

Creative AI for HRI Design Explorations

Hoggenmueller, Marius; Lupetti, Maria Luce; Van Der Maden, Willem; Grace, Kazjon

DOI

[10.1145/3568294.3580035](https://doi.org/10.1145/3568294.3580035)

Publication date

2023

Document Version

Final published version

Published in

HRI 2023 - Companion of the ACM/IEEE International Conference on Human-Robot Interaction

Citation (APA)

Hoggenmueller, M., Lupetti, M. L., Van Der Maden, W., & Grace, K. (2023). Creative AI for HRI Design Explorations. In *HRI 2023 - Companion of the ACM/IEEE International Conference on Human-Robot Interaction* (pp. 40-50). (ACM/IEEE International Conference on Human-Robot Interaction). IEEE. <https://doi.org/10.1145/3568294.3580035>

Important note

To cite this publication, please use the final published version (if applicable). Please check the document version above.

Copyright

Other than for strictly personal use, it is not permitted to download, forward or distribute the text or part of it, without the consent of the author(s) and/or copyright holder(s), unless the work is under an open content license such as Creative Commons.

Takedown policy

Please contact us and provide details if you believe this document breaches copyrights. We will remove access to the work immediately and investigate your claim.

Green Open Access added to TU Delft Institutional Repository

'You share, we take care!' - Taverne project

<https://www.openaccess.nl/en/you-share-we-take-care>

Otherwise as indicated in the copyright section: the publisher is the copyright holder of this work and the author uses the Dutch legislation to make this work public.



Creative AI for HRI Design Explorations

Marius Hoggenmueller
marius.hoggenmueller@sydney.edu.au
Sydney School of Architecture, Design and Planning
The University of Sydney
Sydney, NSW, Australia

Maria Luce Lupetti
m.l.lupetti@tudelft.nl
Faculty of Industrial Design Engineering
Delft University of Technology
Delft, The Netherlands

Willem van der Maden
w.l.a.vandermaden@tudelft.nl
Faculty of Industrial Design Engineering
Delft University of Technology
Delft, The Netherlands

Kazjon Grace
kazjon.grace@sydney.edu.au
Sydney School of Architecture, Design and Planning
The University of Sydney
Sydney, NSW, Australia



Figure 1: We explored whether and how the use of creative AI tools can support design ideation in HRI. To do so, the first three authors used generative text-to-image models to envision robotic artefacts and reflect on robot sociotechnical imaginaries [41]; over four weeks, they exchanged images and textual annotations of the first-person design exploration via digital postcards.

ABSTRACT

Design fixation, a phenomenon describing designers' adherence to pre-existing ideas or concepts that constrain design outcomes, is particularly prevalent in human-robot interaction (HRI), for example, due to collectively held and stabilised imaginations of what a robot should look like or behave. In this paper, we explore the contribution of creative AI tools to overcome design fixation and enhance creative processes in HRI design. In a four weeks long design exploration, we used generative text-to-image models to ideate and visualise robotic artefacts and robot sociotechnical imaginaries. We exchanged results along with reflections through a digital postcard format. We demonstrate the usefulness of our approach to imagining novel robot concepts, surfacing existing assumptions

Permission to make digital or hard copies of all or part of this work for personal or classroom use is granted without fee provided that copies are not made or distributed for profit or commercial advantage and that copies bear this notice and the full citation on the first page. Copyrights for components of this work owned by others than the author(s) must be honored. Abstracting with credit is permitted. To copy otherwise, or republish, to post on servers or to redistribute to lists, requires prior specific permission and/or a fee. Request permissions from permissions@acm.org.
HRI '23 Companion, March 13–16, 2023, Stockholm, Sweden
© 2023 Copyright held by the owner/author(s). Publication rights licensed to ACM.
ACM ISBN 978-1-4503-9970-8/23/03...\$15.00
<https://doi.org/10.1145/3568294.3580035>

and robot stereotypes, and situating robotic artefacts in context. We discuss the contribution to designerly HRI practices and conclude with lessons learnt for using creative AI tools as an emerging design practice in HRI research and beyond.

CCS CONCEPTS

• Human-centered computing → Interaction design process and methods.

KEYWORDS

human-robot interaction, creative AI, generative AI, text-to-image models, ideation, sociotechnical imaginaries, design research

ACM Reference Format:

Marius Hoggenmueller, Maria Luce Lupetti, Willem van der Maden, and Kazjon Grace. 2023. Creative AI for HRI Design Explorations. In *Companion of the 2023 ACM/IEEE International Conference on Human-Robot Interaction (HRI '23 Companion)*, March 13–16, 2023, Stockholm, Sweden. ACM, New York, NY, USA, 11 pages. <https://doi.org/10.1145/3568294.3580035>

1 INTRODUCTION

A driving force in human-robot interaction (HRI) research is the recognition of robots as complex sociotechnical systems [8]. Robots designed to interact with humans are no longer perceived as tools

that simply extend our capabilities, but rather as counterparts with whom we engage in alterity relations [32]. However, designing for robotic otherness [9] and envisioning novel human-robot relations in real-world settings [8] appears to remain a difficult task. Despite recent endeavours towards a paradigm shift (e.g., advocating for designerly [50] and critical research approaches in HRI [68]) and certain design exemplars of alternative robot designs [22, 39, 46, 53], the majority of HRI designs centre on a limited design space of simple anthropomorphism or zoomorphism [32]. Researchers argue that this is due to the fact that the notion of ‘robot’, culturally and historically conditioned, holds many pre-existing assumptions about what a robot should do, look like, and behave like [2, 52, 64]. While HRI design seems particularly prone to what is commonly referred to as ideation or design fixation, this phenomenon has been recognised in many fields and domains of design [40, 55, 70]. Considered a cognitive bias [11], design fixation suggests that designers tend to adhere to pre-existing ideas or concepts that ultimately constrain design outcomes [55]. While design creativity and innovation have been extensively researched, overcoming fixation and stimulating ideation in design is an ongoing research topic that is increasingly being addressed through computational approaches, such as co-creative artificial intelligence (AI) systems [11].

This paper contributes to growing HRI design practices that integrate design research methods and exploratory approaches for ideation and conceptualisation of robotic designs [2, 50, 52, 81]. At the same time, we consider the generative potential of creative AI tools [56, 75], which is predicted to innovate creative processes and design practice at large [36, 37]. Specifically, we examine whether and how the use of generative text-to-image models (cf. [59, 62]) can improve and facilitate design ideation processes in HRI, for example, to overcome design fixation around robot assumptions and to inform divergent and desirable visions of robotic futures.

We frame our investigation around the theoretical concept of ‘sociotechnical imaginaries’. Jasanoff and Kim [41] define sociotechnical imaginaries as “collectively held, institutionally stabilized, and publicly performed visions of desirable futures, animated by shared understandings of forms of social life and social order attainable through, and supportive of advances in science and technology”. As such, sociotechnical imaginaries can be employed as a design resource to inspire the future of technological innovation (e.g., robotics), for example, by drawing on or deliberately counteracting the imaginations of media and popular culture [42]; furthermore, the concept offers an analytical lens to examine how an idea or concept surrounding a sociotechnical assemblage aligns with or challenges a collective vision.

The rapid development and increasing accessibility of generative AI systems, such as Stable Diffusion [62] and DALL-E [59], brought forward a range of examples that demonstrate how these tools can help to visualise sociotechnical imaginaries, share them, before they eventually get materialised. For example, a Brooklyn-based artist used the DALL-E 2 image generator to transform Google Street View images of U.S. streets into pedestrianised areas (e.g., by integrating AI-generated images of public plazas or shady parks into the original footage). Engaging with the sociotechnical imaginary of the ‘car-free city’, the images that the artist posted under the twitter pseudonym “betterstreets.ai” have gone viral, helping people to imagine how their places could look in a pedestrian-centric

world¹. Another example comes from design researcher Dan Lockton and social scientist Josie Chambers who used generative image generation to integrate wind turbines and photovoltaic plants into landscape masterpieces from Van Gogh and Vermeer. Accompanied by provocations in the form of short texts², the researchers engage with the sociotechnical imaginary of ‘energy transition’ by immersing the reader in an alternative world in which humanity shifted to renewable energies centuries ago. Their example demonstrates how sociotechnical imaginaries can also engage with future visions by waiving in and altering narratives from the past. Our work builds on this growing body of creative explorations in which AI is used to surface the sociotechnical imaginary of possible technological futures, to understand what the contributions of these practices would be for the HRI field, as well as the implications for the role of the designer in such processes.

The remaining parts of the paper are organised as follows: we first present current design research in HRI, followed by a review of co-creative AI tools. We then outline our first-person design exploration, in which the first three authors (Hoggenmueller, Lupetti, and Van der Maden) used generative text-to-image AI tools to ideate and visualise robotic artefacts and their surrounding imaginaries; exchanged the resulting images along with textual reflections weekly for four weeks (see Fig. 1); and concluded with a focus group discussion to further examine the generated HRI imageries and the experiences of collaborating with generative AI systems. Next, we present the findings from our design exploration along three themes: the use of text-to-image models to imagine novel robot concepts, surfacing existing assumptions and robot stereotypes, and situating robotic artefacts in context. Reflecting on our approach, we discuss how designers can use creative AI tools as an emerging design practice for HRI research and beyond.

2 BACKGROUND

2.1 Designerly HRI

We recently witnessed a spring of designerly HRI work—human-robot interaction research with a strong orientation towards design [50]. On the one hand, design practices such as participatory design are now regularly adopted as a means of bridging the gap between the technical research interests that drive most engineering HRI approaches and the needs and realities of potential users [6]. On the other hand, the distinctive designerly ways of producing knowledge [18] have been increasingly proposed as a meaningful perspective to understand and learn from robotic artefacts [50].

The HRI field, in fact, is very prolific in generating artefact-centered contributions, yet these tend to have a significantly lower impact on the community (e.g., a lower number of citations) [16]. Thereafter, a plurality of voices is now trying to unpack the value of designerly work, especially artefact-centered contributions, for the HRI field. Zamfirescu-Pereira et al. [81] suggest that design explorations, and exploratory prototyping techniques more specifically, are critical practices for “devising new robot forms, actions and behaviors, and for eliciting human responses to designed interactive features.” Relatedly, Lupetti et al. [50] suggest that “the main contribution and influence of – some – HRI artifact-centered works

¹<https://betterstreets.ai/>, last accessed: November 2022

²https://twitter.com/imaginari_es/status/1553130066829262849, last accessed: November 2022

lies within the way their embodiments and behaviors challenge what we believe a robot should look, act, and be like”, as such “exploring the conceptual implications of novel artifacts can lead to a different understanding and shaping of what our future with robots might be”. In this line of thinking, Alves-Oliveira et al. [2] explored the use of metaphors as a design strategy to generate novel robot ideas, reframe existing HRI problems, and—most importantly—to surface and challenges current assumptions that accompany the term ‘robot’ with regards to its physicality, roles, and behaviors. Learning from this designerly exploration of metaphorical associations and building on Auger’s [3] speculative design work, the authors argue that “design can answer—up to a large extent—to the growing demand for reflections and investigations on the impact and desirability of the robotic futures we propose”.

Hoggenmüller et al. [34] further explored what and how we can learn from robotic artefacts by developing an annotated portfolio of ludic urban robots and inviting other HRI experts to annotate the robotic artefacts. By doing so, they illustrate how this kind of practice enables a form of creative and generative peer critique on the contribution of a work, helping to “capture and disseminate designerly contributions”. Designerly HRI, then, constitutes a body of artefact-centered work that is often difficult to legitimize, yet carries high conceptual value. Especially for its capacity to surface and challenge stereotypical ideas and assumptions, it offers potentially meaningful ways for reframing dominant robot sociotechnical imaginaries, and envisioning novel HRI futures.

2.2 Co-creative AI tools for Designerly Practices

Cascini et al. [11] have discussed the potential of AI in design, noting that it lies in its ability to assist in the conceptual phase of design. Generative models (GMs) play a key role in this, as they are able to learn from and generate data that is similar but not simply a copy of the training data. Due to recent advancements in machine learning (ML), GMs have become radically accessible, meaning that basically anyone with a computer and an internet connection can train and use them. With this in mind, many interaction designers, artists, and researchers have started to use GMs in their work. Example applications are drawing [38], creative writing [73], generating Emmy-worthy song lyrics [47], aiding architects in their design process [7], product design of fashion bags [20] and chairs [57], and the general design process [61].

While there are numerous types of GMs, generative adversarial networks (GANs) are the most predominant AI technique used in the creative domain. These models can be used to accurately replicate a training set with a certain degree of novelty, however, it is often limited to one content type (e.g., faces, classic piano music, spreadsheet data) and adversarial training requires a lot of time and resources. While reviews recognise GANs’ potential for incorporation in design pipelines [35] and HCI researchers have begun to develop AI-assisted ideation tools using GANs (e.g., for news illustrations [45]), the community also critically discusses whether these tools limit designers in exploring more exciting directions [76]. Furthermore, a preliminary review by Hwang [37] examining the product landscape of current AI-empowered co-creative tools found that only few support the early stages of creative processes. They refer to those as the Q&A and wandering stages of the creative process, where the former refers to understanding the creative problem

at hand and gathering information, and the latter to playing around with preliminary ideas and formulating creative strategies.

It is worth noting that the aforementioned studies have all been published before the “AI Art” boom in mid-2022. The rapid development of a technology called Diffusion Models (DMs) fueled this hype, with DMs having a wide range of applications from computer vision to multimodal learning (for a preliminary review see Yang et al. [80] or Croitoru et al. [17]). Text-to-image generation seems to be the most common application of DMs, with examples being OpenAI’s DALL-E2, Google’s Imagen, and StabilityAI’s Stable Diffusion. The public release of Stable Diffusion sparked a lot of interest in the potential of this technology. Because the model was released open source, third-party researchers started building upon the technology, resulting in a range of new applications for DMs, such as video editing [62], 3D rendering and animation [58], and model “personalization” [27, 63]. This is resulting in a growing body of scientific literature on the use of co-creative AI, in design, design education, and other creative practices [24, 28, 69]. Nevertheless, the specific use of DMs for confronting designers’ pre-existing assumptions and related socio-technical imaginaries remains under-explored.

3 A DESIGN EXPLORATION OF CREATIVE AI FOR HUMAN-ROBOT INTERACTION

To investigate the opportunities and implications of using generative AI tools to ideate and visualise robotic artefacts and their surrounding imaginaries, we chose a first-person research through design approach. First-person methods are enjoying growing popularity in HCI [13, 15, 30, 49, 54], with autobiographical accounts recently also being recognised as a rich resource for HRI design [14]. As introspective research practices have been acknowledged to bring a unique value for experience-driven design [79], we decided on a self-reflective, first-person approach as it would allow us to report on the genuine use of creative AI tools over an extended period of time. This approach, as opposed to more targeted and singular evaluations with external experts or users, would also allow for more authentic reporting on the experiences of integrating creative AI into HRI design practice.

The insight to draw on rich qualitative data from individual accounts emerged organically at an early planning phase of this research when the first two authors reached out to each other and exchanged image generations of robotic artefacts from personal exercises working with collaborative AI tools (e.g., Deep Dream, Artbreeder, and Stable Diffusion). The two authors shared their image generation, accompanied by longer text messages explaining their designerly intentions and how the results inspired their thinking about new concepts for robot design, along with self-reflective accounts of using the AI tools. What emerged from these initial conversations was a design exploration of creative AI for HRI that lasted for a total of two months. Below, we outline the different stages of the design exploration, including a preparation workshop, weekly exercises that included the exchange of ‘CreativeAI Postcards’, and a final focus group discussion.

3.1 Preparation

Before commencing with the actual design explorations, we organised a one-hour workshop held via Zoom to discuss the protocol and scope of the study.

3.1.1 Methodological Approach. The three participating design researchers (the first, second, and third authors) agreed on following a methodological approach inspired by the “Dear Data” book [51]. “Dear Data” was a one-year-long project in which information designers Stefanie Posavec and Giorgia Lupi exchanged postcards with hand-drawn data visualisations about their lives on a previously agreed-upon topic. The approach since then has also been adopted in first-person HCI design research, such as that of Friske et al. [26], who exchanged artefacts of hand-crafted physicalizations and sonifications. Importantly, the “Dear Data” project by Posavec and Lupi tested how effectively the two designers, who had never met before, could get to know one another by reading their data visualisations. While we were not opposed to attending to personal accounts, our main goal in adopting this approach was to agree on a more structured way to document our self-reflective practice and, by exchanging our ‘CreativeAI Postcards’, to learn from each other’s approaches and experiences of cooperating with creative AI tools. Given the fact that we would be designing with digital tools and that the three authors live on different continents, we agreed to use a digital format, the online whiteboard Miro, to exchange the postcards. To account for the temporal affordances of sending and receiving postcards, however, we decided on a set day every week in which we would post the ‘CreativeAI Postcards’. By doing so, we could also balance the influence of our collaboration with the autobiographical elements [26]. We designed a postcard template in Miro that included headers and space for the agreed-upon information (see Fig. 1): the image output together with the utilised prompt and three brief paragraphs of up to a maximum of 150 words each about a) the designerly intention, b) the reflection in regards to HRI design (e.g., what is striking about the output; how it aligns with existing or inspires new robot concepts and imaginaries), and c) the reflection in regards to working with the creative AI tool.

3.1.2 Robot Sociotechnical Imaginaries. In the preparation workshop, we also agreed on the theoretical framing of our design exploration around sociotechnical imaginaries. We discussed the concept and identified relevant literature [5, 41, 42], ensuring that all three participating researchers could become familiar with it. Chien and Hassenzahl [15] previously emphasised the importance of interpreting individual accounts of first-person research through design approaches in the light of theoretical knowledge. They argue that only in this way can detailed accounts contribute to existing knowledge of design research rather than simply remaining design attempts. Aiming to design robotic artefacts through the lens of sociotechnical imaginaries, we brainstormed different topics and contexts of robotic futures with high relevance to the HRI field and impact on society and communities at large. We agreed on four topics, including urban robotics, healthcare robotics, collaborative robotics in work contexts, and domestic robotics.

3.1.3 Text-to-Image Models. After the workshop, the first author cloned and prepared Google Colab notebooks to run the previously agreed-upon text-to-image models, namely Stable Diffusion³ and VQGAN+CLIP⁴. We extended the python code for the cloned notebooks in order to automatically save generated images into each of the designer’s private Google Drive folders. We decided on the

³https://colab.research.google.com/github/huggingface/notebooks/blob/main/diffusers/stable_diffusion.ipynb, last accessed: November 2022

⁴[https://colab.research.google.com/github/justinjohn0306/VQGAN-CLIP/blob/main/VQGAN%2BCLIP\(Updated\).ipynb](https://colab.research.google.com/github/justinjohn0306/VQGAN-CLIP/blob/main/VQGAN%2BCLIP(Updated).ipynb), last accessed: November 2022

aforementioned generative AI tools as, at the time of the research project, other tools such as DALL-E 2 were only accessible through a waitlist. However, one of the researchers had access to DALL-E 2 at the time and we agreed that they could use it for the design exploration. It should be noted that one important aspect in the performance of the different models is the dataset used for training. The data quality can greatly impact the model’s ability to generate new and diverse outputs. While OpenAI is opaque about the dataset to train DALL-E 2, merely reporting on its size (650 million text-image pairs) [60], StabilityAI disclosed to have used an aesthetic subset of LAION-5B [66], LAION-2B-en. This set contains 2 billion images with a predicted rating of 5 or higher on the question “How much do you like this image on a scale from 1 to 10?” [66].

3.2 Weekly Exercises

For a total of four weeks, we completed the design explorations as weekly exercises. Each week was dedicated to one of the predetermined topics. We did not specify the number of hours per week that each designer should work with the creative AI tools, but everyone agreed to create at least two ‘CreativeAI Postcards’ every week / per topic. It should be noted that not each generated image would be used for a postcard, but it was up to the discretion of every designer to decide on the images that they would like to reflect on and share with the group. The postcards were added to the Miro board every Friday. Once posted, everyone could read each other’s postcards and also leave comments on the board through post-it notes.

3.3 Focus Group Discussion

After the four-week period of posting weekly creative AI postcards, we conducted a follow-up focus group discussion lasting for two hours in total. The focus group discussion was held via Zoom and supported by the Miro board. In the first half of the session, the discussion focused on the designer’s intentions when developing their prompts, the image results, and what they have learned in regards to conceptual implications and HRI imaginaries. In regards to the image results, we particularly discussed what aspects were surprising or even challenged the designer’s existing assumptions about robots and their imaginations, and what aspects aligned with their expectations. We also began to cluster images and draw connections between postcards that would entail similar ideas or image outcomes. In the second half of the session, the discussion evolved around the designers’ experiences of engaging in a co-creative design process with the generative AI tools, touching on conceptual aspects such as sense of agency and creativity but also more practical findings, e.g., around prompt engineering [44].

3.4 Design Researcher Positionality

The first three authors of this paper were involved in the preparation session, the weekly exercises, and the post-study focus group discussion. All of them refer to themselves as “design researchers” and work in the broader fields of HRI and/or HCI. However, all of them also pursue their individual research interests and bring in different professional, disciplinary, and cultural backgrounds. Stemming from feminist standpoint theory [4], researchers in HCI have increasingly argued for the importance of disclosing designers’ and researchers’ positionality in order to acknowledge their backgrounds, values, beliefs, and biases. This is especially important

when designing with generative AI systems, which are known to amplify human biases, raising ethical concerns about their use [56]. In addition to acknowledging positionality, HCI researchers have also argued that positionality can have a leveraging effect that enriches design outcomes, such as when a diverse group of designers works together on open-ended design explorations [30]. While agreeing on this rough framework (presented above) to carry out the design exploration study, we also ensured to maintain enough flexibility so that each design researcher could bring in their own perspectives and knowledge practices. This also aligned well with our exploratory research approach. Below, we present the researchers' positionality written as a first-person statement. In the course of this paper, we refer to the participating design researchers with acronyms (e.g., 'DR1' for 'Design Researcher 1').

Design Researcher 1. I am an interaction design researcher with a strong interest in designing and prototyping emerging technologies. I have been working in the domain of urban robots for the past four years, from mundane applications such as the interface design of autonomous vehicles to more playful applications studying the experiential qualities of interaction with urban robots. Building on my previous expertise, my recent interest lies in leveraging AI to improve prototyping processes for complex technologies.

Design Researcher 2. I am a researcher working at the intersection of design and HRI/AI fields. I am mostly interested in understanding how socio-technical imaginaries shape the way we introduce advanced technologies in society, and in exploring ways to reframe technological narratives through design. As such, I approached the whole exercise to really unpack what is the robot imaginary 'out there' and to see if and how the generated images would match or challenge what is my own imaginary of robots.

Design Researcher 3. I am interested in exploring the opportunities of generative models in design and beyond, and understanding their current limitations and weaknesses. This may contribute to laying out future directions for research and investigations into AI alignment with human intent. In the present design explorations, this has been particularly addressed through the comparison between models and the use of newly developed extensions of the technology such as DreamBooth and Textual Inversion.

4 FINDINGS

Our weekly exercises culminated in the design and exchange of 40 postcards (see two examples in Fig. 1), with DR1 and DR2 each designing 16 postcards and DR3 designing 8 postcards. Stable Diffusion was mainly used by DR1 and DR2 with a total of 471 image generations. DR3 mainly used DALL-E 2 with a total of 620 image generations. VQ-GAN has only been utilised by DR1 for 18 image generations, one of which resulted in the creation of a postcard.

In this section, we present three themes regarding the usage of generative image models for HRI design that we identified through our first-person design exploration and reflections: imagining novel robot concepts, surfacing existing assumptions and robot stereotypes, and situating robots in context. We identified these themes through reviewing and clustering the generated images, as well as affinity mapping [48] of the textual accounts from the postcards. Transcriptions from the post-study focus group were used to gain further insights into particular patterns and add to the designer's

intentions and reflections. Rather than presenting postcard results for each topic individually, we chose a high-level analysis of results; although insightful and diverse in perspectives, the ideation results of three designers would be insufficient to map out an exhaustive HRI design space. Furthermore, this focus aligns with our overarching goal to examine whether and how the use of generative models can support design ideation in HRI. Fig. 2 shows a larger selection of generated images to contextualise results along specific designs.

4.1 Imagining Novel Robot Concepts

Our design explorations provided a wide range of instances in which the generated images stimulated our thinking towards novel robot concepts and imageries. For example, we noticed some "pleasant surprises" (DR2) when the depicted robotic artefact would yield some form of robot hybridity [46], i.e. the blending of robot characteristics or functionalities with unrelated concepts in novel ways (see Fig. 2, DR1 in Week 3, DR2 & DR3 in Week 4). Other instances include those that we referred to as 'robot/non-robot'-imaginaries, i.e. when the typical robot characteristics (e.g., human- or animal-like) were reduced to the point that the robot became almost invisible, thus challenging our collectively held and stabilised imagination of what a robot should look like. For example, DR1 (see Fig. 2, Week 4), prompting a domestic robot to interact with a lonely person, identified a round disk-like object attached to a long stick as the robotic artefact. Subsequently, with Fig. 3, DR1 expanded the image collection to search for more 'robot/non-robots'. This caused DR1 to reflect how existing products could integrate robot capabilities and how seamless interactions could be established with these robotic objects, feeding the imaginary of pervasive AI that is embodied yet disappears. DR2 made similar observations and commented that "[...] sometimes the robot is indeed missing, but there is a fine line between missing and being somewhere there." Furthermore, our design explorations revealed the potential of using creative AI tools for imagining robotic appearances inspired and shaped through different temporal contexts (see Fig. 1, left and Fig. 4), inviting for robot time travel [53] in the future and the past.

In some instances those unprecedented robot appearances would emerge simply by chance, adopting the prevalent human-AI interaction paradigm of "pressing a button and seeing what happens" [11]; however, we achieved those results more effectively when consciously prompting for serendipity, which often preceded a trial-and-error process to understand and adjust prompt components and, once recognised a pattern to make use of it. Most importantly, however, in order to imagine novel robot concepts, the generated images primarily served the purpose of providing a thought-provoking impulse. As designers we then had to cluster and annotate across generated images in order to derive robot concepts and imaginaries for extending the existing HRI design space and assess not only 'novelty' of the design but also potential 'value' [31]. As such, text-conditioned generative image models offer a strategy to speculate on futuristic designs [23] through visual interpretations.

4.2 Surfacing Existing Assumptions and Robot Stereotypes

Although our design explorations resulted in outcomes with the potential to inspire novel robot concepts and surrounding imaginaries, we also discovered that many existing assumptions and robot




Figure 2: Selected images from the ‘CreativeAI Postcards’ with prompts (in-image) and excerpts of designer’s reflections (below).

stereotypes were reinforced through the generated images with outcomes that were neither ‘novel’ nor ‘surprising’ [31]. For robot appearance and embodiment, this was in particular true when not further specifying the characteristics of the robot within the prompt (e.g., simply using the term ‘robot’ or ‘social robot’) and if the situated context was rather neutral (e.g., ‘at work’), thus not evoking a strong specificity that would influence the appearance of the robotic artefact or its embedding. For example, this would lead to image generation in which the robotic artefact would be depicted as blue-and white-coloured, confirming the predominant racialisation of AI as White [12]. Due to the fact that text-to-image models are trained on real-world datasets, we also found that including in the prompt a context where robots are already widely deployed, such as factories, would lead to more realistic image outputs (see Fig. 2, DR2 in Week 3). On the other hand, robots interacting closely with humans

(see Fig. 1, right & Fig. 2, DR2 in Week 2) frequently resulted in “clearly wrong compositions”, surfacing the stereotypical imaginary that robots are still more often servants than companions [2].

The ability to surface human biases became particular evident through one example in which DR1 prompted ‘a robot that is a sex worker’. Here, even social stereotypes – assuming that sex workers are mostly female – were reinforced and blended into the robot appearance. While highly problematic when not being critically assessed [71], we discussed that from a design perspective, generative image models could help to identify biases and stereotypes, thus providing a first step forward to confront them. For example, through designerly action by considering feminist and non-heteronormative perspectives for the design of robots [21, 78]. As such, we found that the responsible use of text-to-image models as a data-driven approach could complement existing methods in




Figure 3: Domestic robots with minimal anthropomorphic features could inspire the design of pervasive AI that blends with mundane objects.

HRI (e.g., metaphors [2]) for surfacing and challenging robot stereotypes as well as their surrounding and embedded social context.

4.3 Situating Robots in Context

Our design exploration has demonstrated the potential of image generative models to virtually situate robotic artefacts in various contexts, scenarios, and situations at a stage of the design process where the artefact does not yet exist in reality. In a matter of a few seconds we were able to situate companion robots into domestic homes, nursing robots into care homes to interact with the elderly, or robotic water dispensers into dried up urban landscapes, and so on. In this regard, DR2 reflected that despite “limited fidelity in terms of captured complexity” (i.e., some of the image compositions were clearly wrong in particular for depicting interactions and relations between robots and humans), “seeing the robots in context, which doesn’t happen often to [them], helps to visualise and account for the elements, actors, and some of the dynamics of the contexts that we are designing for.”

While our design exploration illustrated an early design ideation phase, we anticipate the potential of virtually situating robotic artefacts in context at various stages of the HRI design process. For example, DR1 and DR3’s experimentation with textual inversion, i.e., training and embedding additional concepts from a small number of personal images (see Fig. 5), could be employed as follows: HRI designers could take pictures of their physical prototypes, extend the pre-trained text-to-image model, and then virtually situate and interpret those prototypes in a variety of new contexts.

5 DISCUSSION

Building on a growing body of designerly HRI practices and methods [50], we explored the use of generative text-to-image models as a way to challenge dominant robot sociotechnical imaginaries and envision novel HRI concepts. We documented a variety of instances that can be considered ‘novel’, ‘valuable’, or ‘surprising’ (some of the key factors identified for computational creativity research [31]). We acknowledge, however, that the notion of creativity we build upon is not universally recognised or defined [74], our design exploration is far from being exhaustive, and other configurations or alternative text-to-image models – especially in light of the fast-paced and steady improvement of text-conditioned generative




Figure 4: Embedding temporal contexts to inspire robot designs: here, a delivery robot and drone by German Renaissance artist Albrecht Dürer (1471–1528).




Figure 5: Experiments by DR3 with textual inversion to visualise the Pepper robot in different contexts.

models [38] – may yield more unexpected results. Nevertheless, our work exemplifies how AI-generated images could serve as inspirations for HRI design instances; especially if looked at collectively, these can be used to conceptualise novel robot appearances and roles [34, 50]. And most importantly, the designer-AI co-creative process becomes itself a site for confrontation with personal and collectively held biases and (limited) imagination. As DR1 noticed “while I was initially aiming to prompt for the ‘novel’ and ‘surprising’, I soon became aware that many generated images delivered expected results and confirmed existing robot stereotypes”. Instead of looking at these results as failures, we rather suggest to embrace those ‘anticipated’ or ‘biased’ outputs. As these models are trained on real-world datasets and thus incorporate human biases [67, 72], they can help to surface, identify, and confront existing assumptions and preconceptions. As such, using generative AI models can be an effective approach to make robot and social stereotypes visible [2] in order to then challenge them through designerly action.

This work also opened up possibilities to anticipate and discuss the possible impact of these technologies in specific contexts, as often happens with video-based and simulation-based studies [1, 81]. As a matter of fact, despite efforts to deploy and test robotic artefacts in the ‘real-world’ (e.g. [25, 33, 65]), the majority of human-robot interaction research is still undertaken in a laboratory setting [43]. This practice is widely discussed in regards to empirical HRI findings, raising questions about their external and ecological validity, i.e., whether they are applicable to other contexts and hold true for natural situations [19]. However, as pointed out by Lupetti et al. [50], the root of the issue resides partially in the conceptualisation and design process itself, in which “most often users and contexts are considered only at later stages of research”. Our design exploration demonstrated the potential of situating robot artefacts in

context. The expected advancements in creative AI tools, such as the ability to generate videos [62] and provide designers with greater control [37], could further simplify and optimise the creation of scenario-based prototype representations already used in HRI evaluation, such as video prototypes [81] and virtual simulations [1].

5.1 Creative AI as Emerging Design Practice: Lessons Learned

Following a first-person research through design approach, we documented self-reflective accounts about the design experiences and practice of working with creative AI tools.

5.1.1 Tool Familiarisation. All three participating designers reported on the importance of *familiarising* with the image generative models. In particular, at the beginning of the weekly exercises, this meant understanding what image outputs to expect. DR2, for example, reported on “micro investigations”, which sometimes meant to just exchange the order of words within a prompt and learn through comparing and interpreting the image results. DR1, on the other hand, wanted to understand whether the AI would share a similar imaginary in relation to terms that are common within HRI to describe the appearance of a robot, e.g., ‘humanoid’ and ‘non-humanoid’, ‘anthropomorphic’ and ‘non-anthropomorphic’, and so on. DR3, instead, regularly explored how stylistic choices could help envisioning and conceptualising future HRI scenarios, such as the use of white background that, if associated to a specific environment in the prompt would help situating robots within a specific environment, but would not facilitate visualising the robot execution of tasks. Understanding the limitations, such as the fact that negations of robot appearances (e.g., ‘non-anthropomorphic’) would not lead to (desired) results of abstract robot appearances (see Fig. 2, DR3 in Week 1); or the difficulty of visualising human-robot interactions and relations, DR1 concluded that “creative and intentional thinking from the designer is not only still needed, but central to this kind of emerging creative practices.” While an iterative and incremental approach to learning and understanding the tools was particularly vital at the outset of the design exploration, even later on, one selected picture for a ‘CreativeAI Postcard’ was often the results of many trials. As such, our reflections connect and add to the broader disciplinary discussions that are currently happening in the related fields of design and HCI, where the probabilistic nature of AI and ML algorithms is asking for a radical rethinking of the role of the designer [28, 29].

5.1.2 Designer as Curator and Interpreter. From the reflections on the postcards and the post-study focus group discussion, two main patterns about the designer’s role in the co-creative AI process emerged: *the designer as curator* and *the designer as interpreter*. The *curator* role was manifested by looking for a “sense of direction” (DR3), e.g., by consolidating HRI design research papers (DR2), newspaper articles (DR1), or simply using google image search (DR1, DR3). This can help to spark inspiration for prompts but also to consciously counteract predominant narratives. DR3 emphasised the importance of “having some vision or initial idea what to prototype and look for [...] in order to visualise your imagination and get the ‘snowball’ running.” As the curator role implies that an initial impulse is necessary, we anticipate positive synergy effects of using

text-to-image models in combination with existing ideation tools, such as new metaphors [2]. Secondly, in the role of the *interpreter*, the designer must look for patterns and add possible lenses in the way we look at the images. Importantly, the generated images considered in isolation did not necessarily hold already a conceptual value for HRI design [50], but it was in combination with the designer’s annotations (e.g., designerly intention and reflection), as well as looking at and interpreting a collection of images. DR2 describes the interpretive aspect along the ‘robot/non-robot’ imagery as follows: “You start looking at what is there as the robot, which challenges the stereotype of robots. [...] You break free from that just because you are looking for where the robot could be.” While initially concerned about reaching a shared understanding with the AI, in the role of the interpreter, DR1 concluded in regards to the ‘robot/non-robots’: “If you would feed this image into an image classifier, probably the AI would not recognize it as a robot? [...] I started to wonder if the AI, when generating this image, ‘thinks’ about a robot or not? But in the end, it doesn’t really matter. What matters is that I get creative input through the generated images.”

5.1.3 Human-human Collaboration. Research on human-AI interaction emphasises the need for more collaborative and dialogic interaction models [10, 38], responding to the fact that the majority of present co-creative AI tools offer limited interaction and intervention opportunities for designers [37]. Some of the limitations discovered through our design exploration and reported throughout this paper certainly support this. However, in addition to collaborative human-machine partnerships, our approach also demonstrated the importance of human-human collaborations during co-creative AI design processes: in our design exploration study, each of the participating designers brought in their own personal references, enriching ideation and design outcomes; at the same time, through our approach of posting ‘CreativeAI Postcards’ (inspired by Lupi and Posavec [51]), we could learn from each others experiences and adopt practices. Human-human collaborations, we argue, can further help to call attention to ethical issues in creative AI, evident in our study when one of the designers pointed out the bias towards Whiteness in depictions of AI [12] but also human representatives, which, admittedly, has not stuck out to the other designers initially. Future co-creative AI tools should therefore not only support improved human-machine interactions but also leverage human-human collaborations.

6 CONCLUSION

Responding to the theme of this year’s Human-Robot Interaction conference, “HRI for all”, and Alt.HRI’s dual perspective, “ALL for HRI”—aiming to expand the ways of building knowledge in our community—in this paper, we explored the potential of creative AI tools as a way to challenge dominant robot sociotechnical imaginaries and envision novel HRI futures. In light of recent advancements in creative AI tools [37], our self-reflective approach provides initial insights on how designerly HRI can leverage those to ideate and design better robots. While in this paper we explored the use of text-to-image generation with a strong focus on conceptual implications, we believe that there is a wide range of other application areas for creative AI tools in HRI, for example, prototyping robot animations and movements by means of computer vision [77].

REFERENCES

- [1] Almohannad Albastaki, Marius Hoggenmüller, Frederic Anthony Robinson, and Luke Hespanhol. 2021. Augmenting Remote Interviews through Virtual Experience Prototypes. In *Proceedings of the 32nd Australian Conference on Human-Computer Interaction* (Sydney, NSW, Australia) (*OzCHI '20*). Association for Computing Machinery, New York, NY, USA, 78–86. <https://doi.org/10.1145/3441000.3441057>
- [2] Patrícia Alves-Oliveira, Maria Luce Lupetti, Michal Luria, Diana Löffler, Mafalda Gamboa, Lea Albaugh, Waki Kamino, Anastasia K. Ostrowski, David Puljiz, Pedro Reynolds-Cuellar, Marcus Scheunemann, Michael Suguitan, and Dan Lockton. 2021. Collection of Metaphors for Human-Robot Interaction. In *Designing Interactive Systems Conference 2021* (Virtual Event, USA) (*DIS '21*). Association for Computing Machinery, New York, NY, USA, 1366–1379. <https://doi.org/10.1145/3461778.3462060>
- [3] James Auger. 2014. Living with Robots: A Speculative Design Approach. *J. Hum.-Robot Interact.* 3, 1 (Feb. 2014), 20–42. <https://doi.org/10.5898/JHRI.3.1.Auger>
- [4] Shaowen Bardzell. 2010. Feminist HCI: Taking Stock and Outlining an Agenda for Design. In *Proceedings of the SIGCHI Conference on Human Factors in Computing Systems* (Atlanta, Georgia, USA) (*CHI '10*). Association for Computing Machinery, New York, NY, USA, 1301–1310. <https://doi.org/10.1145/1753326.1753521>
- [5] Ned Barker and Carey Jewitt. 2022. Future touch in industry: Exploring socio-technical imaginaries of tactile (tele)robots. *Futures* 136 (2022), 102885. <https://doi.org/10.1016/j.futures.2021.102885>
- [6] Christoph Bartneck, Tony Belpaeme, Friederike Eyssele, Takayuki Kanda, Merel Keijsers, and S. Sabanovic. 2020. *Human-Robot Interaction: An Introduction*. <https://doi.org/10.1017/9781108676649>
- [7] Jessica Bielski, Burak Mete, Viktor Eisenstadt, Christoph Langenhan, Frank Petzold, Andreas Dengel, and Klaus-Dieter Althoff. 2022. Chasing the White Rabbit—A case study of predicting design phases of architects by training a deep neural network with sketch recognition through a digital drawing board. *ICCC, [Preprint], Pinhal de Marrocos: Association for Computational Creativity* (2022).
- [8] Andreas Bischof, Eva Hornecker, Antonia Lina Krummheuer, and Matthias Rehm. 2022. Re-Configuring Human-Robot Interaction. In *Proceedings of the 2022 ACM/IEEE International Conference on Human-Robot Interaction* (Sapporo, Hokkaido, Japan) (*HRI '22*). IEEE Press, 1234–1236.
- [9] Laurens Boer and Harvey Bewley. 2018. Reconfiguring the Appearance and Expression of Social Robots by Acknowledging Their Otherness. In *Proceedings of the 2018 Designing Interactive Systems Conference* (Hong Kong, China) (*DIS '18*). Association for Computing Machinery, New York, NY, USA, 667–677. <https://doi.org/10.1145/3196709.3196743>
- [10] Oliver Bown, Kazjon Grace, Liam Bray, and Dan Ventura. 2020. A Speculative Exploration of the Role of Dialogue in Human-ComputerCo-creation.. In *ICCC*. 25–32.
- [11] G. Cascini, Y. Nagai, G. V. Georgiev, J. Zelaya, N. Becattini, J. F. Boujut, H. Casakin, N. Crilly, E. Dekoninck, J. Gero, A. Goel, G. Goldschmidt, M. Gonçalves, K. Grace, L. Hay, P. Le Masson, M. L. Maher, D. Marjanović, D. Motte, P. Papalambros, R. Sosa, Srinivasan V, M. Štorga, B. Tversky, B. Yan-nou, and A. Wodehouse. 2022. Perspectives on design creativity and innovation research: 10 years later. *International Journal of Design Creativity and Innovation* 10, 1 (2022), 1–30. <https://doi.org/10.1080/21650349.2022.2021480> arXiv:https://doi.org/10.1080/21650349.2022.2021480
- [12] Stephen Cave and Kanta Dihal. 2020. The Whiteness of AI. *Philosophy & Technology* 33, 4 (01 Dec 2020), 685–703. <https://doi.org/10.1007/s13347-020-00415-6>
- [13] Alan Chamberlain, Mads Bødker, and Konstantinos Papangelis. 2017. Mapping Media and Meaning: Autoethnography as an Approach to Designing Personal Heritage Soundscapes. In *Proceedings of the 12th International Audio Mostly Conference on Augmented and Participatory Sound and Music Experiences* (London, United Kingdom) (*AM '17*). Association for Computing Machinery, New York, NY, USA, Article 32, 4 pages. <https://doi.org/10.1145/3123514.3123536>
- [14] EunJeong Cheon and Norman Makoto Su. 2018. Futuristic Autobiographies: Weaving Participant Narratives to Elicit Values around Robots. In *Proceedings of the 2018 ACM/IEEE International Conference on Human-Robot Interaction* (Chicago, IL, USA) (*HRI '18*). Association for Computing Machinery, New York, NY, USA, 388–397. <https://doi.org/10.1145/3171221.3171244>
- [15] Wei-Chi Chien and Marc Hassenzahl. 2020. Technology-Mediated Relationship Maintenance in Romantic Long-Distance Relationships: An Autoethnographical Research through Design. *Human-Computer Interaction* 35, 3 (2020), 240–287. <https://doi.org/10.1080/07370024.2017.1401927> arXiv:https://doi.org/10.1080/07370024.2017.1401927
- [16] Nazli Cila, Cristina Zaga, and Maria Luce Lupetti. 2021. Learning from Robotic Artefacts: A Quest for Strong Concepts in Human-Robot Interaction. In *Designing Interactive Systems Conference 2021* (Virtual Event, USA) (*DIS '21*). Association for Computing Machinery, New York, NY, USA, 1356–1365. <https://doi.org/10.1145/3461778.3462095>
- [17] Florinel-Alin Croitoru, Vlad Hondru, Radu Tudor Ionescu, and Mubarak Shah. 2022. Diffusion models in vision: A survey. *arXiv preprint arXiv:2209.04747* (2022).
- [18] Nigel Cross. 1982. Designerly ways of knowing. *Design Studies* 3, 4 (1982), 221–227. [https://doi.org/10.1016/0142-694X\(82\)90040-0](https://doi.org/10.1016/0142-694X(82)90040-0) Special Issue Design Education.
- [19] Kerstin Dautenhahn. 2018. Some Brief Thoughts on the Past and Future of Human-Robot Interaction. *J. Hum.-Robot Interact.* 7, 1, Article 4 (may 2018), 3 pages. <https://doi.org/10.1145/3209769>
- [20] Laura Della Sciuca, Emanuele Balloni, Marco Mameli, Emanuele Frontoni, Primo Zingaretti, and Marina Paolanti. 2022. StyleTrendGAN: A Deep Learning Generative Framework for Fashion Bag Generation. In *International Conference on Image Analysis and Processing*. Springer, 191–202.
- [21] Skyla Y. Dudek and James E. Young. 2022. Fluid Sex Robots: Looking to the 2LGBTQIA+ Community to Shape the Future of Sex Robots. In *Proceedings of the 2022 ACM/IEEE International Conference on Human-Robot Interaction* (Sapporo, Hokkaido, Japan) (*HRI '22*). IEEE Press, 746–749.
- [22] Anthony Dunne and Fiona Raby. 2007. *Technological Dreams Series: No.1, Robots*. <http://dunneandraby.co.uk/content/projects/10/0>, last accessed: June 2020.
- [23] Ylva Fernaeus, Sara Ljungblad, Mattias Jacobsson, and Alex Taylor. 2009. Where Third Wave HCI Meets HRI: Report from a Workshop on User-Centred Design of Robots. In *Proceedings of the 4th ACM/IEEE International Conference on Human-Robot Interaction* (La Jolla, California, USA) (*HRI '09*). Association for Computing Machinery, New York, NY, USA, 293–294. <https://doi.org/10.1145/1514095.1514182>
- [24] Fabio Antonio Figoli, Lucia Rampino, and Francesca Mattioli. 2022. AI in design idea development: A workshop on creativity and human-AI collaboration. In *DRS2022* (Bilbao, Spain), D. Lockton, S. Lenzi, P. Hekkert, A. Oak, J. Sádaba, and P. Lloyd (Eds.). <https://doi.org/10.21606/drs.2022.414>
- [25] Mary Ellen Foster, Manuel Giuliani, and Ron Petrick. 2011–2018. *PubRob.org - The Workshop & Tutorial Series on Public Space Human-Robot Interaction*. <https://pubrob.org>, last accessed: January 2020.
- [26] Mikhaïla Friske, Jordan Wirfs-Brock, and Laura Devendorf. 2020. Entangling the Roles of Maker and Interpreter in Interpersonal Data Narratives: Explorations in Yarn and Sound. In *Proceedings of the 2020 ACM Designing Interactive Systems Conference* (Eindhoven, Netherlands) (*DIS '20*). Association for Computing Machinery, New York, NY, USA, 297–310. <https://doi.org/10.1145/3357236.3395442>
- [27] Rinon Gal, Yuval Alaluf, Yuval Atzmon, Or Patashnik, Amit H Bermano, Gal Chechik, and Daniel Cohen-Or. 2022. An image is worth one word: Personalizing text-to-image generation using textual inversion. *arXiv preprint arXiv:2208.01618* (2022).
- [28] Maliheh Ghajargar, Jeffrey Bardzell, and Love Lagerkvist. 2022. A Redhead Walks into a Bar: Experiences of Writing Fiction with Artificial Intelligence. In *25th International Academic Mindtrek Conference* (Tampere, Finland) (*Academic Mindtrek 2022*). Association for Computing Machinery, New York, NY, USA, 230–241. <https://doi.org/10.1145/3569219.3569418>
- [29] Elisa Giaccardi and Johan Redström. 2020. Technology and More-Than-Human Design. *Design Issues* 36, 4 (2020), 33–44. https://doi.org/10.1162/desi_a_00612
- [30] Phillip Gough, A. Baki Kocaballi, Khushnood Z. Naqshbandi, Karen Cochrane, Kristina Mah, Ajit G. Pillai, Yeliz Yorulmaz, Ainnoun Kornita Deny, and Naseem Ahmadpour. 2022. Co-Designing a Technology Probe with Experienced Designers. In *Proceedings of the 33rd Australian Conference on Human-Computer Interaction* (Melbourne, VIC, Australia) (*OzCHI '21*). Association for Computing Machinery, New York, NY, USA, 1–13. <https://doi.org/10.1145/3520495.3520513>
- [31] Kazjon Grace, Mary Lou Maher, Douglas Fisher, and Katherine Brady. 2015. Data-intensive evaluation of design creativity using novelty, value, and surprise. *International Journal of Design Creativity and Innovation* 3, 3-4 (2015), 125–147. <https://doi.org/10.1080/21650349.2014.943295> arXiv:https://doi.org/10.1080/21650349.2014.943295
- [32] Marc Hassenzahl, Jan Borchers, Susanne Boll, Astrid Rosenthal-von der Pütten, and Volker Wulf. 2020. Otherware: How to Best Interact with Autonomous Systems. *Interactions* 28, 1 (dec 2020), 54–57. <https://doi.org/10.1145/3436942>
- [33] Marius Hoggenmüller, Luke Hespanhol, and Martin Tomitsch. 2020. Stop and Smell the Chalk Flowers: A Robotic Probe for Investigating Urban Interaction with Physicalised Displays. 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, 1–14. <https://doi.org/10.1145/3313831.3376676>
- [34] Marius Hoggenmüller, Wen-Ying Lee, Luke Hespanhol, Malte Jung, and Martin Tomitsch. 2021. Eliciting New Perspectives in Rtd Studies through Annotated Portfolios: A Case Study of Robotic Artefacts. In *Designing Interactive Systems Conference 2021* (Virtual Event, USA) (*DIS '21*). Association for Computing Machinery, New York, NY, USA, 1875–1886. <https://doi.org/10.1145/3461778.3462134>
- [35] Rowan T Hughes, Liming Zhu, and Tomasz Bednarz. 2021. Generative Adversarial Networks—Enabled Human—Artificial Intelligence Collaborative Applications for Creative and Design Industries: A Systematic Review of Current Approaches and Trends. *Frontiers in artificial intelligence* 4 (2021), 604234.
- [36] Rowan T. Hughes, Liming Zhu, and Tomasz Bednarz. 2021. Generative Adversarial Networks—Enabled Human—Artificial Intelligence Collaborative Applications for Creative and Design Industries: A Systematic Review of Current Approaches and Trends. *Frontiers in Artificial Intelligence* 4 (2021). <https://doi.org/10.3389/frai.2021.604234>

- [37] Angel Hsing-Chi Hwang. 2022. Too Late to Be Creative? AI-Empowered Tools in Creative Processes. 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 38, 9 pages. <https://doi.org/10.1145/3491101.3503549>
- [38] Francisco Ibarrola, Tomas Lawton, and Kazjon Grace. 2022. A Collaborative, Interactive and Context-Aware Drawing Agent for Co-Creative Design. <https://doi.org/10.48550/ARXIV.2209.12588>
- [39] Daisuke Iizawa and Shunji Yamanaka. 2022. Face on a Globe: A Spherical Robot That Appears Lifelike Through Smooth Deformations and Autonomous Movement. In *Proceedings of the 2022 ACM/IEEE International Conference on Human-Robot Interaction* (Sapporo, Hokkaido, Japan) (HRI '22). IEEE Press, 502–510.
- [40] David G. Jansson and Steven M. Smith. 1991. Design fixation. *Design Studies* 12, 1 (1991), 3–11. [https://doi.org/10.1016/0142-694X\(91\)90003-F](https://doi.org/10.1016/0142-694X(91)90003-F)
- [41] S. Jasanoff and S.H. Kim. 2015. *Dreamscapes of Modernity: Sociotechnical Imaginaries and the Fabrication of Power*. University of Chicago Press. <https://books.google.com.au/books?id=CbbdrQEACAAJ>
- [42] Carey Jewitt, Sara Price, Kerstin Leder Mackley, Nikoleta Yiannoutsou, and Douglas Atkinson. 2020. *Sociotechnical Imaginaries of Digital Touch*. Springer International Publishing, Cham, 89–106. https://doi.org/10.1007/978-3-030-24564-1_6
- [43] Malte Jung and Pamela Hinds. 2018. Robots in the Wild: A Time for More Robust Theories of Human-Robot Interaction. *J. Hum.-Robot Interact.* 7, 1, Article 2 (may 2018), 5 pages. <https://doi.org/10.1145/3208975>
- [44] Vivian Liu and Lydia B Chilton. 2022. Design Guidelines for Prompt Engineering Text-to-Image Generative Models. 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 384, 23 pages. <https://doi.org/10.1145/3491102.3501825>
- [45] Vivian Liu, Han Qiao, and Lydia Chilton. 2022. Opal: Multimodal Image Generation for News Illustration. In *Proceedings of the 35th Annual ACM Symposium on User Interface Software and Technology* (Bend, OR, USA) (UIST '22). Association for Computing Machinery, New York, NY, USA, Article 73, 17 pages. <https://doi.org/10.1145/3526113.3545621>
- [46] Diana Löffler, Judith Dörrenbächer, Julika Welge, and Marc Hassenzahl. 2020. Hybridity as Design Strategy for Service Robots to Become Domestic Products. In *Extended Abstracts of the 2020 CHI Conference on Human Factors in Computing Systems* (Honolulu, HI, USA) (CHI EA '20). Association for Computing Machinery, New York, NY, USA, 1–8. <https://doi.org/10.1145/3334480.3382832>
- [47] Junlan Lu and Magdalini Eiriniaki. 2021. Can a machine win a Grammy? An evaluation of AI-generated song lyrics. In *2021 IEEE International Conference on Big Data (Big Data)*. IEEE, 4896–4905.
- [48] Andrés Lucero. 2015. Using Affinity Diagrams to Evaluate Interactive Prototypes. In *Human-Computer Interaction – INTERACT 2015*, Julio Abascal, Simone Barbosa, Mirko Fetter, Tom Gross, Philippe Palanque, and Marco Winckler (Eds.). Springer International Publishing, Cham, 231–248.
- [49] Andrés Lucero, Audrey Desjardins, Carman Neustaedter, Kristina Höök, Marc Hassenzahl, and Marta E. Cecchinato. 2019. A Sample of One: First-Person Research Methods in HCI. In *Companion Publication of the 2019 on Designing Interactive Systems Conference 2019 Companion* (San Diego, CA, USA) (DIS '19 Companion). Association for Computing Machinery, New York, NY, USA, 385–388. <https://doi.org/10.1145/3301019.3319996>
- [50] Maria Luce Lupetti, Cristina Zaga, and Nazli Cila. 2021. Designerly Ways of Knowing in HRI: Broadening the Scope of Design-Oriented HRI Through the Concept of Intermediate-Level Knowledge. 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, 389–398. <https://doi.org/10.1145/3434073.3444668>
- [51] Giorgia Lupi and Stefanie Posavec. 2016. *Dear Data*. Princeton Architectural Press, Hudson, NY, United States.
- [52] Michal Luria, Judeth Oden Choi, Rachel Gita Karp, John Zimmerman, and Jodi Forlizzi. 2020. *Robotic Futures: Learning about Personally-Owned Agents through Performance*. Association for Computing Machinery, New York, NY, USA, 165–177. <https://doi.org/10.1145/3357236.3395488>
- [53] Michal Luria, Juliet Pusateri, Judeth Oden Choi, Reuben Aronson, Nur Yildirim, and Molly Wright Steenson. 2020. Medieval Robots: The Role of Historical Automata in the Design of Future Robots. In *Companion Publication of the 2020 ACM Designing Interactive Systems Conference* (Eindhoven, Netherlands) (DIS '20 Companion). Association for Computing Machinery, New York, NY, USA, 191–195. <https://doi.org/10.1145/3393914.3395890>
- [54] Kristina Mah, Lian Loke, and Luke Hespanhol. 2021. Understanding Compassion Cultivation for Design: Towards an Autoethnography of Tonglen. In *Proceedings of the 32nd Australian Conference on Human-Computer Interaction* (Sydney, NSW, Australia) (OzCHI '20). Association for Computing Machinery, New York, NY, USA, 748–754. <https://doi.org/10.1145/3441000.3441065>
- [55] Diana P. Moreno, Lucienne T. Blessing, Maria C. Yang, Alberto A. Hernández, and Kristin L. Wood. 2016. Overcoming design fixation: Design by analogy studies and nonintuitive findings. *Artificial Intelligence for Engineering Design, Analysis and Manufacturing* 30, 2 (2016), 185–199. <https://doi.org/10.1017/S0890060416000068>
- [56] Michael Muller, Lydia B Chilton, Anna Kantosalo, Charles Patrick Martin, and Greg Walsh. 2022. GenAICHI: Generative AI and HCI. 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 110, 7 pages. <https://doi.org/10.1145/3491101.3503719>
- [57] H Na and W Kim. 2022. An Empirical Study on the Use of Generative Design on Product Design. *Archives of* (2022).
- [58] Ben Poole, Ajay Jain, Jonathan T Barron, and Ben Mildenhall. 2022. Dreamfusion: Text-to-3d using 2d diffusion. *arXiv preprint arXiv:2209.14988* (2022).
- [59] Aditya Ramesh, Prafulla Dhariwal, Alex Nichol, Casey Chu, and Mark Chen. 2022. Hierarchical Text-Conditional Image Generation with CLIP Latents. <https://doi.org/10.48550/ARXIV.2204.06125>
- [60] Aditya Ramesh, Prafulla Dhariwal, Alex Nichol, Casey Chu, and Mark Chen. 2022. Hierarchical Text-Conditional Image Generation with CLIP Latents. <https://doi.org/10.48550/ARXIV.2204.06125>
- [61] Lyle Regenwetter, Amin Heyrani Nobari, and Faez Ahmed. 2022. Deep generative models in engineering design: A review. *Journal of Mechanical Design* 144, 7 (2022), 071704.
- [62] Robin Rombach, Andreas Blattmann, Dominik Lorenz, Patrick Esser, and Björn Ommer. 2021. High-Resolution Image Synthesis with Latent Diffusion Models. *CoRR* abs/2112.10752 (2021). [arXiv:2112.10752](https://arxiv.org/abs/2112.10752) <https://arxiv.org/abs/2112.10752>
- [63] Nataniel Ruiz, Yuanzhen Li, Varun Jampani, Yael Pritch, Michael Rubinstein, and Kfir Aberman. 2022. Dreambooth: Fine tuning text-to-image diffusion models for subject-driven generation. *arXiv preprint arXiv:2208.12242* (2022).
- [64] Selma Sabanovic. 2010. Robots in Society, Society in Robots. *International Journal of Social Robotics* 2 (2010), 439–450.
- [65] S. Sabanovic, M.P. Michalowski, and R. Simmons. 2006. Robots in the wild: observing human-robot social interaction outside the lab. In *9th IEEE International Workshop on Advanced Motion Control*, 2006. 596–601. <https://doi.org/10.1109/AMC.2006.1631758>
- [66] Christoph Schuhmann, Romain Beaumont, Richard Vencu, Cade Gordon, Ross Wightman, Mehdi Cherti, Theo Coombes, Aarush Katta, Clayton Mullis, Mitchell Wortsman, et al. 2022. Laion-5b: An open large-scale dataset for training next generation image-text models. *arXiv preprint arXiv:2210.08402* (2022).
- [67] Lucas Nunes Sequeira, Bruno Moreschi, and Vinicius Ariel Arruda dos Santos. 2021. Which Human Faces Can an AI Generate? Lack of Diversity in This Person Does Not Exist. *AoIR Selected Papers of Internet Research* 2021 (Sep. 2021). <https://doi.org/10.5210/spir.v2021i0.12240>
- [68] Sofia Serholt, Sara Ljungblad, and Niamh Ni Bhroin. 2021. Introduction: special issue—critical robotics research. *AI & SOCIETY* (06 2021). <https://doi.org/10.1007/s00146-021-01224-x>
- [69] Luca Simeone, Riccardo Mantelli, and Alfredo Adamo. 2022. Pushing divergence and promoting convergence in a speculative design process: Considerations on the role of AI as a co-creation partner. In *DRS2022* (Bilbao, Spain), D. Lockton, S. Lenzi, P. Hekkert, A. Oak, J. Sádaba, and P. Lloyd (Eds.). <https://doi.org/10.21606/drs.2022.197>
- [70] Ut Na Sio, Kenneth Kotovsky, and Jonathan Cagan. 2015. Fixation or inspiration? A meta-analytic review of the role of examples on design processes. *Design Studies* 39 (2015), 70–99. <https://doi.org/10.1016/j.destud.2015.04.004>
- [71] Ramya Srinivasan and Kanji Uchino. 2021. Biases in Generative Art: A Causal Look from the Lens of Art History. In *Proceedings of the 2021 ACM Conference on Fairness, Accountability, and Transparency* (Virtual Event, Canada) (FAccT '21). Association for Computing Machinery, New York, NY, USA, 41–51. <https://doi.org/10.1145/3442188.3445869>
- [72] Ryan Steed and Aylin Caliskan. 2021. Image Representations Learned With Unsupervised Pre-Training Contain Human-like Biases. In *Proceedings of the 2021 ACM Conference on Fairness, Accountability, and Transparency* (Virtual Event, Canada) (FAccT '21). Association for Computing Machinery, New York, NY, USA, 701–713. <https://doi.org/10.1145/3442188.3445932>
- [73] Claire Stevenson, Iris Smal, Matthijs Baas, Raoul Grasman, and Han van der Maas. 2022. Putting GPT-3’s Creativity to the (Alternative Uses) Test. *arXiv preprint arXiv:2206.08932* (2022).
- [74] Calvin W. Taylor. 1988. *Various approaches to and definitions of creativity*. Cambridge University Press, New York, NY, US, 99–121.
- [75] Justin D. Weisz, Mary Lou Maher, Hendrik Strobel, Lydia B. Chilton, David Bau, and Werner Geyer. 2022. HAI-GEN 2022: 3rd Workshop on Human-AI Co-Creation with Generative Models. In *27th International Conference on Intelligent User Interfaces* (Helsinki, Finland) (IUI '22 Companion). Association for Computing Machinery, New York, NY, USA, 4–6. <https://doi.org/10.1145/3490100.3511166>
- [76] Justin D Weisz, Mary Lou Maher, Hendrik Strobel, Lydia B Chilton, David Bau, and Werner Geyer. 2022. HAI-GEN 2022: 3rd Workshop on Human-AI Co-Creation with Generative Models. In *27th International Conference on Intelligent User Interfaces*. 4–6.
- [77] Chung-Yi Weng, Brian Curless, and Ira Kemelmacher-Shlizerman. 2019. Photo Wake-Up: 3D Character Animation From a Single Photo. In *2019 IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*. 5901–5910. <https://doi.org/10.1109/CVPR.2019.00606>

- [78] Katie Winkle, Gaspar Isaac Melsión, Donald McMillan, and Iolanda Leite. 2021. Boosting Robot Credibility and Challenging Gender Norms in Responding to Abusive Behaviour: A Case for Feminist Robots. 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, 29–37. <https://doi.org/10.1145/3434074.3446910>
- [79] Haiyan Xue and Pieter M.A. Desmet. 2019. Researcher introspection for experience-driven design research. *Design Studies* 63 (2019), 37–64. <https://doi.org/10.1016/j.destud.2019.03.001>
- [80] Ling Yang, Zhilong Zhang, Yang Song, Shenda Hong, Runsheng Xu, Yue Zhao, Yingxia Shao, Wentao Zhang, Bin Cui, and Ming-Hsuan Yang. 2022. Diffusion models: A comprehensive survey of methods and applications. *arXiv preprint arXiv:2209.00796* (2022).
- [81] J.D. Zamfirescu-Pereira, David Sirkin, David Goedicke, Ray LC, Natalie Friedman, Ilan Mandel, Nikolas Martelaro, and Wendy Ju. 2021. Fake It to Make It: Exploratory Prototyping in HRI. 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, 19–28. <https://doi.org/10.1145/3434074.3446909>