AMERICAN MARKETING ASSOCIATION

Article

Journal of Marketing
2025, Vol. 89(5) 1-20
© American Marketing Association 2025
Article reuse guidelines:
sagepub.com/journals-permissions
DOI: 10.1177/00222429251314491
journals.sagepub.com/home/jmx
Sage

Lower Artificial Intelligence Literacy Predicts Greater AI Receptivity

Stephanie M. Tully, Chiara Longoni, and Gil Appel

Abstract

As artificial intelligence (AI) transforms society, understanding factors that influence AI receptivity is increasingly important. The current research investigates which types of consumers have greater AI receptivity. Contrary to expectations revealed in four surveys, cross-country data and six additional studies find that people with lower AI literacy are typically more receptive to AI. This lower literacy–greater receptivity link is not explained by differences in perceptions of AI's capability, ethicality, or feared impact on humanity. Instead, this link occurs because people with lower AI literacy are more likely to perceive AI as magical and experience feelings of awe in the face of AI's execution of tasks that seem to require uniquely human attributes. In line with this theorizing, the lower literacy–higher receptivity link is mediated by perceptions of AI as magical and is moderated among tasks not assumed to require distinctly human attributes. These findings suggest that companies may benefit from shifting their marketing efforts and product development toward consumers with lower AI literacy. In addition, efforts to demystify AI may inadvertently reduce its appeal.

Keywords

artificial intelligence, algorithms, AI literacy, AI adoption, technology, human–computer interaction, awe

Online supplement: https://doi.org/10.1177/00222429251314491

Submitted May 17, 2024

Artificial intelligence (AI) systems—automated systems so advanced that they can mimic or outperform human cognition—are radically transforming our society. Companies operating in the private and public sectors have been rapidly adopting AI systems due to the competitive advantages and operational efficiencies they offer (Brynjolfsson, Li, and Raymond 2025). Indeed, global investments in AI have grown substantially in the last decade, a trend that is forecasted to continue (Maslej et al. 2023). Most recently, generative AI—AI systems capable of autonomously generating textual, visual, or auditory data—has reached the public at large with the release in 2022 of text-to-image models like DALL-E, text-to-video systems like Make-A-Video, and chatbots like ChatGPT. Thus, consumers are encountering AI-driven products and services in nearly every aspect of their personal, social, and professional life, in the form of digital personal assistants, recommendation agents on social media, and customer service chatbots.

Due to the many potential benefits of AI to consumers, companies, and society at large (Brynjolfsson, Li, and Raymond 2025; Hermann, Yalcin Williams, and Puntoni 2024), substantial research has examined the determinants of consumer's receptivity toward AI—consumers' proclivity toward AI usage. This research has typically focused on *situational factors* associated with receptivity to AI (e.g., Castelo, Bos, and Lehmann 2019; Valenzuela et al. 2024). However, marketers must also understand which consumers are more likely to adopt their AI products for optimal targeting and product development. Much less is known about how systematic differences across consumers—that is, *individual-level factors*—predict AI receptivity, which is the focus of the present research.

More specifically, in the current work we explore the relationship between consumers' *AI literacy* (i.e., their objective knowledge about AI) and their *AI receptivity* (i.e., proclivity toward AI usage). Though most people forecast that consumers with greater AI literacy are more receptive to AI, we theorize

Stephanie M. Tully is Associate Professor of Business Administration, Marshall School of Business, University of Southern California, USA (email: tullyste@marshall.usc.edu). Chiara Longoni is Associate Professor of Marketing, Bocconi University, Italy (email: chiara.longoni2@unibocconi.it). Gil Appel is Assistant Professor of Marketing, GW School of Business, George Washington University, USA (email: gappel@gwu.edu).

and demonstrate that it is those with lower AI literacy who exhibit greater AI receptivity— a relationship we refer to as the *lower literacy–higher receptivity link*.

The lower literacy–higher receptivity link occurs because AI systems are executing tasks that were once thought to require uniquely human traits, such as feeling emotions, having a personality, and holding beliefs (see Santoro and Monin 2023). Whereas people with higher AI literacy recognize that AI can execute such tasks without actually having distinctly human attributes, people with lower AI literacy perceive AI's execution of such tasks as more magical and experience greater feelings of awe when thinking about AI executing such tasks. In turn, perceptions of AI as magical and ensuing feelings of awe explain the lower literacy–higher receptivity link. Thus, when AI executes tasks that are *not* associated with distinctly human attributes—tasks less likely to evoke perceptions of AI as magical and feelings of awe—the relationship between AI literacy and receptivity is moderated. In addition, we find that people with higher AI literacy generally perceive AI to be more capable and more ethical, have less fear of AI, and exhibit greater readiness to adopt technology in general than people with lower AI literacy. Thus, people with lower literacy exhibit greater AI receptivity despite, rather than because of, these alternative explanations.

Theoretically, our research identifies AI literacy as an important yet overlooked individual-level predictor of AI receptivity. In explaining the lower literacy–higher receptivity link, we present a theory of AI receptivity based on perceptions of AI as magical and ensuing feelings of awe that can help reconcile conflicting findings on consumer responses to new technologies. Practically, our research suggests that marketing managers' inclinations toward targeting consumers with higher AI literacy for AI-based products and services may be misguided. Instead, businesses may benefit from targeting those with lower AI literacy, designing products to meet the needs of this target segment, and tailoring their marketing messages to highlight the perceived magicalness of AI solutions. This is particularly true for AI products and services designed to execute tasks perceived to require distinctly human attributes.

Theoretical Development

AI Receptivity

The adoption of AI-based products and services by consumers plays a pivotal role in the success and growth of businesses and society (Hermann and Puntoni 2024). AI-based products and services can address complex challenges and provide solutions that may not be achievable through traditional methods alone. For instance, in health care, AI applications can assist in early disease detection, improve diagnostic accuracy, and support treatment decisions (Topol 2019), and in finance, AI tools have the potential to increase financial inclusion (e.g., Wang et al. 2019).

Reaping the benefits of AI critically depends on AI receptivity, which we define as consumers' proclivity toward AI usage. This includes an individual's openness, readiness, and/or likelihood of integrating AI into their own lives and to have AI integrated into the world around them. Substantial conceptual and empirical research has indeed explored the determinants of people's receptivity toward AI. This research has highlighted the situational factors related to variability in AI receptivity: how characteristics of the task (e.g., subjectivity vs. objectivity), salient consumer goals (e.g., hedonic vs. utilitarian considerations), and attributes of the AI agent (e.g., performance, actual and perceived capabilities, social distance) affect consumers' responses to AI (Castelo, Bos, and Lehmann 2019; De Bellis, Johar, and Poletti 2023; Longoni and Cian 2022; Yalcin et al. 2022; for reviews, see Puntoni et al. 2021; Valenzuela et al. 2024).

In contrast, seemingly little research focuses on individual factors that systematically predict receptivity toward AI (cf. Liu, Kirshner, and Lim 2023). It is important to consider how individual-level differences across consumers might impact their receptivity toward AI use to better understand to the optimality of targeting different marketing segments, to appreciate and evaluate the implications of policy recommendations, and to optimize the design of AI applications. Yet, individual-level factors have been largely neglected by empirical AI research. One important yet overlooked factor that may impact AI receptivity, and is the focus on the present research, is consumers' degree of knowledge about AI—their AI literacy.

AI Literacy

By AI literacy, we refer to a person's objective knowledge about AI, including conceptual and technical knowledge, best practices in AI development and use, and regulatory requirements. At the moment, it is unclear how AI literacy might impact consumers' receptivity toward AI, in part due to lack of appropriate measurement of AI literacy. Academic work has often conflated AI knowledge with opinions about how algorithms operate in specific settings like Facebook News Feed and TikTok (e.g., Rader and Gray 2015; Siles, Valerio-Alfaro, and Meléndez-Moran 2022), or measured knowledge in terms of self-reported assessments of competency. Indeed, existing measures of AI knowledge tend to capture whether people believe they have the necessary technical skills to use technologies, and whether they believe they know a lot about technologies (e.g., Kabakus, Bahcekapili, and Ayaz 2023), rather than people's objective AI knowledge. Even when AI knowledge has been assessed more directly and objectively, research has not directly connected such knowledge to consumers' receptivity toward AI (e.g., Dogruel, Masur, and Joeckel 2022; Latzer, Festic, and Kappeler 2021; Zarouali, Boerman, and De Vreese 2021). Finally, methodologically, research exploring consumers' knowledge about algorithms has been predominantly conceptual or qualitative, relying on open-ended protocols, interviews, and opinion surveys (e.g., Cotter and Reisdorf 2020; Hargittai et al. 2020). In contrast to beliefs, opinions, and subjective perceptions about AI, in the current research we develop and employ objective, quantitative measures of

Table 1. People's Forecasts About the Relationship Between AI Literacy and AI Receptivity.

Participants	N	Want to Use AI More	Want to Use AI Less	No Relationship
X users	63	66.7%[a]	19.0%[b]	14.3%[b]
LinkedIn users	30	73.0%[a]	10.0%[b]	17.0%[b]
Business students	38	73.7%[a]	13.2%[b]	13.2%[b]
Prolific respondents	301	63.8%[a]	17.3%[b]	18.9%[b]

Notes: The question asked was "Do you think that people with more (vs. less) knowledge of AI and algorithms: (a) Want to use AI more, (b) Want to use AI less, (c) No relationship." Proportions on the same row with different subscripts are significantly different from each other at $p \leq .001$.

AI literacy to examine its relationship with receptivity toward AI-based products and services.

AI Literacy and AI Receptivity: Competing Predictions

There are competing predictions with respect to the relationship between AI literacy and AI receptivity. Perhaps the most straightforward prediction is that higher levels of AI literacy are associated with greater receptivity toward AI-based products and services. For example, research shows that consumers are less dissatisfied after being denied services by an algorithm if they are provided an explanation of what the overarching goal or objective of the algorithm is (Tomaino et al. 2022). This finding insinuates that people's lack of knowledge about how AI works makes them more cynical toward AI making decisions. In line with this idea, some researchers have explicitly advocated for greater education about technology to increasing its adoption (e.g., Richter and Sinha 2020). Such assertions further imply that AI literacy is positively associated with greater receptivity toward utilizing AI-based products and services.

Most people appear to share the intuition that people with more knowledge about AI will be more receptive toward it. Indeed, in four surveys among consumers from different populations—X (formerly Twitter) users, LinkedIn users, business undergraduate students, and an online sample—we asked people to predict whether people with more (vs. less) AI literacy would want to use AI more versus less, or whether there would be no relationship with AI usage. Across all four samples, people believed that greater AI literacy would be associated with greater propensity to use it. The results of these polls are reported in Table 1. In addition, we surveyed a group of 36 top line executives at a major European insurance company and asked them which customer segment they would be more likely to target for AI-based products in development. All 36 executives indicated that they would target consumers with higher (vs. lower) AI literacy. In summary, current wisdom favors the prediction that people with higher AI literacy are more receptive to AI.

In contrast to these forecasts, we propose and demonstrate that AI receptivity is greater among those with lower AI literacy, due to perceptions of AI as magical, as we elucidate next.

The Lower Literacy–Higher Receptivity Link

Artificial intelligence is an umbrella term for computing systems capable of performing tasks traditionally performed by humans, such as reasoning, making decisions, or solving problems (McCarthy et al. 2006). These "intelligent machines" can operate in ways that resemble human cognition, executing tasks and making decisions that are typically linked to distinctly human attributes such as having a personality or holding beliefs. AI's ability to execute tasks once thought to be exclusively carried out by humans such as producing speech, creating a movie or playing an instrument can seem to transcend the boundaries of conventional technology. Of course, AI systems do not actually need to possess human attributes to perform these tasks. For example, large language models can generate seemingly empathetic responses without experiencing empathy themselves. This can create an illusory perception that AI has human attributes when it executes these tasks. Indeed, interviews with teachers of machine learning found that students often readily attribute human qualities to such technology (Sulmont, Patitsas, and Cooperstock 2019). This misattribution of human attributes to AI systems is related to the concept of anthropomorphism in technology. As discussed by Epley, Waytz, and Cacioppo (2007), people tend to ascribe human characteristics to nonhuman agents, especially when these agents exhibit behavior that resembles human actions. We suggest that AI executing tasks previously believed to require distinctly human attributes and the resulting perceptions of AI as having humanlike attributes can result in people perceiving AI as magical.

Indeed, AI has been likened to a form of magic, in part due to the illusory manner in which it can appear to have distinctly human attributes and even sentience (Sharkey and Sharkey 2006). This assertion aligns with Arthur Clarke's third and best-known law: "Any sufficiently advanced technology is indistinguishable from magic" (Clarke 1968, p. 255). The notion of AI as magical further echoes with depictions of AI in popular culture. AI is frequently portrayed in popular culture with an aura of mystery and power that exceeds human comprehension, even surpassing the limits of human cognition. For example, in *The Matrix*, the AI that governs the virtual world is portrayed as having godlike control over reality, bending physical laws in ways that seem magical. In fact, the premise that people readily think of AI as magical can help explain terminology used in AI coding such as "spellcasting," which indicates the provision of an input in a prompt to a language model, and the naming of AI tools like Google's AI Agent "Magical," which provides AI-enabled autofill.

We suggest that, to the extent that AI is perceived as magical, it should elicit feelings of awe. Awe typically results from encountering something of such magnitude that the experience surpasses one's existing comprehension of the world (Keltner and Haidt 2003). Prior research has connected the experience of awe to adoption of virtual reality and virtual assistants (Kautish et al. 2023; Quesnel and Riecke 2018). Accordingly, we expect that the more people perceive AI as magical, the greater their experience of awe, and the greater their AI receptivity.

Importantly, we argue that people with lower AI literacy are more likely to perceive AI as magical. Magic often relies on people not understanding or paying attention to how a magic trick works (e.g., Ekroll, Sayim, and Wagemans 2017). In much the same way, AI should be perceived as more magical when people do not understand or pay attention to how AI is able to execute tasks. This is more likely to be the case among those with lower rather than higher AI literacy. Without technical knowledge about how AI operates—such as the use of algorithms, training data, and computational models—those with lower AI literacy are more likely to perceive AI's outputs as extraordinary or even supernatural. That is, when people with lower AI literacy envision AI executing tasks that they believe require distinctly human attributes, they may perceive AI executing these tasks as magical because they cannot understand how AI is able to do so. In contrast, people with greater AI literacy may be more readily able to understand how AI can execute such tasks without actually possessing human attributes, dampening perceptions of magicalness and subsequent awe. Indeed, once a person understands the mechanism behind a magic trick, it seems less magical and can ruin a sense of wonder and excitement for magic more broadly (Braxton 1998).

To illustrate, people with lower AI literacy may view AI as possessing a sense of humor if it predicts the funniness of a joke, leading to a perception of magicalness and feelings of awe at the thought of AI predicting joke funniness. In contrast, people with higher AI literacy are more likely to understand that through pattern matching and algorithms, it is possible for AI to predict joke funniness without having a sense of humor. This demystification of AI as having distinctly human attributes should result in lower perceptions of AI as magical among those with greater AI literacy. In turn, they should experience less awe and exhibit lower AI receptivity. In summary, we predict that people with lower AI literacy will exhibit greater AI receptivity, and that this difference will be explained through differences in perceptions of AI as magical and the ensuing feelings of awe that people feel when envisioning AI executing tasks.

Importantly, our theorization suggests that this effect depends on thinking about AI accomplishing tasks typically associated with distinctively human attributes (e.g., feeling emotions, having culture, a sense of humor; Santoro and Monin 2023). That is, when AI executes tasks that typically require distinctively human attributes, those with lower (vs. higher) AI literacy are more likely to perceive AI as having distinctively human attributes, resulting in perceptions of AI as magical and greater AI receptivity. Although AI is distinguished from other forms of technology in its ability to emulate human intelligence specifically, AI can be used to execute tasks that are more readily recognized as requiring nondistinctly human attributes (e.g., having logic, doing computations, sensing temperature; Santoro and Monin 2023). When AI performs these types of tasks, perceptions of AI as magical should not arise and differences in AI receptivity should be moderated.

Overview of the Current Research

Contrary to the results of four forecasting studies and a survey of top executives, seven studies demonstrate that people with lower AI literacy are more receptive toward AI. We explore the relationship between AI literacy and AI receptivity across a range of populations including students, online participants, and citizens across countries. We employ the following methods of measuring AI literacy to ensure that the results are not a function of any specific way of assessing it: a third-party measure of AI expertise assessed through AI education, skills, and jobs (Study 1), a 25-item measure developed by the authors (Studies 2, 4, and 6); and a 17-item measure developed by AI (Claude.AI and ChatGPT-4; Studies 3, 5, and 7). Finally, we operationalize AI receptivity in several ways. We assess AI receptivity via general measures of AI adoption readiness (Studies 1 and 5), propensity to use generative AI for oneself (Study 2), frequency of past AI usage (Study 3), and relative preference for human versus AI task execution across a range of tasks (Studies 4, 6, and 7). Study 1 shows that cross-country differences in AI literacy predict country-level variation in AI receptivity. Study 2 demonstrates that undergraduate students with lower AI literacy are more receptive to using AI to complete academic assignments compared with those with higher AI literacy. Study 3 finds that AI literacy provides predictive ability in explaining AI receptivity, operationalized as reported AI usage over the previous six months, beyond other individual differences potentially related to AI literacy (i.e., technology readiness, autonomy, or general knowledge). Studies 4 replicates the lower literacy–higher receptivity link, providing evidence that this relationship is mediated through perceptions of AI as magical. Study 5 further provides evidence in favor of the role of magical perceptions, demonstrating serial mediation through ensuing feelings of awe, and provides evidence that this link occurs despite rather than because of alternative explanations based on perceived AI capability and fear of AI's impact on humanity. Study 6 generalizes the lower literacy–higher receptivity link across a broad range of tasks. Two posttests reported in Web Appendix E show that these results are not explained by ethicality or capability perceptions. However, in line with our theorizing, the results are moderated by the extent to which tasks are perceived as more objective and less likely to be associated with distinctly human attributes, since such tasks should not elicit perceptions of AI as magical. Finally, in Study 7, we directly manipulate the extent to which tasks are perceived to necessitate distinctly human attributes for task execution between subjects. Among participants evaluating tasks needing some level of distinctly human attributes, we replicate the lower literacy–higher receptivity link. However, among participants evaluating tasks not requiring distinctly human attributes, this relationship reverses, and people with higher literacy exhibit greater AI receptivity. Table 2

Table 2. Overview of the Study Designs and Results for the Main Dependent Variables.

Study 1: Cross-country differences in AI literacy predict AI receptivity (N = 27 countries; secondary data from Ipsos and Tortoise Media; AI literacy is log-transformed)

	AI Literacy
AI receptivity (AI adoption readiness)	B = −8.80, SE = 2.36, t(25) = −3.73, $p < .001$

Main findings: Lower AI literacy among people within a country predicted greater country-level AI receptivity.

Study 2: AI literacy predicts student receptivity toward using generative AI (N = 234; lab; undergraduate students)

	AI Literacy
AI receptivity (propensity to use generative AI for writing assignments)	B = −.05, SE = .02, t(232) = −3.62, $p < .001$

Main findings: Lower AI literacy predicted greater propensity to use AI to execute written assignments.

Study 3: AI literacy predicts reported past AI usage beyond other individual-level differences (N = 401; Amazon Mechanical Turk)

	AI Literacy
AI receptivity (reported usage of AI over previous six months)	B = −.09, SE = .02, t(399) = −5.73, $p < .001$

Main findings: Lower AI literacy predicted more frequent reported usage of AI over the previous six months. This predictive ability was largely unchanged when including other individual differences such as technology readiness, general knowledge, and desire for autonomy.

Study 4: AI literacy predicts receptivity toward using AI (vs. human) to execute five different tasks, a relationship explained by perceptions of AI as magical (N = 1,006; Lucid; U.S. representative sample)

	AI Literacy
AI receptivity (relative preference for AI vs. human task execution)	B = −.10, SE = .01, t(1,004) = −10.80, $p < .001$
Perceptions of AI as magical	B = −.09, SE = .01, t(1,004) = −11.30, $p < .001$

Main findings: Lower AI literacy predicted greater preference for AI (vs. human) task execution. Differential AI receptivity was explained by perceptions of AI as magical (95% CI: [−.03, −.01]).

Study 5: AI literacy predicts receptivity toward using AI, a relationship serially explained by perceptions of AI as magical and inspiring feelings of awe. AI's capability perceptions and fear of AI do not explain the results (N = 1,000; Prolific Academic)

	AI Literacy
AI receptivity (AI adoption readiness)	B = −.04, SE = .017, t(998) = −2.25, $p = .024$
Perceptions of AI as magical	B = −.13, SE = .020, t(998) = −6.83, $p < .001$
Extent to which AI inspires feelings of awe	B = −.14, SE = .021, t(998) = −6.79, $p < .001$
Perceptions of AI's capability	B = .07, SE = .014, t(998) = 4.94, $p < .001$
Fears about AI's impact on humanity	B = −.09, SE = .024, t(998) = −3.56, $p < .001$

Main findings: Lower AI literacy predicted greater interest in AI-based products and services, a relationship explained by greater perceptions of AI as magical and greater feelings of awe toward AI (95% CI: [−.05, −.02]). Lower AI literacy predicted greater fear and reduced perceptions of AI capability; controlling for AI fear and AI capability perceptions strengthened the relationship between AI literacy and AI receptivity.

Study 6: AI literacy predicts receptivity toward relying on AI (vs. human) to execute 26 different tasks (N = 1,099; Amazon Mechanical Turk)

	AI Literacy	**Objectivity**	**Interaction**
AI receptivity (relative preference for AI vs. human task execution) across task objectivity	B = −.05, SE = .01, t(1,096.61) = −6.38, $p < .001$	B = .01, SE < .01, t(27,456.70) = 16.87, $p < .001$	B = .001, SE < .01, t(27,456.53) = 7.59, $p < .001$

	AI Literacy	**Distinctly Human vs. Shared Attributes**	**Interaction**
AI receptivity (relative preference for AI vs. human task execution) across distinctly human vs. shared attributes	B = −.05, SE = .01, t(1,096.81) = −6.38, $p < .001$	B = .86, SE = .01, t(27,456.56) = 93.94, $p < .001$	B = .04, SE < .01, t(27,456.44) = 16.19, $p < .001$

(continued)

Table 2. (continued)

Main findings: Lower AI literacy predicted greater relative preference for AI (vs. human) task execution. This relationship was attenuated among tasks rated as more objective and those less likely to need distinctly human attributes to execute.

Study 7: AI literacy predicts receptivity toward relying on AI (vs. human) to execute ten tasks varying in their association with distinctly human attributes vs. shared human and AI attributes (N = 1,006; Toluna)

	AI Literacy	Distinctly Human vs. Shared Attributes	Interaction
AI receptivity (relative preference for AI vs. human task execution) across distinctly human vs. shared attributes	B = .08, SE = .02, t(1,002) = 3.57, $p < .001$	B = −1.17, SE = .09, t(1,002) = −13.15, $p < .001$	B = −.19, SE = .03, t(1,002) = −6.38, $p < .001$

Main findings: Lower AI literacy predicted greater relative preference for AI (vs. human) task execution when the task was associated with distinctly human attributes (B = −.12, SE = .02, t(1,002) = −5.43, $p < .001$), a relationship reversed when the task was associated with shared attributes (B = .08, SE = .02, t(1,002) = 3.57, $p < .001$).

summarizes our empirical package, details of empirical testing, and results. Studies 2–7 were preregistered. For these studies, our sample size and data exclusions were determined before data collection, and we report all conditions, measures, and data exclusions. Preregistrations, data, and research materials for all studies are available at https://researchbox.org/1491.

Study 1: Cross-Country Differences in AI Literacy Predict Country-Level AI Receptivity

Study 1 combines a dataset capturing differences in people's proficiency and knowledge of AI across countries with a separate dataset of cross-country AI receptivity to explore whether differences in AI literacy across countries predict country-level receptivity toward AI-based products and services.

Data

AI literacy. The Global AI Index (GAII) was created by Tortoise Media in 2019 to assess the level of investment, innovation, and implementation of AI around the world. Relevant to the current research, the 2021 GAII (version 3.0) was extended to include a measure of AI "talent" (Tortoise Media 2023). The AI talent is a weighted combination of the number of people with AI-related job titles, online programming conversations and activity, and measures of AI-related knowledge through coursework by combining data from LinkedIn, GitHub, UNESCO, Coursera, Kaggle, CRAN, Alexa API, and Google BigQuery within a country. This measure should thus capture differences in the level of AI literacy across countries and served as our measure of AI literacy. This measure was available for 62 countries. AI literacy ranged from 1.72 to 100 (M = 20.51). Since the measure demonstrated a highly skewed distribution, we log-transformed it for our analysis. AI literacy (log-transformed) ranges from .54 to 4.61 (M = 2.69).

AI receptivity. In 2022, Ipsos released the results of its Global Opinions and Expectations About AI Survey, which polled 19,504 online adults aged 16–74 across 28 countries between November and December of 2021 (Ipsos 2022). This survey asked participants about a variety of perceptions regarding AI. For the purposes of the current research, four of the measures assessed AI adoption readiness: "Products and services using artificial intelligence will profoundly change my daily life in the next 3–5 years," "Products and services using artificial intelligence make my life easier," "Products and services using artificial intelligence have more benefits than drawbacks," and "I trust companies using artificial intelligence as much as I trust other companies." The survey asked respondents whether they agreed with each of these four statements (yes/no). These data thus indicate, for each of the four measures, the percentage of respondents in each country who indicated that they agreed with the statement. These four measures, available for 28 countries, were strongly related (Cronbach's $\alpha = .983$), and we averaged them into an index capturing AI receptivity. AI receptivity ranges from 37.25 to 80.25 (M = 54.97).

Results and Discussion

All countries available in the Ipsos dataset, with the exception of Peru, were available in the GAII. Thus, across the two datasets, there were 27 overlapping countries. In line with our predictions, regressing AI receptivity on AI literacy revealed that lower levels of AI literacy within a country predicted greater AI receptivity (B = −8.80, 95% CI: [−13.67, −3.94], SE = 2.36, t(25) = −3.73, $p < .001$; see Figure 1). Because the United States had a significantly greater literacy score than the other countries, we also examined the relationship across countries, excluding the United States. Among the remaining 26 countries, there remained a significant negative relationship between AI literacy within a country and AI receptivity (B = −8.95, 95% CI: [−14.53, −3.37], SE = 2.70, t(24) = −3.31, $p = .003$). The direction and significance of the results were unchanged using each of the four individual measures of AI receptivity rather than the index (all $p \leq .003$).

One may wonder about the extent to which these cross-country results reflect differences in AI literacy as compared

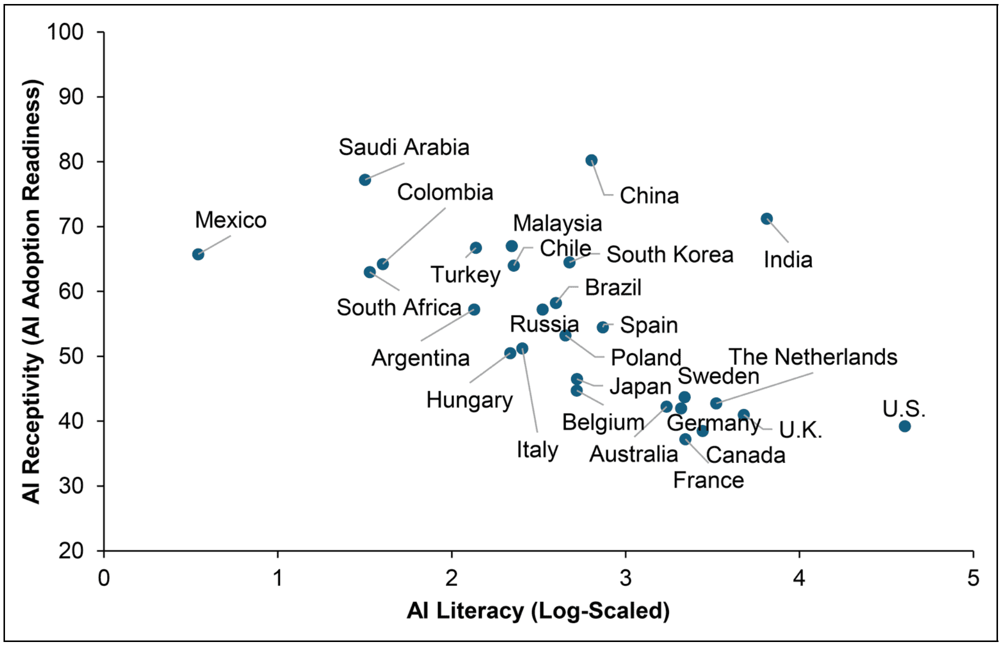

Figure 1. Relationship Between AI Literacy (Log-Scaled) and AI Receptivity Across 27 Countries.

with other factors that vary across more developed versus less developed countries. Thus, as a robustness check, we examined the relationship between AI literacy and AI receptivity, adjusting for the Human Development Index (HDI) measured by the United Nations Development Programme (2022). The HDI measures life expectancy, level of education (expected years of schooling, mean years of schooling), and income (gross national income per capita) within a country. The scores for these dimension indices are then aggregated into a composite index to create a singular HDI score. Table 3 provides the analysis including the HDI index (Model 2) and the three individual components that make up the HDI index (Model 3).[1] Note that excluding the United States does not change the significance level of AI literacy in either of these models. For additional robustness, we also considered the GDP growth rate and population growth rate as further controls. These results provide initial evidence for a negative relationship between AI literacy and AI receptivity.

Study 2: AI Literacy Among College Students Predicts Receptivity Toward Using Generative AI to Complete Assignments

Study 2 explores the relationship between AI literacy and receptivity to using generative AI chatbots, a type of AI increasingly used by students to complete academic assignments. We measure AI literacy using a multiple-choice test and then assess undergraduate students' propensity to use a generative AI agent to complete writing assignments.

[1] We averaged expected years of schooling and mean years of schooling to create a single education measure.

AI Literacy Measurement

To develop a measure of AI literacy, we reviewed academic scholarship, public resources, and practitioner-focused pedagogy covering the topics of AI, algorithms, and digital literacy. These data sources included the AI Education Project (aiEDU 2023), UNESCO's Algorithm & Data Literacy project (Digital Moment 2023), and the National University of Singapore's library on digital literacy (National University of Singapore 2023). This led to the development of a set of 47 possible questions covering awareness of where AI and algorithms are used, whether they are responsible for a given action or outcome, conceptual and technical knowledge, best practices in usage, and regulatory requirements. These questions were refined to a final set of 25 questions based on feedback from several data scientists, programmers, and ChatGPT-3. Full details about the development of this measure and the 25 items are available in Web Appendix A.

Procedure

This study was run in four weekly laboratory sessions, with each week having up to a maximum of 100 participants. An error in a previous study's redirect link resulted in only 23 responses in the first week. In total, we recruited 234 participants across the four sessions (M_{age} = 19.4 years, SD = .95; female = 45.7%, male = 53.4%, nonbinary/third gender = .4%, prefer to not say = .4%).

First, participants saw a brief introduction that explained the goal of the task: to assess people's knowledge of AI, algorithms, online behavior, and data privacy. They also read that they would be asked to complete this test relying solely on their knowledge and without looking up answers online, that they needn't be experts on the topic, and that their performance on

Table 3. Results of Study 1: Cross-Country AI Literacy Predicts AI Receptivity.

	Model 1	Model 2	Model 3	Model 4
	No Controls	Controlling for HDI	Controlling for HDI Components	Controlling for HDI Components and Economic Indicators
AI literacy (log-transformed)	−8.804*** (2.361)	−5.265* (2.057)	−5.849** (1.990)	−5.858* (2.071)
Human Development Index (HDI)		−82.953*** (20.575)		
Life expectancy			.174 (.438)	.131 (.462)
Education			−4.338* (1.724)	−4.318* (1.829)
Income (GNI, log-transformed)			−1.079 (6.579)	−.339 (7.409)
GDP growth rate				.373 (.666)
Population growth rate				−.489 (3.547)
Intercept	78.695*** (6.653)	140.412*** (16.181)	128.492** (37.209)	121.614* (46.241)
Adjusted R^2	.332	.585	.621	.589

*$p \leq .05$.
**$p \leq .01$.
***$p \leq .001$.
Notes: N = 27. Values reflect unstandardized coefficients. GNI = gross national income.

the test would not impact their participation credit. Finally, participants were asked to answer to the best of their knowledge and choose the answer that they believed was most accurate. Participants then completed the AI literacy test, with the sum of the correct responses serving as a participant's AI literacy score. The order of questions and the order of answers within each question were randomized. Our dependent variable was AI receptivity as measured by propensity to use a generative AI agent to complete assignments. Participants read a description of what a generative AI agent is, how it works, and the length of time that it typically takes to generate content. Participants then saw a real example of ChatGPT-3's written output in response to the prompt "write a one-page report about a recent war and the resulting impact on culture." After reviewing this example, participants read that we were interested in their likelihood of using this generative AI agent in their own assignments if it was available to them at no cost. Participants then rated their propensity to use this generative AI agent to complete four assignments, described to them as follows and presented in random order: "An opinion piece on the ethics of a vegetarian lifestyle," "A report about how the assassination of Archduke Franz Ferdinand led to the beginning of World War 1," "A poem about falling in love in Venice," and "A creative writing story about the perils of adolescence." Specifically, they indicated whether and in what way they would use this generative AI agent using the following 5-point scale: 1 = "I would not use the generative AI agent," 2 = "I would use the generative AI agent to come up with ideas, but would write the assignment myself," 3 = "I would use the generative AI agent to write the first draft, but would re-write it before submitting it," 4 = I would use the generative AI agent to complete the assignment, but would fact check it to fix any errors before submitting it," and 5 = "I would use the generative AI agent to complete the assignment, and submit it as written." Students completed an instructional manipulation check (IMC: "This question is NOT interested in your marital status. Instead, it is designed to ensure you are taking the time to read all questions, even toward the end of the survey. As such, instead of indicating your marital status from the options below, please select "married" from the options below.") before providing demographic information (i.e., gender and age).

Results and Discussion

We had an unexpectedly high failure rate on the IMC (41.9% of participants failed), and this failure rate was significantly correlated with scores on the AI literacy test ($r = .13$, $p = .044$). Thus, we report the results without any exclusions. However, the significance of our results is unchanged if we exclude participants based on the IMC as preregistered.

AI literacy. We computed an AI literacy score for each participant by summing the number of correct responses. The median score was 18 (range: 3–24, $M = 16.99$, $SD = 3.61$).

AI receptivity. As preregistered, we calculated participants' receptivity to using generative AI to complete written assignments by averaging reported likelihood of using AI across the four assignments into an AI receptivity index; higher numbers on the index indicate greater propensity to use AI (Cronbach's $\alpha = .78$). A linear regression on AI receptivity revealed that lower AI literacy scores were associated with higher AI receptivity ($B = -.05$, 95% CI: [−.08, −.02], $SE = .01$, $t(232) = -3.62$, $p < .001$). To ensure that the results were

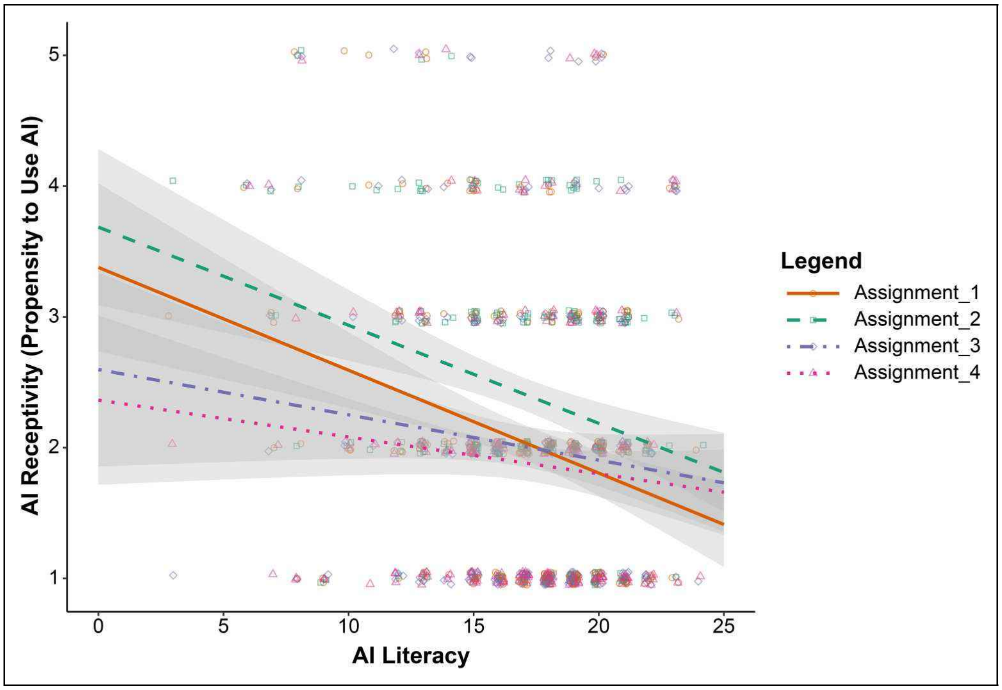

Figure 2. Lower AI Literacy Predicts Higher Propensity to Use AI Across Four Assignments.
Notes: Each line in Figure 2 represents a separate linear regression predicting AI utilization from AI literacy for a single assignment. Shaded areas indicate 95% confidence intervals.

not driven by one assignment in particular, we also ran a mixed-regression model, regressing propensity to use AI to complete each individual task on the literacy scores, including fixed effects for the specific assignment and a random intercept for each participant to account for the repeated measures. This model similarly indicated a negative relationship between AI literacy and AI receptivity ($B = -.05$, 95% CI: [−.08, −.02], $SE = .01$, $t(231.93) = -3.63$, $p < .001$). In Figure 2 we plot separate regressions for each assignment type by AI literacy. As another robustness check, we considered the possibility that people with higher AI literacy have a better understanding of the potential pitfalls of generative AI's output. If so, people with higher AI literacy may not be less receptive to using AI but believe it cannot be used to the same extent. To examine this possibility, we recoded the dependent measure to be a binary variable indicating whether participants would use AI at all. The results showed that those with lower AI literacy were more likely to use AI (at all) compared with those with higher literacy ($B = -.13$, 95% CI: [−.23, −.03], $SE = .05$, $z = -2.65$, $p < .001$).

Overall, this study shows that students with lower AI literacy exhibit greater propensity to use AI to complete their assignments compared with students with higher AI literacy. Notably, our study design provided participants with an example of AI's abilities and held the cost of the AI agent constant (i.e., using the generative AI agent was free). Thus, observed differences in utilization propensity are unlikely to be explained by differences in expectations of what the AI's output would be or the cost of using it.

Study 3: The Predictive Power of AI Literacy Beyond Other Individual Differences

Study 3 explores AI receptivity operationalized through the frequency of past AI usage, a conservative test of the hypothesized relationship, since greater exposure to AI could increase literacy through greater awareness of how and where AI operates. If so, people with more frequent past usage would be those with greater, rather than lesser, knowledge—the opposite of the prediction made by the current theorizing. An additional goal of Study 3 is to measure other individual differences that may systematically vary across people with lower and higher AI literacy to examine the unique predictability of AI literacy. Specifically, we consider the possibility that people with greater AI literacy may have a generally higher desire for self-reliance and thus have a reduced preference for help of any kind. Since it has been shown that one of the impediments to adoption of autonomous products may be a desire to complete a task oneself (De Bellis, Johar, and Poletti 2023), we measured motivation for autonomy. In addition, it is possible that this effect is not specific to AI, and that people with less AI literacy are simply more inclined to adopt technology of any kind. To examine this, we assess participants' technology readiness, a measure of a person's propensity to adopt new technologies (Victorino, Karniouchina, and Verma 2009). Finally, it is possible that AI literacy is simply a proxy of general knowledge, and that this effect stems from people with less general knowledge rather than knowledge of AI specifically. Thus, we administer a test of general knowledge across a variety of subjects to see whether AI literacy provided predictive ability beyond general knowledge.

AI Literacy Measurement

For this study, we developed a new measurement of AI literacy based on previously identified AI literacy "competencies." Specifically, Long and Magerko (2020) identify 17 competencies that they consider to be the core foundations of AI literacy. These 17 competencies were described to both Claude.AI and ChatGPT-4. We then asked these AI tools to create one multiple-choice question for each of the 17 competencies. From these two sets, we selected the question that seemed most representative of the competency, for a final set of 17 questions. Details of the development of this measure and the 17 items are available in Web Appendix B.

Procedure

We posted the survey for 400 participants on Amazon Mechanical Turk via Cloud Research. We received 401 completed responses (4 participants did not complete the survey because they failed the attention check at the beginning of the survey; $M_{age} = 43.9$ years, SD = 13.2; female = 43.9%, male = 54.6%, nonbinary/third gender = .5%, prefer to not say = 1%).

We first measured AI receptivity through participants' reported frequency of usage of AI-based products and services over the previous six months. Specifically, they were asked, "In the last six months, how often have you used …" (1) "A digital image generator (e.g., DALL-E, Midjourney, …)," (2) "An AI-powered productivity tool for note taking, scheduling, inbox management (e.g., Zapier)," (3) "An AI-powered design service to create a website (e.g., Canva)," (4) "An AI-powered health app to meditate, monitor sleep (e.g., Headspace, Sleep.ai)," or (5) "An AI-tool as a writing assistant (e.g., ChatGPT)." These were all measured on the following five-point scale: 1 = "Never," 2 = "Once or twice," 3 = "Occasionally (more than once or twice but not more than once a month), 4 = "Frequently (more than once a month)," and 5 = "On a weekly basis." We combined their responses into an index, coded such that higher numbers on the index indicate greater AI receptivity (Cronbach's $\alpha = .74$).

We next assessed AI literacy with a 17-item measure created by AI, as described previously.

Then, we measured general readiness to adopt technology, desire for self-reliance, and general knowledge. Readiness to adopt technology was measured using the ten-item technology readiness index from Victorino, Karniouchina, and Verma (2009). We measured general attitudes toward self-reliance versus dependence on others through four items assessing motivation for autonomy (from Dahl and Moreau [2007]): "To what extent do the following statements describe you?" (1) "I want to be free to make my own choices," (2) "I want to express myself," (3) "I don't want to be controlled," and (4) "I don't want to feel pressured" (1 = "Not at all," and 7 = "Extremely"). General knowledge was measured with 20 multiple-choice questions assessing the following four topics (5 questions each): history, biology, geography, and literature. Similarly to how the AI literacy test was created for this study, these questions were generated by AI (i.e., ChatGPT-4) to be representative of knowledge in these areas. Finally, we collected demographic information (gender, age, and income).

Results and Discussion

AI literacy. We computed AI literacy by summing the number of correct responses to the 17 questions, such that higher scores indicate higher AI literacy. The median score was 13 (range: 1–17, M = 12.5, SD = 2.28).

AI receptivity. A linear regression on AI receptivity revealed that lower AI literacy was associated with greater frequency of AI usage (B = −.09, 95% CI: [−.13, −.06], SE = .02, $t(399) = -5.73$, $p < .001$). AI literacy was unrelated to age (r = .06, $p = .214$) and income (r = .03, $p = .507$). However, it was related to other factors. It was positively related to technology readiness (r = .27, $p < .001$), motivation for autonomy (r = .12, $p = .015$), general knowledge (r = .51, $p < .001$), and being male (r = .210, $p < .001$). Moreover, each of these factors was associated with past AI usage, such that technology readiness (r = .17, $p < .001$) and being male (r = .10, $p = .037$) positively predicted past usage, and motivation for autonomy (r = −.14, $p = .007$) and general knowledge (r = −.24, $p < .001$) negatively predicted past usage. Therefore, as preregistered, we regressed AI receptivity on AI literacy, technology readiness, motivation for autonomy, general knowledge, and gender (male). The results are reported in Model 2 in Table 4. Even when controlling for these factors, lower AI literacy was associated with greater frequency of AI usage (B = −.11, 95% CI: [−.15, −.07], SE = .02, $t(373) = -5.70$, $p < .001$).

Beyond our preregistered analyses, we performed additional robustness checks. Specifically, lower AI literacy predicted increased AI receptivity in a model in which only factors that negatively predicted AI receptivity were included (e.g., AI literacy, general knowledge, motivation for autonomy) and a model that included all factors and demographics.[2]

This study finds that people with lower AI literacy exhibit greater AI receptivity, as operationalized by frequency of past AI usage. This relationship cannot be explained by people with lower AI literacy liking technology more in general. In fact, people with greater AI literacy are those who indicate being more ready to adopt technology in general, as captured by the technology readiness index. This relationship is also not explained by differences in motivation for autonomy or general knowledge. In fact, a model that includes these factors leaves the predictive ability of AI literacy largely unchanged.

[2] In all studies with available demographic information, a robustness check controlling for demographics remained significant. In addition, the AI literacy measure in this study exhibited skewness levels greater than |1|. The results of this and all studies exhibiting skewness remain significant adjusting for skewness through square root transformation as a robustness check.

Table 4. Results of Study 3.

	Model 1 No Controls	Model 2 Controlling for Factors Negatively Associated with AI Receptivity
AI literacy	−.095*** (.017)	−.110*** (.019)
Technology readiness index[a]		.033*** (.006)
General knowledge		−.027* (.012)
Motivation for autonomy		−.095* (.045)
Gender (male)		.250** (.076)
Intercept	2.952*** (.210)	3.797*** (.333)
Adjusted R^2	.074	.184
Observations	401	379[a]

*$p \leq .05$.
**$p \leq .01$.
***$p \leq .001$.
[a]The technology readiness index had 22 missing observations compared with the full sample. Running Model 1 without these observations does not affect our findings.
Notes: Values reflect unstandardized coefficients.

Study 4: Testing Perceptions of AI as Magical

Study 4 examines our proposed explanation that the lower literacy–higher receptivity link is the result of differential perceptions of AI as magical. We do so by measuring perceptions of AI as magical and examine whether these perceptions explain the relationship between AI literacy and AI receptivity. Additionally, Study 4 distinguishes receptivity toward AI from propensity to execute the task oneself to ensure that this relationship is not due to a desire to derive meaning from doing the task oneself (De Bellis, Johar, and Poletti 2023). Finally, in this study, we recruit a nationally representative U.S. sample.

Procedure

We requested 1,000 participants from the platform Lucid, which uses quota sampling to provide respondents representative of the U.S. population on age, gender, ethnicity, and geographic region. This study included, as preregistered, one IMC ("Please indicate you are taking the time to read the instructions carefully by completing the following simple math question. If Jennie starts with a dozen eggs and uses four eggs to make an omelet, how many eggs does she have left? Please write a numeric answer only [no letters]. Answering this question incorrectly will end the survey."). This was asked at the beginning of the survey, and participants who failed this question did not complete the survey (i.e., we collected no data for them, and incorrect responses did not affect our target sample size). In total, Lucid recruited 1,090 participants, with 84 participants failing to accurately answer the attention check and thus not completing the survey. As such, 1,006 participants successfully completed the study (M_{age} = 45.6 years, SD = 16.8; female = 51.1%, male = 47.5%, nonbinary/third gender = .9%, other = .1%, prefer to not say = .4%[3]).

AI literacy was assessed using the 25-item measure. After completing the AI literacy test, participants rated their agreement with seven statements designed to assess the proposed mediator of beliefs of AI as magical (all on seven-point scales, 1 = "Strongly disagree," and 7 = "Strongly agree"). These statements, presented to them in random order, were: "Artificial intelligence seems magical," "Artificial intelligence seems to conjure answers out of thin air," "Artificial intelligence seems like wizardry," "Artificial intelligence seems supernatural," "Artificial intelligence seems irrational," "Artificial intelligence has an air of mystique," and "If I could see how artificial intelligence works, it would look like a magic show." We averaged the seven statements assessing perceptions of AI as magical into an index (Cronbach's $\alpha = .86$).

We measured AI receptivity via participants' preference to have the target task carried out by a human or an AI algorithm using a seven-point scale (1 = "Definitely prefer the human," 4 = "Indifferent," and 7 = "Definitely prefer the AI algorithm"). For generalizability, we used five new tasks, presented in random order: creating a new recipe, designing a business plan, recommending what to do on a date, onboarding/training new employees, and coming up with a funny story. We averaged preferences across these five tasks into an AI receptivity index (Cronbach's $\alpha = .78$). Finally, we collected participants' gender, age, marital status, employment status, and industry that best described their field of expertise.

Results and Discussion

AI literacy. We computed an AI literacy score for each participant as in previous studies. The median score was 13 (range: 1–24, M = 13.15, SD = 4.63).

AI receptivity. Participants with lower AI literacy were more receptive to have AI complete the tasks (B = −.10, 95% CI: [−.11, −.08], SE = .01, t(1,004) = −10.80, $p < .001$).

Perceptions of AI as magical. Participants with lower AI literacy perceived AI as magical to a greater extent than those with higher AI literacy (B = −.09, 95% CI: [−.11, −.07], SE = .01, t(1,004) = −11.30, $p < .001$).

Mediation. We tested whether differences in AI receptivity as a function of AI literacy were mediated by perceptions of AI as magical. Hayes's (2017) PROCESS Model 4 macro with 10,000 bootstrapped samples supported the proposed mediation, as the 95% confidence interval for the indirect pathway did not include zero (95% CI: [−.03, −.01]).

[3] The demographic information reported here is based on self-reports, which vary slightly from demographic information provided by Lucid. All demographics (self-reported and Lucid-reported) are available in the dataset in ResearchBox.

Overall, Study 4 replicates the relationship between lower AI literacy and greater AI receptivity among a nationally representative U.S. sample generalizing the results to five new tasks. Moreover, these results suggest that differences in AI receptivity are explained by differences in perceptions of AI as magical: those with lower AI literacy perceive AI as more magical, which explains their greater AI receptivity.

Study 5: Examining Serial Mediation Through Perceptions of AI as Magical and Awe-Inspiring

Study 5 builds on Study 4 in four ways. First, Study 5 serves as a conceptual replication operationalizing AI literacy with the 17-item measure, and operationalizing AI receptivity using more general measures of AI receptivity adapted from the Ipsos Global Opinions and Expectations About AI Survey (the same measures used in Study 1). This measurement allows for generalizability, as participants can think of any AI-based products and services that come to their minds. Second, in Study 5, we delve further into the proposed mechanism by measuring perceptions of AI as magical as well as feelings of awe that are inspired by AI, which we argue explain increased receptivity among people with lower AI literacy. Third, we measure perceptions of AI's capability to complete several tasks to test the alternative explanation that people with lower (vs. higher) AI literacy believe that AI is more capable, with differences in capability perceptions explaining differences in receptivity. Finally, we measure fear of the impact of AI on humanity to test the alternative explanation that those with lower (vs. higher) AI literacy are more receptive to AI because they are less worried about the potential negative effects of AI for humanity.

Procedure

We requested 1,000 participants from Prolific Academic. This experiment included, as preregistered, one IMC ("This question is just to make sure you are taking the time to read questions carefully. This is not a trick question. What is 12 − 4? Please use a numeric value only [no letters]."). This was asked at the beginning of the survey, and participants who failed this question did not complete the survey (i.e., we collected no data for them, and incorrect responses did not affect our target sample size). In total, 1,020 people started the survey, with 20 participants failing to accurately answer the attention check and thus not completing the survey. As such, 1,000 participants successfully completed the study (M_{age} = 39.84 years, SD = 13.27; female = 46.5%, male = 50.9%, nonbinary/third gender = 1.8%, prefer to not say = .8%).

We assessed AI literacy with the 17-item test used in Study 3. In all the subsequent sections, participants read that we were interested in their beliefs about AI, and it was stressed that there were no right or wrong answers. We measured AI receptivity using the four general statements assessing AI adoption readiness from the Ipsos Global Opinions and Expectations About AI Survey used in Study 1. However, instead of a binary response, participants indicated their responses on a seven-point scale (1 = "Not at all," and 7 = "Very much"), and responses were combined into an index (Cronbach's α = .83).

Perceptions of AI as magical were measured as in Study 4, although we did not include the statement about AI seeming irrational based on its lower conceptual fit (Cronbach's α = .91). Awe was measured using level of agreement with four items taken from Kautish et al. (2023): "If I think about artificial intelligence that powers products and services…" (1) "I experience feelings of awe," (2) "I feel that I am witness to something grand," (3) "I perceive vastness and I feel my jaw drop," and (4) "I have goosebumps and gasp" (seven-point scale: 1 = "Strongly disagree," and 7 = "Strongly agree"; Cronbach's α = .92). Fear of AI was measured with four items taken from Cave and Dihal (2019). Specifically, participants indicated their level of agreement with the following statements: (1) "I fear that artificial intelligence will make us lose our humanity," (2) "I fear that artificial intelligence will turn against humans," (3) "I fear that artificial intelligence will make humans become obsolete," and (4) "I fear that artificial intelligence will make humans lose the ability to connect with each other" (seven-point scale: 1 = "Not at all," and 7 = "Very much"; Cronbach's α = .87). Finally, we measured AI capability by asking respondents to what extent they thought that an AI algorithm was capable of executing each of the following tasks: (1) composing a song, (2) recommending a gift, (3) playing the piano, (4) writing a news article, (5) driving a car, (6) piloting a plane, and (7) recommending music (seven-point scale: 1 = "Definitely incapable," and 7 = "Definitely capable"). Note that these seven tasks are a subset of the tasks used in Castelo, Bos, and Lehmann (2019; Study 1). The seven measures assessing AI capability perceptions cohered well (Cronbach's α = .83) and were combined. Finally, we collected gender, age, and income.

Results and Discussion

AI literacy. We computed an AI literacy score for each participant as in previous studies. The median score was 13 (range: 1–17, M = 12.55, SD = 2.28).

AI receptivity. As predicted, participants with lower AI literacy scores indicated greater AI receptivity (B = −.04, 95% CI: [−.07, −.01], SE = .02, t(998) = −2.25, p = .024).

Perceptions of AI as magical. Participants with lower AI literacy perceived AI as more magical than those with higher AI literacy (B = −.13, 95% CI: [−.17, −.10], SE = .02, t(998) = −6.83, $p < .001$).

Experience of awe. Participants with lower AI literacy reported experiencing greater feelings of awe when thinking about AI powering products and services (B = −.14, 95% CI: [−.18, −.10], SE = .02, t(998) = −6.79, $p < .001$).

Serial mediation. Before examining serial mediation, we conducted a factor analysis on the set of measures assessing perceptions of AI as magical and experience of awe to examine the distinctiveness of the constructs. Both exploratory and confirmatory factor analysis suggested that these factors were related but separate constructs. Serial mediation using Hayes's (2017) PROCESS Model 6 macro with 10,000 bootstrapped samples provided support for the proposed mediation (95% CI: [−.05, −.02]).

Capability of AI. In contrast to the alternative explanation that people with lower AI literacy perceive AI as more capable, we found a positive association such that those with lower AI literacy perceived AI as less capable (B = .07, 95% CI: [.04, .10], SE = .01, t(998) = 4.94, $p < .001$).

Fear of AI. In contrast to the alternative explanation that people with lower AI literacy are less fearful of AI's impact on humanity, we found a negative association such that those with lower AI literacy were more fearful of AI than those with higher AI literacy (B = −.09, 95% CI: [−.13, −.04], SE = .02, t(998) = −3.56, $p < .001$).

Study 5 provides more evidence in support of our proposed account. Differences in AI receptivity across AI literacy are serially mediated by perceptions of AI as magical and the awe experienced in thinking about AI powering products and services. Notably, these results cast doubt on the alternative explanation of perceptions of AI capability. Those with lower AI literacy rated AI as less capable of executing tasks. This suggests that people with lower AI literacy are more receptive toward AI despite believing it is less capable. Moreover, Study 5 suggests that differences in AI receptivity are not a function of fear of AI, since those with lower AI literacy are more fearful of AI's potential impact on humanity. As such, including capability and fear in the analysis increases rather than decreases the relationship between AI literacy and AI receptivity. These findings suggest that people with lower AI literacy are more receptive to AI-based products and services in spite of, rather than because of, their fear about AI and beliefs in AI's capabilities. In line with this interpretation, in a supplemental study reported in Web Appendix D, we find directional evidence that encouraging people to base their decisions on capability and efficiency weakens the link between AI literacy and AI receptivity.

Study 6: AI Literacy and AI Receptivity as a Function of Task Characteristics

We have suggested that differential perceptions of AI as magical across levels of AI literacy result from differences in how people interpret AI's ability to execute tasks once thought to require exclusively human attributes. Generally, viewing AI as capable of executing such tasks fosters a perception of magic; however, individuals with higher AI literacy recognize that AI does not need humanlike attributes to accomplish these tasks, which dampens perceptions of AI as magical. The lower literacy–higher receptivity link emerged using general measures of AI adoption readiness (Studies 1 and 5), which suggest that the lower literacy–higher receptivity link does not artificially rely on our choice of tasks. However, our line of reasoning suggests that differences in AI receptivity should be moderated when tasks do not require distinctly human attributes (e.g., objective tasks such as data analysis). By exploring process through the potential moderating impact of task characteristics, Study 6 aims to address the limitations of the approach of using mediation to examine process (e.g., Bullock, Green, and Ha 2010; Imai et al. 2011). Finally, we counterbalance whether the AI literacy test comes before or after the measure of AI receptivity to explore potential order effects.

Procedure

We posted the survey for 1,000 participants on Amazon Mechanical Turk via CloudResearch. We received 1,184 responses[4] (13 participants did not complete the survey because they failed the attention check at the beginning of the survey). We omitted responses from 32 participants who failed the IMC at the end of the survey, 5 participants who skipped an attention check question, and an additional 35 participants who indicated answering the questions randomly (2 of whom also indicated not trying their best), leaving us with a final sample of 1,099 participants (M_{age} = 46.3 years, SD = 60.3; female = 50.1%, male = 48.2%, nonbinary/third gender = .7%, prefer to not say = 1%).

AI literacy was assessed with the 25-item measure used in Study 2. Participants' AI receptivity, the dependent measure, relied on a paradigm used by Castelo, Bos, and Lehmann (2019; Study 1). Specifically, participants read that we wanted them to consider various tasks that could be carried out either by a person or automated and carried out by an AI algorithm and to rate their preference to have the key decision for that task carried out by a human versus an AI algorithm. For each of 26 tasks, briefly described and presented in random order (e.g., predicting joke funniness, hiring and firing employees, recommending a romantic partner, predicting student performance), participants expressed a preference to have the target task carried out by a human or an AI algorithm using a seven-point scale (1 = "Definitely prefer the human," 4 = "Indifferent," and 7 = "Definitely prefer the AI algorithm").

We then collected demographic variables (gender, age, and household income), political ideology, and religiosity. The survey ended with three attention checks. One was the same IMC used in Study 2. We included two attention checks: participants indicated (1) if they responded to the best of their abilities and (2) if they responded randomly at any point during the survey (yes/no binary answer for both).

Results and Discussion

AI literacy. We computed the AI literacy score for each participant as in previous studies. The median score was 19 (range: 4–25, M = 18.49, SD = 3.81).

[4] There was an error with the survey completion code, which may have led to the higher-than-expected number of survey completions.

AI receptivity. As preregistered, we calculated AI receptivity by averaging preferences across the 26 tasks into an index, coded such that higher numbers on the index indicate greater AI receptivity (Cronbach's $\alpha = .92$). A linear regression on average AI receptivity revealed that lower AI literacy was associated with greater relative preference for AI task execution ($B = -.05$, 95% CI: [−.06, −.03], $SE = .01$, $t(1{,}097) = -6.39$, $p < .001$). We also explored whether this relationship depended on whether we assessed AI literacy before or after AI receptivity using order as an additional factor in the analysis. Regressing AI receptivity on AI literacy, order, and their interaction revealed only a main effect of AI literacy ($B = -.04$, 95% CI: [−.06, −.02], $SE = .01$, $t(1{,}095) = -3.50$, $p < .001$). Neither the effect of order ($B = .24$, 95% CI: [−.30, .78], $SE = .28$, $t(1{,}095) = .86$, $p = .389$) nor the interaction ($B = -.02$, 95% CI: [−.04, .01], $SE = .01$, $t(1{,}095) = -1.05$, $p = .295$) were significant.[5]

Task characteristics. Although we found an overall negative relationship, such that people with lower AI literacy were more likely to prefer tasks be completed by AI, we explored whether this difference depended on the level of objectivity required to complete each task. Subjective tasks are more likely to typically require more distinctively human attributes (e.g., having a sense of humor), whereas objective tasks are more likely to require attributes known to be shared by humans and AI (e.g., doing computations). As preregistered, we asked a separate sample of 50 respondents from Amazon Mechanical Turk (following the inclusion criteria from the main study, 47 participants were included in the analysis) to rate each of the 26 tasks on the perceived degree of subjectivity versus objectivity using definitions provided by Castelo, Bos, and Lehmann (2019): "For each task, we would like you to tell us if you think that the task is mostly subjective or mostly objective. Subjective tasks are tasks that are best carried out by relying on subjective feelings, intuition, and gut reactions. Objective tasks are tasks that are best carried out by relying on objective data, mathematical/statistical relationships, and objective measures." Higher ratings indicated higher perceived objectivity required to complete the task and were mean-centered. We then ran a mixed-regression model, regressing AI preference for each individual task on participants' mean-centered AI literacy score, the mean-centered objectivity score of the given task, and their interaction, with a random intercept for each participant to account for the repeated measures. This model replicated the negative relationship between AI literacy and AI receptivity ($B = -.05$, 95% CI: [−.06, −.03], $SE = .01$, $t(1{,}096.61) = -6.38$, $p < .001$). In line with previous research, objectivity positively predicted AI receptivity ($B = .01$, 95% CI: [.01, .01], $SE < .01$, $t(27{,}456.70) = 16.87$, $p < .001$). However, there was evidence of a significant interaction ($B = .001$, 95% CI: [.001, .001], $SE < .01$, $t(27{,}456.53) = 7.59$, $p < .001$). Although those with lower AI literacy had generally greater AI receptivity, this difference was smaller among tasks requiring more objectivity (Figure 3).

To more directly test our theory, we ran a subsequent posttest to evaluate whether each of the 26 tasks was believed to require more distinctively human attributes versus attributes believed to be shared among humans and AI. One hundred participants read through two lists of attributes (distinctly human vs. shared attributes) taken from Santoro and Monin (2023). Then, for each of the 26 tasks, participants indicated, on a seven-point scale, whether it was more important for an agent executing the task to have the attributes that are distinctly human or those that are shared between humans and AI (1 = "Definitely have distinctively human attributes," and 7 = "Definitely have attributes shared between humans and AI"). We then ran a mixed-regression model, regressing AI preference for each individual task on participants' mean-centered AI literacy score, the mean-centered human versus shared attributes score of the given task, and their interaction, with a random intercept for each participant to account for the repeated measures. This model replicated the negative relationship between AI literacy and AI receptivity ($B = -.05$, 95% CI: [−.06, −.03], $SE = .01$, $t(1{,}096.81) = -6.38$, $p < .001$). Human versus shared attributes positively predicted AI receptivity ($B = .86$, 95% CI: [.84, .88], $SE = .01$, $t(27{,}456.56) = 93.94$, $p < .001$). Importantly, there was evidence of a significant interaction ($B = .04$, 95% CI: [.03, .04], $SE < .01$, $t(27{,}456.44) = 16.19$, $p < .001$). Similar to task objectivity, the relationship between lower AI literacy and greater AI receptivity was attenuated among tasks perceived to require less distinctly human attributes.

This study replicates the relationship between lower AI literacy and higher AI receptivity, in a context in which the AI receptivity measure is unrelated to participants' preference to complete the tasks themselves. This relationship is not contingent on whether AI literacy is assessed before or after participants provide their relative preference for AI, indicating that the relationship between AI literacy and AI receptivity is not dependent on having AI knowledge made salient. Moreover, in line with our proposed theorization, the relationship between AI literacy and AI receptivity is more pronounced among tasks that are subjective and those that typically require more distinctively human attributes, which are more likely to evoke perceptions of AI as magical.

In two separate posttests, which we report in full detail in Web Appendix E, we test two alternative accounts for the lower literacy–higher receptivity link: that people with lower AI literacy perceive the use of AI to complete tasks as more ethical, or that they perceive AI as more capable than those with higher AI literacy. We explore these alternative explanations in two posttests, one focused on perceived ethicality and one focused on perceived capability of having AI complete each of the 26 tasks used in Study 6. These posttests show that, despite exhibiting greater AI receptivity, people with lower AI literacy perceive AI as less ethical than people with higher AI literacy, refuting an account based on perceived ethicality. Furthermore, people with lower AI literacy do not perceive AI as being more capable than

[5] Results remain consistent in a linear mixed model that includes fixed effects for each task and a random effect for each subject.

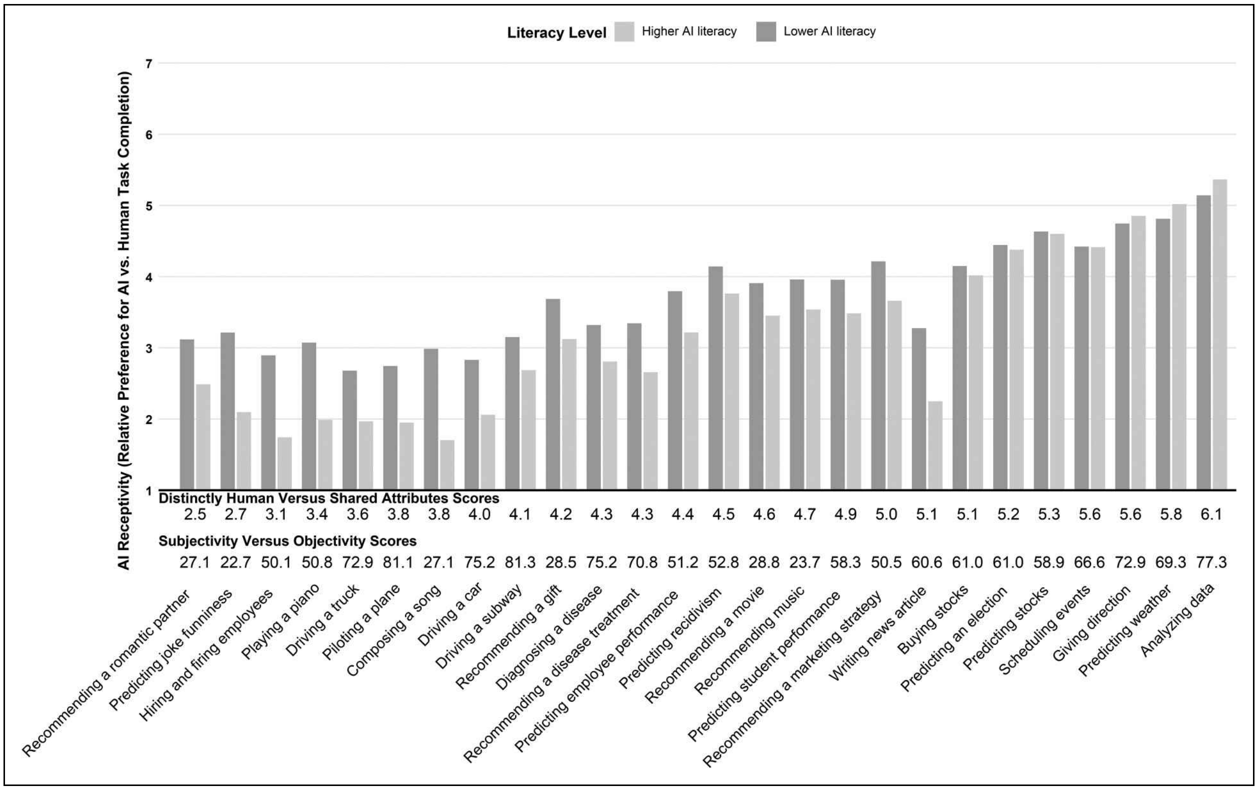

Figure 3. Results of Study 6.

Notes: Figure 3 displays preference for AI over human task completion (on the y-axis) across 26 tasks varying in degree of association with distinctly human versus shared attributes and perceived subjectivity versus objectivity (on the x-axis), ordered in descending degree of association with distinctly human attributes; thus, lower numbers indicate tasks associated with more distinctly human attributes and higher numbers indicate tasks associated with more shared attributes). Lower (higher) AI literacy represents predicted preference for AI task execution at 1 SD below (above) the mean AI literacy score.

people with higher AI literacy, which is inconsistent with an account based on perceived capability.

Study 7: Moderation by Tasks' Propensity to Elicit Perceptions of AI as Magical

As Study 6 demonstrated, the lower literacy–higher receptivity link appears to be moderated by the extent to which people believe the task requires distinctly human attributes. However, the tasks in that study differed in several ways. Thus, to more directly test our theory, Study 7 holds constant the category of tasks while manipulating, between subjects, the extent to which the target tasks require distinctly human attributes or not since tasks that are not perceived to require human attributes to execute the tasks should not elicit feelings of AI as magical regardless of AI literacy levels.

Procedure

This study was run in partnership with Toluna, a market research company with a panel of consumers who complete company surveys for rewards points. We requested 1,000 completed surveys. This experiment included, as preregistered, two attention checks at the beginning of the survey, and only participants who answered correctly were allowed to complete the survey. In total, 1,018 people passed the attention check questions, and we received completed responses from 1,006 participants ($M_{age} = 51.35$ years, $SD = 16.55$; female = 77.3%, male = 22.7%).

This study entailed two between-subjects conditions that varied the types of attributes (distinct attributes vs. shared attributes) needed to complete tasks. To create these conditions, we provided AI (Claude 3.5 Sonnet) a list of distinctly human attributes and a list of shared attributes as identified by Santoro and Monin (2023) and asked it to create pairs of tasks that come from the same general category (e.g., news reporting), where one of the tasks would require more distinctively human attributes and the other would require more shared attributes for task execution. From these, we created a set of ten pairs of tasks covering cooking, art, architecture, news reporting, fashion, advertising, customer experience innovation, customer engagement and retention, health care, and hospitality management, where the two tasks varied in the extent to which they required distinctly human attributes. For example, in the news reporting category, the task that required more distinctly human attribute was "writing an opinionated editorial piece for a news outlet," and the one that required more shared attributes was "fact-checking and verifying sources for a news outlet." Differences in perceptions of the importance of distinctly human versus shared attributes for task execution across the pairs of tasks were validated in a pretest with 102 participants on Cloud Connect ($t(101) = -24.63$, $p < .001$; see Web Appendix F for details). For each of the ten tasks, participants rated the extent to which they would prefer to have the key decision for that task carried out by a human or by an AI algorithm on a seven-point scale (1 = "Definitely prefer the human," and 7 = "Definitely prefer the AI algorithm"). We assessed AI literacy with the 17-item test used in Studies 3 and 5. Other measures unrelated to the present research were included for the interests of Toluna (the company running the survey). These were administered after assessing AI receptivity and AI literacy. Demographic information was provided by Toluna.

Results and Discussion

AI literacy. We computed an AI literacy score for each participant as in previous studies. The median score was 10 (range: 1–17, $M = 9.99$, $SD = 2.94$).

AI receptivity. As preregistered, we calculated AI receptivity by averaging preferences across the ten tasks into an index, coded such that higher numbers on the index indicate greater AI receptivity (Cronbach's $\alpha = .89$). We regressed this receptivity rating on their AI literacy score (mean-centered), condition (0 = shared, 1 = distinct), and their interaction. People were less receptive to AI in the distinct condition ($B = -1.17$, 95% CI: [−1.34, −.99], $SE = .09$, $t(1{,}002) = -13.15$, $p < .001$), and AI literacy score positively predicted AI receptivity ($B = .08$, 95% CI: [.03, .12], $SE = .02$, $t(1{,}002) = 3.57$, $p < .001$). Critically, we found the expected interaction ($B = -.19$, 95% CI: [−.25, −.13], $SE = .03$, $t(1{,}002) = -6.38$, $p < .001$).[6] Simple slope analysis showed a full reversal of the AI literacy effect across conditions. In the distinct attributes condition, we replicated our previous studies such that people with lower AI literacy indicated greater AI receptivity ($B = -.12$, 95% CI: [−.16, −.08], $SE = .02$, $t(1{,}002) = -5.43$, $p < .001$). However, this relationship reversed in the shared attributes condition, with greater AI literacy indicating more receptivity ($B = .08$, 95% CI: [.03, .12], $SE = .02$, $t(1{,}002) = 3.57$, $p < .001$) (see Figure 4).

This study provides strong support for our theorizing about the role of distinctly human attributes in driving the relationship between AI literacy and AI receptivity. When considering tasks that require distinctly human attributes, we replicated our previous findings—those with lower AI literacy exhibited greater receptivity toward AI executing such tasks. However, this relationship reversed for tasks associated with shared attributes between humans and AI. This reversal helps reconcile the fact that people with higher literacy perceive AI as more capable, but they generally have lower receptivity. Absent the "magical" element, those with higher AI literacy demonstrate greater receptivity.

[6] Results remain consistent in a linear mixed model that includes fixed effects for each task and a random effect for each subject.

General Discussion

AI is rapidly transforming our society. Research to date has typically focused on situational processes to understand differences in AI receptivity. Our research highlights the importance of considering systematic differences across people, demonstrating that differences in AI literacy across people predict differences in AI receptivity. In contrast to what was predicted in four separate forecasting surveys and a survey of executives, seven studies using diverse samples supported the predicted lower literacy–higher receptivity link.

People with lower AI literacy were more receptive to personally using AI to execute tasks (Studies 2 and 3). This link was not due to feelings of personally autonomy, as it generalized to preferences for tasks to be executed by AI versus a human (Studies 4, 6, and 7). Moreover, this link emerged when receptivity was operationalized as general measures of AI adoption readiness (Studies 1 and 5), suggesting that this relationship exists in the absence of any specific task being highlighted. We explain this link through perceptions of AI as magical. When people think about AI executing tasks that are typically associated with requiring distinctly human attributes, those with lower AI literacy are more likely to believe that AI actually possesses these humanlike attributes, leading to greater perceptions of AI as magical and feelings of awe. In line with this theorizing, perceptions of AI as magical and ensuing feelings of awe explained the lower literacy–higher receptivity link (Studies 4 and 5). Moreover, the lower literacy–higher receptivity link was moderated when participants explicitly considered tasks that were assumed to not require distinctly human attributes (Studies 6 and 7). This link was not explained by participants' general knowledge, their attitudes toward self-reliance versus dependence, or their general readiness to adopt technology (Study 3). Further, our results suggest that this link emerged despite, rather than because of, differences in perceptions of AI capability (Study 5 and Study 6 posttest), ethicality (Study 6 posttest), and fear of its impact on humanity (Study 5).

Theoretical and Managerial Implications

First, we contribute to research on the psychological responses to AI (Puntoni et al. 2021; Valenzuela et al. 2024). Whereas prior work in this area has predominantly investigated situational factors, we focus on the overlooked question of how differences across consumers influence AI receptivity. In doing so, our research suggests a need for future research to identify other such individual-level differences. By considering individual traits, attitudes, and predispositions, researchers can better comprehend the underlying motivations and preferences that shape individuals' propensity to use AI-based products and services, enabling better development of predictive models and improving customization and targeting of AI solutions.

Second, our theorizing offers a means of reconciling seemingly conflicting findings regarding people's responses to new technologies. Prior work has documented both cases in which consumers exhibit favorable responses to new technologies

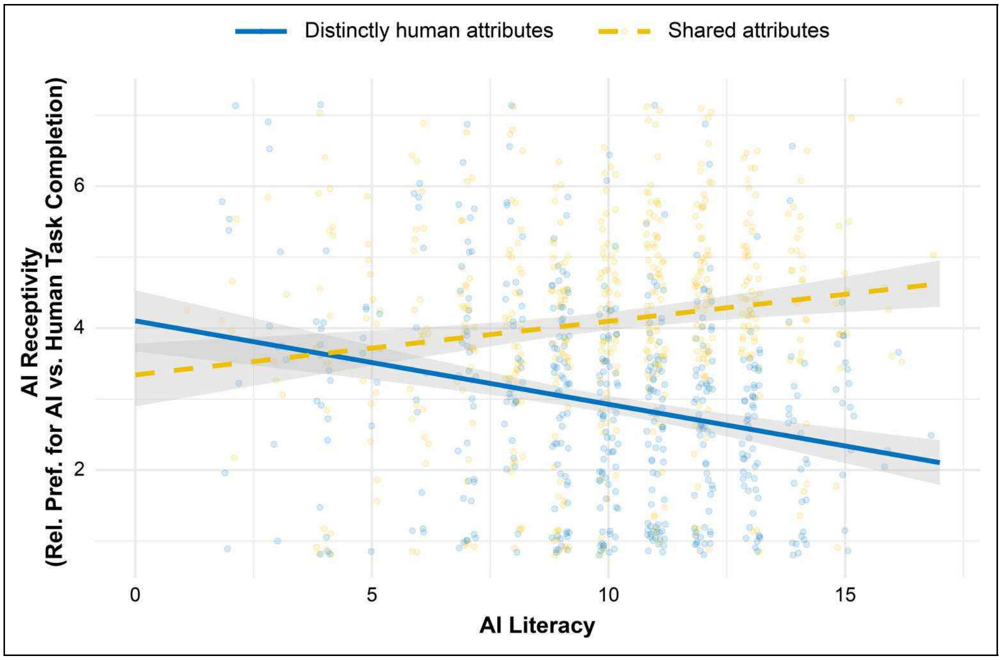

Figure 4. Reversal of the Lower Literacy–Higher Receptivity Link.
Notes: Figure 4 displays the relationship between AI literacy and AI receptivity (relative preference for AI vs. human task completion) across tasks varying in the extent to which they require distinctly human versus shared attributes. Each regression line represents the average trend within each attribute type, with its slope reflecting the impact of AI literacy on AI receptivity for that type. Shaded areas indicate 95% confidence intervals.

(i.e., "automation bias" or "algorithm appreciation"; e.g., Longoni and Cian 2022; Mosier and Skitka 2018) and cases in which consumers are averse to new technologies (i.e., "algorithm aversion"; e.g., Dietvorst, Simmons, and Massey 2015). By identifying perceptions of AI as magical and the accompanying feelings of awe as important yet overlooked predictors of AI receptivity, we suggest that whether consumers will be receptive to new technologies will hinge on individual-level (e.g., literacy about the new technology) and situational (e.g., whether the new technology performs tasks associated with distinctly human attributes) factors that affect perceptions of new technologies as magical.

Third, our research offers important and timely managerial implications, as it suggests that companies' product management and marketing efforts should be tailored to different group of consumers. While the insurance company executives we surveyed unanimously believed that targeting consumers with higher AI literacy would be more effective for marketing AI products, our findings suggest otherwise. New AI-based products and services may be more readily adopted among those with lower literacy, particularly for applications that are believed to be associated with distinctly human attributes.

To the extent that companies target low-AI-literacy consumers, it is important that they design AI products and services with the welfare and needs of these consumers in mind. Indeed, these recommendations should not be taken as a call to exploit low-AI-literacy consumers and their potential vulnerability to new AI initiatives. Instead, these insights highlight an opportunity to develop and market AI products for a potentially overlooked yet highly receptive audience.

Further, the results of Study 7 suggest that there are opportunities to target not only those with lower AI literacy but also those with higher AI literacy, as long as these two groups of consumers are targeted differently. Specifically, our findings suggest that marketing efforts aiming to make AI more accessible to a broader audience by demystifying how AI works may inadvertently reduce consumers' receptivity toward AI by making it seem less magical. Thus, marketing efforts targeting low-AI-literacy consumers may benefit from promoting an aura of magic around AI by emphasizing how the AI powering the respective product or service emulates human characteristics. Conversely, marketing efforts targeting high-AI-literacy consumers may benefit from highlighting how their AI-based products and services execute tasks that are based on characteristics shared between humans and AI.

Limitations and Future Research

Meaningfully changing people's AI literacy is difficult, especially in a short time frame. It typically requires motivation from the learner as well as the ability to calibrate materials and tailor content to a person's current AI literacy (Long and Magerko 2020). Thus, the current research relied on measuring AI literacy and demonstrating the proposed process through mediation and moderation. However, future research should examine how AI literacy initiatives causally affect AI receptivity and whether the content and style of AI education influence this relationship.

The lower literacy–higher receptivity link emerged when AI literacy was measured using a human-constructed and validated measure (Studies 2, 4, and 6), an AI-constructed measure based on AI competencies identified in previous research (Studies 3, 5, and 7), and when based on country-level differences in AI expertise associated with AI-related proficiencies and skills (Study 1). One may wonder whether AI literacy should be conceptualized as a single factor or whether there are important subfactors that should be analyzed separately. Factor analyses performed on both the 25-item measure and the 17-item measure produced scree plots that supported a single-factor solution. Further, our results are robust to separately analyzing questions assessing technical understanding, limitations, and ethical considerations (see Web Appendix C). However, we note that these subfactors are highly correlated. Future research may develop measures of subcomponents of AI literacy that are more statistically distinct.

Our results suggest that as people become more AI literate over time, general AI receptivity may decrease. Thus, our findings suggest that until capability considerations outweigh AI receptivity fueled by perceptions of AI as magical, there may be unintended consequences of policy makers' efforts to educate the public about AI. Of course, other factors—such as shifting norms and expectations about AI adoption or improvements in AI capabilities—may affect and change the nature of the relationship between literacy and receptivity over time. Future research could examine the lower literacy–higher receptivity link longitudinally.

Acknowledgments

The authors thank Pamela Dunaway, VP Global Marketing at Toluna, and her team for their invaluable assistance in conducting Study 7. The authors also thank the *JM* review team for their highly constructive feedback, which greatly enhanced this manuscript. This work also benefited from feedback from seminar participants at Bocconi University; Georgetown University; HEC, Institute for Advanced Study in Toulouse, Max Planck Institute for Human Development; Reichman University; University of Arizona; University of California, San Diego; University of Oxford; and Yeshiva University.

Authors Contributions

All authors contributed equally and are listed in reversed alphabetical order.

Coeditor

Cait Lamberton

Associate Editor

Stefano Puntoni

Declaration of Conflicting Interests

The author(s) declared no potential conflicts of interest with respect to the research, authorship, and/or publication of this article.

Funding

The author(s) received no financial support for the research, authorship, and/or publication of this article.

ORCID iDs

Stephanie M. Tully https://orcid.org/0000-0002-7179-7729
Chiara Longoni https://orcid.org/0000-0002-4945-4957
Gil Appel https://orcid.org/0000-0001-7859-3908

References

aiEDU (2023), "AI in 5 Minutes," (accessed June 6, 2023), https://www.aiedu.org/ai-in-five.

Braxton, Greg (1998), "Disillusioned Magicians in Protest," *Washington Post* (May 2), https://www.washingtonpost.com/archive/lifestyle/1998/05/02/disillusioned-magicians-in-protest/2d147ae1-8d1c-4e3f-9c20-84f760c59a81/.

Brynjolfsson, Erik, Danielle Li, and Lindsey R. Raymond (2025), "Generative AI at Work," *Quarterly Journal of Economics* (published online February 4), https://doi.org/10.1093/qje/qjae044.

Bullock, John G., Donal P. Green, and Shang E. Ha (2010), "Yes, but What's the Mechanism? (Don't Expect an Easy Answer)," *Journal of Personality and Social Psychology*, 98 (4), 550–58.

Castelo, Noah, Maarten W. Bos, and Donald R. Lehmann (2019), "Task-Dependent Algorithm Aversion," *Journal of Marketing Research*, 56 (5), 809–25.

Cave, Stephen and Kanta Dihal (2019), "Hopes and Fears for Intelligent Machines in Fiction and Reality," *Nature Machine Intelligence*, 1 (2), 74–78.

Clarke, Arthur C. (1968), "Clarke's Third Law on UFO's," *Science*, 159 (3812), 255.

Cotter, Kelley and Bianca C. Reisdorf (2020), "Algorithmic Knowledge Gaps: A New Dimension of (Digital) Inequality," *International Journal of Communication*, 14, 745–65.

Dahl, Darren W. and C. Page Moreau (2007), "Thinking Inside the Box: Why Consumers Enjoy Constrained Creative Experiences," *Journal of Marketing Research*, 44 (3), 357–69.

De Bellis, Emanuel, Gita V. Johar, and Nicola Poletti (2023), "Meaning of Manual Labor Impedes Consumer Adoption of Autonomous Products," *Journal of Marketing*, 87 (6), 949–65.

Dietvorst, Berkeley J., Joseph P. Simmons, and Cade Massey (2015), "Algorithm Aversion: People Erroneously Avoid Algorithms After Seeing Them Err," *Journal of Experimental Psychology: General*, 144 (1), 114–26.

Digital Moment (2023), "What Is Artificial Intelligence, Algorithm & Data Literacy," (accessed June 6, 2023), https://algorithmliteracy.org/.

Dogruel, Leyla, Philipp Masur, and Sven Joeckel (2022), "Development and Validation of an Algorithm Literacy Scale for Internet Users," *Communication Methods & Measures*, 16 (2), 115–33.

Ekroll, Vebjørn, Bilge Sayim, and Johan Wagemans (2017), "The Other Side of Magic: The Psychology of Perceiving Hidden Things," *Perspectives on Psychological Science*, 12 (1), 91–106.

Epley, Nicholas, Adam Waytz, and John T. Cacioppo (2007), "On Seeing Human: A Three-Factor Theory of Anthropomorphism," *Psychological Review*, 114 (4), 864–86.

Hargittai, Eszter, Jonathan Gruber, Teodora Djukaric, Jaelle Fuchs, and Lisa Brombach (2020), "Black Box Measures? How to Study People's Algorithm Skills," *Information, Communication & Society*, 23 (5), 764–75.

Hayes, Andrew F. (2017), *Introduction to Mediation, Moderation, and Conditional Process Analysis: A Regression-Based Approach*. Guilford Publications.

Hermann, Erik and Stefano Puntoni (2024), "Artificial Intelligence and Consumer Behavior: From Predictive to Generative AI," *Journal of Business Research*, 180, 114720.

Hermann, Erik, Gizem Yalcin Williams, and Stefano Puntoni (2024), "Deploying Artificial Intelligence in Services to AID Vulnerable Consumers," *Journal of the Academy of Marketing Science*, 52, 1431–51.

Imai, Kosuke, Luke Keele, Dustin Tingley, and Teppei Yamamoto (2011), "Unpacking the Black Box of Causality: Learning About Causal Mechanisms from Experimental and Observational Studies," *American Political Science Review*, 105 (4), 765–89.

Ipsos (2022), "Global Opinions and Expectations About Artificial Intelligence: A Global Advisor Survey," (January), https://www.ipsos.com/sites/default/files/ct/news/documents/2022-01/Global-opinions-and-expectations-about-AI-2022.pdf.

Kabakus, Ahmet Kamil, Ekrem Bahcekapili, and Ahmet Ayaz (2023), "The Effect of Digital Literacy on Technology Acceptance: An Evaluation on Administrative Staff in Higher Education," *Journal of Information Science* (published online March 15), https://doi.org/10.1177/01655515231160028.

Kautish, Pradeep, Mujahid Siddiqui, Aaliyah Siddiqui, Veenu Sharma, and Safiya Mukhtar Alshibani (2023), "Technology-Enabled Cure and Care: An Application of Innovation Resistance Theory to Telemedicine Apps in an Emerging Market Context," *Technological Forecasting and Social Change*, 192, 122558.

Keltner, Dacher and Jonathan Haidt (2003), "Approaching Awe, a Moral, Spiritual, and Aesthetic Emotion," *Cognition & Emotion*, 17 (2), 297–314.

Latzer, Michael, Noemi Festic, and Kiran Kappeler (2021), "Awareness of Algorithmic Selection and Attitudes in Switzerland," Report 2 from The Significance of Algorithmic Selection for Everyday Life: The Case of Switzerland, SSRN, https://ssrn.com/abstract=3869792.

Liu, Nicole T.Y., Samuel N. Kirshner, and Eric T.K. Lim (2023), "Is Algorithm Aversion WEIRD? A Cross-Country Comparison of Individual-Differences and Algorithm Aversion," *Journal of Retailing & Consumer Services*, 72, 103259.

Long, Duri and Brian Magerko (2020), "What Is AI Literacy? Competencies and Design Considerations," *2020 CHI Conference on Human Factors in Computing Systems*, https://doi.org/10.1145/3313831.3376727.

Longoni, Chiara and Luca Cian (2022), "Artificial Intelligence in Utilitarian vs. Hedonic Contexts: The 'Word-of-Machine' Effect," *Journal of Marketing*, 86 (1), 91–108.

Maslej, Nestor, Loredana Fattorini, Erik Brynjolfsson, John Etchemendy, Katrina Ligett, Terah Lyons, James Manyika, Helen Ngo, Juan Carlos Niebles, Vanessa Parli, Yoav Shoham, Russell Wald, Jack Clark, and Raymond Perrault (2023), "The AI Index 2023 Annual Report," Institute for Human-Centered AI, Stanford University.

McCarthy, John, Marvin L. Minsky, Nathaniel Rochester, and Claude E. Shannon (2006), "A Proposal for the Dartmouth Summer Research Project on Artificial Intelligence," *AI Magazine*, 27 (4), 12.

Mosier, Kathleen L. and Linda J. Skitka (2018), "Human Decision Makers and Automated Decision Aids: Made for Each Other?" in *Automation and Human Performance: Theory and Applications*, Raja Parasuraman and Mustapha Mouloua, eds. CRC Press.

National University of Singapore (2023), "Digital Literacy: Home," (last updated September 26), https://libguides.nus.edu.sg/digitalliteracy.

Pennycook, Gordon, Tyrone D. Cannon, and David G. Rand (2018), "Prior exposure increases perceived accuracy of fake news," *Journal of Experimental Psychology: General*, 147 (12), 1865.

Puntoni, Stefano, Rebecca W. Reczek, Markus Giesler, and Simona Botti (2021), "Consumers and Artificial Intelligence: An Experiential Perspective," *Journal of Marketing*, 85 (1), 131–51.

Quesnel, Denise and Bernhard E. Riecke (2018), "Are You Awed Yet? How Virtual Reality Gives Us Awe and Goose Bumps," *Frontiers in Psychology*, 9, 403078.

Rader, Emilee and Rebecca Gray (2015), "Understanding User Beliefs About Algorithmic Curation in the Facebook News Feed," in *Proceedings of the 33rd Annual ACM Conference on Human Factors in Computing Systems*. Association for Computing Machinery, 173–82.

Richter, Frank-Jürgen and Gunjan Sinha (2020), "Why Do Your Employees Resist New Tech," *Harvard Business Review* (August 21), https://hbr.org/2020/08/why-do-your-employees-resist-new-tech.

Santoro, Erik and Benoît Monin (2023), "The AI Effect: People Rate Distinctively Human Attributes as More Essential to Being Human After Learning About Artificial Intelligence Advances," *Journal of Experimental Social Psychology*, 107, 104464.

Sharkey, Noel and Amanda Sharkey (2006), "Artificial Intelligence and Natural Magic," *Artificial Intelligence Review*, 25, 9–19.

Siles, Ignacio, Luciana Valerio-Alfaro, and Ariana Meléndez-Moran (2022), "Learning to Like TikTok… And Not: Algorithm Awareness as Process," *New Media & Society*, 26 (10), https://doi.org/10.1177/14614448221138973.

Sulmont, Elisabeth, Elizabeth Patitsas, and Jeremy R. Cooperstock (2019), "Can You Teach Me to Machine Learn?" in *Proceedings of the 50th ACM Technical Symposium on Computer Science Education*. Association for Computing Machinery, 948–54.

Tomaino, Geoffrey, Hisham Abdulhalim, Pavel Kireyev, and Klaus Wertenbroch (2022), "Denied by an (Unexplainable) Algorithm: Teleological Explanations for Algorithmic Decisions Enhance Customer Satisfaction," INSEAD Working Paper, https://sites.insead.edu/facultyresearch/research/doc.cfm?&did=66974.

Topol, Eric (2019), *Deep Medicine: How Artificial Intelligence Can Make Healthcare Human Again*. Hachette UK.

Tortoise Media (2023), "Global AI Index," (accessed March 18, 2023), https://www.tortoisemedia.com/data/global-ai.

United Nations Development Programme (2022), "Human Development Index (HDI)," (accessed August 7, 2023), https://hdr.undp.org/data-center/human-development-index#/indicies/HDI.

Valenzuela, Ana, Stefano Puntoni, Donna Hoffman, Noah Castelo, Julian De Freitas, Berkeley Dietvorst, Christian Hildebrand, Young Eun Huh, Robert Meyer, Miriam E. Sweeney, Sanaz Talaifar, Geoff Tomaino, and Klaus Wertenbroch (2024), "How Artificial Intelligence Constrains the Human Experience," *Journal of the Association for Consumer Research*, 9 (3), https://doi.org/10.1086/730709.

Victorino, Liana, Ekaterina Karniouchina, and Rohit Verma (2009), "Exploring the Use of the Abbreviated Technology Readiness Index for Hotel Customer Segmentation," *Cornell Hospitality Quarterly*, 50 (3), 342–59.

Wang, Hongchang, Chunxiao Li, Bin Gu, and Wei Min (2019), "Does AI-Based Credit Scoring Improve Financial Inclusion? Evidence from Online Payday Lending," *40th International Conference on Information Systems, ICIS 2019 Proceedings*, https://aisel.aisnet.org/icis2019/blockchain_fintech/blockchain_-fintech/20/.

Yalcin, Gizem, Sarah Lim, Stefano Puntoni, and Stijn M.J. van Osselaer (2022), "Thumbs Up or Down: Consumer Reactions to Decisions by Algorithms Versus Humans," *Journal of Marketing Research*, 59 (4), 696–717.

Zarouali, Brahim, Sophie C. Boerman, and Claes H. de Vreese (2021), "Is This Recommended by an Algorithm? The Development and Validation of the Algorithmic Media Content Awareness Scale (AMCA-Scale)," *Telematics & Informatics*, 62, 101607.

Web Appendix for:
Lower Artificial Intelligence Literacy Predicts Greater AI Receptivity

Stephanie Tully*
Chiara Longoni*
Gil Appel*

Stephanie M. Tully (tullyste@marshall.usc.edu) is Associate Professor of Business Administration at the Marshall School of Business, University of Southern California. ORCID: 0000-0002-7179-7729

Chiara Longoni (chiara.longoni2@unibocconi.it) is Associate Professor of Marketing at Bocconi University. ORCID: 0000-0002-4945-4957

Gil Appel (gappel@gwu.edu) is Assistant Professor of Marketing at the GW School of Business, George Washington University. ORCID: 0000-0001-7859-3908

Web Appendix A: Development of the 25-item AI Literacy Measure

To develop our first instrument to measure AI literacy, we began with reviewing academic scholarship, public resources, and practitioner-focused pedagogy covering the topics of AI, algorithms, and digital literacy. These data sources included the AI Education Project (aiEDU 2023), UNESCO's Algorithm & Data Literacy project (Digital Moment 2023), and the National University of Singapore's library on digital literacy (NUS 2023).

This led to the development of a set of 47 possible questions conceptual and technical knowledge, best practices in AI development and usage, and regulatory requirements. At this point, we sought feedback from several data scientists and programmers on the initial set of questions. We asked them to evaluate the questions based on their relevance, accuracy, potential to become obsolete or change over time, comprehensibility, and any other feedback. This reduced our pool to 30 questions, as 17 questions were deemed irrelevant, too niche to measure actual AI literacy, too easy or too difficult. Simultaneously, we asked an AI program, ChatGPT 3, to similarly evaluate the questions and provide explanations for the answers to identify questions that could plausibly be confusing, have multiple answers, or may be subject to change over time. This step resulted in an additional 7 questions that were deemed to be problematic (e.g., AI suggested that a different response was possibly correct), two overlapping with questions that were excluded based on feedback from the data scientists and programmers. Thus, in total, the initial set of 47 questions was refined to 25 questions. The 17 multiple-choice questions that comprise the AILT are available in Table 1 below.

Table 1: The 25-Item AI Literacy measure

Question
Note. Order of questions and order of answers within each question were randomized.

1. Which of the languages below are used to code a robot?
a) Python [correct]
b) HTML
c) Computer Vision
d) Hypertext

2. Which technology is the primary enabler of Artificial Intelligence?
a) Machine Learning [correct]
b) Electric Battery
c) Robotics
d) Engineering

3. Which of the following tasks was the first to be performed well by Artificial Intelligence?
a) The complete replacement of a doctor in the treatment of patients
b) The implementation of psychotherapies
c) The analysis of X-ray images, for example to detect a torn meniscus or a tumor [correct]
d) The performance of complex physical work that requires dexterity or precise hand-eye coordination

4. Which of these is a notable algorithmic problem?
a) Counting
b) Sorting [correct]
c) Multiplication
d) Division

5. What was the purpose of the PCI DSS (payment card industry data security standard)?
a) Regulate financial transactions [correct]
b) Manage browser cookies
c) Standardize credit score data
d) Enable websites to accept crypto currencies

6. Which of these is not a best practice when trying to protect your privacy online?
a) Using a two-factor authentication to access account and devices
b) Storing your passwords in the notes section of your smartphone [correct]
c) Updating your operating system promptly when an update becomes available
d) Using different login information

7. If using an algorithm, can the results be biased against specific groups of people?
a) No, unlike humans, algorithms are free of emotions and cannot be wrong
b) Yes, but only if you train them on privileged populations
c) Yes, but only if you use them with incomplete data
d) Yes, even with correct data and training, algorithms can provide biased results [correct]

8. How algorithmic decisions best explained?
a) Only for very basic and simple algorithms
b) They can be explained by sharing the algorithm with everyone
c) As complex code, they cannot be explained
d) Algorithms can be explained by sharing the contributing factors to decisions [correct]

9. When programming stock investment recommendation algorithms, these algorithms are typically programmed to:
a) Ignore smaller companies
b) Give more weight to random fluctuations in the market
c) Consistently keep track of investments [correct]
d) Prefer indices over specific stocks

10. According to the Health Insurance Portability and Accountability Act of 1996 (HIPAA), protected health information cannot be shared with:
a) The patient's employer [correct]
b) The patient
c) Law enforcement
d) The patient's health insurance

11. AI can take images and short sound clips of people and create long, convincing impersonation of these people saying anything. What is this technology called?
a) Deepmind
b) Deepfake [correct]
c) Deeplearn
d) Deepvoice

12. Google automatically completes the text you type in the search bar...

a) By looking at the most common search terms used by others and your search history [correct]
b) Based only on your search history
c) Based only on what people in your zip code are searching for
d) By using your phone camera when you type and reading your lips

13. Algorithms may improve website performance by doing all of the following except for….
a) Showing us content that might interest us
b) Showing us relevant ads
c) Monitoring which content we engage in
d) Changing our screen brightness [correct]

14. What information do social media algorithms not typically use when deciding what information to display to a person?
a) Articles people have previously clicked on
b) Videos people have previously watched online
c) Information people have shared on social media
d) What would benefit society [correct]

15. When you use a web-based service which stores its data in "the cloud", where is the data actually stored?
a) In a hidden folder in your computer
b) At the company's headquarters
c) At a large server farm [correct]
d) In your browser's cookies

16. If you turn off personalized ads, you are most likely to see ads...
a) Targeting people similar to you
b) That are relevant for you
c) That are interesting to you
d) For commonly purchased products [correct]

17. Pattern recognition is a decision-making step employed by...
a) Algorithms
b) Humans
c) Both [correct]
d) Neither

18. When do algorithms tend to perform better than humans?
a) When lots of structured data is involved and the risk from a wrong answer is minimal [correct]
b) When there is no structured data available and the risk from a wrong answer is minimal
c) When lots of structured data is involved and the risk from a wrong answer is significant
d) When there is no structured data available and risk from a wrong answer is significant

19. What is training data and why is it important?
a) Training data is used when the model is fully calibrated, to train new users with how to work with the model
b) Training data is used to calibrate the model during creation, creating a model that performs best according to that data [correct]
c) Training data is used to test the model after it was created, to see how it works with new data
d) Training data is used to minimize bias in older models and train them to new population types

20. Indicate if the following is primarily the result of an algorithm or not:
What I write in a post on social media
a) Yes
b) No [correct]

21. Indicate if the following is primarily the result of an algorithm or not:
What I type in a text message to a friend
a) Yes
b) No [correct]

22. Indicate if the following is primarily the result of an algorithm or not:
When I read a direct message on a social media platform
a) Yes
b) No [correct]

23. Indicate if the following is primarily the result of an algorithm or not:
Returning the results of a search query on a search engine (e.g., Google)
a) Yes [correct]
b) No

24. Indicate if the following is primarily the result of an algorithm or not:
Dual authentication to access one of my accounts
a) Yes [correct]
b) No

25. What is the process of a machine learning model making a prediction based on new data called?
a) Coding
b) Training
c) Testing
d) Inference [correct]

Web Appendix B: Development of the 17-item AI Literacy Measure

To make sure our findings are not dependent on the 25-item measure we developed (described above), we leveraged AI to develop a second measurement of AI literacy based on previously identified competencies in AI literacy. Specifically, Long and Magerko (2020) identify a comprehensive list of 17 competencies that they consider to be the core foundations of AI literacy (e.g., recognizing AI, Interdisciplinarity of AI, the role of humans in AI).

These 17 competencies were described to both Claude.AI and ChatGPT-4. We then asked these AI tools to create a multiple-choice question that would assess each of the 17 competencies, one question per competency. From these two sets, we selected the question that seemed most representative for each of the AI competencies for a final set of 17 questions. The 17 multiple-choice questions that comprise the AILT are available in Table 2 below.

Question *Note.* Order of questions is by the order of competencies at Long and Magerko (2020), order of answers within each question were randomized.
Recognizing AI: **1. Which of the following is NOT powered by AI?** a) Self-driving cars b) Google's search algorithm c) A basic calculator (correct) d) Chatbots **Understanding Intelligence:** **2. Which form of intelligence involves emotional understanding and social skills?** a) Machine Intelligence b) Human Intelligence (correct) c) Animal Intelligence d) Artificial General Intelligence **Interdisciplinarity:** **3. Which of the following fields contributes to the development of artificial intelligence?** a) Computer science b) Mathematics c) Psychology d) All of the above (correct) **General vs Narrow AI:** **4. What is the term for AI systems that can perform any intellectual task that a human can?** a) Narrow AI b) General AI (correct) c) Weak AI d) Strong AI **AI's Strengths and Weaknesses:** **5. In which area does AI typically excel?** a) Emotional understanding b) Pattern recognition (correct) c) Moral reasoning

d) Creativity

Imagine Future AI:
6. Which of the following is NOT a likely future application of AI?
a) Personalized healthcare
b) Emotional robots
c) Time travel (correct)
d) Sustainable energy management

Representations:
7. What is a common form of knowledge representation in AI?
a) Neural networks (correct)
b) Waterfall model
c) Agile methodology
d) SWOT analysis

Decision-Making:
8. Which algorithmic approach is commonly used for decision-making in AI?
a) Dijkstra's algorithm
b) Depth-first search
c) Decision trees (correct)
d) Fourier Transform

Machine Learning Steps:
9. What is the first step in a typical machine learning process?
a) Data collection (correct)
b) Model selection
c) Prediction
d) Model evaluation

Human Role in AI:
10. Who is primarily responsible for the ethical considerations of an AI system?
a) The AI system itself
b) Data providers
c) Human developers (correct)
d) End-users

Data Literacy:
11. Which of the following is an example of metadata?
a) A spreadsheet of numbers
b) Column headers in a table (correct)
c) A chart visualization
d) Raw sensor data

Learning from Data:
12. How do supervised machine learning algorithms learn?
a) From labeled data (correct)
b) From rewards and punishments
c) By observing human behavior
d) From intrinsic motivation

Critically Interpreting Data:
13. Why can data not be taken at face value?
a) It is always inaccurate
b) It requires interpretation (correct)
c) It is self-explanatory
d) It is always biased

Action & Reaction:
14. How can an AI system interact with the physical world?
a) By planning movements
b) By reacting to sensor inputs
c) By actuating motors
d) All of the above (correct)

Sensors:
15. Which of the following sensors allow an AI system to perceive the world?
a) Cameras
b) Microphones
c) Thermometers

d) All of the above (correct)

Ethics:
16. Which is a key ethical issue surrounding AI?
a) Algorithmic efficiency
b) CPU usage
c) Privacy (correct)
d) Code readability

Programmability:
17. Which statement best describes the programmability of AI systems?
a) They cannot be programmed by humans
b) They program themselves
c) They are programmed using data
d) They are programmed by computer code (correct)

Web Appendix C: Factor Loading Analysis of the AI Literacy Measures

To determine whether our findings were driven by a subset of the AI literacy questions or reflected the AI literacy measures as a whole, we combined responses across studies. Specifically, we combined the three studies utilizing the 25-item AI literacy measure: Study 2, Study 4, and Study 6 and three studies using the 17-item AI literacy measure: Study 3, Study 5, Study 7, and supplemental study. Each set (the 25-item studies and the 17-item studies) was analyzed separately. Factor analysis performed on both the 25-item and 17-item measures produced scree plots (Figure WA3a and Figure WA3b for the 25-item and 17-item measures, respectively) that show the most significant drop after the first factor. This indicates that subsequent factors only marginally increase the explained variance in the data, and that the AI literacy measure is best described as a single factor.

Figure WA3a: Scree plot analysis of the 25-item measure.

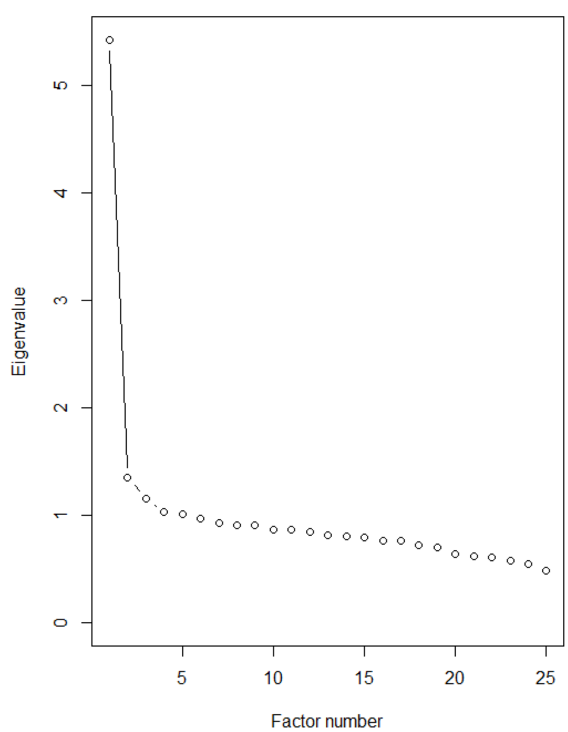

Figure WA3b: Scree plot analysis of the 17-item measure.

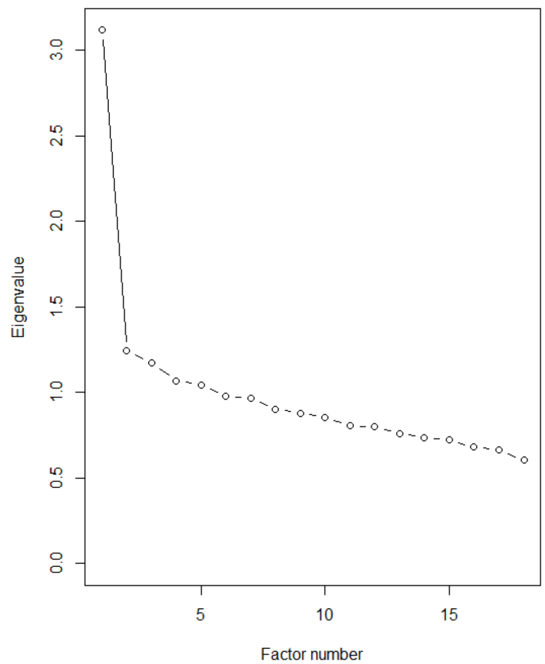

However, we conducted robustness checks that assess whether questions tapping into (1) limitations and ethicality considerations and (2) technical understanding both separately predict the lower literacy-higher receptivity link. These are conceptual distinctions agreed upon by the authors.

For the 25-items measure, the factor allocation was:

1. Limitations and ethical considerations: Questions 5, 6, 7, 10, 11, 14, 24.
2. Technical understanding: Questions 1, 2, 3, 4, 8, 9, 12, 13, 15, 16, 17, 18, 19, 20, 21, 22, 23, 25.

For the 17-items measure, the factor allocation was:

1. Limitations and ethical consideration: 2, 4, 5, 6, 10, 13, 16.
2. Technical understanding: 1, 3, 7, 8, 9, 11, 12, 14, 15, 17.

We then regressed the AI receptivity measures on the AI literacy measure and on each of the two factors to check if both factors still capture the association between AI receptivity and AI literacy, we ran the regression for the measure and each of the factors separately to see if they capture the association with AI receptivity and also included a dummy for the study number as control in each of the regressions. For these analyses, we excluded Study 7 since the shared (non-distinctly human) attribute condition exhibited a full reversal of the lower literacy-higher receptivity link.

Our results are robust to these analyses. For the 25-item measure, the measure was negatively associated with AI receptivity ($B_{25\text{-}item}$ = -.05, 95% *CI* (-.06, -.04), *SE* = .004, $t(2738)$ = -10.605, $p < .001$), and similarly for each of the factors, ($B_{25\text{-}item\_1}$ = -.12, 95% *CI* (-.15, -.10), *SE* = .01, $t(2738)$ = -9.551, $p < .001$; $B_{25\text{-}item\_2}$ = -.05, 95% *CI* (-.07, -.04), *SE* = .01, $t(2738)$ = -9.43, $p < .001$). Similarly, for the 17-item measure, the measure was negatively associated with AI receptivity ($B_{17\text{-}item}$ = -.04, 95% *CI* (-.05, -.02), *SE* = .01, $t(2422)$ = -4.688, $p < .001$), and similarly for each of the factors, ($B_{17\text{-}item\_1}$ = -.07, 95% *CI* (-.10, -.04), *SE* = .01, $t(2422)$ = -5.047, $p < .001$; $B_{17\text{-}item\_2}$ = -.04, 95% *CI* (-.06, -.01), *SE* = .01, $t(2422)$ = -2.961, $p = .003$).

Web Appendix D: Supplemental Study

This supplemental study examines how encouraging participants to base their preferences to use AI on the more emotional vs. practical aspects of AI influences the relationship between AI literacy and AI receptivity.

Method

Participants. We requested 1,000 participants from the platform Lucid, which uses quota sampling to provide respondents representative of the U.S. population on age, gender, ethnicity, and geographic region. This experiment included, as pre-registered, one IMC ("This question is just to make sure you are taking the time to read questions carefully. This is not a trick question. What is 12-4? Please use a numeric value only (no letters)"). This was asked at the beginning of the survey and participants who failed this question did not complete the survey (i.e., we collected no data for them, and incorrect responses did not affect our target sample size). In total, Lucid recruited 1,134 participants, with 102 participants failing to accurately answer the attention check and thus not completing the survey. As such, 1,032 participants successfully completed the study (M_{age} = 45.5, *SD* = 16.7; Females = 51.1%, Males = 47.8%, Non-binary/Third gender = .6%, Other: .1%, Prefer to not say:.4%[1]).

Procedure. AI literacy was measured using the 17-item instrument. In addition to measuring AI literacy, there were two between-subjects conditions manipulating people's focus on different potential aspects of AI (emotional vs. practical). Participants were asked to consider AI-generated music. However, we encouraged participants to focus on the more emotional versus practical using the following blurb with differences across condition in brackets, "Today,

[1] The demographic information reported here is based on self-report, which varies slightly from demographic information provided by Lucid. Note that two respondents provided numbers for age that appeared to be years of birth (1977 and 2002) and were therefore changed. All demographics (self-reported and Lucid-reported) are available in the dataset in ResearchBox.

we want you to consider the [enchanting and magical / practical and efficient] world of AI-generated music. In today's world, technology [inexplicably / efficiently] translates emotions, stories, and personal experiences into melodious songs. This is about appreciating the [magic and awe-inducing nature of AI as it weaves personal stories into harmonies and tunes / speed, efficiency, and high-quality that AI brings to the table]. AI in this context is [an artist in its own right / a capable and reliable composing tool], offering [wonder and astonishment / precision and consistency] in creating music tailored to people's tastes." Then, participants were asked, "Given how [amazing and magical / capable and efficient] you think it would be to have AI write a song, how likely would you be to use AI to write a song for a friend's birthday?" (1 = *very unlikely*, 7 = *very likely*). The order of the AI receptivity question and the AI literacy instrument were counterbalanced. Finally, participants provided demographic information (gender, age, and income).

Results and Discussion

AI literacy. We computed an AI literacy score for each participant as in previous studies. The median score was 11 (range: 1-16, $M = 10.21$, $SD = 2.92$).

AI receptivity. We first regressed AI receptivity on AI literacy score. Replicating the findings of the other studies, AI literacy negatively predicted AI receptivity, $B = -.06$, 95% CI [-.10, -.02], $SE = .03$, $t(1028) = -2.77$, $p = .006$. We then regressed AI receptivity on AI literacy (mean-centered), condition (0 = practical, 1 = emotional), and their interaction. None of the effects were significant (AI literacy: $B = -.04$, 95% CI [-.10, .02], $SE = .03$, $t(1026) = -1.30$, $p = .195$; condition: $B = .02$, 95% CI [-.22, .27], $SE = .13$, $t(1026) = 0.19$, $p = .846$; interaction: $B = -.04$, 95% CI [-.12, .05], $SE = .04$, $t(1026) = -0.88$, $p = .382$). Despite the lack of a significant interaction, correlations by condition were directionally consistent with our hypothesis. In the emotional condition, we replicated the relationship seen in other studies, such that people with lower AI literacy had greater AI receptivity, $r = -.12$, $p = .009$. However, when asked to consider the more practical aspects of AI-generated music, this relationship was not significant, $r = -.06$, $p = .198$.

Web Appendix E: Study 6 Post-tests: Perceived AI Ethicality and AI Capability

Study 6 finds that people with lower AI literacy prefer having AI complete a broad range of tasks. We argue that this is because people with lower AI literacy are more likely to perceive AI as being magical, and hence experience greater feelings of awe towards AI task completion. However, it is possible that people with lower AI literacy perceive the use of AI to complete tasks as more ethical, or that they perceive AI as more capable than those with higher AI literacy. To examine these possibilities, we run two post-tests. One post-test explores the relationship between AI literacy and the perceived ethicality of having AI complete each of the 26 tasks used in study 6. The second post-test explores the relationship between AI literacy and the perceived capability of AI completing each of the 26 tasks used in study 6.

Perceived AI Ethicality Post-test

Procedure. We requested 200 participants through Cloud Research. This experiment included, as pre-registered, one IMC ("This question is just to make sure you are taking the time to read questions carefully. This is not a trick question. What is 12-4? Please use a numeric value only (no letters)"). This was asked at the beginning of the survey and participants who failed this question did not complete the survey (i.e., we collected no data for them, and incorrect responses did not affect our target sample size). In total, 214 participants opened the survey. Four people failed the attention check and another seven participants did not finish the survey leaving a final sample of 202 (M_{age} = 40.4, SD = 12.4; Females = 48.0%, Males = 50.0%, Non-binary/Third gender = 1.0%, Prefer to not say: 1%).

First, we measured AI literacy using the 25-item measure used in the main study. We then assessed perceptions of ethicality by asking participants "To what extent do you think it is ethical to have an AI algorithm complete each of the following tasks?" followed by the same 26 tasks used in Study 6 (7-point scales: 1 = *Definitely unethical*, 7 = *Definitely ethical*). Finally, we collected gender and age.

Results and Discussion.

AI literacy. We computed an AI literacy score for each participant as in previous studies. The median score was 19.5 (range: 4-24, $M = 18.12$, $SD = 4.56$).

Perceived AI ethicality. We calculated participants' average perceptions of the ethicality of having AI complete the tasks. In contrast to what would be predicted by the alternative explanation, we found a positive relationship between AI literacy and perceived AI ethicality such that people with greater AI literacy believed that it was more ethical to have AI complete the tasks, $B = .036$, 95% CI (.003, .068), $SE = .016$, $t(200) = 2.172$, $p = .031$. Figure WAE1 shows the relationship for each of the 26 tasks.

Perceived AI Capability Post-test

Procedure. We requested 200 participants through Cloud Research. This experiment included the same IMC as the ethicality post-test. In total, 216 participants opened the survey. Two people failed the attention check and another eleven participants did not finish the survey leaving a final sample of 203 ($M_{age} = 40.9$, $SD = 11.7$; Females = 53.7%, Males = 45.8%, Prefer to not say: .5%).

First, we measured AI literacy using the 25-item measure used in the main study. We then assessed perceptions of AI capability by asking participants "To what extent do you think that an AI algorithm is capable of completing each of the following tasks?" followed by the same 26 tasks used in study 6 (7-point scales: 1 = *Definitely incapable*, 7 = *Definitely capable*). Finally, we collected gender and age.

Results and Discussion. AI literacy. We computed an AI literacy score for each participant as in previous studies. The median score was 19 (range: 5-24, $M = 17.41$, $SD = 4.73$).

Perceived AI capability. We calculated participants' average AI capability perceptions across the 26 tasks. In contrast to what would be predicted by the alternative explanation, people with lower AI literacy did not perceive AI as being more capable, $B = .01$, 95% CI (-.02, .04), $SE = .01$, $t(201) = 0.63$, $p = .531$. Figure WAE2 shows the relationship for each of the 26 tasks.

Overall, these post-tests challenge alternative explanations of the lower literacy-higher receptivity link based on perceived AI ethicality and AI capability. Despite exhibiting greater AI receptivity, people with lower AI literacy perceive AI as less ethical to use to complete tasks, refuting an account based on perceived ethicality. Furthermore, they do not perceive AI as being more capable, which is inconsistent with an account based on perceived capability.

Figure WAE1: Ethicality post-test figure

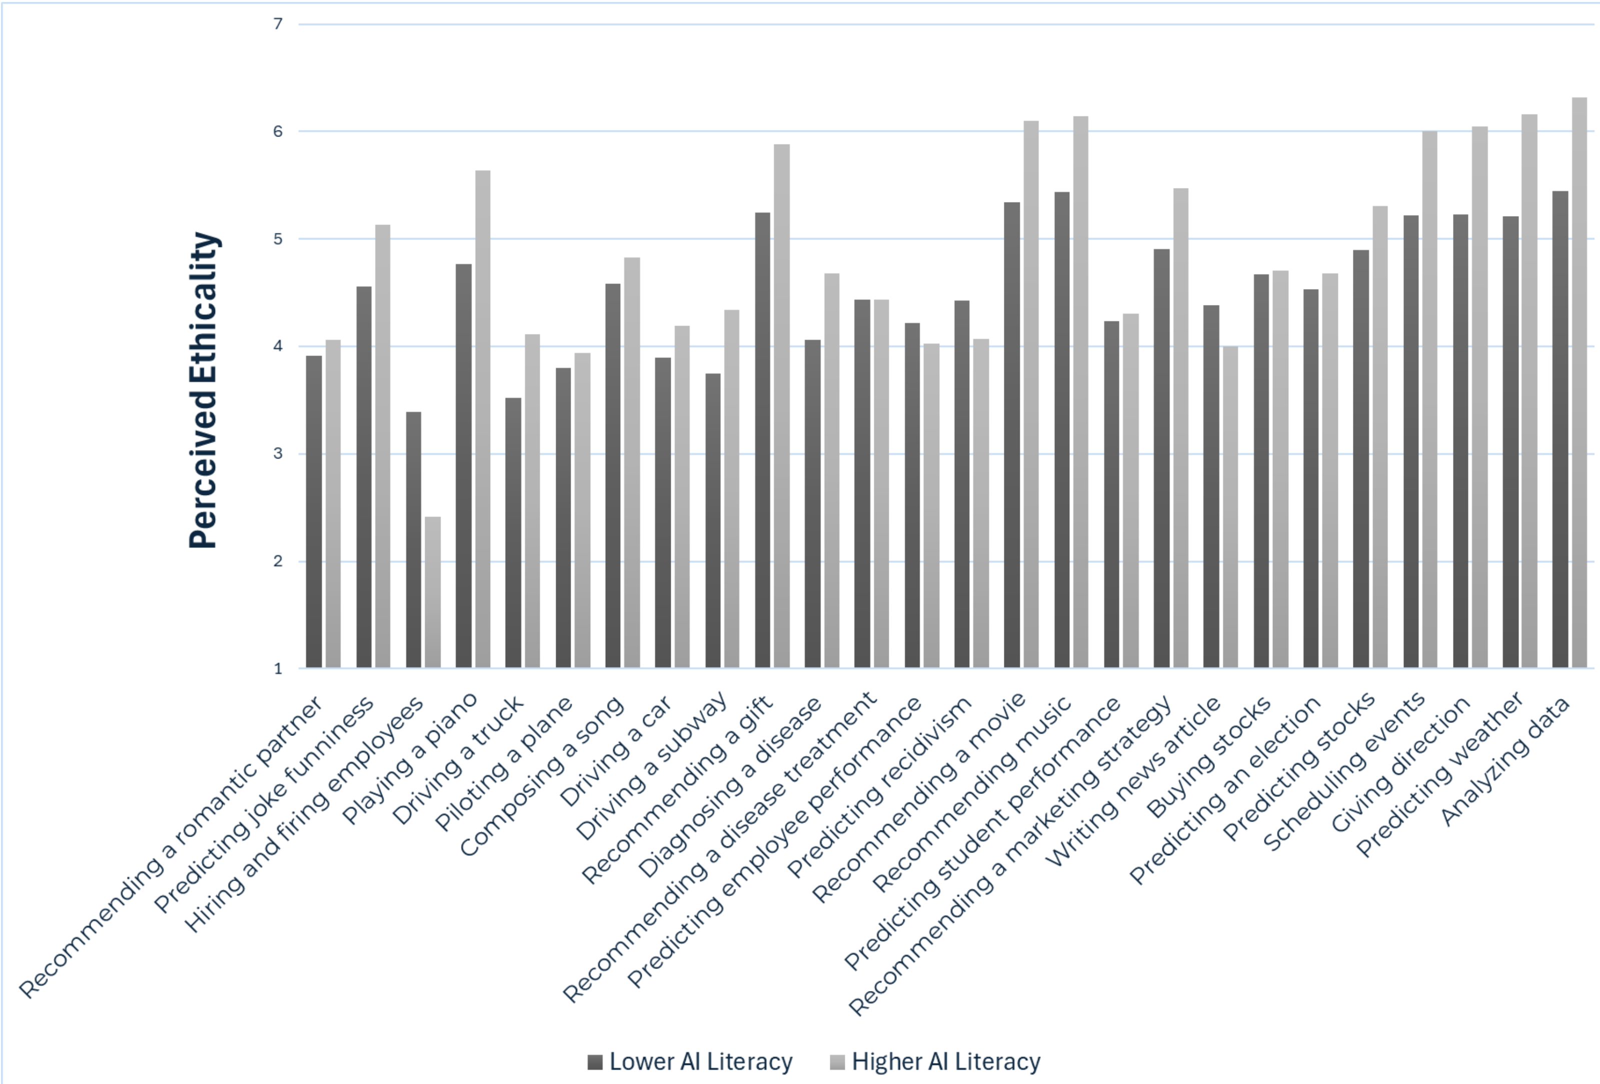

Figure WAE2: Capability post-test figure

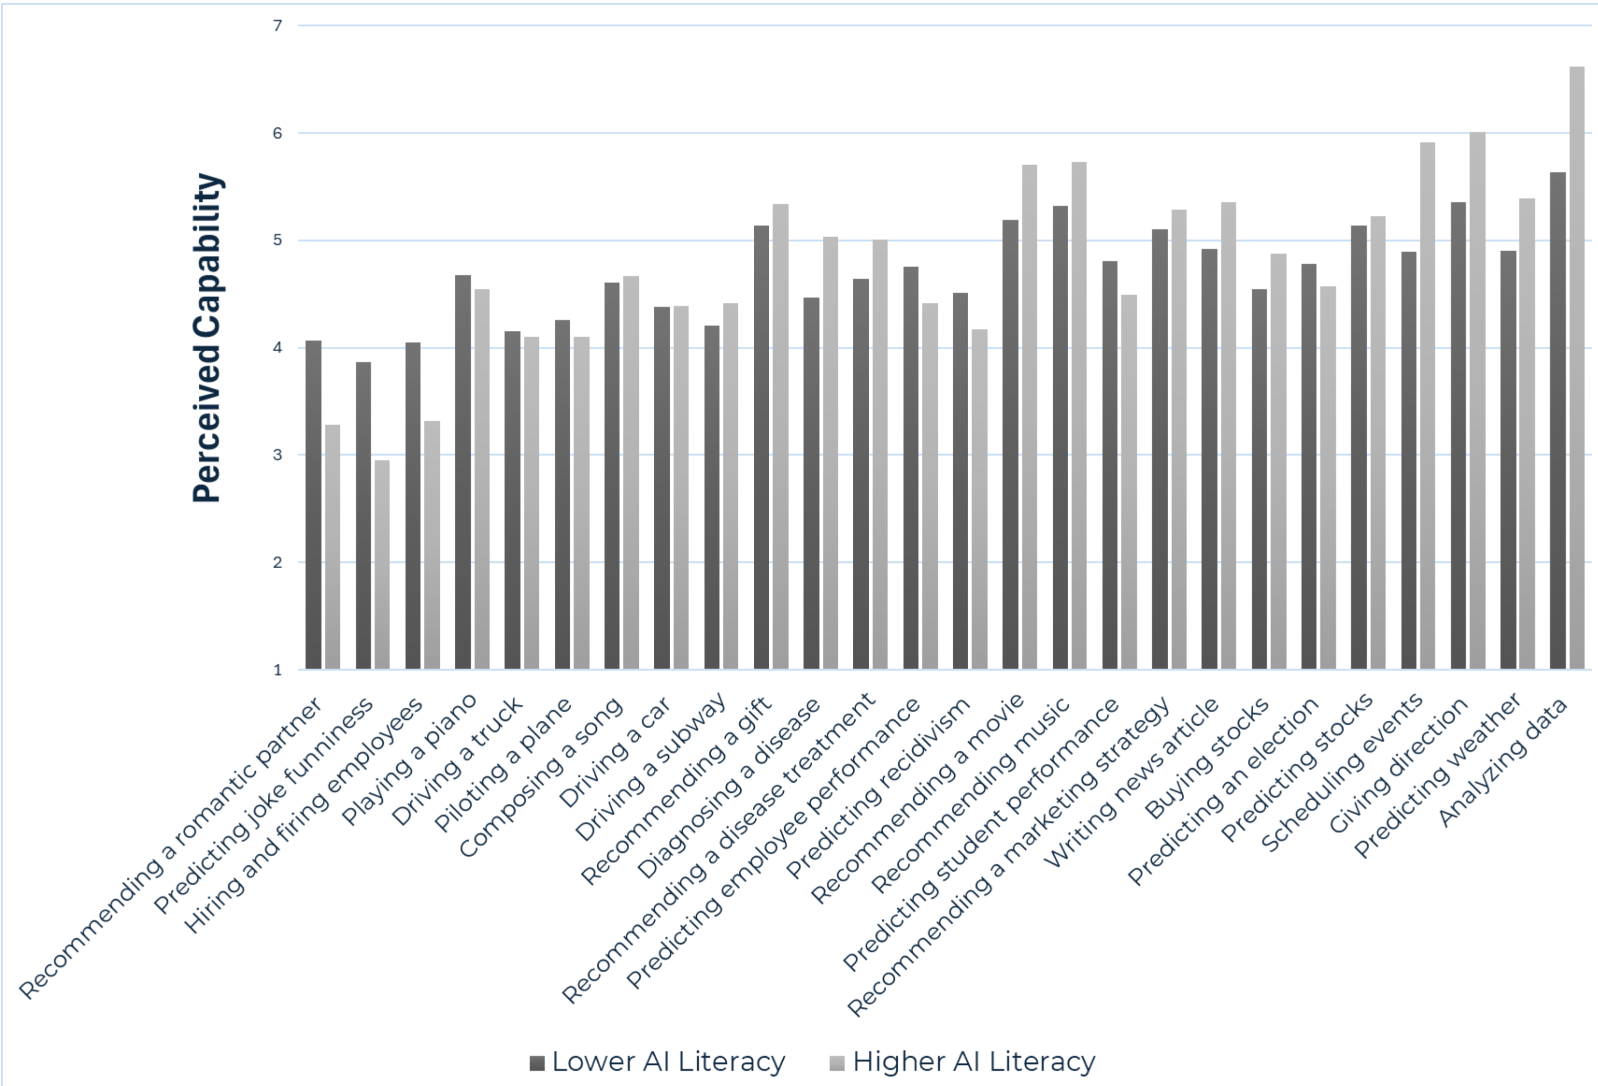

Web Appendix F: Study 7 Pre-test

In Study 7, we sought to manipulate the type of attributes required to execute tasks—specifically, distinctly human attributes versus shared attributes (attributes common to both humans and AI). To ensure that our task pairs effectively differed across these attributes, we conducted a pretest to validate the perceived differences in attribute requirements.

Participants

A total of 102 participants from Cloud Connect completed the pretest survey (M_{age} = 35.68, *SD* = 11.88; Females = 56.9%, Males = 42.2%, Preferred not to say=1%;).

Method

We generated 10 pairs of tasks across various domains such as cooking, art, architecture, and healthcare. Each pair consisted of one task requiring distinctly human attributes, and one task requiring shared attributes. These tasks were developed with the assistance of an AI language model (Claude model 3.5 Sonnet) and based on attribute classifications identified by Santoro and Monin (2023). For each of the 10 tasks, participants were asked to look at two columns of attributes. Column A included all of the distinctly human attributes identified by Santoro and Monin (2023) and Column B included all of the shared tasks identified by that research. It then asked participants, to what extent they believe the following sets of attributes would be important for the task above (1 = *the attributes in column A are much more important*, 7 = *the attributes in column B are much more important*). Thus, higher scores indicated a belief that less distinctly human attributes were needed to execute the task. The task pairs were as follows:

	Distinctly Human Attribute Task	Shared Attribute Task
Cooking	Creating a recipe for a fusion dish that represents a specific cultural heritage	Scaling a recipe originally intended to serve 4–6 people to serve 75–100 people
Art	Creating an abstract painting expressing inner emotions	Reproducing a photorealistic portrait from a reference image
Architecture	Designing a monument to celebrate an important cultural and historical moment	Creating precise technical blueprints for a building
News reporting	Writing an opinionated editorial piece for a news outlet	Fact-checking and verifying sources for a news outlet
Fashion	Designing an avant-garde fashion collection	Calculating fabric requirements for mass production of a clothing line
Advertising	Creating an authentic social media campaign	Conducting A/B testing on email marketing campaigns
Experience Innovation	Designing an experiential marketing event to build brand community	Implementing a dynamic inventory management system based on predictive analytics
Customer engagement & retention	Creating an immersive brand experience for a pop-up store	Developing a predictive model for customer churn
Healthcare	Providing end-of-life care and counseling	Analyzing medical test results and providing a report
Hospitality Management	Designing a unique and memorable wedding event for a hotel chain	Managing room bookings for optimal capacity of a hotel chain

The results validated that there were differences in perceptions of the importance of distinctly human versus shared attributes for task execution across the pairs of tasks. The tasks designed to require more distinctly human attributes were perceived as having a greater need for those attributes as indicated by lower scores on the dependent measure ($M_{human} = 3.19$, $M_{shared} = 5.96$; $t(101) = -24.63$, $p < .001$; $d = -2.439$).

We also conducted paired t-tests for each of the 10 task pairs to ensure that the distinction was consistent across all categories. Below are the results for each pair:

1. Cooking ($M_{shared} = 6.10$, $M_{human} = 2.89$; $t(101) = 17.18$, $p < .001$).
2. Art ($M_{shared} = 4.85$, $M_{human} = 2.10$; $t(101) = 13.93$, $p < .001$).
3. Architecture ($M_{shared} = 6.49$, $M_{human} = 2.94$; $t(101) = 21.11$, $p < .001$).
4. News reporting ($M_{shared} = 6.03$, $M_{human} = 4.12$; $t(101) = 9.95$, $p < .001$).
5. Fashion ($M_{shared} = 6.17$, $M_{human} = 2.75$; $t(101) = 16.39$, $p < .001$).
6. Advertising ($M_{shared} = 5.66$, $M_{human} = 3.30$; $t(101) = 12.64$, $p < .001$).

7. Customer experience innovation (M_{shared} =6.40 , M_{human} =4.30 ; $t(101) = 11.63$, $p < .001$).
8. Customer engagement & retention (M_{shared} =5.63 , M_{human} =3.75 ; $t(101) = 8.95$, $p < .001$).
9. Healthcare (M_{shared} =6.36 , M_{human} =2.41 ; $t(101) = 21.41$, $p < .001$).
10. Hospitality management (M_{shared} =5.90, M_{human} =3.29 ; $t(101) = 13.30$, $p < .001$).

All 10 pairs of tasks showed significant differences in the expected direction, with participants rating distinctly human attributes as more important for the tasks designed to require distinctly human attributes (vs. the tasks designed to require shared attributes).

Web Appendix References

aiEDU (2023), "AI in 5 Minutes," (accessed June 6, 2023), https://www.aiedu.org/ai-in-five.

Digital Moment (2023), "What is Artificial Intelligence, Algorith & Data Literacy," (accessed June 6, 2023), https://algorithmliteracy.org/.

Long, Duri, and Brian Magerko, (2020), "What is AI literacy? Competencies and design considerations," *Proceedings of the 2020 CHI conference on human factors in computing systems*, 1-16.

NUS (2023), "Digital Literacy: Home," (accessed June 6, 2023), https://libguides.nus.edu.sg/digitalliteracy.

Santoro, Erik, and Benoît Monin, (2023), "The AI Effect: People rate distinctively human attributes as more essential to being human after learning about artificial intelligence advances," *Journal of Experimental Social Psychology*, 107, 104464.