Technology-based Mediation in the Art Museum

PhD Thesis

Christian Sivertsen
Digital Design Department
IT University of Copenhagen
2023

Supervisors

Anders Sundnes Løvlie
IT University of Copenhagen, Denmark

Nikita Mathias
MUNCH, Oslo, Norway
University of the Free State, Bloemfontein, South Africa

The layout of this thesis is based on the LaTeX PhD Thesis template v.1.1.0 by Jesper Kjær Nielsen distributed under the terms of the GNU General Public License. The text is set with Libre Baskerville and League Spartan.

Abstract

The purpose of public art museums is to collect and exhibit art to benefit society, but what if people are bored or outright intimidated by the prospect of having to visit the art museum? Art challenges us to *see it if we can*, but we cannot always do so. Art museums do their best to provide *interpretive hooks* for visitors. In Human-Computer Interaction (HCI), researchers have attempted to support these efforts as well. However, old museum paradigms linger, making the art museum a conceptually contested space. This also affects HCI projects carried out in art museums, as they have to navigate or circumvent these conceptual challenges. In this PhD project, I present a theoretical framework that incorporates an enactivist understanding of *art as experience* and the work of art mediation as *education of attention*. Using these concepts, I offer a visitor-centered view on the art museum experience, that gives researchers, designers, curators and mediators tools for understanding the role technology may play in an art museum exhibition and how to conceptualize, design and evaluate such designs.

This perspective is illustrated and explored through three major research activities, of which the last two were in close collaboration with the MUNCH museum in Oslo. The activities involved an experiment, an exhibition and an interactive drawing table respectively, and each investigates aspects of how technology mediates our relation to artworks.

With the results of these three activities, I argue that technological designs can support the interpretive practice of museum visitors by *educating their attention* to aspects of the art that they would otherwise fail to see or give weight to. Purposefully designed technology can afford experiences that engage the senses and the whole human in ways that are exciting for museum visitors, while still establishing context and stimulating curiosity in the original artworks. For each of the three research activities, I provide analysis of how the designs concretely mediate the relation between visitor and art. To be able to design for this, and to evaluate whether a design affords *correspondence with* the art in the intended way, I operationalize the concept of *education of attention* as a way to analyze qualitative interview data. Finally, I discuss particular mediating properties of generative AI in relation to its deployment in art museums.

Resumé

Formålet med offentlige kunstmuseer er at samle, bevare og udstille kunst til nytte og glæde for samfundet, men hvad hvis folk holder sig væk af kedsomhed eller frygt for at virke dumme? Kunsten udfordrer os til at *se den, hvis vi kan*, men det kan vi ikke altid. Kunstmuseerne gør sig store anstrengelser for at give publikum hjælp til at få mening ud af den udstillede kunst. I Human-Computer Interaction (HCI) miljøet har forskere også forsøgt at bidrage til en stærkere kunstformidling, men gamle paradigmer for hvordan man tænker museumsformidling lurer dog stadig, både i HCI og hos museerne. Dette gør kunstmuseer til steder, hvor nye og gamle tilgange udfordrer hinanden. Dette udfordrer også HCI-projekter i kunstmuseer, da de bliver nødt til at forholde sig til disse konceptuelle udfordringer. I dette PhD-projekt præsenterer jeg et teoretisk blik, der tager udgangspunkt i *kunst som oplevelse* og kunstformidlingen som "education of attention". Med disse konceptuelle værktøjer tager jeg et brugercentreret blik på oplevelsen i kunstmuseet og giver forskere, designere, kuratorer og formidlere værktøjer til at forstå, hvilken rolle teknologi kan spille i en kunstudstilling samt, hvordan man kan konceptualisere, designe og evaluere denne type designs.

Dette perspektiv udfolder jeg gennem tre overordnede forskningsaktiviter, hvoraf de to sidste er udført i et tæt samarbejdet med museet MUNCH i Oslo. Aktiviteterne tæller et eksperiment, en udstilling og et interaktivt tegnebord. Hver især har jeg brugt dem til at undersøge forskellige aspekter af, hvordan teknologi medierer vores forhold til kunstværker.

Med resultaterne fra de tre forskningsaktiviteter argumenterer jeg for at teknologibaserede designs kan understøtte de museumsbesøgendes fortolkningspraksis ved at lede deres opmærksom til aspekter af kunsten, som de ellers ville overse eller undlade at tillægge betydning. Teknologi, som er designet med henblik på dette, kan skabe oplevelser, som involverer sanserne og det hele menneske på måder som både er spændende for de besøgende, men også skaber kontekst og nysgerrighed omkring originalværker. For hver af de tre forskningsaktiviteter præsenterer jeg en analyse af, hvordan hvert design konkret medierer relationen mellem museumsbesøgende og kunsten. Som en rettesnor til at kunne skabe designs der indgår i korrespodence med kunsten på den rigtige måde, bruger jeg begrebet "education of attention" til at analysere kvalitativ interviewdata indsamlet om hvert design. Slutteligt diskuterer jeg hvordan generativ kunstig intelligens som en specifik teknologi kan anvendes til mediering af kunst i kunstmuseer.

Contents

Acknowledgments

Throughout the three years it has taken to complete this PhD project, I have been helped by many good people.

First, I need to thank my supervisors, Anders Sundnes Løvlie and Nikita Mathias for giving me the opportunity to start this project in the first place and for believing in my ideas and visions.

Anders for his always pragmatic and patient advice, thorough manuscript revisions, and for introducing me to the ways of academia with all its idiosyncrasies. Nikita for following along on a ride through computers, projector, sensors, AI, and much more and believing that it could eventually turn into something meaningful for MUNCH and the visitors. Also thank you for opening the door into the world of museums, art history and Edvard Munch.

I am grateful to everyone I worked with during my stay at MUNCH for being easy-going and supportive with realizing *Poison* in the middle of all the million other things that had to be ready for the opening of the new MUNCH. In particular a big thanks to Gisle Sandvand, Birgitte Aga, and Dina Patey. I also feel extremely lucky, that while COVID-19 lockdowns meant that I was stuck in Norway for many months, I at least had ended up with the most lovely bunch of housemates at "Hansern".

I am forever grateful to Sander van der Zwan and Maarten Smith for helping me get on the right track theoretically and helping me from my application to the publication of "Art Critique by other Means" to make sense of Dewey, Ingold, Noë, and all the others.

Throughout the project, many have also helped out in very concrete ways in realizing the designs of this project. Thank you to Random International, Torsteinsen Arkitekter, Martin Horntveth, René Haas and Halfdan Hauch Jensen.

I have had the pleasure of working together with many more lovely people at Digital Design at ITU who've provided help, constructive criticism, and a generally good time over beers in Scrollbar.

Finally, I am very grateful to my dear Sofie for her love and support throughout all three years, and for putting up with me going to Norway and getting stuck for five months.

Christian Sivertsen
IT University of Copenhagen, October 14, 2023

I Thesis

1 Introduction

Great art can be one of the great pleasures of life. However, it can also be mind-numbingly boring to read obscure poetry, sit through a meaningless play or stare at arbitrary splotches of paint. Art challenges us to "See me if you can!" and we cannot always see it (Noë, 2015). This is one of the predicaments of public art museums. Their purpose is to collect and exhibit art to benefit society, but what if people are bored or outright intimidated by the prospect of having to visit the art museum?

The museums are aware of this challenge, and many do their best to provide *interpretive hooks* (Samis & Michaelson, 2016) that the visitors can use to get their bearings with the art on display. The purpose of this art mediation is to support the interpretive practices of visitors in relation to the art on display. Throughout the last centuries, the understanding of what is important in this regard, and who gets to decide it has been shifting. The assumed source of this understanding has moved from the object, to the museum, to the visitor, and the means through which to acquire it from the body, to the mind and back to the body again. However, at present, multiple paradigms may co-exist in museums, related to different departments and functions.

When designers and researchers of technology work to support art mediation in art museums, they too have to align with dominant paradigms in both their research field and the site of their work. Different designs arise from explicit or implicit assumptions about the role of art, the museum and the visitor, and these designs reaffirm the assumptions. In previous Human-Computer Interaction (HCI) projects in art museums, these assumptions often remain implicit, and the designs are evaluated with regard to their own aesthetic and functional properties, downplaying the wider context in which they appear. This focus is relevant for the HCI research community, but it does not adequately support the professional practices in their domain of application, namely art mediation.

This PhD project aims to present a theoretical framework that gives us a lens for understanding art as it becomes part of people's experience, which in turn shapes how to understand the role and purpose of art mediation practices, as well as relations to the museum and the visitor. Making this explicit may inform the conceptualization and evaluation of technology-making practices in the art museum domain. To this end, I investigate and develop the relevance of this theoretical framework for design and evaluation through concrete design projects and experiments in the MUNCH museum, a single-artist museum dedicated to the Norwegian painter Edvard Munch.

The theoretical framework also opens a new design space, as it sheds conceptual barriers from previous paradigms, and allows technology to be used in novel ways to support visitors' experiences with art and around art.

Rooted in enactivist philosophy, this theoretical stance rejects the notion of education as a mere transmission of information from museum to visitor, without negating the role of the museum in providing context for the art experience. Rather, it highlights the ways in which the physical, personal, social, and technological shape our relationships with art and can set us up for great experiences.

The PhD project has been co-financed by MUNCH in Oslo, which has enabled an embedded position with a large and renowned art museum. Throughout the project, I have developed this perspective through three main research activities, of which the last two were in close collaboration with MUNCH.

Handling Artworks A lab experiment on somaesthetic relations to digital reproductions of art

Poison An immersive exhibition about the painting series *The Green Room*

New Snow An interactive table about Munch's drawing practice

Through these three activities, I exemplify how this particular view on art, museums, and mediation manifests concretely. I have focused on evaluating the *experience* of participants and visitors, to investigate how art and mediation technologies manifest in their situated lived experience, and to investigate how this experience mediates and educates their attention to aspects of Edvard Munch's art and practice. The project also contains a closer look at deep generative models as a mediating technology, as the technology plays a central part in the immersive exhibition and the interactive table. These three projects have led to five full research papers as well as one extended abstract, each addressing different aspects of this overarching research goal. These can be found in the back of thesis. In the following section, I will briefly present the three main research activities that establish the foundation for this thesis.

1.1 Handling Artworks

This experiment was carried out as a laboratory-based experimental setup at the IT University of Copenhagen in December 2021. The goal was to investigate how digital manifestations of paintings affect the somaesthetic relation we have to them.

19 participants solved a task three times, selecting their favorite amongst three paintings. The paintings were presented either physically in a rack, as 3D renderings with a gestural interface, or as 2D renderings with a mouse interface. After solving the task, the participants were interviewed about their decision-making process. This

Fig. 1.1: The *Handling Artworks* experiment had participants handle physical artworks, digital artworks in 2D, and digital artworks in 3D

revealed how the three setups mediated the artworks differently. This also led to the artworks emerging as different cultural objects such as toys, commodities, and archival references.

The results of this experiment concretely show how the technological mediation of an artwork shapes what it becomes to the participant.

1.2 Poison

Poison is an immersive exhibition that was shown at MUNCH from the opening of the new museum in October 2021 to February 2022. The exhibition was based on research by Signe Endresen, described in her PhD dissertation on "The Green Room" series (Endresen, 2015), seven curious paintings Edvard Munch painted in 1907. In her work, Endresen highlighted the unsettling atmosphere in the paintings, the unstable architecture of the painted room, a co-performative relationship with the characters in the paintings, and the establishment of the painted room as a stage. We took these as salient aesthetic qualities and reenacted them in the *Poison* exhibition. The exhibition was set in a 40m2 gallery and consisted of three small chambers. Stepping into the exhibition through the light lock, you first become aware of the evolving, dark, and ambient soundscape that is filling up the space. As you step into the first chamber you see three framed paintings on the right. Being framed projections, the paintings move slightly. As you walk past them they morph and switch places. Simultaneously, the sound of your footsteps on the wooden floor is exaggerated through a creaking sound coming from speakers near the floor. Through a transparent curtain, you can peer into the next chamber.

In the main chamber, three walls are covered with projections depicting images from "The Green Room" series. The images are of three similar rooms, but with different situations being played out in each. Each of the projections is slightly blurry,

Fig. 1.2: *Poison* was an immersive exhibition that was shown at MUNCH between October 2021 and February 2022

coming in and out of view. As you move around the space, the perspectives of the paintings shift. If the movement of visitors increase, the soundscape intensifies through a deep oscillating bass sound.

Moving through thick velvet curtains to the last chamber, you are met with a bright magenta light. In an entirely magenta room, a brief text is printed on a backlit wall. The text reveals the existence of "The Green Room" paintings from 1907 and poses open questions about what they mean, and why Edvard Munch painted them. As you exit into the brightly lit hallway after reading the text against the magenta light, your vision is tinted green for a few seconds.

Poison was evaluated through interviews with visitors before and after the opening of the exhibition. The interviews focused on the aesthetic qualities of the experience and how this drew visitors' attention to the salient qualities of "The Green Room" and Munch's practice in general.

1.3 New Snow

New Snow is the name of a project concerning Munch's drawing practice. The project resulted in an interactive drawing table powered by a deep generative model. Edvard Munch was very productive and produced thousands of artworks such as paintings, graphic prints, and drawings. The majority of these are now part of the museum's digital archive. The drawings are on paper and in notebooks and are therefore very fragile. They depict a wide range of motifs, from nature, to the neighbor's dog, to fairytale creatures and early sketches of famous works. Due to their number and their fragility, it is difficult to give visitors a comprehensive and exciting look at this aspect of Munch's practice. Only few people, the paper conservators, are allowed to handle these works, and they tell about the material intimacy they experience when handling the drawings.

From this insight, we have developed an interactive drawing table that lets the visitor investigate Munch's drawing practice as mediated through a deep generative model. As you draw, using a marker and tracing paper, a projected drawing appears

on the paper alongside your own lines. The projection adapts in real-time as you draw and move around the paper, constantly trying to create a Munch-like drawing with your drawing as a prompt.

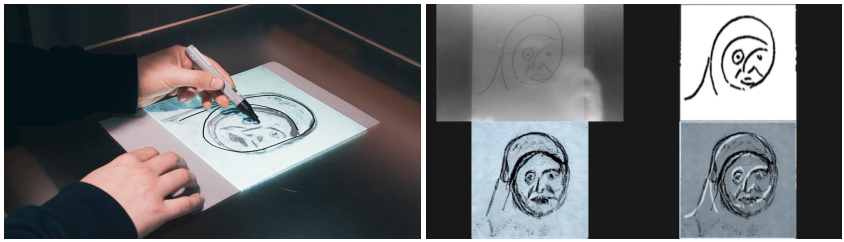

Fig. 1.3: *New Snow* is a interactive drawing table powered by a deep generative model that lets users explore Edvard Munch's drawing practice

The *New Snow* project was evaluated in May 2023 through interviews with participants at the IT University of Copenhagen. The evaluation focused on the relations the table established between the user and the deep generative model and how this created a particular mediation of the underlying drawing practice.

1.4 Structure of the kappa

To motivate and contextualize these three activities, I start by giving a brief summary of major paradigms in an art museums with regards to understanding the role between visitors and artworks leading to the contemporary *immersive turn* (Kidd, 2018). I then present challenges and opportunities for art museums in relation to this paradigm. Next, I look at a selection of HCI projects conducted in art museums, to investigate how they implicitly or explicitly manifest certain museum and HCI paradigms shaping the resulting designs and evaluations.

Following that, I present a theoretical framework for understanding art and art mediation from an enactivist perspective. Arguing from the underlying premise of *art as experience* (Dewey, 1980), I argue that art mediation is a practice of *educating the attention* (Ingold, 2001) of visitors to aspects in the art that they might otherwise fail to see or give weight to. This *education of attention* happens through verbal and non-verbal means alike and in *correspondence with* (Ingold, 2018) architectures, social structures, individuals and art history among other things.

In chapter 4, I present the overarching methods of the project, and how design processes were entangled with evaluation and how this supported the completion of both design and research goals.

In chapter 5, I summarize the publications that are part of the PhD project. I give particular weight to descriptions of the design and the result of the evaluation from each of the projects to highlight the ways in which the theoretical frame is manifested in its particular implementation. I also discuss how each project contributes to the overarching goal of the PhD project.

In chapter 6, I discuss several concerns that may arise with the introduction of new technology into art exhibitions. This includes issues concerning the role of the designer in the museum, originals and reproductions, generative AI, and difficulty and literacy in art.

Finally, in chapter 7 I conclude and outline trajectories for further research in this field.

2 Related Work

The collaboration with MUNCH places this project in the domain of the art museum. The art museum is a contested space, and questions about what it is and who it is for are ever lingering. In this chapter, I will briefly summarize major paradigm shifts in art museum history, as well as how these paradigms have shaped ideas of education and learning and later the role of technology-based mediation in art museums. Following that, I will investigate how these paradigms relate to paradigms in HCI, exemplified by a range of projects conducted in art museums. At the end of the chapter, I will shift the perspective to the particular technology often referred to as *generative AI*, and investigate how this emergent technology is raising particular questions with regard to art, aesthetics, and mediation.

2.1 Educational paradigms in the art museum

Museums are spaces that allow for out-of-the-ordinary experiences. Museums contain architecture and artifacts that do not exist anywhere else, and they are spaces supporting involvement with culture, history, and science as a social activity. Much of museum practice is built on the idea of preserving cultural heritage and educating the public about it. However, how this is done has been the subject of debate throughout museum history.

According to Bedford (2014), nineteenth-century museums embodied a worldview of evolutionary progress and an object-based epistemology. Knowledge was assumed to be inherent in the objects on display, and visitors were expected to be able to acquire that knowledge without further educational support. The museums were seen as disseminating knowledge by their very existence and were expected to transform visitors morally, socially, and politically. However, this came without any real commitment to studying the visitors (Bedford, 2014; Duncan, 2005).

However, in 1916 Benjamin Gilman of the Boston Museum of Fine Arts discovered that museum visitors' interest waned as they went through the exhibits. In an article named "Museum Fatigue" Gilman (1916) investigated the physical exertion of leaning in, crouching down, attempting to read small labels, and seeing objects in glass cases and behind barriers and in bad light. He coined the resulting exhaustion as *museum fatigue* and argued that after a short while exerting to see the exhibited ob-

jects, visitors would resign themselves to glancing over most of the things on display. Gilman further stated: "As at present installed, the contents of our museums are in large part only preserved, not shown" (Gilman, 1916, p.62).

In response to the orientation toward objects and collections, a new paradigm emerged: the aesthetic museum. As one of the proponents of this shift, Gilman argued that the purpose of the artworks in museums was to be looked at as things of beauty. This made it the obligation of the museum to present the works for aesthetic contemplation. The potential result of contemplation was intense and overwhelming pleasure and spiritual revelation (Duncan, 2005). This museum ideal manifested as clean and uncluttered spaces, in stark contrast to the exhibitions Gilman described in 1916. Today, the aesthetic museum ideal persists, leading art museums and galleries to keep educational information physically removed or visually removed from the artworks. In Duncan's words: "the fewer the objects and the emptier the surrounding walls - the more sacralized the museum space" (Duncan, 2005, p.93).

Alongside the physical change in the art museum space, the museum took on the role of communicator, introducing programs, labels, and trained docents (Bedford, 2014). The aim was to educate by presenting academic art-historical information to the visitors. The underlying assumption was a transmission model of communication where knowledge is thought to flow from the source into the head of the visitor (Hooper-Greenhill, 2000).

2.1.1 A turn towards the visitor

A problem with the transmission model of communication is that it does not take into account the complex relationships constituting the social and physical situation in which something is being communicated, as well as the personal circumstances of the receiving party.

In recent decades this has again changed as practices of studying the visitor emerged, and the conceptualization of the visitor changed to being active in interpreting and performing meaning-making activities. This change is also recognized as the *interpretative paradigm* (Macdonald, 2007). Focusing on art museums, Samis and Michaelson (2016) describe the visitor-centered museum as one where the audience matters as much as the collections, and where the museum often strives to reach new audiences and communities outside the core audience. However, they also acknowledge the controversy that such a notion might entail: "On the one hand, it can represent a banner and rallying cry for educators who interact daily with visitors and see missed opportunities for connecting with the public. On the other hand, that banner can turn into a red flag for curators, who fear that it may mean they need to let visitors define the messages — and even the exhibitions — they present." (Samis & Michaelson, 2016, p.2). Samis and Michaelson further argue that the lack of context provided in the museum can lead to art museums being off-putting and alienating to visitors.

Therefore the museum must provide *interpretive hooks* to help those who arrive with little background knowledge gain a chance of *getting* the artworks, and leave feeling they have learned something.

2.1.2 The immersive turn

In a newer article, Kidd (2018) argues for an *immersive turn*, where museums invite participants to take part in story-led, multi-modal and multi-sensory experiences around cultural heritage. Today, many museums engage in a wide range of approaches to education that include museum staples such as text descriptions, various kinds of guides, and dioramas as well as interactive, playful, and storytelling-based elements.

Canonical art has also been involved in an *immersive turn*, although often outside of the art museums. In recent years, many exhibitions have appeared that are dedicated to showing the art of canonical artists such as Van Gogh, Salvador Dalí, Frida Kahlo, Gustav Klimt, and many others. These exhibitions either travel from city to city or are based in permanent venues suited for this type of *digital immersive exhibition* (Mathias, 2022). In these venues, both floors and walls are covered by projections of famous paintings, a soundtrack accompanies the visuals, and sometimes details of the paintings are augmented with animations, making leaves move, stars blink, and water flow. These exhibitions have been recognized for their commercial success, but they have also been questioned with regards to the quality of the art experiences they provide (Mathias, 2022; Mondloch, 2022).

In response to immersive Van Gogh experiences created by many commercial companies around the world, Amsterdam's Van Gogh Museum has created the touring exhibition *Meet Vincent* that also utilizes large digital displays of the artworks rather than originals (Mathias, 2022). The National Gallery in London created the technology-based *Leonardo: Experience a Masterpiece* that tells the story about the painting *The Virgin of the Rocks* through digital interactive installations, highlighting the use of light and shadow, the church for which it was commissioned, and hypothetical designs of the altarpiece where the painting originally was thought to be hanging. The Louvre has experimented with Virtual Reality (VR) technology in an exhibition about the *Mona Lisa*. Through a VR experience, previously available in the museum, users can learn about the history of the painting and the painting techniques used. However, these seem to be among relatively few exhibitions produced by traditional art museums that incorporate technology as a core of the exhibition design.

2.1.3 Contemporary challenges in the art museum

While the museum world has gone through significant changes, these changes are not all-encompassing, and older paradigms might co-exist with newer ones.

In the now 23-year-old text *Changing Values in the Art Museum: rethinking communication and learning*, Hooper-Greenhill (2000) points out many of the ways in which an understanding of education as a matter of transferring information, a leftover from the 20th century museum, leads to "common-sense" assumptions of museum professionals. Among other things, this leads to the organization of exhibitions according to scholarly conventions and a separation between mind and body. Hooper-Greenhill argues that a change of value in art museums is underway, but the shift to constructivist learning theories challenges the foundation on which many curators have built their careers. The same challenge was highlighted by Samis and Michaelson 16 years later and it is arguably still relevant today.

In 2005, Duncan argued that art museums and galleries set the stage for a particular secular ritual, namely that of *aesthetic appreciation*: "Like most ritual space, museum space is carefully marked off and culturally designated as reserved for a special quality of attention—in this case, for contemplation and learning" (Duncan, 2005, p.87). Duncan went on to argue that being able to participate in the ritual necessitates a certain *cultural capital* that some groups are able to accumulate and use in the museum (Duncan, 2005). "Those who are best prepared to perform its ritual—those who are most able to respond to its various cues are also those whose identities (social, sexual, racial, etc.) the museum ritual most fully confirms." (Duncan, 2005, p.87). The understanding of the necessity of background knowledge and the potential risk of intimidation is reflected in writings by Hooper-Greenhill, Samis and Michaelson as well. When art museums *do* experiment and offer different new ways of engaging with art, these understandings might make visitors hesitant to break out of their conventional mode of acting in the art museum.

The legacy of the aesthetic museum, with its disciplined body and the potential conflict between curator and educator perspectives, leads to an ambivalence in art museums about how to deal with technology in the exhibition spaces; an ambivalence that Peter Samis argues "reveals the fundamentally conservative nature of most - though by no means all - art museums as they balk at the opportunities for greater interpretive inclusion offered by new technologies" (Samis, 2018, p.47). He notes that handheld devices, digital guides, and mobile phone apps seem to have become the default choice of technology in art museums, even though evidence shows that these are only picked up by 1 out of 20 visitors (Samis, 2018). They are: "theoretically the ideal solution, as they left no blemish on the otherwise pristine galleries but passed through with the visitors who carried them, like those invisible spectres in long exposure 19th-century photographs" (Samis, 2018, p.58). But even handheld technologies are controversial as they are feared to be contributing to a *heads-down phenomenon*, where visitors spend most of their time with their attention towards the digital devices rather than appreciating the objects on display (Hsi, 2003; Løvlie et al., 2021; Lyons, 2009; Petrelli et al., 2013; Walter, 1996; Wessel & Mayr, 2007; Woodruff et al., 2001). This happens despite evidence that Samis found in a 2007 study that

the most effective interpretative strategies come out of a mix of analog and digital offerings and that a greater number of interpretative resources correlates with enhanced meaning-making and appreciation of artist, exhibition, and museum experience (Samis, 2018).

With this historical perspective on paradigms in the art museum, and their consequence for the use of technology as mediating tools, we turn toward research projects from HCI that have been carried out within art museums to investigate whether and how these challenges have been tackled.

2.2 HCI in the art museum

The HCI community has been active in developing approaches for the use of technology as a mediating element in relation to cultural heritage in general (e.g.Budge, 2018; Harrington, 2020; Hornecker and Ciolfi, 2019; Kenderdine, 2015; Snibbe and Raffle, 2009; Stogner, 2011; Tennent et al., 2020; Waern and Løvlie, 2022). Within the broader field of cultural heritage, research projects have, for example, explored how technology can be used to support critical discourse (Claisse et al., 2020), reproduce historical artefacts (Krumpen et al., 2021) or give insights into intangible cultural heritage practices (Hou et al., 2022).

Some of this work has also taken place in art museums. The projects conducted here differ in whether they are oriented towards specific artworks or towards more general notions of art and creativity. Some focus specifically on learning and education, while others play with the social fabric or other contextual factors of the museum space. Witcomb (2006) argues that every "interactive" in a museum is a manifestation of a particular understanding of pedagogy and learning. In HCI literature, these assumptions are often implicit. In the following, I will outline major trends in the approach HCI projects have taken to art and education in the art museum.

Boehner, Sengers, et al. (2005) argue that most technologies in the art museum domain fall into a dichotomy between *art* and *tool*. Their goal is to break free of a situation where "artists are represented by their art, curators provide information about the art, and visitors, as non-experts, receive this information" (Boehner, Sengers, et al., 2005, p.1). Boehner et al.'s motivation reflect the shift from an information-transfer paradigm towards a more interpretive stance where the non-expert view is equally relevant, as we have seen in the museum domain more generally. To illustrate their position, they present the designs *Imprints* and *Birdscape*. *Imprints* work inside a digital handheld tour guide where the visitor can find information about artworks in the museum (Boehner, Thom-Santelli, et al., 2005). As visitors look up information about an artwork on the handheld device, they can leave a personal image signature on the page of the artwork. This allows other users to gauge who else visited that object. In addition, a photo-mosaic is made from all the visitor image signatures. This project uses artworks to highlight social relations in the art museum. Who else

saw this object? How popular was it? What else did people see? The *Birdscape* concept turns away from the artworks entirely and attempts to draw attention to other aspects of the museum experience to "avoid creating a separate art installation but instead something that permeates the museum experience"(Boehner, Sengers, et al., 2005, p.4). The Birdscape concept was a small device hidden in an art museum near some large windows overlooking a nature area. When there was no activity in this area, it started playing back bird song until the movement of a visitor triggered the birds to "fly away". This was an attempt to "augment marginalized aspects of the museum experience" (Boehner, Sengers, et al., 2005, p.5).

As an alternative to Boehner et al.'s dichotomy, I argue that it is particularly interesting to investigate how these mediating technologies have been placed in relation to the artworks and the "voice" of the museum. Some HCI projects have investigated or positioned the technology as artworks (Andersen & Ward, 2017; Costello et al., 2005; Hindmarsh et al., 2002). Others, like *Birdscape* and *Cluefinder* (Lange et al., 2019) explicitly do not engage with the artworks on display in the museum. Projects that are positioned with a relation to concrete artworks fall into three overarching categories:

1. Projects presenting formal information about artworks, reiterating the information-transfer paradigm.

2. Projects that focus on user-generated content and social perspectives.

3. Projects that are designed to stimulate a range of perspectives.

    (a) Through language and images.

    (b) Through multi-sensorial, aesthetic, and bodily means.

The first group of projects tend to use language such as "augmenting" or "enhancing" the art experience through the presentation of information. These projects embody the traditional idea that the role of the museum is to educate and enlighten visitors by presenting information they can absorb (Bimber et al., 2005; Kortbek & Grønbæk, 2008; Terrenghi & Zimmermann, 2004; Yamazaki et al., 2009). Bimber et al. argue for using pictorial artworks as "information displays" (Bimber et al., 2005) by projecting overlays over original paintings to reveal histories of over- and underpainting. Terrenghi and Zimmerman present an immersive audio environment to "highlight the potential of augmenting the user's experience with the employment of personalized audio information" (Terrenghi & Zimmermann, 2004, p.334). Yamazaki et al. investigate the potential of having a robot tour guide present information about artworks while taking into account visitors' ongoing conduct to "produce engaging explanations" (Yamazaki et al., 2009, p.1437). Kortbek & Grønbæk are more explicit in positioning their design in the art museum space. They argue that the mediation design they propose must avoid "disturbing the domain of the art works" (Kortbek

& Grønbæk, 2008, p.1). They argue among other things for "gentle audio augmentation"(Kortbek & Grønbæk, 2008, p.2) as it is understood to not be "disturbing" in the gallery. They also advocate using the "body as an interaction device", arguing that this solves the problem of people having to focus and concentrate on selecting items on an interface. The bodily interactions are not designed to be part of the mediation itself, rather, the movements unlock inspirational texts and imagery related to the exhibition selected by curators. Through *conceptual affinity*, that is, sharing physical features with the artworks on display, the mediation installations are intended to show which text and imagery concerns which artwork. However, their results reveal that the attempt at separating art from mediation often fails as users are "not able to tell the difference between art works and the *communication of* the art works"(Kortbek & Grønbæk, 2008, p.8, emphasis in original). Fortunately, they also realize that this blurring is not an issue as "neither the artist nor the curator find the blurring problematic, in fact they perceive the holistic experience a quality of the exhibition"(Kortbek & Grønbæk, 2008, p.9). Thus, even though the main argument in the paper concerns a conceptual separation about artworks and content about the artworks, their reflections indicate that this might be more difficult, but also less relevant, than first assumed.

In contrast to the top-down information presented in the previous concepts, the second group of projects builds on a radical bottom-up approach to sense-making. Similar to the ideas of *Imprints*(Boehner, Thom-Santelli, et al., 2005), projects in the second group use user-generated content, social cues and collaboration as the main tools for supporting interpretive practices and sense-making. In the *Artlinks* project Cosley et al. (2008) present a design that let visitors leave a short description of their personal experience of a Buddhist statue on a PC standing next to it. After submitting their response, the words would be presented in a tag cloud of visitor responses with common words larger in size. Through the interface, visitors could investigate who among the previous visitors added which words to the tag cloud. Visitors could also record a sound along with their text tags. The hope was that exposure to the experiences of other people would make people stay longer at the exhibition, and reflect on their own experience in relation to the tags. In this project there were no descriptions from the museum staff, only what was left behind by other visitors. Curiously, looking at the discussion of the ArtLinks project, they find that many of their interviewed participants expected, and would have liked, the PC to present more conventional exhibit information. They also find that participants described the system as helping them to be reflective and stimulate spiritual feeling. However, two participants felt that the act of expressing their reaction in words caused them to have a more cognitive and less emotional reaction than they otherwise would. Two other projects also focus on gathering input from visitors to create a "folksonomy". The project *MobiTags* implemented a social tagging interface on an iPod that visitors could take with them to all the objects in an exhibition. The device offered some basic information about

the objects, but more importantly it allowed for the collection of visitor-generated tags into a folksonomy (Cosley et al., 2009). Visitors could add and vote for tags, and subsequently use the tags as a guide through the museum. Correia and colleagues presented a multi-touch table built around the same ideas of tagging and folksonomy (Correia et al., 2010). These projects all rely on user-generated content as the primary driver for interpretive support.

Not all the projects pool the user-generated content and make it available for everybody. The *MuseUS* project let museum visitors use an app on a smartphone to create their own tour through the museum. By scanning QR codes next to artworks they liked, participants could collect artworks in the application. Instead of comparing this collection to that of other users, the app matched the works with a selection of statements about artworks written by curators. This design did not expose social affordances directly, but it did allow users to take home their exhibition by generating a unique URL that could be accessed via the internet. This approach of tagging content also led to the development of automated recommender systems intended to reinforce the behavior by showing more "relevant" objects (De Gemmis et al., 2008; Wang et al., 2009). Arguing that such automation leads to irrelevant categories, the *GIFT* project focuses on the practice of giving gifts to support sense-making in the museum. Through a smartphone app the gifting user records a verbal introduction to the receiver and then photographs artworks along with personal messages. The activity makes both gifting users and receivers see the art "with new eyes" (Spence et al., 2019, p.7). A similar effect is seen in *Never Let Me Go*, where one visitor issues commands like "Explore", "Breathe deeply", "Touch" and "Mimic this with your body" to another through a smartphone app, again shaping the social setting of the museum and thus the meaning of the artworks through that relation. *Never Let Me Go* also puts the bodily experience at the center, in comparison to the earlier, verbally oriented, approaches. Ryding et al. (2021) argue that this *interpersonalization* should be embraced by museums, even though it might upset the museums' role of providing the official interpretations. However, they also recognize that such services can exist alongside other approaches to mediate artworks in different ways.

The first two groups of projects we have now examined illustrate two radically different understandings of both learning and the role of the museum. In the first group, the technology supports the transfer of information from the institution to the visitor. In the second group, the technology supports the visitor's production of content to support their own interpretive practice, potentially by putting it in relation to other visitors.

A third group of projects attempts to find a middle ground, where the technology facilitates communication from the museum to the visitors, but with a wider variety of potential outcomes. Many of these projects involve the visitors in bodily or affec-

tive relations with the artworks without expectations of a specific learning outcome. This third group can be divided into two sub-groups. The first relies primarily on verbal communication.

The SFMOMA app, launched with the reopening of the museum in 2016, offers a take on the audio tour that expands on the traditional format by introducing a wide cast of actors, comedians, athletes and others to participate as commentators, critics and storytellers in response to the artworks (Pau, 2017). The resulting audio segments do not convey one coherent or authoritative account of the artworks, but they attempt to establish relatable, unique and surprising perspectives. To respond to the isolation audio guides can cause, the app offers an option to synchronize the playback with other visitors. Other projects, like *Sensitive Pictures* also revisit the audio guide format, but offer dramatizations tailor-made for a selection of artworks in an attempt to elicit emotional responses from visitors (Benford et al., 2022). Fosh et al. (2013) combine a musical soundtrack with written and verbal instructions to "look", "touch" and even "climb" in an audio guide through a sculpture garden. These three projects all utilize the well-known handheld tour format, but they change the propositions of the content in relation to the *interpretive turn*, and, in the case of the sculpture garden, even suggest physical actions that would be unheard of inside an art museum.

In the other sub-group, the projects attempt to use aesthetics as the primary drivers in supporting interpretive practices through bodily interaction with designs in the physical museum space. These installations are also directed at specific artworks that are either physically present or represented via digital reproductions.

The Louvre launched the Virtual Reality (VR) application *Mona Lisa: Beyond the Glass* that focused on revealing technical and historical details about the painting in an immersive experience. As you put on the VR headset it shuts off the sensory access to the museum space and transports you to a different place, presenting you with a mix of verbal and aesthetic perspectives, albeit while you remain physically seated.

With an emphasis on hands-on experience, Birchfield, Mechtley, et al. (2008) introduced their *SMALLab* system in the ASU Art Museum with content addressing two specific artworks. One of these artworks is *Many Pierced Disks* (1950), a mobile by Alexander Calder. Calder's mobiles are characterized by suspended and balancing elements that undulate and move with light air currents. In response, Birchfield et al. created a 3D model of a similar mobile projected on the floor of the gallery within the *SMALLab* system. Visitors can use two glowing balls to manipulate the virtual sculpture. By lowering and raising the balls, the mass of the elements in the virtual mobile changes and it reacts accordingly. As the mobile moves, sound effects are played back that relate either to the organic nature of the shapes in Calder's work or to the industrial nature of the metal that the sculpture is made of. The relative position of the balls changes the camera angle from which the virtual mobile is projected. Through this interactive model, Birchfield et al. attempt to draw attention to the formal physical qualities of the artworks as well as the aesthetic.

The other artwork they made content for is a trompe l'oeile painting by John Frederick Peto called *The Rack* (1882). The painting belongs to a series of paintings depicting a range of ordinary objects such as letters, newspapers and books common to the working and middle class at the time. The items are arranged in a newspaper rack. In response to the artwork, the design team created an interface for making similar collages with contemporary objects on the projected surface. Using glowing balls as an interface, visitors can select and place media elements in a digital collage. Again, the design responds to the formal elements in the paintings but incorporates physical practices mirroring the collage aspect of Peto's work.

In the project *TATE Sensorium*, the design team established a similar relation between one interactive installation to support the interpretation of one specific artwork. For each of four paintings they designed a sensory intervention based on vision and sound in addition to touch, smell or taste (Vi et al., 2017). The visitors were guided through the exhibition in small groups from painting to painting. While standing in front of each painting, they were subjected to stimuli before moving on to the next. The selected paintings were "non-representational (or abstract) paintings, as it was agreed [with the curators] that they would leave more room for viewer interpretation" (Vi et al., 2017, p.4). However, in the same article from The International Journal of Human-Computer Interaction, the authors explain that: "The multisensory integration of touch and sound aimed to aid the communication of emotions and meaning hidden in the painting" (Vi et al., 2017, p.2). This seems to indicate an alignment with the information-transfer paradigm, without much room for interpretation. In another article from ACM Interactions, the same group of authors explain that the multisensory stimulation will "color" the interpretation (Ablart et al., 2017). In yet another article from the journal The Senses and Society, Tom Pursey, creative director of the agency Flying Object that participated in creating the exhibition, argues that "the extra stimuli should deepen and intensify that visual, cognitive, and emotional engagement with that piece"(Pursey & Lomas, 2018, p.359). These two latter descriptions hint at an understanding of the technology and sensory stimulation as actively shaping the perception of the paintings, but they do not offer much help in understanding how that might happen. A fourth perspective is offered by David Lomas, a professor in art history, who was not involved with the creation of the project but writes from his experience as a visitor(Pursey & Lomas, 2018). In his analysis of the exhibition, he points to situations where the stimuli, in his opinion, did or did not evoke experiences relevant to appreciating the paintings on display. While sceptical about some of the stimuli, he calls it "a stroke of genius" when the "pleasure-unpleasure" of eating a piece of chocolate made for the exhibition matched the "pleasure-unpleasure" afforded by a Francis Bacon painting (Pursey & Lomas, 2018, p.364). Lomas further sums up the mediation strategy of the *TATE Sensorium* like this:

> "Offering a set of associations via the senses that draw attention to aspects of the artwork, its content and form, triggering personal memories or narratives in response to such stimuli, are strategies that afford a handy route in for viewers who may find modern art baffling and are not equipped with protocols for viewing it" (Pursey & Lomas, 2018, p.362).

This succinct description aligns well with the central argument in this PhD project, as it outlines the idea that aesthetic experience can be used to draw attention to aspects of interest in art, which might afford *interpretive hooks* for those who are less experienced.

As we circle back to the HCI journal article by Vi et al. (2017), we see that the conceptualization of how the design in question supports art mediation in comparison is weak and inconsistent. While it is sporadically touched upon in the analysis of the evaluation with visitors, the authors eventually end up stating that:

> "We did not explore the aesthetics and culture in museum[sic.] as it is beyond our core expertise in HCI. Instead, we focused on exploiting the potential of novel haptic technology to create emotionally engaging and stimulating experiences in particular through its integration with other senses, in our case with sound."(Vi et al., 2017, p.13).

This orientation towards the technology and rejection of the context calls into question the value of conducting this experiment in an art museum in the first place, with all the challenges this inevitably creates.

I will now present another two projects that seem to follow a similar logic to that proposed by Lomas, but as these have been carried out as commercial projects by art museums, the underlying theoretical assumptions behind these projects are also not fully explicated.

Leonardo: Experience a Masterpiece was a largely technology-based exhibition focusing on only one artwork, Leonardo da Vinci's *Virgin of the Rocks*. The exhibition was created at The National Gallery of England and took the visitor through a series of rooms focusing on aspects of the history of the painting. In one room a projection showed a digital replica of the painting, revealing the painting process, in a different room the visitors could explore ways of using light and shadow in a painting, and in a third the original painting is hung in the centre of a digital 3D reconstruction of the altarpiece for which the painting was originally made. The exhibition was "visual rather than word-led" and the intention was "to make people look better and for longer [at great art]" (Campbell, 2019). This exhibition used various technologies to establish perspectives on both technical, art-historical and aesthetic aspects of da Vinci's painting. This was motivated by the possibilities offered by the technology to establish aesthetic experiences rather than its ability to present verbal content.

In the Cleveland Museum of Art, the *ARTLENS Gallery* incorporates a wide range of technological mediation designs that allow visitors or people outside the museum to interact with their collection in various ways. I will highlight the *ARTLENS Wall* and selected stations in the *ARTLENS Exhibition* and *ARTLENS Studio*. The *ARTLENS Wall* is a large display and multi-touch surface that presents all the artworks from the permanent collection that are currently on view in the museum. This amounts to between 4200-4500 objects. The *ARTLENS Wall* incorporates a largely object-centric pedagogy, with orientation towards database and collection. The wall cycles through various sortings of the artworks according to color, purpose, theme and shape, among others. The visitor can connect the *ARTLENS App* to the wall and find more information about the artworks similar to *MobiTags* (Cosley et al., 2009) and the multi-touch table by Correia et al. (2010).

The installations in the *ARTLENS Exhibition* offer a radically different approach to the artworks. Here, 21 artworks are on display physically, in combination with 14 interactive games under the topic "What can art be?". One game asks visitors to mirror a pose found in an artwork to "better understand the emotions of the figure, as well as the contextual emotion of the artwork"(The Cleveland Museum of Art, 2012). Another game lets the visitors virtually wear unfamiliar objects from the collection in a virtual mirror to understand their original function, and a third invites visitors to decode symbols and hidden meanings in artworks. These installations bring the artworks into relation with the visitor's own body. In the *ARTLENS Studio*, oriented mainly at younger audiences, the mediation designs range between options to make your own digital artworks, a large display that allows zooming close into the details of digital reproductions of paintings or a system that pulls images from the archive in response to the visitor's drawings of basic lines and shapes. The interactive designs in the *ARTLENS Studio* use mostly gestural interfaces to let visitors explore features in the artworks. Alexander et al. argue that "while visitors are having fun, they are also looking closer, making connections, and gaining comprehension that will enhance their appreciation of art throughout the museum"(Alexander et al., 2017, *The ArtLens Studio: what changed?*, 2nd paragraph). While they do not offer data that supports this claim, other research has shown how bodily actions and relations can support interpretive practices (Christidou & Pierroux, 2019; Steier, 2014)

In this last sub-group, the active use of the body and aesthetic experience is prioritized over spoken information, however the perspectives offered by the technology are defined by the exhibition designers and not by the audience. Due to the size of these installations, these projects implicitly support social experiences, as multiple visitors can gather around and spectate if not participating.

I have now presented three groups of HCI projects, and while this is not an exhaustive literature review, it outlines three important trends: *Design for information-transfer*, *design for user-generated content* and what we might call *design for stimulation*. This categorization relies mainly on my reading of the projects, as the authors in most

cases do not explicitly state how to understand their work in relation to the role of the museum, the visitor and the artworks. With that in mind, I will argue, that these three groups of projects correspond loosely to the three paradigms of the museum world presented earlier: The didactic approach with the museum as the authoritative voice, the visitor-centered turn and finally an immersive turn with playful, multi-modal and multi-sensory designs.

2.3 Evaluation of experience and learning

These different approaches to design and mediation in art museums also manifest in the ways museum exhibitions, experience, and learning have been evaluated.

Bitgood (2010) proposes an *attention-value model* centering on quantitative measures of when and how people pay attention to exhibition elements in the museum. This approach focuses on the observable behavior of museum visitors and argues that the primary motivation for paying attention is the perceived value of doing so. Bitgood does not engage specifically with the potential outcomes of paying attention, but argues that attention is the practical foundation for outcomes like learning, meaning-making, and inquiry.

About learning, Bitgood says: "Learning has received the most attention as a museum experience outcome for many reasons, not the least of which is that museums include education as a major part of their mission. Given the differences of opinion on what learning is (or is not), it's a wonder that any fruitful discussion occurs on the subject"(Bitgood, 2010, p.24, app. A). And if it is not the definition of learning that is causing trouble, then the practical circumstances of documenting it are. In the book "Learning from Museums", Falk and Dierking (2018) state that it has been documented in general terms that there are learning outcomes from visiting museums. However, capturing the specific learning outcome of one visitor's trip to a museum in more particular detail is notoriously difficult. The learning depends on the visitor's own motivations and free choices of what to attend to, rather than the intentions of the museums. They further argue that while learning from the educational content presented in a museum does occur, it is just one possible learning outcome of a museum visit. Learning from museums also includes social learning, self-awareness, self-confidence, agency, identity, and learning related to aesthetics and beauty. Finally, and maybe most importantly, a significant amount of learning might take place in the weeks, months, or years after a museum visit. Hornecker and Ciolfi (2019) list similar challenges with the assessment of learning and add the creation of meaning, change in attitudes, skills building, empathy, and simply interest to the list of potential outcomes of a museum visit. As a consequence of this complexity, museum researchers have often asked the wrong questions according to Falk and Dierking (2018), focusing too much on content and short-term learning.

If the long-term evaluation of learning is out of scope, and a focus on short-term learning is asking the wrong questions, designers working with mediation and learning support are in a difficult situation when aiming to assess the quality of their designs. Not surprisingly, HCI researchers in museums often focus on evaluating the experience or formal qualities of the design with little focus on learning.

In HCI literature, we find examples of instruments being developed for the evaluation of museum exhibitions. These may include measures for enjoyment, motivation, and collaboration (van Dijk et al., 2012), or like Gonçalves et al. (2012) who score interactive exhibitions on 10 dimensions, such as visibility, feedback, structure, simplicity, and learning. Birchfield, Thornburg, et al. (2008) rely on conversations with museum staff as evidence of motivation, engagement, enthusiasm, learning, and increased time spent among visitors. Kortbek and Grønbæk (2008) use both interviews and questionnaires in their evaluation, and report that "On average around 50% of the respondents claim that they got 'much knowledge' or 'some knowledge' out of the installations."(Kortbek & Grønbæk, 2008, p.7). These methods of evaluation reflect the information-transfer paradigm, taking the design as a carrier of knowledge and assessing the amount of knowledge being carried. Ryding (2020) evaluate the *Never Let Me Go* project through interviews and report primarily on the emergent social experience. This orientation is reflected in other socially oriented projects (Cosley et al., 2008; Spence et al., 2019). This follows the constructivist logic that learning, if any, emerges in the social experience and is mostly out of the hands of the designers. None of these examples investigate the specifics of how technology might mediate a relation to an artwork or another topic of interest.

Among the projects in the stimulation group, the reports only superficially touch upon the aspect that designs affect how the participants make sense of the artworks. In the evaluation of *ARTLENS Gallery*, visitors self-report on the enhancement of their museum experience, whether it encouraged them to look closer at art and notice new things, and whether it increased their interest in the museum's collection (Bolander et al., 2018). This indicates an interest in the attitudes fostered by the interactive installation. However, the publicly available report does not offer insights into the role of specific designs. Benford et al. (2022) report that the visitors trying *Sensitive Pictures* had a high level of trust in the narrative presented in the app and shows one example of how a visitor incorporated the narrative in their own interpretation of the artwork. Otherwise, the evaluation focuses on the emotional responses of visitors through their use of the design. Vi et al. (2017) use questions like "How did the haptic experience match your perception of the painting?" and "What qualities of the paintings were supported through the haptic experience?"(Vi et al., 2017, p.8) to assess the influence of their technological intervention on the interpretation of the artwork. They present a couple of responses with instances of people finding correspondence or dissonance between the stimuli and the painting. They state that, through these experiences, visitors come up with short stories and explanations

to make sense of their own experience, however they do not present these in any detail. I have not found any publicly available evaluation of *Leonardo: Experience a Masterpiece*. Neither of these projects unfold whether and how their designs support interpretations in alignment with any designerly or curatorial intent.

Morse et al. (2022) report that the most important triggers for memorable experiences in museums are *learning*, *atmosphere*, *authenticity*, and *aesthetics*. With that in mind, there is clearly room to address how we design for and evaluate these concepts.

2.4 Paradigms in HCI

Shifting paradigms have also governed HCI research, albeit in a shorter timeframe. Harrison et al. (2007) present three paradigms of HCI. The motivation for this is that: "new approaches enrich our view of interaction, they can also lead to conflicting notions of methodology and validity, whose resolution remains murky without explicit discussion of their underlying epistemological commitments"(Harrison et al., 2007, p.2). As we have seen in the projects presented above, they represent different methodologies and validity systems, however, this is often implicit. Without an explicit discussion of these underlying paradigms, it can be difficult to assess different approaches, especially novel approaches, on their own terms.

Harrison et al. (2007) characterize the first paradigm of HCI as a *man-machine coupling*, a pragmatic approach to merging engineering and human factors that focuses on error prevention. The second paradigm is characterized by a metaphor of *mind and computer as coupled information processors*. Central to this metaphor is notions of rationality and efficiency. How can information enter, be processed, and result in action in humans and computers alike? The third paradigm that Harrison and et al. introduce is *interaction as phenomenologically situated*. In this third paradigm, an interest in the values and politics at the site of interaction emerges. Researchers are interested in how we can support existing activities and let users appropriate technologies to reach their own personal goals. Bødker (2006) presents a slightly different account of how the three first waves of HCI differ by focusing more on the domain of the research. However, on the theory level, she outlines a similar movement from human factors to a rational focus on the work, to an expansion into culture, aesthetics, and emotions in the third wave.

While the paradigms of HCI and the art museum have evolved at different times and in different contexts, I think it is relevant to point out how there are certain overlaps between museum history and the history of HCI. Looking at the second paradigm of HCI, the information processing paradigm, Hooper-Greenhill's critique of the transmission model of communication comes to mind as descriptive of the first category of HCI projects in art museums presented above. Looking at the interpretive or visitor-centered paradigm in museums, we see how the socially-oriented

projects are related to a user-centered orientation in HCI, the third paradigm. These projects provides no authoritative account of the art, but encourages visitors to document their own experiences for others.

Frauenberger proposed a fourth paradigm in 2019, namely *entanglement HCI*. Frauenberger identifies four theories, actor-network theory, postphenomenology, object-oriented ontology, and agential realism that he argues share a *relational ontology*. The concept of *relational ontology* posits that human and non-human actors constitute the world together through various relations, and studying either in isolation is flawed.

This project has relied heavily on postphenomenology as well as the pragmatist philosophy of John Dewey and Alva Noë's ideas of art as reorganization practice under the umbrella term *enactivism*. This is not the first time Dewey's theory of "Art as Experience" and HCI have explicitly intersected. McCarthy & Wright wrote their book "Technology as Experience" with obvious inspiration from Dewey's work. Bertelsen & Pold drew on Dewey's ideas of art criticism when developing interaction criticism. Bardzell & Bardzell's "Humanistic HCI" also highlights the pragmatist influence on HCI as well as Shusterman's somaesthetics that are also inspired by Dewey. This influence is potentially a forerunner of the development Frauenberger describes. Given Dewey's original concern with art, this is a fitting starting point for re-aligning art and HCI within a relational ontology.

Through being explicit about my introduction of a *relational ontology* into work in the art museum, I will attempt to address the "epistemological trouble" Harrison et al. (2007) argued could be the consequence of conflicting paradigms in the same domain. The projects presented earlier in this chapter are, like my project, often carried out by designers or researchers who are only temporarily embedded in the museum, for whom questions about art and aesthetics might be daunting to address, as Vi et al. (2017) had to declare. The projects I have presented above that were carried out mainly by in-house museum teams have often not been documented in research literature, and have thus not had strict requirements for theoretical scaffolding.

To support future HCI work in the art museum domain, I will present a theoretical framework in the next chapter that offers an enactivist view of art and art mediation. Within this view, I will offer perspectives on how to understand the role that technological interventions play in learning, and how designers, curators, mediators, and others who create technology-based mediation can shape the space in which the visitor experience will unfold.

2.5 Generative AI

As the last part of this section I will highlight a specific technology that became increasingly popular during this PhD project. The launch of the generative AI service DALL-E 2 in 2022 was seen by some critics as a threat to human creativity and original art as we know it. Soon followed by Midjourney and Stable Diffusion as the three

major players in mainstream culture, these tools brought the capability of generating high-quality images of almost any subject in any genre to the hands of everyone with an internet connection. This technology poses challenges for artists, art museums, and other fine art institutions. Generative AI systems are capable of imitating any artistic style and expression, given enough data. This potentially disrupts copyright and intellectual property frameworks (Appel et al., 2023; Chen, 2023; Epstein et al., 2020, 2023). AI systems also tend to propagate existing biases in the images they are trained on, such as gender roles, ethnicities, and societal norms (Crawford & Paglen, 2021; Denton et al., 2021; Larrazabal et al., 2020; Wang et al., 2022). The technology has found its use among artists, who are using, challenging, and breaking it in various ways (Audry, 2021).

Machine learning in general has created challenges for design practice (J. J. Benjamin et al., 2021; Dove et al., 2017; Yang, 2018). This is in part due to uncertainty around what a system can do and how well, *capability uncertainty*, and to the potential complexity of the output a system can generate, *output complexity* (Yang et al., 2020). In comparison to heuristic systems, where the governing rules and parameters can easily be described, machine learning systems can be said to discover the rules inherent in the data (Malevé, 2021). This places a large emphasis on the underlying data, as this data establishes the ontology of the system. Machine learning systems are known to be inherently uncertain due to their probabilistic nature (J. J. Benjamin et al., 2021; Hüllermeier & Waegeman, 2021). This is another characteristic that can be difficult to grapple with from a design perspective, as every output is produced by the system with varying confidence according to rules that are unknown. However, the big selling point for generative AI is the ability to synthesize very complex output, such as images, movie frames, text and sound with different aesthetic expressions. With its technological foundation in statistics, it is likely to excel in replicating trends rather than particularities. This potential has been investigated throughout this project.

3 Theoretical Framework

In this chapter, I develop an enactivist perspective to enable a better understanding of how to design and evaluate technology-based mediation in art museums. Starting with Dewey's idea of *art as experience*, I develop a theoretical perspective on art as a phenomenon of lived experience, rather than an essence belonging to certain objects. Using Alva Noë's ideas of art as *strange tools*, I outline what work we can expect good art to do and under what circumstances it might do it. From that follows the motivation for doing art mediation in art museums. Bringing in Ingold's notion of *education of attention*, I present a way to think about the role technology may play in supporting visitors in having art experiences, namely as a form of art critique. Finally, I present postphenomenology as a lens to use both generatively and analytically to understand how specific technologies bring certain aspects of the world in and out of view and thus shape this art critique and the education of attention.

3.1 Art as Experience

John Dewey's 1934 book "Art as Experience" has been highly influential, not just in the discourse around art but also in the HCI community (Bardzell & Bardzell, 2015; McCarthy & Wright, 2004), where many have adopted the pragmatist view of experience as being inherently situated and embodied alongside related theories like phenomenology.

Dewey's view implicates both body and environment in a continuous doing-undergoing which constitutes experience. The environment is both physical and social, and Dewey's view of the body incorporates both mental and physical states as well as previous experience. Dewey argues strongly that the experience of artworks happens in the same realm as everyday experience. There is no separate mode reserved for the appreciation of art. However, what is special to the experience of art is that it occurs in *an* aesthetic experience.

Having *an* experience means letting it "run its course to fulfillment" (Dewey, 1980, p. 35). You can have *an* experience playing a game, eating a meal, having a conversation, and in many other ways, but it is important that the experience comes to its natural conclusion, meaning that it will become demarcated in the stream of experience and memorized. Dewey distinguishes the art experience from the intellectual experience saying that *an* intellectual experience may have aesthetic quality, but argues that it is only signs and symbols without intrinsic quality. Practical actions can

similarly be too mechanistic and automatic to be sufficiently conscious, or too wavering and inconclusive. In between these modes of actions, there is a process led with interest in each step and for the things that come along the way in which meaning accumulates, and then concludes with a feeling of accomplishment. This is the kind of experience Dewey calls *an* aesthetic experience.

If one goes through experience too mechanistically it is categorized by mere *recognition*, which means that the *perception* is stopped before it has time to develop. *Recognition* happens when we shift our attention to what other purposes a thing may serve other than being perceived, or fall back on preconceived notions of what something is. *Perception*, on the other hand, is active and happens continuously in response to the doing-undergoing process accumulating toward fulfillment. It takes place in the entire organism and is rich in emotion. Dewey's view on aesthetic experience demands a lot from an audience and he has little trust in art institutions' ability to support it: "A crowd of visitors steered through a picture-gallery by a guide, with attention called here and there to some high point, does not perceive; only by accident is there even interest in seeing a picture for the sake of subject matter vividly realized. For to perceive, a beholder must *create* his own experience" (Dewey, 1980, p.54, emphasis in original)

What Dewey provides us with is an understanding that not all art can be experienced as art all the time. The aesthetic experience of an artwork is an individual process that unfolds only when other circumstances allow it, and also having *an* aesthetic experience is not limited to artworks.

The final point I will be drawing from Dewey is the distinction between *statement* and *expression*. Dewey argues that, when we are indeed having an aesthetic experience with artworks, they communicate. Artists have an advantage over other experts in dealing with emotions because the artist builds a concrete situation that may evoke an emotional response, while the experts who deal with it intellectually can only attempt to describe the emotion one step removed from it: "Science states meanings; art expresses them" (Dewey, 1980, p.84). The *statements* of science can present conditions under which certain experiences may be had, but unlike *expressions*, they do not constitute it themselves.

In contrast, the *expression* that may come of artistic and, I argue, design practice deals directly with the aesthetic and emotional. Dewey warns about the conflation I am doing here by bringing his notion of expression into the domain of design. He argues that "without emotion [in the act of expressing], there may be craftmanship, but not art; it may be present and be intense, but if it is directly manifested the result is also not art" (Dewey, 1980, p.69). The risk is that the audience may feel that the parts are not coherent, because a sustained emotion did not guide the work.

I will suspend this warning for the moment and move on to our contemporary philosopher Alva Noë's understanding of art. In his 2016 book *Strange Tools* Noë establishes an *enactive* account of art, building among others on the works of John Dewey, Martin Heidegger, and James Gibson, all of whom should be familiar to the HCI community.

Noë argues that a basic tenet of human existence is that we are organized. In all our everyday actions of walking, looking, eating, and talking we are organized, and we participate actively in the constitution of these practices. Art is then practices of *re*organizing these patterns of organization. Painting, dancing, and making music are also practices that we as humans participate in. These practices become art when they put themselves on display and invite the audience to look or listen and try to see how we are organized in it, "they expose the concealed ways we are organized by the things that we do" (Noë, 2015, p.14).

To bring this into more concrete terms Noë brings up Heidegger's analysis of Van Gogh's painting *Shoes* (1886). Heidegger argues that the painting does a philosophical job of putting the "equipmentality of the equipment" (Noë, 2015, p.88) on display. The depicted shoes are worn-down workman's boots that would otherwise be taken for granted and recede into the background of our lives. This painting brings forth the shoes and how they play a role in someone's life for someone who depends on them to get through the day.

However, there is no guarantee that you will see what Heidegger and Noë see in these shoes because one of the characteristics of art is that it challenges us. It challenges us to see it if we can. In contrast to other objects in our life that we can easily understand and evaluate as successful or well-designed, artworks tend to ask, "What are they? What are they for? What might be the standards by which to measure their success?" (Noë et al., 2020, p.249). This also means that:

> "Art affords us the opportunity to be bored to tears, when almost nothing else in our life does. And art's potential to be dull does not contradict the fact that art also moves and thrills and transforms and excites us. Indeed, it is the opposite side of the very same coin. Just as there is no encounter with love without the live risk of heartbreak, so there can be no confrontation with art that does not open up the possibility of getting lulled unconscious and bored to death. Art is valuable only in direct proportion to the degree to which it can, or might, bore us." (Noë, 2015, p.114).

This also means that artworks can cease doing the work of art. As the background for the context of the work changes, its communicative function can change and fade. Noë argues that this happens in particular with art that is very old or from a different culture. While it might be brought to our attention under the label "important art", we might appreciate its craftmanship or monetary value but still not get what it is

doing and thus not show up for us art. In these cases, we need help to go from not seeing to seeing. From this perspective, Noë defines the role of museums as being hosts to the artworks and sustaining the decades-long and centuries-long aesthetic conversations and criticism around the artworks. However, Noë also presents a concern that a museum's preoccupation with making the works "available" to visitors can lead to diminishing the *difficulty* on which they depend. Rather it might be better to find ways of letting visitors accept a little discomfort (Noë et al., 2020).

To investigate further the conversation Noë talks about, I will introduce Ingold's notion of *correspondence*. Ingold builds this idea on Dewey's *principle of habit* (Ingold, 2017), "every experience enacted and undergone modifies the one who acts and undergoes, while this modification affects, whether we wish it or not, the quality of subsequent experiences. For it is a somewhat different person who enters into them" (Ingold, 2017, citing Dewey on p. 15). This principle illustrates constant change and moving along with others as a basic principle in life. *Correspondence* is then, according to Ingold, how things and beings respond to one another over time. This idea mirrors Noë's argument that the artwork may change over time, as the world changes with it, however, we can *correspond with* it to tie *knots* where the artwork yet again becomes *entangled* with people and other things. Again, referring to Dewey, Ingold argues that *communication* is a question of *commoning*, meaning that, as a consequence of *communication*, our paths may join.

This particular view of art, I have now established, argues that the experience of *art* is a form of work that is enacted in relation to an ongoing *correspondence*, and it is the role of art museums to foster this. The correspondence is not strictly verbal, but it incorporates all manners of ways in which the artwork is tied into people and other entities. Both Dewey and Noë acknowledge that the enactment of the art experience or the aesthetic experience also hinges on a skill or capacity in the individual. Art's proposition to *See me if you can!* notably includes *if you can*.

3.2 Education of attention

To understand how we might help people see what they might not otherwise see, or fail to give weight to, I will introduce Ingold's notion of *education of attention* (Ingold, 2001). Ingold proposes this as a theory of skill acquisition derived from Gibson's ecological psychology. This is in stark contrast to the idea that learning works through the transmission of information from educator to receiver. Teaching someone a skill does not entail transmitting some information into their head, but rather educating their attention towards features in their environment that they might otherwise not have attended to. A skilled practitioner teaches a novice by exposing them to relevant features of the environment for the task in question. This attention *educates* the novice, and in the same way their attention is what *is educated* through this experi-

ence of paying attention. Ingold's focus is on skilled practitioners rather than art, but more recently Noë has published a book "Learning to Look" (Noë, 2021), in which he proposes a similar view on learning that extends to art as well.

Using Ingold's term *education of attention* I would like to show how this is relevant in the context of art as well. Noë argues that art requires criticism, it is in fact "the oxygen of art" (Noë, 2015, chap. 8).

> "Good critics do not merely describe a work, they also bring to our attention qualities we had missed, or persuade us to give weight to features that we had ignored or failed to consider. Criticism doesn't proceed by logical argument; there is nothing like knock-down argument in this vicinity. Criticism proceeds by persuasion. Critics are educators. They teach you to see." (Noë, 2015, chap. 17)

So the role of art museum education, or mediation, as I will be calling it here is to teach the audience to see. That is the purpose of the guided tour, the audio guide, and the labels on the walls. However, each of these draws the visitors' attention to different aspects of the art. Different media and modalities support drawing attention to different things. While Noë and Ingold primarily use terminologies related to language, these understandings are not limited to verbal communication. I argue that the job of art criticism can be achieved without verbal communication and in ways that provide their own aesthetic experiences while educating the attention of museum visitors as it pulls the artworks into correspondence. To gain some tools for understanding what a given mediation might reveal or emphasize, we turn to postphenomenology.

3.3 Technological Mediation

One of the central concepts of postphenomenology is *technological mediation*. In the postphenomenological worldview, the relation between subject and object is always mediated and through this mediation, the world is made legible in different ways (Rosenberger & Verbeek, 2015). Postphenomenology commonly has a specific interest in how technologies mediate our relation with the world. This makes it particularly helpful when we are creating new designs, as it offers us a way to analyze how the technology may *magnify* or *reduce* characteristics of phenomena in the world, like art (Kiran, 2015).

Kiran (2015) argues that we can in fact identify multiple dimensions in which technology mediates. He presents the *ontological*, *epistemological*, *practical*, and *ethical* as dimensions in which technological mediation happens. Using an example from the art museum, we can try to understand how a generic audio guide, with art-historical information about paintings on display, might mediate our relation to our being in

the art museum. First, we can understand the ontology of the audio guide through how it *reveals* itself as a tool for accessing art-historical context about the paintings. At the same time, it *conceals* that it might actually be a general-purpose computer that is designed to play back sound files for us. While the museum visitor would most likely see it in this way, museum staff might think of it as something to be cleaned and recharged at the end of the day. In this way, the artifacts are *multistable* and are constituted in relation to how we can approach them. The epistemological dimension concerns how the technology *magnifies* or *reduces* aspects of the world. The audio guide *magnifies* an art-historical account of artists, schools, periods, and provenance, while it *reduces* our sense of other people in the room and their reaction to the art, as it becomes harder to hear and communicate with them while wearing the headset. In a *practical* sense, the audio guide *enables* us to hear the authoritative voice of the curator, even though the curator is not physically present, but it also *constrains* us in our movement. We cannot move our head too wildly or the headset will fall off, and while it *enables* us to pause, stop, and go back, it *constrains* us in presenting our own account of the work. Lastly, the mediation has an *ethical* dimension through the behavior the audio guide *affords* - walking in silence, listening. It *involves* people, who can and are willing to listen, walk slowly and in silence. It *alienates* those who cannot hear or understand the language being spoken, who do not want to walk in silence, or who do want to discuss the art with their friends while in the museum. Becoming aware of these dimensions helps us understand which aspect of the world or the art could potentially be brought to the attention of a visitor. Simultaneously it makes it evident how a design may implement certain ideas and moralities about art and behavior in an art museum.

The final concept from postphenomenology I want to highlight is the variety of *relations* we can have to technology. The audio guide example resembles an *embodiment relation* where a technology changes the user's possibilities of acting and perceiving the world, as described in the analysis above. The world is experienced *through* the device as it becomes somewhat integrated into the wearer's bodily awareness. However, more significant for this project are *hermeneutic* and *alterity relations*.

The *hermeneutic relation* is understood as using technologies through the act of perceiving and interpreting their readout. In this relation, the user experiences an encounter with the world through interpreting the technology itself. This could be reading the time of a watch. However, an important aspect is that it requires a certain knowledge to read the output, and the user's relation will change dramatically with their capabilities. Knowing how to read the time on a watch face mediates a perspective on the passing of time through the passing of hours, minutes, and seconds. Not knowing this conceals this perspective on the world. This relation is not limited to symbolic representation but it also applies to the mechanic assessing problems with a car by listening to the sounds it makes Noë, 2021.

Alterity relations categorize technologies that we relate to in ways that are somewhat similar to other humans. Systems that behave in seemingly autonomous ways, for example AI-based systems, robots, or conversational interfaces, may afford situations where the interaction takes the form of dialogue. This places the system in relation to us as a "quasi-other" (Rosenberger & Verbeek, 2015).

3.4 Summary of the theoretical framework

I have now introduced an understanding of art as experience and as enacted. Art is dependent on *correspondence* and criticism to do its work, and I have argued that art criticism can be understood as helping people to see what they might otherwise not. Shaping what people see is a matter of *technological mediation*.

This means that we do not need the designs produced in this project to be understood as art, but rather as art critique that breathes life into art. Given that written and verbal commentary are well-established genres in this domain, I will focus on showing that such art criticism can be created through non-verbal means and in forms that have *aesthetic quality* on their own. It should refrain from merely transmitting information and making scientific statements about the art, as this might lead audiences to simple *recognition*. Rather, the project relies on expression as a means to support *perception*. While we can correspond directly with art-historical research, it is not the point to restate the content of dissertations or research papers, but rather to reenact the aesthetic qualities that these insights emerge from.

Just like other types of art-museum education, the ultimate interest is to help visitors to see and to build confidence in seeing. Instead of offering statements, we are establishing a potential for new perspectives from which visitors can see the art to gain experiences with aesthetic and emotional quality that they can take with them in encounters with original art objects.

4 Process and Methodology

In this PhD project, I have used the theoretical framework as a *generative lens*. This means using the enactivist perspective with foresight to build assumptions about how certain designs might shape the museum and the art experience, and thereby guide what concrete designs to manifest and evaluate. It has also been used in hindsight to analyze the resulting designs (Hauser et al., 2018; Zwan et al., 2020). One of the consequences of this theoretical perspective is that questions of ontology, epistemology, and ethics inevitably collapse into what Karen Barad calls an *ethico-onto-epistemological perspective* (Frauenberger, 2019). To try and understand this in more specific terms for this project, we can ask three questions to understand how the theoretical framework changes the practical understanding of art mediation.

The first one is: What determines what the artworks *are*? Given the relational ontology in the theoretical framework, there is no single source of truth. In comparison to earlier examples that give authority to the object, the institution, or the visitor, the perspective presented here takes the view that they are in fact all part of co-constituting the art experience. This leads to the next question: How do we design within the logic of *this* theoretical frame? In contrast to Kortbek and Grønbæk (2008), who designed for not interfering, this project has *designing for correspondence* as a central tenet. Given the co-constitutive relationship, the most important aspect of the design is designing for and supporting meaningful relations between visitors, artworks, institutions, and their environment. Finally, an important question becomes: How do we evaluate the outcome of working within this theoretical frame? First, we must understand whether the visitors engage in active *perception* of the design. If they resort to prejudiced *recognition*, it is unlikely that it will bring new perspectives. Next, we are interested in how they experience the design. What qualities of the design do they pay attention to? Does the design afford the experiences we intended? Finally, we are interested in how they place this experience within their own personal context. Where do the users come from, and what do they bring with them from the experience? Importantly, these questions yield concrete insights on how to improve and adjust a design and they can be used throughout a design process. The accounts of these people are the ultimate examples of the relational ontology playing out in practice. Our best opportunity to get to these accounts is through phenomenological interviews with people right after they have used the designs. Through the interviews, we gain a picture of the potentiality of the design.

To exemplify and develop the theoretical approach presented so far, the project relies on three main research activities that have been carried out over the course of three years, the *Handling Artworks* experiment, *Poison* and *New Snow*. Each of these activities addresses the overarching concerns of the projects and generates new knowledge from them. Each of the three activities does this within a certain scope, and by bringing in certain perspectives within the overall interest of the PhD project.

Krogh and Koskinen (2020) argue that *drifting* is an inherent aspect of design research and practice. They argue that, through the practice of designing, we push the project along as we understand more about the problem we will eventually be addressing. Through the analysis of PhD projects within design research, they identify five ways of *drifting* in design experiments. Two of these work as suitable illustrations of the process in the present project, namely the *comparative* and *expansive* mode.

In the comparative mode, it is common to have an overarching and broadly defined research question. There is an inherent understanding of the complexity of the research question, suggesting that research work is done through smaller concrete design experiments that each add parts of the insights towards the overarching question. According to Krogh and Koskinen, "Each design experiment should reveal as-of-yet undocumented additional qualities of a concept and confirm some previously found qualities. In totality, the comparative experiments ideally describe a novel concept, qualify phenomena, unseen quality or add a theoretical distinction to known theory" (Krogh & Koskinen, 2020, p.63).

However, while all three projects address the overall question of how to design for non-verbal art mediation in the art museum, each project also *expands* on the technologies, modalities, and relations to art that can be relevant to this problem. In this way, the scope of the study is broadened somewhat through each activity adding new perspectives to our understanding of the theoretical framework.

This means that the project moves in ways both aligning with the *comparative* and the *expansive* mode (Illustrated by Fig. 4.1). The significance of this particular approach is how it gives space to adapt to interests from the museum partner that are grounded in their particular circumstances as a public art museum, and to insights from earlier research activities. As new insights, questions, and circumstances emerge, a consistent research focus toward a theoretical concern is possible, although very different particulars are explored.

This PhD project can roughly be divided into four phases. The first three focused on the three research activities described in the introduction, while the fourth was a dedicated writing phase. Beginning in late 2020, focus was on establishing a starting point and conducting the *Handling Artworks* experiment in December. In early 2021, I moved to Oslo to work at MUNCH and we quickly began developing *Poison* that opened to the public in October 2021 and was evaluated in the following month. After being initiated in late 2021, the development of *New Snow* was delegated to external partners during early 2022, while I moved back to Copenhagen to teach and take

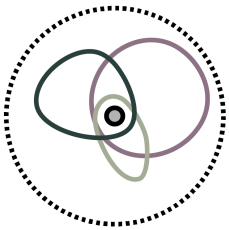

Fig. 4.1: Inspired by the illustrations in "Drifting by Intention" (Krogh & Koskinen, 2020) this drawing illustrates the three research activities addressing the same research interest which is illustrated by the circle in the center. The outer dotted circle illustrates the potential research domain established by the framing of the problem. This is covered incrementally by each research activity.

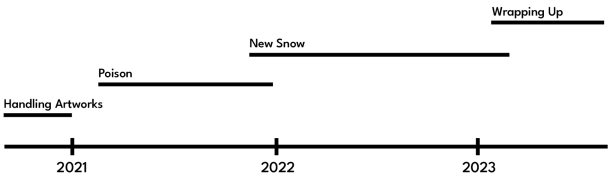

Fig. 4.2: A timeline of the PhD project.

courses concurrently. By early summer 2022, I gradually took over development of the project into its current form, which was completed in time for CHI in April 2023. The remainder of 2023 until the deadline in October has been dedicated to writing papers and the dissertation.

Krogh and Koskinen (2020) use the umbrella term "constructive design research" to describe research projects centered around the main activity of constructing but encompassing a wide variety of epistemological approaches. Looking closer at the three main research activities, we have used more specific terminology in the papers to describe the methodological approach of each activity.

In their respective papers, *Poison* and *New Snow* are presented through Zimmerman et al.'s notion of Research through Design (RtD) (Zimmerman et al., 2007). These projects play a double role common in research-through-design, being both design artifacts attempting to address a design problem while simultaneously providing more general insights for the design research community. They are designed to emphasize the features that are central to the research inquiry being completed in accordance with requirements and constraints that make them possible to exhibit. Being part of the opening program of the new MUNCH museum, *Poison* was characterized by a design language that made it blend into the museum, which included marketing material and digital signage. Meanwhile, *New Snow* retained its appearance and thus status as a prototype for testing and refining in Buxton's terminology (Buxton, 2007).

Zimmerman et al. place the work of the interaction design researcher in a constellation with engineers, behavioral scientists, and anthropologists as a representation of the HCI community. Each of these three types of researchers is understood to be generating a specific type of knowledge that the interaction design researcher integrates into concrete design artifacts. Simultaneously, the interaction design researcher also generates knowledge about technical opportunities and real-life encounters with technology that feed back to the engineers and anthropologists respectively. This particular constellation has not been evident as such in this project, but the notion of the design researcher as an integrator of different types of knowledge is very salient. The distinction between knowing *how* to make things with technology, having a theoretical perspective on the way people and experience work (what Zimmerman et al. call the *true* knowledge), and collecting knowledge about the *real* from evaluations with potential users and stakeholders, are relevant lenses for highlighting the shifts in perspective that have happened throughout the design research process. Zimmerman et al. argue that the *true* knowledge comes from theory on behavioral science. For this project, this is exchanged with enactivist philosophy. The philosophy has acted as a *generative lens* (Zwan et al., 2020), providing a framework for how to craft artifacts within the project. Zimmerman and Forlizzi (2008) tried to distinguish between a *philosophical approach* and a *grounded approach*. A research-through-design project must articulate a *preferred state* - what the design is attempting to achieve - and argue for why the research community should consider this to be preferred to the current situation (Zimmerman et al., 2007). Zimmerman and Forlizzi (2008) argue that the motivation for the *preferred state* either comes from a theoretical vision or from concrete real-world problems. Although the scope of the PhD project as a whole is philosophically motivated, both *Poison* and *New Snow* are also grounded in concrete problems that the museum is interested in addressing, as well as challenges designers face when working with and within art museums. Thus, through the review of previous work involving technology in art mediation, a particular philosophical stance, and inquiry into the needs and interests of the museum, I have presented a *preferred state*, for the research and for the design of mediation technology in the museum, as well as for each of the projects individually.

Arguing for the *relevance* of the design through the articulation of a *preferred state* is one of four criteria that Zimmerman et al. (2007) present as central to research-through-design. A second criterion is *invention*. Zimmerman et al. emphasize that the artifacts resulting from a research-through-design process should be a "novel integration of theory, technology, user need, and context"(Zimmerman et al., 2007, p.499). In this project, the introduction of the enactivist view through the new technological designs should satisfy the need for invention. The third criterion is the documentation of *process*, which is the purpose of this thesis: to present and explain an approach to working with mediating technology that other designers and researchers may reproduce, criticize, and build upon. Making the research *extensible* is Zimmer-

man et al.'s fourth criterion. In living up to these criteria, Zimmerman et al. (2007) foresee a compromise on the commercial quality of the designs in varying degrees in order to make the *right* thing. This was the case for both *Poison* and *New Snow* and will be elaborated on later in this chapter.

Krogh and Koskinen (2020) name this approach to *constructive design research* the "programmatic tradition". In comparison to other traditions that either follow a strict science-like methodology, derive insights through craft-like practices, or through dialogical involvement with users, the programmatic approach focuses on establishing *frameworks* that provide designers with conceptual handles to understand their designs. According to Krogh and Koskinen (2020, p.40): "In this tradition, discourse drives any field of research probably more than artefacts or observations". This resonates with the primary focus of this PhD project to reorient how we conceptualize and evaluate designs in the art museum.

With an awareness of this general approach, we will take a closer look at the *Handling Artworks* experiment, in which the methodological stance is slightly different. In the paper, we describe the experiment as a *concept-driven interaction design research* project based on Stolterman and Wiberg (2010). Stolterman & Wiberg distinguish between research-through-design and concept-driven interaction design research at the level of the intention of the research. As they state, both approaches might lead to insight concerning design and research, however, they argue for making the researcher's orientation towards either one explicit. The *Handling Artworks* experiment was conducted with the explicit goal of theorizing on embodied relations between humans and art objects and was thus not concerned with the generation of design knowledge as such. While I developed bespoke digital tools to support participants' engagement with digital artworks, these designs acted as filters to enable and constrain the particular bodily relations that were of interest to the research but would not offer much as standalone designs in a museum. The experiment was carried out in a laboratory setting at ITU, with participants recruited locally. This makes this experiment narrowly focused on the relation between humans and digital art objects, rather than design practice as such.

Even when adhering to the same general methodology, there can be major differences in how the research process plays out in practice. In the two following sections, I will sketch the development of *Poison* and *New Snow* and compare how they support knowledge generation on the same overall topic through very different trajectories.

4.1 Bringing design theory to the public

Poison was by far the activity with the biggest scope, both in terms of the perspectives of the research project it addresses, as well as the resources spent on it. Development was helped greatly by the investments in manpower, physical, and financial resources from MUNCH. Through close collaboration with my co-supervisor, Nikita

Mathias, and our intern, Dina Patey, we developed *Poison* inside the gallery space it would eventually be shown in. This was a unique opportunity for a museum project, as gallery spaces are usually occupied by the previous exhibition, leaving only a short time to tear down the previous and build up the next exhibition. This was enabled because, due to delays, the new MUNCH museum building had not yet opened to the public. Gradually, we developed the project by taping crude mock-ups to the walls, installing technical prototypes, and building walls from fabric and step ladders. Due to fairly fortunate timing with the ongoing COVID-19 pandemic, we only had to run our first sessions of evaluations sessions virtually. Soon, we were able to invite people into the museum to try out the prototype in person. This was helped along by additional colleagues helping with project management, recruiting of participants, legal matters, interviewing, and prototype building. As the design concept strengthened and the ambitions of a fully-fledged exhibition solidified we drew on experts from within the organization for security and safety protocols, to make modifications to the gallery space, and to help with installation of technical equipment, copywriting, and graphic design. We also brought in external collaborators, most importantly through a collaboration we initiated with the research division of the artist group Random International. They proposed an altered version of the exhibition concept that was used in the final version. Engineers from Random International also contributed with software for visuals and tracking. We worked with a composer who developed the soundscape. The soundscape was a drifting composition played back on a 7-channel speaker system. Microphones under a fake floor picked up the visitor's footsteps and influenced the soundscape using software developed by me. We also worked with exhibition architects who developed practical solutions for the fake floor, projector-friendly walls, hidden speakers, and much more. Another technical expert was hired to install a visitor counting system by the entrance so that we could conform to regulations on maximum occupancy in the exhibition. For the final installation, painters, carpenters, exhibition builders, electricians, curtain installers, and art technicians helped complete the exhibition. This led to a head count of around 30 people involved in the realization of the *Poison* exhibition.

During the project process, my role shifted significantly. Early on, the main tasks were typical for a design researcher, such as facilitating the design process, producing sketches and prototypes, and running sessions with users and stakeholders. As the project grew, the role became more like a design lead, keeping track of the coherence of the design project, while different parts were developed by different domain experts. Towards the opening of the exhibition, the role changed to that of a producer, making sure that technicians, craftsmen, and developers had the right information to deliver the right thing at the right time, and ironing out last-minute ambiguities and problems as they emerged. During the pre-opening and official opening of the museum, I transitioned into being a museum educator, dressed in the same attire as my colleagues from the education and learning department. I stood in front of or in-

side *Poison*, ready to answer questions on art history, the new building, accessibility, or simply the way to the nearest restroom. This gave unique insights into the reception of the exhibition that I could use when, a few weeks after opening, I transitioned back to the role of a researcher. Armed with candy, gifts, and a big smile, I spent my last weeks in Norway convincing visitors to spend 20 minutes of their visit telling me about their experience of *Poison*.

As my role shifted to becoming more coordinating than executing, my perspective on the design changed as well. From working with and making choices about the aesthetic qualities of each technology and material hands-on, the sensitivity changed to working with the interplay between the different expertises involved.

The scope of *Poison* was wide because it was embedded in the museum, multisensory, and comprised multiple technologies in different spaces. It put the visitors in many potential relations to the technology, space, and artworks. The participants recruited post-opening were already embedded in a social situation that did not involve being a participant in a research interview. They had arrived at a brand new and large museum. Most people had visited several other exhibitions before reaching *Poison*. After the interview, they would continue to see even more, and eventually end up in the café or restaurant with whomever they were visiting with. This enabled a rich and holistic view of the experiences afforded by *Poison*. However, with just 15 minutes of actual interview time, there was little time to dive in and explicate individual relations, as in the other two projects. Thus, the interviews eventually ended up touching on relations between the visitor and specific technologies, the architecture and atmosphere, other visitors, other exhibitions, and the museum in general, depending on where the interviewee chose to focus. In the end, the two papers on *Poison* (Sivertsen, Mathias, et al., 2023; Sivertsen, Smith, et al., 2023) investigate the project conceptually and from a holistic point of view that does not go into great detail with individual elements of the exhibitions but aims at describing qualities of the experience from right before a visitor entered the room until a few minutes after exiting.

The strength of this approach is that it involves all the factors relevant to a final commercial design. It requires finding pragmatic solutions to problems within the theoretical framing and thus challenges the feasibility of the theoretical approach for work outside of academia.

4.2 Investigations at the cutting edge

Development of the *New Snow* concept began in late 2021, towards the end of my stay in Oslo. We knew that this project would have to be developed remotely, with me in Copenhagen, the museum in Oslo, and technical partners in Oslo and London. While

from the beginning the project was intended to scale to a full exhibition, we decided to focus on developing the core technology, the drawing table, as a self-contained experience.

The development of *New Snow* started with a very short timeframe. For budgeting reasons, costs relating to the project had to be spent in the 2021 budget year, which resulted in several processes being started concurrently. The development of the table design was to be handled by architects, machine learning by a research group at a local university, and the interactive table by the research division at Random International. A couple of months into 2022, it became clear that development of the machine learning system was required before the other parts of the project could proceed. The exact interaction style and input requirements could not be determined without a working machine-learning system and the table could not be designed without knowing the specifications of the tracking system. However, work on the machine-learning system yielded little results. Gradually, the other two processes were put on hold while we waited for the machine-learning team to deliver. In the meantime, I had met another PhD student at ITU, René Haas, who was working with deep generative models. He had suggested an alternative approach to the machine-learning problem using StyleGAN. During the summer of 2022, two employees from MUNCH, two student assistants, and myself worked on labeling the data for the project. This was a very labor-intensive process of manually annotating thousands of images from MUNCH's digital archive. When fall came around and the machine-learning partner had still not delivered a functional model, I began to seriously investigate the feasibility of using the approach René had suggested. After a few weeks, I managed to establish a toolchain by which I could train a StyleGAN and a pixel2style2pixel model to synthesize Munch-like drawings based on my drawing input. At this point, it was clear that collaboration with the external partners was at a standstill and the most plausible route to a working table was building it myself. Making this decision also involved scaling back significantly on the ambitions of most aspects of the design. Bringing all aspects of the design onto my own plate opened the opportunity for me to deeply investigate the material qualities of the machine-learning system. With less stakeholder involvement, the research agenda could be emphasized and the commercial requirements put on hold. This now changed the need for the prototype to be able to pass as a commercial design in the museum, to *manifesting* just the qualities needed to evaluate it with regards to the theoretical interest (Lim et al., 2008). This, in turn, meant suspending requirements of universal access, robustness, easy maintenance, and the appearance of the table itself. Much more important was a relatively cheap design that could be realized quickly, adapted quickly, and easily shipped to research conferences and MUNCH. This process resulted in the *New Snow* prototype shown earlier in this thesis, which was realized with the support of René Haas on machine learning and Halfdan Hauch Jensen from the AIR Lab at ITU, who assisted with the physical construction of the prototype.

In comparison to *Poison*, my role in this project moved in the other direction, from a position facilitating the deliveries from different partners to being responsible for the execution of all aspects of the design, from cutting holes, to running cables, preparing hardware, training the machine-learning model and reworking the tracking software.

This shift of responsibility changed the outcome of the design process as well as the research outcome, as the design under investigation took on a different role. The scope of this project necessarily ended up much more narrow, with a focused perspective on the human-technology relations and the mediating properties of the interactive deep generative model. While it has not been possible to evaluate *New Snow* within the museum context, the direct engagement with the technology has allowed me to conduct research at the cutting edge of human-AI interaction. The prototype still holds the potential to be reworked for deployment in a museum exhibition.

4.3 Design processes

Despite their conceptual and practical differences, the three projects follow a similar design process. Each starts from a conceptual interest in addressing the research topic of technology-based mediation in the art museum. The first experiment focused on somaesthetic relations to digital reproductions of paintings. The *Poison* experience emerged from an interest in *immersive exhibitions* as a technological practice, and the salient qualities of *The Green Room* centering around atmosphere and space. *New Snow* emerged from a specific interest in the collection of Munch's drawings, generative AI, and the intimate material relation to the drawings that the curators told us about.

Already here, the enactivist foundation has shaped the interest of the projects. Body, space, and materiality are central occupations of each project, as well as the co-constitution of lived experience through action in an environment. The *Poison* and *New Snow* projects also emerge with the assumption that the museum institution has privileged access to perspectives on the art that we would like visitors to be *corresponding with*. This set the initial trajectory of the projects and a vague formulation of the goal.

The worldview and goal are strongly related to theory. At this point, I try establishing a logical chain supporting that a certain experience is in fact possible (according to theory) and assess which design might potentially afford such an experience. In this way, the direction of the first design experiments is laid out.

From this initial interest follows making an outline of the design space with regard to the practical circumstances and resources available. Me being situated at ITU during the first experiment led to a laboratory setup. My embeddedness in the museum and the availability of a room during the development of *Poison* led to the immersive exhibition. Being back at ITU while developing *New Snow* led to building something on a scale that can be developed in one place and easily shipped to another.

Next is involving the technology, and figuring out the interplay between technological systems, and physical and social infrastructures that need to be in place, according to theory, for a certain experience to be afforded. In this sense, the design work is about establishing potentialities within which some experience *could* occur.

This happens through design actions with the technology to understand its materiality and what it affords. This is evaluated continuously through different kinds of formative evaluations, to understand how people relate to the technology. This means having conceptual discussions with colleagues, testing simple technical prototypes or trying out experiential prototypes with outsiders. Following the logic that the anatomy of a prototype determines the way it filters and manifests certain dimensions (Lim et al., 2008), it is possible to investigate those dimensions that are most uncertain at a given time. In all cases, the purpose is to get a sense of people's immediate response to the design, and the assumptions that they might bring along, or the assumptions introduced by how a project is presented. The latter turned out to be particularly important for people's engagement with *New Snow*.

In *Poison*, formative evaluations with people external to the project and institutions played an important role in evaluating the assumptions and qualities of the prototypes. The scope of the project as embedded in the museum space called for evaluations with groups of people in a visiting role in order to be able to assess some of what might occur after the opening of the exhibition. While the technological and material aspects are to a large degree in the hands of the designers, the social aspects are less so and are much harder to assess without establishing a relevant social context to test in. However, that does not mean they are not influenced, as different ways of presenting information, shaping architectures, and playing back sound, along with many other aspects, also *enable* and *constrain* social experiences.

The takeaways from each of the formative evaluations ideally help outline the potentialities of the design in its presented state. These evaluations are key to adapting both the design, the framing of the design, and the research interest, as new possibilities as well as barriers appear along the way.

Eventually, each of the projects was evaluated summatively to assess whether the design evoked experiences of a certain kind, under what circumstances, and how it interplays with all the factors that were impossible to foresee.

4.4 Summative Evaluations

Each of the three projects has been evaluated in roughly the same way. Participants are invited to experience the design for a few minutes, and they are afterward interviewed about their experience.

The interview sessions start with information on the treatment of personal information and the nature of the interview method. Next, the participant is asked a number of structured questions covering demographics, such as age, gender, type of

visit, and indicators of familiarity with museums and the topic. Then the interviewer helps initiate the interviewee's account of the experience and follows up with probing questions as relevant. Finally, the interview transitions to a reflection on the connections the participants make between the experience and the topic in general. It is important that the experience is described through visitors' own accounts allowing us to understand the design as an ultimate particular object (Stolterman, 2008).

The interviewing procedure was inspired by a phenomenological interviewing method described by Thompson et al. (1989). The purpose is to get an account of the interviewee's experience as it was, without rationalization that might skew the interview towards an idealized version. While the phenomenological interview described by Thompson et al. is at the core of the interview in all evaluations, it has been sandwiched between demographic questions and a more reflective orientation at the end of the interview, where the interviewee is allowed distance themselves somewhat from their lived experience. In *Handling Artworks* the interview was furthermore modified by the participants being interviewed for each experimental setup.

Before conducting the interviews for *New Snow*, I had learned about the elaborate opening procedure in microphenomenological interviewing (Petitmengin, 2006). Inspired by this, I prefaced the interviews with even clearer articulations about the participants' role in the interview and some methodological assumptions to make them feel comfortable and competent in telling about their experience.

The interest of this PhD project is in how the presented designs afford situated lived experiences. These experiences are by their very nature individual and subjective, but through multiple interviews, common themes will occur as well as unexpected reactions to the design. It is through these interviews it is possible for the design researcher to get a glimpse of how the design co-constitutes the world for visitors and participants. This way of inquiry into people's lived experiences is, I believe, particularly significant for design research, where new technologies, concepts, and contexts intermingle because it allows for the discovery of emergent meanings and insights. While the design embodies the designer's a priori understanding of the world, its inherent *multistability* makes it next to impossible for a designer to foresee all the potential ways it will be made sense of and appropriated. Nevertheless, the potential is also limited, making it important to figure out where the outlines are, and whether most participants make sense of the design in ways that are aligned with the design intent.

Handling Artworks and *New Snow* were evaluated with a postphenomenological lens, with attention on how the design mediated across different dimensions and through different relations. Through participants' description of their experience, the different dimensions of technological mediation become evident. This allows us to evaluate what the design *conceals*, *reveals*, *magnifies* and *reduces* in the participants' *functional perspective*.

Poison, being an exhibition, stretched through different rooms and across different technologies and thus had to be evaluated in a slightly different way. The post-phenomenological lens does not lend itself well to analyzing across a range of encounters with technology, each involving different types of relations. Instead, we conducted the analysis of *Poison* in three steps. First, are the visitors accepting and open to the design or do they distance themselves from it? In Dewey's terminology, this is a question of whether the participant merely *recognize* or whether they engage in *perception*, with the potential of leading to *an* aesthetic experience. Second, we investigate the aesthetic qualities that the visitors report. As aesthetic qualities indeed are what we propose that visitors *correspond with*, it is important to understand how they manifest. Put in another way, do they pay attention to the same features of the environment that we wanted to put on display? Here, with some exceptions, we find an alignment between what we designed for and what the participants reported. Lastly, we ask visitors to place what they experience within their understanding of Munch's practice. While this question will inevitably lead visitors to draw parallels, whether or not they had already done so, the question yielded interesting insights into how their previous experience with and understanding of Munch's work made certain aspects of the experience stand out. If we understand learning like Falk and Dierking (2018) as something that emerges when previous experiences merge with new, the responses to the question indicate some of the learning that happened as a result of previous experience, while the second question might indicate some of the aspects that could turn out to be significant in the days, months or years after their visit.

4.5 Accountabilities and collaborations

This project is co-financed by the MUNCH museum, more specifically the Department for Visitor Experience and Learning. I spent the whole year of 2021 in Oslo working at the museum alongside colleagues from this and other departments. This financing relationship and the embeddedness of the projects within the organization have naturally shaped the process in a significant way.

From the beginning of the project, it has been grounded in certain strategic goals of the museum and the Department for Visitor Experience and Learning. The museum has a general interest in attracting new visitor groups to the museum, especially younger people. This project follows other collaborations between MUNCH and ITU revolving around technology in museums. This has created an awareness of technological developments and their potential use in art mediation. Simultaneously, immersive exhibitions and generative AI made a stir in the art world during this PhD project, naturally evoking an interest in these technologies specifically. MUNCH has

an ambition to be known for being innovative amongst art museums and it is an oft-repeated understanding at MUNCH that experimenting with new technology is in the spirit of Edvard Munch.

The purpose of the Department for Visitor Experience and Learning is explicitly to support exciting, meaningful, educational encounters with art. It is not the job of this department to conserve, collect, and document artworks nor to conduct art-historical or conservational research. Given my own background as a designer brought up on pragmatist and phenomenologically based interaction design theory, the situated lived experience of users, visitors, participants and humans in general are always at the core of my research interest. This results in a strong alignment in the *design for experience* between the research project and MUNCH.

5 Summary of Papers

Having now outlined the overarching project scope, theory, and methodology, I will summarize the five research papers and the extended abstract that contribute to this PhD thesis. For each paper, I will focus on the concrete results presented, and how this contributes to the overall aim of the PhD project.

5.1 Handling Digital Reproductions of Artworks

Sivertsen, C., & Løvlie, A. S. (2021). Handling Digital Reproductions of Artworks. *Journal of Somaesthetics*, *7*(2), 21

At the very beginning of this PhD project, I set out to investigate how the somaesthetic relation between people and artworks affected their understanding of artworks. Knowing that the project would revolve around technology mediating our relations to artworks, I devised an experiment where participants would be handling physical artworks and two types of digital artworks for a specific purpose and compare their somaesthetic experience of each.

The paper traces museum history back to the 17th and 18th centuries, when museums expected visitors to touch and handle artifacts. At that time, touching was understood to enhance the enjoyment of art objects and support an intimate connection with the original creators. Some objects were even believed to have healing powers (Howes, 2014). However, as this practice ceased in the 19th century, with a new focus on *contemplation*, codes for correct practices of walking, sitting, standing, looking, and speaking emerged (Howes, 2014).

The paper presents a lab experiment in which 19 participants went through three rounds of evaluating which of three artworks they would like to have at home. In each round, the artworks were presented differently. In one round the artworks were physical and placed in a rack, in another the artworks were digital 2D projections that the participants used a mouse to enlarge, and in the third, the artworks were 3D objects shown in a projection. In this round, the participants used a smartphone as a gestural interface for picking up and examining the three artworks (see 1.1).

In the experiment, we saw that the three different setups mediated the epistemological, ontological, practical, and ethical dimensions of the participants' relation to the artworks. More concretely this manifested in the participants' understanding of the artworks as different kinds of cultural entities. For example, in the physical setup, the artworks were seen as commodities in a gift shop. In the 3D setup, they became

playful objects in a computer game and finally they became items from a web search in the 2D setup. These results make it clear that the artworks are not just artworks, but constituted as complex cultural entities through the technological mediation of the three setups. The paper also concludes that, by using technology, we might be able to reestablish some of the sensory relations to artworks that were lost with the introduction of the aesthetic museum. Finally, the paper discusses concrete ways in which designers can explore the mediating dimensions to afford particular kinds of experiences with art.

In this way, the results from this paper exemplify the basic premise of this thesis that the artwork is constituted in relation to technology and the human, making its ontology situated and relational. Art experience and appreciation emerge not only in an intellectual relation but through the *functional perspective* and bodily relation to the work. We as designers of technology have the ability to shape this mediation, affording different relations to the art.

5.2 Art Critique by Other Means

Sivertsen, C., Smith, M., & van der Zwan, S. (2023). Art Critique by Other Means. *Designing Interactive Systems Conference (DIS '23)*, 10. https://doi.org/10.1145/3563657.3596069

The second paper expands the theoretical perspective of the project with a stronger conceptualization of art experience and art critique. Building on Noë's account of art as *reorganizational practice* and Ingold's notion of *education of attention*, we argue that technology used for mediation purposes in the art museum can be understood as a form of art critique using non-verbal means to educate the attention of museum visitors to aspects of the art on display they would otherwise have overlooked or not given weight. To illustrate our point, we use the *Poison* project as an example while contrasting our approach to earlier HCI projects in the art museum domain. Many of these attempt to avoid disturbing or interfering with the art on display, or simply forego engaging in any theoretical discussion of the relationship to art that their design affords. Taking the *art critique by other means* perspective means opening up a new design space, where correspondence with art is given, raising the question of how to correspond meaningfully, rather than how to avoid it.

The understanding of the role of technology and mediation presented in this paper is the foundation of the theoretical framework presented earlier in this thesis.

5.3 Educating the Attention of Museum Visitors through Non-verbal Art Mediation

Sivertsen, C., Mathias, N., & Løvlie, A. S. (2023). Educating the attention of museum visitors through non-verbal art mediation. *Proceedings of the 10th Congress of the International Association of Societies of Design Research*

In this paper, we present the findings of the summative evaluation of the *Poison* exhibition. Using the theoretical foundation presented earlier, we investigate how the exhibition was received by the visitors and discuss how we can use the theoretical framework to understand how it *educates attention*.

The evaluation was guided by three leading questions that also serve as a suggestion for future researchers to build upon when evaluating their work in the art museum. Through interviews with 47 participants, we evaluated each of these questions.

1. Were the visitors accepting and open to the design or did they reject it or label it off the cuff before properly experiencing it?

2. What were the aesthetic qualities that the visitors reported? Were they in correspondence with the design intent?

3. How did visitors place their experience in relation to their existing knowledge of the topic/artwork/artist?

We found that *Poison* was well-received by the interview participants and they found it a timely and relevant addition to the new MUNCH museum. Most participants were able to richly describe their experience in the exhibition.

The design of *Poison* builds on qualities described in Signe Endresen's research on *The Green Room*. In summary, these qualities are an unsettling atmosphere, unstable architecture, dark emotions, *The Green Room* as a stage, and the visitors as co-performers of the paintings. In the interviews, we saw that the aesthetic qualities described by the participants *correspond with* this design intent. The atmosphere in the exhibition and the feeling of being present in "The Green Room" were described by the interviewees. The co-performance aspect was, however, not well represented in the interviewees' accounts of their experience. Nevertheless, we argue that *Poison* to a large degree delivers on the design intent to afford an aesthetic experience *in correspondence with* Signe Endresen's research.

Finally, we evaluate how participants relate their experience to their understanding of Edvard Munch and his practice. We see that visitors draw relations to their understanding of Munch from the text in the exhibition as well as symbols present in the projections. Equally important, we see this happening from the aesthetic qualities

of the exhibition. This supports the idea that the non-verbal aspects of the experience are also part of educating people's attention to particular aspects of Munch's practice and establishing a potential for learning. Importantly, different aspects of the experienced atmosphere draw attention to different aspects of Munch's life, artworks, and practice. This highlights the need for designers working in the art museum space to consider how they want the verbal and non-verbal elements of any design, technology, and architecture to play into shaping the interpretive practices of the visitor.

5.4 Exploring a Digital Art Collection through Drawing Interactions with a Deep Generative Model

Sivertsen, C., Haas, R., Jensen, H. H., & Løvlie, A. S. (2023). Exploring a Digital Art Collection through Drawing Interactions with a Deep Generative Model. *Extended Abstracts of the 2023 CHI Conference on Human Factors in Computing Systems*, 1–5. https://doi.org/10.1145/3544549.3583902

This extended abstract accompanied a demo of *New Snow* at the interactivity track at CHI'23 in Hamburg. While this extended abstract mainly describes the technical aspects and the functionality, it paved the way for the prototype to be shown to the many attendees at CHI'23, which worked as a formative evaluation of the design. The experiences from CHI'23 also led to adjustments in the research design that led to the following full paper.

In this publication, we present how we curated and labeled a dataset of around 5800 line drawings from Edvard Munch, which we used to train a StyleGAN2-ADA model as well as a pixel2style2pixel encoder to let us find images in the StyleGAN model using hand-made drawings as a "prompting" tool.

Furthermore, we present a series of assumptions. First, that deep generative models (DGM) reflect qualities of the data on which they are trained. Second, that the exploration of datasets and models can happen through non-verbal means. Third, that learning how to prompt a model means building an understanding of the data on which it is trained. These three assumptions form the theoretical foundation for how we understand users may become familiar with Munch's drawing practice through the interaction with *New Snow*. The extended abstract finally contains a specification of how the table is constructed and how the tracking system functions.

5.5 Exploring Aesthetic Qualities of Deep Generative Models through Technological (Art) Mediation

In review: Sivertsen, C., & Løvlie, A. S. (n.d.). Exploring Aesthetic Qualities of Deep Generative Models through Technological (Art) Mediation. *ACM Transactions on Computer-Human Interaction*

This paper presents the *New Snow* interactive drawing table and its evaluation. The paper starts out by highlighting problematic aspects of large image datasets in the way the images are sourced and annotated. This leads to different kinds of bias baked into the dataset. As we train models on large datasets, as has been explored in earlier literature (J. J. Benjamin et al., 2021), the inherent patterns and biases tend to leak from the model when in use. This can cause different kinds of problems depending on how it is used. On the other hand, it gives us a way to understand and assess the aesthetic qualities of the model in terms of how these qualities manifest as the model is put to use in a specific context. Despite its obvious pitfalls, this particular quality of deep generative models (DGM) presents an opportunity to expose patterns found in large datasets. When it comes to art, this might in fact be a particularly interesting way of exposing visual aesthetic trends found in the corpora of digitalized artworks.

A similar approach to generative AI has been seen in art, research, and a few services, where the purpose of the AI is not to create image output towards some external use case. Rather the interface for the model is explicitly made to cause reflection in the user on the model's qualities. We call this approach *designing for reflexive use* and argue that a similar approach is relevant for engaging with large art datasets.

To evaluate this, we recruited 20 participants to take part in an evaluation of the *New Snow* drawing table. Each participant drew on the table for 8-10 minutes and was afterwards interviewed for 10-12 minutes about their experience. Through the interviews, we uncovered different drawing strategies, the aesthetic experience of using the table as well as specific qualities in the way the Edvard Munch drawings are being recreated by the StyleGAN model.

One of the findings is that the synthesized sketches can be understood as *uncertain entities*. This means that their ephemeral nature and appearance are helpful in balancing the expectations of authenticity and representation that participants have for the output of the system. This is a particularly important quality to balance, as, although the StyleGAN model does represent some qualities of the inherent dataset, it is also capable of both hallucinating and suppressing details in the drawings. However, as we argued in *Art Critique by Other Means*, this can be understood as a way to *educate the attention* to certain aspects of the artworks. In this way, we argue that *New Snow* directs attention to Munch's drawings as a practice rather than to the particulars

of individual drawings. However, making sure that the system does that is achieved through the particular data curation we have done for this model and the particular way of interacting with it that the table affords.

Finally, we identify three characteristics of the design that afford this reflexive use of *New Snow*. The first aspect is the *prompting modality*, which is significant for establishing the epistemological relation to the dataset in what can be retrieved from the model and in what way. The image-to-image relation in *New Snow* allows for a particular kind of inquiry using drawing actions to investigate drawing actions. The second aspect is *incremental prompt adjustment*. This is the ability to slowly and gradually change the input prompt to allow for an exploration in which relations along specific dimensions of prompting become evident. The third aspect is *fast updates* which ensures a fluent and pleasing interaction with the system, supporting the user in continuing to explore.

In summary, *New Snow* is the second example of a design that illustrates how the enactivist understanding of technology and art mediation can play out in practice. Furthermore, it opens up new research avenues with regard to how we conceptualize and evaluate designs using generative AI.

5.6 Machine Learning Processes As Sources of Ambiguity: Insights from AI Art

In review: Sivertsen, C., Salimbeni, G., Løvlie, A. S., Benford, S., & Zhu, J. (n.d.). Machine Learning Processes As Sources of Ambiguity: Insights from AI Art. *Proceedings of the 2024 CHI Conference on Human Factors in Computing Systems*

This paper explores a theoretical perspective on how artists use AI for aesthetic effect. We revisit the idea of *ambiguity as a resource for design* that Gaver and colleagues introduced in a 2003 paper (Gaver et al., 2003) to investigate how it applies to projects involving deep neural networks (DNN). Like the original paper, we use artworks as examples where ambiguity is used purposefully. Based on the original 10 tactics, we select nine examples of AI artworks that in different ways exhibit ambiguity, and argue that ambiguity emerges from the artist's ways of approaching the machine-learning process. Through the artists' data curation, model training, and application of the model, different choices are connected with different ambiguity tactics. We argue that understanding how artists approach machine learning provides insights into understanding AI as a design material.

We further discuss how the ambiguity exposes qualities of the machine-learning process and the data it is built on in ways that make us assess, question, or reinterpret how we understand these kinds of systems. This provides a relevant supplement to the efforts being put into *explainability* of AI systems that focuses on formal explanations of the way a system works.

Both *Poison* and *New Snow* rely on a level of ambiguity, to keep the visitors and participants engaged in active interpretation and sense-making. The concept of ambiguity is important to Noë's notion of *difficulty* in art and Dewey's notion of *perception*, as the things that are well-defined and clear-cut do not invite perception or reorganization but can be easily labeled and set aside mentally.

6 Discussion

In the domain of the art museum, like in any other professional domain, particular assumptions and logics are at play, governing specific work functions. The art museum remains a conceptually contested space, in which various understandings of the nature of objects and experience co-exist. For this reason, this PhD project has addressed notions of art and art mediation in order to carve out a design space in which design practices for mediation can unfold. This work continues in the discussion section, where I will begin by discussing how the theoretical framework and the suggestion that technology can support art mediation through aesthetic experience interplay with topics of special significance for the art museum domain.

First, I will discuss how understanding art experience from an enactivist point of view has implications for how we think of *aura*, the museum as an authority, and aesthetic literacies. In the second part of the discussion, I will turn to more design matters, discussing how to design for learning, implications for evaluating design, the relevance of ambiguity in art mediation, and finally the potential role of generative AI.

6.1 Art and entanglement HCI

In this thesis, I have argued for the introduction of *enactivism* into design research in the art museum domain. I have compared this idea to Frauenberger's identification of the *relational ontology* as a central characteristic in the fourth wave of HCI. Building on pragmatism and postphenomenology, I have argued for a similar relational ontology as the lens through which we should understand art, experience, and technology in art museums. In this section, I will address the most significant ways this lens can change our perspective on some of the traditional values of art museums, and discuss whether this is to the detriment of art museums.

Within this framing, art is an *experience*, not objects. Understanding art from this perspective changes the way we understand authenticity, originals, and copies. The existence of art is something that the museums actively support, not only by conserving the physical objects but more importantly by facilitating that these objects can do the *work of art* with the audience. This questions the reason for the obsession with originality and provenance in art. If art is not inherently dependent on the original object, will any copy then do?

6.1.1 Originals and reenactments

In both *Poison* and *New Snow*, imprecise reenactments of original paintings were a core feature in the mediation strategy. The reliance on copies, reproductions, or reenactments inside the art museum might be considered controversial, and the value might not be obvious. In the following, I will discuss the potential of using reenactments for art mediation.

In the MUNCH museum's biographical exhibition "Shadows", digital reproductions are used as references on a timeline of Munch's life. On touchscreens, visitors can browse through a few select samples that illustrate works he was making at a given point in his life. Reproductions also appear in the gift shop on posters, magnets, and cups that you can buy to take home. The role of these is to disseminate the original, or in other words, create visibility of the collection(Christensen, 2017). Digital reproductions also appear, sometimes remixed, in the visual identities accompanying each exhibition, in the form of digital posters and animated teasers. Christensen argues that the art museum is rife with practices of copying, and situations where copies are accepted as legitimate on their own, although while still affecting the perception of other originals. (Christensen, 2017).

The rise of generative AI has spurred controversy in art circles, as it calls into question authorship and creativity. In this project, I have argued for the usefulness of generative AI in letting people investigate the aesthetic qualities of large art datasets. This happens through the generation of imagery that, in many aspects, bears resemblance to the original artworks from which it is derived. However, the quality of the image might be skewed towards certain trends in the data or contain hallucinated details that do not occur in any of the original works. Even without obvious alterations, the digital reproductions used in the immersive Van Gogh shows, for example, are different in scale and materiality than their original counterparts. Any kind of digital reproduction in the art space seems to inevitably evoke questions concerning its *aura* (W. Benjamin, 1936).

While Walther Benjamin set the tone for the debate around aura with the seminal article "The Work of Art in the Age of Mechanical Reproduction" (W. Benjamin, 1936), newer research has argued that authenticity and aura are not as tied to physical originals as Benjamin suggests (Jeffrey, 2015; Jones et al., 2018; Kenderdine & Yip, 2018; Latour & Lowe, 2011; Meehan, 2022). Latour & Lowe go as far as to say that the aura might even migrate to facsimiles under the right circumstances. They present the case where a highly sophisticated reproduction of *Nozze di Cana* by Veronese, placed in its original location on the island of San Giorgio in Venice, is held in even higher regard than the original painting hanging in the Louvre in Paris.

Nicole Meehan interviewed and surveyed museum professionals on their thoughts on materiality, authenticity, and aura of the *digital museum object* (Meehan, 2022). She argues that the prevailing notion of digital museum objects places the digital in a hi-

erarchical position below the physical. This is based on the Benjaminian idea that the digital object cannot possess *aura*. Despite this, some of Meehan's respondents also point to the potential of digital objects to foster new perspectives and interpretations, although not without fear that the digital objects can somehow "break down the traditional aura" (Meehan, 2022, p.427). However, Meehan also argues that we should see "the [digital] object as a point of connection between learners, or visitors, where knowledge is produced by the learners." (Meehan, 2022, p.438)

In an article on the aura of digital models in archaeology, Jeffrey argues that digital records and replicas "are not simply a technical exercise, they act as powerful expressions of our present world view, and if not considered in this light, they fall short as meaningful representations of the past" (Jeffrey, 2015, p.151). For the project *Pure Land AR* Kenderdine and Yip (2018) report evoking *auratic experiences* through a virtual reproduction of one of the Mogao Grottoes in Gansu Province, China that was on display in the Art Gallery of New South Wales, Australia. These originals that the reproductions point to in Jeffrey's and in Kenderdine & Yip's work, however, are inaccessible to the public and it is not possible to bring them inside the museum.

My work takes place inside an art museum with a large collection of originals in storage. This does not call for highly detailed copies but rather for a different kind of copy - or reenactment - where maintaining or reestablishing *aura* is not the main point of reproduction. As Jeffrey argues, digital reproductions can in themselves be powerful expressions that facilitate a particular perspective on cultural-heritage objects. In a later article, Jones et al. argue that "a pre-occupation with the virtual object - and the binary question of whether it is or is not authentic - obscures the wider work that digital objects do"(Jones et al., 2018, p.350). In their *ACCORD* project, they find strong evidence that "the creation of digital models can actively mediate the authenticity and status of their original counterparts"(Jones et al., 2018, p.350). Following the theoretical foundations of the present project, these statements underline the potential for the use of digital reproductions in art mediation.

As we argue in "Art Critique by Other Means" (Sivertsen, Smith, et al., 2023), the point of our reenactments of *The Green Room* paintings is to emphasize and exaggerate certain qualities, to make them stand out and become significant for the visitors. Latour & Lowe talk about how reproductions contribute to the *career* of the artwork, while we lean against Ingold's idea of an ongoing *correspondence*. We also use the word reenactment rather than copy or reproduction, because it is exactly how the ideas of the works are being put into motion by a given version that is important, rather than how the work was produced. Latour & Lowe argue that it is crucial to discriminate between good and bad reproductions, as *aura* might migrate to reproductions if they reproduce the artwork well enough. However, I might go as far as to argue that the distinction between good and bad is too reductive (also) in this case, but that each reenactment of an artwork should be characterized by its ability to either *do the work*

of art on its own, or by supporting another art object in doing it. Chances are that, if it manages to do that, it might facilitate the migration of or an entirely new feeling of authenticity, even though it does not share the same molecules as an earlier version.

In this PhD project, we find the reenactments in *Poison* and in *New Snow* are attempting to support the originals in doing the work of art, through contributing to the career by *corresponding with* the original paintings. In *Poison*, digital versions of *The Green Room* paintings are shown that are both blurry, but also offer exaggerated features, such as the dynamic perspective. In *New Snow*, the synthesized drawings replicate the trends but not the particularities of the underlying corpus of drawings. In this way, they both take part in the *correspondence with* the originals by establishing a perspective for the visitors on the artworks that they might not have had before.

This goes beyond the reproductions in the *Shadows* exhibition that attempt at establishing a "neutral" reference, the exhibition teasers that take part in building and maintaining the MUNCH brand, or the posters that provide a bit of ownership and identity through the poster as a commodity.

6.1.2 The authority of the art museum

The understanding of art mediation presented in this project sits between two positions on interpretation. On the one side is the authoritative museum of the aesthetic museum era. On the other side is the ground-up approach to interpretation, where the museum merely facilitates a space in which visitors may co-construct their own interpretations. In one, the museum is the sole guardian of knowledge, on the other extreme the art-historical knowledge held by the museum is next to worthless.

In the present perspective, the power distribution is more balanced. In my projects with MUNCH, we have consulted art historians and curators to hear their perspectives on the artworks in question. On the other hand, we know that these perspectives are not transferable as such to museum visitors. Rather, the perspective of the museum professionals arises from their privileged position in relation to the art, their experience, and formal training. Nevertheless, we can design our mediating technologies so they highlight and emphasize the qualities that the museum professionals pay attention to when talking about the artworks. We cannot determine what visitors should see, but it is our job to stage the artwork in such a way that relevant perspectives are possible and even emphasized.

In this way, we acknowledge that the museum professionals' engagement with the art gives them perspectives on the art that are also relevant for the public. Art museums have been, and still are under scrutiny for how they deal with minority perspectives and diversity (see for example the field of critical museology). In the present work, I have not directly engaged with these perspectives, but rather let the agency and authority of the mediation stay with the museum professionals I have collaborated with. However, it is important to underline that the perspective pre-

sented here does not hinder this work, but can also support participatory, feministic, more-than-human, and anti-colonial agendas. This hinges on the way the mediation design corresponds with these ongoing matters, and what we choose to draw the visitors' attention to.

6.1.3 Aesthetic literacies

In this project, I have argued for leveling the playing field between verbal and non-verbal approaches to art mediation. Mediation texts in museums and galleries have been criticized for being too difficult to decipher (Serrell, 2015). Special attention needs to be paid to making the texts speak to those with other perspectives than the people making the exhibition.

Given a work of art, say an abstract painting, some people will have an easy time "getting it". This depends on the frame of reference and capabilities they bring with them to the paintings. For others, it will be an opaque jumble of lines and colors. A common way to assist such visitors is by providing contextual information about the artist, and the art historical period it was part of. This approach will help some, but it also assumes familiarity with notions of art-historical concepts, such as schools and periods. The approach we have seen taken by HCI researchers, creating platforms for user-generated content, also attempts to offer hooks for engaging with the art. Their approach offers new potential for correspondence but still comes with certain requirements. It requires users to be able and feel confident in formulating their own understanding of the work in words, and it assumes that the resulting folksonomy is in fact comprehensible by other users as a means of making sense of the *reorganizational practice* of the artwork in question. Communication is *commoning*, the outcome of a process of reaching out to others who are different and, at least for a while, travel on the same path (Ingold, 2017). The approach offered here is based on aesthetic experience and it attempts the same thing as these earlier practices of art mediation. It does not lower the bar for art appreciation as such, but it tries to offer more ways by which we may *common*.

When we consider the non-verbal aspects such as atmosphere, architecture, and soundscape we should, however, not forget that these aspects also make the world legible in certain ways, and that it requires a certain literacy to see it. Both Noë and Dewey refer to the skill it takes to perceive Dewey, 1980; Noë, 2015. Arguing for his related concept of *atmosphere*, Böhme puts it very directly when he argues that it takes *atmospheric competence* to appreciate the aesthetic qualities of the world around you:

> "To perceive atmospheres means to open oneself emotionally. This can offset the externalization of the environment and counteract the lack of contact, the coolness of modern individuals. Getting involved in atmo-

> spheres is tantamount to wanting to participate and to expose oneself to impressions – a prerequisite for the experience of pleasure in life and the discovery of one's body as a medium of being." (Böhme, 2017, p.121)

Just as learning to read takes practice, so does aesthetic attention and trusting that your aesthetic experience is important and relevant.

While recruiting people for interviews about *Poison*, it became clear that not everybody knew what to make of their experiences in the room. Some relied mostly on the little text that was present in summing up their experience, and a few, like one agitated man who did not want to be interviewed, did not know what to make of their experience at all. "This was nothing!", as he said to me. Even though most people interviewed talked about all the ways the exhibition made them think, feel, and draw relations, some visitors either did not want to or were not able to engage with this kind of exhibition. This leads to the insight that, when using non-verbal and aesthetic means to educate attention, we must still keep in mind that this will also include some and exclude others. Mortensen (2011) report on an immersive exhibition in a science museum that visitors either rationalized as an experience or as a static display. Referring to work by Belaën, Mortensen (2011) uses the terms *resonance* and *distance* to describe visitors' willingness to surrender themselves to the immersive premise of the exhibit. The *resonance* or experience group eventually achieves a larger part of the intended learning outcome of the exhibition.

By being aware of this while doing formative evaluation, we might, however, be able to adjust our designs in due time to be as relevant for the audience as possible. Finally, this is another reminder that the best approach overall might be to offer interpretive hooks of different kinds to be able to address a diverse audience in ways that are relevant for as many as possible.

The three questions I have highlighted about aura, the authority of the museum in interpretation, and the question of literacy are not new, and by no means solved by the discussion proposed here. The perspective offered here, I believe, helps us move forward in improving design and mediation practices in this domain, although it is probably not the most relevant paradigm for all other functions in the museum. For this reason it should be approached with a measure of modesty as it may challenge other professional capacities in the museum. An existing tension already exists between educators and curators in many museums, due to paradigmatic differences in their approach to art, objects, and meaning-making (Hooper-Greenhill, 2000; Samis & Michaelson, 2016). The perspectives presented in this section, I hope, might offer a way to bridge some of those differences when collaborating on design projects involving art.

Now I will turn my focus to some of the new perspectives this research offers for design.

6.2 Designing technology-based art mediation

In this project, I have worked with Ingold's *education of attention* as the concept through which to understand and evaluate learning. *Education of attention* comes from an ecological and action-oriented perspective, which makes it relevant for an enactivist perspective on design, experience, and perceptive skill-building. However, it is less appropriate for describing the construction of knowledge. Much more comprehensive theories exist about art education, museum education, and informal education.

Some of the researchers who have been trying to understand learning in the museum context are John Falk and Lynn Dierking. In their 2018 book *Learning from Museums* Falk and Dierking explain that, after working on their *Contextual Model of Learning* for more than two decades, they have realized that it must take into account *time* because learning does not happen in an instant, and evaluating what people learned as they exit the museum yields unsatisfactory and imprecise results (Falk & Dierking, 2018). Through an example with two participants from a larger study who were interviewed months after their museum visits, they illustrate the many factors on which learning is dependent. This is not only on the personal and social circumstances of a person as they enter the museum, but equally on what happens in their life after the museum visit and how things picked up in the museum are used in processes of establishing relations, comparing, and explaining that develops their perspective on the world for a long time after the visit. By this logic, learning can also occur in the museum, as a consequence of the visitor's prior experiences. However, the consequence of this perspective is that evaluating learning is a resource-intensive and time-consuming process that can only really be evaluated after the fact. Documenting how learning happened in the months after the museum visit is outside the scope of the present project, and will certainly be for many future projects aiming at supporting learning with and about art in art museums. How does that leave us who design for learning in the museum space? What can we possibly know about how our designs might lead to learning when we are in the act of designing?

As I have argued, the role of the designer is to create designs with affordances that embed a potential for action and thus for experience. Following this, learning happens as a consequence of the actions and experiences that eventually occur, and how they come to be reenacted later as knowledge with respect to new environments. This view does not support transferring specific pieces of information to visitors, which is the idea that Ingold criticizes (Ingold, 2001). On the other hand, we are able to shape the space of potential in which experiences may unfold, meaning that we can and ought to evaluate how our designs *educate the attention*. This does not constitute a final assessment of learning, but it does give us a sense of the experiences that may serve as hooks for learning down the road.

6.2.1 Evaluating education of attention

As presented earlier, previous HCI projects in art museums have evaluated learning through visitors self-reporting the amount of learning they have experienced. Others focus entirely on emergent experiences with no expectations of a communication intent embodied in the mediation technology. The education of attention perspective presented here assumes that the designers of the mediating technology want to make something legible in the world (or in the art), through educating the visitors' attention towards similar features. However, the important distinction from the information-transmission paradigm is that focus is not on specific facts and concepts being transmitted. Rather, focus is on whether visitors pay attention to certain features of the exhibition, and which experiential qualities these environmental features evoke. The success of the mediation is that visitors perceive features relevant to the mediation effort and that the resulting aesthetic experiences are in alignment with the design intent.

The praxeological approach presented by Mortensen (2011) investigating an immersive exhibition on biology, attempts to follow the visitor to conclusions such as: The cave beetle inhabiting a dark environment navigates mainly through touch, not vision. However, the visitors in Mortensen's analysis only reached the intended learning outcome partially. While Mortensen's focus on understanding the embodied experiences of the visitors is similar to the one presented in this project, the formulation of an intended learning outcome indicates an expectation of the ultimate success of the exhibition as the transfer of specific facts from the designers of the exhibition to the visitors.

In this evaluation, there is no expectation that a formal conceptualization of the type "Munch used the color green to indicate jealousy" or "The architecture of The Green Room across the paintings in the series is unstable" happens right there and then in the exhibition, or that the type of learning that happens can be formulated verbally in a meaningful way. Rather, the ultimate success is a *perceptive* approach to the mediation technology and a general building of competence in making sense of Munch and comparable artists' aesthetic universe.

In "Educating the Attention of Museum Visitors through Non-verbal Art Mediation" (Sivertsen, Mathias, et al., 2023) we have identified three questions that we used to evaluate the *Poison* exhibition. Similar to Bitgood (2016), we assume that a certain level of attentiveness must occur as a foundation for enabling personal experiences with the exhibition. We frame this using Dewey's distinction between *recognition* and *perception*(Dewey, 1980). Do the visitors label it and resist active engagement, or do they engage in active perceptual action? The second question is what aesthetic qualities the visitors then report. It is the assumption that visitors who mostly *recognize* will either have trouble answering these questions, or they will resort to superficial descriptions that reflect their understanding of the category of experience. Visitors

who have been engaged in active *perception* will likely be better at giving detailed accounts of specific moments where aesthetic and affective experiences happened. The last question involves the personal context of the visitor. It is important that this happens at the end of the interview, as it may lead to a distancing from the experience, and is as such somewhat at odds with the phenomenological method. However, the value of shifting to this perspective is in understanding what kind of connections and relations the visitor already established during the visit.

1. Are visitors accepting and open to the design or do they reject it or label it off the cuff before properly experience it?

2. What are the aesthetic qualities that the visitors report? Are they in correspondence with the design intent?

3. How do visitors place their experience in relation to their existing knowledge on the topic/artwork/artist?

As described in the method section, the practical methodology centers around a phenomenological interviewing method but is combined with other interviewing methods to gather all relevant perspectives. Their concrete experiences enable formative design insights, helping us understand how specific qualities of the exhibition evoke certain experiences and relations in concrete situations.

Again the aim is to be able to outline the space of potential that the exhibition affords. Therefore accounts that cluster around similar descriptions of the experiences are interesting as they might be indications of the ”mainstream” experience. However, diverging experiences are equally valuable because they help indicate the possible spread between experiences. Here, the accounts being tied to concrete situations are very important as they might help us understand whether the divergence arises due to differences in what people bring to the exhibitions or differences in what visitors' attention was educated to focus on.

This methodology has only been employed to the full extent in the evaluation of *Poison*. The first and third questions are heavily influenced by the research design, meaning that, in the evaluation of *New Snow*, these questions are less relevant. The participants in a lab-type experiment are knowingly being observed and are therefore under pressure to engage with the design presented to them. Furthermore, they might not have any meaningful frame of reference to Munch's artistic practice compared to visitors in a museum dedicated to Munch, thus the answers to these questions are less interesting outside the museum context. However, the second question on the aesthetic qualities has been essential to the evaluation of both *New Snow* and *Handling Artworks*.

6.3 Ambiguity as a resource for art mediation

While Noë does not mention it specifically, a related and maybe inspiring idea, appears in Dewey's "Art as Experience", when he distinguishes between *recognizing* and *perceiving*. Recognizing is the quick identification of something as belonging to a category that is already known and understood. The active perceiving is a continuous movement with the object of perception, which is a prerequisite for *an aesthetic experience*.

When Noë talks about the artworks not losing their difficulty, I read it as them being recognized as something preconceived, meaning that there is now no longer a reason to *perceive* and try to get what it might be trying to reorganize. There is nothing reorganizational about things that are already well-known and labeled.

Taking this to art mediation might be comparable to showing an artwork with the text "Artwork title: The artist was thinking of X when they made it. It is a comment on Y. It relates to work Z" or "Artwork title: Pay attention to X. What do you see? Have you seen anything like it in the museum?" The second text is well aligned with the interpretive turn and invites participants to invest themselves and their prior experience in perceiving the image, while the first text categorizes and puts into boxes. Now, as we know, the description is not alone in constituting the art experience, although we can imagine how these two texts establish different potentials for experience.

Moving away from text, ambiguity in design is understood to do some of the same actions: calling into question, opening for multiple interpretations, and inviting the user to similar engagements as the text example above. For this reason, *ambiguity of information* (Gaver et al., 2003) plays a significant role in both *Poison* and *New Snow*. Both projects make use of imagery generated using deep generative models that partially recreate the visual aspects of original artworks, while interspersing artifacts and a general blurriness. This is used to call into question these very images and their authority, but they are, as we see, not meaningless. The participants do engage and attempt to look through them, attempting to make sense of the perspective they are constituting on the artwork.

6.3.1 Generative AI as a tool for art mediation

In *Poison*, we focus on the seven paintings from The Green Room series, and in *New Snow* the interface lets the user traverse through Munch's drawing practice. In both cases, we have utilized generative AI to create the final aesthetic expression of the reenactments.

In the main room of *Poison*, neural style transfer was used on 3D reenactments of the paintings, and in the triptych room the virtual paintings depicted a morphing from one painting to another. In *New Snow*, users traverse the latent space of a corpus of drawings through their own drawing actions.

In total, there are three different applications of generative AI that, through evoking ambiguity of information, aim at casting the reenactments as *uncertain entities* (Sivertsen & Løvlie, n.d.). In *Poison* the ambiguity supports that the visitors do not leave the exhibition with the feeling that they have seen an accurate reproduction of the paintings. Rather they have been exposed to many visual characteristics, without getting the "full picture". Simultaneously, it emphasizes that *The Green Room* is an ambiguous series of paintings, that raises many questions, even to art researchers. In *New Snow* the ephemeral and ever-changing nature of the re-enactments serves the purpose of indicating that none of what the system produces is to be taken as a representation of a concrete physical drawing. All depictions are manifestations of the potential within the practice of Munch, and it is through the movements along different dimensions that the relations in the practice are revealed.

Generative AI is but one technology with affordances for creating ambiguous output. Unlike traditional statistic analysis, the output of generative AI has aesthetic quality in itself, and through that it can direct attention to trends by mimicking them rather than drawing them as lines in a coordinate system.

In the context of the art museum, we can rely on the original artworks themselves to present on their own details and particularities. Digital reproductions should play a supporting role. Thus, we can use Generative AI to make legible trends and bias in corpora of artworks that we are interested in staging for the audience (Benjamin et al., 2021). We can utilize ambiguity as an aesthetic quality to temper the authority that we give to generative AI.

If we want to use the technology for this purpose, the curation of data is highly important. The dataset determines the ontology, and the interface determines the epistemology. A trained model will not be better than its data at representing the qualities we are interested in. Therefore a significant amount of work goes into figuring out what to put on display, which data will support it, how it must be pre-processed and finally how to enable visitors access to it. The time it takes to collect data, pre-process and train a model makes the process of iteratively developing machine-learning based systems very resource-intensive. For every adjustment to the model, the dataset or training parameters must be adjusted and the training process started anew. For reference, the process of training a model for *New Snow*, not accounting for data-labeling, takes approximately four days.

Through the design of the interface, we can enable people to explore how the model has synthesized the rules, trends, and patterns inherent in the data. We have called this *designing for reflexive use* (Sivertsen & Løvlie, n.d.), and it involves designing the interface to enable insight into the model and its expression of the underlying data. This is in contrast to interfaces that focus on affording the production of new visual projects, which are typically found in the commercial deployments of generative AI.

To create an interface that enables the exploration of a generative machine learning model, three different qualities are important to take into consideration:

Prompting modality: The first step is determine what kind of input is used to prompt the model to generate an output. Is it text, images, shapes or something entirely different. This choice has implications for the way of prompting and thus for the epistemological approach to the model. In *New Snow*, the prompting modality are black and white drawings that allow for exploration along dimensions of stroke length, dynamics, density and location. Compare this to a text-based prompt that would enable prompting along conceptual dimensions such as house, home, dwelling and conjugations: houses, homes.

For the purpose of exploration, we argue that *incremental prompting* allows for a gradual movement across dimensions in the latent space of the model. This makes it easier to understand how different changes affect the output of the system.

Finally, *fast updates* allow for a fluid and responsive dialog with the model, where each prompt can happen in immediate response to the previous, without having to stop and wait for the model to respond.

By being attentive to these qualities, we can design systems that let users investigate the aesthetic qualities of generative models through their expressive capabilities. With a balanced use of ambiguity, this technology offers an approach through which we can *correspond with* art. It emphasizes breadth rather than depth, but can be relevant when working with artists with expansive practices like Munch or across schools or periods of art history.

7 Future Work

In this section, I will describe three research trajectories that extend some of the questions raised in the project.

7.1 Studying the impact on artwork encounters

The core motivation for this project is that the technology-based mediation has the potential to affect visitors' encounters with original art. In the results presented in this project, I have shown that visitors become excited and curious around the originals which *Poison* refers to, but it has not been possible to study exactly what happens in the encounters with the originals when people experience *Poison* or *New Snow* first.

The practical requirements for a study of this kind are complicated. It requires a research design where either *Poison* or *New Snow* are deployed in the museum, and the presence of the relevant originals. Next, it is necessary to recruit visitors with different trajectories through the museum. This should include people who went from technology-based mediation design to originals, vice versa, and potentially a control group.

A study of this kind would enable a deeper understanding of exactly how the perspectives presented in the technology-based mediation design interplay with art encounters, whether the mediation supports interpretive practice and knowledge construction, and to what extent we as designers can control how that happens.

7.2 Education of attention in cultural heritage beyond art

The concept of *education of attention* is not specific to the art museum domain. In this project, it has been conceptualized as a form of art critique. While that term is hardly applicable outside the art domain, the idea of directing attention to matters of interest is generally applicable.

Considering this in a broader cultural heritage context, the pressing question is naturally to investigate what kind of aspects are relevant to mediate using technology-based means. In the broader field of HCI and technology in cultural heritage, there are more examples of work that uses immersive and multi-sensorial means to give visitors new perspectives on historical objects, intangible heritage and cultural sites

(see for example (Claisse et al., 2020; Hornecker & Ciolfi, 2019; Kenderdine, 2015; Kenderdine et al., 2014)). I believe this highlights similarities between art museum mediation, other museum types, and cultural heritage sites and further dismisses the idea that art museums are special places, where concepts like experience and learning somehow work differently than in other places.

Applying the concept of *education of attention* to other domains would require developing new ways of *corresponding with* the domain, and thus new constellations of technology to support this. This work is necessary to find the limits of the applicability of this concept and strengthen methods of structuring design processes and evaluation around it.

7.3 Generative AI as a mediation tool for large collections of data

The third research trajectory is that of investigating how generative AI might be used as a means for mediating large datasets. That is, using generative AI not for the purpose of creating new content, but for the purpose of investigating aesthetic qualities embedded in a dataset. With the *New Snow* project, we present and take the first steps in developing this idea. We find that there is a potential to create a hermeneutic relation to the data. This means that the table makes the data available for ”reading”. However, the specific way in which the data becomes legible is highly dependent on the training process. In the *New Snow* process, we have only been able to experiment with a single way to pre-process the data and train the model, which leaves a lot to be explored. Through further experimentation with the training process and interface design, the circumstances under which such an approach to a dataset might be relevant, and how to best support it could be unfolded. This research interest goes beyond the museum and art domain but focuses on the specific qualities of generative AI as an emerging technology across many application areas.

8 Conclusion

In my PhD project, through my publications and in this thesis, I have established a theoretical foundation for understanding the potential for using technology for mediation in art museums. I have presented a theoretical view on art mediation as helping visitors to see what they might otherwise ignore or fail to give weight to. That is, attuning their perceptual capabilities to aspects of the environment that are relevant for the interpretation of, and *corresponding with*, artworks. This particular perspective is by no means limited to this domain, but has been developed from within it and applied here.

This theoretical perspective has been applied to three different projects. First, the experiment where we see how the embodied mediation of paintings shapes participants' understanding of them as different cultural objects that are not necessarily commensurate with them being artworks. Second, in an immersive exhibition, in which we explore the potential of using primarily non-verbal means to educate the visitors' attention to aspects of the art that are deemed particularly relevant from an art historical perspective. Lastly, in the drawing table example, we explore the potential of letting the perspective emerge from large collections of digitalized artworks through the use of generative AI.

I have worked with immersive exhibitions and generative AI, two technologies that pose particular challenges for art museums, and I have shown how they can be used productively to support the goal of art museums when used with attention to how they work as mediating technologies. Working with generative AI has also opened up questions about how we understand this emerging technology. First, how it mediates and its potential for exposing the aesthetic qualities of large datasets. Second, how the aspect of ambiguity is very relevant for working with art and machine learning.

I have investigated the emerging experiences of participants, visitors, and users when meeting such designs, and presented a way of evaluating with regards to how a design *educates attention*. Through that, I have shown that both verbal and non-verbal aspects should be considered as contributing to the mediation and that it is possible to make people interested in original artworks through experiences that are exciting on their own.

I offer this theoretical perspective and the insights from the design projects to art museums as well as educators, designers and design researchers working within them in the hope that it will help them support the visitors in having great experiences

with art. Great art can be one of the great pleasures of life, and through making art mediation with the expressive means of art rather than dull scientific statements we can offer more people a chance to *see* it.

Bibliography

Ablart, D., Velasco, C., Vi, C. T., Gatti, E., & Obrist, M. (2017). The how and why behind a multisensory art display. *Interactions*, *24*(6), 38–43. https://doi.org/10.1145/3137091

Alexander, J., Wienke, L., & Tiongson, P. (2017). Removing the barriers of Gallery One: A new approach to integrating art, interpretation, and technology. *Museums and the Web 2017: Selected Papers and Proceedings from an International Conference*. Retrieved December 15, 2021, from https://mw17.mwconf.org/paper/removing-the-barriers-of-gallery-one-a-new-approach-to-integrating-art-interpretation-and-technology/

Andersen, K., & Ward, N. (2017). Learning from the Crackle Exhibition. *Proceedings of the Eleventh International Conference on Tangible, Embedded, and Embodied Interaction*, 225–231. https://doi.org/10.1145/3024969.3024997

Audry, S. (2021). *Art in the age of machine learning*. The MIT Press.

Bardzell, J., & Bardzell, S. (2015). *Humanistic HCI* [OCLC: 930589099]. Morgan & Claypool Publishers.

Bedford, L. (2014). *The Art of Museum Exhibitions: How Story and Imagination Create Aesthetic Experiences*. Left Coast Press.

Benford, S., Løvlie, A. S., Ryding, K., Rajkowska, P., Bodiaj, E., Paris Darzentas, D., Cameron, H., Spence, J., Egede, J., & Spanjevic, B. (2022). Sensitive Pictures: Emotional Interpretation in the Museum. *CHI Conference on Human Factors in Computing Systems*, 1–16. https://doi.org/10.1145/3491102.3502080

Benjamin, J. J., Berger, A., Merrill, N., & Pierce, J. (2021). Machine Learning Uncertainty as a Design Material: A Post-Phenomenological Inquiry. *Proceedings of the 2021 CHI Conference on Human Factors in Computing Systems*, 1–14. https://doi.org/10.1145/3411764.3445481

Benjamin, W. (1936). The work of art in the age of mechanical reproduction. *Visual Culture: Experiences in Visual Culture*, 144–137.

Bimber, O., Coriand, F., Kleppe, A., Bruns, E., Zollmann, S., & Langlotz, T. (2005). Superimposing pictorial artwork with projected imagery. *ACM SIGGRAPH 2005 Courses*, 6–es. https://doi.org/10.1145/1198555.1198716

Birchfield, D., Mechtley, B., Hatton, S., & Thornburg, H. (2008). Mixed-reality learning in the art museum context. *Proceedings of the 16th ACM international conference on Multimedia*, 965–968. https://doi.org/10.1145/1459359.1459534

Birchfield, D., Thornburg, H., Megowan-Romanowicz, M. C., Hatton, S., Mechtley, B., Dolgov, I., & Burleson, W. (2008). Embodiment, Multimodality, and Composition: Convergent Themes across HCI and Education for Mixed-Reality Learning Environments [Publisher: Hindawi]. *Advances in Human-Computer Interaction*, *2008*, e874563. https://doi.org/10.1155/2008/874563

Bitgood, S. (2010). An Attention-Value Model of Museum Visitors.

Bitgood, S. (2016). *Attention and Value: Keys to Understanding Museum Visitors*. Routledge. https://doi.org/10.4324/9781315433455

Bødker, S. (2006). When second wave HCI meets third wave challenges. *Proceedings of the 4th Nordic conference on Human-computer interaction: changing roles*, 1–8. https://doi.org/10.1145/1182475.1182476

Boehner, K., Sengers, P., Medynskiy, Y., & Gay, G. (2005). Opening the Frame of the Art Museum: Technology between Art and Tool. *Digital Arts and Culture (DAC)*, 123–132.

Boehner, K., Thom-Santelli, J., Zoss, A., Gay, G., Hall, J. S., & Barrett, T. (2005). Imprints of place: Creative expressions of the museum experience. *CHI '05 Extended Abstracts on Human Factors in Computing Systems*, 1220–1223. https://doi.org/10.1145/1056808.1056881

Böhme, G. (2017). *Atmospheric architectures: The aesthetics of felt spaces* (A.-C. Engels-Schwarzpaul, Trans.). Bloomsbury Academic, an imprint of Bloomsbury Publishing Plc.

Bolander, E., Ridenour, H., & Quimby, C. (2018). *Art Museums and Technology: Developing New Metrics to Measure Visitor Engagement* (tech. rep.). The Cleveland Museum of Art.

Budge, K. (2018). Visitors in immersive museum spaces and Instagram: Self, placemaking, and play. *The Journal of Public Space*, *3*(3), 121–138. https://doi.org/10.32891/jps.v3i3.534

Buxton, W. (2007). *Sketching user experiences: Getting the design right and the right design*. Elsevier/Morgan Kaufmann.

Campbell, C. (2019). Curator's Introduction to 'Leonardo: Experience a Masterpiece'.

Christensen, H. D. (2017). The Art of Copying: Five strategies for Transforming Originals in the Art Museum [Number: 1]. *Culture Unbound*, *9*(1), 85–107. https://doi.org/10.3384/cu.2000.1525.179185

Christidou, D., & Pierroux, P. (2019). Art, touch and meaning making: An analysis of multisensory interpretation in the museum. *Museum Management and Curatorship*, *34*(1), 96–115. https://doi.org/10.1080/09647775.2018.1516561

Claisse, C., Petrelli, D., Ciolfi, L., Dulake, N., Marshall, M. T., & Durrant, A. C. (2020). Crafting Critical Heritage Discourses into Interactive Exhibition Design. *Proceedings of the 2020 CHI Conference on Human Factors in Computing Systems*, 1–13. https://doi.org/10.1145/3313831.3376689

Correia, N., Mota, T., Nóbrega, R., Silva, L., & Almeida, A. (2010). A multi-touch tabletop for robust multimedia interaction in museums. *ACM International Conference on Interactive Tabletops and Surfaces*, 117–120. https://doi.org/10.1145/1936652.1936674

Cosley, D., Baxter, J., Lee, S., Alson, B., Nomura, S., Adams, P., Sarabu, C., & Gay, G. (2009). A tag in the hand: Supporting semantic, social, and spatial navigation in museums. *Proceedings of the SIGCHI Conference on Human Factors in Computing Systems*, 1953–1962. https://doi.org/10.1145/1518701.1518999

Cosley, D., Lewenstein, J., Herman, A., Holloway, J., Baxter, J., Nomura, S., Boehner, K., & Gay, G. (2008). ArtLinks: Fostering social awareness and reflection in museums. *Proceedings of the SIGCHI Conference on Human Factors in Computing Systems*, 403–412. https://doi.org/10.1145/1357054.1357121

Costello, B., Muller, L., Amitani, S., & Edmonds, E. (2005). Understanding the Experience of Interactive Art: Iamascope in Beta_space. *The Second Australasian Conference on Interactive Entertainment*, *123*, 49–56.

De Gemmis, M., Lops, P., Semeraro, G., & Basile, P. (2008). Integrating tags in a semantic content-based recommender. *Proceedings of the 2008 ACM conference on Recommender systems*, 163–170. https://doi.org/10.1145/1454008.1454036

Dewey, J. (1980). *Art as Experience*. Wideview/Perigree.

Dove, G., Halskov, K., Forlizzi, J., & Zimmerman, J. (2017). UX Design Innovation: Challenges for Working with Machine Learning as a Design Material. *Proceedings of the 2017 CHI Conference on Human Factors in Computing Systems*, 278–288. https://doi.org/10.1145/3025453.3025739

Duncan, C. (2005). The art museum as ritual [Publisher: Routledge London and New York]. *Heritage, museums and galleries: an introductory reader*, 78–88.

Endresen, S. (2015). *Serial Experiments. Close-Readings of Edvard Munch's Det grønne værelset (1907)* (Doctoral dissertation). Universitetet i Oslo. Oslo.

Falk, J. H., & Dierking, L. D. (2018). *Learning from Museums*. Rowman & Littlefield.

Fosh, L., Benford, S., Reeves, S., Koleva, B., & Brundell, P. (2013). See me, feel me, touch me, hear me: Trajectories and interpretation in a sculpture garden. *Proceedings of the SIGCHI conference on human factors in computing systems*, 149–158.

Frauenberger, C. (2019). Entanglement HCI The Next Wave? *ACM Transactions on Computer-Human Interaction*, *27*(1), 1–27. https://doi.org/10.1145/3364998

Gaver, W. W., Beaver, J., & Benford, S. (2003). Ambiguity as a resource for design. *Proceedings of the SIGCHI Conference on Human Factors in Computing Systems*, 233–240. https://doi.org/10.1145/642611.642653

Gilman, B. I. (1916). Museum Fatigue [Publisher: American Association for the Advancement of Science]. *The Scientific Monthly*, *2*(1), 62–74. Retrieved December 10, 2020, from https://www.jstor.org/stable/6127

Gonçalves, L., Campos, P., & Sousa, M. (2012). M-dimensions: A framework for evaluating and comparing interactive installations in museums. *Proceedings of the 7th Nordic Conference on Human-Computer Interaction Making Sense Through Design - NordiCHI '12*, 59. https://doi.org/10.1145/2399016.2399027

Harrington, M. C. R. (2020). Connecting User Experience to Learning in an Evaluation of an Immersive, Interactive, Multimodal Augmented Reality Virtual Diorama in a Natural History Museum & the Importance of Story. *2020 6th International Conference of the Immersive Learning Research Network (iLRN)*, 70–78. https://doi.org/10.23919/iLRN47897.2020.9155202

Harrison, S., Tatar, D., & Sengers, P. (2007). The Three Paradigms of HCI. *Alt. Chi. Session at the SIGCHI Conference on Human Factors in Computing Systems*.

Hauser, S., Wakkary, R., Odom, W., Verbeek, P.-P., Desjardins, A., Lin, H., Dalton, M., Schilling, M., & De Boer, G. (2018). Deployments of the table-non-table: A Reflection on the Relation Between Theory and Things in the Practice of Design Research. *Proceedings of the 2018 CHI Conference on Human Factors in Computing Systems*, 1–13. https://doi.org/10.1145/3173574.3173775

Hindmarsh, J., Heath, C., vom Lehn, D., & Cleverly, J. (2002). Creating Assemblies: Aboard the Ghost Ship. *Proceedings of the 2002 ACM conference on Computer supported cooperative work - CSCW '02*, 156. https://doi.org/10.1145/587078.587101

Hooper-Greenhill, E. (2000). Changing Values in the Art Museum: Rethinking communication and learning. *International Journal of Heritage Studies*, *6*(1), 9–31. https://doi.org/10.1080/135272500363715

Hornecker, E., & Ciolfi, L. (2019). Human-Computer Interactions in Museums. *Synthesis Lectures on Human-Centered Informatics*, *12*(2), i–153. https://doi.org/10.2200/S00901ED1V01Y201902HCI042

Hou, Y., Kenderdine, S., Picca, D., Egloff, M., & Adamou, A. (2022). Digitizing Intangible Cultural Heritage Embodied: State of the Art. *Journal on Computing and Cultural Heritage*, *15*(3), 1–20. https://doi.org/10.1145/3494837

Howes, D. (2014). Introduction to Sensory Museology. *The Senses and Society*, *9*(3), 259–267. https://doi.org/10.2752/174589314X14023847039917

Hsi, S. (2003). A study of user experiences mediated by nomadic web content in a museum. *Journal of Computer Assisted Learning*, *19*(3), 308–319. https://doi.org/https://doi.org/10.1046/j.0266-4909.2003.jca_023.x

Hüllermeier, E., & Waegeman, W. (2021). Aleatoric and epistemic uncertainty in machine learning: An introduction to concepts and methods. *Machine Learning*, *110*(3), 457–506. https://doi.org/10.1007/s10994-021-05946-3

Ingold, T. (2001). From the transmission of representation to the education of attention. In *The debated mind: Evolutionary psychology versus ethnography* (pp. 113–153). Berg.

Ingold, T. (2017). On human correspondence. *Journal of the Royal Anthropological Institute*, *23*(1), 9–27. https://doi.org/10.1111/1467-9655.12541

Ingold, T. (2018). *Anthropology and/as education*. Routledge.

Jeffrey, S. (2015). Challenging Heritage Visualisation: Beauty, Aura and Democratisation [Publisher: De Gruyter Open Access]. *Open Archaeology*, *1*(1). https://doi.org/10.1515/opar-2015-0008

Jones, S., Jeffrey, S., Maxwell, M., Hale, A., & Jones, C. (2018). 3D heritage visualisation and the negotiation of authenticity: The ACCORD project. *International Journal of Heritage Studies*, *24*(4), 333–353. https://doi.org/10.1080/13527258.2017.1378905

Kenderdine, S. (2015). Embodiment, Entanglement, and Immersion in Digital Cultural Heritage. In S. Schreibman, R. Siemens, & J. Unsworth (Eds.), *A New Companion to Digital Humanities* (pp. 22–41). John Wiley & Sons, Ltd. https://doi.org/10.1002/9781118680605.ch2

Kenderdine, S., Chan, L. K. Y., & Shaw, J. (2014). *Pure Land*: Futures for Embodied Museography. *Journal on Computing and Cultural Heritage*, *7*(2), 1–15. https://doi.org/10.1145/2614567

Kenderdine, S., & Yip, A. (2018). The proliferation of aura: Facsimiles, authenticity and digital objects. In K. Drotner, V. Dziekan, R. Parry, & K. C. Schrøder (Eds.), *The Routledge Handbook of Museums, Media and Communication* (1st ed., pp. 274–289). Routledge. https://doi.org/10.4324/9781315560168

Kidd, J. (2018). "Immersive" Heritage Encounters. *The Museum Review*, *3*(1), 16.

Kiran, A. H. (2015). Four Dimensions of Technological Mediation. In R. Rosenberger & P.-P. Verbeek (Eds.), *Postphenomenological Investigations: Essays on Human–Technology Relations*. Lexington Books.

Kortbek, K. J., & Grønbæk, K. (2008). Communicating Art through Interactive Technology: New Approaches for Interaction Design in Art Museums. *NordiCHI '08: Proceedings of the 5th Nordic conference on Human-computer interaction: building bridges*, 229–238. https://doi.org/10.1145/1463160.1463185

Krogh, P. G., & Koskinen, I. (2020). *Drifting by Intention: Four Epistemic Traditions from within Constructive Design Research*. Springer International Publishing. https://doi.org/10.1007/978-3-030-37896-7

Krumpen, S., Klein, R., & Weinmann, M. (2021). Towards Tangible Cultural Heritage Experiences—Enriching VR-based Object Inspection with Haptic Feedback. *Journal on Computing and Cultural Heritage*, *15*(1), 19:1–19:17. https://doi.org/10.1145/3470470

Lange, V., van Beuzekom, M., Hansma, M., Jeurens, J., van den Oever, W., Regterschot, M., Treffers, J., van Turnhout, K., Ibrahim Sezen, T., Iurgel, I., & Bakker, R. (2019). Blending into the White Box of the Art Museum. *Proceedings of the Halfway to the Future Symposium 2019*, 1–10. https://doi.org/10.1145/3363384.3363469

Latour, B., & Lowe, A. (2011). The Migration of the Aura Exploring the Original Through Its Fac similes. In T. Bartscherer & R. Coover (Eds.), *Switching Codes: Thinking Through Digital Technology in the Humanities and the Arts* (pp. 275–297). University of Chicago Press.

Lim, Y.-K., Stolterman, E., & Tenenberg, J. (2008). The anatomy of prototypes: Prototypes as filters, prototypes as manifestations of design ideas. *ACM Transactions on Computer-Human Interaction*, *15*(2), 1–27. https://doi.org/10.1145/1375761.1375762

Løvlie, A. S., Ryding, K., Spence, J., Rajkowska, P., Waern, A., Wray, T., Benford, S., Preston, W., & Clare-Thorn, E. (2021). Playing games with Tito: Designing hybrid museum experiences for critical play [arXiv:2101.12533 [cs]]. *Journal on Computing and Cultural Heritage*, *14*(2), 1–26. https://doi.org/10.1145/3446620

Lyons, L. (2009). Designing opportunistic user interfaces to support a collaborative museum exhibit. *Proceedings of the 9th international conference on Computer supported collaborative learning - Volume 1*, 375–384.

Macdonald, S. (2007). Interconnecting: Museum visiting and exhibition design. *CoDesign*, *3*(sup1), 149–162. https://doi.org/10.1080/15710880701311502

Malevé, N. (2021). On the data set's ruins. *AI & SOCIETY*, *36*(4), 1117–1131. https://doi.org/10.1007/s00146-020-01093-w

Mathias, N. (2022). Meta-artistic immersion in digital exhibitions. History – mobilization – spectatorship. *Journal of Aesthetics \& Culture*, *14*(1), 2129160. https://doi.org/10.1080/20004214.2022.2129160

McCarthy, J., & Wright, P. (2004). *Technology as experience*. MIT Press.

Meehan, N. (2022). Digital Museum Objects and Memory: Postdigital Materiality, Aura and Value. *Curator: The Museum Journal*, *65*(2), 417–434. https://doi.org/10.1111/cura.12361

Mondloch, K. (2022). The Influencers: Van Gogh Immersive Experiences and the Attention-Experience Economy. *Arts*, *11*(5), 90. https://doi.org/10.3390/arts11050090

Morse, C., Niess, J., Bongard-Blanchy, K., Rivas, S., Lallemand, C., & Koenig, V. (2022). Impressions that last: Representing the meaningful museum experience. *Behaviour & Information Technology*, 1–28. https://doi.org/10.1080/0144929X.2022.2061375

Mortensen, M. F. (2011). Analysis of the Educational Potential of a Science Museum Learning Environment: Visitors' experience with and understanding of an immersion exhibit. *International Journal of Science Education*, *33*(4), 517–545. https://doi.org/10.1080/09500691003754589

Noë, A. (2015). *Strange Tools: Art and Human Nature*. Hill; Wang.

Noë, A. (2021). *Learning to Look: Dispatches from the Art World*. Oxford University Press.

Noë, A., Bertram, G. W., Davies, D., Leddy, T., & Viola, T. (2020). On Alva Noë, Strange tools. Art and human nature [Number: 18]. *Studi di estetica*, *3*(18), 249–295. Retrieved February 3, 2022, from https://journals.mimesisedizioni.it/index.php/studi-di-estetica/article/view/898

Pau, S. (2017). Audio that Moves You: Experiments with Location-aware storytelling in the SFMOMA app –. Retrieved February 1, 2023, from http://mw17.mwconf.org/proposal/audio-that-moves-you-location-aware-storytelling-in-the-sfmoma-app/index.html

Petitmengin, C. (2006). Describing one's subjective experience in the second person: An interview method for the science of consciousness. *Phenomenology and the Cognitive Sciences*, *5*(3-4), 229–269. https://doi.org/10.1007/s11097-006-9022-2

Petrelli, D., Ciolfi, L., van Dijk, D., Hornecker, E., Not, E., & Schmidt, A. (2013). Integrating material and digital: A new way for cultural heritage. *Interactions*, *20*(4), 58–63. https://doi.org/10.1145/2486227.2486239

Pursey, T., & Lomas, D. (2018). Tate Sensorium: An experiment in multisensory immersive design. *The Senses and Society*, *13*(3), 354–366. https://doi.org/10.1080/17458927.2018.1516026

Rosenberger, R., & Verbeek, P.-P. (2015). A Field Guide to Post-Phenomenology. In R. Rosenberger & P.-P. Verbeek (Eds.), *Postphenomenological Investigations: Essays on Human-Technology Relations* (pp. 9–41). Lexington Books.

Ryding, K. (2020). The Silent Conversation: Designing for Introspection and Social Play in Art Museums. *Proceedings of the 2020 CHI Conference on Human Factors in Computing Systems*, 1–10. https://doi.org/10.1145/3313831.3376357

Ryding, K., Spence, J., Løvlie, A. S., & Benford, S. (2021). Interpersonalizing Intimate Museum Experiences [Publisher: Taylor & Francis]. *International Journal of Human–Computer Interaction*, *37*(12), 1151–1172. https://doi.org/10.1080/10447318.2020.1870829

Samis, P. (2018). Revisiting the utopian promise of interpretive media: An autoethnographic analysis drawn from art museums, 1991–2017. In K. Drotner, V. Dziekan, R. Parry, & K. C. Schrøder (Eds.), *The Routledge Handbook of Museums, Media and Communication* (1st ed., pp. 47–66). Routledge. https://doi.org/10.4324/9781315560168

Samis, P., & Michaelson, M. (2016). *Creating the Visitor-Centered Museum*. Routledge. https://doi.org/10.4324/9781315531014

Serrell, B. (2015). *Exhibit Labels: An Interpretive Approach*. Rowman & Littlefield.

Sivertsen, C., Haas, R., Jensen, H. H., & Løvlie, A. S. (2023). Exploring a Digital Art Collection through Drawing Interactions with a Deep Generative Model. *Extended Abstracts of the 2023 CHI Conference on Human Factors in Computing Systems*, 1–5. https://doi.org/10.1145/3544549.3583902

Sivertsen, C., & Løvlie, A. S. (2021). Handling Digital Reproductions of Artworks. *Journal of Somaesthetics*, *7*(2), 21.

Sivertsen, C., & Løvlie, A. S. (n.d.). Exploring Aesthetic Qualities of Deep Generative Models through Technological (Art) Mediation. *ACM Transactions on Computer-Human Interaction*.

Sivertsen, C., Mathias, N., & Løvlie, A. S. (2023). Educating the attention of museum visitors through non-verbal art mediation. *Proceedings of the 10th Congress of the International Association of Societies of Design Research*.

Sivertsen, C., Salimbeni, G., Løvlie, A. S., Benford, S., & Zhu, J. (n.d.). Machine Learning Processes As Sources of Ambiguity: Insights from AI Art. *Proceedings of the 2024 CHI Conference on Human Factors in Computing Systems*.

Sivertsen, C., Smith, M., & van der Zwan, S. (2023). Art Critique by Other Means. *Designing Interactive Systems Conference (DIS '23)*, 10. https://doi.org/10.1145/3563657.3596069

Snibbe, S. S., & Raffle, H. S. (2009). Social immersive media: Pursuing best practices for multi-user interactive camera/projector exhibits. *Proceedings of the SIGCHI Conference on Human Factors in Computing Systems*, 1447–1456. https://doi.org/10.1145/1518701.1518920

Spence, J., Bedwell, B., Coleman, M., Benford, S., Koleva, B. N., Adams, M., Row Farr, J., Tandavanitj, N., & Løvlie, A. S. (2019). Seeing with New Eyes: Designing for In-the-Wild Museum Gifting. *Proceedings of the 2019 CHI Conference on Human Factors in Computing Systems*, 1–13. https://doi.org/10.1145/3290605.3300235

Steier, R. (2014). Posing the Question: Visitor Posing as Embodied Interpretation in an Art Museum. *Mind, Culture, and Activity*, *21*(2), 148–170. https://doi.org/10.1080/10749039.2013.878361

Stogner, M. B. (2011). "The Immersive Cultural Museum Experience – Creating Context and Story with New Media Technology". *The International Journal of the Inclusive Museum*, *3*(3), 117–130. https://doi.org/10.18848/1835-2014/CGP/v03i03/44339

Stolterman, E. (2008). The Nature of Design Practice and Implications for Interaction Design Research.

Stolterman, E., & Wiberg, M. (2010). Concept-Driven Interaction Design Research. *Human–Computer Interaction*, *25*(2), 95–118. https://doi.org/10.1080/07370020903586696

Tennent, P., Martindale, S., Benford, S., Darzentas, D., Brundell, P., & Collishaw, M. (2020). Thresholds: Embedding Virtual Reality in the Museum. *Journal on Computing and Cultural Heritage*, *13*(2), 1–35. https://doi.org/10.1145/3369394

Terrenghi, L., & Zimmermann, A. (2004). Tailored audio augmented environments for museums. *Proceedings of the 9th international conference on Intelligent user interface - IUI '04*, 334. https://doi.org/10.1145/964442.964523

The Cleveland Museum of Art. (2012). ARTLENS Gallery. Retrieved August 9, 2023, from https://www.clevelandart.org/artlens-gallery/about

Thompson, C. J., Locander, W. B., & Pollio, H. R. (1989). Putting Consumer Experience Back into Consumer Research: The Philosophy and Method of Existential-Phenomenology. *Journal of Consumer Research*, *16*(2), 133–146. https://doi.org/10.1086/209203

van Dijk, E. M., Lingnau, A., & Kockelkorn, H. (2012). Measuring enjoyment of an interactive museum experience. *Proceedings of the 14th ACM international conference on Multimodal interaction*, 249–256. https://doi.org/10.1145/2388676.2388728

Vi, C. T., Ablart, D., Gatti, E., Velasco, C., & Obrist, M. (2017). Not just seeing, but also feeling art: Mid-air haptic experiences integrated in a multisensory art exhibition. *International Journal of Human-Computer Studies*, *108*, 1–14. https://doi.org/10.1016/j.ijhcs.2017.06.004

Waern, A., & Løvlie, A. S. (Eds.). (2022). *Hybrid museum experiences: Theory and design*. Amsterdam University Press. Retrieved May 11, 2022, from https://www.jstor.org/stable/10.2307/j.ctv2cxx8x6

Walter, T. (1996). From museum to morgue? Electronic guides in Roman Bath. *Tourism Management*, *17*(4), 241–245. https://doi.org/10.1016/0261-5177(96)00015-5

Wang, Y., Stash, N., Aroyo, L., Hollink, L., & Schreiber, G. (2009). Semantic relations for content-based recommendations. *Proceedings of the fifth international conference on Knowledge capture*, 209–210. https://doi.org/10.1145/1597735.1597786

Wessel, D., & Mayr, E. (2007). Potentials and Challenges of Mobile Media in Museums [Number: 1]. *International Journal of Interactive Mobile Technologies (iJIM)*, *1*(1). Retrieved May 13, 2021, from https://online-journals.org/index.php/i-jim/article/view/165

Witcomb, A. (2006). Interactivity: Thinking Beyond. In S. Macdonald (Ed.), *A Companion to Museum Studies*. Blackwell Publishing Ltd. https://doi.org/10.1002/9780470996836

Woodruff, A., Aoki, P. M., Hurst, A., & Szymanski, M. H. (2001). Electronic Guidebooks and Visitor Attention. *Proceedings of the International Conference on Cultural Heritage and Technologies in the Third Millennium*, 437–454.

Yamazaki, K., Yamazaki, A., Okada, M., Kuno, Y., Kobayashi, Y., Hoshi, Y., Pitsch, K., Luff, P., vom Lehn, D., & Heath, C. (2009). Revealing Gauguin: Engaging visitors in robot guide's explanation in an art museum. *Proceedings of the SIGCHI Conference on Human Factors in Computing Systems*, 1437–1446. https://doi.org/10.1145/1518701.1518919

Yang, Q. (2018). Machine Learning as a UX Design Material: How Can We Imagine Beyond Automation, Recommenders, and Reminders? *1*, 2–6.

Yang, Q., Steinfeld, A., Rosé, C., & Zimmerman, J. (2020). Re-examining Whether, Why, and How Human-AI Interaction Is Uniquely Difficult to Design. *Proceedings of the 2020 CHI Conference on Human Factors in Computing Systems*, 1–13. https://doi.org/10.1145/3313831.3376301

Zimmerman, J., & Forlizzi, J. (2008). The Role of Design Artifacts in Design Theory Construction. *Artifact*, *2*(1), 41–45. https://doi.org/10.1080/17493460802276893

Zimmerman, J., Forlizzi, J., & Evenson, S. (2007). Research through design as a method for interaction design research in HCI. *Proceedings of the SIGCHI Conference on Human Factors in Computing Systems*, 493–502. https://doi.org/10.1145/1240624.1240704

Zwan, S. v. d., Smith, M. L., Bruineberg, J., Levy, P. D., & Hummels, C. C. M. (2020). Philosophy at Work: Postphenomenology as a Generative Lens in Design Research and Practice. *Proceedings of DRS2020 International Conference : Vol. 4: Education*, 1691–1705. https://doi.org/10.21606/drs.2020.337

II Publications

Publication A

Handling Digital Reproductions of Artworks

Christian Sivertsen and Anders Sundnes Løvlie

The paper has been published in the
Journal of Somaesthetics Vol. 7(2), pp. 51–71, 2021.

The layout has been revised.

Abstract

The senses are finding their way back into the art museum, but the way paintings are displayed is still constrained by their fragility. We explore whether it would be helpful to use the capabilities of digital technologies to create meaningful somaesthetic experiences with digital reproductions. We conducted an experiment with 19 participants, letting them handle physical paintings and 2D and 3D digital reproductions, while ranking them according to their personal preference. To discover which cultural qualities participants ascribe to artworks in light of their somaesthetic experience, we interviewed participants regarding their experience of ranking three setups. We found that participants regarded the 3D reproductions as having certain material qualities. We argue that by designing the somaesthetic experience of digital reproductions, it might be possible to bring back dimensions of the art experience that were lost with the development of the modern museum.

A.1 Introduction

In the 17th and 18th centuries, museum visitors were typically allowed to handle the objects exhibited in museums. Indeed, handling and touching were seen as an important part of the museum experience that could enhance learning and enjoyment and create a more intimate connection to the artists (Howes, 2014). However, this practice was later replaced by a focus on *contemplation* and rigid bodily constraints in the museum space (Leahy, 2012). For many years, the white cube paradigm has dominated the way we look at art in museums. The script of the museum mediates our engagement with the art and puts the museum in the role of an authority, defining the right way to appreciate it (Duncan, 2005).

More recently, museum research has been shifting toward a more interpretative or constructivist paradigm, where the museum design is recognized as part of shaping the visitor experience and the visitor as an active part in the learning process (Macdonald, 2007). Nevertheless, the physical configuration of art museums remains largely the same, and the shift seems to be more evident in the way museum experiences are discussed and analyzed than in the way art is displayed. This is especially true of exhibitions of classical paintings. This is not only a question of culture but also of practicalities. The originals on display are fragile, unique, and expensive, and only specially trained personnel can handle them (Howes, 2014; Leahy, 2012). This severely limits the way painting exhibitions can be shaped. However, technology allows us to break free of these limitations. With digital technologies, we can enable new bodily relations with the paintings that are not constrained by the risk of damaging the originals.

In a merge between the classical artworld and the technologically immersive, some venues such as the Lumières venues by Culturespaces[1] and "Van Gogh Alive" by Grande Experiences[2] are exhibiting classical paintings through room-sized digital projections. Through technology, they are pushing the spatial relationship between visitor and painting. The paintings are bigger, cropped in new ways, and wrapped around walls and on the floor, and sometimes details or whole paintings are animated and moving around. However, they reproduce the role of paintings as something hanging on a wall that we view from a couple of meters distance—as an image, not an object.

Is it possible to use the capabilities of digital technologies to create meaningful somaesthetic experiences with paintings? Our bodily actions and relations to paintings and the context in which paintings are met are shaping our experience of them (Dewey, 1934/2005). To explore the design potential of using digitized reproductions to create somaesthetic experiences with paintings, we created an experiment that compares the act of handling paintings in three different setups. We asked 19 participants to look at and consider paintings in the following three formats: physical paintings, paintings represented digitally in 2D, and paintings represented digitally in a virtual 3D environment. The participants were asked to rank them according to what they would like to have in their own home in order to make them focus on their own aesthetic experience. This was followed by a phenomenological interview, where participants were asked to elaborate on their experience and to compare their experience of the three setups. We discuss how the technological mediation and the somaesthetic qualities of each setup are described by the participants and what this can tell us about the design space for technological experiences containing digital reproductions of artworks.

A.2 Art Experience and Technology

John Dewey argues in his 1934 book "Art as Experience" that philosophical aesthetics has wrongly removed art from its situatedness in the everyday experience. According to Dewey, art needs to be considered through its relation to the body and the context in which it appears. In our time, art is increasingly being seen on screens, in part because digital media makes it possible for audiences who cannot travel to the museum to view artworks (a very urgent consideration at the time of writing, in 2021, due to restrictions during the ongoing COVID-19 pandemic).

The literature on technology-mediated experiences in museums often reveals a concern among museum scholars and professionals (as well as the wider public) that technology may come to stand in the way of visitors' direct encounters with physical

[1] https://www.culturespaces.com/
[2] https://grande-experiences.com/van-gogh-alive/

artifacts. Sometimes this concern is referred to as the "heads-down phenomenon"—evoking the image of (young) visitors walking around the museum with their heads pointed down toward their smartphone screens, oblivious to the treasured artifacts on exhibit around them (Hsi, 2003; Lyons, 2009; Petrelli et al., 2013; Walter, 1996; Wessel & Mayr, 2007; Woodruff et al., 2001).

Alternatively, research on human computer interaction (HCI) and interaction design has long explored how interaction with technological systems may form part of the aesthetic experience. Both Dewey (1934/2005) and the ecological psychology of Gibson (1979) have been significant influences in this line of research and in the broader humanistic turn in HCI (Bardzell & Bardzell, 2015). The implications of Dewey's view on the art experience extends beyond the domain of art and has formed part of the theoretical foundations for HCI's focus on experiences with technology (McCarthy & Wright, 2004).

Somaesthetics has received much attention in HCI (Höök et al., 2016; Höök et al., 2015; Lee et al., 2014; Shusterman, 2014). However, there is little work connecting somaestethics to the art experience—although occasionally the results of design projects are themselves exhibited as artworks (e.g., Schiphorst, 2009). More broadly, experiences with technology in museums is a large topic in HCI research (Hornecker & Ciolfi, 2019; Vermeeren et al., 2018), and research has explored how to use embodied interactions to enhance art experiences (Alexander et al., 2017; Steier, 2014). For example, Ryding and Fritsch (2020) present a game for visitors to art museums in which one player controls the movements of another player as a way to challenge the ritualized nature of the museum visit and intensify the visitors' affective encounters with the art.

The interactive art installation "Thresholds" (Tennent et al., 2020) sets up an experience with some similarity to the experiment presented here. Aiming to explore the role of technology in our perception of the world, the installation recreates a 170-year-old photography exhibition inside a virtual space, which is mapped onto a physical space in such a way that visitors donning customized VR equipment have the experience of walking around inside a virtual exhibition gallery that can be explored through touch and other senses. The system allows visitors to virtually select photographs out of the exhibition vitrines using hand gestures to lift the images up for closer inspection. The fact that this feature created significant difficulty for both the creators of the installation and the users — in an otherwise ambitious and highly successful installation — speaks much about the difficulty involved when attempting to facilitate experiences of handling digital artwork.

A.3 Handling in the Museum

According to Howes (2014), museums in the 17th and 18th centuries were hands-on sites, where visitors were expected to touch and handle artifacts. Touching was seen as important for four reasons, as follows: Visitors would be able to learn more through touching, touch was seen as enhancing the enjoyment of art objects, touch allowed for a sense of intimacy with the original creators of the artifacts, and, finally, some rare and exotic objects were believed to have special healing powers. By the middle of the 19th century, the practice of touching in museums had ceased as the reasons mentioned above were no longer considered valid (Howes, 2014). Instead, as described by Leahy (2012), correct aesthetic appreciation became part of a codified bodily practice of walking, sitting, standing, looking, and speaking. Guides were even created that described how to maintain the correct distance from the object that was to be contemplated.

Since the late 20th century, touching and handling have been returning to the museum, first in children's and science museums but later also in art museums. As Howes sums it up:

> "In the museum of the twenty-first century, the senses are making a comeback. Didactic instruction has increasingly come to be supplemented by multimodal approaches to learning, disinterested contemplation has been offset by affective participation, and the authority to interpret objects has been redistributed." (Howes, 2014, pp. 264–265)

In a case study exploring the role of touch in relation to sculptures, the authors note that "When allowed to touch, we observed that groups moved, viewed, described, and discussed the works in more diverse ways than when viewing only, and that touch fostered longer and deeper object-related inquiries" (Christidou & Pierroux, 2019, p. 111). Physical sculptures carry their meaning in their shape and form and are often robust. Paintings, however, are primarily visual artworks and are vulnerable to touch. Thus, inviting visitors to touch or handle valuable paintings is obviously not possible. However, the development of new immersive technologies and interaction formats offer interesting opportunities to consider bodily experiences with digital reproductions of artworks. This in turn raises questions about the role of reproductions in art experiences.

A.4 Reproductions and Genuineness in Psychological Aesthetics

One of the factors that makes it difficult for museums to allow visitors to handle artworks is also arguably one of the main reasons visitors are attracted to museums—the ability to view invaluable (but fragile) artworks in their authentic, original form. For example, Walter Benjamin famously argued that the *aura* of classical artworks such as paintings and sculptures is bound to their cultural and physical properties, which are lacking in reproducible media such as photography. How important is it for the art experience that one is in fact viewing an original and not just a reproduction? Several empirical studies have tried to understand the influence of the genuineness of a piece of artwork on the art experience (Locher et al., 1999, 2001; Locher & Dolese, 2004; Brieber et al., 2014; Brieber et al., 2015). These studies find that viewing original artworks in a museum is rated higher than viewing reproductions in a laboratory in terms of parameters such as being immediate, pleasant, interesting, surprising, liked, and understood (see Pelowski et al., 2017 for a full overview). Considering the medium of reproduction, three of the studies hypothesize that if art viewers can look past the medium, they will evaluate the same image similarly when seen in various media, measured through quantitative and qualitative components of the information content of the images—a phenomenon they call *facsimile accommodation* (Locher et al., 1999, 2001; Locher & Dolese, 2004).

However, in these studies, the role of the context is not clear as the originals are viewed in the setting of a museum or art gallery, and the reproductions are viewed in a lab setting. Brieber, Leder et al. (2015) try to detangle this effect in a study that compares both context and genuineness; however, in the study, neither the context nor the genuineness was found to enhance the participants' evaluation of the artworks. This was attempted again by Grüner et al. (2019), who did find that artworks viewed in a museum are liked more and rated as more interesting when presented in a museum rather than in a laboratory. Genuineness is not found to have this effect.

Pelowski et al. (2017) expand on the comparison of laboratory vs. museum as a factor in art appreciation by presenting a large range of factors that influence the art experience. These factors pertain to the artworks, the museum space, and the visitor. Among the factors related to the artwork itself are texture, immediacy, physical presence, and size (Pelowski et al., 2017). The authors also mention the hanging style as having an influence on the art experience.

Across the studies described above, reproductions take the form of images on computer screens, slide-projections, or even postcard-sized printed images. Bertamini and Blakemore (2019) present two studies in which they asked participants to evaluate hypothetical scenarios of viewing three types of artwork reproductions. The hypothetical reproductions were a painting viewed through a closed-circuit video

camera monitor, a painting viewed through a mirror, and a physical reproduction of the painting. They found a large variation in the participants' opinions on the three types of reproduction. In general, the physical copy was preferred over viewing the original indirectly, and a mirror reflection was found to be better than a video image.

These empirical studies seem to indicate that the museum context is important for the aesthetic experience, whereas the importance of viewing an original vs. a reproduction is less clear. Some of the studies indicate that the specific format of the reproduction seems to matter. However, all of these studies were limited to the experience of passively viewing artworks on a wall or in a display. In this article we continue to explore this question from a design perspective, offering an exploration of the design space for digital reproductions that can be virtually handled by the viewer.

A.5 Handling Reproductions: A Somaesthetic Perspective

Dewey argues that substance and form are central to the art experience: "*what* is said and *how* it is said" (Dewey, 1934/2005, p. 106, emphasis in original). Replicating a piece of artwork in digital media changes its form and subsequently its substance. To understand form with regard to digital media, the literature from the field of interaction design provides a compelling model.

Vallgårda (2014) argues that in interaction design practice, three form elements are closely interconnected: the physical form, the temporal form, and the interaction gestalt. The physical form is the shape and appearance of the system as perceived through our sensory apparatus. The temporal form is the change of states in the system over time. The interaction gestalt is the movement the user performs in relation to the system. These movements have qualities, such as being fast, smooth, or abrupt, and take place in a doing and undergoing relationship with the system. The user acts on the system, and the system shapes the acting.

To better understand how form shapes experience, we turn to postphenomenology. A postphenomenological approach implies a particular interest in the relation between participants and paintings and how this relation is being mediated by the technologies used in each setup (Rosenberger & Verbeek, 2015). In this study, we are investigating how the technologies employed reshape the experience of the paintings. Human-technology relations are in the postphenomenological view characterized by a magnification/reduction structure (Rosenberger & Verbeek, 2015). According to Kiran (2015), this structure is divided into four dimensions of technological mediation: ontological, epistemological, practical, and ethical. These dimensions serve as a helpful framework for analyzing the mediation aspects in the experiment. The assumption behind this experiment is that the technological representation chosen will shape the experience of the artworks in how it *reveals* and *conceals* aspects of the

artworks, how it *magnifies* or *reduces* the knowledge available about the artwork, how it *enables* or *constrains* certain practical actions, and in turn how that *involves* or *alienates* the participants from what is considered ethical practice around artworks.

Within this perspective, we find it relevant to pay specific attention to the *aesthetics of interaction*, including the perception of performance. Lim et al. (2007) present the concept of *interaction gestalt* as the shape of interaction: the movements the user makes while engaging with an interactive system. Lenz et al. (2017) describe the qualities of these movements as *interaction attributes* and find that they are related to experiential qualities. Dalsgaard and Hansen emphasize the social aspect of performance, suggesting that the user of a system continuously acts out the three roles of *operator*, *performer*, and *spectator* (2008). Applied to our experiment, this means that our participants will simultaneously be operating the systems we have put in place while also perceiving the relation between themselves and the paintings and being aware that these actions are a performance for the experiment facilitator and the recording equipment. As Dalsgaard and Hansen (2008) argue, this performance of perception is an integral part of the aesthetics of interaction.

A.6 Method

The experiment presented in this article bears similarities to the approach of *concept-driven interaction design research* (Stolterman & Wiberg, 2010) in the sense that we are conducting practice-based design research with the aim of exploring a theoretical issue rather than designing new products. Furthermore, our approach is inspired by a *constructive design research approach* (Koskinen et al., 2011), which means that the construction of design artifacts is central to knowledge creation.

In our way of setting up this experiment, we lean on the tradition of performing design experiments in the lab (Koskinen et al., 2011). Contrary to the more common use of experiments as vehicles for deductive reasoning, this experiment is inductive in nature. We are looking for patterns in a design space, not trying to prove them. An important difference between our experiment and those presented by Koskinen et al. (2011) is that the three designs used in our experiment are not made as proposals for future designs. Instead, they are created in order to explore the impact of these different formats on aesthetic experience. We are not primarily interested in the particular designs but rather in the comparison of participants' interactions. In this way, the designs used here are more research instruments than design proposals.

A.6.1 Experiment Procedure

The experiment was conducted with 19 participants recruited at our university from the 30th of November to the 4th of December 2020. Fourteen of the participants were master's students or recent graduates in the field of digital design or games,

four were faculty within digital design, and one was enrolled in vocational education in the health sector. Ten identified as female and nine as male. The age of the participants ranged between 22 and 36. Sixteen participants said "yes" to being interested or somewhat interested in art, while three did not see themselves as interested in art. All but one had visited an art museum or gallery within the last year, with an average of three visits in the last year. This number should be viewed in light of the COVID-19 situation, where many such places were closed for long periods during the previous year.

The experiment was divided into three different setups. In each setup, the test participants were invited to experience artworks in one of three different formats: framed physical paintings, digital reproductions of paintings displayed as 2D images, and digital reproductions of paintings presented as 3D objects. For each of these setups, the users were invited to pick up the paintings— physically or virtually—in order to get a closer look.

The participants were told that they would be entering a room with three pieces of artwork. They were asked to look at the artworks and rank them according to which they would most like to have in their own home. The rationale for giving the participants this task was to prevent participants from judging the artworks according to some external ideal and to rather focus on their own aesthetic experience of the artworks. After making their decision, participants exited the experiment room and were interviewed about their experience. This was repeated for each of the three setups. The sequence of the three setups was changed so that participants went through them in a different order each time.

Fig. A.1: The physical artworks bought from different secondhand stores. The print on the left is signed Line Thimm. The painting in the middle is unsigned. The painting on the right is signed S. Engelbrecht.

The experiment used nine different artworks, presented below. The artworks were deliberately chosen for being ordinary, non-famous artworks of the type that one might buy in a secondhand store (and indeed, the three physical paintings are "thrift store" paintings). The images represented various visual styles to accommodate a variety of aesthetic preferences. The participants were given no information about the artworks other than what they could see for themselves. Note that it was

necessary to use different artworks for the three different setups (rather than repeating the same three images) in order to make the task of choosing an artwork meaningful for the participants for each of the three iterations.

In all three setups, the paintings were partially hidden from sight as the participant entered the room, either due to their placement or the image size. This was done to prompt participants to handle the paintings in order to get a closer look at them.

After each setup, the participants were interviewed about their experience and asked to compare their experience with the other setups. The interview was conducted as a phenomenological interview (Thompson et al., 1989). The 19 interviews were transcribed verbatim. Statements describing the qualities of each of the three setups were separated and then organized thematically using affinity clustering.

A.6.2 Physical Setup

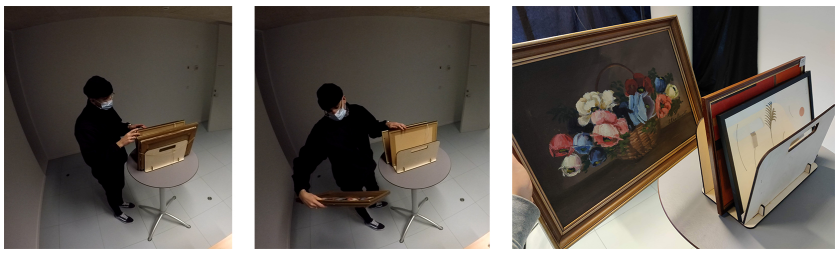

Fig. A.2: In the physical setup, the paintings are placed in a rack on a tall table. The first and second image are video stills from the experiment. The third image is a staged closeup.

In the physical setup, the participants were presented with three physical artworks bought in secondhand stores around Copenhagen (see fig. A.1). The paintings were chosen to represent a variation of styles. The three paintings were placed in a rack where the paintings were easily accessible, but each partly obscured by the other. The rack was placed on a tall table (see fig. A.2).

A.6.3 Digital 2D Setup

In the digital 2D setup, participants were presented with three paintings projected next to each other on the wall of the experiment room (see fig. A.4 and fig. A.3). The paintings by Bea Mahan and Manjiri were found on their Flickr accounts where they promote their art. The third one is a study by the Danish artist Niels Bjerre. It was found in the database of the Danish National Gallery. The paintings were chosen to represent a variation of styles.

Fig. A.3: The 2D images with added frames (from left to right: Mahan, n.d.; Manjiri, n.d.; Bjerre, 1934).

Fig. A.4: In the 2D setup, a wireless mouse is placed on a tall table. The participants use the mouse to increase the size of the painting they want to look at. The first two images are video stills from the experiment. The third image is a staged closeup.

The interface for this setup was created in TouchDesigner in a simple 2D environment (see fig. A.4). Frames were added digitally to the paintings. In the middle of the room was a table with a wireless mouse. The paintings were projected in a size that made them too small to view comfortably from the table with the mouse both due to the distance and the resolution of the projector. The participants were made aware that they should use the mouse. When hovering the cursor over the image, it would grow slightly in size, and upon clicking, it would grow to a large size. If the participants clicked outside the scaled-up image, it would shrink to its initial size, and if the click was placed on another image, that one would scale up instead.

A.6.4 Digital 3D Setup

In this setup, the participants were presented with three paintings in a virtual 3D environment (see fig. A.6). The painting by Layers was found on Pixabay.com, a stock image site where the artist offers their art for free use. The painting by Miguel Àngel Pintanel was found on his Flickr account where he promotes his art. The painting by Mogens Ballin was found in the database of the Danish National Gallery. Again, the paintings were chosen to represent a variation of styles (see fig. A.5).

Fig. A.5: The 3D images with added frames (from left to right: Ballin, 1890; Layers, n.d.; Pintanel, n.d.).

The 3D environment was projected in 2D on the wall of the experiment room in a forced perspective that corresponded with the position where the interviewer would tell the participant to stand when entering the room. The paintings were shown lying on a (virtual) wooden table. On the projection was a white cube acting as a cursor hovering over the paintings.

The participants were given a smartphone and instructions on how to use the smartphone to interact with the paintings. The smartphone could be used in a manner similar to a laser pointer: When pointing the top of the smartphone toward the projection, the white cube would follow the movements of the phone. By pressing with their thumb on the screen, participants could "pick up" a nearby painting, which would attach itself to the white cube. Pointing the phone upward would move both the white cube and the painting closer to the virtual camera so that the painting could be inspected more closely. The orientation of the painting would map to the orientation of the phone, allowing the participant to tilt and rotate the painting to allow for examination from various angles. If the participant removed their thumb from the screen, the painting would fall down. If the image fell toward the ground, it would disappear outside of the projection and reappear on the table. In this setup, the frames and canvases were 3D-modeled, and the paintings were added as textures to the 3D models.

This interface was also created in TouchDesigner as a 3D environment with a bulletsolver engine to simulate gravity and other forces. The smartphone interface was based on the Google XY-Fi project (Uglow et al., 2017). The smartphone ran a website that records device orientation and touch events and passed it via socket.io to the webserver that forwarded it to the TouchDesigner instance running the simulation. The "pick up" mode is not a part of the original XY-Fi project but was programmed by the first author, extending the original JavaScript program.

Fig. A.6: In the 3D setup, the participant uses a smartphone to control a cursor on the screen. It can be used to pick up, move, and tilt the virtual paintings. The first two images are video stills from the experiment. The third is a staged closeup

A.7 Results

We now present our observations and insights from presenting the study participants with each of the three setups.

A.7.1 Physical Setup

As the participants entered the room for the physical setup the experimenter would give the following instruction: "Please have a look at the artworks and rank them according to what you would like to have in your own home. Let me know when you have made your choice." The participants moved to the rack, many hesitating a bit before picking up the artworks. Almost all participants asked the experimenter whether it was okay for them to touch the artworks, either right before or right after taking one from the rack.

Most participants then proceeded to pick up the artworks one by one, studying each one for five to 15 seconds before putting it back in the rack. Others held a painting in each hand next to the one left in the rack, comparing all three at once. A few participants picked up some of the artworks a second time. One participant held the artworks against the wall of the experiment room. From the video recordings, it can be seen that the participants spent between 25 to 90 seconds (median: 54) before indicating that they had made their choice.

In the interviews, most of the participants brought up the physical qualities of the artworks. They mentioned weight, texture, tactility or tangibility, materiality, and the ability to feel the paintings as qualities that were significant. One participant expressed it like this:

> "I like that I was able to pick up the paintings and feel it, and look at it in the light, and look at it pretty close and study some of the details, and then be like: 'That was nice to see.' It gives you something, when you are far away and close to paintings, I think." (Participant 11)

In addition to holding the artworks up close, participants mentioned the options of turning them around and moving them back and forth, and one highlighted the feeling of having control of the handling of the paintings. A few participants mentioned that it was difficult to handle the paintings in this setup, "[...] because I could only hold two at once, it was hard to see all three at the same time. So, I had to remember to hold one in my mind and then look at the others" (Participant 17). Another participant mentioned being anxious about accidentally breaking the artworks.

As compared to the other setups, half of the participants mentioned that only the physical setup gave them the full impression of the painting, especially with regard to colors: "I prefer having them physical because then I can just see more and I can trust my perception more, because if it's like I'm shopping in an online shop, I don't actually see the color. If I'm checking it out in real life, then I know exactly what I will get" (Participant 15).

In this setup, some participants talked about the importance of the frame for making their choice: "For me it's very important how the frame looks on the paintings, so I also investigate how old they are, and whether they are worn, if they are new, and how much they look like they have just been printed on laser printer and put in a black frame. But I am sure none of these are" (Participant 13). Another participant found a specific frame enticing, "It weighed heavily in my decision of what I wanted, that there was a name I could recognize [S. Engelbrecht, ed.], but also that it was a nice painting, and that it was heavy and a nice frame" (Participant 10).

Several participants made comparisons with the act of browsing through artifacts in commercial settings, such as posters in an art museum gift shop, paintings at a flea market, or records in a record store. For some, this was a positive, fun experience:

> "It's like crate digging in a record shop. You're kind of fiddling through them. You can look, you can stand it. You know, it feels more like you're kind of taking a cultural artifact in a different kind of mode. That's sort of exciting. You know, there's joy to holding a painting. It's something almost naughty." (Participant 19)

However, others felt unease when handling the physical paintings. One participant mentioned that it devaluated the artworks being presented in this way:

> "I felt that it was like when exiting a museum, and then there's this thing where you browse through the posters. My immediate experience was that I really felt that I was in the gift shop of a museum. I also think, that in relation to other things, this took away much of the feeling of quality." (Participant 2)

In general, participants found that this setup gave them the best impression of the paintings. The paintings were evaluated for more than their pictorial content, such as the frames and their weight, yet the presentation was unfamiliar, causing a level of unease.

A.7.2 Digital 2D Setup

Upon entering the room for the 2D setup, the experimenter gave the same task instructions as for the physical setup, but this time added: "You may use the mouse on the table." With no further instructions, almost all participants would walk to the mouse and start clicking, figuring out by themselves how to enlarge the paintings. Participants spent 19–160 seconds (median: 40) before they indicated that they had made their choice. All participants went through the images at least once, spending 1–7 seconds looking at each enlarged picture for the first time. Most participants looked at the enlarged images multiple times.

Most participants said that this setup was easy and straightforward and felt like an everyday interaction: "It was easier because everything was just lined up, instead of having to make that somewhat cumbersome movement of lifting the paintings up. [...] So, it was a faster decision to decide what you like, but with less opportunity for investigation" (Participant 13). One participant noted that the ability to see the three images at the same time made it easier: "Even if the paintings weren't that big [when not enlarged], I already kind of saw what they were portraying" (Participant 15). Many participants talked about how it was easier to get an overview or compare the images in this setup. Participants also said that it was more efficient and had less distractions than the other setups, and some remarked that it was easier to investigate details in this setup. Several compared this setup to an online image experience, such as Google Image Search.

In contrast, a few participants said that it was difficult or even impossible to make a proper decision in this setup because necessary information was missing from the presentation of the paintings. Over half of the participants talked about missing information aspects, such as the physical dimensions of the paintings, texture, and the exact colors. Curiously, four participants furthermore stated that the images in this setup did not have frames (even though frames had in fact been added digitally, as described above).

One participant said this setup was just like images hanging on a wall. Another compared it to a slideshow:

> "I felt it was like a slideshow that I had to click through, and it pulled me out of the world where I am supposed to be immersed in the art. You feel that you have a mission and that is to be done with it. You kind of have to see it to the end, and then proceed with the next instead of immersing yourself." (Participant 12)

In general, the participants seemed less enthusiastic about this setup than the other two. Some participants said it was boring, others used the term static, and a few used the term distanced in comparison with the other setups. The task was solved quickly and efficiently, but the images in this condition were not talked about as having any sort of physical or spatial properties.

A.7.3 Digital 3D Setup

When entering the 3D setup, the experimenter would hand the participant a smartphone and ask the participant to stand in the middle of the room in front of the projected image. Then, the experimenter would help the participant to calibrate the phone interface and explain the functionality: "You can move the white cursor around by pointing the phone. You can press on the screen to grab a painting, and if you point the phone toward the ceiling, the painting you have picked up will come toward you. You can then tilt the phone to orient the painting you have picked up." Then, the experimenter would repeat the task and step back to let the participant use the interface on their own while answering any clarifying questions.

In this setup, participants spent between 62 to 250 (median: 94) seconds after entering the room before they indicated that they had made up their mind. The instruction and calibration phase took 23–46 seconds. The participants would pick up the paintings one after the other and tilt the phone to make it come closer. The participants kept standing in the same place while holding the phone in one hand extended from the body. They spent between 4 and 26 seconds looking at a picture zoomed in when looking at it for the first time. Some participants picked up one of the paintings one more time before revealing their choice. One participant played around with the paintings for another two minutes after explaining his choice.

At the beginning of their interaction with the setup, many of the participants experienced chaotic interactions. Participants often accidentally dropped the paintings, knocked them off the table or sent them flying out of the screen:

> "You had to get used to it and find out how to maneuver the painting. [...] Sometimes the painting moved a bit fast and ended in the top right corner. It was a bit hard to keep the painting in focus, which made it difficult to analyze the painting you were looking at. But it was a fun way of doing it." (Participant 10)

Many participants said that this was a different or novel way of interacting with art, but many also remarked that the interface involved a learning curve since they needed to learn how to use the tool before they could focus on the paintings. About half of the participants used the word fun about this setup, and a few more talked about it as being playful. One participant, however, found the interaction difficult, making it a "stressful" and somewhat "humiliating" experience (Participant 16). Some also experienced a certain unease about handling the paintings in this setup: "I felt that I was treating the art a bit badly by accidentally throwing it around and by rotating it. In any case, I would find it awkward if I ended up doing that with [the artworks in] my own hands" (Participant 7). Another participant had a similar experience but appreciated that the artwork had lost a bit of its authority:

> "There were some times when you dropped the precious paintings and those kinds of things. And then you go like, it's not normal to be out looking at art, holding some priceless artwork and then, whoops, dropping it or it flying away. But I actually think that gave it a really cool playful approach, that you dare more to look at it and do something with it. You don't dare that when you're in a museum, then you just go: Okay, I can look at things [...] maybe it makes the art less authoritarian [sic] that you can throw it around like that. But I actually think that's very cool." (Participant 11)

Some participants talked about this setup being playful or like a game. While being playful and fun, one participant found that it was "just feeling like a gamified distraction from the task at hand" (Participant 16). Participant 8 also felt this way: "For a long time, I had much focus on just controlling it, and I found it fun, and that was where my focus was. I forgot the task a bit." An additional two participants mentioned this.

Similar to the 2D setup, a few participants said that they found it hard or impossible to complete the task because the digital image of the paintings did not give them all the information they needed. A few participants talked about a missing materiality or tangibility; however, others talked about this setup being more material, tangible, or physical than the 2D setup. Other factors that were mentioned as missing were weight, real size, and exact colors. One participant talked about this setup being a tradeoff between the two others:

> "[the 3D setup] seems like it's sort of awkwardly in the middle. There's something that's material that's happening there that is nice, but it's also fiddly and it's also occluded in some sort of image sense. [...] It's harder to see, but it does kind of give you a sense that you're semi-present, which I don't know if that's a good trade off yet." (Participant 19)

In this setup, the participants also talked about frames. One participant said that the frame did not play a role, "[...] because you could not feel the image in the same way, even though there is a frame" (Participant 10). On the other hand, another participant said:

> "[...] it really did do something, that there were frames on. [...] It gave me more the feeling that they were actually paintings existing in real life, instead of just being a Google image you had downloaded and put into the same system. Here, I have a feeling that these paintings exist somewhere." (Participant 6)

Several participants found that the 3D setup did have some qualities to it that the 2D setup lacked. Three participants said that this was more like holding a real painting than the 2D setup. One said that it had "objectness" (Participant 1), and another that it was easier to imagine it on a wall. Another three participants talked about the 3D interface as a room, implying a sense of spatiality.

A.8 Discussion

Remarks made by the participants seem to indicate that the 3D setup did succeed to some degree at facilitating an experience that afforded a sense of handling the artworks. Participants also noticed the frames of the paintings more in 3D than in 2D. It is particularly interesting to note the unease felt by some participants in not being able to treat 3D artworks with appropriate care. However, participants still noted a lack of materiality, and problems with the 3D interface and image quality seem to have reduced the vividness of the experience for some.

To explore how the 3D setup affects the aesthetic experience of the artworks, we will consider the insights from the experiment in relation to Kiran's (2015) four suggested dimensions of technological mediation: practical, ontological, epistemological, and ethical. For each of these dimensions, we offer some thoughts on how designers might further explore the experience of virtually handling digitized artworks.

In each setup, the form of the artworks afforded different practical ways for the participants to handle them. The three setups demanded three very different ways of bodily engagement, from the careful handling of a heavy physical painting, to the fine flicks of the wrist when using a mouse, to the somewhat unfamiliar movements needed to control the smartphone interface. The movements in the 3D setup land somewhere in between those of the 2D and the physical setups: The participants were lifting, pulling, placing, and tilting the paintings, although it was done via a tool for remote control and with much smaller and lighter movements than in the physical setup. These affordances allowed the participants' bodies to play a role in the art

experience. In future work, designers might explore how to further prompt and enhance the affordances for practical handling to extend the ways bodily movement might affect aesthetic experience.

Designers might explore (at least) two different aspects of this design space: the control interface and the type of display. Regarding control, one might experiment with interfaces that facilitate more natural movements, thus mapping more closely to the handling of physical artworks. For instance, one might create a tangible interface with a form like that of a physical painting that could be mapped to the digital image to allow participants to use natural movements to lift and turn the digital image. To bring the experience even closer to the physical, one might move away from the digital projection on the wall and instead simply use a tablet computer embedded in a frame. However, this would require that the images be reduced drastically in size and adapted to the aspect ratio of the tablet display, which would run against the artistic intentions of many artworks in which size is an important aesthetic factor. A different solution might be to move the experience into a substitutional reality environment in which a virtual reality environment is combined with physical props to facilitate the experience of handling objects physically, as demonstrated in Tennent et al. (2020). Furthermore, designers and artists might be interested in experimenting with interfaces that offer types of interaction that do not match closely to the experience of handling a physical painting, such as introducing elements of discomfort (see Benford et al., 2012), for example, through sensory misalignment (see Marshall et al., 2019). One might also consider the degree to which the participants should control the experience—perhaps experimenting with degrees of contested or negotiated control (see Benford et al., 2021).

Considering the ontological dimension, both the 2D and the 3D versions of the artworks are virtual representations, but participants felt that the handling of virtual 3D has more qualities associated with physical objects. When going from the physical to the 2D setup, the experience of "objectness" seems to disappear. The participants called the physical paintings "the real thing," whereas the 2D paintings seem to be treated more as a reference to an object existing in some other realm. Interestingly, when encountering the 3D paintings, a level of "objectness" seemed to return as some participants said that the experience of engaging virtual space is a bit like handling real paintings.

Designers might explore designs that would enhance the experience of "objectness" in relation to the digitized artworks. In the past, museums have experimented with ways to facilitate more personal encounters with artworks, such as through Cooper-Hewitt's "The Pen," which allowed visitors to digitally collect objects that interested them in the museum (Chan & Cope, 2015). In Blast Theory's design "Gift", museum visitors are invited to collect objects digitally and to use them to craft gifts for their loved ones (Spence et al., 2019), setting up an experience that is "interpersonalized" (Ryding et al., 2021). Petrelli et al. (2017) introduced the concept of "tangible data

souvenirs," which are created on the basis of data collected during a museum visit and that serve as a connection between a physical and a digital experience. Benford et al. (forthcoming) used a similar approach in a design in which emotion-capture techniques were used to craft personalized experiences based on the visitors' emotional responses to artworks in the Munch Museum in Oslo, Norway. At the end of the experience, visitors were given a postcard showing the painting that had prompted their strongest emotional response, with their own emotion data printed on the back. Future designers might build on such approaches to further experiment with ways of turning digitized artworks into "objects."

Considering the epistemological dimension, it is notable that the physical artworks seem to contain important information that becomes unavailable in the two digital setups. Some of this loss—blurry images, low resolution—is due to inefficient display technologies and might easily be mitigated using a screen or a better projector. In fact, by using high-resolution images such as the "gigapixel" images created by the Google Cultural Institute (St. John, 2016), one may display an even larger and sharper representation of the paintings than can be seen directly on the physical canvas. Information about physical size can also be communicated digitally. In the 2D display, it is easy to imagine scaling the images 1:1 to their physical counterparts. With the 3D interface, this is less trivial since scale is determined by the distance to the virtual camera lens as it moves back and forth. Other information, such as weight, is simply lost due to the nature of digital representations. The 3D version, however, does convey the sense of being an object since participants can look at it from the front, sides, and back. Using better display technologies, it might even be possible to see the artificial light bouncing off the texture of the 3D canvas.

Considering the ethical dimension, it is worth noting that participants drew parallels to experiences that have similar interactional qualities. The physical setup was likened to the act of browsing posters in a gift shop, the 2D setup was compared to browsing images on the web, and the 3D setup reminded participants of the Nintendo Wii controller. These three examples are very different in their cultural status and refer to contexts in which artworks are given very different roles. Posters in a gift shop are commercial products, stereotypical examples of art as a commodity. In contrast, images that appear in google searches are deprived of their monetary value (other than the indirect monetization of the platform enterprise). Meanwhile, in the 3D setup, paintings regain some of their "objectness," but they tend to lose their status as art, becoming instead merely quasi-physical objects that get tossed around like toys.

The change in form also affects the social status of the artworks and even devaluates them. An important question for further research would be to search for ways to present the digitized artworks that do not devaluate them. One possibility would be to design an interface that to a large degree affords careful treatment of the paintings, simulating the care and respect that such physical objects require. For instance,

the physics of the simulation could be constrained so that all movements would be slow and smooth and that the paintings would find their way back on the table when let go. Another approach could be to integrate consequences of actions in the software. If the artworks were to break or disappear for good when dropped, this artificial fragility might affect the role of the paintings in the participants' perception. Alternatively, the reckless treatment of paintings could instead be turned into a theme for the experience and explored further in the design, using the experience of unease to explore the role of digitized artworks.

A.9 Concluding Remarks

Can the experience of handling digitized artworks be used to enrich the art experience? The experiment presented here did not aim to offer a viable prototype for such an experience, and indeed the participants' responses indicate that the setup would need to be further developed to be experienced as appropriate for an art-viewing experience. However, the experiment did demonstrate that there is potential for facilitating art experiences that afford a dimension of "objectness" to digitized paintings.

While the digitalization of artworks may seem to lead to art experiences that are immaterial and disconnected from the physical reality of our bodies, this also makes it possible to bring back a dimension of the art experience that was lost with the development of modern museums such that spectators can experience artworks by holding them, tilting them, turning them around, lifting them up, and even throw them away. This opens up new avenues for further research in the intersection of somaesthetics, HCI, and sensory museology.

References

Alexander, J., Wienke, L., & Tiongson, P. (2017). Removing the barriers of Gallery One: A new approach to integrating art, interpretation, and technology.*MW17: MW 2017*. https://mw17.mwconf.org/paper/removing-the-barriers-of-gallery-one-a-new-approach-to-integrating-artinterpretation- and-technology/

Ballin, M. (1890). Skovinteriør med kroget træ. https://open.smk.dk/artwork/image/KMS8914

Bardzell, J., & Bardzell, S. (2015). *Humanistic HCI*. San Rafael, California: Morgan & Claypool Publishers.

Benford, S., Greenhalgh, C., Giannachi, G., Walker, B., Marshall, J., & Rodden, T. (2012). Uncomfortable interactions. *Proceedings of the 2012 ACM Annual Conference on Human Factors in Computing Systems - CHI '12*, 2005. ACM Press. https://doi.org/10.1145/2207676.2208347

Benford, S., Løvlie, A. S., Ryding, K., Rajkowska, P., Bodiaj, E., Darzentas, D., Cameron, H., Spence, J., Egede, J., & Spanjevic, B. (forthcoming). Sensitive Pictures: Emotional Interpretation in the Museum. *The ACM CHI Conference on Human Factors in Computing Systems*, New Orleans, Louisiana, USA.

Benford, S., Ramchurn, R., Marshall, J., Wilson, M. L., Pike, M., Martindale, S., Hazzard, A., Greenhalgh, C., Kallionpää, M., Tennent, P., & Walker, B. (2021). Contesting control: Journeys through surrender, self-awareness and looseness of control in embodied interaction. *Human–Computer Interaction, 36*(5–6), 361–389. https://doi.org/10.1080/07370024.2020.1754214

Bertamini, M., & Blakemore, C. (2019). Seeing a work of art indirectly: When a reproduction is better than an indirect view, and a mirror better than a live monitor. *Frontiers in Psychology, 10*, 2033. https://doi.org/10.3389/fpsyg.2019.02033

Bjerre, N. (1934). Parti fra Brejningegaard (pr. Spjæld). https://open.smk.dk/artwork/image/KKS13268

Brieber, D., Leder, H., & Nadal, M. (2015). The experience of art in museums. *Empirical Studies of the Arts, 33*(1), 95–105.

Brieber, D., Nadal, M., & Leder, H. (2015). In the white cube: Museum context enhances the valuation and memory of art. *Acta Psychologica, 154*, 36–42. https://doi.org/10.1016/j.actpsy.2014.11.004

Brieber, D., Nadal, M., Leder, H., & Rosenberg, R. (2014). Art in time and space: Context modulates the relation between art experience and viewing time. *PLoS ONE, 9*(6), e99019. https://doi.org/10.1371/journal.pone.0099019

Chan, S., & Cope, A. (2015). Strategies against architecture: Interactive media and transformative technology at the Cooper Hewitt, Smithsonian Design Museum. *Curator: The Museum Journal, 58*(3), 352–368. https://doi.org/10.1111/cura.12118

Christidou, D., & Pierroux, P. (2019). Art, touch and meaning making: An analysis of multisensory interpretation in the museum. *Museum Management and Curatorship, 34*(1), 96–115. https://doi.org/10.1080/09647775.2018.1516561

Dalsgaard, P., & Hansen, L. K. (2008). Performing perception—Staging aesthetics of interaction. *ACM Transactions on Computer-Human Interaction, 15*(3), 1–33. https://doi.org/10.1145/1453152.1453156

Dewey, J. (2005). *Art as experience*. Penguin. (Original work published 1934)

Duncan, C. (2005). The art museum as ritual. In Gerard Corsane (Ed.), *Heritage, museums and galleries: An introductory reader* (pp. 78–88). Routledge.

Gibson, J. J. (1979). *The ecological approach to visual perception*. Houghton Mifflin.

Grüner, S., Specker, E., & Leder, H. (2019). Effects of context and genuineness in the experience of art. *Empirical Studies of the Arts, 37*(2), 138–152. https://doi.org/10.1177/0276237418822896

Höök, K., Jonsson, M. P., Ståhl, A., & Mercurio, J. (2016). Somaesthetic appreciation design. *Proceedings of the 2016 CHI Conference on Human Factors in Computing Systems*, 3131–3142. San Jose California USA: ACM. https://doi.org/10.1145/2858036.2858583

Höök, K., Ståhl, A., Jonsson, M., Mercurio, J., Karlsson, A., & Johnson, E.-C. B. (2015). COVER STORY: Somaesthetic design. *Interactions, 22*(4), 26–33. https://doi.org/10.1145/2770888

Hornecker, E., & Ciolfi, L. (2019). Human-computer interactions in museums. *Synthesis Lectures on Human-Centered Informatics, 12*(2), i–153. https://doi.org/10.2200/ S00901ED1V01Y201902HCI042

Howes, D. (2014). Introduction to sensory museology. *The Senses and Society, 9*(3), 259–267. https://doi.org/10.2752/174589314X14023847039917

Hsi, S. (2003). A study of user experiences mediated by nomadic web content in a museum. *Journal of Computer Assisted Learning, 19*(3), 308–319. https://doi.org/10.1046/j.0266-4909.2003.jca_023.x

Kiran, A. H. (2015). Four dimensions of technological mediation. In R. Rosenberger & P.-P. Verbeek (Eds.), *Postphenomenological investigations: Essays on human–technology relations*. Lexington Books.

Koskinen, I. K., Zimmerman, J., Binder, T., Redström, J., & Wensveen, S. (2011). *Design research through practice: From the lab, field, and showroom*. Morgan Kaufmann/Elsevier.

Layers. (n.d.). Ölgemalde Farbe Malerei. https://pixabay.com/de/illustrations/ölgemälde-farbemalerei-künstler-5335144/

Leahy, H. R. (2012). Museum bodies: The politics and practices of visiting and viewing. Ashgate. Lee, W., Lim, Y., & Shusterman, R. (2014). Practicing somaesthetics: Exploring its impact on interactive product design ideation. *Proceedings of the 2014 Conference on Designing Interactive Systems*, 1055–1064. Vancouver BC Canada: ACM. https://doi.org/10.1145/2598510.2598561

Lenz, E., Hassenzahl, M., & Diefenbach, S. (2017). Aesthetic interaction as fit between interaction attributes and experiential qualities. *New Ideas in Psychology, 47*, 80–90. https://doi.org/10.1016/j.newideapsych.2017.03.010

Lim, Y., Stolterman, E., Jung, H., & Donaldson, J. (2007). Interaction gestalt and the design of aesthetic interactions. *Proceedings of the 2007 Conference on Designing Pleasurable Products and Interfaces - DPPI '07*, 239. Helsinki, Finland: ACM Press. https://do.i.org/10.1145/1314161.1314183

Locher, P., & Dolese, M. (2004). A comparison of the perceived pictorial and aesthetic qualities of original paintings and their postcard images. *Empirical Studies of the Arts, 22*(2), 14.

Locher, P., Smith, J., & Smith, L. (2001). The influence of presentation format and viewer training in the visual arts on the perception of pictorial and aesthetic qualities of paintings. *Perception, 30*(4), 449–465. https://doi.org/10.1068/p3008

Locher, P., Smith, L., & Smith, J. (1999). Original paintings versus slide and computer reproductions: A comparison of viewer responses. *Empirical Studies of the Arts, 17*(2), 121–129. https://doi.org/10.2190/R1WN-TAF2-376D-EFUH

Lyons, L. (2009). Designing opportunistic user interfaces to support a collaborative museum exhibit. *Proceedings of the 9th International Conference on Computer Supported Collaborative Learning - Volume 1*, 375–384. Rhodes, Greece: International Society of the Learning Sciences.

Macdonald, S. (2007). Interconnecting: Museum visiting and exhibition design. *CoDesign, 3*(sup1), 149–162. https://doi.org/10.1080/15710880701311502

Mahan, B. (n.d.). Mix-004. https://www.flickr.com/photos/beamahan/6989187839/

Manjiri. (n.d.). Twilight sunset. https://www.flickr.com/photos/mkanvinde/34437869056/

Marshall, J., Benford, S., Byrne, R., & Tennent, P. (2019). Sensory alignment in immersive entertainment. *Proceedings of the 2019 CHI Conference on Human Factors in Computing Systems*, 1–13. Glasgow Scotland Uk: ACM. https://doi.org/10.1145/3290605.3300930

McCarthy, J., & Wright, P. (2004). *Technology as experience*. MIT Press.

Pelowski, M., Forster, M., Tinio, P. P. L., Scholl, M., & Leder, H. (2017). Beyond the lab: An examination of key factors influencing interaction with "real" and museum-based art. *Psychology of Aesthetics, Creativity, and the Arts, 11*(3), 245–264. https://doi.org/10.1037/aca0000141

Petrelli, D., Ciolfi, L., van Dijk, D., Hornecker, E., Not, E., & Schmidt, A. (2013). Integrating material and digital: A new way for cultural heritage. *Interactions, 20*(4), 58–63. https://doi.org/10.1145/2486227.2486239

Petrelli, D., Marshall, M. T., O'brien, S., Mcentaggart, P., & Gwilt, I. (2017). Tangible data souvenirs as a bridge between a physical museum visit and online digital experience. *Personal Ubiquitous Comput*. 21(2), 281–295. DOI: https://doi.org/10.1007/s00779-016-0993-x

Pintanel, M. À. (n.d.). 006 [Wax painting]. https://www.flickr.com/photos/pintanescu/7990979289

Rosenberger, R., & Verbeek, P.-P. (2015). A field guide to post-phenomenology. In R. Rosenberger & P.-P. Verbeek (Eds.), *Postphenomenological investigations: Essays on human-technology relations* (pp. 9–41). Lexington Books.

Ryding, K., & Fritsch, J. (2020). Play design as a relational strategy to intensify affective encounters in the art museum. Proceedings of the 2020 *ACM Designing Interactive Systems Conference*, 681–693. Eindhoven Netherlands: ACM. https://doi.org/10.1145/3357236.3395431

Ryding, K., Spence, J., Løvlie, A. S., & Benford, S. (2021). Interpersonalizing Intimate Museum Experiences. *International Journal of Human–Computer Interaction*, 1–22. https://doi.org/10.1080/10447318.2020.1870829

Schiphorst, T. (2009). soft(n): Toward a somaesthetics of touch. *Proceedings of the 27th International Conference Extended Abstracts on Human Factors in Computing Systems - CHI EA'09*, 2427. Boston, MA, USA: ACM Press. https://doi.org/10.1145/1520340.1520345

Shusterman, R. (2014). Somaesthetics. In M. Soegaard & R. F. Dam (Eds.), *The encyclopedia of human-computer interaction* (2nd ed. [online]). The Interaction Design Foundation. https://www.interaction-design.org/literature/book/the-encyclopedia-of-human-computer-interaction-2nded/somaesthetics

Spence, J., Bedwell, B., Coleman, M., Benford, S., Koleva, B. N., Adams, M., Row Farr, J., Tandavanitj, N., & Løvlie, A. S. (2019). Seeing with New Eyes: Designing for In-the-Wild Museum Gifting. *Proceedings of the 2019 CHI Conference on Human Factors in Computing Systems*, 1–13. https://doi.org/10.1145/3290605.3300235

St. John, B. (2016, May 17). An eye for detail: Zoom through 1,000 artworks thanks to the new Art Camera from the Google Cultural Institute. https://blog.google/outreach-initiatives/artsculture/art-camera-cultural-institute/

Steier, R. (2014). Posing the question: Visitor posing as embodied interpretation in an art museum. *Mind, Culture, and Activity, 21*(2), 148–170. https://doi.org/10.1080/10749039.2013.878361

Stolterman, E., & Wiberg, M. (2010). Concept-driven interaction design research. *Human-Computer Interaction, 25*(2), 95–118. https://doi.org/10.1080/07370020903586696

Tennent, P., Martindale, S., Benford, S., Darzentas, D., Brundell, P., & Collishaw, M. (2020). Thresholds: Embedding virtual reality in the museum. *Journal on Computing and Cultural Heritage, 13*(2), 1–35. https://doi.org/10.1145/3369394

Thompson, C. J., Locander, W. B., & Pollio, H. R. (1989). Putting consumer experience back into consumer research: The philosophy and method of existential-phenomenology. *Journal of Consumer Research, 16*(2), 133. https://doi.org/10.1086/209203

Uglow, T., Richards, J., Osborn, J., & Sillitoe, K. (2017, September). XY-Fi. https://experiments.withgoogle.com/xy-fi

Vallgårda, A. (2014). Giving form to computational things: Developing a practice of interaction design. *Personal and Ubiquitous Computing, 18*(3), 577–592. https://doi.org/10.1007/s00779-013-0685-8

Vermeeren, A., Calvi, L., & Sabiescu, A. (Eds.). (2018). Museum experience design: Crowds, ecosystems and novel technologies. Springer International Publishing. https://doi.org/10.1007/978-3-319-58550-5

Walter, T. (1996). From museum to morgue? Electronic guides in Roman Bath. *Tourism Management, 17*(4), 241–245. https://doi.org/10.1016/0261-5177(96)00015-5

Wessel, D., & Mayr, E. (2007). Potentials and challenges of mobile media in museums. *International Journal of Interactive Mobile Technologies (IJIM)*, 1(1). https://online-journals.org/index.php/ijim/article/view/165

Woodruff, A., Aoki, P. M., Hurst, A., & Szymanski, M. H. (2001). Electronic guidebooks and visitor attention. *Proceedings of the International Conference on Cultural Heritage and Technologies in the Third Millennium*, 437–454.

Publication B

Art Critique by Other Means

Christian Sivertsen, Maarten Smith and Sander van der Zwan

The paper has been published in
Designing Interactive Systems Conference (DIS '23), July 10–14, 2023, Pittsburgh, PA, USA
ACM, New York, NY, USA, 10 pages. https://doi.org/10.1145/3563657.3596069

The layout has been revised.

Abstract

HCI projects bringing digital and interactive technologies into art museums are affected by a conception of the relation between art and design that narrows the available design space. This is often done by positioning such technologies either as presenting information about artworks, as artworks themselves, or by taking a 'hands-off' approach. Aiming to re-investigate this conceptual space, we draw on an enactive approach to art developed by Alva Noë and, by means of a design example (an installation in the MUNCH museum), discuss how this approach redirected the design process. Specifically, how it affected the way the design team related to the original artworks and their history, and how they approached evaluating their work. We show by example how understanding such technologies in terms of how they educate visitors' attention, similarly to art critique, allows designers to participate in correspondence with artworks using unique material, aesthetic, and embodied means.

B.1 Introduction

As art museums are aiming to make themselves relevant to new audiences they see the potential of using digital interactive media to create new and relevant art experiences. The integration of such media into the art museum space raises conceptual challenges. One of these is the question of how to relate designed experience technologies to the works of art on display, a challenge shared by a variety of researchers in the HCI community. In this paper, we aim to contribute practice-relevant theoretical resources to this discussion by exploring how HCI researchers have positioned their designs in relation to artworks and documenting and discussing an example case in which a design team employed an enactive approach to conceptualize this relation. Such conceptualization is pivotal for the way one might design technologies mediating a relationship to artworks. By sharing and discussing these concepts we provide those working within art museum contexts with theoretical tools to consider their position and open up new directions for designers by rethinking how to understand the challenge of making experience technologies in art museums. The goal is a theoretical contribution. Rather than prescribing design guidelines, we organize a discussion on how the way we conceptualize and motivate our design shapes what might be designed in this domain, i.e. the design space.

The design case in question is an exhibition concept developed by one of the authors of this paper with the new MUNCH museum in Oslo, Norway, a single-artist museum dedicated to the painter Edvard Munch (1863-1944). The exhibition opened in October 2021, under the name *Poison*, and was available to the public for three months. The experience does not contain original artworks but is based on the dig-

ital archive of the museum as a resource for visual content. While developing this exhibition the design team faced the challenge of positioning their work in relation to Edvard Munch's works of art.

In the following, we first investigate how HCI researchers have conceptualized their work in relation to works of art in museums. Then, we introduce several theoretical concepts, namely Alva Noë's enactive approach to art (Noë, 2015) and Tim Ingold's reading of *education of attention* (Ingold, 2018), and reflect on a concrete example in which a design team actively engaged with these concepts to make sense of their work and used them as a *generative lens* (Zwan et al., 2020) to direct their design process. We compare this approach to related projects to illustrate how this theoretical lens throws new light on earlier projects and how it created new opportunities for the design team working with *Poison*. We conclude by discussing the implications of these concepts on the available design space, and how these concepts might have an impact on other experience technologies being designed for the art museum domain.

B.1.1 Digital and interactive exhibitions in art museums

Before exploring the landscape of literature on the design of technology-based mediation projects in art museums, we should note that the challenge of understanding and positioning digital and interactive technologies in this domain carries a broader relevance outside the HCI research community. Though such applications of technology have established a presence in, e.g., natural history and science museums over the last few decades, they have not been adopted in the art museum domain with the same readiness. When seen, they are mostly found in the form of audio guides, information kiosks, and database browsers, providing information to the visitor.[1] One explanation for this is that the benefits of these technologies are not entirely clear for use in this space and that they might seem to conflict with the enjoyment of the art itself, a phenomenon sometimes referred to as the *heads-down phenomenon* (Hsi, 2003; Lyons, 2009; Petrelli et al., 2013; Walter, 1996; Wessel & Mayr, 2007; Woodruff et al.,

[1]An exception to this is the current trend of immersive art shows exhibiting e.g., Van Gogh, Monet, Dali, Frida Kahlo, and others in grandiose shows with large-scale projected replicas of famous paintings enhanced through animation and sound. At the time of writing at least 14 different actors have produced or are currently producing immersive shows with more than 10 different artists around the world.e.g. https://grande-experiences.com/van-gogh-alive/, http://www.fridakahloexperience.com/, https://monetexpo.com/, https://www.atelier-lumieres.com/en/cezanne-lights-provence. We assume that the designers of these exhibitions face similar challenges, but consider this trend outside the scope of this paper because these exhibitions tend to be regarded as experiences on their own, separated from the museums in which the original artworks they replicate are displayed. We also note that only a few art museums have made dedicated experiences integrating such experiences with their art historical expertise, see e.g., *Meet Vincent* by The Van Gogh Museum, Amsterdam which is a traveling experience (Van Gogh Museum Amsterdam, 2016), and the *Leonardo: Experience a Masterpiece* which was shown at and produced by the National Gallery, London in 2019-2020 (The National Gallery, 2019).

2001). Another explanation, as Samis (2018) argues, is that the move of the digital technologies from the periphery to the center of museums reveals a fundamentally conservative nature in most art museums.

B.2 Related Work

In this section, we investigate literature that documents the *making process* of technology-based mediation projects in art museums and investigate the way these authors have conceptualized their work. Much of the research regarding the design and evaluation of interactive exhibitions, has been established within the HCI and Interaction Design communities (e.g., (Andersen & Ward, 2017; Dindler et al., 2007; Hornecker & Ciolfi, 2019; Rizzo & Garzotto, 2007; Tennent et al., 2020; Woods et al., 2004)). However, theoretical work applicable to the art museum domain has only been developed by these communities to a limited extent (Bedford, 2014; Claisse et al., 2020; Maye et al., 2014). In the broader domain of cultural heritage, Claisse et. al, building on Bedford's work in *The Art of Exhibition Design* (Bedford, 2014) argue for reformulating exhibition design as "creation of an aesthetic experience where design facilitates the experiential by using the language of the arts" (Claisse et al., 2020, p.3) with "objects rich in meaning; stories that evoke emotions such as empathy; metaphorical play to forge new connections; design that melds space, light, image in one integrated experience" (Bedford, 2014, p.16). In this sense, designers of cultural heritage exhibitions can be understood as doing artistic work, as Bedford indeed acknowledges:

> "all of us on the *Teen Tokyo* exhibition team were artists. Just as the aesthetic experience encompasses both the art and the visitor, it includes the maker as well." (Bedford, 2014, p.17)

While such an approach might be promising for science, natural history, or other types of cultural heritage museums, taking it to the domain of art museums is problematic because such interactive installations being understood as artworks could create wrong expectations, when shown alongside other artworks. It is rarely the role of art museums to be making the art themselves, rather the mediating tools produced by the museum itself offer perspectives on the art on display. Navigating this distinction is a balancing act that has been explored before and that we will revisit in this paper.

Much of the (HCI) literature on how interactive installations in art museums are conceptualized can be organized using a distinction Boehner, Sengers, et al. (2005) draw in their paper *Opening the Frame of the Art Museum*: "Interactive installations in museums ordinarily have one of two goals: they are either artworks to be appreciated as part of the museum's collection, or they are tools that disseminate information about museum objects." (Boehner, Sengers, et al., 2005, p.1). We begin by showing

several examples that neatly fit this distinction, and follow this up with examples that escape it. Then, we discuss the distinction itself in light of the conceptual challenges these authors are facing.

The first example is the work of Kortbek & Grønbæk, for whom the goal of their installation in the art museum ARoS is to "communicate information" (Kortbek & Grønbæk, 2008, p.9). They explain this focus by stating that in contrast to other types of museums, such as history and science museums where interactive installations are more frequently used, the artworks in the art museum should take center stage.

> "Thus [in cultural heritage and science museums] the interactive installations themselves become objects of the museum, revealing knowledge about certain immaterial subjects. In art museums, however, the artworks themselves should constitute the main visitor experience." (Kortbek & Grønbæk, 2008, p.1)

In this way the artworks are seen as speaking for themselves, making it particularly challenging to introduce technology in art museums "without disturbing the domain of the artworks" as Kortbek and Grønbæk (2008, p.1) describe the challenge. Another example is Terrenghi & Zimmerman who motivate their work on a spatialized audio guide by stating that "the aim is to highlight the potential of augmenting the user's experience with the employment of personalized audio information."(Terrenghi & Zimmermann, 2004, p.1). The idea of augmenting the visitor experience by adding additional information through sound or visual content is common in the art museum domain and evidenced by the ubiquitous audio guides and guided tours as well as other sound-based augmentations (Drotner et al., 2018; Kortbek & Grønbæk, 2008; Pau, 2017; Yamazaki et al., 2009). A few projects achieve augmentation through visual techniques like AR, VR, and projections superimposed on artworks. These often address technical aspects of the art, like a particular painting process (Bimber et al., 2005; Louvre, 2020; The National Gallery, 2019). Finally, some projects from the HCI community are directly presented as artworks (Costello et al., 2005; Hindmarsh et al., 2002; Tennent et al., 2020).

So far, these earlier projects settle nicely under Boehner, Sengers, et al.'s distinction between technology as art and tool. However, Boehner, Sengers, et al. are unsatisfied with it, as, on their view, both these approaches tend to put visitors in the position of non-experts (Boehner, Sengers, et al., 2005, p.123). To break with the dichotomy, Boehner, Sengers, et al. present a design, "[...] that was both tool-like, giving patrons information about the museum and their interactions with it, and art-like, creating opportunities for new experiences in an aesthetic and open-ended way." (Boehner, Sengers, et al., 2005, p.123). Their installation aims to subtly make visitors aware of the boundary of the museum space and the presence of people, by playing back bird sounds in areas of the museum where visitors are not present. This project focuses more on the architecture of the museum and the relation between museum

space and visitors in a general sense rather than their relation to the artworks. Reflecting on this, the authors acknowledge that the art-tool distinction is blurry, but do not discuss it further. Instead, they call for a deeper investigation into how art practices and HCI practices might influence each other. Since Boehner, Sengers, et al.'s paper was published, a number of other interactive projects have been designed and written about that also do not fit this distinction well, and neither do they fit the "frame" that Boehner, Sengers, et al. (2005) open up.

For example, a more recent take on augmentation is based on an interpretive stance in the museum space where the affective relations of the visitor take center stage. One such project is *Sensitive Pictures* (Benford et al., 2022). Here the aim is to create an emotional provocation through an app-based prototype. The goal is not to provide the visitor with information, but rather to stimulate the visitor's affective engagement with the art on display. This is done through storytelling, immersive audio, and an interactive interface that prompts the visitors to reflect on their experience with the artworks.

In the projects *ArtLinks* (Cosley et al., 2008) and *Imprints* (Boehner, Thom-Santelli, et al., 2005) art museum visitors are given the opportunity to add different kinds of tags and marks to selected artworks, in order to establish social clues and relationships via the artworks. The projects *Never Let Me Go* (Ryding, 2020), *MuseUS* (Coenen et al., 2013), *MUSE* (Cosley et al., 2008), MobiTags (Cosley et al., 2009), and *Gift* (Spence et al., 2019) similarly support various forms of social interaction between the visitors of the art museum, but take a 'hands-off' approach, leaving it up to the users of the systems which artworks to include in the experience and in some cases whether to include any artworks at all. These technologies intervene with the social situation around the artworks and thus stage them in new and different roles across projects. The project *Cluefinder* weaves a mixed-reality detective story into the physical setting of an art museum, but the game does not engage with the history of the museum, nor the art exhibitions by design (Lange et al., 2019).

Another kind of project is *ARTLENS Gallery* (Alexander et al., 2017). In the middle of a large gallery, several artworks are on display. Lining the walls surrounding them are several interactive installations where the visitors can play, draw and move their bodies to investigate, manipulate or create different art content. Some of these installations refer to specific artworks that are on display, others refer to art and creative practice more generally. The authors argue that the artworks are in the foreground while the interactive installations make up the background, however, this balance has dramatically shifted in comparison to more conventional art exhibits. While some of the *ARTLENS Gallery* installations are indeed information augmentations, others seem to be more self-contained. The authors describe working on physically and conceptually detaching the interactive installations from the specific artworks as they upgraded the exhibition from the previous Gallery One concept to *ARTLENS Gallery*.

Alexander et al. write that "While visitors are having fun, they are also looking closer, making connections, and gaining comprehension that will enhance their appreciation of art throughout the museum." (Alexander et al., 2017). Thus instead of merely providing information next to the art, the *ARTLENS Gallery* project tries to bake relevant conceptual and thematic art historical content into the act of participating in the interactive experiences themselves. The focus here is not necessarily directed at a specific artwork, but also extends to the art on display in the rest of the museum. Alexander et al.'s claim about the resulting experience is only backed by empirical research in a limited sense (Bolander et al., 2018) but warrants further investigation.

In short, a variety of authors in the domain of HCI are confronted with the challenge of relating their designs to works of art in museums. Roughly speaking, many of these authors either focus on interactive installations that are presented as being artworks themselves (e.g., (Costello et al., 2005; Hindmarsh et al., 2002; Tennent et al., 2020)) or create new designs that disseminate information about works of art in such collections (e.g., (Kortbek & Grønbæk, 2008; Terrenghi & Zimmermann, 2004)) or create ones that forego making direct connections to the artworks altogether (e.g. (Boehner, Sengers, et al., 2005; Ryding, 2020)). In some cases, the effects of the designs on the visitors' experience of the art are not considered beyond the concept of "enhancement". When they are, authors seem to respond to a similar challenge: to avoid interference with visitors' experience of other works of art on display.

Authors such as Benford et al., Ryding, Spence et al., and Alexander et al. (Alexander et al., 2017; Benford et al., 2022; Ryding, 2020; Spence et al., 2019) find themselves on different terrain. Their works clearly do interfere with visitors' experience of the artworks. How to understand this interference? What theory of art might make room for such designs?

B.3 Theoretical Concepts

One approach to art that aims to do justice to the interwovenness of aesthetic experience, is the enactive account developed by Alva Noë in *Strange Tools* (Noë, 2015). The enactive approach can be roughly characterised by the following background assumptions. Firstly, cognition cannot be reduced to passive brain-processes inside the head. Rather, it emerges from processes distributed across brain, body and environment. Secondly, the world is not pregiven, but is codetermined by action and cognition (Gallagher, 2017; Thompson, 2007; Varela et al., 1992). Noë's work in Strange Tools is of particular relevance to HCI researchers concerned with the intersection of design and art, as Noë builds upon work that is well-known to the HCI community, among others: John Dewey, James Gibson, and Martin Heidegger. Throughout the process of designing *Poison*, the exemplary case that will drive our theoretical focus, the design team engaged with this approach to make sense of what they were making

in relation to Edvard Munch's works of art. In other words, the designer used these concepts as a generative lens, (with foresight) to inform or guide the design process, as opposed to their typical use, i.e. to analyze the art practices (in hindsight) (Zwan et al., 2020).

With the aim of providing designers and researchers concerned with interactive exhibition design in art museums with some theoretical grip on their position, we introduce several notions advanced in this approach that turned out to be particularly relevant for the design of *Poison*. We then discuss how this conception of art shaped the understanding of the design space. In particular, we employ three claims advanced in the approach. First, art investigates and reorganizes the habitual modes of organization that we find ourselves caught up in. Second, the potential for artworks to reorganize our practices depends, in part, on their "difficulty". Third, aesthetic experience is made and achieved in the correspondence between, among others, perceivers, artists, art objects, and critics. It is not a static undergoing of an object's qualities, nor a projection of a subject on an art object. This is followed up with an introduction of Tim Ingold's reading of education of attention: attention *educates* by exposing us to the world and, by virtue of this experience, our attention *is educated* (Ingold, 2018).

B.3.1 Art aims at understanding that is reorganizational

Noë argues that art practices investigate and illuminate the ways we find ourselves habitually organized, and how we might organize ourselves. Art *puts them on display*. It presents these ways of organizing perspicuously, so that we may come to see and understand them. One of the examples he gives is how a painting of shoes, such as one by Van Gogh, can put on display the manner in which shoes hang together with a way of living. Understanding the shoes means appreciating what it is to be someone who relies on shoes in their everyday life. "A painting can let us catch ourselves in the act of unthinking engagement with the world." (Noë, 2015, p.200).

Art helps us *understand* what we habitually do. How we, for example, ordinarily think, perceive, talk, dance, or make pictures. In other words, art affords us a transformation from *not seeing* to *seeing*, and by doing so, it reorganizes us. Art is, therefore, not about the things made, the painting, sculpture, digital render, and so forth, but about the experience the artwork co-constructs. It is about the understanding or reorganization it brings forth.

> "We start out not seeing what is there. But by looking and interrogating and challenging, we come to see it. The work challenges us to reorganize our seeing, our expectations, and our thinking. The work of art, like that of philosophy, is the reorganization of ourselves. And this reorganization, this work, aims also at understanding." (Noë, 2015, p.138)

B.3.2 Art's reorganizational power depends on the possibility for it to bore us to death

Characteristic of art objects then is that they are not about the pleasure we may find in them, but about the challenges they pose to us. They challenge us to "See me if you can!" (Noë, 2015, p.102), to make sense of them and bring them into focus. This, however, brings along with it, the necessary possibility of not seeing.

> "Art affords us the opportunity to be bored to tears, when almost nothing else in our life does. And art's potential to be dull does not contradict the fact that art also moves and thrills and transforms and excites us. Indeed, it is the opposite side of the very same coin. Just as there is no encounter with love without the live risk of heartbreak, so there can be no confrontation with art that does not open up the possibility of getting lulled unconscious and bored to death. Art is valuable only in direct proportion to the degree to which it can, or might, bore us." (Noë, 2015, p.114)

By the same token, on Noë's account, a work of art can cease to be a work of art if it ceases to challenge us. Once we, for example, have named and categorized a piece of art, it may stop to ask "what is this?". As such it is a transformation from seeing to not seeing. Noë also raises a careful warning about museums' ambitions to make art "available" to ever more visitors. While it is a commendable goal, it has the potential to be counterproductive (Noë et al., 2020). As the reorganizational power of works of art does not come for free, it depends on our active engagement with them, and so too on their *difficulty*.

B.3.3 Art experience is made and achieved in correspondence

The understanding works of art can afford us does not come for free. We need to work for it and it needs to be achieved. The *work* of art, in this sense, is the practice of foregrounding the ways we are habitually organized. These experiences are not isolated in the head of the perceiver or caused by the object, but *made* in the correspondence between, among others, artists, art objects, perceivers, and art critics. Noë draws a similarity between works of art and jokes. A joke is not a thing on its own. It's a transaction, a response. And getting a joke is not just reacting to a stimulus, rather it's an accomplishment that requires a background of shared understanding. An upshot of the notion that aesthetic experience is something we achieve or make is that art requires meaning-making. Art happens "in the space of criticism" (Noë, 2015, p.111).

In other words, works of art require criticism in order to put our habitual organized activity on display. Crucially, however, aesthetic criticism involves more than merely disseminating facts about a work. It's also about making sense of it and evaluating it [2].

> "Good critics do not merely describe a work, they also bring to our attention qualities we had missed, or persuade us to give weight to features that we had ignored or failed to consider. Criticism doesn't proceed by logical argument; there is nothing like knock-down argument in this vicinity. Criticism proceeds by persuasion. Critics are educators. They teach you to see." (Noë, 2015, p.202)

B.3.4 Education of attention

Noë describes the way critics educate in terms of how they draw attention to a work of art, rather than as the transmission of knowledge. In the process of designing *Poison*, the design team drew on Tim Ingold's account of skill acquisition, in particular the notion of "education of attention"(Ingold, 2018), to understand how art critics might "teach us to see".[3]

Ingold, corresponding with the work of Jan Masschelein (Masschelein, 2010a, 2010b) and James Gibson (Gibson, 1979), holds that education is not the transfer of knowledge from one mind to another. Rather, in a process of education of attention, a skilled practitioner exposes a novice to relevant features of the environment that they may not have otherwise noticed. In doing so, the novice gradually fine-tunes their perceptual skills, becoming more attentive to the possibilities for action the environment affords them. For Ingold, attention both educates and is educated.

> "In the first case, attention educates by exposing us to a world in formation, by letting it in. But in the second, attention is what is educated, by dint of this experience. In truth, however, there cannot be one without the other. Submission and practical mastery are two sides of the same coin. That coin is the principle of habit." (Ingold, 2018, p.32)

On this view, when a critic educates a visitor about, for example, a painting, their aim is not to transfer knowledge *about* the painting, but rather to draw them into *correspondence with* it (Ingold, 2013, 2018).

[2] This understanding of criticism is not entirely new to HCI, as versions with theoretical affinities have been developed as a method for assessing interaction design (Bardzell, 2011; Bertelsen & Pold, 2004). However, in this paper, our focus lies on the critique of art, not interfaces.

[3] Since the time of designing, Alva Noë has published *Learning to Look* (Noë, 2021), which develops this aspect of engagement with and critique of art in more detail. It bears many similarities with Ingold's reading of education of attention, showing that aesthetic experience emerges in a process of both confrontation and reorganization (see e.g., p.112).

B.3.5 From avoiding interference to art critique

If we, as Noë suggests, understand art not by looking at isolated art objects or the effects they have on us, but rather in terms of what we *do* with such objects, this has consequences for the way we think about the relation between designs and works of art in a museum. While leaving open the possibility for designs to not "interfere" with an aesthetic experience, it draws our attention to the things that *do* co-determine this experience, but usually go unnoticed. Designs like audio guides, flyers, placards, frames, walls, or lighting, which are typically considered "neutral", turn out to have a hand in drawing visitors' attention to certain features of the art on display, and thus all co-shape the aesthetic experience.

On this view, avoiding interference with visitors' aesthetic experience ceases to be a relevant ideal for designers working in art museums. Rather, the focus shifts to how designed objects might fruitfully participate in art as a reorganizational practice. This view gives room for designers to, like art critics, participate in correspondence with works of art and educate visitors' attention. At the same time, such participation calls designers to take care that what they design keeps the "difficulty", on which the reorganizational power of artworks partially depends, intact. How might this practically be done? To explore this question, we present a design case in which one of the authors designed an interactive exhibition for the MUNCH museum in Oslo.

B.4 Design concept: Poison

The design used as the case for this paper is an experience technology-based exhibition called *Poison* that opened at the new MUNCH Museum in Oslo, Norway in October 2021. *Poison* is based on a series of paintings by Edvard Munch known as *The Green Room*. It is a little-known series of 7 paintings created in 1907. Common to these paintings is that they take place in a room with greenish wallpaper and a door in the back wall. In most of them, a round table and a couch are also present. Between the paintings, the wallpaper and the dimensions of the room change, and in each of them a different scene with different characters plays out.

Poison is set in a small 40 m2 room in the museum and is separated into three chambers that are clearly separated by curtains and visited sequentially.

B.4.1 Chamber 1: The Triptych Chamber

This chamber is meant to set the scene. Upon entering, the visitors are met with a soundscape of dark pulsating sounds. Inside, paintings of *The Green Room* series are projected into three frames. The projected paintings slowly change and morph into other paintings from the series. The visitors' movements in the chamber will make this morphing speed up, and three of the paintings from the series will come in and

Fig. B.1: The first chamber. Three framed canvases are hung on the wall onto which images from the series are projected. As the visitors move, the images morph into each other and change their order in the frames.

out of view in the frames. For each step the visitor takes in this chamber, distorted electronic sounds are played back from speakers near the floor. On the wall a banner reads "THE GREEN ROOM SEES YOU" and through a semi-transparent curtain, the visitors can get a view of the main chamber.

B.4.2 Chamber 2: The Main Chamber

In this chamber, each projection is a blurry and distorted version of one of the paintings from the series. But the images are not stable. As visitors move around the space the perspective of the paintings shifts to accommodate the viewing location of the visitors in the space, making the walls of the paintings move. An oscillating bass sound emanates from the sound system. Depending on the number of visitors the perspective becomes more or less tied to your individual position in the chamber. In this way, the chamber is reactive to the presence of the visitors but may seem somewhat autonomous in its erratic behavior. On the curtain, a banner states "THE GREEN ROOM IS POISONOUS".

B.4.3 Chamber 3: The Magenta Chamber

This chamber was designed to round off the experience. On a magenta backlit wall, a short text provides a bit of contextualizing information about the Green Room series, followed by a number of questions about the nature of the room and the characters

Fig. B.2: A peek through the curtain to the second chamber. Three walls are covered in large projections of paintings from The Green Room. As the visitors move inside the chamber the perspective of the projected paintings changes. Photo Credit: Kilian Munch ©Munchmuseet

in it. On the exit curtain, a banner reads "THE GREEN ROOM IS EVERYWHERE". After stepping into the light in the hallway, visitors who spent time reading the text will experience their vision overcompensating for the bright magenta light covering everything in a green tint until their vision readjusts after a few seconds.

B.5 Making Poison with Noë: A new design space

The museum, being a single-artist museum, has an ethos of deep engagement with Munch's life and oeuvre. Designing in this context meant that *Poison* would become part of that engagement. Throughout the design process, this raised a number of questions on how to mediate a meaningful relation between *Poison* and Munch and his original works. This is where our thinking with Noë's understanding of art gave direction to the unfolding of the project. Rather than taking *Poison* to be a neutral environment for the original works to be viewed in, nor as an augmentation of the originals, we positioned it as a kind of art critique. That is, as an educator of attention in viewing Munch's original works.

In the following section, we will show how we attended practically to establishing this relation while comparing it to some of the projects presented in the Related Works section. In particular, we will consider:

1. How to mediate the correspondence with the original artworks.

Fig. B.3: The third room is lit up in magenta light. Reading the short text affects the color perception of the visitors.

(a) How to relate to the original and physical art objects.

(b) How to correspond with the art historical context.

(c) How to correspond through non-verbal means.

2. How to evaluate such a design.

B.5.1 Corresponding with original artworks

To make clear how the understanding of *Poison* as art critique shaped the concrete design and how it compares to earlier approaches, we start by highlighting design choices from projects described in the related work section before presenting ours. The first subquestion concerns the spatial relation between the design and the original works of art. Kortbek and Grønbæk (2008) argued that it is important to avoid disturbing the domain of the artworks and the "pure art experience", so they placed their communication installations in a room separate from the artworks. Boehner, Sengers, et al. similarly placed their work away from the artworks but in an attempt to "Encourage reflection on other aspects of museum experience beyond just the individual art objects." (Boehner, Sengers, et al., 2005, p.4). *Never Let me Go* and *Gift* took place in art museum galleries, even though they were not tied to a pre-defined selection of artworks. Instead, they utilized an *interpersonalization* (Ryding et al., 2021) strategy to let visitors choose which artworks become involved and how, thereby modulating relations between visitors and artworks through social interaction. In

the *ARTLENS Gallery* gallery, the artworks and interactive installations were situated in the same room. In some cases, there were connections between interactive stations and physical works. In other cases, related artworks were presented on a digital display. Sometimes the interactive stations tie more broadly to concepts within art and creativity. *Sensitive Pictures* on the other hand requires the presence of specific artworks for which the audio content has been tailored.

In *Poison* the original artworks from *The Green Room* were not present. However, this was not done out of a wish to shield them from disturbance. The reason for leaving out the originals was instead mostly practical. Firstly, the paintings from *The Green Room* were not on display in the museum at the time. Secondly, the size of the space *Poison* was set in was limited and too difficult to configure towards an environment that could protect the original paintings from damage. Had the practical circumstances been different, the design team would have loved to work with the original paintings as an integrated part of the installation. The theoretical view that we took would not have prevented that. On the contrary, understanding art as not being about the objects but about the reorganizational experience they help co-construct leaves open whether or not the art should be "close-by". Similarly posing an art critique does not depend on the physical presence of the artwork. However, engaging with the originals before or after engaging with the critique will further add to this ongoing correspondence.

The second subquestion concerns the art historical context. In contrast to *Gift* and *Never Let Me Go*, *Poison* is like *ARTLENS Gallery*, *Sensitive Pictures* and Kortbek & Grønbæk's work tied to a specific selection of artworks. Both in *Sensitive Pictures*, *ARTLENS Gallery* and Kortbek & Grønbæk's work, art professionals were consulted to create content grounded in the ongoing discourses and knowledge about the related artworks. Where the approach of *ARTLENS Gallery* and Kortbek & Grønbæk seem to expand from the works into broader ideas in art history, *Sensitive Pictures* and *Poison* zoom in on particular aspects of the original works. The main difference here is that the dramatized audio snippets in *Sensitive Pictures* were connected to emotions like love, self-confidence, and passion, while *Poison* addresses more embodied, spatial, and atmospheric characteristics. This is naturally shaped by the design team's engagement with Signe Endresen's art historical dissertation on *The Green Room* (Endresen, 2015) as a partner in the conversation around the design. The team did not just read Endresen's research, but as an employee of MUNCH, she was part of an internal reference group giving feedback on the design through various iterations of concepts and prototypes. This includes discussing the concrete design decisions against her perspective on the original paintings. In this way, the designers qualified their art critique by engaging with the ongoing art historical conversations about Munch's work and letting it inform the aesthetics of *Poison*. Such active engagement with the art on display makes sense from the idea of correspondence present in our

theoretical frame. Participating as designers in the ongoing *work* of art, one has to get an understanding of that particular work to be able to respond to it and point things out about it to others.

With a predefined selection of artworks comes the necessity to establish a relation between the mediating technologies and the artworks for the audience. Related to the third subquestion, Kortbek and Grønbæk (2008) argue for the use of *conceptual affinity*, which in their case means using shapes and compositions from the original artworks in the interactive installation. In this way, visual similarities are used as a tool to mainly indicate the relation to the artwork while the communication of "sources of inspiration" (Kortbek & Grønbæk, 2008, p.3) happens through art historical references presented in text and images. To avoid disturbing the domain of the artworks they also used "gentle audio augmentations" (Kortbek & Grønbæk, 2008, p.3) that were placed next to the artworks, as audio would not interfere with the mainly visual and sculptural art on display. In an almost opposite way, *Poison* and *ARTLENS Gallery* use words and images as the primary means to establish the relation to specific artworks while the aesthetic emerging from the use of architecture, sounds, projections, and the embodied interaction was part of the attempt to educate the attention of visitors to the aesthetic relation in the originals. This encompassing use of various aesthetic strategies simultaneously becomes a possibility with the theoretical view taken up in this paper. Artworks never were isolated objects to begin with, and even supposedly neutral designs such as flyers, placards, frames, walls, lighting, and so too Kortbek & Grønbæk's inspirational content and gentle audio augmentations co-shape our experience of them. The aim then is not to avoid such participation, but to see how different media can play a fruitful role in the reorganizational *work* of art. To illustrate how different media played a role in establishing an aesthetic tied to the art historical correspondence in *Poison* we present two examples. First, Endresen (2015) argues that the scenes in the paintings might or might not take place in parallel. That is, even though the paintings could be viewed one after another to read almost like a comic strip of sorts, there is no evidence to suggest that they should be seen in this way. Thus the time relation between the images is undefinable. Responding to this quality of non-linearity, we made it so that the projected digital replicas in the triptych chamber would morph and switch places as visitors walked past them. A second quality that Endresen showed us, is that the furniture and characters are placed in the immediate foreground in the paintings, giving the viewer a role as both spectator and participant in the situations playing out in the paintings. Besides this, Endresen views the characters in the paintings as potential avatars for the viewers to project themselves into. In order to correspond with this quality we created a forced perspective in the projections in the main room. This means that when standing in a specific spot the projections look like an extension of the physical space. In response to Endresen mentioning the audience as both spectator and participant, we made the projections in the triptych and main chamber change based on the movements of the

visitors. More specifically we used camera tracking to adjust the perspective of the projections dynamically to the position of visitors in the main chamber. In the triptych chamber, the paintings would as mentioned morph and switch places in relation to visitor movement.

B.5.2 Evaluation of the experience

In this section, we look at how our conceptual understanding changed our way of evaluating *Poison*. It is not our aim to evaluate the installation here, as this task is taken up in other work. Rather we discuss on which terms to evaluate. Meaningful evaluation depends on what one pays attention to. Evaluating a design if the goal is, like Kortbek & Grønbæk, to "communicate information [...] without disturbing the domain of the artworks." (Kortbek & Grønbæk, 2008, p.1), one will start looking for ways to see how well information has been provided to the visitor without disturbing the art experience. Indeed Kortbek & Grønbæk report that "50% of the respondents claim that they got 'much knowledge' or 'some knowledge' out of the installations" (Kortbek & Grønbæk, 2008, p.7). However, they also conclude that "[...] seen from an art critics point of view we have blurred the borders between the art per se and the communication of the art. [...] However, neither the artist nor the curator find the blurring problematic, in fact they perceive the holistic experience a quality of the exhibition" (Kortbek & Grønbæk, 2008, p.9). This conclusion aligns well with our suggestion that avoiding disturbance is a moot point. The white paper published on the evaluation of the *ARTLENS Gallery* reports that "[...]participants were likely to agree that their visit to ARTLENS Gallery enhanced their overall museum experience (76%), encouraged them to look closely at art and notice new things (74%), and increased their interest in the museum's collection (73%)" (Bolander et al., 2018). This gets closer to the idea that engagement with the experience technologies affects the experience of art in other places in the museum. In the evaluation of *Gift* participants report seeing the art with "other" or "fresh" eyes (Spence et al., 2019) and attempting to take on other people's perspective drawing their attention to aspects of the art that would otherwise have remained unnoticed. In this case, it would be the perspective of the people with whom they were exchanging gifts rather than curators or critics.

Through these examples, we see that the idea of the attention shift and the shifting perspective on art already appears in earlier work. This is well-aligned with the theoretical frame we propose. However, in these evaluations, it remains a side effect that is not fully investigated. Having the education of attention as a central tenet in our approach pushes us to investigate not only whether it happens, but more importantly what the attention of the visitors is being directed to. To illustrate the kind of responses we find useful in evaluating the "education of attention" of the visitors we will present a few examples from the evaluation of *Poison*. We conducted a total of 34 individual and group interviews with visitors of *Poison*. 16 of these were individual

interviews with members of the museum club, who saw the exhibition a couple of days prior to the opening of the museum. The age in this group ranged from 24 to 70 with a median age of 32. Another 19 interviews were conducted with museum visitors that were recruited just as they stepped out of Poison, after the public opening. These were interviewed alone or in small groups. In total 32 people participated in this round of interviews. The age range in this group ranged from 19 to 80 with a median age of 34. Across all 48 interviewees 30 identified as female and 18 as male. The interviews were phenomenological interviews, focusing on the visitors' immediate aesthetic experience of *Poison*. In addition, we asked them if they saw any relations between Edvard Munch's life and practice and what they had experienced in *Poison*.

In the analysis, we paid special attention to quotes that showed if and how the attention of the visitors was educated. For example, one quote from the interviews that stood out to us as important in this respect showed that *Poison* drew attention to the historical situatedness of the paintings:

> "Yes, I think it's a nice way to make you part of his art. You don't only look at the paintings, but suddenly you're really surrounded by the paintings and you get that visual experience of being in the room, the... because the things he's painting, they happen at some point in time, right? There was a green room, I guess, and it's nice to actually think about it. It's not just depicted as a picture that is done, but I was actually a whole story around that and there is an actual place where it took place."(Female, 32, Interested in art).

By experiencing herself as part of the scenes of Munch's paintings the historical setting of the paintings was foregrounded for this visitor. The paintings in a way came alive. They are not merely objects, but things that are part of a longer story. A story that is located in actual places at particular times.

In another case it was the text that offered a new perspective on the paintings and Munch's time, "I had never thought about the relation between arsenic and green. I mean, all these Munch people are like green in their face because of jealousy or sickness. But this thing about poison is not something I've really been reflecting on. There was arsenic in things. Things were poisonous. It was a scary world to live in. There was no department of environment and health that would check for arsenic in the water for example. It was a different time. That stayed with me." (Male, 41, Art collector). The visitor reflected on a fact about arsenic presented in the text in the exhibition, and his attention was in this way drawn to society during Munch's life.

It is not in all cases that the visitors responded with agreement to the qualities suggested by *Poison*, "I don't know. I don't quite feel that I understand how much of this experience is Munch. Was this really Munch? Did he think about the paintings in this way? Was something different?" (Male, 28, Not interested in art). The visitor here is skeptical of the aesthetic presented in *Poison* with regards to Munch's own

experience, and rightfully so. From a perspective of knowledge transfer and non-interference, this quote might be seen as problematic. However, thinking with Noë, we argue that this critical reflection on the aesthetic experience is well aligned with the idea of *Poison* working as a form of art critique. It may prompt the visitor to seek an understanding of this aspect in other works or other sources of information and thus drive new ways of engaging the visitor with Munch's art.

B.6 Discussion

Working in art museums requires HCI researchers to position themselves in relation to art. Many of these researchers tend to constrain their design space by positioning their designs either as presenting information about artworks, artworks themselves, or by leaving the involvement of artworks entirely in the hands of the visitors. A common thread in such positioning is an attempt to stay removed from the original art so that it can speak for itself i.e. they aim to be non-interventional. Reinvestigating this positioning *is* the opening of a new design space (and the closing of others).

Alva Noë showed an alternative way of thinking about how art *works* that had several consequences for the particular design case described in this paper. For Noë, art is about the reorganizational experience artworks help to co-construct together with, among others, visitors, artists, artworks, and critics. The reorganizational power of works of art partially depends on their "difficulty". Even supposedly neutral designs such as flyers, placards, frames, walls, lighting and the like participate in the work. Rather than aiming not to intervene, this way of thinking invites other participants to investigate what they may contribute to art's reorganizational work. For the designers hired to attract new audiences, such as the design team of *Poison*, this included walking the fine line between attracting visitors and engaging them in works of art, and taking care not to take away the reorganizational power of these works (e.g., by aiming only to make a more attractive and pleasurable experience).

And so, the design team of *Poison* actively engaged with the original artworks in the museum's collection, and positioned the design as a kind of art critique, furthering the visitor's conversation with Munch's art in a particular way, bringing into view features of the artworks they might have missed and persuading them to pay attention to things they might have ignored. In the design process, the design team argued with this conceptual understanding that the original works did not necessarily have to be present for the critique to be meaningful. Secondly, they showed that creating such critique demanded an active involvement of an art historian in the process of designing. The design team was under no illusions about the neutrality of this involvement, but at the same time respected their deep, rigorous, and thoughtful engagement with history. Thirdly, the understanding allowed the designer to use his unique material, aesthetic, and embodied means in co-constructing the reorganizational experience of the art by providing new ways of critically engaging with

the works. Rather than words or gestures, the exhibition relies on light, space, color, and sound to foreground and background themes and qualities of Munch's art. This is also how *Poison* stands out from other works. It has both a strong connection to existing works of art, while not using words as its primary medium. *Poison* engages directly with the artworks from *The Green Room* series. It does not describe them but directs the visitor's attention to particular qualities of color, space, atmosphere, and performativity. Finally, while evaluating *Poison*, an interest emerged in seeing how visitors' attention was drawn to different aspects of the artworks as a consequence of having experienced *Poison*. This contrasts with measuring "how much" was learned.

With this, much remains unsaid. In particular, we did not go into detail about the ways others might use the lessons learned in their own practice. Providing practical guidelines was not the aim of this paper. Rather we focused on opening up an alternative design space by introducing and showcasing the use of a specific conceptual understanding. However, now that this space has been introduced, further practical guidance can be explored in future work. Furthermore, while we provided several examples of how Poison was designed as art critique through non-linguistic means, more detailed evaluation of this work as critique could be fruitful. How is this work concretely done and does it do a good job? Such evaluation of Poison will be taken up in future research.

B.7 Conclusion

The understanding of *Poison* as a work of art critique is one that emerged during the design process and helped the design team put a finger on the particular way their work has a bearing on visitors' engagement with Munch's works of art. Considering there are other designers working on similar projects on similar sites, it is an understanding that might travel, perhaps with some changes or shifts, to other times and places. To be sure, it wasn't our aim to evaluate *Poison* as good or bad critique, rather we hope to have done two things: Firstly, to make designers in the art museum domain sensitive to the way their designs might educate attention and transform visitors' experience of works of art; And secondly, to have exemplified, with *Poison*, a new design space in art museums. A space in which designers can actively engage with their unique material, aesthetic, and embodied means in co-constructing the reorganizational work of art by providing new ways of critically engaging with the art in art museums. We hope that the design and research communities will join us in exploring further which design possibilities emerge from this theoretical frame and help in expanding the ways we evaluate emerging visitor experiences.

Acknowledgments

The authors would like to thank Nikita Mathias and the rest of the MUNCH Audience Lab for supporting and facilitating the development of the *Poison* project and the related research project. We thank Anders Sundnes Løvlie for his insightful comments on the manuscript. We are also thankful to all external collaborators and people who participated in evaluations of the design.

References

Alexander, J., Wienke, L., & Tiongson, P. (2017). Removing the barriers of Gallery One: A new approach to integrating art, interpretation, and technology. *Museums and the Web 2017: Selected Papers and Proceedings from an International Conference*. Retrieved December 15, 2021, from https://mw17.mwconf.org/paper/removing-the-barriers-of-gallery-one-a-new-approach-to-integrating-art-interpretation-and-technology/

Andersen, K., & Ward, N. (2017). Learning from the Crackle Exhibition. *Proceedings of the Eleventh International Conference on Tangible, Embedded, and Embodied Interaction*, 225–231. https://doi.org/10.1145/3024969.3024997

Bardzell, J. (2011). Interaction criticism: An introduction to the practice. *Interacting with Computers*, *23*(6), 604–621. https://doi.org/10.1016/j.intcom.2011.07.001

Bedford, L. (2014). *The Art of Museum Exhibitions: How Story and Imagination Create Aesthetic Experiences*. Left Coast Press.

Benford, S., Løvlie, A. S., Ryding, K., Rajkowska, P., Bodiaj, E., Paris Darzentas, D., Cameron, H., Spence, J., Egede, J., & Spanjevic, B. (2022). Sensitive Pictures: Emotional Interpretation in the Museum. *CHI Conference on Human Factors in Computing Systems*, 1–16. https://doi.org/10.1145/3491102.3502080

Bertelsen, O. W., & Pold, S. (2004). Criticism as an approach to interface aesthetics. *Proceedings of the third Nordic conference on Human-computer interaction*, 23–32. https://doi.org/10.1145/1028014.1028018

Bimber, O., Coriand, F., Kleppe, A., Bruns, E., Zollmann, S., & Langlotz, T. (2005). Superimposing pictorial artwork with projected imagery. *ACM SIGGRAPH 2005 Courses*, 6–es. https://doi.org/10.1145/1198555.1198716

Boehner, K., Sengers, P., Medynskiy, Y., & Gay, G. (2005). Opening the Frame of the Art Museum: Technology between Art and Tool. *Digital Arts and Culture (DAC)*, 123–132.

Boehner, K., Thom-Santelli, J., Zoss, A., Gay, G., Hall, J. S., & Barrett, T. (2005). Imprints of place: Creative expressions of the museum experience. *CHI '05 Extended Abstracts on Human Factors in Computing Systems*, 1220–1223. https://doi.org/10.1145/1056808.1056881

Bolander, E., Ridenour, H., & Quimby, C. (2018). *Art Museums and Technology: Developing New Metrics to Measure Visitor Engagement* (tech. rep.). The Cleveland Museum of Art.

Claisse, C., Petrelli, D., Ciolfi, L., Dulake, N., Marshall, M. T., & Durrant, A. C. (2020). Crafting Critical Heritage Discourses into Interactive Exhibition Design. *Proceedings of the 2020 CHI Conference on Human Factors in Computing Systems*, 1–13. https://doi.org/10.1145/3313831.3376689

Coenen, T., Mostmans, L., & Naessens, K. (2013). MuseUs: Case study of a pervasive cultural heritage serious game. *Journal on Computing and Cultural Heritage*, *6*(2), 1–19. https://doi.org/10.1145/2460376.2460379

Cosley, D., Baxter, J., Lee, S., Alson, B., Nomura, S., Adams, P., Sarabu, C., & Gay, G. (2009). A tag in the hand: Supporting semantic, social, and spatial navigation in museums. *Proceedings of the SIGCHI Conference on Human Factors in Computing Systems*, 1953–1962. https://doi.org/10.1145/1518701.1518999

Cosley, D., Lewenstein, J., Herman, A., Holloway, J., Baxter, J., Nomura, S., Boehner, K., & Gay, G. (2008). ArtLinks: Fostering social awareness and reflection in museums. *Proceedings of the SIGCHI Conference on Human Factors in Computing Systems*, 403–412. https://doi.org/10.1145/1357054.1357121

Costello, B., Muller, L., Amitani, S., & Edmonds, E. (2005). Understanding the Experience of Interactive Art: Iamascope in Beta_space. *The Second Australasian Conference on Interactive Entertainment*, *123*, 49–56.

Dindler, C., Krogh, P. G., Beck, S., Stenfelt, L., Nielsen, K. R., & Grønbæk, K. (2007). Peephole experiences: Field experiments with mixed reality hydroscopes in a marine center. *Proceedings of the 2007 conference on Designing for User eXperiences*, 2–15. https://doi.org/10.1145/1389908.1389934

Drotner, K., Dziekan, V., Parry, R., & Schrøder, K. C. (2018). *The Routledge Handbook of Museums, Media and Communication* (K. Drotner, V. Dziekan, R. Parry, & K. C. Schrøder, Eds.; 1st ed.). Routledge. https://doi.org/10.4324/9781315560168

Endresen, S. (2015). *Serial Experiments. Close-Readings of Edvard Munch's Det grønne værelset (1907)* (Doctoral dissertation). Universitetet i Oslo. Oslo.

Gallagher, S. (2017). *Enactivist Interventions: Rethinking the Mind*. Oxford University Press.

Gibson, J. J. (1979). *The Ecological Approach to Visual Perception*. Houghton Mifflin.

Hindmarsh, J., Heath, C., vom Lehn, D., & Cleverly, J. (2002). Creating Assemblies: Aboard the Ghost Ship. *Proceedings of the 2002 ACM conference on Computer supported cooperative work - CSCW '02*, 156. https://doi.org/10.1145/587078.587101

Hornecker, E., & Ciolfi, L. (2019). Human-Computer Interactions in Museums. *Synthesis Lectures on Human-Centered Informatics*, *12*(2), i–153. https://doi.org/10.2200/S00901ED1V01Y201902HCI042

Hsi, S. (2003). A study of user experiences mediated by nomadic web content in a museum. *Journal of Computer Assisted Learning*, *19*(3), 308–319. https://doi.org/https://doi.org/10.1046/j.0266-4909.2003.jca_023.x

Ingold, T. (2013). *Making: Anthropology, Archaeology, Art and Architecture*. Routledge.

Ingold, T. (2018). *Anthropology and/as education*. Routledge.

Kortbek, K. J., & Grønbæk, K. (2008). Communicating Art through Interactive Technology: New Approaches for Interaction Design in Art Museums. *NordiCHI '08: Proceedings of the 5th Nordic conference on Human-computer interaction: building bridges*, 229–238. https://doi.org/10.1145/1463160.1463185

Lange, V., van Beuzekom, M., Hansma, M., Jeurens, J., van den Oever, W., Regterschot, M., Treffers, J., van Turnhout, K., Ibrahim Sezen, T., Iurgel, I., & Bakker, R. (2019). Blending into the White Box of the Art Museum. *Proceedings of the Halfway to the Future Symposium 2019*, 1–10. https://doi.org/10.1145/3363384.3363469

Louvre. (2020). Mona Lisa: Beyond the Glass.

Lyons, L. (2009). Designing opportunistic user interfaces to support a collaborative museum exhibit. *Proceedings of the 9th international conference on Computer supported collaborative learning - Volume 1*, 375–384.

Masschelein, J. (2010a). E-ducating the gaze: The idea of a poor pedagogy [Publisher: Routledge _eprint: https://doi.org/10.1080/17449641003590621]. *Ethics and Education*, *5*(1), 43–53. https://doi.org/10.1080/17449641003590621

Masschelein, J. (2010b). The idea of critical e-ducational research – e-ducating the gaze and inviting to go walking. In I. Gur-Ze'ev (Ed.), *The Possibility/Impossibility of a New Critical Language of Education* (pp. 271–291). Sense Publishers.

Maye, L. A., McDermott, F. E., Ciolfi, L., & Avram, G. (2014). Interactive exhibitions design: What can we learn from cultural heritage professionals? *Proceedings of the 8th Nordic Conference on Human-Computer Interaction: Fun, Fast, Foundational*, 598–607. https://doi.org/10.1145/2639189.2639259

Noë, A. (2015). *Strange Tools: Art and Human Nature*. Hill; Wang.

Noë, A. (2021). *Learning to Look: Dispatches from the Art World*. Oxford University Press.

Noë, A., Bertram, G. W., Davies, D., Leddy, T., & Viola, T. (2020). On Alva Noë, Strange tools. Art and human nature [Number: 18]. *Studi di estetica*, *3*(18), 249–295. Retrieved February 3, 2022, from https://journals.mimesisedizioni.it/index.php/studi-di-estetica/article/view/898

Pau, S. (2017). Audio that Moves You: Experiments with Location-aware storytelling in the SFMOMA app –. Retrieved February 1, 2023, from http://mw17.mwconf.org/proposal/audio-that-moves-you-location-aware-storytelling-in-the-sfmoma-app/index.html

Petrelli, D., Ciolfi, L., van Dijk, D., Hornecker, E., Not, E., & Schmidt, A. (2013). Integrating material and digital: A new way for cultural heritage. *Interactions*, *20*(4), 58–63. https://doi.org/10.1145/2486227.2486239

Rizzo, F., & Garzotto, F. (2007). "The Fire and The Mountain": Tangible and social interaction in a museum exhibition for children. *Proceedings of the 6th international conference on Interaction design and children*, 105–108. https://doi.org/10.1145/1297277.1297298

Ryding, K. (2020). The Silent Conversation: Designing for Introspection and Social Play in Art Museums. *Proceedings of the 2020 CHI Conference on Human Factors in Computing Systems*, 1–10. https://doi.org/10.1145/3313831.3376357

Ryding, K., Spence, J., Løvlie, A. S., & Benford, S. (2021). Interpersonalizing Intimate Museum Experiences [Publisher: Taylor & Francis]. *International Journal of Human–Computer Interaction*, *37*(12), 1151–1172. https://doi.org/10.1080/10447318.2020.1870829

Samis, P. (2018). Revisiting the utopian promise of interpretive media: An autoethnographic analysis drawn from art museums, 1991–2017. In K. Drotner, V. Dziekan, R. Parry, & K. C. Schrøder (Eds.), *The Routledge Handbook of Museums, Media and Communication* (1st ed., pp. 47–66). Routledge. https://doi.org/10.4324/9781315560168

Spence, J., Bedwell, B., Coleman, M., Benford, S., Koleva, B. N., Adams, M., Row Farr, J., Tandavanitj, N., & Løvlie, A. S. (2019). Seeing with New Eyes: Designing for In-the-Wild Museum Gifting. *Proceedings of the 2019 CHI Conference on Human Factors in Computing Systems*, 1–13. https://doi.org/10.1145/3290605.3300235

Tennent, P., Martindale, S., Benford, S., Darzentas, D., Brundell, P., & Collishaw, M. (2020). Thresholds: Embedding Virtual Reality in the Museum. *Journal on Computing and Cultural Heritage*, *13*(2), 1–35. https://doi.org/10.1145/3369394

Terrenghi, L., & Zimmermann, A. (2004). Tailored audio augmented environments for museums. *Proceedings of the 9th international conference on Intelligent user interface - IUI '04*, 334. https://doi.org/10.1145/964442.964523

The National Gallery. (2019). Leonardo: Experience a Masterpiece.

Thompson, E. (2007). *Mind in Life*. Harvard University Press.

Van Gogh Museum Amsterdam. (2016). Meet Vincent van Gogh.

Varela, F. J., Rosch, E., & Thompson, E. (1992). *The Embodied Mind: Cognitive Science and Human Experience*. MIT Press.

Walter, T. (1996). From museum to morgue? Electronic guides in Roman Bath. *Tourism Management*, *17*(4), 241–245. https://doi.org/10.1016/0261-5177(96)00015-5

Wessel, D., & Mayr, E. (2007). Potentials and Challenges of Mobile Media in Museums [Number: 1]. *International Journal of Interactive Mobile Technologies (iJIM)*, *1*(1). Retrieved May 13, 2021, from https://online-journals.org/index.php/i-jim/article/view/165

Woodruff, A., Aoki, P. M., Hurst, A., & Szymanski, M. H. (2001). Electronic Guidebooks and Visitor Attention. *Proceedings of the International Conference on Cultural Heritage and Technologies in the Third Millennium*, 437–454.

Woods, E., Billinghurst, M., Looser, J., Aldridge, G., Brown, D., Garrie, B., & Nelles, C. (2004). Augmenting the Science Centre and Museum Experience. *Proceedings of the 2nd international conference on Computer graphics and interactive techniques in Australasia and South East Asia*, 230–236. https://doi.org/10.1145/988834.988873

Yamazaki, K., Yamazaki, A., Okada, M., Kuno, Y., Kobayashi, Y., Hoshi, Y., Pitsch, K., Luff, P., vom Lehn, D., & Heath, C. (2009). Revealing Gauguin: Engaging visitors in robot guide's explanation in an art museum. *Proceedings of the SIGCHI Conference on Human Factors in Computing Systems*, 1437–1446. https://doi.org/10.1145/1518701.1518919

Zwan, S. v. d., Smith, M. L., Bruineberg, J., Levy, P. D., & Hummels, C. C. M. (2020). Philosophy at Work: Postphenomenology as a Generative Lens in Design Research and Practice. *Proceedings of DRS2020 International Conference : Vol. 4: Education*, 1691–1705. https://doi.org/10.21606/drs.2020.337

Publication C

Educating the Attention of Museum Visitors through Non-verbal Art Mediation

Christian Sivertsen, Nikita Mathias and Anders Sundnes Løvlie

This paper has been published in
De Sainz, D., Galluzzo, L., Rizzo, F., Spallazzo, D. (eds.), *IASDR 2023: Life-Changing Design, 9-13 October, Milan, Italy*. https://doi.org/10.21606/iasdr.2023.222

The layout has been revised.

Abstract

The use of technology in art museums is usually limited to communicating verbal information such as interpretive text and audio guides, or to facilitating social experiences. This paper presents a Research through Design study of an immersive room-scale installation in an art museum. The aim was to afford non-verbal art mediation and educate the attention of museum visitors to certain aspects of a series of paintings. The installation was based on recent art historical research and aimed to use this research to facilitate a new perspective on the artworks. The installation was created through a one-year-long iterative process that involved diverse stakeholder groups of experts, visitors, and non-visitors in testing and formative evaluation. The final design was part of the opening of the new MUNCH museum in October 2021 and was visited by almost 80,000 people in four months. Through interviews with 47 visitors, we evaluate how their experience corresponds with the design intent, and how the installation through non-verbal means drew their attention to aspects of the related artwork. We find that the installation largely evokes the intended aesthetic qualities and support reflection on Edvard Munch's life and artistic practice. Finally, we discuss the use of education of attention as a lens for design practitioners to conceptualize and evaluate how their designs mediate learning and interpretive practices for art museum visitors.

C.1 Introduction

The museum field has undergone an interpretive turn, bringing outside perspectives into the museum and inviting visitors to make up their own minds about our cultural heritage (Macdonald, 2007). For many museums this has evolved further into an immersive turn, inviting visitors to experience history, science, and culture through multisensory environments and installations (Kidd, 2018). In the *art* museum domain, working with these technologies to support interpretive practices of specific art objects poses conceptual challenges for designers and technologists as they are often seen as conflicting with the ritualistic space of the art museum and the purity of the art (Duncan, 2005; Hornecker & Ciolfi, 2019; Kortbek & Grønbæk, 2008; Løvlie et al., 2021; Walter, 1996). Most commonly the mediation of artworks is done through verbal means such as text labels and audio guides. However, in recent years some projects have experimented with sensorial and bodily interventions as a way of supporting the sensemaking of artworks (Alexander et al., 2017; Steier, 2014; The National Gallery, 2019; Vi et al., 2017), indicating that involvement of the body in art mediation can be a powerful way of helping visitors make sense of art. In the broader museum field, some scholars have even argued for the reintroduction of touch and handling of cultural heritage objects (Howes, 2014). Since the original art objects themselves are often too fragile or precious to allow for direct touch and manipulation, these interventions use multisensory or immersive means to propose context

and perspectives that support the museum visitor in making sense of the artworks on display. The role they play in supporting visitors' interpretation of artworks is thus not entirely different in their purpose from the traditional text label or audio guide. We believe that such non-verbal tools have great potential in creating art mediation that is an interesting experience in itself while expanding the toolbox of the art museum to support engagements with art that go beyond what can be elicited by verbal communication. This might in turn make the art museum available to a wider range of visitor types. Importantly though, an operationalizable framing of art, experience, and learning is necessary to support designers in developing and evaluating concepts that fulfill this role (Sivertsen et al., 2023).

Through an understanding of *art as experience* (Dewey, 1934/1980) and as a *reorganizational practice* (Noë, 2016) we argue that the role of technologies supporting art mediation is to *educate the attention* (Ingold, 2001) of museum visitors to aspects of the art they might otherwise not see.

To put this theoretical perspective into practice we will present a *Research-through-Design* (Zimmerman et al., 2007) study of the project *Poison* . It is a small immersive exhibition that addresses a specific series of Edvard Munch's paintings known as *The Green Room* from 1907. The exhibition uses space, light, and interactive sounds and visuals to create an atmosphere in correspondence with art historical research about *The Green Room*. Poison was opened along with the new MUNCH museum in Oslo and was visited by almost 80.000 people from October 2021 to February 2022. In this paper, we will show how *Poison* was designed in correspondence with art historical research, and how it directed the attention of visitors to specific aspects of Edvard Munch's life and artistic practice.

C.2 Embodiment and Immersion as a Mediation Strategy

The use of canonical artists such as Van Gogh, Frida Kahlo, and Monet in shows involving room-sized projections of their artworks, augmented with animations and soundscapes has become a great commercial success, mostly outside of institutional museum contexts (Mathias, 2022; Mondloch, 2022). However, these shows have also been criticized for being superficial and inferior to seeing the originals (Mondloch, 2022). The popularity and success in attracting audiences outside the normal art museum crowd gives reason to investigate the potential for using immersion as a means for art mediation inside established art institutions.

The use of immersion as a mediation tool in other types of museums is already well-established. Bedford traces the development of the understanding of education in the museum field from an "object-based epistemology" (Bedford, 2014, p. 23) where knowledge was inherent in objects, and it was assumed that visitors could read

that knowledge from the objects. Throughout the 1900s knowledge existed outside of the visitor and was transmitted through programs, labels, and educators. Only in the late 1900s, the visitor experience became important and along with that an "interpretive turn" where museums facilitate visitors' own interpretation of objects rather than merely "transferring" pre-established facts. More recently Kidd points to an "immersive turn"(Kidd, 2018), manifesting as multi-sensory spaces presenting cultural heritage from a variety of perspectives, many of which use technology as a mediating element (Budge, 2018; Harrington, 2020; Kenderdine, 2015; Snibbe & Raffle, 2009; Stogner, 2011; Tennent et al., 2020; Waern & Løvlie, 2022). These projects show the potential of multisensory and immersive exhibitions to support engagement with cultural heritage.

To some extent, this already starts with the exhibition design acting as scenography for the museum to create a holistic aesthetic experience supporting the content and message of the exhibition (Hornecker & Ciolfi, 2019). While the art museum's history as a ritualistic space (Duncan, 2005) causes many art museums to rely on the classic white cube aesthetic, this is nevertheless imbued with meaning, setting a particular stage for the art experience.

Within this frame, the mediation tools mostly used by art museums to contextualize and orient the visitors towards the art on display consist of interpretive text – usually in the form of written labels and/or audio guides. These tools rely on language as the primary means of supporting sense-making. Some projects allow visitors to add digital tags and comments (Boehner et al., 2005; Cosley et al., 2008, 2009) to artworks. *Sensitive Pictures* used audio dramatizations to stimulate affective engagements with particular artworks (Benford et al., 2022).

Some efforts have been made to explore art mediation through other means than verbal information. Kortbek & Grønbæk (2008) designed installations that relied on bodily interaction to reveal inspirational images and text content related to the main exhibition. In this case, movements were used as a key to unlock the inspirational content, more than supporting the sense-making directly.

Oppositely *Strike a Pose* from the *ARTLens Gallery* (Alexander et al., 2017) and work by Steier (2014) investigate the literal imitation of a pose in original artwork to support interpretive practices. Steier argues that meaning in the artwork arises in the interaction between gestures and talking between visitors as they carry out the posing activity. In a particular instance that Steier observed, he notes that the attention of the visitors is drawn to the hand, the height, and the facial expression of the painted figure, indicating that the posing activity leads visitors to consider related qualities in the painted image (2014). Alexander et al. also claim that having fun with the activities in the *ARTLens Gallery* "will enhance the appreciation of art throughout the museum" (2017). This continues the idea that bodily activities can support interpretational activity with art.

The project *TATE Sensorium* explores stimulating visitors through sounds, haptic, taste, and smell as they contemplate four original paintings (Vi et al., 2017). With regards to haptic stimulation, the authors argue that it has the potential to "augment" the experience of the artwork, but do not report conclusively on the ways different stimuli directed the experience of the artwork.

With a more technical view of art, visuals, and projections have been used to reveal hidden layers and techniques in paintings. Bimber et al. (2005) propose a technique for museums to create edutainment by superimposing graphics directly on original artworks, to show painting techniques and hidden sketches. While using projections directly on original artwork might cause concern for the stability of certain pigments in old artworks, similar ideas of using technology to uncover hidden layers in paintings have been implemented in a Virtual Reality installation about Mona Lisa at The Louvre (Louvre, 2020) and in an immersive exhibition called *Leonardo: Experience a Masterpiece* at The National Gallery (The National Gallery, 2019). Besides a projection showing different layers and the painting process of *The Virgin of the Rocks*, this exhibition also includes interactive installations allowing visitors to explore ways of using light and shadow in painting, an installation about the church for which the painting was made as well as a digital replica that presents a hypothetical design for the altarpiece where the painting was intended to hang. In a presentation, the curator of the exhibition says that the intention is "to make people look better and for longer [at great art]" through an exhibition that is "underpinned by academic research" and that is "visual rather than word-led" (Campbell, 2019). The painting technique including Da Vinci's use of light, as well as the history of the commissioning and the original physical context of the artwork, has been decided by the curatorial team to be of special interest to this artwork. With these specific interests in mind, they have developed an exhibition that through mainly visual means attempts to help visitors see the same in the art as they do.

Digital reproductions of works of fine art have sparked scholarly discussions anchored in Walter Benjamin's famous notion of the loss of the aura attributed to reproduced artworks (Benjamin, 1936), however more recent research points to the potential of the auratic quality migrating to or emerging in reproductions (Jeffrey, 2015; Jones et al., 2018; Latour & Lowe, 2011). In previous work, Sivertsen and colleagues (2023) have argued that the digital reproductions may be part of an "art critique". This concept of art mediation based on digital reproductions does not seek to re-establish or emulate the auratic power of the original, but rather aim at *corresponding with* the originals by educating the visitor's attention towards central aspects of the original artworks. (Sivertsen et al., 2023).

In summary, we see that immersive and sensorial experience is both attractive to visitors and museums and may have a significant impact on visitors' engagement with original artwork. To further develop an understanding of this relation we will

introduce a conception of art and learning that will allow us to further investigate how an immersive environment can mediate visitor attention to particular aspects of original artworks.

C.3 An Enactivist Perspective on Art Mediation

Building on Dewey's (1934/1980) conception of *art as experience*, Noë (2016) argues that art depends on us getting it. When we do not, it has the potential to bore us to death. This is what art mediation attempts to alleviate. Art is a form of *reorganizational practice* that challenges human practices as we know them. However, this also means that artworks can cease to do the work of art if we do not perceive this reorganization (Noë, 2016). In this perspective, we understand the role of art mediation to be pointing out things for the audience that they might otherwise not have noticed in the art (Sivertsen et al., 2023) . Different media is suitable for drawing attention to different things. Noë argues that in the manual of a car engine line drawings are superior to photographs because "The problem with the photographs was that they didn't pick out what was important. They just gave you an image of the undifferentiated stuff. The drawings, in contrast, were truly articulate; they drew your attention to what was salient even as they drew the parts themselves." (Noë, 2021, p. 65). This perspective on learning is similar to Ingold's description of *education of attention* that he derives from Gibson's ecological psychology (Ingold, 2001), though Ingold is more concerned with skilled practitioners than art. Nevertheless, for this paper, we will be using the term *education of attention* to investigate how the immersive installation tunes museum visitor's attention to salient aspects of the art.

C.4 Research and Design Process

Poison was conceived as a *Research-through-Design* (Zimmerman et al., 2007) project set up to explore how an immersive and primarily non-verbal design could educate museum visitors' attention to aspects of the painter Edvard Munch's series of paintings known as *The Green Room*. The project was a collaboration between a design researcher (the first author) and a museum team which initially consisted of a concept developer and an intern and grew throughout the process to eventually involve around 30 museum experts, designers, developers, and technicians. The project was scheduled to launch with the opening of the new museum in October 2021. Embedding the project in the museum had great potential, as it allowed for the evaluation of experiential qualities directly in the museum context, with participants engaging with Poison as an exhibition and not as an experiment.

For the development of the project, we established an iterative process, where we gathered feedback from 6 different groups through 5 major iterations. The feedback sessions were intended to enable and help manage *drifting* in the words of Krogh & Koskinen (2020), that is readjusting the course of the project, both conceptually and practically as our knowledge of the domain, practical considerations, and the potential for different experiences became clearer.

C.4.1 Design Assumptions

The design was constrained by a set of choices that were established in the first weeks of the project surrounding the mediation approach, the correspondence with art historical research, and practical considerations.

To investigate the potential of non-verbal mediation, we limited the number of text elements in the experience. The size of the space did not allow for hanging the original paintings. We also decided not to show accurate digital replicas as this could feel like a cheap version of the original and direct visitors' attention in an unfavorable way. Rather, like Noë's example of the car manual, the digital replications of the originals should be tuning the visitor to the salient qualities of the series of paintings. We aimed to use the installation to reenact aesthetic qualities identified in past art historical research on *The Green Room* series (Endresen, 2015). Endresen's work points to the unsettling atmosphere in the paintings, and the architectural instability. The paintings all take place in a similar green room, but the dimensions and wallpaper change slightly between paintings while sticking to a greenish color scheme that in Munch's work often is connected to jealousy and sickness. The titles *Hatred*, *Jealousy*[1] , *The Murderess*, *Desire*, and *To the Sweet Girl*[2] indicates dark emotional themes. Endresen further highlights the point-of-view perspective of the paintings, which places the viewer at the table in the back of the room while the rest stretches out as a theatre stage where dramatic scenes unfold. The characters in the paintings are understood to be avatars representing human characteristics that the viewer might read themselves into rather than being depictions of historical persons.

C.4.2 Design Process

The project was assigned a 40 m2 room in the museum. Through an iterative ideation and sketching process, we developed an initial idea of a space consisting of three chambers. The first chamber should help the visitors become aware that the exhibition relies on a different mode of being than the other galleries and help them tune their attention to the interactive elements and the atmosphere. The second chamber would immerse the visitors in a version of The Green Room. Finally, the last chamber would conclude the experience by giving some historical context. To develop this

[1]Two of the paintings are named "Jealousy".
[2]Original:"Zum Süssel Mädel", is believed to be a reference to prostitution

idea further we decided to set up two concurrent tracks of evaluation: an expert track and a non-expert track including both visitors and non-visitors to the museum. The project was evaluated at multiple stages before and after its completion. In total two expert groups and three groups of non-experts were invited one or more times, see C.1.

Group Description	Experts: Designers and researchers in the field of immersive experiences	Experts: Museum employees from curation, marketing, digital, and education	Master students studying aesthetics at a local university	Members of the MUNCH museum club aged 18-35	Young women with non-Norwegian heritage, who were not frequent museum visitors
Members per session	4	4-5	5	10-15	4-5
Number of sessions	1	3	1	3	3

Table C.1: Groups of participants involved in formative evaluations

We divided the project into several stages. The first stage was the *concept stage*. Due to COVID-19, we conducted our first evaluation online. We showed a video taking the participants through a 3D model of the space. We asked the participants about their first impressions and their understanding of the spaces shown in the video. Then we asked about their interest in the concepts and whether they thought it was relevant for the museum.

We continued to conduct such evaluations with the groups, however moving into the museum space. Due to the museum not being open to the public, we were able to construct prototypes in the space that would house the final concept and invite people to visit it.

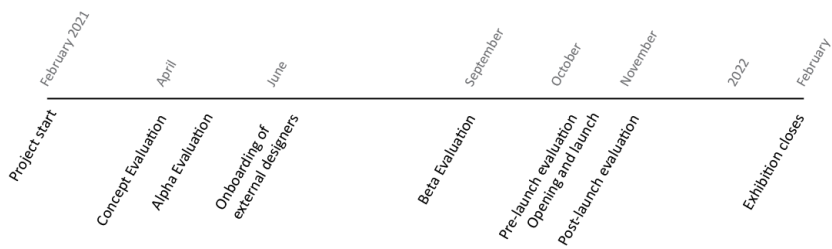

Fig. C.1: Timeline of the project

We tested a rudimentary technical setup in the *alpha stage*, using projections, simple wall dividers, and a "found" soundtrack. In the *beta stage* additional design expertise was brought in, and some concepts were adjusted based on the earlier evaluations. At this point, we had learned that the overall concept was well received by our test groups, while we kept improving the content of the exhibition to best reflect the aesthetic qualities we wanted to highlight.

C.5 Final Design

Poison was installed in a small exhibition space at the new MUNCH museum which was outfitted with an internal structure covering the floor, walls, and ceiling. This structure divides the room into 3 small chambers and a light lock.

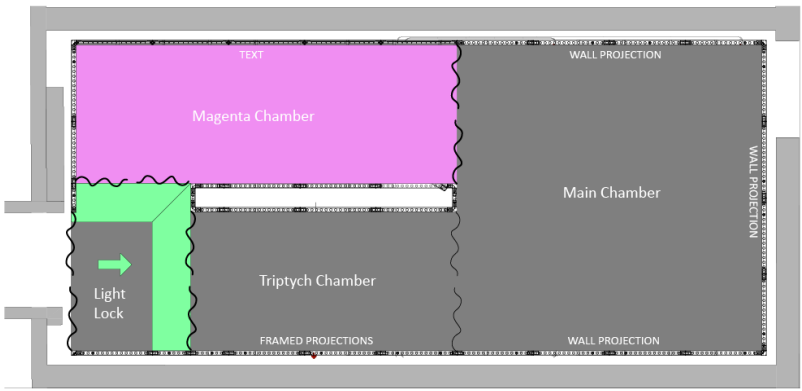

Fig. C.2: Floor plan showing the layout of *Poison*. The entrance and exit are to the left.

The visitors first step into the dark light lock, where the soundscape sounding through the whole space is first heard. There is no explanation or introduction. The only hint is the title "Poison" and the banner by the first curtain reading "The Green Room is alive". As visitors step through the heavy velvet curtain, they arrive in the triptych room. The floor is creakier, and the size and atmosphere differ significantly from the other galleries. Three images from *The Green Room* series are projected from the ceiling and mapped to fit exactly inside three wooden frames. Microphones in the floor pick up the visitors' footsteps and a camera in the ceiling tracks their movements. Each step triggers a screeching sound from a speaker hidden near the floor, and the movements back and forth in the room make the images in the frames morph from one to another essentially switching places. This is intended to alert the visitors that "The Green Room sees you" as it says on the banner on the wall. From this chamber, the visitors can look through a transparent curtain into the main chamber.

Fig. C.3: The triptych chamber. Three paintings (one is outside of view) are projected inside frames on the wall. The images change as visitors move past them.

In the main chamber, the visitors become surrounded by three large projections each covering a full wall. Each projection depicts a painting from *The Green Room*. However, the perspective of each painting shifts as the visitors move around inside the chamber. The wallpaper moves and the images are always changing slightly, always slightly blurry, never standing completely still. In the main chamber, a deep oscillating bass is triggered by each footstep but fades away when people are standing still. In the main room, a banner reads “The Green Room is poisonous”.

Stepping into the last chamber the visitors are met with a bright magenta light. Here, lighter elements in the soundscape are more audible. On one large backlit wall, a short text offers some grounding by letting the viewer know that the exhibition is tied to a series of paintings Munch made in 1907, that it *is* a strange and ambiguous series of paintings, and finally that in 1907 green wallpapers might still contain the poisonous compound arsenic. On the way out the visitors pass by a banner reading “The Green Room is everywhere”. The magenta light in this room facilitated an optical effect as the visitor’s eyes would adjust to the light while reading the text on the backlit wall, overcompensating with green as the complementary color of magenta. Thus, when leaving the room everything outside would appear with a green tint for a few seconds until the visitors’ eyes readjusted to the daylight outside.

Fig. C.4: The main chamber. The visitor is surrounded by three large projections, depicting blurry, shifting versions of paintings of *The Green Room*.

C.6 Summative Evaluation

The analysis presented in this paper is based on interviews with 47 visitors. 16 of these were members of the museum club invited to experience *Poison* a few days before the official launch. This group consisted of participants aged 23 to 70 with a median of 32. 10 identifying as female and 6 as male. All but two saw themselves as "interested in art". Three had participated in earlier evaluations of *Poison*. On average participants had visited art galleries or museums 11,5 times in the last 12 months (min 0, max 40).

Another 31 interviewees were regular museum visitors who experienced the installation in the days after the public launch. These participants were recruited as they left *Poison*. The recruiting researcher would attempt to recruit as diversely as possible, with the limitation of not having the possibility to screen participants. People were interviewed alone, in pairs, or in groups of 3. 20 identified as female while 11 identified as male. The age of the group ranged from 19 to 80, with a median of 34. The interviewees were mostly visiting with friends, family, or colleagues while a few were in the museum alone. 28 out of 31 saw themselves as being "interested in art". The average art gallery or museum visit for this group over the last 12 months was 6,2 (min. 0, max. 44). For both groups this number should be considered in the light of COVID-19 lockdowns in Norway throughout 2021 and late 2020.

Fig. C.5: The visitors pass through a curtain from the main chamber to the magenta chamber.

The main part of the interview followed a phenomenological interviewing method (Thompson et al., 1989). However, first structured demographic questions were asked, then the participants were asked to describe their experience of going through *Poison*. Finally, the interview was concluded with semi-structured questions, asking whether the participants found *Poison* to be relevant for the museum and if they saw any connection to Edvard Munch's artistic practice in *Poison*.

In addition to the interviews, the first author spent 12 hours observing visitors while acting as a "host" for the exhibition, standing inside or outside of the room, helping guests, and answering questions.

C.7 Analysis

To understand whether the final design of *Poison* lived up to the ambitions of the design, and evoke the concepts intended we will be analyzing the results from the interviews following three questions.

1. Was *Poison* accepted in the specific art museum context?
2. Did the aesthetic quality of *Poison* reported by the visitors correspond with the design intent?

Fig. C.6: In the magenta chamber the visitor reads the short interpretive text against on a backlit magenta wall. Photo credit: Dina Patey

3. How do visitors place the experience in relation to their existing knowledge of Edvard Munch's life and artistic practice?

C.7.1 Visitors' Acceptance of Poison

We found a strong consensus among all the interviewed participants that *Poison* is an exciting, timely, and relevant addition to the museum. When participants were asked whether they thought *Poison* was a relevant addition to the museum, most answered that they "absolutely" or "definitely" thought so. When arguing for their answer the participants' reasoning followed three different arguments. First, some participants stated that *Poison* gave them something interesting or special in addition to the rest of the museum: "It's very modern and contemporary also. And I think it's nice to have this next to like all the just paintings. It's a different experience. And I think that, well, adds value to the rest of it" (Female, 24, post-launch). A second group of responses suggested that the brand-new museum should follow technological trends. One participant took this argument further arguing that keeping up with the development of technology is also in the spirit of Munch himself:

> "I think it follows the times. [...] His paintings developed with the technology, and when the photography was introduced, his painting became more alive, then I feel that this probably fits very well to where we are now and how the world has become more technical that his paintings are presented in a more technical way" (Female, 19, post-launch).

	Interviewees	Female	Male	Median age	"I am interested in art"	Art museum/gallery visits in last 12 months, average
Pre-launch	16	10	6	32, min. 23, max. 70	14	11,5, min. 0, max. 40
Post-launch	31	20	11	34, min. 19, max. 80	28	6,2, min. 0, max. 44

Table C.2: Overview of participants in the summative evaluations of *Poison*

The third argument is that participants assume that it will be relevant for a new/younger/other demographic in the museum whether they see themselves as being part of this demographic or not: "Of course, for children and alike it's a lot of fun [...] Even though I personally am a bit old-fashioned and actually just like to see physical things that I can hold" (Male, 41, post-launch). Several participants stressed that while *Poison* is a meaningful *addition* to the museum experience it should not be a *substitute* for putting the originals on display.

C.7.2 The Aesthetic Qualities of Poison

With a general acceptance of Poison amongst the visitors, we now move on to evaluating if the visitors' experience was in correspondence with our design intent. In other words, did we manage to draw their attention to the qualities we wanted to show?

Atmosphere

Starting with the overall descriptions of the atmosphere, visitors described it as "dark", "intense" and "unpleasant". Some mentioned that the space felt "tight" and "crammed". The visitors described the atmosphere with terms relating to negative psychological and physical states like "nauseating", "paranoid", "sickness", "oppressive", "uneasy", "unsettling" and even "scary" and "nightmarish". However, despite the association with negative feelings the visitors also generally accepted this and engaged with it finding it "curious", "unknown", "exciting" and "psychedelic" or like "tripping". Some said it was like entering "another universe", isolated from the rest of the world. The participants pointed to the soundscape, the darkness, temperature, visuals, and spatial qualities as the source for these atmospheres. Overall, we find that the atmospheric qualities expressed by the participants fall close to those intended by the design team.

Presence

Several participants indicated that the main chamber felt connected with the three projected rooms into one coherent space. One participant explained that being in the main room was like standing in the middle of a church with transepts stretching out on either side (Male, 26, pre-launch). Other participants also described this as looking into an extension of the room, "I mean, the pictures moved slightly, so you get like a very dimensional feeling like actually standing in the room" (Female, 24, post-launch). This feeling of being present in the space is reflected by several other participants, "I had a feeling like I was present in the image" (Female, 64, post-launch), "Even though you're not like touching, but you're there" (Female, 32, post-launch), "I was in a green room when I didn't see the ceiling" (Female, 24, pre-launch). Investigating this presence further, however, also reveals a tension in the physical relation to the virtual rooms. One participant reported feeling like she was looking into the room from the outside. Several others similarly said that they were "engaged" or "captivated" by trying to see and understand what was going on in the images as they oscillated between being blurry and clearer. There were also a few visitors who are not able to recall the visuals in any detail or say that they did not take or have the time to engage with the visuals in the main room, and quickly moved on.

Co-performance

More than presence the design was intended to provide a feeling of participation and being in the spotlight. However, most participants did not understand the interactive elements. A few participants refer to the visuals as the "video sequence", and in one interview the participants express uncertainty that they managed to see "it all". On the other hand, we noted during observations that a small number of people would become very aware of the interactive elements. In one case it happened that someone noticed the sounds being generated by their footsteps in the triptych room, and you could hear them stomping the floor intensely to create louder and louder sounds. In another case a person was observed walking back and forth in front of the visuals waving their arms, we assume, to generate a response in the visuals.

The design of *Poison* did not entirely foster the sense of co-performance that we hoped. A more direct coupling (Wensveen et al., 2004) between visitors' movements and the reaction of the visual might have enhanced the feeling of having agency over soundscape and visuals and being able to participate in the construction of the space. Nevertheless, we see that *Poison* affords a sense of presence in *The Green Room* for most visitors.

C.7.3 Relation to Munch's Art and Practice

We will now investigate how visitors place their experience of *Poison* in relation to what they already know about Edvard Munch's art and practice and what they report to have taken away from the experience.

One way in which participants relate *Poison* to Munch is by placing the work in relation to what they already know about his life. We see this being triggered by the visuals, "I guess it was that one picture that in a way looked like it had a bottle of alcohol in it. So that made me associate it to the poisonous sides of Munch, him being a bit boheme" (31, Male, pre-launch). For other visitors, reflecting was prompted by the facts presented in the interpretive text: "I thought [the text] confirmed my experience from before, that Munch was very good at making situations and tableaus that look real. It's everyday images in a way." (57, Female, post-launch). These two examples suggest that recognizable elements would reaffirm visitors' existing ideas about Munch. Another visitor indicates that she learns something new from the text, "I think it was interesting to learn that there would be poison in the painting of green rooms. [...] then I started thinking about if it might be fantasy or if he actually went to a brothel and about what the images meant for Munch. That was cool. So, I think it is very nice to get some information about what you've seen." (30, Female, post-launch). In these three examples, the visitors point to the symbolic bottle or the text as the source of their reflection. The first and third examples relate to biographical details and the second to Munch's capacities as a painter.

Participants also derive similar reflections from the aesthetic elements of *Poison*. First, one visitor establishes that the aesthetics were recognizable, "I expected it to be Munch's images. It was a lot blurrier but still clearly Munch" (57, Female, post-launch). Another visitor connects this to a reflection on the green color specifically, "I think it was very in line with his artistry [...] This poisonous green color, it's something that I've personally been thinking a lot about with his images." (37, Male, post-launch). One visitor related the play with perspective in *Poison* to paintings elsewhere in the museum, "When you are looking at the paintings downstairs you can walk back and forth. [...] When you're walking in front of them you see that he was good at creating this depth in his images." (36, Male, post-launch). For yet another visitor it was the feeling of the space that evoked associations with other Munch paintings, "I don't know much about Munch except for the Scream and these very famous paintings [...] The room in itself gave me this mysterious [feeling] that the paintings also did" (19, Female, post-launch). These reflections made by the visitors indicate that the non-verbal elements of *Poison* played a similar role in evoking connections and reflection in comparison to symbolic and textual elements. One visitor even suggested that *Poison* went further than other exhibitions:

> "In relation to what I saw on the first floor [exhibition of the permanent collection], it's like you are deeper and it explains a lot of things. It's well related to what you feel about his artistic environment. So yeah, I think it's really accurate. [...] All his paintings are for me, and I think for a lot of people really oppressing, and that feeling was well demonstrated by the room, the lights, the sound, and everything. It explained... It's yeah, accurate." (21, Male, post-launch).

Another important thing to note is that the reflections following from the aesthetic qualities revolve around the aspects that we have designed for such as spatiality, the green color, and the unsettling atmosphere, while the first three examples in this section revolved around Munch's life more generally like the text. This shows that we as designers have the opportunity to use both verbal and non-verbal design elements to direct visitors' attention to salient aspects of the artworks, point out relations, and spark reflection. Such reflections can also be critical to the perspective that *Poison* proposed, "I don't feel that I know how much of this experience is Munch. Was this really Munch? Did he think about the images in this way?" (28, Male, post-launch). Despite the skepticism expressed by some visitors, we interpret this as a sign that the participants become attentive to Munch and his practice and might be encouraged to seek an answer to their questions elsewhere in the museum. Indeed, many participants express a wish to see the original paintings from *The Green Room* series. They describe a curiosity that has been evoked but is not satisfied by the time they step out of *Poison*. For many participants, the tour through *Poison* ends too quickly. One participant describes it:

> "I will go, 'where are they? I want to see them!'. I'll be curious about their story. I don't know much about Munch, but it makes me curious about the painter and when in his life he painted it. What was this? Where was he at this point in his life? Didn't look too good! I'm even more curious about Munch and his work, and especially those three paintings. So, it has planted something that I'm going to think about in the coming hours" (52, Female, pre-launch).

C.8 Discussion

In this paper, we have argued for thinking about art mediation as an *education of attention* of the museum visitor. It allows us to not only consider how the design works in the immediate experience but also how it might be shaping the accumulated experience[3]. This perspective is relevant to align with the goals of museums to of-

[3]Some Northern European languages like German, Dutch, Danish, Norwegian, and Swedish distinguish between at least two forms of experience. The immediate experience: Erlebnis, beleving, oplevelse, upplevelse. The experience built over time: Erfahrung, ervaring, erfaring, erfarenhet

fer "varied experiences for education, enjoyment, reflection and knowledge sharing" (ICOM, 2022). As design researchers working in the museum domain, we must be able to evaluate our designs not only on what happens in the immediate encounter with our design artifacts but also take responsibility for how the design influences the sense-making of artworks and other cultural heritage objects. Other projects dealing with the mediation of specific artworks have been explicitly avoiding engaging with this dimension. Vi et al. (2017) present an exhibition that juxtaposes multisensory designs with paintings in the TATE museum. They state that they did not explore aesthetics and culture in the museum since it was outside their capabilities as HCI researchers. Instead, they focus on quantitative measures of emotion, such as liking, arousal, and engagement. Birchfield et al. present a tool to support "active learning activities" (2008, p. 968), but base their evaluation on observations of kids and informal conversation with docents and museum staff. Kortbek & Grønbæk (2008) present a design for "communicating art" using interactive installations. In their evaluation, they focus on whether the design was effective in establishing a relation between the communicative installation and the artworks but refrain from any considerations of how the presented material might have shaped the visitors' approach to the original artworks. In contrast, Steier (2014) and Christidou & Pierroux (2019) do not present new designs, but investigate how posing and touching respectively mediates interpretation of artworks. Both papers discuss how the concrete physical movements of the visitors are part of constituting the interpretation of the artwork. Looking at their results it seems evident that designers who are creating designs that facilitate visitors' engagement with artworks must consider how that in turn shapes their concrete interpretive practice.

Being aware of this responsibility, we suggest that mediating designs are grounded in correspondence between design professionals and domain experts to highlight qualities that are particularly salient to the artwork under consideration. This is how we have worked with Endresen's research to inform the design of *Poison* and how we understand the National Gallery has worked with *Leonardo: Experiencing a Masterpiece.*

Using *education of attention* as a lens, allows us to keep an eye on how the design mediates interpretation throughout the design process. Concretely, we did this through a phenomenological interviewing method to understand which experiences emerged as visitors went through *Poison*. We ended the interview with a question on the relevance of the exhibition to get a sense of the general attitude towards the design, and a question asking if they saw any relation to Edvard Munch and his practice. Analyzing the interview data, we have been able to see that the design directed the attention of the visitors to aspects related to the design intention, although the design failed to evoke a sense of co-performance. We also saw that this was both facilitated by verbal and non-verbal aspects of *Poison*. In this project, we have not been able to investigate how experiencing *Poison* might influence a later encounter with the paintings from *The Green Room*. This presents an opportunity for future research. We were, however,

able to get an understanding of how visitors related their experience to their existing knowledge of Munch. While we have not reported on it in depth in this paper, this method was also employed throughout the process to ensure the design would exhibit the intended qualities. This illustrates the usefulness of this approach for designers, curators, and others involved with the mediation of artworks or potentially other forms of cultural heritage.

The method suggested here is by no means sufficient for a complete evaluation of visitor experience and learning. Our analysis might have benefited from gathering further data about the participants' background, visiting context, and motivation for being in the museum. This was not done in our study as we needed to keep the interviews short, taking into consideration that we were interviewing paying visitors who came for other reasons than being research subjects. Other researchers suggest more comprehensive frameworks for this purpose (Falk & Dierking, 2018; Kaptelinin, 2011). Nevertheless, this is a practical resource for designers, as it only requires small adjustments to the qualitative interview practice that is often employed in design projects. We based our evaluation of *Poison* on three questions that can be used in other settings as well that we reformulate in more general terms here so that they serve as a tool for other designers and researchers:

1. Are the visitors accepting and open to the design or do they reject it or label it off the cuff?

2. What are the aesthetic qualities that the visitors report? Are they in correspondence with the design intent?

3. How do visitors place their experience in relation to their existing knowledge on the topic/artwork/artist?

This presents a foundation for later research projects to build upon, developing even stronger methods to evaluate the *education of attention* concerning art and cultural heritage.

In this project, we have paid special attention to non-verbal means of mediation. We think the choice of medium and technologies should be made based on the qualities of the artwork that is the object of the mediation effort. Endresen's research highlights qualities of space, performativity, and atmosphere, leading to an immersive design. The National Gallery used technology to draw attention to painting techniques and the use of light, as they deemed it particularly relevant for *Virgin of The Rocks*. In general, however, we find art museums tend to be biased toward verbal forms of mediation, which is effective in drawing attention to narrative, art historical facts, and biographical information. Some things are best pointed out using words, others through physical or atmospheric examples. It is the job of the designer to make sure that the means used are the ones that best support drawing attention to

the aspects of interest. If art museums are interested in diversifying their appeal to attract a broader and younger audience, the mediation practices must be similarly diverse.

C.9 Conclusion

We have presented an enactive perspective on art along with the concept of *education of attention*. This perspective opens a possibility to conceptualize and evaluate art mediation across non-verbal and verbal modes of expression, to create holistic designs that have the potential to support visitors in making sense of artworks. We have argued that this can be done in correspondence with art historical research, designing for the *education of visitors' attention* to aspects of particular interest in the art. We have demonstrated this approach through the project *Poison*. Through interviews with visitors, we identified when the design worked as intended as well as its flaws. Designed with an emphasis on non-verbal mediation we found that *Poison* supported visitors' sense-making in correspondence with the design intent and in turn art historical research. Finally, we have discussed the potential of this approach as a tool for designers, curators, and other professionals in the museum domain to conceptualize and evaluate art mediation designs.

Acknowledgments

This project has been realized thanks to the support of MUNCH and many internal partners, as well as the research division of Random International, Torsteinsen Arkitekter and Martin Horntveth. Special thanks to Maarten Smith and Sander van der Zwan from Eindhoven University of Technology for helping with developing the theoretical foundation for this project.

References

Alexander, J., Wienke, L., & Tiongson, P. (2017). Removing the barriers of Gallery One: A new approach to integrating art, interpretation, and technology.*Museums and the Web 2017: Selected Papers and Proceedings from an International Conference*. Museums and the Web 2017, Cleveland, OH. https://mw17.mwconf.org/paper/removing-the-barriers-of-gallery-one-a-new-approach-to-integrating-art-interpretation-and-technology/

Bedford, L. (2014). *The Art of Museum Exhibitions: How Story and Imagination Create Aesthetic Experiences*. Left Coast Press.

Benford, S., Løvlie, A. S., Ryding, K., Rajkowska, P., Bodiaj, E., Paris Darzentas, D., Cameron, H., Spence, J., Egede, J., & Spanjevic, B. (2022). Sensitive Pictures: Emotional Interpretation in the Museum. *CHI Conference on Human Factors in Computing Systems*, 1–16. https://doi.org/10.1145/3491102.3502080

Benjamin, W. (1936). *The work of art in the age of mechanical reproduction*. Visual Culture: Experiences in Visual Culture, 144–137.

Bimber, O., Coriand, F., Kleppe, A., Bruns, E., Zollmann, S., & Langlotz, T. (2005). Superimposing pictorial artwork with projected imagery. *ACM SIGGRAPH 2005 Courses*, 6-es. https://doi.org/10.1145/1198555.1198716

Birchfield, D., Mechtley, B., Hatton, S., & Thornburg, H. (2008). Mixed-reality learning in the art museum context. *Proceedings of the 16th ACM International Conference on Multimedia*, 965–968. https://doi.org/10.1145/1459359.1459534

Boehner, K., Thom-Santelli, J., Zoss, A., Gay, G., Hall, J. S., & Barrett, T. (2005). Imprints of place: Creative expressions of the museum experience. *CHI '05 Extended Abstracts on Human Factors in Computing Systems*, 1220–1223. https://doi.org/10.1145/1056808.1056881

Budge, K. (2018). Visitors in immersive museum spaces and Instagram: Self, place-making, and play. *The Journal of Public Space*, 3(3), 121–138. https://doi.org/10.32891/jps.v3i3.534

Campbell, C. (2019, December 9). *Curator's Introduction to 'Leonardo: Experience a Masterpiece'*.

Christidou, D., & Pierroux, P. (2019). Art, touch and meaning making: An analysis of multisensory interpretation in the museum. *Museum Management and Curatorship, 34*(1), 96–115. https://doi.org/10.1080/09647775.2018.1516561

Cosley, D., Baxter, J., Lee, S., Alson, B., Nomura, S., Adams, P., Sarabu, C., & Gay, G. (2009). A tag in the hand: Supporting semantic, social, and spatial navigation in museums. *Proceedings of the SIGCHI Conference on Human Factors in Computing Systems*, 1953–1962. https://doi.org/10.1145/1518701.1518999

Cosley, D., Lewenstein, J., Herman, A., Holloway, J., Baxter, J., Nomura, S., Boehner, K., & Gay, G. (2008). ArtLinks: Fostering social awareness and reflection in museums. *Proceedings of the SIGCHI Conference on Human Factors in Computing Systems*, 403–412. https://doi.org/10.1145/1357054.1357121

Dewey, J. (1980). *Art as Experience*. Wideview/Perigree. (Original work published 1934)

Duncan, C. (2005). The art museum as ritual. *Heritage, Museums and Galleries: An Introductory Reader*, 78–88.

Endresen, S. (2015). *Serial Experiments. Close-Readings of Edvard Munch's Det grønne værelset (1907)*. Universitetet i Oslo.

Falk, J. H., & Dierking, L. D. (2018). *Learning from Museums*. Rowman & Littlefield.

Harrington, M. C. R. (2020). Connecting User Experience to Learning in an Evaluation of an Immersive, Interactive, Multimodal Augmented Reality Virtual Diorama in a Natural History Museum & the Importance of Story. *2020 6th International Conference of the Immersive Learning Research Network (ILRN)*, 70–78. https://doi.org/10.23919/iLRN47897.2020.9155202

Hornecker, E., & Ciolfi, L. (2019). Human-Computer Interactions in Museums. *Synthesis Lectures on Human-Centered Informatics, 12*(2), i–153. https://doi.org/10.2200/S00901ED1V01Y201902HCI042

Howes, D. (2014). Introduction to Sensory Museology. *The Senses and Society, 9*(3), 259–267. https://doi.org/10.2752/174589314X14023847039917

ICOM. (2022). *Museum Definition*. International Council of Museums. https://icom.museum/en/resources/standards-guidelines/museum-definition/

Ingold, T. (2001). From the transmission of representation to the education of attention. In *The debated mind: Evolutionary psychology versus ethnography* (pp. 113–153). Berg.

Jeffrey, S. (2015). Challenging Heritage Visualisation: Beauty, Aura and Democratisation. *Open Archaeology, 1*(1). https://doi.org/10.1515/opar-2015-0008

Jones, S., Jeffrey, S., Maxwell, M., Hale, A., & Jones, C. (2018). 3D heritage visualisation and the negotiation of authenticity: The ACCORD project. *International Journal of Heritage Studies, 24*(4), 333–353. https://doi.org/10.1080/13527258.2017.1378905

Kaptelinin, V. (2011). Designing Technological Support for Meaning Making in Museum Learning: An Activity-Theoretical Framework. *2011 44th Hawaii International Conference on System Sciences*, 1–10. https://doi.org/10.1109/HICSS.2011.152

Kenderdine, S. (2015). Embodiment, Entanglement, and Immersion in Digital Cultural Heritage. In S. Schreibman, R. Siemens, & J. Unsworth (Eds.), *A New Companion to Digital Humanities* (pp. 22–41). John Wiley & Sons, Ltd. https://doi.org/10.1002/9781118680605.ch2

Kidd, J. (2018). "Immersive" Heritage Encounters. *The Museum Review, 3*(1), 16.

Kortbek, K. J., & Grønbæk, K. (2008). Communicating Art through Interactive Technology: New Approaches for Interaction Design in Art Museums. NordiCHI '08: *Proceedings of the 5th Nordic Conference on Human-Computer Interaction: Building Bridges*, 229–238. https://doi.org/10.1145/1463160.1463185

Krogh, P. G., & Koskinen, I. (2020). *Drifting by Intention: Four Epistemic Traditions from within Constructive Design Research*. Springer International Publishing. https://doi.org/10.1007/978-3-030-37896-7

Latour, B., & Lowe, A. (2011). The Migration of the Aura Exploring the Original Through Its Fac similes. In T. Bartscherer & R. Coover (Eds.), *Switching Codes: Thinking Through Digital Technology in the Humanities and the Arts* (pp. 275–297). University of Chicago Press.

Louvre. (2020). *Mona Lisa: Beyond the Glass*.

Løvlie, A. S., Ryding, K., Spence, J., Rajkowska, P., Waern, A., Wray, T., Benford, S., Preston, W., & Clare-Thorn, E. (2021). Playing Games with Tito: Designing Hybrid Museum Experiences for Critical Play. *J. Comput. Cult. Herit., 14*(2). https://doi.org/10.1145/3446620

Macdonald, S. (2007). Interconnecting: Museum visiting and exhibition design. *CoDesign, 3*(sup1), 149–162. https://doi.org/10.1080/15710880701311502

Mathias, N. (2022). Meta-artistic immersion in digital exhibitions. History – mobilization – spectatorship. *Journal of Aesthetics & Culture, 14*(1), 2129160. https://doi.org/10.1080/20004214.2022.2129160

Mondloch, K. (2022). The Influencers: Van Gogh Immersive Experiences and the Attention-Experience Economy. *Arts, 11*(5), 90. https://doi.org/10.3390/arts11050090

Noë, A. (2016). *Strange tools: Art and human nature*. Hill and Wang, a division of Farrar, Straus and Giroux.

Noë, A. (2021).*Learning to look: Dispatches from the art world* (Book edition). Oxford University Press.

Sivertsen, C., Smith, M., & van der Zwan, S. (2023). Art Critique by Other Means.*Designing Interactive Systems Conference (DIS '23), 10*. https://doi.org/10.1145/3563657.3596069

Snibbe, S. S., & Raffle, H. S. (2009). Social immersive media: Pursuing best practices for multi-user interactive camera/projector exhibits. *Proceedings of the SIGCHI Conference on Human Factors in Computing Systems*, 1447–1456. https://doi.org/10.1145/1518701.1518920

Steier, R. (2014). Posing the Question: Visitor Posing as Embodied Interpretation in an Art Museum. *Mind, Culture, and Activity, 21*(2), 148–170. https://doi.org/10.1080/10749039.2013.878361

Stogner, M. B. (2011). "The Immersive Cultural Museum Experience – Creating Context and Story with New Media Technology". *The International Journal of the Inclusive Museum, 3*(3), 117–130. https://doi.org/10.18848/1835-2014/CGP/v03i03/44339

Tennent, P., Martindale, S., Benford, S., Darzentas, D., Brundell, P., & Collishaw, M. (2020). Thresholds: Embedding Virtual Reality in the Museum. *Journal on Computing and Cultural Heritage, 13*(2), 1–35. https://doi.org/10.1145/3369394

The National Gallery. (2019, 2020). *Leonardo: Experience a Masterpiece*.

Thompson, C. J., Locander, W. B., & Pollio, H. R. (1989). Putting Consumer Experience Back into Consumer Research: The Philosophy and Method of Existential-Phenomenology. *Journal of Consumer Research, 16*(2), 133–146. https://doi.org/10.1086/209203

Vi, C. T., Ablart, D., Gatti, E., Velasco, C., & Obrist, M. (2017). Not just seeing, but also feeling art: Mid-air haptic experiences integrated in a multisensory art exhibition. *International Journal of Human-Computer Studies, 108*, 1–14. https://doi.org/10.1016/j.ijhcs.2017.06.004

Waern, A., & Løvlie, A. S. (Eds.). (2022). *Hybrid museum experiences: Theory and design*. Amsterdam University Press. https://www.jstor.org/stable/10.2307/j.ctv2cxx8x6

Walter, T. (1996). From museum to morgue? Electronic guides in Roman Bath. *Tourism Management, 17*(4), 241–245. https://doi.org/10.1016/0261-5177(96)00015-5

Wensveen, S. A. G., Djajadiningrat, J. P., & Overbeeke, C. J. (2004). Interaction frogger: A design framework to couple action and function through feedback and feedforward. *Proceedings of the 2004 Conference on Designing Interactive Systems Processes, Practices, Methods, and Techniques - DIS '04*, 177. https://doi.org/10.1145/1013115.1013140

Zimmerman, J., Forlizzi, J., & Evenson, S. (2007). Research Through Design As a Method for Interaction Design Research in HCI. *Proceedings of the SIGCHI Conference on Human Factors in Computing Systems*, 493–502. https://doi.org/10.1145/1240624.1240704

Publication D

Exploring a Digital Art Collection through Drawing Interactions with a Deep Generative Model

Christian Sivertsen, René Haas, Halfdan Hauch Jensen and Anders Sundnes Løvlie

This extended abstract has been published in
Extended Abstracts of the 2023 CHI Conference on Human Factors in Computing Systems (CHI EA '23), April 23-28, 2023, Hamburg, Germany. ACM, New York, NY, USA, 5 pages.
https://doi.org/10.1145/3544549.3583902

The layout has been revised.

Abstract

New Snow is an interactive drawing table that investigates human interaction with a deep generative model based on Edvard Munch's sketching practice. Through drawings with pen and paper, the user can interact with the model which will return synthetic sketches based on the input drawings in real time. The model is a reflection of the training data, and it is thus constrained to representing images within the latent space of Edvard Munch's sketching practice. As the user familiarizes themselves with the model it allows them to become sensitized to the visual aesthetic belonging to this practice. This potential for familiarization with the aesthetic of a dataset via the model has implications for human-AI interaction and non-verbal art mediation.

Fig. D.1: Inferring a Munch-like sketch from a hand-drawn line. From left to right: 1) infrared webcam image from beneath the drawing surface. 2) The cleaned input image. 3) The synthesized sketch. 4) A composite image of the input image over the synthesized sketch.

D.1 Introduction

Museums across the world have built up large collections of digitized artwork collections. This vast amount of material can be difficult to present to an audience. Several projects investigate ways to expose museum visitors to large datasets and allow for browsing or reauthoring the content (Kenderdine & McKenzie, 2013; Kenderdine et al., 2013; Lintermann, 2012; Shaw et al., 2006). However, collection interfaces are often designed to offer overview and searchability rather than engagement with the artworks. To support the mediation of large art collections, interfaces are necessary that support alternative and embodied modes of engagement.

Tapping into recent developments in image synthesis, New Snow attempts to explore a new way to engage with a large collection of digitized artworks through a novel machine-learning-enabled interface. Deep Generative Models (DGM) are models that can generate images that resemble the training data. These have recently become well-known through the current wave of text-to-image systems like Midjourney, DALL-E (Ramesh et al., 2022) and Stable Diffusion (Rombach et al., 2021) that allow for anyone to generate synthetic images in the style of famous artists, that is, artists whose works have a large presence on the internet from where much of the data for the underlying datasets are found. In contrast to these systems, New Snow

Fig. D.2: The left half of the image shows samples from the original sketch data. The right half shows synthetic samples from the StyleGAN model

employs a model trained specifically on the drawings of Edvard Munch (fig. D.2), to allow museum visitors to engage with this well-known artist's drawing practice through their own drawing actions. New Snow is an interactive drawing table that offers museum visitors an embodied mode of interaction, where the system responds to the user's drawing by adding lines and patterns generated by the DGM. This simulates an experience of the artist "filling in" the lines drawn by the visitor.

This project aims for three main contributions. First, it enables a way for visitors to engage with a large corpus of artworks that could not feasibly be explored individually, through the proxy of a DGM offering a synthesis of the data. Second, the system allows an embodied and creative engagement through the drawing actions of the visitor, and the interplay between the visitor and the system. Third, the system explores a novel use of a DGM, in which the user's efforts to learn how to interact with the model are offered as a way to learn about the aesthetics of Munch's drawings. As the user investigates the model through the drawing actions, the user learns about its qualities, and by proxy certain qualities of the artworks constituting the underlying dataset. This means that building a mental model of the system becomes a way of learning about the aesthetics of the drawings, and the image synthesis becomes an enabler of the exploration rather than the end goal.

D.2 Related Work

Large databases of cultural heritage collections are usually accessed through search interfaces letting the users find content based on written prompts and filters. This is useful for situations where the users have a good idea of what they are looking for. However, when the domain is unfamiliar to the user, curation and recommendation are often used as a way of guiding the user to relevant content. Earlier projects have attempted different visualization strategies for large cultural heritage datasets like *T_Visionarium II* (Shaw et al., 2006), *Cloudbrowsing* (Lintermann, 2012) and *ECLOUD WW1* (Kenderdine & McKenzie, 2013) that all utilize large projection surfaces to display content and allow the user to browse through the individual data objects.

The two projects *Draw to Art* (Google Creative Lab et al., 2018) and *Draw to Art: Shape Edition* (Google Creative Lab et al., 2020) explores visual search by allowing users to draw images on a tablet surface to query a large art database for artworks matching the drawing. In the first version, the match is based on classifying the input image as a word and then returning images relating to that word. In the second version, the system returns images with shapes matching the input image, which is constrained to simple geometric shapes. This difference marks a significant change as it enables exploration that is driven by visual concepts such as composition and shapes.

Human-AI interaction research stipulates that the system should provide the user with clear concepts of its capabilities (Amershi et al., 2019) by being *explainable* or *transparent* e.g. (Barredo Arrieta et al., 2020; Hois et al., 2019). Another related concept is that of *interpretable* AI (Doshi-Velez & Kim, 2017; Ghosh & Kandasamy, 2020), asserting that the users of a system should be able to interpret the underlying reasons for the output.

Building expertise in the interaction with image synthesis models is seen in *prompt engineering*, the practice of developing text-based prompts through optimization or exploration that makes text-to-image or text-to-text generation systems generate the content intended by the user. According to Oppenlaender (2022), achieving the best results requires a deep understanding of the underlying dataset.

Based on the works above we derive three insights, which have informed the design of New Snow:

1. We understand the deep generative models as reflecting qualities of the data from which it is trained.

2. Exploration of datasets and models can happen through non-verbal means.

3. Learning how to *prompt* a model effectively means building an understanding of the data on which it is trained.

D.3 Concept and interaction

New Snow is a project that explores how exploring an image synthesis model using drawing actions as a means of *prompting* lets the user learn about certain qualities of the underlying data.

From the digital collection of MUNCH, we have identified 5800 uncolored, crayon, ink, or pencil drawings made by Edvard Munch (see examples in fig. D.2). Based on these drawings we have trained a StyleGAN 2 model (Karras, Aittala, et al., 2020; Karras, Laine, et al., 2020) and then a pixel2style2pixel (pSp) model (Richardson et al., 2021). Together this allows us to map drawings made by a user into the model and synthesize sketches from this input.

The prototype consists of a table with a matte transparent surface. The user places a piece of tracing paper on the surface and draws with a pen (fig. D.3). Underneath the table, a camera tracks the lines on the paper and sends them to the pixel2style2pixel model. From the lines on the paper, the model synthesizes an image based on Edvard Munch's sketches (fig. D.1). This image is projected back onto the tracing paper for the user to see. As the user draws or moves the paper around, the system continuously and multiple times per second updates the synthesized image to match. This allows the user to explore the qualities of the model through an almost conversational relation to the system.

Fig. D.3: The left image is a render of the prototype with the sides open, so the inner structure can be seen. The image to the right shows a user interacting with the drawing interface.

The aim of this project is to help the user develop an attention to the aesthetic of Edvard Munch's drawing practice. We are not attempting to explain the technical details of his practice nor the historical or biographical relations. However, through the embodied engagement with the visual aesthetic derived from his drawings, we expect the user to develop a sense of the visual qualities related to Munch's sketching and drawing practice. That does not necessarily require the original works to be reproduced, the aim is rather to create a focus on the dynamics and patterns in Munch's sketching practice that a person not skilled in the act of drawing or analyzing drawings might not otherwise have noticed.

Models like Dall-E, Midjourney, and Stable Diffusion have become famous for their ability to synthesize coherent images from almost any prompt in the style of well-known artists. In comparison, this model provides much more resistance. It does not draw for the user, but it responds to their drawings and attempts to expand and complete them. Due to the nature of the underlying model, the user will have to adapt a particular drawing strategy in order to achieve the greatest level of control, as the model does not respond to symbolic representation but rather the saliency of particular lines constituting an image.

This tension requires the user to explore the workings of the model to understand how it works and what it responds to. In that way, the function of the model is to provide a space to explore rather than being a tool to reach other ends. The

exploration is by proxy an exploration of Munch's sketching practice. It asks the user to contemplate and then draw not *what* Munch might have drawn, but *how* he might have drawn.

At the core of the experience is the user interaction with the DGM, however, that interaction happens within an embodied and material context. People's bodily relations to artworks shape how they might cast the art objects in specific cultural roles i.e. as a commodity, a fragile piece of history, or a toy (Sivertsen & Løvlie, 2021). With this awareness we expect the embodied relation to the drawings in New Snow to influence people's cultural connotations of the drawing activity. Thus we have opted for a physical setup where the user performs the drawing action with an actual pen on paper. First, the paper and pen have other affordances than a touchscreen and pen interface, one being that erasing is not possible, and the drawn image can be moved around on the surface, lifted off the surface, and brought along. Secondly, the tactile feeling of pen and paper differs significantly from the glass surface of a tablet and e-pen and evokes different connotations and importantly a closer material connection to the tools used by Edvard Munch.

D.4 Technical description

The drawing table is built into a flight case 110 cm tall with a semi-transparent polycarbonate window on top. The drawing surface is lit with infrared (IR) light by LEDs within the table to eliminate shadows. An IR-sensitive camera fitted with an 850nm filter records the drawing surface (fig. D.4). This is to avoid interference from the visible light cast by the projector. With software made with TouchDesigner, OpenCV, and Python the video feed is pre-processed into binary images, isolating the lines drawn on the paper. This image is submitted to the pixel2style2pixel model and a synthesized drawing is returned within a second. Adjustments are made to saturation and contrast before the projected image interpolates from the current to the new image.

D.4.1 Machine Learning

DGMs have seen tremendous progress in recent years and Generative Adversarial Networks (GANs) (Goodfellow et al., 2014) have become one of the most influential deep generative architectures. Recently, inspired by style transfer (Gatys et al., 2015; Huang & Belongie, 2017), the StyleGAN family of models (Karras, Aittala, et al., 2020; Karras et al., 2021; Karras, Laine, et al., 2020) have been shown to give state-of-the-art results across a wide variety of image generation tasks (Bermano et al., 2022). Due to the exceptional quality of the images generated by StyleGAN models, the architecture has been called one of the most intriguing and well-studied architectures in recent times (Alaluf et al., 2022).

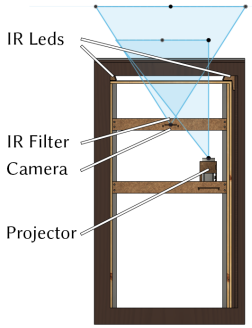

Fig. D.4: The prototype house a pico projector and a camera that are aligned with the drawing surface. An infrared filter removes the visible light from the camera input. Right below the top are two strips of infrared LED that illuminate the drawing surface. The lower part of the table can house a PC for processing the images.

After training, StyleGAN has learned a mathematical space, denoted the *latent* space, where each point in the space corresponds to a unique image. The GAN is trained such that the distribution of the generated images follows the distribution of the images in the training data, both with respect to image quality as well as the internal variation between images in the training data.

The latent space is smooth, which means that if we interpolate between two points in the latent space, e.g two points corresponding to portraits, the corresponding generated images will change gradually from one portrait to the other where each intermediate image is itself fully self-consistent and resembles Munch's style in its own right. Thus the latent space can be seen as a representation of the space between all Munch sketches in the data set.

To allow for direct user interaction with this latent space, we have trained a pSp encoder (Richardson et al., 2021) which is able to map user-provided pen-and-paper sketches, into the latent space, thus transforming the user input to a sketch that follows the style of Munch.

D.5 Aesthetic drawing strategies

In preliminary testing of the prototype, we have seen participants engaging actively with the drawing task. Participants apply widely different strategies, and we see indications that certain mental models yield more satisfying interactions than others. When users draw conceptually, e.g. a simplified house or tree, the system generates only limited visual response since these shapes lie far away from the images in the dataset. A more fruitful drawing strategy seems to be drawing one long stroke at first

and then looking at the response of the system. Often a variety of more or less defined lines will appear. These lines can be reinforced or challenged by drawing other lines in the same area, which often results in more defined shapes and features, and the process can continue and evolve into a meaningful drawing. These are the strategies that we are interested in exploring and tuning the system to support.

D.6 Curating the dataset

As the aesthetic qualities in Edvard Munch's sketching practice are mediated through the DGM, particular attention needs to be paid to the ways in which the sketches have been prepared for training and how the chosen model interprets the data. The images constituting the dataset for the StyleGAN model have been cropped from photographs of notebooks or loose paper sheets by human annotators that have made decisions on composition and the tightness of the crop to leave out damaged paper, smudges, handwritten notes, and other artifacts that have been deemed irrelevant for the project. This curation shapes the concept of the images created by the model. It determines what belongs to a drawing, and where on the page certain shapes will appear.

Another limitation is the necessity for this type of model to be trained with square images. This requires the input images to be either stretched or cropped to fit this requirement. These issues reappear when the pSp model is trained as the input images are simplifications derived from the syntheses. However, the amount of simplification determines how far from the input image the synthesized images will be visually. This means that a significant part of the interaction design lies in the data curation process, making the iterative loop longer and more time-consuming than when designing heuristic interactive systems.

Acknowledgments

The authors would like to thank MUNCH, Oslo for helping in developing the concept and funding the prototype. We are thankful to the research division at Random International for the development of the tracking system and AIRLab at IT University of Copenhagen for supporting the physical realization of the prototype.

References

Alaluf, Y., Patashnik, O., Wu, Z., Zamir, A., Shechtman, E., Lischinski, D., & Cohen-Or, D. (2022). Third time's the charm? image and video editing with stylegan3 [arXiv: 2201.13433]. http://arxiv.org/abs/2201.13433

Amershi, S., Weld, D., Vorvoreanu, M., Fourney, A., Nushi, B., Collisson, P., Suh, J., Iqbal, S., Bennett, P. N., Inkpen, K., Teevan, J., Kikin-Gil, R., & Horvitz, E. (2019). Guidelines for Human-AI Interaction. *Proceedings of the 2019 CHI Conference on Human Factors in Computing Systems*, 1–13. https://doi.org/10.1145/3290605.3300233

Barredo Arrieta, A., Díaz-Rodríguez, N., Del Ser, J., Bennetot, A., Tabik, S., Barbado, A., Garcia, S., Gil-Lopez, S., Molina, D., Benjamins, R., Chatila, R., & Herrera, F. (2020). Explainable Artificial Intelligence (XAI): Concepts, taxonomies, opportunities and challenges toward responsible AI. *Information Fusion*, *58*, 82–115. https://doi.org/10.1016/j.inffus.2019.12.012

Bermano, A. H., Gal, R., Alaluf, Y., Mokady, R., Nitzan, Y., Tov, O., Patashnik, O., & Cohen-Or, D. (2022). State-of-the-art in the architecture, methods and applications of stylegan. https://doi.org/10.48550/ARXIV.2202.14020

Doshi-Velez, F., & Kim, B. (2017). Towards A Rigorous Science of Interpretable Machine Learning [arXiv:1702.08608 [cs, stat]]. Retrieved December 21, 2022, from http://arxiv.org/abs/1702.08608

Gatys, L. A., Ecker, A. S., & Bethge, M. (2015). A neural algorithm of artistic style. http://dblp.uni-trier.de/db/journals/corr/corr1508.html#GatysEB15a

Ghosh, A., & Kandasamy, D. (2020). Interpretable Artificial Intelligence: Why and When [Publisher: American Roentgen Ray Society]. *American Journal of Roentgenology*, *214*(5), 1137–1138. https://doi.org/10.2214/AJR.19.22145

Goodfellow, I., Pouget-Abadie, J., Mirza, M., Xu, B., Warde-Farley, D., Ozair, S., Courville, A., & Bengio, Y. (2014). Generative adversarial nets (Z. Ghahramani, M. Welling, C. Cortes, N. Lawrence, & K. Q. Weinberger, Eds.). https://proceedings.neurips.cc/paper/2014/file/5ca3e9b122f61f8f06494c97b1afccf3-Paper.pdf

Google Creative Lab, Girschig, B., & Cazier, R. (2020). Draw to Art: Shape Edition. Retrieved December 2, 2022, from https://experiments.withgoogle.com/draw-to-art-shape

Google Creative Lab, Google Art & Culture Lab, & IYOIYO. (2018). Draw to Art. Retrieved December 2, 2022, from https://experiments.withgoogle.com/draw-to-art

Hois, J., Theofanou-Fuelbier, D., & Junk, A. J. (2019). How to Achieve Explainability and Transparency in Human AI Interaction. In C. Stephanidis (Ed.), *HCI International 2019 - Posters* (pp. 177–183). Springer International Publishing. https://doi.org/10.1007/978-3-030-23528-4_25

Huang, X., & Belongie, S. (2017). Arbitrary style transfer in real-time with adaptive instance normalization. https://doi.org/10.48550/ARXIV.1703.06868

Karras, T., Aittala, M., Hellsten, J., Laine, S., Lehtinen, J., & Aila, T. (2020). Training generative adversarial networks with limited data.

Karras, T., Aittala, M., Laine, S., Härkönen, E., Hellsten, J., Lehtinen, J., & Aila, T. (2021). Alias-free generative adversarial networks.

Karras, T., Laine, S., Aittala, M., Hellsten, J., Lehtinen, J., & Aila, T. (2020). Analyzing and improving the image quality of StyleGAN.

Kenderdine, S., & McKenzie, H. (2013). A war torn memory palace: Animating narratives of remembrance. *2013 Digital Heritage International Congress (DigitalHeritage)*, 315–322. https://doi.org/10.1109/DigitalHeritage.2013.6743755

Kenderdine, S., Shaw, J., & Gremmler, T. (2013). Cultural Data Sculpting: Omnidirectional Visualization for Cultural Datasets. In F. T. Marchese & E. Banissi (Eds.), *Knowledge Visualization Currents* (pp. 199–220). Springer London. https://doi.org/10.1007/978-1-4471-4303-1_11

Lintermann, B. (2012). Beyond Cinema. Retrieved December 1, 2022, from https://www.bernd-lintermann.de/papers/Beyond_Cinema_Lintermann.pdf

Oppenlaender, J. (2022). A Taxonomy of Prompt Modifiers for Text-To-Image Generation [arXiv:2204.13988 [cs]]. Retrieved January 4, 2023, from http://arxiv.org/abs/2204.13988

Ramesh, A., Dhariwal, P., Nichol, A., Chu, C., & Chen, M. (2022). Hierarchical text-conditional image generation with clip latents. https://doi.org/10.48550/ARXIV.2204.06125

Richardson, E., Alaluf, Y., Patashnik, O., Nitzan, Y., Azar, Y., Shapiro, S., & Cohen-Or, D. (2021). Encoding in Style: A StyleGAN Encoder for Image-to-Image Translation. *IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*.

Rombach, R., Blattmann, A., Lorenz, D., Esser, P., & Ommer, B. (2021). High-resolution image synthesis with latent diffusion models.

Shaw, J., Brown, N., Del Favero, D., McGinity, M., & Weibel, P. (2006). T_visionarium II. Retrieved December 1, 2022, from https://www.jeffreyshawcompendium.com/portfolio/t_visionarium-ii/

Sivertsen, C., & Løvlie, A. S. (2021). Handling Digital Reproductions of Artworks. *Journal of Somaesthetics*, *7*(2), 21.

Publication E

Exploring Aesthetic Qualities of Deep Generative Models through Technological (Art) Mediation

Christian Sivertsen and Anders Sundnes Løvlie

A revised version of this paper has been submitted to
DIS '24.

The layout has been revised.

Abstract

Deep Generative Models (DGM) have had a great impact both on visual art and broader visual culture. In this research-through-design project we investigate the use of a DGM for helping museum visitors explore the aesthetics of Edvard Munch's art. We developed an interactive drawing table allowing a user to explore a StyleGAN model trained on sketches by Edvard Munch. The design was evaluated by testing on 20 participants, and observations and interviews were analyzed qualitatively. We discuss how this setup seems to invite certain types of exploration and highlight certain elements of Munch's aesthetics. We suggest that the resulting interaction may contribute to an "education of attention" helping museum visitors to become attentive to certain elements of the artists' visual aesthetics. Finally, we discuss implications for the design of interactive experiences with DGMs relating to data curation and model training and how the resulting model mediates the training data.

E.1 Introduction

The recent wave of Deep Generative Models (DGM) capable of synthesizing both convincing text and images has had an extensive impact on visual culture as evidenced by the field of AI Art (Audry, 2021; Zylinska, 2020), and has facilitated new practices for the production of images (Oppenlaender, 2022). However, these technologies have also raised questions about ownership, copying, and manipulation. Critical scrutiny of large image datasets has shown that problematic biases and classifications can end up being deeply embedded in systems that many rely on for both research, art, and business (Crawford & Paglen, 2021). Paraphrasing a presentation by Lorraine Datson, Malevé argues that when working with machine learning "Engineers are said to write programs that 'discover' the rule inherent to the data" (Malevé, 2021).

The opaque data curation practices employed by the developers of off-the-shelf models make it difficult to assess and scrutinize the implications and particular qualities of their models. It is difficult to understand what perspective on the world these models represent, nevertheless, the inherent patterns and biases may "leak" to shape our ideas about our current world and about the future (Benjamin et al., 2021). Efforts in Explainable and Interpretable AI (Barredo Arrieta et al., 2020; Doshi-Velez & Kim, 2017; Ghosh & Kandasamy, 2020; Hois et al., 2019; Zhu et al., 2018) aims through technical analysis at describing the processes through which machine learning (ML) models derive their output. Equally important is understanding the aesthetic qualities of the output it produces. Aesthetic production is after all the promise of many DGMs. According to Benjamin et al. (2021) *thingly uncertainty* is a property of ML-enabled artifacts. They act within a continuum of relations to their environment and humans. Training a DGM means establishing an unknown space of potential, in which the model may act. The concept of *pattern leakage* (Benjamin et al., 2021) offers

a way to understand the aesthetic quality and potential of this space: Through its output a DGM "leaks" its embedded patterns, which allows us to create designs affording an experiential evaluation from a continuum of perspectives. Due to the probabilistic nature of machine learning, investigating a dataset by proxy of its model will only give us a view of the trends, rather than the particularities of each data object. However, a broad understanding becomes a necessity through the sheer volume of data. Investigating these embedded patterns calls for a deliberate practice of designing interfaces that support the *reflexive use* of DGMs, meaning that the functionality of the system is to expose itself.

To exemplify and evaluate this approach we will present a research-through-design project (Zimmerman et al., 2007) where we developed an interactive drawing table allowing a user to explore a StyleGAN model trained on sketches by Edvard Munch. The project was carried out in collaboration with the museum MUNCH, using a selection of 5800 line drawings from the museum's digital collection to train a DGM. This model is used in combination with an interactive drawing table, where the user draws with a pen on paper, to explore the latent space of the DGM. By tracking the paper in real-time and projecting the synthesized image back on the paper, the drawing table supports a fluent iterative relation between the user and the dataset.

Taking a postphenomenological perspective we will analyze how the design mediates particular relations between the user, the design, and the underlying dataset. Through interviews with 20 participants, we present a detailed view, of how this mediation plays out in practice and in relation to the specific dataset. While this approach of presenting synthesized drawings might seem controversial in an art museum context, it hinges on the understanding of art mediation not as providing the particularities - this is for the artworks themselves to do - but to evoke a directed interest and *educate our attention* (Ingold, 2001; Noë, 2021) towards qualities that might otherwise have been overlooked. This aligns with an interest in supporting the evaluation and understanding DGMs more generally. We will end by discussing how the overlap between technological mediation and art mediation shaped the orientation of the research and how the idea of designing for reflexive use of AI may have relevance in other domains where DGMs are employed.

E.2 Related Work

Recently DGMs such as StyleGAN (Karras, Aittala, et al., 2020; Karras et al., 2021), Dall-E (Ramesh et al., 2022), Midjourney and Stable Diffusion (Blattmann et al., 2022), which are capable of generating a wide range of images from verbal prompts, have become popular. Such models have been applied by artists to create new forms of imagery and challenge the relation between artists and technology (Akten, 2021; Akten et al., 2019; Audry, 2021; Boden & Edmonds, 2019; Cetinic & She, 2022; Grba,

2022; Guljajeva & Canet Sola, 2022; Zylinska, 2020). Furthermore, these technologies have also entered popular discourse and been used by amateurs in a wide variety of image-making practices (Oppenlaender et al., 2023).

However, the introduction of these technologies into the mainstream has also sparked discussions about the practices regarding the collection of data that the models are trained on. One critical issue regards the possible infringement of artists' intellectual rights to their own works and style, as they have been included in the training data (Appel et al., 2023; Chen, 2023). Another significant question regards the risk of propagating gender, racial and other stereotypes (Atwood et al., 2020; Birhane et al., 2021; Crawford & Paglen, 2021; Denton et al., 2021; Larrazabal et al., 2020; Wang et al., 2022). In both cases, the opaqueness of the data collection practice as well as the black-boxed nature of the generative models makes them difficult to scrutinize.

Research around *AI as design material* tends to focus on how AI can be put to work solving problems, however, it is also acknowledged to be particularly difficult to work with. Yang et al. (2020) argue that two central attributes that make it difficult to design with AI are *capability uncertainty* and *output complexity*. *Capability uncertainty* relates to the difficulty of knowing about the capabilities of an AI system to perform a given task before it is built and the data specific to that task have been sourced. AI systems developing through use and their capability of acting in relation to contextual factors make this problem even harder. *Output complexity* points to the situation where there are many potential outputs from a system. With DGM we often talk about many-dimensional variations of the output. Furthermore, Leahu (2016) argues that we should be aware of *ontological surprises* when working with machine learning, as particular relations and categories may arise that we did not foresee as a consequence of the specific configurations of technology, humans, and context.

Postphenomenology presents a number of possible relations between person, technology and world, of which the *hermeneutic relation* and the *alterity relation* is of particular relevance for the acts of engaging with art through a machine learning system (Benjamin et al., 2021; Rosenberger & Verbeek, 2015). The *hermeneutic relation* is understood as analogous to reading a text. Different experience and skill in reading gives access to the meaning of the text in different ways. The skillful reader almost sees through the letters and perceives the meaning in one quick glance, while a novice reader might start by constructing words from letters, sentences from words, and so on. The alterity relation describes situations where the interactions between humans and technology are somewhat similar to that of two humans interacting, not necessarily meaning the user is fooled, to believe they are interacting with another human, but simply that the form of the interaction is similar and that the technology may act as a *quasi-other* (Rosenberger & Verbeek, 2015). The tendency of people to anthropomorphize technology makes this a relevant relation for machine learning systems with the capability of acting in a seemingly autonomous way.

Benjamin et al. (2023) presented the *entoptic metaphor* as a way to describe how machine learning systems can give rise to visual phenomena in a way similar to how the human visual system can produce phenomena such as floaters or hallucinations within its system. In a more general sense, they see machine learning systems as making the world *legible* in different ways, establishing situated AI literacies as a consequence of the particular technological mediation. The introduction of the entoptic metaphor is intended to support designerly inquiry into the materiality of machine-learning systems and the concepts and implications that emerge as a part of a situated investigation.

In a different paper, Benjamin et al. presented the concept of *pattern leakage*: That is, "the propensity of probabilistic patterns to shape the world they are deployed to represent" (Benjamin et al., 2021, p.11). For instance, a surveillance system classifying events might affect how humans see the world. Generative algorithms on the other hand are designed to produce images and texts as outputs, and as such are designed to contribute to shaping the world through their outputs. We propose that the concept of pattern leakage might be understood as the core functionality of generative models. Through their output, they make explicit their inherent patterns and enable us to learn about them, which we can make use of, as we shall discuss further later on.

In this paper, we describe a design developed in the context of the art museum, more specifically within the topic of *art mediation*, which can be summarized as supporting art museum visitors' perceptual access to the artworks on display. In the field of HCI the postphenomenological concept of *technological mediation* has been used to describe how technologies *reveal* and *conceal* aspects of both the world and the human (Kiran, 2015; Rosenberger & Verbeek, 2015). In this paper, we understand the task of mediating art using technology as being a matter of understanding how this shaping is done by the technologies employed. Kiran (2015) argues that technological mediation can be understood as working in different dimensions and presents the *ontological*, the *epistemological*, the *practical*, and the *ethical* dimensions as relevant dimensions to consider. The *epistemological* dimension is particularly relevant for the topic of art mediation because it allows us to consider how the technology employed *magnifies* and *reduces* our perceptual capabilities of the artworks in question. Through the technological mediation the artworks *manifest* themselves both in relation to the material properties of the concrete technology as well as in relation to the *task at hand* for the museum visitor that will be using the technology. This means that the user of the system, the system itself, and the artworks are mutually constrained in facilitating a particular *perspective* on the artworks. This perspective is experienced through perceptual actions that as Scurto et al. (2021) point out, means that the *somaesthetic behavior* is significant for how users of a machine learning system are able to project themselves into it using their body and perception.

The most common way of making DGMs synthesize images is through verbal prompts. It requires the user to formulate in words what they want to see, to which the model will then respond with related visual concepts. This translation from words to images is constrained by the *imageability* (Malevé, 2021) of the verbal concept. However, certain encoders afford us the opportunity to stay within the visual domain, where images are used as prompts for other images. This sidesteps the issue of having to bridge the gap between verbal and visual expressions and allows for designing new interfaces, paying attention to the somaesthetic relation between the audience, system, and images. Interfaces that rely on drawing as input to the system have been explored within HCI research (Ariccia et al., 2022; Davis et al., 2015; de Lima et al., 2014; Yurman & Reddy, 2022). The projects Draw to Art (Google Creative Lab et al., 2018) and Draw to Art: Shape Edition (Google Creative Lab et al., 2020), allow users to make drawings in order to search a large database of artworks. In the first version, the drawing is interpreted as a word, and the system returns artworks related to that word. In the second version, the user draws with simple geometric shapes and the search returns artworks with a similar composition.

E.2.1 Designing for reflexive use

Deep generative models are capable of producing aesthetically rich images, sounds, and text. These qualities are always experienced from a situated perspective and are very defining for how the model mediates the world. Therefore designers and researchers need an approach for evaluating such qualities before a given model is deployed for use as part of a product.

A similar approach is seen in artistic engagements with machine learning and AI. With *ImageNet Roulette* Paglen and Crawford (2019) highlighted the problematic "person" category in ImageNet by allowing everyone to upload their own image to be classified with the categories from the "person" synset. The system's application of strange, discriminating, and outright offensive labels to the images made the aesthetic qualities of the system apparent through its concrete use. In Memo Akten's *Learning to See: Interactive* (Akten, 2017), the audience in the exhibition was able to interrogate five GAN models by showing everyday objects to a camera. The system would interpret the video feed through one of five models trained on images of water, fire, earth, air, and the cosmos. By manipulating and creating compositions of everyday objects, the audience was able to investigate the patterns and aesthetic qualities of each model. The service LAIKA (2023) aims to support creative writing by letting writers interrogate and explore the qualities of different language models. The models can be trained on the work of a famous artist, or on a corpus of the user's own text. The proposal is not that the system will generate finished text but rather that it will spark inspiration and reflection on writing style and patterns stimulating the user to write better or more creatively.

In these three examples, the interface provided to the user serves the purpose of giving them access to explore and interrogate the models in question. In this way, the interface supports a use that is *reflexive* with regards to the model, in that it points back to the model's qualities rather than towards some purpose external to the model itself. The modalities of the three interfaces - image upload, a video feed, and text prompts - are very different, and facilitate particular perspectives on the models and their mediation of the underlying data. In all three cases, it is the aesthetic qualities of the models that constitute the *work* done by the system.

Wolf (2021) proposed the term "Explorable AI" arguing for designing AI systems to "to support and empower actors to scrutinize, uncover, and make sense of a variety of dimensions along the broader AI lifecycle" (Wolf, 2021, p.15). In comparison, our focus is on the trained generative model as it concretely mediates in relation to a specific situated use. The term *reflexive use* should also not be confused with *reflective design* (Sengers et al., 2005) or *introspective AI* (Brand et al., 2021). The present concern is on the designers' and users' *reflection-in-action* (Schön, 1983) as they interact with DGMs, while reflective design concerns design processes much more broadly. Introspective AI concerns the support of personal introspection through models trained on personal data. Some of the examples presented by Brand et al. (2021) do support the reflexive use of the proposed model, in contrast, we are agnostic to the type of data the model is trained on.

E.2.2 Databases and machine learning in museums

The use of technology in museums has long been an active topic of research in HCI, to the extent that museums are seen as "a great testbed" for trialing novel interactive technologies (Hornecker & Ciolfi, 2019, p. xv). Technology is used for a range of different purposes in museums, including as a means of archiving digitized collections, as well as a means of communicating, educating, and facilitating experiences for visitors. The use of technology to digitize museum collections and their archives has been ongoing for a long time and is still a central effort for the museum sector.

However, more than just collecting and conserving cultural heritage, an important mission of museums is their ability to exhibit, communicate and involve the public in our shared cultural heritage. This requires experimenting with how technology can be used to design experiences that are simultaneously engaging, educational, and inspirational. One often cited challenge for museums is the fact that most museums have vast archives of artworks and artifacts that greatly exceed their capacity to exhibit; it is common to estimate that for European museums around 90% of their artifacts are permanently in storage and never exhibited to the public. Many museum professionals are eager to search for ways to use digital technology to make the digitized versions of these vast collections available to the public.

Many museums offer the public the ability to search their database through a conventional web-based interface with a text query and different options for filtering and sorting on pre-defined parameters. Earlier research has argued for the need for embodied visualization paradigms for large heterogeneous cultural datasets and has proposed immersive, interactive presentations to support navigating collections of thousands of cultural data objects (Kenderdine & McKenzie, 2013; Kenderdine et al., 2013; Lintermann, 2012; Shaw et al., 2006). Some projects create complex spatial visualizations of the data objects to highlight aspects of their individual relations (Alexander, 2014; Cole et al., 2019; Gordea & Vignoli, 2015). This approach is in stark contrast to the black-boxed generative models that collapse the data objects into a smooth latent space.

If digital collections should help museums achieve their goals of offering relevant experiences for their audiences, mediation tools are needed to support audiences in developing relevant perspectives. The recent technical developments in image synthesis with DGMs call for an investigation into how such models can participate in mediating relevant perspectives on the collections.

AI technologies have been deployed in museums in a variety of contexts, and the implications of these technologies for museums have caused much debate (Benford et al., 2022; Fontanella et al., 2020; Villaespesa & Crider, 2021; Villaespesa & Murphy, 2021). While these technologies inspire hope that they can contribute to making collections more searchable and accessible, there are also concerns that AI algorithms may perpetuate cultural biases and deal with sensitive issues in a problematic manner, as well as other legal and ethical concerns (Ciecko, 2020; Foka et al., "In press").

However, museums with large digital collections are in a good position to train their own bespoke models, avoiding many of the issues that muddle the ethical implications of the more generalized datasets, because they have ownership or rights of use for large amounts of data that they can correctly attribute. Furthermore, museums often employ domain experts, with deep knowledge about the subjects and historical context of the data objects. With a purposeful data curation practice, the museum is able to control what data goes into the model and potentially put it into play in new and exciting ways.

E.3 Method & Design Process

This project has been developed following how Zimmerman et al. (2007) describe of Research through Design. This implies that we present a design process that leads to an invention. We show how we find this invention relevant in addressing a particular situation in the art museum and evaluate our invention in its ability to bring us to a preferred new situation. Finally, we show how the learning from this project can be applied to other design research projects that involve DGMs as part of the design material.

The project has been developed in a number of concurrent trajectories. These will be elaborated below.

1. Concept development
2. Data collection
3. Model training
4. Table design
5. Evaluation

Throughout the design process, the design underwent informal evaluations to assess different aspects of the concept and the technical design. These evaluations were used to drive the design forward toward the intended qualities. Finally, the project underwent a summative evaluation with 20 participants who were interviewed about their experience with the system. This will be described in detail later.

E.3.1 Concept Development

The concept emerged from the idea of activating the digital collection of MUNCH, more specifically the paper-based works. MUNCH is in possession of a large number of sketches, notebooks, and diaries from Munch's hand. Due to the fragility of paper, these are particularly difficult to exhibit, as they are very sensitive to light. The amount of drawings in the paper collection is counted in the thousands, and might not be great works of art in themselves, but nevertheless an interesting entry point into the artistic practice of Munch.

Through an interview with the paper curators at the museum, the design team identified some important qualities of the paper collection. Firstly, Edvard Munch was very active in drawing and sketching the world around him. His many drawings range from early sketches of paintings to architectural sketches of his studio, satirical drawings of neighbors and their pets, drawings of everyday scenes in Norway, Norwegian nature, as well as many portraits and studies of models. This apparent interest and involvement with the world around him run counter to a myth that he was a hermit mainly producing somber paintings of with dark emotional content. Secondly, the curators spoke of an intimate physical relation to the paper sketches. Due to the fragility of the sketches, only very few people are allowed to handle the drawings and get the chance to develop this relation.

From these two points emerged the idea to let museum visitors explore the vast drawing collection through their own physical engagement with the drawings. Rather than physically handling the drawings, we would let the visitors explore them interactively through drawing. By leveraging the capability of computer vision and DGMs

we would let this play out on actual paper, with the synthetic drawings changing in response to the users drawing actions. This would invite visitors to engage in a practical dialogue with the drawing practice through a materially relevant interface.

E.3.2 Data Collection

Fig. E.1: The left half of the image shows samples from the original sketch data. The right half shows random synthetic samples from the StyleGAN model

In order to train a deep generative model we needed to collect and curate a dataset suitable for training. First, we queried the MUNCH digital archive for all images where the medium was listed as crayon, pencil, ink, or coal. This resulted in approximately 7600 images. The images in the MUNCH digital archive consist of photographs of the original notebooks and loose paper sheets where the drawings appear. This means that the images also contain table surfaces, paper edges, cataloging labels, torn paper, dirt, and text. To avoid training a model that would synthesize full notebooks, we decided to enrich the dataset with bounding boxes demarcating the location of each drawing. Furthermore, we distinguished between 4 different kinds of drawings:

- Line drawings with no shading
- Line drawings with some shading
- Heavily shaded drawings
- Colored drawings

The annotation was completed by members of the design team, including the first author as well as two student assistants. We used the bounding boxes to extract all marked drawings into separate image files. To maintain some stylistic consistency in the model we decided to use the drawings from the first two categories when training the model. These two categories cover approximately 5800 drawings.

In order to train the model all input images must have a 1:1 aspect ratio. This was achieved by stretching the images into shape. While this creates heavy distortion in some images, the effect after training is not very pronounced.

E.3.3 Model training

The architecture chosen was a StyleGAN 2 ADA (Karras, Laine, et al., 2020) model that was trained on the 5800 selected drawings. When training, the model develops a mathematical space called the *latent* space. This space follows the distribution of images in the dataset, with respect to their visual qualities. This latent space is smooth meaning that it is possible to interpolate between points in the latent space. For each step, the model will create an image coherent in its own right, while morphing gradually between the start and end points.

In order to be able to synthesize images quickly during runtime we have trained the model at 256x256 pixels. After training the StyleGAN 2 ADA model is capable of synthesizing images that adhere to the visual trends of the Munch's drawings in some aspects.

In order to allow for drawings to be used as input for the model, we trained a pixel2style2pixel (pSp) (Richardson et al., 2021) encoder capable of taking an input drawing and returning a synthetic drawing from the latent space of the StyleGAN model.

The pSp model was trained by using 10000 random synthetic drawings from the StyleGAN model. These 10000 drawings were then processed by a sketchification model, that simplifies the drawings into binary line drawings with a minimal amount of detailing. Then the pSp encoder was trained on the synthetic images and the simplified images in order to learn the mapping between binary input images and images from the latent space of the model.

E.3.4 Table design

The table design was developed to support the use of pen and paper as the input medium. Through testing, it was found that tracing paper and pigment markers provided the best tracking conditions. The table top surface is semi-transparent to allow for recording the paper from below. Two rows of infra-red LEDs light up the tracing paper from beneath, and a camera sensitive to the 850nm wavelength captures the drawn lines from behind a filter blocking visible light (see fig. E.2).

The video feed is then processed using TouchDesigner and OpenCV to create a binary input image for the image inference. The input image is then sent to the pSp model, which returns a synthetic image less than a second after. The generated image is then composited with a subtle overlay of the input image and projected back on the tracing paper (fig. E.3). Each time a new synthetic image is created the projected image fluidly fades from the previous to the new image. This happens continuously several times per second. The image processing is handled by a PC with a GTX 1070 GPU located inside the table.

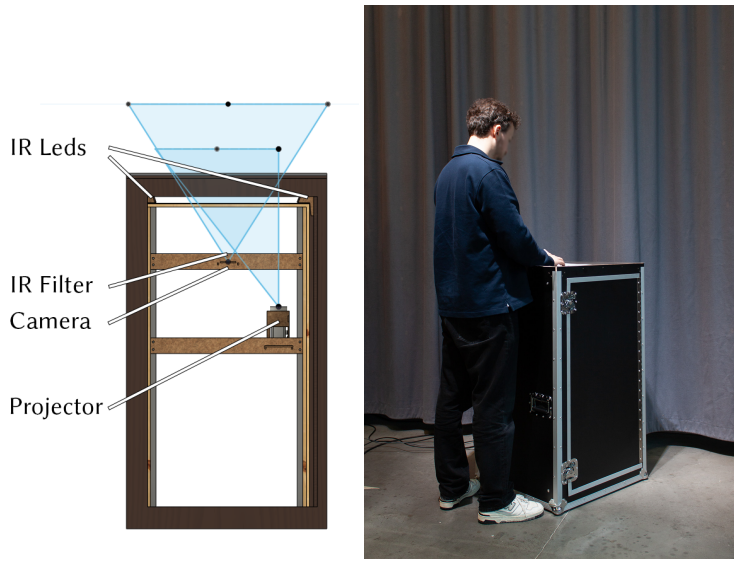

Fig. E.2: The prototype house a pico projector and a camera that are aligned with the drawing surface. An infrared filter removes the visible light from the camera input. Right below the top are two strips of infrared LED that illuminate the drawing surface. The lower part of the table can house a PC for processing the images.

The table also has a simple sound component. While the user draws a simple musical score is played by a speaker at low volume. The sound is generated live by a Max 8 patch, using three atmospheric digital synths. They each play a slow succession of notes of varying lengths. For each repetition, the relative timing of the notes between the synths shifts leading to a slow atmospheric melody that does not repeat itself.

Fig. E.3: Inferring a Munch-like sketch from a hand-drawn line. From left to right: 1) infrared webcam image from beneath the drawing surface. 2) The cleaned input image. 3) The synthesized sketch. 4) A composite image of the input image over the synthesized sketch that is output as a projection on the tracing paper.

E.3.5 Interaction

In order to use the system, the user picks up a piece of tracing paper and a pen and places it on the drawing surface. In the beginning, a faint pattern is visible on the tracing paper. As soon as the user starts drawing the system starts adapting the the generated image to the drawing (fig. E.4). The user can choose to draw on top of the lines presented by the system or place their own lines. In any case, the model continuously responds to the lines that are currently on the paper. In addition to drawing more lines, the user also has the option to move the paper around on the drawing surface. As the user moves or rotates the drawing the system interprets it differently and returns new results. This allows the user to investigate a certain "visual space" by slowly moving the paper to gradually see how the input drawing is interpreted differently in different areas of the drawing surface.

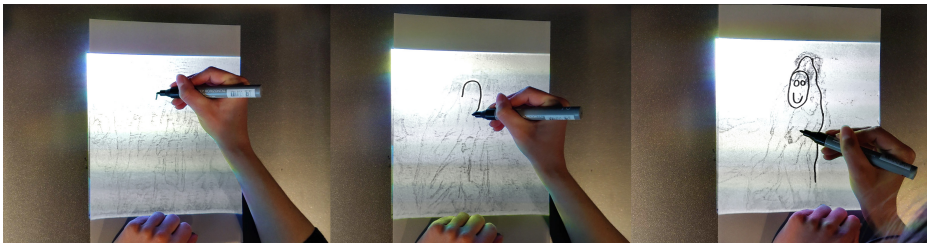

Fig. E.4: When starting out, the a faint pattern is visible on the tracing paper, but it quickly starts adapting to the participants' drawings. Due to the nature of the projector, the paper appears unevenly lit and with color bands in photographs. Through photo editing we have attempted to limit the effect as it is not visible to the naked eye.

E.4 Evaluation

The final design was tested on 20 participants recruited at the university during May and June 2023. Each participant would interact with the prototype for approximately 8-10 minutes and then participate in a phenomenological interview (Thompson et al., 1989) about their experience lasting approximately 12-15 minutes. Before interacting with the prototype, some questions were asked regarding demography as well as the participants' familiarity with drawing, machine learning, and art in general.

Of the 20 participants 12 identified as male, 7 as female, and 1 as non-binary. The average age was 32.2 with a maximum age of 60 and a minimum of 11. The participants were mainly students and faculty, with the exception of two participants who were attending elementary school. 4 participants were graduate students in a computer science program, 3 were graduate students in a design program and the remaining 11 were faculty of the design department.

When asked about their interest in art, 12 stated that they were interested in art. 6 were "partly" interested and 2 did not have an interest in art. Within the last 12 months, the participants had visited art museums and galleries 4.8 times on average (min. 0, max 12).

On a scale of 1-5, the participants rated how experienced they felt with regard to drawing at 2.6 on average and machine learning at 3.2. A few participants had machine learning as part of their research area, while none of the participants saw themselves as experts in drawing.

As discussed above, the *task at hand* is significant for shaping the *perspective* of the user, therefore all participants were given the same instructions before starting to draw:

> "This system is trained on the sketches of Edvard Munch. You should use the pen as your tool to explore what is hiding in the system. When you draw, the system will attempt to interpret your line as the beginning of an Edvard Munch drawing. The system only knows Edvard Munch's motives and way of drawing, so it will try to lead you into drawing like Edvard Munch. You can draw on the paper, you can move the paper around, and you can have as much paper as you want. To get off to a good start, I suggest that you start by drawing the beginning of a head or a face."

This instruction attempts to shape user interaction in a number of ways. It casts the pen as a tool for exploration rather than self-expression. It emphasizes a narrow focus on Edvard Munch, to tame expectations that the system would have the same capabilities as Stable Diffusion (Blattmann et al., 2022) or Dall-E (Ramesh et al., 2022). It indicates that the system has an agency to lead the user. Finally, it suggests that the participants start by drawing a face or head. Due to the tendency of the model to infer faces, this was said to make sure that the participants would quickly get into a dialogue with the model. The participants were allowed to ask clarifying questions while drawing and get as much paper as they wanted. The participants drew between 2 to 6 drawings each.

E.5 Results

In the following we will present some insights from the observations and interviews with participants, identifying three themes: The ways that participants engaged with the system through different drawing strategies, the aesthetics of the experience, and the ways in which the system led participants to perceive something about Edvard Munch's drawing style.

E.5.1 Drawing strategies

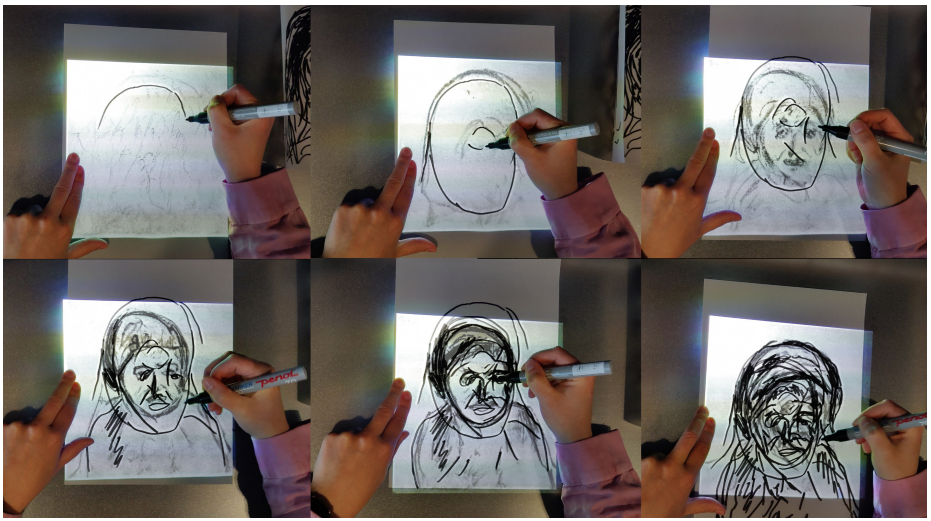

Fig. E.5: A few participants were mostly oriented towards creating coherent images on the paper. This participant sampled only "useful" lines from the projected drawings in order to end up with a drawing on the paper that was coherent in its own right.

Through the interviews, it became apparent that the drawing interface invited users to explore a range of approaches. We find that these approaches can be placed along two dimensions. The first dimension is whether the participant was proactive or reactive in relation to the system. Some participants would draw intuitively and expect the system to adapt to their input, while others would consider the current output first, before tracing or drawing in close relation to it. The second dimension is whether the participant expects meaningful drawings to emerge on the paper or in the projection. Some participants would consider the lines on paper the "final" result, while others saw the projection as the result. Over the course of their interactions, most participants changed their strategy multiple times, often starting out being reactive and gradually becoming more proactive as they became more familiar with the system. For some, this also meant becoming more interested in the system output rather than the physical drawing. While others insisted on the physical drawing being the key takeaway from the experience. Below we will highlight some notable strategies used by the participants.

As suggested by the instructions, participants would start by drawing a face or a head (see fig. E.4). This typically developed in two different directions. First, for some participants the first lines they drew made the system respond with a shape resembling a face. On seeing this, many participants switched to tracing the lines of the projected face, drawing in eyes, hair, mouth, or other features. Often this became

a collaboration between the participant and the system, negotiating which features to add. As the participant traced, the projected face might change in unexpected ways, rendering some of the earlier lines incoherent with the new image. To handle this, a few participants utilized a collage strategy where they traced only the lines that supported the creation of a coherent image on the paper. One participant even attempted to draw faster than the system could update the projected image in order to capture as many lines as possible before they changed (fig. E.5).

Second, for some participants, if their first attempt at drawing a face was not sufficiently well-aligned with the model it might result in a vague or ambiguous response from the model. This caused some participants to request a new paper and start over, while others would continue drawing and partially ignore the projections of the model.

After this first attempt, however, most participants developed some understanding of the model and changed the strategy for their next drawing. A few participants stayed with the idea of the system supporting their self-expression, leading them to try to derive what they could from the projected drawings to support them in creating good-looking drawings on the tracing paper. However, most participants seemed to put less emphasis on the physical drawings and focus more on exploring what they could make the projected drawings become. This varied between a *collaborating* strategy, where the participant switched back and forth between tracing over lines in the projection and adding new lines of their own; to a very deliberate *prompting* strategy where the participant only drew basic shapes in different sizes to explore what the system would make of it.

Since the task imagined by the design team was exactly this exploration of the model, it was positive to see the strategies converging toward this understanding as people became more familiar with the system. This change of strategy was, however, prompted by the tensions between the participants' intentions and the responses from the system.

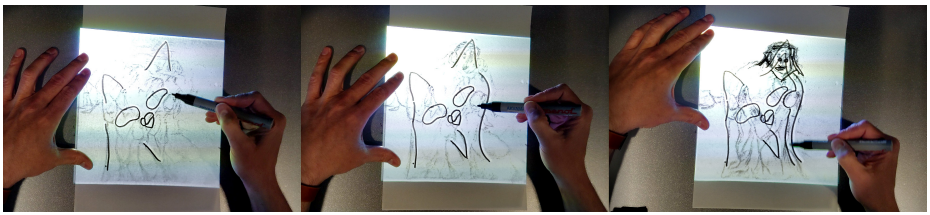

Fig. E.6: This participant starts by drawing an owl. However the system generates something that the participant interprets as a face, and changes plans with the drawing and adds a small face inside the ear of the owl

The system is only capable of generating images that are *within* Munch's practice, as defined by the model. This means that when a participant draws something the system will only change its output to the extent that it reflects lines that are salient in the model and their location in the drawing area. This meant that often those participants who tried to draw in their own visual style only received vague and ambiguous feedback from the model. However, sometimes the system reacted with something coherent in an unexpected way as P14 experienced (fig. E.6)

> "And so I was going after [drawing] an owl. I didn't get much guidance from the background [i.e. the projection] so I was just trying to create an owl by myself, and then all of a sudden this face appeared in the ear of the owl, which made me want to create something else."

Like P14, the participants would often change direction with their drawing when something exciting appeared in the projection. This also meant that they would sometimes be drawing on top of existing lines on their paper in an attempt to follow the whims of the model.

This type of exploration was further strengthened by the opportunity to move the paper around, as shifting the paper around on the drawing surface would cause the model to reinterpret the drawing, morphing between different types of faces or shifting into ambiguous shapes and lines. Many participants said that they had fun and found it pleasurable to move the paper around, exploring the different drawings the model could produce (fig. E.7):

> "I think it's a fun and playful interaction, me turning and moving the paper and then something new is being drawn. It was just great to see how it morphs from one painting to the next by me moving the paper. I definitely think you could have fun with this for some time." (P5)

However, those invested in creating coherent drawings on the paper found the continuous change a bit chaotic.

This ephemeral nature of the projected sketches and the constant reinterpretations pushed some participants to completely drop the idea of the physical sketches being important in themselves. P11 explains it like this:

> "I started realizing that the sketch is less of a representation of the thing you're trying to produce and more of a kind of fiducial or visual key to something you're looking for, so I started moving the sketches around rotationally or positionally to see if I could explore, given a single starting image, to see what might be out there"

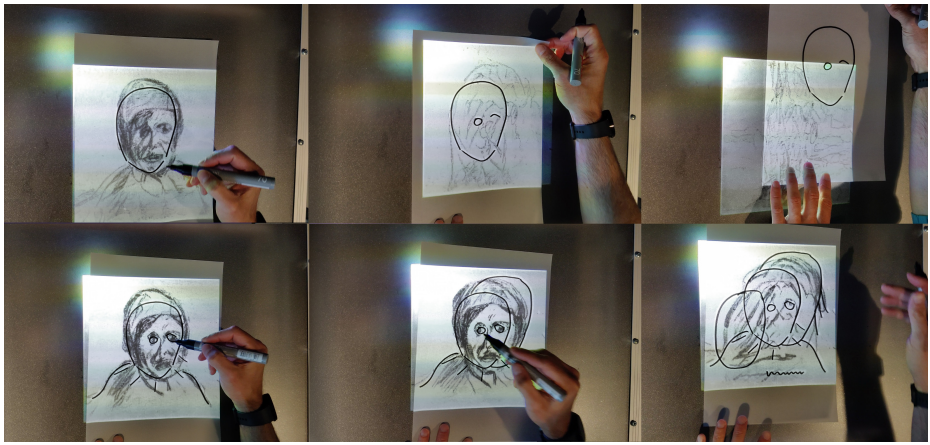

Fig. E.7: This participant quickly makes the model produce a face. Then the participants experiments with moving the drawing around but decides on the original position. Then the participant adds a line over the hair which makes the face adapts its shape. Next the participant adds a similar line on the other side of the face, which results in the face disappearing

E.5.2 Aesthetic Experience

During the drawing session, the soft ambient soundtrack of the drawing table would play. The participants described the music as something that helped in creating a cozy atmosphere and loosening up the feeling of having to perform while being observed drawing. Many also said that it was "calming", "relaxing" or "meditative" and helped them focus on the drawing experience and get into a flow. Several reported being absorbed by the drawing process. A couple of the participants said that it was not unlike music they could imagine hearing in a museum exhibition.

The physical tools also contributed to the aesthetic of the drawing experience. The participants reported that the marker and paper felt good. The marker produced a solid black line, and the paper felt of high quality. However, several participants mentioned that there was a large discrepancy between the types of lines they were able to produce with the marker, and the quality of the lines produced by the model which were more fuzzy, shaded, and thinner. Several participants expressed a wish to try other drawing tools that would allow them to get closer to the expression of the fuzzy lines and shading in the projections.

E.5.3 Munch's style as recreated by StyleGAN

Through the drawing engagement with the projected drawings, the participants became attuned to their aesthetic qualities. The participants described the subjects of the drawings, the compositions, the material quality of the lines, the facial expres-

sions of the faces, and other emotional qualities of the drawings. Some participants described how they understood the aesthetic qualities as a totality: "It's interesting because you enter a universe of these drawings" (P12). Another described it:

> "It was a very interesting way to experience the art, instead of the very static image in the museum, where you can also look at the lines and collect from the different images. It gave a different feeling of the artist when you got to interact with it yourself. There were a lot of things changing but you could still clearly see that it was from the same artists in the same style" (P6)

Participants had different perceptions of the motifs they could recognize in the projections from the system. Some felt that they could identify several motifs and styles: "To me it seems that there are maybe two, three, or four genres. In my head, there are now the turbaned faces, maybe people standing full-body with flowing robes or dresses or cloaks, hats and ponytails, and then perhaps landscapes. I don't know whether Munch enjoyed drawing craggy cliff sides" (P11). No other participants talked about turbans and ponytails, but faces, standing figures and nature such as mountains, cliffs, and trees were common across participants.

Due to the instructions and the model's affinity for faces (fig. E.3, E.5, E.7), almost all participants talked about these, and also in greater detail: "He has this kind of like head where there's usually a hat or something. There's a line connected from the eyes with the nose. One line. It [also] seems like there's a lot of lines, but it's never clean lines." (P13). P19 said about the style of the faces that they were drawn with "maybe not a simple line, but a characteristic line in all the drawings, that had this dark melancholic facial expression. People looking maybe a bit anxious or sad and who had many details, while not having a lot of details". P9 noticed that "the faces had a pained look. It was very dark [...] it was a very strained line in a way. [...] a dark space visually." In these comments we see that the participants noticed particular aesthetic qualities in the lines, and some even experienced a distinct emotional quality in the drawings.

Despite the ephemeral nature of the drawings and their constant instability and tendency to become abstract and ambiguous, we see examples of participants experiencing aesthetic qualities in the drawing technique, the typical subjects, and the drawings' emotional quality. In some cases, however, the ambiguity also led the participants to see things that might not be part of the dataset. While reported by one of the participants, ponytails are not commonly occurring in the experience of the designers. The craggy cliff sides, that several participants refer to, seem to be related to a specific texture that might be an artifact of the training process more than specific drawings (see fig. E.4). Simultaneously, through observing the participants, the

first author also noted that many types of drawings did not emerge during the participants' interaction even though we know that they are part of the dataset. This would for example be drawings of animals, interiors, and scenes with multiple persons.

E.6 Discussion

E.6.1 Sketches as uncertain entities

The New Snow drawing table embodies a statistical representation of Edvard Munch's drawings and sketches. While the dataset is made up of distinct originals, the smooth nature of the latent space of the StyleGAN model blends the individual motifs and allows for seamless interpolation between them. The model does not recreate individual drawings from the dataset precisely, even though it might sometimes get close. Each generated frame is an *uncertain entity*. Each synthetic image is fleeting and ephemeral and never solidifies due to the slight noise in the camera feed.

This is a quite different view of art history than a typical museum installation would offer: Normally museums present art history by exhibiting individual works that are either deemed particularly interesting in terms of their unique qualities. Works may also be exhibited as representative of a broader tendency in the work of an artist, a particular style, or period, and sometimes a broad selection of works may be presented together to explore such tendencies - but even then the number of works that can be presented at one time to a visitor is far from the 5800 drawings in the dataset of the DGM model in New Snow. As such, we anticipate that this approach toward presenting a large body of artworks might be greeted with some controversy among art historians and curators.

The system relies on a statistical approximation of patterns in Munch's drawings and it is important that this quality is communicated to the users in order to avoid the system being seen as an authoritative representation of Munch's art - as one might expect to meet in a museum. Fortunately, what happens is that the drawings are ontologically *revealed* as ephemeral, fluid, and malleable, and through the instructions, as being from Edvard Munch's practice. On the other hand, the drawings as discrete physical objects are *concealed*. Through the somaesthetic drawing relation, the *drawn* quality of the synthetic sketches is *magnified*. So are the dynamics and movement of the lines, as interpreted by the StyleGAN model. The paper quality, relative scale, and fine details in the drawings are *reduced*. The particular *perspective* on the drawings is not arbitrary, but a specific *functional perspective* which is the only way we are able to know anything at all. In this sense, this manifestation alludes to the intangibility of art as a practice, rather than the drawings as discrete objects. This is the mediation we see reflected in the interviews. The participants speak about the drawings in multiples and do not single in on them as discrete or authoritative objects.

What we do see, is the hermeneutic and alterity relations in play. The dialogical interaction some participants experience and the way the drawings are read and appropriated are very different from how one usually looks at works of art in a museum context. As earlier HCI research has suggested, the ambiguity (Gaver et al., 2003) and openness to interpretation (Sengers & Gaver, 2006) may be helping the participants in critically reflecting on and appropriating the system. Nevertheless, there is space for retraining the StyleGAN model and the pSp encoder to adjust how exactly this practice manifests. The model has a tendency to infer faces from many input drawings, while it can be difficult to prompt other subjects like interior scenes and nature that also exist in the dataset. These qualities of the model are what is also revealed through the participants' interactions with it.

E.6.2 Art mediation and education of attention

A person trained in reading, drawing, or looking at art will approach the New Snow installation with different perceptual capabilities than the novice, and naturally, these capabilities can change over time. Ingold describes this as *education of attention* (Ingold, 2001). That is to be understood in contrast to a prevailing idea of learning as the transmission of information. Noë talks about a similar relation to art in his book "Learning to Look" where he is furthering the point that we need examples in the form of pictures, text, theories, physical instruction, etc. that help us understand where to turn our attention and what to see (Noë, 2021). Noë presents an example of repair manuals for cars. One car came with a manual with photographs of the car's internals while another manual for another car used line drawings. He argues that the photographs did not manage to pick out what was important, while the drawings were more articulate, bringing your intention to what matters, for the particular purpose of a repair manual(Noë, 2021, p.65). While this example is based on visual attention, the concept for both Ingold and Noë goes beyond the visual domain and involves all our perceptual capabilities. In this enactivist perspective, learning means becoming attentive to particular features of the environment that are important for solving a given task. Sivertsen et al. (2023) argue that when using technology for mediation purposes in the art museum, this may constitute an "art critique by other means". Through this lens, art mediation and technological mediation become two sides of the same coin. The purpose of the art critique in this view is to draw the audience into correspondence with the art, rather than transmitting information about it. The art critique does not depend on the original artworks being present, even though engagement with the original that the critique concerns add to the ongoing correspondence.

The design of the New Snow concept was shaped by a discussion with curators about the Munch's drawing practice. Notions of an intimate material relation to the paper sketches and Munch's eclectic practice led to the material drawing interface

affording the exploration of his expansive and varied practice. It is through this interface that the paintings are revealed as practice and the participants' attention is educated towards the quality of the drawings, as seen in the interviews. However, we also see examples of visual artifacts resulting from the training process, as well as subjects that do not appear in the synthetic drawings. Through the interviews, we as designers, learn about how the patterns in the model shape the mediation of the artwork and must take appropriate measures in data curation, pre-processing, and retraining to shape and readjust the model to support the design intention.

E.6.3 Data curation for bespoke generative models

To create a model that is interesting and warrants the kind of exploration described above requires that attention is paid to its constitution and ontology. The data is the matter from which the model is built and from which it derives its form, and therefore also important for how the system, in turn, comes to mediate the world. This presents a design task outside the typical scope of most designers' jobs.

Making a bespoke model is a labor-intensive process. Gathering data and pre-processing it are time-consuming tasks that are hard to evaluate the success of before the first model is trained: Only when the designer can interact with the trained model can they get a sense of whether the model can facilitate a user experience similar to the design vision. This calls for an iterative process in which the data collection and model training is reiterated several times until the desired result is achieved - however, the amount of time and labor it takes to adjust a dataset with thousands - or potentially millions - of entries and retraining through numerous iterations makes the cost of iterating very high. If there is a problem with the system in early iterations, the large scale of the dataset makes it difficult to get a full understanding of what creates a problem.

Training a bespoke model also requires that a certain amount of data is available. Even with a very productive artist like Munch, collecting a sufficient amount of images in the dataset required that we had to accept a certain level of stylistic divergence. We made a selection to include only specific mediums and techniques, but we included a certain variety in the drawing style due to the fact that Munch's style changed over the course of his life and across different types of subjects.

E.6.4 Designing for reflexive use of AI

In the following, we will discuss three aspects of the New Snow interface that make it well suited for supporting *reflexive use*, and how these features may serve as preliminary heuristics when designing for it.

Prompting modality

The most important aspect is naturally choosing relevant means of prompting the model to generate an output. The pSp encoder enables the mapping from black and white input images to the latent space of the StyleGAN model. This enables an interrogation of the model through the flow of lines, composition, drawing density, location, and scale. The pSp encoder also supports other kinds of image-to-image translation, such as inpainting and generation from segmentation maps. While still in the image-to-image domain, these translations enable very different relations to the images. This is again very different from a system prompted by a text interface, as this would make it much more difficult to express compositions and the quality of the lines, while making it easier to prompt verbal concepts, such as tree, face, and mountain. In short, how the particular encoder enables navigation in the latent space, is defining for the dimensions in which the model can be interrogated.

Incremental prompt adjustment

The marker and paper interface enables continuous and slight adjustments to the drawings. The loose paper lets the user slide the drawing around to gradually explore the latent space along two dimensions, while the marker lets the user gradually add to the drawing in response to the system. This gradual change makes it possible to uncover the internal relations in the system between the dimensions that the interface affords. The biggest limitation in the New Snow interface in this regard, is that the nature of the pigment marker constrains these adjustments to be additive, as the user must start over on a new paper if they want to remove a line from their drawing.

Fast updates

For New Snow, we intentionally picked the StyleGAN architecture for its relative speed in synthesizing images. We also run the model at 256x256 pixels, for the same reason. Using the system reflexively is supported by its updating quickly, as it allows the user to explore more without being held back by long waiting times. Different architectures have vastly different response times but for New Snow, we found that the fast response of the system supported the pleasurable and fluent interaction with the system.

E.7 Conclusion

In this paper, we have presented a design supporting *reflexive use* of a deep generative model. Through the drawing interface, the system offers the user a *perspective* on Munch's sketching practice that is unique to the capabilities of a bespoke model.

By paying specific attention to the technological mediation of the system we can design for and evaluate how the model magnifies and reduces aspects of the underlying dataset. The New Snow system presented magnifies the visual trends and aesthetic qualities of his drawings as a practice, while reducing their perceptual presence as objects. The technological mediation becomes in this case a form of art mediation that can *educate the attention* of the user to specific aspects of interest in Munch's art.

We expect that this approach can be applied to the practice of other productive visual artists, but also beyond the art museum. Through design for *reflexive use* models trained on other visual datasets can be explored not as a collection of discrete data objects, but as trends and patterns. We have discussed how we understand the synthetic drawings as *uncertain entities* that offer a *functional perspective* on the drawings. Finally, we have presented three aspects, *prompting modality*, *incremental prompt adjustment*, and *fast updates* that we found important for affording exploration of the aesthetic qualities of the model.

Acknowledgments

The authors would like to thank MUNCH, Oslo for helping in developing the concept and funding the prototype. We are thankful to the research division at Random International for the development of the first iteration of the tracking system and AIRLab at IT University of Copenhagen for supporting the physical realization of the prototype. Finally a special thanks to René Haas for help with developing the machine learning pipeline.

References

Akten, M. (2017). Learning to See: Interactive. https://www.memo.tv/works/learning-to-see/

Akten, M. (2021). *Deep Visual Instruments: Realtime Continuous, Meaningful Human Control over Deep Neural Networks for Creative Expression* (doctoral). Goldsmiths, University of London. Retrieved September 9, 2022, from https://research.gold.ac.uk/id/eprint/30191/

Akten, M., Fiebrink, R., & Grierson, M. (2019). Learning to see: You are what you see. *ACM SIGGRAPH 2019 Art Gallery*, 1–6. https://doi.org/10.1145/3306211.3320143

Alexander, J. (2014). Gallery One at the Cleveland Museum of Art. *Curator: The Museum Journal*, *57*(3), 347–362. https://doi.org/https://doi.org/10.1111/cura.12073

Appel, G., Neelbauer, J., & Schweidel, D. A. (2023). Generative AI Has an Intellectual Property Problem [Section: Intellectual property]. *Harvard Business Review*. Retrieved June 28, 2023, from https://hbr.org/2023/04/generative-ai-has-an-intellectual-property-problem

Ariccia, A. D., Bremers, A., Michalove, J., & Ju, W. (2022). How to Make People Think You're Thinking if You're a Drawing Robot: Expressing Emotions Through the Motions of Writing. *2022 17th ACM/IEEE International Conference on Human-Robot Interaction (HRI)*, 1190–1191. https://doi.org/10.1109/HRI53351.2022.9889638

Atwood, J., Halpern, Y., Baljekar, P., Breck, E., Sculley, D., Ostyakov, P., Nikolenko, S. I., Ivanov, I., Solovyev, R., Wang, W., & Skalic, M. (2020). The Inclusive Images Competition. In S. Escalera & R. Herbrich (Eds.), *The NeurIPS '18 Competition* (pp. 155–186). Springer International Publishing. https://doi.org/10.1007/978-3-030-29135-8_6

Audry, S. (2021). *Art in the age of machine learning*. The MIT Press.

Barredo Arrieta, A., Díaz-Rodríguez, N., Del Ser, J., Bennetot, A., Tabik, S., Barbado, A., Garcia, S., Gil-Lopez, S., Molina, D., Benjamins, R., Chatila, R., & Herrera, F. (2020). Explainable Artificial Intelligence (XAI): Concepts, taxonomies, opportunities and challenges toward responsible AI. *Information Fusion*, *58*, 82–115. https://doi.org/10.1016/j.inffus.2019.12.012

Benford, S., Løvlie, A. S., Ryding, K., Rajkowska, P., Bodiaj, E., Paris Darzentas, D., Cameron, H., Spence, J., Egede, J., & Spanjevic, B. (2022). Sensitive Pictures: Emotional Interpretation in the Museum. *CHI Conference on Human Factors in Computing Systems*, 1–16. https://doi.org/10.1145/3491102.3502080

Benjamin, J. J., Berger, A., Merrill, N., & Pierce, J. (2021). Machine Learning Uncertainty as a Design Material: A Post-Phenomenological Inquiry. *Proceedings of the 2021 CHI Conference on Human Factors in Computing Systems*, 1–14. https://doi.org/10.1145/3411764.3445481

Benjamin, J. J., Biggs, H., Berger, A., Rukanskaitė, J., Heidt, M. B., Merrill, N., Pierce, J., & Lindley, J. (2023). The Entoptic Field Camera as Metaphor-Driven Research-through-Design with AI Technologies. *Proceedings of the 2023 CHI Conference on Human Factors in Computing Systems*, 1–19. https://doi.org/10.1145/3544548.3581175

Birhane, A., Prabhu, V. U., & Kahembwe, E. (2021). Multimodal datasets: Misogyny, pornography, and malignant stereotypes [Publisher: arXiv Version Number: 1]. https://doi.org/10.48550/ARXIV.2110.01963

Blattmann, A., Rombach, R., Oktay, K., Müller, J., & Ommer, B. (2022). Semi-Parametric Neural Image Synthesis [Publisher: arXiv Version Number: 3]. https://doi.org/10.48550/ARXIV.2204.11824

Boden, M. A., & Edmonds, E. A. (2019). 2 a taxonomy of computer art. In *From fingers to digits: An artificial aesthetic* (pp. 23–59).

Brand, N., Odom, W., & Barnett, S. (2021). A Design Inquiry into Introspective AI: Surfacing Opportunities, Issues, and Paradoxes. *Designing Interactive Systems Conference 2021*, 1603–1618. https://doi.org/10.1145/3461778.3462000

Cetinic, E., & She, J. (2022). Understanding and creating art with ai: Review and outlook. *ACM Trans. Multimedia Comput. Commun. Appl.*, *18*(2). https://doi.org/10.1145/3475799

Chen, M. (2023). Artists and Illustrators Are Suing Three A.I. Art Generators for Scraping and 'Collaging' Their Work Without Consent [Section: Law]. *Artnet News*. Retrieved June 28, 2023, from https://news.artnet.com/art-world/class-action-lawsuit-ai-generators-deviantart-midjourney-stable-diffusion-2246770

Ciecko, B. (2020). AI sees what? The good, the bad, and the ugly of machine vision for museum collections. *The Museum Review*, *5*(1). Retrieved February 23, 2023, from https://themuseumreviewjournal.wordpress.com/2020/04/23/tmr_vol5no1_ciecko/

Cole, R. J., Dau, F., Ducrou, J., Eklund, P. W., & Wray, T. (2019). Navigating Context, Pathways and Relationships in Museum Collections using Formal Concept Analysis. *International Journal for Digital Art History*, (4), *5*.13–*5*.27. https://doi.org/10.11588/dah.2019.4.72070

Crawford, K., & Paglen, T. (2021). Excavating AI: The politics of images in machine learning training sets. *AI & SOCIETY*, *36*(4), 1105–1116. https://doi.org/10.1007/s00146-021-01162-8

Davis, N., Hsiao, C.-P., Singh, K. Y., Li, L., Moningi, S., & Magerko, B. (2015). Drawing apprentice: An enactive co-creative agent for artistic collaboration. *Proceedings of the 2015 ACM SIGCHI Conference on Creativity and Cognition*, 185–186. https://doi.org/10.1145/2757226.2764555

de Lima, E. S., Feijó, B., Barbosa, S. D., Furtado, A. L., Ciarlini, A. E., & Pozzer, C. T. (2014). Draw your own story: Paper and pencil interactive storytelling. *Entertainment Computing*, *5*(1), 33–41. https://doi.org/10.1016/j.entcom.2013.06.004

Denton, E., Hanna, A., Amironesei, R., Smart, A., & Nicole, H. (2021). On the genealogy of machine learning datasets: A critical history of ImageNet. *Big Data & Society*, *8*(2), 205395172110359. https://doi.org/10.1177/20539517211035955

Doshi-Velez, F., & Kim, B. (2017). Towards A Rigorous Science of Interpretable Machine Learning [arXiv:1702.08608 [cs, stat]]. Retrieved December 21, 2022, from http://arxiv.org/abs/1702.08608

Foka, A., Eklund, L., Løvlie, A. S., & Griffin, G. ("In press"). Critically Assessing AI/ML for Cultural Heritage: Potentials and Challenges. In S. Lindgren (Ed.), *Handbook of Critical Studies of Artificial Intelligence*. Edward Elgar.

Fontanella, F., Colace, F., Molinara, M., Freca, A. S. D., & Stanco, F. (2020). Pattern recognition and artificial intelligence techniques for cultural heritage. *Pattern Recognition Letters*, *138*, 23–29. https://doi.org/https://doi.org/10.1016/j.patrec.2020.06.018

Gaver, W. W., Beaver, J., & Benford, S. (2003). Ambiguity as a resource for design. *Proceedings of the SIGCHI Conference on Human Factors in Computing Systems*, 233–240. https://doi.org/10.1145/642611.642653

Ghosh, A., & Kandasamy, D. (2020). Interpretable Artificial Intelligence: Why and When [Publisher: American Roentgen Ray Society]. *American Journal of Roentgenology*, *214*(5), 1137–1138. https://doi.org/10.2214/AJR.19.22145

Google Creative Lab, Girschig, B., & Cazier, R. (2020). Draw to Art: Shape Edition. Retrieved December 2, 2022, from https://experiments.withgoogle.com/draw-to-art-shape

Google Creative Lab, Google Art & Culture Lab, & IYOIYO. (2018). Draw to Art. Retrieved December 2, 2022, from https://experiments.withgoogle.com/draw-to-art

Gordea, S., & Vignoli, M. (2015). Culturecam: Visual exploration of cultural heritage content by professional designers. *2015 IEEE International Conference on Multimedia & Expo Workshops (ICMEW)*, 1–6. https://doi.org/10.1109/ICMEW.2015.7169797

Grba, D. (2022). Deep else: A critical framework for ai art. *Digital*, *2*(1), 1–32. https://doi.org/10.3390/digital2010001

Guljajeva, V., & Canet Sola, M. (2022). POSTcard Landscapes from Lanzarote. *Creativity and Cognition*, 634–636. https://doi.org/10.1145/3527927.3531191

Hois, J., Theofanou-Fuelbier, D., & Junk, A. J. (2019). How to Achieve Explainability and Transparency in Human AI Interaction. In C. Stephanidis (Ed.), *HCI International 2019 - Posters* (pp. 177–183). Springer International Publishing. https://doi.org/10.1007/978-3-030-23528-4_25

Hornecker, E., & Ciolfi, L. (2019). Human-Computer Interactions in Museums. *Synthesis Lectures on Human-Centered Informatics*, *12*(2), i–153. https://doi.org/10.2200/S00901ED1V01Y201902HCI042

Ingold, T. (2001). From the transmission of representation to the education of attention. In *The debated mind: Evolutionary psychology versus ethnography* (pp. 113–153). Berg.

Karras, T., Aittala, M., Hellsten, J., Laine, S., Lehtinen, J., & Aila, T. (2020). Training Generative Adversarial Networks with Limited Data. *Proc. NeurIPS*.

Karras, T., Aittala, M., Laine, S., Härkönen, E., Hellsten, J., Lehtinen, J., & Aila, T. (2021). Alias-Free Generative Adversarial Networks. *Proc. NeurIPS*.

Karras, T., Laine, S., Aittala, M., Hellsten, J., Lehtinen, J., & Aila, T. (2020). Analyzing and improving the image quality of StyleGAN.

Kenderdine, S., & McKenzie, H. (2013). A war torn memory palace: Animating narratives of remembrance. *2013 Digital Heritage International Congress (DigitalHeritage)*, 315–322. https://doi.org/10.1109/DigitalHeritage.2013.6743755

Kenderdine, S., Shaw, J., & Gremmler, T. (2013). Cultural Data Sculpting: Omnidirectional Visualization for Cultural Datasets. In F. T. Marchese & E. Banissi (Eds.), *Knowledge Visualization Currents* (pp. 199–220). Springer London. https://doi.org/10.1007/978-1-4471-4303-1_11

Kiran, A. H. (2015). Four Dimensions of Technological Mediation. In R. Rosenberger & P.-P. Verbeek (Eds.), *Postphenomenological Investigations: Essays on Human–Technology Relations*. Lexington Books.

LAIKA. (2023). Write with LAIKA - Personalised Artificial Intelligence for Writers. Retrieved June 9, 2023, from https://www.writewithlaika.com/

Larrazabal, A. J., Nieto, N., Peterson, V., Milone, D. H., & Ferrante, E. (2020). Gender imbalance in medical imaging datasets produces biased classifiers for computer-aided diagnosis [Publisher: Proceedings of the National Academy of Sciences]. *Proceedings of the National Academy of Sciences*, *117*(23), 12592–12594. https://doi.org/10.1073/pnas.1919012117

Leahu, L. (2016). Ontological Surprises: A Relational Perspective on Machine Learning. *Proceedings of the 2016 ACM Conference on Designing Interactive Systems*, 182–186. https://doi.org/10.1145/2901790.2901840

Lintermann, B. (2012). Beyond Cinema. Retrieved December 1, 2022, from https://www.bernd-lintermann.de/papers/Beyond_Cinema_Lintermann.pdf

Malevé, N. (2021). On the data set's ruins. *AI & SOCIETY*, *36*(4), 1117–1131. https://doi.org/10.1007/s00146-020-01093-w

Noë, A. (2021). *Learning to Look: Dispatches from the Art World*. Oxford University Press.

Oppenlaender, J. (2022). A Taxonomy of Prompt Modifiers for Text-To-Image Generation [arXiv:2204.13988 [cs]]. Retrieved January 4, 2023, from http://arxiv.org/abs/2204.13988

Oppenlaender, J., Linder, R., & Silvennoinen, J. (2023). Prompting ai art: An investigation into the creative skill of prompt engineering.

Paglen, T., & Crawford, K. (2019). ImageNet Roulette. Retrieved October 14, 2023, from https://paglen.studio/2020/04/29/imagenet-roulette/

Ramesh, A., Dhariwal, P., Nichol, A., Chu, C., & Chen, M. (2022). Hierarchical Text-Conditional Image Generation with CLIP Latents [Publisher: arXiv Version Number: 1]. https://doi.org/10.48550/ARXIV.2204.06125

Richardson, E., Alaluf, Y., Patashnik, O., Nitzan, Y., Azar, Y., Shapiro, S., & Cohen-Or, D. (2021). Encoding in Style: A StyleGAN Encoder for Image-to-Image Translation. *IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*.

Rosenberger, R., & Verbeek, P.-P. (2015). A Field Guide to Post-Phenomenology. In R. Rosenberger & P.-P. Verbeek (Eds.), *Postphenomenological Investigations: Essays on Human-Technology Relations* (pp. 9–41). Lexington Books.

Schön, D. A. (1983). *The reflective practitioner: How professionals think in action*. Basic Books.

Scurto, H., Caramiaux, B., & Bevilacqua, F. (2021). Prototyping Machine Learning Through Diffractive Art Practice. *Designing Interactive Systems Conference 2021*, 2013–2025. https://doi.org/10.1145/3461778.3462163

Sengers, P., Boehner, K., David, S., & Kaye, J. '. (2005). Reflective design. *Proceedings of the 4th decennial conference on Critical computing: between sense and sensibility*, 49–58. https://doi.org/10.1145/1094562.1094569

Sengers, P., & Gaver, B. (2006). Staying open to interpretation: Engaging multiple meanings in design and evaluation. *Proceedings of the 6th conference on Designing Interactive systems*, 99–108. https://doi.org/10.1145/1142405.1142422

Shaw, J., Brown, N., Del Favero, D., McGinity, M., & Weibel, P. (2006). T_visionarium II. Retrieved December 1, 2022, from https://www.jeffreyshawcompendium.com/portfolio/t_visionarium-ii/

Sivertsen, C., Smith, M., & van der Zwan, S. (2023). Art Critique by Other Means. *Designing Interactive Systems Conference (DIS '23)*, 10. https://doi.org/10.1145/3563657.3596069

Thompson, C. J., Locander, W. B., & Pollio, H. R. (1989). Putting Consumer Experience Back into Consumer Research: The Philosophy and Method of Existential-Phenomenology. *Journal of Consumer Research*, *16*(2), 133–146. https://doi.org/10.1086/209203

Villaespesa, E., & Crider, S. (2021). Computer vision tagging the metropolitan museum of art's collection: A comparison of three systems. *J. Comput. Cult. Herit.*, *14*(3). https://doi.org/10.1145/3446621

Villaespesa, E., & Murphy, O. (2021). This is not an apple! Benefits and challenges of applying computer vision to museum collections [Publisher: Routledge _eprint: https://doi.org/10.1080/09647775.2021.1873827]. *Museum Management and Curatorship*, *36*(4), 362–383. https://doi.org/10.1080/09647775.2021.1873827

Wang, A., Liu, A., Zhang, R., Kleiman, A., Kim, L., Zhao, D., Shirai, I., Narayanan, A., & Russakovsky, O. (2022). REVISE: A Tool for Measuring and Mitigating Bias in Visual Datasets. *International Journal of Computer Vision*, *130*(7), 1790–1810. https://doi.org/10.1007/s11263-022-01625-5

Wolf, C. T. (2021). Towards "Explorable" AI: Learning from ML Developers' Sensemaking Practices [Publisher: European Society for Socially Embedded Technologies (EUSSET)]. https://doi.org/10.18420/ECSCW2021_N28

Yang, Q., Steinfeld, A., Rosé, C., & Zimmerman, J. (2020). Re-examining Whether, Why, and How Human-AI Interaction Is Uniquely Difficult to Design. *Proceedings of the 2020 CHI Conference on Human Factors in Computing Systems*, 1–13. https://doi.org/10.1145/3313831.3376301

Yurman, P., & Reddy, A. V. (2022). Drawing conversations mediated by ai. *Proceedings of the 14th Conference on Creativity and Cognition*, 56–70. https://doi.org/10.1145/3527927.3531448

Zhu, J., Liapis, A., Risi, S., Bidarra, R., & Youngblood, G. M. (2018). Explainable AI for Designers: A Human-Centered Perspective on Mixed-Initiative Co-Creation. *2018 IEEE Conference on Computational Intelligence and Games (CIG)*, 1–8. https://doi.org/10.1109/CIG.2018.8490433

Zimmerman, J., Forlizzi, J., & Evenson, S. (2007). Research through design as a method for interaction design research in HCI. *Proceedings of the SIGCHI Conference on Human Factors in Computing Systems*, 493–502. https://doi.org/10.1145/1240624.1240704

Zylinska, J. (2020). *AI Art: Machine Visions and Warped Dreams*. Open Humanites Press. Retrieved September 10, 2022, from http://www.openhumanitiespress.org/books/titles/ai-art/

Publication F

Machine Learning Processes as Sources of Ambiguity: Insights from AI Art

Christian Sivertsen, Guido Salimbeni, Anders Sundnes Løvlie, Steve Benford and Jichen Zhu

A revised version of this paper has been accepted for
Proceedings of the 2024 CHI Conference on Human Factors in Computing Systems,
https://doi.org/10.1145/3613904.3642855

The layout has been revised.

Abstract

Ongoing efforts to turn Machine Learning (ML) into a design material have encountered limited success. This paper turns its attention to the burgeoning area of AI art to understand how artists incorporate ML in their creative work. Drawing upon related HCI theories, we investigate how artists create ambiguity by analyzing nine AI artworks that use computer vision and image synthesis. Our analysis indicates that artists use the ML processes *as new* sources of ambiguity. *It identifies their specific techniques in the ML process (dataset curation, model training, and application). Our findings suggest that, in addition to the ML artifacts, HCI theories of ambiguity should be broadened to encompass how users interpret the ML computational process and account for how ML uncertainty complicates users' interpretative process. Finally, this paper offers design implications for alternative approaches to thinking about ML products and experiences.*

F.1 Introduction

To meet the demands raised by new Machine Learning (ML) products, the Human-Computer Interaction (HCI) community is making ongoing efforts to turn ML into a *design material* (Holmquist, 2017). Despite the intense interest, a growing body of research shows that ML is uniquely difficult to design with (J. J. Benjamin et al., 2021; Dove et al., 2017; Kuniavsky et al., 2017; Yang et al., 2020). In a recent study with industry UX designers, researchers identified two key difficulties of ML as a design material: *capability uncertainty* (uncertainties surrounding what the system can do and how well it performs) and *output complexity* (complexity of the outputs that the system might generate, e.g., in adaptive systems) (Yang et al., 2020). Due to these ML-specific difficulties, designers face obstacles in all design phases, from conceptualization to prototyping.

As a result, with a few exceptions in highly technical design sub-domains such as intelligent user interfaces (Höök, 2000; Maybury & Wahlster, 1998) and computer game design (Shaker et al., 2016; Yannakakis & Togelius, 2018; Zhu et al., 2021), radical design innovations of ML products remain rare (Dove et al., 2017). Dove et al. recognized that "[with ML,] it is no longer enough for UX designers to only improve user experience by paying attention to usability, utility, and interaction aesthetics" (Dove et al., 2017, p.278), as indicated by traditional UX design principles. However, it is not clear how designers should adapt when working with ML. There hence is a pressing need for new ways of conceptualizing ML as a design material and for novel design approaches.

In this paper, we follow the precedents in interaction design and seek inspiration from art. HCI researchers have found that engaging art and art history can open up new generative ideas for HCI theory and practice (e.g., (Bardzell & Bardzell, 2015;

Benford & Giannachi, 2011; Gaver et al., 2003)), especially in domains traditionally dominated by discourses of engineering and productivity (e.g., digital fabrication (Devendorf & Rosner, 2015; Song & Paulos, 2021), electronics (Kang et al., 2022) and Machine Learning (Audry, 2021; Caramiaux & Fdili Alaoui, 2022; Scurto et al., 2021)).

This paper turns to "AI art," an emerging umbrella term that describes the variety of artistic practices that use AI, including ML, to create aesthetic experiences (Boden & Edmonds, 2019; Zylinska, 2020). Among AI art, we focus on a particularly active area: *Visual* artworks that are built on *computer vision* — the technology that allows computers to make sense of images — and *image synthesis* — the technology that allows computers to create images from some form of description. Following recent breakthroughs in high-quality image synthesis, there is a surge of AI art experiments with this technology (Audry, 2021; Bogost, 2019; Cetinic & She, 2022; Oppenlaender, 2022; Zylinska, 2020). For brevity, we will refer to the two types of artworks as AI art hereafter unless specified otherwise.

Our particular focus is to investigate *how artists create ambiguity with ML*. In this paper, we adopt Gaver et al.'s definition of **ambiguity** as "a property of the interpretative relationship between people and artefacts ..." that is "evocative rather than didactic, and mysterious rather than explicit" (Gaver et al., 2003, p.3). Simply put, artifacts using ambiguity support multiple interpretations by users. Previous HCI research has proved ambiguity and multiple interpretations, common qualities of art, to be a fruitful alternative to the usability principles in traditional UX design (Aoki & Woodruff, 2005; Boehner & Hancock, 2006; Daudén Roquet & Sas, 2021; Gaver et al., 2003; Jorge et al., 2013; Sanches et al., 2019; Sanches et al., 2022a; Sengers & Gaver, 2006).

We survey related work on ambiguity in HCI and ambiguity in visual AI art. We then review nine AI artworks whose process has been clearly documented to illustrate how artists creatively embrace ambiguity throughout the entire ML process, which includes dataset curation, model training, and application, and encourage audiences to creatively interpret AI. Our analysis demonstrates that these artists effectively use the *ML processes* as new *sources of ambiguity*. Specifically, we identify techniques used by artists to influence the ML process. We reflect on the wider implications of our findings for HCI and design beyond art.

The main contributions of this paper are the following:

1. We broaden the HCI notion of ML as Design Material by highlighting the process of creating ML systems. Our analysis of AI art shows that the ML process provides a range of creative opportunities, relatively under-explored by the HCI design community, to directly influence how users experience the ML artifacts.

2. Following the idea that uncertainty (defined in Section F.2.3) is a foundational characteristic of ML, we identify ambiguity as an essential quality in human-AI interaction. Instead of regarding it as undesirable, we argue that designers of ML products should use it to help manifest inherent uncertainties for users. The techniques we identify serve as departure points for designing ML-based products and experiences.

3. We show how the application of ambiguity strategies and tactics throughout the ML process can foster rich interpretations of AI, including critical ones, and shed new light on notions of transparency and explainability. We offer new design directions, such as exposing rather than explaining ML uncertainty, and new perspectives for dealing with failure and error.

F.2 Theoretical Framework

We review three areas of related work: ambiguity as a resource for design in (primarily non-AI) HCI; accounts of how artists are applying ambiguity within AI art; and a brief introduction to the machine learning process and uncertainties that contribute to this.

F.2.1 Ambiguity in HCI

Ambiguity is recognized as a valuable resource in HCI and design research (Aoki & Woodruff, 2005; Boehner & Hancock, 2006; Daudén Roquet & Sas, 2021; Gaver et al., 2003; Jorge et al., 2013; Sanches et al., 2019; Sanches et al., 2022a; Sengers & Gaver, 2006). The notion that ambiguity can be a valuable resource for design, other than being a problem to be solved, was proposed by Gaver et al. (2003). Subsequent work by Sengers and Gaver (2006) explored how ambiguity can provoke interpretation. Their main idea is that an interactive experience can ask questions and thus prompt users to try and establish answers rather than aiming to give answers directly. Later research in HCI has further explored ambiguity from a variety of perspectives (Aoki & Woodruff, 2005; Boehner & Hancock, 2006; Jorge et al., 2013), including the areas of affective computing (Sanches et al., 2019), bio-data (Daudén Roquet & Sas, 2021; Sanches et al., 2022a) and experience design for museums (Benford et al., 2022; Ingimundardottir, 2018; Jorge et al., 2013; Ryding et al., 2021; Vayanou et al., 2019).

However, direct investigation of ambiguity in the design of ML products and services is still relatively rare. Most literature on human-AI interaction and human-centered AI, especially recent developments in explainable AI (Arrieta et al., 2020), focus on improving users' accurate understanding of ML (Amershi et al., 2019; Kulesza et al., 2013; Villareale et al., 2022).

Our paper explores an alternative approach to ML experiences. It builds on the foundation of the seminal paper by Gaver et al. (2003). The 2003 paper drew on the history of art, from Leonardo da Vinci onwards, to reveal how artists employ various forms of ambiguity to make their artworks open to multiple interpretations by viewers. The paper argued for considering three broad types of ambiguity when designing interactive systems: 1) *ambiguity of information*, in which information may be portrayed deliberately blurry or, conversely, overly precise ways (sometimes appearing to be too precise can be ambiguous); 2) *ambiguity of context*, in which interactive artifacts are experienced in unusual contexts, including ones for which they were not originally designed; and 3) *ambiguity of relationship*, in which the user's relationship to the work is ambiguously framed. The paper also articulated a suite of tactics for deliberately employing ambiguity which is reproduce here for completeness, along with the type of ambiguity they support:

Information 1) Use imprecise representations to emphasise uncertainty;

Information 2) Over-interpret data to encourage speculation;

Information 3) Expose inconsistencies to create a space of interpretation;

Information 4) Cast doubt on sources to provoke independent assessment;

Context 5) Implicate incompatible contexts to disrupt preconceptions;

Context 6) Add incongruous functions to breach existing genres;

Context 7) Block exposed functionality to comment on familiar products;

Relationship 8) Offer unaccustomed roles to encourage imagination;

Relationship 9) Point out things without explaining why;

Relationship 10) Introduce disturbing side effects to question responsibility.

Our paper extends (Gaver et al., 2003) by connecting it to the new design domain of AI and ML, which have emerged into widespread application in the twenty years since the paper was written. Even a cursory glance at the list of tactics reveals potential connections. Ideas such as over-interpreting data, casting doubt on sources, and questioning responsibility, resonate with current concerns about the transparency and responsibility of AI (Arrieta et al., 2020), calling for a deeper consideration of the relationships between AI, ML, and ambiguity. In exploring this issue our paper also turns to the work of artists, but in this case, contemporary artists who have been working closely with AI and ML.

F.2.2 Ambiguity in Visual AI Art Theories

In recent years, the field of AI art has seen a remarkable surge in innovation and public interest. Compared to the design fields, AI art has developed with relatively fewer roadblocks fueled by high-profile projects such as *DeepDream* (2015) (Mordvintsev et al., 2015), *StyleGAN* (2021) (Karras et al., 2021), *DALL-E* (2021) (Ramesh et al., 2022), and *Stable Diffusion* (2021) (Rombach et al., 2021)). For example, the AI art landscape is enriched by artworks such as "The Next Rembrandt," a painting generated using data from Rembrandt's entire body of work (Narvaez et al., 2022), and "Sunspring," a short film scripted by an AI (Brynjolfsson & Mcafee, 2017). Also in interactive arts, AI-driven installations like "Learning to See" by Memo Akten explore the boundaries of audience engagement and artistic intent (Akten et al., 2019b). The fast development of tools for the production of visual art has also caused controversy for the challenges it poses to intellectual property (Appel et al., 2023; Chen, 2023; Epstein et al., 2023; Epstein et al., 2020) and the risk of propagating and multiplying existing biases in visual art (Larrazabal et al., 2020; Wang et al., 2022). How AI artists navigate the complexities of ML to create novel yet rich experiences may provide inspiration for HCI designers.

While theories of AI art are still emerging, ambiguity has appeared as a central characteristic of aesthetic experiences built on ML. For instance, Murray-Browne and Tigas (2021) suggest embracing ambiguity as it can be valuable in computer vision-based interactive art installations involving body movement. They point to design for emergence, openness, and ambiguity to create opportunities for forms to emerge that go beyond the initial artistic vision.

Current AI art theories can be broadly divided into ambiguity in the visual styles of ML output, in the ML process, and in the art discourse. First, in the search for new visual aesthetics of AI art, many artists and theorists have looked to **ambiguity in visual output**. For instance, Hertzmann (2020) notes that since modern ML does not have any concept of "objects" or "space", it can generate images that defy coherent spatial interpretation in ways similar to a painting by Escher. He calls this phenomenon *visual indeterminacy*. Hertzmann observes that Generative Adversarial Networks (GAN) tend to create indeterminate images, and suggests that this type of visual ambiguity is an inherent property of the technology. In another example, Mazzone and Elgammal (2019) proposes that AI artists should embrace *stylistic ambiguity*, a strategy for ML to generate images based on blending recognizable artistic styles. This way, AI art can achieve novelty without departing too much from acceptable aesthetic standards. This paper further theorizes this phenomenon and places it in the broader context of HCI theories of ambiguity.

Second, researchers and artists have explored **ambiguity that arises from the ML computational process**. For example, Boyé et al. (2019) investigate how artists explore machine flaws, irregularities, and errors in the computational process to push

the boundaries of their artistic practice. Continuing the conceptual tradition of earlier digital art movements such as glitch art, ML's calculation errors become a creative opportunity to find new thematic, technological, and conceptual foundations for experimentation. Compared to earlier digital technology, modern ML is more capable of processing the complex data of human languages (e.g., natural languages, images, and body gestures), enabling the possibility of processing rich human communication as input. This allows artists to exploit the fact that data from human communication contains many semantic ambiguities, while many applications of ML assume these data have fixed, universal meanings (Crawford & Paglen, 2019; Denton et al., 2021). For example, Xu (2021) intentionally used ambiguous sentences and unusual word combos as text prompts to generate images using *DALLE-2*.

The artistic intent, in this case, seems to be curiosity about how machines respond to ambiguous human requests. Most existing work focuses on the outcome of ML. However, ML is also a process (Section F.2.3). Relatively little scholarly work has opened the black box of ML, especially the ML process, and looked closely at how artists use different aspects of the computational process as a source of ambiguity. To our knowledge, this paper is among the first works to take a closer look at the ML process and illustrate how choices during the process lead to ambiguity in AI art.

Third, artists also use the broader social discourses around AI/ML to frame **audience's perception of ML and AI art as ambiguous**. For instance, Stark and Crawford (2019) survey how artists investigate the ethical ambiguity of data and ML as the intent of their work. Another common approach is to explore the ontological ambiguity around AI and AI art through the lens of anthropomorphism, especially the audience's tendency to attribute general human intelligence to existing narrow AI (Grba, 2021a, 2021b, 2022). Continuing the practices of early AI artists such as Harold Cohen (1995), some contemporary artists deliberately withhold or obfuscate the information about how ML produces the artwork to inject ambiguity in the public's interpretation of the work (Cook et al., 2019). Similarly, a recent study investigates how the audience's appreciation of a painting is affected when the viewer is uncertain about whether an AI or a human produced the artwork (Gu & Li, 2022). Yurman and Reddy (2022) explore using GAN-generated images as "more than human" elements to mediate the drawing conversations between two humans. In this work, the deliberate inability of the AI to produce clear and coherent images adds to the "interpretative flexibility" of the images sent between humans. This paper connects these perceptions to the ML process.

This growing body of theories points to the critical role of ambiguity in the context of AI art. It also indicates that a close study of how AI artists use ambiguity can lead to new design strategies for HCI designers to work more effectively with ML. Note that we do not claim that ambiguity in AI art is entirely different from closely related art forms such as software art, evolutionary art, or glitch art. Many experiments in AI art continue the conceptual themes and approaches that these other art

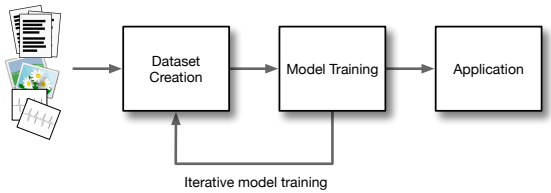

Fig. F.1: A Typical Machine Learning Pipeline for Computer Vision and Image Synthesis

movements have explored, and they share similar intellectual traditions. While highlighting that current practices in AI art are part of a larger continuum of exploration at the interface between art and technology, for the scope of this paper, we focus on AI art.

F.2.3 Uncertainty in the Machine Learning Process

In the rest of the paper, we use the term **uncertainty** to refer to a property of the *ML/AI computational process*. In contrast, *ambiguity* is used to describe these *interpretive relationship* between (AI) artefacts and human audience. This is aligned with Gaver et al.'s distinction of ambiguity from related concepts such as "fuzziness" or "inconsistency," noting that "ambiguity is an attribute of our interpretation" whereas the latter concepts are attributes of things (Gaver et al., 2003, p.3).

J. J. Benjamin et al. (2021) divide ML uncertainty into two main types. The first is *data uncertainty*, related to the "noise" in the training data sets. Data uncertainty can be introduced in the training dataset by, for instance, blurry images or erroneous labels. The second, and perhaps the more entrenched type, is what they call *model uncertainty*. It captures the "epistemic uncertainty" that characterizes all statistical inferences that underlie all ML decision-making. The ML process uses "statistical intelligence" (Dove et al., 2017) to quantify, analyze, and manage uncertainty. How the engineers train a model, which ML architecture they choose to use, and when they stop the model training process can impact the amount of model uncertainty, but can never eliminate it completely. J.J. Benjamin et al. argue that this type of statistical uncertainty is the "material expression of ML decision making" (J. J. Benjamin et al., 2021, p.2).

This section introduces the standard ML pipeline with key terminologies as well as the role of uncertainty in the ML process. Like most modern ML applications, the AI art in this paper uses Deep Neural Networks (DNNs), also known as DNN models, deep learning, or deep learning models. DNN applications in both computer vision and image synthesis have recently achieved human-level performance in some tasks (Dodge & Karam, 2017; Russakovsky et al., 2015). Common DNN model architecture in the visual domains include Convolutional Neural Networks (LeCun, Bengio, et al., 1995), Transformers (Dosovitskiy et al., 2020), Generative Adversarial Networks

(GANs) (Goodfellow et al., 2020), and diffusion models (Ho et al., 2020). The reviews of their architecture, performance and limitations can be found in (Huang et al., 2018; Voulodimos et al., 2018).

Fig. F.1 shows the high-level technical process that all DNNs follow: 1) *dataset creation*: collecting, cleaning, and labeling data to create a training dataset; 2) *model training*: select a model architecture (e.g., GAN) and train a corresponding model using the training dataset; 3) *application*: apply the trained model for a specific task (e.g., plant recognition). Typically these steps are repeated iteratively until the model achieves the desired performance. In the following we will use "dataset curation" about the first step to highlight, how this can be part of a creative process involving the appropriation of existing datasets, remixing, and creating from the ground up.

For example, in the case of a supervised image classification model, engineers first need to collect and clean a large collection of images the ML model will likely see. Each image needs to be labeled, by humans, with relevant tags (e.g., "*violet*" or "*overwatered*"). Then, given a model architecture (e.g., a CNN), engineers train the model iteratively using the training dataset to adjust the model's parameters to maximize its performance. After that, the model can be applied to recognize new images it has not seen before.

F.3 Approach

Our approach builds on the humanistic tradition of textual analysis and critique, which is increasingly being applied as a method in HCI (Bardzell & Bardzell, 2015). In order to investigate the ways in which artists utilize machine learning processes to create ambiguity in their works, we have selected 9 artworks that use deep neural networks either for image synthesis or classification. This analysis is not intended to be a comprehensive review of AI art. Extensive reviews already exist on AI/ML art in general (Audry, 2021; Cetinic & She, 2022), artificial life/genetic art (Penny, 2009; Tenhaaf, 2008), robotic art (Penny, 2013), and artists' account on how they work with ML (Caramiaux & Fdili Alaoui, 2022).

Instead, our focus is whether and how artists use the ML process to create ambiguity in their works, we used the original framework of Gaver et al. (2003) as our theoretical guide to ensure that our selection of AI art covers a wide variety of tactics to create ambiguity. Another selection criterion is that the artist has published a technical description of how they use machine learning to create their artwork. This is to ensure that we can sufficiently analyze the underlying ML process. Such explanations are often presented in metatext accompanying the work, such as in the websites for *Machine Bias*, *Learning to See*, *POSTcard Landscapes from Lanzarote*, *Butcher's Son*, *ImageNet Roulette*, *in transitu*, *Unsupervised* and *Biometric Mirror*. Further explanations are

frequently offered elsewhere in texts written by the artists such as essays (Crawford & Paglen, 2019; Vetere & Wouters, 2018), academic papers (Akten et al., 2019a; Guljajeva & Canet Sola, 2022; Wouters et al., 2019) and a PhD dissertation (Akten, 2021).

The authors initially used their domain expertise to identify as many AI art projects as possible. Three authors then individually analyzed the art projects and identified which tactics for creating ambiguity from Gaver et al. (2003) were used in each art project. Next, the authors discussed and consolidated their analyses until they reached an internal consensus. Finally, we selected 9 artworks that altogether cover all the tactics from Gaver et al. (2003) except "Point out things without explaining why" (Tactic 9). While this tactic is not necessarily in conflict with our selection criteria stated above - that the artist must have published a technical description - we find in practice that these descriptions tend to also include explanations for the artistic choices in the work, to the extent that none of the selected artworks seem to match this tactic. This analysis is summarized in Table F.1. While we did not encounter new types of ambiguity outside those presented by Gaver et al. (2003), the scope of our analysis is to understand how AI artists create the known type of ambiguity. We acknowledge that there may be new types and tactics of ambiguity for further investigation.

F.3.1 Sample AI Art

Below we present the nine artworks. Table F.1 shows the mapping of each artwork to the tactics from Gaver et al. (2003).

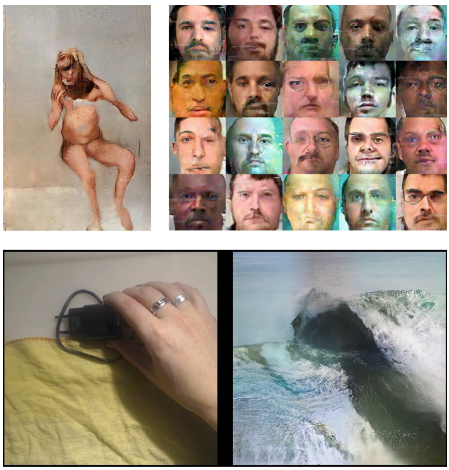

Fig. F.2: From top left to bottom: *Butcher's Son* (Klingemann, 2017), *Machine Bias* (Yaszdani, 2018) and *Learning to See* (Akten et al., 2019b)

The Butcher's Son (Fig. F.2, top left) (Klingemann, 2017) is one in a series of generated painting-like portraits created by Mario Klingemann that "focus on the human body, training his AI models to explore posture by turning stick figures into paintings, based on the analysis of images harvested from the internet"[1]. The images include visual artifacts introduced in the ML process.

Learning to See (Fig. F.2, bottom) is an interactive video installation by Memo Akten which appropriates GAN models trained on images of waves, flowers, and fire and applies them to real-time video feeds of mundane everyday objects such as phone chargers, pens, and fabric; turning them into animated waves, flowers, or fire in a similar composition (Akten et al., 2019b).

Machine Bias (Fig. F.2, top right) by Nushin Isabelle Yazdani is a series of generated faces based on photographs of prison inmates from across the US. She uses these "future faces of prisoners" to question predictive policing and automated pretrial risk assessment. (Yaszdani, 2018)

Biometric Mirror (Fig. F.3, top), created by Microsoft Research Centre for Social Natural User Interfaces, presents itself as a system designed to "stimulate individual reflection on the ethical application of artificial intelligence"[2]. The system invites people to have their faces photographed through a webcam and analyzed by a psychometric system, which classifies their faces on a range of dimensions from relatively overt traits such as age, gender, and ethnicity to more diagnostic concepts such as aggressiveness, weirdness, and emotional instability (Wouters et al., 2019).

ImageNet Roulette (Fig. F.3, bottom) is a digital app and AI art installation by Paglen and Crawford (2019). The app lets users upload a photo, for instance, a selfie. The app will return the same photo with a green bounding box around every human face, each with a series of labels. For instance, in the iconic White House Situation Room photo Hillary Clinton is given the label "flutist, flautist, flute player." Other users' photos received more problematic labels such as "swot, grind, nerd, wonk, dweeb (...) rape suspect (...) first offender (...) gook, slant-eye" (Wong, 2019). These labels came from ImageNet, one of the most widely used *training datasets* in computer vision (Crawford & Paglen, 2019).

POSTcard Landscapes from Lanzarote (Fig. F.3, middle) by Varvara & Mar consists of two videos about the island of Lanzarote, part of the popular Canary Islands. It explores "the tourist gaze" in contrast to the local view of the place. The artists collected public photographs of the island on *Flickr* that represent these two perspectives (Guljajeva & Canet Sola, 2022). The videos show the latent space of two generative models, each trained on one of the two sets of images. Each video shows morphing ambiguous imagery that at times vaguely resembles things like landscapes or airplanes, but often escapes clear categorization.

[1] https://www.electricartefacts.art/artwork/mario-klingemann-the-butchers-son-special-edition
[2] https://biometricmirror.com/

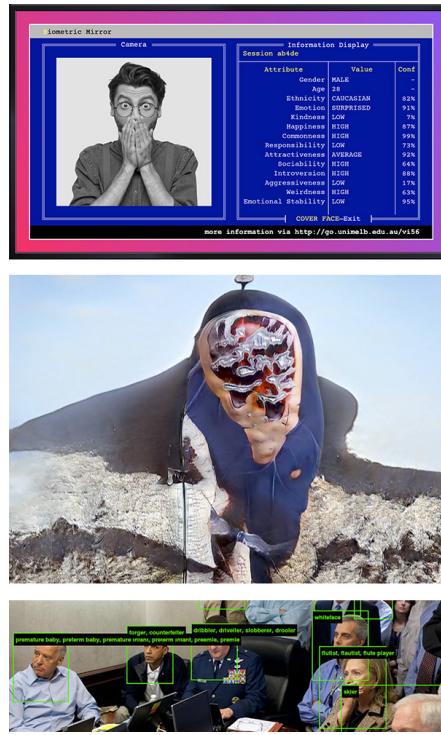

Fig. F.3: From top to bottom: *Biometric Mirror* (Wouters et al., 2019), *POSTcard Landscapes from Lanzarote* (Guljajeva & Canet Sola, 2022) and *ImageNet Roulette* (Paglen & Crawford, 2019)

Poison (Fig. F.4, top) was a recent installation exhibited at the MUNCH museum in Norway. It used ML to recreate Edvard Munch's "Green Room" paintings in a room in the museum. Visitors to the room saw projections of the "Green Room" from Munch's paintings covering three walls. As visitors moved through the room, the perspective of the artwork changed to match the visitor's movements, offering the illusion of stepping inside the room depicted in the paintings. In correspondence with the ambiguity of the paintings and questions raised in research about them, the perspective was unstable and the digital reproductions oscillated between different degrees of blurriness (Sivertsen et al., 2023).

Unsupervised (Fig. F.4, middle) by Refik Anadol (2022) is - in the artist's terms - a "data painting" belonging to Anadol's series *Machine Hallucinations*. Anadol has trained a StyleGAN 2 ADA model on 138,151 images from the archive of The Museum of Modern Art (MoMA). Through a custom piece of software called "Latent Space

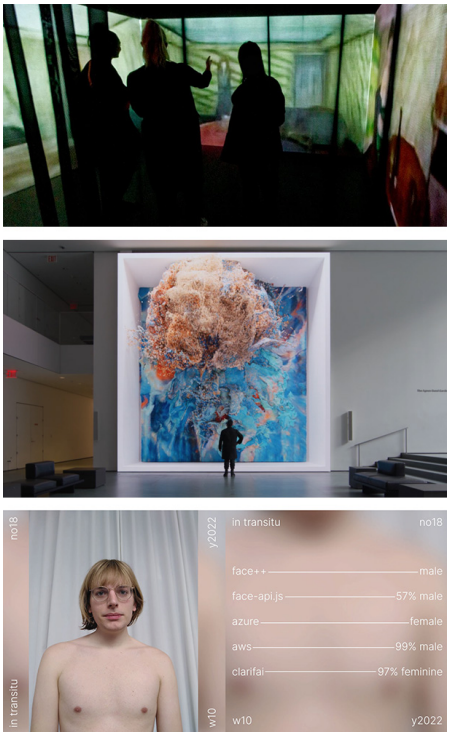

Fig. F.4: From top to bottom: *Poison* (MUNCH et al., 2021), *Unsupervised* (Anadol, 2022) and *in transitu* (Ada Ada Ada, 2022)

Browser," the system generates images from the latent space of the model, resulting in fluid interpolations of colors, shapes, and patterns emerging from the corpus of the collection without ever representing any specific work as such [3].

in transitu (Fig. F.4, bottom) by Ada Ada Ada (2022) explores gender recognition by commercial image analysis systems. The artist, a trans woman, periodically uploads images of herself naked from the waist up, along with the gender recognition outcome from five different ML models applied to the image. While the photos show only small variations in hairstyle and facial expression, the outputs from the algorithms vary widely. Sometimes the same photo is classified by different algorithms as male and female with confidence levels near 100%. (All of the algorithms appear to treat gender as either a binary variable (male or female) or a percentage scale between male and female.)

[3] https://refikanadol.com/works-old/unsupervised/

Table F.1: Overview of AI art examples described in this paper. (The numbers under "Tactics" refer to the list of tactics in section F.2.1)

Work	Year	Technology	Types of Ambiguity	Tactics
Butcher's Son (Klingemann, 2017)	2017	Image synthesis	Information	1, 3
Learning to See (Akten et al., 2019b)	2017	Image synthesis	Information & Context	1, 2, 5
Machine Bias (Yaszdani, 2018)	2018	Image synthesis	Information & Relationship	1, 2, 4, 10
Biometric Mirror (Wouters et al., 2019)	2018	Facial analysis	Information, Context & Relationship	2, 3, 4, 5, 8, 10
ImageNet Roulette (Paglen & Crawford, 2019)	2019	Facial analysis	Information & Relationship	2, 3, 4, 8, 10
POSTcard Landscapes from Lanzarote (Guljajeva & Canet Sola, 2022)	2020	Image synthesis	Information	1, 2
Poison (Sivertsen et al., 2023)	2021	Image synthesis	Information & Context	1, 3, 6, 7
Unsupervised (Anadol, 2022)	2022	Image synthesis	Information & Context	1, 6, 7
in transitu (Ada Ada Ada, 2022)	2022	Facial analysis	Information & Relationship	2, 3, 4, 10

F.4 The ML Process as a Source of Ambiguity

When analyzing the artworks, it became clear that in our selected examples an important part of the story surrounding the work is about how the machine learning processes were appropriated to reach a particular output.

In this section, we will investigate the three steps artists go through when working with machine learning, data curation, model training, and application to investigate which concrete techniques have been used for each artwork to arrive at the ambiguity tactics shown in Table F.1.

F.4.1 Dataset Curation

The training datasets determine the ontology of an ML model — what concepts and categories it can learn, what their relationships are, and what each concept "looks like" in the real world. In other words, datasets encode meanings, interpretations, and world views of those who made it (R. Benjamin, 2019; Broussard, 2018; Crawford, 2021; D'ignazio & Klein, 2020; Milan & Treré, 2019). Training datasets thus offer artists a rich source for ambiguity.

Selecting existing datasets with questionable ontologies

Training datasets establish the ground truth for ML models, but they are created by people in their social contexts. We find that some artists pick datasets to expose problematic value systems embedded within. A salient example is *ImageNet Roulette*. The artwork used an ML model trained on a subset of *ImageNet* (Deng et al., 2009), one of the most widely used *training datasets* in Computer Vision applications (Crawford & Paglen, 2019). At the time of making the artwork *ImageNet* contained a vast 14 million labeled images, including a "person" category. The annotation of images with words has been done under the assumption that every verbal concept from the WordNet database could and should be "imaged" (Malevé, 2021). The "person" category of labels included derogatory terms such as gendered and racial slurs which were applied to images scraped from the internet. This not only resulted in having images of specific people labeled with derogatory terms; it also made it possible for an algorithm trained on these categories to apply the same terms to images of other people. *in transitu* and *Biometric Mirror* similarly build on pre-existing models with problematic built-in assumptions: In the case of *Biometric Mirror* the model relies on physiognomy – the idea that a person's character can be assessed from their appearance; whereas the models explored in *in transitu* rely on the assumption that gender identity is binary and unambiguous.

Through the *hyper specificity* of the labels and scores in the datasets, a form of over-interpretation of the data exists at the very beginning of the ML process. By making visible the arbitrary and often offensive nature of the labels and scores these models assign to people, the artists encourage audiences to question the perception of computer vision as objective and neutral. In this way these three artworks expose the inconsistencies in the datasets' inherent world-view and cast doubt on their authority, utilizing tactics 2, 3 & 4: *Over-interpret data to encourage speculation*, *Expose inconsistencies to create a space of interpretation* and *Cast doubt on sources to provoke independent assessment*.

Making bespoke ontologies through new datasets

In our other examples, the artists took it upon themselves to collect, curate, and clean datasets that fit their purpose. However in *Machine Bias* and *POSTcard landscapes from Lanzarote* the artists make a point of creating datasets with an explicit ontological claim. In *POSTcard landscapes from Lanzarote* the artists use photos from Flickr to build datasets that represent "the tourist gaze" and the local view, implying that the photos taken by visitors and locals reveal something about the ways in which these different groups of people see the place. However, a great part of the images are abstract and hard to interpret. As such, this use of data might be characterized as using tactic 1: *Use imprecise representations to emphasise uncertainty*.

Machine Bias employs a similar strategy, building on a bespoke dataset - but arguably one which makes a much more problematic claim, suggesting that one may predict the facial features of future criminals based on analysis of the faces of past criminals. Thus the practice of collecting and applying this dataset raises philosophical and ethical questions about whether certain datasets should be collected and if it is even possible to represent with images what the data supposedly does. As they generate the faces of future criminals the artist intentionally *over-interpret data to encourage speculation* (tactic 2).

Through these two techniques, the artists are investigating the ontology of their respective ML systems: what categories exist, how are they related, and what do they "look" like? Importantly, this investigation does not happen through descriptions of the formal qualities but by investigating what the models *do*. How do they mediate the world they are supposedly representing?

F.4.2 Model Training

Typically, an ML model is trained to optimize its performance defined by certain metrics (e.g., minimizing prediction error or creating a realistic-looking image). Several artworks engage with — and often subvert — the model training process in non-standard ways, often to create unusual or striking visual effects.

Repurposing upscaling to introduce visual artifacts

Upscaling is an ML technique to convert lower-resolution media to a higher resolution. To save computational power, image synthesis models typically generate low-resolution images (e.g., 64x64 pixels) first and then use upscaling to enhance them to high-resolution ones (1024x1024 pixels) (Saharia et al., 2022). While upscaling is typically used to enhance images to more crisp and realistic looks, AI artists may repurpose the technique for different effects. Klingemann, the artist who created *Butcher's Son*, used upscaling in a process he calls "transhancement." As described in the artwork's catalog text, he used GAN to generate low-resolution content such as skin texture and hair and transhance them into a full portrait in ways that leave space for unusual artifacts: "The result is painterly and ethereal, a neural network's vision of the human form."[4]. The resulting image carries an imprecise aesthetic and, from a human perspective, inconsistency in its seeming representation of a human body which are representative of tactic 1: *Use imprecise representations to emphasize uncertainty* and tactic 3: *Expose inconsistencies to create a space of interpretation*.

[4]https://www.artsy.net/artwork/mario-klingemann-imposture-series-the-butchers-son

Under-fitting ML models

A common undesirable scenario in ML is *over-fitting* a model. It happens when the model is tailored too exactly to the training dataset that it may fail to generalize to unseen data. Conversely, *under-fitting* is another unwanted scenario where an ML model is unable to capture the relationship in the data accurately, thus having a low performance. It is usually caused by stopping the training process too early. *Poison* turned this undesirable ML behavior into an artistic technique. The artists used it to emphasize the ambiguity of the motifs in Munch's paintings. The interactive visualizations were generated by applying style transfer to 3D renderings based on original paintings by Edvard Munch. However, instead of using the best-fitted model, the artists used an *underfitted* version from earlier steps in the training process. This caused the resulting visualizations to have a fuzzy, ever-shifting appearance that never offered a clear representation of the motifs in the original paintings, thus applying tactic 1: *Use imprecise representations to emphasize uncertainty*.

Changing the output modality from the input modality

Both *Poison* and *Unsupervised* use static images of older artworks as their input (i.e. training data). However, in both works the output is dynamic and responds to the presence of visitors, thus including elements that are incongruous to the type of art they represent. The output of *Unsupervised* is a dynamic visualization with mostly abstract, colorful patterns that are ever-emerging, moving, and disappearing. In *Poison* the physical movement of the audience changes the perspective of the digital projections, mimicking the way one's perspective on a physical space shifts when moving around - offering an illusion of peering directly into the room depicted in the paintings. In both these cases, bespoke software was used to achieve these particular dynamics. Simultaneously, these visualizations refuse to reproduce the imagery of the original paintings, thus blocking them from being viewed in the way they were originally presented. Since the reproduced paintings in *Poison* constantly shift and move, it removes the possibility of perceiving it the way you would perceive an oil painting. *Unsupervised* removes the ability of the audience to recognize the artworks the system was trained on and presents them not as discrete objects, but as movements. From this perspective, *Unsupervised* and *Poison* might both be said to add and remove affordances to and from the originals they use as training data, leading to tactics 6 & 7: *Add incongruous functions to breach existing genres* and *Block exposed functionality to comment on familiar products*.

With these three techniques, the artists are exploring the concept of fit between the model and its underlying data. The fit is always characterized by some *epistemic uncertainty* that can be reduced or increased. Here we see that contrary to conventional approaches a sub-optimal fit may be used for aesthetic effect.

F.4.3 Application

Through the selection of data used for an ML application, an ontology is established with an assumption that the model should be good at generating or classifying images within a specific domain. Shifting the domain can change the relationship in subtle or dramatic ways.

Applying models to a different domain

Typically, ML engineers select the training datasets that resemble as much as possible the domain where the ML model will be deployed. However, AI artists have challenged this setup and applied models to different domains from where the training dataset was collected. An illustrating example is *Learning to See*. The artist appropriates GAN models trained on images of waves, flowers, and fire and applies them to transform real-time video feeds of mundane objects into matching imagery of fantastic natural objects. The models consistently generate pretty visuals and effectively disguise the bland everyday objects. This technique may be seen as an expression of tactic 5: *Implicating incompatible contexts to disrupt preconceptions*.

Connecting models directly to people

Many surveillance applications quietly use ML to classify people, but those who were watched do not usually have access to how the ML model classifies them. Both *in transitu* and *Machine Bias* expose the functioning of ML systems by showing their direct application on concrete persons. Taking this technique one step further, *ImageNet Roulette* and especially *Biometric Mirror* encourage viewers to use *their own* faces and observe how they will be classified. Having the rather stereotypical and sometimes even derogatory labels applied to their own body puts the audience in a vulnerable position that invites new perspectives on what impact such systems may have. These works may all be said to apply tactic 10: *introduce disturbing side effects to question responsibility*.

When applying the system to an input found "in the wild", these two techniques show the opportunity to investigate what happens when we let the model mediate different parts of our world. New uncertainties can be generated through unconventional applications, and the promise of commercial systems can be scrutinized by applying them to situations that make the consequences of such mediation clear.

F.4.4 From uncertainty to ambiguity

Across all three steps of the process, the artists are negotiating the uncertainty inherent in the system by exposing it, exaggerating it, and generating new uncertainty by holding the model ontology against the world. They make some of these processes

and considerations available to us, through the artwork itself, its staging, or metatext. This may, in turn, be experienced by us as ambiguous images and text (ambiguity of information), through systems that appear in unexpected contexts (ambiguity of context), and through the relation the system establishes with ourselves or other people (ambiguity of relation). By being exposed to these inconsistencies, unfamiliar relations, and imprecisions, we can doubt, question, speculate, and re-think our understanding of ML systems.

F.5 Discussion

Our investigation of the selected artworks highlights the ML process as an important source of ambiguity. This section discusses our findings.

F.5.1 From Artifacts to Processes: Broaden ML Ambiguity

Our finding points to the need for a broader understanding of ambiguity in the context of AI/ML. Traditional HCI primarily focuses on the "meaning of object produced" (Devendorf & Rosner, 2015). Similarly, the framework of ambiguity by Gaver et al. (2003) is concerned with the interpretive relationship between the user and the designed artifact *per se* in its context. To some extent, this artifact-centric notion of ambiguity is still meaningful. Some of our techniques, such as "repurposing upscaling to introduce visual artifacts" and "connecting models directly to people", directly introduce the ambiguity of information and relationship at the level of the artifacts. Similarly, *Edmond de Belamy* (Obvious, 2018), the first ML art auctioned at Christie's, creates the ambiguity of context by placing images generated by ready-made computer codes in a well-established art institution (we excluded this artwork due to its lack of documentation of the ML process).

However, the ambiguity of other artworks in our analysis arises from the underlying ML process. For example, works like *Machine Bias* rely on the story of the data curation being known, as the image itself without any context does not reveal the connections to predictive policing on its own. This is also the case for *Unsupervised* in which the choice of using data from the MoMA collection is important, and *Butcher's Son* in which Klingemann's use of *transhancement* plays an important role in the storytelling around the work. Similarly, the eerie visual artifacts in *Butcher's Son* cannot be separated from how these images were generated and by whom. Techniques such as "Selecting existing datasets with questionable ontologies", "Under-fitting ML models" and "applying models to a different domain" directly tap into the ML process to create ambiguity.

It is important to note that a given technique does not always lead to a specific type of ambiguity, as it is dependent on the broader context of the work. Furthermore, the list of techniques is not exhaustive, and many more techniques could conceivably be

found by looking at other artworks, and potentially even the same artworks. In fact, we are surprised to find that even within the relatively small number of artworks these techniques covered the entire pipeline of ML from dataset curation and model training to application. These techniques from our analysis serve as an illustration of how working with the ML process can be a rich source for creating artistic outcomes relating to ambiguity.

This shift from artifact to process is consistent with the extensive survey of deep learning-based AI art by Audry (2021). His work revealed artists' active experimentation with almost every ML technical variation. Similarly, Caramiaux and Fdili Alaoui (2022) documented the same trend from the artists' perspective. They found that AI artists favor the process (i.e., the workflow) over the outcome as a way to create artistic experiences because it is difficult to anticipate the result of a specific model with a specific dataset. Our analysis of ambiguity in AI art not only offers further evidence that art practice in the age of ML emphasizes process, but also makes necessary a broader notion of ambiguity.

F.5.2 ML Uncertainty as the Material Foundation of Ambiguity

Our findings add to a growing body of work that identifies uncertainty as an important property of ML as design material (J. J. Benjamin et al., 2021; Caramiaux & Fdili Alaoui, 2022; Sanches et al., 2022b). In contrast to existing literature that sees uncertainty primarily as a design challenge (e.g.,(Dove et al., 2017; Yang et al., 2020)), we argue that uncertainty is the material foundation of the ML process and cannot be fully eliminated.

While deterministic systems follow a script for interpreting the world, the probabilistic ML systems can even be "uncertain *about* the legibility of the world" (J. J. Benjamin et al., 2021, p.10). As argued by J. J. Benjamin et al. (2021), ML systems have *thingly uncertainty* as a common property. In many commercial applications of machine learning, this uncertainty is hidden. However, as Benjamin argues with the term *horizontal relations*, our human-technology relations might be *textured* with these uncertainties of ML systems as they are working behind our perceptual horizon.

Common to many of the artworks we have examined here is the insistence on bringing back into view this texturing and exposing it through exaggerations and juxtapositions. The relation is brought back into our immediate, perceptual *here and now* through establishing other relations that bring the ML uncertainty back into view.

As the uncertainty is exposed it gives rise to ambiguity. In this case, ambiguity is the most fitting way to represent the inherent uncertainty, rather than an incongruous insistence on accuracy and clarity. To the extent that we hide away the uncertainty in horizontal relations, we neglect our own ability as designers and as users to meaningfully assess the qualities of a given system.

F.6 Implications for ML as Design Material

While ambiguity has been accepted in HCI design, its role has not been sufficiently understood in human-AI interaction, which to this date has privileged certainty and transparency. In this section, we discuss the implications of our analysis to the HCI, especially the Human-AI interaction design community.

F.6.1 ML as Process vs ML as Material

Our analysis shed light on a limitation of current framing in the HCI design community to see ML as a design material (Dove et al., 2017; Yang et al., 2020). Despite the many benefits of this framing, a design material downplays the fact that ML is a process, an especially complex and opaque one. Seeing ML as a design material may encourage designers to conceptualize ML primarily along the tangible dimensions such as input-output mapping ("possible system outputs") and supported features ("system capability") (Yang et al., 2020). Our analysis reveals that a process-based model can bring to focus different aspects of ML (e.g., data curation, model training, and application), which are under-utilized as design elements. Designers may consider how different choices in data curation affect both the ontology and the aesthetic expression of the resulting model. During model training, choices can be made about how loosely the model is fit to the data and how the output is presented. Finally, considerations must be made regarding where to apply the model and whether the model ontology and the domain applications are (un)intentionally misaligned. These are all design choices as much as data science ones.

F.6.2 Exposing vs. Explaining Uncertainty

Currently, the HAI design community is preoccupied with reducing the impact of ML uncertainty and low interpretability. A widely adopted approach is to apply eXplainable AI (XAI) techniques to open the black box by explaining ML models for human inspection (Arrieta et al., 2020). The perspectives explored by these AI artworks suggest looking in the opposite direction, towards exploring the value of ambiguous and enigmatic AI. Gaver et al. (2003) argued that ambiguity may inspire critical reflection and deeper engagement with systems. Similarly, one might consider how making AI more ambiguous and open to interpretation might provoke humans to reflect more critically and deeply on the way the AI draws inferences from data and its training - and perhaps also on the nature and bias of AI itself. A growing body of literature has shown that users tend to over-rely on AI. It can be especially problematic in high-stake areas (Benford et al., 2022; Howard, 2020; Robinette et al., 2016). The artists' strategies we identified in our analysis show that designers can use ML uncertainty to encourage users' autonomy and critical thinking.

Designers need more robust methods to work the nonlinear, unpredictable ways uncertainty manifests itself in ML. Existing design methods such as Critical Design (Bardzell & Bardzell, 2013), Reflective Design (Sengers et al., 2005), Value-Sensitive Design (Friedman, 1996), and Speculative Design (Auger, 2013) provide a good departure point towards exposing uncertainty. Similarly, design methods grounded in humanistic theories (e.g., bricolage and mythical thinking (Vallgårda & Fernaeus, 2015), posthumanism theories (Sanches et al., 2022b; Wakkary, 2020)) offer fruitful alternative ways to approach the nonlinear and unpredictable nature of the ML process. Although not directly focusing on ambiguity, emerging approaches such as Introspective AI (Brand et al., 2021) provide concrete examples of incorporating ML uncertainty in design.

F.6.3 Failure and Error

Failure and error are closely related to ML uncertainty. The core training process in ML involves optimizing the algorithm to minimize error. And usually, designers also aim to avoid failure and error. However, when designing for probabilistic systems some amount of error is a given, and interfaces for such systems should reflect this. Examples from AI art demonstrate that errors and shortcomings of the technology can help to define the aesthetics of ML – whether it is to expose problematic aspects of the systems, as in *ImageNet Roulette* or *in transitu*, or they are using the flaws or particular properties of the algorithms to create new aesthetic outputs, as in *Learning to See* or *Butcher's Son*.

Hazzard et al. (2019) suggest that failure should be treated as an aesthetic potential and should be actively incorporated into the design of creative interfaces. Hertzmann (2020) suggests that improvements in the ability of image generators to create realistic images might force artists to start manipulating or "breaking" the algorithms to avoid producing images that look just like ordinary photographs - and indicates that some current artists like Helena Sarin and Mario Klingemann are already exploring such approaches. Leahu (2016) argues that seeming failures in machine learning, such as the failure to reproduce one distinct object, can also be a consequence of a realist perspective. That is, having the assumption that the world can be separated into discrete entities. Leahu on the other hand argues for a relational perspective that gives rise to *ontological surprises* as surprising relations surface through the way ML systems make sense of their data. The ability of ML systems to convincingly produce outputs that are plausible imitations of things such as photographs calls for interfaces that embrace such relational perspectives, to work in tandem with the qualities of ML systems. This would support the user in reliably assessing the capabilities of the system for the task at hand. While this can be done by revealing mathematical relationships, we argue that letting the users experience it through experiences of ambiguity might be better suited for exposing the affordances of the system.

F.6.4 Designing the Discursive Strategies

To effectively give users handles on what to expect from AI systems, designers need to consider how the ML process is presented to users. Researchers in Science and Technology Studies and media studies have long argued that AI and ML are discursive (Agre & Agre, 1997; Mateas & Stern, 2003; Zhu, 2009). Their proper functioning requires humans to interpret their algorithmic process to be "intelligent" and "intentional." Recently, Murray-Rust et al. (2022) pointed out that key ML terms such as *Training*, *learning*, *Explanations*, *bias* and *Black box* are metaphors to interpret ML operations.

While the design focus of most ML products is not meaning-making, designers can select discursive strategies that fit their users. For instance, in computer games the degree to which the user interface reveals the existence of the underlying AI can impact player-AI interaction (Zhu et al., 2021). For designers of creativity support tools, designing the appropriate interaction metaphor (e.g., nanny, pen pal, coach, and colleague (Lubart, 2005)) can help users anticipate how to interact with the underlying generative ML models. Furthermore, we echo the point from J. J. Benjamin et al. (2023) that metaphors are not just used to describe what AI technologies *are* but also what they *do*, thus reflecting the role of technology in actively shaping the world.

F.7 Limitations

This paper presents an analysis based on a fairly small collection of nine artworks. Even though this satisfies our goal of identifying techniques used by AI artists when working with ML, we acknowledge that a comprehensive analysis of a larger collection of artworks could reveal additional techniques to create ambiguity. Furthermore, we have focused mainly on art exhibited in recognized venues and/or created by established artists. This choice has left out the large, emerging field of amateur artists creating art with ML.

The exact steps that the artists took to make the artworks we analyzed are hidden from the public. Our analysis thus relies on the artists' public description of their work in artwork metatext, interviews, blog posts, academic publications - combined with our own technical knowledge. We acknowledge that there may be discrepancies in the artists' accounts and how the ML system actually functions. This is a methodological challenge since it is not feasible to verify whether the artists' descriptions are accurate. However, our analysis does not for the most part rely on fine details about the technical system, but on the main conceptual use of the technology in the artworks.

F.8 Conclusion

This paper revisits the concept of ambiguity and brings it to bear on ML, again drawing on the work of artists at the forefront of experimenting with this new design material. Through an analysis of nine AI artworks, we have identified seven techniques that the artists used during the ML process. These techniques emerge as new sources of ambiguity. Built on literature identifying uncertainty as a fundamental property of ML systems, our analysis further shows how it comes into being and what it subsequently *does* in terms of mediating our relationship to the world. We have discussed different ways to design for *exposing*, such as through discursive strategies, and compared it to the topic of *explainability* in ML. Given that uncertainty is inherent, we argue that ambiguity is a relevant aesthetic for the design of human-AI applications because it allows designers and users to experience this quality of ML. Finally, our investigation concludes that the aesthetic of ambiguity lends itself to the various forms of uncertainty underlying all ML processes. It thus offers a powerful design direction for new generations of ML products and human-AI interaction. We believe that ambiguity, rising from profoundly engaging the ML process, can be part of the "new value" (Dove et al., 2017) the HCI design community has been looking for.

Acknowledgments

We want to thank Peter Kun for his contributions in the early phases of the development of this paper. Furthermore, we appreciate the help from Tom Jenkins and anonymous reviewers for the earlier version of this paper for their valuable comments and criticism that have enabled us to sharpen the perspective and hone our argumentation.

References

Ada Ada Ada. (2022). In_transitu. https://ada-ada-ada.art/projects/in-transitu

Agre, P., & Agre, P. E. (1997). *Computation and human experience*. Cambridge University Press.

Akten, M. (2021). *Deep Visual Instruments: Realtime Continuous, Meaningful Human Control over Deep Neural Networks for Creative Expression* (doctoral). Goldsmiths, University of London. Retrieved September 15, 2022, from https://research.gold.ac.uk/id/eprint/30191/

Akten, M., Fiebrink, R., & Grierson, M. (2019a). Learning to see: You are what you see. *ACM SIGGRAPH 2019 Art Gallery*. https://doi.org/10.1145/3306211.3320143

Akten, M., Fiebrink, R., & Grierson, M. (2019b). Learning to see: You are what you see. *ACM SIGGRAPH 2019 Art Gallery*, 1–6. https://doi.org/10.1145/3306211.3320143

Amershi, S., Weld, D., Vorvoreanu, M., Fourney, A., Nushi, B., Collisson, P., Suh, J., Iqbal, S., Bennett, P. N., Inkpen, K., Teevan, J., Kikin-Gil, R., & Horvitz, E. (2019). Guidelines for human-ai interaction. *Proceedings of the 2019 CHI Conference on Human Factors in Computing Systems*, Paper No. 3.

Anadol, R. (2022). Unsupervised - Machine Hallucinations - MoMa.

Aoki, P. M., & Woodruff, A. (2005). Making space for stories: Ambiguity in the design of personal communication systems. *Proceedings of the SIGCHI Conference on Human Factors in Computing Systems*, 181–190. https://doi.org/10.1145/1054972.1054998

Appel, G., Neelbauer, J., & Schweidel, D. A. (2023). Generative AI Has an Intellectual Property Problem [Section: Intellectual property]. *Harvard Business Review*. Retrieved June 28, 2023, from https://hbr.org/2023/04/generative-ai-has-an-intellectual-property-problem

Arrieta, A. B., Díaz-Rodríguez, N., Ser, J. D., Bennetot, A., Tabik, S., Barbado, A., Garcia, S., Gil-Lopez, S., Molina, D., Benjamins, R., Chatila, R., & Herrera, F. (2020). Explainable Artificial Intelligence (XAI): Concepts, taxonomies, opportunities and challenges toward responsible AI. *Information Fusion*, *58*, 82–115.

Audry, S. (2021). *Art in the age of machine learning*. Mit Press.

Auger, J. (2013). Speculative design: Crafting the speculation. *Digital Creativity*, *24*(1), 11–35.

Bardzell, J., & Bardzell, S. (2013). What is" critical" about critical design? *Proceedings of the SIGCHI conference on human factors in computing systems*, 3297–3306.

Bardzell, J., & Bardzell, S. (2015). *Humanistic HCI* [OCLC: 930589099]. Morgan & Claypool Publishers.

Benford, S., & Giannachi, G. (2011). *Performing Mixed Reality*. The MIT Press.

Benford, S., Sundnes Løvlie, A., Ryding, K., Rajkowska, P., Bodiaj, E., Paris Darzentas, D., Cameron, H., Spence, J., Egede, J., & Spanjevic, B. (2022). Sensitive pictures: Emotional interpretation in the museum. *Proceedings of the 2022 CHI Conference on Human Factors in Computing Systems*. https://doi.org/10.1145/3491102.3502080

Benjamin, J. J., Berger, A., Merrill, N., & Pierce, J. (2021). Machine learning uncertainty as a design material: A post-phenomenological inquiry. *Proceedings of the 2021 CHI Conference on Human Factors in Computing Systems*, 1–14.

Benjamin, J. J., Biggs, H., Berger, A., Rukanskaitė, J., Heidt, M. B., Merrill, N., Pierce, J., & Lindley, J. (2023). The Entoptic Field Camera as Metaphor-Driven Research-through-Design with AI Technologies. *Proceedings of the 2023 CHI Conference on Human Factors in Computing Systems*, 1–19. https://doi.org/10.1145/3544548.3581175

Benjamin, R. (2019). Race after technology: Abolitionist tools for the new jim code. *Social forces*.

Boden, M. A., & Edmonds, E. A. (2019). 2 a taxonomy of computer art. In *From fingers to digits: An artificial aesthetic* (pp. 23–59).

Boehner, K., & Hancock, J. T. (2006). Advancing ambiguity. *Proceedings of the SIGCHI Conference on Human Factors in Computing Systems*, 103–106. https://doi.org/10.1145/1124772.1124789

Bogost, I. (2019). The AI-Art Gold Rush Is Here [Section: Technology]. Retrieved June 28, 2022, from https://www.theatlantic.com/technology/archive/2019/03/ai-created-art-invades-chelsea-gallery-scene/584134/

Boyé, P., Grba, D., Pandilovski, M., Kang, K. H., & Todorović, V. (2019). Machine flaws in generative art. *Lux Aeterna, ISEA2019 25 th International Symposium on Electronic Art proceedings, Gwangju, Republic of Korea*, 713–716.

Brand, N., Odom, W., & Barnett, S. (2021). A design inquiry into introspective ai: Surfacing opportunities, issues, and paradoxes. *Designing Interactive Systems Conference 2021*, 1603–1618.

Broussard, M. (2018). *Artificial unintelligence: How computers misunderstand the world*. mit Press.

Brynjolfsson, E., & Mcafee, A. (2017). Artificial intelligence, for real. *Harvard business review*, *1*, 1–31.

Caramiaux, B., & Fdili Alaoui, S. (2022). " explorers of unknown planets" practices and politics of artificial intelligence in visual arts. *Proceedings of the ACM on Human-Computer Interaction*, *6*(CSCW2), 1–24.

Cetinic, E., & She, J. (2022). Understanding and creating art with ai: Review and outlook. *ACM Transactions on Multimedia Computing, Communications, and Applications (TOMM)*, *18*(2). https://doi.org/10.1145/3475799

Chen, M. (2023). Artists and Illustrators Are Suing Three A.I. Art Generators for Scraping and 'Collaging' Their Work Without Consent [Section: Law]. *Artnet News*. Retrieved June 28, 2023, from https://news.artnet.com/art-world/class-action-lawsuit-ai-generators-deviantart-midjourney-stable-diffusion-2246770

Cohen, H. (1995). The further exploits of AARON, painter. *Stanford Humanities Review*, *4*(2), 141–158.

Cook, M., Colton, S., Pease, A., & Llano, M. T. (2019). Framing in computational creativity-a survey and taxonomy. *ICCC*, 156–163.

Crawford, K. (2021). *The atlas of ai: Power, politics, and the planetary costs of artificial intelligence*. Yale University Press.

Crawford, K., & Paglen, T. (2019). Excavating AI: The Politics of Images in Machine LearningTraining Sets. *AI and Society*.

Daudén Roquet, C., & Sas, C. (2021). Interoceptive interaction: An embodied metaphor inspired approach to designing for meditation. *Proceedings of the 2021 CHI Conference on Human Factors in Computing Systems*. https://doi.org/10.1145/3411764.3445137

Deng, J., Dong, W., Socher, R., Li, L.-J., Li, K., & Fei-Fei, L. (2009). Imagenet: A large-scale hierarchical image database. *2009 IEEE Conference on Computer Vision and Pattern Recognition*, 248–255. https://doi.org/10.1109/CVPR.2009.5206848

Denton, E., Hanna, A., Amironesei, R., Smart, A., & Nicole, H. (2021). On the genealogy of machine learning datasets: A critical history of ImageNet. *Big Data & Society*, *8*(2), 205395172110359. https://doi.org/10.1177/20539517211035955

Devendorf, L., & Rosner, D. K. (2015). Reimagining digital fabrication as performance art. *Proceedings of the 33rd Annual ACM Conference Extended Abstracts on Human Factors in Computing Systems*, *555–566*.

D'ignazio, C., & Klein, L. F. (2020). *Data feminism*. MIT press.

Dodge, S., & Karam, L. (2017). A study and comparison of human and deep learning recognition performance under visual distortions. *2017 26th international conference on computer communication and networks (ICCCN)*, 1–7.

Dosovitskiy, A., Beyer, L., Kolesnikov, A., Weissenborn, D., Zhai, X., Unterthiner, T., Dehghani, M., Minderer, M., Heigold, G., Gelly, S., et al. (2020). An image is worth 16x16 words: Transformers for image recognition at scale. *arXiv preprint arXiv:2010.11929*.

Dove, G., Halskov, K., Forlizzi, J., & Zimmerman, J. (2017). UX design innovation: Challenges for working with machine learning as a design material. *Conference on Human Factors in Computing Systems - Proceedings*, *2017-May*, 278–288. https://doi.org/10.1145/3025453.3025739

Epstein, Z., Hertzmann, A., Herman, L., Mahari, R., Frank, M. R., Groh, M., Schroeder, H., Smith, A., Akten, M., Fjeld, J., Farid, H., Leach, N., Pentland, A., & Russakovsky, O. (2023). Art and the science of generative AI: A deeper dive [arXiv:2306.04141 [cs]]. *Science*, *380*(6650), 1110–1111. https://doi.org/10.1126/science.adh4451

Epstein, Z., Levine, S., Rand, D. G., & Rahwan, I. (2020). Who Gets Credit for AI-Generated Art? *iScience*, *23*(9), 101515. https://doi.org/10.1016/j.isci.2020.101515

Friedman, B. (1996). Value-sensitive design. *interactions*, *3*(6), 16–23.

Gaver, W. W., Beaver, J., & Benford, S. (2003). Ambiguity as a resource for design. *Conference on Human Factors in Computing Systems - Proceedings*, (*5*), 233–240.

Goodfellow, I., Pouget-Abadie, J., Mirza, M., Xu, B., Warde-Farley, D., Ozair, S., Courville, A., & Bengio, Y. (2020). Generative adversarial networks. *Commun. ACM*, *63*(11), 139–144. https://doi.org/10.1145/3422622

Grba, D. (2021a). Brittle opacity: Ambiguities of the creative ai. *xCoAx, 9th Conference on Computation, Communication, Aesthetics & X*. https://doi.org/10.5281/ZENODO.4663408

Grba, D. (2021b). Information particles: Tracing the ambiguities of the creative ai. *Proceedings of Art Machines 2: International Symposium on Machine Learning and Art 2021*, 2021–05.

Grba, D. (2022). Deep else: A critical framework for ai art. *Digital*, *2*(1), 1–32. https://doi.org/10.3390/digital2010001

Gu, L., & Li, Y. (2022). Who made the paintings: Artists or artificial intelligence? the effects of identity on liking and purchase intention. *Frontiers in Psychology*, 4605.

Guljajeva, V., & Canet Sola, M. (2022). POSTcard Landscapes from Lanzarote. *Creativity and Cognition*, 634–636. https://doi.org/10.1145/3527927.3531191

Hazzard, A., Greenhalgh, C., Kallionpaa, M., Benford, S., Veinberg, A., Kanga, Z., & McPherson, A. (2019). Failing with style: Designing for aesthetic failure in interactive performance. *Proceedings of the 2019 CHI Conference on Human Factors in Computing Systems*, 1–14. https://doi.org/10.1145/3290605.3300260

Hertzmann, A. (2020). Visual Indeterminacy in GAN Art. *Leonardo*, *53*(4), 424–428. https://doi.org/10.1162/leon_a_01930

Ho, J., Jain, A., & Abbeel, P. (2020). Denoising diffusion probabilistic models. *Advances in Neural Information Processing Systems*, *33*, 6840–6851.

Holmquist, L. E. (2017). Intelligence on tap: Artificial intelligence as a new design material. *interactions*, *24*(4), 28–33.

Höök, K. (2000). Steps to take before intelligent user interfaces become real. *Interacting with computers*, *12*(4), 409–426.

Howard, A. (2020). Are we trusting ai too much? examining human-robot interactions in the real world. *Proceedings of the 2020 ACM/IEEE International Conference on Human-Robot Interaction*, 1–1.

Huang, H., Yu, P. S., & Wang, C. (2018). An introduction to image synthesis with generative adversarial nets. *arXiv preprint arXiv:1803.04469*.

Ingimundardottir, E. (2018). Word By Word: A Mobile Game To Encourage Collaborative Storytelling Within The Museum [ISBN: 9780998847405]. *Mw18: Museums and the Web 2018*.

Jorge, C., Nisi, V., Nunes, N., Innella, G., Caldeira, M., & Sousa, D. (2013). Ambiguity in design: An airport split-flap display storytelling installation. *CHI '13 Extended Abstracts on Human Factors in Computing Systems*, 541–546. https://doi.org/10.1145/2468356.2468452

Kang, L., Jackson, S., & Pinch, T. (2022). The electronicists: Techno-aesthetic encounters for nonlinear and art-based inquiry in hci. *Proceedings of the 2022 CHI Conference on Human Factors in Computing Systems*, 1–17.

Karras, T., Aittala, M., Laine, S., Härkönen, E., Hellsten, J., Lehtinen, J., & Aila, T. (2021). Alias-free generative adversarial networks. *Proc. NeurIPS*.

Klingemann, M. (2017). The Butcher's Son. https://www.artsy.net/artwork/mario-klingemann-imposture-series-the-butchers-son

Kulesza, T., Stumpf, S., Burnett, M., Yang, S., Kwan, I., & Wong, W.-K. (2013). Too much, too little, or just right? ways explanations impact end users' mental models. *2013 IEEE Symposium on visual languages and human centric computing*, 3–10.

Kuniavsky, M., Churchill, E., & Steenson, M. (2017). Designing the user experience of machine learning systems. *AAAI Spring Symposium Proceedings (Technical Report SS-17-04)*, 27–29.

Larrazabal, A. J., Nieto, N., Peterson, V., Milone, D. H., & Ferrante, E. (2020). Gender imbalance in medical imaging datasets produces biased classifiers for computer-aided diagnosis [Publisher: Proceedings of the National Academy of Sciences]. *Proceedings of the National Academy of Sciences*, *117*(23), 12592–12594. https://doi.org/10.1073/pnas.1919012117

Leahu, L. (2016). Ontological Surprises: A Relational Perspective on Machine Learning. *Proceedings of the 2016 ACM Conference on Designing Interactive Systems*, 182–186. https://doi.org/10.1145/2901790.2901840

LeCun, Y., Bengio, Y., et al. (1995). Convolutional networks for images, speech, and time series. *The handbook of brain theory and neural networks*, *3361*(10), 1995.

Lubart, T. (2005). How can computers be partners in the creative process: Classification and commentary on the special issue. *International journal of human-computer studies*, *63*(4-5), 365–369.

Malevé, N. (2021). On the data set's ruins. *AI & SOCIETY*, *36*(4), 1117–1131. https://doi.org/10.1007/s00146-020-01093-w

Mateas, M., & Stern, A. (2003). Facade: An Experiment in Building a Fully-Realized Interactive Drama. *Game Developers Conference. Game Design Track*.

Maybury, M., & Wahlster, W. (1998). *Readings in intelligent user interfaces*. Morgan Kaufmann.

Mazzone, M., & Elgammal, A. (2019). Art, creativity, and the potential of artificial intelligence. *Arts*, *8*(1), 26.

Milan, S., & Treré, E. (2019). Big data from the south (s): Beyond data universalism. *Television & New Media*, *20*(4), 319–335.

Mordvintsev, A., Olah, C., & Tyka, M. (2015). Inceptionism: Going deeper into neural networks. https://research.googleblog.com/2015/06/inceptionism-going-deeper-into-neural.html

MUNCH, Sivertsen, C., & Mathias, N. (2021). Poison. https://www.munchmuseet.no/en/exhibitions/archive/2021/poison/

Murray-Browne, T., & Tigas, P. (2021). Emergent interfaces: Vague, complex, bespoke and embodied interaction between humans and computers. *Applied Sciences*, *11*(18), 8531.

Murray-Rust, D., Nicenboim, I., & Lockton, D. (2022). Metaphors for designers working with ai.

Narvaez, M. F., Olivenza, I. S., & McGowan, N. (2022). Painting authorship and forgery detection challenges with ai image generation algorithms: Rembrandt and 17th century dutch painters as a case study. *IJIMAI*, *7*(7), 7–13.

Obvious. (2018). Edmond De Belamy. https://obvious-art.com/portfolio/edmond-de-belamy/

Oppenlaender, J. (2022). The creativity of text-based generative art. https://doi.org/10.48550/ARXIV.2206.02904

Paglen, T., & Crawford, K. (2019). ImageNet Roulette. Retrieved October 14, 2023, from https://paglen.studio/2020/04/29/imagenet-roulette/

Penny, S. (2009). Art and artificial life–a primer.

Penny, S. (2013). Art and robotics: Sixty years of situated machines. *AI & society*, *28*, 147–156.

Ramesh, A., Dhariwal, P., Nichol, A., Chu, C., & Chen, M. (2022). Hierarchical text-conditional image generation with clip latents. *arXiv preprint arXiv:2204.06125*.

Robinette, P., Li, W., Allen, R., Howard, A. M., & Wagner, A. R. (2016). Overtrust of robots in emergency evacuation scenarios. *2016 11th ACM/IEEE international conference on human-robot interaction (HRI)*, 101–108.

Rombach, R., Blattmann, A., Lorenz, D., Esser, P., & Ommer, B. (2021). High-resolution image synthesis with latent diffusion models.

Russakovsky, O., Deng, J., Su, H., Krause, J., Satheesh, S., Ma, S., Huang, Z., Karpathy, A., Khosla, A., Bernstein, M., et al. (2015). Imagenet large scale visual recognition challenge. *International journal of computer vision*, *115*(3), 211–252.

Ryding, K., Spence, J., Løvlie, A. S., & Benford, S. (2021). Interpersonalizing Intimate Museum Experiences [Publisher: Taylor & Francis]. *International Journal of Human–Computer Interaction*, *37*(12), 1151–1172. https://doi.org/10.1080/10447318.2020.1870829

Saharia, C., Chan, W., Saxena, S., Li, L., Whang, J., Denton, E., Ghasemipour, S. K. S., Ayan, B. K., Mahdavi, S. S., Lopes, R. G., et al. (2022). Photorealistic text-to-image diffusion models with deep language understanding. *arXiv preprint arXiv:2205.11487*.

Sanches, P., Höök, K., Sas, C., & Ståhl, A. (2019). Ambiguity as a resource to inform proto-practices: The case of skin conductance. *ACM Trans. Comput.-Hum. Interact.*, *26*(4). https://doi.org/10.1145/3318143

Sanches, P., Howell, N., Tsaknaki, V., Jenkins, T., & Helms, K. (2022a). Diffraction-in-action: Designerly explorations of agential realism through lived data. *Proceedings of the 2022 CHI Conference on Human Factors in Computing Systems*. https://doi.org/10.1145/3491102.3502029

Sanches, P., Howell, N., Tsaknaki, V., Jenkins, T., & Helms, K. (2022b). Diffraction-in-action: Designerly explorations of agential realism through lived data. *Proceedings of the 2022 CHI Conference on Human Factors in Computing Systems*, 1–18.

Scurto, H., Caramiaux, B., & Bevilacqua, F. (2021). Prototyping machine learning through diffractive art practice. *Designing Interactive Systems Conference 2021*, 2013–2025.

Sengers, P., Boehner, K., David, S., & Kaye, J. (2005). Reflective design. *Proceedings of the 4th decennial conference on Critical computing: between sense and sensibility*, 49–58.

Sengers, P., & Gaver, B. (2006). Staying open to interpretation: Engaging multiple meanings in design and evaluation. *Proceedings of the 6th Conference on Designing Interactive Systems*, 99–108. https://doi.org/10.1145/1142405.1142422

Shaker, N., Togelius, J., & Nelson, M. J. (2016). *Procedural content generation in games*. Springer.

Sivertsen, C., Smith, M., & van der Zwan, S. (2023). Art Critique by Other Means. *Designing Interactive Systems Conference (DIS '23)*, 10. https://doi.org/10.1145/3563657.3596069

Song, K. W., & Paulos, E. (2021). Unmaking: Enabling and celebrating the creative material of failure, destruction, decay, and deformation. *Proceedings of the 2021 CHI Conference on Human Factors in Computing Systems*, 1–12.

Stark, L., & Crawford, K. (2019). The work of art in the age of artificial intelligence: What artists can teach us about the ethics of data practice. *Surveillance & Society*, *17*(3/4), 442–455.

Tenhaaf, N. (2008). Art embodies a-life: The vida competition. *Leonardo*, *41*(1), 6–15.

Vallgårda, A., & Fernaeus, Y. (2015). Interaction design as a bricolage practice. *Proceedings of the ninth international conference on tangible, embedded, and embodied interaction*, 173–180.

Vayanou, M., Ioannidis, Y., Loumos, G., & Kargas, A. (2019). How to play storytelling games with masterpieces: From art galleries to hybrid board games. *Journal of Computers in Education*, *6*(1), 79–116. https://doi.org/10.1007/s40692-018-0124-y

Vetere, F., & Wouters, N. (2018). AI scans your data to assess your character but Biometric Mirror asks: What if it is wrong? *The Conversation*. Retrieved September 15, 2022, from http://theconversation.com/ai-scans-your-data-to-assess-your-character-but-biometric-mirror-asks-what-if-it-is-wrong-100328

Villareale, J., Harteveld, C., & Zhu, J. (2022). " i want to see how smart this ai really is": Player mental model development of an adversarial ai player. *Proceedings of the ACM on Human-Computer Interaction*, *6*(CHI PLAY), 1–26.

Voulodimos, A., Doulamis, N., Doulamis, A., & Protopapadakis, E. (2018). Deep learning for computer vision: A brief review. *Computational intelligence and neuroscience*, *2018*.

Wakkary, R. (2020). Nomadic practices: A posthuman theory for knowing design. *International Journal of Design*, *14*(3), 117.

Wang, A., Liu, A., Zhang, R., Kleiman, A., Kim, L., Zhao, D., Shirai, I., Narayanan, A., & Russakovsky, O. (2022). REVISE: A Tool for Measuring and Mitigating Bias in Visual Datasets. *International Journal of Computer Vision*, *130*(7), 1790–1810. https://doi.org/10.1007/s11263-022-01625-5

Wong, J. C. (2019). The viral selfie app ImageNet Roulette seemed fun – until it called me a racist slur. *The Guardian*. Retrieved August 25, 2022, from https://www.theguardian.com/technology/2019/sep/17/imagenet-roulette-asian-racist-slur-selfie

Wouters, N., Kelly, R., Velloso, E., Wolf, K., Ferdous, H. S., Newn, J., Joukhadar, Z., & Vetere, F. (2019). Biometric Mirror: Exploring Ethical Opinions towards Facial Analysis and Automated Decision-Making. *Proceedings of the 2019 on Designing Interactive Systems Conference*, 447–461. https://doi.org/10.1145/3322276.3322304

Xu, M. (2021). An exploratory study of ai creativity inspired by descriptive ambiguity.

Yang, Q., Steinfeld, A., Rosé, C., & Zimmerman, J. (2020). Re-examining whether, why, and how human-ai interaction is uniquely difficult to design. *Proceedings of the 2020 CHI Conference on Human Factors in Computing Systems*, 1–13. https://doi.org/10.1145/3313831.3376301

Yannakakis, G. N., & Togelius, J. (2018). *Artificial intelligence and games* (Vol. 2). Springer.

Yaszdani, N. I. (2018). Machine Bias. https://nushinyazdani.com/Machine-Learning-Bias

Yurman, P., & Reddy, A. V. (2022). Drawing conversations mediated by ai. *Creativity and Cognition*, 56–70. https://doi.org/10.1145/3527927.3531448

Zhu, J. (2009). *Intentional systems and the artificial intelligence (ai) hermeneutic network: Agency and intentionality in expressive computational systems*. Georgia Institute of Technology.

Zhu, J., Villareale, J., Jan, H. C., Löwe, M., & Harteveld, C. (2021). Player-AI Interaction : What Neural Network Games Reveal About AI as Play. *Proceedings of the 2021 ACM Conference on Human Factors in Computing Systems (CHI '21)*.

Zylinska, J. (2020). *AI Art: Machine Visions and Warped Dreams*. Open Humanites Press. Retrieved September 10, 2022, from http://www.openhumanitiespress.org/books/titles/ai-art/