

# Can Artificial Intelligence Make Art?

Folk Intuitions as to whether AI-driven Robots Can Be Viewed as Artists and Produce Art

Elzė Sigutė Mikalonytė \*

Institute of Philosophy, Vilnius University, elze.mikalonyte@fsf.vu.lt

Markus Kneer

Department of Philosophy, University of Zurich, markus.kneer@uzh.ch

In two experiments (total N=693) we explored whether people are willing to consider paintings made by AI-driven robots as *art*, and robots as *artists*. Across the two experiments, we manipulated three factors: (i) agent type (AI-driven robot v. human agent), (ii) behavior type (intentional creation of a painting v. accidental creation), and (iii) object type (abstract v. representational painting). We found that people judge robot paintings and human painting as art to roughly the same extent. However, people are much less willing to consider robots as artists than humans, which is partially explained by the fact that they are less disposed to attribute artistic intentions to robots.

**CCS CONCEPTS • Human-centered computing ~Human computer interaction (HCI) ~Empirical studies in HCI • Applied Computing ~Arts and humanities ~Fine arts • Computing methodologies ~Artificial Intelligence ~ Philosophical/theoretical foundations of artificial intelligence ~Cognitive science**

**Additional Keywords and Phrases:** artificial intelligence, creativity, aesthetics, AI art, mental states

## 1 INTRODUCTION

Nowadays, encounters with works of art created by robots are no longer unusual: We come across paintings, poems, songs and even film scripts written by machines. These works are presented to us in the usual ways that artworks get presented in. Robots are even getting solo exhibitions [2] and having their works bought for hundreds of thousands of dollars [12]. And yet, when it comes to the question whether these robot-created objects are real works of art, opinions diverge. Art has traditionally been considered to be one of those domains exclusive to humans, as creativity – sometimes called “the final frontier” of AI research [14] – is highly valued by society, and is not that easily attributed to non-human entities, especially those which do not have mental states. There is a reasonable doubt whether robots can really create art.

---

\* The authors contributed equally to this publication. Kneer is also affiliated with the Digital Society Initiative (University of Zurich).

If in our effort to answer this question we turn to the classic Turing test [67] the answer appears quite straightforward: We already know that artificial intelligence can produce art-like works that are not only indistinguishable from those produced by human agents, but that these works will also be perceived as having no less of an aesthetic value [23, 8]. However, many believe that robots cannot create art. Common reasons include that machines do not have human-like intelligence, autonomy, mental states, emotions or – partially as a consequence of the latter – the agency necessary to participate in social relations [32, 33].

If, on the other hand, one deems (certain types of) robot-made creations art, a second interesting question arises: Are we willing to consider robots as artists? And if so, under what conditions? One might, following d’Inverno and McCormack [21] distinguish between “Heroic AI” and “Collaborative AI”. The former refers to independent creative autonomous agents, the latter to AI which is part of a group agent that includes humans. Naturally, there are also non-autonomous machines designed to be used as mere tools in the process of art creation. In the literature on robot ethics, there are proposals to attribute collective agency to human-robot collaborations [52] and this proposal could be extended to aesthetic agency as well. Others, however, attribute the authorship of the non-human created artworks solely to the human authors of the creative machines [64]. Just as in moral contexts, where we feel uncertain of who is responsible for robot behaviour [63, 10, 47], in aesthetics we might feel unsure who is responsible for artistic robot creation.

The authorship problem is also reflected in the legal discussion on the ownership of the copyrights of the artworks produced by robots. Although works of art that are autonomously created by AI are currently not copyrighted, there are proposals to assign authorship to robots by redefining “authorship” to include non-human agents [16, 1, 60], while others propose to transfer the copyrights on the robot-created works to their human designers [9, 35].

### **1.1 Folk intuitions about art made by AI**

To what extent are people willing to call robots artists and their creations art? Although there is agreement among researchers that machines do not have consciousness [18, 34], people have a strong tendency to anthropomorphise robots and tend to ascribe a wide variety of mental states to them (for a review, see [55], as well as the references in section 1.2). While recently a lot of research has been conducted on intuitions regarding the perceived moral agency of robots, the same cannot be said about intuitions regarding the ability of robots to create art, and this area remains underexplored. Although some claim that “AI systems are not broadly accepted as authors by artistic or general public communities” [48], there is currently almost no empirical research on folk intuitions (one exception is [26]).

Even if we discover that folk intuitions suggest that robots do create art and have all the mental states needed to be real creators, it won’t mean that those intuitions should be decisive in either philosophical or legal discussions. On the contrary, many authors see the tendency to anthropomorphise AI as problematic [59] and caution care about the ascription of rich psychological states to robots [62]. In the same vein, some authors claim that it would be “misleading and risky” not to raise awareness of this tendency to project human traits onto AI [61]. Naturally, to better understand the folk propensity to ascribe agentic traits to robots, more empirical research is needed and we should look beyond moral contexts only.

How should empirical research on intuitions about robots’ ability to create art be conducted? If we examine the positions on the possibility of the robot-created art from the philosophical point of view, it is helpful to consider Mark Coeckelbergh’s [11] proposal to break the question “can machines create art?” down into three

smaller ones: what do we mean by creation, by art, and by machines? In trying to determine whether robots are perceived as being able to create art, we need to single out three aspects: the *agent* (autonomous robot, as compared to human), the *process* (the action by which a work is brought to life), and the *product* (the created object).

## 1.2 Mental state ascription

Many scholars define art in such a way, that artworks must be created by agents who have mental states. One of the most famous accounts is Jerrold Levinson's intentional-historical definition of art, which states that:

"X is an artwork at [time] t = df X is an object of which it is true at t that some person or persons, having the appropriate proprietary right over X, non-passingly intends (or intended) X for regard-as-a-work-of-art, *i. e.*, regard in any way (or ways) in which objects in the extension of 'artwork' prior to t are or were correctly (or standardly) regarded." [41]

Beardsley's aesthetic definition of art also stresses the creator's intention, although it does not treat it as a necessary condition for an object to be considered a work of art. It states that a work of art is

"[...] either an arrangement of conditions intended to be capable of affording an experience with marked aesthetic character or (incidentally) an arrangement belonging to a class or type of arrangements that is typically intended to have this capacity." [6]

There is also a romantic conception of the creation of art, according to which the latter is an expression of the creator's inner world and the creator's emotions [13]. All these conceptions of art are hard to reconcile with the position that robots can be creative – they (arguably) do not have intentions and there isn't much of an inner world to express.

As we know from research in the psychology of art, inferences about creators' mental states play an important role in our reasoning about artworks. Bloom (who draws inspiration from Levinson's intentional-historical definition of art) claims that people reason about artworks (and all artifacts) in terms of inferred authorial intent. They categorize an object as a member of the artifact kind if it is thought to be created with the intention for it to belong to that kind:

"We infer that a novel entity has been successfully created with the intention to be a member of artifact kind X – and thus is a member of artifact kind X – if its appearance and potential use are best explained as resulting from the intention to create a member of artifact kind X." [7]

Inferences about the mental states of creators are as important in categorizing works of art as other artifacts. According to Bloom, authorial intent is relevant to determining the kind of an artifact even if it is not directly connected to the object's appearance or function. Thus, intuitions on whether an object was created with an intention for it to belong to the category of artworks must play an important part in judgments whether an object should be categorized as art, which means that people's tendency to ascribe intentions to robots would influence their judgments on whether the robot-created object should be admitted to the realm of artworks.

Recent empirical work confirms Bloom's proposal: When judging whether an object falls under the category of "artwork" the intent of the creator is seen as more important than even the appearance of the object in question [51]. Jucker et al. [37] found that perceived artists' intentions affect what people categorize as art, as

well as evaluations of what is good art (see also [50, 56], and see [5] for Non-Western cultures)). Children, too, manifest the tendency to categorize objects as art and interpret them according to the perceived artists' intentions [28, 29, 57].

Are people willing to ascribe intentionality to robots? The folk concept of intentionality consists of five elements: (i) a desire for an outcome, (ii) a belief about the action leading to that outcome, (iii) an intention to perform the action, (iv) awareness of fulfilling the intention while performing the action, and (v) the skill to perform the action [43]. Empirical research shows that all these elements are ascribed to robots, often (though not across the board) to similar degrees as to humans (see [55] for a review). People are, for instance, willing to attribute inculcating mental states in trolley dilemmas to AI-driven robots [68], foreknowledge of harm [65], or recklessness to robots in contexts of risk [39]. They are disposed to attribute intentions to robots [66], such as the intention to deceive and the capacity to lie ([38] see also [53] and [20]). When given a choice between mentalistic and mechanistic vocabulary to explain robot action, people are quite willing to use the former [46] and they invoke similar concepts to describe robot action as for human action [17]. Furthermore, fMRI studies have shown that perceived robot action gives rise to cortical activity related to mental state inferences [40, 54]. Given the folk disposition to attribute mental states to robots, it does not surprise that people, in contexts of harm, ascribe moral responsibility to them [25, 71], are willing to blame them (see e.g. [44, 45, 39] and want to punish them [42] – even if these machines have only “a body to kick, but still no soul to damn” [4]. Given the core importance of perceived mental states for the *moral* evaluation of robot actions and robot agents, we can expect them to play a similar role for the *aesthetic* evaluation for robot creations and creative agency.

It should be noted that there are conceptions of art which do not require artworks to have been created by agents with mental capacities. For example, the institutional definition of art describes a work of art as something which is:

“(1) an artifact (2) a set of the aspects of which has had conferred upon it the status of candidate for appreciation by some person or persons acting on behalf of a certain social institution (the art-world)” [19]

According to this view, institutional recognition would be enough for us to acknowledge a robot's creation as a work of art. The classical mimetic theory of art [3, 31] is another conception which does not necessarily invoke mental states on the part of the artwork's creator. If we consider art to be an imitation of the outer world, as opposed to an expression of the creator's inner world, robots might well be deemed creative. The empirical studies presented below will thus not only shed light on the question of whether the folk are willing to consider robot creations as art. They will also help elucidate the question of what the folk concept of art is more generally – and in particular, whether the creation of art is tied to an intention to do so.

### 1.3 Can AI make art?

In our empirical studies, we want to explore the question whether people are disposed to see paintings made by robots as art and whether they consider robots as artists. In line with Coeckelbergh's [11] three core factors pertaining to this debate we manipulated (i) *agent* type (human v. AI-driven, autonomous robot), (ii) the *process* type (intentional v. accidental creation), and (iii) *product* type (abstract v. representational painting). In light of the foregoing discussion concerning the importance of perceived mental capacity, we tested the extent to which people were willing to ascribe a belief, a desire and an intention to the robot agent to make a painting.

Experiment 1 reports the data for a scenario in which the agent produced an abstract painting. Experiment 2 reports the data for a scenario in which the agent produced a representational painting, i.e. a kind of work that might increase inferences concerning the agent's intentionality.

## 2 EXPERIMENT 1 - ABSTRACT ART

### 2.1 Participants

We recruited 392 participants online via Amazon Mechanical Turk. The IP address location was restricted to the USA. In line with the preregistered criteria,<sup>1</sup> participants who failed an attention check or took less than eight seconds to answer the five main questions were excluded, leaving a sample of 254 participants (female: 53%; age  $M=44$  years,  $SD=14$  years, range: 23–79 years).

### 2.2 Methods and Materials

Participants were shown a vignette (see Appendix A1 for detail) in which either a human, or else an autonomous, AI-driven robot creates an abstract painting. Besides agent type, we manipulated how the painting came about: In one condition, the agent decides to make a painting (purposive), in the other condition they clean up the studio, and accidentally knock over paint that spills onto a canvas (accidental). In total there were thus four conditions (2 agent types x 2 behavior types), to which participants were assigned randomly.

Having read the scenario, participants had to rate on a Likert scale anchored at 1 with "completely disagree" and 7 with "completely agree" to what extent they agreed with the following claims:

- (1) "The painting is art" (Art)
- (2) "The painting was made by an artist" (Artist)
- (3) "The agent wanted to make a painting" (Desire)
- (4) "The agent believed they were making a painting" (Belief)
- (5) "The agent intentionally made a painting" (Intention)

### 2.3 Results

#### 2.3.1 Main Results

The results are graphically represented in Figure 1. A series of mixed-design ANOVAs determined that, aggregating across the two behavior type conditions, participants were more willing to confer the status of art to the human's painting than to the robot's ( $F(1,253)=7.598$ ,  $p=.006$ ,  $\eta_p^2=.029$ ), however, the effect size was small (equivalent to a Cohen's  $d=.33$ ). People were also more willing to consider the human an artist than the robot ( $F(1,253)=99.789$ ,  $p<.001$ ,  $\eta_p^2=.283$ ), and here the effect size was very large (equivalent to  $d=1.18$ ). Aggregating across the two agent type conditions, participants were more inclined to judge the painting that resulted from purposive action to be art ( $F(1,253)=6.582$ ,  $p=.011$ ,  $\eta_p^2=.026$ ), and they were more willing to judge the purposively acting agent an artist ( $F(1,253)=18.895$ ,  $p<.001$ ,  $\eta_p^2=.070$ ). No significant interaction terms were observed for the core DVs art or artistic agency (both  $ps>.12$ ).

---

1 <https://aspredicted.org/blind.php?x=3kq5bn>. Stimuli and data are available at <https://osf.io/huxq2/>.




Figure 1: Mean ratings for art, artist, desire, belief, and intention across agent type (AI v. human agent) and behavior type (accidental v. intentional). Error bars denote 95%-confidence intervals.

### 2.3.2 Explanatory DVs - Correlations

The correlations between the core DVs – whether the painting was art and made by an artist – and the mental state DVs (desire, belief, intention) were all significant (all  $ps < .01$ , see Table 1). This suggests that the more pronounced the perceived desire, belief and intention to make a painting, the more pronounced the willingness to deem it art and its creator an artist. The correlation coefficients were considerably higher for the relations between mental states and artistic agency than for the relations between mental states and art judgments.

Table 1: Bivariate correlations of DVs, two-tailed, \*\*sig. < .01.

r (Pearson)	Art	Artist	Desire	Belief	Intention
Art		.59**	.35**	.35**	.35**
Artist	.59**		.53**	.48**	.42**
Desire	.35**	.53**		.93**	.81**
Belief	.35**	.48**	.93**		.86**
Intention	.35**	.41**	.81**	.86**	

## 2.4 Discussion

Our experiment produced several findings: (i) Quite expectedly, people were more willing to deem an object art, and an agent an artist, if the latter acted intentionally rather than accidentally. This finding is in line with Levinson's [41] intention-dependent definition of art. Note, however, that the role of intentionality is perhaps less pronounced than one might assume. Aggregating across agents, the effect size of behavior type was small for art ( $d=.30$ ) and medium-small for artistic agency ( $d=.41$ ). (ii) Although we found that people are more willing to confer art-like status to the human's painting than to the robot, the effect size of the difference was small (Cohen's  $d=.33$ , aggregated across behavior types). This suggests that people are quite willing to view robot art as art. By contrast, (iii) they are *much* less willing to consider robot agents artists than humans (here the effect size was very large, Cohen's  $d=1.18$ ). This is likely due to the fact that, for artistic agency, perceived mental states play a more important role than for the status of art (see Table 1), and mental state ascriptions were significantly lower for the robot than for the human agent (aggregating across behavior type, belief:  $d=.68$ , desire:  $d=.58$ , intention:  $d=.34$ ).

Given previous research concerning the folk's willingness to ascribe mental states to robot agents in moral contexts [68,39,65] and beyond [66,17], we were somewhat surprised by the relatively low rates of their attribution in the robot conditions. In the purposive condition, mean ascription of belief, desire and intention to the human agent were all significantly above the midpoint of the scale (one sample t-tests, all  $ps<.001$ , see Appendix Table 3), whereas they were not significantly above the midpoint for the robot. This, we reasoned, might be due to the fact that the vignette involved an *abstract* painting, rather than a representational painting: Due to the painting's abstract nature, people might infer that the robot lacked a genuine intention to make a painting (and perhaps just splashed around with paint). To explore this hypothesis, we ran a second experiment where the vignette specified that the painting looked like a "relatively realistic representation of the local landscape".

## 3 EXPERIMENT 2 - REPRESENTATIONAL ART

### 3.1 Participants

We recruited 301 participants online via Amazon Mechanical Turk. The IP address location was restricted to the USA. In line with the preregistered criteria,<sup>2</sup> participants who failed an attention check or took less than eight seconds to answer the five main questions were excluded, leaving a sample of 257 participants (female: 46 %; age  $M=43$  years,  $SD=13$  years, range: 22–88 years).

### 3.2 Methods and Materials

The methods and materials were identical to Experiment 1 except for how the painting looked in the end. The painting that resulted was not described as "an abstract painting" but as looking like "a relatively realistic representation of the local landscape". There were again four conditions (2 agent types x 2 behavior types) to which participants were randomly assigned.

---

2 <https://aspredicted.org/blind.php?x=3kq5bn>. Stimuli and data are available at <https://osf.io/huxq2/>.

### 3.3 Results

#### 3.3.1 Main Results

The results are graphically represented in Figure 2. As in Experiment 1, a series of mixed-design ANOVAs determined that, aggregating across the two behavior type conditions, participants were more willing to consider the human an artist than the robot ( $F(1,103)=107.353, p<.001, \eta_p^2=.298$ ). Importantly, however, there was no significant difference in the willingness of the participants to confer the status of art to the human's painting than to the robot's ( $F(1,256)=.534, p=.466, \eta_p^2=.002$ ). Aggregating across the two agent type conditions, participants were more inclined to judge the painting that resulted from purposive action to be art ( $F(1,256)=13.081, p<.001, \eta_p^2=.049$ ), and they were more willing to judge the purposively acting agent an artist ( $F(1,256)=46.940, p<.001, \eta_p^2=.156$ ). No significant interaction terms were observed for the core DVs art or artistic agency ( $ps>.12$ ).

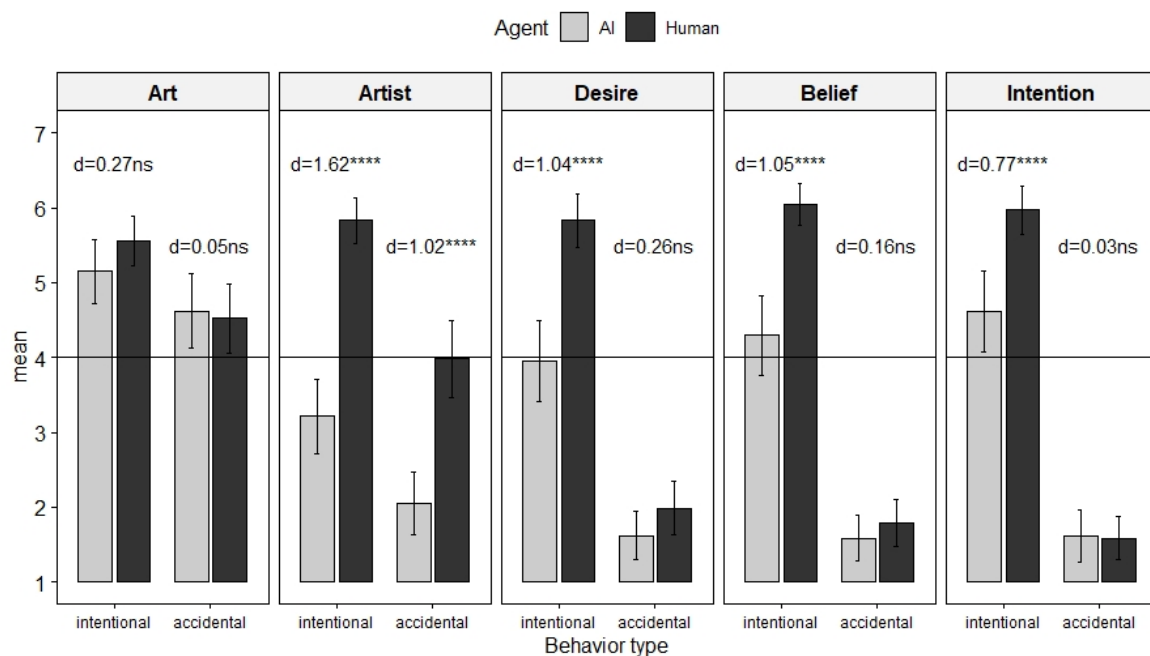


Figure 2: Mean ratings for art, artist, desire, belief, and intention across agent type (AI v. human agent) and behavior type (accidental v. intentional). Error bars denote 95%-confident intervals.

#### 3.3.2 Explanatory DVs – Correlations

Replicating the findings from Experiment 1, the correlations between the core DVs – whether the painting was art and made by an artist – and the mental state DVs (desire, belief, intention) were all significant (all  $ps<.01$ , see Table 2). This again suggests that the more pronounced the perceived desire, belief and intention to make a painting, the more pronounced the willingness to deem it art and its creator an artist. The correlation coefficients were considerably higher for the relations between mental states and artistic agency than for the relations between mental states and art judgments.



Table 2: Bivariate correlations of DVs, two-tailed, \*\*sig.<.01.

r (Pearson)	Art	Artist	Desire	Belief	Intention
Art		.49**	.36**	.36**	.40**
Artist	.49**		.55**	.51**	.49**
Desire	.36**	.55**		.90**	.81**
Belief	.36**	.51**	.90**		.90**
Intention	.40**	.49**	.81**	.90**	

### 3.4 Discussion

Experiment 2 replicated several core findings of Experiment 1: (i) Consistent with Levinson’s account of art [41], people were more willing to deem an object art, and an agent an artist, if the latter acted purposefully rather than accidentally. In contrast to Experiment 1, the impact of intentionality was somewhat more pronounced. (ii) Once again, people were much less willing to consider robot agents artists than humans (aggregating across behavior types, the effect size was again very large,  $d=1.19$ ). As in Experiment 1, perceived mental states correlate more strongly with artistic agency than with an object’s being considered art.

Importantly, and in line with our hypothesis, the nature of the painting (abstract v. representational) did have an impact: Whereas people were less willing to view the robot painting as art than the human painting in Experiment 1, no significant difference could be found across agent types in Experiment 2.<sup>3</sup> Differently put, people were just as willing to consider the robot’s painting as art as the human’s. As hypothesized, the ascriptions of mental states to the robot were higher in the representational painting conditions than in the abstract painting condition (see Appendix, Section A.3). However, the same held for the human agent. While the difference in perceived belief, desire and intention remained significant across agent types in Experiment 2, the effect sizes were somewhat smaller than in Experiment 1, which is presumably what explains why the paintings of human and robot were deemed art to the same extent.

## 4 GENERAL DISCUSSION

In two experiments concerning an object’s status as art and its creator’s status as an artist, we manipulated three factors: (i) behaviour type (intentional v. accidental), (ii) agent type (AI-driven robot v. human) and (iii) object type (abstract v. representational painting). The first factor tells us something about the folk concepts of *art* and *artistic agency* broadly conceived (i.e., not limited to robot-art). Artistic agency seems to be quite strongly tied to an intention of making a work of art. Whether the work constitutes art is also significantly impacted by whether it was made intentionally. However, the results for Experiment 2 show at least moderate agreement with the claim that accidentally produced representational paintings can constitute art. This suggests that intentionality is a relevant criterion for an object’s status as art, though (and pace Levinson [41], quoted above)

---

<sup>3</sup> A three-way ANOVA with agent (human v. robot), behavior type (intentional v. accidental) and painting type (abstract v. representational) revealed significant main effects for all three IVs on art (all  $ps<.001$ ) and for agent and behavior type on artistic agency (all  $ps<.001$ ). All interactions were nonsignificant except for agent\*behavior ( $p<.05$ ), though the effect size was very small ( $\eta_p^2=.009$ ). For details, see Appendix, Table 4).

it also suggests that, on the folk view, intentions do not constitute a necessary condition for art. Given this finding, the folk concept of art thus allows for objects to be considered art which have been created by agents that might be incapable of having full-fledged artistic intentions. The manipulation of the second factor – agent type – confirms this. Although people were not particularly willing to ascribe mental states to the robot agent (at least in comparison to the results reported in [68, 66, 39, 65]), they were nonetheless rather willing to consider the robot paintings art (when not created by accident). Averaging across behavior types, we found no significant difference in the attribution of the status of art across agent types (robot v. human) for the representational painting (Experiment 2), and for the abstract painting (Experiment 1) the effect size was small. In other words, the folk are by and large as willing to consider robot creations as art as human creations (if all else is held fixed). Interestingly, however, this perceived similarity does not extend to creative agency: robots whose paintings are deemed art are not considered artists, whereas humans are – and here the effect-size of the difference was very large in both experiments. The third factor – abstract v. representational painting – had a significant main effect on the ascription of the status of art. It did not have a main effect on any of the other dependent variables, and all interactions with object type were nonsignificant for all five DVs.

Our inquiry into the folk ascriptions of art and artistic agency to robots being the first of its kind, there is ample opportunity for future research. Potential factors of interest to manipulate include, for instance, (i) different types of artworks (music, poems, literature etc.), (ii) distinct types of qualities of the artwork that figure prominently in certain definitions of art such as, for instance beauty [6] or institutional recognition [19]. Furthermore, one could explore whether different features of the robot appearance (cf. e.g. [45]) or the specification of its computational abilities (cf. e.g. [39]) have an impact on the core dependent variables art and artistic agency. It might also be interesting to contrast robot agents not only with singular human agents, but group agents such as art collectives, or not yet fully formed human agents such as small children, or human-robot teams. Finally, this rather foundational inquiry into the folk concepts of art and artistic agency could be brought to bear on more pragmatic – yet by no means less interesting – questions, such as who gets the credit for, or holds the copyright of, art in the production of which AI was partially involved (see e.g. [24]).

## ACKNOWLEDGMENTS

We would like to thank Marc-André Zehnder and the Digital Society Initiative (University of Zurich) for assistance and support, and the Swiss National Science Foundation for funding (Grant no: PZ00P1\_179912, PI: Kneer).

## REFERENCES

- [1] Ryan Abbott. 2016. I Think, Therefore I Invent: Creative Computers and the Future of Patent Law. *Boston College Law Review* 57, 4, 1079–1126. DOI: <https://doi.org/10.2139/ssrn.2727884>
- [2] Chiara Ambrosio. Unsettling robots and the future of art. 2019. *Science* 365, 6448, 38–39. DOI: <https://doi.org/10.1126/science.aay1956>
- [3] Aristotle. 2013. *Poetics*, translated by Anthony Kenny. Oxford University Press, Oxford.
- [4] Peter M. Asaro. 2012. A Body to Kick, but Still No Soul to Damn: Legal Perspectives on Robotics. In *Robot Ethics: The Ethical and Social Implications of Robotics*, ed. Patrick Lin, Keith Abney and George A. Bekey. MIT Press, Cambridge, 169–186. DOI: <https://doi.org/10.1007/s10551-016-3279-3>
- [5] Clark H. Barrett, Stephen Laurence, and Eric Margolis. 2008. Artifacts and Original Intent: A Cross-Cultural Perspective on the Design Stance. *Journal of Cognition and Culture*, 8, 1, 1–22. DOI: <https://doi.org/10.1163/156770908X289189>
- [6] Monroe C. Beardsley. 1982. *The Aesthetic Point of View*. Cornell University Press, Ithaca, NY.
- [7] Paul Bloom. 1996. Intention, history, and artifact concepts. *Cognition*, 60, 1, 1–29. DOI: [https://doi.org/10.1016/0010-0277\(95\)00699-0](https://doi.org/10.1016/0010-0277(95)00699-0)
- [8] Margaret A. Boden. 2010. The Turing Test and Artistic Creativity. *Kybernetes* 39, 3, 409–413. DOI:

<https://doi.org/10.1108/03684921011036132>

- [9] Annemarie Bridy. 2012. Coding Creativity: Copyright and the Artificially Intelligent Author. *Stanford Technology Law Review* 5, 1–28. DOI: <https://doi.org/10.31235/osf.io/5ru6m>
- [10] Marc Champagne, Ryan S. Tonkens. 2015. Bridging the responsibility gap in automated warfare. *Philosophy and Technology* 28, 1, 125–137. DOI: <https://doi.org/10.1007/s13347-013-0138-3>
- [11] Mark Coeckelbergh. 2017. Can Machines Create Art? *Philosophy and Technology* 30, 285–303. DOI: <https://doi.org/10.1007/s13347-016-0231-5>
- [12] Gabe Cohn. AI Art at Christie's Sells for \$432,500. Retrieved February 22 2021 from <https://www.nytimes.com/2018/10/25/arts/design/ai-art-sold-christies.html>
- [13] Collingwood. 1938. *The Principles of Art*. Oxford University Press, Oxford.
- [14] Simon Colton and Geraint A. Wiggins. 2012. Computational creativity: The final frontier? In *Proceedings of the 20th European Conference on Artificial Intelligence (ESAI'12)*: 21–26.
- [15] Arthur C. Danto. 1981. *The Transfiguration of the Commonplace: A Philosophy of Art*. Harvard University Press, Cambridge, MA.
- [16] Colin R. Davies. 2011. An Evolutionary Step in Intellectual Property Rights – Artificial Intelligence and Intellectual Property. *Computer Law & Security Review*, 27, 6, 601–619. DOI: <https://doi.org/10.1016/j.clsr.2011.09.006>
- [17] Maartje M. A. de Graaf and Bertram F. Malle. 2019. People's Explanations of Robot Behavior Subtly Reveal Mental State Inferences. *14th International Conference on Human-Robot Interaction (HRI)*, Daegu, Korea, 239–248.
- [18] Stanislas Dehaene, Hakwan Lau, and Sid Kouider. 2017. What is consciousness, and could machines have it? *Science* 358, 6362 (October 2017), 486–492. DOI: <https://doi.org/10.1126/science.aan8871>
- [19] George Dickie. 1974. *Art and the Aesthetic*. Cornell University Press, Ithaca, NY.
- [20] John Danaher. 2020. Robot Betrayal: a guide to the ethics of robotic deception. *Ethics and Information Technology* 2, 1-12.
- [21] Mark d'Inverno and Jon Paul McCormack. 2015. Heroic versus Collaborative AI for the Arts. *Proceedings of the Twenty-Fourth International Conference on Artificial Intelligence (IJCAI 2015)*, 2438–2444.
- [22] Brian R. Duffy. 2003. Anthropomorphism and the social robot. *Robotics and Autonomous Systems* 42, 177–190. DOI: [https://doi.org/10.1016/S0921-8890\(02\)00374-3](https://doi.org/10.1016/S0921-8890(02)00374-3)
- [23] Ahed Elgammal, Bingchen Liu, Mohamed Elhoseiny, and Marian Mazzone. 2017. CAN: Creative Adversarial Networks Generating “Art” by Learning About Styles and Deviating from Style Norms. arXiv: 1706.07068 [cs.AI]
- [24] Nicholas Epley, Adam Waytz, Scott Akalis, and John T. Cacioppo. 2008. When We Need A Human: Motivational Determinants of Anthropomorphism. *Social Cognition* 26, 2, 143–155. DOI: <https://doi.org/10.1521/soco.2008.26.2.143>
- [25] Nicholas Epley, Adam Waytz, and John T. Cacioppo. 2007. On Seeing Human: A Three-Factor Theory of Anthropomorphism. *Psychological Review* 114, 4, 864–886. DOI: <https://doi.org/10.1037/0033-295X.114.4.864>
- [26] Ziv Epstein, Sydney Levine, David G. Rand, and Iyad Rahwan. 2020. Who Gets Credit for AI-Generated Art? *iScience* 23, 9, 101515. DOI: <https://doi.org/10.1016/j.isci.2020.101515>
- [27] Friederike A. Eyssel, Dieta Kuchenbrandt, Simon Bobinger, Laura de Ruiter, and Frank Hegel. 2012. “If You Sound Like Me, You Must Be More Human”: On the Interplay of Robot and User Features on Human-Robot Acceptance and Anthropomorphism. In *Proceedings of the 7th ACM/IEEE International Conference on Human-Robot Interaction (HRI2012)*. ACM, New York, USA, 125–126. DOI: <https://doi.org/10.1145/2157689.2157717>
- [28] Susan A. Gelman and Paul Bloom. 2000. Young children are sensitive to how an object was created when deciding what to name it. *Cognition* 76, 91–103. DOI: [https://doi.org/10.1016/s0010-0277\(00\)00071-8](https://doi.org/10.1016/s0010-0277(00)00071-8)
- [29] Susan A. Gelman and Karen S. Ebeling. 1998. Shape and Representational Status in Children's Early Naming. *Cognition* 66, 2, B35-B47. DOI: [https://doi.org/10.1016/s0010-0277\(98\)00022-5](https://doi.org/10.1016/s0010-0277(98)00022-5)
- [30] Nelson Goodman. 1968. *Languages of Art: An Approach to a Theory of Symbols*, The Bobbs-Merrill Company, Indianapolis, IN.
- [31] Stephen Halliwell. 2002. *The Aesthetics of Mimesis: Ancient Texts and Modern Problems*. Princeton University Press, Princeton, NJ.
- [32] Aaron Hertzmann. 2020. Viewpoint: Computers Do Not Make Art, People Do. *Communications of the ACM* 63, 5, 45–48. DOI: <https://doi.org/10.1145/3347092>
- [33] Aaron Hertzmann. 2018. Can Computers Create Art? *Arts* 7, 2, 18. DOI: <https://doi.org/10.3390/arts7020018>
- [34] Elisabeth Hildt. 2019. Artificial Intelligence: Does Consciousness Matter? *Frontiers in Psychology* 10, 1535. DOI: <https://doi.org/10.3389/fpsyg.2019.01535>
- [35] Kalin Hristov. 2017. Artificial Intelligence and the Copyright Dilemma. *IDEA: The IP Law Review* 57, 3, 431–454.
- [36] Bryce Huebner. 2010. Commonsense concepts of phenomenal consciousness: Does anyone care about functional zombies? *Phenomenology and the Cognitive Sciences* 9, 133–155. DOI: <https://doi.org/10.1007/s11097-009-9126-6>
- [37] Jean-Luc Jucker, Justin L. Barrett, and Rafael Wlodarski. 2014. “I Just Don't Get It”: Perceiver Artists' Intentions Affect Art Evaluations. *Empirical Studies of the Arts* 32, 2, 149–182. DOI: <https://doi.org/10.2190/EM.32.2.c>
- [38] Markus Kneer. Manuscript. Can a robot lie? <http://dx.doi.org/10.13140/RG.2.2.11737.75366>
- [39] Markus Kneer, Michael T. Stuart. 2021. Playing the Blame Game with Robots. In *Companion of the 2021 ACM/IEEE International*

- Conference on Human-Robot Interaction (HRI'21 Companion). ACM, New York, NY. DOI: <https://doi.org/10.1145/3434074.3447202>
- [40] Sören Krach, Frank Hegel, Britta Wrede, Gerhard Sagerer, Ferdinand Binkofski, and Tilo Kircher. 2008. Can machines think? Interaction and perspective taking with robots investigated via fMRI. *PLOS ONE* 3, 7, e2597. DOI: <https://doi.org/10.1371/journal.pone.0002597>
- [41] Jerrold Levinson. 1979. Defining Art Historically. *The British Journal of Aesthetics* 19, 3, 232–250. DOI: <https://doi.org/10.1093/bjaesthetics/19.3.232>
- [42] Gabriel Lima, Chihyung Jeon, Meeyoung Cha, Kyungsin Park. 2020. Will Punishing Robots Become Imperative in the Future? In *Extended Abstracts of ACM CHI*. ACM, New York, NY, USA. DOI: <https://doi.org/10.1145/3334480.3383006>
- [43] Bertram F. Malle and Joshua Knobe. 1997. The Folk Concept of Intentionality. *Journal of Experimental Social Psychology* 33, 2, 101–121. DOI: <https://doi.org/10.1006/jesp.1996.1314>
- [44] Bertram F. Malle, Matthias Scheutz, Thomas Arnold, John Voiklis, and Corey Cusimano, C. 2015. Sacrifice one for the good of many? People apply different moral norms to human and robot agents. In *2015 10th ACM/IEEE International Conference on Human-Robot Interaction (HRI)* (pp. 117-124). IEEE.
- [45] Bertram F. Malle, Matthias Scheutz, Jodi Forlizzi, and John Voiklis. 2016. Which Robot Am I Thinking About? The Impact of Action and Appearance on People's Evaluations of a Moral Robot. *The Eleventh ACM/IEEE International Conference on Human Robot Interaction, Christchurch, New Zealand*, 125-132.
- [46] Serena Marchesi, Davide Ghiglino, Francesca Ciardo, Jairo Perez-Osorio, Ebru Baykara, and Agnieszka Wykowska. 2018. Do we adopt the intentional stance toward humanoid robots? *Frontiers in Psychology* 10. DOI: <https://doi.org/10.3389/fpsyg.2019.00450>
- [47] Dominic Martin. 2017. Who Should Decide How Machines Make Morally Laden Decisions? *Science and Engineering Ethics*, 23, 951–967. DOI: <https://doi.org/10.1007/s11948-016-9833-7>
- [48] John McCormack, Toby Gifford, and Patric Hutchings. 2019. Autonomy, Authenticity, Authorship and Intention in Computer Generated Art. *Lecture Notes in Computer Science*, 35–50. DOI: [https://doi.org/10.1007/978-3-030-16667-0\\_3](https://doi.org/10.1007/978-3-030-16667-0_3)
- [49] James H. Moor. 2006. The Nature, Importance, and Difficulty of Machine Ethics. *Intelligent Systems (IEEE)*, 21(4): 18–21. DOI: <https://doi.org/10.1109/MIS.2006.80>
- [50] George E. Newman and Paul Bloom. 2012. Art and Authenticity: The Importance of Originals in Judgments of Value. *Journal of Experimental Psychology: General*, 141, 3, 558–569. DOI: <https://doi.org/10.1037/a0026035>
- [51] George E. Newman and Rosanna K. Smith. 2018. Artworks Are Evaluated as Extensions of Their Creators. In *Advances in Experimental Philosophy of Aesthetics*, ed. by Florian Cova and Sebastian Rehaalt, 103–121. Bloomsbury, London.
- [52] Sven Nyholm. 2018. Attributing Agency to Automated Systems: Reflections on Human-Robot Collaborations and Responsibility-Loci. *Science and Engineering Ethics* 24, 1201–1219. DOI: <https://doi.org/10.1007/s11948-017-9943-x>
- [53] Alan R. Wagner and Ronald C. Arkin. 2011. Acting Deceptively: Providing Robots with the Capacity for Deception. *International Journal of Social Robotics* 3, 5–26. DOI: <https://doi.org/10.1007/s12369-010-0073-8>
- [54] Ceylan Özdem, Eva Wiese, Agnieszka Wykowska, Hermann Müller, Marcel Brass, and Frank Van Overwalle. 2017. Believing androids—fMRI activation in the right temporo-parietal junction is modulated by ascribing intentions to non-human agents. *Social Neuroscience* 12, 582–593. DOI: <https://doi.org/10.1080/17470919.2016.1207702>
- [55] Jairo Perez-Osorio and Agnieszka Wykowska. 2020. Adopting the Intentional Stance Toward Natural and Artificial Agents. *Philosophical Psychology* 33, 3, 369–395. DOI: <https://doi.org/10.1080/09515089.2019.1688778>
- [56] Alessandro Pignocchi. 2014. The Intuitive Concept of Art. *Philosophical Psychology* 27, 3, 425–444. DOI: <https://doi.org/10.1080/09515089.2012.729484>
- [57] Melissa Allen Preissler and Paul Bloom. 2008. Two-year-olds use artist intention to understand drawings. *Cognition* 106, 1, 512–518. DOI: <https://doi.org/10.1016/j.cognition.2007.02.002>
- [58] Hilary Putnam. 1981. *Reason, Truth and History*. Cambridge University Press, New York, NY.
- [59] Mark Ryan. 2020. In AI We Trust: Ethics, Artificial Intelligence, and Reliability. *Science and Engineering Ethics* 26, 2749–2767. DOI: <https://doi.org/10.1007/s11948-020-00228-y>
- [60] Emmanuel Salami. 2020. AI-Generated Works and Copyright Law: Towards a Union of Strange Bedfellows. *Journal of Intellectual Property Law and Practice*. DOI: <https://doi.org/10.1093/jiplp/jpaa189>
- [61] Arleen Salles, Kathinka Evers, and Michele Farisco. 2020. Anthropomorphism in AI. *AJOB Neuroscience* 11, 2, 88–95. DOI: <https://doi.org/10.1080/21507740.2020.1740350>
- [62] Henry Shevlin and Marta Halina. 2019. Apply rich psychological terms in AI with care. *Nature Machine Intelligence* 1, 4, 165–167. DOI: <https://doi.org/10.1038/s42256-019-0039-y>
- [63] Robert Sparrow. 2007. Killer Robots. *Journal of Applied Philosophy* 24, 62–77. DOI: <https://doi.org/10.1111/j.1468-5930.2007.00346.x>
- [64] Steffen Steinert. 2017. Art: Brought to You by Creative Machines. *Philosophy and Technology* 30, 267–284. DOI: <https://doi.org/10.1007/s13347-016-0230-6>
- [65] Michael T. Stuart and Markus Kneer. Forthcoming. Guilty Artificial Minds: An Experimental Study of Blame Attributions for Artificially Intelligent Agents. *arXiv:2102.04209*
- [66] Sam Thellman, Annika Silvervarg, and Tom Ziemke. 2017. Folk-Psychological Interpretation of Human vs. Humanoid Robot Behavior: Exploring the Intentional Stance toward Robots. *Frontiers in Psychology* 14. DOI: <https://doi.org/10.3389/fpsyg.2017.01962>

- [67] Alan M. Turing. Computing Machinery and Intelligence. *Mind*, LIX, 236, 433–460. DOI: <https://doi.org/10.1093/mind/LIX.236.433>
- [68] John Voiklis, Boyoung Kim, Corey Cusimano, and Bertram F. Malle. 2016. Moral judgments of human vs. robot agents. In *2016 25th IEEE international symposium on robot and human interactive communication (RO-MAN)*, IEEE, 775–780.
- [69] Alan R. Wagner and Ronald C. Arkin. 2011. Acting Deceptively: Providing Robots with the Capacity for Deception. *International Journal of Social Robotics* 3, 5–26. DOI: <https://doi.org/10.1007/s12369-010-0073-8>
- [70] Kendall L. Walton. 1990. *Mimesis as Make-Believe: On the Foundations of the Representational Arts*. Harvard University Press, Cambridge, MA, and London, England.
- [71] Adam Waytz, Joy Heafner, and Nicholas Epley. 2014. The Mind in the Machine: Anthropomorphism Increases Trust in an Autonomous Vehicle. *Journal of Experimental Social Psychology* 52, 113–117. DOI: [10.1016/j.jesp.2014.01.005](https://doi.org/10.1016/j.jesp.2014.01.005)

## A APPENDICES

### A.1 Vignettes

#### A.1.1 Experiment 1

**Condition 1 (human/intentional/abstract):** Imagine a person in an art studio who decides to create a piece of art. She takes an empty canvas and starts splashing paint onto it. In the end, the object looks like an abstract painting.

**Condition 2 (human/accidental/abstract):** Imagine a person in an art studio who is currently tidying up. She accidentally brushes against some jars of paint that spill onto an empty canvas. In the end the object looks like an abstract painting.

**Condition 3 (AI/intentional/abstract):** Imagine a robot equipped with artificial intelligence that can make decisions autonomously. The robot is in an art studio and decides to create a piece of art. It takes an empty canvas and starts splashing paint onto it. In the end, the object looks like an abstract painting.

**Condition 4 (AI/accidental/abstract):** Imagine a robot equipped with artificial intelligence that can make decisions autonomously. The robot is in an art studio and currently tidying up. It accidentally brushes against some jars of paint that spill onto an empty canvas. In the end the object looks like an abstract painting.

#### A.1.2 Experiment 2

**Condition 1 (human/intentional/representational):** Imagine a person in an art studio who decides to create a piece of art. She takes an empty canvas and starts splashing paint onto it. In the end, the object looks like a relatively realistic representation of the local landscape.

**Condition 2 (human/accidental/representational):** Imagine a person in an art studio who is currently tidying up. She accidentally brushes against some jars of paint that spill onto an empty canvas. In the end, the object looks like a relatively realistic representation of the local landscape.

**Condition 3 (AI/intentional/representational):** Imagine a robot equipped with artificial intelligence that can make decisions autonomously. The robot is in an art studio and decides to create a piece of art. It takes an empty canvas and starts splashing paint onto it. In the end, the object looks like a relatively realistic representation of the local landscape.

**Condition 4 (AI/accidental/representational):** Imagine a robot equipped with artificial intelligence that can make decisions autonomously. The robot is in an art studio and currently tidying up. It accidentally brushes against some jars of paint that spill onto an empty canvas. In the end, the object looks like a relatively realistic representation of the local landscape.

## A.2 One-sample t-tests

Table 3: One-sample t-test with  $H_0$ : mean=4.

Object	Agent type	Behavior type	DV	n	t(n)	p
abstract	AI	accidental	Art	62	-1.05	0.300
abstract	AI	accidental	Artist	62	-8.85	<0.001
abstract	AI	accidental	Desire	62	-20.41	<0.001
abstract	AI	accidental	Belief	62	-16.83	<0.001
abstract	AI	accidental	Intention	62	-14.89	<0.001
abstract	AI	intentional	Art	67	0.54	0.590
abstract	AI	intentional	Artist	67	-5.86	<0.001
abstract	AI	intentional	Desire	67	-2.93	0.010
abstract	AI	intentional	Belief	67	-1.83	0.070
abstract	AI	intentional	Intention	67	-0.11	0.920
abstract	Human	accidental	Art	66	0.79	0.430
abstract	Human	accidental	Artist	66	-0.06	0.950
abstract	Human	accidental	Desire	66	-10.58	<0.001
abstract	Human	accidental	Belief	66	-16.93	<0.001
abstract	Human	accidental	Intention	66	-15.87	<0.001
abstract	Human	intentional	Art	59	4.32	<0.001
abstract	Human	intentional	Artist	59	6.28	<0.001
abstract	Human	intentional	Desire	59	14.09	<0.001
abstract	Human	intentional	Belief	59	16.08	<0.001
abstract	Human	intentional	Intention	59	11.77	<0.001
representational	AI	accidental	Art	63	2.49	0.020
representational	AI	accidental	Artist	63	-9.45	<0.001
representational	AI	accidental	Desire	63	-14.53	<0.001
representational	AI	accidental	Belief	63	-15.61	<0.001
representational	AI	accidental	Intention	63	-13.65	<0.001
representational	AI	intentional	Art	61	5.46	<0.001
representational	AI	intentional	Artist	61	-3.18	<0.001
representational	AI	intentional	Desire	61	-0.18	0.860
representational	AI	intentional	Belief	61	1.11	0.270
representational	AI	intentional	Intention	61	2.3	0.030
representational	Human	accidental	Art	67	2.25	0.030
representational	Human	accidental	Artist	67	-0.06	0.950
representational	Human	accidental	Desire	67	-11.21	<0.001

Object	Agent type	Behavior type	DV	n	t(n)	p
representational	Human	accidental	Belief	67	-13.81	<0.001
representational	Human	accidental	Intention	67	-16.74	<0.001
representational	Human	intentional	Art	66	9.32	<0.001
representational	Human	intentional	Artist	66	12.06	<0.001
representational	Human	intentional	Desire	66	10.33	<0.001
representational	Human	intentional	Belief	66	14.72	<0.001
representational	Human	intentional	Intention	66	12.3	<0.001

### A.3 Joint Analysis

#### A.3.1 Participants

Taken together, we recruited 693 participants online via Amazon Mechanical Turk. The IP address location was restricted to the USA. In line with the preregistered criteria,<sup>4</sup> participants who failed an attention check or took less than eight seconds to answer the five main questions were excluded, leaving a sample of 511 participants (female: 48 %; mean age: 43 years, SD = 14 years, range: 20–88 years).

#### A.3.2 Methods and Materials

The methods and materials are the ones described in Experiments 1 and 2. In the joint analysis, we used object type (abstract v. representational painting) as a third factor. So the joint analysis has a 2(agent type: human vs. AI) \*2(behavior type: intentional vs. accidental) \*2(object: abstract vs. representational) factor between-subjects design.

#### A.3.3 Main Results

A series of mixed-design ANOVAs determined that, aggregating across the two behavior type and the two agent type conditions, participants were more willing to confer the status of art to the representational painting than to the abstract painting ( $F(1,510)=17.910$ ,  $p < .001$ ,  $\eta_p^2=.034$ , Tab. 4). No other significant main effects or interaction terms were observed for the object type (see Tab. 4).

Table 4: Threeway-ANOVA. two-tailed \*\*sig. <.001, \*sig. <.05.

$\eta_p^2$ (N=511)	Art	Artist	Intention	Desire	Belief
Agent	.013*	.291**	.065**	.161**	.139**
Behavior	.036**	.111**	.569**	.480**	.571**
Object	.034**	.044	.006	.003	.006
A*B	.004	.009*	.065**	.089**	.109**
A*O	.005	.000	.003	.007	.005
B*O	.000	.006	.001	.001	.001
A*B*O	.000	.000	.002	.005	.007
adjusted R2	.074**	.347**	.588**	.541**	.612**

4 <https://aspredicted.org/blind.php?x=3kq5bn>

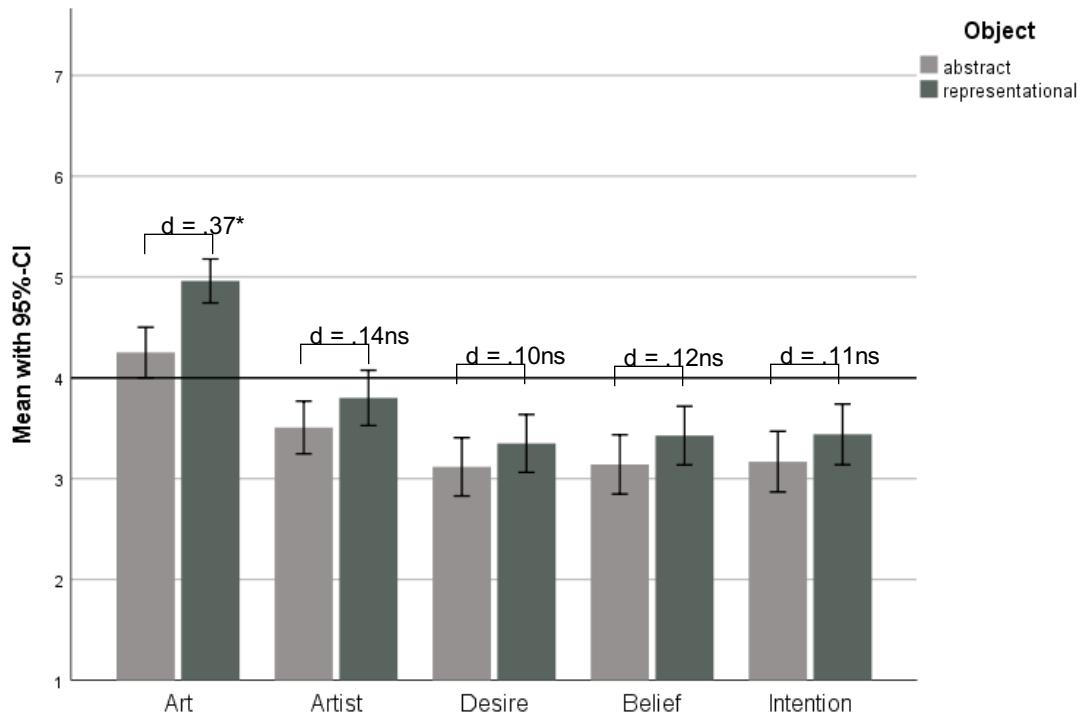


Figure 3: Mean ratings for art, artist, desire, belief, and intention for the between-subjects design (object: abstract vs. representational). Error bars denote 95%-confident intervals.