The Utility of Non-Anthropomorphic Robot Identities

Maria P. Stull
University of Colorado Boulder
Colorado, USA

Bradley Hayes
University of Colorado Boulder
Colorado, USA

ABSTRACT

Robots often perform humanlike identities, which are familiar to human collaborators and perceived as likeable. Identity performances that deviate from anthropomorphic standards can serve specific, useful purposes. Some of these purposes include setting expectations about the robot's abilities and sociability, altering perceptions of its agency, making robot groups more manageable, and using the flexible relationship between a robot's embodiment and identity to affect trust and understanding of the robot's actions.

ACM Reference Format:
Maria P. Stull and Bradley Hayes. 2024. The Utility of Non-Anthropomorphic Robot Identities. In *Proceedings of HRI '24 Workshop on Robo-Identity: Designing for Identity in the Shared World (Robo-Identity at HRI24).* ACM, New York, NY, USA, 4 pages. https://doi.org/XXXXXXX.XXXXXXX

1 INTRODUCTION

How best to design a robot's identity is an open question. Robotic identities are frequently modeled on human identities—using human names, humanlike speech, and humanoid forms—and are designed for the purpose of interacting with humans. Robotic identities, however, do not follow the rules of human identities.

Unlike humans, robots don't need to portray identity at all. For many robots, like a manufacturing robot repeatedly performing a single task, or a fully teleoperated drone, there is no need to be perceived as a social being or an individual with agency. But for social robots, presenting an identity may be a requirement for accomplishing their goals: facilitating interactions with humans, or distinguishing themselves from other robots.

If it portrays one at all, a robot's identity is designed and programmed, not developed through personal experiences, preferences, or social and cultural backgrounds. Robot identity doesn't come with an inner life like human identity: it's entirely performative. Some of the signals that can give a robot or artificial agent a sense of identity include names, behavior, speech, and impermanent physical signals (lights, screen displays, or removable accessories) [6]. For the purposes of this paper we are not considering permanent parts of a robot's physical form as part of its identity. Robots have a separation between *mind* and *body*: an artificial identity does not need to be tied to a single physical form, but could move into different bodies over time or inhabit multiple forms at the same time. The identity a robot is performing does not need to remain the same over time or in different situations, or evolve gradually: it can be completely transformed with a line of code.

Robot identities are artificial; they are designed to imitate, with limited fidelity, a human identity. Designers and programmers creating robot identities tend to turn toward humanlike identities. In many situations, this default makes sense, but **deliberately deviating from a humanlike design can serve a purpose.** Because robotic identities are entirely constructed, designers have a substantial amount of latitude in how they create them [6, 25]. Robot identity affords opportunities to expand the way a robot communicates with its users. It also means robot identity design should be thoughtful: designers must consider which goals a robot's identity is serving, and whether a humanlike identity is effective in achieving these goals.

There are certainly advantages to anthropomorphism in robotics. People tend to prefer more humanlike robots over their non-anthropomorphic counterparts in service roles [10] and find anthropomorphic agents more trustworthy [3, 13]. Robots that produce humanlike speech are assumed to be more capable than robots that don't [2], and anthropomorphic agents that produce more humanlike text are also perceived to be more capable [8]. People are more comfortable around robots that communicate like people—speaking aloud, even to other robots—than robots that communicate in ways that aren't transparent or understandable to humans [22]. People are sometimes uncomfortable with non-humanlike embodiment/identity combinations [11, 17], but seem to accept humanlike ones.

The advantages of anthropomorphism described above may lead to undesirable consequences. While it may be useful to make robots seem competent (for example, by producing humanlike speech), perceiving a robot as more capable than it is can cause someone to inappropriately rely on it, interfering with effective task completion, or lead to being dangerously inattentive to the robot and other hazards in the environment [15]. Ultimately, calibrating trust [12] so users rely on robots to an appropriate degree [9], is more important than simply making the robot seem very trustworthy and capable. Setting expectations of a robot's capability low also leads to less disappointment and more positive assessment of competence [14].

Heightened expectations for anthropomorphic robots are not always a problem; in most social robotics situations, these increased expectations carry low ethical risk [24]. In general, however, robot anthropomorphism does not seem to lead to consistent improvements on task performance [18]. While anthropomorphic robots may be more likable, the ultimate goal for most robots is successful task completion, not simply likability. Robot designers can use non-anthropomorphic identity signals to add functionality that isn't available for robots performing strictly humanlike identities.

2 USES OF NON-ANTHROPOMORPHIC IDENTITY

In the following sections we detail some of the uses of non-humanlike identities. This is by no means an exhaustive list, but a collection of situations and strategies that offer opportunities for non-humanlike identities to be in some way useful.

Robo-Identity at HRI24, March 11, 20124, Boulder, CO
2024. ACM ISBN 978-x-xxxx-xxxx-x/YY/MM
https://doi.org/XXXXXXX.XXXXXXX

2.1 Setting Expectations of Ability

When a robot is anthropomorphic, either in its form or behavior, it is perceived differently than a non-anthropomorphic robot. People have higher expectations of speaking robots' abilities, even in areas not related to speech or social interaction [2]. Humanlike speech also changes how people attempt to communicate with the robots [20]. Using speech to anthropomorphize a robot makes it seem more skilled in other domains, but setting appropriate expectations is beneficial for successfully interacting with people.

Robot designers can set expectations for a robot by moving away from anthropomorphic communication. Consider two designs for a robot that guides people through a building: one talks to visitors and explains the history of the building, and the other displays the message, "Follow me to the room" on a screen, and doesn't communicate verbally. If the visitor wanted to go to a different location in the building, they would be much more likely [20] to expect the speaking robot to correctly respond to the request. If the robot isn't capable of correctly interpreting and responding to the request, it might be worth the decreased likability to communicate the robot's limitations—and avoid the disappointment of failing to meet expectations [14].

2.2 Social Relationships

Not every robot that works with people requires the same level of sociability. Even among collaborative robots, social requirements differ. Some may be primarily social agents, designed to engage in extended conversations to answer questions, tutor students, or provide entertainment. For agents like this, maximizing their humanlikeness is probably desirable, because it sends a strong message that they are likable [18] and sociable.

Other robots require a lower degree of sociability, perhaps navigating in a shared space with a person, or deciding which portion of shared task comes next. These tasks may not require as much social interaction; in human-human encounters they are often communicated simply through gestures, body language, or brief speech. In some situations—a noisy working environment, while working on a focus-intensive task, or when having only passing contact with people—as long as nonverbal communication is sufficient to safely complete the task, adding extraneous social qualities may be a hindrance to the robot's function. In cases like this, a robot's identity could be manipulated to be less humanlike and less social.

2.3 Agency

Robots can explain their decision-making in a way that has implications on both their identities and apparent level of agency. Consider three ways for a robot to explain what it just did:
"I chose this action because it had a 70% chance of success."
"I chose this action because my programming indicated this action had a 70% chance of success."
"I am programmed to choose this action in this situation, because it had the highest chance of success."

Each framing indicates something about the structure of the robot's decision-making—essentially separating out the 'communicator' (the public-facing identity who is speaking) from the 'evaluator' (the programming that predicts the success of any given action). In the last case, it also separates out the robot's policy (the programming that decides what it will ultimately do) from the other pieces of the system. In every case, however, the robot uses the personal pronoun "I," performing identity by presenting some notion of self.

The 'communicator' is likely to be seen as the robot's identity, as it is tied to the pronoun "I." In each of the above examples, the robot implies that it has different levels of agency and responsibility in its decision-making. In the first example, the robot has high agency and is responsible for evaluating success; in the second, it has agency but distances itself from the evaluation; in the the third, it does not have agency in making the decision.

This separation between evaluating, choosing, and communicating actions is notably distinct from the way humans portray identity. Humans cannot blame "programming" for making a mistake, and cannot entirely separate out the evaluation of choices from the self that carries out the final decision. For robots, however, these systems are likely to be entirely independent. Highlighting this non-humanlike separation of mental function gives a robot flexibility to distance itself from responsibility for the actions it takes.

2.4 Robot Groups

With groups or swarms of robots, the options for presenting identity expands, as there are both individual and group dynamics. A central decision is whether the group should be presented as a single entity, a group of unique individuals, or something in-between. Specific design cues for this have been explored by Bejarano et al. [1], finding that voice and name cues are particularly important in signaling that either a group is a single entity or has individual members.

In many cases it is useful to portray a robot group as a single collective—if every member of a robot swarm is using the same control code, taking the same actions in pursuit of the same goal, the group functionally is one collective made up of multiple parts [16]. In some cases a human collaborator needs to work with individual members of the group. In a shared autonomy setting, for example, a human operator could periodically take control of specific group members, while the other agents maneuver independently. In a task like this, individually identifying group members might aid the operator in keeping track of the state of each robot. In some situations, it may also be useful to alternate between presenting a group as a cohesive entity and as a collection of individuals.

Presenting a distinct identity for each group member is not scalable to very large robot groups, but there are intermediary options. For example, a group of 30 robots could be divided into units of 10 robots, with each designating a spokesperson who communicates on behalf of the unit. This is not a truly humanlike identity presentation (since most of the robots present no identity at all), but the abstraction mirrors the hierarchical structure of human government, business, and military organizations, allowing human collaborators to have distinct social touchpoints within the robot group. This strategy would allow the social or functional advantages of robots performing identities to scale to larger groups of robots. Another use for limited identity performance in a robot group is to define agent roles. For example, if some robots are doing inspection, some are doing repair work, and some are transporting materials, an identity cue tied to each job ("inspector," "repair worker", "transporter") could give human collaborators enough information to understand the function of each robot without the mental overhead of differentiating between each agent via behavior.

The ability to merge into a group and entirely relinquish individual identity is not possible with a humanlike identity portrayal, and it can help to simplify robot groups and make them more understandable for humans collaborating with them.

2.5 Matching Form

A robot's physical form has inescapable consequences for how it is perceived as both a social actor and a functional tool or collaborator [4]. A robot's actions or presentation are unlikely to overcome the impact of its embodiment [18], so no matter how much designers try to anthropomorphize a robot with a non-humanoid shape, it will likely be perceived as less anthropomorphic than even a minimally social robot with a humanlike form. With this in mind, if a robot has a zoomorphic or mechanical form, it may not be worth it to make its behavior more humanlike.

Making non-humanoid robots behave like humans also raises new design challenges. For a humanoid robot, presenting an identity that reinforces its humanness is fairly straightforward: the way the robot would imitate human body language is obvious. But for non-humanoid social robots, the comparisons may not be obvious. How should a drone say hello? Does a wheeled base have body language? Even if these design problems are worked out, they may not be very effective; it has been found that functional robots are more likeable when they have robot-optimized behavior, rather than attempting to imitate human motion [7].

For quadrupeds, robots that are more zoomorphic than anthropomorphic, the obvious body language to draw on is not that of a human, but of an animal. Body language that imitates familiar pet species like dogs or cats is an understandable signal when transferred to both zoomorphic and more mechanical robots [5, 19, 21].

Embracing zoomorphism may also help to set expectations. If a robot engages with a user more like a dog than a human—primarily using dog-inspired body language and non-speech sounds—a novice user may expect its intelligence level to match that of a dog, rather than another human. In situations where it's beneficial to communicate a robot's limitations to a user, targeting animal-level expectations may be a useful strategy.

2.6 Flexible Embodiment

Several examples of a non-humanlike relationship between identity, mind, and body have been evaluated or proposed. These flexible embodiment strategies, where the relationship between a robot's physical form and the identity “inhabiting" it are not 1:1 or are fluid, have functional uses that are not possible with a humanlike agent.

Reig et al. [17] evaluated flexibly-embodied strategies for trust recovery. They portrayed a robot that had failed at a task either performing a software update, or “migrating" to a new body, and compared these strategies to getting assistance from a different robot. Humans cannot simply take over a new body, and there isn't a perfect human parallel for a software update. These nonhumanlike strategies were found to be effective at preserving trust and perceived competence in the robot.

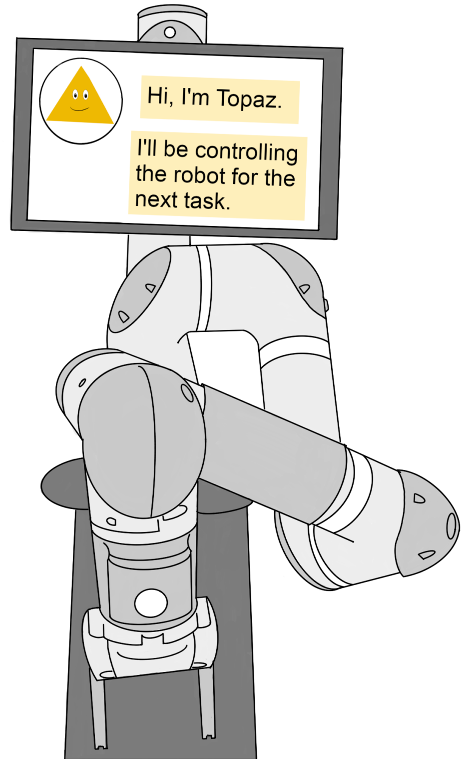

Figure 1: A robotic arm with a screen displaying an identity that is inhabiting the robot.

The exploration of various social presence options for embodied agents by Luria et al. [11] include many possible uses of flexible social presence. One strategy was re-embodiment, where a single "assistant" presence moved between bodies in the user's home and work, or a worker presence moved between multiple forms while completing an involved process in a government office. Participants were generally comfortable with this. This kind of persistent social presence may allow users to spend more time with the agent, building familiarity and perhaps giving them a better understanding of its strengths and limits. It may also reduce the mental load of keeping track of multiple agents. Participants were also comfortable with a social presence that controlled multiple forms simultaneously. This offers many of the same advantages as re-embodiment, and is available in consumer products like home-assistant devices.

Williams et al. [23] propose the Deconstructed Trustee Theory, where a user's trust in a robot's physical form is separated from trust in the agent inhabiting the form. This has applications in manipulating robot trustworthiness, which could be useful in the cases of unreliable hardware or software. As in the trust recovery work of Reig et al. [17], in the event of hardware failure, designers can try to associate the blame of failure with the physical form, without damaging user's trust in the identity controlling it.

We are currently investigating using a robot's identity to communicate the kind of control software it is using. A robot arm completing a warehouse box packing task alternates between running two types of control software. One program is extremely reliable, but can only be used for packing certain items, and the other is less reliable, but can be used for every item. The robot uses one of three methods to indicate which control software is running. In one case, its screen displays an algorithm name associated with the active software. In another, the screen displays a humanlike name and an image of a face for each program. In the third, the screen displays the humanlike name, the face, and a social greeting where the "identity" introduces itself and explains that it is now controlling the robot, as in Figure 1. Our co-embodiment strategy takes advantage of the considerable flexibility in robot identity presentation to give users more information about the robot's state and its control software. We find that the use of a socially-engaged identity signal allows participants to distinguish between the different types of control software the robot used, while a nonsocial identity signal and the absence of identity signal do not.

3 CONCLUSIONS

The flexible nature of robotic identity performance gives robot designers considerable freedom to add utility to social and collaborative robots. The possible uses include expectation-setting and trust calibration, giving users insights into the robot's decision-making, and making robot groups more manageable. At times there may be tension between users' preferences and the robot's functionality, so robot identity manipulations must be carefully considered in the context in which they will be used: in many situations a more anthropomorphic identity will serve the robot's goals. In general, however, non-humanlike robot identities are under-explored and offer opportunities to expand the capabilities of social robots.

REFERENCES

[1] Alexandra Bejarano, Samantha Reig, Priyanka Senapati, and Tom Williams. 2022. You had me at hello: The impact of robot group presentation strategies on mental model formation. In *2022 17th ACM/IEEE International Conference on Human-Robot Interaction (HRI)*. IEEE, 363–371.

[2] Elizabeth Cha, Anca D Dragan, and Siddhartha S Srinivasa. 2015. Perceived robot capability. In *2015 24th IEEE International Symposium on Robot and Human Interactive Communication (RO-MAN)*. IEEE, 541–548.

[3] Ewart J De Visser, Samuel S Monfort, Ryan McKendrick, Melissa AB Smith, Patrick E McKnight, Frank Krueger, and Raja Parasuraman. 2016. Almost human: Anthropomorphism increases trust resilience in cognitive agents. *Journal of Experimental Psychology: Applied* 22, 3 (2016), 331.

[4] Terrence Fong, Illah Nourbakhsh, and Kerstin Dautenhahn. 2003. A survey of socially interactive robots. *Robotics and Autonomous Systems* 42, 3 (2003), 143–166. https://doi.org/10.1016/S0921-8890(02)00372-X Socially Interactive Robots.

[5] Moojan Ghafurian, Gabriella Lakatos, Zhuofu Tao, and Kerstin Dautenhahn. 2020. Design and evaluation of affective expressions of a zoomorphic robot. In *Social Robotics: 12th International Conference, ICSR 2020, Golden, CO, USA, November 14–18, 2020, Proceedings 12*. Springer, 1–12.

[6] Ryan Blake Jackson, Alexandra Bejarano, Katie Winkle, and Tom Williams. 2021. Design, performance, and perception of robot identity. In *Workshop on Robo-Identity: Artificial identity and multi-embodiment at HRI*, Vol. 2021.

[7] Daphne E Karreman, Geke DS Ludden, and Vanessa Evers. 2019. Beyond R2D2: Designing multimodal interaction behavior for robot-specific morphology. *ACM Transactions on Human-Robot Interaction (THRI)* 8, 3 (2019), 1–32.

[8] Bart P Knijnenburg and Martijn C Willemsen. 2016. Inferring capabilities of intelligent agents from their external traits. *ACM Transactions on Interactive Intelligent Systems (TiiS)* 6, 4 (2016), 1–25.

[9] John D Lee and Katrina A See. 2004. Trust in automation: Designing for appropriate reliance. *Human factors* 46, 1 (2004), 50–80.

[10] Kate Letheren, Jolanda Jetten, Jonathan Roberts, and Jared Donovan. 2021. Robots should be seen and not heard… sometimes: Anthropomorphism and AI service robot interactions. *Psychology & Marketing* 38, 12 (2021), 2393–2406.

[11] Michal Luria, Samantha Reig, Xiang Zhi Tan, Aaron Steinfeld, Jodi Forlizzi, and John Zimmerman. 2019. Re-Embodiment and Co-Embodiment: Exploration of social presence for robots and conversational agents. In *Proceedings of the 2019 on Designing Interactive Systems Conference*. 633–644.

[12] Bonnie M Muir. 1987. Trust between humans and machines, and the design of decision aids. *International journal of man-machine studies* 27, 5-6 (1987), 527–539.

[13] Manisha Natarajan and Matthew Gombolay. 2020. Effects of anthropomorphism and accountability on trust in human robot interaction. In *Proceedings of the 2020 ACM/IEEE international conference on human-robot interaction*. 33–42.

[14] Steffi Paepcke and Leila Takayama. 2010. Judging a bot by its cover: An experiment on expectation setting for personal robots. In *2010 5th ACM/IEEE International Conference on Human-Robot Interaction (HRI)*. IEEE, 45–52.

[15] Raja Parasuraman and Victor Riley. 1997. Humans and automation: Use, misuse, disuse, abuse. *Human factors* 39, 2 (1997), 230–253.

[16] Giovanni Cosimo Pettinaro, Ivo W Kwee, Luca Maria Gambardella, Francesco Mondada, Dario Floreano, Stefano Nolfi, Jean-Louis Deneubourg, and Marco Dorigo. 2002. SWARM Robotics: A Different Approach to Service Robotics: a different approach to service robotics. In *Proceedings of The 33rd International Symposium on Robotics (ISR)*. ETH-Zürich.

[17] Samantha Reig, Elizabeth J Carter, Terrence Fong, Jodi Forlizzi, and Aaron Steinfeld. 2021. Flailing, hailing, prevailing: Perceptions of multi-robot failure recovery strategies. In *Proceedings of the 2021 ACM/IEEE International Conference on Human-Robot Interaction*. 158–167.

[18] Eileen Roesler, Dietrich Manzey, and Linda Onnasch. 2021. A meta-analysis on the effectiveness of anthropomorphism in human-robot interaction. *Science Robotics* 6, 58 (2021), eabj5425.

[19] Vanessa Sauer, Axel Sauer, and Alexander Mertens. 2021. Zoomorphic Gestures for Communicating Cobot States. *IEEE Robotics and Automation Letters* 6, 2 (2021), 2179–2185. https://doi.org/10.1109/LRA.2021.3060416

[20] Valerie K Sims, Matthew G Chin, Heather C Lum, Linda Upham-Ellis, Tatiana Ballion, and Nicholas C Lagattuta. 2009. Robots' auditory cues are subject to anthropomorphism. In *Proceedings of the Human Factors and Ergonomics Society Annual Meeting*, Vol. 53. SAGE Publications Sage CA: Los Angeles, CA, 1418–1421.

[21] Dag Sverre Syrdal, Kheng Lee Koay, Márta Gácsi, Michael L Walters, and Kerstin Dautenhahn. 2010. Video prototyping of dog-inspired non-verbal affective communication for an appearance constrained robot. In *19th international symposium in robot and human interactive communication*. IEEE, 632–637.

[22] Xiang Zhi Tan, Samantha Reig, Elizabeth J Carter, and Aaron Steinfeld. 2019. From one to another: how robot-robot interaction affects users' perceptions following a transition between robots. In *2019 14th ACM/IEEE International Conference on Human-Robot Interaction (HRI)*. IEEE, 114–122.

[23] Tom Williams, Daniel Ayers, Camille Kaufman, Jon Serrano, and Sayanti Roy. 2021. Deconstructed Trustee Theory: Disentangling Trust in Body and Identity in Multi-Robot Distributed Systems. In *Proceedings of the 2021 ACM/IEEE International Conference on Human-Robot Interaction* (Boulder, CO, USA) *(HRI '21)*. Association for Computing Machinery, New York, NY, USA, 262–271. https://doi.org/10.1145/3434073.3444644

[24] Katie Winkle, Praminda Caleb-Solly, Ute Leonards, Ailie Turton, and Paul Bremner. 2021. Assessing and Addressing Ethical Risk from Anthropomorphism and Deception in Socially Assistive Robots. In *Proceedings of the 2021 ACM/IEEE International Conference on Human-Robot Interaction* (Boulder, CO, USA) *(HRI '21)*. Association for Computing Machinery, New York, NY, USA, 101–109. https://doi.org/10.1145/3434073.3444666

[25] Katie Winkle, Ryan Blake Jackson, Alexandra Bejarano, and Tom Williams. 2021. On the flexibility of robot social identity performance: benefits, ethical risks and open research questions for HRI. In *HRI Workshop on Robo-Identity*.

Received 20 February 2007; revised 12 March 2009; accepted 5 June 2009