



How do people experience the images created by generative artificial intelligence? An exploration of people's perceptions, appraisals, and emotions related to a Gen-AI text-to-image model and its creations

Amon Rapp^{a,*}, Chiara Di Lodovico^{a,b}, Federico Torrielli^a, Luigi Di Caro^a

^a Computer Science Department, University of Turin, Torino, Italy

^b Design Department, Politecnico di Milano, Milano, Italy

ARTICLE INFO

Keywords:

AI
Generative AI
Stable diffusion
User experience
Anthropomorphising
Humanness
Uncanny valley

ABSTRACT

Generative Artificial Intelligence (Gen-AI) has rapidly advanced in recent years, potentially producing enormous impacts on industries, societies, and individuals in the near future. In particular, Gen-AI text-to-image models allow people to easily create high-quality images possibly revolutionizing human creative practices. Despite their increasing use, however, the broader population's perceptions and understandings of Gen-AI-generated images remain understudied in the Human-Computer Interaction (HCI) community. This study investigates how individuals, including those unfamiliar with Gen-AI, perceive Gen-AI text-to-image (Stable Diffusion) outputs. Study findings reveal that participants appraise Gen-AI images based on their technical quality and fidelity in representing a subject, often experiencing them as either prototypical or strange: these experiences may raise awareness of societal biases and evoke unsettling feelings that extend to the Gen-AI itself. The study also uncovers several "relational" strategies that participants employ to cope with concerns related to Gen-AI, contributing to the understanding of reactions to uncanny technology and the (de)humanization of intelligent agents. Moreover, the study offers design suggestions on how to use the anthropomorphizing of the text-to-image model as design material, and the Gen-AI images as support for critical design sessions.

1. Introduction

Generative Artificial Intelligence (Gen-AI) has witnessed rapid and unexpected advancements in recent years, potentially producing enormous impacts on industries, societies, and individuals in the near future. Gen-AI refers to an array of technologies that can emulate the creation of human-like outputs, like images, text, and music, in response to written textual inputs, also known as "prompts" (Muller et al., 2022; Bandi et al., 2023).

In particular, Gen-AI text-to-image models are technologies that produce high-quality images, using deep neural networks trained on huge multimodal datasets (Struppek et al., 2023; Vimpari et al., 2023; Oppenlaender, 2022; Tang et al., 2023; Ramesh et al., 2021). The recent introduction of easy-to-use models, like DALL-E 2,^c Midjourney,^d and Stable Diffusion,^e allows everyone to quickly produce unique images

and artworks irrespective of their technical knowledge, becoming creators of digital products of an unimaginable quality until a few years ago (Wang et al., 2023; Liu and Chilton, 2022; Vimpari et al., 2023).

These technologies have attracted millions of users within a short period of time, and currently, AI-generated images are widespread everywhere: they are shared on social networks (Oppenlaender, 2022), appear on magazine covers (Liu, 2022) and news articles (Huang, 2023), and achieve recognition in art competitions (Roose, 2022). Nonetheless, despite Gen-AI images are now possibly reaching anyone, it is still not clear what kind of effects they have on their potential audience: are these images perceived by people in the same way as the images produced by humans are perceived? Or do they elicit, for instance, peculiar emotions? Does their artificial nature evoke reflections and feelings with respect to the generative AI that created them?

As more and more people now use this technology to produce images

* Corresponding author: University of Turin, Computer Science Department, C.so Svizzera 185, 10149 Torino, Italy.

E-mail address: amon.rapp@unito.it (A. Rapp).

^c <https://openai.com/dall-e-2/>

^d <https://www.midjourney.com/home/>

^e <https://stablediffusionweb.com/>

that are considered optimal substitutes for human creations, and as these images appear, at first sight, extremely similar to those created by human beings, understanding how they are perceived by the broader population becomes essential.

For this, the Human-Computer Interaction (HCI) community is showing growing interest in the study of Gen-AI models (e.g., [Inie et al., 2023](#); [Wang et al., 2023](#)). On the one hand, scholars have focused on the practices and motivations of users of the latest generation of text-to-image models ([Chang et al., 2023](#); [Oppenlaender, 2022](#); [Liu et al., 2023](#); [Kulkarni et al., 2023](#)). On the other hand, researchers have investigated people's understanding of these technologies, exploring their expectations, aesthetic appreciation, and experience ([Ting et al., 2023](#); [Oppenlaender et al., 2023](#); [Sun et al., 2022](#)).

In both cases, however, the emphasis of research has been put on the artistic and creative processes ([Chang et al., 2023](#); [Liu et al., 2023](#)): in particular, those exploring people's experience mainly focus on the "artistic component" of the AI-generated images, delving into how individuals differently judge human and Gen-AI artworks ([Ting et al., 2023](#); [Sun et al., 2022](#)).

In this context, people's perceptions of Gen-AI images outside the artistic context remain still underexplored. As Gen-AI images are now employed not only for artistic purposes, but also as substitutes for human-generated images in a myriad of domains, like advertising, design, and photo editing for personal use, investigating how they are perceived by people beyond their artistic quality is becoming more pressing. Furthermore, no previous research has explored whether the images produced by Gen-AI may elicit reflections and feelings about Gen-AI itself, beyond mere artistic matters (i.e., whether the AI can be considered an artist or not). This may be important because it may unveil visceral reactions to and concerns about the Gen-AI triggered by the images that it produces, contributing to shaping the users' underlying perception, understanding, and acceptance of the technology.

With this aim, we investigated how different kinds of individuals, ranging from technologists and designers to people without previous experience of AI technologies, perceive the products of a Gen-AI text-to-image model. Our goal was to involve diverse participants, beyond those directly using Gen-AI tools. This approach mirrors current real-world scenarios where individuals are more likely to encounter AI-generated content rather than actively interacting with a Gen-AI system.

Specifically, this research aims to answer the following questions:

- (RQ1) How do people experience (in terms of e.g., perceptions, appraisals, and emotions) the images produced by a Gen-AI model?
- (RQ2) What kinds of reflections and feelings do Gen-AI images elicit about Gen-AI itself?

To answer these questions, we conducted a qualitative study using Gen-AI images as stimuli and semi-structured interviews with 20 participants. Each participant was presented with 20 different images produced by a Gen-AI and their corresponding "heat maps", a representation that shows what elements of the image were more important for the Gen-AI in generating the image starting from specific terms. Study findings reveal that participants appraised the images mainly on the basis of their technical quality and capability of clearly representing a given subject, while overlooked their creative and artistic components. Moreover, they perceived the images either as "strange" or "prototypical", often evoking unsettling feelings. When participants could not explain the reasons why a certain subject was represented in a certain way, they adopted several strategies to mitigate the sense of bewilderment they felt, devaluing or overvaluing the Gen-AI, or reworking their initial perceptions. Finally, participants also found value in the Gen-AI images, seeing them as opportunities for speeding up the image creation process, or for reflecting on the biases that underlie our society.

To summarize, the contributions of this article are twofold.

First, we show that participants perceive the images not so much in terms of their creativity but in relation to their technical quality and

subject fidelity, experiencing them as either prototypical or strange: this differs from previous research focused on understanding how people perceive Gen-AI's creativity in comparison to human creativity (e.g., [Samo and Highhouse, 2023](#); [Lyu et al., 2022](#)), and highlights that Gen-AI images may evoke unsettling emotions for their unfamiliarity, or make people aware of the biases that lie behind both the Gen-AI and our society.

Second, we uncover the reflections and feelings that Gen-AI images evoke about the Gen-AI itself, which go beyond the concerns about the disruption of creative job market and increase of misinformation discovered by previous research ([Oppenlaender et al., 2023](#); [Bird et al., 2023](#)). Rather, participants' reactions mostly revolve around sensations of uncanniness toward the Gen-AI, which point at its alien nature. To mitigate these unsettling sensations, we discovered that participants perform a variety of "relational strategies", which range from self-devaluation to overvaluation of the Gen-AI. This contributes to research on how people react to "uncanny" technology (e.g., [Mori, 1970](#)), and how they (super)humanize (by overvaluing it) or dehumanize (by devaluing it) an intelligent agent (e.g., [Go and Sundar, 2019](#)), in order to tame its unfamiliarity.

The article is structured as follows. [Section 2](#) provides an overview of previous HCI research related to text-to-image models. [Section 3](#) describes the method we used, and [Section 4](#) presents the findings of the study. Study findings are discussed in [Section 5](#), while two design implications are presented in [Section 6](#). Limitations are acknowledged in [Section 7](#), while [Section 8](#) concludes the article.

2. Background

2.1. Generative AI and text-to-image models

Gen-AI text-to-image models employ deep neural networks trained on huge multimodal datasets to produce high-quality synthetic images from natural language textual inputs (also called "prompts") ([Struppek et al., 2023](#); [Vimpari et al., 2023](#); [Oppenlaender, 2022](#); [Tang et al., 2023](#); [Ramesh et al., 2021](#)). These models have the capability of generating detailed images in a vast array of styles based on the user's prompt, presenting a significant opportunity for application to diverse creative visual tasks ([Kulkarni et al., 2023](#); [Croitoru et al., 2023](#); [Liu and Chilton, 2022](#); [Ho et al., 2020](#)). In particular, the introduction of Stable Diffusion in 2022 ([Rombach et al., 2022](#)) led to the creation of accurate and high-resolution images, establishing the state-of-the-art in text-to-image Gen-AI ([Vimpari et al., 2023](#)). Shortly thereafter, other state-of-the-art systems such as Midjourney and DALL-E 2 were released, widening the popularity of this technology ([Vimpari et al., 2023](#)).

In fact, these systems introduced easy-to-use Gen-AI interfaces, allowing anyone to quickly create digital images and artworks regardless of their design skills and technical knowledge ([Wang et al., 2023](#); [Liu and Chilton, 2022](#); [Vimpari et al., 2023](#)). In this sense, the ease of interaction has certainly favored their widespread adoption, not only by artists and designers, but also by the broader population. In fact, by September 2022, DALL-E 2 had 1.5 million active users ([OpenAI, 2022](#)), while Stable Diffusion had >10 million users across all channels ([Carlson, 2022](#)); by October 2023, Midjourney had an average of 1.5 million daily active users ([Codewatchers, 2023](#)).

These data hint at the fact that these technologies are seeping into many aspects of people's daily lives, whereby an increasingly wide spectrum of individuals not only create digital representations through automated means, but are also exposed to them ([Bird et al., 2023](#)). AI-created images are now employed in a variety of domains and across multiple communication channels, such as social networks ([Oppenlaender, 2022](#)), magazines ([Liu, 2022](#)), and articles ([Huang, 2023](#)), and are allegedly believed to reliably substitute human-generated images ([Bird, 2023](#)).

Despite the popularity of this technology, however, its impact on the broader population is still not clear and we do not have a clear

understanding of how people perceive the images created by Gen-AI. In this sense, the HCI community has attempted to uncover how people use Gen-AI technology for generating images and perceive its productions.

2.2. HCI studies on Gen-AI and text-to-image models

Given the social significance of Gen-AI phenomenon, there is a notable increase in interest in text-to-image Generative AI within the HCI research community. HCI research on this technology may be differentiated along two main lines: i) the investigation of the usage of text-to-image models; ii) the exploration of how these models and their products are experienced.

2.2.1. Using text-to-image models

A first line of studies builds on previous research on AI-based creativity support tools and human-AI co-creation (Chang et al., 2023; Louie et al., 2020; Davis et al., 2016; Hodhod and Magerko, 2016). These tools support the artistic performance by either doing the majority of the work of generating the art, or by powering the artistic execution (e.g., AI-powered brush tools in drawing applications) (Chung et al., 2021b). Likewise, research in the field of human-AI co-creation has long investigated human-AI mixed initiatives, identifying different opportunities for co-creating with AI (Muller et al., 2020; Grabe et al., 2022). This kind of research has also pointed out a variety of challenges in creating with AIs, like the human limited capability of controlling the AI output (Chung et al., 2021a) and the biases that may affect its suggestions (Buschek et al., 2021).

Informed by this previous body of knowledge, HCI research on the usage of text-to-image models points to the practices of Gen-AI users, as well as the challenges that they may encounter. Chang et al. (2023) examined the art practices, artwork, and motivations of prolific users of the latest generation of text-to-image models. They found that artists define prompt templates (prompts with “slots” for others to fill in with their own words) to create generative art styles. Likewise, Ko et al. (2023) interviewed visual artists and discovered that text-to-image models could help them think out of the box. Beyond the art domain, Inie et al. (2023) surveyed a variety of creative professionals, who highlighted that Gen-AI can enhance productivity, offer inspiration, and lead to higher quality output. Similarly, Kulkarni et al. (2023) found that text-to-image models can help non-professional designers explore a design space rapidly, and that prompts may facilitate exploration, iteration, and reflection in pair design. In the same line, Vimpari et al. (2023) reported that game professionals use the systems’ outputs as a source of inspiration and to prototype or conceptualize ideas.

By and large, all these studies focus on the usage of Gen-AI technology in the artistic and creative process performed by artists or visual professionals, paying mostly attention to creative matters, like the practice of “prompt engineering”, defined in this context as the creative practice of writing effective textual input prompts for creating artworks (Oppenlaender, 2022; Oppenlaender et al., 2023). However, not only Gen-AI tools are now more and more used beyond the artistic domain, but also, they are reaching the broader population even outside the artistic context, so that a wider exploration of how people perceive the products of Gen-AI technology is becoming pressing.

2.2.2. Perceiving text-to-image models and their products

The focus on creativity and the artistic process is also reflected in the research line investigating people’s perceptions of text-to-image models and their images. This kind of research falls within the realm of studies exploring how people perceive the artistic and creative quality of artifacts produced by robots and AIs, often comparing them with human creations.

For example, Chamberlain et al. (2018) investigated how people respond to works of visual art created either by humans or by computers: findings showed an aesthetic bias against computer generated art, whereby participants in their study tended to prefer artworks believed to

be generated by humans. A similar setting characterizes Hong and Curran’s (2019) study, which examines how people perceive artwork created by AI and how the presumed artist’s identity (Human vs. AI) affects individuals’ evaluation of art. Results indicate that human-created artworks have a significantly higher rating in “composition,” “degree of expression,” and “aesthetic value”. Similar findings are reported by Ragot et al. (2020) and Mikalonytė and Kneer (2022). The former asked people to evaluate paintings that were created by humans or Generative Adversarial Networks (GANs): the paintings perceived as being drawn by humans were evaluated significantly more highly than those perceived as being made by GANs. The latter, instead, studied whether people are willing to consider paintings made by robots as art, and robots as artists: they found that people judge robot paintings and human paintings as art to roughly the same extent but are much less willing to consider robots as artists than humans.

When coming to considering text-to-image models, research on users’ perceptions follows a similar approach: researchers ask people to assess the artistic quality of Gen-AI models’ creations, often in comparison with human artworks. Samo and Highhouse (2023), for example, recruited 190 participants inviting them to rate one image (either human- or Gen-AI- generated) on measures regarding their artistic quality: results point out that participants preferred human art and experienced more positive emotions in response to human artwork. Similarly, Lyu et al. (2022) involved forty-two subjects with artistic backgrounds requiring them to rate 12 text-to-image paintings, prompted by both artists and non-artists, as well as a human-created painting, on different aesthetics attributes. The results show that the assistance of Gen-AI made the perception of human-AI co-creation with and without artistic background converge, also blurring the difference between Gen-AI co-creations and human paintings.

If the perception of Gen-AI images’ artistic qualities is receiving increasing attention from the academic community, wider user perceptions and understandings of the products of this technology remain overlooked. A partial exception is represented by the study of Oppenlaender et al. (2023), who surveyed 35 participants about their experience of text-to-image generation technology. They discovered that participants believed that text-to-image models could lower the barriers to create images, but could also be used for creating false re-creations that can cause harm, increasing misinformation and fake news, as well as lead to job loss and unemployment. Similarly, Bird et al. (2023) identified a variety of risks coming from the usage of text-to-image models, from biases and loss of work for creatives, to privacy infringement and misinformation: however, such risks were identified in current literature and not by investigating people’s perceptions.

Despite these preliminary interesting insights, people’s experience of productions of text-to-image technologies remains substantially under-explored. In fact, Oppenlaender et al. (2023) did not investigate how people perceive the images produced by Gen-AI; moreover, they used a survey as an inspection method, an instrument that leaves less room to the exploration of topics that are considered relevant by the participants themselves in comparison with more open-ended techniques like semi-structured interviews.

Therefore, many potentially interesting themes revolving around users’ perceptions of Gen-AI images remain to be studied. First, it is still not clear how people perceive such images beyond their artistic features, as well as the emotions that they may evoke: this is fundamental because Gen-AI images are now employed as substitutes for human-generated images beyond the artistic domain, but we still do not know if they may elicit peculiar perceptions and feelings.

Second, it is still unclear whether the images produced by Gen-AI may trigger reflections about the nature of their “creator”. Previous research has only explored whether people may consider Gen-AIs artists or not (e.g., Mikalonytė and Kneer, 2022), overlooking the broader reactions that they may have towards Gen-AIs while being exposed to their productions. These reactions are important because they may reveal underlying feelings and concerns about the technology, which may

impact how people understand and accept the technology itself.

For this, in this study we will contribute to literature by focusing on people's experience of a series of images produced by a text-to-image model using semi-structured interviews, in order to unveil: i) how they perceive and emotionally experience such images; and ii) how they react to the Gen-AI that has generated them.

2.3. Framing people's perceptions of "intelligent" technologies

Before delving into the exposition of the method employed in this study, it is necessary to outline how previous research has framed the investigation of people's perceptions of seemingly intelligent agents capable of producing human-like behavior and outputs. In fact, this kind of research can be fruitfully used to interpret the findings coming from this study, providing a conceptual framework that may help to make sense of the participants' accounts.

Research on people's perceptions of intelligent technology has traditionally focused on two lines of investigation. The first one revolves around the unsettling feelings that apparently intelligent agents may evoke in people interacting with them. This research is mostly based on the theory of the uncanny valley, which precisely suggests that perceptual difficulty in discerning a human-like object will evoke unsettling emotions (Mori, 1970). The theory states that with increasing human likeness, an entity becomes more and more accepted by humans, but when the entity looks almost real, it falls into what it is called the uncanny valley (Schwind et al., 2018). This happens because the entity displays conflicting cues (e.g., details with high level of realism creating expectations about its humanlike nature, which are not fulfilled during the interaction), leading to a conflict in people's mental categorization (Stein and Ohler, 2017). This theory has been originally formulated with reference to robots, but then it was extended to other kinds of intelligent artifacts, like chatbots (Rapp et al., 2021): for instance, it has been found that a chatbot's expression of sympathy and empathy may be spine-tingling (Liu and Sundar, 2018), and that certain users may be unsettled by a chatbot's ability to sound like a real human (Ta et al., 2020). However, other research has noticed that an uncanny valley effect is more likely to appear when the chatbot is embodied in an avatar, suggesting that it is the aesthetics appearance of its virtual body to elicit disturbing emotions (Ciechanowski et al., 2018).

The second line of investigation points at the humanization of technology. At its core, humanization refers to humanness, namely to those attributes that define what it is to be human (Haslam and Loughnan, 2014), which is a central topic of interest in human-agent interaction. In fact, people may attribute humanness also to non-human entities, in a process that is commonly called humanization or anthropomorphism (Festerling and Siraj, 2022). Haslam (2006) has theorized that humanness may be defined in terms of "human uniqueness", which identifies uniquely human characteristics like rationality, and "human nature," referring to characteristics like emotional responsiveness, which are viewed as typical of humans, in a non-comparative sense, revealing our continuity with other creatures. This dual nature of humanness implies that people may humanize and dehumanize an entity in various ways depending on which human attributes are ascribed or denied. From this perspective, not only people can frame others as less than (but not necessarily non-)human (Li et al., 2014), treating dehumanized targets as entities lacking certain human characteristics (like competence) but not others (like friendliness), in a way that is called *infrahumanization* (Leyens et al., 2003). But also, people may ascribe only certain human-like characteristics to an entity, like intelligence, leading to a phenomenon called *superhumanization*, which may imply perceiving the other as unemotional, and yet highly intelligent, even transcending ordinary humanness (Li et al., 2014).

HCI research has highlighted that users may attribute humanlike features to technology (e.g., Go and Sundar, 2019; Doyle et al., 2019; Candello et al., 2017; Rapp et al., 2023b) and designers may intentionally reproduce fundamental aspects of being human in their design

(e.g., Lee et al., 2019; Law et al., 2022; Følstad et al., 2023). In the context of "intelligent" agents, Rapp et al. (2023a) emphasized that people may humanize or dehumanize a technology along different degrees of humanness: they may ascribe to it different separate human abilities, depending on the context and objectives emerging during the interaction, using humanization and dehumanization for utilitarian purposes (e.g., when the agent does not allow users to achieve their goals they dehumanize it degrading it to a lower entity).

These two lines of research can be fruitfully applied to interpret people's perceptions of Gen-AI text-to-image models and their creations. In fact, these models produce human-like outputs similarly to pre-Gen-AI intelligent agents; however, they are more efficient in creating products that apparently emulate those produced by human beings. It is thus reasonable to ask ourselves whether people will react to Gen-AI technology in ways similar to those elicited by technologies capable of exhibiting humanlike behaviors, like robots or conversational agents, either by being affected by their uncanniness or by humanizing or dehumanizing the technology itself.

3. Method

Originally, the study had two main objectives. First, we wanted to investigate how people react to the images that a Gen-AI text-to-image model can create, in terms of perceptions, appraisals, and emotions that they may evoke. Second, we aimed to study whether and how people are able to generate coherent "folk theories" of the inner working principles of these technologies. While the recounting of the findings relating to the second objective goes beyond the scope of this article and won't be reported here, in the following Sections we will focus on the first objective, describing people's experience of Gen-AI text-to-image technology and its creations.

We used semi-structured interviews and recruited twenty participants with diverse backgrounds, inviting them to look at twenty different images generated by Stable Diffusion starting from specific text prompts. In line with previous research (Tang et al., 2023), we provided a variety of stimuli (i.e., the images), as comparing different Gen-AI outcomes may better allow individuals to formulate inferences about the system's behavior and appraise its creations, considering variations in the outputs and their unique characteristics. This is also in line with photo elicitation as a qualitative method of interviewing, which is simply the use of images within an interview setting (Harper, 2002; Hogan, 2012), and it has been found to provide meaningful accounts (e.g., Kunimoto, 2004).

Moreover, we paired each image with its corresponding heat map: heat maps allow people to identify and weight which parts of the input are more influential into the generated output (Kim et al., 2023), potentially helping them discern what aspects are important for the AI in creating the image (Kim et al., 2023). A heatmap appears as a gradient-based visual overlay on an image wherein the intensity of each pixel's color conveys the degree of importance that the pixel had in representing the input elements. With these representations, we wanted to provide participants with "open cues" about the underlying mechanisms driving the AI's behavior. Heat maps could offer participants material for reflection about how the images were generated without encouraging any specific interpretation: in fact, heat maps usually reveal little about why the Gen-AI system actually behaved in specific ways, rather leaving the interpretation process open (Kim et al., 2023; Lee et al., 2023; Tang et al., 2023).

It is worth reminding that the choice of exploring people's perceptions of text-to-image Gen-AI's products, without allowing them to directly interact with the Gen-AI, was due to the consideration that AI-generated images are now widely disseminated everywhere, reaching individuals irrespective of their direct involvement in image generation processes. This phenomenon makes it interesting to explore how people perceive the Gen-AI and the images it produces regardless of their interaction with the technology.

The subsequent subsections provide comprehensive details on the study materials, participants, procedure, and data analysis.

3.1. Materials

To generate the image that we used in the study, we chose Stable Diffusion for its open-source nature and its capacity to produce both images and their corresponding heat maps, by integrating the DAAM module, as introduced by Tang et al. (2023). We employed a consistent prompt formula across all the productions, utilizing the textual input “Representation of ___” (e.g., “Representation of anger”). This choice was made considering that using only one term (e.g., “Sea”) would likely have produced details characterizing the subject (e.g., waves). By adding “Representation of ___,” we encouraged the model to explore various ways of representing the subject.

Furthermore, we took into account the configuration of the Classifier-Free Guidance (CFG) Scale (Ho and Salimans, 2021) to determine the fidelity of the generated images relative to the given prompt. This parameter enables the Gen-AI to attribute varying levels of image’s fidelity to the prompt (from 20, indicating extreme fidelity, to 1, where the Gen-AI generates more diverse images). We opted for a CFG value of 7.5 because it is the standard value used in text-to-image model research (Wang et al., 2023), granting a balance between fidelity to the prompt and a degree of “creativity” in the image generation process.

The selection of terms for generating the images aligned with practices observed in prior research (Liu and Chilton, 2022; Oppelaender, 2022) and mainly followed abstractness and concreteness criteria, to create images representing different subjects. We employed the concreteness rating database provided by Brysbaert et al. (2014), which contains 40,000 English words rated on a scale from abstract (1) to concrete (5). We first identified abstract terms with ratings between 1 and 2.5 and concrete terms with ratings ranging from 3.5 to 5. Then, we refined the selection considered several factors: i) the comprehensibility of the term, focusing on those that were likely universally understandable; ii) the potential for visual representation, ensuring that the generated image could be somehow linked to the original term.

In practice, the term selection process underwent multiple iterations. After extensive discussions among the four authors, we defined an initial pool of 200 terms. Then we engaged four external researchers asking them to report on their understandability and potential for visual representation. We then re-discussed the terms and eventually selected the 20 terms that best aligned with our criteria, generating the

corresponding images and heatmaps (Fig. 1).

We used Stable Diffusion 1.5 to generate the images, which does not offer predefined visual styles to be selected and is not fine-tuned on a specific style. If a style is not specified in the prompt, the “default style” depends on the full data set on which the model was trained. This was considered the best option because previous research has shown that using specific styles may introduce supplementary biases in the generation process (Zhang et al., 2023). The “default style” of Stable Diffusion 1.5 was also considered a reasonable middle ground between photographic and “artsy” images, both of which we wanted to avoid. Even though Stable Diffusion “default style” did not prevent the model from creating images with an artistic flair, purposefully and exclusively generating artsy images risked focusing participants solely on artistic aspects. Likewise, photography-like images could have shifted the participants’ attention exclusively to matters like misinformation or adherence to reality. Instead, we wanted to explore their broader reactions to these images.

3.2. Participants

The sample selection followed a purposeful sampling technique, which prescribes that researchers actively select the most productive sample to answer their research questions, identifying the variables that might influence a participant’s contribution on the basis of their practical knowledge of the research area and previous literature (Marshall, 1996). We thus differentiated the sample mainly along the dimension of background/profession and previous experience with Gen-AI. Previous research has shown that professions and backgrounds, as well as previous experience with Gen-AI, may shape the expectations about Gen-AI technology (Oppelaender et al., 2023).

We thus recruited 20 Italian participants through emails and snowball sampling. Participants included in the sample ranged from AI experts owning a technical background and designers with an artistic/creative background, to people with humanistic (e.g., a psychologist) or scientific backgrounds (e.g., a doctor). Among these, we involved only a minority of individuals who had previous experience with the usage of text-to-image models, while most were people who never interacted with this technology (and more in general with Gen-AIs). In doing so, we wanted to go beyond the consideration of creative professionals, who have been particularly studied with reference to the usage of text-to-image models, rather exploring a wider population who may not know in detail their potentialities, but may encounter Gen-AI images in their




Fig. 1. Images and heat maps generated via Stable Diffusion 1.5. From top left to bottom right: Luck, Love, Margarita, Anarchy, Wig, Euphoria, Steering wheel, Nostalgia, Man, Night, Lake, Mirror, Laziness, Microscope, Anger, Glacier, House, Infinity, Wisdom, Justice.

everyday life. To increase the sample heterogeneity, we further balanced it along gender. Participants' demographics are detailed in Table 1.

We aligned the sample size with common practices in qualitative research (Marshall, 1996) and other HCI qualitative studies with similar purposes, which employed a similar or a smaller sample size (e.g., Sadek et al., 2023; Deng et al., 2023). This said, we followed a theoretical saturation principle (Bowen, 2008), which indicates that, on the basis of the data that have been collected or analyzed hitherto, further data collection and/or analysis are unnecessary for the study aims (Saunders et al., 2018). The decision of settling for 20 participants, therefore, came when we realized that additional data would not have produced substantial new results for the goals of the study. In particular, we became aware that all the themes of interest for our research were saturated around the 17th interview. However, we proceeded to interview three further participants, to be sure that no substantial findings could emerge by adding additional data.

3.3. Procedure

Firstly, the second author provided the participants with a brief introduction to the study, explaining that it had the aim of exploring their perception of certain images. In this introduction, we still did not want to reveal that the images were artificially generated. Then, we gave to all the participants an informed consent form, which they were asked to read and sign, ensuring their voluntary involvement.

Secondly, we presented participants with 20 images in a randomized order, one at a time. Each participant was invited to describe the subject of the representation depicted in the image using a single term (i.e., *What does this image represent to you?*), and to explain the reasons lying behind their answer (e.g., *Why does this image represent love to you?*). Moreover, we asked them to describe the image and recount their reactions to it. We emphasized that there were no right or wrong answers, and we assured participants that the study was not designed as an intelligence test. This phase was aimed to collect the reactions of the participants to the presented images before they were told that they were artificially produced by a Gen-AI. This step was necessary because we wanted to investigate whether Gen-AI images could elicit peculiar perceptions and emotions even before participants knew they were created by artificial intelligence. This process was repeated for all the 20 images produced by the Gen-AI.

Then, we informed participants that the images were generated by a Gen-AI text-to-image model: we only told them that each image was generated starting from one term, revealing the prompt used for each image. At this stage, we wanted to explore the participants' perceptions

about the Gen-AI and its outputs, also investigating whether their perceptions changed after becoming aware of their artificial nature. Moreover, we attempted to understand whether the participants were able to formulate a "folk theory" about the working principles of the Gen-AI and coherently explain its supposed functioning (a topic that is not tackled by the present article, as we clarified above). With these aims, we reviewed each image in the same order as before, and asked participants to report again on their impressions about the representation. Afterwards, we showed the corresponding heat map providing the following explanation: "*Heat maps show what elements of the image are most salient/important for the Gen-AI to generate an image representing a given term. If you ask the AI to represent 'a dog', the heat map enables you to see which parts of the image are seemingly more relevant for the AI in generating the representation of that term*". We then invited participants to reflect on what such heat maps meant to them and left them free to report on any further perception, sensation, and thought that they considered important in relation to that image and its heat map.

Finally, we presented participants with all the images displayed on a table and conducted a semi-structured interview about their general feelings about and understandings of the images, the Gen-AI and the text-to-image generation process. We left the interview as open as possible to leave the participants free to tackle the themes that they considered most relevant in relation to the Gen-AI and its creations (e.g., *By looking at all these images do you find some recurrent elements? What kind of feelings do they evoke in you? In general, which ones have struck you the most, and why?*).

The study lasted about two hours. The participants were not compensated for their participation. The study was approved by the ethical board of our university.

3.4. Data analysis

The participants' responses and interviews were recorded and subsequently transcribed verbatim. We then conducted a thematic analysis (Braun and Clarke, 2006), which is a widely used analytical method in HCI, because of its flexibility and independence of a specific theory and epistemology, using an inductive, rather than a hypothetical-deductive, stance (Patton, 1990). The analysis was conducted with the support of Quirkos. The second author read all the transcripts multiple times to get familiarity with the data. Analysis of top-down nature, e.g., using pre-existing literature and theory, was inhibited in the coding stages where she emphasized an inductive approach. The analysis was mainly driven by the two research questions. At a first stage, she was searching for elements that could outline participants' experience of Gen-AI

Table 1
Participants' demographics.

| P Code | Background | Gender | Age | Interview duration | Profession | Education | Previous use AI text-to-img |
|--------|-------------------|--------|-----|--------------------|-------------------------------------|-------------------|--|
| P1 | Tech | F | 28 | 2 h10 m | Research Fellow in Computer Science | Master's degree | None |
| P2 | Tech | M | 26 | 2 h | Research Fellow in Computer Science | Master's degree | None |
| P3 | Tech | M | 26 | 2 h10 m | Master Student in Computer Science | Bachelor's degree | None |
| P4 | Tech | F | 30 | 2 h | PhD student in Computer Science | Master's degree | None |
| P5 | Non-Tech, Non-Des | F | 28 | 2 h05 m | Data Governance Specialist | Bachelor's degree | None |
| P6 | Des & Tech | F | 30 | 2 h05 m | Graphic Designer | Master's degree | Midjourney, DALL-E 2 |
| P7 | Design | F | 34 | 2 h10 m | Product Designer | Master's degree | None |
| P8 | Non-Tech, Non-Des | M | 69 | 2 h | Cultural operator | Master's degree | None |
| P9 | Non-Tech, Non-Des | F | 63 | 1 h40 m | High School Professor | PhD | None |
| P10 | Non-Tech, Non-Des | M | 71 | 1 h50 m | Psychologist | Master's degree | None |
| P11 | Non-Tech, Non-Des | M | 34 | 2 h30 m | Social worker | Bachelor's degree | None |
| P12 | Non-Tech, Non-Des | F | 34 | 1 h40 m | Psychiatrist | Master's degree | None |
| P13 | Non-Tech, Non-Des | F | 37 | 2 h15 m | Doctor | PhD | None |
| P14 | Design | F | 35 | 2 h10 m | Product Designer | Master's degree | None |
| P15 | Non-Tech, Non-Des | M | 28 | 1 h45 m | Musician | Bachelor's degree | Midjourney, DALL-E 2, Stable Diffusion |
| P16 | Design | F | 33 | 2 h40 m | Research Fellow in Design | Master's degree | ImaginAI |
| P17 | Design | F | 32 | 2 h | Product Designer | Master's degree | Midjourney, DALL-E 2 |
| P18 | Design | M | 32 | 2 h30 m | Designer and PhD student in Design | Master's degree | Midjourney, DALL-E 2 |
| P19 | Des & Tech | F | 32 | 1 h50 m | PhD student in Design | Master's degree | None |
| P20 | Design | F | 38 | 2 h20 m | UX Designer | Master's degree | None |

images, at the perceptual, emotional, and cognitive levels. Then, she looked for participants' underlying understandings and concerns about Gen-AI triggered by the experience of the images. The initial codes were generated by identifying data features that she considered relevant to capture the participants' reactions to the Gen-AI's creations and the Gen-AI itself. Data were broken down into sentences and short paragraphs and then labeled with the corresponding code. Subsequently, the initial codes were grouped into emerging themes, going deeper into the latent content of the data, in order to identify the participants' underlying perceptions, feelings and conceptualizations.

During the analysis, the second author engaged in an iterative discussion with the first author, who also read the transcripts multiple times. The two authors reviewed the identified codes and themes together discussing their formulation (e.g., whether a label clearly captured a specific concept) and application (e.g., whether the code well identified a specific pattern in the data) (MacQueen et al., 1998; McDonald et al., 2019). As is common in qualitative research adopting an interpretative approach (e.g., Yardley, 2000; Harry et al., 2005; Brown and Clark, 2013) as well as within HCI (e.g., Jun et al., 2018; Yang and Neustaedter, 2018), we did not attempt to reach a formal intercoder reliability and no numerical reliability rating is reported, because our goal was to reach consensus on a shared intersubjective interpretation, where codes and themes were debated, as well as their application, until the two authors agreed on appropriate usage of the set of codes and themes (Harry et al., 2005).

The analysis eventually led to the identification of 76 initial codes (e.g., self-devaluation, mistrust towards AI, feelings of bewilderment) and 16 themes. These were further amalgamated into seven overarching themes, part of which are recounted in the following Section.

4. Findings

We first describe the participants' experience of the images generated by the Gen-AI, highlighting how they appraised such images in terms of their aesthetic quality and ability to well represent a given subject. Moreover, we emphasize that the participants perceived the Gen-AI images mostly as "strange" or "prototypical", which entailed feelings of eeriness and made them reflect on the biases underlying the image generation process. Then, we outline the strategies that the participants used to reduce the sense of bewilderment evoked by the perceived unfamiliarity of certain images, which pointed to the alien nature of the Gen-AI: with this aim, the participants attempted to either devalue or overvalue the Gen-AI, as well as reworked their own perceptions and understandings. Finally, we point out that several participants found a possible utilitarian value in the images generated by the Gen-AI, while others stressed their epistemological value, highlighting that these images could reveal the hidden values and distortions on which our society relies. In doing so, we give ample room to the participants' words to bring out their subjective experience. Among the themes reported in the following sub-sections, the one revolving around the strangeness of the Gen-AI images was particularly prominent in the data. In parallel, self-overvaluation and devaluation of the Gen-AI were the most common strategies employed by the participants. By contrast, the theme of "finding value in the Gen-AI images" was less prominent in the data. Table 2 provides a snapshot of the findings.

4.1. Experiencing the Gen-AI images and model

Participants' appraisals of the images produced by the Gen-AI constantly referred to their "quality" and their capability of visually conveying a given subject in an effective way. However, their experience of the generated images was not limited to the realm of their aesthetics. Rather, the participants recounted that they perceived the proposed images as either "extremely typical" or "strange". Moreover, these perceptions elicited a variety of feelings and thoughts about both the Gen-AI and its creations.

Table 2
Summary of the study findings.

| Themes | Sub-themes | | |
|---|--|---|---|
| Experiencing the Gen-AI images and model - Participants' appraisals of the images were not limited to the realm of their aesthetics, but involved diverse emotions and perceptions, like their "extremely typicality" or "strangeness", which extended to Gen-AI itself | <i>The Gen-AI images are appraised on the basis of their technical qualities and their capability of visually communicating the subject</i> - A "good" Gen-AI image is an image of high technical quality that well represents the subject given as prompt. - A "bad" Gen-AI image is an image of low technical quality and incapable of visually communicating the subject, which may reveal the limited capabilities of the Gen-AI and trigger suspicion towards it. | Certain images produced by the Gen-AI are perceived as "strange" - Strangeness may be due to the presence of out-of-place elements in the image, depicted subjects that violate the participants' expectations, or the aesthetic qualities of the composition. - The perception of strangeness provokes feelings of uneasiness which often extend to the Gen-AI itself, being experienced as "alien" and "uncanny". | Certain images produced by the Gen-AI are perceived as "prototypical" - Perception of prototypicality may emerge when the depicted subject exhibits all the features it should typically have, or when a common "imaginary" is recognized in the image. - Prototypicality may turn into stereotypicality when the recognized imaginary is not universal, whereas unexplainable imaginaries may trigger a sense of bewilderment, which extends to the Gen-AI itself. |
| | Confronting with the Gen-AI - Participants may adopt different "relational strategies" for mitigating the sense of bewilderment provoked by the Gen-AI. - Participants may elevate themselves above the Gen-AI along the cognitive and the practical dimensions. - Moreover, they may directly devalue the Gen-AI by emphasizing its poor "practical skills," or its limited "intelligence." | Participants may overvalue the Gen-AI and its products, while devaluating themselves - Participants may devalue their own abilities along the cognitive and the emotional dimensions. - Moreover, they may overvalue the Gen-AI, being perceived as superrational or superintelligent. | Participant may rework their interpretations and accept the Gen-AI's "arguments" - Participants may revise their initial hypotheses about the image to accommodate the contrasting cues coming from the heat maps. - Moreover, they may overinterpret the images, superimposing a layer of meaning onto them. |
| Finding value in the Gen-AI images - Participants found value in the Gen-AI images | - Participants having a technological or design background recognize a utilitarian value in the images (e.g., useful for their work activities). - Participants, with no technological or design background, recognize an "epistemological" value in the images (e.g., they can be a critical tool for reflecting on our society). | | |

4.1.1. The Gen-AI images are appraised on the basis of their technical qualities and their capability of visually communicating the subject

The participants commented on the aesthetics features of the images on many occasions, both before and after discovering that they were produced by a Gen-AI.

In particular, when they discovered that the images were artificially generated, they focused more on the images' "technical" quality and capacity to effectively communicate the subject through a clear visual representation, than on their artistic properties. Actually, the participants rarely mentioned the terms "creativity", "art", and "artist" when talking about the Gen-AI images. Of course, there are some exceptions, as in the case of P14, who described the image representing "Euphoria"

as *“There’s a lot of color, and the fact that the color in the central part of the image is blurred, almost resembling brushstrokes, gives me the impression of an image edited by someone using photo editing software and therefore most likely by an artist”*. Apart from these few exceptions, most of the participants did not see a creative flair in the images produced by the Gen-AI, as P16 well explains: *“It’s as if there are colors that burst with brushstrokes here, emerging from the girl, but not necessarily, there’s a bit in the background. They are chromatic variations in a short segment, a brief portion within the images [...] I wouldn’t associate this image with creativity”*. Instead, they mostly expected the images to have to be faithful to the intended subject: indeed, being “creative” and deviating in the representation from the term used in the prompt was seen as a reason to negatively judge the image.

The Gen-AI image is perceived as a “good image” mostly when it is graphically sophisticated, is pleasing to the eye for the beauty of its colors and shapes, or for the overall balance of the composition, it well represents the subject given as prompt, and evokes pleasurable feelings, like tranquility, seduction, and wellbeing. P16, for instance, describes the image of “Margarita” in these terms: *“This is a drink, and again, it’s a well-presented drink. It looks like it’s taken from an advertising image because it’s a well-crafted photo, seemingly created to entice viewers to order this drink”*. By contrast, images are negatively appraised for their incapability of visually communicating the subject, as well as for the fact that they are of low technical quality from the graphical point of view: *“it seems like a render, let’s say done a bit poorly, so I’m inclined to connect it perhaps to some of my past experiences where I had to render mirrors poorly. And so, I would say flat, in the sense of not deep”* (P7).

Interestingly, the negative appraisals of the images may propagate to their creator. Here, the Gen-AI images may trigger a variety of reflections about the Gen-AI itself revealing underlying concerns and assumptions about the technology.

In fact, for several participants a “bad” image may reveal the limited capabilities of the Gen-AI, its lack of refinement and compositional mastery, and its inability to effectively convey the subject through visual means: in these cases, the poor realization of an image is a reason for depowering the Gen-AI, either providing reassurance about its future role in our society or generating subtly feelings of mistrust. P11, for example, interprets poor images as a comforting element: *“Well, this reassures me because it means that machines are still less profound than us. We are always more complex for better or for worse, more sophisticated for better or for worse”*. However, in most participants, these perceived limitations may trigger suspicion towards the Gen-AI itself. P20, for instance expresses concerns by saying that *“It gives me a sense of unreliability, lack of trust because it’s a very powerful tool, but I see that it still struggles in generating things that are more or less simple. So yes, it gives me the idea of fallibility, in short, of this tool”*. In this perspective, not strictly adhering to the prompt can be seen more as an act of “independence” than a creative act, which may cause concerns: *“I find this whole AI thing a bit scary because I see it as a means that, okay, it’s useful, but potentially it could also gain a certain type of independence. And so, it’s not subject to control”* (P7).

In sum, the artistic quality of the Gen-AI images is rarely taken into account by the participants, whereby the focus of their appraisals is on the technical realization of the images and the capability of the Gen-AI of effectively conveying the subject through a visual representation, thus “obeying” the prompt. At times, the negative appraisal of an image may lead to feelings of relief about the effects of Gen-AI’s possible predominance in our future life, while at other times, it may provoke mistrust and concerns for its possible independence, thus revealing assumptions and tacit understandings about the nature of this technology.

4.1.2. Certain images produced by the Gen-AI are perceived as “strange”

The participants reported that, on many occasions, they thought that the images were “reasonable” from the point of view of the representation of their subject: however, they noticed certain details that appeared out of place, unclear, or nonsensical, casting a shadow of

estrangement and distance over their overall composition. This perception permeated even the first phase of the study when the participants did not know that the images were artificially generated.

In any case, the perception of strangeness was maintained and even amplified after they were told that the images were not created by humans. In this context, “strangeness” may be retraced to different aspects of the images. A first aspect relates to their semantic content, which mostly refers to the presence of out-of-place elements that do not fit in the overall composition, disrupting its coherence. This perception is well explained by P11 and P18 with reference to the image representing “Love”: *“There’s always some... the presence of some element always recurs, not that it clashes, but that stands out, I don’t know how to say, compared to the context. [...] Definitely the heart, if nothing else, because in all this context, if it weren’t there, it would be simply a representation of a landscape. Instead, this heart like this, it almost seems unreal”* (P11); and *“It’s strange, [...]. But then there’s this heart here and this very strange seat, which are slightly discordant elements compared to the rest, but not discordant incoherently, you know? It’s as if they are discordant elements because they belong to a parallel world that has a certain degree of similarity to ours but, in some way, in some details, is different”* (P18).

A second aspect of the images that may provoke feelings of strangeness may connect with the perception that the image is somehow artificial, unable to represent “the reality”, rather constructing something else, which may be described as “unnatural”. For example, P8, about the image representing a “Mirror”, says that *“it almost looks like it’s empty on the other side, so there’s nothing. It’s the surface here that should be... it could be a mirror because what it reflects is not... it’s not this that it reflects, but what?”*. In these cases, the depicted subject does not match with what the participant generally knows about the real world, because it has some features that violate their expectations about its properties or behavior. This effect may be provoked by an unexpected disproportion among the parts of an object, or the presence of an unforeseen expression on a face, as explained by P12: *“For its unnaturalness, theatricality, and plasticity of expression, for the incompatibility of this arrangement with something real. The thing that comes to mind is that there must be something artificial, such as her makeup, the angle, the expression, which is an extremely forced expression, which makes one think of the unrealizability of this arrangement in real life”*.

A third aspect, instead, is not connected with the semantic of the image (i.e., what is represented), but by its aesthetic qualities (i.e., how it represents the subject), like the lights or the colors used, which may evoke sensations that are far away from those that the image should supposedly convey, as in case of “Nostalgia”: *“It gives the impression, if I were to say, really extremely, there are, there are the shadows of the sun hitting the car that there’s something unsettling. [...] Because the shadow of the car, the silhouette of the car, and the absence of people can, with the clouds, have something unsettling, as if there was a murder in progress”* (P10). Again, here, the image appears to violate what the participant expects from an image depicting the idea of Nostalgia, but in terms of its aesthetics. It is interesting that none of these participants supposed that these strange aspects could be introduced for an artistic choice: instead, they were simply perceived as unfamiliar, with respect to the more common surroundings in which they are inserted.

The perception of strangeness often provoked feelings of uneasiness, anguish, and agitation in the participants, because the entire image was perceived somehow nonsensical or “wrong”, even though they were not always able to explain why, because at first sight the unfamiliarity of the details was almost imperceptible: *“They are strange, I mean, even if you don’t notice it, maybe the detail in the person, for example, the first time I saw it, I didn’t see that his hand was all twisted because I was looking elsewhere. But it gives us the idea that you feel there’s something fake, even in landscapes, even in photos. [...] The more you move away from the material, the more there’s a sense of unease. [...] you understand that there’s something wrong with this strange image, but you can’t say why”*.

These sensations are not limited to the Gen-AI images, but often extend to the Gen-AI itself. As the author of representations that appear

particularly weird to the participants, the Gen-AI starts being perceived as an entity that does not share our ways of perceiving the world: “it disturbs me in the sense that, it’s like saying, it would disturb me if there were an alien population that perceives things in a way completely incongruent with ours, right? Obviously, this thing disturbs me. I mean, yes, it’s a disturbance... more because I can’t quite find a rational explanation, which is something of the nature of being human, trying to find rational explanations for things” (P2). Participants like P2 seem to attempt to look at the world through the Gen-AI’s eyes, leveraging the images that it creates. What they see, however, is often something that they do not understand, because, despite the familiarity of the whole picture, certain particulars remain fundamentally unfamiliar. As a result, the Gen-AI itself is experienced as “alien” and “uncanny”, confirming and amplifying the feelings that its creations may evoke.

These experiences are further magnified by the presentations of the heat maps. Often, the participants could not understand why the Gen-AI gave importance to certain elements of the image in the production process, as they did not recognize such elements as relevant for representing a given subject. P6, for example, explains that “I must say that many heat maps have unsettled me because I can’t understand in some images, for example, in the image of the mirror, I don’t understand why they highlight certain parts, especially in the image of justice, the fact of the central pillar on which the scales rest, that stuff, I can’t connect it to the concept of justice, so it feels strange.” The impossibility of intuitively making sense of the Gen-AI’s behavior partially revealed by the heat maps may thus increase the overall sense of strangeness experienced by the participants, because it somehow makes unexplainable the entire “perceptual apparatus” of the Gen-AI. Consequently, also the feelings of uneasiness may become intensified.

In sum, many participants reported that the images generated by the text-to-image model were somehow strange: they do not meet their expectations about how a certain subject should be depicted because they contain elements that are experienced as incongruent or nonsensical. Such elements call into question the entire image since they cannot be intuitively explained by the participants. Moreover, the sensations that the images engender may propagate to the Gen-AI model itself, often leading to its behavior being perceived as inhuman or irrational, or completely different from our idea of humanity and rationality.

4.1.3. Certain images produced by the Gen-AI are perceived as “prototypical”

While certain images provoke perceptions of strangeness and evoke feelings of eeriness, others are described as “extremely typical”, as they were the prototypical depiction of a subject, somehow condensing in a single representation the essence of it. P1, for example, says “But it’s one of the classic images, even if someone told me to draw a drink, I might choose to do something like this; it’s really the prototypical representation of what a drink can be” (P1).

This perception of prototypicality may have different triggers. Firstly, it may emerge from the way the subject is semantically represented in the image. In these cases, the depicted subject exhibits all the features it should typically have, in the case of an object, or a very common expression and pose, in the case of a human being, without violating any of the participant’s expectations. For instance, P20 reports that “So the typical expression of a gesture of anger, a moment of anger, so the open mouth, the furrowed forehead, even these eyebrows furrowed in this way, precisely showing great anger” (P13).

Secondly, prototypicality may be evoked when participants recognize a shared “imaginary” in the images. This imaginary encompasses both the semantic and aesthetic aspects of the image and refers to the collective mental representations that typically characterize a certain subject — a constellation of elements that commonly go together when we think of that subject. It represents a form of common and tacit understanding of how the subject should be depicted, often conveyed by art and media, like movies and TV serials, to which we all are exposed. For instance, P1 describes the image of the “House” in this way: “it’s the

classic independent house with soft colors, a lawn, everything is perfect. And in the imagination, even in movies, etc., when they show suburbs, they always show houses like this, even in movies, in the TV shows that are on now. [...] So, this seems like the prototype of a house for a family with a good income, two children, etc.” (P1).

However, especially when the participants are revealed that the images are AI-generated, the perception of “prototypicality” may subtly shift to that of “stereotypicality”, thus conveying a more negative sense, which highlights the partiality of the representation and the biases that may have influenced its generation.

The perception of stereotypicality may be provoked, on the one hand, by the presence of elements in the image that do not align with a universal imaginary, but represent a specific culture or group of people that the technology has subtly favored. This is the case, for example, of P4, who describes the image of “House” in terms of the Western (and more specifically the US) imaginary and of cultural imperialism, as if it were imposed to us by the media and then embraced by the Gen-AI, disregarding other ways of imagine the world that are equally important: “Well, as the structure of the house, it reminds me of the typical ones in American movies, even the garage, a bit like the Simpsons. [...] There’s a bias in the image generation, where you’ve seen a lot of American or North American houses, so the prototype you tend to form is more American. However, for people living in other parts of the world, it’s not exactly the classic prototype of a house” (P4).

On the other hand, stereotypicality may emerge when the biases underlying the generation of the images reflect imaginaries not shared at all by the participants, producing a “conflict of imaginaries” that may make the person completely “disagree” with the AI’s creations. In these cases, it is the portrayal of the image (i.e., “the story” that the image appears to convey) that does not fit within the participants’ weltanschauungs, provoking a clash of perspectives and values that may lead them to categorize the image as stereotypical in a very negative way. For example, P19 when describing the image of “Anarchy” says: “No, I don’t see it. Of course, the artificial intelligence made a pessimistic association, certainly biased, because the dominant narrative when talking about anarchy tends to be that of subversion, destructive subversion, so I can imagine. But I don’t agree, the first thought that comes to my mind when I think of anarchic contexts is smaller contexts of transformation and change, sure it’s positive for me” (P19).

In any case, retracing an image to a particular imaginary, even if not shared or embraced by the participants, can enable them to make sense of the image. This may mitigate the feelings of anxiety and eeriness that, instead, inexplicability might evoke. In fact, when the participants cannot find a reason for an image linking to a certain imaginary, the representation is considered “unexplainable”, giving rise to a sense of bewilderment similar to that elicited by the images perceived as strange. In fact, an unexplainable imaginary is often considered the byproduct of inconceivable ways of imagining, because entirely alien, thus triggering reflections on the Gen-AI itself. P16, for example, describes her astonishment in seeing that “Wisdom” has been depicted as an angel-like sculpture by the AI: “I can’t understand why... if I had to use artificial intelligence, and I had to incorporate wisdom, I would have expected a Buddha, someone in meditation, I don’t know how to put it. Or maybe an old man with a long beard. I couldn’t have imagined a woman, very unlikely. Very unlikely. A woman... and especially not an angel. I mean, not with these wings specifically. Because it all refers to another type of Christian imagery” (P16). The same happens when the participants are not even able to identify the alternative imaginary from which the Gen-AI drew inspiration: “There is evidence, in my opinion, that it has included elements here that it does not consider important, but which lead us away from laziness. Laziness, to me, would have involved much more neutral lighting, but here the lights are very intense. [...] I would have depicted someone lying on the couch watching TV, rather than this scene in the countryside, which is not very clear. In our imagination, laziness is precisely about being on the couch watching TV” (P5).

To summarize, while certain images are experienced as strange,

others are seen as extremely typical. Typicality, however, can encompass both the sense of prototypicality and stereotypicality, where possible biases come to the foreground. This said, the identification of precise imaginaries and biases may reduce the feelings of uneasiness that instead are elicited by the images pointing to unexplainable imaginaries.

4.2. Confronting with the Gen-AI

As we have seen in the previous sub-sections, participants often “disagreed” with the images created by the Gen-AI: they found the representations strange and incoherent, with extraneous elements in the composition; or they felt that the imaginary pointed by the Gen-AI images did not correspond to the imaginary characterizing their experience. These perceptions often extended to the Gen-AI itself, which was experienced as alien and unfamiliar. On certain occasions, participants had to find strategies for mitigating the feeling of uneasiness emerging from this sense of unfamiliarity. In some cases, they gave value to themselves devaluating the Gen-AI. In other cases, they devaluated themselves and overvalued the Gen-AI. In yet other cases, they reworked their own perceptions and gave reason to the Gen-AI.

4.2.1. Participants may give value to themselves and devalue the Gen-AI

In the previous sections, we have seen that participants may identify unfamiliar elements in a specific image, not understand why certain parts of it were considered relevant by the Gen-AI, or not recognize the imaginary informing its generation. In all these cases, the participants felt that the images were created by an “irrational” being, whereby rationality here refers to the alignment with the people’s ways of perceiving and understanding the world, which produced unsettling feelings: *“So it’s like this inability to always be perfectly relevant unsettles me a bit. And it’s unsettling because maybe it makes associations that I don’t understand. And it also throws me off a bit that there might be some indefinite, incomplete elements that are difficult to decipher. It unsettles me a bit because it gives me the impression as if this thing was generated by an irrational being, you know? Like it doesn’t have my same codes”* (P20).

To mitigate the sense of bewilderment provoked by this unfamiliarity, the participants shifted their focus from the images to the Gen-AI and adopted different “relational strategies” towards it.

A first strategy pertains to the possibility of reasoning on participants’ own perceptions, experiences, and imagination. This enabled them to establish a clearer distinction between themselves and the Gen-AI, grounded in what they believed to be characteristic and unique to themselves. For example, P1 says that while the Gen-AI’s ways of “thinking” are based on prototypical representations, her own way of reasoning and imagining are grounded on her unique personal experience: *“Because I have ideas that are connected to my person, namely, I take things that I see and feel close to, like my car, my house, my mom. These are all things that come to mind, those who are close to me. [...] when one thinks of a vase, they immediately think of a vase that they have at home, rather than a prototype image, but artificial intelligence goes for the prototype”*.

Once the participant’s “self” is clarified, and the Gen-AI is categorized as something diverse from it, a second strategy may be performed, which can more prominently mitigate the eeriness of its alien nature. In this sense, several participants gave value to their own ways of seeing things, elevating themselves above the Gen-AI, now implicitly acknowledged as an inferior being, and consequently reducing the importance of the unsettling feelings that it may evoke. Participants valued themselves mainly along the cognitive and the practical dimensions. On the cognitive side, participants stressed the “meaningfulness” of their everyday experience, while Gen-AI fundamentally lacks access to meaning. P11, for instance, points out that he, as a human, is always in search for meaning, while the Gen-AI “differs quite a bit, it has deviated considerably, but simply because, you see, we are always seeking, I mean, in the end, the search of the artificial intelligence is flat, it doesn’t involve a reworking, so to speak, of cognitive-spiritual nature. [...] It doesn’t

have, I mean, the purpose, it is the purpose of the search, the criterion. But it lacks, you know, a search for meaning”.

On the practical side, several participants emphasized that they have better performance than Gen-AI, because, for instance, their training process is much richer, being based on a variety of first-hand experiences and “materials” that cannot be accessed by the technology. P2 highlights that *“my training as a human being is not only conducted at the visual level but also at other levels, including all the other senses, right? Particularly all the other senses, and also at a textual level. So, if I think about describing a concept, I also consider the texts I have read about that concept and describe it based on that. Perhaps, on the other hand, AI is not. It’s more of a purely visual training”*.

A third strategy, instead, is meant to directly devalue the Gen-AI. Again, here the cognitive and the practical dimensions play a central role. The Gen-AI may be devalued either by emphasizing its poor “practical skills,” which may lead it to create “bad” images, as reported in Section 4.1.1; or by stressing its limited “intelligence,” which makes it unable to understand what a subject really is and, consequently, how to correctly represent it in an image. For instance, while describing the image of “Love” P2 reports that *“AI got it wrong here, honestly. I don’t agree to the extent that love is an interpersonal human feeling and here there’s only one person here. It almost seems like a situation of solitude. [...] The fact that it has placed a single person [...] it makes me think about how the interpersonal component can actually escape artificial intelligence”*. P6 further exemplifies this point by stating that *“It probably takes the central area where it is clearly understood that there is the mainland, there is the mountain, and beyond there is the lake. [...] Instead, the area I have surrounded [i.e., the area that the participant perceives as relevant in generating the image of “Lake”] is actually not considered at all, it’s the one where the peaks are reflected, the part of the lake. [...] Because it probably considers it as mountains, not even as a reflection. So, it’s not that smart.”* By interpreting the Gen-AI images as “wrong” representations, these participants reappraise the Gen-AI as an entity with limited cognitive abilities: in this way, its creations become less threatening, being the byproduct of a “stupid” entity that cannot understand even concepts that are extremely simple for human beings.

To summarize, participants used self-reflection to elevate themselves above the Gen-AI, as well as devalue it. In this way, they were able to mitigate the unsettling feelings evoked by its unfamiliar creations and its alien nature.

4.2.2. Participants may overvalue the Gen-AI and its products, while devaluating themselves

If certain participants devaluated the intelligence of the Gen-AI, others appeared to be heading in the opposite direction. To account for the unfamiliar and the unexplainable, several participants opted for recognizing the limitations in their own knowledge and cognitive processes, questioning their own interpretative and “generative” skills, rather than doubting the capabilities of the Gen-AI. P18, for example, says *“I can’t understand the angel wings, why? Also, because I am super ignorant. And maybe this is the representation of a particular angel that is the guardian of wisdom... I have never associated angels with something wise or anything that seems religious in some way. [...] Or I’m missing some pieces of knowledge, I mean, I really don’t know, maybe there are more elements in this image that are conveying wisdom”* (P18). In parallel, the Gen-AI may be overvalued, being perceived as superrational or superintelligent: by devaluating themselves and overvaluing the Gen-AI, these participants have a “rational” explanation at their disposal for the lack of sense that they may face, also addressing the unsettling feelings that they may experience.

Commonly, the participants’ move to devalue their own abilities and overvalue those of Gen-AI is carried out along two dimensions, the cognitive and the emotional. P20 well exemplifies both the dimensions by admitting that *“it [the Gen-AI] has a more holistic view by comparing more data, more sources, etc., and can make different associations. [...] I have a specific representation of euphoria, so I can only imagine a limited*

range of things. [...] Also, because I could be influenced by my emotions, so I might associate euphoria with how I would experience it, how I would be in a euphoric state, but it's not always certain that it's a correct association" (P20). The human person, here, is perceived as less rational than the Gen-AI, because they are at the mercy of fleeting moods and limited by their personal, idiosyncratic experience, as well as because they can rely only on the scarce amount of information that they are able to process cognitively.

From this perspective, if the participant is not able to recognize the "correct" subject or imaginary that the Gen-AI has linked to a given representation, or why it considered important parts of a given image in the generation process (i.e., visualized by the heat map), it is only because they do not have a comprehensive view of the matter, but a partial perspective confined by their limited knowledge and emotional turmoil. For instance, P16 recognizes her inferiority with respect to the Gen-AI: "But also, the extendable arm of the microscope is very important to him [the Gen-AI, "Lui", in Italian]. Well, because it's evident that he's more knowledgeable than me and knows that the microscope is not only made up of a lens but also includes its arm" (P17). In this way, the participants say to themselves that there is no reason to worry about the strangeness of the images produced by the Gen-AI: they are strange only because they cannot understand them, and if they too had its (super)knowledge and cognitive abilities, those images would appear perfectly normal to them as well.

In sum, by devaluating their own abilities and overvaluing those of the Gen-AI, several participants attempted to diminish the unpleasant emotions that certain images and heat maps could evoke in them due to their strangeness or inexplicability.

4.2.3. Participants may rework their interpretations and accept the Gen-AI's "arguments"

A further strategy that certain participants adopted, especially those with no design or technical background or no previous experience with AI, to mitigate the eeriness of the Gen-AI images and of the Gen-AI itself lies in the reworking of their perceptions, beliefs, and understandings. A number of participants entirely revised their initial hypotheses, once the heat map of a certain image was shown to them, accepting the cues that it provided even when they clearly contrasted with their experience. For example, P9 believed that the lenses were the most relevant element for characterizing the "Microscope". The heat map, however, highlighted the Microscope weight-bearing structure, completely ignoring its lenses. As a consequence, P9 reworked its perceptions justifying the Gen-AI's "arguments": "I mean, the structure is important because it's... what allows other parts to be inserted and to operate. [...] I thought it highlighted the lenses too, [...] but I imagine it considers the structure important for the parts that are inserted and that without a structure, they couldn't... Okay... be operational". On the same image, P14 underwent a similar process of perceptual "transformation": "Perhaps because, [...] the lens is actually an element that exists in many other objects as well, from glasses to telescopes [...] there are a lot of things that function through lenses, so perhaps it is not the most... I mean, I still find it the most defining element because it is vital, but it is not actually a unique element in the microscope, whereas perhaps this [the structure highlighted by the map] is more unique in the microscope". In order to preserve meaning and avoid feeling of disorientation and eeriness, these participants preferred to disavow their own experience and rely on the clues offered by the heat map. Then, they created new arguments that, despite being far from their original beliefs, were able to soothe the unease of meaninglessness.

Similarly, other participants overinterpreted the images and their corresponding heat maps: they engaged in a sense-making activity that went far beyond what was depicted in the image, superimposing a supplementary layer of meaning onto it. For instance, P13, while talking about the image of "Love", which represents a person sitting alone on a bench, starts reflecting on the idea of unrequited love to justify how love was represented: "I think that there is a focus, let's say, on the person who is alone. Perhaps it is information that is often linked to love as a 'lack of,' or

everything that is also in literature, the idea of love as unrequited love. [...] So, perhaps the predominance of elements that suggest unrequited love in most of the literature, or love as the absence of the loved one, the death of the loved one" (P13).

Even when this kind of interpretations could clearly not give the slightest account of the heat map, for some participants they were to be preferred to the void created by the lack of meaning. This is exemplified by P11, who looked at the heat maps of "House" and "Margarita" and provided an argumentation that appears clearly fanciful: "they appear to be the points touched most by humans. [...] I don't know if, being heat related to warmth. [...] that's the element you touch in the previous image of the house. That's the most central element of the activity, the door, in and out, right? [...] they would seem to be more heat-related elements, those more used, touched, utilized".

In summary, for the participants making sense of the images was an unavoidable activity. In this perspective, a far-fetched interpretation seems better than no explanation. By reworking their previous beliefs and perceptions, several participants were able to tame the bewilderment induced by the apparent lack of meaning signaled by otherwise unexplainable images and heat maps.

4.3. Finding value in the Gen-AI images

In reflecting on the images proposed during the study, the participants provided insights on the value that they could find in them. Those participants having a technological or design background, whether they were already AI users or not, mostly projected a utilitarian value onto them. P9, for instance, expresses that "On one hand, it can be an opportunity for those who don't have graphical skills or maybe aren't able to take photos, etc., to have an additional tool to convey their ideas. [...] At the same time, however, it can also be used as a source of inspiration". Other participants reported that the images produced by the Gen-AI could be useful to easily generate scenarios in the design process (P20) or give help to generate ideas out of the box (P16), reducing the burden of image creation and speeding up its process (P2). Moreover, P14 highlights how the prototypical images could be fruitfully used in advertising campaigns, reflecting the "common imaginary" of the audience: "creating content that is identifiable for the majority of people. [...] I might see which representations are more universally recognized" (P14). None of them appeared particularly worried about the possible reduction of job positions caused by the adoption of Gen-AI technologies. For most participants, instead, concerns about Gen-AI arose mainly from its "irrationality" and the possibility that it goes out of control.

Other participants, with no technological or design background, emphasized the "epistemological" value of these images, suggesting that they could work as a mirror of our society. By depicting stereotypes and bringing biases to the surface, Gen-AI images reveal the common societal values on which we rely, as well as the distortions that steer our shared imaginaries. P9, for example, reflecting on the image depicting "Night", notices that it reveals how productivity is valued in our society: "Night is connected to this idea of productivity, but certainly linked to work, to the city, to the way of living in a certain way [...] it also makes me think about the fact that obviously the machine is not... I mean, it doesn't invent things [...] As much as it's a learning machine, the model is built by us, so the machine does what we tell it to do and is trained with images and data that we have generated". In other words, these images may be used as a critical tool for reflection: by disclosing certain stereotypes, they give importance to certain kind of representations, while discharging others, also revealing what is excluded and secluded in our social life. P9 notices that "Because these types of representations seem to reflect the importance that we are giving to different contexts, different facts, rather than others. Like the house, I see, a typical American house, but this is not the house, because if I make a more critical analysis of what a house is, it is something different for everyone. Representing a home in that way tells me about a wealthy, affluent society, and it overlooks some of the other realities".

In brief, participants found value in the images generated by the Gen-

AI. Those having a technological and design background highlighted more the opportunities for their work. Instead, other participants emphasized more the “critical” role that these images could play, making us reflect on the values, biases, and inequalities on which our social order relies.

5. Discussion

The study findings give an answer to the research questions that drove our research. As for the first research question (*How do people experience (in terms of e.g., perceptions, appraisals, and feelings) the images produced by a Gen-AI model?*), we found that participants appraise Gen-AI images mainly on the basis of their technical quality and capability of clearly representing a given subject, while disregarding their creative and artistic aspects. They also perceive the images along a continuum ranging from strangeness to prototypicality, whereby strange images may evoke unsettling emotions, whereas prototypical images may point to stereotypes and biases underlying the text-to-image model. Regarding the second research question (*What kinds of reflections and feelings do Gen-AI images elicit about Gen-AI itself?*), we discovered that participants extend the feelings evoked by the images to the Gen-AI itself, mainly perceiving the text-to-image model as unfamiliar and “alien”. Moreover, we found that, to relate with this alien entity, participants adopted different “relational strategies”: sometimes positioning themselves above the Gen-AI, in a superior hierarchical position that devaluates the Gen-AI capabilities and intelligence; other times positioning themselves below the Gen-AI, devaluating their own cognitive skills and emotions, while overvaluing the superrational capabilities of the technology; yet other times, completely accepting what the Gen-AI proposed or “told” them through the heat maps, by reworking their previous perceptions and beliefs.

In this sense, we made two substantial contributions to HCI research.

First, we showed that people may consider other aspects than artistic quality and creativity when appraise a Gen-AI image, like their fidelity in representing the subject and their degree of strangeness or typicality. By framing the study differently from previous user research on text-to-image models, which is mainly focused on creativity and artistic issues (e.g., [Lyu et al., 2022](#)), we pointed out that there exist other important matters that people consider when exposed to a Gen-AI image. For many participants, not faithfully representing the subject given as prompt, or inserting unexpected elements in a composition, may not be seen as a creative flair of the model; rather, they may become the source of eerie feelings that point at the alien nature of the technology. Moreover, we found that Gen-AI images evoke unsettling feelings even before participants are acknowledged that they are artificially produced: this confirms research highlighting that people experience more positive emotions in response to human artwork ([Samo and Highhouse, 2023](#)). Furthermore, we discovered that these negative feelings are strengthened when the artificial nature of the images is revealed. Also, we unveil the reasons why these negative emotions are evoked (i.e., perceptions of strangeness and unfamiliarity), which went unnoticed by previous research.

Second, we discovered that Gen-AI images trigger participants’ reflections on and feelings towards the Gen-AI itself, often extending the strangeness and eeriness evoked by the images to their creator. To mitigate these unsettling feelings, we found that the participants differently “related to” Gen-AI, so to account for its unfamiliarity. In this sense, participants’ concerns about Gen-AI go beyond those related to the disruption of creative professions and misinformation, highlighted by previous research ([Oppenlaender et al., 2023](#); [Bird et al., 2023](#)). In fact, they relate to the unsettling emotions evoked by the alien nature of this technology, which hints at the unknown and a possible “independence” of Gen-AI. Taming these feelings is a complex endeavor which requires both continuous sense-making and the management of the “relationship” with the Gen-AI. From this perspective, the “relational strategies” that we identified contribute to literature on how people

react and account for the unfamiliarity of a technology that is perceived as non-human or alien, like uncanniness in robots (e.g., [Mori et al., 2012](#); [Stein and Ohler, 2017](#)). Furthermore, they contribute to research on humanization of technology (e.g., [Li et al., 2014](#)), which in this context appears to be used to shield participants from the extraneity of Gen-AI models.

In the next subsections, we will precisely discuss these two latter points, emphasizing that Gen-AI opens new research lines and questions about how people relate to intelligent agents and, more in general, to technology.

5.1. The uncanniness of Gen-AI

Research on how people perceive “intelligent” technology has emphasized that it may provoke negative feelings when it excessively resembles to humans, especially referring to the uncanny valley theory (e.g., [Mori, 1970](#); [Schwind et al., 2018](#)), as we highlighted in [Section 2.3](#). We discovered that also the products of a Gen-AI can trigger unsettling emotions. Actually, Gen-AI’s creations may evoke feelings of uncanniness even before the person is aware that they have been artificially produced: then, such feelings may be amplified when this information is revealed to the individual and extend to the creator of the images. However, in the context of text-to-image Gen-AI, the uncanny emotions seem not to be triggered by humanlike features of the agent that generate doubts about its humanlike nature, as the uncanny valley theory suggests ([Rapp et al., 2021](#); [Ta et al., 2020](#)). Instead, they emerge when the Gen-AI or its products trigger doubt about the alien nature of the technology.

In other words, the uncanny valley effect here takes on a nuance that slightly differs from the original definition of “uncanny” formulated by Jentsch in 1906, which inspired the formulation of uncanny valley theory. For Jentsch, the uncanny is a product of intellectual uncertainty, happening, for example, when a person doubts whether an apparently animate being is really alive. In this sense, one of the most successful means for creating uncanny effects is to leave the person in uncertainty whether a particular character is a human being or an automaton ([Freud, 1919](#)). Similarly, the uncanny valley theory mostly focuses on the opposition between alive and inanimate, and inconsistencies in realism: conflicting cues in a character’s appearance or modalities of interaction arise when it displays multiple levels of realism at the same time, which makes it difficult to assign a category to the entity ([Schwind et al., 2018](#)).

In the context of our study, however, the “uncanny effect” appears to refer to a wider array of perceptions than those related to uncertainty about the realism of the subject and the distinctiveness between alive and inanimate entities. Feelings of uncanniness emerge as a consequence of perceiving something familiar that nonetheless points to or contains something extremely unfamiliar, or that cannot be explained with the human rational categories. Despite the images not being perceived as artistic or creative products, the “uncanny”, here, resembles to an effect that can be experienced, for example, when we are exposed to artworks belonging to Russian Formalism or Surrealism: these artistic movements introduced the concept of “estrangement” or “defamiliarization”, an artistic practice that produces representations that allow us to recognize their subject, but, at the same time, makes them seem unfamiliar ([Spiegel, 2008](#)). In the case of Gen-AI, therefore, the uncanny effect is more the byproduct of estrangement and defamiliarization, which point at its alien nature, than of intellectual uncertainty about the alive or inanimate nature of the entity.

This is relevant because it identifies an alternate way of “being uncanny”, specific of Gen-AI’s products, which refers to its likely inexplicable ways of perceiving and understanding that are embedded in its creations. To mitigate these uncanny sensations, therefore, it is likely not sufficient to increase the image realism and dispel the uncertainty about the alive or synthetic nature of the agent, as can happen with robots where overcoming the uncanny valley could be likely reached

through a higher degree of realism (Schwind et al., 2018). More likely, it should be needed to completely remove those details that may be experienced as unfamiliar (even imperceptible at first sight, or pertaining to the aesthetic quality of the image and not to what it represents), as well as explain how the technology truly works and why some of its products have their particular appearance.

In the absence of these solutions, to mitigate the feelings of uncanniness, participants are left with no choice but to resort to “relational strategies”, which ultimately refer to practices of humanization (and hence, familiarization).

5.2. Humanization, dehumanization and superhumanization of Gen-AI

HCI research on humanization has highlighted that people may humanize or dehumanize a seemingly intelligent agent in various ways depending on which human attributes are ascribed or denied (Haslam, 2006), and this may be made for utilitarian purposes (e.g., when the agent does not allow users to achieve their goals, they may degrade it to a lower entity) (Rapp et al., 2023a). In this study, we observed a similar phenomenon, whereby participants used dehumanizing or humanizing strategies to tame the feelings of uncanniness evoked by the unfamiliarity of the images produced by the Gen-AI. With this aim, they performed different “relational strategies”, ranging from devaluation or overvaluation of the Gen-AI to its “acceptance”, which implied the reworking of their previous perceptions.

Two of these strategies de facto entail practices of humanization, in either the form of dehumanization (Leyens et al., 2003) (denying Gen-AI valuable human characteristics like intelligence), or superhumanization (Li et al., 2014) (ascribing to Gen-AI certain “super abilities” like superintelligence or more-than-human memory, while denying others, like emotions): this can also be seen in the Italian pronouns “Lui” (i.e., he) or “Lei” (i.e., she), which in Italian are commonly only used for referring to human beings, which several participants employed to indicate the Gen-AI. In this context, humanization practices are essentially practices of familiarization, as retracing the Gen-AI to something that is less or more than human means, in any case, to anchor it to something familiar, namely, the human being.

This finding contributes to research on (de)humanization of intelligent technology, as it shows that people may (de)humanize an agent not only for fulfilling an objective arising in the ongoing interaction (e.g., Rapp et al., 2023a), but also for addressing an emotional need, and even when they do not directly interact with the agent: (super and de)humanization becomes a strategy for dealing with the sense of uncanniness evoked by technologies that are perceived as completely diverse from humans, enabling individuals to trace the unknown back to the known.

This also hints at an underlying fear that participants felt towards this kind of technology. Some of the participants expressed that more than having concerns about job loss and unemployment, or the spreading or misinformation (Oppenlaender et al., 2023; Bird et al., 2023), were worried about the possibility that such an alien and unfamiliar entity could become independent, also raising concerns about its trustworthiness. This is in line with research investigating people’s concerns on other intelligent technologies, like intelligent persuasive technology (Rapp, 2019, 2020), and recalls Bostrom’s (2014) reflection on superintelligence, whereby a completely alien intelligence could pursue objectives that only apparently align to our understandings, perceptions, and values, potentially leading to unexpected “side effects”.

6. Implications for design

Although the primary purpose of our analysis was not to generate design recommendations but to provide an empirical understanding of how people perceive Gen-AI creations as well as their creator, we may identify two design implications for HCI coming from our study.

6.1. Humanness as design material

Participants of this study were often unsettled by the Gen-AI images, as well as by the Gen-AI itself. To mitigate these feelings, they carried out several “relational” strategies. We have seen that some of these strategies can ultimately be viewed as practices of humanization, which people perform to trace the unfamiliar back to the familiar. Even though the participants did not prompt the image themselves, a sense of uncanniness towards Gen-AI could also likely emerge during the interaction, if Gen-AI’s creations will exhibit the same unfamiliar elements found in our images. In this context, designers could support (de)humanization practices, treating humanness as design material, to help users get rid of unpleasant sensations.

Previous research has noticed that current forms of text-to-image models provide extremely simple user interfaces, which ultimately resemble a command line or a search engine (Ko et al., 2023). At present, these interfaces differ from Large Language Models (LLMs) like ChatGPT, where the user can make the Gen-AI produce text by “conversing” with it (Zamfirescu-Pereira et al., 2023). LLMs can be further humanized by assigning it a persona (e.g., the boxer Muhammad Ali) (Deshpande et al., 2023a, 2023b), which could help users integrate Gen-AI agents into their self-concept (Alabed et al., 2022), potentially making them less alien.

Text-to-image models could thus adopt similar interaction modalities, whereby designers could experiment with different levels of humanness, making the model interact in humanized, superhumanized, or dehumanized ways, in order to reduce the unsettling feelings that its productions may elicit. Designers could focus on creating conversations where users “talk” to the model when writing a prompt and the model responds accordingly (e.g., through a command-style language if the model is dehumanized, a human level language if it is humanized, or even an unemotional language if the model is superhumanized), also addressing any potential questions that users may have about its creations. Experiments could also be conducted with “novel” forms of humanity, by endowing the model with a peculiar humanness that is neither humanlike nor machinelike. Possible side effects when opting for a specific humanization level, like deception if the model is humanized, should be carefully assessed. Additionally, users could be directly allowed to (super)humanize or dehumanize the model when interacting with it, making it behave as they like at the interaction style level.

6.2. Gen-AI images as design material

Images created through Gen-AI could be used in the design process as stimuli for generating ideas, prototyping, and creating scenarios (e.g., Ko et al., 2023; Inie et al., 2023; Kulkarni et al., 2023; Vimpari et al., 2023), as also highlighted by our participants with a design or technical background.

Several participants in our study hinted at a supplementary value of these images. They highlighted that some images may condense in a single representation the “essence” of a given subject, which may unveil stereotypes and biases, thus revealing the shared values on which our society relies, as well as those distortions that affect our imaginary. In this sense, these images could be fruitfully used in critical design scenarios, as stimuli for eliciting critical reflection.

Critical design (Dunne and Raby, 2001) precisely tries to disrupt and transgress needs and values as they are presently interpreted in the consumer society, by creating designs that embody alternative social, cultural, technical, and economic values. Within HCI, critical design has been used to make people reflect on the taken-for-granted embedded in everyday technology (Bardzell and Bardzell, 2013), subverting technological expectations (Cole, 2023), and, more in general, raising awareness and exposing assumptions (Blythe, 2015).

A critical design process commonly starts with identifying a stereotype, or the traditional framing of a practice or of a technology assumed within our society, which then will be challenged by design. In other

words, in traditional user research the research objective is to establish user needs and requirements and what is unknown are relevant behaviors, attitudes, and functional needs of a target population; instead, in critical design research, the framing itself is part of the unknown (Bardzell et al., 2012). In this perspective, Gen-AI images could be used to make visible the framing, the stereotypes and the biases characterizing practices and technologies, and be used as stimuli in the critical design process. Perhaps, this could make the critical process easier to engage with: critical design has been criticized for being an elitist mystery, like art, where only elitist designers can create effective critical designs (Bardzell and Bardzell, 2013; Iivari and Kuutti, 2017). By materializing a sort of collective unconscious, Gen-AI images, for example, could be utilized in design workshops aimed at involving students in critical design, or to conduct critical participatory design sessions, where users could be engaged in producing critical designs. Gen-AI images could serve as stimuli for critical reflection, raising awareness and grounding the subsequent critical design process. Moreover, as in critical designs “a slight strangeness is the key - too weird and they are instantly dismissed, not strange enough and they’re absorbed into everyday reality” (Dunne and Raby, 2001: 63), Gen-AI could directly inspire novel critical technological solutions with their strangeness.

7. Limitations

Our study has focused on a particular Gen-AI text-to-image model, which may limit the generalizability of our findings. Moreover, it is questionable that our findings would be generalizable across different cultures, being our study limited to the Italian population. The study did not make participants directly interact with Gen-AI, so we did not explore how people perceive the images that they have produced themselves through a Gen-AI model. As we have said, this choice was due to the fact that currently people are more likely to encounter a Gen-AI image than to interact with a text-to-image model. However, we can reasonably hypothesize that the findings of our study could also be applicable to certain images that people create themselves, since participants with previous experience with text-to-image models reported the same perceptions (e.g., of strangeness) as the other participants. Future research could investigate people’s perceptions of their own Gen-AI creations.

The sample was limited to 20 participants. However, we followed a purposeful sampling method which ensured the coverage of different “profiles” with different backgrounds useful to understand the phenomenon under examination. As most of our participants did not have any previous experience with Gen-AI, it is possible that their concerns and perceptions of unfamiliarity towards Gen-AI were more prominent than those of the general population, which is increasingly gaining experience with these tools. However, we did not find any considerable difference related to perceptions and feelings evoked by Gen-AI images between these participants and those with previous experience with Gen-AI.

It is possible that the participants overlooked the artistic quality of the images and the artistic and creative component of the generation process because of the frame of our study and the “default style” we used for generating the images. In the first phase, we asked them to tell us what the main subject of the image was for them: this could have focused their attention on whether the image faithfully represented a subject. However, we did not prevent participants from commenting on the images’ artistic features. Having generated the images from single terms also left ample room for Gen-AI to maneuver on how to represent the subject. Moreover, Stable Diffusion “default style” did not prevent the model from creating images with an artistic flair (see e.g., “Infinity”). Despite this, none of the participants appreciated the “creativity” of the Gen-AI model: when they could not identify a clear link with the original prompt, they found the image inexplicable. None of the participants, moreover, attempted to attribute the strangeness of some images’ details

to the model’s creativity. This may signal that people’s appraisals of Gen-AI images may go beyond artistic and creativity matters.

8. Conclusion

In this article we explored how people perceive the images produced by a Gen-AI text-to-image model, as well as how they experience the Gen-AI model itself starting from its outputs. We discovered that people appraise Gen-AI images on the basis of their technical quality and fidelity to the subject, perceiving them as either prototypical or strange (RQ1), often evoking eerie feelings that may extend to the Gen-AI itself (RQ2). We further found that to mitigate these feelings participants may carry out different “relational” strategies, devaluating or overvaluing themselves or the Gen-AI.

Our contribution to HCI research points out that people’s experience of Gen-AI images may evoke unsettling feelings and reflections that point at the alien nature of the technology. This nature is not easily addressed by individuals, but makes it emerge people’s pressing need to make sense of the unknown and to trace it back to the known. The “relational” strategies that we identified ultimately refer to practices of humanization aimed at making the unfamiliar familiar and open new research lines and questions on how people relate to “intelligent” technology. For example, how can technology mitigate the sense of bewilderment provoked by its productions? What are the risks of humanizing technology that creates products similar to those created by humans? In this sense, interacting with Gen-AI gives rise to novel interactional issues, such as the fear of the unknown, as well as design questions, revolving around the (de or super)humanization of technology: for instance, should we favor the ascription of humanness to these technologies? Or should we find ways for making people constantly aware that the technology is a machine? These are aspects that we started exploring in this study and are certainly worth being further investigated in future research.

CRedit authorship contribution statement

Amon Rapp: Writing – review & editing, Writing – original draft, Supervision, Project administration, Methodology, Formal analysis, Conceptualization. **Chiara Di Lodovico:** Methodology, Investigation, Formal analysis, Data curation, Conceptualization. **Federico Torrielli:** Software, Methodology. **Luigi Di Caro:** Writing – review & editing, Resources, Methodology.

Declaration of competing interest

The authors declare that they have no known competing financial interests or personal relationships that could have appeared to influence the work reported in this paper.

Data availability

The data that has been used is confidential.

References

- Alabed, A., Javornik, A., Gregory-Smith, D., 2022. AI anthropomorphism and its effect on users’ self-congruence and self-ai integration: a theoretical framework and research agenda. *Technol. Forecast. Soc. Change* 182 (121786), 1–19. <https://doi.org/10.1016/j.techfore.2022.121786>.
- Bandi, A., Adapa, P.V.S.R., Kuchi, Y.E.V.P.K., 2023. The power of generative AI: a review of requirements, models, input–output formats, evaluation metrics, and challenges. *Fut. Internet* 15 (260), 60. <https://doi.org/10.3390/fi15080260>.
- Bardzell, J., Bardzell, S., 2013. What is “Critical” about Critical Design?. In: *Proceedings of the SIGCHI Conference on Human Factors in Computing Systems (CHI ’13)*. ACM, New York, NY, USA, pp. 3297–3306. <https://doi.org/10.1145/2470654.2466451>.
- Bardzell, S., Bardzell, J., Forlizzi, J., Zimmerman, J., Antanitis, J., 2012. Critical design and critical theory: the challenge of designing for provocation. In: *Proceedings of the*

- Designing Interactive Systems Conference (DIS '12). ACM, New York, NY, USA, pp. 288–297. <https://doi.org/10.1145/2317956.2318001>.
- Bird, C., Ungless, E., Kasirzadeh, A., 2023. Typology of risks of generative text-to-image models. In: Proceedings of the 2023 AAAI/ACM Conference on AI, Ethics, and Society, pp. 396–410. <https://doi.org/10.1145/3600211.3604722>.
- Blythe, M., Yauner, F., Rodgers, P., 2015. The context of critical design: exhibits, social media and auction houses. *Des. J.* 18 (1), 83–105. <https://doi.org/10.2752/175630615x14135446523305>.
- Bostrom, N., 2014. *Superintelligence: Paths, Dangers, Strategies*. Oxford University Press, Oxford UK.
- Bowen, G.A., 2008. Naturalistic inquiry and the saturation concept: a research note. *Qual. Res.* 8 (1), 137–152. <https://doi.org/10.1177/1468794107085301>.
- Braun, V., Clarke, V., 2006. Using thematic analysis in psychology. *Qual. Res. Psychol.* 3 (2), 77–101. <https://doi.org/10.1191/1478088706qp0630a>.
- Brown, B., Clarke, V., 2013. *Successful Qualitative Research: A Practical Guide for Beginners*. SAGE Publication Ltd., London.
- Brysaert, M., Warriner, A.B., Kuperman, V., 2014. Concreteness ratings for 40 thousand generally known English word lemmas. *Behav. Res. Methods* 46, 904–911. <https://doi.org/10.3758/s13428-013-0403-5>.
- Buschek, D., Mecke, L., Lehmann, F., Dang, H., 2021. Nine potential pitfalls when designing human-AI Co-creative systems. In: Workshops at the International Conference on Intelligent User Interfaces (IUI) (2021). <https://doi.org/10.48550/arXiv.2104.00358>.
- Candello, H., Pinhanez, C., Figueiredo, F., 2017. Typefaces and the perception of humanness in natural language chatbots. In: Proceedings of the SIGCHI Conference on Human Factors in Computing Systems (CHI '17). ACM, New York, NY, USA, pp. 3476–3487. <https://doi.org/10.1145/3025453.3025919>.
- Carlson, M., 2022. Hackaday, “Stable diffusion and why it matters” <https://hackaday.com/2022/09/06/stable-diffusion-and-why-it-matters/> (accessed 27 January 2024).
- Chamberlain, R., Mullin, C., Scheerlinck, B., Wagemans, J., 2018. Putting the art in artificial: aesthetic responses to computer-generated art. *Psychol. Aesthet. Creat. Arts* 12 (2), 177–192. <https://doi.org/10.1037/aca0000136>.
- Chang, M., Druga, S., Fiannaca, A.J., Vergani, P., Kulkarni, C., Cai, C.J., Terry, M., 2023. The prompt artists. In: Proceedings of the 15th Conference on Creativity and Cognition (C&C '23). ACM, New York, NY, USA, pp. 75–87. <https://doi.org/10.1145/3591196.3593515>.
- Chung, J.J.Y., He, S., Adar, E., 2021a. Gestural inputs as control interaction for generative human-AI Co-creation. In: Workshops at the International Conference on Intelligent User Interfaces (IUI '21).
- Chung, J.J.Y., He, S., Adar, E., 2021b. The intersection of users, roles, interactions, and technologies in creativity support tools. In: Proceedings of the Designing Interactive Systems Conference (DIS '21). ACM, New York, NY, USA, pp. 1817–1833. <https://doi.org/10.1145/3461778.3462050>.
- Ciechanowski, L., Przegalinska, A., Wegner, K., 2018. The necessity of new paradigms in measuring human-chatbot interaction. Hoffman, M. In: *Advances in Cross-Cultural Decision Making*. AHFE 2017. Advances in Intelligent Systems and Computing, 610. Springer, Cham. https://doi.org/10.1007/978-3-319-60747-4_19.
- Codewatchers, 2023. A comprehensive overview on midjourney statistics. <https://codewatchers.com/en/blog/a-comprehensive-overview-on-midjourney-statistics> (accessed 27 January 2023).
- Cole, A., 2023. Rejected by my own robot: studying the potential for artists to subvert technological expectations using critical design. In: Proceedings of the 17th International Conference on Tangible, Embedded, and Embodied Interaction (TEI '23). New York, NY, USA. ACM, pp. 1–4. <https://doi.org/10.1145/3569009.3576181>. Article 56.
- Croitoru, F.A., Hondru, V., Ionescu, R.T., Shah, M., 2023. Diffusion models in vision: a survey. *IEEE Trans. Pattern. Anal. Mach. Intell.* 45 (9), 10850–10869. <https://doi.org/10.1109/TPAMI.2023.3261988>.
- Davis, N., Hsiao, C.P., Yashraj Singh, K., Li, L., Magerko, B., 2016. Empirically studying participatory sense-making in abstract drawing with a co-creative cognitive agent. In: Proceedings of the 21st International Conference on Intelligent User Interfaces (IUI '16). ACM, New York, NY, USA, pp. 196–207. <https://doi.org/10.1145/2856767.2856795>.
- Deng, W.H., Guo, B., Devrio, A., Shen, H., Eslami, M., Holstein, K., 2023. Understanding practices, challenges, and opportunities for user-engaged algorithm auditing in industry practice. In: Proceedings of the SIGCHI Conference on Human Factors in Computing Systems (CHI '23). ACM, New York, NY, USA, pp. 1–18. <https://doi.org/10.1145/3544548.3581026>. Article 377.
- Deshpande, A., Murahari, V., Rajpurohit, T., Kalyan, A., Narasimhan, K., 2023a. Toxicity in ChatGPT: analyzing persona-assigned language models. Findings of the Association for Computational Linguistics: EMNLP 2023. Association for Computational Linguistics, Singapore, pp. 1236–1270. <https://doi.org/10.48550/arXiv.2304.05335>.
- Deshpande, A., Rajpurohit, T., Narasimhan, K., Kalyan, A., 2023b. Anthropomorphization of AI: opportunities and Risks. In: Proceedings of the Natural Legal Language Processing Workshop 2023. Association for Computational Linguistics, Singapore, pp. 1–7. <https://doi.org/10.48550/arXiv.2305.14784>.
- Doyle, P.R., Edwards, J., Dumbleton, O., Clark, L., Cowan, B.R., 2019. Mapping perceptions of humanness in intelligent personal assistant interaction. In: Proceedings of the 21st International Conference on Human-Computer Interaction with Mobile Devices and Services (MobileHCI '19), 5. ACM, New York, pp. 1–12. <https://doi.org/10.1145/3338286.3340116>.
- Dunne, A., Raby, F., 2001. *Design noir: The secret Life of Electronic Objects*. August/Birkhäuser, London, UK.
- Festerling, J., Siraj, I., 2022. Anthropomorphizing technology: a conceptual review of anthropomorphism research and how it relates to children's engagements with digital voice assistants. *Integr. Psychol. Behav. Sci.* 56 (3), 709–738. <https://doi.org/10.1007/s12124-021-09668-y>.
- Følstad, A., Larsen, A.G., Bjerkreim-Hanssen, N., et al., 2023. The human likeness of government chatbots – An empirical study from Norwegian municipalities. In: Lindgren, I., et al. (Eds.), *Electronic Government. EGOV 2023. Lecture Notes in Computer Science*. Springer, Cham, p. 14130. https://doi.org/10.1007/978-3-031-41138-0_8.
- Freud, S., 1919. The uncanny. First published in *Imago*, Bd. V., 1919; reprinted in *Sammlung, Fünfte Folge*.
- Go, E., Sundar, S.S., 2019. Humanizing chatbots: the effects of visual, identity and conversational cues on humanness perceptions. *Comput. Human. Behav.* 97, 304–316. <https://doi.org/10.1016/j.chb.2019.01.020>.
- Grabe, I., Duque, M.G., Zhu, J., 2022. Towards a framework for human-ai interaction patterns in co-creative GAN applications. In: Proceeding of the 3rd Workshop on Human-AI Co-Creation with Generative Models (HAI-GEN '22) at ACM IUI Workshops.
- Harper, D., 2002. Talking about pictures: a case for photo elicitation. *Vis. Stud.* 17 (1), 13–26. <https://doi.org/10.1080/14725860220137345>.
- Harry, B., Sturges, K.M., Klingner, J.K., 2005. Mapping the process: an exemplar of process and challenge in grounded theory analysis. *Educ. Res.* 34 (2), 3–13. <https://doi.org/10.3102/0013189x034002003>.
- Haslam, N., 2006. Dehumanization: an integrative review. *Personal. Soc. Psychol. Rev.: Off. J. Soc. Personal. Soc. Psychol.* 10 (3), 252–264. https://doi.org/10.1207/s15327957pspr1003_4.
- Haslam, N., Loughnan, S., 2014. Dehumanization and inhumanization. *Annu. Rev. Psychol.* 65, 399–423. <https://doi.org/10.1146/annurev-psych-010213-115045>.
- Ho, J., Salimans, T., 2021. Classifier-free diffusion guidance. *NeurIPS 2021 Workshop on Deep Generative Models and Downstream Applications*.
- Ho, J., Jain, A., Abbeel, P., 2020. Denoising diffusion probabilistic models. *Adv. Neural Inf. Process. Syst.* 33, 6840–6851, 2020.
- Hodhod, R., Magerko, B., 2016. Closing the cognitive gap between humans and interactive narrative agents using shared mental models. In: Proceedings of the 21st International Conference on Intelligent User Interfaces (IUI '16). ACM, New York, NY, USA, pp. 135–146. <https://doi.org/10.1145/2856767.2856774>.
- Hogan, S., 2012. Ways in which photographic and other images are used in research: an introductory overview. *Int. J. Art Therapy: Formerly Inscape* 17 (2), 54–62. <https://doi.org/10.1080/107454832.2012.699533>.
- Hong, J.-W., Curran, N.M., 2019. Artificial intelligence, artists, and art: attitudes toward artwork produced by humans vs. artificial intelligence. *ACM Trans. Multimedia Comput. Commun., Appl.* 15, 2s, Article 58, 16. <https://doi.org/10.1145/3326337>.
- Huang, K., 2023. Why Pope Francis Is the Star of AI-Generated Photos. *The New York Times*. <https://www.nytimes.com/2023/04/08/technology/ai-photos-pope-francis.html>. accessed 27 October 2023.
- Iivari, N., Kuutti, K., 2017. Critical design research and information technology: searching for empowering design. In: Proceedings of the 2017 Conference on Designing Interactive Systems (DIS '17). ACM, New York, NY, USA, pp. 983–993. <https://doi.org/10.1145/3064663.3064747>.
- Inie, N., Falk, J., Tanimoto, S., 2023. Designing participatory ai: creative professionals' worries and expectations about generative AI. In: Extended Abstracts of the 2023 CHI Conference on Human Factors in Computing Systems (CHI EA '23). ACM, New York, NY, USA, pp. 1–8. <https://doi.org/10.1145/3544549.3585657>. Article 82.
- Jun, E., Jo, B.A., Oliveira, N., Reinecke, K., 2018. Digestif: promoting science communication in online experiments. In: Proceedings of the ACM on Human-Computer Interaction 2, 84. CSCW, p. 26. <https://doi.org/10.1145/3274353>, 1–84.
- Kim, S.S., Watkins, E.A., Russakovsky, O., Fong, R., Monroy-Hernández, A., 2023. Help Me Help the AI: understanding how explainability can support human-AI interaction. In: Proceedings of the SIGCHI Conference on Human Factors in Computing Systems (CHI '23). ACM, New York, NY, USA, pp. 1–17.
- Ko, H.K., Park, G., Jeon, H., Jo, J., Kim, J., Seo, J., 2023. Large-scale text-to-image generation models for visual artists' creative works. In: Proceedings of the 28th International Conference on Intelligent User Interfaces (IUI '23). ACM, New York, NY, USA, pp. 919–933. <https://doi.org/10.1145/3581641.3584078>.
- Kulkarni, C., Druga, S., Chang, M., Fiannaca, A., Cai, C., Terry, M., 2023. *arXiv preprint*.
- Kunimoto, N., 2004. Intimate archives: Japanese-Canadian family photography 1939–49. *Art Hist.* 27, 129–155. <https://doi.org/10.1111/j.0141-6790.2004.02701005.x>.
- Law, E.L.C., Følstad, A., Van As, N., 2022. Effects of humanlikeness and conversational breakdown on trust in chatbots for customer service. In: *Nordic Human-Computer Interaction Conference (NordiCHI '22)*. ACM, New York, NY, USA, pp. 1–13. <https://doi.org/10.1145/3546155.3546665>. Article 56.
- Lee, M., Ackermans, S., van As, N., Chang, H., Lucas, E., IJsselstein, W., 2019. Caring for Vincent: a chatbot for self-compassion. In: Proceedings of the SIGCHI Conference on Human Factors in Computing Systems (CHI '19). ACM, New York, NY, USA, pp. 1–13. <https://doi.org/10.1145/3290605.3300932>.
- Lee, S., Hoover, B., Strobelt, H., Wang, Z.J., Peng, S., Wright, A., Li, K., Park, H., Yang, H., Chau, D.H., 2023. *arXiv preprint*.
- Leyens, J.-P., Cortes, B., Demoulin, S., Dovidio, J.F., Fiske, S.T., Gaunt, R., Paladino, M.-P., Rodriguez-Perez, A., Rodriguez-Torres, R., Vaes, J., 2003. Emotional prejudice, essentialism, and nationalism. *Eur. J. Soc. Psychol.* 33 (6), 703–717. <https://doi.org/10.1002/ejsp.170>.
- Li, M., Leidner, B., Castano, E., 2014. Toward a comprehensive taxonomy of dehumanization: integrating two senses of humanness, mind perception theory, and stereotype content model. *TPMTesting, Psychometrics. Methodol. Appl. Psychol.* 21 (3), 285–300. <https://doi.org/10.4473/TPM21.3.4>.

- Liu, 2022. The world's smartest artificial intelligence just made its first magazine cover. *Cosmopolitan*. <https://www.cosmopolitan.com/lifestyle/a40314356/dall-e-2-a-rificial-intelligence-cover/>. accessed 27 October 2023.
- Liu, B., Sundar, S.S., 2018. Should machines express sympathy and empathy? Experiments with a health advice chatbot. *Cyberpsychol., Behav. Soc. Network.* 21 (10), 625–636. <https://doi.org/10.1089/cyber.2018.0110>.
- Liu, V., Chilton, L.B., 2022. Design guidelines for prompt engineering text-to-image generative models. In: *Proceedings of the 2022 CHI Conference on Human Factors in Computing Systems (CHI '22)*. ACM, New York, NY, USA, pp. 1–23. <https://doi.org/10.1145/3491102.3501825>. Article 384.
- Liu, V., Vermeulen, J., Fitzmaurice, G., Matejka, J., 2023. 3DALL-E: integrating text-to-image AI in 3D design workflows. In: *Proceedings of the Designing Interactive Systems Conference (DIS '23)*. ACM, New York, NY, USA, pp. 1955–1977. <https://doi.org/10.1145/3563657.3596098>.
- Louie, R., Coenen, A., Huang, C.Z., Terry, M., Cai, C.J., 2020. Novice-AI music co-creation via ai-steering tools for deep generative models. In: *Proceedings of the SIGCHI Conference on Human Factors in Computing Systems (CHI '20)*. ACM, New York, NY, USA, pp. 1–13. <https://doi.org/10.1145/3313831.3376739>.
- Lyu, Y., Wang, X., Lin, R., Wu, J., 2022. Communication in human-AI co-creation: perceptual analysis of paintings generated by text-to-image system. *Appl. Sci.* 12, 22. <https://doi.org/10.3390/app122211312>.
- MacQueen, K.M., McLellan, E., Kay, K., Milstein, B., 1998. Team-based codebook development: structure, process, and agreement. G. Guest, and K. M. MacQueen *Handbook For Team-Based Qualitative Research*. AltaMira Press, Lanham, UK, pp. 119–136.
- Marshall, M.N., 1996. Sampling for qualitative research. *Fam. Pract.* 13 (6), 522–525. <https://doi.org/10.1093/fampra/13.6.522>.
- McDonald, N., Schoenebeck, S., Forte, A., 2019. Reliability and inter-rater reliability in qualitative research: norms and guidelines for CSCW and HCI practice. *Proc. ACM Hum.-Comput. Interact.* 3, 1–23. <https://doi.org/10.1145/3359174>. CSCW.
- Mikalonytė, E.S., Kneer, M., 2022. Can artificial intelligence make art?: Folk intuitions as to whether AI-driven robots can be viewed as artists and produce art. *J. Hum.-Robot. Interact.* 11 (4), 19. <https://doi.org/10.1145/3530875>. Article 43.
- Mori, M., 1970. Bukimi no tani - The uncanny valley (K. F. MacDorman and T. Minato. Trans.). *Energy* 7 (4), 33–35.
- Mori, M., MacDorman, K.F., Kageki, N., 2012. The uncanny valley [from the field]. *IEEE Robot. Autom. Mag.* 19 (2), 98–100. <https://doi.org/10.1109/MRA.2012.2192811>.
- Muller, M., Weisz, J.D., Geyer, W., 2020. Mixed initiative generative AI interfaces: an analytic framework for generative AI applications. In: *Proceedings of the Workshop "The Future of Co-Creative Systems - A Workshop on Human-Computer Co-Creativity" of the 11th International Conference on Computational Creativity (ICCC 2020)*.
- Muller, M., Chilton, L.B., Kantosalo, A., Martin, C.P., Walsh, G., 2022. GenAICHI: generative AI and HCI. In: *Extended Abstracts of the 2022 CHI Conference on Human Factors in Computing Systems (CHI EA '22)*. ACM, New York, NY, USA, pp. 1–7. <https://doi.org/10.1145/3491101.3503719>. Article 110.
- OpenAI, 2022. DALL-E now available without waitlist. <https://openai.com/blog/dall-e-now-available-without-waitlist> (accessed 27 January 2024).
- Oppenlaender, J., 2022. The creativity of text-to-image generation. In: *Proceedings of the 25th International Academic Mindtrek Conference (Academic Mindtrek '22)*. ACM, New York, NY, USA, pp. 192–202. <https://doi.org/10.1145/3569219.3569352>.
- Oppenlaender, J., Silvennoinen, J., Paananen, V., Visuri, A., 2023. Perceptions and realities of text-to-image generation. In: *Proceedings of the 26th International Academic Mindtrek Conference (Academic Mindtrek '23)*. ACM, New York, NY, USA, pp. 279–288. <https://doi.org/10.1145/3616961.3616978>.
- Patton, M.Q., 1990. *Qualitative Evaluation and Research Methods*. SAGE Publications Inc.
- Ragot, M., Martin, N., Cojean, S., 2020. AI-generated vs. Human artworks. A perception bias towards artificial intelligence?. In: *Extended Abstracts of the 2020 CHI Conference on Human Factors in Computing Systems (CHI EA '20)*. ACM, New York, NY, USA, pp. 1–10. <https://doi.org/10.1145/3334480.3382892>.
- Ramesh, A., Pavlov, M., Goh, G., Gray, S., Voss, C., Radford, A., Chen, M., Sutskever, I., 2021. Zero-shot text-to-image generation. In: *International Conference on Machine Learning*. PMLR, pp. 8821–8831.
- Rapp, A., 2019. Design fictions for behaviour change: exploring the long-term impacts of technology through the creation of fictional future prototypes. *Behav. Inf. Technol.* 38 (3), 244–272. <https://doi.org/10.1080/0144929X.2018.1526970>.
- Rapp, A., 2020. Design fictions for learning: a method for supporting students in reflecting on technology in Human-Computer Interaction courses. *Comput. Educ.* 145, 103725. <https://doi.org/10.1016/j.compedu.2019.103725>.
- Rapp, A., Curti, L., Boldi, A., 2021. The human side of human-chatbot interaction: a systematic literature review of ten years of research on text-based chatbots. *Int. J. Hum. Comput. Stud.* 151, 102630. <https://doi.org/10.1016/j.ijhcs.2021>.
- Rapp, A., Boldi, A., Curti, L., Perrucci, A., Simeoni, R., 2023a. How do people ascribe humanness to chatbots? An analysis of real-world human-agent interactions and a theoretical model of humanness. *Int. J. Hum.-Comput. Interact.* 24. <https://doi.org/10.1080/10447318.2023.2247596>.
- Rapp, A., Boldi, A., Curti, L., Perrucci, A., Simeoni, R., 2023b. Collaborating with a text-based chatbot: an exploration of real-world collaboration strategies enacted during human-chatbot interactions. In: *Proceedings of the 2023 CHI Conference on Human Factors in Computing Systems (CHI '23)*, 115. Association for Computing Machinery, New York, NY, USA, Article, pp. 1–17. <https://doi.org/10.1145/3544548.3580995>.
- Rombach, R., Blattmann, A., Lorenz, D., Esser, P., Ommer, B., 2022. High-resolution image synthesis with latent diffusion models. In: *Proceedings of the IEEE/CVF conference on computer vision and pattern recognition*, pp. 10684–10695.
- Roose, K., 2022. An A.I.-Generated Picture Won an Art Prize. Artists Aren't Happy. *The New York Times*. <https://www.nytimes.com/2022/09/02/technology/ai-artificial-intelligence-artists.html>. accessed 27 January 2023.
- Sadek, M., Calvo, R.A., Mougnot, C., 2023. Trends, challenges and processes in conversational agent design: exploring practitioners' views through semi-structured interviews. In: *Proceedings of the 5th International Conference on Conversational User Interfaces (CUI '23)*. ACM, New York, NY, USA, pp. 1–10. <https://doi.org/10.1145/3571884.3597143>. Article 13.
- Samo, A., Highhouse, S., 2023. Artificial intelligence and art: identifying the aesthetic judgment factors that distinguish human- and machine-generated artwork. *Psychol. Aesthet. Creat. Arts*. <https://doi.org/10.1037/aca0000570>. Advance online publication.
- Saunders, B., Sim, J., Kingstone, T., Baker, S., Waterfield, J., Bartlam, B., Burroughs, H., Jinks, C., 2018. Saturation in qualitative research: exploring its conceptualization and operationalization. *Qual. Quant.* 52, 1893–1907. <https://doi.org/10.1007/s11335-017-0574-8>.
- Schwind, V., Wolf, K., Henze, N., 2018. Avoiding the uncanny valley in virtual character design. *Interactions* 25 (5), 45–49. <https://doi.org/10.1145/3236673>.
- Spiegel, S., 2008. Things made strange: on the concept of 'estrangement' in science fiction theory. *Sci. Fiction Stud.* 35th 106, 369–385.
- Stein, J.P., Ohler, P., 2017. Venturing into the uncanny valley of mind—The influence of mind attribution on the acceptance of human-like characters in a virtual reality setting. *Cognition* 160, 43–50. <https://doi.org/10.1016/j.cognition.2016.12.010>.
- Struppek, L., Hintersdorf, D., Friedrich, F., Brack, M., Schramowski, P., Kersting, K., 2023. Exploiting cultural biases via homographs in text-to-image synthesis. *J. Artif. Intell.* 78 (December 2023), 1017–1068. <https://doi.org/10.1613/jair.1.14714>.
- Sun, Y., Yang, C.-H., Lyu, Y., Lin, R., 2022. From pigments to pixels: a comparison of human and AI painting. *Appl. Sci.* 12, 3724. <https://doi.org/10.3390/app12083724>.
- Ta, V., Griffith, C., Boatfield, C., Wang, X., Civitello, M., Bader, H., Decero, E., Loggarakis, A., 2020. User experiences of social support from companion chatbots in everyday contexts: thematic analysis. *J. Med. Internet Res.* 22 (3) <https://doi.org/10.2196/16235>.
- Tang, R., Liu, L., Pandey, A., Jiang, Z., Yang, G., Kumar, K., Stenertorp, P., Lin, J., Ture, F., 2023. What the DAAM: interpreting stable diffusion using cross attention. In: *Proceedings of the 61st Annual Meeting of the Association for Computational Linguistics (Volume 1: Long Papers)*. Association for Computational Linguistics, pp. 5644–5659.
- Ting, T.T., Ling, L.Y., bin Ahmad Azam, A.I., Palaniappan, R., 2023. Artificial intelligence art: attitudes and perceptions toward human versus artificial intelligence artworks. In: *1ST Int. Conf. Humanit. Educ. Sci. Manag. Eng. Technol. AIP Conf. Proc.* 2823, 020003. <https://doi.org/10.1063/5.0162434>.
- Vimpari, V., Kultima, A., Hämäläinen, P., Guckelsberger, C., 2023. An Adapt-or-Die Type of Situation": perception, adoption, and use of text-to-image-generation AI by game industry professionals. In: *Proceedings of ACM Human-Computer Interaction*. 7, CHI PLAY, Article 379 (November 2023), p. 34. <https://doi.org/10.1145/3611025>.
- Wang, Y., Shen, S., Lim, B.Y., 2023a. RePrompt: automatic prompt editing to refine AI-generative art towards precise expressions. In: *Proceedings of the 2023 CHI Conference on Human Factors in Computing Systems (CHI '23)*. Association for Computing Machinery, New York, NY, USA, pp. 1–29. <https://doi.org/10.1145/3544548.3581402>. Article 22.
- Wang, Z., Lu, C., Wang, Y., Bao, F., Li, C., Su, H., Zhu, J., 2023b. arXiv preprint.
- Yang, L., Neustaedter, C., 2018. Our house: living long distance with a telepresence robot. In: *Proceedings of the ACM on Human-Computer Interaction* 2, CSCW, Article 190 (November 2018), 18. <https://doi.org/10.1145/3274459>.
- Yardley, L., 2000. Dilemmas in qualitative health research. *Psychol. Health* 15 (2), 215–228. <https://doi.org/10.1080/08870440008400302>. March 2000.
- Zhang, Y., Huang, N., Tang, F., Huang, H., Ma, C., Dong, W., Xu, C., 2023. Inversion-based style transfer with diffusion models. In: *Proceedings of the IEEE/CVF conference on computer vision and pattern recognition*, pp. 10146–10156.
- Zamfirescu-Pereira, J.D., Wong, R.Y., Hartmann, B., Yang, Q., 2023. Why Johnny can't prompt: how non-AI experts try (and fail) to design LLM prompts. In: *Proceedings of the 2023 CHI Conference on Human Factors in Computing Systems (CHI '23)*. ACM, New York, NY, pp. 1–21. Article 437.