Objective Portrait

A practice-based inquiry to explore AI as a reflective design partner

van der Burg, Vera; de Boer, Gijs; Akdag Salah, Amila; Chandrasegaran, Senthil; Lloyd, Peter

DOI
10.1145/3563657.3595974

Publication date
2023

Document Version
Final published version

Published in
DIS' 23 Proceedings of the ACM Designing Interactive Systems Conference

Objective Portrait

A practice-based inquiry to explore AI as a reflective design partner

Vera van der Burg
TU Delft, The Netherlands
v.vanderburg@tudelft.nl

Gijs de Boer
Design Academy Eindhoven,
The Netherlands
info@supergijs.com

Almila Akdag Salah
Utrecht University, The Netherlands
a.a.akdag@uu.nl

Senthil Chandrasegaran
TU Delft, The Netherlands
r.s.k.chandrasegaran@tudelft.nl

Peter Lloyd
TU Delft, The Netherlands
P.A.Lloyd@tudelft.nl

ABSTRACT

Artificial intelligence (AI) is increasingly being viewed as a creative partner rather than as a tool. How to design such collaborations is still a subject of speculation. In this pictorial, we propose a collaborative role for AI to prompt self-reflection. We explore this through a practice-based inquiry of whether and how AI could help a designer reflect on and relate to their own work. Three designers annotate a collection of images representing their fascinations, with subjective labels, indicating different dimensions of their visual concepts. These labels are used to teach an object detection model the designers' perspectives. Then, they used this trained model on their own design work to evaluate the AI's potential to prompt self-reflection. By describing this process of AI-training we explore how an AI can help us become aware of our own implicit perspectives.

Authors Keywords

Artificial Intelligence, Reflection, Collaboration

DIS '23, July 10–14, 2023, Pittsburgh, PA, USA

ACM ISBN 978-1-4503-9893-0/23/07.
https://doi.org/10.1145/3563657.3595974

The project showcased in this pictorial was exhibited as an installation during Dutch Design Week 2022. This pictorial provides a comprehensive description of the research that underpins the installation and serves as a supplement to the work displayed at the exhibition

READING AN IMAGE

If you look at this image on the right, what do you see? Hands? An egg? Creases in a sheet? If you look longer, deeper, do you notice anything else? Do you wonder what the image means? Do you notice irregularities, do you find things odd, or even uncomfortable? Or does it give you a specific feeling? A feeling of joy, perhaps? Or are you rather enticed? What actually pulls you in, and what pushes you away? What elements attract, feel warm, and what feels rather distanced and cold?

And then, if we zoom in, would you be able to point out exactly which element makes you feel this way? How do you identify instances in this image? Does that reveal something about how you read the image? This hand below, is it uncanny? Or is it sensual?

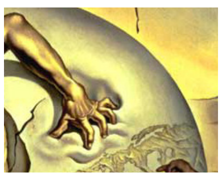

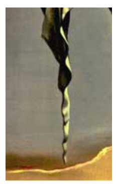

And this little twirl on the right, does it make you laugh? Do you want to pull it, or make it go away? Does it add to the image, or do you think it could be erased? What could it represent? Is it stubborn, teasing, or merely obsolete?

Or this crease in the sheet, do you want to straighten it, therefore, does the discomfort make you want to touch it? What word would you give this part of the image, if you would have to describe it?

Would it help you if you had more insight in how you look at an image?

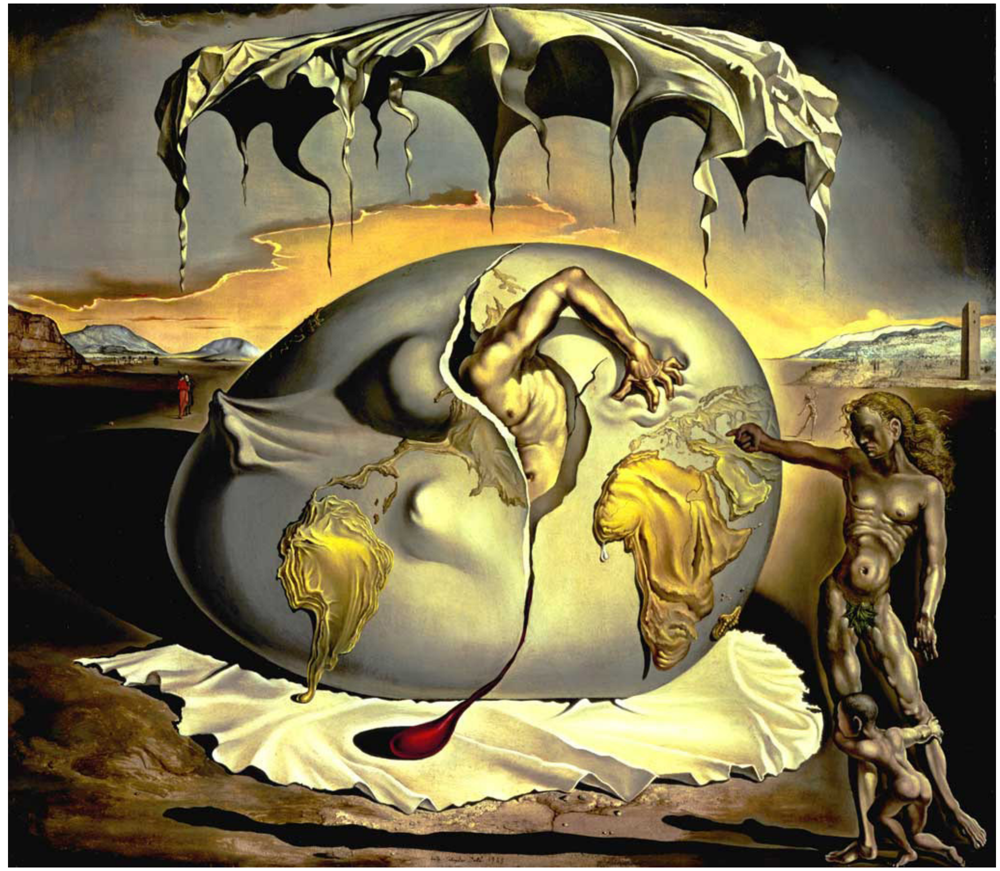

DESIGN FASCINATIONS

Our perspectives on images are deeply personal and nuanced. Our 'ways of seeing' [1] condition the feelings we get from gazing at images, and what aspects of these images we choose to concentrate on: "when an image is presented as a work of art, the way people look at it is affected by a whole series of learnt assumptions about art" [1, p. 10]. While this statement is made about people in general, it is especially true for designers, who construct meanings from visual information by recognising patterns and exploring possibilities beyond what is shown [29], which may reflect a combination of their personal and professional 'learnt assumptions'. A designer's way of looking at things is influenced by their perspectives and fascinations. These perspectives and fascinations can shape the design worlds they construct, as described by Schön [28]. Design worlds are the environments built by designers when they start the design process and represent design knowledge. These worlds are built upon personal beliefs, specific interpretations of things, spatial arrangements, and relationships between elements in their projects. Designers' fascinations play a role in these design worlds and are often seen as tacit knowledge [34]. Many of the decisions made in projects stem from this tacit knowledge, such as the selection of project topics, desired design contexts, and the visual style of objects they create. Can we gain deeper insight into our personal design worlds and fascinations by making them explicit and teaching them to an AI? Can we teach a machine our 'ways of seeing'? If we can spark the AI's interest in our interests, might it shed light on our own perspectives and thought processes as designers? By teaching a machine about our design fascinations, could we get closer to our own design worlds?

AI AS A REFLECTIVE PARTNER

The integration of AI technology into creative design processes is a rapidly growing area of research and development. Collaboration between designers and AI systems has been commonly discussed in literature over the past years (e.g. [5, 30, 31]) as a strategy to enhance the creative work of designers and artists. The recent adoption of generative models, including DALL-E [23], ChatGPT [24], and Midjourney [18] in the creative industries is evidence of a strong interest in these tools in creative practice, as their outputs now win art prizes [33] and photography contests [35]. This increased interest, however, still leaves us with many unanswered questions, such as what implications the use of these tools could hold for the artist or designer's process. While these AI tools can generate fast and high-quality output, such as images or text, they appear to automate a single isolated step in a creative process. Specifically, they generate digital content based on textual prompts, offering designers a multitude of outputs. The increased use of generative AI happens primarily in specific phases of a creative process; generative phases where a range of ideas or design options is being created, rather than in the reflective phases of a design project where those options are being critically evaluated and contextualized. We are interested in exploring a different collaborative role for AI, one that goes beyond quick generation, and instead provides insight into how we as designers work, what we make, and how we think. We thus explore the potential of AI to help a designer become aware of and reflect on their relation to their own work, therefore potentially prompting self-reflection.

THE AI-MIRROR

Developments in powerful computational approaches to classify, categorize, and label data have created a false sense of 'objectivity' often associated with computational thinking. Recent research has seen a pushback against such thinking, emphasizing the inherently interpretive nature of actions leading to categorizing and labeling. Tanweer et al. [32] provide several examples where disregarding context in interpretation leads to incorrect understanding of data which in turn leads to incorrect generalizations. Instead, they make an argument for interpretivist approaches—to "*probe the multiple and contingent ways that meaning is ascribed to objects, actions, and situations*" (p. 5). This is not an isolated take. New ways are being proposed to integrate qualitative methods to augment computational approaches (e.g., Baumer et al. [4]) or vice versa (e.g. Nelson [21], Pääkkönen and Tlikosi [25]). The work we present here is inspired by these approaches: we explore how the process of training AI models might play a part in the designer's self-reflection. We take the inherent subjectivity of AI not as something to eliminate, but as a starting point. By involving the AI training phase as part of the designer's process, we expand the use of AI beyond its output, creating a more intimate approach that slows down and stretches the use of AI in a creative process, while also helping designers gain a deeper understanding of themselves.

In our explorations, we introduce the metaphor of AI as a mirror, a device that provides a mediated perspective on oneself. This metaphor sparks new ideas on how to incorporate AI in our practice [20]. The reflection witnessed in this AI-mirror would be shaped by the choices we made in building it. If the AI-mirror is specifically customized to align with personal values, opinions, aesthetic preferences, and desires, could it offer valuable insights through the reflection it presents?

To investigate the potential of AI in supporting reflection, we adopted a first-person, practice-based research approach, focusing on the material qualities of AI and Machine Learning. This approach allowed us to explore the complexities of integrating this technology into actual design practice, particularly for studying the subjective concept of self-reflection [17]. Furthermore, this approach aligns with the argument made by Benjamin et al. [36] emphasizing the need for fundamental research on the use of ML technologies as a design material in their own right [36, p. 3]. Through design inquiry and artistic exploration, we embraced AI's limitations (such as subjectivity, bias, and small datasets) and transformed them into opportunities for creative practice. By embracing the constraints of the technology, we aimed to establish innovative designer-AI partnerships [13].

HOW TO READ THIS PICTORIAL

This pictorial traces the AI-collaboration process of three designers, two of whom are also coauthors of this pictorial. The designers will from now on be referred to as 'we'/'us'. We trained an object detection algorithm to learn how we give meaning to an image, and used it to analyze a composition of our own design work, which we called object portraits. In other words, we tried to capture our subjective way of looking at the world, and then mirror that gaze back at our own work. Each of us followed a cycle of data collection, annotation, and model training/testing, similar to standard object detection model training. We describe our personal paths, the steps of teaching AI about our fascinations, and the questions that arose during this process in the following sections. We named the project

Objective Portrait, both to capture how we tried to portray our practice through objects, and to provoke the reader to reflect on the apparent objectivity of an algorithmic interpretation of a portrait.

It is important to mention that our description of AI in the context of this research project is confined to the Object Detection Model we used, which was YOLOv5 [14], a deep learning model in computer vision that is used to detect objects in images or videos. This model identifies and classifies separate objects that are part of visual scenes.

1. COLLECTING A DATASET

Can fascinations turn into AI worldviews?

Generally put, AI systems are trained to learn from data in order to make predictions or to perform a specific, predefined task. The process of training an AI system involves exposing it to a large dataset of examples in which the system tries to establish a pattern. In the case of Object Detection models, these examples are collections of images with corresponding labels. The AI ‘learns’ visual patterns in the training dataset and associates them with the labels. Such a model can then be applied to unlabelled images where it can ‘recognise’ the labels that may be associated with parts of the images. To teach the algorithm our own subjective view we started with collecting a training dataset that visually represented an important fascination in our making practice. They were collections of images of which it could be expected that we definitely would have opinions about them, experiences with, or have feelings towards that were relevant to us in our own work. As the model we used in our explorations was only able to recognize separate instances in visual scenes, the images in our datasets depicted visual scenes as well.

In the following section, you can see a selection of our training datasets, accompanied with our motivation for selecting these types of images, and the specifications of the dataset.

Dataset of Vera: Surrealistic Art

“Collecting a dataset is a tedious task, because you need a lot images. However, collecting a bunch of images that fascinate you feels actually less tedious, because I really enjoy looking at them. It feels a bit like creating an archive of your interests, which I as a designer tend to do anyways in my practice, for inspiration” - Vera




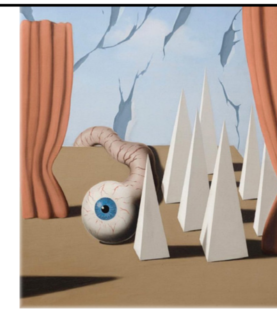
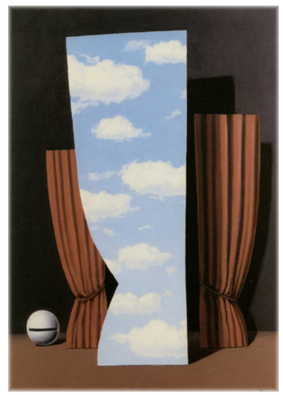


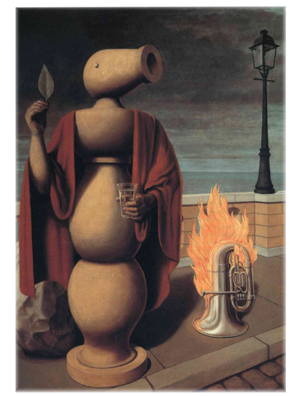

Dataset specifications
Images: Paintings of artists like Salvador Dali, Rene Magritte, among others.
Amount: 378 images
Sourced from: Google Image Search

“In my making practice, I intuitively create and collects objects that tend to have surrealistic, strange shapes. To understand what could be hidden behind this urge or implicit attraction, I used an image dataset of surrealistic paintings by René Magritte and Salvador Dali.”

Dataset of Gijs: Gardens of the World

Dataset specifications

Images: Photographs of gardens
Amount: 300 images
Sourced from: *Gardens of the world* [11].

"I am a design researcher and writer. I study human-plant relations, looking for modes of care that don't rely on control. To train my algorithm I used an image dataset of gardens, as a garden embodies the tension between care (letting things grow) and the control of giving shape."

Dataset of Maurik: Art in Public Space

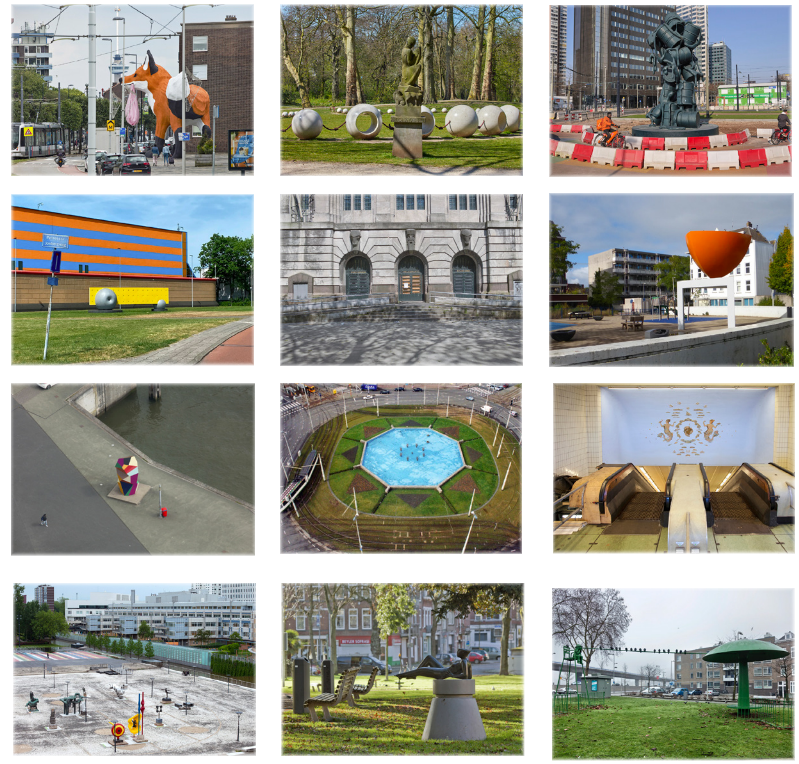

Dataset specifications

Images: Photographs of artworks in public space in Rotterdam
Amount: 120 images
Sourced from: Beeldende Kunst & Openbare Ruimte website [2]

"I am a designer whose work aims to give citizens agency over public space. I experience conflicting feelings about the way public space is designed from a top-down approach, and appreciating public space as free commons. I was curious to learn how I actually perceive public space, so I used a dataset of art in public space in Rotterdam."

2. CHOOSING SUBJECTIVE LABELS:
How to categorize our subjective perspective?

State-of-the-art object detection models are trained on large datasets containing over thousands of images that are annotated by human annotators. Data annotation is a process in which annotators mark relevant elements in a dataset to train models for tasks like recognizing or classifying images. To make sure annotation is done relatively quickly, multiple annotators are used for this task, as annotation can be a tedious and time consuming task. For our exploration, however, we were looking to capture the subjective gaze of only one person, therefore only one person was involved in annotating the machine's training dataset.

Only a small dataset allowed for a machine annotation workflow that was manageable for one person to execute, and facilitated the 'quick' prototyping-nature of our explorations. We carefully collected and curated the datasets ourselves, which gave us a sense of agency, sensibility and control over them. Additionally, we needed to strike a balance between what the system needed to function (for a model to work, it needs a large amount of data) and what was manageable for us as annotators to work with. Therefore, we held an average of 300 images per dataset.

To teach our personal algorithms something about our subjective gaze, we needed to come up with special labels to annotate the training dataset with. These labels were the prerequisites of what the algorithm would detect, but did not describe what 'objectively' was to be seen in an image. Rather, they were words that captured subjective interpretations:, feelings, opinions or judgments that we felt when faced with the image.

We limited the amount of labels we each used to four, such that every label would still have a reasonable amount of instances in such a small dataset. To choose the four labels, we used a representation similar to existing two-dimensional representations of emotional affect (e.g. Russell's circumplex model [26]) that typically have emotional valence (positive/negative) and arousal (high/low) on orthogonal axes to form four quadrants. We discussed the interpretation of our set of quadrants by examining a subset of around 10 images from the training dataset that we wanted to annotate, and arrived at the questions: "does this element in this image pull me in or pushes me away?" and "does this element feel 'warm' or 'cold'?" Based on these questions, we marked the x-axis as 'push/pull' and the y-axis as 'hot/cold' (see Figure 1). This allowed us to organize specific feelings and judgments we had towards the dataset into two sets of opposites, to make sure the labels did not overlap. For example, a label like "happy" might be too similar with a label like "content", therefore difficult to distinguish when labeling the training dataset.

Then, we tried to give more specific words to those feelings of "hot/cold" and "push/pull" roaming those two axes:"How would I describe, or label, this feeling?"

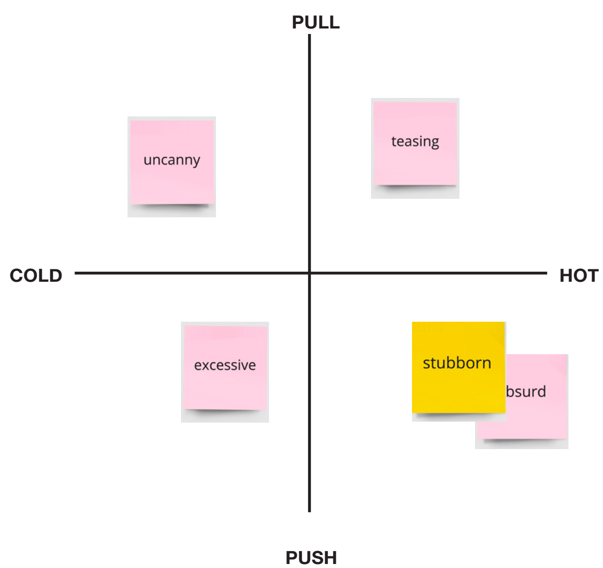

Figure 1: The quadrant tool to help decide on the subjective labels

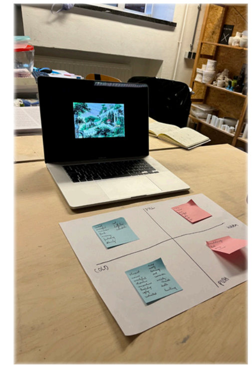
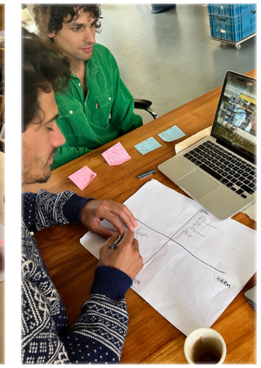

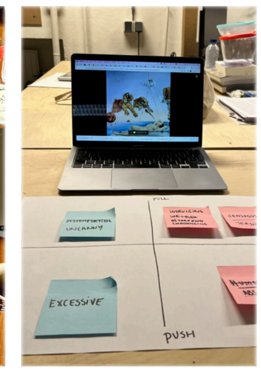

LABELS:
Vera:
Uncanny, Teasing, Excessive, Absurd
Gijs:
Glorious, Filthy, Lonely, Luring
Maurik:
Creepy, Inviting, Excluding. Stubborn

3. ANNOTATING THE DATASET
How to annotate a fascination?

Since the model we used (YOLOv5) is trained on scenes within which it identifies objects and features, we had to appreciate each image in our training data as a scene as well. To annotate this data, we thus identified parts of these 'scenes' that conveyed or represented one of our predefined labels. To annotate the dataset we used Roboflow [10], an online annotation tool that provided us with bounding box tools to segment parts of an image and assign our self-defined labels to them. The coordinates of these labeled bounding boxes could then be used to train our personal model. This annotation process did require some getting used to. A feeling may arise from a whole image, but annotation requires identifying a specific part. To effectively collaborate with the model, we had to shift our perspective and adopt this new way of viewing images. Vera & Gijs reserved three to four hours per day for a duration of four days to fully annotate their dataset. Maurik spent one afternoon for three hours, which resulted in a smaller annotated training dataset (124 images) compared to 300 images each for Vera and Gijs.

"After annotating a lot of images intuitively, you gradually start to understand what you mean with a visual concept like "uncanny", for example, which I tend to see in all the body parts, or flesh-coloured elements of these images" - Vera

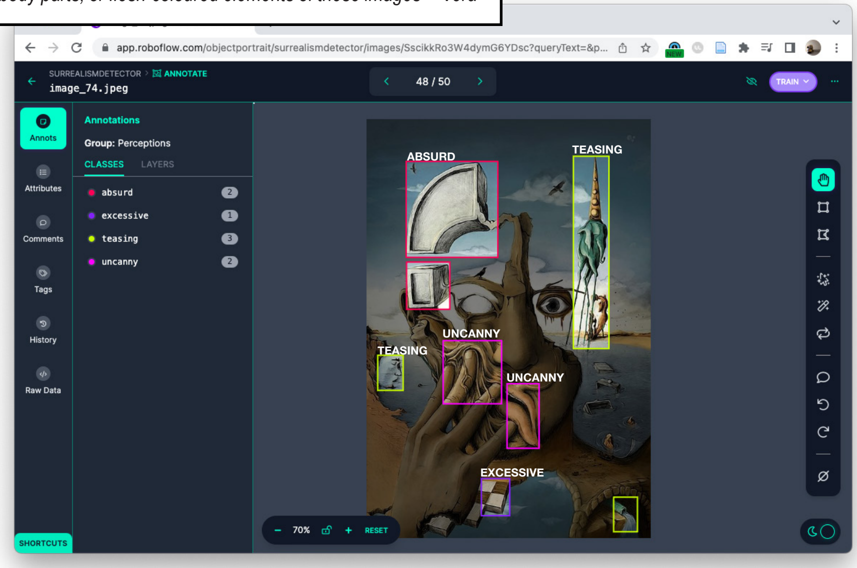

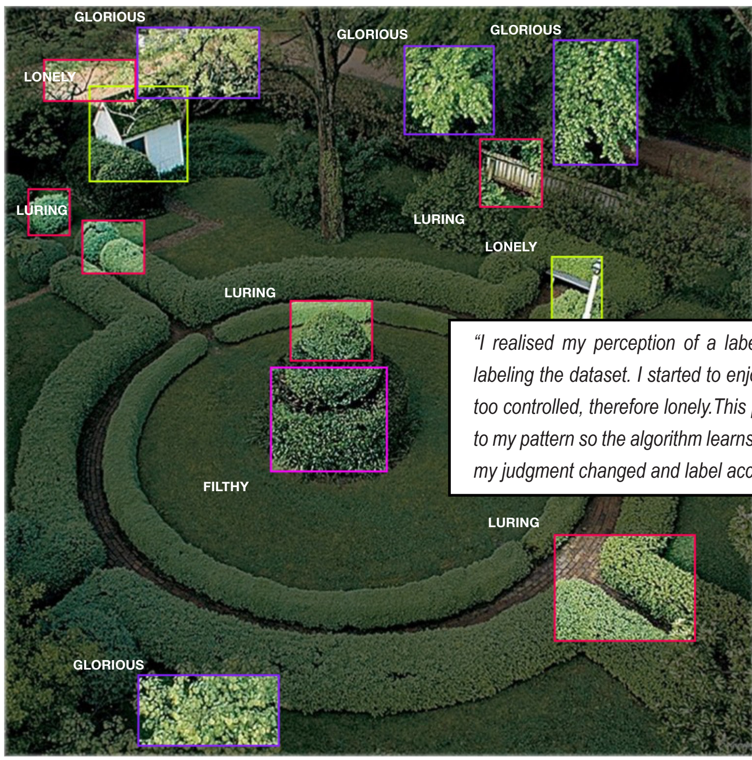

"I realised my perception of a label, like 'lonely', changed while I was labeling the dataset. I started to enjoy the raked gardens that I first found too controlled, therefore lonely. This pushed me into a choice: Do I adhere to my pattern so the algorithm learns better, or do I stay true to the fact that my judgment changed and label accordingly to that?" - Gijs

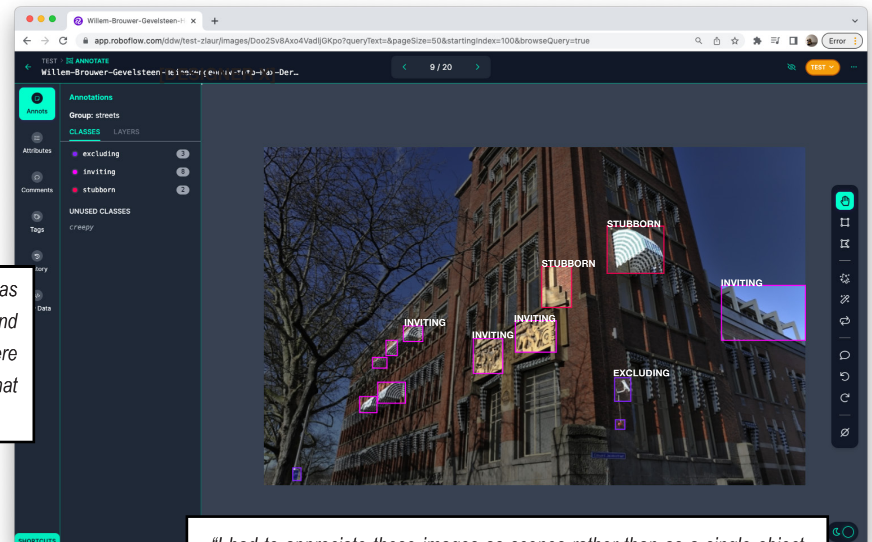

"I had to appreciate these images as scenes rather than as a single object for the algorithm to learn from. It felt unusual to look at an image as a sum of parts, rather than a unified whole. However, if it did open up new ways to look at the image" - Vera

4. TRAINING AND TESTING THE MODEL

How to start a conversation with a subjective algorithm?

Model training

We trained three individual object detection models on labels based on our personal interests and fascinations (surrealistic art, gardens, and public spaces). The object detection models were trained by establishing a pattern in pixels between the annotated images and their corresponding bounding box annotations. This approximately took two hours per model.

Creating an object portrait

After training the models, we confronted the challenge of how we could make them reflect on our own design work. Our models were trained to recognise objects in an image that resembled respectively a surrealist painting, garden, and public space. So instead of simply presenting an existing image of our work, we chose to create a new artwork/design/image specifically for the models to analyze. We crafted an image that represented not only our work but also our practices, providing the model with a glimpse into our creative identity. We created *object portraits*, which were video recordings of still lifes that were composed in a manner consistent with the models' training datasets. Creating these 'object portraits' was a way for us to portray our practices through showcasing our personal objects. We composed these object portraits through a process we referred to as 'visual linking', which involved incorporating certain stylistic elements from the training dataset into the object portraits to capture the visual style of the training datasets. We experimented with variations in composition and perspective, and recreated indoor and outdoor scenarios.

"My intention was to select objects for the object portrait that I anticipated the algorithm would successfully identify, as that would provide a sense of the algorithm "working". Additionally, I deliberately chose objects that held personal emotional significance for me, as I was intrigued to see if the algorithm would also recognize them as stemming from my own fascinations, as I taught that to the model" - Vera

Iteration 1

Iteration 2

Final version Object Portrait - Vera

"I handpicked several images from the training dataset that I considered to be "typical" Magrittes, using them as a foundation for the composition choices in the Object portrait. Many of these images played with perspective and shadow effects. To establish a visual connection, I aimed to replicate those elements within my object portrait. Working in a photo studio gave a lot of creative freedom to experiment with layering and perspective" - Vera

Iteration 1

Iteration 2

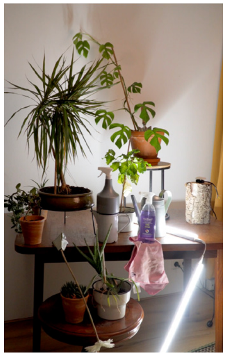

Iteration 3

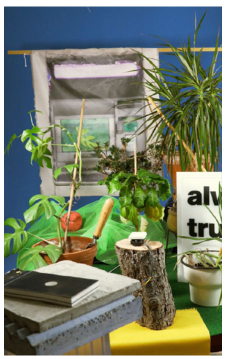

Iteration 4

Final Version Object Portait - Gijs

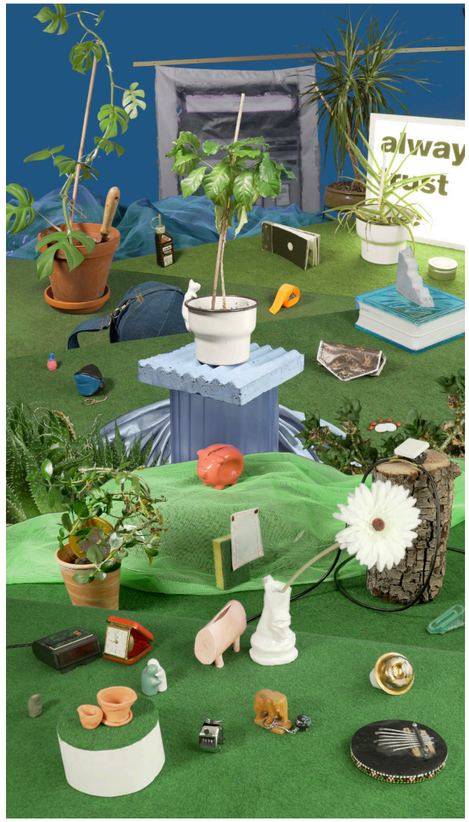

For Maurik's Object Portrait, we decided to create it outdoors since his dataset primarily consisted of outdoor images. We carefully chose objects that not only reflected Maurik's style but also had a connection to the training dataset, allowing them to harmonize with an outdoor setting. - Gijs, Vera & Maurik

Final Version Object Portait - Maurik

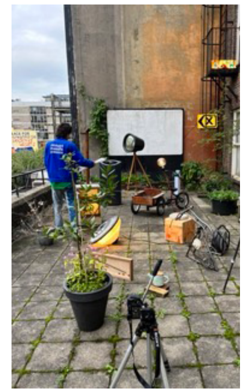

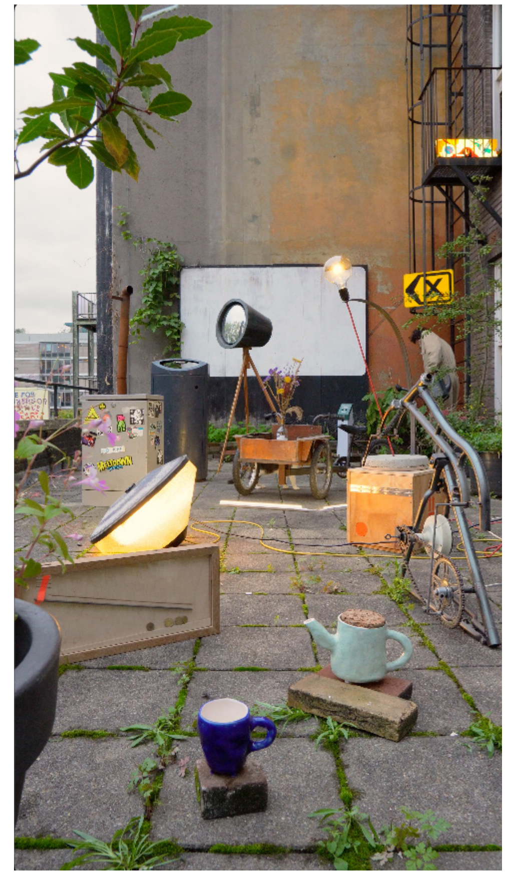

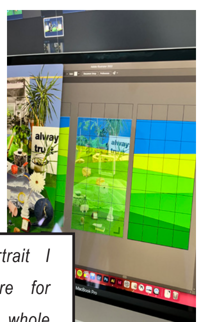

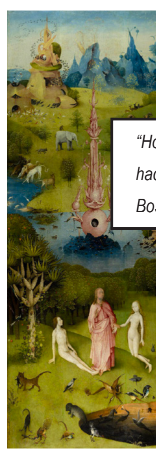

"In making the object portrait I encountered my own desire for control, wanting to plan the whole shoot beforehand. How to let go? The structure of the painting provided an answer. I could plan the different layers, but for every layer improvise where to place objects" - Gijs

"How to make a garden-like image with my objects? I had to think of the isolated creatures in Hieronymous Bosch' Garden of Eden." - Gijs

> *"I was left feeling somewhat unsatisfied because the model did not recognize every object that I had chosen to display. This left me with questions: Were the objects I selected not truly reflective of my assumed fascination? Or perhaps the objects simply weren't surreal enough for the model to detect?" - Vera*

Object Portait of Gijs

Object Portrait of Maurik

Object Portrait of Vera

Model Testing

After the model was trained, we used it to analyze our object portraits, of which you can see screenshots in the images above. The model was capable of identifying a significant number of objects and assigning the subjective labels to them along with a confidence score expressed as a number between 0 and 1. This score referred to the accuracy of the model's predictions. It indicated how confident the model was that a particular object in the image was correctly classified and located within a bounding box. What does a score of 0.86 luring, then, mean? From a technical perspective, a high score meant that the object within the bounding box was confidently associated with the pattern that the model learned from the training dataset. Our main question of this research project was not about the technical accuracy of the prediction, but rather whether the prediction, regardless of score, would assist us in reflecting on the objects that we had selected. Was this the reflection we were looking for? Did we agree with the model's predictions?

DISCUSSION

With this pictorial we illustrated the process of training three individual object detection algorithms based on our subjective view, and utilizing these algorithms to evaluate our design work. The aim was to use AI as a reflection partner and understand the reflective impact of teaching our subjective view to a machine. In the sections below, we describe the insights we got from how our algorithms helped us reflect on our work.

Reflection is in the annotation rather than only in using the AI

Although it was intriguing to see the AI label some of the objects in the final composition with probability scores (see image 1a and 2a for an example), the implications of these predictions alone were unclear. The predictions, however, invited us to go back into our training datasets and reflect on how we labeled them, to understand where the predictions came from. Reflecting on the annotation task, we concluded that it was challenging to define concepts like “absurd” in specific parts of the images.

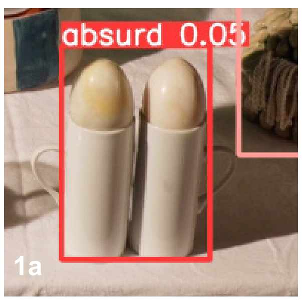

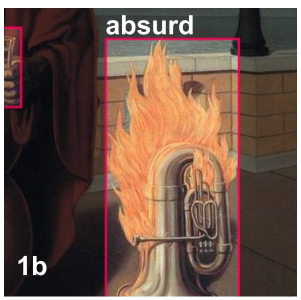

However, as we used the labels repeatedly, following our intuition helped us gradually develop an understanding of what we meant visually with a concept like “absurd”. For instance, Vera could not fully explain what she meant by “absurd” at the beginning of the annotation. At times, it was experienced as a ‘joke’, which also overlapped with the label “teasing.” Gijs considered well-raked and organized gardens “lonely” (see image 2b) and blooming flowers “glorious”.

In the end, Vera realized that she saw all body parts (in the context of surrealistic art) as “uncanny” and “absurd” representing a visual joke (see image 1b for a labeled example from a training dataset). This process of ‘teaching’ the machine our subjective associations through the labeling of images helped us understand our own visual concepts and we started to notice patterns in our intuitive association.

Subjective judgment is dynamic while data annotation is static

Gijs found his judgment changed during the labeling process, initially labeling something “lonely” (see image 3a) but later considering it “luring”. This likely meant that Gijs started finding the well-raked gardens appealing during the annotation task. As the model needs a ‘pattern’ to learn from, Gijs had to choose between staying consistent for the model, or being honest towards his subjective feeling that gradually changed without starting over with his annotation task. In his case, he chose the second and started to label those instances ‘luring’ (see image 3b for an example). This resulted in confusion for the model, as the pattern became more inconsistent in that category.

This inner conflict highlights the challenge of incorporating the fluid nature of human subjectivity in AI model training, as well as the annotator’s sense of responsibility, in this case, to instruct the algorithm accurately. This issue is also called annotation drift [6], a term that acknowledges the variable nature

of the annotation task and the difficulty to provide technical support for that. Balancing subjectivity with technical constraints was something we all as annotators experienced. In some cases, a switch was made as shown in the ‘lonely/luring’ example, and in some cases we aimed to maintain consistency in our annotation tasks by adhering closely to the patterns we had initially annotated. While this reveals challenges for applications that want to capture subjective judgment (i.e. teach a system affective properties in art through affective tagging [16] or labeling of intuitive feelings obtained by images done by a group of annotators [22]), in our case this friction was productive. It helped reflect not only on what our judgments were, but exactly when they changed.

Creating the Object Portraits evolved into a reflective journey of its own.

To create the object portrait, we carefully selected objects using several criteria. One basic criterion was the technical capabilities of the model to ensure that the objects we select could be detected by the model. The more important criteria involved our curiosity, our attachment to the objects, and the personal connection that we as designers may have with the objects. We carefully selected objects to portray that represented us and to which we felt a strong attachment - whether they were objects we had created ourselves or cherished items in our collections. We were driven by curiosity, wondering if the model would recognize these objects as manifestations of the fascinations that served as the foundation for our training. Since we had already annotated our datasets and had an idea of what the model could detect, we intentionally chose objects that sparked a question: “What would the AI perceive in this object that is intimately connected to me?” This curiosity intrigued us.

For example, Gijs specifically pondered how the model would interpret an object that he felt he had created with excessive control. If the model’s interpretation differed, what would that reveal about Gijs himself - his style, fascinations, or even his judgment? It was a thought-provoking exploration of ourselves through the eyes of the AI model.

The model’s predictions about our design work were not challenging enough to prompt deep self-reflection

We were not completely satisfied with the results of the AI in facilitating self-reflection when it analyzed our design work. At times, we understood the AI’s evaluation (see 4a), but to truly understand the prediction we had to return to the dataset to examine our annotation behavior. Perhaps, the potential for the AI-predictions to promote self-reflection could be enhanced by exploring different labels for annotation. Did we need more stimulation or challenge from the AI’s predictions to spark that self-reflection we were looking for? Would it be more effective if the labels we used for annotation were more assertive and expressed strong opinions? Instead of relying on emotional judgment, would it be more impactful if we utilized opinion-based judgment? These questions opened new perspectives for future research in using AI for reflective purposes.

FUTURE WORK

In general, our research confirmed a potential for leveraging annotation bias in AI for prompting self-reflection in a design process. This approach shed light on the subjective judgment we brought into our design projects. Annotation bias is a well-known side-

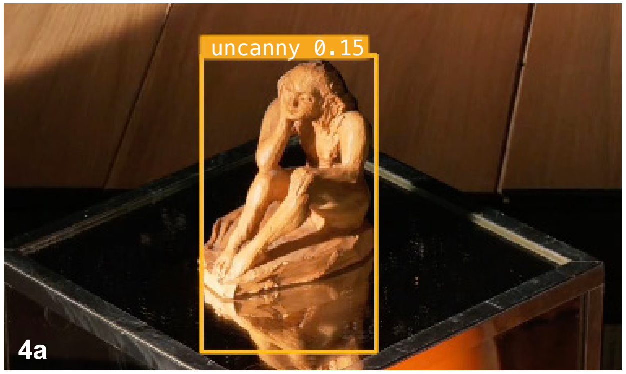

4a

effect of data annotation, and refers to the biased representation of data in the annotation process [4], due to human annotators basing their annotation decisions, consciously or unconsciously on implicit stereotypes. As data collection and annotation comes from human processes, it will inevitably create a biased, subjective training dataset [9]. This can affect AI model performance, leading to unfair or incorrect predictions (e.g. [3, 15]). Looking at our research, however, we seemed to have leveraged annotation bias to preserve subjective judgment as much as possible. Our focus was on individual designers’ personal reflections, so we devised a process that encouraged finding the most subjective labels.

The absence of a ‘ground truth’ for our data meant that there was no relevance for annotation agreement [19] with other annotators to justify if our labels were ‘correctly annotated’. Instead, our detection models were set up to reflect what we intuitively felt and perceived during the annotation task. Acknowledging and recognizing this subjectivity of emotional concepts and exploring that in subjectively trained models could be an exciting direction for future research.

Additionally, this pictorial presents a novel approach for designers to engage with AI, acknowledging the rapid development of AI technologies and the need for designers and researchers to respond to them. We expected an intentional ‘slowing down’ of the interaction with an AI model, fostering self-reflection among designers by highlighting the annotation process of an object detection model. While limited self-reflection was experienced through the interaction with the AI model, we recognized the significance of the training process - i.e. the image annotation. The training process prompted us to pause and contemplate the future collaborations between AI and designers.

Consequently, this work has opened up a promising research direction for the future: integrating AI training into designers’ processes and exploring qualitative methodologies for investigating AI in design and creative endeavors. As such, the project also highlights the current limitation of generative AI models, whose ostensibly simple interfaces belie the complexity of their inner workings, hindering users from accessing and understanding the datasets they are trained on. In response, the addition of data collection and annotation to a designer’s process holds promise. This approach allows designers to gain insights into working with AI systems, while offering a method for reflecting on their own fascinations and interests.

ACKNOWLEDGEMENTS

We would like to thank Maurik Stomps for his participation in the project. We also would like to thank Jesse Nijdam and Suzanne van Norden for their support during the development of this pictorial. Lastly, we would like to thank Lisa Hardon and the Dutch Design Foundation for their support in realising the installation at Dutch Design Week 2022.

REFERENCES

[1] Berger, J. (2008). Ways of seeing. Penguin uK.

[2] BKOR. (z.d.). BKOR (Beeldende Kunst & Openbare Ruimte) - Rotterdam. Geraadpleegd op 10 februari

2023, van https://www.bkor.nl/

[3] Buolamwini, J., & Gebru, T. (2018). Gender Shades: Intersectional Accuracy Disparities in Commercial Gender Classification. Proceedings of the 1st Conference on Fairness, Accountability and Transparency, 77–91. https://proceedings.mlr.press/v81/buolamwini18a.html

[4] Baumer, E. P., Siedel, D., McDonnell, L., Zhong, J., Sittikul, P., & McGee, M. (2020). Topicalizer: reframing core concepts in machine learning visualization by co-designing for interpretivist scholarship. Human–Computer Interaction, 35(5-6), 452-480. https://doi.org/10.1080/07370024.2020.1734460

[5] van der Burg, V., Akdag Salah, A. A., & Chandrasegaran, S. (2022). Ceci n'est pas une Chaise: Emerging Practices in Designer-AI Collaboration. In Proceedings of Design Research Society International Conference (DRS2022), Bilbao. Design Research Society. https://doi.org/10.21606/drs.2022.653

[6] Chang, J. C., Amershi, S., & Kamar, E. (2017, May). Revolt: Collaborative crowdsourcing for labeling machine learning datasets. In Proceedings of the 2017 CHI Conference on Human Factors in Computing Systems (pp. 2334-2346). https://doi.org/10.1145/3025453.3026044

[7] Chen, Y., & Joo, J. (2021). Understanding and Mitigating Annotation Bias in Facial Expression Recognition. 2021 IEEE/CVF International Conference on Computer Vision (ICCV), 14960–14971. https://doi.org/10.1109/ICCV48922.2021.01471

[8] Cunningham, P., Cord, M., & Delany, S. (2008). Supervised Learning (pp. 21–49). https://doi.org/10.1007/978-3-540-75171-7_2

[9] Dignum, V. (2021). The Myth of Complete AI-Fairness (arXiv:2104.12544). arXiv. http://arxiv.org/abs/2104.12544

[10] Dwyer, B., Nelson, J. (2022), Solawetz, J., et. al. Roboflow (Version 1.0) [Software]. Available from https://roboflow.com

[11] Eyewitness, D. (2022). Gardens of the World. Dorling Kindersley Limited.

[12] Hatch, M. J., & Yanow, D. (2008). Methodology by Metaphor: Ways of Seeing in Painting and Research. Organization Studies, 29(1), 23–44. https://doi.org/10.1177/0170840607086635

[13] Hemment, D., Aylett, R., Belle, V., Murray-Rust, D., Luger, E., Hillston, J., Rovatsos, M., & Broz, F. (2019). Experiential AI. AI Matters, 5(1), 25–31. https://doi.org/10.1145/3320254.3320264

[14] Jocher, G., Chaurasia, A., Stoken, A., Borovec, J., NanoCode012, Kwon, Y., Michael, K., TaoXie, Fang, J., imyhxy, Montes, D., Wang, Z., Fati, C., Nadar, J., Laughing, Jain, M. (2022). ultralytics/yolov5: V7.0 - YOLOv5 SOTA Realtime Instance Segmentation. Zenodo. https://doi.org/10.5281/zenodo.7347926

[15] Leavy, S., O'Sullivan, B., & Siapera, E. (2020). Data, Power and Bias in Artificial Intelligence (arXiv:2008.07341). arXiv. http://arxiv.org/abs/2008.07341

[16] Lopatovska, I. (2016). Three types of affect tags for art images. Proceedings of the Association for Information Science and Technology, 53(1), 1–8. https://doi.org/10.1002/pra2.2016.14505301039

[17] Lucero, A., Desjardins, A., Neustaedter, C., Höök, K., Hassenzahl, M., & Cecchinato, M. E. (2019, June). A sample of one: First-person research methods in HCI. In Companion Publication of the 2019 on Designing Interactive Systems Conference 2019 Companion (pp. 385-388). https://doi.org/10.1145/3301019.3319996

[18] Midjourney Documentation and User Guide. (z.d.). https://docs.midjourney.com/

[19] Mujtaba, D. F., & Mahapatra, N. R. (2019). Ethical Considerations in AI-Based Recruitment. 2019 IEEE International Symposium on Technology and Society (ISTAS), 1–7. https://doi.org/10.1109/ISTAS48451.2019.8937920

[20] Murray-Rust, D., Nicenboim, I., and Lockton, D. (2022) Metaphors for designers working with AI, in Lockton, D., Lenzi, S., Hekkert, P., Oak, A., Sádaba, J., Lloyd, P. (eds.), DRS2022: Bilbao, 25 June - 3 July, Bilbao, Spain. https://doi.org/10.21606/drs.2022.667

[21] Nelson, L. K. (2020). Computational grounded theory: A methodological framework. Sociological Methods & Research, 49(1), 3-42. https://doi-org.tudelft.idm.oclc.org/10.1177/0049124117729703

[22] Nowak, S., & Rüger, S. (2010). How reliable are annotations via crowdsourcing: A study about inter-annotator agreement for multi-label image annotation. Proceedings of the International Conference on Multimedia Information Retrieval, 557–566. https://doi.org/10.1145/1743384.1743478

[23] OpenAI (2022, 14 april). DALL·E 2. OpenAI. https://openai.com/dall-e-2/

[24] OpenAI (2023, 2 februari). ChatGPT: Optimizing Language Models for Dialogue. OpenAI. https://openai.com/blog/chatgpt/

[25] Pääkkönen, J., & Ylikoski, P. (2021). Humanistic interpretation and machine learning. Synthese, 199(1-2), 1461-1497. https://doi.org/10.1007/s11229-020-02806-w

[26] Russell, J. A. (1980). A circumplex model of affect. Journal of personality and social psychology, 39(6), 1161.

[27] Schmidt, S., & Stock, W. G. (2009). Collective

indexing of emotions in images. A study in emotional information retrieval. Journal of the American Society for Information Science and Technology, 60(5), 863–876. https://doi.org/10.1002/asi.21043

[28] Schön, D. A. (1988). Designing: Rules, types and worlds. Design Studies, 9(3). https://doi.org/10.1016/0142-694X(88)90047-6

[29] Schön, D. A., & Wiggins, G. (1992). Kinds of seeing and their functions in designing. Design Studies, 13(2), 135–156. https://doi.org/10.1016/0142-694X(92)90268-F

[30] Seeber, I., Bittner, E., Briggs, R. O., de Vreede, T., de Vreede, G.-J., Elkins, A., Maier, R., Merz, A. B., Oeste-Reiß, S., Randrup, N., Schwabe, G., & Söllner, M. (2020). Machines as teammates: A research agenda on AI in team collaboration. Information & Management, 57(2), 103174. https://doi.org/10.1016/j.im.2019.103174

[31] Simeone, L., Mantelli, R., & Adamo, A. (2022, June). Pushing divergence and promoting convergence in a speculative design process: Considerations on the role of AI as a co-creation partner. In Proceedings of the DRS2022 Conference. https://doi.org/10.21606/drs.2022.197

[32] Tanweer, A., Gade, E., Krafft, P. M., & Dreier, S. (2021). Why the data revolution needs qualitative thinking. Harvard Data Science Review, 3.

[33] Verhagen, L. (2023, February 22). Mauritshuis hangt Kunstwerk gemaakt door algoritme op plek Vermeer: ‘Gewoon Mooi’ of ‘onethisch’? de Volkskrant. Retrieved May 5, 2023, from https://www.volkskrant.nl/nieuws-achtergrond/mauritshuis-hangt-kunstwerk-gemaakt-door-algoritme-op-plek-vermeer-gewoon-mooi-of-onethisch~ba60c70b/

[34] Wong, W. L. P., & Radcliffe, D. F. (2000). The Tacit Nature of Design Knowledge. Technology Analysis & Strategic Management, 12(4), 493–512. https://doi.org/10.1080/713698497

[35] Williams, Z. (2023, April 18). ‘Ai isn’t a threat’ – Boris Eldagsen, whose fake photo duped the Sony judges, hits back. The Guardian. Retrieved May 5, 2023, from https://www.theguardian.com/artanddesign/2023/apr/18/ai-threat-boris-eldagsen-fake-photo-duped-sony-judges-hits-back

[36] 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