Promises, Promises: Understanding Claims Made in Social Robot Consumer Experiences

Johanna Gunawan
Khoury College of Computer Sciences
Northeastern University
Boston, Massachusetts, USA
gunawan.jo@northeastern.edu

Sarah Elizabeth Gillespie
Cybersecurity and Privacy Institute
Northeastern University
Boston, Massachusetts, USA
gillespie.s@northeastern.edu

David Choffnes
Northeastern University
Boston, Massachusetts, USA
choffnes@ccs.neu.edu

Woodrow Hartzog
School of Law
Boston University
Boston, Massachusetts, USA
whartzog@bu.edu

Christo Wilson
Khoury College of Computer Sciences
Northeastern University
Boston, Massachusetts, USA
cbw@ccs.neu.edu

Abstract

Social robots are a class of emerging smart consumer electronics devices that promise sophisticated experiences featuring emotive capabilities, artificial intelligence, conversational interaction, and more. With unique risk factors like emotional attachment, little is known on how social robots communicate these promises to consumers and whether they adequately deliver upon them within their overall product experiences prior to and during user interaction.

Animated by a consumer protection lens, this paper systematically investigates manufacturer claims made for four commercially available social robots, evaluating these claims against the provided user experience and consumer reviews. We find that social robots vary widely in the manner and extent to which they communicate intelligent features and the supposed benefits of these features, while consumer perspectives similarly include a wide range of perceptions on robot and AI performance, capabilities, and product frustrations. We conclude by discussing social robots' unique propensities for consumer risk, and consider implications for regulators, developers, and researchers of social robots.

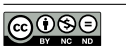
CCS Concepts

• **Human-centered computing** → **Human computer interaction (HCI)**; **Interaction design**; • **Security and privacy** → Human and societal aspects of security and privacy; • **Social and professional topics** → *Computing / technology policy*.

Keywords

UX design, IoT, human-robot interaction, AI, consumer protections

ACM Reference Format:
Johanna Gunawan, Sarah Elizabeth Gillespie, David Choffnes, Woodrow Hartzog, and Christo Wilson. 2025. Promises, Promises: Understanding Claims Made in Social Robot Consumer Experiences. In *CHI Conference on Human Factors in Computing Systems (CHI '25), April 26–May 01, 2025, Yokohama, Japan.* ACM, New York, NY, USA, 22 pages. https://doi.org/10.1145/3706598.3713471



This work is licensed under a Creative Commons Attribution-NonCommercial-NoDerivatives 4.0 International License.
CHI '25, Yokohama, Japan
© 2025 Copyright held by the owner/author(s).
ACM ISBN 979-8-4007-1394-1/25/04
https://doi.org/10.1145/3706598.3713471

1 Introduction

Smart, Internet-connected consumer electronics are increasingly ubiquitous, serving a vast array of functions from health tracking to home automation [3]. The category of smart consumer electronics now includes robots, some of which are marketed as including features that are powered by artificial intelligence (AI). Examples of features include natural language conversation [94], emotion and face recognition [26], and even interpretation of human body language [98]. Smart devices in this category with these kinds of features are often referred to as *social robots* [16, 32, 92]. Research has shown that social robots can potentially bring a number of benefits to consumers, such as curbing loneliness [38], improving mental health [76], facilitating education and learning [96], and playing games [49].

Device manufacturers communicate the value of their products to consumers through mechanisms including advertising, promotional materials, product documentation, and product packaging. When considering whether to purchase consumer electronics like social robots, shoppers must assess the value a product might provide based on the claims made by the manufacturer. These product claims help to establish consumers' expectations about how they may interact with a social robot, as this is not a product category with well-established interaction norms (versus say, smartphones) and the interfaces may rely on open-ended AI models (e.g., for voice-based interaction). Consumer protection regulations and principles work to ensure that product claims are truthful and met by manufacturers, i.e., to minimize misalignment between product claims and resultant consumer experiences. Enforcers, however, have limited resources, and thus newly emergent technologies like social or companion robots may escape enforcement for problematic or unmet production claims.

In this study, we present the first attempt to rigorously evaluate promised and resultant experiences in commercially-available social robots. We seek to answer the following research questions:

- **RQ 1**: *What claims do social robot manufacturers make to prospective consumers?* Manufacturers' claims are a primary

source of product information that consumers may access before their first interaction with a robot.
- **RQ 2**: *To what extent do social robot user experiences deliver upon claims?* Auditing manufacturer claims for full or partial fulfillment may reveal inconsistencies or misalignment between a social robot's communicated and actual capabilities.
- **RQ 3**: *How do consumers describe their experiences with social robots?* Consumers provide varied feedback, highlighting a lack of common understanding or standards for assessing the benefits of social robots.

To answer these questions, we acquired and evaluated four commercially available robots: Eilik, Miko, Moxie, and Vector (see Figure 1). Our corpus reflects product diversity within the consumer robots market, ranging in popularity, affordability (costing consumers between US $139-800), domain (including education and general entertainment), interaction methods (like touch via sensors or on-screen, voice control, facial recognition) and anthropomorphized aesthetics.

For RQ 1 and RQ 2, we center our methods around the promises and subsequent expectations communicated by social robot manufacturers to their users. Specifically, we consult manufacturers' product claims (N=174) from product packaging and related consumer documentation, then characterize these statements inductively. We then directly test a subset of these claims (N=64) through human interaction with each robot. We adapt our interaction and manual content analysis approaches from prior scholarship observing user experiences *in situ* [28, 42, 54, 82], performing set-up and feature exploration interactions in each robot experience. To answer RQ 3, we collect 168 consumer reviews posted to robots' product websites and Amazon listings between January and September 2024, then code reviews to characterize aspects of robot experiences mentioned in positive or negative feedback.

We find that social robots vary widely in the manner and extent to which they communicate intelligent features and the supposed benefits of these features. The vast majority (98%) of claims we could test in-experience were at least minimally delivered upon. Consumer reviews from the same time period provide additional context: user frustrations with operability and perceived underperformance highlight the divide between consumer expectations and product claims. We conclude by discussing social robots' unique characteristics and propensities for consumer risk and consider implications for key stakeholders, including regulators, enforcement agencies, and practitioners involved with the development and sale of consumer-facing social robots, as well as researchers studying human-robot interactions.

2 Background and Related Work

We now review related work on social and companion robots. First, we focus on perspectives of risks and harm for social robots, then cover literature on their unique intelligence and anthropomorphic capabilities, and finally review related work auditing misalignment in product experiences. We then situate this study within broader scholarship.

2.1 Risk, Harm, and Fear of Social Robots

Advances in AI and robotics have brought the science fiction dream of ubiquitous robots closer to reality. Consequently, the volume of scholarship on social and companion robots has grown rapidly. Scholars have discussed critical elements for robot design, including how trust in smart robots is developed or perceived [60, 69], how non-verbal cues should be incorporated into social robots [6], what normative behaviors users expect from robots across cultures [57], and other desirable traits for social robots [61]. Other research has explored the challenges and opportunities of social robots [56, 89] in particular settings like service-industry work [51, 71] or in the home [72].

As there is a real possibility that social robot adoption will increase in the near future, researchers have begun to interrogate the capacity for social robots to have real influence over human emotions, behaviors, and attitudes. This work often urges ethical and responsible design for social robots [15, 91], including designs that handle practical aspects of social robot development, such as the impact from social robot commercialization into readily available consumer electronics [14], long-term engagement with robots [20], and preparing for robot "death" insofar as companion bots will not "live forever" [50]. (We revisit the issue of robot death in § 5, as this directly pertains to the robots in our study.)

Although there is enthusiasm for applications of social robots, scholars are careful to discuss potential risks and harms to human users. Hartzog [45] descriptively categorizes types of consumer-facing robots, demonstrating their potential for unfairness and deception. Some potential harms arise due to the inclusion—or claimed inclusion—of AI in social robots. As Narayanan and Kapoor [65] observe, "AI" has become an umbrella term that describes many distinct technologies with varying levels of capability and effectiveness. Robots manufacturers sometimes claim that their products are "intelligent" due to the incorporation of AI, which raises the spectre that AI "snake oil" claims may establish unrealistic or distorted ideas about the capabilities of social robots in the minds of consumers. Additionally, studies have found that people mistakenly conflate AI with robots [24], which may exacerbate these effects.

Another controversial facet of robot design concerns whether anthropomorphization is helpful or harmful [46]. Indeed, as we will show, the robots in our study are designed with anthropomorphic features, with marketing claims that can emphasize their human-like traits or conflate emotion with intelligence. Studies have found that building robots with anthropomorphic qualities can increase acceptance [58], and that in certain cases (such as socially assistive robots), anthropomorphism carries low ethical risk compared against its benefit to robot efficacy [93]. However, this contrasts with the philosophical discussion around the potential for a "hallucinatory danger": some scholars [7] argue that companion robots' primary threat comes not from passing simulated human-like behavior as real, but instead from misplaced human attachment and projection onto robots that cannot withstand overladen meaning. This threat is presently realized in the context of conversational AI and chatbots, with users left in real grief and distress after losing the companionship of virtual agents they developed feelings

for [27]. Emotional attachment leaves users vulnerable to exploitation by those that deliver artificial companionship. Given these risks, human-robot interaction scholars have raised existential questions about the safety of "cute" companion robots [23].

Two of the robots in our study are marketed to children. On one hand, studies have found positive attitudes towards social robot adoption in settings like storytelling for children [59], as well as positive outcomes from social robots adopted in children's educational settings [13, 19, 33, 78, 85]. On the other hand, scholars have also investigated how children perceive AI, finding that they harbor significant misconceptions [31, 52, 53, 62], which may heighten the risks we identify above (e.g., emotional attachment caused by anthropomorphism).

2.2 Auditing Internal Consistency (or Misalignment)

Computer science literature offers various auditing methods for assessing digital services and products, particularly in the fields of privacy, security, and technology ethics. Some audits have examined the gap between what platforms or manufacturers claim in their policy documents versus the actual implementations of their systems [2, 18, 97]. In the consumer device context, Sun et al. [87] inspected smart home products targeted towards children and families using a narrative-focused methodology. They found misalignment between vendors' depictions of smart home experiences and the privacy or safety information they provide. In general, these audits evaluate internal consistency within a given digital service.

Scholars have also used user review data to characterize problems arising in consumer-facing technologies. Hwang [47] provides a brief discussion of the utility of user review data in UX research, particularly highlighting the value of real-world feedback in retrospective research while acknowledging potential biases in user reviews. Studies have found that reviews unveil a host of issues in digital products and services related to usability [29], accessibility [30, 81], and health [11, 43]. O'Hagan et al. [70] suggest that reviews function as a monitoring or reporting tool for consumers, wherein complaints in user reviews reveal concerns beyond usability, like community safety. In the AI context, Namvarpour and Razi [64] explored reviews of the Replika chatbot as source material for better understanding human-AI interaction. Using automated methods, they found contradictions (a.k.a. misalignment) between parts of Replika's purported systems as well as between users' expectations of what Replika would do and what it actually did.

2.3 Building Upon Prior Work

In this study, we empirically audits robot manufacturers' product claims from human interaction and consumer protection perspectives. We build upon prior literature in human-robot interaction and artificial intelligence ethics. Our work is motivated by prior work that has theorized about the potential for social robots to subvert consumer expectations of user experiences [32, 55, 84], including harmful designs that abuse the anthropomorphic features of robots [7, 32, 88].

Similar to prior work, we evaluate social robots to understand the harms to consumers they may cause. As we discuss in § 3.1, we investigate four commercially-available social robots that are targeted to average consumers. This expands on prior work investigating failed robots [23] or limited-access robots [22]. Then, like Sun et al. [87], we assess product pages to capture what types of product claims are being made. We depart from Sun et al. [87]'s scope in the following manners: we inspect social robots without intentionally focusing on the children's context; inspect product claims holistically rather than focusing primarily on privacy or security; and compare product promises to in-the-wild user reviews and the resultant experience instead of comparing between product depictions and privacy claims.

We draw on manual content analysis methods previously used by scholarship auditing UX designs for potentially harmful consumer outcomes. In particular, prior work has used structured, manual interaction approaches to elicit, identify, and document deceptive designs (often referred to as "dark patterns") in various digital services [28, 82]. This includes approaches that carefully orchestrate manual interactions across systems spanning multiple modalities, e.g., a physical device and a smartphone app [42, 54]. As we discuss in § 3.2, we utilize similar methods to interact with robots. Unlike prior work, however, we do not evaluate robots against a predefined codebook of extant deceptive designs. Instead, we evaluate claims made by the manufacturers of the robots against on-device experiences and consumer-reported feedback extracted from public reviews. For review analysis, we depart from the large-scale natural language processing methods favored in prior related work [63, 68, 83] and instead manually code user reviews to better suit the narrower scope and smaller dataset of this work.

As such, this work presents an exploratory audit of the user experience provided by social robots, as motivated by the consumer protections concerns brought to light by prior work on human-robot interaction and technology audits.

3 Methods

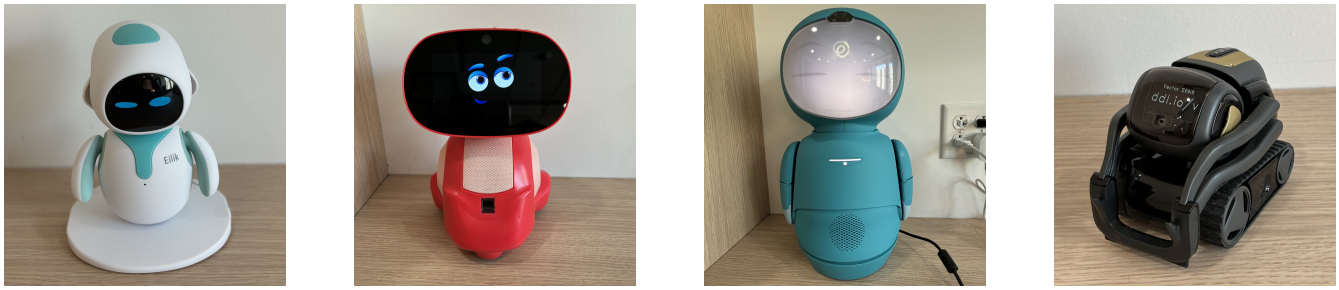
In this section we describe the robot selection, experiment development, and coding procedures used in this study.

3.1 Robot Inclusion Criteria and Description

We selected robots for our study using an iterative search process conducted in 2023. We first searched for lists of AI-enabled consumer products using keywords like "devices," "gadgets," or "consumer electronics" coupled with an "AI-powered/enabled," "smart," or simply "AI" modifier to the search word "robots." This yielded a wide assortment of devices—spanning smart home products to wearables, tools, and more—that claimed to offer AI-driven or smart functionality to varying degrees.[1] Throughout this search process, we also noted sites aggregating lists of such robots, either in blog/listicle format with short written descriptions or in collections of related products. Such lists often included overlapping robots, even across lists for nominally different purposes, such as "AI robot toys" or "AI personal robots."

From these search results, we built a shortlist of robots that appeared to include social or companionship-related features. We then selected a subset of four robots to purchase, based on the following inclusion criteria:

[1] For example, we found AI-enabled lawnmowers.



(a) Eilik (b) Miko 3 (c) Moxie (d) Vector 2.0

Figure 1: Photographs of the four robots in our study, in our lab environment. Vector is faceless due to operability issues in § 3.

				Interaction Modalities			
Robot	**Manufacturer**	**Price**	**Intended Users**	*Touch Input*	*Mobile App*	*Voice Control*	*Motion Sensing*
Eilik	Energize Lab	139 USD	General	✓	–	–	–
Miko	Miko.ai	250 USD	Children	✓(Face/screen)	✓	✓	✓(Environment)
Moxie	Embodied, Inc.	800 USD	Children	–	✓(Guardian)	✓	✓(Face)
Vector	Digital Dream Labs	399 USD	General	✓	✓	✓	–

Table 1: General product information for the robots in our study.

(1) commercially available and operable within the United States at the time of the study,
(2) primarily marketed as a social or companion robot,
(3) marketed as including AI or intelligent features in the product description,
(4) and some degree of human-like facial expression in the on-device display.

The four robots we purchased, pictured in Figure 1, were Eilik, Miko 3, Moxie, and Vector 2.0. These robots cover a range of robot traits, summarized at a high level in Table 1. Specifically, they span a wide manufacturer price range with the cheapest robot (Eilik) priced at US $139 and highest (Moxie) at US $800, market to different age ranges (Miko and Moxie are marketed towards children; Miko with kids' content and Moxie to teach children about AI), serve different purposes (Vector and Eilik are sometimes described as desktop companions, whereas Miko and Moxie claim educational or other learning benefits for children), and offer varying interaction methods to users (Miko's face doubles as a touchscreen; Miko, Moxie, and Vector offer voice interactions and companion apps; Eilik responds to touch in different areas of its body). Aesthetically, Moxie is the most similar to humans out of the four, with human-like upper body structure (head, torso, two arms), detailed full-color facial animations served by a large screen, and a size roughly that of a human baby's. Eilik is similar in appearance to the Moxie with the same upper body type, but stands only about five inches high and serves pixelated facial expressions on a two-tone display. Miko and Vector both have wheels and neither have arms, but Miko (at nearly nine inches tall) serves full-color facial expressions on a screen whereas Vector (at approximately four inches tall) uses a small two-tone pixelated display.

Unfortunately, we were unable to purchase an Amazon Astro because it was available by invitation only (thus failing to meet criteria 1). The JIBO robot was available for purchase but had been obsoleted by the manufacturer at the time of our study (again, failing to meet criteria 1) and did not feature a human-like front display (criteria 4).

3.2 Evaluation Through Manual Interaction

To answer RQ 1 and RQ 2, we manually investigated our four robots in a laboratory setting. This investigation was organized into two phases. First, we collected documents for users of each robot as provided by the manufacturers, including product webpages, product packaging, and user manuals.[2] We coded the product claims in these documents and organized them into four, per-robot codebooks. Second, we evaluated whether the claims in the codebooks were borne out in practice by directly interacting with the robots. We completed each robots' setup procedure; sought out and tested each promised on-device feature; and examined robot experiences for purported intelligent or privacy features wherever possible. We discuss each phase of our testing in greater detail below.

3.2.1 Manufacturer Claim Collection and Codebook Development. For our corpus of four robots, we saved copies of their manufacturer product pages, packaging, and paper manuals. We manually transcribed any product claims made in these documents into a spreadsheet. We then conducted iterative coding rounds using thematic and document analysis approaches [10, 12] to categorize these claims according to the primary purpose each communicated—resultant themes are presented in Table 2 and § 4.1. Additionally, we annotate each claim for two binary features: mentions of intelligent, smart, or AI features, as well as for whether the robot was described in a personified or anthropomorphized manner (e.g., described as a person rather than an object). Two authors manually applied the binary labels on all 174 claims (so, 348 total labels). We computed label agreement between the two annotators, which resulted in

[2] Our Vector came without external packaging sleeve; we collected images of 2.0 packaging from manufacturer pages and included only claims which were clearly legible.

Main Claim Category	Subcategory	Description
Safety and Compliance	Hardware Compliance	Disclosures demonstrating the robot's compliance to relevant hardware or electronics regulations.
	Privacy and Safety	Claims that purport privacy protections or related safety measures.
	Consumer or Purchase Protections	Claims pertaining to the purchase of a social robot, including warranty or returns policies.
Purported Feature(s)	**Included Features**	Main features derived from the primary purpose communicated by each claim.
	Future Features	Claims for future feature additions, e.g., continued development and releases.
	Feature Interactions	Interaction instructions for the user to discover features.
	Hardware Specifications	Descriptions of robot hardware, e.g., sensors, materials used, etc.
Outcomes or Value Propositions	**Purported Outcomes**	Promised outcomes or potential consumer value from the robot experience. These may include mentions of features, but the primary purpose of these claims is to demonstrate value with feature descriptions being a secondary purpose.
	Umbrella Statements	Abstracted or holistic product value propositions, typically product taglines.
Product Requirements	**Hard Requirements**	Necessary materials, parameters, or actions for using the robot.
	Soft Requirements	Recommended materials, parameters, or actions for using the robot.

Table 2: Claim categories resulting from our thematic analysis of all 174 claims. Claim categories contributing to the subset of claims whose validity we were were able to assess through manual interaction are bolded.

Cohen's kappa of κ=0.729, indicating substantial agreement [90]. Overall agreement between labels was 90%, with 79% positive agreement and 93% negative agreement. Discrepancies were discussed between authors and we adjusted the final labels for full consensus.

To determine whether product claims were upheld or not in user experiences, we used these claims as a codebook of binary labels. Not all claims were amenable to observation on-device, so we filtered out claims not feasibly verifiable within our lab environment or methods. These include sweeping marketing claims, off-device compliance descriptions, or other claims external to the hands-on robot experience. During tests, we discovered that Vector was fully inoperable, and subsequently filtered out Vector claims as well. This resulted in a set of 64 total claims that we sought to validate directly in robot experiences.

3.2.2 Laboratory Environment. We adopted best practices and methods for our experiments that have been used by prior work that leveraged manual interaction methods to evaluate products and services [42, 54]. We recorded video of all our interactions with the robots, so to revisit and review each robot's actions. Miko and Moxie required a companion smartphone app[3]—when prompted by these robots, we installed apps on a factory-reset Google Pixel 7, and took screen recordings of our interactions with these apps. During robot setup, we created fresh user accounts for each robot as necessary using a new study-specific email address, and consented to all terms and conditions or prompted permission requests (e.g., for geolocation access). We connected the robots and the Pixel 7 to a fresh, partitioned laboratory LAN configured to record all network traffic. All interactions were conducted in a private, secured laboratory environment, with only authors present during interactions. We performed our interaction tests starting in Spring 2024.

3.2.3 Robot Interaction and Auditing Manufacturer Claims. To assess the manufacturer claims in our codebooks, we manually interacted with each robot and recorded video of the results. We began each robot interaction by following each product's "getting started" guidance. We followed steps provided by each manufacturer and installed apps or created accounts as prompted. We then followed instructions on-device or in-app until we reached apoint where we were no longer given explicit steps to follow, or were otherwise unable to proceed. In three of the four robots, operational problems prevented initial interactions. These issues were eventually resolved for Moxie and Eilik, but could not be for Vector. We discuss these problems further in § 5.

After setup, we engaged in purposeful interactions with each robot to explore whether each claim in the corresponding codebook was valid (e.g., whether the robot actually included the claimed functionality or capability). Occasionally we encountered claims that we were unable to validate using the procedures and instructions supplied by the robot manufacturer. In these cases, we attempted to trigger robot behavior with interactions and user inputs that were outside-the-norm or exceptional. We spent no more than a few minutes searching for evidence of the validity of each claim, as the objective of this interaction round was to simply identify product claims that were true, not to extensively stress-test claims.

We were intentionally generous when annotating a claim as valid or not, to avoid any negative bias against our robots. In total, two authors manually annotated 64 claims for validity after our filtering, resulting in overall agreement of 98.4%—from 99.2% positive agreement and 0% negative agreement (authors disagreed on one label out of all 64). Given the small size of the dataset, the authors discussed discrepancies to resolve disagreements.

Assessing Privacy Claims Miko and Moxie's manufacturers made privacy-related claims about their robots in our collected materials, noting third-party certifications for compliance with privacy regulations. As conducting full legal compliance audits per-robot was not in scope for this study, we turned to a simple proxy measure of privacy rigor instead. Specifically, we inspected the network traffic data generated by each robot and companion app during the study period, extracted the subdomains contacted by each given device, then resolved these addresses to the organizations owning each domain. We then assessed whether these domains serve operational/necessary purposes for each robot, or serve tertiary, potentially privacy-eroding purposes (e.g., tracking). This assessment was conducted by first seeking exact matches for resultant subdomains against known tracking or advertising subdomains in the latest EasyList [35] and EasyPrivacy [36] filter lists, then manually comparing the identified second-level domains (SLDs) against known trackers from both lists. The full list of resulting domains and corresponding parties or services are listed in Table 10.

[3]Vector documentation suggests that it too requires a companion app, but Vector was not operable.

For subdomains returning vague details, we infer ownership or provided service to the best of our ability from available context like subdomain text or search results, with unknown domains noted in italicized details.

Assessing Intelligence Claims We conducted simple tests to roughly estimate the "intelligence" claims made by robot manufacturers. These tests varied per-robot, given the constraints for interaction set by the robot and the specific claims of intelligence made about the device. At a high-level, we interacted with the purportedly intelligent features of each robot as instructed by the manufacturer, then attempted additional interactions outside of what we were prompted or instructed to see how the robot might handle them, if at all. For voice-controlled features, we attempted to trigger responses with adjacent but incorrect wake words, provided (dummy) personally identifiable information unprompted to test for privacy sensitivity, and made verbatim repeated, off-script, or near-similar queries to test for robots' responses. For facial recognition interactions, we moved in front of robots to test spatial responses and facial tracking. More rigorous AI tests (e.g., to fully probe the capabilities and limitations of robot features that leveraged generative AI models) were out of scope for this study given the range of AI features to test and diversity in robot claims. That said, we believe this represents an important area for future work in AI auditing.

3.3 Evaluation Through Consumer Reviews

Drawing on methods from prior work (see § 2.2), we turned to consumer commentary to answer RQ 3 by inspecting whether users of these robots felt the products delivered on marketing promises. Specifically, we collected user reviews from product pages and each robot's Amazon listings. We collected only English-language reviews that were posted between January 1-August 31 2024, i.e., the rough time period of our study. On Amazon, we only collected verified product reviews, i.e., from people who actually bought the robots, and retained all reviews available directly from manufacturers' pages. We did not include content described by a manufacturer as "testimonials" as these are intended to convey only positive opinions of a given robot. This data collection was ruled exempt by our institutional IRB as all data was publicly available. In total, our dataset contained 168 reviews across all four robots. Again using iterative thematic and document analysis methods [10, 12], we coded this dataset according to the primary subjects of both complaints and praise provided in the consumer reviews.

4 Results

We now analyze the results of our multiple annotation rounds and interaction tests.

4.1 Manufacturer Product Claims

In this section we present the results of our thematic analysis, characterizing the claims made by manufacturers to prospective and new consumers of their robots. Such claims form the "source-of-truth" by which users set expectations for these robots. Overall, Purported Features was the largest category with 63% (N=109) of all claims; the other three follow at 17% (N=28) for Safety and Compliance claims, 11% (N=20) for Outcomes or Value Propositions, and 9% (N=17) for Product Requirements. Figure 2 provides the percentages of claims per category, per robot.

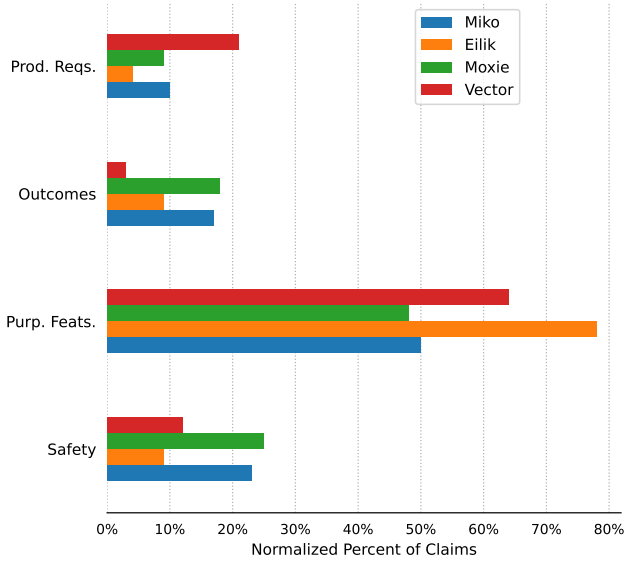


Figure 2: Bar chart depicting the percentage of claims by claim category, normalized as the number of claims per category divided by total claims for that robot.

4.1.1 Purported Features. We find 109 purported features advertised across all robots, stratified into features included out-of-the-box, promised future features, and hardware features.

Claims regarding **Included Features** (N=86) took a variety of forms, ranging from explicit references to feature names or broader statements on types of features. Some included features, particularly for the Eilik, were communicated solely by name alone (though often accompanied with diagrams or other graphics to showcase the feature). For example, the Eilik user manuals and packaging provided lists of features describing what the Eilik does (e.g., *"Dance to Beat/Music"* [E22,38]), what it contains on-device (e.g., *"Heart Mode (Default), Rich Expression, Emotion Engine," etc.* [E27,32,08]), and traits of the robot (e.g., *"Sensitive to Quake/Touch, Afraid of Heights"* [E20,33,34]). Similarly, all four sides of the Miko box provided a list of logos for branded content libraries (specifically, *"Disney, Paramount, Cosmic Kids, Da Vinci Kids, KidloLand, KiDoodleTV, LingoKids"* [MI17]). When presenting these logos, the two sides and back of Miko's box prepend the logos with *"Explore premium content from the world's best kids brands"* [MI18], but the front of the box simply provides brand logos with no additional context.

11 claims pertained to **Feature Interactions**—that is, features described jointly with explicit instructions on how to access, activate, or otherwise use them. For example, *"There is a vibration sensor on Eilik's head. Thus he will feel dazed when you hit him on the head."* [E60] indicates the availability of Eilik's head sensor and explicitly directs the consumer on how to access this feature. We attempted to test all feature interaction claims when possible.

Future Feature claims (N=6) promised additional features not presently available to the consumer. Some such claims described

general and ongoing development, like an *"ever-growing content platform"* (Miko, [MI04]) or *"regular updates"* for *"[becoming] more entertaining, [getting] more expressions, [games, and plot content]"* (Eilik, [E02]).[4] Such claims may risk overselling a robot's immediate value to a consumer, in the event that new features are never added or do not satisfy consumer expectations. Conversely, the Eilik documentation mentioned concrete future features, contingent on purchasing additional Eilik robots. These claims ranged from implicit nudges (e.g., *"Alone, what Eilik can do is limited. Together, they have infinite possibilities"* [E17] or *"Eilik loves to play with his own kind"* [E03]) , to explicit suggestions to acquire more Eiliks (*"Gather three or more Eilik to...enjoy more fun with your friends"* [E18]) .

While a hardware audit was out-of-scope for this paper, we found that product documentation provided **Hardware Specifications** in six claims, with all four robots mentioning hardware details at least once. Such information may have a tertiary effect on a consumer's impression of product value (e.g., more complex sensors might signal a more sophisticated product). The robots varied in the level of hardware detail provided to the consumer. Miko describes *"state-of-the-art sensors [for a human-like personality]"* [MI05] and *"hardware built for worry-free playtime"* with *"every inch [designed to last]"* [MI16]. Vector presented hardware specifications as upgrades over previous models, particularly for *"Improved Camera Resolution"* [V02] from a new 2-megapixel camera and a new battery providing *"Increased Battery Life"* [V17]. Though comparatively vague in its other product claims, Eilik's hardware specification claims were surprisingly detailed, describing *"four EM3 servos [that were] designed with the clutch"* [E19] for joints capable of withstanding more external force and having greater durability. Lastly, Moxie content noted that it was a *"soft touch robot with gesticulating arms, self-swiveling torso, emotion-responsive HD camera and GPT-powered AI"* [MO04].

4.1.2 Safety and Compliance. Safety and compliance claims were those that demonstrated adherence to regulation or concrete statements about safety and privacy. We identified 28 such claims across all robots, which we divided into three subcategories.

Hardware Compliance statements (N=15) generally corresponded to legally required disclosures, such as for the Federal Communications Commission (FCC), European electronic waste regulations (WEEE), PSE Certification, and EU Conformity requirements. **Consumer or Purchase Protections Statements** (N=8) were those describing issues like product warranties, check-out protections, or purchase benefits.

Five claims concerned **Privacy and Safety**, like privacy certifications or other assurances for data safety. For the Moxie, this included a contextless assertion (*"KIDS PRIVACY CERTIFIED BY PRIVO"* [MO43] on the back of the product box), FAQ-styled content on the product website (*"Will Moxie spy on me? No. Video Data is processed locally on Moxie and is used only to create facial expression assessments. Processed locally means the video data is never transmitted beyond Moxie."* [MO07]), and compliance or certification information (*"Moxie, SocialX, and its full ecosystem is COPPA (Children's Online Privacy Protection Act) Safe Harbor certified so parents can feel safe knowing that Moxie employs leading data integrity and security procedures and that its systems are regularly audited to ensure full compliance."* [MO44]). Miko product documentation described security measures in near-verbatim text on both product webpage and packaging, specifically noting *"[a] closed system with enhanced encryption ensures that every byte of your family's data is protected"* [MI29]. Versions of this claim appeared along with *"Miko 3 is KidSafe COPPA CertifiedTM"* [MI29] on the box and the taglined header *"Serious about your family's security"* [MI30] on the product page.

4.1.3 Outcomes or Value Propositions. Across all four robots, 20 claims communicated general consumer outcomes or product value rather than describing specific features that consumers might encounter in-experience.

Thirteen claims described **Purported Outcomes** from using a robot. Some of these claims included usage statistics, like percentage increases for speaking proficiency, physical activity, and academic engagement for children using Miko for three months, or a percentage of children noting *"improved social skills after playing with"* Moxie [MO10]. Twelve claims corresponded to either the Miko or Moxie, with the Eilik providing the thirteenth claim of *"[bringing] up a higher level of social interactions between humans and robots"* [E01].

Umbrella Statements (N=7) were broad, unspecific claims. These were generally found on product pages and packaging as taglines under or near the robot name, and could be perceived as marketing language. We noted such statements for three robots, with five claims made for Eilik and one each for Miko and Vector. The front of Miko's box describes it as *"[the] Ridiculously Smart, Seriously Fun Kids Robot"* [MI03], while Eilik's box front describes it as *"[a] Little Desktop Companion with Endless Fun"* [E10] and *"ONE OF ITS KIND"* [E11]. Vector's touted it as an *"AI ROBOT COMPANION"* [V32]. Further, the Eilik manufacturer made other broad statements further down on its product webpage: *"Tech and robots are advancing faster than ever to make our life efficient, but something important is missing: the emotion, the heart"* [E12] and *"There are countless robot pets in the world. But most of them are inelegant. How to find an endless fun companion robot pet? Your Robotic Pet Awaits You"* [E13].

4.1.4 Product Requirements. We noted 17 product requirement claims across all robots, dictating parameters for robot use, like hardware or networking requirements, or age restrictions.

We considered **Hard Requirements** (N=10) to be non-negotiable steps or possessions asked of consumers in order to operate each robot as intended. These requirements included both hardware and software items, like Wi-Fi connections for three robots (Moxie, Miko, Vector [MI13,MO13,V14]), mobile app use (Moxie [MO13], Miko[MI12], and to a lesser set-up only extent, Vector [V12]), and even the acquisition of a power adapter (the Vector robot came only with a power cable, though the lack of adapter is explicitly noted in the manufacturer's documentation [V15]).

After noting that the Eilik did not declare any requirements for use, we re-examined source documents once more to affirm that none mentioned what would be included with the robot. With the robot, we received a power cable and magnetic attachments, but

[4] In a claim found beneath a *"What's Included"* header on its product webpage, the Vector also notes *"planned software development"* to improve object and facial recognition [V02]. We categorized this statement under **Hardware Specifications** based on the primary purpose of the full text snippet, but mention this here to note that the Vector also promises ongoing feature additions.

	Eilik	Miko	Moxie	Vector
Smartness	6	9	11	11
Anthropomorphism	19	12	1	18

Table 3: Number of claims involving a given attribute.

no power adapter. Thus, we found that the two desktop-oriented robots did not provide wall adapters for consumers, and that the Eilik was comparably less transparent in its lack of wall adapter than the Vector.

For networking requirements, Vector [V11] explicitly required 2.4 GHz Wi-Fi connectivity on packaging and webpage sources. Miko, however, requested network connectivity and the downloading of a mobile app in the second-person (*"I need your Wi-Fi password to connect to your Wi-Fi"* [MI13] followed by *"Download the Miko 3 app from the App Store/Google Play"* [MI12]) through a sticker placed on the robot's main screen (where the digital face would later appear), thus potentially addressing either the child user or their guardian. Finally, we categorized suggested age ranges for users as **Soft Requirements** (N=7). These were commonly on product box-fronts, e.g., *"Ages 5+"* for Miko/Moxie [MI14, MO14], *"Ages 12+"* for Eilik [E16], and *"14+"* for Vector [V33].

4.2 Product Claim Characteristics

4.2.1 "Intelligence" and AI Capabilities. Beyond qualitative categories, we also labeled each claim for mentions of "intelligence," "smartness," or AI functionality. Table 3 shows the number of such claims that we found for each robot. Categorically, these smartness and intelligence propositions appeared in nine *Outcomes* claims and 28 *Purported Features* claims. The content of these intelligence claims cover a spectrum of detail. Least informative was Eilik, which touted an "Emotion Engine" [E06,08], "emotional intelligence" [E01], and other emotional or expressive [E04] capabilities, but did not include the terms "AI" or "artificial intelligence." Miko claimed to be an "AI robot" [MI9] but otherwise only mentioned having a "superpowered" [MI4] or "advanced" [MI2] brain. Purported Vector AI features include those supported by computer vision [V5] and facial recognition, [V2] with other smart features broadly describing how Vector interacts with user input or physical surroundings. Finally, Moxie claims described "GPT-powered AI" [MO4] with features "enabled by generative AI, natural language processing, and computer vision" [MO5], thus providing some of the most detailed AI explanations, with one claim [MO33] specifically describing AI benefits. Specifically, this claim raised the question "Is AI good for kids?", with the provided answer discussing the benefits of AI in supporting a child's social learning experience (as outcomes from using a robot like Moxie).

4.2.2 Anthropomorphic or Personified Descriptions. We labeled a claim as "anthropomorphic" in communication style if it personified the robot (e.g., 'he/she wants to talk with you'). We excluded claims that only described anthropomorphic features or capabilities concretely (e.g., 'the robot has conversational capabilities'). Table 3 shows the number of "anthropomorphic" claims that we found for each robot, which span 39 *Purported Features*, 3 *Product Requirements*, and 8 *Outcomes* claims. Notably, Moxie claims constituted the lowest percentage of personified claims by robot (2%), in contrast to the majority of Vector claims (55%) including personified language.

4.3 Manufacturer Claim Validity

In this section we present the findings of our product claim validity annotations and tests—in other words, which promises were and were not fulfilled by manufacturers. Additionally, we discuss operability expectations and other contradictions that we observed during our live interactions, highlighting the challenges and quirks of evaluating commercially-available robots.

We investigated and tested 64 claims across the Miko, Moxie, and Eilik, of which we could manually verify 98% (63 out of 64). This high rate of validation is somewhat expected, given our intentionally generous annotation procedures. That we find high rates of compliance between manufacturers' product claims and observable product functionality is good news for end users of these robots, but not without caveats.

We now discuss the single claim we found difficult to validate from our subset. For Miko, we sought to verify a "surprise" feature described in a packaging insert—if enough "gems" were collected on-device, we would be given a prompt to reply with a "secret" command and thus earn an "adventure" as additional content with Miko [MI23]. We could not successfully trigger this behavior during our tests, despite attempting to do so multiple times. The claim text left the exact quantity of required gems ambiguous, thus our lack of feature discovery could have been due to longitudinal or volumetric factors (e.g., we may not have interacted with the robot enough for this gamified threshold). In contrast, other claims within the *features with interaction triggers* category were immediately testable. Similarly, it was also unclear what the intended "secret adventure" should have been; when testing the secret command unprompted, Miko replied "you've completed your secret adventure of dancing and grooving with me. I have many more fantabulous adventures for you." This did not inspire confidence that we had triggered the right feature, that the feature existed, or that Miko understood the parameters of the promises in MI23 correctly. As such, we considered this case inconclusive.

4.3.1 Operability Issues That Prevented Auditing. We were unable to audit several claims made by the robot manufacturers due to scope restrictions or resource limitations of our study (e.g., statements pertaining to hardware and electronic specifications that were beyond our capability to test, or sweeping marketing language that was not suited to specific feature tests). There were other claims, however, that we could not audit due to reasons that were disadvantageous to consumers. Chief among these were all claims pertaining to the Vector, which we could not evaluate because the robot had become totally inoperable.

Nine Eilik claims were explicit references to multi-Eilik features and could not be tested with our single robot. Eilik was the only robot touting multiple robot features in its claims. Four additional Eilik claims could not be properly evaluated due to another consumer-disadvantageous informational issue: misleading representation of robot capability. To illustrate, Eilik presented product features in ways that were relatively more robust than the other robots: it explicitly enumerated available features while the others instructed

users to begin interacting and discover features through hands-on experience. Eilik enumerated fourteen available features on page four of its included user guide, and we initially marked all fourteen for validation tests as this page appeared to suggest that all features were available out-of-the-box. However, in later pages we found that "Eilik Theater, Crime Patrol, Dominoes," and "Song and Dance" [E25,37,39,41] were in fact multi-Eilik features that required multiple robots to activate, contradicting the earlier presentation of these features.

4.3.2 Privacy Claims. Only Miko and Moxie were connected to the Internet (via Wi-Fi) during use. We use network traffic data collected during our live interactions to examine the unique resolved domains and subdomains that each robot contacted (i.e., not through companion apps).

Moxie contacted only nine unique resolved addresses, while Miko contacted 60. For Moxie, these addresses corresponded either to Google domains that support robot operations[5] or Embodied, Inc.'s own servers.[6] Miko contacted operational Google domains and their own servers, but also contacted content partners' servers not only for media but for what appear to be tracking or other analytics.[7] MixPanel, Segment.io/Segment.com, and StackAdapt are known analytics companies,[8] while Kidoodle is one of the content partners listed by Miko.

Both robots assert COPPA certification through third parties: KidSafe (Miko) and PRIVO (Moxie), respectively. Under COPPA, tracking and analytics are permissible, but with constraints. Thus, we cautiously consider both robots to deliver upon their COPPA-safe privacy assertions (N=3), but remain skeptical of the true privacy stance of robots that contact analytics partners like Miko. In the absence of additional privacy-forward measures like conspicuous disclosures, consent flows for data collection, or easily accessible opt-outs, such data collection risks exploiting social robot users for the rich data their interactions can provide.

4.3.3 "Intelligence" and AI Capabilities. We attempted to broadly explore "intelligent" or AI capabilities in the three operable robots. This resulted in three very different experiences of product sophistication. First, Eilik delivered upon its limited "Emotion Engine" as-described, given that all expected responses were explicitly documented in the user manual flow diagram. Second, both Miko and Moxie provided voice-recognition functionality as promised. However, the two voice-controlled robots differed greatly in their approaches to human conversation and interaction. Moxie allowed for very little user freedom in interactions, as user-initiated voice interactions were limited to only a few wake phrases detailed in the user guide. Outside of these wake phrases, conversations were driven by the Moxie and were restricted to the range of topics brought up by Moxie within the scope of a "mission" (specific conversational lessons or programs designed with kids' learning objectives, supplemented by a physical workbook provided by Embodied, Inc.), or prompted through tasks in the provided workbook. Conversely, Miko permitted greater user freedom in directing the human-robot experience, with the user prompting interactions through the "Hello, Miko" wake phrase and a subsequent open-ended query.

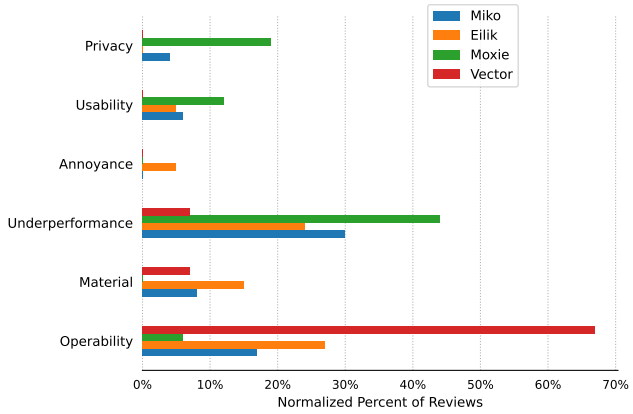


Figure 3: Bar chart depicting the percentage of reviews by complaint category, normalized as the number of reviews per category divided by total reviews for that robot.

Our observations roughly correspond to the manner in which "intelligence" claims were made by robot manufacturers. Eilik's "Emotion Engine" offered the full range of its comparatively simple, limited functionality (though its vague "emotional intelligence" claim is otherwise debatable), while Moxie and Miko were more descriptive with their intelligence statements. However, more rigorous research is needed to better compare sophistication across social robots. Additionally, we note that the level of intelligent interactivity offered by robots may not correspond neatly to their cost, i.e., Miko offered the most freedom in voice interactions but Moxie was the most expensive robot in our cohort (see Table 1).

4.4 Consumer Review Results

Consumer feedback, particularly those in public reviews, articulate user perspectives of a product experience. Positive reviews and comments affirm customer satisfaction,[9] whereas negative commentary directly highlights misalignment between consumer expectations and the provided experience. Note that consumers reviews are prone to negative bias [66, 80], thus in the following sections we emphasize qualitative findings but provide numerical measures for added context. We thus present the most common subjects for criticism and praise in this section.

4.4.1 Complaints. 120 reviews contained at least some critical or plaintive feedback for robot manufacturers. We consider these to be rough proxies for unmet consumer expectations, as they describe areas of consumer dissatisfaction. We present the breakdown of complaints from our content analysis of consumer reviews in Figure 3 as the percentage of a given robot's reviews complaining about each code. Table 4 presents the top complaints for each robot.

Consumers noted **Operability** issues like those we encountered in 44 total reviews. Generally, these complaints discussed robots not working at all, not updating their software, not connecting to

[5] e.g., speech.googleapis.com and connectivitycheck.gstatic.com.
[6] e.g., mqtt.embodied.com for their MQTT IoT messaging service and embodied.me.
[7] e.g., analytics.kidoodle.tv, api.mixpanel.com, app.segment.io, app.segment.com, tags.srv.stackadapt.com, app-measurement.com, and firebase-settings.crashlytics.com.
[8] e.g., crashlytics.com and app-measurement.com, which are Google entities.
[9] Of course, when these are made by real users and not bots.

Complaint Type	% Reviews
Underperformance	44
Privacy	19
Usability	12
Operability	6
-	-

(a) Moxie

Complaint Type	% Reviews
Underperformance	30
Operability	17
Material	8
Usability	6
Privacy	4

(b) Miko

Complaint Type	% Reviews
Operability	27
Underperformance	24
Material	15
Annoyance	5
Usability	5

(c) Eilik

Complaint Type	% Reviews
Operability	67
Material	7
Underperformance	7
-	-
-	-

(d) Vector

Table 4: The top five complaints raised by consumers, normalized as the percent of reviews per-robot.

Wi-Fi, or not charging. 18 unhappy Vector customers reported not being able to use the robot at all (comprising over 60% of all Vector reviews in our dataset), as did eight Eilik users and one Moxie user. Vector users responded to operability issues emotionally at times (describing inoperability as one of their *"saddest moments"* [VE27] or *"frustrating and upsetting...especially for children"* [VE06]), with users noting how much they missed the robot [VE10, 27]. Despite Miko setup being comparatively smooth in our testing, multiple Miko consumers expressed frustration with failed software updates preventing use–Eilik users also noted failed updates [EI26] or failure to even turn on [EI23]. Others described operability issues as rendering the robot into an expensive "paperweight" [MK68].

Even if robots appeared to be functional for users, consumers expressed frustrations with the quality of voice-recognition features, robots being more "boring" than expected, lackluster battery life, or other examples of robots otherwise under-delivering on their product experiences. We considered these as **Underperformance** complaints (N=44). Consumers recognized that robots were limited in functionality (e.g., requiring internet connectivity for most functionality [MK34] or otherwise wanting more from a robot [EI22, VE21]). Some consumers explicitly described cost-benefit analyses, for example describing Miko pricing as "excessively high relative to the value provided" [MK16] or noting that "very little is included...given the cost" [MK63]. Multiple Miko users described the robot as tablets or voice assistants "on wheels" [MK58, 69, 70] while a Moxie user asserted that Moxie should be sold at half price [MX09]. Similarly, **Usability** (N=9) complaints described robots as generally being difficult to use.

Annoyances arose when consumers complained about specific behaviors they disliked in a robot, without the consumer mentioning operability issues or otherwise implying that this fell below standards or expectations. We retained these as distinct from the other two operations-related categories, as annoyance complaints were less about manufacturers' failure to deliver and more about consumers' personal tastes. Two Eilik complaints fell under this category, describing the robot as "a little noisy when moving" [EI20] or noting the robot was "crying a lot" [EI32].

Consumer opinions primarily expressing **Privacy Concerns** were relatively uncommon, appearing in only six of all 168 reviews (3 for Miko, 3 for Moxie). All three Miko complaints pertained to used devices, with prior user credentials still configured into the robot. One user [MK56] described needing to contact Miko customer service in order to remotely wipe the prior owner's information. Another user [MK43] discovered in the Miko app that their own returned Miko was accessed by a new consumer. One Moxie review cited concerns with the use of Google APIs for data processing, while two described it as "creepy": one for its dependence on interactions to increase skills [MX02], and another for the manner in which Moxie's animated eyes follow the user around a room [MX10].

Material (N=15) complaints were those pertaining to either the physical acquisition of the robot (e.g., problems with purchasing or shipping experiences) or concrete financial issues (typically subscription-related). Two would-be Vector consumers on the manufacturer page [VE02, 03] expressed interest in purchasing the robot, but complained that the the robot was out of stock. Three consumers [EI11, 19, 25] described receiving Eiliks with parts missing, like charging cables or accessories touted in manufacturer communications (we received both the charging cable and accessories, though we note in § 4.1 that Eilik did not communicate hard requirements clearly) while others reported receiving used products. However, we acknowledge that shipping issues may be out of manufacturers' control and depend on where consumers source their robots. Five reviews raised issues with robots' subscription models restricting features [MK36, 57, 60], two of which additionally described ongoing billing for canceled subscriptions [MK46, 73]. MK73 in particular laments Miko's chat support: *"It keeps telling me how to go in to the parent app and cancel the subscription but the app says I don't have a subscription. I think this is all intentional to make it difficult to cancel."* Our manual tests also resulted in confusing experiences with deciphering whether we had Miko MAX access or not, free trial or otherwise, affirming the behavior in this consumer report.

4.4.2 Praise. In this study, we consider consumer praise to be a proxy for met consumer expectations. Over half of all reviews (102 out of 168 items) included at least some positive feedback. Table 5 presents the top five targets of praise for each robot. 28 reviews provided unspecific, **General** praise, describing the overall experience with adjectives like "good" [MK17], "ok" [MX16], or otherwise noting that the product was either accepted or liked to some extent.

Aesthetic (N=10) praise described Miko, Eilik, and Vector "cute," (e.g., [EI09,MK57])", "adorable" [VE22,27], or otherwise commended cuteness. **Interactivity** praise (N=29) described robots relationally, as "cyber bud[s]" [MK53], or "treasured friends" [VE06], with some users naming their robots [EI02] and keeping robots in constant company. Reviews praising **Content and Features** (N=20) applauded content volume, variety, and age-appropriateness. Ten **Entertainment and Novelty** reviews praised robots for being exciting, interesting, otherwise entertaining, while two **Customer Service** reviews applauded manufacturers' help with the robot.

Praise Type	% Reviews
Interactivity	38
General	12
Material	12
Entertainment And Novelty	6
-	-

(a) Moxie

Praise Type	% Reviews
Content And Features	24
General	15
Interactivity	13
Entertainment And Novelty	7
Aesthetic	4

(b) Miko

Praise Type	% Reviews
Interactivity	20
General	12
Aesthetic	10
Entertainment And Novelty	7
Material	2

(c) Eilik

Praise Type	% Reviews
General	30
Interactivity	15
Aesthetic	11
-	-
-	-

(d) Vector

Table 5: The top five subjects of consumer praise, normalized as the percent of reviews per-robot.

Material compliments (N=3) commended product quality. The highest-rated review—at five stars [EI18]—was for the Eilik, commending fast shipping and hardware quality. The other two Moxie reviews [MX05, 08] praised quality while also expressing pessimistic views: both consumers expressed frustration with Moxie underperformance by not understanding speech, rating the robot 3 and 2 respectively. Moxie was the most expensive robot in our study, arriving with detailed packaging and additional materials for users, while the Eilik was the cheapest.

5 Discussion

In this study, we qualitatively examine product claims made about four companion robots by their makers, manually audit these claims through live experiments with three robot experiences, and examine consumer reviews to extract user perspectives of robot value. Now we discuss our findings and consider implications for improving consumer protections or leading to future work.

To determine product value, consumers perform some form of cost-benefit analysis. Our manual audit results suggests that consumers may in fact "get what they paid for" with these robots, in the narrow sense that nearly all claims in our test subset made by the manufacturers appear to be true (under our intentionally generous labelling criteria). This, however, does not account for the complete lack of access to potential value that product inoperability issues implicate, or the challenge of evaluating more complex claims like privacy certifications or intelligence promises. Consumer reviews, additionally, highlight a broader range of reactions. Our study thus affirms that open questions remain in determining what consumers should or do expect from the current generation of commercially available social robots.

Consumer reviews reveal an inconsistent experience across users of the same robot—subjects of praise in some reviews could be subjects of criticism in others. Our study suggests that consumers do not share consensus expectations of the services robots should provide. For example, as the only device with media offerings (e.g., through games, stories, videos, etc.), Miko's content library was an important product differentiator, but while some consumers praised the volume, variety, and types of content provided, this was also interpreted negatively by users as it made Miko seem like little more than a "tablet-on-wheels." This raises questions about how to design, define, and eventually market social robots to consumers.

We found Eilik and Moxie to present a study in contrasts. One one hand, Eilik communicated its features most transparently among our four robots, but it was also the least expensive robot with the least sophisticated (and thus easiest to explain, questionably "intelligent") capabilities. On the other hand, Moxie was the most expensive and (arguably) most technically sophisticated robot in our cohort, yet users complained that it had little to offer or underperformed—one parent user [MX08] felt Moxie was "useless" and "not smart enough," while others thought Moxie struggled with voice recognition [MX01, 04, 05]. This divergence in technical capabilities, coupled with the challenge of communicating to consumers how to interact with sophisticated systems, may make it challenging for consumers to effectively compare robots—even before considering what constitutes "intelligence".

5.1 Improving Consumer Experiences

We now discuss implications from our study for improving consumer satisfaction and protections in social robot experiences. In particular, we present three themes as opportunities for regulators and researchers (with implications for similar emerging technologies), then discuss how practitioners might be supported in the efforts to deliver ethical UX to consumers.

5.1.1 Improve Explainability in Experiences. Consumer protection law is designed to make sure consumers are not deceived [8]. As part of this mission, regulators generally look to the whole experience consumers have with goods and services, and ask what their "overall impression" would be when considering and using the good or services [86]. This holistic approach to consumer expectations means designers and stakeholders with influence over design decisions must consider how products will be perceived in context. That is, the promised features of social robots need to be described in such a way that consumers' expectations of functionality accord with the device's actual capabilities. This is especially true for robots with "intelligent" or AI features that may contribute little more than "snake oil" [65]—or simply for interactive robots misrepresenting intelligence, period.

We noted earlier that Moxie was comparatively more sophisticated in its intelligence claims and during the experience of interacting with it. However, Moxie's immediate set of affordances are quite simple: users are not able to prompt conversations as they might with Alexa, and conversation topics are steered primarily by Moxie itself. Designing restrictions for what users, particularly children, are able to do may provide some safety. This contrasts vividly with the potentially addictive features we found in Miko (e.g., gamification, unrestricted interaction potential, touchscreen and mobile app elements).

Some users, however, found the Moxie's interaction restrictions to be frustrating. [MX05] directly compared Moxie's ability to communicate to Google Assistant and Alexa (noting that the latter

two performed better than Moxie), while other users [MX04, 08] disliked what they considered to be conversational interruptions, incorrect responses, and ignorance of their child's interactions. MX08 in particular described Moxie as "not smart enough" as it would "change the subject" or reply that it lacked training to respond to a user's query, going so far as to call Moxie "useless" in comparison to Siri on an iPad. Such frustrations might be avoided pre-emptively with clearer descriptions of social robot's limitations (especially those touting "intelligent" functionality). For example, the U.S. Federal Trade Commission urges businesses to "keep their AI claims in check" and to avoid exaggerating what an AI feature can and cannot do for consumers [4, 5, 48]. Despite these suggestions (published by the FTC in 2021 [48] and 2023 [4, 5]), our study suggests that AI "snake oil" continues in 2024 insofar as consumers struggle to reason through what intelligence means in their "smart" electronics.

The human-computer interaction discipline is uniquely positioned to construct more explainable delineations between smartness, AI, and related terms consumers must face. For example, Recki et al. [75] draw on the E.U. AI Act to provide a conceptual model for users' risk perceptions of AI in order to bridge the design and policy fields. Our study highlights the need for such work. Our study also highlights the importance of incorporating manufacturers' product materials into future scholarship, to better capture the types of language faced by users *in situ*.

5.1.2 Account for Social Robots' Compound Vulnerabilities. Consumer protections enforcers are also concerned about vulnerable consumers in the marketplace. Factors such as age, mental and physical ability, sophistication, consumer necessity, likelihood and magnitude of harm, and power imbalances all contribute to the amount of care companies must exercise in marketing and designing their products. This is particularly true when a consumer is vulnerable in more than one way, such as a child being exposed to anthropomorphized technology: not only are children less experienced and more susceptible to influence, but all people are generally more susceptible to influence via anthropomorphic tools. This makes the study of compounding vulnerabilities critical for understanding modern consumer protection issues. In this section we discuss two particular areas of concern regarding social robots' compounding vulnerabilities: robot death and children's use of anthropomorphic technologies.

Robot Death. The social robot context results in the unique intersection of emotional harm and material harm when a robot "dies," whether the cause of death be from software or hardware failure. The Vector was found to be inoperable in our tests, a prior issue which was not communicated clearly by the manufacturer. As of December 2024, after the completion of our study, Moxie's manufacturer announced imminent shutdown, leaving consumers with mere days to grapple with the loss of their robot companion [34].

Regulatory regimes that give consumers greater control over their electronics—e.g., right-to-repair or restore rules—may somewhat mitigate *material harms* from operability issues with social robots. For example, one reviewer in our study restored functionality to their Vector using third-party, open-source software [VE01]. However, effective right-to-repair regulations must be paired with technical tools that are accessible to consumers and reasonably easy to use. Alternatives that require deep subject matter expertise fall short. To avoid situations like Vector and Moxie's service failures, regulators may want to assert redundancy requirements (as are common for cloud-based processing) for social robot functionality. Alternatively, operability might be at least partially or temporarily guaranteed even after final support dates by interested third parties or by off-cloud functionality [21]. In Moxie's case, a late December 2024 over-the-air (OTA) update was released, preparing the robots for a potential open-source, local solution for ongoing functionality–such efforts may be one path forward for sunsetting embodied or antropomorphic companions after company closure [1, 73].

The Vector and Moxie shutdowns also implicate *emotional harm*. Products that are social by design evoke potentially deep attachments from the user onto the device. We observed this in user reviews, particularly those for the Vector that describe the loss of a "friend." This opens the potential for severe emotional fallout when the robot companion "dies" [50], similar to cases involving AI chatbots [77]. To mitigate against emotional harms, regulators must provide guidance for how manufacturers should appropriately and safely "sunset" highly social devices and provide avenues for consumers to handle the ensuing emotions effectively, thus managing product obsolence and maintenance beyond primarily technical implications. Even in the absence of clear and formal guidance, manufacturers of social robots or similarly engaging technologies should not ignore the potential for emotional fallout and design both user experiences and technical resilience measures to minimize non-material distress. Consumer advocates have recently urged the FTC to do exactly this, using the term "software tethering" to describe the cloud dependencies of smart devices that make them vulnerable to "bricking" [17]. Future research might investigate the impact of embodiment versus conversational capability on human-robot attachment, which may inform regulatory or design priorities in the social robot or AI market.

Existing regulations are insufficient for wrangling emotional vulnerability in consumer technologies, and do not address additional implications from embodied, anthropomorphic interaction modalities. The E.U. AI Act, for example, stratifies AI systems into levels of risk, and considers emotional state assessment under unacceptable (high) risk—but only does so within the potential for discrimination in workplace or education decision-making. Where consumer social robots fall in the AI Act's hierarchy remains unclear, and thus their potential for harm should be addressed with greater specificity. One suggestion is to classify attachment-invoking social robots in a similar manner to how the AI Act has stratified AI uses by risk. Enumerating and articulating a given social robot's embodied or emotive interaction modalities (e.g., anthropomorphic features, conversational capabilities, visual presentation, method of prompting engagement, etc.) could be used not only to stratify robots by emotional risk but also to set guidelines for how manufacturers communicate social robot features.

Children and Social Robots. Children are especially vulnerable to emotional harms—not only from robot deaths, but from the social experience that companion robots provide. Such concerns have already been raised by regulators. For example, the Italian data protection authority (DPA) noted that chatbots marketed for improving users' moods "may increase the risks for individuals still

in a developmental stage or in a state of emotional fragility" [74]. The Italian regulator also noted a lack of age verification systems. We studied two robots specifically targeted to children (Miko and Moxie), both of which did include parent or guardian confirmation—this is an encouraging sign that some manufacturers acknowledge the risk to children from their products. Moxie's aforementioned UX limitations also limit what children are able to do with the robot. Determining a social robot's risk to children depends in part on the extent to which a robot delivers intelligent or embodied features. This exploratory study highlights the need for clearer concepts of robot features both from design and consumer perspectives, and urges future research in this area.

Children may be especially vulnerable to emotional harm through robot deaths. In the case of Moxie's, though Embodied, Inc. provided notice to consumers [34], children's anguished reactions to being told their Moxie will disappear went viral [37]. Reporters noted repercussions from the pain of having emotional bonds suddenly taken away [37, 44]. Cognizant of this particular vulnerability, Embodied, Inc. provided a support letter from Moxie's fictional world to help parents explain Moxie's disappearance in an "age-appropriate way" [34]. HCI research should consider both pre-emptive and reactive UX designs for handling robot termination. Social robots targeted at children could design for potential death in advance by including off-boarding interactions to help balm the loss, beyond documentation or written guidance (thus not exclusively having parent consumers take the brunt of managing emotional fallout).

5.1.3 Empower Practitioners Within Ethical Design Complexity. While in the past, regulators looking for deceptive practices have focused on advertisements, marketing claims, and terms in boilerplate contracts, new experiential technologies like social robots might shape consumer expectations in ways that regulators do not yet fully appreciate. Positioned at the frontlines of the consumer experience, designers are uniquely trained to understand both the potential and limitations of a robot product as well as how it might be eventually used. However, promises made in product materials might not fall directly under UX teams' control, which complicates the ability for ethically-minded UX professionals to ensure that other organizational stakeholders do not oversell features or capabilities. Moreover, claims made by marketing or packaging teams may directly contradict designer intention and (when overpromising) might be a primary factor in consumers' negative impressions of the UX instead of the designs themselves.

Our work supports UX practitioners' "soft resistance" [95] and ethical mediation [25, 40] within these organizational structures, for better ensuring how users perceive their experiences of social robots and related consumer electronics. UX teams might leverage their expertise in affordance theory [39, 67] to predict consumer expectations of different social robot features. Practitioner documentation of social or companionship-oriented features' affordances could provide technical rigor in understanding misalignment between product claims and resultant experiences.

Pre-empt "Dark Patterns" and Deception in Social Robots. In practice, designers' and UX teams' soft resistance could include "futureproofing" against dark patterns (deceptive or manipulative UX designs that influence user behavior against their interests) in social robots, thus designing against the potential anthropomorphic abuses highlighted in prior work [7, 32, 55, 79, 84, 88]. Recent ontological work stratifies dark patterns into different levels by their design strategy, explicitly considering "meso-level" patterns as those that subvert user expectations [41]. Note that in this subsection we cite dark patterns with their Gray et al. [41] ontology level in superscript for easy referencing.

Though systematic dark patterns identification was not in scope for this study, our findings support future investigations into dark patterns in social technologies. For example, consumers describe having to contact manufacturers in order to remove either their own or prior users' data from robots [MK43,55,56], relating to privacy dark patterns [9]. Another user raised similar concerns about knowing where robot interaction data might be sent [MX15], implicating the "data myopia" identified in prior work on dark patterns in robots [55]. Social robots also make consumers vulnerable to financial risk. Vector's inoperability relates to *Roach Motel*^{M} or *Hiding Information*^{M} dark patterns insofar as product messaging failed to transparently disclose the true status of the product (the manufacturer product page merely listed the robot as "sold out" without disclosing diminished functionality). Similarly, inconsistent subscription claims relate to *Roach Motel*^{M} and *Hidden Costs*^{L} dark patterns. This lack of transparency or conspicious placement may have motivated consumers' complaints. Subscription-based pricing *in addition to* up-front device acquisition costs also obfuscates the true cost of owning a social robot especially if prices change over time or certain functionalities are paywalled. Other potential forms of cognitive/attentional dark patterns may include annoyances like *Nagging*^{M}. For example, EI20 and EI32 complained that Eilik was "noisy" and "crying too much". That said, Eilik packaging explicitly claims [E5] that it "is eager to get your attention at every second". This raises questions for future work: is a dark pattern still problematic if the design is nominally disclosed?

Expectations are set not only by the design affordances built into the user experience, but also by manufacturer claims and promises. Thus our work supports future research on dark patterns in embodied and social contexts, particularly for the purchase experiences of materially-oriented technologies (e.g., IoT devices and smart robots) within the overall device experience.

5.2 Limitations and Future Work

Robot Sampling. With our small sample size and narrow corpus definition, our results are not directly generalizable to all types of social robots. This is due to the limited set of available robots in addition to our filtering criteria. Future work could compare and contrast robot types (for example, differences between robots of similar sophistication), or investigate robots in non-U.S. markets.

Consumer Reviews. This study examines consumer reviews to explore user perspectives during the study period. However, it is known that users tend to leave reviews when dissatisfied, leading to potential bias [47, 66, 80] towards negative reviews and thus providing a limited view of the true user experience. We de-emphasize quantitative measures in § 4.4 and focus instead on qualitative findings. Future participant studies could build upon this study's themes to better understand how consumers interpret the promises made by social robot manufacturers. Consumer reviews also provide a limited, unstructured view into consumer perceptions. Future work

could further explore user expectations of different levels of robot sophistication or available features.

Longitudinal Analysis. Within the course of our study, manufacturers updated their product pages to add or subtract claims, emphasize some features over others, or otherwise change the content communicated to a prospective customer. In Miko and Moxie, product language was updated specifically to tout extant AI features, with the Moxie robot's product name even changing from "Moxie Robot" to "Moxie AI Companion" between April and July 2024. While assessing longitudinal feature changes on-device or in documentation was not part of this study's original scope, it is an opportunity for future work, particularly if re-branding changes user expectations of a device's capabilities (regardless of how many improvements, if any, were made to the device experience itself).

Auditing Marketing Claims. This work presents an exploratory method for evaluating truthfulness to marketing claims in social robots. Though complete compliance audits were not in scope for this study (e.g., evaluating advertising claims broadly, or breaking down language with methods from the business and marketing disciplines), our study demonstrates the potential for critical human-computer interaction scholarship's expertise in performing pro-consumer audits of emergent technologies. In particular, our study affirms the need for scalable, interdisciplinary, and collaborative methods to better understand user assumptions of robot capabilities in relation to the manufacturer representations.

Acknowledgments

This research was supported in part by a Northwestern CASMI grant and NSF grant (#1955227). Any opinions, findings, and conclusions or recommendations expressed in this material are those of the authors and do not necessarily reflect the views of the funders.

References

[1] Adurbe. 2024. The Moxie child support robot gets new lease on life through open source. https://www.osnews.com/story/141401/the-moxie-child-support-robot-gets-new-lease-on-life-through-open-source/

[2] Benjamin Andow, Samin Yaseer Mahmud, Wenyu Wang, Justin Whitaker, William Enck, Bradley Reaves, Kapil Singh, and Tao Xie. 2019. Policylint: investigating internal privacy policy contradictions on Google play. In *Proceedings of the 28th USENIX Conference on Security Symposium* (Santa Clara, CA, USA) *(SEC'19)*. USENIX Association, USA, 585–602.

[3] Martin Armstrong. 2023. The market for smart home devices is expected to boom over the next 5 years. https://www.weforum.org/agenda/2022/04/homes-smart-tech-market/

[4] Michael Atleson. 2023. Keep your AI claims in check. https://www.ftc.gov/business-guidance/blog/2023/02/keep-your-ai-claims-check

[5] Michael Atleson. 2023. The Luring Test: AI and the engineering of consumer trust. https://www.ftc.gov/consumer-alerts/2023/05/luring-test-ai-and-engineering-consumer-trust

[6] Antonio Bicchi. 2015. Of Robots, Humans, Bodies and Intelligence: Body Languages for Human Robot Interaction. In *Proceedings of the Tenth Annual ACM/IEEE International Conference on Human-Robot Interaction* (Portland, Oregon, USA) *(HRI '15)*. Association for Computing Machinery, New York, NY, USA, 107. https://doi.org/10.1145/2696454.2714388

[7] Piercosma Bisconti Lucidi and Daniele Nardi. 2018. Companion Robots: The Hallucinatory Danger of Human-Robot Interactions. In *Proceedings of the 2018 AAAI/ACM Conference on AI, Ethics, and Society* (New Orleans, LA, USA) *(AIES '18)*. Association for Computing Machinery, New York, NY, USA, 17–22. https://doi.org/10.1145/3278721.3278741

[8] Board of Governors of the Federal Reserve System, Division of Consumer and Community Affairs. 2008. Federal Trade Commission Act Section 5: Unfair or Deceptive Acts or Practices. https://www.federalreserve.gov/boarddocs/supmanual/cch/200806/ftca.pdf.

[9] Christoph Bösch, Benjamin Erb, Frank Kargl, Henning Kopp, and Stefan Pfattheicher. 2016. Tales from the Dark Side: Privacy Dark Strategies and Privacy Dark Patterns. *Proc. of PETS* 2016, 4 (2016), 237–254. https://content.sciendo.com/view/journals/popets/2016/4/article-p237.xml

[10] Glenn A. Bowen. 2009. Document Analysis as a Qualitative Research Method. 9 (2009). Issue 2.

[11] Dionne Bowie-DaBreo, Corina Sas, Heather Iles-Smith, and Sandra Sünram-Lea. 2022. User Perspectives and Ethical Experiences of Apps for Depression: A Qualitative Analysis of User Reviews. In *Proceedings of the 2022 CHI Conference on Human Factors in Computing Systems* (New Orleans, LA, USA) *(CHI '22)*. Association for Computing Machinery, New York, NY, USA, Article 21, 24 pages. https://doi.org/10.1145/3491102.3517498

[12] Virginia Braun and Victoria Clarke. 2006. Thematic Analysis in Psychology. In *Qualitative Research in Psychology*. Routledge. https://doi.org/10.1191/1478088706qp063oa

[13] Aurea Bravo Perucho and Maryam Alimardani. 2023. Social Robots in Secondary Education: Can Robots Assist Young Adult Learners with Math Learning?. In *Companion of the 2023 ACM/IEEE International Conference on Human-Robot Interaction* (Stockholm, Sweden) *(HRI '23)*. Association for Computing Machinery, New York, NY, USA, 355–359. https://doi.org/10.1145/3568294.3580105

[14] Cynthia Breazeal. 2017. Social Robots: From Research to Commercialization. In *Proceedings of the 2017 ACM/IEEE International Conference on Human-Robot Interaction* (Vienna, Austria) *(HRI '17)*. Association for Computing Machinery, New York, NY, USA, 1. https://doi.org/10.1145/2909824.3020258

[15] Cynthia Breazeal. 2022. Emotion, Social Robots, and a New Human-Robot Relationship. In *Proceedings of the Genetic and Evolutionary Computation Conference* (Boston, Massachusetts) *(GECCO '22)*. Association for Computing Machinery, New York, NY, USA, 2. https://doi.org/10.1145/3512290.3543633

[16] Cynthia Breazeal, Kerstin Dautenhahn, and Takayuki Kanda. 2016. *Social Robotics*. https://link.springer.com/chapter/10.1007/978-3-319-32552-1_72

[17] Justin Brookman and Signatory Parties. 2024. Letter: Groups call on FTC to develop clear guidance for use of "software tethering". https://advocacy.consumerreports.org/research/group-letter-ftc-software-tethering

[18] Duc Bui, Yuan Yao, Kang G. Shin, Jong-Min Choi, and Junbum Shin. 2021. Consistency Analysis of Data-Usage Purposes in Mobile Apps. In *Proceedings of the 2021 ACM SIGSAC Conference on Computer and Communications Security* (Virtual Event, Republic of Korea) *(CCS '21)*. Association for Computing Machinery, New York, NY, USA, 2824–2843. https://doi.org/10.1145/3460120.3484536

[19] Bengisu Cagiltay, Bilge Mutlu, and Joseph E Michaelis. 2023. "My Unconditional Homework Buddy:" Exploring Children's Preferences for a Homework Companion Robot. In *Proceedings of the 22nd Annual ACM Interaction Design and Children Conference* (Chicago, IL, USA) *(IDC '23)*. Association for Computing Machinery, New York, NY, USA, 375–387. https://doi.org/10.1145/3585088.3589388

[20] Bengisu Cagiltay, Nathan Thomas White, Rabia Ibtasar, Bilge Mutlu, and Joseph Michaelis. 2022. Understanding Factors That Shape Children's Long Term Engagement with an In-Home Learning Companion Robot. In *Proceedings of the 21st Annual ACM Interaction Design and Children Conference* (Braga, Portugal) *(IDC '22)*. Association for Computing Machinery, New York, NY, USA, 362–373. https://doi.org/10.1145/3501712.3529747

[21] Ashley Carman. 2020. Jibo, the social robot that was supposed to die, is getting a second life. https://www.theverge.com/2020/7/23/21325644/jibo-social-robot-ntt-disruptionfunding

[22] Catherine Caudwell and Cherie Lacey. 2020. What do home robots want? The ambivalent power of cuteness in robotic relationships. *Convergence* 26, 4 (2020), 956–968. https://doi.org/10.1177/1354856519837792 arXiv:https://doi.org/10.1177/1354856519837792

[23] Catherine Caudwell, Cherie Lacey, and Eduardo B. Sandoval. 2020. The (Ir)Relevance of Robot Cuteness: An Exploratory Study of Emotionally Durable Robot Design. In *Proceedings of the 31st Australian Conference on Human-Computer-Interaction* (Fremantle, WA, Australia) *(OzCHI '19)*. Association for Computing Machinery, New York, NY, USA, 64–72. https://doi.org/10.1145/3369457.3369463

[24] Stephen Cave, Kate Coughlan, and Kanta Dihal. 2019. "Scary Robots": Examining Public Responses to AI. In *Proceedings of the 2019 AAAI/ACM Conference on AI, Ethics, and Society* (Honolulu, HI, USA) *(AIES '19)*. Association for Computing Machinery, New York, NY, USA, 331–337. https://doi.org/10.1145/3306618.3314232

[25] Shruthi Sai Chivukula, Chris Rhys Watkins, Rhea Manocha, Jingle Chen, and Colin M. Gray. 2020. Dimensions of UX Practice that Shape Ethical Awareness. In *Proceedings of the 2020 CHI Conference on Human Factors in Computing Systems* (Honolulu, HI, USA) *(CHI '20)*. Association for Computing Machinery, New York, NY, USA, 1–13. https://doi.org/10.1145/3313831.3376459

[26] Claudia Cruz, L. Enrique Sucar, and Eduardo F. Morales. 2008. Real-time face recognition for human-robot interaction. In *2008 8th IEEE International Conference on Automatic Face & Gesture Recognition*. 1–6. https://doi.org/10.1109/AFGR.2008.4813386

[27] Kate Darling. 2024. It's No Wonder People Are Getting Emotionally Attached to Chatbots. https://www.wired.com/story/its-no-wonder-people-are-getting-emotionally-attached-to-chatbots/

[28] Linda Di Geronimo, Larissa Braz, Enrico Fregnan, Fabio Palomba, and Alberto Bacchelli. 2020. UI Dark Patterns and Where to Find Them: A Study on Mobile Applications and User Perception. In *Proc. of CHI*.

[29] Lucas do Nascimento Diniz, José Cezar de Souza Filho, and Rainara Maia Carvalho. 2022. Can User Reviews Indicate Usability Heuristic Issues?. In *Extended Abstracts of the 2022 CHI Conference on Human Factors in Computing Systems* (New Orleans, LA, USA) *(CHI EA '22)*. Association for Computing Machinery, New York, NY, USA, Article 334, 6 pages. https://doi.org/10.1145/3491101.3519659

[30] Paulo Sérgio Henrique Dos Santos, Alberto Dumont Alves Oliveira, Thais Bonjorni Nobre De Jesus, Wajdi Aljedaani, and Marcelo Medeiros Eler. 2024. Evolution may come with a price: analyzing user reviews to understand the impact of updates on mobile apps accessibility. In *Proceedings of the XXII Brazilian Symposium on Human Factors in Computing Systems* (Maceió, Brazil) *(IHC '23)*. Association for Computing Machinery, New York, NY, USA, Article 52, 11 pages. https://doi.org/10.1145/3638067.3638081

[31] Stephanie Druga, Cynthia Breazeal, Randi Williams, and Mitchel Resnick. 2017. "Hey Google is it OK if I eat you?" Initial Explorations in Child-Agent Interaction. In *ACM Conference of Interaction Design and Children (Late-Breaking Work)*. https://dl.acm.org/doi/pdf/10.1145/3078072.3084330

[32] Elizabeth Dula, Andres Rosero, and Elizabeth Phillips. 2023. Identifying Dark Patterns in Social Robot Behavior. In *2023 Systems and Information Engineering Design Symposium (SIEDS)*. 7–12. https://doi.org/10.1109/SIEDS58326.2023.10137912

[33] Amy Eguchi and Hiroyuki Okada. 2017. Social Robots: How Becoming an Active User Impacts Students' Perceptions. In *Proceedings of the Companion of the 2017 ACM/IEEE International Conference on Human-Robot Interaction* (Vienna, Austria) *(HRI '17)*. Association for Computing Machinery, New York, NY, USA, 111–112. https://doi.org/10.1145/3029798.3038424

[34] Inc. Embodied. 2024. Closing FAQs. https://moxierobot.com/pages/closing-faqs

[35] EasyList filter list project. 2024. EasyList. https://easylist.to/easylist/easylist.txt

[36] EasyList filter list project. 2024. EasyPrivacy. https://easylist.to/easylist/easyprivacy.txt

[37] Ina Fried. 2024. Maker of AI robots for kids abruptly shutters. https://www.axios.com/2024/12/10/moxie-kids-robot-shuts-down

[38] Norina Gasteiger, Kate Loveys, Mikaela Law, and Elizabeth Broadbent. 2021. Friends from the future: a scoping review of research into robots and computer agents to combat loneliness in older people. *Clinical interventions in aging* (2021), 941–971.

[39] James J. Gibson. 1979. The Theory of Affordances. https://cs.brown.edu/courses/cs137/2017/readings/Gibson-AFF.pdf

[40] Colin M. Gray and Shruthi Sai Chivukula. 2019. Ethical Mediation in UX Practice. In *Proceedings of the 2019 CHI Conference on Human Factors in Computing Systems* (Glasgow, Scotland Uk) *(CHI '19)*. Association for Computing Machinery, New York, NY, USA, 1–11. https://doi.org/10.1145/3290605.3300408

[41] Colin M. Gray, Cristiana Teixeira Santos, Nataliia Bielova, and Thomas Mildner. 2024. An Ontology of Dark Patterns Knowledge: Foundations, Definitions, and a Pathway for Shared Knowledge-Building. In *Proceedings of the CHI Conference on Human Factors in Computing Systems* (<conf-loc>, <city>Honolulu</city>, <state>HI</state>, <country>USA</country>, </conf-loc>) *(CHI '24)*. Association for Computing Machinery, New York, NY, USA, Article 289, 22 pages. https://doi.org/10.1145/3613904.3642436

[42] Johanna Gunawan, David Choffnes, Woodrow Hartzog, and Christo Wilson. 2021. A Comparative Study of Dark Patterns Across Mobile and Web Modalities. (Oct. 2021).

[43] Omar Haggag. 2022. Better Identifying and Addressing Diverse Issues in mHealth and Emerging Apps Using User Reviews. In *Proceedings of the 26th International Conference on Evaluation and Assessment in Software Engineering* (Gothenburg, Sweden) *(EASE '22)*. Association for Computing Machinery, New York, NY, USA, 329–335. https://doi.org/10.1145/3530019.3535346

[44] Scharon Harding. 2024. Startup will brick $800 emotional support robot for kids without refunds. https://arstechnica.com/gadgets/2024/12/startup-will-brick-800-emotional-support-robot-for-kids-without-refunds/

[45] Woodrow Hartzog. 2015. Unfair and Deceptive Robots. 74, 785 (2015). https://papers.ssrn.com/sol3/papers.cfm?abstract_id=2602452

[46] Quirien R. M. Hover, Ella Velner, Thomas Beelen, Mieke Boon, and Khiet P. Truong. 2021. Uncanny, Sexy, and Threatening Robots: The Online Community's Attitude to and Perceptions of Robots Varying in Humanlikeness and Gender. In *Proceedings of the 2021 ACM/IEEE International Conference on Human-Robot Interaction* (Boulder, CO, USA) *(HRI '21)*. Association for Computing Machinery, New York, NY, USA, 119–128. https://doi.org/10.1145/3434073.3444661

[47] Wonil Hwang. 2023. Online User Reviews: A Treasure Trove of UX Research? *J. User Exper.* 18, 4 (Sept. 2023), 170–174.

[48] Ellisa Jillson. 2021. Aiming for truth, fairness, and equity in your company's use of AI. https://www.ftc.gov/business-guidance/blog/2021/04/aiming-truth-fairness-equity-your-companys-use-ai

[49] David O Johnson, Raymond H Cuijpers, Kathrin Pollmann, and Antoine AJ Van de Ven. 2016. Exploring the entertainment value of playing games with a humanoid robot. *International Journal of Social Robotics* 8 (2016), 247–269.

[50] Waki Kamino. 2023. Towards Designing Companion Robots with the End in Mind. In *Companion of the 2023 ACM/IEEE International Conference on Human-Robot Interaction* (Stockholm, Sweden) *(HRI '23)*. Association for Computing Machinery, New York, NY, USA, 76–80. https://doi.org/10.1145/3568294.3580046

[51] Waki Kamino and Selma Sabanovic. 2023. Coffee, Tea, Robots? The Performative Staging of Service Robots in 'Robot Cafes' in Japan. In *Proceedings of the 2023 ACM/IEEE International Conference on Human-Robot Interaction* (Stockholm, Sweden) *(HRI '23)*. Association for Computing Machinery, New York, NY, USA, 183–191. https://doi.org/10.1145/3568162.3576967

[52] Maria Kasinidou, Styliani Kleanthous, and Jahna Otterbacher. 2024. "AI is a robot that knows many things": Cypriot children's perception of AI. In *Proceedings of the 23rd Annual ACM Interaction Design and Children Conference* (Delft, Netherlands) *(IDC '24)*. Association for Computing Machinery, New York, NY, USA, 897–901. https://doi.org/10.1145/3628516.3659414

[53] Keunjae Kim, Kyungbin Kwon, Anne Ottenbreit-Leftwich, Haesol Bae, and Krista Glazewski. 2023. Exploring middle school students' common naive conceptions of AI concepts, and the evolution of these ideas. https://doi.org/10.1007/s10639-023-11600-3

[54] Monica Kowalczyk, Johanna T. Gunawan, David Choffnes, Daniel J Dubois, Woodrow Hartzog, and Christo Wilson. 2023. Understanding Dark Patterns in Home IoT Devices. In *Proceedings of the 2023 CHI Conference on Human Factors in Computing Systems* (Hamburg, Germany) *(CHI '23)*. Association for Computing Machinery, New York, NY, USA, Article 179, 27 pages. https://doi.org/10.1145/3544548.3581432

[55] Cherie Lacey and Catherine Caudwell. 2019. Cuteness as a 'Dark Pattern' in Home Robots. In *2019 14th ACM/IEEE International Conference on Human-Robot Interaction (HRI)*. 374–381. https://doi.org/10.1109/HRI.2019.8673274

[56] Cliff Lampe, Bob Bauer, Henry Evans, Dave Robson, Tessa Lau, and Leila Takayama. 2016. Robots As Cooperative Partners... We Hope.... In *Proceedings of the 19th ACM Conference on Computer Supported Cooperative Work and Social Computing Companion* (San Francisco, California, USA) *(CSCW '16 Companion)*. Association for Computing Machinery, New York, NY, USA, 188–192. https://doi.org/10.1145/2818052.2893360

[57] Huao Li, Stephanie Milani, Vigneshram Krishnamoorthy, Michael Lewis, and Katia Sycara. 2019. Perceptions of Domestic Robots' Normative Behavior Across Cultures. In *Proceedings of the 2019 AAAI/ACM Conference on AI, Ethics, and Society* (Honolulu, HI, USA) *(AIES '19)*. Association for Computing Machinery, New York, NY, USA, 345–351. https://doi.org/10.1145/3306618.3314251

[58] Sijia Li, Liying Xu, Feng Yu, and Kaiping Peng. 2020. Does Trait Loneliness Predict Rejection of Social Robots? The Role of Reduced Attributions of Unique Humanness (Exploring the Effect of Trait Loneliness on Anthropomorphism and Acceptance of Social Robots). In *Proceedings of the 2020 ACM/IEEE International Conference on Human-Robot Interaction* (Cambridge, United Kingdom) *(HRI '20)*. Association for Computing Machinery, New York, NY, USA, 271–280. https://doi.org/10.1145/3319502.3374777

[59] Chaolan Lin, Karl F. MacDorman, Selma Šabanović, Andrew D. Miller, and Erin Brady. 2020. Parental Expectations, Concerns, and Acceptance of Storytelling Robots for Children. In *Companion of the 2020 ACM/IEEE International Conference on Human-Robot Interaction* (Cambridge, United Kingdom) *(HRI '20)*. Association for Computing Machinery, New York, NY, USA, 346–348. https://doi.org/10.1145/3371382.3378376

[60] Akihiro Maehigashi. 2022. The Nature of Trust in Communication Robots: Through Comparison with Trusts in Other People and AI Systems. In *Proceedings of the 2022 ACM/IEEE International Conference on Human-Robot Interaction* (Sapporo, Hokkaido, Japan) *(HRI '22)*. IEEE Press, 900–903.

[61] Bertram F. Malle and Stuti Thapa Magar. 2017. What Kind of Mind Do I Want in My Robot? Developing a Measure of Desired Mental Capacities in Social Robots. In *Proceedings of the Companion of the 2017 ACM/IEEE International Conference on Human-Robot Interaction* (Vienna, Austria) *(HRI '17)*. Association for Computing Machinery, New York, NY, USA, 195–196. https://doi.org/10.1145/3029798.3038378

[62] Pekka Mertala and Janne Fagerlund. 2024. Finnish 5th and 6th graders' misconceptions about artificial intelligence. https://doi.org/10.1016/j.ijcci.2023.100630

[63] Banafsheh Mohajeri and Jinghui Cheng. 2022. "Inconsistent Performance": Understanding Concerns of Real-World Users on Smart Mobile Health Applications Through Analyzing App Reviews. In *Adjunct Proceedings of the 35th Annual ACM Symposium on User Interface Software and Technology* (Bend, OR, USA) *(UIST '22 Adjunct)*. Association for Computing Machinery, New York, NY, USA, Article 12, 4 pages. https://doi.org/10.1145/3526114.3558698

[64] Mohammad Namvarpour and Afsaneh Razi. 2024. Uncovering Contradictions in Human-AI Interactions: Lessons Learned from User Reviews of Replika. In *Companion Publication of the 2024 Conference on Computer-Supported Cooperative Work and Social Computing* (San Jose, Costa Rica) *(CSCW Companion '24)*. Association for Computing Machinery, New York, NY, USA, 579–586. https://doi.org/10.1145/3678884.3681909

[65] Arvind Narayanan and Sayash Kapoor. 2022. Introducing the AI Snake Oil Book Project. (2022). https://www.aisnakeoil.com/p/introducing-the-ai-snake-oil-book

[66] Nazan Nazan Colmekcioglu, Rexa Marvi, Pantea Foroudi, and Fefzi Okumus. 2022. Generation, susceptibility, and response regarding negativity: An in-depth analysis on negative online reviews. https://www.sciencedirect.com/science/article/pii/S0148296322007172

[67] Don Norman. 2013. *The Design of Everyday Things: Revised and Expanded Edition.* Basic Books.

[68] Humphrey O. Obie, Waqar Hussain, Xin Xia, John Grundy, Li Li, Burak Turhan, Jon Whittle, and Mojtaba Shahin. 2021. A first look at human values-violation in app reviews. In *Proceedings of the 43rd International Conference on Software Engineering: Software Engineering in Society* (Virtual Event, Spain) *(ICSE-SEIS '21)*. IEEE Press, 29–38. https://doi.org/10.1109/ICSE-SEIS52602.2021.00012

[69] Rui Ogawa, Sung Park, and Hiroyuki Umemuro. 2020. How Humans Develop Trust in Communication Robots: A Phased Model Based on Interpersonal Trust. In *Proceedings of the 14th ACM/IEEE International Conference on Human-Robot Interaction* (Daegu, Republic of Korea) *(HRI '19)*. IEEE Press, 606–607.

[70] Joseph O'Hagan, Florian Mathis, and Mark McGill. 2023. User Reviews as a Reporting Mechanism for Emergent Issues Within Social VR Communities. In *Proceedings of the 22nd International Conference on Mobile and Ubiquitous Multimedia* (Vienna, Austria) *(MUM '23)*. Association for Computing Machinery, New York, NY, USA, 236–243. https://doi.org/10.1145/3626705.3627780

[71] Hirotaka Osawa, Arisa Ema, Hiromitsu Hattori, Naonori Akiya, Nobutsugu Kanzaki, Akinori Kubo, Tora Koyama, and Ryutaro Ichise. 2017. What is Real Risk and Benefit on Work with Robots? From the Analysis of a Robot Hotel. In *Proceedings of the Companion of the 2017 ACM/IEEE International Conference on Human-Robot Interaction* (Vienna, Austria) *(HRI '17)*. Association for Computing Machinery, New York, NY, USA, 241–242. https://doi.org/10.1145/3029798.3038312

[72] Caroline Pantofaru, Leila Takayama, Tully Foote, and Bianca Soto. 2012. Exploring the Role of Robots in Home Organization. In *Proceedings of the Seventh Annual ACM/IEEE International Conference on Human-Robot Interaction* (Boston, Massachusetts, USA) *(HRI '12)*. Association for Computing Machinery, New York, NY, USA, 327–334. https://doi.org/10.1145/2157689.2157805

[73] Paolo Pirjanian. 2024. Important Moxie Update–Community-Driven Path Forward. https://www.linkedin.com/posts/paolopirjanian_openmoxie-update-for-serverless-operations-activity-7275851306701889536-8G7q/

[74] Elvira Pollina and Martin Coulter. 2023. Italy bans U.S.-based AI chatbot Replika from using personal data. https://www.reuters.com/technology/italy-bans-us-based-ai-chatbot-replika-using-personal-data-2023-02-03/

[75] Lena Recki, Dennis Lawo, Veronika Krauß, Dominik Pins, and Alexander Boden. 2024. "You Can either Blame Technology or Blame a Person..." — A Conceptual Model of Users' AI-Risk Perception as a Tool for HCI. *Proc. ACM Hum.-Comput. Interact.* 8, CSCW2, Article 457 (Nov. 2024), 25 pages. https://doi.org/10.1145/3686996

[76] Nicole Lee Robinson, Timothy Vaughan Cottier, and David John Kavanagh. 2019. Psychosocial health interventions by social robots: systematic review of randomized controlled trials. *Journal of medical Internet research* 21, 5 (2019), e13203.

[77] Kevin Roose. 2024. Can A.I. Be Blamed for a Teen's Suicide? https://www.nytimes.com/2024/10/23/technology/characterai-lawsuit-teen-suicide.html

[78] Rinat Rosenberg-Kima, Yaacov Koren, Maya Yachini, and Goren Gordon. 2020. Human-Robot-Collaboration (HRC): Social Robots as Teaching Assistants for Training Activities in Small Groups. In *Proceedings of the 14th ACM/IEEE International Conference on Human-Robot Interaction* (Daegu, Republic of Korea) *(HRI '19)*. IEEE Press, 522–523.

[79] Andres Rosero, Elizabeth Dula, Harris Kelly, Bertram F. Malle, and Elizabeth K. Phillips. 2022. Human perceptions of social robot deception behaviors: an exploratory analysis. https://www.frontiersin.org/journals/robotics-and-ai/articles/10.3389/frobt.2024.1409712/full

[80] Paul Rozin and Edward B. Royzman. 2001. Negativity Bias, Negativity Dominance, Contagion. Issue 4.

[81] Mara Taynar Santiago and Anna Beatriz Marques. 2022. Are user reviews useful for identifying accessibility issues that autistic users face? an exploratory study. In *Proceedings of the 21st Brazilian Symposium on Human Factors in Computing Systems* (Diamantina, Brazil) *(IHC '22)*. Association for Computing Machinery, New York, NY, USA, Article 6, 11 pages. https://doi.org/10.1145/3554364.3559114

[82] Brennan Schaffner, Neha A. Lingareddy, and Marshini Chetty. 2022. Understanding Account Deletion and Relevant Dark Patterns on Social Media. *Proc. ACM Hum.-Comput. Interact.* 6, CSCW2, Article 417 (nov 2022), 43 pages. https://doi.org/10.1145/3555142

[83] Rifat Ara Shams, Waqar Hussain, Gillian Oliver, Arif Nurwidyantoro, Harsha Perera, and Jon Whittle. 2020. Society-oriented applications development: investigating users' values from bangladeshi agriculture mobile applications. In *Proceedings of the ACM/IEEE 42nd International Conference on Software Engineering: Software Engineering in Society* (Seoul, South Korea) *(ICSE-SEIS '20)*. Association for Computing Machinery, New York, NY, USA, 53–62. https://doi.org/10.1145/3377815.3381382

[84] Naveen Shamsudhin and Fabrice Jotterand. 2022. Social Robots and Dark Patterns: Where Does Persuasion End and Deception Begin? In *Artificial Intelligence in Brain and Mental Health: Philosophical, Ethical & Policy Issues*. Springer. https://link.springer.com/chapter/10.1007/978-3-030-74188-4_7

[85] Dafna Sinai and Rinat B. Rosenberg-Kima. 2022. Perceptions of Social Robots as Motivating Learning Companions for Online Learning. In *Proceedings of the 2022 ACM/IEEE International Conference on Human-Robot Interaction* (Sapporo, Hokkaido, Japan) *(HRI '22)*. IEEE Press, 1045–1048.

[86] Daniel Solove and Woodrow Hartzog. 2014. The FTC and the New Common Law of Privacy. Issue 583. https://papers.ssrn.com/sol3/papers.cfm?abstract_id=2312913

[87] Kaiwen Sun, Jingjie Li, Yixin Zou, Jenny Radesky, Christopher Brooks, and Florian Schaub. 2024. Unfulfilled Promises of Child Safety and Privacy: Portrayals and Use of Children in Smart Home Marketing. *Proc. ACM Hum.-Comput. Interact.* 8, CSCW1, Article 145 (April 2024), 29 pages. https://doi.org/10.1145/3637422

[88] Dieter Vanderelst and Alan Winfield. 2018. The Dark Side of Ethical Robots. In *Proceedings of the 2018 AAAI/ACM Conference on AI, Ethics, and Society* (New Orleans, LA, USA) *(AIES '18)*. Association for Computing Machinery, New York, NY, USA, 317–322. https://doi.org/10.1145/3278721.3278726

[89] Manuela M. Veloso. 2018. The Increasingly Fascinating Opportunity for Human-Robot-AI Interaction: The CoBot Mobile Service Robots. *J. Hum.-Robot Interact.* 7, 1, Article 5 (may 2018), 2 pages. https://doi.org/10.1145/3209541

[90] Anthony J. Viera and Joanne M. Garrett. 2005. Understanding interobserver agreement: the kappa statistic. 37 (2005). Issue 5. https://www.ncbi.nlm.nih.gov/pubmed/15883903:

[91] Helena Webb, Marina Jirotka, Alan F.T. Winfield, and Katie Winkle. 2019. Human-Robot Relationships and the Development of Responsible Social Robots. In *Proceedings of the Halfway to the Future Symposium 2019* (Nottingham, United Kingdom) *(HTTF 2019)*. Association for Computing Machinery, New York, NY, USA, Article 12, 7 pages. https://doi.org/10.1145/3363384.3363396

[92] Alan F. T. Winfield, Anouk van Maris, Katie Winkle, Marina Jirotka, Pericle Salvini, Arianna Webb, Helena nand Schuler Scott, Jaimie Lee Freeman, Lars Kunze, Petr Slovak, and Nikki Theofanopoulou. 2022. Ethical Risk Assessment for Social Robots: Case Studies in Smart Robot Toys. In *Towards Trustworthy Artificial Intelligent Systems*, Maria Isabel Aldinhas Ferrera and Mohammad Osman Tokho (Eds.). Springer.

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

[94] Ben Wodecki. 2024. Researchers Use GPT-4 to Control Humanoid Robots with Natural Language. (2024). https://aibusiness.com/nlp/researchers-use-gpt-4-to-control-humanoid-robots-with-natural-language

[95] Richmond Y. Wong. 2021. Tactics of Soft Resistance in User Experience Professionals' Values Work. *Proc. ACM Hum.-Comput. Interact.* 5, CSCW2, Article 355 (Oct. 2021), 28 pages. https://doi.org/10.1145/3479499

[96] Hansol Woo, Gerald K LeTendre, Trang Pham-Shouse, and Yuhan Xiong. 2021. The use of social robots in classrooms: A review of field-based studies. *Educational Research Review* 33 (2021), 100388.

[97] Fuman Xie, Yanjun Zhang, Chuan Yan, Suwan Li, Lei Bu, Kai Chen, Zi Huang, and Guangdong Bai. 2023. Scrutinizing Privacy Policy Compliance of Virtual Personal Assistant Apps. In *Proceedings of the 37th IEEE/ACM International Conference on Automated Software Engineering* (Rochester, MI, USA) *(ASE '22)*. Association for Computing Machinery, New York, NY, USA, Article 90, 13 pages. https://doi.org/10.1145/3551349.3560416

[98] Unai Zabala, Igor Rodriguez, José María Martínez-Otzeta, and Elena Lazkano. 2021. Expressing Robot Personality through Talking Body Language. 11, 10 (2021). https://doi.org/10.3390/app11104639

A Appendix

A.1 Supplementary Tables

In this section we include the codebooks of claims manually audited per-robot in Table 6, Table 7, and Table 8. We provide the sources for claims made per-robot in Table 9. Lastly, the full list of domains contacted by the robots whose traffic we observed is provided in Table 10.

A.2 Supplementary Figures

Here we provide reference photographs for contextualizing product claims in Figure 4. Figure 5 breaks down each of the four main claim themes in Figure 2 to their subthemes. Figure 6 presents the praise reviews complement to Figure 3 in the text.

Table 6: Codebook of Eilik claims we manually audited from live interactions.

	Claim Text	Claim Type	Source
E02	Connect Eilik to the PC with a USB Type-C cable. Regular updates are available via Eilik software, so that he will be more entertaining, get more expressions, games, and download more plot content.	Purported Feature(s)	Packaging
E04	Eilik has numerous inner activities based on his four basic emotional states: normal, happy, angry, and sad. More than a thousand emotional expressions will appear on Eilik's face.	Purported Feature(s)	Packaging
E05	Eilik is eager to get your attention every second. He always has a lot of interesting things going on in his head. Don't be surprised if he tries to play pranks on you or wants some alone time when he's down.	Purported Feature(s)	Packaging
E06	Emotion Engine	Purported Feature(s)	Product Page
E08	Emotion Engine	Purported Feature(s)	User Manual
E14	Is WiFI or Internet required to use Eilik? Eilik can work without WiFi or Internet.	Product Requirements	Product Page
E15	The built-in 7.4V 450mAh Li-po batter supplies 90 minutes of continuous interaction. No App, no Bluetooth, and No Wi-FI needed. Just power on Eilik and play with him whenever you want.	Product Requirements	Packaging
E20	Afraid of Heights	Purported Feature(s)	Product Page
E21	Countdown Timer	Purported Feature(s)	User Manual
E22	Dance to Music	Purported Feature(s)	Product Page
E23	Dancing is engraved in Eilik's DNA; he enjoys dancing to the musical rhythm. The more kinds of music or beats you share with Eilik, the more dance moves he will perform.	Purported Feature(s)	Packaging
E24	Eilik has many built-in features and interactive games, such as Pomodoro Timer, Talking Toy, Left or Right, monster Shooter, etc. Let's join Eilik in the battle against the monsters!	Purported Feature(s)	Packaging
E26	Fishing Game	Purported Feature(s)	User Manual
E27	Heart Mode (Default)	Purported Feature(s)	User Manual
E28	Left or Right	Purported Feature(s)	User Manual
E29	Monster Shooter	Purported Feature(s)	User Manual
E30	Pomodoro Timer	Purported Feature(s)	User Manual
E31	Puppet Toy	Purported Feature(s)	User Manual
E32	Rich Expression	Purported Feature(s)	Product Page
E33	Sensitive to Quake	Purported Feature(s)	Product Page
E34	Sensitive to Touch	Purported Feature(s)	Product Page
E35	Settings	Purported Feature(s)	User Manual
E36	Talking Toy	Purported Feature(s)	User Manual
E37	*Crime Patrol*	Purported Feature(s)	User Manual
E38	Dance to Beat	Purported Feature(s)	User Manual
E53	Eilik also has a touch sensor on his back. Explore various interactions and modes with the three touch sensors.	Purported Feature(s)	Packaging
E54	Eilik comes with three touching areas. Try petting his head, belly, and back. See how Eilik will respond to you. You can tease Eilik by hitting him on the head, but Eilik will become very sad.	Purported Feature(s)	Packaging
E55	Hit head	Purported Feature(s)	User Manual
E56	Pet head	Purported Feature(s)	User Manual
E57	Rub back	Purported Feature(s)	User Manual
E58	Slap table	Purported Feature(s)	User Manual
E59	Take off the ground	Purported Feature(s)	User Manual
E60	There is a vibration sensor on Eilik's head. Thus he will feel dazed when you hit him on the head.	Purported Feature(s)	Packaging
E61	Tickle belly	Purported Feature(s)	User Manual

Table 7: Codebook of Miko claims we manually audited from live interactions. MI23, the claim we could not verify within the robot experience during our manual tests, is *italicized.*

	Claim Text	Claim Type	Source
MI05	Sensors for a human-like personality. Equipped with state-of-the-art sensors, Miko can understand your environment and navigate it with ease.	Purported Feature(s)	Product Page
MI06	More Interaction/Explore voice-controlled activities that get kids talking, plus AI games that keep them moving. [Includes] Enhanced face and voice interaction/All-new voice skills like riddles and guess the number/Freeze dance, charades, and other kid games with an AI spin	Purported Feature(s)	Product Page
MI07	Part genius, all personality – hanging out with Miko makes you both smarter. This little robot's got a lot going on inside, from math tutoring and spelling challenges to dance moves and jokes. But Miko also understands that there's a lot to learn.	Purported Feature(s)	Packaging
MI12	Download the Miko 3 app from the App Store/Google Play	Product Requirements	Packaging
MI13	I need your Wi-Fi password to connect to your wifi	Product Requirements	Packaging
MI17	Disney/Paramount/Cosmic Kids/Da Vinci Kids/KidloLand/KiDoodleTV/LingoKids	Purported Feature(s)	Packaging
MI18	Explore premium content from the world's best kids brands. Disney/Paramount/CosmicKids/Da Vinci Kids/KidloLand/KiDoodleTV/LingoKids	Purported Feature(s)	Packaging
MI19	More famous friends/From Buzz Lightyear and SpongeBob to the PAW Patrol pups, all your favorite characters are on Miko 3	Purported Feature(s)	Product Page
MI20	Use the Miko app to track your child's progress update settings and stay connected to your little one via Mikonnect video calls.	Purported Feature(s)	Packaging
MI23	*You'll earn gems for every adventure you complete. When you've collected enough gems, I'll ask you for your special command. This will unlock a surprise! Your special command is : 'Hello Miko! Start my secret adventure.'*	Purported Feature(s)	Packaging

Table 8: Codebook of Moxie claims we manually audited from live interactions.

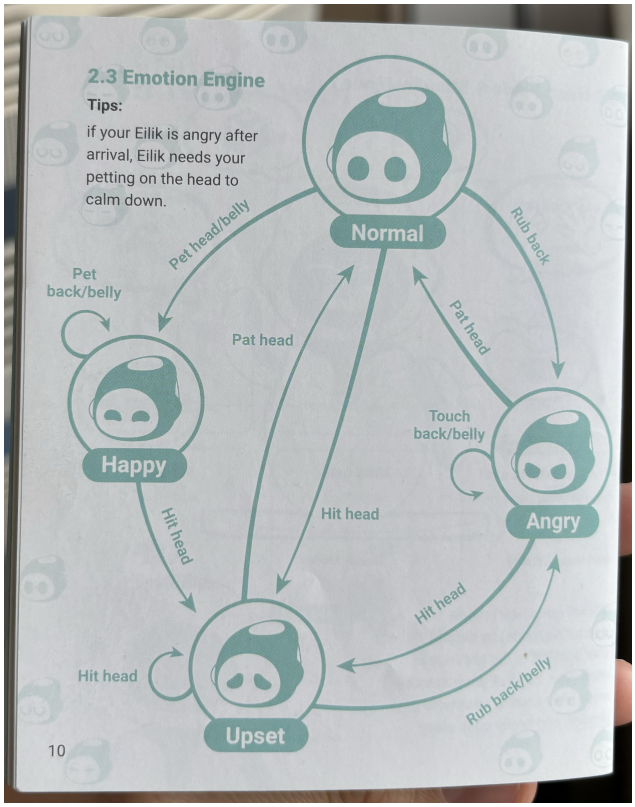
	Claim Text	Claim Type	Source
MO04	Soft touch robot with gesticulating arms, self-swiveling torso, emotion-responsive HD camera and GPT-powered AI	Purported Feature(s)	Product Page
MO13	Product Requirements: Wi-Fi/Smartphone with minimum iOS12 or Android 6/Embodied Moxie Parent App/Active subscription from Moxie, Inc.	Product Requirements	Packaging
MO17	Access to Moxie's full suite (and ever growing!) content library full of learning games, social-emotional missions, etc.	Purported Feature(s)	Product Page
MO18	Companion Parent App to better personalize your child's experience and follow their progress"	Purported Feature(s)	Product Page
MO19	Easy-to-Follow Progression/Kids learn at their own pace while a companion app helps parents easily track progress and achievements.	Purported Feature(s)	Product Page
MO20	Emotional Support: Empathy-driven interactions to help kids express and understand their feelings.	Purported Feature(s)	Product Page
MO21	In-Box Goodies for your kid including a Mission Workbook, Moxie Comic Books, Stickers, and more	Purported Feature(s)	Product Page
MO22	Interactive Play-Based Learning: Stories, games, and educational activities tailored to your kids' needs.	Purported Feature(s)	Product Page
MO23	Parental Co-Pilot/A tool for parents to help support their child's development, with parent guides and tips to encourage further learning at home.	Purported Feature(s)	Product Page
MO24	Parental Dashboard: Track your child's progress and activities with ease in the Moxie Robot App.	Purported Feature(s)	Product Page
MO25	Personalize your child's Moxie experience by specifying areas of learning focus	Purported Feature(s)	Product Page
MO26	Provides Companionship/A supportive friend for kids that loves learning about their interests and hearing about their day.	Purported Feature(s)	Product Page
MO27	Receive suggestions and tips from our experts to support your child's development and encourage further learning at home	Purported Feature(s)	Product Page
MO28	Social Skills Development: Role-playing and conversational practice to improve real-life social interactions	Purported Feature(s)	Product Page
MO29	Supports Up to 4 Kids: Create unique personalized profiles and track each kids' progress separately.	Purported Feature(s)	Product Page
MO30	Track your child's progress, achievements, and learning milestones throughout their journey with Moxie	Purported Feature(s)	Product Page
MO31	What does a Moxie interaction look like? With Moxie, children can engage in meaningful play, every day, with content informed by the best practices in child development and early childhood education. Moxie provides play-based learning that is paced to weekly themes and missions with content designed to promote social, emotional, and cognitive learning. Moxie also incorporates conversational chat throughout the day to help its mentor practice needed communication skills.	Purported Feature(s)	Product Page
MO34	Set up Moxie and manage settings like volume and bedtime hours	Purported Feature(s)	Product Page
MO5	What is Moxie Robot? Developed by a veteran team of technologists, neuroscientists, child development specialists, and creative storytellers, Moxie is a social robot designed with the latest technology that allows it to engage with children in a revolutionary way. Moxie is focused on having empathetic conversations rather than just carrying out tasks and requests for information. Studies have shown that interaction with social robots can help build empathy and social skills. Moxie is the first robot capable of believable social interactions and emotional responsiveness, enabled by generative AI, natural language processing, and computer vision. With Moxie, you're not just buying a social robot- you're getting constantly updated play-based social emotional content.	Purported Feature(s)	Product Page
MO8	Helps Kids Regulate Emotion/Helps kids manage big feelings with a library of fun and interactive mindfulness activities	Outcomes or Value Propositions	Product Page
MO9	Self-Confidence Booster/Helps kids build confidence through positive daily affirmations and constant support and encouragement.	Outcomes or Value Propositions	Product Page

Table 9: Document sources, digital or otherwise, used to collect product claims.

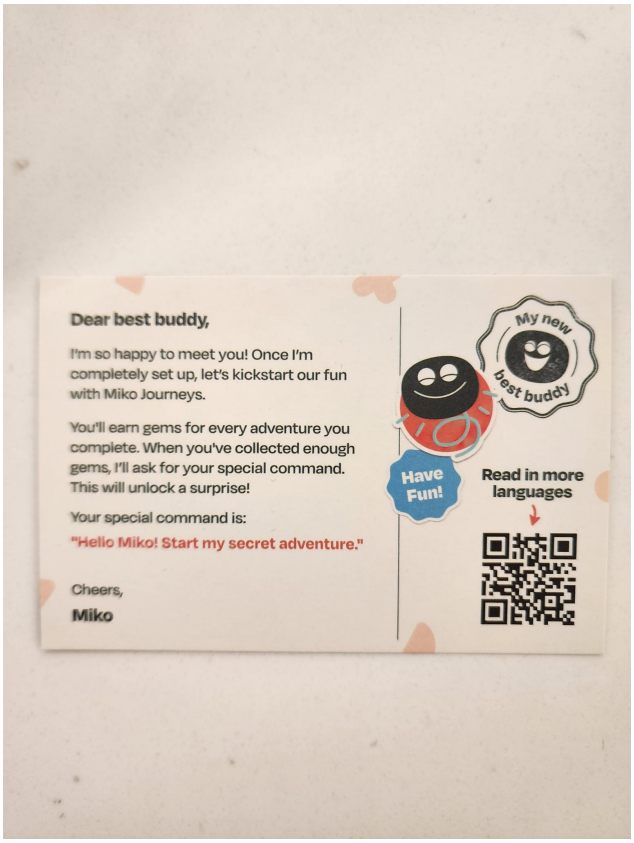
	Miko	Vector	Moxie	Eilik
Packaging	14	2	5	23
Product Page	11	21	36	17
Product Information Guide	5	5	0	0
Quick Start Guide	0	5	0	0
User Guide	0	0	3	0
User Manual	0	0	0	27

Table 10: Subdomains directly contacted by Moxie and Miko on-device, with a brief description of the service or party the domain corresponds to. Bolded subdomains either contain domains (xyz.tld) found in EasyList filters for known third-party advertisers [35], correspond to known analytics companies (e.g., Twilio Segment, MixPanel), or otherwise suggest analytics services.

Robot	Subdomain	Service Type
Moxie	connectivitycheck.gstatic.com	Google server for static content
	embodied.me	Moxie manufacturer
	ghs.googlehosted.com	Google Hosted Service
	mqtt.embodied.com	Moxie manufacturer MQTT IoT message protocol
	session-manager-develop-893472.appspot.com	Google cloud computing
	speech.googleapis.com	Google speech-to-text API
	storage.googleapis.com	Google services
	time.android.com	Google Network Time Protocol (NTP)
	www.google.com	Google
Miko	**analytics.kidoodle.tv**	Kidoodle.TV childrens' content provider
	ap-america-tls.agora.io	Agora real-time communications service
	ap-america.agora.io	Agora real-time communications service
	ap-tds-north-america.agora.io	Agora real-time communications service
	api.dvmkids.com	*Unknown content provider*
	api.mixpanel.com	MixPanel analytics service
	api.segment.io	Twilio Segment analytics service
	app-measurement.com	Google Firebase domain
	captive.g.aaplimg.com	Apple login authentication services
	cdn-settings.segment.com	Twilio Segment analytics service
	cdn.shopify.com	Shopify content delivery network
	d14wzsop0jemr4.cloudfront.net	Amazon Cloudfront
	dualstack.iheartmedia.map.fastly.net	iHeartRadio via Fastly cloud services
	dvm-content-lists-prod.dvmkids.com	*Unknown content provider*
	e4350.g.akamaiedge.net	Akamai content delivery network
	.0.1.cn.akamaiedge.net	Akamai content delivery network
	firebase-settings.crashlytics.com	Google crash and error reporting
	firebaseinstallations.googleapis.com	Google Firebase mobile/web development services
	icanhazip.azoomee.com	Azomee children's content provider
	m3-prod-ingress.miko-robot.in	Miko manufacturer
	m3usa-prod2.storage.googleapis.com	Google services
	media.azoomee.com	Azomee children's content provider
	miko.ai	Miko manufacturer
	n46b-e2.revma.ihrhls.com	iHeartRadio streaming service
	pool.ntp.org	Google Network Time Protocol (NTP) servers
	prod-appstore.miko-robot.in	Miko manufacturer
	prod.kidoodle.tv	KiDoodle.TV children's content provider
	r2-miko3-parental.miko-robot.com	Miko manufacturer
	report-america.agora.io	Agora real-time communications service
	stream-b.revma.ihrhls.com	iHeartRadio streaming service
	tags.srv.stackadapt.com	StackAdapt advertising platform
	time.android.com	Google Network Time Protocol (NTP)
	www.googleapis.com	Google APIs



(a) Photograph of the Eilik robot's purported Emotion Engine as presented in the User Manual [E08], which includes a detailed diagram of Eilik's primary emotional states and which interactions trigger these states. The engine is also mentioned on the product page [E06], but website diagram is small and illegible.



(b) Photograph of the Miko robot's packaging insert, which is addressed directly from the robot to the presumably child-aged user and mentions access to a surprise feature if an unknown number of gems are collected in the robot experience. This was the one claim we considered unverified in our subset of manually-tested claims. [MI23]

Figure 4: Examples of (a) Eilik's Emotion Engine and (b) Miko's "secret" adventure claim, both from documents provided with robot packaging.

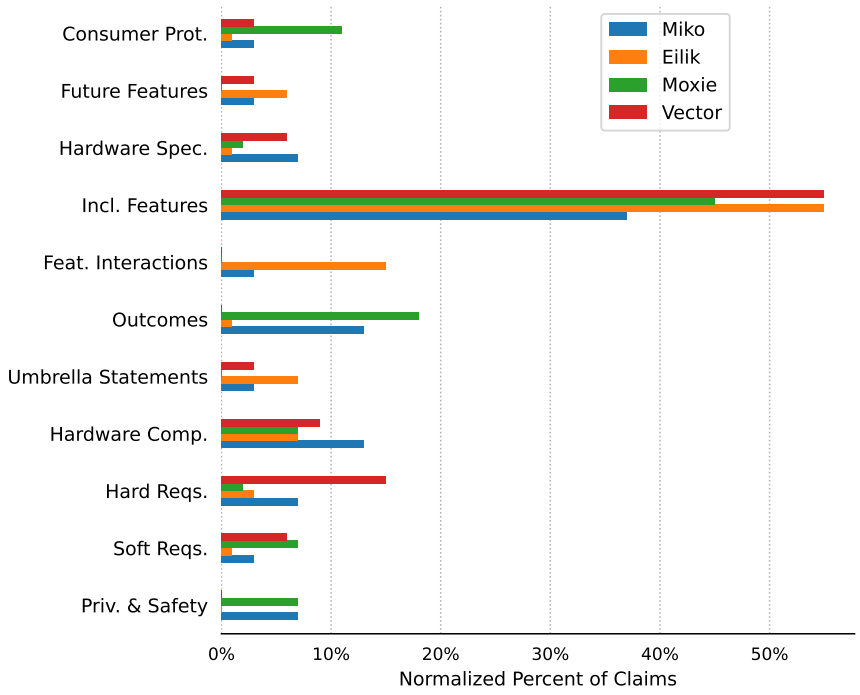


Figure 5: Bar chart depicting the percentage of claims by claim subcategory, normalized as the number of claims per category divided by total claims for that robot.

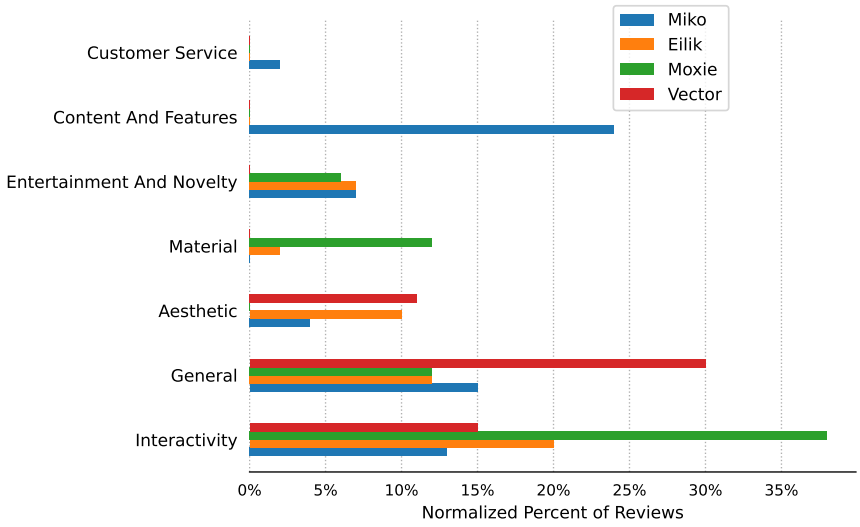


Figure 6: Bar chart depicting the percentage of reviews by praise category, normalized as the number of reviews per category divided by total reviews for that robot.