Title

How People Infer a Humanlike Mind from a Robot Body

Authors

Xuan Zhao,^{1,2*} Elizabeth Phillips,³ Bertram F. Malle¹

Affiliations

¹Department of Cognitive, Linguistic, and Psychological Science, Brown University. 190 Thayer Street, Providence, RI 02912. Email: xuan_zhao@xbrown.edu.

²Center for Decision Research, University of Chicago Booth School of Business. 5807 S. Woodlawn Street, Chicago, IL 60637.

³Department of Psychology, George Mason University. 4400 University Dr, Fairfax, VA 22030.

Abstract

Robots are entering a wide range of society's private and public settings, often with a strikingly humanlike appearance and emulating a humanlike mind. When do laypeople attribute a humanlike mind to these machines? In this paper, we reveal a systematic relationship between variations in robots' humanlike appearance and the mental capacities people infer. In a programmatic, bottom-up investigation across 251 real-world robots, we first decompose humanlike appearance into three distinct appearance dimensions—body-manipulators, face, and surface. Then, we show that experimentally manipulating specific appearance constellations leads people to make specific inferences about a robot's mental capacities. In particular, while the presence of body-manipulators shapes people's inferences about a robot's perception-action capacities, facial and surface appearance shapes people's inferences about a robot's affective and moral capacities. This research provides theoretical and methodological clarity on the powerful impact that a robot's body has on human perceptions of its mind.

Introduction

Society is entering a time that could only be imagined in science fiction decades ago: From supermarket aisles to children's playrooms, robots are reaching a wide range of society's private and public settings and are expected to interact with people on an increasingly frequent basis (1–4). These artificial agents capture significant public attention, partly because they often resemble the human species in two prominent ways: the body and the mind. Bounded not by biological rules or governmental regulations but only by a designer's imagination, the physical appearance of a robot can vary greatly—from the mere possession of an eye-like structure to a full-blown humanoid body with hair and skin (5, 6). And even though scholars have yet to fully understand the nature and scope of “machine intelligence” (7, 8), decades of psychological literature have shown that laypeople readily attribute humanlike mental states such as intention, thoughts, and emotions to nonhuman agents—a phenomenon often referred to as “anthropomorphism” (9–13). Attributing mental capacities to robots is often consequential because it guides people's beliefs on what kind of jobs robots are suitable for (14), raises questions on what moral rights and obligations we should grant these machines (15, 16), and may even threaten humans' sense of uniqueness as individuals and as a species (17, 18).

Perhaps most strikingly, people need not observe direct evidence of mental activities in a robot to form their opinions on the presence of a humanlike mind. To date, a handful of studies have explored how robots' mere appearance can influence how much “mind” people subsequently attribute. The predominant view so far suggests that the more physically human-looking a robot is, the more humanlike people will infer its mind to be (19–22). However, a simple and straightforward “more humanlike appearance ⇒ more humanlike mind” hypothesis faces at least two major challenges: It rests on a limited understanding of what constitutes both a humanlike body and a humanlike mind.

First, given the considerable variability in robots' physical appearance, understanding what makes a robot look humanlike and how to experimentally manipulate it for systematic investigation has been challenging. Most of the current literature rests on the assumption that physical humanlikeness is a unitary construct that spans the spectrum of “non-humanlike” (or “machine-like”) to highly “humanlike” (19–26). As a result, prior research tends to adopt one of two approaches to manipulate how “humanlike” a robot's physical appearance is: to designate a few robots with vastly different appearance as examples of mechanical robots, humanoid robots, or android robots and compare people's perceptions across them (20, 25, 26); or to single out and manipulate the humanlikeness of specific features in particular robot models and observe people's subsequent reactions (19, 22, 27). However, given the vast and ever-expanding landscape of humanlike robots, it is unclear how we can derive generalizable insights from a few robot models or idiosyncratic manipulations of appearance features.

Second, people form a wide variety of inferences about a robot's mind, so systematically capturing such variations with a comprehensive, validated measurement of mental capacities is critical. When such tools were not available, some studies singled out specific mind attributes relevant to particular contexts or tasks, whereas others focused on overall trait and personality perception, such as to what extent a robot seems “intelligent”, “capable”, or “smart”, or more generally, whether the robot seems “human” or “has a mind” (21, 27–30). Seminal research in the mind perception literature recently revealed that human mind perception is actually multi-dimensional. While an initial framework

identified two dimensions (10), other work using comprehensive collections of mental capacity items later suggested a three-dimensional (13) or even a five-dimensional structure (31). Yet regardless of the precise numbers, all these multi-dimensional models were developed by contrasting different agent types such as adult humans, God, dogs, or frogs, with “robots” being an undifferentiated type. It remains unclear what kind of multi-dimensional structure of mind perception applies to people’s inferences of distinct robots, such as those differing in humanlike appearance.

In light of these challenges, we took a systematic and multi-dimensional approach to investigate how people infer a robot’s mind from its physical appearance. Specifically, using the largest image collection of real-world robots to date and a bottom-up, data-driven approach, we first identified multi-dimensional structures for both humanlike appearance and robot mind inference. Then, we revealed how specific dimensions of humanlike appearance lead people to make inferences about specific aspects of a robot’s mind.

Study 1: Decomposing “Humanlike” Appearance in Modern Robots

What constitutes humanlike appearance in robots is, to date, both theoretically and methodologically underspecified. To address this problem, we took a feature-based approach to systematically document and analyze the appearance of real-world humanlike robots. As the first step, we created a comprehensive image collection of such robots. Our systematic search process, detailed in the Materials and Methods section, yielded an image collection of 251 extant robots built for research or commercial purposes. All images are standardized and archived as part of the Anthropomorphic roBOT (ABOT) Database—a publicly available online resource we have created for the study of robots’ humanlike appearance (<http://www.abotdatabase.info>).

With this image collection, we sought to understand what kind of high-level structure exists among this huge variety of humanlike robots. Based on prior literature on robot appearance (6) and a pilot study with independent coders, we identified 16 unique humanlike appearance features that are visible in static images of robots and represent a wide range of common features in contemporary robots, including: apparel, arms, fingers, eyes, eyebrows, eyelashes, face, genderedness (i.e., being perceived to have a gender), hair, hands, head, legs, mouth, nose, skin, and torso. To measure laypeople’s common judgment about the presence of these features in robots, in Study 1, we recruited 1,284 human coders and asked them to make binary judgments on the presence (coded as “1”) or absence (coded as “0”) of one specific humanlike feature (see Materials and Methods in the Supplemental Materials). To ensure a consistent criterion across judges on each feature, we provided participants with definitions that were adapted from the Oxford English and Merriam Webster’s dictionaries (see Table S1). Participants’ responses showed a high level of inter-judge agreement (see Table S2), and we calculated a “feature-presence score” for each feature and each robot to represent the proportion of participants who rated the specific feature as present in that robot. These feature-presence scores ranged from 0 (i.e., no judges saw a specific appearance feature in a given robot) to 1 (i.e., all judges saw a specific appearance feature in a given robot).

Next, we examined whether the 16 features could be reduced to a few high-level dimensions. A Principal Component Analysis (PCA) on the 251 robots × 16 feature-

presence score matrix identified a clear dimensional structure, with three distinct and interpretable principal components accounting for 73.9% of the total variance of all features (see loadings in Table 1). The first component, which we labeled *Surface*, accounted for 29.4% of the total variance and included seven high-loading features common to android robots: eyelashes, head hair, skin, genderedness, apparel, nose, and eyebrows (listed from highest to lowest factor loading). Intriguingly, although these features fulfill, in the human body, concrete biological functions related to physical protection, reproduction, and respiration, their static depiction on a robot only serve to approximate the human form without performing any similar functions.

The second component, which we labeled *Body-Manipulators*, accounted for 26.0% of the total variance and included five high-loading features defining the robot's trunk and limbs: hands, arms, torso, fingers, and legs. Different from surface features, body-manipulator features often fulfill similar functions in humans and robots, serving the essential faculty of navigating and manipulating the physical environment.

The last component, which we termed *Face*, accounted for 18.4% of the total variance and included four features: face, eyes, head, and mouth. Working together, they often serve communicative functions in a robot, such as indicating head/face orientation and gaze direction, expressing facial emotions, and visualizing mouth movement during speech, yet they rarely serve biological functions such as seeing or eating for robots.

The discovery of three dimensions of humanlike appearance (see Fig. 1A) thus provides a methodological tool to systemically and parsimoniously delineate in what way any robot physically resembles a human body. Specifically, by averaging the feature-presence scores for the highest-loading features ($> .50$) on each principal component, we computed three *dimension scores* for each robot that quantify how strongly a given robot embodies the defining features of each appearance dimension—with a score close to 1 suggesting that all defining features of that dimension are salient in that robot, and a score close to 0 suggesting that none of the features are salient (see Fig. S1 in the Supplemental Materials for an example). Finally, we confirmed that these dimension scores have strong internal consistency ($\alpha = .89$ for surface, $.92$ for body-manipulators, and $.84$ for face) and are highly representative of the regression-based, standardized principal component scores (surface: $r(249) = .95$; body-manipulators: $r(249) = .99$; face: $r(249) = .92$, p 's $< .001$).

Study 2: Validating a Three-Dimensional Structure of Humanlike Appearance

Seeing that a clear three-dimensional structure emerged from a dizzying array of 16 humanlike features among 251 robots, we next examined whether this three-dimension structure can effectively capture the vague and seemingly unitary construct of “humanlike appearance”. To this end, we conducted Study 2 to uncover how these three appearance dimensions can give rise to the overall impression of a robot's physical humanlikeness.

We recruited a new group of participants ($n = 147$) and asked them to report how physically humanlike the appearance of each robot seemed to them, using a single-item measure ranging from 0 (“*not humanlike at all*”) to 100 (“*just like a human*”) (see Materials and Methods). These ratings showed a high level of inter-judge agreement (see Table S3), and we averaged participants' ratings to yield an overall physical humanlikeness score for each robot. As expected, all three appearance dimensions

strongly correlated with people's overall rating of a robot's physical humanlikeness, Surface: $r(249) = .72$; Body: $r(249) = .66$; Face: $r(249) = .54$, $ps < .001$. A multiple linear regression model further showed that, together, the three appearance dimensions accounted for 80.2% of the total variance of robots' physical humanlikeness ($R = .90$, $F(3,247) = 333$, $p < .001$). Importantly, each of the dimension scores made a unique contribution to this impressive predictive power: the Body-Manipulators dimension uniquely explained 25.7% of the total variance ($b = 28.65$; $t = 17.91$, $p < .001$); the Surface dimension uniquely explained 21.6% of the total variance ($b = 51.34$; $t = 16.40$, $p < .001$); the Face dimension uniquely explained 0.5% of the total variance ($b = 6.85$; $t = 2.62$, $p = .009$); and 32.3% was shared explained variance, reflecting the naturally occurring correlations among the three appearance dimensions (see Fig. S2 in the Supplemental Materials).

At a theoretical level, the strong predictive power of this multiple linear regression suggests that the broad concept of "humanlikeness appearance" is not unidimensional but can be further decomposed into three constituent appearance dimensions: Body-Manipulators, Surface, and Face. This finding provides a methodological foundation for quantifying robots' humanlike appearance along three ecologically valid and psychologically meaningful dimensions (see Fig. 1B). As we showcase below, this finding enables systematic research into the psychological impact of robots' humanlike appearance on people's mind inferences.

Study 3: From Humanlike Appearance to Robot Minds

After revealing that robots' humanlike appearance is multi-dimensional, we experimentally manipulated appearance to examine its systematic impacts on people's mind inferences.

To survey a wide range of appearances, we identified subsets of robots in the ABOT Database that were located at the high and low ends of the three appearance dimensions, which amounted to eight possible combinations in the appearance space, yet two cells contained no extant robots (see Fig. 1C). From each of the remaining six cells, we then chose four state-of-the-art robot models that represented typical robots that cell but also maintained variability of appearance within and across cells, resulting in 24 robots populating the entire humanlikeness spectrum (see Fig. S3 for all 24 robots; see Fig. S4 for results of a hierarchical cluster analysis on 251 robots, which confirmed the diversity of our selected robot subsets). And in order to capture as many potential mental capacities as possible, we adopted a 20-item instrument that was developed through an iterative process of item analysis from 54 mental capacities (31) and consists of five reliable subscales: positive affect (e.g., feeling happy, loving specific people), negative affect (e.g., feeling pain, experiencing fear), moral cognition (e.g., telling right from wrong, praising moral actions), social cognition (e.g., understanding others' minds, setting goals), and reality interaction (e.g., moving on its own, learning from instruction) (see Table 2). In a PCA provided in Table S5, we confirmed a three-dimensional structure of this instrument, where the subscales of positive affect and negative affect hang together, and the subscales of social cognition and moral cognition hang together. For the current analysis, we retained the five subscales to allow for maximum differentiation on the predictive relationships between appearance and mind inference.

Participants in Study 3 ($n = 510$) each judged one robot on 20 mental capacities using 8-point rating scales from 0 (“*not at all capable*”) to 7 (“*completely capable*”). To understand how variations of humanlike appearance gave rise to specific inferences of a robot’s mental capacities, we first inspected to what extent the 20 mental capacity items were correlated with the 16 humanlike appearance features across 24 robots. As shown in Fig. 2, a heatmap of correlation coefficients highlighted two areas of strong relationships: In the upper-right region, features from the Surface and Face dimensions consistently predicted capacities of affect and moral cognition; in the lower-left region, features of the Body-Manipulators dimension consistently predicted reality interaction capacities.

We next assessed more systematically whether, and to what extent, robots’ three physical appearance dimensions influenced people’s mind inferences on the five mental capacity subscales. We computed each robot’s five mental capacity subscale scores by averaging the ratings of participants who judged that robot, and with these subscale scores as the outcome variables, we constructed multiple linear regression models to examine how well the three appearance dimension scores (established in Study 1) can predict people’s mind inferences of a robot. As indicated by the total R^2 of each model in Table 3, a robot’s standing on the three appearance dimensions, taken together, strongly affected people’s perceptions of the robot’s capacities of positive affect, negative affect, moral cognition, and reality interaction, explaining a remarkable 40% to 60% of the variance in such inferences. By contrast, variations in appearance had only a negligible impact on inferences of social cognition capacities, suggesting that people may have derived these inferences from other sources. Turning to the distinct impact of each appearance dimension (see unique and shared variance components in Table 3 and Fig 3.), we found that only the face dimension made significant *unique contributions* to inferences of positive affect, negative affect, and moral cognition, whereas only the body-manipulators dimension made a significant unique contribution to inferences of reality interaction. The surface dimension, by contrast, made negligible unique contributions to robot mind inferences but substantial *shared contributions*—mostly with the face dimension—to inferences of negative affect, positive affect, and moral cognition. These results corroborated the systematic and reliable patterns of appearance-to-mind relationship revealed by the item-level analysis in Fig 2.

Discussion

Leveraging the increasing prevalence of robots in the modern society, the current research presents a novel perspective on how people infer the presence of a mind from the sight of a humanlike body. Our results show that such inferences can systematically stem from specific dimensions of a robot’s humanlike appearance. In particular, people infer a robot’s capacities to perceive, learn from, and interact with the environment in response to its body-manipulators appearance (e.g., arms, legs, fingers); they infer affective capacities such as being able to feel love, gratitude, stress, or pain based in response to a robot’s facial features (e.g., eyes, mouth) and surface appearance (e.g., skin, hair, apparel); and they even infer a robot’s moral competence from its facial and surface features. These results reject the oversimplified assumption of “more humanlike appearance $\Rightarrow$ more humanlike mind” by contributing a more accurate and nuanced understanding of the relationship between human resemblance and mind inference, and

they also affirm that robots’ “humanlikeness”—in both the body and the mind—encompasses multiple dimensions.

This intricate relationship between appearance and mind inference is unlikely to be unique to robots. Because most participants had no prior exposure to the robots they judged, inferences about robot mind are more likely to derive from experiences with familiar agents such as humans and animals. Consequently, our findings also promise novel insights into how people may systematically infer humans’ mental capacities from their appearance. Prior research, by considering the human body as a unitary construct, has suggested that body focus leads to overall objectification (32, 33) or to decreased agency attribution yet increased attribution of emotional capacities (34). However, in light of the distinct dimensions of humanlike appearance and their dissociable impact on mind inference, we speculate that highlighting different dimensions of the human body may create differential psychological impact on human observers. For instance, the famous “We Can Do It!” poster—featuring a woman with a strong arm (i.e., emphasizing the body-manipulator dimension)—may enhance people’s beliefs of her capacity to interact with and change her environment. By contrast, the movie poster of “The Silence of the Lambs”—with an obscure face contour and occluded mouth (i.e., reducing the face dimension)—may lead to specific inferences of reduced affective and moral capacities. Similar uncharitable inferences may also arise when a human’s appearance is altered by accessories or garments such as dark sunglasses, veils, or masks. Future research should explore how the saliency of each appearance dimension—in both robots and humans—systematically affect the kind of mind people infer in an agent.

In this paper, we have focused on how static images of robots influence people’s perception of a humanlike mind. Such mind inferences are not only ecologically valid—as people encounter robots in news articles, on commercial websites, or even on a display shelf or in public deployment—but likely consequential due to the powerful influence of first impressions (35). However, dynamic features that can be observed only in videos or live human-robot interactions, such as a robot’s movement (36, 37), voice (38), or task performance (39, 40), will also critically affect inferences of robot minds. We hope that our work will inspire researchers to investigate such dynamic features by employing a similar bottom-up approach to reveal how static and dynamic features interact over time.

The decision to design robots with more or less humanlike features comes with practical and economic considerations but also with social and ethical significance. An anthropomorphic form can create a range of favorable impressions of social robots, such as a sense of familiarity/affinity (23), trustworthiness (24), likeability (36), and as demonstrated here, the impression of a multi-faceted humanlike mind. However, an anthropomorphic form can also lead to exaggerated expectations, misplaced trust, and unhealthy attachments, followed perhaps by shock and disappointment when those exaggerated expectations are laid bare (40). Robot design is therefore at a crossroads. With the increasing presence of robots in private and public domains, knowledge of the psychological impact of humanlike appearance must lead to more informed decisions about what kind of robots—with what kind of appearance—we should introduce into our society.

ABOT Database: The Image Collection of 251 Robots. In order to uncover a systematic relationship across the broad landscape of anthropomorphic robots, we set out to gather as many robots as possible that varied in both number and type of humanlike appearance features. We obtained candidate images from a variety of sources between January 2017 and October 2018, including academic outlets (e.g., journal publications, conference proceedings), technology-focused media (e.g., online magazines, newsletters, blogs), websites of robotics companies and university laboratories, online robotics communities and discussion forums, robotics community listservs, direct correspondence with developers, and general internet searches. From this search and solicitation process, we generated an initial collection of 329 real-world robots developed for research or commercial purposes, each with one or more plausible humanlike features. Next, we employed a three-step exclusion procedure: We first excluded models that appeared clearly zoomorphic; then, to avoid redundancies, we excluded images of robots that were highly physically similar due to multiple versions or models of the same robot; finally, we reviewed the robot images to ensure that each robot had at least one photograph with no visual obstruction or motion blur, depicted the robot's whole body (whenever possible), and without other robots or humans in the image—a robot was excluded only when no such images could be obtained on the Internet or by contacting the developers directly. For the remaining 251 robots, we identified the best photograph of each robot by taking three factors into account: reasonable image resolution, a standing, neutral, forward-facing posture (when available), and a neutral or mildly positive facial expression (when applicable). For every chosen image, we then removed cluttered or colored backgrounds to standardize the presentation of all robots against a white background.

Selecting and Defining 16 Common Humanlike Appearance Features. To identify appearance features common across our robots, we developed an initial coding scheme of 29 possible robot appearance features based on previous coding schemes (36). After pilot studies, we excluded 11 of the 29 features because they were either rare (e.g., wings, antennas), often invisible or unidentifiable in static images (e.g., toes, ears, eyelids), or too equivocal to be classified (e.g., digital vs. physical faces, neck, wrist, grippers). We also collected data on the presence of wheels and treads/caterpillar tracks because prior research suggests that locomotion is an important source of variation in robot design (36). However, given that the current research focus is on humanlike appearance, we excluded these two non-humanlike features (e.g., wheels, treads/tracks) in the analyses reported here.

Recruitment, Participants, and Exclusions. We recruited participants (with IP addresses within the United States) via the crowdsourcing platform Amazon Mechanical Turk and paid either \$1.00 (Studies 1 and 2) or \$0.35 (Study 3) in exchange for their participation. Across all three studies, 1,871 participants met our predetermined data inclusion criteria (Study 1: $n = 1,214$; $M_{age} = 35.66$, $SD_{age} = 11.74$, 53% female, 46% male, 1% unreported; Study 2: $n = 147$; $M_{age} = 33.75$, $SD_{age} = 9.73$, 55% male, 45% female; Study 3: $n = 510$; $M_{age} = 38.07$, $SD_{age} = 12.18$, 54% female, 46% male). We set systematic criteria to exclude participants whose responses (i) contained multiple incorrect answers to catch trials, (ii) showed little to no variance across all robots, and/or

(iii) substantially deviated from the group average. Full details on target sample sizes, exclusion criteria, and rating reliability are reported in *supplemental materials*.

Study 1.

Experimental Design. The first 200 robots were randomly assigned into one of four blocks, consisting of 50 robots each. For the later 51 robots, their images were entered as a single block (which contained 54 robots—the additional three included one new robot to replace a zoomorphic robot in the first 200 robots, and two robots from the first 200 that were accidentally omitted from data collection due to a technical error). Participants were either randomly assigned to one of these four blocks (August 2017) or received the fifth block (January 2019) and judged one of the 16 appearance features (e.g., eyes). Thus, each participant judged whether one humanlike feature was present for each of the 50 (or 54) robots in their assigned block. In addition, each participant also judged whether that same feature was present in 16 “catch trials”: eight images of humans (varying in age, gender, and ethnicity) and eight images of featureless smart home devices (e.g., the Amazon Echo, the Google Home). These additional trials were used to establish clear anchors and to detect possibly confused or careless raters (see *supplemental materials, Methods*).

Procedure. Participants learned that the goal of the study was to investigate how people perceive the appearance of robots and other entities. After receiving the definition of one of the features (*supplemental materials, Table S1*), they judged whether the entity in each image possessed that feature by clicking “yes” or “no”. The order of images was fully randomized, and participants made judgments without time constraint. Halfway through the task, they were required to take a 10-second break and were reminded of the feature definition. At the end of the study, participants completed a brief demographic questionnaire (i.e., level of knowledge and experience with robots, age, gender, education, language) and were prompted to provide additional feedback to the researcher. The entire study took approximately 5-7 minutes to complete.

Statistical Analysis. We first calculated inter-rater reliability using a variant of Cohen’s kappa (κ) supported by Monte Carlo simulations (see *Supplemental Materials*). We then conducted a Principal Component Analysis (PCA) with Varimax rotation on the 251 robots $\times$ 16 feature-presence score matrix and computed reliabilities (Cronbach’s α) for each dimension. We next calculated Pearson’s correlation coefficients between the standardized regression-based composite scores and the averaged dimension scores of each appearance dimensions. Finally, for the hierarchical cluster analysis reported in the *Supplemental Materials*, we used Ward’s minimum variance method with Euclidean distance as the dissimilarity measure and visualized our results in a dendrogram that represents to what extent two robots are similar.

Study 2.

Experimental Design. Similar to Study 1, 251 robots were divided into a total of 5 blocks, and participants were randomly assigned to either one of the first four blocks or the fifth block.

Procedure. At the beginning of the study, participants were informed that we were interested in “how people perceive robots and other entities that look like humans” and learned that they would “judge several photographs on how much each entity in the photograph looks like a human in its physical appearance.” To help them calibrate the full range of humanlike appearance, we provided seven images of entities juxtaposed in a row, ranging from a feature-less loudspeaker to a real human male. Then, participants were given as much time as needed to rate images of robots in their respective block (plus the 16 catch trials identical to those in Study 1) on the question “*Does this look physically humanlike?*”. They indicated their response by moving a slider on a continuum between “*not humanlike at all*” (left-most point of the slider’s scale) to “*just like a human*” (right-most point of the slider’s scale). At the end of the study, participants were asked to fill out a brief demographic questionnaire identical to that in Study 1. Once complete, participants were provided with a payment code, thanked, and debriefed. The entire study took approximately 7 minutes to complete.

Statistical Analysis. We first calculated inter-rater reliability using two intraclass coefficients: ICC(2,1) estimates what the expected correlation would be between any two randomly selected members of the population and ICC(2,*k*) estimates what the expected correlation would be between any two randomly selected groups of our sample size (see Supplemental Materials). We conducted Pearson’s correlations between the overall physical humanlikeness score and each appearance dimension score. We then conducted a multiple linear regression with the overall physical humanlikeness score as the outcome variable and the three appearance dimension scores as predictors. In the Supplemental Materials, we also report a multiple linear regression with the overall physical humanlikeness score as the outcome variable and the 16 individual appearance features as predictors.

Study 3.

Experimental Design. We selected robots on both the high-end (dimension score $\geq .65$) and the low-end (dimension score $\leq .35$) of each appearance dimension while retaining a wide range of humanlikeness scores (range: 1.44–93.44; $M = 46.54$; $SD = 28.75$).

Note that the present studies provided a conservative estimate of the relationship between physical appearance and mind inferences for two reasons. First, the appearance scores and the mind inference scores came from independent groups of participants who were unaware of the other group’s judgments. Therefore, the strong appearance-to-mind relationship we identified is not the result of within-participant covariance of ratings (which may inflate correlations). Second, we manipulated robots’ appearance by showing only a single robot to each participant, guaranteeing that they would not make their mind inferences as a function of obvious appearance differences across multiple presented robots.

Procedure. At the beginning of the study, participants were reminded of the prevalence of robots and the variation in their appearance and saw an image panel of eight robots from the ABOT collection that were not in the target sample of 24 robots. Then, they were told that we were interested in “how people perceive robots”, and their task was to indicate—following their best guess—to what degree they thought that each

robot had various mental capacities. Each participant was then randomly assigned to view one robot (image on top of each page) and responded to the question “*To what degree do you think this robot is capable of...*” on 8-point rating scales from 0 (*not at all capable*) to 7 (*completely capable*) regarding 20 mental capacity items (presented on 4 pages, with item order randomized within each page). After all mental capacity judgments, participants reported what entity they thought the robot seemed to be modeled after (i.e., human infant, human child, human adolescent, human adult, animal, other), to what extent they perceived the robot to be elegantly designed, and whether the robot presented a high-quality finish. Finally, participants reported whether they had seen and/or interacted with the robot they judged prior to the study. Once completed, participants filled out a demographic questionnaire similar to the one used in Study 1. They were then provided with a payment code, thanked, and debriefed. The entire study took approximately 2 to 3 minutes to complete.

Statistical Analysis. Similar to previous studies, we started with calculating inter-rater reliability for the purpose of data exclusion (see Supplemental Materials). Then we conducted Pearson correlations between 16 appearance features and 20 mental capacity items. Next, we conducted univariate and multiple linear regression with each of the five mental capacity subscale scores as the outcome variable and appearance dimension scores as the predictors. We also conducted a Principal Component Analysis (PCA) with Varimax rotation on a 20 mental capacity items $\times$ 510 participants matrix. In the Supplemental Materials, we reported multiple ANOVA (analysis of variance) tests comparing mental capacity subscale scores across six robot categories, and Pearson correlations between three appearance dimensions and five mental capacity subscales.

References

1. T. Belpaeme, J. Kennedy, A. Ramachandran, B. Scassellati, F. Tanaka, Social robots for education: A review. *Sci. Robot.* **3**, eaat5954 (2018).
2. E. Broadbent, Interactions With Robots: The Truths We Reveal About Ourselves. *Annu. Rev. Psychol.* **68**, 627–652 (2017).
3. J.-F. Bonnefon, A. Shariff, I. Rahwan, The social dilemma of autonomous vehicles. *Science* (6293). **352**, 1573–1576 (2016).
4. I. Rahwan *et al.*, Machine behaviour. *Nature*. **568**, 477–486 (2019).
5. E. Phillips, X. Zhao, D. Ullman, B. F. Malle, What is Human-like?: Decomposing Robots’ Human-like Appearance Using the Anthropomorphic roBOT (ABOT) Database. *HRI ’18 Proc. 2018 ACM/IEEE Int. Conf. Human-Robot Interact.*, 105–113 (2018).
6. E. Phillips, D. Ullman, M. M. A. De Graaf, B. F. Malle, What does a robot look like?: A multi-site examination of user expectations about robot appearance. *Proc. Hum. Factors Ergon. Soc.*, 1215–1219 (2017).
7. A. M. Turing, Computing machinery and intelligence. *Mind*. **59**, 433–460 (1950).
8. D. C. Dennett, *Content and consciousness* (Routledge, 2010).
9. N. Epley, A. Waytz, J. T. Cacioppo, On seeing human: A three-factor theory of anthropomorphism. *Psychol. Rev.* **114**, 864–886 (2007).

- 475 10. H. M. Gray, K. Gray, D. M. Wegner, Dimensions of mind perception. *Science* (80-.).
476 **315**, 619 (2007).
- 477 11. J. Piaget, *The child's conception of the world* (Routledge and Kegan Paul, Ltd., London,
478 UK, 1929).
- 479 12. J. L. Jipson, S. A. Gelman, Robots and rodents: Children's inferences about living and
480 nonliving kinds. *Child Dev.* **78**, 1675–1688 (2007).
- 481 13. K. Weisman, C. S. Dweck, E. M. Markman, Rethinking people's conceptions of mental
482 life. *Proc. Natl. Acad. Sci.* **114**, 11374–11379 (2017).
- 483 14. A. Waytz, M. I. Norton, Botsourcing and outsourcing: Robot, british, chinese, and
484 german workers are for thinking-not feeling-jobs. *Emotion.* **14**, 434–444 (2014).
- 485 15. Zara Stone, Everything you need to know about Sophia, the world's first robot citizen.
486 *Forbes* (2017), pp. 1–8.
- 487 16. S. R. R. Nijssen, B. C. N. Müller, R. B. van Baaren, M. Paulus, Saving the robot or the
488 human? Robots who feel deserve moral care. *Soc. Cogn.* **37**, 41–56 (2019).
- 489 17. E. Leung, G. Paolacci, S. Puntoni, Man Versus Machine: Resisting Automation in
490 Identity-Based Consumer Behavior. *J. Mark. Res.* **55**, 818–831 (2018).
- 491 18. J. Złotowski, K. Yogeewaran, C. Bartneck, Can we control it? Autonomous robots
492 threaten human identity, uniqueness, safety, and resources. *Int. J. Hum. Comput. Stud.*
493 **100**, 48–54 (2017).
- 494 19. E. Broadbent *et al.*, Robots with display screens: A robot with a more humanlike face
495 display is perceived to have more mind and a better personality. *PLoS One.* **8**, 1–10
496 (2013).
- 497 20. S. Krach *et al.*, Can machines think? Interaction and perspective taking with robots
498 investigated via fMRI. *PLoS One.* **3**, e2597 (2008).
- 499 21. C. E. Looser, T. Wheatley, The tipping point of animacy: How, when, and where we
500 perceive life in a face. *Psychol. Sci.* **21**, 1854–1862 (2010).
- 501 22. M. C. Martini, C. A. Gonzalez, E. Wiese, Seeing minds in others - Can agents with
502 robotic appearance have human-like preferences? *PLoS One.* **11**, e0146310 (2016).
- 503 23. M. Mori, The uncanny valley. *Energy.* **7**, 33–35 (1970).
- 504 24. M. B. Mathur, D. B. Reichling, Navigating a social world with robot partners: A
505 quantitative cartography of the Uncanny Valley. *Cognition.* **146**, 22–32 (2016).
- 506 25. F. Ferrari, M. P. Paladino, J. Jetten, Blurring human-machine distinctions:
507 Anthropomorphic appearance in social robots as a threat to human distinctiveness. *Int. J.*
508 *Soc. Robot.* **8**, 287–302 (2016).
- 509 26. K. Gray, D. M. Wegner, Feeling robots and human zombies: Mind perception and the
510 uncanny valley. *Cognition.* **125**, 125–130 (2012).
- 511 27. A. Powers, S. Kiesler, The advisor robot: Tracing people's mental model from a robot's
512 physical attributes. *HRI '06 Proceeding 1st ACM SIGCHI/SIGART Conf. Human-robot*
513 *Interact.*, 218–225 (2006).
- 514 28. C. Bartneck, D. Kulić, E. Croft, S. Zoghbi, Measurement instruments for the

- anthropomorphism, animacy, likeability, perceived intelligence, and perceived safety of robots. *Int. J. Soc. Robot.* **1**, 71–81 (2009).
29. C. F. DiSalvo, F. Gemperle, J. Forlizzi, S. Kiesler, All robots are not created equal: the design and perception of humanoid robot heads. *DIS' 02 Conf. Des. Interact. Syst. Process. Pract. methods Tech.*, 321–326 (2002).
 30. A. Abubshait, E. Wiese, You look human, but act like a machine: Agent appearance and behavior modulate different aspects of human-robot interaction. *Front. Psychol.* **8**, 1–12 (2017).
 31. B. F. Malle, How many dimensions of mind perception really are there? *Proc. 41st Annu. Meet. Cogn. Sci. Soc.*, 2268–2274 (2019).
 32. N. A. Heflick, J. L. Goldenberg, Seeing eye to body: The literal objectification of women. *Curr. Dir. Psychol. Sci.* **23**, 225–229 (2014).
 33. S. Loughnan *et al.*, Objectification leads to depersonalization: The denial of mind and moral concern to objectified others. *Eur. J. Soc. Psychol.* **40**, 709–717 (2010).
 34. K. Gray, J. Knobe, M. Sheskin, P. Bloom, L. F. Barrett, More than a body: Mind perception and the nature of objectification. *J. Pers. Soc. Psychol.* **101**, 1207–1220 (2011).
 35. A. Todorov, *Face value: The irresistible influence of first impressions* (Princeton University Press, Princeton, NJ, 2017).
 36. Á. Castro-González, H. Admoni, B. Scassellati, Effects of form and motion on judgments of social robots' animacy, likability, trustworthiness and unpleasantness. *Int. J. Hum. Comput. Stud.* **90**, 27–38 (2016).
 37. C. K. Morewedge, J. Preston, D. M. Wegner, Timescale bias in the attribution of mind. *J. Pers. Soc. Psychol.* **93**, 1–11 (2007).
 38. J. Schroeder, N. Epley, Mistaking minds and machines: How speech affects dehumanization and anthropomorphism. *J. Exp. Psychol. Gen.* **145**, 1427–1437 (2016).
 39. M. De Graaf, S. Ben Allouch, J. Van Dijk, Why do they refuse to use my robot?: Reasons for non-use derived from a long-term home study. *HRI '17 Proc. 2017 ACM/IEEE Int. Conf. Human-Robot Interact.*, 224–233 (2017).
 40. B. G. Hoffman, Anki, Jibo, and Kuri: What we can learn from social robots that didn't make it. *IEEE Spectr.* (2019), pp. 1–4.

Acknowledgments

General: We thank Daniel Ullman, Salomi Aladia, McKenna Cisler, Fuei Yue, Broderick Allan, and Maya Menefee for their research assistance, Jeremiah Zhe Liu for feedback on statistical analyses, as well as all the researchers who responded to our email requests and contributed their robot images to our research.

Funding: This project was supported by a grant from the Office of Naval Research (ONR), No. N00014-16-1-2278. The opinions expressed here are our own and do not necessarily reflect the views of ONR.

Author contributions: All authors jointly developed the research idea and designed the study. E.P. and X.Z. curated the robot image collection. E.P. collected and B.F.M. cleaned the survey data on robot appearance. X.Z. and E.P. collected and cleaned the survey data on robot mind inference. All authors analyzed the data. X.Z. and B.F.M. visualized the data. X.Z. wrote the initial draft. All authors edited and revised the manuscript.

Competing interests: The authors declare no competing interests.

Data and materials availability: Experimental stimuli, surveys, and raw data are available in the Open Science Framework repository: https://osf.io/wksda/?view_only=f9c1c8b5b73f4ea18345880554324b60. In addition, the robot images and appearance profiles are archived in the ABOT Database (<http://www.abotdatabase.info>). The 20-item instrument to measure people's inferences of mental capacities is available at: <http://research.clps.brown.edu/SocCogSci/Measures/MI5.html>.

570

571

572

573

574

575

576

577

578

579

580

581

582

583

584

Fig 1. (A) Humanlike appearance constellations. Three appearance dimensions emerged from 16 appearance features with a bottom-up approach among 251 robots in the ABOT Database. **(B) Humanlike appearance space.** A total of 251 robots are depicted in the three-dimensional space of humanlike appearance based on their dimension scores of Body-Manipulators, Face, and Surface in Study 1. Their degree of physically humanlikeness, measured by the single-item measure in Study 2, is indicated by the filled color of each point, from low (blue) to high (red). **(C) Examples of representative robots.** Combining the low and high ends of each of the three appearance dimensions yields eight cells. Among them, two contain no robots, and the other six cells are each represented by one robot in this graph (for illustrative purposes).

Fig. 2. Correlation matrix between 16 humanlike appearance features (X-axis) and 20 mental capacity items (Y-axis). Colors indicates zero-order correlation coefficients, with lighter colors indicating higher correlation coefficients. Sixteen appearance features (BD1 = Arms, BD2 = Fingers, BD3 = Hands, BD4 = Torso, BD5 = Legs; FC1 = Eyes, FC2= Face, FC3 = Head, FC4 = Mouth; SF1 = Apparel, SF2 = Eyebrow, SF3 = Eyelashes, SF4 = Genderedness, SF5 = Head hair, SF6 = Nose, SF7 = Skin) are grouped by three appearance dimensions indicated in grey boxes. The mental capacity items come from a 20-item instrument of mind inference (31), and the items are grouped into the instrument’s five recommended subscales (see Table 2 for specific meaning of mental capacity item codes).

Fig. 3. Inferences of robots’ mental capacities predicted by three dimensions of the robots’ humanlike appearance. The squares represent the total variance across 24 robots on each mental capacity subscale. Within a square, the size of a circle represents the proportion of the total variance of mental capacity explained by that appearance dimension—either uniquely (nonoverlapping areas) or jointly with other appearance dimensions (overlapping areas).

605
606

Table 1. Principal components loading matrix of feature-presence scores in Study 1 (after varimax rotation). Loadings greater than or equal to 0.50 are in boldface type.

Feature	PC1: Surface	PC2: Body- Manipulators	PC3: Face
Eyelashes	0.88	-0.05	0.16
Skin	0.87	0.10	0.07
Headhair	0.87	0.07	0.03
Genderedness	0.80	0.30	0.22
Apparel	0.70	0.26	0.10
Nose	0.69	0.06	0.38
Eyebrows	0.68	-0.14	0.40
Hands	0.13	0.93	0.09
Arms	0.02	0.92	0.10
Torso	0.04	0.88	0.21
Fingers	0.13	0.86	0.03
Legs	0.05	0.70	-0.09
Face	0.26	0.13	0.91
Eyes	0.15	-0.01	0.87
Head	0.11	0.49	0.73
Mouth	0.49	-0.06	0.60
<i>% of variance explained</i>	29.4	26.0	18.4

607
608

Table 2. Twenty mental capacity items grouped into five subscales in Study 3. The first two columns list five subscales that can be further grouped into three overarching dimensions according to previous research and a PCA reported in the Supplemental Materials. The third and fourth columns show mental capacity items and their item codes corresponding to those in Fig 2. The final two columns show internal consistency measures for each subscale, calculated for robot-level and participant-level mental capacity ratings.

Mental Capacity Subscales	Item Code	Mental Capacity Item	Cronbach's α	
			24 robots	N = 510
Affect	Positive	P1 Feeling pleasure	.94	.86
		P2 Feeling happy		
		P3 Experiencing gratitude		
		P4 Loving specific people		
	Negative	N1 Experiencing fear	.94	.85
		N2 Feeling pain		
		N3 Feeling tired		
		N4 Feeling stress		
Social-Moral Cognition	Social Cognition	S1 Understanding others' minds	.86	.78
		S2 Inferring what a person is thinking		
		S3 Setting goals		
		S4 Planning for the future		
	Moral Cognition	M1 Telling right from wrong	.92	.88
		M2 Disapproving of immoral actions		
		M3 Praising moral actions		
		M4 Upholding moral values		
Reality Interaction		R1 Moving on their own	.78	.68
		R2 Seeing and hearing the world		
		R3 Communicating verbally		
		R4 Learning from instruction		

Table 3. Relationship between appearance dimensions and mental capacity subscales according to univariate and multivariate linear regressions. For multiple linear regression, each robot's average ratings on each mental capacity subscale was simultaneously regressed on three humanlike appearance dimensions, and the unstandardized regression coefficients (*b*), 95% CI, and *p* values for each predictor, as well as the *R*² and *p* value for each model are shown.

Predicted Mental Capacity	<i>Appearance Predictor</i>	Zero-order correlation <i>r</i>	<i>Multiple Regression</i>				Overall Model <i>R</i> ²
			Semi-partial correlation <i>r</i> _{<i>sp</i>}	Regression weights <i>b</i> [95% CI]	Unique + shared variance		
Positive Affect	<i>Body</i>	.10	-.01	-0.01 [-0.30, 0.28]	.00 + .01	.58***	
	<i>Face</i>	.75***	.48**	0.74 [0.31, 1.17]	.23 + .33		
	<i>Surface</i>	.59**	.15	0.22 [-0.20, 0.63]	.02 + .33		
Negative Affect	<i>Body</i>	.01	-.08	-0.05 [-0.27, 0.16]	.01 + -.01	.44**	
	<i>Face</i>	.63***	.37*	0.37 [0.04, 0.69]	.14 + .26		
	<i>Surface</i>	.54**	.18	0.17 [-0.14, 0.48]	.03 + .26		
Moral Cognition	<i>Body</i>	.24	.15	0.17 [-0.20, 0.54]	.02 + .04	.47**	
	<i>Face</i>	.64***	.37*	0.64 [0.09, 1.20]	.14 + .28		
	<i>Surface</i>	.55**	.18	0.29 [-0.24, 0.83]	.03 + .28		
Social Cognition	<i>Body</i>	.24	.19	0.21 [-0.23, 0.65]	.04 + .02	.15	
	<i>Face</i>	.32	.21	0.35 [-0.32, 1.02]	.05 + .06		
	<i>Surface</i>	.23	.03	0.04 [-0.60, 0.68]	.00 + .06		
Reality Interaction	<i>Body</i>	.66***	.63***	0.91 [0.44, 1.38]	.40 + .03	.45**	
	<i>Face</i>	.21	.06	0.13 [-0.58, 0.84]	.00 + .04		
	<i>Surface</i>	.21	.06	0.13 [-0.54, 0.81]	.00 + .04		

p* < .05, *p* < .01, ****p* < .001, two-tailed.