specific Agent and Task-unspecific App Expertise Beliefs* Mental Model condition are reported first (see also **Figure 5a**). In Step 1 of the

[9] Perceived competence was measured right before engaging in the task but after reading the metacognitive descriptions.

mediation model, "Task Type x Cognitive Environment"[10] was found to predict offloading preferences (*c* pathway or *total effect*: *M* = .76, 95% CI = [.63 .89]). Step 2 showed that "Task Type x Cognitive Environment" predicted competence rating differences (*a* pathway*: M* = .82, 95% CI = [.71 .93]). Please recall that "competence rating difference" is defined as the task-specific difference in perceived proficiency between the available offloading options (i.e., between both agents for trials in the *Agents* Cognitive Environment and between both apps the *Apps* Cognitive Environment). Step 3 of the mediation process showed that the mediator (competence rating difference), when controlling for "Task Type x Cognitive Environment", predicted offloading preference (*b* pathway*: M* = .26, 95% CI = [.05 .46]). Step 4 showed that "Task Type x Cognitive Environment", when controlling for the mediator (competence rating differences) still predicted offloading preferences (*c'* pathway*: M* = .48, 95% CI = [.28 .68]). Note that step 3 and 4 are relying on the same regression equation. To test for the indirect effect, R's bmlm package also computes Bayesian estimates and confidence intervals for the complete indirect pathway (i.e., *a***b*) as well as the percentage mediated (i.e., *a***b* / *c*). Both the indirect effect (*M* = .28, 95% CI = [.10 .48]) as well as the percentage mediated (*M* = .37, 95% CI = [.13 .62]) were significantly greater than zero. These results support partial mediation and therefore suggest that task-specific competence ratings are a relevant part of a human problem solver's mental model of an agent or an app that

---

[10] Please note that the predictor "Task x Environment" was realized not by including the actual factor levels into the model but by using the offloading preferences as predicted by a multiple regression with the predictors *Task*, *Environment*, *Task* x *Environment,* and *Participant*. Predicted scores were computed with participants as a categorical variable in a regular multiple regression rather than as a random factor in a multilevel model because adding participants a random factor would turn the multilevel model singular. The corresponding R code can be accessed via the OSF repository linked in this manuscript.

affords help with cognitive tasks. For details on the Bayesian parameter estimation, consult Vuorre and Bolger (2018). Mean rating data is provided for the curious reader in **Figure S1**[11].

In this paragraph, results of the mediation model in the *Task-unspecific Agent and Task-specific App Expertise Beliefs* Mental Model condition are reported; compare **Figure 5b**. The identical procedure as applied to the *Task-specific Agent and Task-unspecific App Expertise Beliefs* Mental Model data was used. In Step 1 of the mediation model, "Task Type x Cognitive Environment" was found to predict offloading preferences (*c* pathway or *total effect*: M = .66, 95% CI = [.52 .79]). Step 2 showed that "Task Type x Cognitive Environment" predicted competence rating differences (*a* pathway: M = .76, 95% CI = [.64 .88]). Step 3 of the mediation process showed that the mediator (competence rating differences), when controlling for "Task Type x Cognitive Environment", still predicted offloading preferences (*b* pathway: M = .23, 95% CI = [.09 .38]). Step 4 showed that "Task Type x Cognitive Environment", when controlling for the mediator (competence rating differences) still predicted offloading preferences (*c'* pathway: M = .50, 95% CI = [.35 .65]). Both the indirect effect (M = .16, 95% CI = [.05 .28]) as well as the percentage mediated (M = .24, 95% CI = [.07 .42]) were significantly greater than zero. Just as in the *Task-specific Agent and Task-unspecific App Expertise Beliefs* Mental Model condition, these results support partial mediation. Specifically, these results cross-validate the findings of the first model in that task-specific competence ratings are likely a relevant part of a human problem solver's mental model of an agent or app that affords help with cognitive tasks.

[11] Note that, for exploratory purposes, we obtained perceived competence ratings before and after participants engaged in the task. However, also note that, as indicated by the means, both ratings seem to be highly correlated.

756 *Perceived competence ratings: Descriptive statistics*

757 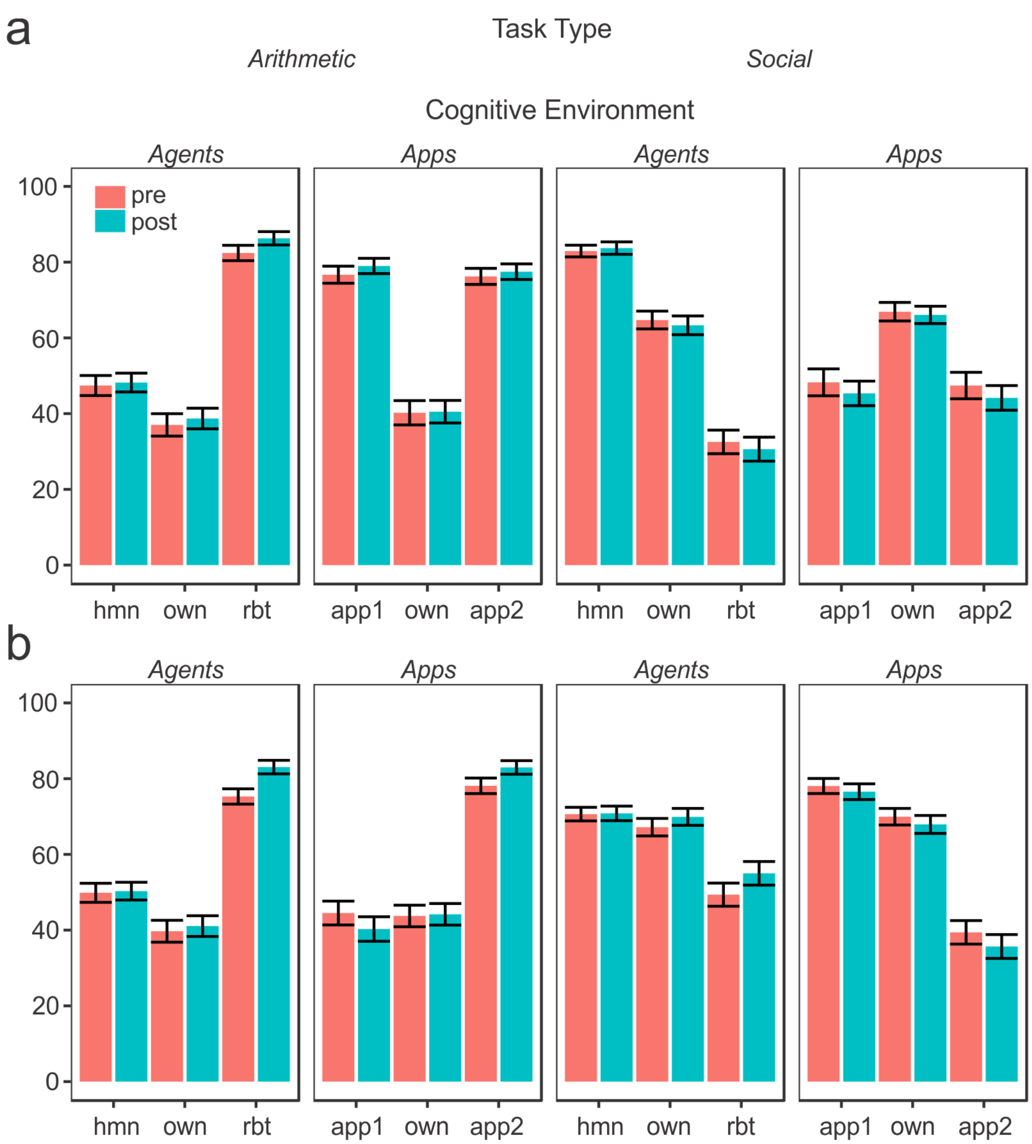

758 *Figure S1: Perceived Competence Ratings*. Ratings for the *Task-specific Agent and Task-unspecific App Expertise Beliefs* (a) and
759 the *Task-unspecific Agent and Task-specific App Expertise Beliefs* (b) Mental Model conditions are depicted on the y-axis. Error
760 bars represent +/- 1 SEM. Env.: Cognitive Environment, hmn: human agent ("Michael"), rbt: robot agent ("Meka), app1: Om-
761 nilearn app, app2: Pattern Analytics app, pre: before engaging in task, post: after engaging in task.

762

Task responses: Numeric values corresponding to Figure 3

Table S1

Response counts split by Mental Model, Task Type, and Cognitive Environment

Mental Model	Task Type	Cognitive Environment	Response [count]		SD
Task-specific Agent and Task-unspecific App Expertise Beliefs	Arithmetic	Agents	Own	6.3	6.7
			Human	0.8	1.5
			Robot	10.9	6.7
		Apps	Own	7.4	7.0
			Pattern Analytics	5.3	4.9
			Omnilearn	5.3	4.7
	Social	Agents	Own	9.8	6.0
			Human	7.8	5.9
			Robot	0.4	0.9
		Apps	Own	12.8	6.3
			Pattern Analytics	2.3	3.9
			Omnilearn	2.9	4.2
Task-unspecific Agent and Task-specific App Expertise Beliefs	Arithmetic	Agents	Own	6.3	6.7
			Human	1.7	3.0
			Robot	10.0	6.7
		Apps	Own	6.5	7.1
			Pattern Analytics	9.3	7.1
			Omnilearn	2.3	3.9
	Social	Agents	Own	11.4	6.4
			Human	4.2	4.9
			Robot	2.5	4.6
		Apps	Own	10.0	7.0
			Pattern Analytics	1.4	3.2
			Omnilearn	6.5	6.7