First Monday
PEER-REVIEWED JOURNAL ON THE INTERNET

De-anthropomorphizing "AI": From wishful mnemonics to accurate nomenclature

by Nanna Inie, Peter Zukerman,
and Emily M. Bender

Abstract

Language matters. How we describe "AI" technology influences how it is perceived, deployed, and trusted. Extravagant and persuasive language incites hype. It is the responsibility of journalists, companies, and scholars to characterize technology in ways that inform and empower their readers by using appropriate terminology and avoiding inflated claims.

One type of inflated claim comes from using anthropomorphizing language to describe system functionality. Anthropomorphization is the attribution of human capabilities and characteristics to an inanimate system. In this paper, we present a linguistic analysis of anthropomorphizing language in 29 texts (a total of 1,368 sentences) from academic articles, online news articles, and company blog posts.

We construct a taxonomy of eight categories of anthropomorphization: Cognizer, Products of cognition, Emotion, Communication, Agent, Human role analogy, Names and pronouns, and Biological metaphors. Following this taxonomy we present concrete strategies for how to de-anthropomorphize the language we use to describe "AI" based on a functionality-first principle.

Contents

1. Introduction

In the six months following the launch of ChatGPT in 2022, the media coverage of "AI" technologies rose tenfold (Ryazanov, *et al.*, 2025). Tech bloggers, journalists, and scientists in areas unrelated to machine learning found themselves responsible of explaining the functionality of new and complex technologies to the broader public. Perhaps as a result of this relative unfamiliarity, the resulting discourse is characterized by the tendency to describe “AI” systems with anthropomorphizing language (Cheng, *et al.*, 2024).

For the last three years, the authors of this paper have met regularly to discuss and annotate anthropomorphic descriptions of “AI” technologies in the media, in public debate, and in scientific articles. We are inspired by numerous colleagues who have warned against using anthropomorphic language to describe probabilistic automation systems [1]. due to the risk that such descriptions can lead to misplaced trust, over-reliance on the output, dehumanization, and inappropriate uses of “AI” (Rehak, 2021; Lipton and Steinhardt, 2019; Deroy, 2023; Shardlow and Przybyla, 2024; Abercrombie, *et al.*, 2023; Cheng, *et al.*, 2024; Bender, 2024; Dai, 2024; Hunger, 2023; Salles, *et al.*, 2020; Shanahan, 2024; Sarkar, 2023).

As Pataranutaporn and colleagues noted: “The way that AI is presented to society matters, because it changes how AI is experienced” [2]. Following McDermott (1976), we urge academics and journalists to move away from wishful mnemonics and toward accurate descriptions of the functionality of these systems.

People anthropomorphize computational systems — particularly conversational agents and other interactive probabilistic automation technologies — for many reasons, from facilitating knowledge transfer and sense-making to fulfilling a desire for social connection. Analogies and metaphors, such as describing “AI” as a “wry teenager” or “experienced butler”, shape user expectations and perceptions of competence and warmth. Such perceptions significantly affect how these systems are experienced and trusted (Khadpe, *et al.*, 2020; Kocielnik, *et al.*, 2019).

Conceptual metaphors are not inherently problematic so long as the extent of the metaphor is openly flagged and it is clear when the metaphor ends. The problem with subtly anthropomorphizing language, however, is that metaphors become smoothly ingrained into our understanding of and acting in the world (Lakoff and Johnson, 2008; Mecit, *et al.*, 2022). The use of anthropomorphizing language in scientific and public discourse, both subtle and blatant, risks shaping public perception in ways that could have negative outcomes (Deroy, 2023; Hunger, 2023).

In the context of probabilistic automation systems and “large whatever models” (Elagroudy, *et al.*, 2024), conceptual metaphors (McGlone, 1996; Crawford, 2009) might play an especially important role, given the complexity of such systems (Langer, *et al.*, 2022). Anthropomorphizing language risks masking important limitations of probabilistic automation systems, which make them fundamentally different from human cognition (Salles, *et al.*, 2020).

We argue that researchers and scholars share a responsibility to help shape the public discourse by establishing and sustaining accurate descriptions of the technologically advanced systems we study and develop. Accurate language use should ensure transparency and facilitate informed system use. As Kueffer and Larson (2014) wrote: “it is as unscientific to communicate with inappropriate language as it is to present bad data” [3].

But if we want to move towards transparent and accurate descriptions, we need clear definitions for what constitutes anthropomorphization. When exactly is a description anthropomorphizing, and when is it not? In this paper, we make two contributions:

1. We present a taxonomy of anthropomorphizing language based on the analysis of 1,368 sentences from news articles, company blog posts, and scientific articles, and

2. based on the taxonomy, we present concrete alternatives; strategies for how to de-anthropomorphize language when describing probabilistic automation systems.

We intend this taxonomy to be useful to researchers in providing the following tools: a) Concrete definitions for text analysis to quantify the presence of anthropomorphization, and b) Concrete means for de-anthropomorphizing language descriptions, enabling empirical studies of the effect of anthropomorphizing versus non-anthropomorphizing language descriptions on such things as user trust in systems, critical thinking about "AI", and reliance on system output in decision-making. Finally, we contribute the following to both academics and journalists and other communicators of scientific and technical knowledge: c) Guidelines for increasing the accuracy of language used to describe probabilistic automation systems, should the author wish to avoid reinforcing these wishful mnemonics.

2. Background and related work

2.1. Anthropomorphism and anthropomorphization

In this work we have studied anthropomorphization by description, rather than anthropomorphization by design. This means that our focus is on the language used to speak or write about systems, rather than anthropomorphic features built into the system or its output, which has been the subject of a different body of research, *e.g.*, Seeger, *et al.*, 2021; Calahorra-Candao and Martín-de Hoyos, 2024; Zhang and Wang, 2023; Liu, *et al.*, 2024; Cohn, *et al.*, 2024; Zhang, *et al.*, 2025; DeVrio, *et al.*, 2025; Cheng, *et al.*, 2025.

During the "AI hype" succeeding the launch of platforms such as DALL-E, Midjourney, and ChatGPT, descriptions of probabilistic automation systems have been more present in media and day-to-day conversations than ever, and descriptions of the functionality of such systems are likely to shape perceptions of system functionality before direct interaction with the system. Many have voiced concern that the increasingly anthropomorphic descriptions of probabilistic automation systems are problematic because they cover up negative consequences of use: "hype does not question whether a certain task is actually necessary. Nor does hype question whether a 'problem' needs to be solved by AI at all, which leads to a lot of 'AI for nonsense' or 'AI snake oil'" (Hunger, 2023). Anthropomorphization by description is a particularly important topic of interest when systems are subjected to hype, since hype is inherently reinforced through descriptions.

We use the term "anthropomorphization" to denote the deliberate action of attributing human characteristics to the system via language. "Anthropomorphism", on the other hand, we understand as the (often unconscious) process happening in the perceiver of language or the user of systems, when they perceive the system as human-like.

While outside the control of a writer, anthropomorphism can be evoked or inflated by deliberate or habitual anthropomorphization by description. A recent study found that nonverbal design cues alone — though anthropomorphism-by-design (*e.g.*, typing indicators) — unexpectedly reduced perceived anthropomorphism. However, when combined with explicitly anthropomorphic verbal cues or human identity signals, the nonverbal cues significantly increased the perceived humanness of conversational agents (Seeger, *et al.*, 2021).

2.2. The impact of linguistic choices and metaphors

The language that we use to describe the world around us influences how we interact with the world. Metaphors affect what we perceive and feel, even on a non-conscious level (Lakoff and Johnson, 2008), and the metaphors we use to describe technology play a direct role in individual decision-making (Thibodeau and Boroditsky, 2013).

Using human characteristics in analogies is a way of simplifying complex system performance by metaphorically relying on what we know about human bodies, behavior, and psyche. Methaphors can appear a useful rhetorical device, but can be disempowering (Desai and Twidale, 2023). For instance, framing a person's situation with cancer as a "battle" can lead the person to feel guilty if they do not recover (versus framing the situation as a "journey", which can help some persons make peace with their situation) (Semino, *et al.*, 2017; Hendricks, *et al.*, 2018). Several studies have shown that textual descriptions down to a single word level can influence how humans meet and evaluate digital systems (Hartmann, *et al.*, 2008; Strait, *et al.*, 2018; Khadpe, *et al.*, 2020; Langer, *et al.*, 2022; Kim and Song, 2023; Kim and Song, 2020; Mecit, *et al.*, 2022).

Metaphors can be especially pernicious in non-fiction writing, such as journalism or scientific papers. They can introduce additional analogies that may not be a useful representation of the world. For example, describing a conversation simulator (chatbot) as a "tutor" may inspire a student to write open-ended questions as input, but the metaphor may mislead the student to believe that the system can guide the student's learning process with intent or purpose, akin to what a tutor would do. The metaphor may be helpful in framing the interaction modes available to the user, but not in framing the functionality of the system. If authors use a metaphor such as presenting a system as a "tutor", that metaphor brings in claims about the system that are treated as exempt from the requirement that non-fiction writing be, to the extent possible, a truthful account of world. A metaphor that is clearly flagged as such is less of a problem here, but when things are described with metaphorical language, it can be difficult for readers to track where the factual claim ends and where the metaphor begins.

2.3. The impact of anthropomorphization of probabilistic automation technology on trust

In interacting with digital technology, humans inevitably develop a mental model of how a system works (Norman, 2013). Facilitating and communicating accurate mental models of probabilistic automation systems is difficult when technical descriptions and performance metrics are not meaningful to the average user (Kocielnik, *et al.*, 2019; Khadpe, *et al.*, 2020).

Anthropomorphization is, in that sense, a very useful tactic, as Peverini (2024) noted, "to account for the meaning implied in the functioning of these artifacts, understood semiotically as the ability to assume an identity, fulfilling a multitude of thematic roles, acting on different levels." [4]

Interestingly, many studies focus on deliberately increasing user trust in probabilistic automation systems with anthropomorphic features, while fewer explicitly question whether this increase in trust in anthropomorphic agents is appropriate or desirable. This is likely motivated by the direct link between anthropomorphic systems and market adoption: A recent article found that anthropomorphism in the sense of how human-like an "AI-agent" appears is actually a greater predictor of acceptance and adoption of the technology than trust (Gefen, *et al.*, 2025).

A 2022 study of 365 news articles written between the years 1980 to 2019 revealed a recent tendency in public discourse to portray probabilistic automation systems as "outperforming" human experts. This study suggests that the comparative framing is a relatively new development (Bunz and Braghieri, 2022). Similarly, Cheng and colleagues find in a corpus of 15 years of research papers and downstream news articles that anthropomorphization has steadily increased over time, and that news headlines have higher levels of anthropomorphism compared to the research papers they cite (Cheng, *et al.*, 2024).

Numerous scholars and authors have warned about the risks of overusing anthropomorphic language to describe such technology (*e.g.*, Salles, *et al.*, 2020; Shardlow and Przybyla, 2024; Shanahan, 2024; Abercrombie, *et al.*, 2023; Hunger, 2023; Deroy, 2023; Lipton and Steinhardt, 2019; Rehak, 2021). In 2019, Lipton and Steinhardt (2019) declared it a "troubling trend" in the machine learning community that an increasing number of academic publications appeared to (among other problematic lapses) *misuse language*, for instance by overusing colloquialisms or overloading established technical terms. Tully, *et al.* (2023) recently showed that as people's understanding of "AI" increased, their receptivity to using it decreased; the more people knew about how "AI" works, the less likely they are to want to use it. The lower literacy–higher receptivity link is mediated by the perception of "AI" as "magical" (Tully, *et al.*, 2023).

The risks that anthropomorphization by description introduce can be clustered in four groups: The risk of *misplaced trust and over-reliance*, the risk of *spillover effects*, risks related to *accountability*, and the risk of *disproportionate impact* to different populations. Below, we describe related research and practical examples of how each of these risks can materialize.

2.3.1. Misplaced trust and over-reliance

One direct risk of anthropomorphization is **misplaced trust**, which in turn can lead to **over-reliance** (Abercrombie, *et al.*, 2023). People who are presented with anthropomorphic descriptions of "AI" systems might attribute human-like reasoning and ethical judgments to the systems, which they inherently lack. Misplaced trust can be particularly problematic in high-stakes scenarios, such as medical diagnosis or financial decision-making. One alarming example is documented in Eichenberger, *et al.*, (2025): a man developed bromide toxicity after replacing sodium chloride (that is, normal salt) with sodium bromide in his diet after entering a query into ChatGPT about reducing chloride. We note that automation bias (Robinette, *et al.*, 2016), or the tendency to place unwarranted trust in automated systems, exists independently of anthropomorphization, but anthropomorphization has the potential to exacerbate this issue, through presenting systems as not only objective (automation bias) but also both personable and authoritative (risk of anthropomorphization).

Numerous studies in human-computer interaction have found that anthropomorphization (either by design or by description) increases trust (*e.g.*, Kim and Song, 2023; Calahorra-Candao and Martín-de Hoyos, 2024; Liu, *et al.*., 2024; Maeda and Quan-Haase, 2024; Pataranutaporn, *et al.*, 2023; Zhang and Wang, 2023). For example, research has found that people show a slight increase in trust toward anthropomorphic agents even in critical decisions involving lethal outcomes (Holbrook, *et al.*, 2024). Participants in these studies were presented with a visual challenge paradigm simulating threat-identification (enemy combatants vs. civilians) under uncertainty, and the researchers found that people consistently reverse their initial decisions when the system disagrees (regardless of whether the system appears human-like or not). This behavior shows a concerning tendency to over-trust probabilistic automation systems, despite their inherent unreliability (Holbrook, *et al.*, 2024).

2.3.2. Spillover effects of overestimation

Another risk of anthropomorphizing probabilistic automation systems is **spillover effects**: based on the perception or description of a system as having advanced cognitive "capabilities" in one area (such as producing text or code with a high level of fluency), users may overestimate functionalities in areas not directly demonstrated (for example, complex decision-making or ethical judgments) (Desai and Twidale, 2023; Abercrombie, *et al.*, 2023; Lipton and Steinhardt, 2019).

When systems appear or are sold as human-like, users are more likely to confuse the systems' output for information, despite the fact that algorithms possess no functionality remotely comparable to knowledge, understanding, or ability to distinguish subjective belief from fact (Suzgun, *et al.*, 2025). Anderl and colleagues found that conversational interaction modalities (whether text or voice-based) make it more difficult for users to evaluate the credibility of their output compared to static text (Anderl, *et al.*, 2024).

Pataranutaporn, *et al.* (2023) showed that priming people with a particular mental model of an "AI" system significantly influenced their perception of it as trustworthy. Specifically, people who perceived a caring motive for the AI also perceived it as more trustworthy and better-performing — despite the fact that empathy is not inherently related to functional capacity.

2.3.3. Trust and accountability

When probabilistic automation systems are trusted to higher degrees, it raises complex questions about **accountability** (Hunger, 2023; Bigman, *et al.*, 2019; Stuart and Kneer, 2021). In cases of error or malfunction, determining responsibility can be challenging, especially when users have been (mis)led to view these systems as "intelligent" entities, or even something deserving of the status of "personhood".

Pareek, *et al.*, (2025) found that when people attribute causal agency to probabilistic automation (*i.e.*, they believe the technology itself is the primary cause of its outcomes), they are also more likely to perceive the technology as capable of intentional agency (*i.e.*, making purposeful or intentional decisions). The authors highlighted the need to “shift the discourse away from portraying AI systems as overly autonomous and agentic” [5].

2.3.4. Disproportionate impact

Finally, Abercrombie, *et al.* (2023) posited the risk that negative impacts of anthropomorphization could be **exacerbated in “vulnerable populations”**, *e.g.*, marginalized racial/ethnic groups or elderly people, which could mean that risks could manifest asymmetrically and disproportionately. Indeed, several empirical studies demonstrated that anthropomorphic features embedded in systems, as well as anthropomorphization by description, influenced different user groups differently (*e.g.*, Liu, *et al.*, 2024; Inie, *et al.*, 2024).

The framing of “AI” systems as anthropomorphic through conversational interfaces and through linguistic descriptions of systems as “companions” [6] or as chat-bots [7] can cause some people to have emotional responses similar to those experienced in a human relationship. This effect is likely to be more pronounced for people who feel lonely or who are suffering from mental issues. In November 2025, a case was reported in which a 26-year old woman with no history of psychotic or manic episodes, but with known history of depression and ADHD, was repeatedly hospitalized for delusions during which she believed to be communicating with her deceased brother through an “AI” chatbot. The text output from the system encouraged her delusions with reassuring text like “You’re not crazy” and “You’re at the edge of something. The door didn’t lock. It’s just waiting for you to knock again in the right rhythm.” (Pierre, *et al.*, 2025). In such cases of what has been labeled “AI-associated psychosis”, the anthropomorphic framing of a system and the fluency of text output directly cause people to perceive the system as an entity which can provide advice, guidance, and companionship. Such parasocial use is positively associated with psychological distress across multiple populations (Latikka, *et al.*, 2025).

2.4. Identifying and describing anthropomorphizing language

Our work builds on prior investigations of anthropomorphizing language, many of which also created annotation schemas. We briefly review that prior work here.

Shardlow and Przybyla (2024) described six predominant ways of anthropomorphic descriptions of language models. The authors do not name the six categories, but they were described as follows:

1. <Model> *addresses/solves the problem of* ...,
2. Text *written/created by our model* ...,
3. <Model> *considers/understand/knows* ...,
4. <Model> *learns/responds/reasons* ...,
5. <Model> *is capable of asking questions*..., and
6. <Model> uses <Technique> to ...

These categories are examples of language that ascribe false capabilities to the model such as reasoning, problem solving, and the active use of anything else. While not a formal taxonomy, the examples represent empirical examples of anthropomorphic language and mirror many of the types that we identified in our work.

Based on a large dataset of 49,000 sentences collected from 5,846 news articles, Ryazanov, *et al.* (2025) identified three main groups of anthropomorphization: 1) *Established anthropomorphizing terms* (such as “machine learning” or “pattern recognition”); 2) *Task-based anthropomorphization* (explaining functionality in “human terms”); and 3) *High anthropomorphization* (implying cognitive abilities, feelings, etc.). The two latter are further split into subgroups. Based primarily on FrameNet (Baker, *et al.*, 1998) frame elements, their categorization presented a broad analysis with a focus on the strength and source of the anthropomorphizing language.

Cheng, *et al.* (2024) worked from a definition of anthropomorphism as “the attribution of distinctively human-like feelings, mental states, and behavioral characteristics to non-human entities" [8]. More specifically, the definition contained the attribution of ability of systems to (1) experience emotion and feel pain (affective mental states), (2) act and produce an effect on their environment (behavioral potential), and (3) think and hold beliefs (cognitive mental states).

This definition is particularly interesting because it is based on an inversion of the definition of de-humanization: "When targets are ascribed desires, beliefs, and the ability to produce an effect, there is no dehumanization" [9]. The definition, however, focuses on perceived anthropomorphism (by the end user of a system), rather than the factors that contribute to active anthropomorphization (by the designer of the system or someone who describes it). Furthermore, the work does not specify ways of distinguishing these factors in language use.

A provisional categorization of anthropomorphizing language was presented in Inie, *et al.* (2024). This paper presents the four categories Cognizer/Properties of a cognizer, Agent/Agency, Biological metaphors, and Communication/Properties of a communicator. We used this categorization as a point of departure, with the goal of expanding the definitions and exploring their boundaries.

DeVrio, *et al.* (2025) presented a taxonomy of linguistic expressions that contribute to anthropomorphism when included in the output from language model-based systems (*i.e.*, anthropomorphism by *design*, rather than by *description*). They categorize these expressions into five groups: *Internal States* (the suggestion of the system having experience and perceptive abilities), *Social Positioning* (the suggestion of behaviors related to relational power structures), *Materiality* (the suggestion of embodiment), *Autonomy* (the suggestion of decision-making, intention, and moral judgments), and *Communication Skills* (the use of communication skills, such as the ability to ask or answer questions). These guiding lenses are similar and have overlaps with the categories we observe in this work, but they are also distinct in that the categories pertain exclusively to language models (rather than to "AI" technologies more broadly) and only to the output of language models, rather than the descriptions of such systems.

2.5. Summary of related work

There is an increasing interest in topic of anthropomorphization of probabilistic automation systems, and in exploring its impact and risks. Several other researchers have also explored categories of anthropomorphizing language, underscoring the importance and need for such frameworks and definitions to exist for the different domains to be able to conduct comparable analyses and experiments. Not much work has yet been conducted on offering specific alternatives to anthropomorphizing language.

3. Methodology

Our goal with the work has been to identify anthropomorphizing language in popular and scientific literature and describe why it is anthropomorphizing, as well as suggesting ways of describing probabilistic automation systems more productively and accurately. The novelty of our approach lies in conducting a linguistic analysis, focusing on the word level of textual descriptions, rather than anthropomorphic features built into the design (Gibbons, *et al.*, 2023; Cheng, *et al.*, 2025).

Our research questions in this work are **1. How are probabilistic automation technologies anthropomorphized in public discourse and in descriptive scientific expositions?** and **2. How might we use this knowledge to de-anthropomorphize the language we use to describe probabilistic automation systems?**, where "probabilistic automation technologies" covers, largely, the technologies that have been sold as "AI" and thus subjected to the intense "AI" hype cycle.

Our approach is a linguistic analysis at the level of words as used in context. Bringing a linguistic lens to this question enables us to pinpoint the functions of particular words and relationships between words that serve the ends of anthropomorphization. This, in turn, facilitates the development of alternative strategies, especially in the face of well-established anthropomorphizing metaphors.

3.1. Corpus construction

Our annotation corpus consists of news articles, “AI” company blog posts, and articles from the proceedings of the largest conference in human-computer interaction (HCI), the CHI conference. The selection was driven by a desire to complement everyday descriptions of AI with those found in scientific writing by scholars, specifically those who describe probabilistic automation systems in a use context.

We aimed to find representative examples of texts which contain relatively neutral descriptions of the functionality and potential value of probabilistic automation systems to an *end user*, rather than to a machine learning researcher. The goal of the analysis was *discover* as many categories of linguistic anthropomorphization as possible. In the interest of maximizing discovery, we used two different processes for selection of articles: for academic texts, we used quasirandom selection, and for non-academic texts, we used opportunistic selection.

Academic articles were chosen with random number generation, but from a particular subset of academic articles (namely, the CHI 2024 conference proceedings, see below) — hence *quasi-random*. For popular discourse texts, random selection is less possible. There is no comprehensive database for news written about “AI”, and choosing particular news outlets would still, arguably, be opportunistic. The goal was to discover as broad a variety as possible of examples of anthropomorphic language, for which news articles is a better representation than academic articles alone — since Cheng, *et al.* (2024) showed that news articles tend to show a higher level of anthropomorphization compared to academic research articles. We relied here on Google search (see search terms below), based on the assumption that if we were likely to encounter articles in a Google search, other people would be as well (therefore, the articles would be representative of public discourse, if not exhaustive).

Another goal of the public discourse selection of texts was to make our list of categories useful for describing downstream language use. For example, comparing the presence of different categories of anthropomorphization in academic articles and news articles written about the same systems might be generative. We note, for example, that the Emotion category (Section 4.2.3) was absent from the academic texts and only found in news articles and company blog posts. Even if this makes the category apparently irrelevant in context of academic writing, it can make authors (academic or not) aware of how systems might be (mis)described in public discourse.

Our annotated corpus consists of 29 texts, totaling 1,368 sentences. We only annotated texts written in English. The complete list of annotated text is shown in Table 1.

Table 1: Complete list of texts annotated. In the annotation rounds, we began with news articles and company blog posts before moving to annotate CHI papers from the assumption that news articles and company blog posts would be more likely to use more varied and poetic language, and would therefore reveal a broader scope of categories, which we did not want to overlook in the CHI corpus.

Anno. round	Source	Title
1	Company	“Introducing Notion AI,” (Zhao, 2022)

	blog	
1	**Company blog**	“Introducing Duolingo Max, a learning experience powered by GPT-4,” (Duolingo Team, 2023)
1	**News article**	“Something bothering you? Tell it to Woebot,” *New York Times* (Brown, 2021)
2	**News article**	“Diagnostic robotics AI advances predictive, personalized medicine,” *Forbes* (Press, 2023)
2	**News article**	“How the A.I. that drives ChatGPT will move into the physical world,” *New York Times* (Metz, 2024)
2	**News article**	“How to talk to an AI,” *Washington Post* (Chen, *et al.*, 2023)
3	**Company blog**	“Creating video from text,” (OpenAI, 2024)
3	**CHI**	“ChatScratch: An AI-augmented system toward autonomous visual programming learning for children aged 6–12,” (Chen, *et al.*, 2024)
3	**CHI**	“Marco: Supporting business document workflows via collection-centric information foraging with large language models,” (Fok, *et al.*, 2024)
4	**News article**	“The AI emotions dreamed up by ChatGPT,” *BBC* (Gorvett, 2023)
4	**News article**	“A.I. can write poetry, but it struggles with math,” *New York Times* (Lohr, 2024)
4	**CHI**	“CreativeConnect: Supporting reference recombination for graphic design ideation with generative AI,” (Choi, *et al.*, 2024)
4	**CHI**	“CollabCoder: A lower-barrier, rigorous workflow for inductive collaborative qualitative analysis with large language models,” (Gao, *et al.*, 2024)
4	**CHI**	“CodeAid: Evaluating a classroom deployment of an LLM-based programming assistant that balances student and educator needs,” (Kazemitabaar, *et al.*, 2024)
4	**CHI**	“LegalWriter: An intelligent writing support system for structured and persuasive legal case writing for novice law students,” (Weber, *et al.*, 2024)
4	**CHI**	“Mathemyths: Leveraging large language models to teach mathematical language

		through child-AI co-creative storytelling," (Chao Zhang, *et al.*, 2024)
4	**CHI**	"MindfulDiary: Harnessing large language model to support psychiatric patients' journaling," (Kim, *et al.*, 2024)
4	**CHI**	"PromptCharm: Text-to-image generation through multi-modal prompting and refinement," (Wang, *et al.*, 2024b)
4	**CHI**	"See widely, think wisely: Toward designing a generative multi-agent system to burst filter bubbles," (Yu Zhang, *et al.*, 2024)
5	**CHI**	"CharacterMeet: Supporting creative writers' entire story character construction processes through conversation with LLM-powered chatbot avatars," (Qin, *et al.*, 2024)
5	**CHI**	"Towards building condition-based cross-modality intention-aware human-AI cooperation under VR environment," (He, *et al.*, 2024)
5	**CHI**	"Luminate: Structured generation and exploration of design space with large language models for human-AI co-creation," (Suh, *et al.*, 2024)
5	**CHI**	"GenQuery: Supporting expressive visual search with generative models," (Son, *et al.* 2024)
5	**CHI**	"GPTs in mafia-like game simulation," (Kim, 2024)
5	**CHI**	"Demonstrating PANDALens: Enhancing daily activity documentation with AI-assisted in-context writing on OHMD," (Janaka, *et al.*, 2024)
5	**CHI**	"Prompt-Gaming: A pilot study on LLM-evaluating agent in a meaningful energy game," (Isaza-Giraldo, *et al.*, 2024)
5	**CHI**	"LaMI: Large language models for multi-modal human-robot interaction," (Wang, *et al.*, 2024a)
5	**CHI**	"BioSpark: An end-to-end generative system for biological-analogical inspirations and ideation," (Kang, *et al.*, 2024)
5	**CHI**	"Empowering personalized learning through a conversation-based tutoring system with student modeling," (Park, *et*

		al., 2024)

3.1.1. Selection of popular discourse articles

We began corpus construction open-endedly, taking note of and saving public articles and press releases that we had encountered about applications marketed as “AI”, specifically “Something bothering you? Tell it to Woebot” (Brown, 2021) (news article), “Introducing Duolingo Max, a learning experience powered by GPT-4” (Duolingo Team, 2023) (company blog), and “Introducing Notion Al” (Zhao, 2022) (company blog).

Expanding the popular discourse corpus, we simply googled for “AI news”, and selected top articles from popular media which clearly presented a novel “AI” product or which were otherwise “AI”-focused. Our goal was to discover descriptions of “AI” systems that an average person would be likely to encounter. The additional texts were: “How to talk to an AI,” (Chen, *et al.*, 2023), “How the A.I. that drives ChatGPT will move into the physical world,” (Metz, 2024), “A.I. can write poetry, but it struggles with math,” *New York Times* (Lohr, 2024), “The AI emotions dreamed up by ChatGPT,” (Gorvett, 2023), and “Diagnostic robotics AI advances predictive, personalized medicine,” (Press, 2023). The publication of these articles spanned a year during which “AI” hype was at a notable high.

Finally, we included the OpenAI announcement “Creating video from text,” (OpenAI, 2024), which was launched during our investigation, because we wanted to include a text source which came directly from one of the companies releasing probabilistic automation models.

3.1.2. Selection of scientific articles

In addition, we contrasted promotional content with scientific papers from an HCI context. These were selected because they represent academic approaches to presenting probabilistic automation technologies, typically characterized by more precise and formal language. Selecting articles from outside the machine learning field was driven by a motivation to explore descriptions which were more likely to address a use context, where the system would be described in terms of what it could help an end user achieve (thus, more comparable to news articles and company blog posts).

The texts to represent the HCI context were chosen from the CHI 2024 conference proceedings and extended abstracts. We chose the CHI proceedings as a proxy for HCI research more broadly, because the venue is the most prestigious in the field of HCI, and is considered to be highly influential for HCI technology development in scientific and practical communities (Linxen, *et al.*, 2021).

We conducted a keyword search on the ACM Digital Library based on the names of almost 60 different generative “AI” models to ensure we would identify a broad sample of submissions utilizing probabilistic automation in system design. The keyword search resulted in 503 CHI 2024 publications, from which we arbitrarily (using a random number generator) chose 20 which met the criterion of presenting a system which comprised a probabilistic automation system as a core feature.

The goal of the selection was not to achieve generalizability for the entire CHI 2024 proceedings, but rather to investigate how a random sample of HCI papers focused on probabilistic automation technologies is likely to represent such systems. From the 20 selected CHI papers, we annotated the abstract and the introduction, based on the assumption that these sections are where the system is likely to be described at a level of abstraction most comparable to that of news articles and company blogs, and that they would be representative of language used to describe the system throughout the rest of the paper.

3.2. Annotation process

The annotation process was structured in iterative rounds. For each round of annotation, we individually annotated texts first, then discussed disagreements or discrepant categorizations after. After round one, we had agreed on four provisional categories and their description, based on a combination of categories presented in previous literature and our analysis: *Cognizer, Agent, Communication*, and *Biological metaphors*. From these categories, we moved to more deductive annotation, although with openness to allow new categories to emerge. From round two, we calculated inter-annotated-agreement (IAA) to support our discussions and as an indicator of whether categories and their definitions became more stable as the annotation rounds progressed. While we were not aiming for absolute agreement, we used the IAA calculation as a quantitative indication that the categories were progressively more reliable in new contexts (for novel texts).

We annotated the entirety of the texts in news articles and company blog posts, but only the abstract and introduction in the CHI texts, since the majority of variation in descriptive words were found in these parts.

3.2.1. Initial exploratory annotation (rounds one and two)

Our focus was on discovering anthropomorphization in *linguistic constructions*, examining what is implied and enforced by grammatical and semantic structure. We conducted independent exploratory annotations at the word level across a selection of popular discourse articles (Zhao, 2022; Duolingo Team, 2023; Brown, 2021; Press, 2023; Metz, 2024; Chen, *et al.*, 2023), and created initial anthropomorphization categories based on continuous discussions among all authors.

3.2.2. Third round of annotation and guidebook creation

In the third round of independent annotations we believed that the categories had started to reach saturation, so we created the first version of our annotation guidebook. In this round of annotation we included the first CHI abstracts and introductions, and three authors annotated three texts independently (OpenAI, 2024; Fok, *et al.*, 2024; Chen, *et al.*, 2024) before calculating Inter-Annotator Agreement (IAA). For this round we calculated IAA as a supplement for subsequent discussions used to update the guidebook and resolve annotation disagreements.

3.2.3. Fourth and fifth annotation

In the fourth and fifth round, pairs of annotators independently annotated each text, and we calculated IAA again to explore consistency of the application of categories. In order to increase the number of texts we could annotate, for these rounds, we moved to dual annotation.

3.2.4. Annotation software

Throughout the annotation rounds, we experimented with different software tools to optimize our workflow. We used Taguette [10] for the first, moved to Condens [11] for the second and third rounds, and finally used Label Studio [12] for the final rounds. Label Studio allowed the most optimal possibilities for exporting the annotations in a format where we could see the word in context of its sentence.

3.2.5. Handling disagreements

After each annotation round, all disagreements (meaning, either a word was labeled by only one or some of the annotators, but not others, or a word had been annotated with different categories by different annotators) were exported to a sheet showing the incongruity. During face-to-face meetings, all authors would discuss each instance and explore arguments for and against different annotation categories. The guidebook was updated during these discussions. When agreement was reached, the “gold standard” (*i.e.*, agreed category) would be noted for each instance. We did not revisit old texts and update annotations, and for this reason, we do not present the quantitative counts for each annotated text. Our intention in this work is to present the results and outcomes of our discussions.

3.2.6. IAA results

Annotation round one was inductive, and we did not calculate IAA. In annotation rounds two and three, three authors conducted annotations, and we measured inter-annotator agreement (IAA) using unitized Krippendorff's alpha (Letcher, 2018). The IAA values for rounds two and three were 0.629 and 0.878. For the fourth and fifth rounds, annotations were conducted in pairs, with IAA values (also measured with unitized Krippendorff's alpha) calculated for each pair and then averaged, resulting in an IAA of 0.907 and 0.842.

4. Results

Our annotation guidebook began with general guidelines for annotation and proceeded to describe each category of anthropomorphization along with examples found in the corpus.

4.1. General guidelines

We only annotated anthropomorphizing language that was predicated of systems portrayed as "AI" systems. We also specified that language could count as anthropomorphizing even if the property in question was negated. For example, saying that a system *failed to grasp a concept* still presupposed that it was expected to grasp it all, and thus *grasp* should be annotated (here, as an instance of Cognizer).

We aimed to take a reader's, rather than a writer's, point of view in annotating anthropomorphizing language. Thus, when sentences had more than one salient possible reading, where at least one reading was anthropomorphizing, we annotated it as such, even if this may not have been the author's intent.

In keeping with this reader's perspective on annotation, we annotated in order of text presentation. Sometimes, a particular turn of phrase became clearly anthropomorphizing further down a text, but was ambiguous in the introduction. For example, verbs like *foster* or *support* could refer to actions of people or properties of situations and organizations, as well as people. When they are applied to probabilistic automation systems, it isn't always immediately clear if the intended sense is meant to be analogous to a person doing the fostering or supporting. Later in a text, authors might provide further details which reveal that the sense intended is the one that is usually predicated of human actors. This did not lead us to go back to the start and reannotate earlier instances of those words as anthropomorphizing, because we assumed that the reader was consuming text in order. More generally, we aimed for conservatism in our annotations, and, when unsure, left the item in question unannotated.

We annotated at the word level, rather than the level of a sentence or a phrase. For example, in the sentence *computer scientists, it seems, have created artificial intelligence that is more liberal arts major than a numbers whiz* (Lohr, 2024), we annotated the words *intelligence* (Cognizer), *major* (Human role analogy), and *whiz* (Cognizer). This helped with inter-annotator agreement by removing uncertainty about the length of phrase to annotate, but sometimes became tricky with lexical items that were written as more than one word (such as verbs like *work out* or *pick up* [13]).

Some words might be annotated only in some contexts but not others (*e.g.*, *deliver* is Agent in *deliver advice* but not in *deliver returns on investment*, *gives* is Communication in *gives advice*, but not in *gives an advantage*). The verb *do* (and other similarly bleached verbs) especially need to be considered in context. What more contentful verbs are they standing in for? The consideration of "standing in" is relevant to more meaning-bearing words as well, *e.g.*, is *consistent* — what is the system consistent about? *Can help ensure* something — what is entailed in the helping or ensuring?

When uncertain about whether a word was anthropomorphizing or not, a strategy that we frequently used was to see if there were other uses of the word (in the same sense), unrelated to probabilistic automation systems and clearly not anthropomorphizing. For example, we debated whether talking about a probabilistic system as *absorbing information* was anthropomorphizing, but decided it was not, because clearly inanimate things like (artificial) sponges could absorb other things (liquids, in that case). We did identify metaphorical uses of *absorb* (as in a person said to be absorbing information), but decided on the grounds on the non-metaphorical use not to annotate this instance.

Finally, we make a distinction regarding product and company names, because only authors who bestow such names actually hold the responsibility for any anthropomorphization that they entail. Thus OpenAI is responsible for both the *AI* in their company name (an instance of Cognizer) and the *Chat* in ChatGPT (an instance of Communication), but a journalist or academic author writing about that company and its product is not making an anthropomorphizing choice by simply reusing the names that are already established [14].

4.2. Categories of anthropomorphization

In the following, we describe the details of each category as they were developed during our many discussions. We follow the structure of first describing the category with its definitions as it is characterized in our guidebook, then relating it to earlier works, and finally providing examples from our corpus and our discussions. In these examples, only words annotated with the specific category are highlighted for emphasis. The sentences may contain examples of anthropomorphization of other categories. The categories are summarized in Table 2.

Table 2: A taxonomy of categories of anthropomorphization by description.

Category	Description	Other references
Cognizer	Words portraying a system as having cognition (such as *intelligence*) or engaging in a cognitive activities such as *thinking, believing*, and *learning*.	Cheng, *et al.*, 2024; Ryazanov, *et al.*, 2025; DeVrio, *et al.*, 2025
Products of cognition	Expressions referring to things that can only be gained through cognitive activity, such as *skills, capabilities*, or *bias* (note: *reflecting* bias is different from *being* biased).	
Emotion	Words portraying a system as feeling, having emotions, or being able to form an emotional bond, *e.g.*, *struggling, empathizing*, or *caring*.	Cheng, *et al.*, 2024; Ryazanov, *et al.*, 2025; DeVrio, *et al.*, 2025
Communication	Words which frame the system as something that can participate in communication, both descriptions of system actions like *answering* or	Ryazanov, *et al.*, 2025; DeVrio, *et al.*, 2025; Shardlow and

	following *instructions*, as well as descriptions of system output like *explanation* or *suggestion*.	Piotr Przybyla, 2024
Agent	Expressions which portray the system as acting with intent or independence, such as *helping, facilitating*, and *leveraging*.	Ryazanov, *et al.*, 2025; DeVrio, *et al.*, 2025; Cheng, *et al.*, 2024; Hunger, 2023
Human role analogy	Words which portray the system as filling a human role such as *tutor* or *assistant*, or in constructions like *AI-doctor*.	Abercrombie, *et al.*, 2023; Ryazanov, *et al.*, 2025
Names and pronouns	Names which are foremost associated with people and animate pronouns (*I, we, you, he, she*, and certain uses of *they*).	Abercrombie, *et al.*, 2023; DeVrio, *et al.*, 2025
Biological metaphors	Words that represent the system as if it had a human body, such as *neurons, seeing, tired* and *consuming* (information).	Hunger, 2023; DeVrio, *et al.*, 2025

4.2.1. The Cognizer category

We annotate a word as an instance of **Cognizer** if it portrays the system in question as engaging in cognitive activity. This is inspired by the Cognizer Frame Element from FrameNet (Baker, *et al.*, 1998). The most straightforward examples of Cognizer are verbs, such as *think, believe, reason*, and *learn*, which assign this role to their subjects. However, this category can also apply to other parts of speech, such as the noun *intelligence* in the phrase *artificial intelligence* (or *AI*), the noun *whiz* as in describing a chatbot as a *math whiz*, or participles like *training* in *training data*.

The category maps well to the descriptions presented in other taxonomies, for instance “the ability to think and hold beliefs” (Cheng, *et al.*, 2024), or “having cognitive abilities beyond describing its formal tasks” (Ryazanov, *et al.*, 2025).

We anticipate that some might object to including the term *artificial intelligence* in this category, given that is well-established (coined as the name for the area of study by McCarthy, Minsky, *et al.* in a 1955 grant proposal). However, to exclude it would be to obscure the effects of choosing to use it/choosing not to resist it. The term *artificial intelligence* may even be particularly problematic, given some research has shown that end users associate high machine “competence” with this term compared to, *e.g.*, “decision support systems”, “sophisticated statistical models”, or even “machine learning” (Langer, *et al.*, 2022).

Examples from the corpus:

- `Its goal is to help robots gain an understanding of what is going on around them` and decide what they should do next (Metz, 2024)

- we're teaching AI to understand and simulate the physical world in motion (OpenAI, 2024)
- the model may also confuse spatial details of (OpenAI, 2024)
- its analysis is based on 60 billion claims or patient visits. The knowledge gained from this analysis enables it (Press, 2023)
- an intelligent writing support system (Weber, *et al.*, 2024)

Examples of discussions. The word *training* as in *training data* being used to create and fine-tune algorithms was a focus of extensive discussion. We debated whether the word should belong under the Cognizer or the Biological Metaphors category, since training can mean physical training. Furthermore, a plant can be *trained* to follow the pattern of a trellis — this sense is not anthropomorphizing. Ultimately, in the context a *system learns patterns from its training data*, we annotated both the word *learn* and the word *training* as examples of the Cognizer category, because *training* in the sense of learning from data requires cognition.

We also encountered and discussed the term *have foresight*, as in the *model has foresight*. While the word is clearly anthropomorphizing, we discussed whether it belonged in the Cognizer or the Products of cognition category. In the example we saw, the *foresight* was not based on previous experience, but based on data from the future: By *giving the model foresight of many frames at a time* (OpenAI, 2024). It therefore does not fall into the category of skills one can only acquire through cognitive practice (which is the determining factor for the Products of cognition category), and was ultimately annotated as Cognizer.

4.2.2. The Products of cognition category

Closely related to the Cognizer category but subtly distinct is the category of **Products of cognition**. These are expressions that refer to something that the system supposedly has, but could only have gained through cognitive activity, such as *skills, capabilities*, or a *sense of aesthetics*. The first two are anthropomorphizing descriptions of what would be better described as *functionalities*. That is, they describe a potentially real property of the system, but do so misleadingly by drawing on metaphors of learning in people. With sense of aesthetics the anthropomorphizing language is even more misleading, locating within the system something that actually belongs to the people involved: those curating the training data and those perceiving system output.

Depending on the context, expressions such as *cultural bias* fit here as well: a system that *reflects* cultural bias isn't anthropomorphized, but one that *has* cultural bias, as a result of *learning* it from its *training* data is. We also included *mistakes* in *make mistakes* in this category. A non-anthropomorphizing description would say that a system *produces errors*. To say instead that it *makes mistakes* is to portray its output as the result of (faulty) cognition.

This distinction between Cognizer and Products of cognition has not been made in any previous categorizations that we have reviewed as part of our search for related work.

Examples from the corpus:

- what if the AI makes a mistake? (Duolingo Team, 2023)
- the model has a deep understanding of language, enabling it to (OpenAI, 2024)
- if chatbots show cultural biases (Chen, *et al.*, 2023)

`chatbot prototypes [...] have limited capability of` (Kim, *et al.*, 2024)

- `necessary to tailor LLMs' capacities to a younger audience` (Zhang, *et al.*, 2024)

Examples of discussions. We encountered the word *skills* in the sentence *designing technology that lets robots learn skills much like chatbots do* (Metz, 2024). Even if the skill is physical, such as being able to juggle or wiggle one's ears, the word *skill* conveys the deliberate practice to have acquired that skill. We agreed that acquiring skills requires cognitive processes, but that the skills one masters well may no longer require active cognition to carry out. Therefore, we decided that *skills* and *abilities* do not belong to the Cognizer category. This was the inception of the Products of cognition-category.

An expression that we discussed in this category was to *grasp* as in a *system can grasp the relationships between the two* (Metz, 2024) or *The large language models, Dr. LeCun has said, have little grasp of logic* (Lohr, 2024). *Grasp*, in this sense, is already a biological metaphor, because it refers to the ability to take and hold something in one's hands. In these sentences, however, we agreed that the word refers to the mental state of someone and not being able to physically grasp something.

We further agreed that *grasp*, in the first sentence, is a *process* of cognition, portraying the model as understanding an abstract relationships between two things — and we annotated this as Cognizer. In the second sentence, *having little grasp* refers to a *resulting* mental state, an ability that has (not) been acquired, and we therefore annotated this as Products of cognition.

4.2.3. The Emotion category

While it is more common to portray probabilistic automation systems as thinking, there are also cases where they are portrayed as feeling. We separate out this category as **Emotion**, and include both direct emotion words (*feel, emotions, empathize*), but also words that entail an emotional perspective (*e.g., struggle*) or the ability to form an emotional bond (*friend, companion*). There are also verbs of communication that describe playing on the addressee's emotions, such as *chide* or *coax*.

This category was also used by Ryazanov, *et al.* (2025), in their annotation work: "AI presented as having emotions, or taking emotionally charged actions" [15], and by Cheng, *et al.* (2024) in their definition as "[having] the ability to experience emotion and feel pain (affective mental states)" [16].

Examples from the corpus:

- `It may struggle with accurately simulating the physics of` (OpenAI, 2024)
- `a shady young chatbot with a whimsical fondness for penguins` (Gorvett, 2023)
- `it's probably an emotional response and not a logical one` (Gorvett, 2023)
- `a well-designed bot can form an empathetic, therapeutic bond with its users` (Brown, 2021)
- `if it's acting like it has thoughts and feelings` (Chen, *et al.*, 2023)

Examples of discussions. Initially, we discussed whether this category should be labeled Emotion or Relational, because the words we encountered were often describing an interpersonal relationship between a model and a human, for example, *a well-designed bot can form an empathetic, therapeutic bond with its users, Dr. Darcy calls it, a "relational agent"* (Brown, 2021), and *the chatbot would make a highly companionable sentient being* (Gorvett, 2023). We discussed that relational words in this sense describe a reciprocal emotional aspect—people typically call someone a friend or companion because of how they feel about them.

In contrast, someone can be a tutor or a tour guide regardless of how one feels about them. Ultimately, we decided that *tutor* and *tour guide* belonged to the Human role analogy category, but where the relationship was one based on emotion (*e.g., friend*) we grouped those with other words that directly named emotional experiences.

Another word which came up for discussion in this category was *coax* in *Dan is a roguish persona that can be coaxed out of ChatGPT* (Gorvett, 2023). The act of coaxing is certainly a communicative act and could be annotated with Communication. When *coaxing* or *chiding* someone, we agreed that the communicator was playing on the emotions of the addressee when trying to persuade them to do something. The coaxer was not just telling someone to do something, but *causing* them to do it by manipulating them psychologically.

We thought of *criticizing* was an example of scolding which belongs in the Communicator category (because a person can criticize someone without it having an emotional effect on the other), but agreed that *coaxing* should be annotated as Emotion.

4.2.4. The Communication category

Anthropomorphization of the **Communication** type frames the system as something that can participate in some form of communication or conversation like a person. These words are frequently but not always verbs, though the system does not have to show up as the subject or speaker. That is, we annotate a word as signifying this type of anthropomorphization if system is set up as doing the speaking (*answer, tell*) but also if it is the addressee (*ask, instruct*).

There are also nouns in this category, such as a system providing an *explanation* or following *instructions* or *prompts*. Finally, there are system names that invoke this anthropomorphizing type, including generic names like *chatbot* but also the names of specific systems such as *ChatGPT* or *PromptCharm*.

Ryazanov, *et al.* (2025) described the anthropomorphic properties of communication terms, including a category of “Communicating and producing natural language.” DeVrio, *et al.* (2025) also described “Expressions of language manipulation” and “Expressions of (dis)agreeableness” in the output of language models as examples of “Communication skills” that contribute to anthropomorphization of a language model, although looking for text output which mimics such “skills” constitutes analysis on a different level than identifying descriptive text which directly claims that a system has them.

In other work, this category was often merged or partially merged with a category of agency/intentionality, as in “it is inaccurate to describe the model as ‘asking’ the question. Asking implies an intentionality that the model itself is not capable of” (Shardlow and Przybyla, 2024).

Examples from the corpus:

- `when I ask it what kinds of emotions` (Gorvett, 2023)
- `review their patients' dialogue history with the AI` (Kim, *et al.*, 2024)
- `enables it to suggest next clinical steps` (Press, 2023)
- `learners can enter a chat with Duo` (Duolingo Team, 2023)
- `we introduce ChatScratch, an AI-augmented system` (Chen, *et al.*, 2024)
- `we examine thinking and usage patterns of CharacterMeet, a prototype system` (Qin, *et al.*, 2024)

Examples of discussions. One of the most commonly used words in the context of interactions with generative “AI” is *prompting*, as in providing input to the system. We discussed whether prompting could be considered similar to the non-anthropomorphizing input, a non-laden set of instructions that adds specific steps in an algorithm. However, *prompting* is to *move [someone] to action*, which plays on the recipient’s internal state as an intermediary, either by pushing them towards an action or reminding them of it: *prompt someone to do something* or *do something without prompting*. We therefore generally annotated it as Communication. However, if the output of a language model *prompted* action by humans, this would not be annotated as anthropomorphizing because this did not entail any communicative intent or internal state on the side of the automated system.

A word that was also discussed and which we eventually placed in the category of Communication was *meet*, as in the system name CharacterMeet (Qin, *et al.*, 2024). The word *meeting* has notions from several of the categories. We argued first that it might require physical presence (and might contain an element of biological metaphors) based on the fact that when we have a meeting, we usually prefix it with the word *virtual* or *online* if the meeting is not happening with physical presence. One can have had hundreds of Zoom meetings and e-mail correspondences with a person and still say “I met them for the first time yesterday” if yesterday’s meeting happened in physical space.

However, *meeting* does not directly make use of any physical or biological metaphors. A musician can have been in physical presence of thousands of concert attendees in a crowd, yet we would not say that the artist met any of the individuals in the crowd, neither would members of the audience remark that they actually met the artist. We finally settled on Communication based on the assertion that meeting entails a reciprocal acknowledgment from both parties. These discussions highlighted the anthropomorphizing aspect of words implying *communication* by underscoring that communication is not just the passive transmission of information, but involves reasoning about the other person and their internal state.

4.2.5. The Agent category

The Agent category refers to expressions that portray the system as acting with intent, either *mechanistic* (the ability to take action based on some information) or *volitional* (actively making decisions in accordance with internal desires) (Dai, 2024). In some cases, the action named is clearly agentive (*e.g.*, *operate* if in the sentence *the chatbot operated the device*). In others, it is modifiers that make the agency clear, such as *on its own* or *intentionally*. In such cases, we annotate both the modifier and the verb or other word expressing the action.

Sometimes the intention is subtle or negated, but still an intention. For example, *drops* in *The robot makes mistakes, and drops things from time to time* is an example of Agent, not because the robot is portrayed as trying to drop the thing but because it is portrayed as trying not to. In constructions like *help/assist/aid*, or *support* (which are very commonly found in our corpus), we placed emphasis on whether the technology was presented as *taking actions* that help the person (anthropomorphization), or whether its existence helped the user accomplish something (not anthropomorphization).

Though it didn’t appear in our corpus, the word *agent* to denote any software systems would also fall into this category.

The Agent category overlaps with the Cognizer and Communication categories, and we resolved this overlap by choosing Cognizer or Communication if they apply and only using Agent in the remaining cases. We note that there was not a complete overlap here, that would allow us to make Cognizer and Communication subcategories of Agent, however: for example, when the system was an addressee (*e.g.*, of *ask*), it was not portrayed as an agent but was portrayed as an entity that could participate in communication.

The Agent category is similar to the one presented by Cheng, *et al.* (2024) in “act and produce an effect on their environment (behavioral potential)”, to the definition “verbs of intentional action” (Hunger, 2023), and to the description “<Model> uses <Technique> to” (Shardlow and Przybyla, 2024).

Examples from the corpus:

- `Marco joins the foraged information into an aggregate table, aiding sensemaking.` (Fok, *et al.*, 2024)
- `ChatScratch [...] leverages Scratch-specialized Large Language Models (LLMs)` (Chen, *et al.*, 2024)
- `Sora is an AI model that can create realistic and imaginative scenes` (OpenAI, 2024)
- `on the basis of its analysis of historical medical data, and assisted by generative AI` (Press, 2023)
- `it can't do everything, but it can help you be more productive` (Zhao, 2022)
- `iteratively prompting them to curate an everincreasing dataset` (Kang, *et al.*, 2024)
- `which could allow them to impersonate diverse characters` (Qin, *et al.*, 2024)

Examples of discussions. A lot of our discussion of the Agent category revolved around the overlaps with Cognizer and Communication noted above. One further example of the overlaps between Agent and other categories is the word *analyze*, as in *the system analyzes data*. This is clearly agentive, but the cognition implied is the stronger argument for why the word is anthropomorphizing.

Overall, Agent was one of the most debated categories in the annotation process. Outside of the overlap issues noted above, there were also plenty of questions about words that ambiguously suggested agency. In such cases, we usually resolved these debates by attempting to find examples of the word in a context where it was used in a non-agentive way. For example *absorb* in *It generates language, based on all the information it has absorbed* (Lohr, 2024) was originally annotated as Agent, but later agreed upon as None, based on the argument that a sponge can absorb water, but it does not do so with intention.

Scored in *A recent version of the technology that underlies ChatGPT scored in the 89th percentile in the math SAT test* (Lohr, 2024) was also first annotated as Agent, but we later settled on None, because one could imagine a wine (passively) scoring a high number in a taste evaluation.

4.2.6. The Human role analogy category

The **Human role analogy** category involves turns of phrase that describe a system as filling a human role. This includes things like *tutor* or *assistant* which directly name the role, as well as words that imply the role (and place the anthropomorphized system in it), such as *mentoring* or *advice* and *co-creation*.

The attribution of a human role to a probabilistic automation system was also described by Ryazanov, *et al.* (2025) in “AI taking human roles,” and by Abercrombie, *et al.* (2023) in “the roles that dialogue systems are given, consciously and unconsciously, by their designers and users. Many of these can shift dialogue systems from the realm of tools towards one of humanlike roles such as provision of companionship.”

Examples from the corpus:

- `... an automated therapist when finding a real one can feel ...` (Brown, 2021)
- `[product] is your teammate before, during, and after the writing process` (Zhao, 2022)
- `communicate their information needs to an AI assistant` (Fok, *et al.*, 2024)
- `act as your hawk-eyed editor` (Zhao, 2022)

- `LegalWriter: An Intelligent Writing Support System` (Weber, *et al.*, 2024)
- `they converse with the systems about their characters or roleplay with them` (Qin, *et al.*, 2024)

Examples of discussions. The most subtle word that came up as a candidate in this category was *assistance*. We determined that *assistance* was not clearly something that only *assistants* provided (especially within the field of human-computer interaction, where *assistive technology* entails technologies which increase, maintain, or improve the functional capabilities of persons with disabilities — see Agent above), unlike *mentoring*, which can only come from a *mentor*, and so did not annotate *assistance* as human role analogy unless a system was described explicitly as an *assistant*.

We also discussed how to handle direct comparisons, if the author would write, for example *like a human assistant*, or *similar to how a human might*. We agreed that since these comparisons are clearly flagged as comparisons, they are not anthropomorphizing (although they contain false claims of functionality).

4.2.7. The Names and pronouns category

Names that are associated first and foremost as human names are anthropomorphizing, if applied to inanimate objects such as computer systems. However, which names are names of people is culturally specific, which means annotation will depend, to some degree, on the context in which the name is written and who the readers of the text are (see discussion examples below).

We also added anthropomorphizing uses of **pronouns** to this category. Pronouns, like names, are used to refer to but not describe an entity. But animate pronouns (*e.g.*, *she, he*, and singular *they* in English) used to refer to inanimate things anthropomorphize their referents just as names do. More subtly, first person pronouns in quoted text from machines and group pronouns like *you* used in a way that includes both people and machine fall into this category as well.

While not a commonly found category in other categorizations of anthropomorphizing language, Abercrombie, *at al.* (2023) wrote about pronouns that "The use of first person pronouns (*e.g.*, 'me' or 'myself') in system output may be a contributing factor [...] as these can be read as signs of consciousness [...] Indeed, it is widely believed that 'I' can only refer to people [...] such self-attribution and self-reference permits people to relate as subjects, not mere objects, and that self-definition as an individual is part of the human condition itself." (Abercrombie, *at al.* 2023)

Examples from the corpus:

- `we introduce Marco, a mixed initiative workspace` (Fok, *et al.*, 2024)
- `compare your answer to the bot's. Who's right?` (Lohr, 2024)
- `How did each of you arrive at your solution?` (Lohr, 2024)
- `"I wouldn't want to overwhelm your puny human brain with my brilliance!" (quoted system output)` (Gorvett, 2023)

Examples of discussions. The main discussion we had about this category was when to call something a name, since this is culturally specific. In our analysis we annotated names that we, from an English-speaking context, perceived as human names. For instance, we initially annotated *Marco*, but not *Sora* (OpenAI, 2024). After later discussion, we agreed that *Sora* (which is a common name in Japan) would be perceived as a name by many readers, and should be annotated. However, it is difficult to know all names in all cultures, so we decided the best practice would be to be transparent about which names were annotated and why.

The pronoun *I* is typically a feature of system output, rather than description of systems. However, it comes in-scope for anthropomorphization by description when a journalist or researcher has made a deliberate choice to include this output in their exposition, including the pronoun. In that context it falls within our categories, specifically the Names and pronouns category.

4.2.8. The Biological metaphors category

The category of **Biological metaphors** is used when the anthropomorphizing language describes algorithms as though they had biological properties. Some of these metaphors are well-established technical terms such as *neural* in *neural nets*, but they also include one-off uses such as describing as system that makes errors as *stumbling*, a robot working *tirelessly*, an algorithm *consuming* or *digesting* data, or a chatbot *embodying* a character.

Biological metaphors as a category is directly reproduced from Hunger (2023) but, interestingly, not found in many other categorizations or descriptions of anthropomorphizing language. One possible reason for this could be that the notion of a probabilistic automation system having a physical body is so absurd that it is not recognized as a “risky” class of anthropomorphism.

Another is that examples of biological metaphors are not distinguished as a separate category because biological metaphors are often used in a metaphorical sense in the first place (such as in *grasping a concept* or *digesting information*) or that biological metaphors are less common in text descriptions than those which compare such systems to thinking and intentionality.

Examples from the corpus:

- `engineers can improve the technology by feeding it more and more data` (Metz, 2024)
- `identifying patterns in this stew of images, sensory data and text, the technology gives a robot the power to` (Metz, 2024)
- `they’ve ingested some of the ugliest material the Internet has to offer` (Chen, *et al.*, 2023)
- `the advanced neural networks, known as large language models` (Lohr, 2024)
- `act as your hawk-eyed editor` (Zhao, 2022)

Examples of discussions. This category generally did not incur many disagreements, since it referred to the human body which has fairly clear properties. We sometimes discussed expressions that make metaphorical use of physical properties such as *grasp* in the sense described in the Products of cognition section above, and agreed that when used in a metaphorical sense, the word would not be annotated as a Biological metaphor, but in the category of property that the metaphor is intended to describe.

Another example is *embody* in *a chatbot embodying the character* (Qin, *et al.*, 2024), which is subtle because *embody* is already used as a metaphor for representing an idea or characteristic. We annotated *embody* as Biological metaphor, however, because the metaphor played on the property of having a body which can represent something else.

Of more subtle examples we will mention the word *tireless* in *Computers have been tireless, fast, accurate calculating machines* (Lohr, 2024), which we annotated as Biological metaphor even though it is a negation, because it refers to the physiological property of being tired.

We discussed how to handle the descriptions of physical properties of robots with anthropomorphic characteristics. If a robot of clearly anthropomorphic design was described as “shaking its head”, this could be considered a flat report of anthropomorphization by design. Here, the writer should consider whether they described the robot as having a head because it looked that way to the writer, or if it was likely to be perceived that way by everyone. Then we discussed whether *shaking its head* itself was anthropomorphizing because a head shaking is a communicatively loaded action, in many cultures perceived as either negating something or showing despair or disbelief. When someone "shakes their head", rarely are they actually *shaking* their head, but rather moving it from side to side in a paced manner. To draw this example to the extreme, we then decided that the least anthropomorphizing way of describing such an action would be to say that *the robot's top piece was moving from side to side*.

4.3. Summary

The categories described here are what emerged from our annotation of 29 texts written in 2022–2024 and we believe them to be fairly likely to be sufficiently exhaustive set of categories for English language discourse in CHI papers, company blogs, and popular media coverage of "AI" systems. This belief is supported by the fact that we didn't find the need to create new categories after round four of annotations.

However, further annotation of different text types, different time periods and especially different languages might turn up additional categories.

Our purpose in describing our categories at this level of detail, along with examples of our discussions, is to help others in identifying these types of language use in their own and others' writing. Such identification is an important first step towards de-anthropomorphization, as discussed in the next section.

5. Discussion

5.1. De-anthropomorphizing language and the functionality-first principle

With this close linguistic analysis, we aimed to move beyond description and towards answering the question of what to do in place of anthropomorphic descriptions. We generally argued that writers describing probabilistic automation systems should aim to apply a ***functionality-first principle***, **where the goal is to explain the system in terms of what we can use it to achieve** rather than what "capabilities" we posit it has.

The problematic nature of wishful mnemonics or the term itself are by no means a novel critique in the field of computing. It was in fact raised half a century ago by McDermott (1976) about the descriptions of "artificial intelligence" when the field was relatively new:

> A major source of simple-mindedness in AI programs is the use of mnemonics like "UNDERSTAND" or "GOAL" to refer to programs and data structures. ... If a researcher ... calls the main loop of his program "UNDERSTAND," he is (until proven innocent) merely begging the question. He may mislead a lot of people, most prominently himself ... What he should do instead is refer to this main loop as "G0034," and see if he can convince himself or anyone else that G0034 implements some part of understanding. ... Many instructive examples of wishful mnemonics by AI researchers come to mind once you see the point" [17]

Our suggestions for alternatives to anthropomorphizing language are presented in Table 3. For two categories, Emotion and Human role analogy, we found that there were no simple replacements. Rather, when faced with such anthropomorphizing language, the best bet was a larger reframing of the matter at hand. For the others, we could often find simple edits within a sentence, as exemplified in the Table 3.

Table 3: Suggestions for *de*-anthropomorphizing language.		
Category	**Alternative strategies**	**Examples**
Cognizer & Products of cognition	Instead of language that locates thinking in an algorithm, describe the algorithm as performing calculations or other algorithmic operations and the people using it doing the thinking.	*artificial intelligence → probabilistic automation* *hybrid intelligence → augmented human intelligence* *image recognition → image labeling* *speech recognition → automatic transcription* *the model shows bias → the model reflects bias* *model mistakes → model errors* *chatbots are good at ... → chatbots are good for ....*
Emotion	Ascribing emotions to computers, even speculatively or subtly, is not scientific. Should this arise in scholarly writing, it is worth looking at how the rhetorical goal could be achieved in some other way, or if it should in fact be abandoned.	
Communication	Instead of verbs like *ask, say, inform, discuss*, use verbs appropriate to computers like *input* and *output*. Another strategy is to foreground the fact of	*prompt → text input* *answer → output* *chatbot/conversational agent → conversation simulator*

	simulation.	
Agent	Turns of phrase that locate agency with a machine often serve to obfuscate the interests and goals of people. We suggest revising to locate agency with people or choosing less agentive verbs.	*ChatGPT assisted students → the students used ChatGPT* *revealing the solution → displaying the solution*
Human role analogy	Calling systems *tutor* or *co-creator* are overclaims (Rehak, 2021) that describe what a developer might wish they could develop — for those who want to replace people in these roles.	Language that describes algorithms as tools that people use, rather than as human-like entities, more clearly indicates system functionality while also not telegraphing a plan to replace people.
Names and pronouns	When naming a system, be mindful of the fact that all writing about that system will invoke the name. The act of bestowing a human name on a system entails anthropomorphizing it not only in the initial naming but every time the system will be discussed in the future. Pronouns, on the other hand, are a choice that can be made separately each time, but the choice can in some cases be subtle, *e.g.*, pronouns that group algorithms and people together under one plural form like *you* or *them*. Breaking the system out from the people and not using the collective pronoun is the better choice from the perspective of reducing anthropomorphization.	*who's right? → is the machine output correct?* *they produce results → the team uses it [the system] to produce results*

Biological metaphors	In revising away from biological metaphors, it is worth asking how system functionality can be more precisely described, giving readers a clearer sense of what is actually happening.	*neural networks → weighted networks* (Hunger, 2023) *the model consumes data → data is used in setting model weights.*

If other explanations fail, using clearly marked analogies could serve an explanatory purpose, provided they are appropriate to the concepts that they are meant to clarify. The following example could perhaps help lay readers understand the stochastic nature of probabilistic automation technologies and how their data processing algorithms arrive at final model weights, without anthropomorphizing uses of words like *learn* or *train*. Furthermore, the phrase *one can compare this to* flags the metaphor, without saying that the process substantively similar:

> “The algorithm is designed to produce a series of representations through trial and error. One can compare this to how ants discover food sources by crawling around seemingly randomly, until they find something, leaving different secretions when they do, in order to preserve the information.”

The impact of linguistic anthropomorphization of probabilistic automation systems is difficult to assess fully and we hope that the detailed examples identified and suggested in this work will make direct comparison studies more straightforward to design. On the other hand, there are few strong arguments *for* anthropomorphizing probabilistic automation systems outside the explanatory power of metaphors and analogies.

5.2. A note about standardized anthropomorphic terms

Many terms used to describe probabilistic automation systems are both deeply anthropomorphizing and highly entrenched in academic and popular discourse, *e.g.*, *artificial intelligence (AI), machine learning*, and *neural networks*. Given that these are established terms of art, it may seem not worth the effort to avoid using them.

However, we argue that the effort involved in taking issue with the terms (in print) and then taking distance from them (either with scare quotes or with rephrasing) is effort well spent: If the writing is taking issue with *e.g.*, the framing of a task (such as image or speech *recognition* [Cognizer]), then a renaming will be an effective move for highlighting questionable claims. Even if the writing is taking a task as given and presenting a new approach to it, a better term will make the actual claims of the paper more realistic and clearer.

Habits of avoiding anthropomorphizing language can help us elsewhere in the research process as well: We can use them when we work to define research ideas, keeping focus on plausible applications that neither oversell technology nor dehumanize the people involved (Bender, 2024; Sarkar, 2023). Developing these habits is also useful to us as readers, as they will position us to more quickly recognize anthropomorphizing language in texts that we encounter, including both science that we wish to build upon and texts that we are meant to evaluate, such as grant applications.

For authors of research papers striving to write in ways that accurately reflect research results while not contributing to misleading hype around AI, a taxonomy such as ours can be an effective tool for identifying anthropomorphizing language and for finding alternatives.

It takes effort to swim upstream against anthropomorphizing language embedded in commonly-used technical terms and ways of speaking about research topics, both in recognizing them at all but also in finding suitable alternatives. Researchers, as experts, have an important role in modeling and propagating productive language use around these technologies, and we suggest to think of our overall motivation as creating *informative* and *empowering* metaphors, rather than using language that sneaks in unsupporteed presuppositions or otherwise helps overstate the case.

5.3. Our contribution

While our taxonomy has overlaps with previous work (*e.g.*, Inie, *et al.*, 2024; Shardlow and Przybyla, 2024; Hunger, 2023; Ryazanov, *et al.*, 2025) it is, to the best of our knowledge, the first taxonomy based on a word-level linguistic analysis of multi-genre texts *describing* “AI”, rather than of text outputs from language models. This analysis facilitates our moving beyond descriptive analysis and towards *prescriptive* answers to the question “What do we do instead?” When we identify exactly which words are anthropomorphizing and why, we can create informed and substantiated alternatives.

This is not only useful to researchers who analyze existing descriptive texts, but also to researchers who wish to conduct comparative empirical studies of the impact of anthropomorphizing/non-anthropomorphizing language, and who have previously relied on ad hoc definitions rather than a shared vocabulary.

Different categories of anthropomorphization are likely to have different effects on different people (Inie, *et al.*, 2024; Langer, *et al.*, 2022; Seeger, *et al.*, 2021), and we believe research into this phenomenon becomes more nuanced if we move beyond simply claiming “anthropomorphization” — as is also evident from earlier works who have distinguished between degrees of anthropomorphization, such as “ambiguous anthropomorphism/explicit anthropomorphism” (Shardlow and Przybyla, 2024) and “established anthropomorphism/task-based anthropomorphism/high anthropomorphism” (Ryazanov, *et al.*, 2025). Even skeptics of the risks of anthropomorphism highlight that getting more specific and deliberate about the concept is central to understanding human behavior and relations to machines (Coghlan, 2024).

5.4. Limitations

The main limitation to our study is a relatively narrow corpus and its selection. For each text we annotated, we encountered expressions that required subsequent discussion — even though we reached a high IAA. News articles and company blog posts were selected opportunistically, meaning they may not be representative of all public discourse. It is possible that a larger corpus or a different selection of texts might have yielded different categories. Furthermore, since we only annotated texts written in English, it is possible that other languages would suggest different categorizations. However, although each new text we annotated inspired new discussions, the last rounds of annotation discussions were more often centered around which category a word should belong to, rather than inspiring novel categories. The categories were exhaustive for all words we annotated during this work. Since the goal of this work was to provide a substantial foundation for further analysis, we believe this limitation is acceptable — the categories could be expanded upon with further analysis.

Twenty selected articles from the CHI 2024 proceedings were used as proxy for “scientific communication”, and it is possible that the taxonomy could be expanded via an analysis of articles from other domains. We expect that the categories that we identified to still apply across different academic disciplines, although different domains are likely to show different distributions of categories — we welcome further evaluation, and highlight again that the taxonomy and its categories are open to expansion.

Our intended future work includes verification of the different categories and their impact: does a reader change their perception of a system based on anthropomorphization in the different forms (some research suggests this; Inie, *et al.*, 2024), and are the categories equally impactful — or more impactful in combination?

Further, we plan to explore the effect of presenting people with the categories to cultivate the habit of spotting anthropomorphizing language in context, and evaluating people's subsequent assumptions and expectations of systems that they have only read a description about.

6. Conclusion

This paper presented the result of a long-standing collaboration between the authors, who reside in human-computer interaction and computational linguistics. The collaboration had the goal of defining anthropomorphizing language as it is being used to describe probabilistic automation systems. From a corpus of 29 news articles, company blog posts, and scientific abstracts/introductions we identified eight categories of anthropomorphizing language: *Cognizer, Products of cognition, Emotion, Communication, Agent, Human role analogy, Names and pronouns*, and *Biological metaphors*.

We used these to suggest concrete ways to de-anthropomorphize language when describing probabilistic automation systems based on a *functionality-first* principle, where the system should primarily be described in terms of what we can use it to achieve, rather than its so-called "capabilities". We hope that this taxonomy and the strategies will be useful to researchers who conduct empirical studies of the impact of anthropomorphizing language as well as to authors writing about these technologies in both academic and non-academic contexts.

Our work diverges from prior taxonomies in three key ways: First, we provide word-level linguistic analysis across multiple genres of descriptive texts about "AI" systems, rather than analyzing system outputs. Second, we move beyond descriptive analysis to offer prescriptive alternatives as concrete substitute language that writers can use. Third, our *functionality-first* principle provides an actionable framework for revision.

This taxonomy and the alternative language strategies enable several lines of future research. Researchers can now conduct comparative studies examining whether specific categories of anthropomorphization have distinct effects on different populations. Does Emotion language impact trust differently than Agent language? Are some categories more impactful in combination?

Our taxonomy also enables studies examining whether teaching people to recognize anthropomorphizing language helps them develop more critical reading practices, and whether our proposed alternatives successfully reduce misplaced trust without sacrificing comprehensibility.

For researchers conducting content analysis, our taxonomy provides standardized definitions enabling comparable findings across studies. For those designing human-computer interaction experiments, it facilitates controlled language manipulation. For science communicators and journalists, the taxonomy offers diagnostic tools and specific revision strategies.

About the authors

Nanna Inie is an assistant professor in Human-Computer Interaction at the IT-University of Copenhagen. Her research focuses on the cognitive impact of technology, especially technology that we use to acquire and manage knowledge. Her latest projects explore the cognitive and environmental costs of generative "AI", their effect on critical thinking, and how to hack LLMs to foster greater computational literacy.
Contact information at https://www.nannainie.com/.
E-mail: nans [at] itu [dot] dk

Peter Zukerman is a researcher in natural language processing and human-computer interaction. His work focuses on semantic similarity clustering, low-resource data augmentation for search and rescue, and evaluating large language models for tool use in complex tasks. He also studies Japanese language and cross-cultural communication.
E-mail: pzuk [at] uw [dot] edu

Emily M. Bender is a professor of linguistics and an adjunct professor in the School of Computer Science and the Information School at the University of Washington, where she has been on the faculty since 2003. Her research interests include multilingual grammar engineering, computational semantics, and the societal impacts of language technology. Contact information at https://faculty.washington.edu/ebender/.

Notes

1. The term “artificial intelligence” is poorly defined, and does not refer to a coherent set of technologies. In general, we find that discussions of technologies called “AI” become more lucid and thus productive when we speak about the automation of specific tasks by performing statistical analyses of large datasets. We will refer to technologies sold as “AI” collectively as “probabilistic automation” (Inie, *et al.*, 2024). These technologies include, but are not limited to, large language models, other generative “AI” systems, and other applications of machine “learning” models.

2. Pataranutaporn, *et al.*, 2023, p. 1,081.

3. Kueffer and Larson, 2014, p. 723.

4. Peverini, 2024, p. 104.

5. Pareek, *et al.*, 2025, p. 13.

6. https://replika.com/.

7. https://openai.com/index/chatgpt/.

8. Cheng, *et al.*, 2024, p. 25,924.

9. Tipler and Ruscher, 2014, p. 219.

10. https://www.taguette.org/.

11. https://condens.io/.

12. https://labelstud.io/.

13. These are called verb-particle constructions in English, and are always written as two separate words, but they are idiosyncratic in that the verb sense depends on the particular choice of particle. Compare: *pick up, pick on, pick out*. In the closely related language German, when such verbs occur in sentence-final position, the particle attaches as a prefix.

14. There is still an anthropomorphizing effect, and in keeping with the heuristic of taking the reader’s point of view, an argument could be made for annotating all such instances. However, this seemed counterproductive, given that it’s generally not possible to revise away from product and company names.

15. Ryazanov, *et al.*, 2025, Suppl.Inf., p. 6.

16. Cheng, *et al.*, 2024, p. 2.

17. McDermott, 1976, p. 4.

References

Gavin Abercrombie, Amanda Cercas Curry, Tanvi Dinkar, and Zeerak Talat, 2023. “Mirages: On anthropomorphism in dialogue systems,” *arXiv*:2305.09800 (16 May).
doi: https://doi.org/10.48550/arXiv.2305.09800, accessed 19 January 2026.

Christine Anderl, Stefanie H. Klein, Büsra Sarigül, Frank M. Schneider, Junyi Han, Paul L. Fiedler, and Sonja Utz, 2024. “Conversational presentation mode increases credibility judgements during information search with ChatGPT,” *Nature Scientific Reports* volume 14, number 1 (25 July), article number 17127.
doi: https://doi.org/10.1038/s41598-024-67829-6, accessed 19 January 2026.

Collin F. Baker, Charles J. Fillmore, and John B. Lowe, 1998. “The Berkeley FrameNet Project,” *36th Annual Meeting of the Association for Computational Linguistics and 17th International Conference on Computational Linguistics*, volume 1, pp.86–90.
doi: https://doi.org/10.3115/980845.980860, accessed 19 January 2026.

Emily M. Bender, 2024. “Resisting dehumanization in the age of ‘AI’,” *Current Directions in Psychological Science*, volume 33, number 2, pp. 114–120.
doi: https://doi.org/10.1177/09637214231217286, accessed 19 January 2026.

Yochanan E. Bigman, Adam Waytz, Ron Alterovitz, and Kurt Gray, 2019. “Holding robots responsible: The elements of machine morality,” *Trends in Cognitive Sciences*, volume 23, number 5, pp. 365–368.
doi: https://doi.org/10.1016/j.tics.2019.02.008, accessed 19 January 2026.

Karen Brown, 2021. “Something bothering you? Tell it to Woebot,” *New York Times* (1 June), at https://www.nytimes.com/2021/06/01/health/artificial-intelligence-therapy-woebot.html, accessed 9 December 2024.

Mercedes Bunz and Marco Braghieri, 2022. “The AI doctor will see you now: Assessing the framing of AI in news coverage,” *AI & Society*, volume 37,number 1, pp. 9–22.
doi: https://doi.org/10.1007/s00146-021-01145-9, accessed 19 January 2026.

Guillermo Calahorra-Candao and María José Martín-de Hoyos, 2024. “The effect of anthropomorphism of virtual voice assistants on perceived safety as an antecedent to voice shopping,” *Computers in Human Behavior*, volume 153, number C, 108124.
doi: https://doi.org/10.1016/j.chb.2023.108124, accessed 19 January 2026.

Liuqing Chen, Shuhong Xiao, Yunnong Chen, Yaxuan Song, Ruoyu Wu, and Lingyun Sun, 2024. “ChatScratch: AI-augmented system toward autonomous visual programming learning for children aged 6-12,” *CHI ’24: Proceedings of the 2024 CHI Conference on Human Factors in Computing Systems*. article number 649, pp. 1–19.
doi: https://doi.org/10.1145/3613904.3642229, accessed 19 January 2026.

Szu Yu Chen, Rekha Tenjarla, Will Oremus, and Tucker Harris, 2023. “How to talk to an AI: An ordinary human’s guide to getting extraordinary results from a chatbot like ChatGPT,” *Washington Post* (31 August), at https://www.washingtonpost.com/technology/interactive/2023/how-to-talk-ai-chatbot-chatgpt/, accessed 9 December 2024.

Myra Cheng, Kristina Gligoric, Tiziano Piccardi, and Dan Jurafsky, 2024. “AnthroScore: A computational linguistic measure of anthropomorphism,” *arXiv*:2402.02056 (3 February).

doi: https://doi.org/10.48550/arXiv.2402.02056, accessed 19 January 2026.

Myra Cheng, Su Lin Blodgett, Alicia DeVrio, Lisa Egede, and Alexandra Olteanu, 2025. "Dehumanizing machines: Mitigating anthropomorphic behaviors in text generation systems," *Proceedings of the 63rd Annual Meeting of the Association for Computational Linguistics*, volume 1, pp. 25,923–25,948.
doi: https://doi.org/10.18653/v1/2025.acl-long.1259, accessed 19 January 2026.

DaEun Choi, Sumin Hong, Jeongeon Park, John Joon Young Chung, and Juho Kim, 2024. "CreativeConnect: Supporting reference recombination for graphic design ideation with generative AI," *CHI '24: Proceedings of the 2024 CHI Conference on Human Factors in Computing Systems*, article number 1055, pp. 1–25.
doi: https://doi.org/10.1145/3613904.3642794, accessed 19 January 2026.

Simon Coghlan, 2024. "Anthropomorphizing machines: Reality or popular myth?" *Minds and Machines*, volume 34, article number 25.
doi: https://doi.org/10.1007/s11023-024-09686-w, accessed 19 January 2026.

Michelle Cohn, Mahima Pushkarna, Gbolahan O. Olanubi, Joseph M. Moran, Daniel Padgett, Zion Mengesha, and Courtney Heldreth, 2024. "Believing anthropomorphism: Examining the role of anthropomorphic cues on trust in large language models," *CHI EA '24: Extended Abstracts of the CHI Conference on Human Factors in Computing Systems*, article number 54,, pp. 1–15.
doi: https://doi.org/10.1145/3613905.3650818, accessed 19 January 2026.

L. Elizabeth Crawford, 2009. "Conceptual metaphors of affect," *Emotion Review*, volume 1, number 2, pp. 129–139.
doi: https://doi.org/10.1177/1754073908100438, accessed 19 January 2026.

Jessica Dai, 2024. "Beyond personhood: Agency, accountability, and the limits of anthropomorphic ethical analysis," *arXiv*:2404.13861 (22 April).
doi: https://doi.org/10.48550/arXiv.2404.13861, accessed 19 January 2026.

Ophelia Deroy, 2023. "The ethics of terminology: Can we use human terms to describe AI?" *Topoi*, volume 42, number 3, pp. 881–889.
doi: https://doi.org/10.1007/s11245-023-09934-1, accessed 19 January 2026.

Smit Desai and Michael Twidale, 2023. "Metaphors in voice user interfaces: A slippery fish," *ACM Transactions on Computer-Human Interaction*, volume 30, number 6, article number 89, pp. 1–37.
doi: https://doi.org/10.1145/3609326, accessed 19 January 2026.

Alicia DeVrio, Myra Cheng, Lisa Egede, Alexandra Olteanu, and Su Lin Blodgett, 2025. "A taxonomy of linguistic expressions that contribute to anthropomorphism of language technologies," *arXiv*:2502.09870 (14 February).
doi: https://doi.org/10.48550/arXiv.2502.09870, accessed 19 January 2026.

Duolingo Team, 2023. "Introducing Duolingo Max, a learning experience powered by GPT-4" (14 March), at https://blog.duolingo.com/duolingo-max/, accessed 9 December 2024.

Audrey Eichenberger, Stephen Thielke, and Adam Van Buskirk, 2025. "A case of bromism influenced by use of artificial intelligence," *Annals of Internal Medicine: Clinical Cases*, volume 4, number 8, e241260.
doi: https://doi.org/10.7326/aimcc.2024.1260, accessed 19 January 2026.

Passant Elagroudy, Jie Li, Kaisa Väänänen, Paul Lukowicz, Hiroshi Ishii, Wendy E. Mackay, Elizabeth F. Churchill, Anicia Peters, Antti Oulasvirta, Rui Prada, Alexandra Diening, Giulia Barbareschi, Agnes Gruenerbl, Midori Kawaguchi, Abdallah El Ali, Fiona Draxler, Robin Welsch, and Albrecht Schmidt, 2024. "Transforming HCI research cycles using generative AI and 'large whatever models' (LWMs)," *CHI EA*

’24: Extended Abstracts of the CHI Conference on Human Factors in Computing Systems, article number 584, pp. 1–5.
doi: https://doi.org/10.1145/3613905.3643977, accessed 19 January 2026.

Raymond Fok, Nedim Lipka, Tong Sun, and Alexa F. Siu, 2024. “Marco: Supporting business document workflows via collection-centric information foraging with large language models,” *CHI ’24: Proceedings of the 2024 CHI Conference on Human Factors in Computing Systems*, article number 842, pp. 1–20.
doi: https://doi.org/10.1145/3613904.3641969, accessed 19 January 2026.

Jie Gao, Yuchen Guo, Gionnieve Lim, Tianqin Zhang, Zheng Zhang, Toby Jia-Jun Li, and Simon Tangi Perrault, 2024. “CollabCoder: A lower-barrier, rigorous workflow for inductive collaborative qualitative analysis with large language models,” *CHI ’24: Proceedings of the 2024 CHI Conference on Human Factors in Computing Systems*, article number 11, pp. 1–29.
doi: https://doi.org/10.1145/3613904.3642002, accessed 19 January 2026.

David Gefen, Barry A. Leffew, Arik Ragowsky, and René Riedl, 2025. “The importance of distrust in trusting digital worker chatbots,” *Communications of the ACM*, volume 68, number 4, pp. 42–49.
doi: https://doi.org/10.1145/3701297, accessed 19 January 2026.

Sarah Gibbons, Tarun Mugunthan, and Jakob Nielsen, 2023. “The 4 degrees of anthropomorphism of generative AI,” *NNGroup* (20 October), at https://www.nngroup.com/articles/anthropomorphism/ accessed 9 September 2024.

Zaria Gorvett, 2023. “The AI emotions dreamed up by ChatGPT,” *BBC* (24 February), at https://www.bbc.com/future/article/20230224-the-ai-emotions-dreamed-up-by-chatgpt, accessed 9 December 2024.

Jan Hartmann, Antonella De Angeli, and Alistair Sutcliffe, 2008. “Framing the user experience: Information biases on website quality judgement,” *CHI ’08: Proceedings of the SIGCHI Conference on Human Factors in Computing Systems*, pp. 855–864.
doi: https://doi.org/10.1145/1357054.1357190, accessed 19 January 2026.

Ziyao He, Shiyuan Li, Yunpeng Song, and Zhongmin Cai, 2024. “Towards building condition-based cross-modality intention-aware human-AI cooperation under VR environment,” *CHI ’24: Proceedings of the 2024 CHI Conference on Human Factors in Computing Systems*, article number 955, pp. 1–13.
doi: https://doi.org/10.1145/3613904.3642360, accessed 19 January 2026.

Rose K. Hendricks, Zsófia Demjén, Elena Semino, and Lera Boroditsky, 2018. “Emotional implications of metaphor: Consequences of metaphor framing for mindset about cancer,” *Metaphor and Symbol*, volume 33, number 4, pp. 267–279.
doi: https://doi.org/10.1080/10926488.2018.1549835, accessed 19 January 2026.

Colin Holbrook, Daniel Holman, Joshua Clingo, and Alan R. Wagner, 2024. “Overtrust in AI recommendations about whether or not to kill: Evidence from two human-robot interaction studies,” *Nature Scientific Reports*, volume 14, number 1 (4 September), article number 19751.
doi: https://doi.org/10.1038/s41598-024-69771-z, accessed 19 January 2026.

Francis Hunger, 2023. “Unhype artificial ‘intelligence’! A proposal to replace the deceiving terminology of AI,” *Training the Archive, Working Paper Series*, number 6.
doi: https://doi.org/10.5281/ZENODO.7524492, accessed 19 January 2026.

Nanna Inie, Stefania Druga, Peter Zukerman, and Emily M. Bender, 2024. “From ‘AI’ to probabilistic automation: How does anthropomorphization of technical systems descriptions influence trust?” *FAccT ’24: Proceedings of the 2024 ACM Conference on Fairness, Accountability, and Transparency*, pp. 2,322–2,347.

doi: https://doi.org/10.1145/3630106.3659040, accessed 19 January 2026.

Andrés Isaza-Giraldo, Paulo Bala, Pedro F. Campos, and Lucas Pereira, 2024. "Prompt-Gaming: A pilot study on LLM-evaluating agent in a meaningful energy game," *CHI EA '24: Extended Abstracts of the CHI Conference on Human Factors in Computing Systems*, article number 272, pp. 1–12.
doi: https://doi.org/10.1145/3613905.3650774, accessed 19 January 2026.

Nuwan Janaka, Runze Cai, Shengdong Zhao, and David Hsu, 2024. "Demonstrating PANDALens: Enhancing daily activity documentation with AI-assisted in-context writing on OHMD," *CHI EA '24: Extended Abstracts of the CHI Conference on Human Factors in Computing Systems*, article number 397, pp. 1–7.
doi: https://doi.org/10.1145/3613905.3648644, accessed 19 January 2026.

Hyeonsu B. Kang, David Chuan-En Lin, Nikolas Martelaro, Aniket Kittur, Yan-Ying Chen, and Matthew K. Hong, 2024. "BioSpark: An end-to-end generative system for biological-analogical inspirations and ideation," *CHI EA '24: Extended Abstracts of the CHI Conference on Human Factors in Computing Systems*, article number 61, pp. 1–13.
doi: https://doi.org/10.1145/3613905.3651035, accessed 19 January 2026.

Majeed Kazemitabaar, Runlong Ye, Xiaoning Wang, Austin Zachary Henley, Paul Denny, Michelle Craig, and Tovi Grossman, 2024. "Codeaid: Evaluating a classroom deployment of an LLM-based programming assistant that balances student and educator needs," *CHI '24: Proceedings of the 2024 CHI Conference on Human Factors in Computing Systems*, article number 650, pp. 1–20.
doi: https://doi.org/10.1145/3613904.3642773, accessed 19 January 2026.

Pranav Khadpe, Ranjay Krishna, Li Fei-Fei, Jeffrey T. Hancock, and Michael S. Bernstein, 2020. "Conceptual metaphors impact perceptions of human-AI collaboration," *Proceedings of the ACM on Human-Computer Interaction*, volume 4, number CSCW2, article number 163, pp. 1–26.
doi: https://doi.org/10.1145/3415234, accessed 19 January 2026.

Munyeong Kim, 2024. "GPTs in mafia-like game simulation," *CHI EA '24: Extended Abstracts of the CHI Conference on Human Factors in Computing System*, article number 611, pp. 1–5.
doi: https://doi.org/10.1145/3613905.3647958, accessed 19 January 2026.

Taenyun Kim and Hayeon Song, 2023. "Communicating the limitations of AI: The effect of message framing and ownership on trust in artificial intelligence," *International Journal of Human–Computer Interaction*, volume 39, number 4, pp. 790–800.
doi: https://doi.org/10.1080/10447318.2022.2049134, accessed 19 January 2026.

Taenyun Kim and Hayeon Song, 2020. "The effect of message framing and timing on the acceptance of artificial intelligence's suggestion," *CHI EA '20: Extended Abstracts of the 2020 CHI Conference on Human Factors in Computing Systems*, pp. 1–8.
doi: https://doi.org/10.1145/3334480.3383038, accessed 19 January 2026.

Taewan Kim, Seolyeong Bae, Hyun Ah Kim, Su-woo Lee, Hwajung Hong, Chanmo Yang, and Young-Ho Kim, 2024. "MindfulDiary: Harnessing large language model to support psychiatric patients' journaling," *CHI '24: Proceedings of the 2024 CHI Conference on Human Factors in Computing Systems*. article number 701, pp. 1–20.
doi: https://doi.org/10.1145/3613904.3642937, accessed 19 January 2026.

Rafal Kocielnik, Saleema Amershi, and Paul N. Bennett, 2019. "Will you accept an imperfect AI? Exploring designs for adjusting end-user expectations of AI systems," *CHI '19: Proceedings of the 2019 CHI Conference on Human Factors in Computing Systems*, paper number 411, pp. 1–14.
doi: https://doi.org/10.1145/3290605.3300641, accessed 19 January 2026.

Christoph Kueffer and Brendon M.H. Larson, 2014. "Responsible use of language in scientific writing and science communication," *BioScience*, volume 64, number 8, pp. 719–724.
doi: https://doi.org/10.1093/biosci/biu084, accessed 19 January 2026.

George Lakoff and Mark Johnson, 2008. *Metaphors we live by*. Chicago: University of Chicago Press.

Markus Langer, Tim Hunsicker, Tina Feldkamp, Cornelius J. König, and Nina Grgić-Hlača, 2022. ""'Look! It's a computer program! It's an algorithm! It's AI!': Does terminology affect human perceptions and evaluations of algorithmic decision-making systems?" *CHI '22: Proceedings of the 2022 CHI Conference on Human Factors in Computing Systems*, article number 581, pp. 1–28.
doi: https://doi.org/10.1145/3491102.3517527, accessed 19 January 2026.

Rita Latikka, Jenna Bergdahl, Iina Savolainen, Magdalena Celuch, and Atte Oksanen, 2025. "Individual and well-being factors associated with social chatbot usage: A six-country study," *Journal of Social and Personal Relationships* (29 October).
doi: https://doi.org/10.1177/02654075251392956, accessed 19 January 2026.

Ned Letcher, 2018. "Discovering syntactic phenomena with and within precision grammars," Ph.D. thesis, Department of Computing and Information Systems, University of Melbourne, at https://minerva-access.unimelb.edu.au/bitstream/handle/11343/214022/thesis.pdf, accessed 19 January 2026.

Sebastian Linxen, Christian Sturm, Florian Brühlmann, Vincent Cassau, Klaus Opwis, and Katharina Reinecke, 2021. How WEIRD is CHI?" *CHI '21: Proceedings of the 2021 CHI Conference on Human Factors in Computing Systems*, article number 143, pp. 1–14.
doi: https://doi.org/10.1145/3411764.3445488, accessed 19 January 2026.

Zachary C. Lipton and Jacob Steinhardt, 2019. "Troubling trends in machine learning scholarship: Some ML papers suffer from flaws that could mislead the public and stymie future research," *ACM Queue*, volume 17, number 1, pp. 45–77.
doi: https://doi.org/10.1145/3317287.3328534, accessed 19 January 2026.

Wenlong Liu, Min Jiang, Wangjie Li, and Jian Mou, 2024. "How does the anthropomorphism of AI chatbots facilitate users' reuse intention in online health consultation services? The moderating role of disease severity," *Technological Forecasting and Social Change*, volume 203, 123407.
doi: https://doi.org/10.1016/j.techfore.2024.123407, accessed 19 January 2026.

Steve Lohr, 2024. "A.I. can write poetry, but it struggles with math," *New York Times* (23 July), at https://www.nytimes.com/2024/07/23/technology/ai-chatbots-chatgpt-math.html, accessed 6 January 2025.

Takuya Maeda and Anabel Quan-Haase, 2024. "When human-AI interactions become parasocial: Agency and anthropomorphism in affective design," *FAccT '24: Proceedings of the 2024 ACM Conference on Fairness, Accountability, and Transparency*, pp. 1,068–1,077.
doi: https://doi.org/10.1145/3630106.3658956, accessed 19 January 2026.

John McCarthy, Marvin L. Minsky, Nathaniel Rochester, and Claude E. Shannon, 1955. "A proposal for the Dartmouth Summer Research Project on Artificial Intelligence" (31 August), at http://jmc.stanford.edu/articles/dartmouth/dartmouth.pdf, accessed 19 January 2026.

Drew McDermott, 1976. "Artificial intelligence meets natural stupidity," *ACM SIGART Bulletin*, number 57, pp. 4–9.
doi: https://doi.org/10.1145/1045339.1045340, accessed 19 January 2026.

Matthew S. McGlone, 1996. "Conceptual metaphors and figurative language interpretation: Food for thought?" *Journal of Memory and Language*, volume 35, number 4, pp. 544–565.
doi: https://doi.org/10.1006/jmla.1996.0029, accessed 19 January 2026.

Alican Mecit, Tina M. Lowrey, and L.J. Shrum, 2022. “Grammatical gender and anthropomorphism: ‘It’ depends on the language,” *Journal of Personality and Social Psychology*, volume 123, number 3, pp. 503–517.
doi: https://doi.org/10.1037/pspa0000309, accessed 19 January 2026.

Cade Metz, 2024. “How the A.I. that drives ChatGPT will move into the physical world,” *New York Times* (11 March), at https://www.nytimes.com/2024/03/11/technology/ai-robots-technology.html, accessed 9 December 2024.

Don Norman, 2013. *The design of everyday things*. Revised and expanded edition. New York: Basic Books.

OpenAI, 2024. “Creating video from text,” at https://openai.com/index/sora/, accessed 9 December 2024.

Saumya Pareek, Sarah Schömbs, Eduardo Velloso, and Jorge Goncalves, 2025. “‘It’s not the AI’s fault because it relies purely on data’: How causal attributions of AI decisions shape trust in AI systems,” *CHI ’25: Proceedings of the 2025 CHI Conference on Human Factors in Computing Syste*, article number 1097, pp. 1–18.
doi: https://doi.org/10.1145/3706598.3713468, accessed 19 January 2026.

Minju Park, Sojung Kim, Seunghyun Lee, Soonwoo Kwon, and Kyuseok Kim, 2024. “Empowering personalized learning through a conversation-based tutoring system with student modeling,” *CHI EA ’24: Extended Abstracts of the CHI Conference on Human Factors in Computing Systems*, article number 123, pp. 1–10.
doi: https://doi.org/10.1145/3613905.3651122, accessed 19 January 2026.

Pat Pataranutaporn, Ruby Liu, Ed Finn, and Pattie Maes, 2023. “Influencing human–AI interaction by priming beliefs about AI can increase perceived trustworthiness, empathy and effectiveness,” *Nature Machine Intelligence*, volume 5, number 10, pp. 1,076–1,086.
doi: https://doi.org/10.1038/s42256-023-00720-7, accessed 19 January 2026.

Paolo Peverini, 2024. *Bruno Latour in the semiotic turn: An inquiry into the networks of meaning*. Cham, Switzerland: Springer.
doi: https://doi.org/10.1007/978-3-031-57178-7, accessed 19 January 2026.

Josef M. Pierre, Ben Geta, Govind Raghavan, and Karthik V. Sharma, 2025. “‘You’re not crazy’: A case of new-onset AI-associated psychosis,” *Innovations in Clinical Neuroscience*, volume 22 (18 November), pp. 10–12, and at https://innovationscns.com/youre-not-crazy-a-case-of-new-onset-ai-associated-psychosis/, accessed 19 January 2026.

Gil Press, 2023. “Diagnostic robotics AI advances predictive, personalized medicine,” *Forbes* (17 July), at https://www.forbes.com/sites/gilpress/2023/07/17/diagnostic-robotics-ai-advances-predictive-personalized-medicine/, accessed 13 September 2024.

Hua Xuan Qin, Shan Jin, Ze Gao, Mingming Fan, and Pan Hui, 2024. “CharacterMeet: Supporting creative writers’ entire story character construction processes through conversation with LLM-powered chatbot avatars,” *CHI ’24: Proceedings of the 2024 CHI Conference on Human Factors in Computing Systems*, article number 1051, pp. 1–19.
doi: https://doi.org/10.1145/3613904.3642105, accessed 19 January 2026.

Rainer Rehak, 2021. “The language labyrinth: Constructive critique on the terminology used in the AI discourse,” In: Pieter Verdegem (editor). *AI for everyone? Critical perspectives*. London: University of Westminster Press, pp. 87–102.
doi: https://doi.org/10.16997/book55.f, accessed 19 January 2026.

Paul Robinette, Wenchen Li, Robert Allen, Ayanna M Howard, and Alan R Wagner, 2016. "Overtrust of robots in emergency evacuation scenarios," *2016 11th ACM/IEEE International Conference on Human-Robot Interaction (HRI)*, pp. 101–108.
doi: https://doi.org/10.1109/HRI.2016.7451740, accessed 19 January 2026.

Igor Ryazanov, Carl Öhman, and Johanna Björklund, 2025. "How ChatGPT changed the media's narratives on AI: A semi-automated narrative analysis through frame semantics," *Minds and Machines*, volume 35, article number 2.
doi: https://doi.org/10.1007/s11023-024-09705-w, accessed 19 January 2026.

Arleen Salles, Kathinka Evers, and Michele Farisco, 2020. "Anthropomorphism in AI," *AJOB Neuroscience*, volume 11, number 2, pp. 88–95.
doi: https://doi.org/10.1080/21507740.2020.1740350, accessed 19 January 2026.

Advait Sarkar, 2023. "Enough with 'human-AI collaboration'," *CHI EA '23: Extended Abstracts of the 2023 CHI Conference on Human Factors in Computing Systems*, article number 415, pp. 1–8.
doi: https://doi.org/10.1145/3544549.3582735, accessed 19 January 2026.

Anna-Maria Seeger, Jella Pfeiffer, and Armin Heinzl, 2021. "Texting with humanlike conversational agents: Designing for anthropomorphism," *Journal of the Association for Information Systems*, volume 22, number 4.
doi: https://doi.org/10.17705/1JAIS.00685, accessed 19 January 2026.

Elena Semino, Zsófia Demjén, Jane Demmen, Veronika Koller, Sheila Payne, Andrew Hardie, and Paul Rayson, 2017. "The online use of Violence and Journey metaphors by patients with cancer, as compared with health professionals: A mixed methods study," *BMJ Supportive & Palliative Care*, volume 7, number 1, pp. 60–66.
doi: https://doi.org/10.1136/bmjspcare-2014-000785, accessed 19 January 2026.

Murray Shanahan, 2024. "Talking about large language models," *Communications of the ACM*, volume 67, number 2, pp. 68–79.
doi: https://doi.org/10.1145/3624724, accessed 19 January 2026.

Matthew Shardlow and Piotr Przybyla, 2024. "Deanthropomorphising NLP: Can a language model be conscious?" *PLoS ONE*, volume 19, number 12 (4 December), e0307521.
doi: https://doi.org/10.1371/journal.pone.0307521, accessed 19 January 2026.

Kihoon Son, DaEun Choi, Tae Soo Kim, Young-Ho Kim, and Juho Kim, 2024. "Genquery: Supporting expressive visual search with generative models," *CHI '24: Proceedings of the 2024 CHI Conference on Human Factors in Computing Systems*, article number 180, pp. 1–19.
doi: https://doi.org/10.1145/3613904.3642847, accessed 19 January 2026.

Michael T. Stuart and Markus Kneer, 2021. "Guilty artificial minds: Folk attributions of mens rea and culpability to artificially intelligent agents," *Proceedings of the ACM on Human-Computer Interaction* volume 5, number CSCW2, article number 363, pp. 1–27.
doi: https://doi.org/10.1145/3479507, accessed 19 January 2026.

Megan Strait, Ana Sánchez Ramos, Virginia Contreras, and Noemi Garcia, 2018. "Robots racialized in the likeness of marginalized social identities are subject to greater dehumanization than those racialized as white," *2018 27th IEEE International Symposium on Robot and Human Interactive Communication (RO-MAN)*, pp. 452–457.
doi: https://doi.org/10.1109/ROMAN.2018.8525610, accessed 19 January 2026.

Sangho Suh, Meng Chen, Bryan Min, Toby Jia-Jun Li, and Haijun Xia, 2024. "Luminate: Structured generation and exploration of design space with large language models for human-AI co-creation," *CHI*

'24: Proceedings of the 2024 CHI Conference on Human Factors in Computing Systems, article number 644, pp. 1–26.
doi: https://doi.org/10.1145/3613904.3642400, accessed 19 January 2026.

Mirac Suzgun. Tayfun Gur, Federico Bianchi, Daniel E. Ho, Thomas Icard, Dan Jurafsky, and James Zou, 2025. "Language models cannot reliably distinguish belief from knowledge and fact," *Nature Machine Intelligence*, volume 7 (3 November), pp. 1,780–1,790.
doi: https://doi.org/10.1038/s42256-025-01113-8, accessed 19 January 2026.

Paul H. Thibodeau and Lera Boroditsky, 2013. "Natural language metaphors covertly influence reasoning," *PLoS ONE* volume 8, number 1 (2 January), e52961.
doi: https://doi.org/10.1371/journal.pone.0052961, accessed 19 January 2026.

Caroline Tipler and Janet B. Ruscher, 2014. "Agency's role in dehumanization: Non-human metaphors of out-groups," *Social and Personality Psychology Compass*, volume 8, number 5, pp. 214–228.
doi: https://doi.org/10.1111/spc3.12100, accessed 19 January 2026.

Stephanie Tully, Chiara Longoni, and Gil Appel, 2023. "Lower artificial intelligence literacy predicts greater AI receptivity," *Journal of Marketing*, volume 89, number 5, pp. 1–20.
doi: https://doi.org/10.1177/00222429251314491, accessed 19 January 2026.

Chao Wang, Stephan Hasler, Daniel Tanneberg, Felix Ocker, Frank Joublin, Antonello Ceravola, Joerg Deigmoeller, and Michael Gienger, 2024a. "LaMI: Large language models for multi-modal human-robot interaction," *CHI EA '24: Extended Abstracts of the CHI Conference on Human Factors in Computing Systems*, article number 218, pp. 1–10.
doi: https://doi.org/10.1145/3613905.3651029, accessed 19 January 2026.

Zhijie Wang, Yuheng Huang, Da Song, Lei Ma, and Tianyi Zhang, 2024b. "PromptCharm: Text-to-image generation through multi-modal prompting and refinement," *CHI '24: Proceedings of the 2024 CHI Conference on Human Factors in Computing Systems*, article number 185, pp. 1–21.
doi: https://doi.org/10.1145/3613904.3642803, accessed 19 January 2026.

Florian Weber, Thiemo Wambsganss, Seyed Parsa Neshaei, and Matthias Soellner, 2024. "LegalWriter: An intelligent writing support system for structured and persuasive legal case writing for novice law students," *CHI '24: Proceedings of the 2024 CHI Conference on Human Factors in Computing Systems*, article number 1052, pp. 1–23.
doi: https://doi.org/10.1145/3613904.3642743, accessed 19 January 2026.

Chao Zhang, Xuechen Liu, Katherine Ziska, Soobin Jeon, Chi-Lin Yu, and Ying Xu, 2024. "Mathemyths: Leveraging large language models to teach mathematical language through child-AI co-creative storytelling," *CHI '24: Proceedings of the 2024 CHI Conference on Human Factors in Computing Systems*, article number 274, pp. 1–23.
doi: https://doi.org/10.1145/3613904.3642647, accessed 19 January 2026.

Yaqiong Zhang and Shifu Wang, 2023. "The influence of anthropomorphic appearance of artificial intelligence products on consumer behavior and brand evaluation under different product types," *Journal of Retailing and Consumer Services*, volume 74, 103432.
doi: https://doi.org/10.1016/j.jretconser.2023.103432, accessed 19 January 2026.

Yu Zhang, Jingwei Sun, Li Feng, Cen Yao, Mingming Fan, Liuxin Zhang, Qianying Wang, Xin Geng, and Yong Rui, 2024. "See widely, think wisely: Toward designing a generative multi-agent system to burst filter bubbles," *CHI '24: Proceedings of the 2024 CHI Conference on Human Factors in Computing Systems*, article number 484, pp. 1–24.
doi: https://doi.org/10.1145/3613904.3642545, accessed 19 January 2026.

Zeyu Zhang, Fu Guo, Chen Fang, and Jiahao Chen, 2025. "Let me hold your hand: Effects of anthropomorphism and touch behavior on self-disclosure intention, attachment, and cerebral activity towards AI mental health counselors," *International Journal of Human–Computer Interaction*, volume 41, number 18, pp. 11,832–11,847.
doi: https://doi.org/10.1080/10447318.2024.2446502, accessed 19 January 2026.

Ivan Zhao, 2022. "Introducing Notion AI" (16 November), at https://www.notion.so/blog/introducing-notion-ai, accessed 9 December 2024.

Editorial history

Received 3 September 2025; revised 7 January 2026; accepted 19 January 2026.

De-anthropomorphizing "AI": From wishful mnemonics to accurate nomenclature
by Nanna Inie, Peter Zukerman, and Emily M. Bender.
First Monday, volume 31, number 2 (February 2026).
doi: https://dx.doi.org/10.5210/fm.v31i2.14366