RESEARCH-ARTICLE

We Cannot Outsource What We Value Most: Toward Deployable Research Products in HRI

KAYLA MATHEUS, Yale University, New Haven, CT, United States

BRIAN SCASSELLATI, Yale University, New Haven, CT, United States

Published: 16 March 2026

HRI '26: 21st ACM/IEEE International Conference on Human-Robot Interaction
March 16 - 19, 2026
Scotland, Edinburgh, UK

Conference Sponsors:
SIGAI
SIGCHI

HRI '26: Proceedings of the 21st ACM/IEEE International Conference on Human-Robot Interaction (March 2026)
https://doi.org/10.1145/3757279.3788816
ISBN: 9798400721281

We Cannot Outsource What We Value Most: Toward Deployable Research Products in HRI

Kayla Matheus
Yale University
New Haven USA
kayla.matheus@yale.edu

Brian Scassellati
Yale University
New Haven USA
brian.scassellati@yale.edu

Abstract

Human-Robot Interaction (HRI) continues to rely on commercial social robot platforms to support academic research. Yet again and again, these systems prove short-lived, inaccessible, or misaligned with research needs. We argue that this is not an industry problem – the goals, needs, and constraints of industry are inherently distinct. Instead, this is a fundamental structural problem in HRI research, and one that must be solved from within. In short, HRI researchers must build their own products. In this paper, we trace the recent problems of industry-supplied robots and frame a new type of HRI research artifact in response: Deployable Research Products (DRPs), which bridge the gap between lab prototypes and commercial products. Drawing on mental models from business and innovation theory, we outline the mindset shifts that HRI must embody to move towards DRPs. We conclude with three emerging examples of this alternative path in the HRI community. These projects differ in scope and approach but share a common thread: to ensure the longevity of our science, we cannot outsource what we value most.

CCS Concepts

• **Human-centered computing** → **HCI theory, concepts and models**; • **Software and its engineering** → *Open source model*; • **Computer systems organization** → **Robotics**; • **Social and professional topics** → *Computing industry*.

Keywords

Human–Robot Interaction, Deployable Research Products, Open Source Robotics, Reproducibility, Social Robots

ACM Reference Format:
Kayla Matheus and Brian Scassellati. 2026. We Cannot Outsource What We Value Most: Toward Deployable Research Products in HRI. In *Proceedings of the 21st ACM/IEEE International Conference on Human-Robot Interaction (HRI '26), March 16–19, 2026, Edinburgh, Scotland, UK*. ACM, New York, NY, USA, 10 pages. https://doi.org/10.1145/3757279.3788816

1 Introduction

A primary goal of Human–Robot Interaction (HRI) research is to understand how robots, and the experiences they enable, can be designed to support human flourishing. Ours is an inherently applications-driven field, with the expectation that scientific insights translate into meaningful real-world impact. To study and generate knowledge at this level, we rely on robotic systems that reflect these real-world contexts. More and more recently, researchers have turned to commercial social robots to do so.

However, this reliance has come at a cost. A recent review of longitudinal social robotics studies from 2013 to 2023, most of which were conducted in the wild to reflect everyday contexts, found that the vast majority of studies relied on off-the-shelf robots [52]. Nearly half of the 120 studies analyzed used one of five specific platforms: NAO [9], Keepon [8], Jibo [33], Paro [10], or Pepper [11]. Yet today, four of these platforms – NAO, Keepon, Jibo, and Pepper – are no longer actively manufactured. In the last decade, we have witnessed the collapse of multiple industry-supplied robots. The consequences have been painfully tangible to the HRI community: hardware has become inaccessible or prohibitively expensive; SDKs, cloud servers, and developer portals have been shut down with little notice; and long-term studies have been thrust into uncertainty. At a wider lens, they have revealed a crack in our pursuit of knowledge: **our ability to generate scientific insight depends on misaligned infrastructures.**

In this paper, we argue that the faltering of off-the-shelf robots for research is not an industry problem – the goals, needs, and constraints of industry are inherently distinct. Rather, this is a fundamental structural problem in HRI research, and one that must be solved with effort from within. Instead of remaining dependent, HRI research must find its own way forward towards reusable, real-world, research-oriented platforms. In other words, **HRI research must begin building its own products**. These are not products as in commercial ventures, but as in research products deliberately designed for real-world inquiry going beyond typical bench prototypes.

This is admittedly a daunting task — and one that requires reimagining how HRI builds and starting to treat engineering efforts as scientific imperatives. Traditionally, most custom-created social robots from academia are one-off artifacts or unsuitable for real-world deployments. **To meet this new responsibility, we propose a new type of HRI artifact that bridges industry pragmatism with the distinct needs of academic research.** We call this a *Deployable Research Product* (DRP). For HRI, a DRP is the simplest form of social robot platform that is robust enough for real-world deployment, replicable enough for scientific inquiry, and durable enough for longitudinal use. While many smaller open-source HRI projects abound, DRP artifacts go beyond baseline replicability toward product-like stability and real-world interaction potential. Already, multiple groups have begun in this pursuit, and their stories reveal the need for a reorientation of HRI around sustainability beyond novelty.

HRI '26, Edinburgh, Scotland, UK

ACM ISBN 979-8-4007-2128-1/2026/03
https://doi.org/10.1145/3757279.3788816

To guide the reader through our argument, this paper unfolds in four segments. Section 2 is diagnostic, exposing the field's long-standing dependence on commercial robots and the precarity that follows when these systems vanish. Section 3 is a provocation – we must respond to this clear history not by repeating it, but by taking responsibility for it. Sections 4 and 5 provide constructive scaffolding towards this responsibility. We first define the DRP artifact for real-world HRI study, building off recent HCI concepts and identified needs from HRI use of industry robots. We next discuss how three frameworks from business innovation theory prompt mindset shifts required for HRI to develop DRPs. Finally, Section 6 is motivational. We present current examples from the community that embody DRP characteristics and principles. They showcase multiple diverse paths that can be pursued in the face of industry instability. Our primary aim in this work is simple: to call for ensuring the science of HRI prospers when the market does not.

2 The Trials and Tribulations of Recent Commercial Social Robots

The allure of commercial social robots for HRI research is clear: they arrive polished, with hardware designed for real-world contexts, often including SDKs or APIs to support software development without starting from scratch. Their scale also provides a platform for replication and comparison studies. It is no wonder that we, as a research community, have flocked to using industry-supplied robots for scientific inquiry. However, a familiar story persists: these robots keep disappearing. Or, more precisely, their parent companies do. When corporate funding runs dry, consumers fail to materialize, or product lines are reshuffled, the research programs built around them repeatedly become caught in the middle.

This is not industry's problem to solve for us. The needs and constraints of industry are fundamentally different from those of academia, and the HRI community must acknowledge this fact. Thus, in this section, we present four brief chronological cases for industry-based social robot companies that illustrate HRI research's problematic dependency. Admittedly, discussion of business needs and challenges may feel foreign to HRI researchers more accustomed to technical and design dialogue. However, we retain, similar to Mandel and Ju's call to the HCI community [50], that research designers and developers must understand the underlying structures that define their ability to produce good work.

We must be clear: this is not intended as a critique of industry at-large, the robots, their creators, or the corporations surrounding them. We also acknowledge that 'industry' is not a singular monolith and that there exists a spectrum of openness of platform and level of research partnership. There are current on-market examples that are used in research today (e.g., Misty [7], Furhat [6], QT Robot [2], TIAGo [12]), and we hope that they can continue to fill scientific needs. What we wish to illustrate here is a sobering reality of the underlying vulnerabilities that industry-developed social robots carry for the research community.

2.1 Keepon (BeatBots, Inc., Wow! Stuff)

Created at Japan's National Institute of Information and Communications Technology (NICT), Keepon (Figure 1a) was designed by

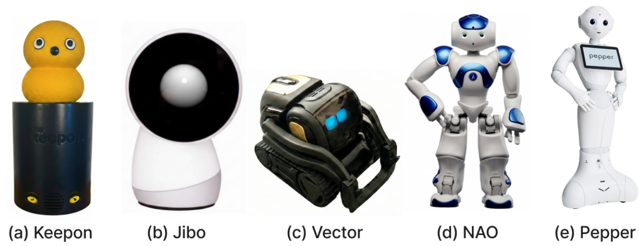

Figure 1: Five industry-supplied social robots in HRI research that are now discontinued or unsupported. (a) *Keepon*, (b) *Jibo*, (c) *Vector*, (d) *NAO*, and (e) *Pepper*.[2] Note: Robots are not to scale.

Hideki Kozima and Marek Michalowski to study social engagement with children on the autism spectrum [47]. Its simple, yellow, snowman-like body and rhythmic "dancing" gestures made it instantly recognizable. A viral 2007 YouTube video turned the lab tool into a pop-culture icon, sparking the creation of BeatBots, Inc. and partnerships with major toy retailers (including Toys 'R Us).

BeatBots, Inc. first released the Keepon Pro [17] for research, with onboard cameras, microphones, and teleoperation. Translating a lab prototype into a professional system, however, proved prohibitively costly; each unit sold for the equivalent of $48,000 today. In 2013, consumer demand drove a collaboration with Wow! Stuff to produce a $40 toy version. The toy preserved Keepon's iconic bouncing but stripped away the sensors and connectivity key to HRI researchers. As Michalowski noted in an interview [37], "You can do a lot when you're building one research robot…It's a very different thing to put something in a toy store…we ran into physics and economics challenges." This same notion is echoed by other HRI researchers who commercialized a research project (Purrble [42]): "We realized that certain key features of an ideal long-term research product needed to be jettisoned from the commercial prototype, in the interest of cost management and durability."

To bridge these capability gaps, some researchers turned to hardware hacking and reverse-engineering. With scattered instructions for tapping into the toy's I^2C bus, adding microcontrollers, and overriding closed firmware, it was technically possible to make a "franken-keepon" [55]. Cao et al. from the HRI community further published detailed hacking instructions and a control GUI for researchers [21], but adoption was limited outside of the developers' institutions. In our own deployments using hacked Keepons [citations removed for anonymity], failures were common, from broken fixtures, uncontrollable motors, and vanishing replacement parts once production ceased just a year after market release.

Keepon's story encompasses a fundamental truth: *translating a research robot into an affordable commercial product demands trade-offs that can undermine scientific fidelity.* Researchers looking to fill these gaps on their own can struggle in this leftover landscape.

2.2 Jibo (Jibo, Inc)

Founded in 2012 as a spin-out from MIT Media Lab researchers, Jibo was designed as "the world's first family robot" [3]. Standing about

[2]Image credits: (a)(c) Author; (b) Wikimedia Commons, Cynthiabreazeal, CC BY-SA 4.0 (background removed); (d) Wikimedia Commons, Softbank Robotics Europe, CC BY-SA 3.0; (e) Wikimedia Commons, Softbank Robotics Europe, CC BY-SA 4.0 (crop).

28 centimeters tall, Jibo featured a three-axis body, an expressive circular display, far-field microphones, cameras, touch sensors, and Wi-Fi connectivity. Part of the late 2010s wave of consumer social robots, Jibo quickly became a high-profile example: it broke records as the most-funded technology Indiegogo campaign of its time [4] and appeared on TIME's 2017 list of "25 Best Inventions" [72]. Excitement from the HRI community at the time was understandably high. A developer SDK for the robot was released early on, and the robot's capabilities were advanced while packaged in a sleek form factor that could be easily deployed into participant homes at scale.

Yet in 2018, the company shut down after shipping initial domestic pre-orders but canceling international ones [74]. Eisenmann has detailed the contextual and structural challenges leading to the company's shutdown [28]. For one, like Keepon, development and unit costs were substantially greater than anticipated when translating from a bench prototype. Perhaps more fateful, however, was the arrival of the Amazon Echo in 2014, which reframed public expectations of "social" technology at a fraction of the price. To HRI researchers, Jibo's embodiment, expressivity, and sensor suite were incomparable; to consumers, it appeared newly redundant and lacked a clear use case [13, 39, 57].

While the public pre-order campaign was not fulfilled in completion, multiple research groups had already obtained Jibo hardware and the SDK. The shutdown of the parent company led to the closure of cloud servers and prevented critical updates to the early and iterative technology stack. As a result, researchers had to modify the robot's hardware and software due to limited access to core sensing and difficulty networking the robot to additional devices and cloud services [43, 67]. The Jibo story is similar to Keepon in that both robots came from HRI research efforts and faced difficulty with cost logistics and short lifespans on the market. However, it reveals a further truth: *consumer attention is a fickle beast and can quickly diverge from what researchers hold dear.*

2.3 Vector (Anki, Digital Dream Labs)

Released in 2018 by Anki, Vector [34] quickly became a recognizable face in HRI studies—from long-term home deployments [79], to companionship development [76], and user perception investigations [46]. The palm-sized robot featured a rolling base, animated OLED eyes, a camera and microphone, and cloud-connected speech and AI capabilities. Anki itself, founded in 2013 by Carnegie Mellon Robotics Institute alumni, aimed to bring artificial intelligence to everyday life through approachable consumer products [23].

For researchers, Vector was a gift: affordable, expressive, easily portable, and well-documented through an open SDK that made larger-scale studies theoretically feasible. But in 2019, barely a year after launch, Anki abruptly shut down due to fundraising challenges [70]. This unexpected collapse created immediate uncertainty for HRI researchers. For instance, Tsoi et al. were unsure how long support for the robot would continue and if robots deployed longitudinally in participant homes would suddenly cease to function [77]. The team also faced issues with networking, the mobile application, and hardware modifications, stating: "We attribute these challenges to the robot being designed as a consumer product, not as a development platform."

To its credit, Anki temporarily kept cloud servers online and communicated with researchers even after closing [24], a rare act of goodwill. Hope briefly resurfaced in 2020 when Digital Dream Labs (DDL) acquired Anki's assets and pledged to revive manufacturing, restore cloud access, and expand the SDK [25]. Reality was harsher: cloud services went offline for over a year, the SDK was nonfunctional, and customer complaints eventually culminated in a 2024 lawsuit by the Pennsylvania Attorney General for fraudulent charges and unfulfilled promises [38]. New developments from DDL have stalled as a result.

Like Keepon and Jibo, Vector came and went on the consumer market rapidly. However, unlike single-product startups, Anki has two previous successful product lines and established manufacturing and marketing pipelines. We note that *even seemingly successful consumer robots on the market can disappear abruptly*. Furthermore, even the best laid plans to support researchers can also quickly crumble in the face of financial and corporate volatility.

2.4 NAO and Pepper (Aldebaran, Softbank Robotics, United Robotics Group)

In contrast to Vector, Jibo, and Keepon, whose market lives were brief, NAO and Pepper have been fixtures of HRI for over a decade. NAO, a small tabletop humanoid introduced in 2007, offers a small but expressive platform with cameras, microphones, tactile sensors, and speech synthesis [31]. Pepper, introduced in 2014, is a larger humanoid with expressive motions, a touchscreen chest, and cloud-based AI integration, marketed as an "emotional robot" for public-facing roles [63]. Together, they have deeply shaped an era of HRI research, appearing in hundreds of studies across education, eldercare, autism therapy, cultural adaptation, and workplaces [16, 48, 52]. Yet, in June 2025, parent company Aldebaran announced that it was shutting down due to bankruptcy, leaving the field's most widely used robots without support [26].

Originally founded in France in 2005, Aldebaran developed NAO in response to Sony's discontinuation of the Aibo robot and the need for a new standard RoboCup platform [31]. Its success made NAO a laboratory staple: small, programmable, and relatively affordable. In 2012, SoftBank acquired Aldebaran and launched Pepper to commercialize humanoid interaction at scale. Pepper appeared in airports, banks, and advertisement campaigns, but struggled with unclear use cases and unmet interaction expectations [41, 56, 80]. By 2021, production had halted due to financial losses from Pepper, and SoftBank sold its robotics division to United Robotics Group (URG), which rebranded back to Aldebaran. URG/Aldebaran promised to continue support for NAO and Pepper until 2026 [29] – a promise that would go unfilled.

Despite this long tenure, NAO and Pepper saw very few product upgrades over 15 years, leading to many known issues stagnating. Amirova et al. [16] reviewed more than 300 NAO studies between 2010 and 2020, finding recurring issues such as outdated sensors, poor speech recognition, overheating, and fragile batteries. Others have reported similar fragility limiting deployment (e.g., [18, 66]). For researchers, this creates a tension between aging and lagging systems, but ones that are deeply embedded in research practices. Under more normal circumstances, the parent companies of these robots would likely have iterated and released upgrades. However, multiple corporate reshufflings and profit pressures became higher priorities. Without open systems that developers can upgrade themselves, HRI researchers have been left waiting.

The story of NAO and Pepper is quite distinct from Vector, Jibo, and Keepon in its longevity and reach. What we can now see is that *longevity and scale in industry do not guarantee stability or innovation for future HRI research.* For the HRI community, the question now is not which company will fill the gap left behind – but whether we should keep waiting for one at all.

3 The Responsibility of the HRI Research Community

Together, the above stories make one thing clear: we cannot be surprised when the next industry-provided social robot goes under. The question we must ask ourselves is: how much are we willing to be complacent in the continued repeat of history?

We could, of course, remain with the status quo and continue to rely on whatever systems the market happens to offer at the time. Despite volatility and divergent goals, these products are tempting. What better way to study real-world HRI than with real-world devices designed for real-world contexts? And who is best equipped with the resources and know-how to build such products? HRI researcher time is far better spent focusing on fieldwork than on building new hardware and software from the ground up. It is well established just how high effort and difficult [35, 65] such endeavors are. Let our community focus on insight generation and let others deal with such difficulties.

Yet, this mindset turns a blind eye to the histories of Keepon, Jibo, Vector, NAO, and Pepper. They remind us that the foundations of our field are often built on borrowed platforms optimized more for sales cycles rather than scientific continuity. With such divergent needs, academia could turn entirely away from industry. However, this approach ignores the inherent nuances of how the two entities are inter-twined. Industry has already mastered much of what research needs for scalable, real-world robot development. It also plays a heavy role in consumer perception of social robots, which impacts downstream HRI research. Symbiotically, much of HRI research seeks to provide bottom-up insights that can drive better integration and adoption of real-world social robots. Fundamentally, both industry and academia share a common vision: enabling robots to play meaningful roles in everyday life.

Instead, we argue for an approach that radically acknowledges these divergent needs and realigns responsibility. The overarching lesson from the cases in Section 2 is not that industry has failed, but that research cannot depend on commercial trajectories alone. To generate durable knowledge about robots in daily life, we must take responsibility for building research-ready systems ourselves. It bears stating again: *HRI research must start building its own products.* The challenge and responsibility before us is not to become industry (or avoid it), but to learn from and adapt its lessons to our needs.

By "products," we do not mean commercial offerings themselves but research artifacts that strive towards similar levels of durability, completeness, and usability. At the same time, these artifacts should be grounded in what we value most as researchers. More than just outlining the fragilities of social robots in industry, the stories from Section 2 highlight what HRI researchers actually require:

(1) *Long-term accessibility* to ensure continuity beyond the lifespan of any one company.
(2) *Affordability at scale* so that multi-site and longitudinal studies are feasible.
(3) *Full control across the stack* to adapt and instrument systems for scientific goals.
(4) *Robustness and reliability in-the-wild* for extended use in homes, schools, and clinics.
(5) *Hardware openness and modifiability* to support creativity, hacking, and replication.
(6) *A community of support* that can come together for maintenance and best practices.

Many of these needs are echoed across prior HRI literature investigating or reflecting on tools for longitudinal, in-the-wild study. For instance, Fraune et al. note that the available set of robot functionality limits what researchers study with users, and that restrictions on closed hardware and software make in-the-wild study especially difficult for long-term applications [30]. Bu et al. describe the required robustness and infrastructural support systems required for in-the-wild deployments [20]. Jung and Hinds discuss the need for more study with a larger array of real-world, in-context situations [44]. Gunes et al. describe, through the lens of reproducibility, how commercial robots are expensive and rigid and generally suggest using more modular robot kits and custom designed robots in research [35]. It is clear that the status quo is not working and that the outstanding needs of HRI researchers are repeatable and pervasive.

In the following section, we propose a way forward through the framing of a new type of HRI artifact: *Deployable Research Products* (DRPs). Building on theoretical and practical work from HCI, we position the DRP between open-source robots and commercial systems, striving for a balance between researcher control and real-world robustness. Building such artifacts, however, is a challenging and exploratory pursuit. In Section 5, we explore how established mental models from innovation and business theory can help HRI adopt product-like thinking without commercializing its aims.

4 Framing a New HRI Artifact: Deployable Research Products

When we speak toward HRI researchers building their own products, we do not mean a product sold for a profit. Rather, we mean a robotic system that is intentionally *product-like*: self-contained, robust in the field, and with enough perception and interaction fidelity to be treated by participants in real-world ways.

The concept of a *research product* has been recently explored in HCI, but it is scarce in HRI, limited to describing the development of a commercial product from research [42]. Odom et al. originally differentiated between a *research prototype* and a *research product*, with the product as a natural and mature extension of the prototype [60]. A research product is described as an artifact that no longer merely represents a concept but instead embodies a design inquiry and can circulate within everyday contexts. Odom et al. outline the following four qualities in HCI research products: (1) *Inquiry-driven* – created to investigate a research question rather than to serve an immediate user or market need; (2) *Finish* – built with a completeness and robustness that allows authentic use in everyday life; (3) *Fit* – situated within real social or material contexts, engaging actual practices and environments; (4) *Independent* – able to persist and produce meaning outside the presence of the researcher. Boucher has additionally discussed the role of *scale*, including the ways that production volume, fabrication method, and distribution strategy affect what a research product can reveal [19].

These notions have inspired numerous physical HCI research products (e.g., [32, 40, 61]) and discussion of artifacts that span varying levels of fidelity and real-world applicability [69]. Overarchingly, these HCI practices and concepts are very well aligned with the new responsibility that the HRI community faces (Section 3). The characteristics of finish, fit, independence, and scale directly address the needs HRI researchers seek in consumer robots, while remaining inquiry-driven describes the core tensions in using them.

However, there are also distinct factors to consider when building social robotic research products. While HRI shares conceptual roots with HCI, a robot's embodiment, social interactions, and autonomous systems introduce unique and complex technical and design challenges [15, 49]. The mode of inquiry with a research product also becomes distinct. A research product in HCI is often described as the end goal or outcome of the Research through Design process or other designerly ways of knowing [60]. In HRI, the research product need is for reusable and replicable systems to generate a wide range of scientific insights. Rather than the research being done *through* the creation of the product, the research becomes *enabled by* it. The *openness* of the system thus becomes a crucial component for long-term applicability.

Building upon the principles of HCI research products and addressing the unique needs of the HRI community, we propose a new class of HRI research artifact: *deployable research products* (DRPs). DRPs operate in a unique balancing act between shorter-term, lower-fidelity lab prototypes and closed commercial social robots (Figure 2). They strive towards research-grade systems intentionally engineered for sustained, real-world use, while remaining accessible for adaptation and study by other researchers. HRI DRPs encourage this complex, real-world use by being:

(1) *Robust*: The hardware and software stacks are durable enough for autonomous, in-the-wild, longitudinal use in everyday environments.
(2) *Unified*: The robot system is packaged as a unified, enclosed system for social interaction, rather than as a collection of sensors or I/Os.
(3) *Contextual*: The robot system provides common baseline features and functionality that may be expected by users in everyday contexts.

While remaining flexible for research needs and enabling community support for longevity:

(4) *Open*: The hardware and software are fully open and functionally accessible beyond the initiating research group.
(5) *Adaptable*: The robot system is not only open, but also is designed to be adaptable by researchers through modularity, hackability, or peripheral tools.
(6) *Reproducible*: The robot system can be easily built or obtained by external groups and includes core infrastructure for deploying robots en masse for study.

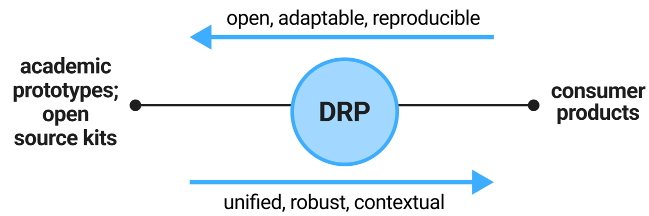

Figure 2: Deployable Research Products (DRPs) occupy a conceptual space between academic prototypes, open source kits, and consumer products. They are research-grade robotic systems intentionally engineered for sustained, unsupervised, real-world use, while remaining adaptable by researchers.

Together, these qualities delineate a differentiated design and development space for HRI artifacts. In practice, DRPs will emphasize more or less of these characteristics, but together they represent a conceptual ideal. HRI DRPs can be distinguished from other types of social robot artifacts through this lens. Prototype robots are typically faster to produce, lower-fidelity, and aimed at one or a small handful of studies. They are also more likely to be deployed with external components for data collection than as a single, unified package. Open source robotics kits are also less likely to be unified but seek to reach larger audiences through reproducibility. However, they may not always carry the robustness or support needed for more longitudinal, in-the-wild deployments. On the other end of the spectrum, commercial products used in HRI research often carry this robustness, but they are closed and less research-flexible. The vision with DRPs is for open and adaptable systems that remain robust and designed for real-world contexts.

5 Toward DRPs with Mental Models from Industry

Building towards DRPs is non-trivial. The pursuit demands the synthesis of engineering rigor, design sensitivity, and scientific framing – within academic systems that often reward novelty more than longitudinality. Doing so raises a practical question: how can a research community with limited resources, short project timelines, and high turnover adopt the practices of long-term product development without losing sight of core scientific goals?

To make this shift more attainable, we look beyond academia and borrow conceptual tools from innovation and product-development theory. In short: how can HRI think like an industry company that builds for the real world—without becoming one? We highlight three mental models (Figure 3): *exploration versus exploitation*, the *minimum viable product*, and *crossing the chasm*. These frameworks were originally developed to help companies navigate uncertainty, balance investment strategies, and scale technologies beyond early adopters. We situate each model within a recurring tension in HRI, highlighting needed mindset shifts for developing sustainable DRPs.

5.1 Exploration vs. Exploitation

Similar to the *exploration-exploitation* paradigm in artificial intelligence, the *exploration-exploitation* framework in industry describes

Figure 3: Three conceptual framings adapted from business and innovation theory to guide DRP development: balancing *Exploration vs. Exploitation*, *(2)* developing a *Minimum Viable Product*, and *(3) Crossing the Chasm* in the community.

the balance between pursuing new ideas and markets versus refining and scaling what already works [51]. Mature companies often struggle to balance the two: over-exploring leads to endless prototypes with no revenue, while over-exploiting leads to stagnating offerings. Finding the balance depends on the stage of a company and the resources available to pursue an *ambidextrous* [36] approach to continue exploration during exploitation.

Academic HRI has historically over-indexed on exploration, given the need for novelty in scientific contribution. This has led to a proliferation of one-off prototypes, bespoke software stacks, and custom study-specific implementations that are rarely maintained or shared beyond initial data collection. This mindset also leads to the under-valuing of the core, necessary process elements of exploitation. Namely: investing in robust infrastructure, applying modern engineering best practices, and creating and maintaining truly reusable artifacts beyond a single lab.

Recent initiatives toward open-source artifacts and replicability mark important progress. However, these efforts often remain exploratory with limited uptake, and they are often viewed as smaller contributions. For DRPs to thrive, the HRI community must shift towards an *ambidextrous academic mindset.* This is one that **values and incentivizes the lower-level, mundane infrastructure work required to transform exploratory prototypes into reliable research platforms.** Without such engineering, the HRI community will continue to find itself in a cycle of dependence on unstable commercial products.

5.2 Minimum Viable Product (MVP)

The *minimum viable product* (MVP) concept emerged from lean-startup and agile-development movements and is common both in industry circles and out. It refers to the simplest version of a product that delivers core value and enables rapid feedback loops [68]. Its goal is not polish but validated learning and fast iteration. Successful companies treat MVPs as living experiments that evolve toward market fit while minimizing wasted effort. In effect, the MVP ensures exploration feeds exploitation efficiently.

While this concept may sound like it already aligns well with existing HRI research practices, the reverse is true for DRP development. **Whereas MVPs in industry address a tendency to over-build before release, academia often underbuilds**. One-off prototypes developed in the lab are often able to compensate for lower levels of fidelity with quick hardware fixes and researcher oversight. DRPs, however, require a higher bar of autonomy and must consider the needs of everyday contexts. For instance, a DRP must consider elements like how easily the robot can be plugged in, how a user can monitor the robot's status or reset it, or how users can move the robot around in their homes or workplaces.

In these ways, the HRI community must consider carefully what ”minimum” means in specific contexts. To preserve resources, DRP development should aspire to "just robust enough" to be deployed in-the-wild for long enough to produce meaningful data – but need not seek multi-year warranty levels. At the same time, the bar must be raised to deploy systems that can appear and interact more like finished products than most open-source robots.

5.3 Crossing the Chasm

In innovation theory, Moore's "crossing the chasm" describes the gap between early adopters and mainstream users of a product [58]. Many promising technologies fail here, not because they lack merit, but because the expectations on either side differ radically. Early novelty-seekers tolerate imperfections, whereas mainstream users expect reliability, polish, and support. The way to cross this chasm is to focus on building a ”beachhead market." Rather than trying to appeal to everyone, the company should identify a small, well-defined segment where the product can achieve clear success. Then, that success story buoys further expansion. This approach requires intentionally building out the ”whole product"; not just the product offering itself, but the documentation, onboarding, and support ecosystem to make it accessible.

With regard to crossing the *academic chasm*, the approaches described in industry are directly translatable. There are two main takeaways for developing HRI DRPs. First, **researchers must carefully consider the layers beyond just the core robot and human–robot interaction**. Peripheral needs such as onboarding processes, remote maintenance, and preference controls are vital to DRPs success. The same is true for infrastructure that enables other researchers to adapt and sustain the system (e.g., documentation, templates, dashboards, data logging). In many ways, DRP creators must have as much empathy for developers as for end robot users.

Secondly, crossing the academic chasm requires **a deliberate focus on community-building and shared stewardship.** It is very challenging for a single lab to have the time, funding, or continuity to maintain a complex social robot platform over years. Sustaining a DRP therefore demands a distributed model of collaboration. This likely includes collaborating with industry experts, co-creating within the greater HRI community, and finding early traction in the open-source community.

6 Emerging Efforts in HRI DRPs

Developing HRI DRPs should not follow a one-size-fits-all template, even while following industry mental models. Real-world contexts are plentiful, and thus our artifacts must be too. Already, multiple groups are testing this terrain, navigating trade-offs among openness, reproducibility, and durability in different ways. In this section, we discuss three example projects as inspiration that paint this diversity. For instance, Quori (Section 6.1) shows what can emerge from collective design toward a shared, replicable humanoid platform. Ommie (Section 6.2) traces the opposite path—translating a single-lab prototype of a unique HRI modality into a deployable, field-ready system. The Jibo Research Platform (Section 6.3), in contrast, demonstrates how defunct commercial hardware can be reclaimed and extended for academic use.

Each of these examples embodies working towards DRPs in their motivation and approaches. Their stories can certainly serve as inspiration to those eager to build DRPs. Just as important, though, is sparking reflection on how each member of the HRI community might support the ecosystem that can enable these types of aims. We also note that there are other active, adjacent efforts (e.g., Blossom [73], Shutter [75], FLEXI [14]) that share similar goals of accessibility and community reuse. While these efforts don't meet all of the criteria described in our proposed artifact framing, their future work goals describe moving towards these characteristics, further showcasing the applicability of DRP framing and pursuit.

6.1 Quori 1 and Quori 2

Quori, introduced originally in 2015, is one of the first explicit attempts to build a social robot as a collective rather than as a single lab. Funded by the U.S. National Science Foundation and developed across several universities, the project was directly motivated by the lack of affordable, replicable, and expressive robots available for HRI research [71]. To inform the platform's design, the team garnered early survey responses from over 80 researchers across the community, followed by four research workshops that engaged dozens more with iterative co-design. Together, these efforts form an unusually transparent, community-driven process.

To meet their identified needs, the team made multiple development trade-offs [71]. For instance, eliminating expensive actuation where possible while preserving key expressive capabilities through projection and lightweight degrees of freedom. In line with the MVP mindset, these reductions sought to preserve the illusion of social presence while ensuring replicability and field robustness and minimizing costs. Support from a U.S. National Science grant enabled the replication of ten Quori 1 units sent to ten different research labs [5] – the project's own "beachhead market." The grant also supported the production of a suite of software tools built in collaboration with industry partners. These "whole product" tools modularized core capabilities such as speech, gesture, attention, and animation through a developer API and included tools for non-programmers to lower barriers to entry.

These efforts, along with further workshops and documentation efforts [1], have led to multiple research studies performed with the Quori 1 robot [27, 45, 59, 78]. At the time of writing, the team is currently developing an updated Quori 2 robot in collaboration with industry partners to apply core lessons from the first generation. The project's trajectory as a whole is a unique example of crossing the academic chasm through shared stewardship. Achieving a decade-long lifespan across multiple institutions and grant cycles is no easy feat within academic contexts.

6.2 Ommie Prototype to Deployable System

Unlike Quori's multi-institutional, grant-funded effort, the *Ommie* robot began as a lab experiment at Yale University that the authors first introduced in 2022 to support individuals with anxiety [54]. The robot features a unique HRI modality: users place their hands on the robot's body during haptically-guided deep breathing.

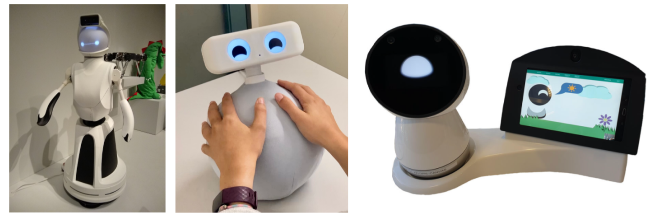

(a) The Quori 1 Robot (b) The Ommie Robot (c) The Jibo Research Platform

Figure 4: Three emerging examples of DRP development within HRI. *(a) Quori*, a collectively-designed humanoid emphasizing replicability; *(b) Ommie*, a single lab prototype evolved into a field-ready system; and *(c) The Jibo Research Platform*, which extends discontinued hardware.[4]

Initial studies [22, 53] have shown in-the-wild therapeutic benefit with single sessions. To show sustained behavior change and other longitudinal impacts, however, the robot required significant upgrades across both hardware and software. Over the past year, with industry guidance sponsored by a nonprofit organization (Semio Community), we have translated the initial prototype into a fleet of home-deployable systems. Effort has been taken to develop not just for the next study, but also toward an adaptable and robust DRP.

Notably, the exterior form factor and core interactions of the upgraded robot remained, ablating any traditional HRI notions of novelty. Instead, all changes were *exploitation*-minded developments in the mechanical, electrical, architectural, and deployment systems. For instance, the robot's breathing mechanism was redesigned to specifically withstand sustained use, while the robot's assembly and disassembly process was reconfigured for ease of replication by non-professionals. A comprehensive autonomy system modularized and configured all lower-level I/Os, middle-layer behaviors, and higher-level states. The team also future-proofed the entire software stack by upgrading to the latest standards and best practices (e.g., ROS2, Zenoh middleware, Raspberry Pi 5, Android Kotlin, etc.). While these efforts were not strictly necessary for the robot's next study, this was an active investment in longevity and accessibility for future researchers.

Deeper consideration of home contexts and product expectations was a part of DRP consideration. For instance, sensory intrusiveness became an unexpectedly important factor (i.e., how errant, or even proactive, lights and sounds might negatively disrupt a user in private space). Minimizing light and sound pollution from off-the-shelf components required low-level firmware modifications, part swaps, and physical light gasketing. The team also sought to provide users with expected methods of consumer product control, requiring ground-up development of status indicators and reset buttons, edge case communications, and a wireless companion application. Not all these features are perfect in their execution, but they enable MVP-level, more real-world user interactions.

The described process has been slow and unglamorous. It is also at a much earlier stage compared to a project like Quori. To turn

[4]Photo credits: (a) Wikimedia Commons, Mary Mark Ockerbloom, CC BY-SA 4.0. (b) "Ommie: A Social Robot for Deep Breathing and Anxiety Reduction" [54], CC BY 4.0. (c) "Deploying a Robotic Positive Psychology Coach to Improve College Students' Psychological Well-Being" [43], CC BY 4.0.

Ommie into a DRP for outside labs requires technical refinement, community engagement, and individual commitments to stewardship. However, these efforts have the potential to shepherd a robot with a unique interaction modality into a sustained platform.

6.3 Jibo Research Platform

While Quori and Ommie represent custom hardware systems developed fully within academia, the Jibo Research Platform [64] illustrates an alternative path towards DRPs — reclaiming dormant industry artifacts. Following Jibo Inc.'s 2018 shutdown, members of an MIT team negotiated access to a research license and undertook a multiyear effort to re-engineer the robot's software and hardware for academic longevity. This approach aligns with what Mandel and Ju describe in HCI as "designing with what remains" [50]: leveraging existing consumer hardware for new practical purposes. This method also avoids the need for extensive ground-up hardware engineering, which is a rare skill in the current HRI community.

Park et al. [64] describes how the research team has extended multiple parts of the robot's hardware and software, inspired by internal work showcasing that a major barrier to collaborative research is the lack of a shared community platform [62]. On the software side, the group migrated core functionality to ROS, added configurable local and cloud integrations, and introduced graphical interfaces that allowed non-programmers to author behaviors. To extend the robot's hardware for research needs, the team designed a physical docking station with an attached touchscreen, additional higher-resolution camera, and extended compute capacity. This effort created a self-contained research rig, which, paired with fleet management tooling and documentation, enabled easier longitudinal, in-the-wild deployments.

These changes are the *exploitation* mindset at full work: refining and extending a proven artifact rather than inventing a new one. Equally important, and similar to the Quori team, the Jibo team has sought intentional ways to engage with community building. A recent conference workshop provided "whole product" tools, documentation, and knowledge sharing for researchers outside of the internal Jibo team [64]. One remaining hurdle is providing access to the Jibo hardware at-scale. Future work is planned to tackle such distribution, as well as provide further support for the HRI community. If this becomes possible, what was initially a source of pain and whiplash will be transformed into a platform of resilience and persistence. We can observe from all of the examples in this section, but particularly from that of Jibo – perseverance and endurance are paramount qualities when it comes to the development of DRPs.

6.4 A Note on Practical Precursors

The emergent examples described above chart three distinct paths toward DRP creation. Yet, they share two practical precursors that cannot be ignored: funding and engineering. All three cases navigated funding in nontypical ways: Quori with a unique, large NSF grant targeting this kind of work, Ommie with nonprofit support in tandem to internal funding, and the Jibo Research Platform with free access to large amounts of spare hardware. Similarly, each of these cases cultivated unique access to engineering resources. Quori hired a professional robotics development firm, Ommie's primary researcher holds a decade of industry experience, and the Jibo Research Platform internal team includes expert staff technicians beyond student researchers.

In practice, traditional HRI lab funding alone is unlikely to be able to support robust DRP creation. Furthermore, the lower-level engineering skills required are rarely present or incentivized in academic environments. We acknowledge these facts as real, immediate constraints. However, we find it critical that rather than waiting for our environments to change, we must seek to change them ourselves. The first step is in acknowledging this responsibility (Section 3). The second step is in framing a clear, actionable goal (Section 4). Acting on this goal requires critical, collective dialogue on the creative and entrepreneurial ways in which we can solve practical needs for DRP creation – whether by following the models presented above, or in forging new ones. Such investment will benefit not only the community but also individual labs themselves; building a DRP grants the developer the scientific privilege of designing a robot to fit research questions, rather than fitting research questions to the constraints of what already exists.

7 Conclusion

The field of HRI continues to build its science atop unstable foundations. Our reliance on volatile commercial robots has left our methods, data, and knowledge precariously tied to market cycles. We argue that this is not an industry failure but a research responsibility — and one that demands we start building our own products. To meet that charge, we introduced the concept of the *Deployable Research Product (DRP)*: a class of research artifact designed to bridge prototype and product, embodying real-world interaction capacity with openness and reproducibility as scientific values. To develop towards DRPs, we must look outside ourselves and bring lessons from industry within.

In the preceding section, we showed how three projects each have begun to take up this call in distinct ways. Across them, we see common threads of endurance, openness, and community stewardship emerging as forms of innovation. Yet, we are just beginning to explore what it takes to build and maintain DRPs and many open questions remain. How might DRPs be accelerated, or pushed aside, by rapid advances in AI? Is engaging with industry always necessary for DRP development? How can research groups build DRPs when funding for the sciences has recently precipitously dropped?

While the questions are many, what is clear is that endurance cannot be left to chance. The future of HRI depends on how deliberately we build and sustain the tools of our own science. Ultimately, DRPs are not merely sturdier robots; they are a statement of values and a collective refusal to let our methods vanish when the market moves on. In short, we cannot outsource what we value most.

Acknowledgments

This work was funded in part by NSF IIS-2106690, an "Envisioning Artificial Intelligence at Yale" Seed Grant, and Semio Community. The views and opinions expressed in this paper are those of the authors and do not necessarily reflect the views of the funding agencies. We would also like to acknowledge Hae Won Park, Ross Mead, and David Porfirio for conversations contributing to the development and refinement of concepts in this paper.

References

[1] [n. d.]. Introduction - Quori v1 Documentation. https://quori-robot.github.io/quori_v1_documentation/. [Accessed 17-10-2025].

[2] [n. d.]. QTrobot, an engaging educational robot for children with autism and special needs education — luxai.com. https://luxai.com/robot-for-teaching-children-with-autism-at-home/. [Accessed 24-09-2025].

[3] 2014. JIBO, The World's First Social Robot for the Home by Cynthia Breazeal. https://www.indiegogo.com/en/projects/cynthiabreazeal/jibo-the-world-s-first-social-robot-for-the-home. [Accessed 24-09-2025].

[4] 2014. JIBO, The World's First Social Robot for the Home by Cynthia Breazeal - Indiegogo. https://www.indiegogo.com/en/projects/cynthiabreazeal/jibo-the-world-s-first-social-robot-for-the-home#/section/project-story. [Accessed 17-10-2025].

[5] 2025. Community. https://quori.org/community#community-update-home. [Accessed 17-10-2025].

[6] 2025. Furhat Robotics — furhatrobotics.com. https://www.furhatrobotics.com/. [Accessed 24-09-2025].

[7] 2025. Home | Misty Robotics — mistyrobotics.com. https://www.mistyrobotics.com/. [Accessed 24-09-2025].

[8] 2025. Keepon - ROBOTS: Your Guide to the World of Robotics — robotsguide.com. https://robotsguide.com/robots/keepon. [Accessed 24-09-2025].

[9] 2025. Nao - ROBOTS: Your Guide to the World of Robotics — robotsguide.com. https://robotsguide.com/robots/nao. [Accessed 24-09-2025].

[10] 2025. Paro - ROBOTS: Your Guide to the World of Robotics — robotsguide.com. https://robotsguide.com/robots/paro. [Accessed 24-09-2025].

[11] 2025. Pepper - ROBOTS: Your Guide to the World of Robotics — robotsguide.com. https://robotsguide.com/robots/pepper. [Accessed 24-09-2025].

[12] 2025. TIAGo Pro - Empowering Mobile Manipulation — pal-robotics.com. https://pal-robotics.com/robot/tiago-pro/. [Accessed 12-01-2026].

[13] Evan Ackerman. 2018. Jibo Is Probably Totally Dead Now. https://spectrum.ieee.org/jibo-is-probably-totally-dead-now. *IEEE Spectrum* (December 2018). [Accessed 25-09-2025].

[14] Patrícia Alves-Oliveira, Matthew Bavier, Samrudha Malandkar, Ryan Eldridge, Julie Sayigh, Elin A. Björling, and Maya Cakmak. 2022. FLEXI: A Robust and Flexible Social Robot Embodiment Kit. In *Proceedings of the 2022 ACM Designing Interactive Systems Conference* (Virtual Event, Australia) *(DIS '22)*. Association for Computing Machinery, New York, NY, USA, 1177–1191. doi:10.1145/3532106.3533541

[15] Patrícia Alves-Oliveira, Alaina Orr, Elin A Björling, and Maya Cakmak. 2022. Connecting the dots of social robot design from interviews with robot creators. *Frontiers in Robotics and AI* 9 (2022), 720799.

[16] Aida Amirova, Nazerke Rakhymbayeva, Elmira Yadollahi, Anara Sandygulova, and Wafa Johal. 2021. 10 years of human-nao interaction research: A scoping review. *Frontiers in Robotics and AI* 8 (2021), 744526.

[17] BeatBots. [n. d.]. Keepon Pro. https://beatbots.net/keepon-pro. [Accessed 24-09-2025].

[18] Lauriane Blavette, Anne-Sophie Rigaud, Salvatore Maria Anzalone, Clément Kergueris, Baptiste Isabet, Sébastien Dacunha, and Maribel Pino. 2022. A robot-mediated activity using the Nao robot to promote COVID-19 precautionary measures among older adults in geriatric facilities. *International Journal of Environmental Research and Public Health* 19, 9 (2022), 5222.

[19] Andy Boucher. 2023. Research Products at Scale: Learnings from Designing Devices in Multiples of Ones, Tens, Hundreds and Thousands. In *Proceedings of the 2023 CHI Conference on Human Factors in Computing Systems*. 1–15.

[20] Fanjun Bu, Alexandra WD Bremers, Mark Colley, and Wendy Ju. 2024. Field notes on deploying research robots in public spaces. In *Extended Abstracts of the CHI Conference on Human Factors in Computing Systems*. 1–6.

[21] Hoang-Long Cao, Greet Van de Perre, Ramona Simut, Cristina Pop, Andreea Peca, Dirk Lefeber, and Bram Vanderborght. 2014. Enhancing My Keepon robot: A simple and low-cost solution for robot platform in Human-Robot Interaction studies. In *The 23rd IEEE International Symposium on Robot and Human Interactive Communication*. IEEE, 555–560.

[22] Aiden Chun, Ellie Mamantov, Ursula Stahl, Brian Scassellati, and Stephanie Leeds. 2025. Breathe Easy: Harnessing Robots for Stress Reduction During Pediatric Oral Challenges. *Journal of Allergy and Clinical Immunology* 155, 2 (2025), AB49.

[23] Carnegie Mellon University (CMU). 2013. Off To The Races. https://www.cmu.edu/homepage/computing/2013/spring/anki.shtml. [Accessed 24-09-2025].

[24] Steve Crowe. 2019. Anki addresses shutdown, ongoing support for robots. https://www.therobotreport.com/anki-addresses-shutdown-ongoing-support-for-robots/. *The Robot Report* (May 2019). [Accessed 24-09-2025].

[25] Steve Crowe. 2021. Digital Dream Labs gearing up for Cozmo, Vector relaunch. https://www.therobotreport.com/digital-dream-labs-relauch-cozmo-vector-robots/. *The Robot Report* (March 2021). [Accessed 24-09-2025].

[26] Eugene Demaitre. 2025. Aldebaran, maker of Pepper and Nao robots, put in receivership. *The Robot Report* (June 2025). https://www.therobotreport.com/aldebaran-pepper-nao-robots-receivership/

[27] Nathaniel Dennler, Mina Kian, Stefanos Nikolaidis, and Maja Matarić. 2025. Designing robot identity: The role of voice, clothing, and task on robot gender perception. *International Journal of Social Robotics* (2025), 1–22.

[28] Tom Eisenmann. 2021. *Why startups fail: A new roadmap for entrepreneurial success*. Crown Currency.

[29] Sam Francis. 2025. RobotLAB Pledges 'Zero Disruption' to Servicing Aldebaran Robots Despite Liquidation of Business. https://roboticsandautomationnews.com/2025/06/09/robotlab-pledges-zero-disruption-to-servicing-aldebaran-robots-despite-liquidation-of-business/91629/. [Accessed 17-10-2025].

[30] Marlena R Fraune, Iolanda Leite, Nihan Karatas, Aida Amirova, Amélie Legeleux, Anara Sandygulova, Anouk Neerincx, Gaurav Dilip Tikas, Hatice Gunes, Mayumi Mohan, et al. 2022. Lessons learned about designing and conducting studies from HRI experts. *Frontiers in Robotics and AI* 8 (2022), 772141.

[31] David Gouaillier, Vincent Hugel, Pierre Blazevic, Chris Kilner, Jerome Monceaux, Pascal Lafourcade, Brice Marnier, Julien Serre, and Bruno Maisonnier. 2008. The nao humanoid: a combination of performance and affordability. *arXiv preprint arXiv:0807.3223* (2008).

[32] Bruna Goveia da Rocha, Oscar Tomico, Panos Markopoulos, and Daniel Tetteroo. 2020. Crafting Research Products through Digital Machine Embroidery. In *Proceedings of the 2020 ACM Designing Interactive Systems Conference* (Eindhoven, Netherlands) *(DIS '20)*. Association for Computing Machinery, New York, NY, USA, 341–350. doi:10.1145/3357236.3395443

[33] IEEE Robots Guide. 2025. Jibo - ROBOTS: Your Guide to the World of Robotics. https://robotsguide.com/robots/jibo. [Accessed 24-09-2025].

[34] IEEE Robots Guide. 2025. Vector - ROBOTS: Your Guide to the World of Robotics. https://robotsguide.com/robots/vector. [Accessed 24-09-2025].

[35] Hatice Gunes, Frank Broz, Chris S Crawford, Astrid Rosenthal-von der Pütten, Megan Strait, and Laurel Riek. 2022. Reproducibility in human-robot interaction: Furthering the science of HRI. *Current Robotics Reports* 3, 4 (2022), 281–292.

[36] Zi-Lin He and Poh-Kam Wong. 2004. Exploration vs. exploitation: An empirical test of the ambidexterity hypothesis. *Organization science* 15, 4 (2004), 481–494.

[37] Brian Heater. [n. d.]. *TechCrunch* ([n. d.]).

[38] COMMONWEALTH OF PENNSYLVANIA By ATTORNEY GENERAL MICHELLE A. HENRY. 2024. Plaintiff, v. DIGITAL DREAM LABS, INC., and HAROLD JACOB HANCHAR, Individually and as CEO of Digital Dream Labs, Inc. https://www.attorneygeneral.gov/wp-content/uploads/2024/09/DDL-Combined-Complaint-for-FIling-1.pdf

[39] Guy Hoffman. 2019. Anki, Jibo, and Kuri: What We Can Learn from Social Robots That Didn't Make It. https://spectrum.ieee.org/anki-jibo-and-kuri-what-we-can-learn-from-social-robotics-failures. *IEEE Spectrum* (May 2019). [Accessed 17-10-2025].

[40] Sarah Homewood, Laurens Boer, and Anna Vallgårda. 2020. Designers in White Coats: Deploying Ovum, a Fertility Tracking Device. 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. doi:10.1145/3313831.3376528

[41] Miho Inada. 2021. Humanoid Robot Keeps Getting Fired From His Jobs. *The Wall Street Journal* (July 2021). https://www.wsj.com/business/humanoid-robot-softbank-jobs-pepper-olympics-11626187461

[42] Katherine Isbister, Peter Cottrell, Alessia Cecchet, Ella Dagan, Nikki Theofanopoulou, Ferran Altarriba Bertran, Aaron J Horowitz, Nick Mead, Joel B Schwartz, and Petr Slovak. 2022. Design (not) lost in translation: A case study of an intimate-space socially assistive "robot" for emotion regulation. *ACM Transactions on Computer-Human Interaction (TOCHI)* 29, 4 (2022), 1–36.

[43] Sooyeon Jeong, Laura Aymerich-Franch, Kika Arias, Sharifa Alghowinem, Agata Lapedriza, Rosalind Picard, Hae Won Park, and Cynthia Breazeal. 2022. Deploying a robotic positive psychology coach to improve college students' psychological well-being. *User Modeling and User-Adapted Interaction* 33, 2 (2022), 571.

[44] Malte Jung and Pamela Hinds. 2018. Robots in the wild: A time for more robust theories of human-robot interaction. *ACM Transactions on Human-Robot Interaction (THRI)* 7, 1 (2018), 1–5.

[45] Roshni Kaushik and Reid Simmons. 2021. Perception of Emotion in Torso and Arm Movements on Humanoid Robot Quori. In *Companion of the 2021 ACM/IEEE International Conference on Human-Robot Interaction* (Boulder, CO, USA) *(HRI '21 Companion)*. Association for Computing Machinery, New York, NY, USA, 62–66. doi:10.1145/3434074.3447129

[46] Merel Keijsers, Christoph Bartneck, and Friederike Eyssel. 2022. Pay Them No Mind: The Influence of Implicit and Explicit Robot Mind Perception on the Right to be Protected. *International Journal of Social Robotics* 14, 2 (2022), 499–514. doi:10.1007/s12369-021-00799-1

[47] Hideki Kozima, Marek P Michalowski, and Cocoro Nakagawa. 2009. Keepon: A playful robot for research, therapy, and entertainment. *International Journal of social robotics* 1, 1 (2009), 3–18.

[48] Andre W Kushniruk, Hrishikesh S Nair, and Elizabeth M Borycki. 2025. The Pepper Robot in Healthcare: A Scoping Review. *Envisioning the Future of Health Informatics and Digital Health* (2025), 101–105.

[49] Iolanda Leite, Carlos Martinho, and Ana Paiva. 2013. Social robots for long-term interaction: a survey. *International Journal of Social Robotics* 5, 2 (2013), 291–308.

[50] Ilan Mandel and Wendy Ju. 2024. Designing with What Remains. In *Proceedings of the 2024 ACM Designing Interactive Systems Conference*. 3002–3015.

[51] James G March. 1991. Exploration and Exploitation in Organizational Learning. *Organization Science* 2, 1 (1991), 71–87.

[52] Kayla Matheus, Rebecca Ramnauth, Brian Scassellati, and Nicole Salomons. 2025. Long-Term Interactions with Social Robots: Trends, Insights, and Recommendations. *ACM Transactions on Human-Robot Interaction* 14, 3 (2025), 1–42.

[53] Kayla Matheus, Marynel Vázquez, and Brian Scassellati. 2022. A social robot for anxiety reduction via deep breathing. In *2022 31st IEEE international conference on robot and human interactive communication (RO-MAN)*. IEEE, 89–94.

[54] Kayla Matheus, Marynel Vázquez, and Brian Scassellati. 2025. Ommie: The design and development of a social robot for anxiety reduction. *ACM Transactions on Human-Robot Interaction* 14, 2 (2025), 1–34.

[55] Marek Michalowski, Kyle Machulis, Mark Gasson, and Thiago Hersan. 2013. Hack "My Keepon" With an Arduino Brain - Make: — makezine.com. https://makezine.com/projects/my-franken-keepon. *Makezine* (November 2013). [Accessed 24-09-2025].

[56] Deepti Mishra, Guillermo Arroyo Romero, Akshara Pande, Bhavana Nachenahalli Bhuthegowda, Dimitrios Chaskopoulos, and Bhanu Shrestha. 2023. An exploration of the pepper robot's capabilities: unveiling its potential. *Applied Sciences* 14, 1 (2023), 110.

[57] Oliver Mitchell. 2018. Jibo social robot: where things went wrong. https://www.therobotreport.com/jibo-social-robot-analyzing-what-went-wrong/. [Accessed 25-09-2025].

[58] Geoffrey A. Moore. 2014. *Crossing the Chasm, 3rd Edition: The Updated Version of the Insightful Guide on Bringing Cutting-Edge Products to the Mainstream*. Harper Business, New York, NY.

[59] Caio Mucchiani, Pamela Cacchione, Michelle Johnson, Ross Mead, and Mark Yim. 2021. Deployment of a Socially Assistive Robot for Assessment of COVID-19 Symptoms and Exposure at an Elder Care Setting. In *2021 30th IEEE International Conference on Robot Human Interactive Communication (RO-MAN)*. 1189–1195. doi:10.1109/RO-MAN50785.2021.9515551

[60] William Odom, Ron Wakkary, Youn-kyung Lim, Audrey Desjardins, Bart Hengeveld, and Richard Banks. 2016. From research prototype to research product. In *Proceedings of the 2016 CHI conference on human factors in computing systems*. 2549–2561.

[61] William Odom, Jordan White, Samuel Barnett, Nico Brand, Henry Lin, Minyoung Yoo, and Tal Amram. 2024. Capra: Making Use of Multiple Perspectives for Capturing, Noticing and Revisiting Hiking Experiences Over Time. In *Proceedings of the 2024 CHI Conference on Human Factors in Computing Systems* (Honolulu, HI, USA) *(CHI '24)*. Association for Computing Machinery, New York, NY, USA, Article 804, 27 pages. doi:10.1145/3613904.3642284

[62] Anastasia K Ostrowski. 2023. *How Do We Design Robots Equitably? Engaging Design Justice, Design Fictions, and Co-Design in Human-Robot Interaction Design and Policymaking Processes*. Massachusetts Institute of Technology.

[63] Amit Kumar Pandey and Rodolphe Gelin. 2018. A mass-produced sociable humanoid robot: Pepper: The first machine of its kind. *IEEE Robotics & Automation Magazine* 25, 3 (2018), 40–48.

[64] Hae Won Park, Cynthia Breazeal, Sharifa Alghowinem, Anastasia K. Ostrowski, Jon Ferguson, Xiajie Zhang, and Dong Won Lee. 2024. Jibo Community Social Robot Research Platform @Scale. In *Companion of the 2024 ACM/IEEE International Conference on Human-Robot Interaction* (Boulder, CO, USA) *(HRI '24)*. Association for Computing Machinery, New York, NY, USA, 1346–1348. doi:10.1145/3610978.3638171

[65] Vatsal V Patel, Minas V Liarokapis, and Aaron M Dollar. 2022. Open robot hardware: progress, benefits, challenges, and best practices. *IEEE Robotics & Automation Magazine* 30, 3 (2022), 123–148.

[66] Vid Podpečan. 2023. Can you dance? A study of child–robot interaction and emotional response using the nao robot. *Multimodal Technologies and Interaction* 7, 9 (2023), 85.

[67] Rebecca Ramnauth, Emmanuel Adéníran, Timothy Adamson, Michal A. Lewkowicz, Rohit Giridharan, Caroline Reiner, and Brian Scassellati. 2022. A Social Robot for Improving Interruptions Tolerance and Employability in Adults with ASD. In *Proceedings of the 2022 ACM/IEEE International Conference on Human-Robot Interaction* (Sapporo, Hokkaido, Japan) *(HRI '22)*. IEEE Press, 4–13.

[68] Eric Ries. 2011. *The lean startup: How today's entrepreneurs use continuous innovation to create radically successful businesses*. Crown Currency.

[69] Raquel B Robinson, Anya Osborne, Chen Ji, James Collin Fey, Ella Dagan, and Katherine Isbister. 2024. " That's Not Good Science!": An Argument for the Thoughtful Use of Formative Situations in Research Through Design. In *Extended Abstracts of the CHI Conference on Human Factors in Computing Systems*. 1–8.

[70] Matt Simon. 2019. R.I.P., Anki: Yet Another Home Robotics Company Powers Down. https://www.wired.com/story/rip-anki-yet-another-home-robotics-company-powers-down/. *Wired* (April 2019). [Accessed 24-09-2025].

[71] Andrew Specian, Ross Mead, Simon Kim, Maja Matarić, and Mark Yim. 2022. Quori: A Community-Informed Design of a Socially Interactive Humanoid Robot. *IEEE Transactions on Robotics* 38, 3 (2022), 1755–1772. doi:10.1109/TRO.2021.3111718

[72] TIME staff. 2017. The 25 Best Inventions of 2017. *TIME Magazine* 190, 22 (2017). https://time.com/5023212/best-inventions-of-2017/

[73] Michael Suguitan and Guy Hoffman. 2019. Blossom: A handcrafted open-source robot. *ACM Transactions on Human-Robot Interaction (THRI)* 8, 1 (2019), 1–27.

[74] Daniel Terdiman. 2016. Once hot social robot Jibo said to be canceling foreign orders. https://www.fastcompany.com/4016727/once-hot-social-robot-jibo-said-to-be-canceling-foreign-orders. [Accessed 20-01-2026].

[75] Sydney Thompson, Austin Narcomey, Alexander Lew, and Marynel Vázquez. 2024. Shutter: A Low-Cost and Flexible Social Robot Platform for In-the-Wild Deployments. In *Companion of the 2024 ACM/IEEE International Conference on Human-Robot Interaction*. 94–96.

[76] Christiana Tsiourti, Anna Pillinger, and Astrid Weiss. 2020. Was Vector a Companion During Shutdown? Insights from an Ethnographic Study in Austria. In *Proceedings of the 8th International Conference on Human-Agent Interaction*. 269–271. doi:10.1145/3406499.3418767

[77] Nathan Tsoi, Joe Connolly, Emmanuel Adéníran, Amanda Hansen, Kaitlynn Taylor Pineda, Timothy Adamson, Sydney Thompson, Rebecca Ramnauth, Marynel Vázquez, and Brian Scassellati. 2021. Challenges Deploying Robots During a Pandemic: An Effort to Fight Social Isolation Among Children. 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, 234–242. doi:10.1145/3434073.3444665

[78] Akash Vidyadharan, Jennifer L Tennison, and Jenna L Gorlewicz. [n. d.]. Late Breaking Work: An Investigation into Autonomy Levels Across Telerobotic Platforms With and Without Manipulators. ([n. d.]).

[79] Astrid Weiss, Anna Pillinger, and Christiana Tsiourti. 2021. Merely a Conventional 'Diffusion' Problem? On the Adoption Process of Anki Vector. In *2021 30th IEEE International Conference on Robot Human Interactive Communication (RO-MAN)*. 712–719. doi:10.1109/RO-MAN50785.2021.9515369

[80] Rian Whitton. 2018. Whiz robot from SoftBank arrives as Pepper renewals lag. *The Robot Report* (November 2018). https://www.therobotreport.com/renewals-softbank-pepper-robot-lag

Received 2025-10-17; accepted 2025-11-21