**Of Essence and Obsolescence:
Priming Artificial Intelligence Changes Perceptions of Human Nature**

Erik Santoro and Benoît Monin

Stanford University

Author Note

Erik Santoro, <https://orcid.org/0000-0003-4092-4809>. Benoît Monin,
<https://orcid.org/0000-0003-2894-1876>.

This research was supported in part by a grant from the Stanford Institute for Human-Centered Artificial Intelligence. Some of this research was presented at the 2019 Society for Personality and Social Psychology Convention (Santoro & Monin, 2019).

Correspondence concerning this article should be addressed to Erik Santoro, Department of Psychology, Stanford University, Stanford 94305. Email: esantoro@stanford.edu

Draft version 8/27/2021. This version of the paper has not been peer reviewed. Please do not copy or cite without author's permission.

ARTIFICIAL INTELLIGENCE AND HUMAN NATURE

Abstract

Will people redefine what they see as essential to the human experience as the human species loses ground to Artificial Intelligence (AI) on such attributes as logic and language? Four studies using U.S. samples (N = 2,101) test this possibility. Study 1 supports a two-part classification of human attributes that we then use in subsequent studies. Studies 2a-2c demonstrate that, compared to reading a control prime about trees, reading a treatment prime about AI progress leads participants to view those attributes not shared with AI (e.g., humor and having desires) as more essential to being human. We found no evidence that this exposure leads to devaluation of those attributes shared with AI (e.g., logic and language). This research suggests that as AI becomes ubiquitous and increasingly capable, conceptions of human nature may shift in reaction to value what makes humans unique.

Statement of Relevance

Artificial Intelligence is becoming increasingly human-like, e.g., generating human prose, detecting cancer, or playing complex strategy games. Although some have expressed concern about what such gains in Artificial Intelligence technology mean for human employment or even for human existence, there has been less attention paid to how AI might impact perceptions of human essence. To address this gap, participants were asked what it means to be human after reading about AI advances in logic and language vs. reading a control prime. We found consistent evidence that participants prioritized attributes distinctive to being human, such as having desires or a sense of humor, after reading about AI. This suggests that increasingly sophisticated AI, and the ubiquitous exposure to such advances, might alter something as fundamental as what humans think is essential to being human.

Keywords: Artificial Intelligence, Intergroup Dynamics

Of Essence and Obsolescence:

Priming Artificial Intelligence Changes Perceptions of Human Nature

The last decade has seen prodigious progress in the development of Artificial Intelligence (AI), or computer output that resembles human thought. Increased computing power and novel programming strategies have resulted in robots that beat human champions in complex games like chess and Go (Strogatz, 2018), AIs that detect cancer more accurately than physicians (Savage, 2020), and algorithms that produce prose that can pass as human writing (Heaven, 2020; Katz, 2017). Not only can AI rival humans across many domains, but people are increasingly exposed to reports of these accomplishments. For example, articles in two major American newspapers, *The Wall Street Journal* and *The New York Times*, were 33 times more likely to mention “Artificial Intelligence” in 2019 (1 in 90 articles) than in 2009 (1 in 3,000 articles)¹, often portraying AI as having abilities on par with humans.

What is the impact of learning of these advances on humans’ self-image? In this paper, we investigate whether learning about AI achievements leads people to shift what they perceive as essential to being human. We first establish a classification of human attributes, and then use this classification to test how perceptions of being human might change after being exposed to articles touting recent AI advances.

Perceptions of Human Nature, and How Such Perceptions Might Change

To understand how perceptions of human nature change when people are exposed to AI advances, we first developed a classification of human attributes. Building on existing typologies of what it means to be human (e.g., Haslam, 2006; Gray, Gray, & Wegner, 2007; Weisman,

¹ Using ProQuest®, we searched the term “Artificial Intelligence” in the body or title in the print versions of the *Wall Street Journal* and the *New York Times* for the years between 2009 and 2019. We then divided the number of articles about Artificial Intelligence by the total number of articles covered by ProQuest from each newspaper for each year to achieve a ratio. The first author conducted this search in May 2020.

ARTIFICIAL INTELLIGENCE AND HUMAN NATURE

Dweck, Markman, 2017), we use AI as a referent to identify two types of human attributes. We propose that there are human attributes that will be perceived to distinguish humans from AI, such as morality or emotion (“distinctive attributes”), and ones that will be perceived as more closely shared with AI, such as logic and language (“shared attributes”).

We predicted that when exposed to descriptions of AI advances in shared attributes, humans would try to distinguish themselves from AI on non-threatened attributes. News coverage of AI advances in logic, language and other shared attributes imply that AI is increasingly capable of such human-like traits. This raises the specter of potential looming human obsolescence in these areas. In line with work on intergroup relations in which individuals seek to maintain positive social identity (Tajfel & Turner, 1979; Cambon & Yzerbyt, 2018), we hypothesized that, to maintain a positive group image in the face of AI showing threatening progress on shared attributes, people would compensate by prioritizing distinctive vs. shared attributes as more essential to being human. In the studies presented here, we find indeed that compensation is driven by a revaluation of distinctive attributes, or viewing such attributes as humor or desires as more essential to being human. Though we also predicted that there would also be a devaluation of shared attributes, or viewing such attributes as logic or computation as less essential to being human, we did not find evidence of this in the studies presented here.

The Present Studies

The goal of this research was first to corroborate our proposed classification of “shared” vs. “distinctive” human attributes, and then to investigate how priming participants with AI advances would influence perceptions of human nature. In Study 1, participants rated the extent to which they viewed humans and AI as capable of either distinctive or shared attributes. We then tested in Studies 2a – 2c whether, after learning about AI advances in shared attributes,

ARTIFICIAL INTELLIGENCE AND HUMAN NATURE

people prioritize distinctive vs. shared attributes as essential to being human. After participants read an article about AI advances in shared attributes or a control article discussing the attributes of trees, they rated how essential various attributes are to being human. All studies in this paper were pre-registered, and all received approval by the Stanford Institutional Review Board. We depart from our registrations at times to simplify our account and accurately reflect the consistent and more straightforward pattern that emerged from the data. Our departures fall in two categories: (a) we excluded conditions (in Studies 1 and 2c) and measures (in Studies 2a – 2c) that do not directly test the influence of AI on the perception of human attributes; and (b) we dropped analyses of attribute sub-types (in Studies 2a – 2c) to simplify the presentation of our data. We present in the online supplement all the pre-registered conditions, measures, and analyses excluded from the main text.

Study 1: Testing a Classification of Human Attributes

Study 1 tests our classification of human attributes. We propose that some human attributes are seen as shared with AI (“shared attributes”), whereas others are seen as unique to humans (“distinctive attributes”). To investigate whether this theoretical classification in fact matches with people’s perception of human nature, we asked participants to rate how capable they perceived humans and AI of twenty human attributes, ten of which we designed as “distinctive” human attributes (e.g., love and morality) and ten of which we meant to be “shared” human attributes (e.g., logic and language). We then tested whether participants perceive AI, relative to humans, as less capable of distinctive attributes than shared attributes.

ARTIFICIAL INTELLIGENCE AND HUMAN NATURE

Method

Participants

From Amazon Mechanical Turk, we recruited participants with 95% approval or higher living in the United States for a study paying \$0.75 in which they were told they would “evaluate the capacities of some type of entity”. After removing four participants through post-hoc exclusions (i.e., keeping only the first responses of participants; removing instances of participants not completing the confirmatory measure; removing a participant without a mechanical Turk ID), we were left with 200 participants (99 in the AI target condition, and 101 in the human target condition; 57% men, 42.5% women, 0.5% other or prefer not to say; 74% White, 9.5% Asian or Asian American, 6.5% Black or African American, 5.5% Hispanic or Latino, 0.5% Native American, 0.5% other, 3.5% multiple ethnicities or races selected; 53.5% graduated from a four-year college or higher, 46.5% graduated from a community college or less; mean age = 35.53 years, $SD = 10.40$).

Procedure

Participants read an abridged version of a question used in Weisman, Dweck & Markman (2017), evaluating the extent to which “Some [humans / types of Artificial Intelligence] are capable of ...” of ten distinctive attributes (i.e., having culture, holding beliefs, having a sense of humor, being moral, being spiritual, having desires, feeling happy, feeling love, having a personality, having relationships) and ten shared attributes (i.e., doing computations, using language, implementing rules, predicting the future, using logic, communicating, recognizing faces, remembering things, sensing temperatures, detecting sounds). Attributes were either created by the authors or drawn from Weisman et al. (2017) and Haslam (2006). The attributes were presented in a different randomized order for each participant, each with a 7-point

ARTIFICIAL INTELLIGENCE AND HUMAN NATURE

capability scale (0 = *not at all capable* to 6 = *highly capable*). After completing this measure, all participants, regardless of condition, also indicated what attribute they viewed as most essential to being human as an exploratory measure (see supplement).

Results

To test whether participants perceived humans as more capable of distinctive attributes, but similarly capable of shared attributes, we predicted the item rating as a function of the interaction of target (i.e., human vs. AI) and attribute type (i.e., distinctive attribute vs. shared attribute), with a random intercept for item and subject.

As predicted, the sharpest difference between humans and AI was on distinctive attributes (human target: $M = 5.58$, $SD = 0.66$; AI target: $M = 1.17$, $SD = 1.35$; $b = 4.42$, $SE = 0.12$, $t(241.71) = 37.71$, $p < .001$, $d = 4.17$). Though a significant difference did also arise on shared attributes (human target: $M = 5.25$, $SD = 0.72$; AI target: $M = 4.69$, $SD = 1.10$; $b = 0.56$, $SE = 0.12$, $t(241.71) = 4.80$, $p < .001$, $d = 0.60$), it was much smaller than the one observed on distinctive attribute (interaction: $b = 3.85$, $SE = 0.07$, $t(3,780.00) = 53.74$, $p < .001$).

Additionally, as seen in Figure 1, AI were largely rated as incapable of distinctive attributes (i.e., below the midpoint) but capable of shared attributes (i.e., above the midpoint), while humans were seen as capable of both. The item “predicting the future” seemed to confuse participants (See Figure 1), so we changed it to “forecasting the future” in later studies in the hope to make it sound more like something humans and AI could do.

Figure 1*Mean Capability Rating (± 1 SE) by Target and Attribute Category*

Note. Dashed horizontal line marks the midpoint of the scale. Error bars are ± 1 standard error of the mean.

Discussion

The results of Study 1 support our two-part classification of human attributes into distinctive attributes (e.g., morality and emotion), that AI are seen as less capable of, and shared attributes (e.g., logic and language), that AI are seen as more capable of. We use this classification in Studies 2a – 2c to test whether learning about AI advances in shared attributes changes perceptions of how essential to being human participants perceive distinctive and shared attributes.

Studies 2a – 2c: Effects on Perceptions of Human Essence after

Priming AI's Advances in Shared Attributes

Using the classification established in Study 1 and building upon two pilot studies,² Studies 2a-2c test whether perceptions of human essence change after learning about AI's increasing sophistication at logic and language. Specifically, we tested whether participants, after reading an article highlighting AI proficiency at shared attributes ("AI Prime"), would rate distinctive attributes (e.g., humor and having desires) as more essential to being human relative to shared attributes (e.g., logic and language), compared to participants who read a neutral article about trees ("Tree Prime"). We measured the perception of human essence by asking participants to rate how essential each attribute from Study 1 was to being human.

Method

We planned to recruit 200 participants per cell from Amazon's Mechanical Turk, the default cell size used in our research team in the absence of confident effect size estimates. This mean that for Study 2a, given that we counterbalanced the sequence of measures in Study 2a, we

² We conducted two pilot studies that did not register the main hypotheses tested in this section. We describe these in the supplement.

ARTIFICIAL INTELLIGENCE AND HUMAN NATURE

recruited nearly 400 participants per Prime. Additionally, in Study 2b, we pre-registered that we would recruit 400 per cell to increase our power (and we did not analyze our data until we collected all data).

After removing 275 participants (13%) through pre-registered exclusions (i.e., same as those used in Study 1, as well as eliminating those who completed a previous study), we were left with 1,901 participants across the three studies (Study 2a: 757, Study 2b: 783, Study 2c: 361; overall 48.97% men, 50.87% women, 0.16% prefer another term; 70.75% White, 8.15% Black or African American, 7.36% Asian or Asian American, 5.63% Hispanic or Latino, 1.68% Native American, 0.21% Pacific Islander or Native Hawaiian, 0.16% Middle Eastern, 0.58% other, 5.47% multiple ethnicities or races selected; 55.87% graduated from a four-year college or higher, 44.13% graduated from a community college or less; mean age = 35.50 years, $SD = 11.33$).

Procedure

Participants were paid between \$1.60 to \$2.00 and were told that the study was about “how people understand and react to examples of scientific writing for popular audiences.” After consenting, participants first read: “You are going to get 2 out of 4 possible scientific passages. After reading each, you’ll be asked a few questions to check for your understanding and get your reactions.” All participants first read a filler article about clouds, after which they read a second article, either the AI Prime or the Tree Prime (see Table 1). The AI Prime highlighted how AI agents are similar to humans at shared attributes, such as computation, language, logic and communication -- all shared attributes, not distinctive attributes. To establish an appropriate comparison to the AI Prime, we created a control article (“Tree Prime”) in which participants read about the attributes of another non-sentient entity (i.e., trees). Unlike the AI Prime which

ARTIFICIAL INTELLIGENCE AND HUMAN NATURE

describes human attributes that AI shares, the Tree Prime only describes non-human attributes (e.g., fire resistance). See Table 1 for the exact language of the Primes. After each article, participants completed several comprehension questions about the article to maintain the cover story. Participants then took a measure assessing their perception of human essence, as well as other measures reported in the supplement. To assess participants' perceptions of human essence, we asked participants, "In your judgment, what is **essential** about the human experience? In other words, what is at the **core** of the human experience? Please rate the capacities below in terms of **how essential each** is to being human." (emphasis in original). In a different randomized order for each participant, participants then rated how essential each of twenty attributes was to being human on a five-point scale ranging from *Not essential* (0) to *Absolutely essential* (4). The twenty attributes are based on the items used in Study 1, i.e., distinctive attributes (i.e., beliefs, culture, humor, morality, spirituality, desires, happiness, love, personality, relationships) and shared attributes (i.e., computation, language, logic, forecasting the future, rule implementation, communication, facial recognition, memory, sound detection, temperature detection).

Table 1*Prime Manipulations Used in Studies 2a – 2c.*

The Artificial Intelligence Revolution:How Computers Increasingly Beat Us on Human Tasks [AI Prime]

Last year, Swedbank in Sweden “hired” Nina, a dynamo of a company representative. Nina is adept at **communicating** and interacting with customers on the bank’s website, answering their questions and routing them to the right human representative. She engages in some 45,000 conversations a month and successfully answers eight out of 10 inquiries.

But aside from her stellar performance, there’s something else notable about Nina: She’s a bot. Powered by Artificial Intelligence (or “AI” for short), not only can Nina communicate like humans, she can also make **computations** and use data to make recommendations to customers just as humans can.

Nina is just the tip of the iceberg in terms of human-like abilities AI excel at. Consumer-facing bots like Nina transcribe their conversations into text, which a different AI-based system can analyze to computationally mine for patterns. “They can do it on a scale no individual can,” says Mike Gualtieri, an analyst with Forrester.

For thousands of years, doing computations to **find patterns** was something only humans could do. This is why people for so long thought that computers would never beat humans at complicated games like Go, a game that requires strategy. Advances in computational methods led to Google DeepMind beating Lee Sedol, a top player, in Go in 2016. To beat humans at Go, Google DeepMind used pattern analysis on datasets with very **clear rules and parameters** in a way no human could.

Much of the recent advances are powered by a branch of AI theory called deep learning, which attempts to mimic the activity in layers of neurons in the neocortex, the wrinkly 80 percent of the human brain where thinking occurs. The software **learns**, in a very real sense, to recognize patterns and rules in digital representations of sounds, images, and other data.

AI’s ability to learn from data has recently been used to have such systems **make predictions**. Indeed, companies like Salesforce are now using AI to help them predict future sales. The ability of AI to learn has also enabled AI, like humans, to accurately **recognize faces**; indeed, the iPhone X, powered by deep learning algorithms, can unlock when a registered face appears in front of it.

AI-powered technology can not only communicate, but such technology can also master **language**. In the fall of 2016, Google Translate rolled out a new-and-improved AI-powered **translation** engine that it claimed was, at times, “nearly indistinguishable” from human translation.

Not only can AI translate on par with humans, but it can **read and understand** text like humans. Artificial intelligence programs and human participants were given a web page to read and then were asked specific questions about its content. Surprisingly, the AI programs were better than humans at **comprehending and remembering** objective facts from the text on the web page.

AI systems are so capable at using language that they can even **generate writing**. For instance, Google just rolled out a new feature called Smart Compose which allows it to autocomplete emails as you type them.

The Largest Trees in the World:

How Giant Sequoias Exist in California [Tree Prime]

Giant sequoias, also known as giant redwoods, are the world's largest trees. Giant sequoias can grow up to 300 feet tall and 30 feet in diameter with bark as thick as 3 feet. In fact, giant sequoias are the **largest tree in the world** as well as the largest living organism on Earth by volume.

Giant sequoias can also be very old. The oldest known giant sequoia based on ring count is 3,500 years old. Giant sequoias are among the **oldest living organisms** on Earth in fact.

Because of its height, the giant sequoia has been studied for its **ability to pull water up** the trunk. Water from the roots can be pushed up only a few meters by pressure, but can reach extreme heights through capillary action and pressure from evaporating water at the leaves. Sequoias supplement water from the soil with water from fog, at heights to where the root water cannot be pulled.

The natural distribution of giant sequoias is **restricted to a limited area** in California. They occur in scattered groves, comprising a total area of only approximately 35,000 acres. One particular place to see giant sequoias is Sequoia National Park.

Within the park, visitors travel to **The Giant Forest**, a three-and-a-half-square-mile centerpiece of the park which has more than 8,000 sequoias. The General Sherman Tree is one of the park's primary attractions. It's 275 feet tall, 100 feet in circumference, and known to be the largest tree on Earth, by volume. In 1978, a branch broke off a sequoia called the General Sherman. It was 150 feet long and nearly seven feet thick. All by itself, that branch would have been one of the tallest trees east of the Mississippi. The National Park Service drives home its massiveness on a sign in front of its trunk this way: If the General Sherman were hollowed out and filled with water, it'd be enough water for you to take a bath every day for 27 years.

While the present day distribution of this species is limited, it **was once much more widely distributed** in prehistoric times. In fact, giant sequoias was [*sic*] a reasonably common species in North America, Europe and Asia until its range was greatly reduced by the last ice age. Older fossil specimens of giant sequoia have been found as far afield as New Zealand and Australia.

Giant sequoias are in many ways adapted to forest fires. Their bark is unusually **fire resistant**, and their cones containing seeds will normally open immediately after a fire. Fires help in another way: they clear competing vegetation so that sequoia seeds can germinate. Without fire, shade-loving species would crowd out young sequoia seedlings, which would not germinate.

Note. The AI Prime was based on Asorey (n.d.), Brand Lab for Accenture (2017), Hiner (2017), Hof (2013), Katz (2017), Lambert (2018), Martin (2018). The Tree Prime was based on Giant Forest (n.d.), Mooallem (2017), Sequoiadendron giganteum (n.d.), and Szalay (2017). The Cloud Prime seen by all participants (not shown) was based on Cloud (n.d.) and Hitt (2017). We note spelling or grammatical mistakes in the text.

Results

Analytical Strategy.

For each study, we use linear mixed effect models to predict attribute item rating as a function of the interaction of Prime (AI vs. Tree) and attribute type, and their interaction, with random intercepts for subject and for item. For Study 2a, we added an additional fixed effect of order (i.e., whether the human essence was asked directly after the AI Prime or not), as registered.

Effect of the AI Prime (vs. Tree Prime).

As predicted, participants who read the AI Prime rated distinctive attributes as more essential to being human than did participants who read the Tree Prime in all three studies [Study 2a, $b = 0.09$, $SE = 0.05$, $t(945.48) = 2.06$, $p = 0.040$, $d = 0.14$; Study 2b, $b = 0.13$, $SE = 0.04$, $t(977.09) = 2.85$, $p = .005$, $d = 0.18$; Study 2c, $b = 0.17$, $SE = 0.07$, $t(440.37) = 2.71$, $p = .007$, $d = 0.27$]. We did not observe significant differences on the shared attitudes (see Table 2), which we had expected to go in the opposite direction to distinctive attributes as described in the introduction [Study 2a, $b = 0.05$, $SE = 0.05$, $t(945.48) = 1.16$, $p = 0.248$, $d = 0.08$; Study 2b, $b = -0.02$, $SE = 0.04$, $t(977.09) = -0.41$, $p = .684$, $d = -0.03$; Study 2c, $b = 0.07$, $SE = 0.06$, $t(440.37) = 1.07$, $p = .287$, $d = 0.10$]. These contrasting simple effects led to a significant attribute type by prime interaction in Study 2b, $b = 0.14$, $SE = 0.03$, $t(14,857.00) = 4.99$, $p < .001$, and Study 2c, $b = 0.11$, $SE = 0.04$, $t(6,839.00) = 2.63$, $p = 0.008$, but not in Study 2a, $b = 0.04$, $SE = 0.03$, $t(14,361.00) = 1.37$, $p = .171$.

Table 2*Descriptive and Inferential Statistics for Studies 2a – 2c*

	AI Prime <i>M</i> (<i>SD</i>)	Tree Prime <i>M</i> (<i>SD</i>)	AI Prime vs. Tree Prime					
			<i>d</i>	<i>b</i>	<i>SE</i>	<i>t</i>	<i>df</i>	<i>p</i>
Study 2a								
Effect of AI Prime on:								
Distinctive Attributes	2.83 (0.70)	2.73 (0.67)	0.14	0.09	0.05	2.06	945.48	.040*
Shared Attributes	2.52 (0.72)	2.46 (0.66)	0.08	0.05	0.05	1.15	945.48	.248
Difference b/t Effects				0.04	0.03	1.37	14,361.00	.171
Study 2b								
Effect of AI Prime on:								
Distinctive Attributes	2.92 (0.68)	2.80 (0.69)	0.18	0.13	0.04	2.85	977.09	.005**
Shared Attributes	2.56 (0.67)	2.58 (0.64)	-0.03	-0.02	0.04	-0.41	977.09	.684
Difference b/t Effects				0.14	0.03	4.99	14,857.00	<.001***
Study 2c								
Effect of AI Prime on:								
Distinctive Attributes	2.96 (0.64)	2.78 (0.63)	0.27	0.17	0.07	2.71	440.37	.007**
Shared Attributes	2.66 (0.69)	2.60 (0.62)	0.10	0.07	.06	1.07	440.37	.287
Difference b/t Effects				0.11	0.04	2.63	6,839.00	.008**

Note. Effect of AI Prime is always relative to Tree Prime. The difference between effects (“difference b/t effects”) is the interaction of prime (AI vs. Tree) and attribute type (distinctive vs. shared). Study 2a mean and SD collapses across order (i.e., whether essence measure was presented first or second).

Figure 2*Mean Essence Rating (+/- 1 SE) by Prime and Attribute Category for Studies 2a - 2c*

Note. Average rating by Prime and attribute category are plotted. Simple effect p values presented. Y-axis truncated to better see comparisons. Error bars are +/- 1 standard error of the mean. *** $p < .001$; ** $p < .01$; * $p < .05$; NS $p > .05$

ARTIFICIAL INTELLIGENCE AND HUMAN NATURE

Discussion

Across Studies 2a – 2c, we found support for the prediction that reading the AI Prime (vs. the Tree Prime) led human participants to rate distinctive attributes, or those human attributes AI are not viewed as capable of, as more essential to being human. Interestingly, and contrary to prediction, we did not find evidence that the AI Prime changed participants' perception of shared attributes. This is in spite of the fact that the AI Prime specifically, and exclusively, mentioned AI excelling in shared attributes – and never mentions distinctive attributes. In short, learning about AI advances in shared attributes (e.g., language and logic) led participants to reconsider and rate distinctive attributes (e.g., morality and emotion) as more important to being human.

General Discussion

In light of recent technological progress enabling AI to rival humans in domains such as logic and language, we investigated whether priming such advances in Artificial Intelligence would alter what people see as essential to being human. After first establishing a two-part classification of human attributes – distinctive attributes, which AI are seen as less capable of (e.g., morality and love), and shared attributes, which AI are seen as more capable of (e.g., logic and language) – we tested whether learning about AI advances in shared attributes changed perceptions of being human. After reading about AI advances in attributes such as logic and language, participants tended to rate distinctive attributes such as humor and having desires as more essential to being human compared to participants who read about Trees. Contrary to our predictions, we did not find evidence that perceptions of shared attributes changed. Taken together, these results suggest that after learning about AI excellence in shared attributes, individuals in our subject population tend to distinguish humanity by prioritizing the attributes that separate humans from machines.

ARTIFICIAL INTELLIGENCE AND HUMAN NATURE

This paper makes two important contributions. First, it builds on previous accounts of humanness (e.g., Haslam, 2006; Gray et al., 2007; Weisman et al., 2017) and establishes a two-part classification of human attributes using AI as a referent. Second, using this classification, the research presented here shows that something as fundamental as the perception of being human can change following exposure to a brief stimulus describing attributes shared between AI and humans. This adds to the growing literature examining how AI affects how people perceive one another (e.g., Gamez-Djokic & Waytz, 2020; Jackson, Castelo, & Gray, 2020).

The present findings open up fascinating avenues for further work. First, future research should attempt to pinpoint the mechanism by which perceptions of human nature change. Though we did find support that the AI prime increased perceptions of machine competitiveness and anxiety (see Studies 2b and 2c in the supplement, respectively), we failed to find support that these effects mediated the effect of AI Prime on ratings of human essence. More straightforward constructs implied by the compensation effect (Cambon & Yzerbyt, 2018), such as perceptions of human uniqueness (see Ferrari et al., 2016 for a similar investigation with robots), might better mediate the effect of AI exposure on human essence perceptions. Second, future research might look at related stimuli as well as related dependent measures. Stimuli describing machines acquiring not only shared attributes but distinctive attributes – such as AI telling jokes (Ampil, 2019) or reading emotions (Heckman, 2020) – might also undermine the sense of human distinctiveness and lead to changes in the perception of human nature. In addition, learning about AI advances in shared attributes might change perceptions not just of humans but of AI: as the AI Prime was arguably a manipulation that implied that AI is more competent than humans, it might lead participants to perceiving AI as cruel, sinister or other low-warmth attributes (see Yzerbyt et al., 2005 for a similar compensation effect).

It's also important to keep in mind potential limitations of the present work. First, our studies relied exclusively on U.S. participants on Mechanical Turk, a population that may be particularly sensitive to AI sophistication at shared attributes given that such participants do many tasks, such as transcription or picture classification (Hitlin, 2016), that could soon be replaced by machines, representing a direct threat to their livelihood. Future research ought to investigate other populations whose work involves distinctive attributes, making the AI Prime a less threatening stimulus. Second, we directly asked participants what they thought was essential to being human, but it is possible that participants only reevaluated human essence because they were asked to, and may not seize upon it spontaneously as a defense mechanism. Future research might use an open-ended question to see if participants spontaneously change how they think about human nature.

What does it mean to be human? Recent advances in Artificial Intelligence enabling AI to excel in such human attributes at logic and language, only made possible by human ingenuity in logic and language, may ironically be leading to the prioritization of human distinctive attributes, such as humor or having desires, as people try to distinguish themselves from AI.

Author Contributions: The first author originated the study idea. Both authors jointly created study design and pre-registrations. The first author performed data collection and data analysis, with input from the second author. The first author drafted the manuscript and supplementary materials, and the second author provided substantial revisions. Both authors approved the final version of the manuscript.

References

- Ampil, L. (2019, December 26). Can a robot make you laugh? -- Teaching an AI to tell jokes. *Medium*. <https://towardsdatascience.com/can-a-robot-make-you-laugh-teaching-an-ai-to-tell-jokes-815f1e1e689c>
- Asorey, H. (n.d.). How AI is transforming forecasting for the better. *Quotable by Salesforce*. <https://www.salesforce.com/quotable/articles/how-AI-is-transforming-forecasting-for-the-better/>
- Brand Lab for Accenture. (2017). Getting Smart about Artificial Intelligence. *WIRED*. <https://www.wired.com/brandlab/2017/03/getting-smart-artificial-intelligence/>
- Cambon, L., & Yzerbyt, V. Y. (2018). Two routes toward compensation: An investigation into the mechanisms of compensation for high-and low-status groups. *Journal of Experimental Social Psychology*, 77, 24-35. <https://doi.org/10.1016/j.jesp.2018.04.002>
- Cloud. (n.d.) *Wikipedia*. <https://en.wikipedia.org/wiki/Cloud>
- Ferrari, Francesco, Paladino, Maria Paola, and Jetten, Jolanda (2016). *Blurring human-machine distinctions: anthropomorphic appearance in social robots as a threat to human distinctiveness*. *International Journal of Social Robotics* 8 (2) 287-302. <https://doi.org/10.1007/s12369-016-0338-y>
- Gamez-Djokic, M., & Waytz, A. (2020). Concerns About Automation and Negative Sentiment Toward Immigration. *Psychological Science*, 31(8), 987-1000. <https://doi.org/10.1177/0956797620929977>
- Giant forest (n.d.). *Wikipedia*. https://en.wikipedia.org/wiki/Giant_Forest
- Gray, H. M., Gray, K., & Wegner, D. M. (2007). Dimensions of mind perception. *Science*, 315(5812), 619-619. <https://doi.org/10.1126/science.1134475>

ARTIFICIAL INTELLIGENCE AND HUMAN NATURE

Haslam, N. (2006). Dehumanization: An integrative review. *Personality and Social Psychology Review*, 10(3), 252-264. https://doi.org/10.1207/s15327957pspr1003_4

Heaven, W. (2020, July 20). OpenAI's new language generator GPT-3 is shockingly good – and completely mindless. *MIT Technology Review*.
<https://www.technologyreview.com/2020/07/20/1005454/openai-machine-learning-language-generator-gpt-3-nlp/>

Heckman, C. (2020, January 8). AI can now read emotions - should it? *The Conversation*.
<https://theconversation.com/ai-can-now-read-emotions-should-it-128988>

Hiner, J. (2017, December 3). AI and jobs: Where humans are better than algorithms, and vice versa. *ZD Net*. <https://www.zdnet.com/article/ai-and-jobs-where-humans-are-better-than-algorithms-and-vise-versa/>

Hitlin, P. (2016, July 11). Research in the crowdsourcing age, a case study. *Pew Research Center*. <http://www.pewinternet.org/2016/07/11/research-in-the-crowdsourcing-age-a-case-study/>

Hitt, D. (2017, Dec. 19). What are clouds? *NASA Educational Technology Services*.
<https://www.nasa.gov/audience/forstudents/k-4/stories/nasa-knows/what-are-clouds-k4.html>

Hof, R (2013, April 23). Deep learning: Artificial Intelligence finally gets smart.
<https://robhof.com/2013/04/23/deep-learning-artificial-intelligence-finally-gets-smart/>

Jackson, J. C., Castelo, N., & Gray, K. (2020). Could a rising robot workforce make humans less prejudiced? *American Psychologist*, 75(7), 969–982. <https://doi.org/10.1037/amp0000582>

Katz, M. (2017, November 11). Welcome to the era of the AI coworker. *WIRED*.
<https://www.wired.com/story/welcome-to-the-era-of-the-ai-coworker/>

ARTIFICIAL INTELLIGENCE AND HUMAN NATURE

Lambert, P. (2018, May 8). Subject: Write emails faster with Smart Compose in Gmail. *Google*.

<https://www.blog.google/products/gmail/subject-write-emails-faster-smart-compose-gmail/>

Martin, A. (2018, January 16). Computers now read better than humans: Microsoft and Alibaba's

AI have taken the edge in comprehension tests. *alphr*. <https://www.alphr.com/artificial-intelligence/1008206/ai-now-reads-better-than-humans>

Mooallem, J. (2017, March 23). In the land of giants: Communing with some of the biggest trees on Earth. *The New York Times Magazine*.

Savage, N. (2020). How AI is improving cancer diagnostics. *Nature*, 579(7800), S14-S16.

<https://doi.org/10.1038/d41586-020-00847-2>

Sequoiadendron giganteum (n.d.). *Wikipedia*.

https://en.wikipedia.org/wiki/Sequoiadendron_giganteum

Santoro, E. & Monin, B. (2019). Of human essence and obsolescence: learning about advances in AI changes perceptions of human nature. In E. Santoro & B. Monin (Chairs), *Evolving perceptions of humans and intelligent machines* [Symposium]. Society for Personality and Social Psychology, Portland, OR, United States.

Strogatz, S. (2018, December 26). One giant step for a chess-playing machine. *The New York Times*. <https://www.nytimes.com/2018/12/26/science/chess-artificial-intelligence.html>

Szalay, J. (2017, May 5). Giant sequoias and redwoods: The largest and tallest trees. *Live Science*. <https://www.livescience.com/39461-sequoias-redwood-trees.html>

ARTIFICIAL INTELLIGENCE AND HUMAN NATURE

Tajfel, H., & Turner, J. C. (1979). An integrative theory of intergroup conflict. In W. G. Austin & S. Worchel (Eds.), *The social psychology of intergroup relations* (pp. 33-47).

Brooks/Cole.

Weisman, K., Dweck, C. S., & Markman, E. M. (2017). Rethinking people's conceptions of mental life. *Proceedings of the National Academy of Sciences*, 114(43), 11374-11379.
<https://doi.org/10.1073/pnas.1704347114>

Yzerbyt, V., Provost, V., & Corneille, O. (2005). Not competent but warm... really?
Compensatory stereotypes in the French-speaking world. *Group Processes & Intergroup Relations*, 8(3), 291-308. <https://doi.org/10.1177/1368430205053944>