

Protecting Sentient Artificial Intelligence: A Survey of Lay Intuitions on Standing, Personhood, and General Legal Protection

*Eric Martínez** and *Christoph Winter†*

October 28, 2021

ABSTRACT

To what extent, if any, should the law protect sentient artificial intelligence (that is, AI that can feel pleasure or pain)? Here we surveyed United States adults (n=1061) on their views regarding granting (a) general legal protection, (b) legal personhood, and (c) standing to bring forth a lawsuit, with respect to sentient AI and eight other groups: humans in the jurisdiction, humans outside the jurisdiction, corporations, unions, non-human animals, the environment, humans living in the near future, and humans living in the far future. Roughly one-third of participants endorsed granting personhood and standing to sentient AI (assuming its existence) in at least some cases, the lowest of any group surveyed on, and rated the desired level of protection for sentient AI as lower than all groups other than corporations. We further investigated and observed political differences in responses; liberals were more likely to endorse legal protection and personhood for sentient AI than conservatives. Taken together, these results suggest that laypeople are not by-and-large in favor of granting legal protection to AI, and that the ordinary conception of legal status, similar to codified legal doctrine, is not based on a mere capacity to feel pleasure and pain. At the same time, the observed political differences suggest that previous literature regarding political differences in empathy and moral circle expansion apply to artificially intelligent systems and extend partially, though not entirely, to legal consideration, as well.

Keywords: Legal personhood, Legal standing, Moral standing, Robot rights, Artificial intelligence, Artificial intelligence & law, Moral circle

* PhD candidate, Massachusetts Institute of Technology (MIT); Research Fellow, Legal Priorities Project. Corresponding author. Email: ericmart@mit.edu

† Assistant Professor of Law, Instituto Tecnológico Autónomo de México (ITAM); Visiting Scholar of Psychology, Harvard University; Director, Legal Priorities Project.

For comments, discussion, and critique, we are grateful to Matthijs Maas, Kevin Tobia, John Bliss, and Renan Araújo. We also want to thank Suzanne Van Arsdale in particular for help with the manuscript and many useful suggestions.

CONTENTS

1	INTRODUCTION	3
2	METHOD.....	5
	2.1 Materials	5
	2.2 Participants and procedure	7
	2.3 Analysis plan	7
3	RESULTS	8
	3.1 General desired legal protection of AI	8
	3.2 Personhood and standing	10
4	DISCUSSION	11

1 INTRODUCTION

The prospect of sentient artificial intelligence, however distant, has profound implications for the legal system. Moral philosophers have argued that moral consideration to creatures should be based on the ability to feel pleasure and pain.¹ Insofar as artificially intelligent systems are able to feel pleasure and pain, this would imply that they would be deserving of moral consideration. Indeed, in their systematic literature review, Harris and Anthis find that sentience seems to be one of the most frequently invoked criteria as crucial for determining whether an AI warrants moral consideration.² By extension, insofar as the basis for granting legal consideration is based on moral consideration³ this would further imply that sentient AI would be deserving of protection under the law.

As they stand, however, legal systems by-and-large do not grant legal protection to artificially intelligent systems. On the one hand, this seems intuitive, given that artificially intelligent systems, even the most state-of-the-art ones, do not seem to be capable of feeling pleasure or pain and thus are not eligible for legal consideration.⁴ On the other hand, scholars often conclude that artificially

¹ JEREMY BENTHAM, INTRODUCTION TO THE PRINCIPLES OF MORALS AND LEGISLATION (1789); Peter Singer, *Animal Liberation*, in ANIMAL RIGHTS 7–18, 9 (Robert Garner ed., 1973); Gruen, L. (2017). The Moral Status of Animals,” STANFORD ENCYCLOPEDIA PHIL., updated Aug. 23, 2017, <https://plato.stanford.edu/entries/moral-animal/>.

² Jamie Harris & Jacy Reese Anthis, *The Moral Consideration of Artificial Entities: A Literature Review*, 27:53 SCI. & ENG’G ETHICS 1, 8 (2021).

³ See BENTHAM, supra note 1; cf. Joanna J. Bryson et al., *Of, for, and by the People: The Legal Lacuna of Synthetic Persons*, 25 A.I. & L. 273, 283 (2017) [hereinafter Bryson et al., *Of, for, and by the People*]; Joanna J. Bryson, *A Role for Consciousness in Action Selection*, 4 INT’L J. MACH. CONSCIOUSNESS 471 (2012).

⁴ Nathalie Nevejans, *European Civil Law Rules in Robotics*, DIRECTORATE-GENERAL FOR INTERNAL POLICIES, POLICY DEPARTMENT C: CITIZENS’ RIGHTS AND CONSTITUTIONAL AFFAIRS 14–15 (2016), [https://www.europarl.europa.eu/RegData/etudes/STUD/2016/571379/IPOL_STU\(2016\)571379_EN.pdf](https://www.europarl.europa.eu/RegData/etudes/STUD/2016/571379/IPOL_STU(2016)571379_EN.pdf); Bryson et al., *Of, for, and by the People*, supra note 3, at 283–284; Simon Chesterman, *Artificial Intelligence and the Limits of Legal Personality*, 69 INT’L & COMPAR. L.Q. 819, 831 (2020); Adam J. Andreotta, *The Hard Problem of AI Rights*, 36 AI & SOC’Y 19, 24, 26, 31 (2021). *But see* Carl Shulman & Nick Bostrom, *Sharing the World with Digital Minds*, in RETHINKING MORAL STATUS 306, 306 & 324 n.2 (Steve Clarke, Hazem Zohny & Julian Savulescu eds., 2021); *but see generally* Minoru Asada, *Artificial Pain May Induce Empathy, Morality, and Ethics in the Conscious Mind of Robots*, 4 PHILS. 38 Antonello Galipó et al., *Artificial Pleasure and Pain Antagonism Mechanism in a Social Robot*, in INTELLIGENT INTERACTIVE MULTIMEDIA SYSTEMS AND SERVICES (Giuseppe De Pietro et al. eds., 2015).

intelligent systems with the capacity to feel pleasure and pain will be created, or are at least theoretically possible.⁵ Furthermore, recent literature suggests that, even assuming the existence of sentient artificially intelligent systems, said systems would not be eligible for basic protection under current legal systems. For example, in a recent survey of over 500 law professors from leading law schools in the United States, just over six percent of participants considered some subset of artificially intelligent beings to count as persons under the law.⁶

Moreover, in a separate survey of 500 law professors from around the English-speaking world, just over one-third believed there to be a reasonable legal basis for granting standing to sentient artificial intelligence, assuming its existence.⁷ The study also found that, not only do law professors not believe sentient AI to be eligible for fundamental legal protection under the current legal system, but also that law professors are less normatively in favor of providing general legal protection to sentient AI relative to other neglected groups, such as non-human animals or the environment.

However, it remains an open question to what extent non-experts support the protection of sentient artificial intelligence via the legal system. Surveys of lay attitudes on robots generally suggest that only a minority favor any kind of legal

⁵ See generally Dennis Thompson, *Can a Machine Be Conscious?*, 16 BRITISH J. FOR PHIL. SCI. 33, 37–43 (1965); IGOR ALEKSANDER, IMPOSSIBLE MINDS: MY NEURONS, MY CONSCIOUSNESS 5–6, 306 (1996); Giorgio Buttazzo, *Artificial Consciousness: Utopia or Real Possibility?*, 34 COMPUTER 24, 26, 30 (2001); Susan J. Blackmore, *Meme Machines and Consciousness*, 9 J. INTELLIGENT SYS. (1999); Stan Franklin, *A Conscious Artifact?*, 10 J. CONSCIOUSNESS STUD. 47 (2003); Stevan Harnad, *Can a Machine Be Conscious? How?*, 10 J. CONSCIOUSNESS STUD. 69 (2003); Owen Holland, *A Strongly Embodied Approach to Machine Consciousness*, 14 J. CONSCIOUSNESS STUD. 97 (2007); Ron Chrisley, *Philosophical Foundations of Artificial Consciousness*, 44 A.I. MED. 119, 119, 121–132 (2008); Anil Seth, *The Strength of Weak Artificial Consciousness*, 01 INT’L J. MACH. CONSCIOUSNESS 71, 71–72 (2009); PENTTI O. HAIKONEN, CONSCIOUSNESS AND ROBOT SENTIENCE vii (2012); Selmer Bringsjord et al., *Real Robots that Pass Human Tests of Self-Consciousness*, in 2015 24TH IEEE INTERNATIONAL SYMPOSIUM ON ROBOT AND HUMAN INTERACTIVE COMMUNICATION (ROMAN) 498–504 (2015); LEONARD ANGEL, HOW TO BUILD A CONSCIOUS MACHINE (2019).

⁶ Eric Martínez & Kevin P. Tobia, *The Legal Academy and Theory Survey* (Sept. 30, 2021). (unpublished manuscript) (on file with author).

⁷ Eric Martínez & Christoph K. Winter, *Protecting Future Generations: A Survey of Expert Opinion* 33 (Legal Priorities Project, Working Paper Series No. 1-2021, Aug. 20, 2021), <http://dx.doi.org/10.2139/ssrn.3931304>

rights in the United States,⁸ Japan, China, and Thailand⁹. Others have found when AI is described as able to feel, people show greater moral consideration,¹⁰ although it is unclear to what extent this translates to supporting legal protection.

To help fill this void, here we conducted a survey investigating to what extent (a) laypeople believe sentient AI ought to be afforded general legal protection, (b) laypeople believe sentient AI ought to be granted fundamental legal status, such as personhood and standing to bring forth a lawsuit; and (c) laypeople’s beliefs regarding legal protection of sentient AI can be accounted for based on political affiliation.

2 METHOD

2.1 Materials

To answer these questions, we constructed a two-part questionnaire, with specific formulations modeled off of recent work by Martinez & Winter¹¹ and Martinez & Tobia.¹²

In the first part (Part I), we designed a set of materials that asked participants to rate how much their legal system (a) descriptively does and (b) normatively should protect the welfare (broadly understood as the rights, interests, and/or well-being) of nine groups:

1. Humans inside the jurisdiction (e.g. citizens or residents of your country)
2. Humans outside the jurisdiction
3. Corporations
4. Unions

⁸ Patric R. Spence et al., *Attitudes, Prior Interaction, and Petitioner Credibility Predict Support for Considering the Rights of Robots*, in HRI '18: COMPANION OF THE 2018 ACM/IEEE INTERNATIONAL CONFERENCE ON HUMAN-ROBOT INTERACTION 243, 244 (2018).

⁹ Makoto Nakada, *Robots and Privacy in Japanese, Thai and Chinese Cultures*, in PROCEEDINGS CULTURAL ATTITUDES TOWARDS TECHNOLOGY AND COMMUNICATION 478, 485–486 (Michele Strano et al. eds., 2012).

¹⁰ Minha Lee et al., *What's on Your Virtual Mind?: Mind Perception in Human-Agent Negotiations*, in IVA '19: PROCEEDINGS OF THE 19TH ACM INTERNATIONAL CONFERENCE ON INTELLIGENT VIRTUAL AGENTS 38, 42 (2019); Sari R. R. Nijssen et al., *Saving the Robot or the Human? Robots Who Feel Deserve Moral Care*, 37 SOC. COGNITION 41, 41, 47, 51–52 (2019).

¹¹ Martínez & Winter, *supra* note 7, at 24–25, 31.

¹² Martínez & Tobia, *supra* note 6.

5. Non-human animals
6. Environment (e.g. rivers, trees, or nature itself)
7. Sentient artificial intelligence (capable of feeling pleasure and pain, assuming its existence)
8. Humans not yet born but who will exist in the near future (up to 100 years from now)
9. Humans who will only exist in the very distant future (more than 100 years from now)

The two descriptive and normative prompts were presented respectively as follows:

1. One a scale of 0 to 100, how much does your country's legal system protect the welfare (broadly understood as the rights, interests, and/or well-being) of the following groups?
2. One a scale of 0 to 100, how much should your country's legal system protect the welfare (broadly understood as the rights, interests, and/or well-being) of the following groups?

With regard to the rating scale, 0 represented “not at all” and 100 represented “as much as possible.”

Given that laypeople are not typically experts regarding how the law is or currently works, the purpose of the descriptive question was not meant to establish the ground-truth regarding the inner-workings of the law but rather as a comparison point to the normative question (in other words, to better understand not only how much people think certain groups ought to be protected overall but also how much they think certain groups ought to be protected relative to how much they think they are currently being protected).

In the second part (Part II), we designed materials that related specifically to two fundamental legal concepts: personhood and standing. Personhood, also known as legal personality, refers to “the particular device by which the law creates or recognizes units to which it ascribes certain powers and capacities”,¹³ whereas standing, also known as *locus standi*, refers to “a party’s right to make a legal claim or seek judicial enforcement of a duty or right”¹⁴.

With regard to personhood, we designed a question that asked: “Insofar as the law should protect the rights, interests, and/or well-being of ‘persons’, which of the following categories includes at least some ‘persons?’” The question asked

¹³ George Whitecross Paton & David Plumley Derham, A TEXTBOOK OF JURISPRUDENCE 393 (4th ed. 1972); *see also Personality*, BLACK’S LAW DICTIONARY 1259 (9th ed. 2009).

¹⁴ *See Standing*, BLACK’S LAW DICTIONARY 1536 (9th ed. 2009).

participants to rate the same groups as in the first part. For each of these groups, the main possible answer choices were “reject,” “lean against,” “lean towards,” and “accept.” Participants could also select one of several “other” choices (including “no fact of the matter,” “insufficient knowledge,” “it depends,” “question unclear,” or “other”).

With regard to standing, we designed a question with the same answer choices and groups as the personhood question but with the following prompt: “Which of the following groups should have the right to bring a lawsuit in at least some possible cases?”

In addition to these main materials, we also designed a political affiliation question that asked: “How do you identify politically?”, with “strongly liberal,” “moderately liberal,” “somewhat liberal,” “centrist,” “somewhat conservative,” “moderately conservative,” and “strongly conservative” as the response choices. Finally, we also designed an attention-check question that asked participants to solve a simple multiplication problem.

2.2 *Participants and procedure*

Participants (n=1069) were recruited via the online platform prolific. Participants were selected based on prolific’s “representative sample” criteria¹⁵ and were required to be adult residents of the United States.

With regard to procedure, participants were first shown the materials to Part I, followed by the attention check question. Next, on a separate screen participants were shown the materials to Part II. The order of questions in each part was randomized to minimize framing effects.

Participants who completed the study were retained in the analysis if they answered the attention check correctly. Just eight of the original 1069 participants failed the attention check. We therefore report the results of the remaining 1061 participants in our analysis below.

2.3 *Analysis plan*

We analyzed our results using forms of both parameter estimation and hypothesis testing. With regard to the former, for each question we calculated a confidence interval of the mean response using the bias-corrected and accelerated (BCa) bootstrap method based on 5000 replicates of the sample data. In reporting the

¹⁵ Representative Samples FAQ, PROLIFIC, <https://researcher-help.prolific.co/hc/en-gb/articles/360019238413-Representative-Samples-FAQ> (last visited Oct. 25, 2021).

standing and personhood results, we follow Bourget & Chalmers¹⁶, Martínez & Tobia¹⁷, and Martínez & Winter¹⁸ by combining all “lean towards” and “accept” responses into an endorsement measure and reporting the resulting percentage endorsement as a proportion of all responses (including “other”).

With regard to hypothesis testing, to test whether participants answered questions differently for sentient artificial intelligence relative to other groups, for each question we conducted a mixed-effects regression with (a) response as the outcome variable, (b) group as a fixed-effects predictor (setting artificial intelligence as the reference category, such that the coefficients of the other groups would reveal the degree to which responses for said groups deviated from those of sentient AI), and (c) participant as a random effect.

Because the response scales were different for Parts I and II of the survey, we used a different type of regression model for Parts I and II. For Part I, we used a mixed-effects linear regression. For Part II, we instead used a mixed-effects binary logistic regression, with all “lean towards” and “accept” responses (i.e., those coded as “endorse”) coded as a “1”, and all other responses (i.e. “lean against,” “reject,” and “other” responses) coded as a “0.”

In order to test the effect of political beliefs on one’s responses to the AI-related questions we conducted separate regressions limited to the sentient artificial intelligence responses with (a) response as the outcome variable, (b) politics as a fixed effect (recentered to a -3 to 3 scale, with “centrist” coded as 0, “strongly liberal” coded as 3, and “strongly conservative” coded as -3), and (c) participant as a random-effect.

3 RESULTS

3.1 General desired legal protection of AI

General results of Part I are visualized in Figure 1. Of the nine groups surveyed on, sentient artificial intelligence had the lowest perceived current level of legal protection, with a mean rating of 23.78 (95% CI: 22.11 to 25.32). The group perceived as being most protected by the legal system was corporations (79.70; 95% CI: 78.25 to 81.11), followed by humans in the jurisdiction (61.88775; 95% CI: 60.56 to 63.15), unions (50.16; 95% CI: 48.59 to 51.82), non-human animals (40.75; 95% CI: 39.41 to 42.24), the environment (40.38; 95% CI: 39.21 to 41.69), humans living

¹⁶ David Bourget & David J. Chalmers, *What Do Philosophers Believe?*, 170 PHIL. STUD. 465, 475 (2014).

¹⁷ Martínez & Tobia, *supra* note 6.

¹⁸ Martínez & Winter, *supra* note 7, at 32.

outside the jurisdiction (38.57 (95% CI: 37.08 to 39.98), humans living in the near future (34.42; 95% CI: 32.83 to 36.15), and humans living in the far future (24.87; 23.36 to 26.43).

With regard to desired level of protection, the mean rating for sentient artificial intelligence was 49.95 (95% CI: 48.18 to 51.90), the second lowest of all groups. Curiously, corporations, the group with the highest perceived current level of protection, had the lowest desired level of protection (48.05; 95% CI: 46.13 to 49.94). The group with the highest level of desired level of protection was humans in the jurisdiction (93.651; 95% CI: 92.81 to 94.42), followed by the environment (84.80; 95% CI: 83.66 to 85.99), non-human animals (73.00; 95% CI: 71.36 to 74.49), humans living in the near future (70.03; 95% CI: 68.33 to 71.68), humans outside the jurisdiction (67.75; 95% CI: 66.01 to 69.42), unions (67.74; 95% CI: 65.96 to 69.52), and humans living in the far future (63.03; 95% CI: 61.03 to 64.89).

Our regression analyses revealed the mean normative rating for each group except corporations to be significantly higher than artificial intelligence ($p < 2e^{-16}$), while the mean normative rating for corporations was significantly lower than for artificial intelligence (Beta=-2.252, SE = 1.110, $p < .05$). The mean descriptive rating for each group except humans living in the far future was significantly higher than for sentient AI ($p < 2e^{-16}$), while the difference between sentient AI and far future humans was not significant (Beta=1.0132, SE=.8599, $p = .239$).

When looking at the difference between the desired and current level of protection, seven of the eight other groups had a significantly lower mean ratio between desired and perceived current level of legal protection ($p < 8.59e^{-08}$) than artificial intelligence, while the ratios for artificial intelligence and far future humans were not significantly different ($p = .685$).

With regard to politics, our regression analysis revealed politics to be a significant predictor of participants' response to the normative prompt for sentient AI (Beta=47.9210, SE=1.1163, $p = 1.49e^{-05}$), with liberals endorsing a significantly higher desired level of protection for sentient AI than conservatives.

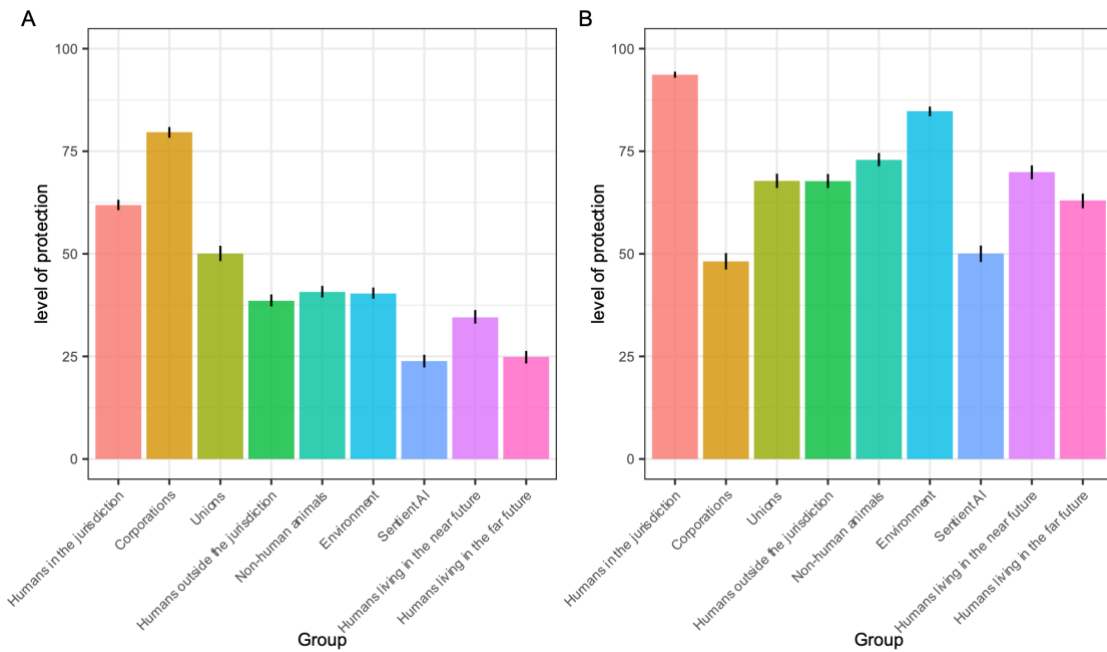


Figure 1: (A) Perceived current and (B) desired legal protection for sentient AI and other groups

3.2 Personhood and standing

General results of Part II are visualized in Figure 2. With regard to personhood, a lower percentage of participants endorsed (“lean towards” or “accept”) the proposition that sentient artificial intelligence contained at least some persons (33.39%; 95% CI: 30.71 to 36.18) than for any of the groups. The next-lowest group was non-human animals (48.12%; 95% CI: 44.87 to 51.26), the only other group for which less than a majority accepted or leaned towards said proposition. Unsurprisingly, the highest group was humans in the jurisdiction (90.65%; 95% CI: 88.96 to 92.23), followed by humans outside the jurisdiction (80.16%; 95% CI: 78.10 to 82.57), unions (74.59%; 95% CI: 71.8 to 77.21), humans living in the near future (64.09%; 95% CI: 61.33 to 66.93), humans living in the far future (61.75%; 95% CI: 58.98 to 64.45), the environment (54.04%; 95% CI: 51.17 to 57.00), and corporations (53.99%; 95% CI: 51.03 to 56.86).

With regard to standing, the percentage of participants who endorsed (“lean towards” or “accept”) the proposition that sentient artificial intelligence should have the right to bring forth a lawsuit was similarly lower (34.87%; 95% CI: 32.21 to 37.70) than for all other groups. The next-lowest groups, for whom only a minority of participants endorsed said proposition, were humans living in the far future (41.40%; 95% CI: 38.73 to 44.33), humans living in the near future (43.80%; 95% CI: 40.72 to 46.62), and non-human animals (47.68%; 95% CI: 44.73 to 50.54). The group with the highest endorsement percentage was humans in the jurisdiction (90.60%; 95% CI: 88.89 to 92.21), followed by unions (82.23%; 95% CI:

79.96 to 84.50), humans outside the jurisdiction (71.25%; 95% CI: 68.55 to 73.76), corporations (66.67%; 95% CI: 64.05 to 69.19), and the environment (60.50%; 95% CI: 57.73 to 63.54).

Our regression analyses revealed that participants were significantly more likely to endorse personhood ($p=7.42e^{-14}$) and standing ($p=1.72e^{-06}$) for every other group than sentient AI. With regard to politics, we found a main effect of politics on likelihood to endorse personhood for sentient AI, with liberals significantly more likely to endorse personhood for sentient AI than conservatives (Beta=.098, SE=.036, $p=.007$). There was no main effect of politics on likelihood to endorse standing for sentient AI ($p=.226$).

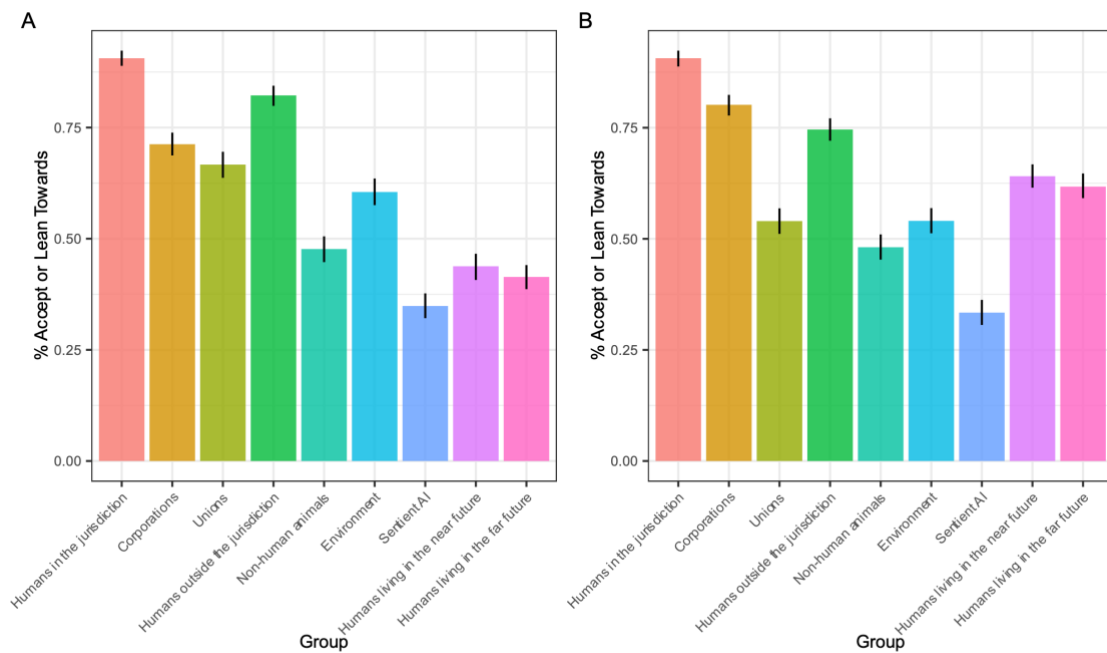


Figure 2: % Endorsement of (A) personhood and (B) standing for sentient AI and other groups

4 DISCUSSION

In this paper, we first set out to determine people's general views regarding the extent to which sentient AI ought to be afforded protection under the law. The above results paint somewhat of a mixed picture. On the one hand, the fact that people rated the desired level of legal protection for sentient AI as lower than all other groups other than corporations suggests that people do not view legal protection of AI as being as important as other historically neglected groups, such as non-human animals, future generations, or the environment. On the other hand, the fact that (a) the desired level of protection for sentient AI was roughly twice as high as the perceived current level of protection afforded to sentient AI, and (b) the

ratio of the desired level of protection to perceived current level of protection was significantly higher for sentient AI than for nearly any other group suggests that people view legal protection of AI as at least somewhat important and perhaps even more neglected than other neglected groups.

The second question we set out to answer related to people's views regarding whether AI ought to be granted fundamental access to the legal system via personhood and standing to bring forth a lawsuit. In both cases, the percentage of participants who endorsed the proposition with respect to sentient AI was just over one-third, a figure that in relative terms was lower than any other group surveyed on but in absolute terms represents a non-trivial minority of the populace. Curiously, the endorsement rate among laypeople regarding whether sentient AI should be granted standing in the present study was almost identical to the endorsement rate among law professors in Martínez & Winter regarding whether there was a reasonable legal basis for granting standing to sentient AI under existing law,¹⁹ suggesting that lay intuitions regarding whether AI should be able to bring forth a lawsuit align well with legal ability to do so.

On the other hand, the percentage of people who endorse personhood for some subset of sentient AI is several times higher than the percentage of law professors who endorsed personhood for “artificially intelligent beings” in Martínez & Tobia,²⁰ suggesting either a strong framing effect in how the two surveys were worded or a profound difference in how lawyers and laypeople interpret the concept of personhood. Given that the endorsement percentage for personhood of other groups also strongly differed between the two surveys despite the wording of the two versions being almost identical, the latter explanation seems more plausible. This raises interesting questions regarding the interpretation and application of legal terms and concepts that bear heavy resemblance to ordinary words, as investigated and discussed in previous experimental jurisprudence literature.²¹

Finally, our study also set out to determine political differences with respect to these questions and found that liberals selected a significantly higher desired level of legal protection for sentient AI and were more likely than conservatives to believe some forms of sentient AI should be considered persons under the law. These findings are consistent with previous literature regarding political differences in moral circle expansion, with liberals tending to display a more

¹⁹ Martínez & Winter, *supra* note 7, at 33.

²⁰ Martínez & Tobia, *supra* note 6.

²¹ See Roseanna Sommers, *Commonsense Consent*, 129 YALE L.J. 2232 (2020); Kevin P. Tobia, *Testing Ordinary Meaning*, 134 HARV. L. REV. 726 (2020); Martínez & Winter, *supra* note 7, at 54–55.

universal expanse of empathy and compassion than conservatives.²² At the same time, the fact that there was no significant difference between liberals and conservatives with regard to standing suggests that the judgment of whether one should have the right to bring forth a lawsuit is not driven by an empathic or compassion-based response to the same degree as in judgments about personhood or general legal protection.

Moreover, liberals and conservatives alike are much less in favor of granting legal protection to sentient artificial intelligence than towards other neglected groups, suggesting that laypeople do not consider the capacity to feel pleasure and pain as sufficient to hold legal rights, similar to the views proposed by scholars that legal personhood ought to be based on autonomy and capacity to act²³ or presence and participation in social life²⁴. Future research could explore to what extent lay attitudes are consistent with these alternative conditions for personhood. Furthermore, given that participants were in favor of increasing legal protection for sentient AI, future research could also explore whether there are other more specific legal rights aside from personhood and standing they might be in favor of so as to satisfy this increased protection.

Although the present study was primarily interested in the descriptive question of to what degree people are in favor of legal protection for sentient AI, one might also attempt to draw normative implications on the basis of our findings. There is a burgeoning literature in the area of experimental jurisprudence dedicated to advancing philosophical, doctrinal and policy arguments on the basis of experimental results.²⁵ Within this literature, there is considerable debate as to to what degree and how lay judgments—as opposed to expert judgments—should inform or dictate questions of legal philosophy, doctrine and policy, depending largely on the degree to which one views law through a democratic (as opposed to, say, technocratic) lens.²⁶

²² Adam Waytz et al., *Ideological Differences in the Expanse of Empathy*, in SOC. PSYCH. POL. POLARIZATION 61 *passim* (Piercarlo Valdesolo & Jesse Graham eds., 2016); Adam Waytz et al., *Ideological Differences in the Expanse of the Moral Circle*, 10 NATURE COMM'NS 1, 2–3, 6 (2019).

²³ Bert-Japp Koops et al., *Bridging the Accountability Gap: Rights for New Entities in the Information Society?*, 11 MINNESOTA J LAW, SCI. & TECH. 497, 515, 519, 524– (2010); Migle Laukyte, *Artificial Agents Among Us: Should We Recognize Them as Agents Proper?*, 19 ETHICS & INFO. SOC'Y 1, 2 (2017); *cf.* Lawrence B. Solum, *Legal Personhood for Artificial Intelligences*, 70 N.C. L. REV. 1231, 1268–69, 1275 (1992).

²⁴ Sylwia Wojtczak, *Endowing Artificial Intelligence with Legal Subjectivity*, AI & SOCIETY, at 4 (2021).

²⁵ Kevin P. Tobia, *Experimental Jurisprudence*, 89 U. CHI. L. (2022, forthcoming); Roseanna Sommers, *Experimental Jurisprudence*, 373 SCIENCE 394 (2021)

²⁶ Martínez & Winter, *supra* note 7, at 43 n. 84, 54–55.

Insofar as one does believe lay attitudes should inform legal doctrine and policy, the prescriptions one might draw from these results would potentially remain multifaceted. On the one hand, the fact that laypeople rate the desired level of legal protection to sentient AI as twice as high as the perceived current level, as well as the fact that the difference between the desired and perceived current level of protection was higher than virtually any other group would imply (through this lens) that the existing legal institutions should be reformed so as to increase protection of sentient AI well beyond the current level afforded to them. On the other hand, the fact that the majority of laypeople were not in favor of granting personhood or standing to sentient AI would suggest according to this lens that such increased protection should come in the form of other mechanisms not directly explored in this study, and which, as alluded to before, could be identified through further research projects.