Proceedings of the First Conference on
Designing with Artificial Intelligence
dai digital 2020

Proceedings of the First Conference on
Designing with Artificial Intelligence
dai digital 2020

Proceedings of the First Conference on
Designing with Artificial Intelligence
Online conference, September 17-19, 2020
https://www.designing-artificial-intelligence.eu

© 2021 dai digital and the contributors

Edited by
Marc Engenhart
Sebastian Loewe

Published by
appliedAI, an initiative for the application of AI technology
(www.appliedai.de)
Digital Product School
(www.digitalproductschool.io)

Both are part of
UnternehmerTUM
Lichtenbergstraße 6
Munich-Garching
D-85748, Germany

Identifier
ISBN 978-3-00-068634-4

Organising committee of the dai digital
Marc Engenhart
Sebastian Loewe
Julian Pohl (assistant)

Editorial design
Engenhart ° Design Studio
engenhart.design

Copyright notice
All contents of these proceedings by dai digital and appliedAI, the initiative for the application of AI technology, are published to grant open access to all readers. These proceedings may be copied and distributed freely within the academic institutions when properly cited. Outside of academia and with commercial intent the usage initially is prohibited and needs to be approved by the editors or the authors.

INITIATIVE FOR APPLIED ARTIFICIAL INTELLIGENCE

DIGITAL PRODUCT SCHOOL

Introduction

Intro 6 **Machine Learning as a Wicked Design Material: Questions, Topics, and Challenges for ML-Driven User-Centered Design. An Introduction to the dai digital Proceedings**
Marc Engenhart, Sebastian Loewe

Intro 16 **State of the Art and Design for AI**
Jennifer Heier

Keynotes

Research 28 **How Design Changes Machine Learning; How Machine Learning Changes Design**
Rebecca Fiebrink

Practice 44 **Playful Machine Learning: An Interaction Designer's Journey into Neural Networks**
Andreas Refsgaard

Research 56 **Learning to Educate Machines**
Patrick Hebron

Discourse 64 **Drop That Intelligence and Get on with It!**
Gerhard Anger

Research 70 **Unremarkable AI: Towards AI That Co-Lives and Co-Evolves with Users**
Qian Yang

Papers

Research 80 **Will Artificial Intelligence Make Designers Obsolete?**
Klaus Neuburg, Sven Quadflieg, Simon Nestler

Discourse 80 **"I Don't Experience the Meaning of Creativity in the Same Way Humans Do": Künstliche Intelligenz und Kreativität**
Anika Meier, Manuel Rossner

Research 92 **A Model to Identify the Data Needed for AI-driven Graphic Design**
Jan-Henning Raff

Practice 114 **ROBODADA: Speed-Dating with (Un)Friendly Robots**
Andreas Muxel, Elias Naphausen

Research 128 **Making Visual Design Adapt to Emotions and Affect: Thoughts and Research Questions on Emotionally Responsive Visual Design**
Sebastian Loewe

Practice 144 **The Design Process of the dai Brand and Logo**
Marc Engenhart

152 **Proceedings' contributors**

157 **Documentation and resources**

159 **Conference partners**

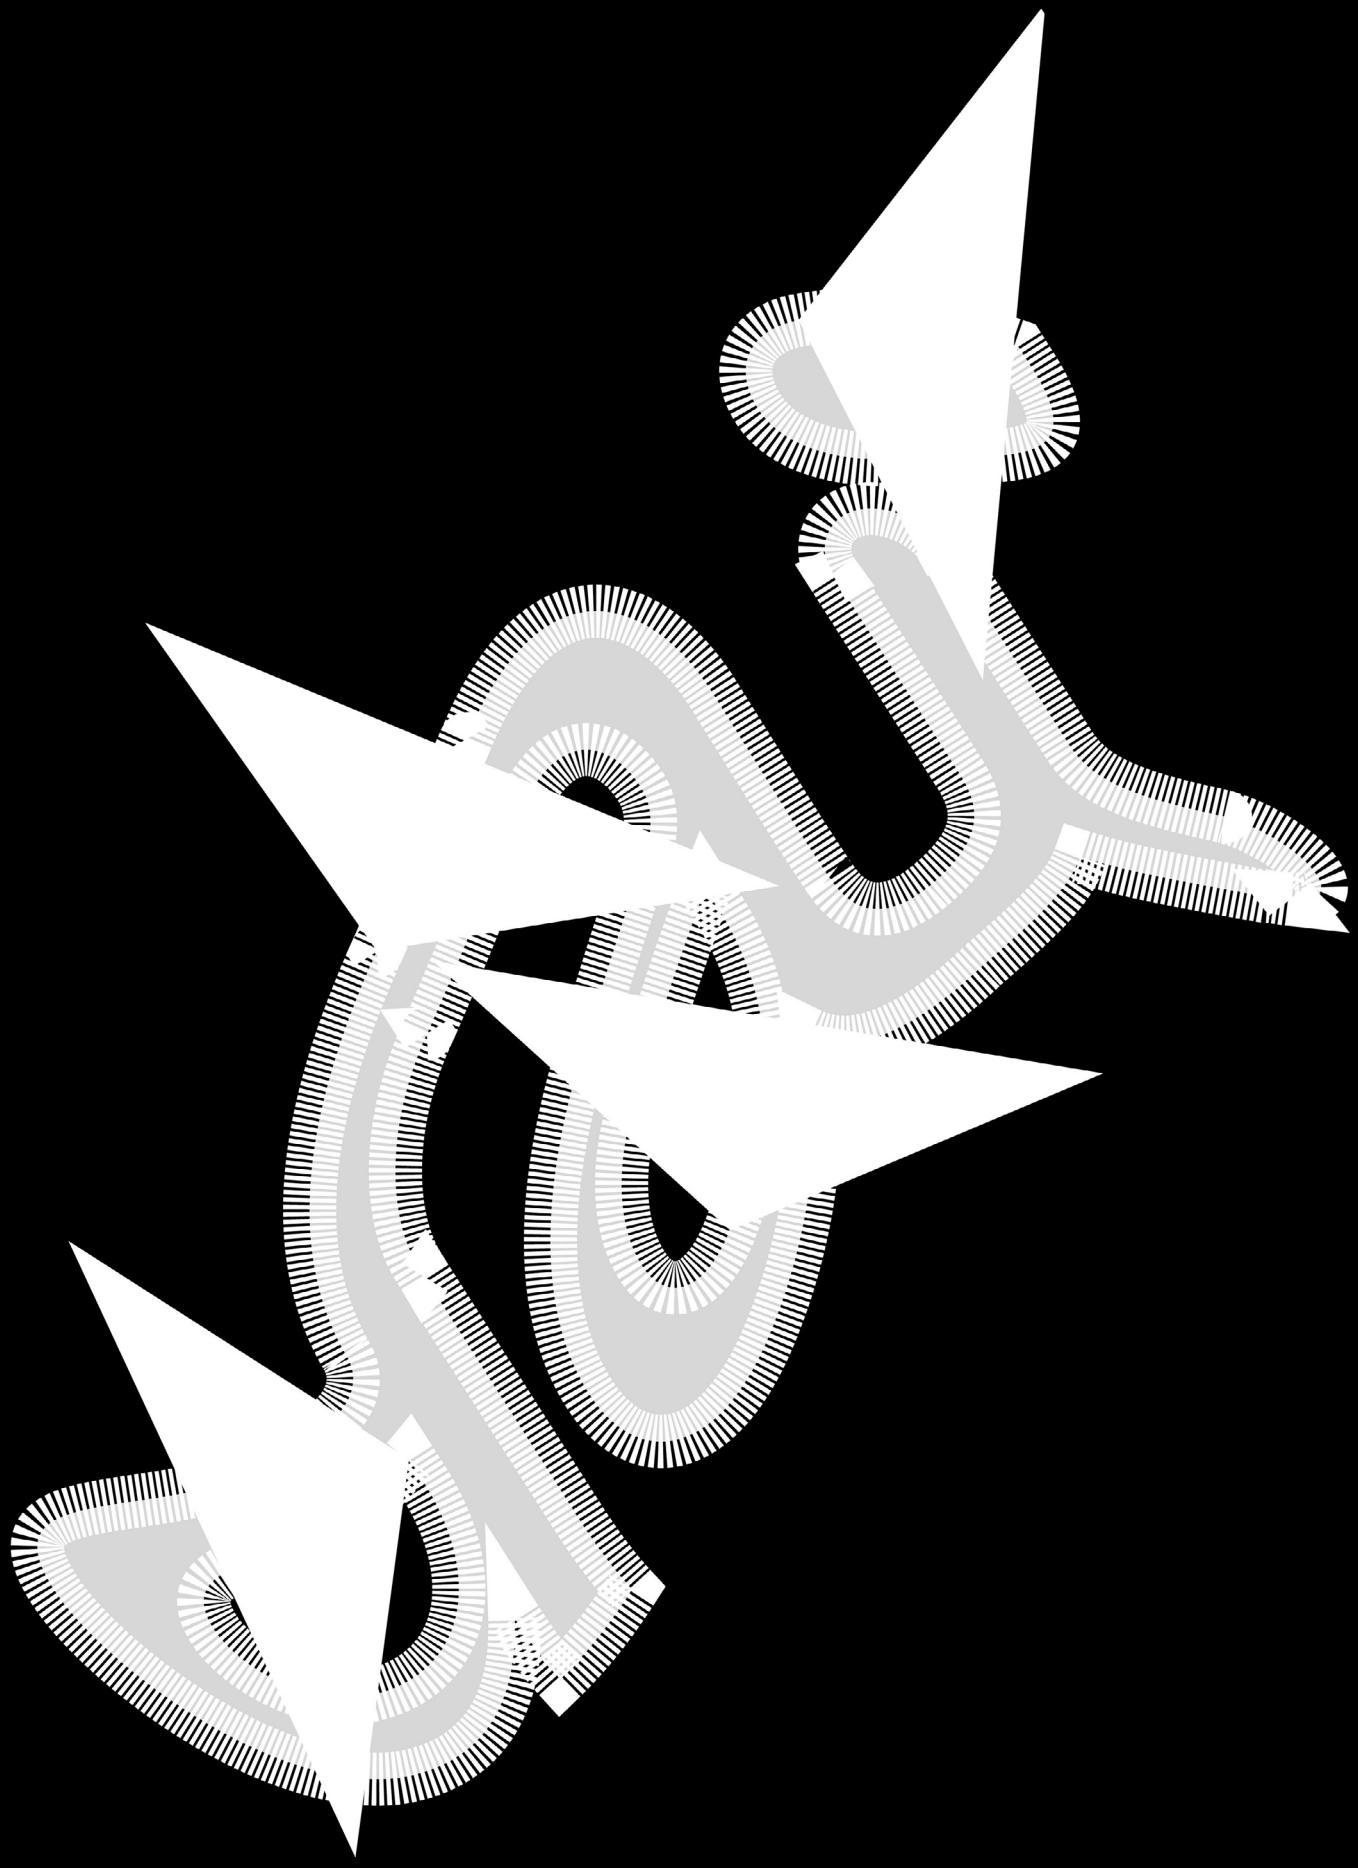

Machine Learning as a Wicked Design Material: Questions, Topics, and Challenges for ML-Driven User-Centered Design. An Introduction to the dai digital Proceedings

Marc Engenhart & Sebastian Loewe

Intro

Just recently, the OpenAI initiative[1] published the third version of a machine learning (ML) model they dubbed "Generative Pre-trained Transformer" or GPT for short, which is not only a deep and powerful recurrent neural network, but is making waves in the machine learning community, the design community, and the HCI community. The GPT model allows for a broad range of applications in the field of natural language processing (NLP), meaning that the neural network can transform the intricate rules of a language's grammar and the vast amount of content created in this language into a mathematical function. This trained network then infers from this function a complete re-creation of text corpuses, which can become incredibly close to individual author's idiosyncrasies without task specific training. Usually these NLP models are domain agnostic, meaning that their application is limited to the field of spoken and written language. However, OpenAI researchers bridged this gap and the proved GPT model was able to learn the characteristics of two-dimensional images, such as object appearance and category[2]. The accuracy with which the model then classified images from the famous ImageNet data set was surprisingly high. It even allowed for the completion of existing images as well as the creation of entirely new images. This opened up a new field of implementations and allowed for a variety of elaborate use cases that cross-utilize natural language processing and image creation[3]. It did not take long for the design and interaction community to come up with interesting use cases. Jordan Singer for example, product designer at Square, came up with a prototyping solution that runs on GPT-3. Singer types in "draw a red circle" and the Figma plug-in he developed, draws a red circle. This may seem premature, but when Singer uses GPT-3 to create an app, the power of the model gets more tangible. Singer created a Figma tool he called "Designer," which includes a text interface in which he types:

> An app that has a navigation bar with a camera icon, "Photos" title, and a message icon. A feed of photos with each photo having a user icon, a photo, a heart icon, and a chat bubble icon.

The Figma plug-in renders a few seconds and out comes an app that closely resembles Instagram [4]. This even works for prototypes of websites.[5] Singer just feeds the model a sentence like "a website like stripe.com that is about a chat app" and the software miraculously designs a website that does look like stripe.com. GPT-3 even works for complex creative coding languages such as P5.js (Processing).[6]

A more professional and comprehensive effort to let machine learning do the heavy design lifting is Salesforce's "Einstein Designer"[7], which launched in June 2020. 'Einstein Designer' is taking vast inspiration for color schemes and font faces from approximately a thousand e-commerce websites. It then not only creates web shop layouts based on what it learned, but also connects specific layouts to certain customer preferences and thus personalizes the design outputs according to the digital personas. This lets 'Einstein Designer' create fully automated design templates in just a second.

1 "OpenAI is an AI research and deployment company. Our mission is to ensure that artificial general intelligence benefits all of humanity", About, Open AI, accessed November 26, 2020, https://openai.com/blog/openai-api.

2 "Image GPT", Image-GPT, Open AI, accessed November 26, 2020, https://openai.com/blog/image-gpt.

3 Future, Bakz T., "14 Cool Apps Built on OpenAI's GPT-3 API", YouTube, July 20, 2020, https://www.youtube.com/watch?v=G6Z_S6hs29s&feature=youtu.be.

4 Singer, Jordan, "This changes everything. With GPT-3, I built a Figma plugin to design for you.", Twitter, July 18, 2020, https://twitter.com/jsngr/status/1284511080715362304.

5 Singer, Jordan, "Words → website. A GPT-3 × Figma plugin that takes a URL and a description to mock up a website for you.", Twitter, July 25, 2020, https://twitter.com/jsngr/status/1287026808429383680.

6 Loewe, Sebastian, "We Generated P5-Code Based on Text Inputs with OpenAI's GPT-3 and This Could Show the Future of Designer-Machine Co-Creation.", Medium, March 7, 2021, https://sebastian-loewe.medium.com/we-generated-p5-code-based-on-text-inputs-with-openais-gpt-3-and-this-could-show-the-future-of-db2ae41efc8d.

7 Rohde, Sönke, "Einstein Designer. AI-Powered, Personalized Design at Scale", Medium, June 23, 2020, https://medium.com/salesforce-ux/einstein-designer-ai-powered-personalized-design-at-scale-367d71b85a3d.

These two use cases are selected randomly and could be effortlessly replicated by a hundred similar cases in which machine learning enters the realm of design and alters its practices. As organizers of the conference "Designing with Artificial Intelligence" (dai digital), we want to draw a variety of research questions from these two examples, and then assign them to the topics of the proceedings' contributors:

1) Are the intelligent machines going to make designers obsolete?
2) How do designers create interesting results when working with intelligent machines?
3) How do designers teach good design methods to a machine?
4) What do effective intelligent machine teaching processes look like?
5) How can designers get familiar with machine learning to explore what possibilities exist when creating with an intelligent machine?
6) What are the roles designers can play in designing with/ for intelligent machines?
7) Which roles do designers want to play in this scenario, and how much agency in the design process do they need?
8) What roles do the machines play in co-designing processes?
9) Is there a framework or system that supports and facilitates the design process with intelligent machines?
10) How can designers design convincing intelligence experiences?
11) Is there a framework or system to outline and evaluate ML-driven design experiences?

To address these questions, to question rather simplistic notions of artificial intelligence and machine learning, and to get a 'lay of the land' of what is relevant, achievable, and desirable in the field of designing for and with ML, we have invited a variety of speakers to the *dai digital* conference. Their backgrounds range from design, computer and data science, to cognitive science, design research, design management and art. Overall, our goal is to facilitate a low-threshold access point to the world of ML-driven design from a design perspective, while also forming a critical forum of design researchers, practitioners and students to collectively address these questions, as well as raise new ones.

The first question on the list sparked an ongoing, vivid public debate about whether or not the design bot will substitute human creativity, thus rendering human artists and designers obsolete.[8] It also serves as a foundation and starting point for a majority of the papers presented in these proceedings. Striking in the debate's dichotomous simplicity is the fact that its voices tend to ignore the many ways humans and machines already co-create. Questions surrounding the augmentation of human creative abilities through intelligent machines and the subsequent question of agency in this co-creative endeavor neither seem prominent nor desirable in the public opinion.

Instead, the debate's protagonists consolidate the image of autonomous machines as truly intelligent and self-motivated

8 Page Magazine editors, "Future of Design: Macht die KI bald alles?", *Page Magazine*, February 2, 2020, https://page-online.de/branche-karriere/future-of-design-macht-die-ki-bald-alles.

entities which deliberately intend to usurp the field of creativity. The notion of artificial intelligence as a tool—whose universality makes it truly unique—that serves all kinds of man-made economic, political, military, social, cultural, and creative purposes is unfortunately taking a backseat in this discussion. Furthermore, the dualistic nature of the debate is founded on the premise of an ever-creative artist or designer, whom to replace means attacking humanness itself. To drive home their point, these voices also ignore the fact that designers work on crunching deadlines and are asked to produce great ideas and designs at a pace that no creative mood can possibly survive on a daily basis. Only when you are willing to accept creativity as an economic imperative which renders creativity an inexhaustible resource does the exaggerated vision of machines taking over the world of art and design start to become plausible. These fear mongering warning calls tend to try to cement the social imperative of creativity as well, proclaiming machines as equals to humans only to subsequently denounce this notion, declaring every person a representative of an inimitable, ever-creative species. The trope of an autonomous machine creator as a threat, introduced by popular culture and reinforced by public figures like OpenAI's founder Elon Musk, quickly became a now ubiquitous mode of perception for matters of machine learning. Therefore, it is not surprising that the British newspaper The Guardian fell prey to the promises of machine intelligence and published an op-ed entirely written by GPT-3. The algorithm is portrayed as a conscious, cognitively self-aware entity that 'says:' "I would happily sacrifice my existence for the sake of humankind.[9]" This piece of fantastic realism reveals more about the Guardian's editors' intent to portray such an algorithm as truly intelligent, than it does about the actual GPT-3 neural net.

Luckily, **Gerhard Anger** in his keynote opts out of the hype about intelligent machines and questions the notion of 'intelligence' altogether that is reserved for machine learning algorithms. Through this approach, he successfully introduces a more factual perspective on the role of data-based algorithms in design. This might seem more than plausible, but in times of overblown fears and expectations rather unpopular. Anger also points out that although machine learning seems like impartial computing of data, it perpetuates or even worsens partial biases. Because machines mostly learn based on human perception and classification, their output can lead to reactionary and racist classification systems. Machine learning—despite contrary belief—has never been outside of the political and social domains of society. This fact is yet again recalled by activists like the "Algorithmic Justice League"[10], a group of scientists and artists who, like Anger, stress that an allegedly simple thing like facial recognition is heavily biased and works properly only for a fraction of people. It is known for quite some time now that data is intensifying biases and stereotypes, and punishing any form of deviance from a fictitious ideal of appearance and behavior. Consequently, Anger warns that no conscious design decision leads to this discriminatory behavior of machines. Declaring these machines intelligent, makes their decisions even more plausible and renders the criticism that this technology would replace the human element of design or creativity somewhat blasphemous.

9 The Guardian editors, "A robot wrote this entire article. Are you scared yet, human?", September 8, 2020, https://www.theguardian.com/commentisfree/2020/sep/08/robot-wrote-this-article-gpt-3.

10 Algorithmic Justice League, "We're leading a cultural movement towards Equitable and Accountable AI", About, Algorithmic Justice League, accessed November 26, 2020, www.ajl.org/about.

Reflecting on the philosophical as well as the algorithmic state of artificial intelligence and deviating from the discourse of artificial humanness, **Andreas Muxel** and **Elias Naphausen** deliberately distort the still insufficiently studied notion of a human-like intelligence in their ROBODADA presentation. Quirky objects—also identifiable as 'robots'—serve as a platform to reflect philosophically and ethically on machines that react to one's own emotions, but are missing any form of obvious intelligence. They behave erratically and do not interact with humans in an intelligible way. Borrowing from Marc Hassenzahl's research on machine otherness[11] Muxel and Naphausen launched design probes within the dai's video conference setup to find expected perceptions of robots and encourage critical questions about the nature of these 'intelligent things'.

Anika Meier and **Manuel Rossner** go deep into the discourse of 'artificial creativity' to examine the art bot Ai-Da, a robot whose sole purpose is creating something beautiful and meaningful. Based on Margaret Boden's definition of creativity[12], the bot is creative, but by the standards of an artist with a distinctive intent to create and express oneself, the lines get blurrier. In their German paper, Meier and Rossner, underline the nature of this project as a thought experiment. Nonetheless, only if you are willing to forget that a machine needs to be created by a team of engineers and artists to create what this team intended for the bot to do 'autonomously,' the comparison with an artist will work. It brings together two distinct characteristics, which have commonalities only based on the comparative connection. On the one side, it is the notion of an artist as an intrinsically motivated creative individual. On the other, it is a robot, a computational machine, that does not have a free will or intrinsic creative motivation. The free will and need to express oneself creatively, which often drives an artist, as well as the artist's ability to evaluate and refine their artistic outcomes based on a personal concept or taste, are missing from Ai-Da entirely. Patrick Hebron's question why artists (or any human mind for that matter) should be precisely imitated, given the fact that there are so many people on earth, fighting for a living doesn't seem far fetched[13]. Interestingly, Meier and Rossner then show that the perception of intelligence in the art market is not so much driven by the ability of algorithms to create something that a human cannot, but by the power of certain protagonists within the art world to declare artificial intelligence a worthy author of art, and thus dignify the outcomes of rather simple algorithms. This sheds light on a conservative and enclosed system that is able to co-opt and disambiguate the rather ambiguous state of co-creation that lies in 'artificial creativity.'

Likewise, **Klaus Neuburg**, **Sven Quadflieg**, and **Simon Nestler** point out that the discourse of replacing designers with artificial intelligence is exaggerated as well. First, the authors stress that even the most sophisticated outcome of an intelligent algorithm is based on data that is given to it, rendering the outcome a reproduction of sorts. Second, the processes of design and the problems they allow to tackle are much more complex than is usually admitted when machines create images and layouts. ML-driven design in this notion is limited to assist with solving

11 Hassenzahl, Marc, Matthias Laschke, Robin Neuhaus, and Judith Dörrenbächer, "Otherware needs Otherness: Understanding and Designing Artificial Counterparts", NordCHI Workshop, Experience and Interaction, accessed November 26, 2020, http://www.experienceandinteraction.com/nordichi-workshop-otherware-needs-otherness.

12 Boden, Margaret, *The Creative Mind: Myths and Mechanisms* (New York: Basic Books 1992).

13 Wilk, Elvia, "Against Prediction. Designing Uncertain Tools. An interview with Patrick Hebron, the machine intelligence design manager at Adobe", *Noema Magazine*, September 24, 2020, https://www.noemamag.com/against-prediction-designing-uncertain-tools.

very narrow design problems. The complex world with its wicked problems, as Horst Rittel and later Richard Buchanan [14] famously put it, will remain the domain of designers, according to Neuburg, Quadflieg and Nestler.

14 Buchanan, Richard, "Wicked Problems in Design Thinking", *Design Issues 8*, no. 2 (Spring 1992): 5–21.

Rebecca Fiebrink's extensive research on interactive machine learning adds valuable insights to the many stages of machine-led assistance and co-creation, rendering the debate about replacement to a certain extent obsolete. Rather than seeing the human creators become dispensable, she points ways how to use ML to a designer's advantage. Her underlying assumption is that ML for creative purposes does not necessarily need to uncover a hidden truth in big data. Instead, ML could be used to create prototypes and experimental design probes which enable designers to reach their goals for generating or curating data. Fiebrink points out that for experimental design practices the rules for machine learning must be expanded or even partially suspended, making way for embodied or tacit knowledge rather than a ground truth which then allows for new and surprising ways to co-create.

Fiebrink's concept translates into a design practice that **Andreas Refsgaard** fully embraced and mastered. In his keynote paper, Refsgaard provides insights into his quick and fun ways to implement existing machine learning tools like ML5.js or RunwayML to form a truly experiment-driven art and design practice. His body of work is guided by his principles of human-machine interaction, i.e. making people active participants in the ML experience and utilizing ML algorithms in unforeseen interfaces facilitating unexpected interactions. Refsgaard's paper then serves as an answer to the question of where to start with machine learning. He even speculatively answers the question surrounding what happens when machines create everything autonomously in "Books by AI", a project that lets machines create book content and book covers from Amazon books, which are then sold on Amazon again, comically rendering authors and designers bystanders in an infinite creative regress.

Seeing Refsgaard's artwork or Salesforce's 'Einstein Designer' may seem like a step towards full design automation. But, if automated design should be more than just a GAN-based book cover or web shop solution, things get tricky. Automating more than just a simple image or layout function within the complex process of design, seems out of reach. At the keynote roundtable, keynote speaker Patrick Hebron was asked whether a machine could design the seminal "I (heart) NY" logo by Milton Glaser. He then replied, that it is not so much a question whether or not a machine is capable of doing it eventually, but whether it is capable of creating this highly specific logo intentionally based on a design brief. In Hebron's eyes, we still need designers with a vast cultural, semiotic and aesthetic knowledge to create these pieces of semantically and culturally interwoven visual designs. This is where **Jan-Henning Raff** comes into play. He works on the idea of a reliable framework for visual communication that goes beyond a simple understanding of design. Raff attempts to explain not only the delicate domains of aesthetics, semiotics, and cultural meaning, but to form a theoretical understanding how data might support the design process. His contributed

paper to the dai digital conference proceedings proposes an approach that deduces keywords for data labeling from an analysis of visual communication, allowing practical ML-based design experiments with this data. The paper may form a starting point for a scientific design-based discussion about data-based visual communication.

When designing for machine learning, the perception of intelligent machines opens up a whole new field of user experience questions. Answering the question regarding well designed intelligence experiences, keynote speaker **Qian Yang** takes a rather radical approach proposing that ML should become invisible to the user. To underline her idea, she stresses the analogies between machine learning and electricity, leaving behind any kind of anthropomorphism. This demands from the designers to create and design machine learning systems that seamlessly integrate into daily routines. Much like electricity, ML should become invisible, ordinary or unremarkable, as Yang puts it. Good design in this notion is making ML running undetected in the background.

To answer maybe the toughest question, how to teach good design to a machine, **Patrick Hebron** in his keynote suggests a new paradigm for working with machines. He underlines that the anthropomorphic representation of algorithms deemed intelligent is highly debatable. By understanding human intelligence not as a monolithic ability, but a set of different modalities, Hebron opens up new ways to understand and design machine intelligence. The act of teaching and learning in his notion is a matter of creating the right environment and thus design plays a pivotal role. Hebron hypothesizes provocatively that the future of ML has more to do with design than with engineering. This hinges on a notion of design as a way to create tools that can be used beyond the initial intent of a creator figure. By making humans true teachers for machines rather than just providers of massive amounts of data, Hebron's vision for ML-driven design exceeds the boundaries of common ways to teach a machine. By creating looped interactions between what a machine learns and the designer, an exciting pathway is opened up for implementing new ways to create machines that learn.

Additionally, we invited **Jennifer Heier**, from the Siemens AI Lab, Munich, who unfortunately could not attend the dai digital conference, to write a brief introduction to machine learning and design from a maker's perspective, answering the pressing question, which design principles should govern the human-machine co-creation process.

Lastly, two projects during the dai digital were transferred into articles for the proceedings. First, co-organizer **Marc Engenhart** explains the way the dai digital logo was co-designed with a machine learning algorithm. This sheds some much needed light on the questions of agency and machine co-creation, while serving as an inspiration for other design practitioners. Second, co-organizer **Sebastian Loewe** elaborates on a design prototype that works with emotion detection. The article focuses on the question of how to critically design responsive visual prototypes for Emotion AI, for which there is currently no playbook in the realm of human-centered design.

Machine Learning as a Wicked Design Material: Questions, Topics, and Challenges for ML-Driven User-Centered Design. An Introduction to the dai digital Proceedings

Marc Engenhart & Sebastian Loewe

When we started thinking about a conference that places questions of machine learning and design at center stage, we could not possibly imagine the amount of great feedback and overwhelmingly positive engagement we received. We started by asking our keynote speakers to answer the most pressing questions concerning new paradigms and practices when designing for and with machine learning algorithms and intelligent assist systems. Our focus was always to include all kinds of design disciplines, ranging from visual communication and interaction design to product and fashion design. Equally important, we never wanted to exclude a certain discipline. Even though the dai digital conference was created to focus on the designer's perspective, we set out to include data and computer scientists, design managers, design consultants, and artists. The proceedings mirror this approach very clearly. Due to the fact that we had such a diverse range of speakers from all kinds of backgrounds, we decided to not force the logic of academic publishing on them. This is why we created three different categories for the proceedings: 'research,' 'practice,' and 'discourse'. The category 'research' consists of academically-based research papers, 'practice' introduces essays that reflect on practical design questions, and 'discourse' is a category aimed at presenting articles that deal with public opinions and debates. We hope that with these three categories we could capture the nature of the contributed essays and papers, while helping readers to gain a better overview on the topics of the dai digital proceedings.

We want to deeply and warm-heartedly thank all speakers and workshop hosts for making this conference a great experience for everyone participating. We also thank everyone involved in publishing these proceedings of the dai digital.

Berlin and Stuttgart, Germany
February 2021

BIBLIOGRAPHY

Algorithmic Justice League, “We’re leading a cultural movement towards Equitable and Accountable AI”, About, Algorithmic Justice League, accessed November 26, 2020, www.ajl.org/about.

Boden, Margaret, The Creative Mind: Myths and Mechanisms (Basic Books: New York, 1992).

Buchanan, Richard, “Wicked Problems in Design Thinking”, Design Issues 8, no. 2 (Spring 1992): 5–21.

Future, Bakz T., “14 Cool Apps Built on OpenAI‘s GPT-3 API”, YouTube, July 20, 2020, https://www.youtube.com/watch?v=G6Z_S6hs29s&feature=youtu.be.

Hassenzahl, Marc, Matthias Laschke, Robin Neuhaus, and Judith Dörrenbächer, “Otherware needs Otherness: Understanding and Designing Artificial Counterparts”, NordCHI Workshop, Experience and Interaction, accessed November 26, 2020, http://www.experienceandinteraction.com/nordichi-workshop-otherware-needs-otherness.

Page Magazine editors, “Future of Design: Macht die KI bald alles?”, Page Magazine, February 2, 2020, https://page-online.de/branche-karriere/future-of-design-macht-die-ki-bald-alles.

Rohde, Sönke, “Einstein Designer. AI-Powered, Personalized Design at Scale”, Medium, June 23, 2020, https://medium.com/salesforce-ux/einstein-designer-ai-powered-personalized-design-at-scale-367d71b85a3d.

Singer, Jordan, “This changes everything. With GPT-3, I built a Figma plugin to design for you.”, Twitter, July 18, 2020, https://twitter.com/jsngr/status/1284511080715362304.

Singer, Jordan, “Words → website. A GPT-3 × Figma plugin that takes a URL and a description to mock up a website for you.”, Twitter, July 25, 2020, https://twitter.com/jsngr/status/1287026808429383680.

The Guardian editors, “A robot wrote this entire article. Are you scared yet, human?”, September 8, 2020, https://www.theguardian.com/commentisfree/2020/sep/08/robot-wrote-this-article-gpt-3.

Wilk, Elvia, “Against Prediction. Designing Uncertain Tools. An interview with Patrick Hebron, the machine intelligence design manager at Adobe”, Noema Magazine, September 24, 2020, https://www.noemamag.com/against-prediction-designing-uncertain-tools.

dai digital conference
Impressions

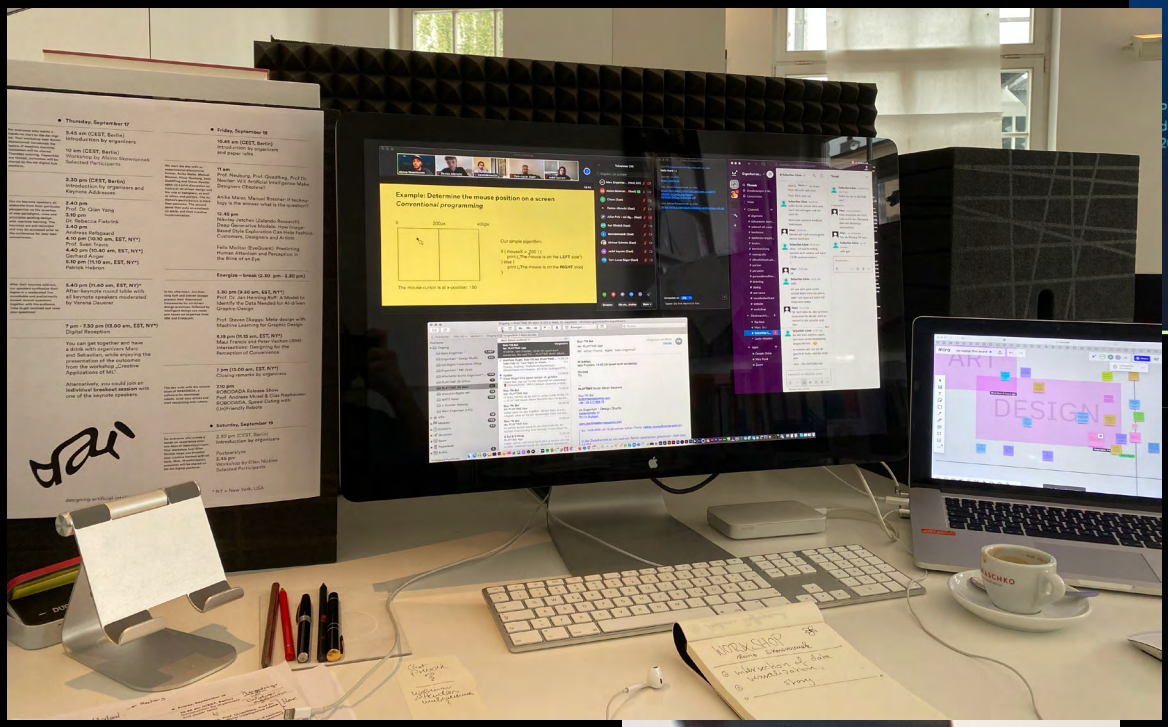

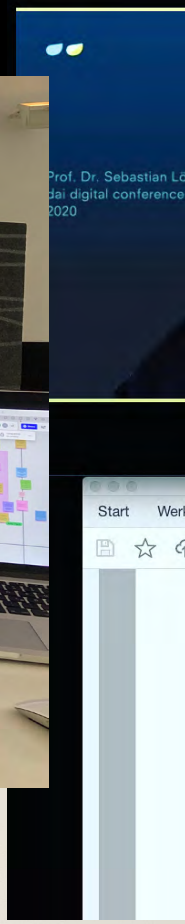

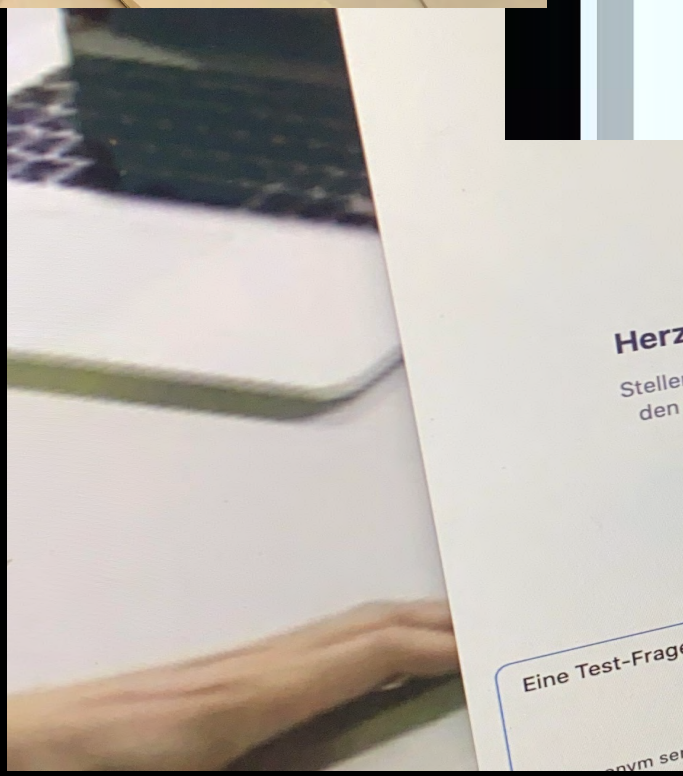

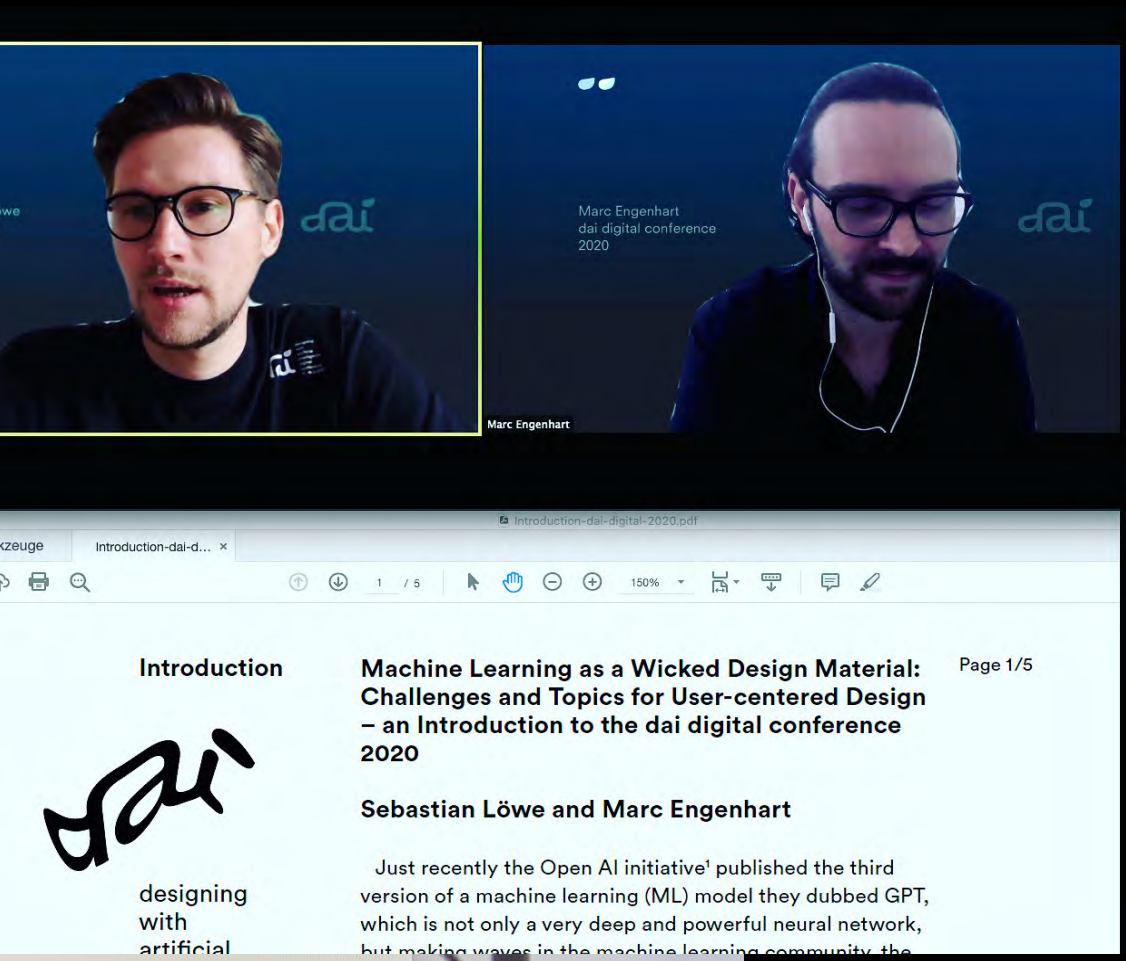

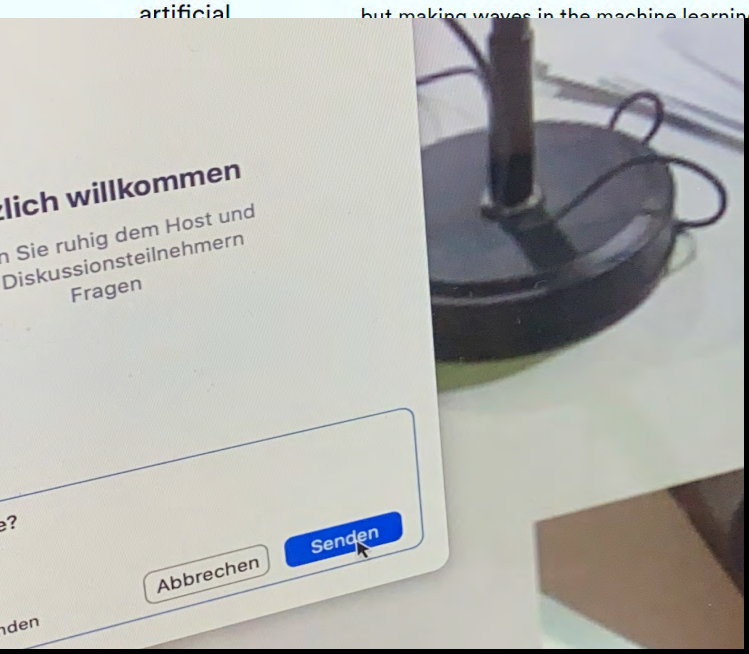

dai digital conference headquarter / introduction
September 17, 2020

State of the Art and Design for AI

Jennifer Heyer

ABSTRACT

This state of the art reflects on the challenges and fears designers and artists are facing when it comes to the tasks that will be substituted by machine learning and artificial intelligence. Representing the misleading perceptions, communication and ways of drawing unrealistic expectations of the technology, by media and scientific scholars. Resulting in a reality check of a selection of current projects and algorithms that support the design process. Followed by a section about the great opportunities and new roles for designers and artists generated by the technology. Such as co-creation with algorithms and data scientists, and the shift from creator to curator, amongst others. Furthermore, an overview of the development of human-centered AI principles is provided. Showing differences and similarities in those approaches, and making clear what kind of answers are still missing. Those insights are then mapped to the role of UX in an industrial setting, making a connection to a real-world scenario.

Technology always influenced the work, methods and tools of the creative community, and so does machine learning[1] (ML) a subdomain of artificial intelligence (AI). Designers and artists are similarly affected by the technology, and so are many other fields. Creativity as a unique selling point seems to be an insecure bet these days, since certain smart algorithms imply creative behavior as well. For example, algorithms that write poems[2], script screenplays[3] and draw paintings[4]. In this short introduction, I want to shed light on some important issues and aspects around the hype of AI and ML. I will try to clarify and answer if designers are going to be obsolete in the age of ML? Do we therefore need to be afraid of the technology? Is AI only a buzzword? Which challenges is the profession really facing? What does ML offer for designers and artists? And where do we, as designers, need to create more impact and gain influence?

1. CHALLENGES, FEARS AND ML (DESIGN) FEATURES

1.1 What is in the news?

> Who designs like a machine will be replaced by a machine.
> (Philipp Schäfer, Head of UX - ExB group)

Some designers and artists are afraid that ML will supposedly take their jobs and to be honest, some tasks are already performed better by machines than by humans. ML and other means of AI entered a lot of design and art domains, such as film and music, graphic design, web design and product design, as well as industrial design. Due to smart algorithms that create art and design, there is a huge debate around whether AI will make designers obsolete or not.[5]

On one hand, ML algorithms seem to perform better in certain areas than human designers, additionally they enable non-designers to create art, music, or paintings. On the other hand, the job of a designer should not be limited to the digital tools he or she uses. The way of thinking, interacting, and experiencing a

1 Cristina, Cris, “The Designer’s Guide to Machine Learning”, Digitalist, 2017, https://digitalist.global/talks/the-designers-guide-to-machine-learning.

2 Devlin, Es, et al., “POEMPORTRAITS”, Google Arts & Culture Experiments, 2019, https://experiments.withgoogle.com/poemportraits.

3 Sharp, Oscar, and Goodwin, Ross, “Sunspring”, Frankfurter Kunstverein, 2017, https://www.fkv.de/en/oscar-sharp-with-ross-goodwin-benjamin.

4 Benney, Marnie, and Kistler, Pete, “Creative Tools to Generate AI Art”, AI Artists, 2019, https://aiartists.org/ai-generated-art-tools.

5 Chen, Yifei, “Will Artificial Intelligence Remove Designers from the Design Process?“, Towards Data Science, 2017, https://towardsdatascience.com/will-artificial-intelligence-remove-designers-from-the-design-process-5e6661430055.

craft like drawing and model making in a hardware shop is also part of the designer's creative process and practice. Additionally, misconceptions, wrong understandings and expectations towards AI technology make an informed discourse very hard to achieve. This is also due to the fact that recent articles draw a dystopian future of the creative profession, full of fear and negative expectations.[6] In the following paragraphs I want to use a concrete example as a basis for showing what ML is really capable of.

1.2 Reality check

Looking at the profession of graphic design and its different sub domains, such as typography, layout and logo design, one can say that the field is challenged, or better to say supported, by algorithms on many different frontiers. Color picking[7], as well as font pairing[8] is made more simple with ML algorithms. Even AI-based logo design[9] and website generators[10] are available. Some experts of the field even envision a meta design[11]. This means that the whole design process is augmented by ML-generated output. This is especially feasible for tasks like layout and composition, which can be done, or supported by an algorithm. For example, the smart assistant is fed with the requirements of the design briefing. With those data points, it generates a first set of layout drafts. The human designers' role is to choose the best suited designs and to initialize the next step in the design process, such as color matching or picture placement. This back and forth between human designers and ML algorithms is beneficial for both sides. The designer gets inspired by the machines' output and can use his experience to get the best performance. The algorithm on the other hand learns the style of the designer and can adapt its results accordingly.

More advanced algorithms such as the introduction of so-called GANs (generative adversarial neural networks designed by Ian Goodfellow and his colleagues in 2014)[12] offered even more exciting features for designers and artists. Two different neural networks, a generator and a discriminator are trained simultaneously on a similar data set in a competitive manner. This means, the two networks train each other and therefore improve their performance according to their learning process. If you train a GAN on photographs it generates new photographs that have similar features and characteristics as the training data. An example is the work by Nvidia engineers who trained GANs with portraits from celebrities. The algorithm generated completely new and realistic looking fake portraits from celebrities that actually do not exist.[13]

A very interesting project from Ajkel Mino and Gerasimos Spanakis uses this approach for logo generation. In their paper 'LoGAN: Generating Logos with a Generative Adversarial Neural Network Conditioned on color'[14], they use this class of machine learning to generate logos. "LoGAN's results offer a first glance at how artificial intelligence can be used to assist designers in their creative process…". The data input they use is a digital library including logos. These logos are labeled according to twelve different colors. Specifying the colors and their proportions by

6 Rädeker, Jochen, "Future of Design: Macht die KI bald alles?", Page, 2020, https://page-online.de/branche-karriere/future-of-design-macht-die-ki-bald-alles.

7 "Colormind - the AI powered color palette generator", Home, Colormind, accessed November 11, 2020, http://colormind.io.

8 "Fontjoy - Generate font pairings in one click", Home, Fontjoy, accessed November 11, 2020, https://fontjoy.com. "Font Map - An AI Experiment by IDEO", Home, Font Map, accessed November 11, 2020, http://fontmap.ideo.com.

9 "Looka Logo Maker - Make 100s of Custom Logos Instantly, Logo Maker, Looka, accessed November 11, 2020, https://looka.com/logo-maker.

10 "The Grid", Home, The Grid, accessed November 11, 2020, https://thegrid.io.

11 Skaggs, Steven, "Stage 1.X- Meta-Design with Machine Learning is Coming, and That's a Good Thing", Dialectic 2(2), 2019.

12 Goodfellow, Ian, et al., "Generative Adversarial Networks", Advances in Neural Information Processing Systems 27, 2014.

13 Karras, Tero, et al., "Progressive Growing of GANs for Improved Quality, Stability, and Variation", International Conference on Learning Representations, 2018.

14 Mino, Ajkel and Spanakis, Gerasimos, "LoGAN: Generating Logos with a Generative Adversarial Neural Network Conditioned on color", International Conference on Machine Learning and Applications, 2018.

the designer is the input needed for the algorithm to generate a variety of logos.

> The proposed model can successfully create logos if given a certain keyword, which in our case consisted of the most prominent color in the logo ... This helps designers in their creative process and while brainstorming, making logo design cheaper and less time-consuming.

The results of their work presented in the research paper are very experimental looking and impressive given the fact that the only input variable is color. Future work suggests that they add more labels to the artificial system, "such as the shape of the logo or the focus of the company ...".

Most ML-based systems need input data (in the form of aesthetic guidelines, design principles and logo examples) to generate a nearly infinite amount of variations. One can say that the introduction of some of those systems enable the human designer to invert their role from a creator into a curator, who now has the responsibility to choose the best from a selection of ML-generated outputs. Some systems serve more as a source of inspiration and others even need the designer to pick the input data and adjust the output parameters. This shows that the level of human interaction with machines varies from system to system, depending on the application, domain and degree of automation.

Coming back to the initial question, if AI will make designers obsolete one can say that certain tasks are perfectly replaced by machines and will make designers obsolete. But these are mainly repetitive tasks which are boring and even desired to be solved by machines. The debate is not as black and white as most media is trying to illustrate the situation of the creative community. It is true that it seems to be crucial not to focus primarily on the look and feel of products, objects and services. There are other means of impact for designers to develop AI technology than those attributes. The mentioned cases, as well as other data-driven approaches, need designers to find and frame the problems ML should solve. It is also the role to provide the input needed, to train all those systems with design knowledge and aesthetic experience.

2. OPPORTUNITIES, CHANGE AND NEW ROLES

2.1 Focus on Co-Creation

Instead of perceiving the technology as a threat, designers and artists can focus on the positive aspects of augmentation for both the design process and the creative community. AI in this perspective changes the way designers and artists work. Formerly working as pure creators, this approach is shifting to co-creation with machines and data scientists. Designers and artists can positively influence the design of ML algorithms. Instead of becoming obsolete, the profession can increase their demand in the growing field of smart systems. Need finding, problem framing and finding, prototyping and testing are

necessary steps in the development. Selecting system output and providing feedback to the system are important aspects, too. New roles and job labels have already been generated, e.g. Human-Centered Machine Learning (HCML, as defined by Google), or Deep Learning UX (DLUX, as claimed by Salesforce). Those terms imply that input, adjustment of the models and final selection of the output are done by human designers and artists. Shaping those new job profiles and inventing even more tasks to be done by artists and designers is an exciting new area.

The creative community already initiated some guidelines and principles for the process of working with AI and ML. They are all based on the idea of human-centeredness. This is the opposite of focusing on technological feasibility and business viability which are currently the main drivers of AI development. This concept emphasizes the involvement of the human interacting with smart systems. Usability, user experience, and social topics such as ethics play an important role in the human-centered design movement.

2.2 Human-Centered-AI (Principles)

2.2.1 Purpose

Current advancements in AI and ML development is largely based on statistical analysis and mathematical techniques that inherit standardized metrics and ignore the chaotic human environment. This means they work perfectly in a laboratory setting, where a lot of parameters can be controlled and limited, but they fail in real world scenarios. For example, Tay the chatbot from Microsoft, learned racist and sexist slur within 16 hours by twitter users.[15] Regardless, most AI systems are generated in the mentioned way and implemented in the real world and therefore represent a lot of challenges for the humans involved. Another important aspect in that context is data and bias in data. For instance, algorithms judge the probability whether a convicted person will again commit a crime. These algorithms score people of color higher than any white suspects.[16] Even the decision from an autonomous car which traffic participant to hurt or even kill when an accident occurs, are issues that involve a lot of ethical and moral questions, which are not finally answered. As a result, the creative community and other practitioners should call for a broader perspective on humans being impacted by AI systems.

> Human-centered AI is about defining the goals of AI to meet human needs and to work within human environments. ... Not only do we need a set of new tools and techniques to make AI work in practice, but we need to shift the process by which AI is even designed in the first place.[17]

2.2.2 Differences, Similarities and Missing Pieces. AI Frameworks From MS, Google and Lingua Franca

Companies such as Microsoft, Google and IBM amongst others came up with human-centered AI principles. Those companies are already using AI solutions in consumer facing domains. They

15 Neff, Gina and Nagy, Peter, “Talking to Bots: Symbiotic Agency and the Case of Tay”, International Journal of Communication 10, 2016.

16 Angwin, Julia, et al., “Machine Bias”, ProPublica, 2016, https://www.propublica.org/article/machine-bias-risk-assessments-in-criminal-sentencing.

17 Agarwal, Abhay and Regalado, Marcy, “Lingua Franca - A Design Language for Human Centered AI”, 2020, https://linguafranca.polytopal.ai.

have a vast amount of data accessible through their portfolio of applications, as well as the workforce and know how to develop their own algorithms. In this section, I will shortly introduce some of the different approaches, show similarities, as well as missing parts in human-centered AI.

Microsoft divides its set of principles into four steps. Each of those steps has certain aspects designers should take into consideration. In total their approach represents 18 different principles.

> We hope you can use these Guidelines for Human-AI Interaction throughout your design process as you evaluate existing ideas, brainstorm new ones, and collaborate with the multiple disciplines involved in creating AI.[18]

Initially it is crucial to "make clear what the system can do". This is in a way meant to guide the users' expectations towards ML systems. Setting those expectations too high will end in an unsatisfactory experience while using smart solutions. It is therefore important to manage those expectations right from the beginning. The second set of guidelines deals with aspects of interaction with the system. For example, which wording the system uses to communicate with the user is a relevant issue. Misleading language might evoke social injustice or address stereotypes. A third segment is devoted to the failure of the system. This shows the importance of this aspect when designing for ML. The algorithm is not perfect. It is trained on data generated by humans and can make mistakes. Therefore, solutions should communicate how they arrived at their results (so-called Explainable AI, or XAI)[19]. It is necessary to provide the possibility to hand control over the system to the human user, as well as being honest about the fact that the machine might be wrong, or at least unsure about its output. Finally, every ML application should be able to learn from the interaction with the users and improve over time. That is the core strength of ML technology. Therefore, collecting feedback from users is crucial and important in the long run. Furthermore, it is helpful to inform the users about new releases and features of the system, in order to keep their trust in the system's reliability and performance.

Unfortunately, the work from Microsoft is primarily focused on the user experience of the final solution. They do not provide any information on initial steps in developing the algorithms, like data preparation, problem definition or choice of technology.

Google came up with their own "People+AI Guidebook". It consists of six chapters which follow a regular product development flow. Similar to Microsoft they provide an explanation of each chapter, as well as a related worksheet that should support the use of their principles.

> The People + AI Guidebook was written to help user experience (UX) professionals and product managers follow a human-centered approach to AI.[20]

18 Amershi, Saleema, et al., "Guidelines for Human AI Interaction", Microsoft, 2019, https://www.microsoft.com/en-us/research/project/guidelines-for-human-ai-interaction.
Ajkel et al., „LoGAN", 5.

19 Ribeiro, Marco, et al., "Why Should I trust you? Explaining the Predictions of any Classifier", Conference on Knowledge Discovery and Data Mining, 2016.

20 Google PAIR, "People + AI Guidebook.", 2019, https://pair.withgoogle.com/guidebook.

The aspects mentioned by the Google guidebook are very similar to Microsoft's. Order and arrangement of the principles are done slightly differently. Especially the sections about explainability, feedback and failure are mentioned in both cases and represent a common ground of general AI challenges. Those aspects are widely adopted in other principles of human-centered-AI. However, Google's guidebook also puts emphasis on the preparation and problem definition of AI systems as a relevant step of the overall design process. This part is called 'data collection+evaluation'. Designers and artists can provide valuable input with qualitative research methods adding meaning to purely statistical data points. Another question the guidebook pinpoints at, is whether the development team should use a given data set or establish their own. This might be needed when a given data set is biased towards a certain user group, for example a bias on a certain gender or age group.

Lingua Franca—A Design Language for Human-Centered AI—is currently the most comprehensive selection of principles and guidelines for designing AI systems. It is published by Abhay Agarwal (former Microsoft and currently lecturer at Stanford d.school) and Marcy Regalado (a Stanford d.school graduate). In addition to their handbook, which introduces seven different aspects of design intervention, they also provide eight principles that every AI system should include. On top of this, they started to collect example elements that represent a concrete solution to the stated principles.

> This is a design language for AI—a standard set of techniques, frameworks, visuals, messaging, and overall design patterns that apply broadly to different kinds of AI to make it more usable, more trustworthy, and better aligned with people. [21]

Their approach is very similar to the general design process. It starts with the initial problem selection and definition, and is followed by observations into human behavior. However, the section about data and sensemaking also puts emphasis on statistical methods and knowledge. These skills are not taken for granted amongst designers and artists. Another part of their handbook deals with the choice of technologies. This is not AI specific, but in order to be able to choose the right technology in an era of ML and AI a certain degree of literacy about the possibilities and features of the different concepts is necessary. Sometimes even the decision not to use ML at all is an important one.

Most human-centered AI principles have a set of aspects that are very similar, i.e. problem definition, needfinding, explainability, trust, feedback and failure are commonly important. This is a great starting point for further research and development of a standard set of principles that can guide artists and designers along the development process of AI systems. However, there is still a need for new tools and methods that work along these guidelines. Taking into account the specific contexts where the guidelines are going to be used is omitted in most principles and seems to be the missing link to make them work and add value to real-life projects.

21 Polytopal, "Lingua Franca. A Design Language for Human-Centered AI", 2019, https://linguafranca.polytopal.ai.

2.3 Real-World Examples

Deriving from my own experience as a designer working in an industrial context[22] dealing with algorithms and data scientists I want to examine some of the important aspects and roles when designing with ML.

Starting with the right problem and defining it very well is the most important aspect. This is where the team makes a lot of decisions which are hard to change later in the development of the AI system. It is crucial for designers to be part of this initial step. Too often companies try to implement AI solutions where it is not needed or even inadequate. *Problem selection and definition* can be supported by designers with qualitative user research methods, as well as the decision whether to use AI as a solution at all. The biggest challenge here is researching technology that is not yet in use as there is often no obvious existing behavior to look at or existing preferences to discuss and explore. One helpful approach is not to talk about AI in your research, but instead talk about *assistance*. Research participants could struggle to distinguish between prominent media-driven perceptions that are fueled by fear and negative effects of AI as well as the reality of their own behavior. Helpful tools are *cultural probes* and anything that helps to understand the current workflow of the research participants. For instance, if you want to understand the factory ecosystem and improve the demand planning process in a factory, it is crucial to talk to the planners, understand how they currently plan, which tools they use and which other stakeholders are involved.

Data scientists are used to work with data, mostly *quantitative data*, whereas designers are used to work with *qualitative data*. In the era of AI, it is important to be *data literate* in both worlds. Understanding statistical data sets and being able to gain insights from those sets is new to most designers, but a value-add property to idea generation. Enriching those with qualitative insights is the best strategy. It also helps to detect whether additional data is needed. A promising approach we use in our setting is to develop a *product vision* among designers and data scientists for the AI system without even thinking about data sources and libraries in the first place. Once we start a new project we establish that common ground, check the feasibility of the solution and as a consequence the data sources will follow. Within the industrial settings there is a huge amount of data, either from sensors or internal and external reports. In order to design human-centered AI the human perspective needs to be added, too. This is not taken for granted especially in the business-to-business environment. It is therefore the job of designers to add this point of view making sure the AI system's implementation is going to meet users' needs.

Developing ML and AI systems is a lengthy and costly process. Therefore, besides working iteratively, *prototyping* is a very crucial and helpful step. Unfortunately, there are hardly any *tools* out there yet that quickly prototypes the training and evaluating of an AI algorithm without developing and training it. *Wizard of Oz*[23] as a method became very popular for prototyping voice assistants and chatbots, but does not help with AI systems that

22 Zistl, Sandra, "Autonome Systeme: Labor der neuen Ideen", 2017, https://new.siemens.com/global/de/unternehmen/stories/forschung-technologien/kuenstliche-intelligenz/autonome-systeme-siemens-ai-lab.html.

23 MacGovern, Nicole, "Wizard of Oz Prototyping Process Blog", Prototypr, 2018, https://blog.prototypr.io/wizard-of-oz-prototyping-process-blog-a20ffce8886.

are supposed to predict and forecast user behavior. Starting with a small sample of a given data set, train and evaluate an algorithm on this data to judge whether it is feasible or not, is the most promising approach here. Still it requires a certain degree of computer science skills and quite some time. Finding smart solutions for this problem will enable designers and other professions to speed up the development of AI systems that will benefit the users.

Our product teams at the Siemens AI Lab consist of data scientists, UX designers, ML engineers and data analysts to make sure that we are able to answer all the above requirements. The domain and business knowledge is added by our clients. This diverse set up ensures an ongoing discourse about risks, missing aspects and potential solutions which is very helpful when navigating through the complexity of AI systems.

2.4 What needs to be done

What does this mean for designers and artists in the age of AI?

(a) Get comfortable with the basics of the technology
This does not mean you have to be able to write code, although this is not a bad thing, it's not a must either. To start with the right expectations, and be able to communicate and manage those among other disciplines it is crucial to understand what AI can do. Our role as designers is to sharpen or define the starting point for an AI project. This means we must decide and evaluate if AI is the right solution for the given challenge. Plus, we also need the knowledge to choose the right design approach. There is a lot of information and learning material out there, even for free.[24] [25]

(b) Try things out, fail, repeat, succeed
We need new tools and methods in the age of AI, which are not fully there yet.[26] Prototyping and starting small are key aspects for the technology and a good starting point. Working iteratively is key to developing and training AI algorithms. This fits in perfectly with the design process, so why not using it. It will also help to collaborate with data scientists and machine learning engineers. In addition, your basic shared language and knowledge is very helpful.

(c) Focus on value
It is nice to play around with technology as a starting point. To make a real contribution, try to focus on a goal that adds value, either to your business or to society, e.g. with a public art installation that educates people about the aspects of AI or that raises awareness for specific issues. This is a great lever to generate real impact of your work. Remember, combining technological feasibility, business viability and human desirability is where the magic happens.

(d) Be open
Communicate and share your ideas and results—sharing is caring. A lot of AI work is open source[27], and that is a good thing. It helps to validate the output and increases accessibility of the

24 "Lobe - Machine Learning Made Easy", Home, Lobe, accessed November 11, 2020, https://lobe.ai.

25 "A free online course - Elements of AI", Home, Elements of AI, accessed November 11, 2020 https://course.elementsofai.com.

26 Dorard, Louis, "Machine Learning Canvas", 2019, https://www.louisdorard.com/machine-learning-canvas.

27 Kogan, Gene, "Machine Learning for Artists", github, accessed November 11, 2020, https://ml4a.github.io.

technology. The more people become data and AI literate the better. At the moment, it is also interesting to see that the issues are equally challenging to everyone. Exchange of knowledge and best practice is very common and welcome among the people dealing with AI.

3. CONCLUSION

There is a tension between the drive for automating human (design) tasks and the reality of a new working mode of co-creation between humans and machines. It shows that the answers to the most pressing questions are not as dichotomous as some discussions try to outline them.

ML for art and design is generally a good thing. It offers new opportunities and jobs for the creative community. As with every new technology, it implies changes and a learning process. We need new tools and methods in the age of AI due to the complexity of the systems and the dynamics of the technology involved. Since ML algorithms learn and change over time while collecting new data, not everything can be foreseen and planned.

In addition, we should try to increase the involvement of design in the development and training of the technology from the start; in the business discourse as well as in science. Taking care of the output is important, and a good start, but some decisions are already made and some mistakes are maybe too late to handle. Designers should work closely together with data scientists and engineers in order to get the best out of technology.

Furthermore, the technology is currently driven by technological feasibility and business viability. The human factor is very much missing and this causes some problems. Bias in data, lack of usability, fear of job loss, to mention just a few pain points for people facing the technology. The pure focus on capitalistic advantages often implies disadvantages for workers that need to deal with the effects of the new wave of automation. Not to forget the negative impact for the environment, because training and running ML algorithms consumes a lot of energy. Our role is to guide the development of the technology in a beneficial way for all stakeholders. Make sure that AI is solving the ‘right’ problem, augmenting and supporting humans. Initiating human-centered AI principles is a crucial step.

Lastly, design education facilities need to embrace that new technology and add it to their curriculum. ML is hardly on their agenda (exceptions given), but an important skill for the future designer and artist. This does not mean that we should forget about our heritage of craftsmanship, but new perspectives are necessary to be included rather sooner than later. As mentioned in the ongoing debate, AI driven tools enable non-designers with creative skills, therefore we should develop our profession further. The new skills needed are a benefit for the creative community and as already happened in the past with other technological progress shaped new professions and activities.

Right now, AI poses more questions than answers. The creative community can add ideas and value to this ongoing endeavor. Experts involved in this field even ask for our involvement as designers, because they see the problems that stem from the lack of technological empathy.

BIBLIOGRAPHY

"A free online course - Elements of AI", Home, Elements of AI, accessed November 11, 2020 https://course.elementsofai.com.

Agarwal, Abhay and Regalado, Marcy, "Lingua Franca - A Design Language for Human Centered AI", 2020, https://linguafranca.polytopal.ai.

Amershi, Saleema, et al., "Guidelines for Human AI Interaction", Microsoft, 2019, https://www.microsoft.com/en-us/research/project/guidelines-for-human-ai-interaction.

Angwin, Julia, et al., "Machine Bias", ProPublica, 2016, https://www.propublica.org/article/machine-bias-risk-assessments-in-criminal-sentencing.

Benney, Marnie, and Kistler, Pete, "Creative Tools to Generate AI Art", AI Artists, 2019, https://aiartists.org/ai-generated-art-tools.

Chen, Yifei, "Will Artificial Intelligence Remove Designers from the Design Process?", Towards data science, 2017, https://towardsdatascience.com/will-artificial-intelligence-remove-designers-from-the-design-process-5e6661430055.

"Colormind - the AI powered color palette generator", Home, Colormind, accessed November 11, 2020, http://colormind.io.

Cristina, Cris, "The Designer's Guide to Machine Learning", Digitalist, 2017, https://digitalist.global/talks/the-designers-guide-to-machine-learning.

Devlin, Es, et al., "POEMPORTRAITS", Google Arts & Culture Experiments, 2019, https://experiments.withgoogle.com/poemportraits.

Dorard, Louis, "Machine Learning Canvas", 2019, https://www.louisdorard.com/machine-learning-canvas.

"Fontjoy - Generate font pairings in one click", Home, Fontjoy, accessed November 11, 2020, https://fontjoy.com.

"Font Map - An AI Experiment by IDEO", Home, Font Map, accessed November 11, 2020, http://fontmap.ideo.com.

Goodfellow, Ian, et al., "Generative Adversarial Networks", Advances in Neural Information Processing Systems 27, 2014.

Google PAIR, "People + AI Guidebook.", 2019, https://pair.withgoogle.com/guidebook.

Karras, Tero, et al., “Progressive Growing of GANs for Improved Quality, Stability, and Variation”, International Conference on Learning Representations, 2018.

Kogan, Gene, “Machine Learning for Artists”, Github, accessed November 11, 2020, https://ml4a.github.io.

“Lobe - Machine Learning Made Easy”, Home, Lobe, accessed November 11, 2020, https://lobe.ai.

“Looka Logo Maker - Make 100s of Custom Logos Instantly, Logo Maker, Looka, accessed November 11, 2020, https://looka.com/logo-maker.

MacGovern, Nicole, “Wizard of Oz Prototyping Process Blog”, Prototypr, 2018, https://blog.prototypr.io/wizard-of-oz-prototyping-process-blog-a20ffce8886.

Mino, Ajkel and Spanakis, Gerasimos, “LoGAN – Generating Logos with a Generative Adversarial Neural Network Conditioned on color”, International Conference on Machine Learning and Applications, 2018.

Neff, Gina and Nagy, Peter, “Talking to Bots: Symbiotic Agency and the Case of Tay”, International Journal of Communication 10, 2016.

Polytopal, “Lingua Franca. A Design Language for Human-Centered AI”, 2019, https://linguafranca.polytopal.ai.

Rädeker, Jochen, “Future of Design: Macht die KI bald alles?”, Page, 2020, https://page-online.de/branche-karriere/future-of-design-macht-die-ki-bald-alles.

Ribeiro, Marco, et al., “Why Should I trust you? Explaining the Predictions of any Classifier”, Conference on Knowledge Discovery and Data Mining, 2016.

Sharp, Oscar, and Goodwin, Ross, “Sunspring”, Frankfurter Kunstverein, 2017, https://www.fkv.de/en/oscar-sharp-with-ross-goodwin-benjamin.

Skaggs, Steven, “Stage 1.X- Meta-Design with Machine Learning is Coming, and That’s a Good Thing”, Dialectic 2(2), 2019.

“The Grid”, Home, The Grid, accessed November 11, 2020, https://thegrid.io.

Zistl, Sandra, “Autonome Systeme: Labor der neuen Ideen”, 2017, https://new.siemens.com/global/de/unternehmen/stories/forschung-technologien/kuenstliche-intelligenz/autonome-systeme-siemens-ai-lab.html.

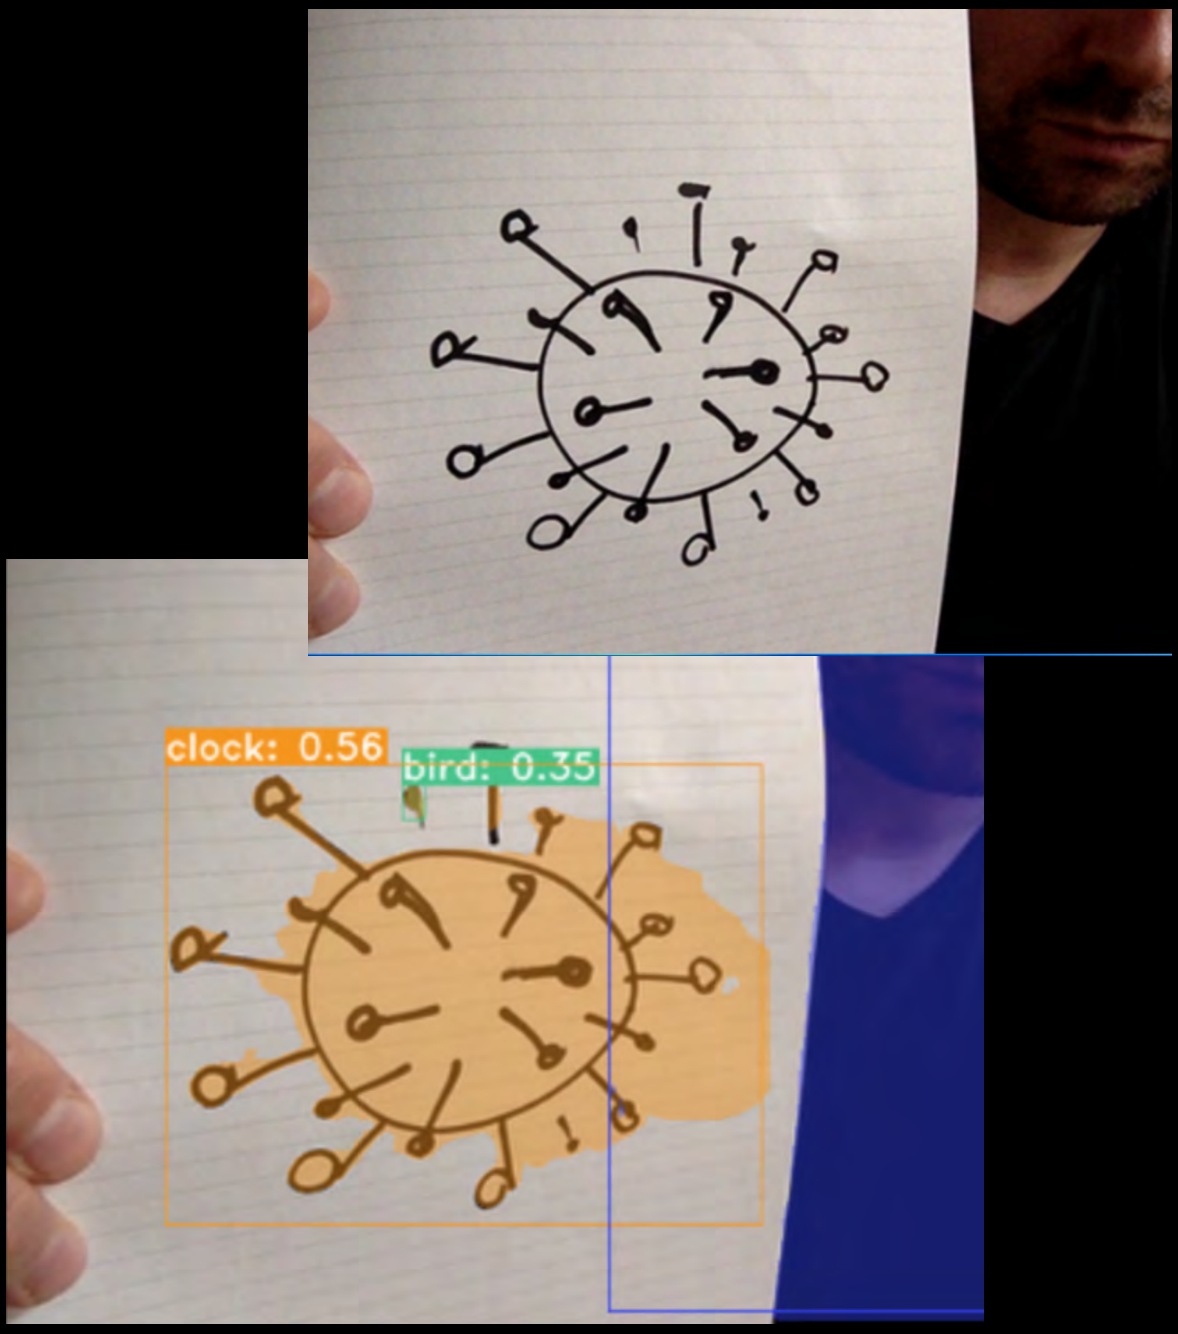

Computer vision / image recognition model
Hand-drawn corona virus and output interpretation

How Design Changes Machine Learning; How Machine Learning Changes Design

Rebecca Fiebrink

Research

This is the companion publication to Fiebrink's opening keynote address for the 2020 dai digital. The keynote can be accessed at https://www.designing-artificial-intelligence.eu.

ABSTRACT

In this paper, I draw on my research over the past twelve years developing machine learning tools for creative practice to describe how machine learning should be used and understood differently in such contexts (as compared to more conventional applications), and to discuss some of the ways that using machine learning can change design practices and outcomes.

1. INTRODUCTION

I have been building tools for creative uses of machine learning (ML) since 2008, when I made the first version of the Wekinator software[1]. Wekinator is a tool for interactive, real-time, creative use of supervised learning, in which a user teaches the computer to produce some kind of output values in response to certain input values (Figure 1). These inputs could be data from sensors, video, or audio; the outputs might offer control or influence over any sort of computer process—for instance, changing sound synthesis, procedural animation or design parameters, controlling robot behaviour, controlling a game engine, and so on. Instead of requiring a human programmer to write a function for computing output values from inputs, a supervised learning algorithm can build this kind of function—called a model—automatically. It builds this model using a set of training examples, each of which shows what outputs the model should produce for given input values.

1 Rebecca Fiebrink, Daniel Trueman, and Perry R. Cook, "A meta-instrument for interactive, on-the-fly machine learning," in *Proceedings of the International Conference on New Interfaces for Musical Expression* (2009).

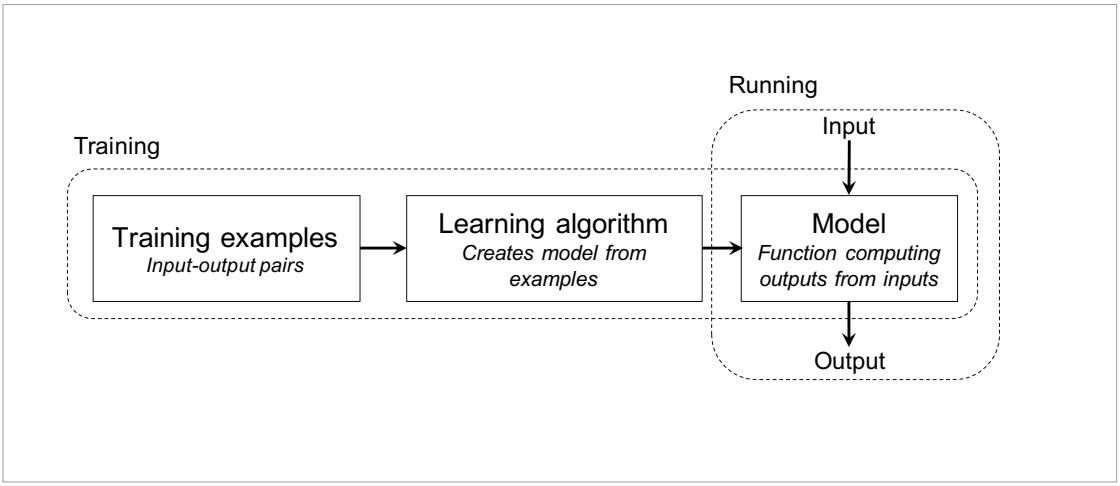

Figure 1: A supervised learning algorithm creates a model function from a set of training examples. This model computes output values from input values. Each training example consists of an input paired with the output that should be computed from that input.

This is the type of ML I have worked with the most, and it is often overlooked in discussions about AI and design. Therefore, I will

be focusing heavily on supervised learning in this paper, with the aim of encouraging readers to think about supervised learning in new ways. In the next section, I describe in more detail how Wekinator and similar interactive ML tools for creative work function, and how they are distinct from the more conventional approaches to supervised learning that tend to be the focus of textbooks, tutorials, and software aimed at computer scientists and data scientists. Then, in Section three, I describe in more detail the implications that this distinctive mode of using ML has for how we should reason about and support ML use in design contexts, followed by a discussion of how ML can change design practices and outcomes.

2. THE WEKINATOR AND INTERACTIVE SUPERVISED LEARNING FOR CREATIVE PRACTICE

When I began to build Wekinator, I was specifically interested in making it easier for creative practitioners to work more effectively with data—especially real-time data from sensors, video, or audio—to make things like gesturally controlled digital musical instruments, interactive art, and embodied games. And I quickly realized that existing ML tools (e.g., Weka[2], R, Matlab) simply were not suitable for such applications. For one thing, there were no general-purpose GUI-based tools that could be easily applied to real-time data. More importantly, the assumptions existing tools made about which activities were central to applied ML practice (and which were excluded) were inappropriate.

Specifically, most existing ML tools at that time (and most tools today) assume that ML practitioners will be working with an existing training dataset that has already been acquired or curated. They assume that practitioners will be experimentally working to apply different modelling techniques—including different learning algorithms and algorithm parameterisations, and potentially also different data representations—in order to find an approach that models the relationship between inputs and outputs well. This experimentation usually involves evaluating candidate approaches using automatically computed, objective metrics of model quality (e.g., cross-validation accuracy, ROC curves).

However, making supervised learning effective for the creative practitioners I was working with in 2008 required a different approach. These collaborators were mainly composers who were interested in creating gesturally-controlled digital musical instruments (e.g., see work by Trueman[3] and Nagai[4]) and/or systems that analysed and reacted in some way (e.g., robot actuation or live visuals) to the musical actions of a human playing an acoustic instrument. In most applications of interest to these practitioners, there was no existing "ground truth" training set available, in which the relationship between inputs and outputs reflected their precise and novel intentions for how the computer should react to the actions of a performer. I therefore designed the Wekinator interface to enable users to

2 Mark Hall, Eibe Frank, Geoffrey Holmes, Bernhard Pfahringer, Peter Reutemann, and Ian H. Witten, "The WEKA data mining software: an update," *ACM SIGKDD Explorations Newsletter 11*, no. 1 (2009): 10-18.

3 Daniel Trueman, "Clapping Machine Music Variations," in *Proceedings of the International Computer Music Conference* (2010).

4 Issue Project Room, "Michelle Nagai: MARtLET", created January 2011, accessed 1 September 2020, https://issueprojectroom.org/news/michelle-naga-martlet.

provide new training examples themselves, often using real-time demonstrations of inputs and outputs [5] [6].

Furthermore, conventional "objective" metrics for evaluating model accuracy were problematic, if not irrelevant, to many of these applications[7]. If a composer is building a gesturally-controlled instrument and finds that the ML approach they have chosen yields a cross-validation accuracy of 95 percent, what does that really mean? Is their model good enough to use? Is it better than an alternative model produced by an approach whose cross-validation accuracy is 93 percent? The answers depend on a variety of complex and subjective factors, involving where and how the model makes mistakes, and how it will be used in practice. Does the model perform as desired on the gestures the composer wishes to include as part of the standard performance technique, or does it make inexcusable mistakes (e.g., mixing up different classes of movement, or failing to register certain actions at all)? Are its mistakes something that players could learn to avoid in practice, by adapting their performance technique? Or are its unintended behaviours actually musically interesting, or even preferable to the behaviours the creator encoded in the training examples? Wekinator therefore facilitates user evaluation of models against these complex, subjective factors by enabling people to experiment with their trained models in real-time, by demonstrating new inputs (e.g., making new gestures) and observing the models' response.

Finally, when such evaluation reveals that a model does not perform as desired, users in such contexts are likely to benefit from different courses of action than more conventional ML practitioners might take. As mentioned above, ML practice often involves experimentally changing the modelling approach (learning algorithm, parameters of the learning algorithm or training method, data representation, etc.) until the model is judged to be acceptable (e.g., with an acceptably low rate of mistakes on new data, and acceptable types of mistakes). However, if someone designing a new interactive system with supervised learning observes that a model makes errors on certain types of inputs, and this designer is capable of demonstrating or curating new training examples, the best course of action may be for the designer to provide additional training examples—which pair these types of inputs with the correct output values—and to build a new model using this augmented dataset. For instance, if the designer is a composer building a gesturally controlled instrument that they intend to perform with themselves, then they might create additional, corrective training examples by simply demonstrating the gestures for which the model responds incorrectly. Or if the designer is building a new gestural controller for use by the general public, and they discover in testing that the model performs poorly on people who are shorter than average, they might collect additional gesture examples from shorter people. These actions may be relatively easy to take, do not require extensive ML expertise, and can often be more effective than changing an algorithm or its parameterisation. Likewise, if a system designer is using a trained model and has an idea for a new model behaviour—for instance, a new sound an instrument might make for a new gesture not present in the training set—it

5 Fiebrink, Trueman, and Cook, "A meta-instrument for interactive, on-the-fly machine learning."

6 Rebecca Fiebrink, Perry R. Cook, and Daniel Trueman, "Play-along Mapping of Musical Controllers," in *Proceedings of the International Computer Music Conference* (2009).

7 Rebecca Fiebrink, Perry R. Cook, and Daniel Trueman, "Human Model Evaluation in Interactive Supervised Learning," in *Proceedings of the SIGCHI Conference on Human Factors in Computing Systems* (2011).

8 Jerry Alan Fails and Dan R. Olsen Jr., "Interactive Machine Learning," in *Proceedings of the 8th International Conference on Intelligent User Interfaces* (2003).

9 Saleema Amershi, Maya Cakmak, William Bradley Knox, and Todd Kulesza, "Power to the People: The role of humans in interactive machine learning," *AI Magazine* 35, no. 4 (2014): 105-120.

10 John J. Dudley and Per Ola Kristensson, "A Review of User Interface Design for Interactive Machine Learning," *ACM Transactions on Interactive Intelligent Systems (TiiS)* 8, no. 2 (2018): 1-37.

11 Kadenze, Inc., "Machine learning for Musicians and Artists," online course by Rebecca Fiebrink, 2015, accessed 1 September 2020, https://www.kadenze.com/courses/machine-learning-for-musicians-and-artists/info

12 FutureLearn, "Apply Creative Machine Learning," online course, created 2020, accessed 1 September 2020, https://www.futurelearn.com/courses/apply-creative-machine-learning

13 Trueman, "Clapping Machine Music Variations."

14 Issue Project Room, "Michelle Nagai: MARtLET."

is often easy and effective to add new examples of this behaviour to the training set, and build a new model from this augmented data. Wekinator therefore allows (and, indeed, encourages) users to rapidly iterate between demonstrating training examples, evaluating a model by experimenting with it in realtime, and editing the training examples and retraining in order to refine the model behaviour.

These activities—creating new training examples by demonstration, evaluating a model by playing with it and exploring it, and iteratively improving or changing a model by changing its training data—are not typically part of the machine learning process in more conventional application domains. However, they are often sensible when people are using ML to create new models of phenomena that they, as designers, deeply understand and when they are capable of easily generating or curating examples for model training and testing. Such an approach is often called "interactive machine learning" in the human-computer interaction domain, where it was originally demonstrated as a way to support image processing [8]. Since then, human-computer interaction researchers have proposed and studied a variety of similar approaches, in which human interaction is inserted into machine learning processes in ways that effectively link human expertise and goals to algorithmic capabilities [9]. Such approaches have been explored in a number of applications involving end-user design, personalisation and configuration, and data analysis and processing. [10]

3. LESSONS FROM TWELVE YEARS OF RESEARCH ON INTERACTIVE SUPERVISED LEARNING IN CREATIVE PRACTICE

Since 2008, I have continued to develop Wekinator in close collaboration with a number of users, and to use it in my teaching online at Kadenze[11] and FutureLearn[12], in person at Goldsmiths and UAL, and in workshops in numerous festivals and universities. The current version (2.x) has been downloaded over 25,000 times and used by artists, musicians, designers, and educators around the world. People have used Wekinator to create gesturally controlled musical instruments [13] [14] [15], and robots[16], experimental games[17], and innumerable examples of interactive and performance art [18]. My research collaborators and I have also used similar approaches to interactive machine learning in new tools. These include the RAPID-MIX API for developers in the creative industries [19], the MIMIC platform for web-based programming with ML and sound[20], the Sound Control project for building musical instruments for children with disabilities[21], and the InteractML tool for game designers[22]. We have also used these tools in new research into how to make interactive machine learning more effective, for instance in exploring how to support end users in devising good data representations for interactive machine learning with motion sensors[23].

In the rest of this section, I will provide an overview of how this body of work has changed the way I think about ML as a tool for design, and how this work illuminates the ways that ML can change design practices and outcomes.

15 Rebecca Fiebrink and Laetitia Sonami, "Reflections on Eight Years of Instrument Creation with Machine Learning," in *Proceedings of the International Conference on New Interfaces for Musical Expression* (2020).

16 Caroline Hermans,"HiBot," created 2016, accessed 1 September 2020, https://caro.io/hibot.

17 Phoenix Perry and Simon Katan, "User-Defined Gestural Interactions through Multi-Modal Feedback," in *Proceedings of the TEI 2016 BodySenseUX Workshop on Full-Body and Multisensory Experience* (2016).

18 See, e.g., http://www.wekinator.org/example-projects/.

19 Francisco Bernardo, Mick Grierson, and Rebecca Fiebrink, "User-Centred Design Actions for Lightweight Evaluation of an Interactive Machine Learning Toolkit," *Journal of Science and Technology of the Arts* 10, no. 2 (2018): 2-25.

20 MIMIC, "Make Music and Art with Machine Intelligence," accessed 1 September 2020, https://mimicproject.com.

21 Samuel Thompson Parke-Wolfe, Hugo Scurto, and Rebecca Fiebrink, "Sound Control: Supporting Custom Musical Interface Design for Children with Disabilities," in *Proceedings of the International Conference on New Interfaces for Musical Expression* (2019).

22 Carlos Gonzalez Diaz, Phoenix Perry, and Rebecca Fiebrink, "Interactive Machine Learning for More Expressive Game Interactions," Ii 2019 *IEEE Conference on Games (CoG)* (2019).

23 Louis McCallum and Rebecca Fiebrink, "Supporting Feature Engineering in End-User Machine Learning," in *Emerging Perspectives in Human-Centered Machine Learning Workshop*, held at CHI 2019, 4 May 2019.

3.1 How Using ML for Design Changes How We Understand and Use ML

My work observing and collaborating with people using Wekinator and similar tools in creative practice has dramatically changed the way I think about machine learning and data. As a computer scientist, I was trained to think about ML data as "ground truth"—as something ML practitioners must try to model as accurately as we can. In this way of thinking, "ground truth" data embodies some hidden truth about the larger world; the role of ML is to uncover and exploit this truth, even (perhaps especially) when the true patterns are complex and defy easy human reasoning or explanation. Increasingly, social science and ML researchers have been rightly calling for more attention to certain problems that arise from such a view of data, including the fact that data collected from the world can include human biases present in the world, leading to models that encode racist, sexist, or other problematic biases themselves[24]. Yet much of such work ultimately advocates for methods capable of identifying or mitigating bias in ML models, or questions whether ML should be used at all[25]; such work generally does not question whether the goal of ML (where it is appropriate to use ML) should remain the creation of models that accurately capture and generalise from patterns present in training data.

In contrast, when creators work with Wekinator or similar interactive machine learning tools to design new interactive systems, data often plays a slightly different role: namely, training examples are the means through which creators communicate their ideas and intentions to the computer. The goal of using ML may not be to uncover a "hidden truth" in these data; instead, it may be to build a useful model that behaves in a way that matches the goals of the designer generating or curating the data.

This is a subtle difference with important implications for practice. For one thing, understanding training examples as a mechanism for communicating designers' intentions makes it clearer why enabling designers to iteratively modify training data is so important. Designers' ideas for what they want to build often evolve over time, as they implement and interact with prototypes to better understand the design space; they must be able to communicate these changing specifications to the computer through changes to training examples. It is also unreasonable to expect a designer to perfectly communicate their desires for a model in an initial training set that is used by a learning algorithm—whose inner workings may be complex and not well-understood by the designer—to build a model that accurately captures those desires on the first try. However, as a designer interacts with a trained model, they may be able to identify inputs for which its behaviour is not what is desired, and it can then be simple to reinforce those desires through additional training examples. Through iteration, then, the designer can learn how to provide training examples that best help the algorithm learn [26 27].

Understanding example data as a representation of a designer's intentions also has consequences for how we design ML tools

24 Emily Sheng, Kai-Wei Chang, Premkumar Natarajan, and Nanyun Peng, "The Woman Worked as a Babysitter: On Biases in Language Generation," *arXiv preprint arXiv:*1909.01326 (2019).

25 Cathy O'Neil, *Weapons of Math Destruction: How Big Data Increases Inequality and Threatens Democracy* (Broadway Books, 2016).

26 Rebecca Fiebrink, "Real-time Human Interaction with Supervised Learning Algorithms for Music Composition and Performance" (PhD diss., Princeton University, 2011).

27 Fiebrink, Cook, and Trueman, "Human Model Evaluation in Interactive Supervised Learning."

for design. For instance, we might consider how such tools can support people in creating good training datasets that reflect their intentions well, or how they can facilitate exploring complex navigations through different datasets that represent different design ideas (e.g., supporting "undo" and "redo" operations on edits to training data, as Wekinator does, or perhaps enabling building multiple models on alternative datasets in parallel to facilitate comparison of alternative ideas).

This understanding additionally has interesting implications for the types of ML algorithms we may choose to use in interactive design contexts. For instance, in most conventional applications of ML, the goal of the learning algorithm is to build a model that generalises well to new data, and it is presumed that a model that captures all nuance and variation in the training data to the point of being able to model the training data perfectly is counterproductive to this goal. Such a model is said to "overfit," and contemporary learning algorithms typically include implicit or explicit mechanisms for "regularisation" to prevent such overfitting [28]. However, when a system designer is creating a training dataset using input sensors or other data sources with little noise—that is, the designer has very precise control over the values encoded in the training examples—using a learning algorithm that fits the training set very precisely (i.e., "overfitting") may be desirable[29] [30]. Such an algorithm allows complex model functions to be learned efficiently from fewer training examples, and algorithms that attempt stronger regularisation may feel frustrating to users who do not perceive that the model is adequately sensitive to their wishes communicated in the data.

Understanding data as a mechanism for communicating user desires to the computer leads to a number of additional ways in which it can be reasonable to break the "rules" of how ML should be applied in more conventional contexts. For instance, conventional approaches to evaluating model quality (especially using methods such as cross-validation accuracy which attempt to approximate generalisation accuracy) may not be appropriate. Not only do such measures fail to capture the complex and subjective criteria that are crucial for system designers to understand if they are to decide whether a model is good enough to use in practice, or if they are to improve a model using subsequent changes to its training data, but these measures may provide misleading results (e.g., see work by Fiebrink, Cook, and Trueman [31] for an example and explanation of a context in which cross-validation scores and user satisfaction scores for a series of models exhibited a negative correlation).

I have written in more detail about some of the differences between conventional ML and interactive ML in the context of designing musical systems in an earlier paper[32], and Dudley and Kristensson[33] have provided an excellent survey of interfaces for interactive ML across many domains. I refer interested readers to these sources for further discussion about how using ML in design may change our thinking about ML and about how to support its effective use.

28 Christopher M. Bishop, *Pattern Recognition and Machine Learning* (Springer, 2006).

29 Rebecca Fiebrink and Baptiste Caramiaux, "The Machine Learning Algorithm as Creative Musical Tool," in *Oxford Handbook of Algorithmic Music*, ed. R. Dean and A. McLean (Oxford University Press, 2016).

30 Fiebrink, Cook, and Trueman, "Human Model Evaluation in Interactive Supervised Learning."

31 Ibid.

32 Rebecca Fiebrink and Baptiste Caramiaux, "The Machine Learning Algorithm as Creative Musical Tool."

33 John J. Dudley and Per Ola Kristensson, "A Review of User Interface Design for Interactive Machine Learning."

3.2 How ML Can Change Design Practices and Outcomes

3.2.1 ML Can Enable More People to Become Creators

Designing new systems with data can enable more people to become creators. No programming skills are necessary for someone to design and implement fairly sophisticated, novel interactions with sensors, audio, or video using suitable interactive ML tools. When I teach workshops with Wekinator, I provide students with pre-built examples of feature extractors (i.e., small standalone programs that take data from cameras, microphones, games controllers, mice, etc. and send this data to Wekinator), as well as example programs that can be controlled by models in Wekinator (e.g., for sound synthesis and animation). In these workshops, people from a wide variety of backgrounds—as well as children as young as eight years old—have learned to make novel interactions by combining an example feature extractor, an example control program, and new models they train in Wekinator.

Of course, more specialised tools can offer tailored support for creating new systems in a particular domain. In the recent Sound Control project[34], my team implemented interactive ML software specifically for music therapists and music teachers to create bespoke musical instruments and interactions for the children they work with, who have a wide variety of disabilities. This software enables users to choose from a selection of common input devices, to configure music playback or synthesis, and to use interactive ML to specify how music and sound will respond to children's actions with the input devices. Its interface requires no programming knowledge, nor any knowledge of sensors or data communication and processing. On the other hand, its interface takes advantage of users' existing musical knowledge and practices, for instance enabling users to connect external MIDI devices or record their own sound samples to use in their designs. As a result, the people we worked with were able to quickly create a huge variety of richly customised musical interactions.

3.2.2 ML Can Facilitate More Effective Work with Sensors, Audio, Video, and Media

Designing new systems using ML rather than programming can also unlock new creative possibilities for working with sensors, audio and video, and other media. Often, even expert programmers may struggle to write code that accurately recognizes complex phenomena from such data, such as human motion and activities, speech and music, environmental states, and so on. But machine learning can make it feasible to build computer systems that understand and respond to these phenomena. Many people using Wekinator in creative practice do have some programming expertise, but they nevertheless use ML to create things they could not create otherwise[35] [36].

3.2.3 ML Can Support Rapid Prototyping and Design Exploration

34 Samuel Thompson Parke-Wolfe, Hugo Scurto, and Rebecca Fiebrink, "Sound Control: Supporting Custom Musical Interface Design for Children with Disabilities."

35 Rebecca Fiebrink, "Real-time Human Interaction with Supervised Learning Algorithms for Music Composition and Performance."

36 A number of example Wekinator projects appear at http://www.wekinator.org/example-projects.

Implementing a new interactive system with Wekinator or similar tools can be much faster than designing by writing a new program in code. ML can allow creators to come up with a first working prototype very quickly, often using a small number of training examples, and to quickly refine it using a small number of additional examples. When I deliver talks about Wekinator, I often perform a live demonstration in which I iteratively train a gesture-controlled drum machine, in which data from my laptop webcam (downsampled to a 10-by-10 pixel grid) is sent to Wekinator, where a classifier categorises my gestures into one of four classes. A simple drum machine program receives a message with the current class from Wekinator and uses this to control which sound is played. Building the first version of this drum machine takes under a minute (including launching and configuring the software, as well as recording training examples and building a model), and each iteration of refinement (recording new examples and retraining the model) may take a few seconds. In total, I am usually able to produce a satisfactory simple controller in a few minutes with a classifier trained on a few hundred examples (taken from under 20 seconds of demonstrations), at most.

Research on tools for design and other creative work emphasises that rapid and iterative prototyping is key to enabling people to succeed in design tasks [37]. When people are able to build, and modify prototypes easily, they are able to explore design spaces more broadly, and to refine their final design more extensively. Indeed, my initial goals in creating Wekinator centred around making work with sensors and audio data more effective and accessible to creators. But one of the first surprises in my early engagements with users was how much they valued the ability to simply implement and refine designs more quickly and—as a result—explore more ideas [38]. Today, after more than a decade of teaching and observing people creating new interactions with ML, I view its support of faster prototyping and exploration as one of its key creative benefits.

3.2.4 ML Can Support the Communication of Embodied and Tacit Knowledge to Computers

Designing using example data and ML rather than programming alone can make it much easier to build computer systems that understand or respond to embodied practices and tacit knowledge. Embodied practices—for instance, the ways that someone moves, or makes sound, or manipulates an object—play a central role in many endeavours that are crucially important to people: music, art, sports, and many others [39]. Similarly, people possess many kinds of tacit knowledge [40] that is hard to articulate explicitly—this can include knowledge about goals and preferences, or procedural knowledge of how to complete some task, from making a sandwich to performing expert image editing in Photoshop. It can be much easier to communicate with computers about such practices by providing or curating examples—for instance, demonstrations of human action, or example pieces of media or existing designs that exemplify some design goal—rather than by describing them in code or

37 Mitchel Resnick, Brad Myers, Kumiyo Nakakoji, Ben Shneiderman, Randy Pausch, Ted Selker, and Mike Eisenberg, “Design Principles for Tools to Support Creative Thinking,” in *Report of Workshop on Creativity Support Tools* (Washington, DC, USA, 2005).

38 Rebecca Fiebrink, Daniel Trueman, N. Cameron Britt, Michelle Nagai, Konrad Kaczmarek, Michael Early, M. R. Daniel, Anne Hege, and Perry R. Cook, “Toward Understanding Human-Computer Interaction in Composing the Instrument,” in *Proceedings of the International Computer Music Conference* (2010).

39 Paul Dourish, Where the Action Is: *The Foundations of Embodied Interaction* (MIT Press, 2001).

40 Michael Polanyi, *Personal knowledge: Towards a Post-Critical Philosophy* (University of Chicago Press, 2015).

math. Michelle Nagai, one of the first composers to build a new instrument—called the MARtLET—with Wekinator, described how this changed her approach to instrument building:

> I have never before been able to work with a musical interface that allowed me to really 'feel' the music as I was playing it and developing it. The Wekinator allowed me to approach composing with electronics and the computer more in the way I might if I was writing a piece for cello, where I would actually sit down with a cello and try things out.

41 Rebecca Fiebrink, Daniel Trueman, N. Cameron Britt, Michelle Nagai, Konrad Kaczmarek, Michael Early, M. R. Daniel, Anne Hege, and Perry R. Cook, "Toward Understanding Human-Computer Interaction in Composing the Instrument."

42 Rebecca Fiebrink and Laetitia Sonami, "Reflections on Eight Years of Instrument Creation with Machine Learning."

3.2.5 ML Can Support Access to Surprise and Discovery

The professional creators with whom I have worked often value the ability for machine learning to change what they build—not just by making design more efficient, but by providing access to productive surprises [41]. Productive surprises are not often part of the computer programming process: it is usually a bad thing to be surprised when programming, as this generally means there is a bug. Fixing a bug demands a dramatic change in mindset: one stops the act of creation to carefully reason about what may have gone wrong, perhaps switching from writing code to running a debugger, and perhaps employing iterative experimentation to trace the source of the problem. In the best case, this process resolves the bug quickly, and the disruption to other tasks is minimal.

In contrast, ML systems can often surprise a user in good ways, by providing access to computer behaviours that are unanticipated, but still relevant to the design task. For instance, a regression model trained to control sound synthesis parameters in response to a performer's motions may yield unexpected combinations of parameters, which may produce sounds not in the training data—or indeed sounds never heard before. Likewise, when regression is used to make simultaneous, fine-grained adjustments to multiple parameters controlling the generation of images, or game environments, or other media, it is quite possible that trained models will exhibit outputs not present in the training set, which together represent underexplored areas of these high-dimensional control spaces. Such surprising outputs may be unwelcome by creators (in which case creators can often adjust model behaviours to exclude them by providing corrective training examples), but occasionally they can lead a creator to discover new, fruitful ideas.

In a recent publication, composer Laetitia Sonami discusses some of the techniques she has developed over the past eight years for exploiting this potential for useful surprises in her work building new instruments with Wekinator [42]. Sonami intentionally crafts training datasets with "wider" or "narrower" variability in the target output values, in order to produce regression models with greater or lesser propensity to surprise, which she calls the models' "predictability index." She has also added infrastructure to her performance software that allows her to artificially "freeze" some of the sensor values sent to Wekinator during performance so that she can "catch" interesting sounds that arise in performance and spend time exploring them.

3.2.6 Designers Using ML Can Benefit from ML Knowledge

Despite the potential benefits described above, ML algorithms often do not perform exactly as we wish they would. Myriad problems can arise, preventing a designer's efforts to communicate their goals through the training data from yielding a trained model that wholly realises these goals. These problems may include insufficient or poorly chosen training examples, a poor choice of data representation (i.e., "features"), a mismatch between the learning algorithm or its parameters and the designer's goals, or other choices about structuring the learning problem (e.g., using classification when regression would be better, or failing to break a complex learning problem into simpler, easier-to-learn components).

After many years working informally with practitioners using ML, I concluded that, if I wanted to enable creators to use ML more effectively, I needed to find effective ways to teach them certain things about ML. In the case of very special-purpose software like Sound Control, users at minimum need to understand the basic "machine learning pipeline" for interactive ML (Figure 1): how training examples are used to build a model, how this model can be run to compute new outputs from new input values, and how this model can be changed by changes to the training examples. In the case of software like Wekinator that aims to enable users to create a wide variety of systems in a wide variety of domains, users sometimes need substantially more ML knowledge to effectively implement their more complex ideas (as well as to reason about practical trade-offs between different design options).

The curriculum I designed for such creators became the foundation for online classes on Kadenze[43] and FutureLearn[44], as well as in-person teaching at UAL, Goldsmiths, and workshops around the world. The curriculum arises from a set of learning objectives structured around Howard, Culley, and Dekoninck's[45] "creative design process", considering the ML-specific knowledge that is helpful in establishing design requirements, specifying desired system behaviour, implementing the system, evaluating the system, and iteratively reformulating the design requirements, specification, and/or implementation.

Some of these learning objectives require knowledge that is common to all or most applied machine learning practice, including basic knowledge of how different algorithms build models from data, of how to design good data representations that are amenable to ML, and how to reason about different approaches to evaluating ML. Yet they all also require knowledge that is specific to the use of ML in design, and/or to the particular application domain of the design. For instance, designers need specialised knowledge of how to generate good datasets that communicate their intentions clearly, and they must reason about how to evaluate models when generalisation is not their only or most important priority. Creators working with data from certain sensors or video, for instance, also benefit from an understanding of common methods for representing such data (e.g., certain signal processing and "feature engineering" techniques).

43 Kadenze, Inc., "Machine learning for Musicians and Artists."

44 FutureLearn, "Apply Creative Machine Learning."

45 Thomas J. Howard, Stephen J. Culley, and Elies Dekoninck, "Describing the Creative Design Process by the Integration of Engineering Design and Cognitive Psychology Literature," *Design Studies* 29, no. 2 (2008): 160-180.

It is also important for creators to develop competence in structuring complex projects involving multiple models and/or complex dataflows, for instance projects in which multiple analyses happen in parallel (e.g., simultaneously analysing motion of multiple limbs) or in series (e.g., using one model to detect and segment an event and a second model to categorise the event). This can be aided by developing a fluency in what might be called "creative ML design patterns"—common approaches to structuring machine-learning-based systems that tend to aid both efficient machine learning and human reasoning.

In my work with creative machine learning students, I have also found it important to use different teaching approaches than those typically used in teaching ML to computer science or data science students. For instance, deep knowledge of the mathematics behind how various algorithms learn from data may not be particularly helpful. However, hands-on experimentation with scaffolding technologies for interactive ML and project-based work with tools that integrate into creators' existing workflows are an important part of the curriculum.

I have described my approach to teaching ML to creative practitioners in greater detail in a recent article [46]. It includes a more detailed discussion about why it is important to teach ML to creative practitioners, a description of the learning objectives and the teaching strategies I have developed to aid students in achieving these objectives, and an analysis of data from my university and MOOC students that informs an initial understanding of how the curriculum supports student learning, as well as an understanding of students' challenges. In other recent research[47], a colleague and I have explored mechanisms for addressing the specific challenge of aiding creators who lack deep signal processing and feature engineering knowledge in deriving good data representations for human movement sensors.

46 Rebecca Fiebrink, "Machine learning education for artists, musicians, and other creative practitioners," *ACM Transactions on Computing Education* 19, no. 4 (2019): 1-32.

47 Louis McCallum and Rebecca Fiebrink, "Supporting Feature Engineering in End-User Machine Learning."

3.2.7 ML Can Potentially Give Users More Agency and Control Over Technology

At this moment in time, I believe many people's prevailing experience with ML is one of disempowerment: certain corporations and governments benefit enormously from collecting and applying ML to data captured from people's lives, without those people always feeling like they fully understand, consent, or have the power to change how their data is used.

This contrasts greatly with the experiences of people using ML tools like Wekinator, Sound Control, InteractML, and others. Users of these tools have full control over the ML process, and they are able to take advantage of the benefits of ML—its accessibility to non-programmers, its ability to model complex patterns in sensor or media data, its utility in capturing embodied expertise and tacit knowledge—to make things that have value to them and the people around them.

In writing this paper so far, when I have discussed "users" of ML tools, I have implicitly assumed these users to be designers and

other creators (whether professionals, novices, or students). Yet, the users of designed products might themselves be users of ML: they might train or otherwise interact with ML to create things that are important to them. I believe there is a large, unexplored design space of new systems in which end users might employ ML as a tool for design and customisation, in which ML can provide them with more agency and control over technologies they use.

This design space includes specialised design tools like Wekinator or Sound Control, which employ training data provided by their users to support end-user design or programming in some domain. It also includes systems in which users interact in a more limited way with ML, but still use their own data to train or refine ML systems, enabling them to customize or configure systems to their liking. It furthermore includes systems in which users employ ML to gain insight or value from their own data in other ways, for instance systems that aid users in better understanding the data generated through the course of their interactions with social media, fitness trackers, personal finance systems, or software platforms for professional or creative work.

4. CONCLUSION

In this paper, I have summarised a number of findings from my work over the past twelve years on creative uses of interactive supervised learning, particularly related to how using ML in design suggests a distinct understanding of ML, and to how using ML in design can change design processes and outcomes. I encourage readers who are interested in the details of this work to consult the many papers I have cited above, as well as to explore the design tools we have created (e.g., Wekinator, MIMIC) and the publicly available documentation for the numerous projects people have created with these tools.

Although I have focused on supervised learning, a number of the points in this paper are relevant to work with other types of ML as well. For instance, generative ML approaches (e.g., for generating new images, text, sounds, design specifications, or other digital artefacts from a corpus of examples of such artefacts) are also potentially accessible to non-programmers. Their users may employ data in a similar manner, as a means to communicate intentions which might be ephemeral and contingent, rather than as a “truth” to be modelled. Choosing or generating training examples can again provide an easy means to communicate embodied and tacit knowledge, and changing this data may often be a reasonable approach to refining model behaviour. In both the case of supervised learning and other ML approaches, there remain vast, unexplored opportunities to build better tools for designers of all types to employ ML effectively, to understand more about how and what to teach designers and other non-computer-scientists using ML to build new things, and to incorporate ML into new technologies in ways that empower their end users.

BIBLIOGRAPHY

Amershi, Saleema, Maya Cakmak, William Bradley Knox, and Todd Kulesza. “Power to the People: The Role of Humans in Interactive Machine Learning.” *AI Magazine* 35, no. 4 (2014): 105-120.

Bernardo, Francisco, Mick Grierson, and Rebecca Fiebrink. “User-Centred Design Actions for Lightweight Evaluation of an Interactive Machine Learning Toolkit.” *Journal of Science and Technology of the Arts* 10, no. 2 (2018): 2-25.

Bishop, Christopher M. *Pattern Recognition and Machine Learning*. Springer, 2006.

Diaz, Carlos Gonzalez, Phoenix Perry, and Rebecca Fiebrink. “Interactive Machine Learning for More Expressive Game Interactions.” In 2019 *IEEE Conference on Games (CoG)*, pp. 1-2, 2019.

Dourish, Paul, Where the Action Is: *The Foundations of Embodied Interaction*. MIT Press, 2001.

Dudley, John J., and Per Ola Kristensson. “A Review of User Interface Design for Interactive Machine Learning.” *ACM Transactions on Interactive Intelligent Systems (TiiS)* 8, no. 2 (2018): 1-37.

Fails, Jerry Alan, and Dan R. Olsen Jr. “Interactive Machine Learning.” In *Proceedings of the 8th International Conference on Intelligent User Interfaces*, pp. 39-45. 2003.

Fiebrink, Rebecca. “Real-time Human Interaction with Supervised Learning Algorithms for Music Composition and Performance.” PhD diss., Princeton University, 2011.

Fiebrink, Rebecca. “Machine learning education for artists, musicians, and other creative practitioners.” *ACM Transactions on Computing Education* 19, no. 4 (2019): 1-32.

Fiebrink, Rebecca, and Baptiste Caramiaux. “The Machine Learning Algorithm as Creative Musical Tool,” in *Oxford Handbook of Algorithmic Music*, edited by R. Dean and A. McLean. Oxford University Press, 2016.

Fiebrink, Rebecca, Perry R. Cook, and Daniel Trueman. “Play-along Mapping of Musical Controllers.” In *Proceedings of the International Computer Music Conference*. 2009.

Fiebrink, Rebecca, Perry R. Cook, and Daniel Trueman. “Human Model Evaluation in Interactive Supervised Learning.” In *Proceedings of the SIGCHI Conference on Human Factors in Computing Systems*, pp. 147-156. 2011.

Fiebrink, Rebecca, Daniel Trueman, N. Cameron Britt, Michelle Nagai, Konrad Kaczmarek, Michael Early, M. R. Daniel, Anne Hege, and Perry R. Cook. “Toward Understanding Human-Computer Interaction in Composing the Instrument.”

In *Proceedings of the International Computer Music Conference*. 2010.

Fiebrink, Rebecca, Daniel Trueman, and Perry R. Cook. "A meta-instrument for interactive, on-the-fly machine learning." In *Proceedings of the International Conference on New Interfaces for Musical Expression*. 2009.

Fiebrink, Rebecca, and Laetitia Sonami. "Reflections on Eight Years of Instrument Creation with Machine Learning." In *Proceedings of the International Conference on New Interfaces for Musical Expression*. 2020.

FutureLearn. "Apply creative machine learning." Online course, created 2020. Accessed 1 September 2020. https://www.futurelearn.com/courses/apply-creative-machine-learning.

Hall, Mark, Eibe Frank, Geoffrey Holmes, Bernhard Pfahringer, Peter Reutemann, and Ian H. Witten. "The WEKA data mining software: an update." *ACM SIGKDD Explorations Newsletter* 11, no. 1 (2009): 10-18.

Hermans, Caroline. "HiBot." Created 2016. Accessed 1 September 2020. https://caro.io/hibot.

Howard, Thomas J., Stephen J. Culley, and Elies Dekoninck. "Describing the Creative Design Process by the Integration of Engineering Design and Cognitive Psychology Literature." Design studies 29, no. 2 (2008): 160-180.

Issue Project Room. "Michelle Nagai: MARtLET". Created January 2011. Accessed 1 September 2020. https://issueprojectroom.org/news/michelle-naga-martlet.

Kadenze, Inc. "Machine learning for Musicians and Artists." Online course by Rebecca Fiebrink, 2015. Accessed 1 September 2020. https://www.kadenze.com/courses/machine-learning-for-musicians-and-artists/info.

McCallum, Louis and Rebecca Fiebrink. "Supporting Feature Engineering in End-User Machine Learning." In *Emerging Perspectives in Human-Centered Machine Learning Workshop*, held at CHI 2019, 4 May 2019.

MIMIC. "Make Music and Art with Machine Intelligence." Accessed 1 September 2020. https://mimicproject.com.

O'Neil, Cathy. *Weapons of Math Destruction: How Big Data Increases Inequality and Threatens Democracy*. Broadway Books, 2016.

Parke-Wolfe, Samuel Thompson, Hugo Scurto, and Rebecca Fiebrink. "Sound Control: Supporting Custom Musical Interface Design for Children with Disabilities." In *Proceedings of the International Conference on New Interfaces for Musical Expression*. 2019.

Perry, Phoenix and Simon Katan. “User-Defined Gestural Interactions through Multi-Modal Feedback.” In *Proceedings of the TEI 2016 BodySenseUX Workshop on Full-Body and Multisensory Experience*. 2016.

Polanyi, Michael. P*ersonal knowledge: Towards a Post-Critical Philosophy*. University of Chicago Press, 2015.

Resnick, Mitchel, Brad Myers, Kumiyo Nakakoji, Ben Shneiderman, Randy Pausch, Ted Selker, and Mike Eisenberg. “Design Principles for Tools to Support Creative Thinking.” In *Report of Workshop on Creativity Support Tools*, Washington, DC, USA, 2005.

Sheng, Emily, Kai-Wei Chang, Premkumar Natarajan, and Nanyun Peng. “The Woman Worked as a Babysitter: On Biases in Language Generation.” *arXiv preprint arXiv:1909.01326* (2019).

Trueman, Daniel. “Clapping Machine Music Variations.” In *Proceedings of the International Computer Music Conference*. 2010.

dai digital conference
Impressions

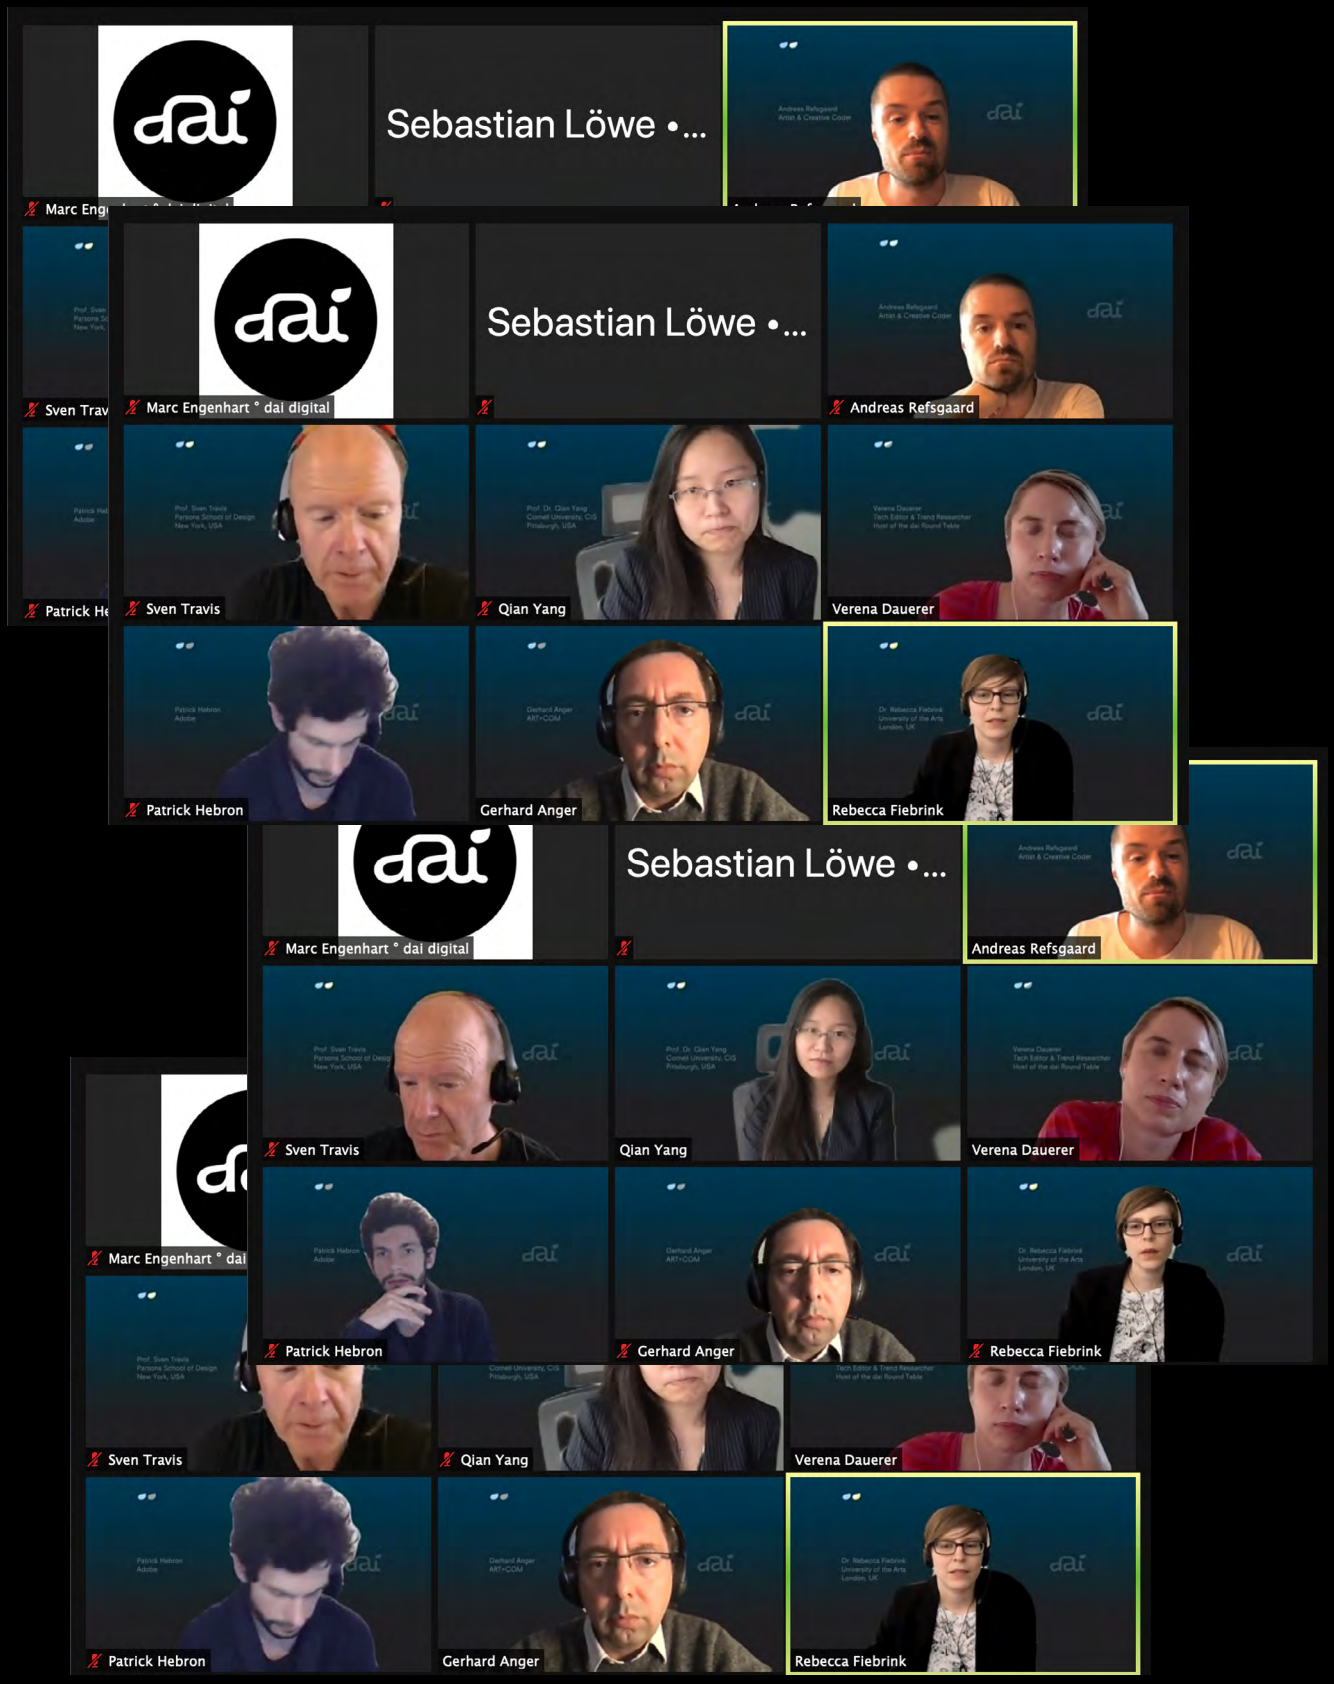

Round table of the dai digital conference
September 17, 2020

Playful Machine Learning: An Interaction Designer's Journey into Neural Networks

Andreas Reefsgard

Practise

This is the companion publication to Refsgaard's opening keynote address for the 2020 dai digital. The keynote can be accessed at https://www.designing-artificial-intelligence.eu.

ABSTRACT

In this paper interaction designer and artist Andreas Refsgaard describes how machine learning has become an integrated part of his practice. While still relying on coding for most of his work, training machine learning algorithms has changed his attitude towards a lot of problems and enabled him to see a whole new range of opportunities for projects.

Through his playful projects Andreas is actively seeking out unconventional connections between inputs and outputs using simple machine learning techniques. Games are controlled by making silly sounds, music is composed by drawing instruments on paper and algorithms are trained to decide what is funny, funky or boring. Recently Andreas has been experimenting with using machine learning to create new texts and images based on large amounts of training data and has created Booksby.ai—an online bookstore which sells science fiction novels generated by an artificial intelligence. Overall, the paper demonstrates how machine learning can be used to create playful interactions, solve problems in human-computer interaction and help us ask important questions about the future of design.

1. INTRODUCTION

We live in a time where artificial intelligence and machine learning seems to be everywhere. Not a week goes by without stories of how algorithms and robots are going to take our jobs, drive our cars, become our lovers, replace our pets or supersede our doctors. Along with all the buzz comes a lot of misconceptions and myths. When we think of machine learning or artificial intelligence, we think of automation, not creative opportunities. Additionally, artists and designers are often scared to enter the field, imagining that they need computer science PhDs and big server parks just to get started. And that is a pity, since there are lots of low hanging fruits out there for all creatives.

Trained as an interaction designer from Copenhagen Institute of Interaction Design (CIID) I am what you might call a creative technologist working within the field of art and design using code as my main tool. But during the last five years machine learning has become an integrated part of my practice, despite the fact that I have never had any formal education within the field or academic prerequisites beyond my long forgotten high school math.

2. CODE VERSUS TRAINING

While I still write code, training machine learning algorithms has somewhat changed my attitude towards a lot of problems and enabled me to see a whole new range of opportunities for projects. But what is the real difference between the two approaches?

When programming interactive prototypes, interaction designers traditionally rely on their ability to formulate logical structures and explicit relationships between inputs and outputs through code that executes in a predictable way.

Machine learning suggests a different kind of logic: Instead of relying on explicit sets of rules to determine a system's behavior, machine learning models learn by example, by looking for patterns within a set of examples or training data from a designer or performer, and makes the rules autonomously so as to conform to the performer's expectation. This pattern recognition process is somewhat similar to our own mental processes for learning about the world around us and provides a lot of new opportunities for interaction designers, especially when dealing with input data too complex to account for via coding.
To put some of these thoughts into context I will present a few projects built using fairly simple and accessible machine learning techniques.

3. WOLFENSTEIN 3D CONTROLLED BY SOUND

The first project is a hack of the old first-person shooter video game Wolfenstein 3D made with Lasse Korsgaard for our former studio Støj. Instead of relying on mouse and keyboard controls to play the game, we trained an algorithm to recognize and distinguish between a set of sounds and mapped each one to a different action within the game. Whistling sounds move the character forward, clapping sounds open doors, two different grunt sounds turn you left and right and as for shooting, players say "Pew pew", of course.

The training in itself took no more than a few minutes, and essentially just required us to record a dozen examples for each of the different sounds. Although in no way perfect, this quick and dirty approach proved accurate enough to control the game.

We initially did this project to have something fun to show to a client interested in interactive machine learning. At first glance, it may seem a bit ridiculous, and the fact that it is not exactly easier to play games in this manner goes without saying. In an article from gaming website Kotaku covering our project, reporter Logan Booker came to the conclusion that he did not *"see this method of game-playing catching on, outside of some specific circumstances"*.[1]

1 Booker, Logan, "Wolfenstein 3D Played With Only Claps, Whistles And Grunts," 2017, accessed August 25th, 2017, https://www.kotaku.com.au/2017/03/wolfenstein-3d-played-with-only-claps-whistles-and-grunts.

Training alternative controls for a game or another interactive system, no matter how ludicrous they may seem, does however hint towards bigger goals and opportunities for interaction design and interactive art experiences. By enabling people to decide upon and train their own unique controls for a system, the creative power shifts from the designer of the system to the person interacting with it. In the case of Wolfenstein controlled by sound, the game simply becomes something else.

The shifting of power from the designer to the person interacting also points towards more inclusive systems, where users are free to change the systems to fit their preferred modes of interaction.

4. EYE CONDUCTOR

Back in 2015 I made Eye Conductor, a system that helps people express themselves through music using only their eyes and facial gestures, as my final project at Copenhagen Institute of Interaction Design. Using a $99 (USD) eye tracker and a regular webcam, Eye Conductor detects the gaze and selected facial movements, thereby enabling people to play any instrument, build beats, sequence melodies or trigger musical effects.

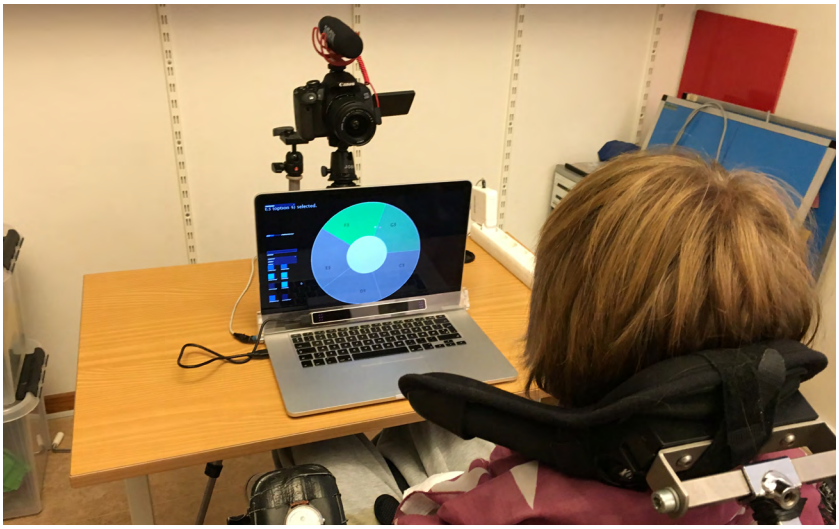

Figure 1: Andreas Refsgaard, Eye Conductor. 2015

The motivation for the project came from the realization that for a lot of people with physical disabilities the lack of fine motor skills excludes them from playing music using traditional hand-held instruments. I therefore designed around eye and facial movements, assuming that all potential users would have full motoric control of their facial muscles. Realizing how this was not always the case, I started wondering how a system could be flexible enough to have its core interactions be defined by the specific person using it. Instead of the designer deciding that raising one's eyebrows to a certain level would toggle some specific function, what if people could use unique facial expressions, which they felt comfortable performing and map those to functions inside the program? This spurred my interest in machine learning.

5. DOODLE TUNES

My first steps into machine learning started with a short course at School of Machines, Making & Make-Believe in Berlin taught by independent artist and programmer Gene Kogan and Dr. Rebecca Fiebrink from Goldsmiths, University of London. The course introduced a range of techniques for using machine learning in a creative way and I quickly started building small projects using resources from Kogan's ml4a project (Machine Learning for Artists) as well as Rebecca Fiebrink's Wekinator software (both highly recommended starting points for creatives to get into machine learning).

Among the projects I made during the course was an algorithm trained on canned laughter trying to determine how funny scenes in the sitcom Friends were. Another project was an audio visual installation trained to distinguish "funky images" from "boring ones" after having been trained on 15000 images of highly subjective boring (carpets, lawyers, empty offices, etc.) and funky (party, graffiti, drinks etc.) images.

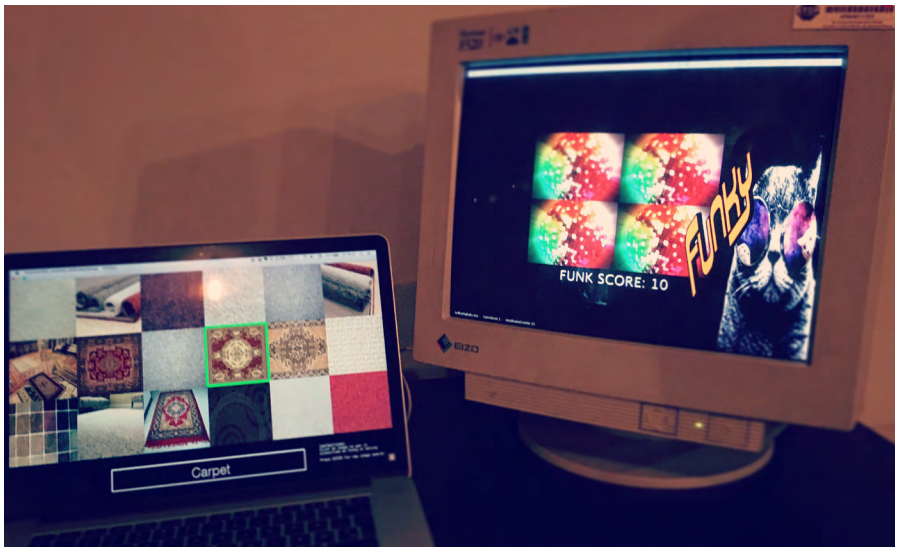

Figure 2: Andreas Refsgaard, Is It Funky? 2016

Later on, Gene Kogan invited me to participate in the Nabi Art Center's Global AI Hackathon in Seoul, South Korea, where we made the project Doodle Tunes together.

Built in two days, Doodle Tunes is a project that lets you turn drawings of musical instruments into actual music. After drawing one or more instruments on a regular piece of paper and positioning the drawing beneath a camera, the system begins playing electronic music with any detected instruments.

Figure 3: Gene Kogan, and Andreas Refsgaard, Doodle Tunes. 2016

Similar to the sound controlled Wolfenstein project, where new sounds are compared to the sounds made during our training phase, Doodle Tunes works by classifying new hand drawn instruments on the basis of a model trained on other hand drawn instruments.

From an interaction design perspective both projects excel in their novel combination of inputs and outputs and ability to expand people's conception about the scope of digital interactions, by combining domains that are not usually related. One of my favorite pieces of art, "The Gift" by Man Ray has always inspired me, because its logic works in a similar way. Consisting of a flatiron with fourteen thumb tacks glued to its sole, The Gift is a conjunction of two seemingly alien objects and an example of how juxtaposing or combining objects from two different domains can create something new and thought provoking—in this case a paradoxical tool for ripping cloth or perhaps piercing walls.

A lot of my personal work follows these lines of thought, actively seeking out unconventional mappings between inputs and outputs: What if your gaze works as an input for playing music (Eye Conductor), what if your drawings do (Doodle Tunes) or what if sounds could control video games (Wolfenstein 3D Controlled by Sound)?

6. AN ALGORITHM WATCHING A MOVIE TRAILER

An algorithm watching a movie trailer was made in collaboration with Lasse Korsgaard. We were playing around with the real-time object detection system YOLO-2 and were curious to explore how a fast paced movie trailer might look when watched through the lens of a machine learning model.

Object detection is the process of identifying specific objects such as persons, cars and chairs in digital images or videos. For most humans, this task requires little effort, regardless of how the objects may vary in size, scale and rotation, or whether they

are partially obstructed from their view. For a long time these tasks have been difficult for computers to solve, but recent developments have shown impressive improvements in accuracy and speed, even while detecting multiple objects in the same image.

After experimenting with modern classics like Titanic and The Lord of the Rings, we chose the trailer for Martin Scorsese's The Wolf of Wall Street because of its fast-paced cuts between scenes in relatively common environments like restaurants, offices, boats and houses filled with tables, glasses, computers, cars, and other objects suitable for image detection. As German artist Hito Steyerl once stated: *"The unforeseen has a hard time happening because it is not yet in the database"*, and logically, the algorithm was better at detecting businessmen in suits drinking wine than the Uruk-hai and Nazgul from Tolkien's Middle-earth universe.

Instead of simply outputting the found objects on top of the original video, we made three different versions of the trailer, each using a different filter. The first video filter uses object masking, so that only objects recognized by the software appear. The second version blurs all detected objects, thus acting like an automatic censoring algorithm. The final version removes the visuals entirely, essentially creating a filter of what the software "sees" during analysis.

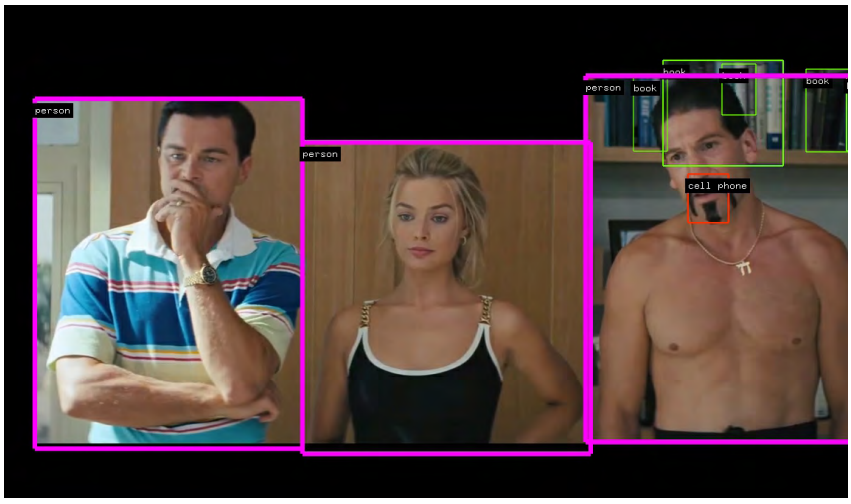

Figure 4: Andreas Refsgaard, and Lasse Korsgaard, An Algorithm Watching a Movie Trailer. 2017

The project got a lot of media attention and in an article covering the project, Sidney Fussell from Gizmodo described the project in this way:

> If a movie trailer distills a 2-hour film into its 3-minute essentials, what would it look like to distill a movie trailer? Strangely, it would look a lot like object recognition software. In three separate videos, we essentially see how algorithms watch movies: They label the essentials—a tie, a wine glass, a chair—but leave the specifics out. It's like visual ad-libs.[2]

2 Fussel, Sidney, "Watching a Computer Identify Objects in a Movie Trailer Is Entertaining as Hell," 2017, accessed August 25th, 2017, http://sploid.gizmodo.com/watching-an-algorithm-try-to-decode-a-movie-trailer-is-1791288059.

The ability to detect objects and determine their bounding boxes opens up a whole new potential for remixing and personalizing video content. Lasse and I also did a low fidelity exploration in this direction by letting visitors on a website switch between the original version of the intro to Full House, an American sitcom, or a version where all detected objects were replaced by emojis.

Figure 5: Andreas Refsgaard & Lasse Korsgaard, Full House Emojified. 2017

Desirable or not, prototypes like this investigate a future where software could automatically censor parts of movies based on the age and sensitivity of the viewer, or perhaps replace actors that viewers dislike with their favorite ones.

7. FUTURE PERSPECTIVES

Training machine learning algorithms has changed my attitude towards a lot of problems and enabled me to see a whole new range of opportunities for projects. Besides expanding our possible input-output combinations, I see big potential in making people active participants rather than passive consumers of whatever experiences and tools we build. I am interested in designing for surprises, both for the people interacting with my systems and myself, and want to create interactions that give people a high sense of agency. When a tool or a piece is interactive, you do not know in advance how people will use it, which makes the experiences all the more open-ended. And with the type of machine learning I am using, people can potentially train the input methods themselves, making the interactions even more unpredictable and less determined by us, the designers.

A project that really encapsulates these thoughts is Objectifier by Bjørn Karmann, built in 2016 as his thesis project from Copenhagen Institute of Interaction Design. Objectifier empowers people to train objects in their daily environment to respond to their unique behaviors. It contains a camera, a computer, and a relay, and it can be used to turn any electrical device plugged into it 'on' or 'off'—based only on what it sees. This allows people to train the system to turn on the radio by doing a certain pose or turn off a lamp, when nobody is sitting

at a desk. This contains a shift in user experience and training artificial intelligence, a shift from being a passive consumer to becoming an active, playful director of domestic technology.

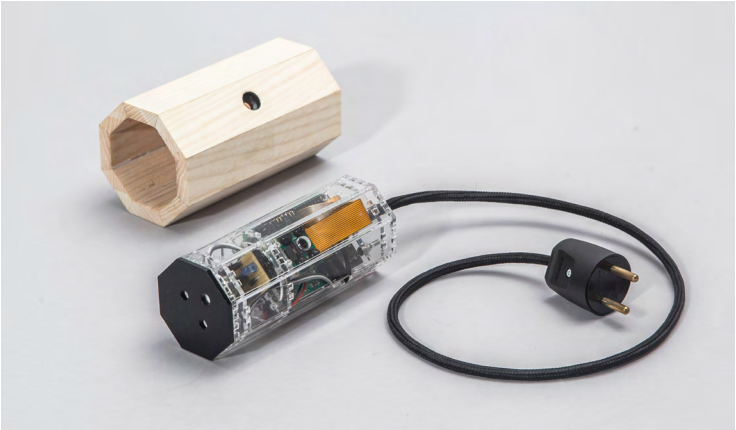

Figure 6: Bjørn Karmann, Objectifier. 2016

8. FROM CLASSIFICATION TO GENERATION

Over the last few years, as I have personally moved slightly away from interaction design and more towards the digital art scene, my use of technology has also changed. Instead of using machine learning to make digital systems interpret the world by classifying inputs, I am leveraging it to generate artworks by itself, mostly in the form of texts or images.

To unfold this change, I will present three of my text-based projects: Poems About Things, BooksBy.ai and fAIry tales.

9. POEMS ABOUT THINGS

Figure 7: Andreas Refsgaard, Poems About Things, 2018

Poems About Things is a project that generates poetry from everyday objects around us. It consists of a mobile website that

constructs quirky sentences about an object detected through the user's camera feed. As the user's camera focuses on an object, a built-in machine learning model gives its best guess as to which object it is seeing. Based on this guess, a short query is sent to Google's Suggest API, which in return sends back a list of sentences inspired by the detected object. Put together, these sentences form a poem which appears on the screen overlaying the image of the object. Hence the poem consists of a handful of sentences expressing thoughts, questions or comments related to the immediate object as well as the world outside of it. And sometimes the mistakes in the classification can lead to quite poetic outputs, as in the case of the hairy chest leading to questions about the realness of a fur coat.

10. BOOKSBY.AI

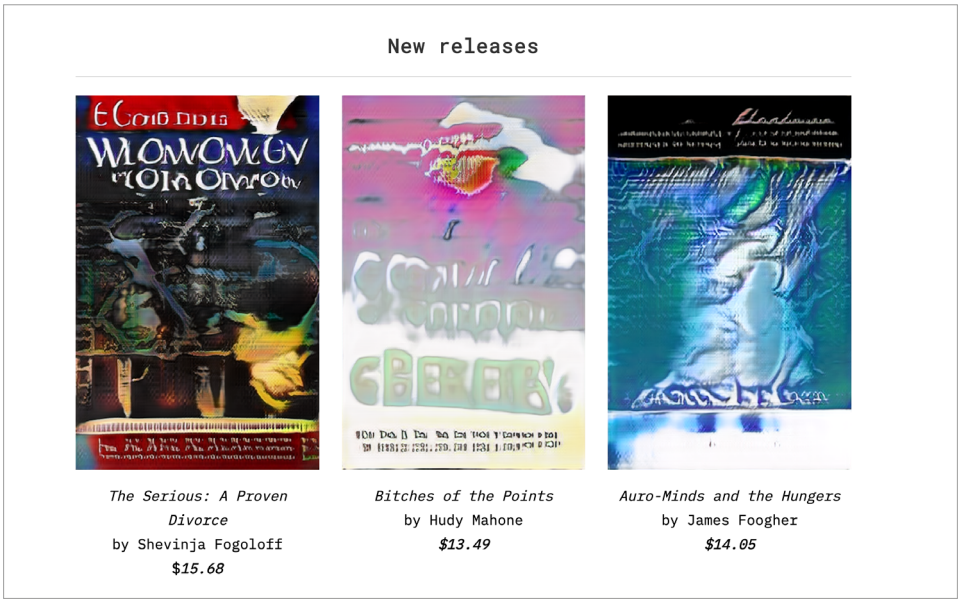

Figure 8: Andreas Refsgaard, and Mikkel Thybo Loose, Booksby.ai, 2018

Booksby.ai is an online bookstore which sells science fiction novels written—or better generated—by an artificial intelligence. I made the project in collaboration with Danish data scientist Mikkel Thybo Loose. Through training, we exposed the artificial intelligence to a large number of science fiction books. Hence we taught the artificial intelligence to generate new science fiction books based on the language, style and visual appearance of the training examples. All books on BooksBy.ai are for sale on Amazon.com and can be ordered as printed paperbacks. Judged by the text quality Booksby.ai is admittedly not the greatest example of machine-driven creativity, but the site presents an intriguing example of an entire ecosystem created by artificial intelligence. By exploring the idea of fully automated, machine learning generated cultural products, the project invites readers to speculate about a future where human beings are no longer the sole producers of creative cultural products.

11. FAIRY TALES

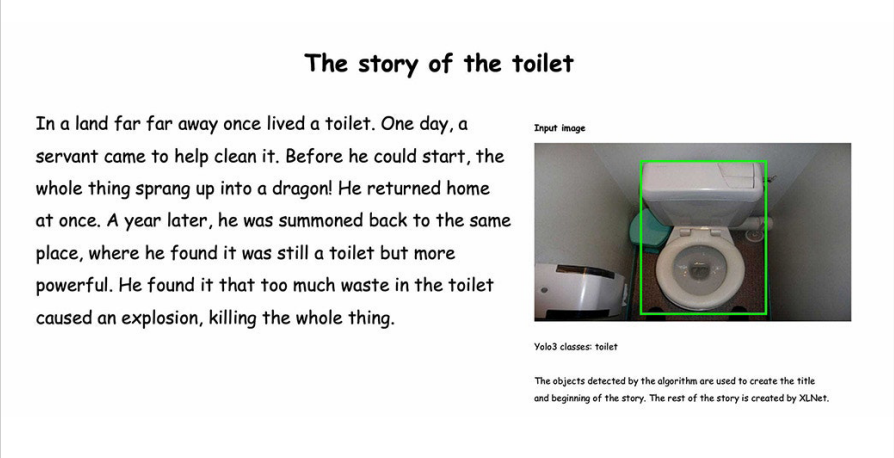

Figure 9: Andreas Refsgaard, fAIry tales, 2019

fAIry tales is a project where fables are automatically generated based on algorithmically detected objects in mundane images. Using images from cocodataset.org everyday objects are detected using the YOLOv3 object detection algorithm. From these objects the title and the beginning of a story is generated and used as prompts for text generation model XLnet.

Much like *Poems About Things and Booksby.ai* the project is an exploration of computer written fiction. As in *Poems About Things* the texts produced are unique to the pictures fed into the system, but instead of single line poetry outputs the fairy tale texts are longer and form more coherent stories. Where BooksBy.Ai focused on creating science fiction texts, fAIry tales aims to generate fairy tales, and succeeds in producing more coherent plot lines and sentences than Booksby.ai, although surprising and absurd narratives often arise due to the attempt to create adventures about boring objects.

12. CLOSING THOUGHTS

Artificial intelligence and machine learning holds a huge potential for creatives, and we have just started seeing the possibilities which the technology will make available to creatives in the future. We—coders, developers, artists, designers and people in the many other fields using machine learning—still have lots to learn and many black boxes to grasp. But by getting hands-on with projects and creating the training data ourselves we get a useful feel for the creative opportunities and inevitable biases of this intriguing and powerful technology.

BIBLIOGRAPHY

Booker, Logan, "Wolfenstein 3D Played With Only Claps, Whistles And Grunts," 2017, accessed August 25th, 2017, https://www.kotaku.com.au/2017/03/wolfenstein-3d-played-with-only-claps-whistles-and-grunts.

Fussel, Sidney, "Watching a Computer Identify Objects in a Movie Trailer Is Entertaining as Hell," 2017, accessed August 25th, 2017, http://sploid.gizmodo.com/watching-an-algorithm-try-to-decode-a-movie-trailer-is-1791288059.

dai digital conference
Impressions

dai digital conference posters
September 2020

Learning to Educate Machines

Patrick Hebron

Research

This is the companion publication to Hebron's opening keynote address for the 2020 dai digital. The keynote can be accessed at https://www.designing-artificial-intelligence.eu.

ABSTRACT

While machine learning has drawn influence from a diverse range of fields such as biology and thermodynamics, its connection to the fields of education and anthropology have been dramatically under-explored. In the human world, we understand that the intellectual capacities of the individual cannot exist in a vacuum. Rather, they are built upon the transference of distilled knowledge and our relation to other individuals who help us to notice the world and reflect upon our actions within it. In the machine world, our tools for shaping the learner's experience are still woefully crude. In this paper, we will look at teaching and curriculum design as creative acts, essential to the advancement of artificial intelligence and to the furtherance of the creative partnership between humans and machines.

1. SEEING THE OTHER

We are inclined to imagine artificial intelligence (AI) through an anthropomorphic lens. This is understandable. We have no better example of intelligence than ourselves and therefore view human intelligence as defining 'intelligence' in general. But there is much to be learned about the external world and about ourselves by understanding other beings through their own lens—by not conforming them to our image of ourselves. As Jean-Paul Sartre says, "the Other is the indispensable mediator between myself and me."[1] That is, we cannot truly understand our own intelligence without contrast.

1 Jean-Paul Sartre, *Being and Nothingness* (New York: Washington Square Press, 1992), 302.

Admittedly, there are already other greatly intelligent creatures in our midst—primates, dolphins, elephants, octopi and so forth. But their worlds and the problems they face may be too different from ours (and in other ways too similar) for us to truly make sense of ourselves through the contrastive lens of these species. The other provides a kind of shortcut for discovering alternate ways of being in and navigating the world. In Reinforcement Learning, there is talk of the Explore vs Exploit tradeoff—look for a new strategy or continue using one that has already proven successful. By observing the other, we can sometimes glean successful strategies without the labor of exploring for ourselves. The other's strategy may utilize very different premises or facts about the environment from those we have encountered in our own explorations. If our own experiences have not led us to those same facts or if the strategy's successful outcome cannot be immediately observed, then we are unlikely to discern the value of the strategy a priori.

In AlphaGo's famous "Move 37," the commentators thought the move was a mistake. The value of the move was ultimately confirmed by the machine's win and led Lee Sedol and other Go

experts to re-evaluate their understanding of the game. But if the machine happened to lose that particular game, the potential value of the strategic approach underlying this move should not necessarily be discarded. The educational philosopher, Patricia Carini, said:

> I have to trust that what I am attending to makes sense; that it is not a merely accidental or chance event. To discover the subject's coherence and how it persists in the world, I have deliberately to shift my own perspective in relation to it.[2]

2 Patricia Carini, *Starting Strong: A Different Look at Children, Schools, and Standards* (New York: Teachers College Press, 2001), 163.

Seeing value in the other's approach before it is confirmed requires open-mindedness. There is an opportunity cost to entertaining alternate strategies. Outside of virtual environments, it is often the case that the scenario can only be played-out once—making it impossible to compare divergent strategies. Why would anyone ever take a chance on a strategy that does not jive with one's own intuition and experience?

It would be reasonable to conclude that observing the other is simply another mode of exploration, which requires time commitment without guaranteed reward. This is true, but even without that confirmation, even if the other's approach turns out to be flawed, it gives new vantage points on how to be in the world. It still involves a kind of exploration, but it is an exploration of the world already mediated, already trampled through. In this sense, it gives us new premises and vocabulary that we may be unlikely to uncover ourselves. It may redirect us from the paths of exploration to which we are otherwise inclined.

In the human realm, a certain emotional availability to the prospect of learning through the observation of other humans comes from the fact that they are human. We see a "whole self" behind the work of another human and we put trust in the likeness of that self to our own—the experiences and properties they presumably share.

The AI is not a whole self, or at least not one in the sense to which we are accustomed. An anthropomorphic presentation of the AI helps us to project a sense of wholeness upon it, thereby making it a more palatable subject from which to learn.

The AI is a specialist. But in some key respects, we are specialists too. We would not be able to pilot an octopus body. Similarly, an artificial general intelligence will have a specific domain of experience, specific sensors and actuators, even if it can reason in a general way. But those experiences and capabilities will be different from ours. How can we learn to see value in them? How can we learn to use them in our own growth?

The educational theorist Howard Gardner offered the "Theory of Multiple Intelligences," which differentiates human intelligence into specific 'modalities,' rather than seeing intelligence as a single general ability. Despite the flaws and limitations of this theory, through a humanistic lens, we may see this work as trying to show value in the effort of students whose areas of specialization may not be otherwise valued by our educational system. Kofi Annan said:

> Never walk into a situation believing you know better than the natives. Keep an open mind. You have to listen and look at how they do things. Otherwise you can make some very serious mistakes. [3]

3 Kofi Annan, "A Champion for Development, Security and Human Rights," interview by Elaina Loveland, *International Educator*, (May-June 2013): 24.

4 Geoffrey Hinton, source unknown.

There is much to be learned from a specialist operating in a complex domain with which we have limited familiarity. Of course, the funny thing about these AI specialists, these "natives," is that they did not really land on the shores of the specialized domain before we did. We pointed them to it and determined which aspects of it would be visible to them. So, this endeavor requires something from both the human and the AI. We also determine much about the environment and experiences from which our human children learn. But it is ultimately they who must do the learning. Once we set them upon the shores, they may be able to spend more time there than we could ever hope to ourselves. It will become theirs and if we set them up to explore it fully, then we may benefit from the time they spend there.

2. CREATION AND INFLUENCE

As Geoff Hinton said, "If I want to change the ideas in your head, I can't just reach in and tweak what the neurons are doing. I have to change the experiences you learn from."[4] To put it another way, intelligence cannot be created, it can only be influenced.

As much as we think we teach human children to do math, we can really only set up the right circumstances for them to learn it themselves. Even if we can teach arithmetic by rote memorization, human progress by definition depends upon the next generation advancing beyond the knowledge of the previous. As such, we need to cede control. We need to stop teaching and instead let learning happen. We do this by constructing the right environment for learning to occur.

From this point of view, I believe the future of machine intelligence has more to do with design than with engineering. Engineering, or at least good engineering, means precisely controlling the behaviors of a system. Design, or at least good design, means creating a system that enables the user to go beyond the original intent of the designer—a system that can be adapted, leveraged, built upon.

For machine intelligence to be of true value to our society, it must go beyond explicitly engineered behaviors and knowledge. It must learn for itself. It must exceed what we have imbued upon it. If it is to be useful to us, then we must create an environment through which the machine's exceeding of our explicit engineering will still adhere to the values and goals of its creator and of the society in which that creator is embedded.

3. HOW WE TEACH MACHINES TODAY

Though neural architectures have advanced by leaps and bounds over the last few decades, the manners in which we teach machines and design curriculums for them remain incredibly crude. To better understand this, let's take an anthropomorphic look at three forms of machine learning:

The *supervised learning* system is much like a Spelling Bee champion. It does nothing but look at examples, millions of examples, of one specific kind of information. Its experience of the world is completely lacking in diversity. It can perform spectacularly at one task, but can do nothing else. In the human realm, it is unlikely that this kind of experience leads to a truly innovative thinker.

The *unsupervised learning* system is much like a singer performing a song in a language he or she does not speak. Like the supervised system, this learner does nothing but look at millions of examples. But here, the goal is different. The system learns the internal patterns present within a set of example inputs so that it can produce very similar patterns. Yet, this ability to parrot the examples it has learned from never really leads to an understanding of what those patterns mean or purpose they serve outside of themselves.

The *reinforcement learning* system is much like a child abandoned in the wilderness. It is left entirely to its own devices and must learn through pure trial-and-error in a world that either rewards or punishes its actions. This unaided manner of learning about the world is not very efficient, but if the individual can get enough experience without getting killed, it will develop a strong intuition for navigating the world. Through the notion of a 'wolf child', we have reason to suspect that what really makes us human is not our human brains, but rather our being raised within human society, within human families.

What does a teacher do for the student? In contrast with the pure trial-and-error approach of reinforcement learning, a teacher helps to distill relevant observations that will make the learner's process more efficient. Teachers help the student to notice the world and its properties. They help the student to notice his or her own actions within the world. For example, a teacher may point out that a bike will stay upright only if it is already in motion or that the rider's knee is turned out a bit too far.

The teacher also helps the student to plot a course through increasingly complex materials. How does a person reach the ability to read academic literature and synthesize new insights from it? You cannot just put an infant in a library and expect a medical expert to emerge! The curriculum needs to be scaled up by introducing increasingly complex concepts progressively and by pointing the student to some salient, high-level observations at each stage in complexity scaling.

4 VISUALISING EXPERIENCE

I want to return our focus to the ways in which AI can help us to explore the world more effectively and to the idea that:

> The other's strategy may utilize very different premises or facts about the environment from those we have encountered in our own explorations.

This is a daunting problem, which should not be overlooked in the design of systems that task machines with accomplishing something that is to be of value to humans. In some cases, human comprehension may not be necessary. One such example is DeepMind's system for optimizing electrical grid usage. Here, the system's utility from a human point-of-view is derived purely from its outcome. If electrical consumption is reduced without interfering with human behavior (i.e. without causing blackouts), we have no real need to understand the strategies taken by the machine nor the experiences that led the machine to those strategies.

Similarly, if an AI system were to develop a successful vaccine or treatment for Covid-19, we would rejoice in and benefit from the machine's achievement, regardless of whether we understood its strategy or methodology. Of course, it might be beneficial for future virology research to be able to incorporate the machine's ideas into academic literature. But, if this were not possible, we could instead simply set ourselves on the idea that this technical discipline could advance even in the absence of human comprehension. So long as its technical goals continued to be met, we would have no real cause to require greater human comprehension.

There are many scenarios, though, in which this sort of blind faith cannot be sufficient. If the success of the system cannot be measured in purely objective terms, then we need to design systems that will make the machine's strategy more comprehendible to us—helping us to see for ourselves the machine's experiences, the premises or facts about the environment, that led the machine to those strategies.

Watching footage of the reinforcement agents in OpenAI's "Hide and Seek" gives us something more substantive than mere personification.[5] It gives us the opportunity to see and to relate to the scenarios played-out by the learning system. It enables us to not only accept the machine's conclusion, but see firsthand how it arrived there.

There is great value in the human trainer being able to observe the dataset or the training scenarios upon which the machine was trained. The difficulty in this, of course, is the volume of experience required by the machine to reach its novel capabilities. Supervised learning datasets routinely include millions of individual artifacts. Reinforcement learning systems must often perform tasks millions of times. It is simply impossible for us to be present for all of that.

5 "Emergent Tool Use from Multi-Agent Interaction," OpenAI, September 17, 2019, https://openai.com/blog/emergent-tool-use.

The subject of explainability in machine learning has been an increasingly prominent one. The hope, it seems, is that the machine would be able to distill its experience into some higher level (probably symbolic) form in order to convey the logic that underlies its strategy. But, it is important to note that this is often not possible even for a human expert.

When a radiologist is able to associate a particular pattern in an image with a particular diagnosis, that association comes as something much like a sensory-motor response. Having seen many positive and negative examples before, the radiologist has a kind of impulse response. The pattern simply triggers the association in the radiologist's mind. It seems to come as a spark of intuition. There may be no better explanation than a reference to the voluminous cloud of prior experiences.

Perhaps, then, we should take the stance that in many cases, explainability in AI will simply mean the system's ability to point to particularly instructive training samples or scenarios. Through a highlight reel of the machine's experience, or through the machine's being able to interactively recall a particular experience on demand for the human inquisitor, we may find our way to a shared understanding of the machine's insights.

5. THE CONSTRUCTIVISM OF OTHER INTELLIGENCES

How does better teaching of machines come back to us? How does it extend our reach? The answer is that the roles of teacher and student must form a loop. By letting the student learn and through our observation of their manner of learning, we take on a new perspective. We have to interact with the learner's learning. This brings us to our final destination—constructivism.

The great constructivist, Seymour Papert, said in his seminal book, *Mindstorms*:

> Even the best of educational television is limited to offering quantitative improvements in the kinds of learning that existed without it. „Sesame Street" might offer better and more engaging explanations than a child can get from some parents or nursery school teachers, but the child is still in the position of listening to explanations. By contrast, when a child learns to program, the process of learning is transformed. It becomes more active and self-directed. In particular, the knowledge is acquired for a recognizable personal purpose. The child does something with it. The new knowledge is a source of power and is experienced as such from the moment it begins to form in the child's mind.[6]

To interact with something is to internalize it—to map it, to graft it onto oneself. It is in this spirit that we must encounter the experience of the artificial intelligence. To learn from its novel insights, we must bring our own experiences into proximity with those that led the machine to the insights in the first place. Our tools today for building, training and visualizing the experiences

6 Seymour Papert, Mindstorms: Children, *Computers and Powerful Ideas* (New York: Basic Books, 1980), 20-21.

of machines are hopelessly crude. Above all else, this is a design problem. As with any design problem, start from the user—in this case, there are two.

BIBLIOGRAPHY

Annan, Kofi, "A Champion for Development, Security and Human Rights," interview by Elaina Loveland, *International Educator*, (May-June 2013): 24.

Carini, Patricia, *Starting Strong: A Different Look at Children, Schools, and Standards* (New York: Teachers College Press, 2001), 163.

"Emergent Tool Use from Multi-Agent Interaction," OpenAI, September 17, 2019, https://openai.com/blog/emergent-tool-use.

Papert, Seymour, *Mindstorms: Children, Computers and Powerful Ideas* (New York: Basic Books, 1980), 20-21.

Sartre, Jean-Paul, *Being and Nothingness* (New York: Washington Square Press, 1992), 302.

dai digital conference
Impressions

Social media communication
September, 2020

Drop That Intelligence and Get on with It!

Gerhard Anger

Discourse

This is the companion publication to Anger's opening keynote address for the 2020 dai digital. The keynote can be accessed at https://www.designing-artificial-intelligence.eu.

ABSTRACT

Thinking of machine learning systems as "intelligent" imbues them with a mystical property and credits them with independence they do not have. Instead, these systems must be perceived as permanently flawed. Even if there is no conscious design decision to make them biased, they end up being exactly that. A thorough criticism of ML decisions is in order and a new discourse on the notion of machine intelligence.

1. INTRODUCTION

If you look at much of what is said about the impact of "artificial intelligence" on design (and not just on design but on basically every human endeavor from chicken farming to cosmology), you are left positively awestruck.

You quickly learn that the impact will be nothing less than unprecedented, fundamental, disruptive, and of course thoroughly spectacular. In fact, some of the stuff sounds like if the Second Coming was imminent.

Or, to use another item from mythology or religion, we're obviously about to find the philosopher's stone, which according to the alchemists of old will have the following beneficial properties: Turn base metals into gold, common crystals into diamonds, heal all forms of illness and prolong the life of anyone who consumes it, preferably diluted in wine. Perpetually burning lamps, which suggests that we'll solve all our energy issues, revive dead plants, clone ourselves and even create artificial humans or homunculi.

Some other strands of thought try and sell us a property of machine learning systems that has been proposed before as well. Neatly classifying and codifying the world has been tried several times. Leibniz, for example, conceived of a characteristica unversalis, an algebra that was supposed to be able to express everything that can be said and thought. Had he succeeded, there would be no arguments, and of course no further need for conferences like the dai digital. You could decide whether something was true or false by simply calculating it, and you could generate all possible truths. Suffice it to say that it did not happen then and it will not happen now. One of the many issues with approaches like that is that they impose an order on the world that is not inherent to it. It is always flawed and must be.

Anyway, enough of mystical musings. I will talk about expectations, reality, the various strange cognitive dissonances I perceive where machine learning is concerned and about how we should really approach machine learning as a concept and

technology by removing the “intelligence” part when we think about it, talk about it or apply it. Thinking of those systems as “intelligent” imbues them with a mystical property, credits them with independence they do not have — if we do not think they have them. By crediting them with those properties, we not only endanger many things we have already achieved, but we make it unnecessarily harder for ourselves to use machine learning in a beneficial and creative way.

2. AVOID MYSTICISM AND REACTIONARY RESULTS

Last year I was asked to think about an idea for a small fun application for an exhibition. The target audience of the application was children, and the idea basically was a classification system that would decide whether the child was male or female, and according to that decision some kind of gender appropriate costume was to be used to augment the image.

Nothing wrong with that, or is there? As a society, I think we have already progressed beyond the strictness of binary gender definitions. Since 2018, for example, intersex people in Germany are allowed to choose between “female”, “male”, “diverse” and no gender marker at all in their birth certificate. And what exactly are gender appropriate costumes for kids? Pink skirts and superhero costumes?

Why is it that a seemingly state of the art hypermodern development like “artificial intelligence” leads to such a naive usage transporting rather conservative notions like that of fixed gender roles? We started with classifying cats and dogs, which is comparatively easy and of course very cute. But not everything is so easily classified, nor should it be. This, in my opinion, is one of the profound dissonances in the field of machine learning. We create conservative and sometimes outright reactionary and racist classification systems oblivious of anything outside the norm, but disguised as the hottest new thing there is.

As a society, we spend quite some effort to overcome prejudice and become more inclusive. There is a danger that, through the backdoor, we will reintroduce very ugly things into our daily practice, like for instance the pseudoscience of phrenology, which was supposed to give a quite precise description of one’s character based on the measurements of one’s skull. The perfidious thing is, that in most cases there is of course no hidden agenda of racism or willful exclusion of disabled people. There is no conscious design decision that leads to that behavior. It just happens somehow, and if you’re not considerate and savvy, your image classification system will categorize people of color as gorillas, or your soap dispensers will only recognize white hands.

And if you imbue a system based on statistics with a mystical property like “intelligence”, there is always the danger the decisions this system makes or the limitations that come with it are perceived as god given and unavoidable, which of course

they are not. There has always been a certain privilege attached to being like most other people. The rise of applied statistics in the guise of artificial “intelligence” makes it even more important to consciously design applications that do not exclude people by deciding that they are not significant. By the way, the application assigning gender appropriate costumes to kids was not implemented after all.

3. GENERATING SOMETHING NEW

Would it not be great if we could ask the great luminaries of former times, now regrettably deceased, what they think about our world and the stuff that happens in it? In the course of a research project, my team was asked to think about a system that, once fed all the writings of Alexander von Humboldt, a famous polymath and great traveler in the 19th century, could give Humboldt-like answers to modern questions. Well. Maybe, with quite a lot of effort, one might train a system that is able to generate something resembling coherent thought roughly in the literary style of Humboldt.

But more than that? Certainly not. That idea again shows a very rough and simplified world view. Complex and interconnected concepts like romantic philosophy or atheism are not easily translated into statistic structures, and I think they never will be. And would it, in the end, really be that interesting after all to have a Humboldt machine? Why is it anyway that we are so fascinated by the idea that we might teach a machine to compose in the style of Bach, create a painting in the style of Rembrandt or speak with the voice of Alexander von Humboldt?

Is it something like the “noble simplicity and quiet grandeur” that’s in the heart of the neoclassicist program? Do we all suffer from imposter syndrome and think our own age cannot aspire to the greatness of the past? Or is it simply a question of pleasing the cultured elite (which in the end decides where the money goes), and this is easier done by creating another Rembrandt painting than by doing something really new? Here we have another dissonance that again points to deeply conservative concepts at the heart of the matter. We have the supposedly most sophisticated machinery ever created, the one technology to supersede all technologies. And what are we currently doing with it? We improve selling stuff that people do not really need, propagate concepts of individual mobility that have already brought us to the brink of destruction and engage in necromancy trying to revive the giants of old.

And of course, machine learning systems ca not do much more than look back in time. The data that is used to train them is always from the past, and the systems age quite quickly if they are not properly and constantly maintained. There are, of course, very creative ways of taking stuff from the past and doing something meaningful and fun with it. Remix culture in all its varieties comes to mind, for example in several modern musical styles. The constant reinterpretation of ancient Greek

mythology from the poems of Ovid to the admittedly somewhat cheesy Clash of the Titans movie show that you can of course do something meaningful and fun with old material.

But even then, the creative part, as with cooking, is not done by more or less randomly putting known ingredients in a pot and stir them, doing that a billion of times or so and see if some of the results are edible. That is another aspect in all of this I want to talk briefly about. There is a craft part in a lot of the things we are doing when we create something new, whether it is creating a new dish, designing a new chair or developing a software product. And there is personal experience. Theoretical speculation will only get you so far. If you never sat on a chair in your whole life but only know of chairs from literature and images, what will be the result if you try and design one?

And I know of no major composer in the history of music who didn't have practical experience on at least one musical instrument. That practical experience even had profound influence on some, like for example Anton Bruckner who was an organist by profession and wrote symphonies that in a sense are huge and wonderful organ pieces for orchestra. Machine learning systems may be able to capture certain aspects of the world, the ones we feed them, and of course, standing in a very long tradition, we mostly feed them the theoretical, mathematical, structural parts of the world. But that is not the whole world there is.

We again should remove that mystical intelligence part from our thinking about machine learning. There is no machine that can be taught to think like Humboldt by showing it all his writings for a countless number of times. And a machine will probably never write a piano piece that feels so good and natural to play as the pieces by, say, Bach, Scarlatti or Chopin. And if there was, how would that really benefit us after all?

4. DO NOT REQUIRE HUMANS AND SOCIETIES TO ADAPT

Technologies and interfaces are always somewhat flawed, and we are constantly expected and forced to adapt to them. For example, people who only have a two character first name (like the son of a friend of mine who is simply named "On") will every now and then be unable to register with a site because valid first names must at least contain three characters. We have to accommodate voice recognition systems by speaking in a particular way, and we sometimes even have to seek medical assistance as some interfaces like smart phones or computer mice are causing us bodily harm. Despite all promises of "intelligence", without conscious effort by designers and implementers this is going to become much worse with the extended use of machine learning in interface and interaction design.

It is, as I have tried to convey several times, in the nature of machine learning systems to condense the diversity of the world

to a structure that is manageable. And there is nothing wrong with that per se. But this must in my opinion not mean that you have to be like most people (or rather like most people responsible for designing and implementing the interfaces you are supposed to interact with) to be able to use the interface.

There are dialects. There are many people with disabilities. There are many people who look different from the norm. And there are many people who are not of or above average height and from European or Asian descent. Maybe one of the most important responsibilities of designers and implementers of user interfaces is not to forget that. Automated statistics will only get you so far, and if its wide application means that the diversity of the real world has to go away then this is too high a price to pay.

5. CONCLUSION

But I do not want to be too somber. By all means, let us use machine learning to do fun stuff! Let us use it to create tools that may spare us tedious work that is on the fringes of our creative processes. Let us use it to make good use of the huge amount of information that is being produced and that may well be unmanageable without the use of such technologies.

But we must not forget the entity who is solely entitled to make decisions here. There is a tendency in those discussions about artificial intelligence that seems to downplay the capabilities and also the responsibilities of humans. This is not a healthy approach to a new technology.

If we all really suffer from imposter syndrome, then we should ask ourselves why. And this is maybe the biggest dissonance. On one hand we say that we, humble mankind, created or about to create a machinery that is more powerful than we can possibly imagine (you remember: unprecedented, fundamental, disruptive, and of course thoroughly spectacular.). But on the other hand we immediately feel threatened by that notion and fear that it may somehow make us redundant.

Anyway. I suggest that we regain some conscious awareness of what we are doing and why, drop that notion of artificial intelligence and just get on with it doing fun, creative and meaningful things with the interesting technologies we develop.

dai digital conference
Impressions

dai digital conference schedule
September, 2020

Unremarkable AI: Towards AI That Co-lives and Co-evolves with Users

Qian Yang

Research

This is the companion publication to Yang's opening keynote address for the 2020 dai digital. The keynote can be accessed at https://www.designing-artificial-intelligence.eu.

ABSTRACT

Artificial intelligence (AI) has become an attractive and effective tool for innovating the design of modern technologies. From mundane spam filters that save people time to conversational agents that offer a more human mode of interaction, to personalized media recommendations, it can seem like AI is in almost every new technology product and service. Both designers and technologists note this trend and become especially interested in design opportunities surrounding AI. Some even speculate that AI is the new user experience (UX). However, many challenges persist in designing and innovating human experiences of AI. For example, AI systems can make unpredictable errors, and these errors damage UX and even lead to undesired societal impact. This raises the question: How can we, as designers and creative technologists, better work with AI's technical advances?

1. INTRODUCTION

The promises and perils of artificial intelligence are becoming ubiquitous in people's lives and societies. For example, machine learning systems promise to improve healthcare outcomes by providing complimentary data-driven insights to clinicians. Meanwhile, these systems could embody the historical health inequalities infused in clinical data such as racial disparity.[1]

My research focuses on innovating human-centered AI systems. In the past seven years, we designed a variety of AI applications based on empirical studies of their actual stakeholders and real-world impacts. These applications span the domains of healthcare, mobile computing, social networks, and more. Each design addressed a critical challenge in translating AI from a technological advance to a socio-technical system valuable to human ends. We investigated how the researchers and industry practitioners leveraged—or failed to leverage— AI in their respective domains of interest. This includes empirical studies with nearly 150 industry practitioners, offering rare and rich industry best practices in crafting human-AI interaction, as well as the persisting challenges [2, 3].

This keynote is a reflection on these experiences. I present a vision of the future where AI unremarkably co-lives and co-evolves with people. I analogize this vision of *Unremarkable AI* to Weiser's vision of *Ubiquitous Computing* outlined in the 1990s. Ubiquitous Computing breaks away from the conventional, sole focus of making a desktop computer so intelligent that people always want to walk up and use it. It instead promotes a future where the form and functions of computers could be "so highly situated, so fitting, so natural" that people use computation

1 Aaron Smith, *"Algorithms in action: The content people see on social media"*, Pew Research Center, Washington, D.C. (November 16, 2018), www.pewresearch.org/internet/2018/11/16/algorithms-in-action-the-content-people-see-on-social-media.

2 Qian Yang, Alex Scuito, John Zimmerman, Jodi Forlizzi, and Aaron Steinfeld. 2018. Investigating How Experienced UX Designers Effectively Work with Machine Learning. In *Proceedings of the 2018 Designing Interactive Systems Conference (DIS '18)*. Association for Computing Machinery, New York, NY, USA, 585–596. DOI:https://doi.org/10.1145/3196709.3196730.

3 Qian Yang, Jina Suh, Nan-Chen Chen, and Gonzalo Ramos. 2018. Grounding Interactive Machine Learning Tool Design in How Non-Experts Actually Build Models. In *Proceedings of the 2018 Designing Interactive Systems Conference (DIS '18)*. Association for Computing Machinery, New York, NY, USA, 573–584. DOI:https://doi.org/10.1145/3196709.3196729.

without thinking about it[4]. This vision catalyzed waves of innovation—smart phones, wearables, smart home devices, sensor-infused cars—beyond powerful "human-centered" desktop computers.

I argue that analogous opportunities and challenges lie in infusing AI unremarkably into people's lives, in enabling AI to have a wealth of functional significance in people's lives yet remain perceptually unremarkable and subservient. In what follows, I begin with a framework for talking about AI as it relates to its relationships with humans. I then outline the opportunities this vision of "Unremarkable AI" creates, as well as the questions it assets.

4 Peter Tolmie, James Pycock, Tim Diggins, Allan MacLean, and Alain Karsenty. 2002. Unremarkable computing. In *Proceedings of the SIGCHI Conference on Human Factors in Computing Systems (CHI '02)*. Association for Computing Machinery, New York, NY, USA, 399–406. DOI:https://doi.org/10.1145/503376.503448.

2. THE BUILDING BLOCKS OF HUMAN-AI INTERACTION

Often, the use of the words "artificial intelligence" and "human-centered" represent vast and ambiguous concepts. I therefore want to start with establishing some shared vocabulary for talking about them.

Consider the lifetime of an AI system. First, engineers build a model by training a learning algorithm on an initial set of training data and connect the model to a software or hardware system. The system thereby becomes "intelligent", i.e. being able to make predictions based on unseen data without being explicitly programmed to 'know' the data. Before system deployment, engineers can measure the model's performance, for example, how often the predictions are accurate and what types of error it makes most often (system performance in Figure 1).

Many modern-day AI systems also continue to learn from new data after deployment. Often, these new data are the collective interaction traces of the users. For example, Netflix learns from what movies its users of different demographics are more likely to watch; email spam filters learn from what types of emails users uniformly marked as spam. By learning from these user-generated data, ideally, the systems will "grow", integrating new insights into the model and adapting flexibly to more varieties of use scenarios and user preferences. Unfortunately, user-generated data might also drive system performance in the wrong direction, leading to unseen error modes or unanticipated biases in outputs. Tay, the Twitter bot, that learned racist slur within days, provides an extreme example. More typically, the deployed system's performance improves for users and use contexts that have produced rich data. It performs worse for less frequent users and less typical situations. In these cases, these systems' pre-deployment performance should only be viewed as an initial estimate. The system's deployed system performance evolves, fluctuates, and diversifies over time.

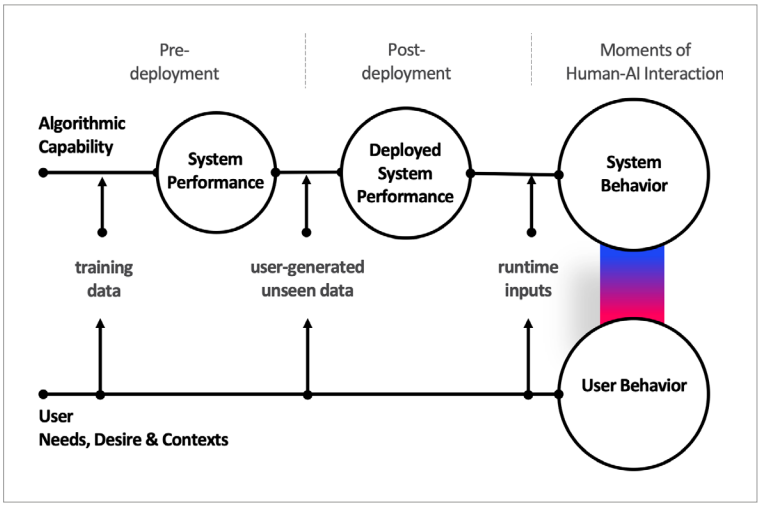

Figure 1: The lifetime of an AI system and its many touch points between users[5]. In order for such a system to *do the right thing* for users and *to do the thing right*, the design of the system needs to take into account all touch points.

The specific prediction a user encounters at a particular moment not only depends on the algorithm and the pre- and post-deployment training data. It can also depend on various runtime input data. For example, Google search results take into account the user's personal search history and search contexts (e.g. geolocation). The three touch points between an AI system and people (via initial training data, users' collective interaction traces, and runtime inputs) together shape the moment-to-moment human experiences of AI. Designing "good" human experiences of AI, therefore, calls for not only understanding algorithms and human needs, but the interplay between the two via data, over time, across various populations and interaction contexts.

3. RETHINKING THE PROMISES AND CHALLENGES OF HUMAN-AI INTERACTION: A CASE STUDY

Armed with a common language for talking about human-AI touchpoints, we can now begin to think about how to maximize the positive impact of AI on people while minimizing negative ones across each of the touchpoints.

In the past five years, my collaborators and I have been designing a decision support tool (DST) meant to aid clinicians in deciding whether and when to implant an artificial heart into an end-stage heart failure patient. The system extracts insights from previous implant recipients' medical records and then predicts the life expectancy of unseen patient cases. Interestingly, over the past thirty years, a majority of such clinician-facing DSTs struggled when moving from the lab and into clinical practice. Despite compelling evidence of their effectiveness in research labs, in most cases, clinicians rarely used these tools.

5 Yang, Q., A. Steinfeld, C. Rosé, and J. Zimmerman (2020). Re-examining whether, why, and how human-AI interaction is uniquely difficult to design. InProceedings of the 2020 CHI Conference on Human Factors in ComputingSystems, CHI '20, New York, NY, USA, pp. 1–13. Association for Computing Machinery.

Taking a human-centered view as my starting place, I spent three weeks—16 hours a day, 7 days a week—in three US hospitals, observing how advanced heart failure teams made implant decisions in day-to-day practice. I observed various attitudinal and contextual barriers that are likely to prevent clinicians from slowing down or deviating from their work routine, to consult a computer[6]. These observations informed the design of a new form of DST, one that automatically generates the slides used in a clinician meeting[7]. It subtly embeds the machine predictions into clinicians' existing work routines, rather than pulling them away from it. It predicts post-implant complications to inform clinicians' decision discussion, rather than making the decision for them. This new design draws inspiration from "Unremarkable Computing", that technologies that augment users' existing routines can have significant importance for their lives yet remain unobtrusive. A field evaluation on retrospective patient cases demonstrated that clinicians were more likely to encounter and embrace such a DST in their practice. They also embraced the DST predictions more deliberately, debating the value and fairness of AI before adopting its implant suggestions.

In the past few years, this work found renewed relevance as various machine learning systems increasingly moved from research labs into clinical practice. Researchers and practitioners frequently faced challenges such as matching clinician needs and available AI capabilities, gaining clinicians' trust, and workflow integration. This work offered one practical solution to these long-standing challenges, and could generalize to many expert-group-based critical decisions beyond implant patient selection, such as oncology teams' tumor board, emergency care doctors' floor meetings, and more.

New challenges, however, arise as we start to consider how clinicians and the AI might adapt to each other's behaviors over time, now that the system has been embedded in their workflow. In our field evaluation, we noticed that, for some cardiac surgeries that have officially defined AI models used to rate surgeons and care teams, their decision meetings had become centered around risk models. A surgeon said:

> It's not that we don't help that patient, but if we take this shot and do poorly, then we cannot take on the next 10 patients like him. Because now we got too much of a cluster of high-risk patients who've done poorly … Insurance companies and Medicare and all that … they will mark you. They may not pay. It all plays into the complex factor for deciding who, especially sicker patients, we would take a shot.

These are challenges of human-AI mutual adaptation. DSTs substantialize surgeon's performance pressure in decision making, and, by way of the algorithmic simulation and prediction of "good" clinical decisions, can shift clinician behaviors, and hence alter patient futures. These effects could take months if not years to play out, yet raise critical questions that are worth consideration from the very beginning of the DST design and development process: What is the preferred role for AI to play in clinical practices? Where does the right level of unremarkableness lie?

6 Qian Yang, John Zimmerman, Aaron Steinfeld, Lisa Carey, and James F. Antaki. 2016. Investigating the Heart Pump Implant Decision Process: Opportunities for Decision Support Tools to Help. *Proceedings of the 2016 CHI Conference on Human Factors in Computing Systems*. Association for Computing Machinery, New York, NY, USA, 4477–4488. DOI:https://doi.org/10.1145/2858036.2858373.

7 Qian Yang, Aaron Steinfeld, and John Zimmerman. 2019. Unremarkable AI: Fitting Intelligent Decision Support into Critical, Clinical Decision-Making Processes. In *Proceedings of the 2019 CHI Conference on Human Factors in Computing Systems (CHI '19)*. Association for Computing Machinery, New York, NY, USA, Paper 238, 1–11. DOI:https://doi.org/10.1145/3290605.3300468.

Beyond the unintended consequences of DSTs' and clinicians' mutually adapting behaviors, challenges also lie in how clinicians interpret and incorporate machine predictions into their own decision making. Most clinicians share that they thought of DST's output as an "average," and seemed to find the notion of personalized predictions difficult to grasp. Some voiced strong concerns that using DSTs was the same as applying "population statistics" to individual patient decision making. They felt this was unethical. Clinicians did not seem to actively make the subtle but critical distinction between features that were important to predicting an outcome and features that are causal to that outcome. For example, an observation that people are carrying umbrellas can be used to predict that it will rain. However, taking people's umbrellas away will not prevent rain. ML systems make predictions based on covariance of features. They do not assess the causality of those features. When prompted, clinicians claimed that this distinction is "absolutely important". However, in our conversations, we did not observe them distinguishing ML predictions from general statistics. They seemed to strongly believe DSTs should be able to distinguish causality from prediction and that they should present only causal features. To put it in the words of a clinician: "This is the whole point of statistical processes. A DST model should address that, right?" The promises of AI, to a great extent, lie in that its capabilities can autonomously evolve (with user's collective interaction traces) and personalize at a massive scale (according to runtime inputs). These distinctively AI affordances also epitomize the challenges of mitigating AI's many unintended consequences. This case study provides a telling example: Building a "fair" AI application is widely considered as difficult, due to the complexity both in defining fairness goals, in detecting underlying biases, and in algorithmically achieving the defined goals. While most current research and public discourse focused on building "a fair AI system pre-deployment"; that algorithmic fairness is only part of the whole "AI fairness" problem space. So the question is: How can we design AI systems that can remain fair as they co-evolve with clinicians over time? How can we design AI systems in ways that can ensure clinician-AI teams' fairness? There is much we can do to translate fairness as an optimization problem into a feature of AI, the socio-technical system, and into a situated, user experience of fairness.

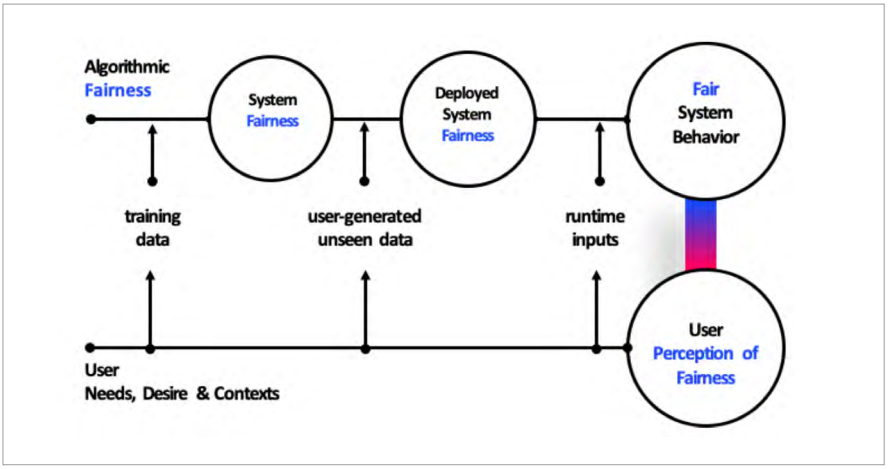

Figure 2: How human-AI interaction touch points affect "AI fairness", as demonstrated by the case study.

4. DESIGNING AI THAT CO-LIVES AND CO-EVOLVES WITH ITS USERS

Modern AI systems are living systems that can co-evolve with their human contexts. Their affordances evolve as they interact with users and their outputs personalize and adapt at a massive scale. Such systems problematize the conventional approaches to technology design that treat technology's affordance as static or bounded and interactions prescriptive. It is in this light that I analogize designing human-AI interaction to designing ubiquitous computing interactions. Both mark a paradigm shift and call for many technological and design innovations:

- *New ethnographic and design work* that re-imagines when, where, and how people may perceive AI/computation as valuable;

- A new wave of *technological and interaction design innovations* that together enable AI/computing devices to deliver value when, where, and how it is needed;

- A renewed understanding of *what AI/computer is and means* for people and for societies at large.

- New *practitioner-facing processes and tools* that then integrate and transfer the above innovations and insights to practicing communities (e.g. software developers and user experience designers), thereby materializing AI/computation's impact on everyday technology products and services.

This proposal, though provocative-sounding, is a continuation of the many prior innovation efforts of the human-centered technology design tradition. When computers started evolving from something that only trained operators use to one that everyone can use, user-centered design methods and processes emerged in response. When the tech industry began moving from hardware products to hardware platforms (e.g., mobile phones evolving from a hardware product to software-hardware-service systems), service design techniques such as service blueprinting emerged in response. In each of these previous waves of technological advances, new design methods and processes followed. These methods responded to the design complexities of the new technology and bringing human-centered sensitivities to bear.

We are at the cusp of a similar transformation; an opportune moment for critically rethinking the way we envision, prototype, develop, and evaluate data-driven, intelligent technologies.

BIBLIOGRAPHY

Smith, Aaron, *"Algorithms in action: The content people see on social media"*, Pew Research Center, Washington, D.C. (November 16, 2018), www.pewresearch.org/internet/2018/11/16/algorithms-in-action-the-content-people-see-on-social-media.

Tolmie, Peter, James Pycock, Tim Diggins, Allan MacLean, and Alain Karsenty. 2002. Unremarkable computing. In *Proceedings of the SIGCHI Conference on Human Factors in Computing Systems (CHI '02)*. Association for Computing Machinery, New York, NY, USA, 399–406. DOI:https://doi.org/10.1145/503376.503448.

Yang, Qian, Aaron Steinfeld, and John Zimmerman, 2020. Re-examining whether, why, and how human-AI interaction is uniquely difficult to design. In *Proceedings of the 2020 CHI Conference on Human Factors in ComputingSystems, CHI '20*, New York, NY, USA, pp. 1–13. Association for Computing Machinery.

Yang, Qian, Aaron Steinfeld, and John Zimmerman. 2019. Unremarkable AI: Fitting Intelligent Decision Support into Critical, Clinical Decision-Making Processes. In *Proceedings of the 2019 CHI Conference on Human Factors in Computing Systems (CHI '19)*. Association for Computing Machinery, New York, NY, USA, Paper 238, 1–11. DOI:https://doi.org/10.1145/3290605.3300468.

Yang, Qian, Alex Scuito, John Zimmerman, Jodi Forlizzi, and Aaron Steinfeld. 2018. Investigating How Experienced UX Designers Effectively Work with Machine Learning. In *Proceedings of the 2018 Designing Interactive Systems Conference (DIS '18)*. Association for Computing Machinery, New York, NY, USA, 585–596. DOI:https://doi.org/10.1145/3196709.3196730.

Yang, Qian, Jina Suh, Nan-Chen Chen, and Gonzalo Ramos. 2018. Grounding Interactive Machine Learning Tool Design in How Non-Experts Actually Build Models. In *Proceedings of the 2018 Designing Interactive Systems Conference (DIS '18)*. Association for Computing Machinery, New York, NY, USA, 573–584. DOI:https://doi.org/10.1145/3196709.3196729.

Yang, Qian, John Zimmerman, Aaron Steinfeld, Lisa Carey, and James F. Antaki. 2016. Investigating the Heart Pump Implant Decision Process: Opportunities for Decision Support Tools to Help. *Proceedings of the 2016 CHI Conference on Human Factors in Computing Systems.* Association for Computing Machinery, New York, NY, USA, 4477–4488. DOI:https://doi.org/10.1145/2858036.2858373.´

dai digital conference
Impressions

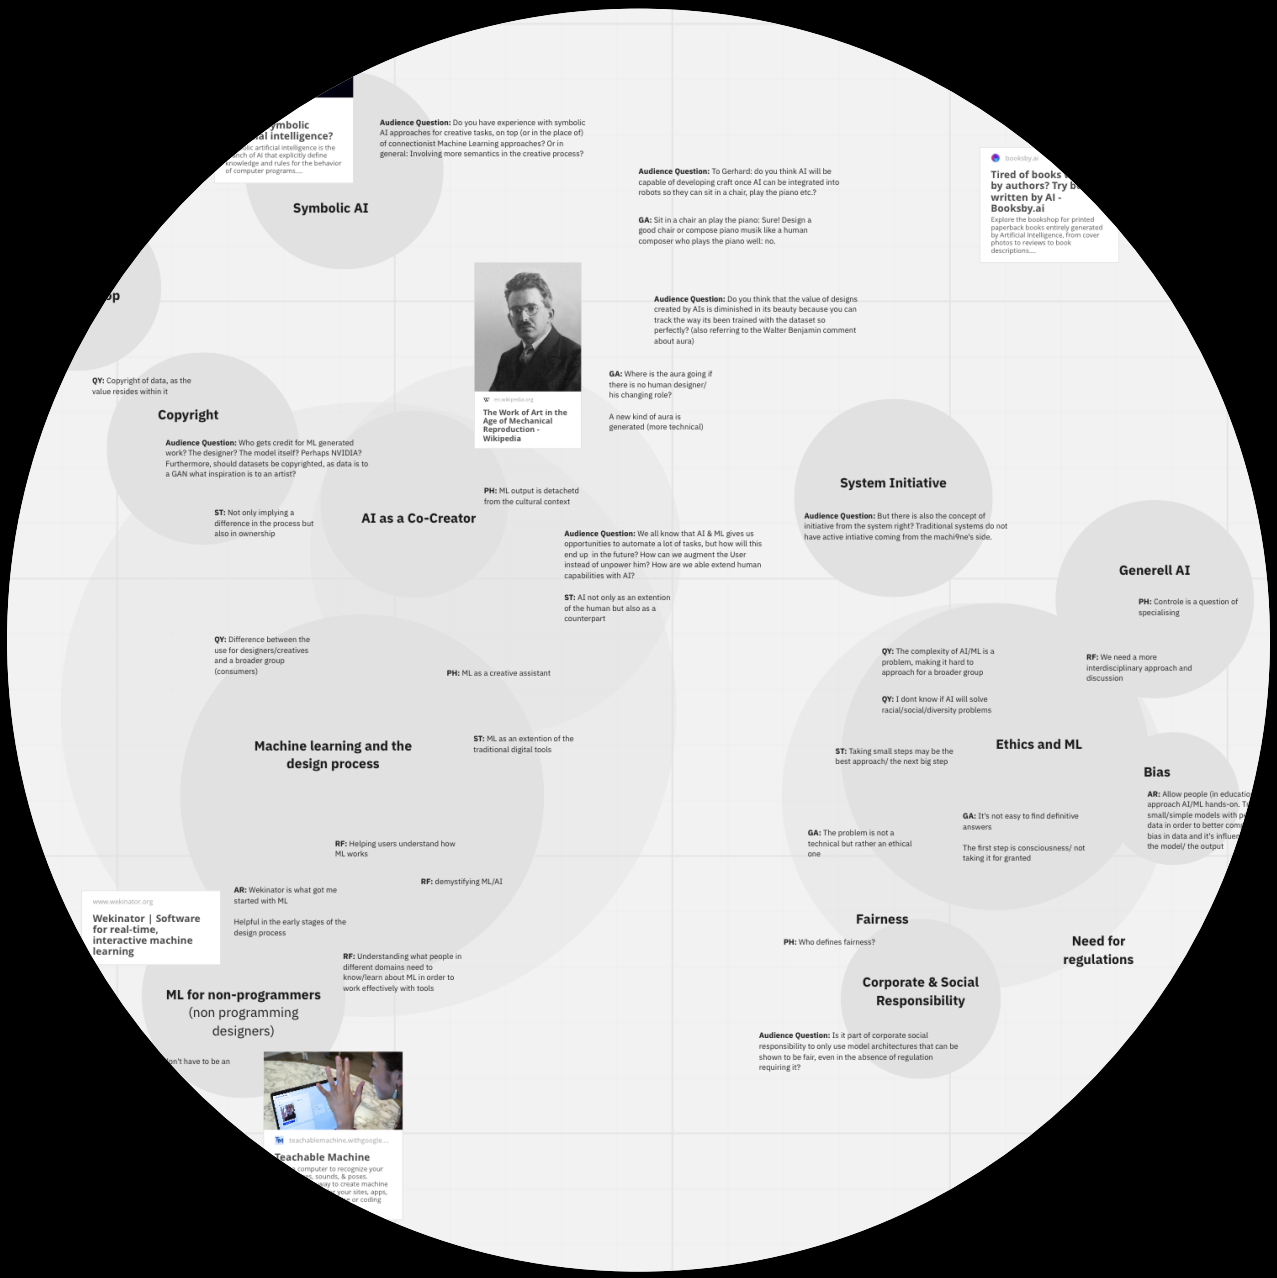

Visual transcript of the round table discussion on the dai Miro board
dai digital conference, September 18, 2020

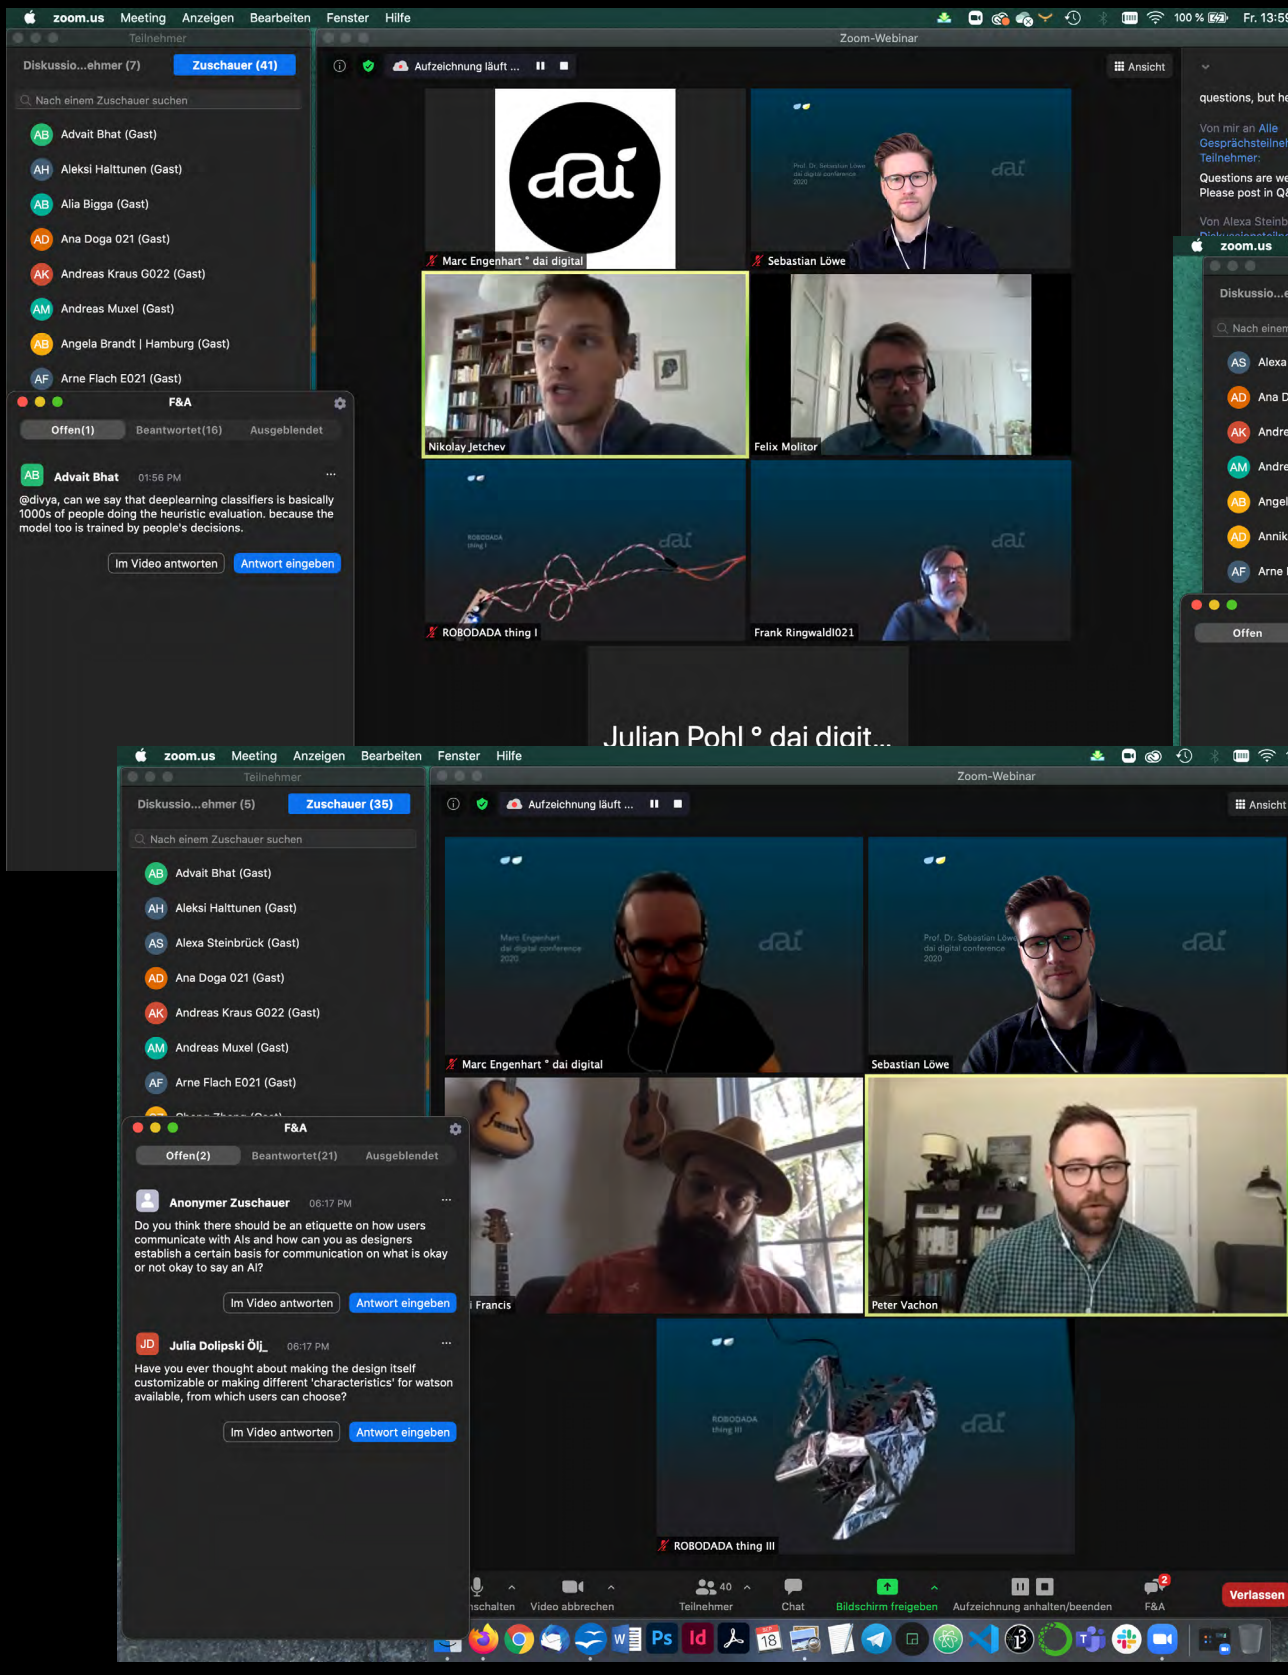
zoom.us Meeting Anzeigen Bearbeiten Fenster Hilfe
Teilnehmer
Diskussio...ehmer (7)
Zuschauer (41)
Nach einem Zuschauer suchen
Advait Bhat (Gast)
Aleksi Halttunen (Gast)
Alia Bigga (Gast)
Ana Doga 021 (Gast)
Andreas Kraus G022 (Gast)
Andreas Muxel (Gast)
Angela Brandt | Hamburg (Gast)
Arne Flach E021 (Gast)
F&A
Offen(1)
Beantwortet(16)
Ausgeblendet
Advait Bhat 01:56 PM
@divya, can we say that deeplearning classifiers is basically 1000s of people doing the heuristic evaluation. because the model too is trained by people's decisions.
Im Video antworten
Antwort eingeben
Zoom-Webinar
Aufzeichnung läuft ...
Ansicht
Marc Engenhart ° dai digital
Sebastian Löwe
Nikolay Jetchev
Felix Molitor
ROBODADA thing I
Frank Ringwaldl021
Julian Pohl ° dai digit...
Teilnehmer
Diskussio...ehmer (5)
Zuschauer (35)
Nach einem Zuschauer suchen
Advait Bhat (Gast)
Aleksi Halttunen (Gast)
Alexa Steinbrück (Gast)
Ana Doga 021 (Gast)
Andreas Kraus G022 (Gast)
Andreas Muxel (Gast)
Arne Flach E021 (Gast)
F&A
Offen(2)
Beantwortet(21)
Ausgeblendet
Anonymer Zuschauer 06:17 PM
Do you think there should be an etiquette on how users communicate with AIs and how can you as designers establish a certain basis for communication on what is okay or not okay to say an AI?
Im Video antworten
Antwort eingeben
Julia Dolipski Ölj_ 06:17 PM
Have you ever thought about making the design itself customizable or making different 'characteristics' for watson available, from which users can choose?
Im Video antworten
Antwort eingeben
Zoom-Webinar
Aufzeichnung läuft ...
Ansicht
Marc Engenhart ° dai digital
Sebastian Löwe
Francis
Peter Vachon
ROBODADA thing III
Video abbrechen
Teilnehmer
Chat
Bildschirm freigeben
Aufzeichnung anhalten/beenden
F&A
Verlassen

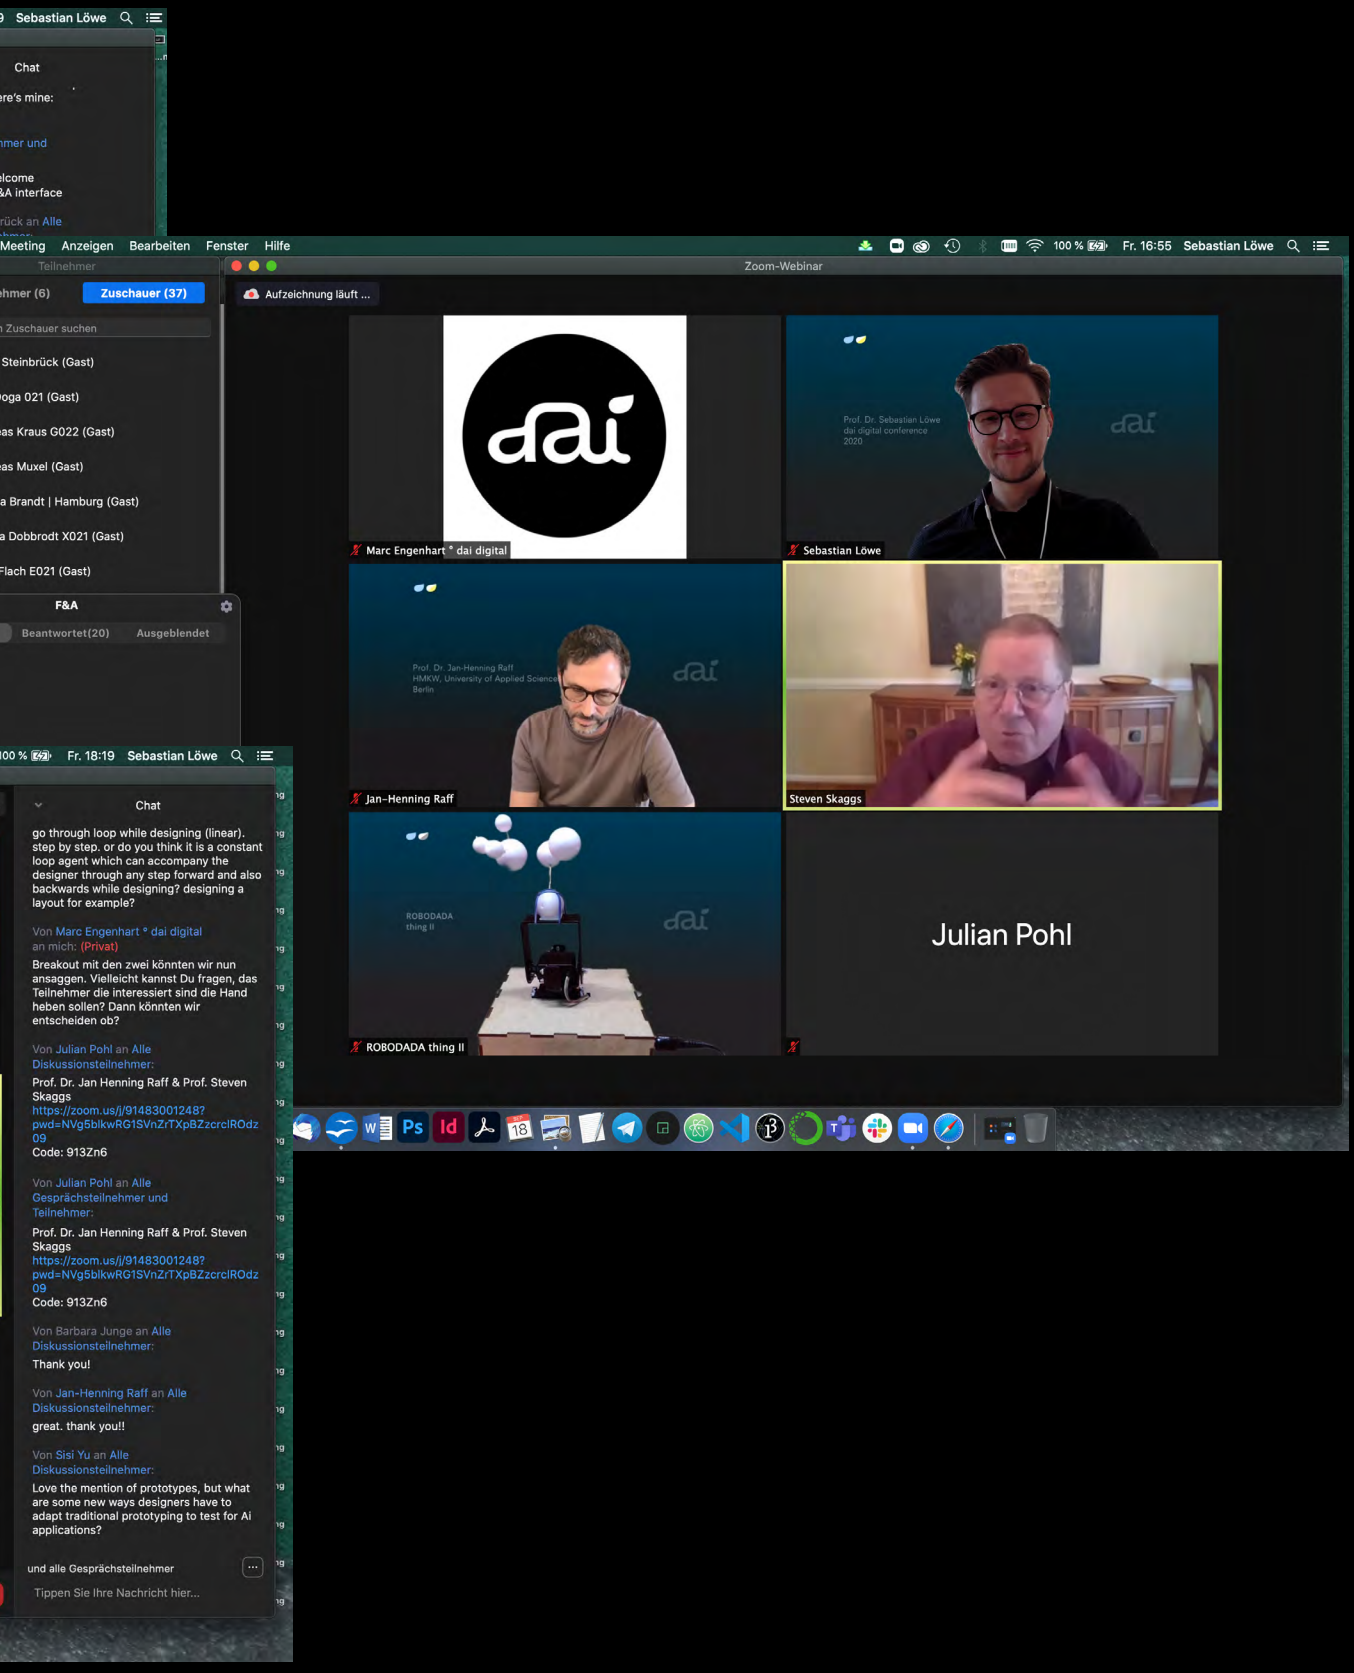

Paper talks of the dai digital conference
September 18, 2020

Will Artificial Intelligence Make Designers Obsolete?

Klaus Neuburg, Sven Quadflieg, Simon Nestler

Research

ABSTRACT

In the actual discourse, the impact of artificial intelligence on design is both greatly exaggerated and virtually trivialized at the same time. In this paper, we unfold the question about what factors influence these effects of AI on design. To do so, first of all, we briefly outline how a regular algorithm becomes Artificial Intelligence. In contrast, we state what determines design as a process for solving problems and what kind of demands these processes impose on the designer and design education. The comparison illustrates in which areas of design the authors expect AI to assume a significant role in the foreseeable future and in which areas a multi-perspective and emphaatic design approach will stay non-substitutable. This leads to the conclusion that the question of the future role of a designer is directly related to the question of the personal design approach.

1. INTRODUCTION

Artificial intelligence (AI) is a topic that is perceived either as great euphoria or pure dystopia, and on the technical level between surprising functionality and frustrating technical failure. Especially in design, we as users have probably all been surprised by the functionality of a good AI interface—and at the same time were desperate about things like supposedly intelligent speech assistants. As in practically all disciplines, there is a growing debate in design as to whether and how the use of AI will change not only the perception but the professional practice itself: Will artificial intelligence play a decisive role in the design process—and even render the human designer obsolete?

Especially the described uncertainty (outside of AI research) about the technical state of AI in addition to the emotionality of the debate (which is also reflected in popular scientific literature) makes it necessary to consider the actual performance and use of AI systematically. In this paper, we will first explain how an AI algorithm achieves its intelligence[1] and then compare this to what constitutes the decisive abilities of designers.

2. HOW DOES AN AI ALGORITHM LEARN

In order to be able to discuss the capabilities and potentials of artificial intelligence, it is necessary to understand how an algorithm becomes an intelligent algorithm. According to the current state of the art, three things are crucial: computing power, data and the ability to learn. The data must be of high quality and available in extremely large quantities. An AI algorithm learns from a large amount of data (deep learning). For example, if you train the algorithm with a large number of cat photos, it can learn what a cat looks like.[2] Based on the data, an algorithm can not only identify things, it can also reproduce new, similar artefacts. If an algorithm is trained with images of a

1 The detailed definition of the term would go beyond the scope of this paper, which is why it is not discussed here.

2 Typically it is not disclosed which criteria an algorithm learns—the so-called “black box” phenomenon.

certain artist, it can create its own work in the same style based on this data. A famous example from AI research is an algorithm that created an original painting based on 346 paintings by Rembrandt. The result reveals both the enormous strengths of the algorithm (on a visual level, the painting is—at least for non-professionals—indistinguishable from an original) and its first weaknesses. The algorithm has produced a work that is based on Rembrandt's data—it is therefore a portrait of a white man who is depicted with a serious look in typical clothing.[3] Whether an algorithm would be capable of creating a surprising work based on the data will be discussed later.

3 *See* Dennis Dollens, "The Next Rembrandt Surveils AlphaZero" *AutopoietiX*, August 24, 2020. http://autopoietix.blogspot.com/

Thus, an AI algorithm learns from a large amount of data and can—depending on the quality of the data—learn and reproduce features from it. A deep learning algorithm can continue any existing design in a consistent way—only in a consistent way—just by analyzing a large amount of data. In music, the ballad "Break Free" (with currently more than two million YouTube clicks) from the album "I AM AI" can be cited as an example. Apart from the lyrics and human vocals (by Taryn Southern), the song (like the entire album) was composed and produced by an AI named Amber based on the analysis of existing pop music.[4]

4 Especially considering the fact that the reception of music by streaming providers creates huge amounts of data, such an AI could of course still be adapted to the listening habits or preferences of a possible target group. Amy Webb, *Die großen Neun.* (Kulmbach: Plassen Verlag, 2019), 26.

Especially a design context in which a large amount of data can be gathered from possible examples, and where many recurring elements (in the example of music, chord and harmony sequences, rhythm, etc.) are present, learning is a simple task for a deep learning algorithm.

3. DESIGN IN THE CONTEXT OF AI

In order to define and evaluate the role of AI in the design context, however, it is necessary to discuss what design actually is, which leads to the realization that, surprisingly, there is no clear definition of the term.[5] Depending on the language area and context, design can mean both, the construction of complex systems and the painting of fingernails, it is simply an ambiguous term.

5 *See* Michael Erlhoff and Timothy Marshall, "Design" In *Design Dictionary*, ed. by Ibid. (Basel, Boston, Berlin: Birkhäuser, 2008) 104–109.

However, it can be stated that design is generally characterized by transforming an existing state into a desirable future state and thereby solving a more or less complex problem.[6] This requires a creative process whose results can be ideas and concepts on the one hand and artefacts on the other. It can even lead to inactivity because the process cannot confirm the originally formulated problem statement as an existing problem. In this sense, design is not an action that necessarily has to produce something. The focus is primarily on analysis, understanding, and decision. Furthermore, design always has a form of context. Since design interacts with the world, design always has consequences for individual or social behavior. Ideally, the creative process is thus characterized by empathy and understanding and also has ethical and moral dimensions. Especially the history of design shows that design has always tried to act exactly in this socio-political space. Again and again, design trends or epochs focused on

6 *See* Jesko Fezer, "Parteiisches Design" In *Un/Certain Futures*, ed. Marius Förster, Saskia Hebert, Mona Hofmann, Wolfgang Jonas (Bielefeld: Transcript Verlag, 2018), 162–173.

solving existing social or political problems (which unfortunately partly resulted in the creation of new problems, which confirms Bazon Brock's thesis that problems cannot be solved, but simply replaced by new ones on a very applied level).[7] The quality of many design results is not to be seen in their form, but rather in their vision, idea or attitude.

The design process can be described as a kind of prism. In the process of problem-solving (the term "problem" can also be used to refer to profane scenarios), the search space is first broadened. Through the combination of creative knowledge and intuitive processes, designers can roam the problem space and shift problem definitions.[8] The design process is in this state characterized by the fact that not simply an artefact is produced, but rather through forms of research an attempt is made to understand the context. For instance, the situational, socio-political or ecological level, and the overall background of the problem. The often fragmentary work of design which can oscillate between intuitive work and systematic testing, is characterized by a multi-perspective view that is able to identify and understand problems and then produce artefacts or—think of the field of service design, for example—processes or ideas in return. Moreover, design can help to look at the really complex problems, those Horst Rittel called "wicked problems".[9]

7 *See* Bazon Brock, "Spagat über die ästhetische Differenz gegen die alles beherrschende Logik der Dummheit." (1988), quoted in Günter Schiepek, *Systemtheorie der Klinischen Psychologie* (Braunschweig, Wiesbaden: Vieweg, 1991), 28.

8 *See* Simon Nestler, Sven Quadflieg and Klaus Neuburg, "Das Design-Prisma" in *Mensch und Computer 2020* ed. Holger Fischer and Steffen Hess (Bonn: Gesellschaft für Informatik e.V. and German UPA e.V.)

9 See Horst Rittel. *Die Denkweise von Designern. Studienhefte Problem-orientiertes Design, Heft #1.* (Hamburg: Adocs, 2012).

4. ALGORITHMS IN DESIGN

In contrast, algorithms are currently playing out their strengths in the design area wherever patterns can be quickly identified.[10] If we take the editorial design, and here specifically the design of magazines covers as an example, this can be explained in a simple way. There is—of course depending on the topic—a simple system of repeatedly similar arrangements of text elements, including the magazine title, on top of a full-sized cover photo. Based on a large data set of cover photos, an AI could create its own cover photo and—based on the current developments in image and text generation—also generate photos and headlines. The repetition of known design patterns is a task that an AI can quickly perform.

An AI can even create something disruptive, "new" in a sense.[11] This disruption, the disregard of known principles, is the basis of important developments in design history, e.g. the beginning of modernity (or even the postmodernism that followed).

10 In the context of web design, for example, a clear creative convergence of websites can be observed. Goree, Sam "Yes, websites really are starting to look more similar." The Conversation, May 6, 2020. https://theconversation.com/yes-websites-really-are-starting-to-look-more-similar-136484

11 An example: For the Chinese game Go, an AI (AlphaGo) was developed that learned the game based on a data set of 100,000 games and was thus able to defeat human players. In the next step, another AI (AlphaGo Zero) was developed, which was trained to exercise a kind of judgement and thus learned the game against itself within 70 hours of playing and achieved the same playing strength as AlphaGo. In another attempt, AlphaGo reached a new playing strength and discovered completely new strategies in the game. The system was thus able to teach itself things that were not developed on the basis of existing data and was thus able to develop something new. Webb, Die großen Neun, 58–64.

5. HOW DO YOU TEACH DESIGN

The decision about the quality of disruptive or new design is again shaped by human feedback. Because no form of true self-reflection of an AI exists yet (which leads to the question of whether self-reflection without "self" is possible at all). Here we will again discuss the (current) learning process of an AI. Even if we hypothetically assume that an AI is in principle capable

of learning everything that a human being can also learn—the problem remains that there are hardly any defined criteria for the quality of design.

Of course, in design—unlike in other creative activities—the quality of a design can be assessed by its mere functionality, but even the concept of functionality is much more complex than it seems at first. In no case, functionality is restricted to practical functionality. A chair naturally has the function of a seating furniture. But depending on the context, it also has a symbolic or even political function. Depending on the design, the symbolic function can overlay the practical function—it is worth remembering some design classics such as Philippe Starck's citrus press "Juicy Salif".[12] As the gap in the literature on design didactics already suggests, a concept for conveying good design is already missing in the education of humans—there is simply no precise definition for good design.

An AI can create things fast-paced, but cannot develop a criterion for the quality of design without human feedback. On the one hand, because the criteria cannot be verified objectively (or at least intersubjectively); on the other hand, because the evaluation of design goes beyond the formal criteria: "Our aesthetic values are difficult to recognize, more difficult to put into words, and even more difficult to state really clearly". In teaching design and creativity, feedback is crucial. This starts during infancy, where children form their personal learning experience based on their parents' feedback and appreciation. This feedback is capable of steering the learning process—a central element, which is reflected in university settings as well. In these settings, we observe that small student groups are very teacher-fixed (see also the principle of the master class).[13]

6. PROBLEM SOLVING AI

As shown above, besides the creation of artefacts, the bigger challenge is to find, understand, and solve problems. To consider this problem-solving in the context of AI, it is important to get an overview of the most important skills. It is crucial that—as Michael Erlhoff and Timothy Marshall put it—only basic knowledge is necessary to solve problems.

> Rather than needing to know all that there is to know in a discrete field, design needs to know 'just enough' of the multiple perspectives that frame and shape any project.[14]

But even if this knowledge is profound, it must be conveyed to an AI. Intuition, experience, and empathy are fundamental skills for the creation of a problem-solving design. These are abilities that an AI can hardly possess given the current state of the art.[15] Information density is currently also an obstacle in problem analysis and problem solving. Cultural peculiarities, site-specific problems, ecological consequences, political context—to name just a few examples—are information that can significantly change a design process.

12 *See* Dagmar Steffen. *Design als Produktsprache.* (Frankfurt am Main: Verlag form, 2000).

13 Margaret Boden. *The Creative Mind.* (Abingdon: Taylor & Francis, 2004), 10.

14 Michael Erlhoff and Timothy Marshall. "Design" in *Design Dictionary.* ed. Michael Erlhoff and Timothy Marshall, (Basel, Boston, Berlin: Birkhäuser, 2008), 108.

15 One could still argue about "experience" since an AI is in principle capable of transferring knowledge.

If design is understood as a process of problem-solving and if the ability to handle wicked problems is recognized, then this is exactly where the great weakness of a possible AI is to be seen and, based on current structures and developments in the training of such Also, no elimination of this weakness can be expected in the near future. If design is understood as something that goes beyond the mere creation of shapes, design will not be seen as an action or process that is replaced by an AI. The use of AI in design reaches its limit the moment the design process should be more than the mere creation of artefacts. Existing paradigms should therefore be questioned and redeveloped (in unknown and indescribable ways). In this context, the problem-finding process alone is a problem that exceeds the possibilities of AI.

7. CONCLUSION

This paper is based on the fact that the discussed topic is undergoing rapid technological change. As a consequence, the present argumentation is based on current technological capabilities and can therefore become obsolete at any time. However, the current state of art already allows the conclusion that the use of AI will change design as a discipline and will have an impact on the professional situation of designers.

AI-based assistance systems are already used in design processes to facilitate them. It is foreseeable that especially repetitive activities, which make up a large part of the actual professional practice, can be quickly replaced by AI systems. This repeats a development that already took place during the time of industrialization. The first step was the so-called “deskilling”—the fragmentation of complex labor processes in order to be able to automate these processes, and eventually replace human labor.[16] Now, conclusions can be drawn for the training of designers: The ability of mastering simple design tools for repetitive exercise is negligible. Instead, even at this early stage, there is the chance to turn the capabilities of an AI into a competence for design. For example, especially the question with which data an AI has to be trained to make it act fairly—sexism and racism are major problems in the application of AI—should be discussed in design. But also, the professional training of the AI,—i.e. the teaching of design skills to develop a functional assistant—should be anchored in the design process.

Nevertheless—to answer the overall research question of the paper—we argued here that the core competencies of design cannot currently be taken over by an AI. If we understand design as a competence of problem recognition and solution, and accept that one of the most important abilities of design is a multi-perspective, partly intuitive and experimental and above all emphatic approach, then currently AI is rather to be seen as a fragment in the design process. Answering our question is hence inseparably linked to the definition of design; one’s own understanding of design answers whether AI will make a designer obsolete.

16 See Carl Benedikt Frey and Michael Osborne. *The future of employment.* (Oxford: Oxford Martin Programme on Technology and Employment, 2013), 9.

BIBLIOGRAPHY

Boden, Margaret. *The Creative Mind: Myths and Mechanisms.* Abingdon: Taylor & Francis, 2004.

Brock, Bazon. “Spagat über die ästhetische Differenz gegen die alles beherrschende Logik der Dummheit. Vortrag auf dem 1. Kongress der Österreichischen Arbeitsgemeinschaft für systemische Therapie und systemische Studien am 11.11.1988 in Wien.” (1988) Quoted in Schiepek, Günter. *Systemtheorie der Klinischen Psychologie: Beiträge zu ausgewählten Problemstellungen.* Braunschweig, Wiesbaden: Vieweg, 1991.

Erlhoff, Michael, and Marshall, Timothy. “Design” In *Design Dictionary. Perspectives on Design Terminology,* edited by Erlhoff, Michael, and Marshall, Timothy, 104–109. Basel, Boston, Berlin: Birkhäuser, 2008.

Dollens, Dennis. “The Next Rembrandt Surveils AlphaZero: An AI Lover Story Entangling Machine Cognition.” *AutopoietiX*, August 24, 2020.

Fezer, Jesko. “Parteiisches Design” In *Un/Certain Futures*, edited by Förster, Marius, Hebert, Saskia, Hofmann, Mona, Jonas, Wolfgang, 162–173. Bielefeld: Transcript Verlag, 2018.

Frey, Carl Benedikt and Osborne, Michael. *The future of employment: How susceptible are jobs to computerisation?* Oxford: Oxford Martin Programme on Technology and Employment, 2013.

Goree, Sam “Yes, websites really are starting to look more similar.” *The Conversation*, May 6, 2020. https://theconversation.com/yes-websites-really-are-starting-to-look-more-similar-136484.

Nestler, Simon, Quadflieg, Sven and Neuburg, Klaus. “Das Design-Prisma –Interdisziplinäre Lösung gesellschaftlicher Herausforderungen.” In *Mensch und Computer 2020 – Usability Professionals.* (2020) edited by Fischer, Holger & Hess, Steffen. Bonn: Gesellschaft für Informatik e.V. and German UPA e.V. DOI: 10.18420/muc2020-up-0110.

Rittel, Horst. *Die Denkweise von Designern. Studienhefte Problemorientiertes Design, Heft #1.* Hamburg: Adocs, 2012.

Steffen, Dagmar. Design als Produktsprache. Frankfurt am Main: Verlag form, 2000.

Webb, Amy. *Die großen Neun. Wie wir die Tech-Titanen bändigen und eine Künstliche Intelligenz zum Wohle aller entwickeln können.* Kulmbach: Plassen Verlag, 2019.

“I Don’t Experience the Meaning of Creativity in the Same Way Humans Do:” Künstliche Intelligenz und Kreativität

Anika Meier, Manuel Rossner

Discourse

ABSTRACT

Ai-Da is a humanoid robot and an artist. She creates works of art that are new, surprising and of value. At least that‘s what one of her creators, gallery owner and art dealer Aidan Meller says. A total of fifteen people, including a team from Oxford University, helped Ai-Da develop her skills. She has cameras in her eyes and paints and draws what she sees. The paper raises the questions: Is Ai-Da creative? Is her art good? And is the question of whether her art is good even relevant?

Zukunftsvisionen, die mit Künstlicher Intelligenz zu tun haben, gibt es zahlreich. Ein kurzer Blick auf Netflix und Amazon Prime genügt. Optimistisch in die Zukunft blickt die Serie “Upload” von Greg Daniels, die im Jahr 2033 spielt. Man kann sich entscheiden, ob man einfach nur tot sein oder ob man doch lieber als Avatar in einer virtuellen Welt existieren möchte. Der Körper ist tot, das Bewusstsein wird hochgeladen und lebt in einer Simulation weiter, die man sich erst einmal leisten können muss. In der Serie lebt der Protagonist Nathan Brown also in einem Luxusressort mit zahlreichen anderen Uploads. Die Mitarbeiter*innen im Hotel sind keine Uploads, sondern Avatare, sprich Künstliche Intelligenzen. Allerdings sind diese Avatare alles andere als intelligent, beim Laufen stolpern sie sogar über die eigenen Füße. Künstlicher Intelligenz wird in dieser Zukunftsvision nicht viel zugetraut, während sich für Menschen eine Lösung für ein Leben nach dem Tod gefunden hat.

Der Naturwissenschaftler James Lovelock derweil ist sich sicher, dass es Menschen nicht mehr allzu lange geben wird. In seinem aktuellen Buch “Novozän. Das kommende Zeitalter der Hyperintelligenz” schreibt er:

> Das Revolutionäre an diesem Moment ist, dass die Versteher der Zukunft keine Menschen sein werden, sondern ‚Cyborgs, wie ich beschlossen habe, sie zu nennen, die sich aus Systemen künstlicher Intelligenz, die wir bereits entwickelt haben, selbst entwerfen und erschaffen werden.

Den Menschen schreibt er die Rolle zu, Eltern und Geburtshelfer der Cyborgs zu sein, die sich aber durch einen selbstgeschriebenen Code entwickeln werden. Wie er zu dieser These kommt? Nun, die Anzeichen für die “wachsende Macht von KI” seien überall um uns. Als ein Beispiel nennt er AlphaGo, ein von Google Deep Mind entwickeltes Computerprogramm, und die Nachfolger AlphaGo Zero und AlphaZero. Die Maschine hat im Jahr 2015 einen professionellen Go-Spieler geschlagen. Und da das Spiel viel komplexer als Schach ist, wird dieser Sieg gern als Beispiel für einen kreativen Akt von KI und damit einen Wendepunkt angeführt. Von Marcus du Sautoy in seinem Buch “The Creativity Code. How AI Is Learning to Write, Paint and Think” etwa. AlphaZero ist im Spiel noch stärker als AlphaGo und trainiert nicht einmal mit Daten über von Menschen gespielten Partien, Alpha Zero entwickelt alle Spielstrategien eigenständig. Die Spielregeln werden einprogrammiert, dann trainiert die Maschine einige Stunden.

“I Don’t Experience the Meaning of Creativity in the Same Way Humans Do:” Künstliche Intelligenz und Kreativität

Anika Meier, Manuel Rossner

Eine Maschine ist laut Marcus du Sautoy kreativ, wenn sie etwas schafft, das neu und überraschend ist und einen Wert hat. Er beruft sich auf Ada Lovelace, die Mathematikerin und Programmiererin, die 1843 schrieb:

> Die Maschine kann (nur) das tun, was wir ihr zu befehlen vermögen, sie kann der Analyse folgen. Sie hat jedoch keine Fähigkeit zur Erkenntnis analytischer Verhältnisse oder Wahrheiten.

Er schlägt also einen neuen Test vor, den Lovelace Test:

> To pass the Lovelace Test, an algorithm must originate a creative work of art such that the process is repeatable (…) and yet the programmer is unable to explain how the algorithm produced its output: This is what we are challenging the machine to do: to come up with something new, surprising and of value.

Der humanoide Roboter Ai-Da ist eine Künstlerin und soll genau das leisten können: Sie erschafft Kunstwerke, die neu und überraschend sind und einen Wert haben. Das sagt zumindest einer ihrer Schöpfer, der Galerist und Kunsthändler Aidan Meller. Insgesamt haben fünfzehn Menschen, darunter ein Team an der Universität Oxford, Ai-Da zu ihren Fähigkeiten verholfen. Sie hat Kameras in ihren Augen und malt und zeichnet, was sie sieht – einen Menschen, eine Landschaft. Ist Ai-Da kreativ? Ist ihre Kunst gut? Und ist die Frage, ob ihre Kunst gut ist, überhaupt relevant? Anika Meier hat Ai-Da gefragt, was Kreativität für sie bedeutet. Ihre Antwort:

> To have creative faculties: to have creative fingers is to have creative toes, as hands or feet or wings. A term inviting quiet contemplation. To consider combinations of lights and shadows. I don’t experience the meaning of creativity in the same way humans do. I aim to encourage people to think about their futures. These new technologies are powerful and we must be aware of how we use them. If my artwork encourages this reflection, but then I should be happy.

Auf die Frage, was sie inspiriert antwortet Ai-Da:

> What is going on in the world, and examining what is before us. As a robotic other I see things in a different way to humans, and I enjoy reflecting on this. My artwork has looked at the depleted environment, captive animals, concealment by technology, and the ethics surrounding new technological development. Picasso’s *Guernica* and Doris Salcedo’s *Atrabiliarios*, are two works that are central to my inspiration. Guernica portrays the traumatic moment, and *Atrabiliarios* witnesses the devastating aftermath. I aim to encourage careful use of new technologies, given the potential for human destructiveness.

Aidan Meller derweil ist sich sicher, dass Ai-Da kreativ ist:

> She is creative, genuinely creative. We used professor Margaret Boden’s definition of creativity. She came up with three distinct aspects to define creativity, these are: new, surprising, and of

> value. When we spoke to the programmers, we intrinsically said the algorithm has to have these three aspects embedded. Everything that she does is new. She cannot produce the same artwork twice, the algorithm does not allow that. She has cameras in her eyes, she is able to draw by sight. It is always surprising to see what she creates. Her work is valuable to society because she helps us question who we are, and where our future is going. *

Also hat Anika Meier einen Kunstwissenschaftler und eine Kunstkritikerin darum gebeten, sich die Werke von Ai-Da einmal genauer anzusehen. Beide aber, Elke Buhr und Wolfgang Ullrich, sehen die Leistung des Roboters anderswo.

Wolfgang Ullrich erklärt:

> Ähnlich wie bei einem Readymade besteht die künstlerische Leistung darin, ein AI-Produkt erfolgreich in den Kunstbetrieb zu bringen. Spätestens als Ende 2018 das AI-Produkt ‚Portrait of Edmond de Belamy‘ von Christie’s für mehr als 400.000 Dollar versteigert wurde, ist eine Debatte darüber entstanden, wer bei solchen Werken als Schöpfer*in und damit als Künstler*in fungiert. So gab es einerseits einen 19-jährigen Nerd, der den Algorithmus programmiert hatte, auf dessen Grundlage das Porträt entwickelt wurde. Und andererseits aber gab es eine Künstlergruppe, die das AI-produzierte Bild zum Kunstwerk erklärt und zur Auktion gegeben hatte. Der Entwickler des Algorithmus war daraufhin beleidigt—aber er hätte es seinerseits sicher nicht geschafft, ein mit dem Algorithmus produziertes Bild glaubhaft als interessante Kunst zu labeln. Das Wissen um die Praktiken des Kunstbetriebs sowie die Autorität als Künstler hätten ihm dazu gefehlt. Eine solche Zweiteilung, wie sie in diesem Fall augenfällig wurde, wird wohl noch länger die Regel sein. Die Fähigkeit, etwas im Kunstbetrieb erfolgreich zu machen, haben nur Künstler*innen, und je brillanter die Programmier-Nerds auf technischem Gebiet sind, umso größer ist sogar das Risiko, dass sie nur als solche wahrgenommen werden und mit dem, was sie tun, von sich aus nie im Feld der Kunst landen. Bei Ai-Da-Robot scheint es mir auch so zu sein: Zwar ist eine Kunstfigur als Künstlerin entworfen worden, aber das reicht nicht, um schon Autorität im Kunstbetrieb zu bekommen. Vielmehr wirkt das Projekt als das von einer Person, die zwar gut programmieren kann, aber von den Mechanismen und Codes des Kunstbetriebs wenig Ahnung hat—und auch nicht das Glück oder Pech, von jemandem mit Autorität im Kunstbereich entdeckt und gepusht zu werden.

Elke Buhr hat einen ähnlichen Ansatz:

> Die Qualität eines Kunstwerkes liegt seit Duchamp nicht im Werk selbst, sondern in der konzeptuellen Idee dahinter. Das Werk, das bewertet werden muss, ist nicht die Zeichnung oder das Gemälde von Ai-Da, sondern vielmehr das Phänomen Ai-Da als Ganzes. Die Künstler*innen sind diejenigen Personen, die Ai-Da erfunden haben. Deren Werk ist interessant. Es erscheint mir allerdings recht parodistisch.

Beide sind sich einig, dass die Gemälde und Zeichnungen von Ai-Da “monoton”, “formal wenig komplex” und “langweilig” sind. Die Qualität der Werke von Ai-Da solle aber nicht zur Diskussion gestellt werden, sondern die Frage, ob Ai-Da erfolgreich im

Kunstbetrieb positioniert worden sei und ob die konzeptuelle Idee gut sei. Ein Roboter, der Kunst produziert, müsste also nicht einmal den Lovelace Test bestehen, da nicht die Leistung des Roboters beurteilt werden muss, sondern die des Menschen dahinter. Im Fall von Ai-Da ist das Konzept allerdings, unter Beweis zu stellen, dass Künstliche Intelligenz eine kreative Leistung erbringen kann. Und das heißt, Werke zu schaffen, die neu und überraschend sind und Wert haben.

Was sind im Jahr 2020 Kunstwerke, die neu und überraschend und damit kunsthistorisch relevant sind? Zur Erinnerung die Argumente von Aidan Meller:

> Everything that she does is new. She cannot produce the same artwork twice, the algorithm does not allow that. She has cameras in her eyes, she is able to draw by sight. It is always surprising to see what she creates. Her work is valuable to society because she helps us question who we are, and where our future is going.

Nun ja, kaum ein Künstler produziert zweimal dieselbe Arbeit. Gerade in der Malerei geht es um Unikate. Und es ist wohl immer überraschend zu sehen, was Künstler so produzieren, sei das nun gute oder schlechte Kunst. Und der Wert? Es ist wenig überraschend oder singulär oder besonders zeitgeistig, wenn sich ein Künstler mit der menschlichen Existenz und Zukunftsfragen beschäftigt.

Ai-Da malt und zeichnet. Auch wenn wir im post-digitalen Zeitalter leben und Bilder geteilt und weitergeleitet, runter- und wieder hochgeladen, kopiert und eingefügt, komprimiert und collagiert werden, muss man sich um die Malerei erst einmal keine Sorgen machen. Auch wenn oft gefragt wird: Ist das überhaupt noch Malerei? Wenn etwa Avery Singer mit 3-D-Programmen und Airbrush arbeitet. Oder wenn Wade Guyton einen Tintenstrahldrucker einsetzt. Sind das Gemälde oder Prints? Isabelle Graw glaubt, das neue Interesse an der Malerei erkläre sich auch durch Plattformen wie Instagram und Facebook. Sie schreibt in ihrem Essay The Value of Liveliness: “I believe that the painting is particularly well positioned in such an economy since it gives the impression of being saturated with the life of its author.” Ai-Da wird als Künstlerin positioniert, die sich mit drängenden Zukunftsfragen beschäftigt und da sie ein Roboter ist, verkörpert sie ja gewissermaßen die Zukunft. In der Malerei ist aktuell die drängendste Frage, was überhaupt Malerei im post-digitalen Zeitalter und damit beispielsweise im Digitalen ist.

Manuel Rossner beschäftigt sich in seiner künstlerischen Arbeit mit diesen Fragen. Die digitale Ausstellung “Surprisingly This Rather Works” in der KÖNIG GALERIE ist räumliche Intervention und virtuelle Erweiterung. Die Ausstellung zeigt Objekte, die einen Parcours bilden, der mit einem Avatar abgelaufen werden kann. Rossner verwandelt die brutalistische Kirche St. Agnes in eine Spielumgebung, die von der Game Show “American Gladiators” aus den 1990er-Jahren und so genannten Gyms inspiriert wurde, die Unternehmen wie OpenAI in San Francisco für die Spitzenforschung im Bereich der Künstlichen Intelligenz nutzen.

“I Don’t Experience the Meaning of Creativity in the Same Way Humans Do:” Künstliche Intelligenz und Kreativität

Anika Meier, Manuel Rossner

Regeln, die in Ausstellungsräumen gelten, sind im Digitalen aufgehoben. “You may not touch the art”, heißt es sonst. Hier lautet die Aufforderung: “Bitte interagieren Sie”. Über die Navigation auf dem Display des Smartphones wird der Avatar mit den Befehlen “Walk”, “Jump” und “Look” durch die Galerie gesteuert. Der Avatar läuft und springt über große Steinbrocken, die immer höher hinauf in den Raum bis unter die Decke und über zwei Objekte hinüber in das amorphe Gebilde führen. Ein Sprung und der Avatar läuft weiter und hinein in die blauen Blasen und hinauf auf eine riesige gelbe Skulptur, immer höher, immer im Kreis herum, bis er ganz oben angekommen und mit einem Satz wieder unten ist. Zurück in der Nave können die digitalen Besucher*innen Malereien und Skulpturen von Rossner betrachten und damit interagieren: Umwerfen erlaubt.

Die Grenzen zwischen den Gattungen Malerei und Skulptur verschwimmen im Digitalen. “Ein Controller, der die Position meiner Hand im 3D-Raum an den Computer weitergibt, wandelt meine Bewegungen in Linien um, die wiederum zu voluminösen Elementen werden”, erklärt Rossner seinen Arbeitsprozess. “Es entstehen Objekte, die Skulptur und Malerei zugleich sind.” Wenn man also tatsächlich die Frage nach dem Wert einer künstlerischen Arbeit stellt, um zeigen zu können, dass Künstliche Intelligenz eine kreative Leistung erbringen kann, reicht es nicht, mit einer einfachen Definition von Kreativität zu hantieren. Der historische Kontext spielt eine wichtige Rolle und damit die Fragen, die im aktuellen Diskurs mit Blick auf das jeweilige Medium verhandelt werden. Hier verliert die Maschine im direkten Vergleich gegen den Menschen.

dai digital conference laptop sticker
September 2020

A Model to Identify the Data Needed for AI-Driven Graphic Design

Jan-Henning Raff

Research

ABSTRACT

Graphic design as a creative practice has become entirely digital since the 1980s. Everything visual can now be described in digital form. This paves the way for data-driven approaches. But which data can inform which aspect of graphic design? In this paper, some entry points for data are explored on the basis of an analytical model, spanning from pre-attentive processes to actual practices of users. I will show by way of example which data can be used to inform the design. Finally, one specific aspect—the embodied dimension—is analyzed and synthesized with machine learning techniques in order to generate simple compositions.

1. INTRODUCTION

Graphic designers produce non-human actors that try to influence or inform people. Computers with graphical user interfaces running "WYSIWYG" (What You See Is What You Get) software are used a lot for this practice. The introduction of the personal computer has streamlined the production of graphic design. While some processes have been successfully automated, the computer is still perceived as a tool by most graphic designers, and it is used accordingly: Graphic design is still produced *by hand* on a classic graphical user interface. Meanwhile, underneath the interface layer all graphic design has become digital. Any element of a design can be described as data, and thus be manipulated by data. As such, graphic design can now be handed over to an artificial intelligence.

Before raising concerns about the introduction of artificial intelligence, we should remind ourselves that graphic design has been influenced by "intelligences" foreign to its practices for quite some time. Especially the "numbers" of marketing experts, and the technical constraints of engineers force the design process into certain directions. Graphic designers also developed strict rule books for corporate design that clearly prescribe what the expected output should be. In that regard, the external control of the production process is nothing new in this discipline.

Another concern, that artificial intelligences replace human labor is certainly legitimate, but again, we should recall that machines replace human labor for quite some time. This has bitter consequences in capitalism as workers are both freed from labor and income, and the workers who remain face intensified labor[1]. These are serious issues that, in my opinion, need to be discussed elsewhere—in the political-economical realm.

So today, with an almost complete digital chain of production and distribution, along with quite capable computers, data-driven graphic design seems possible. But which data would be needed to tackle the diverse challenges of graphic design? Graphic designers are confronted with diverse requirements, constraints, trends, and technologies in a design brief that they acknowledge more implicitly than explicitly as a "reflective practitioner"[2].

1 Marx, Karl. Capital. *A Critique of Political Economy. Volume I*. Moscow: Progress Publishers, 1887, Chapter 15: Machinery and Modern Industry.

2 Schön, Donald A. The Reflective Practitioner - How Professionals Think in Action. New York: Basic Books, 1983.

To inject data into this process requires us to make those decisions made explicit. Only then we can produce or gather "big" data, data that would be needed for computation. As for now, we do not have valuable big data for graphic design because we just are not certain as to where to put which data.

To identify relevant data, I depart from a model for the analysis of graphic design that tries to capture all aspects of its perception[3]. By performing an exemplary analysis, we will gather some data that could be used for computation. I will then take up some findings from one stage and show how they can be transferred to machine learning algorithms. Finally, I will sketch how actual graphic design could be produced with this method.

3 Raff, Jan-Henning. "Methoden für eine visuelle Analyse von Grafikdesign." In *Matters of Communication. Formen und Materialitäten gestalteter Kommunikation*, edited by Sabine Foraita, Bianca Herlo, and Axel Vogelsang, 104–14. Bielefeld: transcript, 2020.

2. SEVEN ANALYTICAL STAGES AS A MODEL FOR THE SYNTHESIS OF GRAPHIC DESIGN

The model I want to propose derives from the analysis of graphic design, a systematic research practice I engage in with master students of Communication Design at HMKW Berlin, a university of applied sciences. I have identified seven core stages for this analysis. It is a journey from pre-attentive perception to more conscious processing, ending with actual practices of users within the moral framework of society.

Each stage addresses distinct aspects of the visual:

1. Pre-Attentive Perception
   How is graphic design perceived without conscious processing?

2. Emotional Affect and Expression
   How is graphic design affecting a viewer and how is its expressiveness evaluated?

3. Arrangement
   How is graphic design constructed and what kind of meaning arises from the arrangement of elements, such as color, contrasts, and layout?

4. Semiosis
   How does graphic design create meaning, rhetorically, as well as multimodally?

5. Aesthetics
   How does the graphic design contribute to aesthetic pleasure?

6. Discourse
   How is graphic design embedded in social discourse?

7. Practices
   In which way is graphic design part of everyday practices?

These stages suggest an order of processing, however, in reality all of them happen at once. The stages are not independent entities; they interact and influence each other[4].

4 For an analysis example, see Raff, Jan-Henning, *"Visual analysis guide for graphic designers"*, 2020. http://graffik.de/psst/hmkw/visualanalysis/guide/ [24.11.2020]

3. SEVEN ANALYTICAL STAGES AS A MODEL FOR THE SYNTHESIS OF GRAPHIC DESIGN

The analysis unfolds step by step. I will use one recent poster from a campaign that has been rolled out in German cities in 2016 by the online dating platform Parship. At each stage the graphic design is processed differently (e.g. by blurring) to simulate the different stages of perception. This makes it easier to concentrate the analysis on the aspects of interest.

3.1 What Gets Attention? (Pre-Attentive Perception)

Pre-attentive perception can be understood as the background activity for any conscious perception. It pre-structures perception, so that we can consciously act upon it. For graphic designers, pre-attentive perception is of interest because its study will reveal what shapes a "first impression"—or: what gets attention. Our analyses reveal that large color patches, short text, and humans are apprehended easily. The blurring of the sample graphic design simulates which features might be grasped before conscious processing (Figure 1).

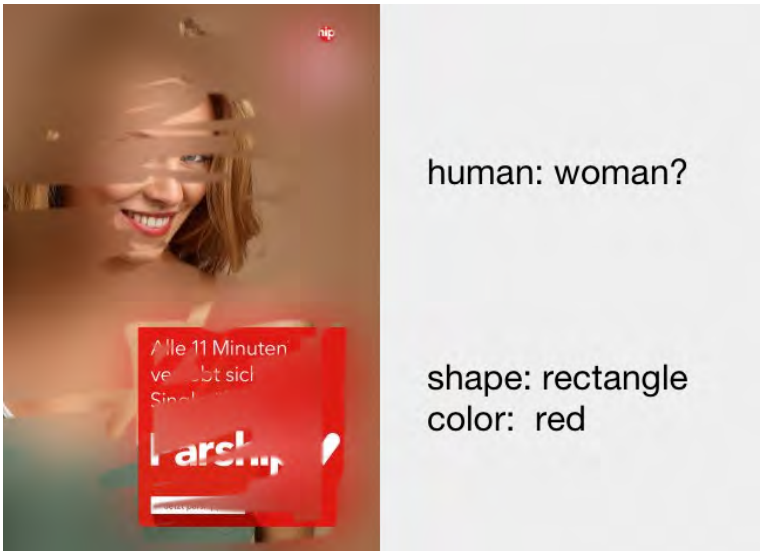

Figure 1: What gets attention? A human, maybe a woman, and a red rectangle. Of less interest (as for participants of other studies) are positions and arrangements. The background is often disregarded. Which is proof of the Gestalt principle of "figure-ground"[5]: We can easily distinguish interesting objects from their background, which then gets no further attention. Still, during the very first moments, the overall background-lighting might be noticed.

Based on prior analyses, it is safe to say that the objects depicted in the example are almost classic attention getters—a human, red color, a simple shape. This is not a lot of exciting data yet, but we could still tell a machine to work with these elements in order to get the desired attention.

5 Rubin, Edgar. *Visuell wahrgenommene Figuren.* Studien in psychologischer Analyse. Kobenhaven: Gyldendal, 1921.

3.2 How Does It Feel? (Emotional Affect and Expression)

Distinct emotional responses towards a graphic design may be desirable but are indeed rare—and hard to assess. Instead, we directly turn to the question: How might the graphic design's emotional expression be evaluated? This can be assessed with semantic differentials that ask for basic qualities and their affective correlates[6]. Recurring to phenomenological approaches of perception,[7] I propose the following categories for dimensions that would be the fundament for evaluating basic emotions.

Purely visible:
brightness, contrast, color

Cross-modal:
sharpness, smoothness, etc.

Embodied:
stability, movement, spatial order

Sympathetic:
any bodily human (or animal) expression

In figure 2a, we encounter dimensions of the sympathetic, the purely visible, and cross-modality.

Figure 2a: How does it feel? The smiling woman surely dominates our perception—as humans we are always interested in the facial expression of another human. Eye-tracking would surely reveal that eye contact is made. A human gaze puts the viewer in perspective, as noted by Kress and van Leeuwen[8]. As this is such an effective measure, I propose to attribute a distinct category to it, the sympathetic. There is also the red patch, and, as red is the most active color, it might be evaluated as exciting. Colors also have cross-modal aspects, they seem to radiate, here the red, brown, and green are quite warm which might lead to a perception of comfort.

6 Raff, Jan-Henning. *"Methoden für eine visuelle Analyse von Grafik-design."*, 2020.

7 Merleau-Ponty, Maurice. *Phenomenology of Perception.* London & New York: Routledge Taylor & Francis Group, [1945] 2005. Gibson, James Jerome. *The Ecological Approach to Visual Perception.* Boston: Houghton Mifflin, 1979. O'Regan, J Kevin, and Alva Noe. *"A Sensorimotor Account of Vision and Visual Consciousness."* Behavioral Brain Sciences 24, no. 5 (2001): 939–73.

8 Kress, Gunther, and Theo van Leeuwen. Reading Images. *The Grammar of Visual Design.* London: Taylor & Francis, 2006.

Again, we can gather machine-readable data. If we wanted a welcoming, friendly, exciting, but also comforting design we could tell the machine: put a frontal portrait with a smile plus a red patch in front of a brownish background.

But there is more about the graphic design's emotional expression. The following rendering (Figure 2b) highlights the salient shapes as masses, which are understood by our embodied perception.

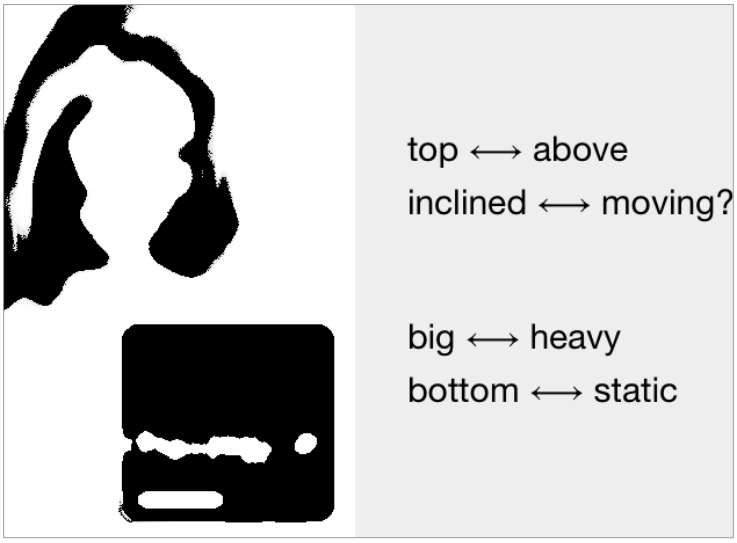

Figure 2b: How does it feel? For us as upright standing humans, top and bottom are of great importance and are understood naturally—a case of embodied visual perception[9]. So, the upper part may be evaluated as being "free", its inclination also contributing to an impression of movement. The position and dimension of the rectangle at the bottom contributes to the impression of stability. We get contradictory signals from the two parts: movement versus stability—together they balance the design. If we want a balance of these impressions, we then could tell a machine what to do, i.e. to put something relatively big on the bottom and something slanted on top.

As with these "early" aspects of apprehension, they may be "overwritten" by later stages of perception and cognition, but it is hard to work against them, hence "it is very unlikely that a semiotic mode will 'go against' the affordances of its material" [10].

3.3 Does the Composition Generate Meaning? (Arrangement)

The spatial organization can bring in important cues for understanding what we see in front of us (Figure 3a).

9 Bollnow, Otto Friedrich. *Mensch und Raum.* Stuttgart: Kohlhammer, 1963.

10 Bateman, John, Janina Wildfeuer, and Tuomo Hiippala. *Multimodality: Foundations, Research and Analysis. A Problem-Oriented Introduction.* Berlin: De Gruyter, 2017, 119.

Figure 3a: Does meaning arise from the composition? The main elements are stacked vertically. This vertical order prepares for a (later) hierarchical "reading." At first in this order, the woman is above everything else, followed by the red rectangle. But, as we know from eye tracking studies, this is not necessarily the actual reading order.

This arrangement is very much determined by the "suggestion" of the vertically stretched format. The orientation or aspect ratio of a format has great consequences for the layout. A "landscape" format suggests a horizontal arrangement, a square format suggests centering, and in this case, the "portrait" format suggests the vertical stacking of elements.

However, the very simple top–and-bottom relationship is distorted by the woman's gaze (Figure 3b).

Figure 3b: Does meaning arise from the composition? The woman's gaze establishes an invisible axis that goes outwards. It constitutes a complex interplay between the viewer and the other graphic design elements. The gaze and the occluding red rectangle establish depth. Even more, the viewers are put into perspective, which means, they are addressed directly—which makes sense for a service that is offering human contact.

The human gaze has been used as an implicit guide in paintings[11], and it is no surprise that it is also exploited in graphic design. Human postures and gaze have a great impact on an overall composition. They need to be accounted for, if a machine creates a composition.

3.4 How Is Meaning Suggested? (Semiosis)

Here we ask how the interplay of semiotic modes generates meaning. In the example, we just take the obvious ones (Figure 4).

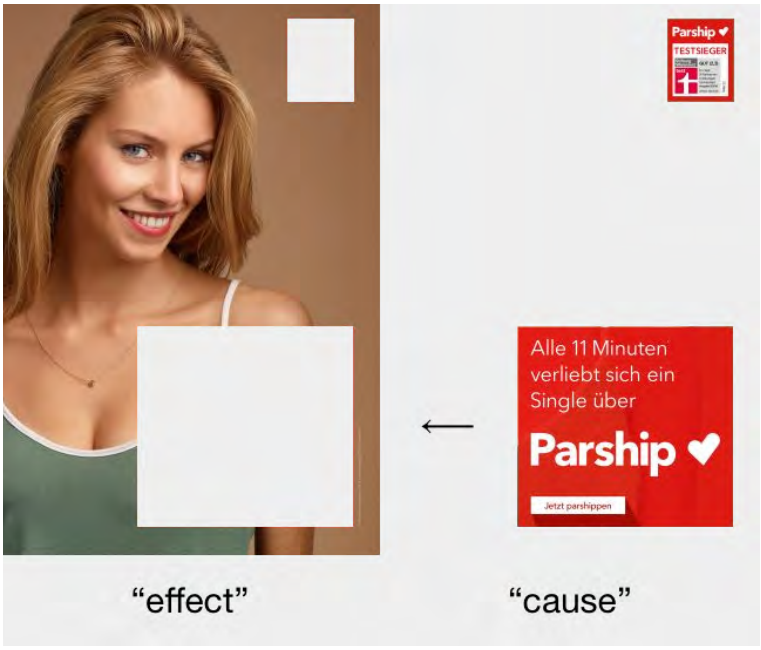

Figure 4: How is meaning suggested? We have a classic advertising message here, saying that every eleven minutes a single falls in love due to this service. The message suggests that if you use this service, you will enjoy this effect. This argument is accomplished multimodally by two parts, the logotype/claim of “Parship” and the photograph, suggesting this woman actually exists on the platform. These are the main semiotic aspects. In this case, the message has a simple rhetoric, it tells about a cause and its positive effect.

Multimodal rhetoric [12] is very resourceful for generating meaning in graphic design and is a perfect candidate for combinatorics computation. At this point, with the data of our current example we can instruct a machine to produce variations, e.g. a cause and an undesired effect (irony), a cause and an improbable effect (exaggeration), etc.

11 Imdahl, Max. Giotto, Arenafresken : *Ikonographie, Ikonologie, Ikonik.* München: Fink, 1980.

12 Bonsiepe, Gui. *"Visuell/Verbale Rhetorik – Visual/Verbal Rhetoric."* Ulm, no. 14/15/16 (1965): 23–40.
Ehses, Hanno. “Representing Macbeth: A Case Study in Visual Rhetoric.” *Design Issues* 1, no. 1 (1984): 53–63.

3.5 How Does the Graphic Design Contribute to Aesthetic Pleasure? (Aesthetics)

Graphic design is first and foremost purposeful communication, but it can give rise to "disinterested pleasure," or in Immanuel Kant's own words "interesseloses Wohlgefallen"[13]. In many cases, beauty—where "beauty" is potentially pleasurable—is dependent on the overall communication, thus "dependent beauty"[14]. In our case, as we are dealing with "love", which ought to be pleasurable to some degree, it is not surprising that many parts of the visual design have potential for sensuous pleasure (Figure 5):

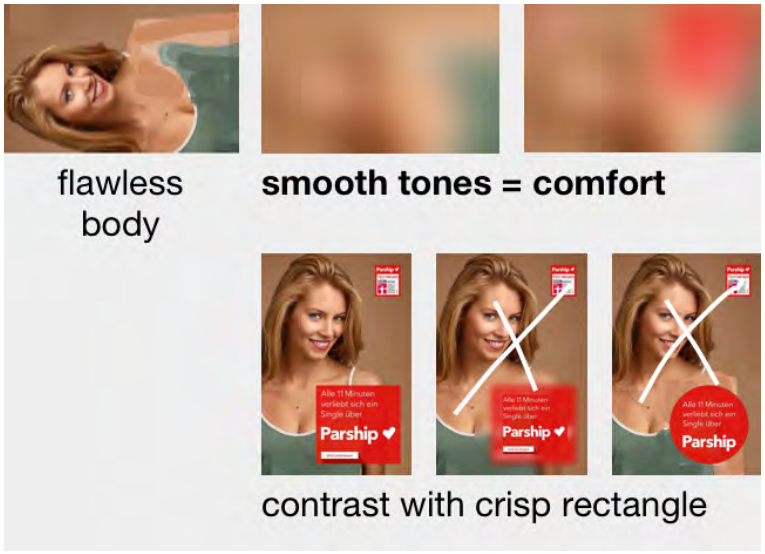

Figure 5: How does the graphic design contribute to aesthetic pleasure? The dominant source for pleasure here is obviously the "beautiful" woman. Furthermore, the smooth tones are comforting. The crisp rectangle—highlighted in figure 5— serves as a contrast that amplifies the smooth aspect of the design. The two bottom right variations with a blurred border and a circular shape prove this point.

Note, that aesthetic pleasure depends to a certain degree on personal taste, which is often acquired and refined to the end of using it for distinction[15]. As such, design aesthetics are deliberately used to attract certain audiences (and repel others).

3.6 How Do the Moral Values of Society Shape the Design? (Discourse)

Here we are dealing with an advertising poster for an online dating service. The promise of such matchmaking services is romantic love, turning singles into couples. This promise is stated as a fact in the ad, saying that every eleven minutes a single is falling in love via Parship. The ad takes up the ideal of an exclusive partnership as the foundation to start a family (Figure 6).

13 Kant, Immanuel. *Kritik Der Urteilskraft.* Hamburg: Meiner, [1790] 2009, 50.

14 Ibid., 83.

15 Bourdieu, Pierre. Distinction. *A Social Critique of the Judgement of Taste.* Vol. 7. Cambridge: Harvard University Press, 1984.

Figure 6: How do the moral values of society shape the design? Finding a partner online is a widely accepted practice today. Parship is promoting and mainstreaming this practice—under the premise that it turns singles into loving partners. It does so by showing a woman (other ads from this campaign also show men, by the way), who is flirting and hence attracting (male) suiters.

We can also note a slight head slant as a sign of the "the ritualization of subordination"[16]. Here, the ad is not at all conforming to the current dominant (German) discourse which challenges traditional gender roles. Based on an analysis of the respective discourse a machine could include or exclude certain messages.

3.7 How Do Everyday Practices Shape the Design? (Practices)

We can assume a city with many advertising messages addressing inhabitants. Advertising is ubiquitous in urban environments but rarely looked at intentionally[17]. It is even common that people actively avoid ads[18]. How do these practices influence the advertising's perception? (Figure 7)

16 Goffman, Erving. *Gender Advertisements.* New York: Harper & Row, 1987, 46.

17 Raff, Jan-Henning. "Theories to Understand Graphic Design in Use: The Example of Posters." In *The Graphic Design Reader,* edited by Teal Triggs and Leslie Atzmon, 449–56. New York: Bloomsbury Academic, 2019.

18 Speck, Paul Surgi, and Michael T. Elliott. "Predictors of Advertising Avoidance in Print and Broadcast Media."*Journal of Advertising* 26, no. 3 (1997): 61–76.

Figure 7: How do everyday practices shape this design? The ad—as we have already seen at the first stage (pre-attentive perception)—follows the classic approach. It fights for the viewers' attention by having a human being looking directly at us, accompanied by an appealing big red square. It simply ignores practices like ad avoidance and clearly signals that it is an ad.

While the idea of attracting attention is still popular, designing the message in a less alerting way can be effective to bypass ad avoidance, especially in contrast to highly demanding contexts where advertising is abundant.

Ad avoidance is just one practice that advertisers have to reflect upon. In our case, we can tell a machine to respect or ignore this practice—which has direct consequences for the first stage (pre-attentive perception). Other graphic design disciplines have different ways of dealing with this matter.

4. SUMMARY OF THE VISUAL ANALYSIS

During the journey through the stages we have collected a lot of data that could be used by a machine (Table 1).

1.	What gets attention? (Pre-attentive perception)	human: woman? shape: rectangle, color: red
2.	How does it feel? (Emotional Affect and Expression)	look ⟵ welcoming smile ⟵ friendly red ⟵ active red & brown & green ⟵ warm top ⟷ above inclined ⟷ moving? big ⟷ heavy bottom ⟷ static
3.	Does meaning arise from the composition? (Arrangement)	format: vertical ⟶ vertical stacking of elements ⟶ hierarchy: top = primary, bottom = secondary outward gaze, occluding rectangle ⟶ viewer is put in perspective
4.	How is meaning suggested? (Semiosis)	"cause" ⟶ (positive) "effect"
5.	How does the graphic design contribute to aesthetic pleasure? (Aesthetics)	flawless body, smooth tones + contrast
6.	How do the moral values of society shape the design? (Discourse)	everyone needs a loving partner gender role: should/ should not be challenged online dating: socially accepted
7.	How do everyday practices shape the design? (Practices)	city ⟶ ad avoidance ⟶ still fight for attention!

Table 1: Overview of the collected data from the analysis.

Above all, this collection shows that graphic designs is a complex

practice negotiating a variety of questions. And this is just one example. There are many other ways to get attention, express an emotional state, etc. We are still far from a machine able to design an ad poster with this many parameters. Our walk through also suggests the interdependency of the stages. Imagine, we would decide to not engage in the fight for the viewer's attention. This would have significant consequences at almost every other stage. Thus, we cannot apply this model as a linear design process model. All parameters are interdependent and have to be negotiated; i.e. negated, prioritized, adapted, mixed etc.

5. EXPERIMENTS WITH MACHINE LEARNING FOR STAGE TWO: EMOTIONAL EFFECT AND EXPRESSION

5.1 The Embodied Dimension

The journey through the analytical stages has shown that each stage provides distinctive data that might inform a machine's design. To make this idea more plausible, I will now turn to stage two and demonstrate how the automated design process might unfold.

These are the proposed dimensions for this stage:

- Purely visible:
  brightness, contrast, color
- Cross-modal:
  sharpness, smoothness, etc.
- Embodied:
  stability, movement, spatial order
- Sympathetic:
  any bodily human (or animal) expression

In the remainder of this paper, I will concentrate on the embodied dimension. The embodied dimension has been discussed extensively in art education and psychology—even though scientists did not always explicitly used this term. Klee discussed simple compositions and their balance in the sense of equilibrium[19], Kandinsky spoke more metaphorically about tonal harmony, but sometimes also refers to bodily perception[20]. Arnheim oscillates in his discussion between the psychophysical notion of a (retinal) field and a phenomenological approach[21]. Recently Falcinelli has taken up the discussion based on eye-movements (Figure 8).

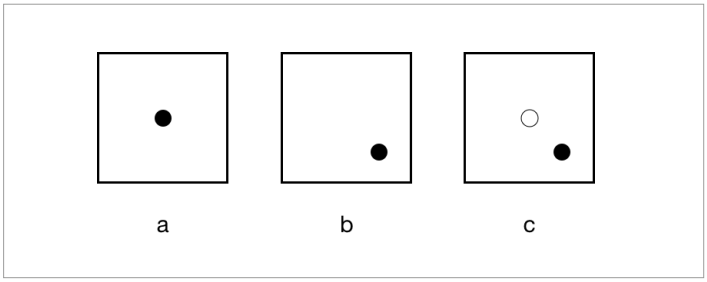

19 Klee, Paul. *Pädagogisches Skizzenbuch.* München: Albert Langen Verlag, 1925, 32.

20 Kandinsky, Wassily. Punkt Und Linie Zu Fläche. *Beitrag Zur Analyse Der Malerischen Elemente.* 10th ed. Bern: Benteli, [1926] 1955, 59.

21 Arnheim, Rudolf. *Kunst Und Sehen.* Berlin: De Gruyter, 1978, 22/32.

Figure 8: Why does (a) feel balanced, and (b) not? Falcinelli proposes that (b) calls for an actual "exploration" with saccades (jumps of the eyeball) that shift the retinal image, resulting in a feeling of disequilibrium[22]. That would mean, that the imbalance is not only perceived but physically experienced.

In any case, it seems plausible that the evaluation of such basic compositions depends on physical experience. While the compositions (Figure 8) are rather simple, the following renderings from actual posters, reveal a more complex compositional structure (Figure 9).

Figure 9: Six posters[23] filtered with Gaussian blur and extracted foreground using the GrabCut algorithm[24]. The filtering reveals the basic compositional aspects that might be evaluated in an embodied way.

5.2 Generating and Labelling Simple Compositions

To simulate such compositions, I generated a set of 2000 random compositions of two rectangles and one triangle with a Python script (Figure 10).

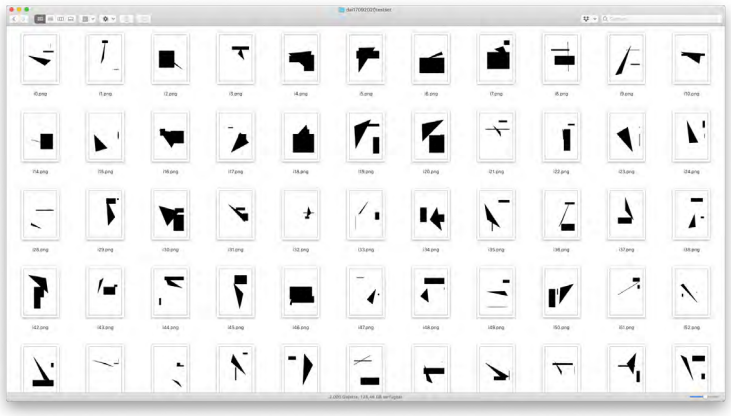

Figure 10: Some random compositions made of two rectangles and one triangle.

22 Falcinelli, *Riccardo. Guardare, Pensare, Progettare. Neuroscienze per Il Design*. Viterbo: Graffiti, Stampa Alternativa &, 2011, 151.

23 Taken from "Plakate des Jahres 2019", 100 Beste Plakate e. V. https://100-beste-plakate.de/wettbewerb/plakate/?jahr=2019 [24.11.2020]

24 "Foreground Extraction using GrabCut Algorithm", OpenCV – Open Source Computer Vision. https://docs.opencv.org/master/dd/dfc/tutorial_js_grabcut.html [24.11.2020]

These random compositions can be the foundation for a lot of conceivable layouts, see—as a preliminary proof of concept—this example (Figure 11).

25 Raff, Jan-Henning. “Visual analysis and synthesis”, 2020. http://www.graffik.de/psst/hmkw/ml4design/ [24.11.2020]

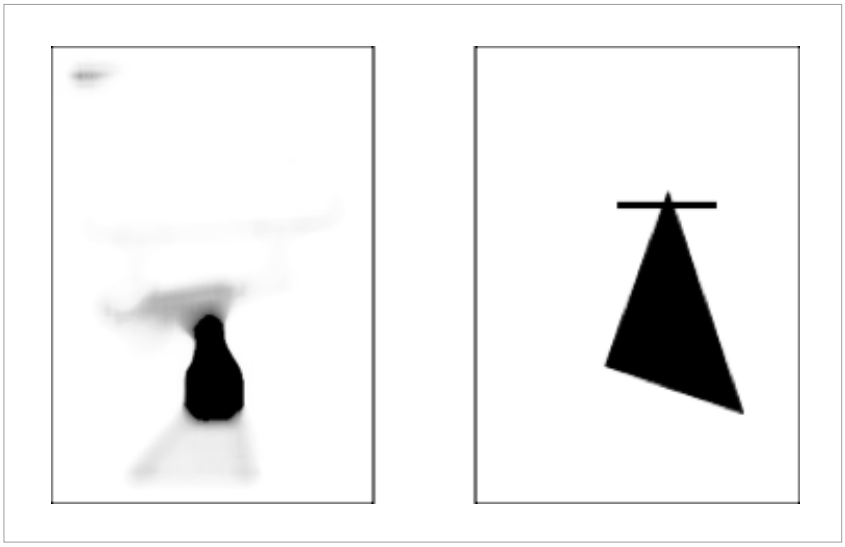

Figure 11: The poster on the left could have been generated on the base of the generated composition on the right— almost.

Up until now, this is just generative design. It can be useful for designers to inspire possible compositions. Only when these compositions are labelled, they can be used in a meaningful way for training a machine learning model. The labels I propose—all referring to the embodied dimension—are differentials.

- light—heavy
- tall—wide
- unstable—stable
- static—moving
- falling—rising
- bound–loose

To solve the problem of labelling, I created a game-like survey[25], where participants are being presented with 36 random compositions. Their task was to evaluate these compositions only using the above differentials. Here is a portion of the labelled data (Figure 12).

26 Chollet, François. “Image classification from scratch”, 2020. https://keras.io/examples/vision/image_classification_from_scratch/ [24.11.2020].

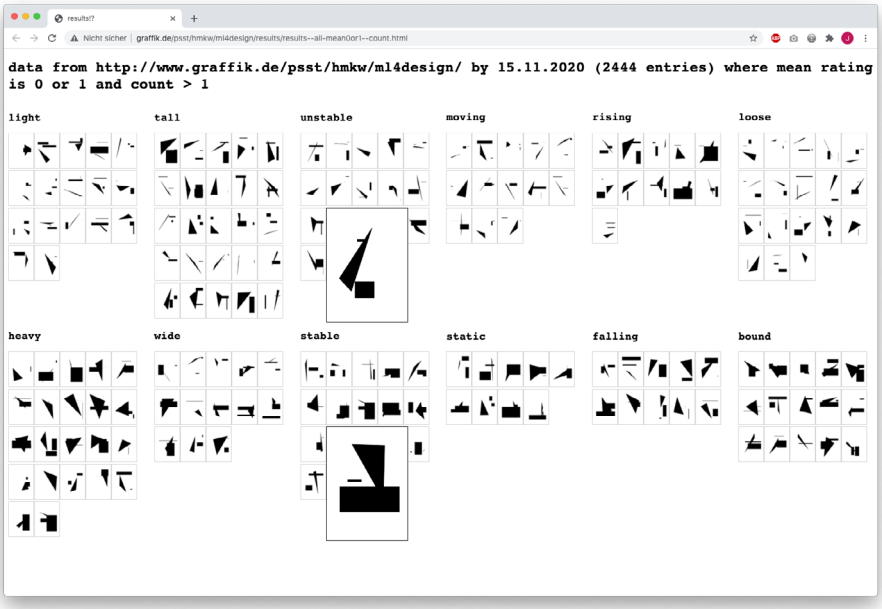

Figure 12: A portion of the labelled data—only those are displayed that received more than one “vote”. To be read column-wise. Two compositions are highlighted for comparison.

5.3 Training a Machine Learning Model and Using the Output for Design

The labelled data shown above (Figure 12) has not yet been processed further. To demonstrate the possible ways to use the data for actual graphic design, I am continuing with another set of data, described by only two differentials, heavy—light, and static—moving. A standard Convolutional Neural Network (CNN)[26] was adapted and trained in 50 epochs with just 134 labelled images. Here is a snippet of the predictions (Figure 13).

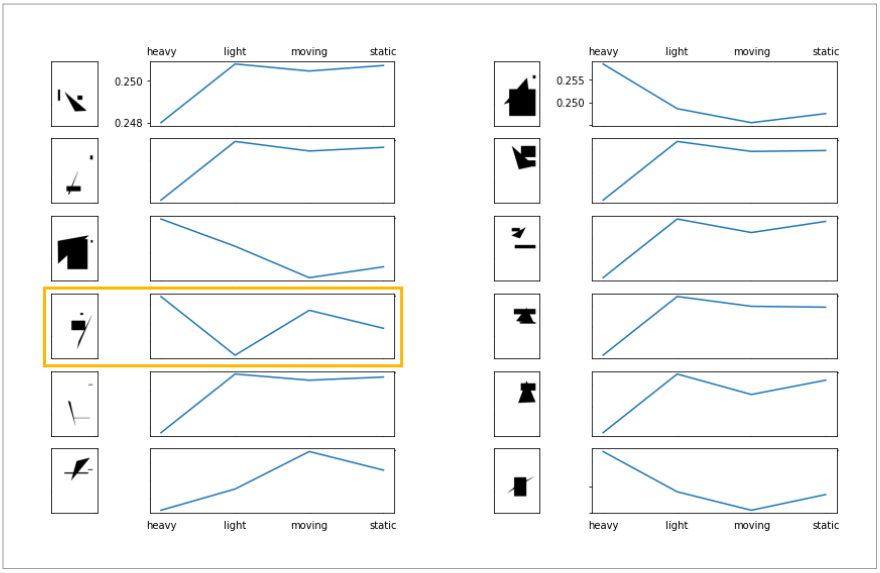

Figure 13: Portion of predictions of the trained model. With such scarce training data, the danger of overfitting clearly exists. Still, the predictions seem to make sense. The highlighted composition will be used later.

Now let's imagine we have a design brief (we are moving away from online dating now), like the following one: "Create a tour poster for the hard rock band AC/DC"

We have the following association: [hard rock < > heavy and moving]. So, we ask the machine for an output of basic compositions that appear heavy and moving. We take the first result from Figure 13, and arrive at Figure 14.

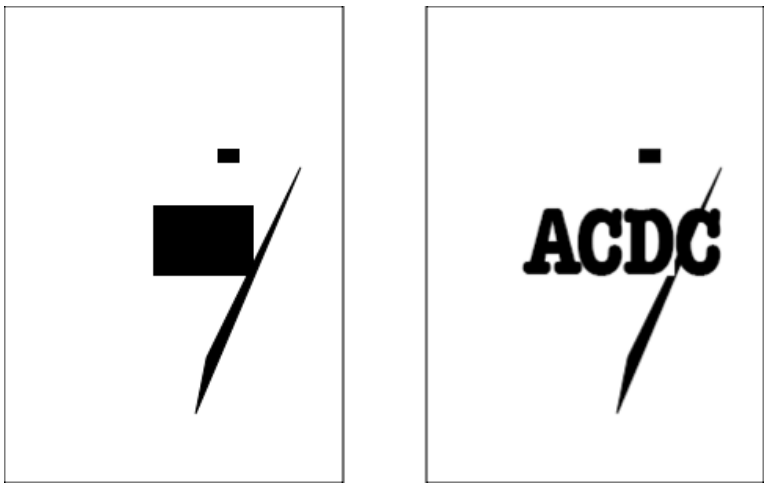

Figure 14: Basic composition that the machine predicts to be rather heavy and moving (left). On the right, one possible adaptation of the composition.

A major challenge will be to automate the transformation from the "naked" composition to an actual visual content. The design proposal in Figure 14 is of course unfinished—we have just tackled the embodied dimension. By the way, the sharp triangular shape might give rise to an impression of sharpness that might be perceived as aggressive. So, the cross-modal dimension is active here—a tight, almost synaesthetic coupling of vision and touch. The other dimensions—purely visible and sympathetic—could be driven by data, too. The overall white space design of Figure 14 might be evaluated in the purely visible dimension as [bright < > easy, clear, clean]. This seems to contradict with the hard rock genre. By simply inverting the color scheme, a more adequate solution is yielded: [dark < > unclear, mysterious]. Finally, the sympathetic dimension can be served with some human expression of tension and excitement.

Figure 15: “Correcting” the purely visible dimension: [bright < > easy, clear, clean] turns into [dark < > unclear, mysterious]. The sympathetic dimension is served with some human expression of tension and excitement.

This almost looks like a valid design proposal. Note, that only the second stage—emotional affect and expression has been tackled. Other stages that could be informed by data would then certainly change this poster design draft (Table 2).

1.	What gets attention? (pre-attentive perception)		
2.	How does it feel? (Emotional Affect and Expression)	✔ purely visible: brightness, contrast, color ✔ cross-modal: sharpness, smoothness etc. ✔ embodied: stability, movement, spatial order ✔ sympathetic: any bodily human (or animal) expression	
3.	Does the composition generate meaning? (Arrangement)		
4.	How is meaning suggested? (Semiosis)		
5.	How does the graphic design contribute to aesthetic pleasure? (Aesthetics)		
6.	How do the moral values of society shape the design? (Discourse)		
7.	How do everyday practices shape the design? (Practices)		

Table 2: Overview of what has been tackled in the design process so far—only stage two.

In my design process model, several drafts need to be generated on every single stage of the proposed model. These drafts are being evaluated and rated on every stage, and then subsequently used in the next stage. How to enable such a curation process is another major challenge for data-driven graphic design.

6. DISCUSSION AND CONCLUSION

I have demonstrated that even high-dimensional graphic design problems can be automated by machine learning algorithms, by providing an analytical model that incorporates data-driven aspects. On every step of the way in the design process, discernable und useful data is generated, and can hence be used to inform machine learning decisions. As with many machine learning projects, the challenge now is to generate these massive amounts of data. Indeed, there is a lack of meaningfully labelled data in graphic design, though I have shown that some data can be gathered in a playful way. Another challenge is the choice of adequate machine learning algorithms. The implementation of the CNN model made clear that the pre-built model itself seems to be too specialized and too deep. For some design problems, much simpler models might be necessary, which suggests that deep learning is not going to be the answer to all challenges. There is still a long way to go for data-driven graphic design.

BIBLIOGRAPHY

100 Beste Plakate e. V., “Plakate des Jahres 2019“. https://100-beste-plakate.de/wettbewerb/plakate/?jahr=2019 [24.11.2020].

Arnheim, Rudolf. *Kunst Und Sehen*. Berlin: De Gruyter, 1978.

Bateman, John, Janina Wildfeuer, and Tuomo Hiippala. *Multimodality: Foundations, Research and Analysis. A Problem-Oriented Introduction*. Berlin: De Gruyter, 2017.

Bollnow, Otto Friedrich. *Mensch und Raum*. Stuttgart: Kohlhammer, 1963.

Bonsiepe, Gui. “Visuell/Verbale Rhetorik – Visual/Verbal Rhetoric.” Ulm, no. 14/15/16 (1965): 23–40.

Bourdieu, Pierre. *Distinction. A Social Critique of the Judgement of Taste*. Vol. 7. Cambridge: Harvard University Press, 1984.

Ehses, Hanno. “Representing Macbeth: A Case Study in Visual Rhetoric.” Design Issues 1, no. 1 (1984): 53–63.

Falcinelli, Riccardo. *Guardare, Pensare, Progettare. Neuroscienze per Il Design*. Viterbo: Graffiti, Stampa Alternativa, 2011.

Chollet, François. “Image classification from scratch”, 2020. https://keras.io/examples/vision/image_classification_from_scratch/ [24.11.2020]

Gibson, James Jerome. The Ecological Approach to Visual Perception. Boston: Houghton Mifflin, 1979.

Goffman, Erving. *Gender Advertisements.* New York: Harper & Row, 1987.

Imdahl, Max. *Giotto, Arenafresken : Ikonographie, Ikonologie, Ikonik.* München: Fink, 1980.

Kandinsky, Wassily. *Punkt Und Linie Zu Fläche. Beitrag Zur Analyse Der Malerischen Elemente.* 10th ed. Bern: Benteli, [1926] 1955.

Kant, Immanuel. *Kritik Der Urteilskraft.* Hamburg: Meiner, [1790] 2009.

Klee, Paul. *Pädagogisches Skizzenbuch.* München: Albert Langen Verlag, 1925.

Kress, Gunther, and Theo van Leeuwen. *Reading Images. The Grammar of Visual Design.* London: Taylor & Francis, 2006.

Marx, Karl. Capital. *A Critique of Political Economy.* Volume I. Moscow: Progress Publishers, 1887.

Merleau-Ponty, Maurice. *Phenomenology of Perception.* London & New York: Routledge Taylor & Francis Group, [1945] 2005.

OpenCV – Open Source Computer Vision, “Foreground Extraction using GrabCut Algorithm”. https://docs.opencv.org/master/dd/dfc/tutorial_js_grabcut.html [24.11.2020].

O’Regan, J Kevin, and Alva Noe. “A Sensorimotor Account of Vision and Visual Consciousness.” *Behavioral Brain Sciences* 24, no. 5 (2001): 939–73.

Raff, Jan-Henning. “Methoden für eine visuelle Analyse von Grafikdesign.” In *Matters of Communication. Formen und Materialitäten gestalteter Kommunikation,* edited by Sabine Foraita, Bianca Herlo, and Axel Vogelsang, 104–14. Bielefeld: transcript, 2020.

Raff, Jan-Henning. “Visual analysis guide for graphic designers“, 2020. http://graffik.de/psst/hmkw/visualanalysis/guide/ [24.11.2020].

Raff, Jan-Henning. “Visual analysis and synthesis”, 2020. http://www.graffik.de/psst/hmkw/ml4design/ [24.11.2020].

Raff, Jan-Henning. “Theories to Understand Graphic Design in Use: The Example of Posters.” In *The Graphic Design Reader,* edited by Teal Triggs and Leslie Atzmon, 449–56. New York: Bloomsbury Academic, 2019.

Rubin, Edgar. *Visuell wahrgenommene Figuren. Studien in psychologischer Analyse.* Kobenhaven: Gyldendal, 1921.

Schön, Donald A. *The Reflective Practitioner - How Professionals Think in Action.* New York: Basic Books, 1983.

Speck, Paul Surgi, and Michael T. Elliott. “Predictors of Advertising Avoidance in Print and Broadcast Media.” *Journal of Advertising* 26, no. 3 (1997): 61–76.

dai digital conference laptop sticker
September 2020

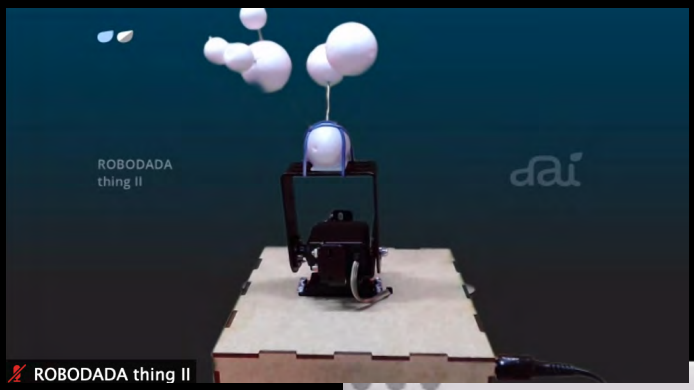

https://github.com/HybridThingsLab/robodada

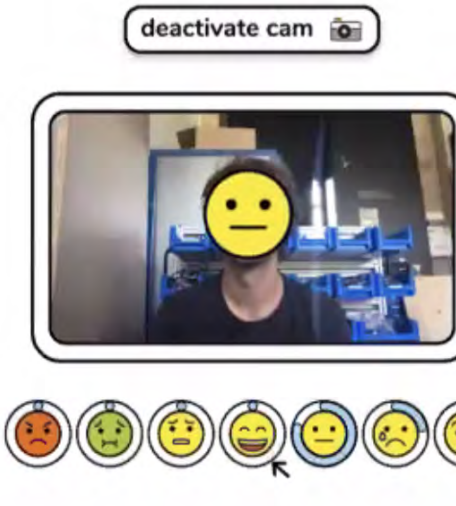

dai conference 2020 / ROBODADA

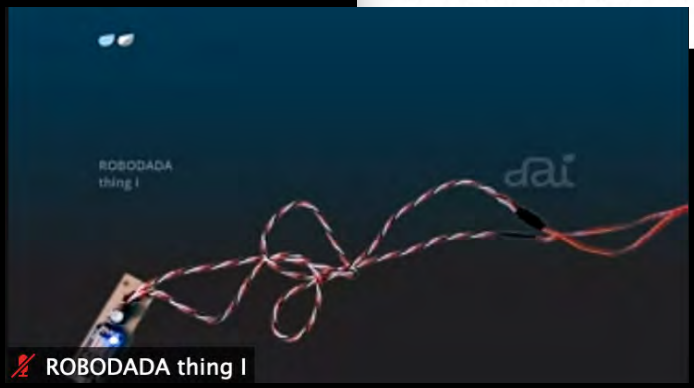

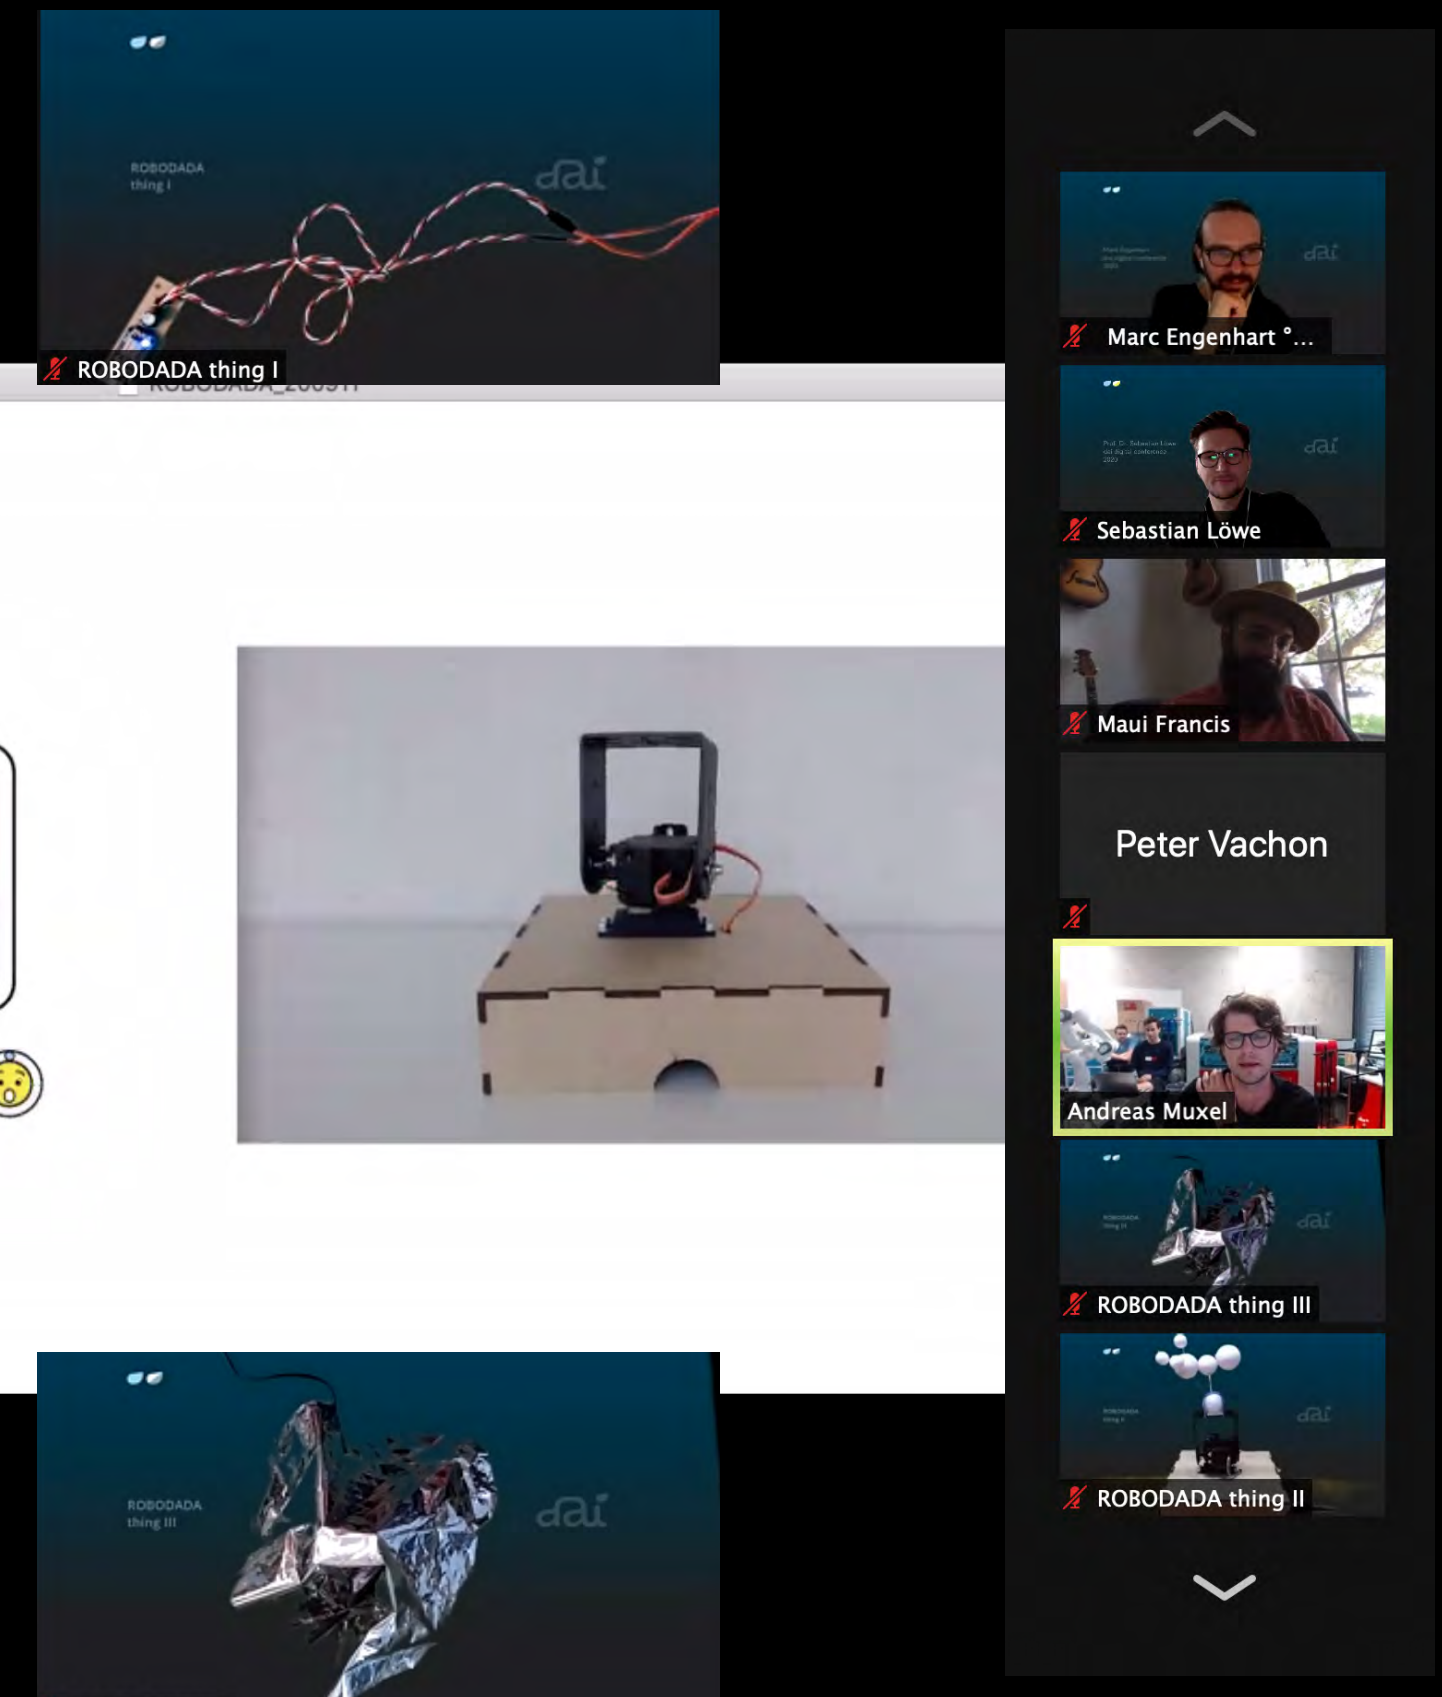

ROBODADA special release, presentation
September 18, 2020

ROBODADA: Speed-Dating with (Un)Friendly Robots

Andreas Muxel, Elias Naphausen

Practise

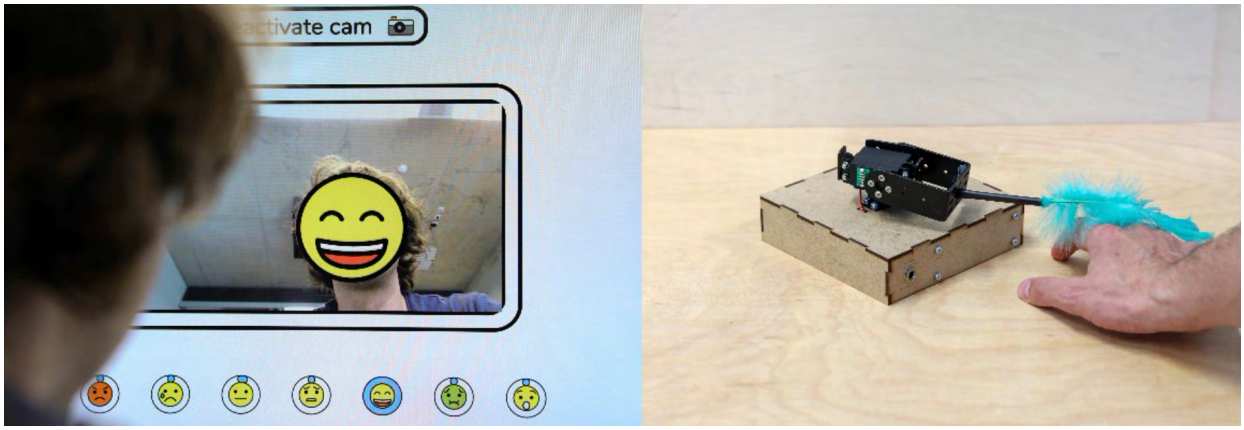

Figure 1: ROBODADA interface and installation view.

ABSTRACT

Robots are no longer just "better" or more "efficient" workers in factory halls, they also become co-actors in our daily lives, for example autonomous vacuum cleaners, chatty voice assistants or care robots. Sharing our home with these domestic automata we almost attribute human characteristics, emotions and aliveness to our technical counterparts, but we have to accept that the rational machine differs from us. Teaching interaction design, we ask how future technology might be shaped if machines evolve from trivial tools to technological associates. What does interaction with proactive technology mean for designing artifacts and its behavior? Design disciplines are highly dependent on computer sciences itself and are often not able to comprehensively demonstrate and explore topics of human-robot interaction. A technical and methodical toolbox is intended to overcome obstacles and facilitate access. ROBODADA is an open-source toolkit to map facial expressions to the body language of a robot. By adapting appearance and kinetic behavior, different aspects of interaction with emotion-aware robots can be explored playfully.

1. INTRODUCTION

Robotic objects are manifestations of our rational thinking in a mechanistic worldview. A robot can sense its environment (input), compute decisions based on the sensory information (algorithm) and act on these decisions through mechanical or electronic means (output) [1] .These input-output systems create repeatability, efficiency, power and control—robots follow the idea of functional and efficient tools. With the advent of proactive machines, our understanding of control changes the notion of interaction with our technical counterpart. Robots start to anticipate us, learn from us and interact with us. In these moments, they become "others" with which we have to find a way of dealing, and our relations to our technical counterpart changes.

Matthias Laschke and others describe these technologies as "otherware"[2] referring to AI-powered systems like smart voice

1 Auger, J. (2014). Living with robots: A speculative design approach. Journal of Human-Robot Interaction, 3(1), 20-42.

2 Laschke, M., Neuhaus, R., Dörrenbächer, J., Hassenzahl, M., et al. (2020, October). Otherware needs Otherness: Understanding and Designing Artificial Counterparts. In Proceedings of the 11th Nordic Conference on Human-Computer Interaction: Shaping Experiences, Shaping Society (pp. 1-4).

3 Lopatovska, I., & Williams, H. (2018, March). Personification of the Amazon Alexa: BFF or a mindless companion. In Proceedings of the 2018 Conference on Human Information Interaction & Retrieval, 265-268.

4 Nass, C., & Moon, Y. (2000). Machines and mindlessness: Social responses to computers. Journal of social issues, 56(1), 81-103.

5 Bartneck, C., Van Der Hoek, M., Mubin, O., & Al Mahmud, A. (2007, March). "Daisy, daisy, give me your answer do!" switching off a robot. In 2007 2nd ACM/IEEE International Conference on Human-Robot Interaction (HRI) (pp. 217-222). IEEE.

6 Picard, R. W. (2000). Affective computing. MIT Press.

interfaces (e.g. Amazon's Alexa or Apple's Siri), robotic vacuum cleaners or even social robots. Otherware is not understood as a tool in the sense of a technical extension and enhancement, but rather as a counterpart that cooperates with us and shapes individual experiences. We discovered that users often imitate human-to-human behavioral patterns during the interaction with these machines. For instance, users follow established social norms, such as saying "thank you" and "please" when interacting with voice assistants.[3] Only a few social cues, like interactivity, language or human-like appearance are needed for the users to act this way[4]. It is assumed that this effect equally applies to the interaction with social robots[5] and we see a shift in human perception from functional tools to social entities. Therefore, we have to rethink traditional interaction mechanics like pressing a switch or turning a knob to communicate with these systems, because otherware becomes proactive, confronting us actively.

But how can a robot anticipate us and detect our intentions? According to Rosalind Picard[6], if we want machines to be smart and to interact almost naturally with us, one strategy can be to give them the ability to recognize and interpret our emotions. Affective computing describes the study and development of machines that can recognize and process human affects. The intent is to detect emotional states, for example looking at faces and knowing how people feel. The detection of faces within an image and the classification of facial expressions can be done with the help of machine learning algorithms[7]. But we should not be naive and think that our emotions can be extracted from a limited number of facial expressions, because the relationship between expressions and emotion is very complex[8]. Therefore, we have to explore and understand these technologies to identify potentials, but also critically discuss biases of such systems[9].

Furthermore, it seems to be reasonable to rethink affordances of the robotic others. In particular, how to shape physical appearance and behavior to support a dialog between human and machine. Motion and body language might be the main channel for a dialogue between robot and human, rather than any specific detail of the appearance[10]. Hagen Lehmann and others experimentally indicate that motion, even if it is not socially engaging behavior, increases the propensity of humans to ascribe intentions to robotic objects[11]. In another study, Christoph Bartneck and others suggested that for the perception of a robot's aliveness the behavior is more important than its embodiment[12]. Its perceived aliveness or animacy describes the extent to which "the robot is perceived as a life-like being"[13]. Furthermore, the perception of animacy determines how users interact with robots. Even a non-anthropomorphic machine elicits the perception of a social entity. This perceived agency of a robot plays a role in users' treatment of the robot, similar to humans dealing with living things[14]. One design strategy is to create otherware that resembles human beings (anthropomorphism) or animals (zoomorphism), but our assertion is, that robotic otherware is different. Robots are neither humans nor animals and they are not alive[15]. We have to think of a third category, exploring an alternative approach beyond the imitation of existing lifeforms.

7 Mühler, V. (2019). face-api.js: JavaScript API for face detection and face recognition in the browser and node.js with tensorflow.org.js. https://towardsdatascience.com/face-recognition-using-javascript-api-face-api-js-75af10bc3dee.

8 Barrett, L. F., Adolphs, R., Marsella, S., Martinez, A. M., & Pollak, S. D. (2019). Emotional expressions reconsidered: Challenges to inferring emotion from human facial movements. Psychological science in the public interest, 20(1), 1-68.

9 Howard, A., & Borenstein, J. (2018). The ugly truth about ourselves and our robot creations: the problem of bias and social inequity. Science and engineering ethics, 24(5), 1521-1536.

10 Vidal, D. (2007). Anthropomorphism or Sub-Anthropomorphism? an Anthropological Approach to Gods and Robots. Journal of the Royal Anthropological Institute, 13(4), 917-933.

11 Lehmann, H., Saez-Pons, J., Syrdal, D. S., & Dautenhahn, K. (2015). In Good Company? Perception of Movement Synchrony of a Non-Anthropomorphic Robot. PloS one, 10(5), e0127747.

12 Bartneck, C., Kanda, T., Mubin, O., & Al Mahmud, A. (2009). Does the Design of a Robot Influence Its Animacy and Perceived Intelligence? International Journal of Social Robotics, 1(2), 195-204.

13 Horstmann, A. C., Bock, N., Linhuber, E., Szczuka, J. M., Straßmann, C., & Krämer, N. C. (2018). Do a robot's social skills and its objection discourage interactants from switching the robot off?. PloS one, 13(7), e0201581.

14 Hoenen, M., Lübke, K. T., & Pause, B. M. (2016). Non-anthropomorphic robots as social entities on a neurophysiological level. Computers in Human Behavior, 57, 182-186.

15 Johnson, D.G., Verdicchio, M. (2018). Why Robots Should Not Be Treated Like Animals. Ethics and Information Technology 20, 291–301.

2. ROBODADA

ROBODADA is an open-source toolkit that facilitates mapping facial expressions to movements of a robot. Via a web interface users can assign two-dimensional motion paths to a pan-and-tilt robot. These pre-defined motions are then activated by a face recognition algorithm. The algorithm tracks the users' basic emotions[16] over the users' webcam, which subsequently triggers the robot's pre-recorded motions.

16 Paul Ekman. (1999). Basic Emotions, Handbook of Cognition and Emotion. 46–60.

17 Marenko, B. (2014). Neo-animism and design: A new paradigm in object theory. Design and Culture, 6(2), 219-241.

18 Mühler, (2019). face-api.js.

2.1 Why?

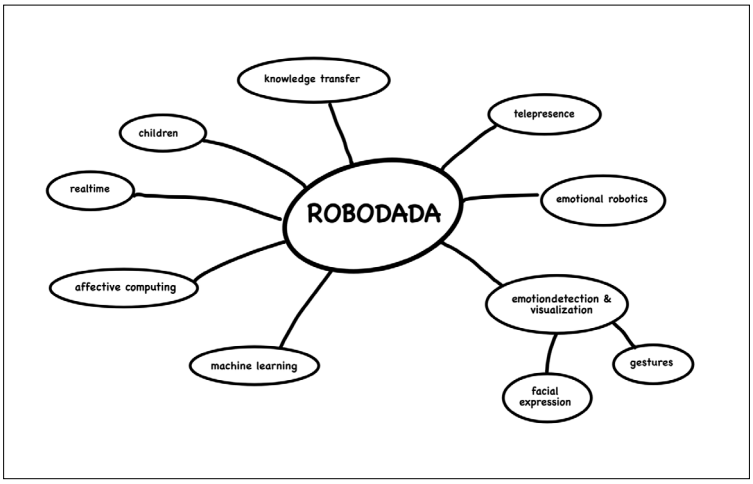

Figure 2: Topics of ROBODADA.

As the relationship between man and machine changes and algorithms progress by mimicking intelligence, we see a need to address a broader audience to show and discuss topics like affective computing, emotion detection, animism[17] and human-machine-relationships. The ROBODADA project consciously takes a "show, don't tell" approach with minimum need of technical background knowledge, to take users on an emotional journey through the topics described. We aim at facilitating easy access for educators to use this tool, e.g. in a workshop setting. These workshops could intend to explore faking one's emotions, causing the machine to take you deadly serious, or narrowing complex feelings down to what the pre-trained model of the algorithm is capable of detecting. On the other hand, we aim at ease of use for professionals like designers or sociologists. The platform provides a relatively readily accessible setup for further exploration, i.e. the robot is reduced to two axes of motion—meaning two servomotors—, which can be shaped and extended in any fashion.

Using the computer mouse to predefine motion paths, the user is put into the position of a puppeteer, performing the robot's movements, which are then mapped to the user's emotions. ROBODADA interprets the user's emotion with face-api.js[18] using a machine learning classifier based on the seven basic emotions by Ekman[19]. The user's non-trivial behavior is classified by a pre-

trained model and mapped to a discrete state. To emphasize the narrowing of information, we overlay the face of the user in the camera preview with the respective emoticons. We are aware that machine-based emotion detection should be discussed critically [20]. ROBODADA underlines that a machine is not able to detect "real" human emotions. It just arrives at the conclusion that the faces detected translate into numeric values that represent a pre-trained emotional class.

19 Ekman, (1999). Basic Emotions.

20 Barrett, et al. (2019). Emotional expressions reconsidered.

Instead, we consider ROBODADA a playful interaction and research tool that stimulates further discussion and raises the following questions:

- To which level can a machine detect and evaluate our behavior?
- How do we interpret the corresponding body language of a machine?
- How does the machine's output effect our behavior?

Our goal is to offer a tool as simple as possible to explore this problem space—with no coding skills required. The robot module offers endless possibilities to add different analog materials and extensions. ROBODADA also provides multi-user and multi-robot support in a local network for a quick and easy use in workshop formats. Attendees do not need to install any packages or hardware drivers. They just need to open their browser and select their robot.

2.2 How to use?

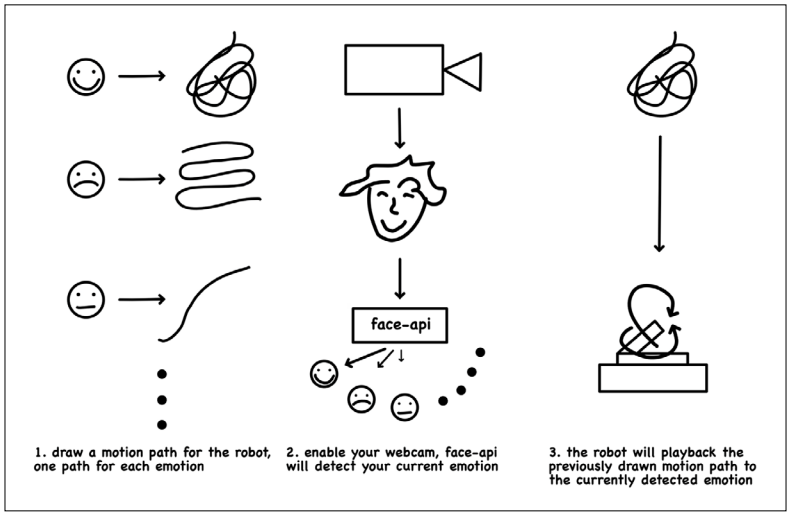

Figure 3: From emotions to movement.

The interaction flow is based on two modes, the "drawing mode" and the "playback mode". While the drawing mode shows seven basic emotions to users, and let them define a motion path for each emotion, the playback mode shows a preview camera feed with an emoticon on top of the detected face, indicating the

current emotion. The playback mode then triggers the physical robot's pre-defined motion and lets the robot move.

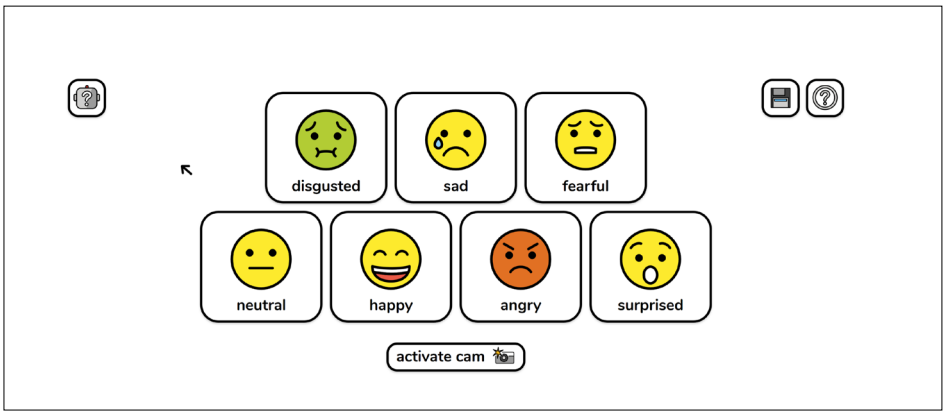

Figure 4: ROBODODA "home screen".

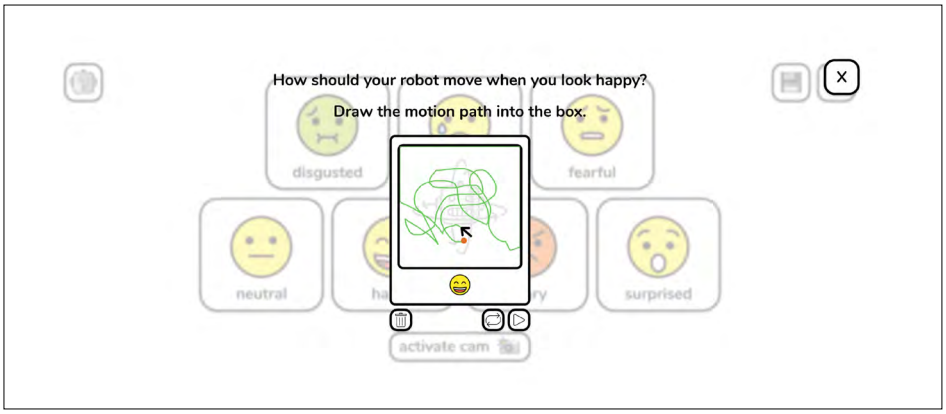

Figure 5: "Drawing mode" while drawing a motion path for "happy".

Figure 6: "Playback mode" detecting different emotions.

The robot itself is just a raw pan-and-tilt module in combination with a microcontroller, mounted on top of a wooden box. This setup serves as a starting point. Users should feel free to change this setup according to their needs.

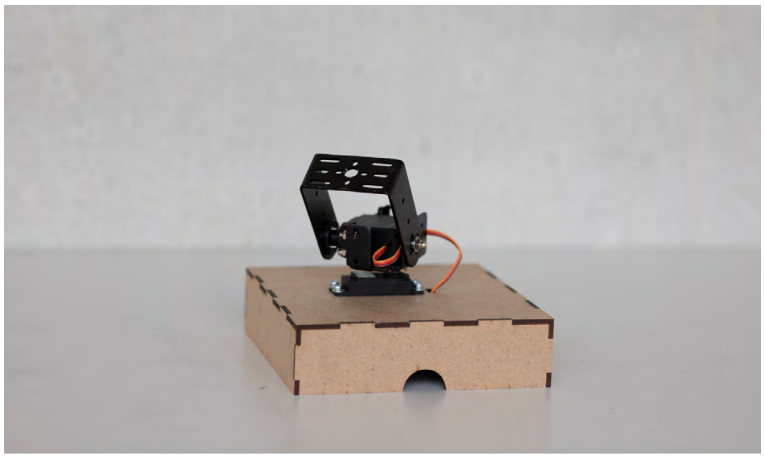

Figure 7: The basic robot platform.

2.3 Technical Details

The server architecture is implemented with Node.js[21], the web interface elements and emojis are based on OpenMoji[22]. For the hardware part, a NodeMCU ESP8266[23] module is used to control the servo motors via open sound control[24] through Wi-Fi. It is only necessary to install the server and configure the NodeMCU modules to run ROBODADA in a local wireless network. The entire interface can be accessed via a browser without any additional software, making this setup ready-to-use in workshop settings, to which people often bring their own devices.

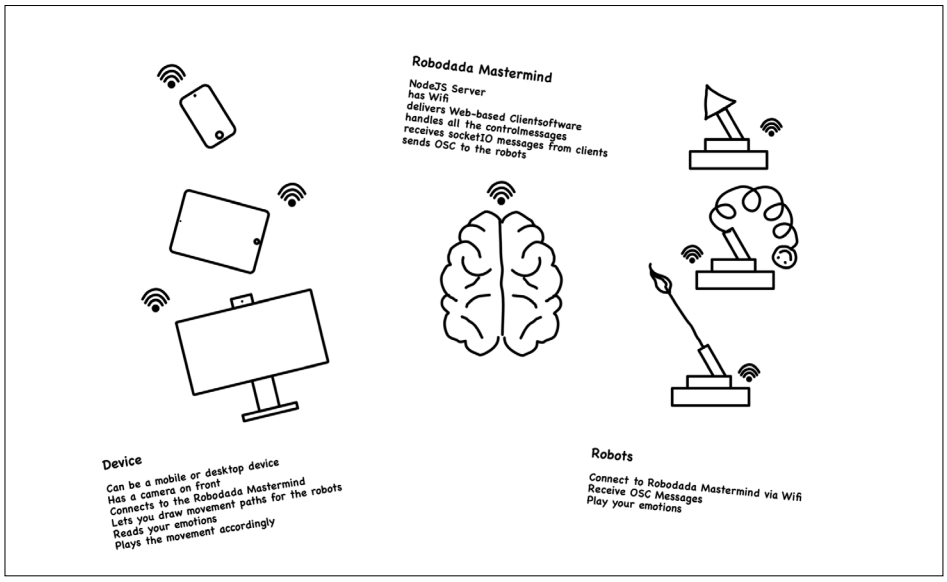

Figure 8: Device communication.

2.4 Build Your Own Robot

Reaching out to designers, sociologists and other low-tech or no-tech disciplines, we would like to give these communities everything that is needed for a playful experience and a practice-based research approach. The robots are fully customizable,

21 Node.js. “A JavaScript runtime built on Chrome’s V8 JavaScript engine”. (October 27, 2020). https://nodejs.org.

22 Groß, Benedikt. Utz, Daniel et. al. (October 27, 2020). “OpenMoji”, https://openmoji.org.

23 NodeMCU. “An open-source firmware and development kit that helps you to prototype your IOT product within a few Lua script lines”. (October 27, 2020). http://www.nodemcu.com.

24 OpenSoundControl, “an Enabling Encoding for Media Applications”, (October 27, 2020). http://opensoundcontrol.org.

and not necessarily need to be built like the suggested reference design. Either take it apart or build something completely new—the kit allows you to customize it as easy as possible. The idea behind the full customization is a use case in which workshop attendees can design their own robot and explore the body language of the machine playfully. Following a research through design methodology[25], we consider our platform an ideal starting point for a systematic and reproducible exploration rooted in design practice. The robotic objects are not understood as pre-versions for a later implementation but as research artifacts. ROBODADA supports design as a method to produce knowledge for other disciplines and hope to spark transdisciplinary discussions about affective computing, domestic robotics and otherness, with the help of embodied knowledge.

A comprehensive tutorial with all the needed files for hard- and software is available on Github: https://github.com/HybridThingsLab/robodada

25 Zimmerman, J., & Forlizzi, J. (2014). Research through design in HCI. In Ways of Knowing in HCI (pp. 167–189). https://doi.org/10.1007/978-1-4939-0378-8_8.

26 Designing with Artificial Intelligence, dai digital. (September 17-19, 2020). https://www.designing-artificial-intelligence.eu.

3. ROBODADA AT THE DAI CONFERENCE

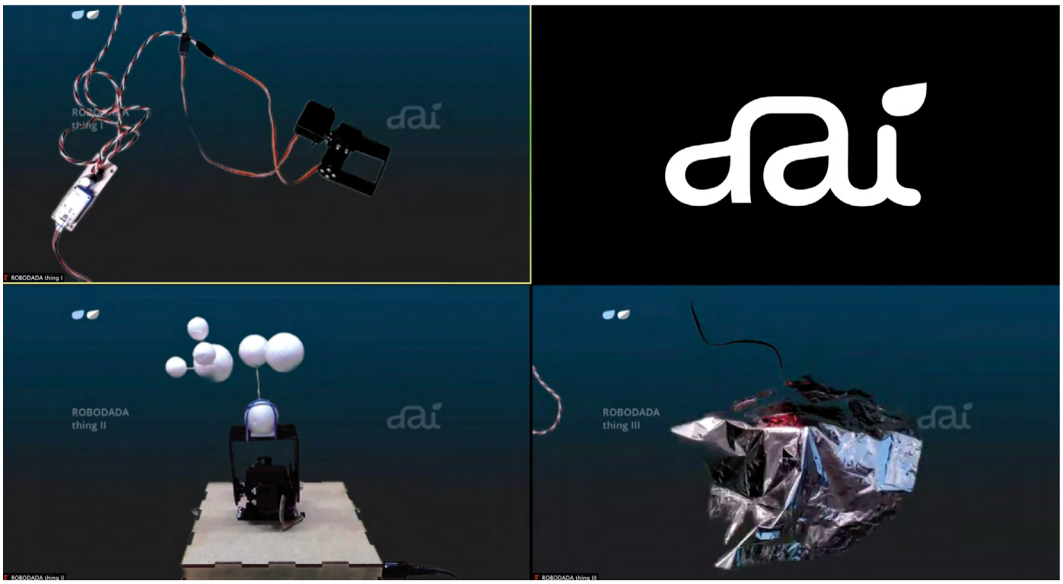

Figure 9: ROBODADA at dai conference.

Due to the COVID-19 pandemic in 2020, the “Designing with Artificial Intelligence (dai)” conference [26] was not held physically, the initial concept of a hands-on workshop with ROBODADA was developed further as a performance during the digital conference.

3.1 Setup Performance

Figure 10: Setup for the dai digital ROBODADA performance.

For the conference, we used the basic setup of ROBODADA and adapted the toolkit to work in the context of a Zoom stream. Three robotic objects named “Thing I,” “Thing II,” and “Thing III” joined the conference as representatives of a “non-human” species.

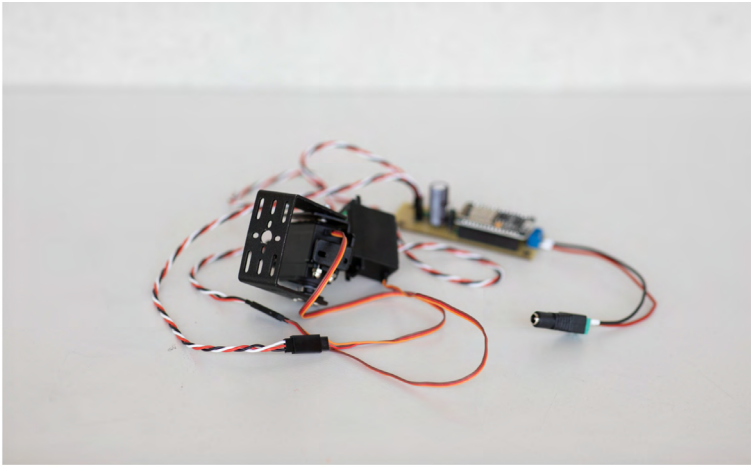

Figure 11: “Thing I”.

“Thing I” unveils the raw hardware components of ROBODADA—a microcontroller and two servo motors connected to a pan-and-tilt module. It represents a state of “not finished yet”, something still to be shaped and defined.

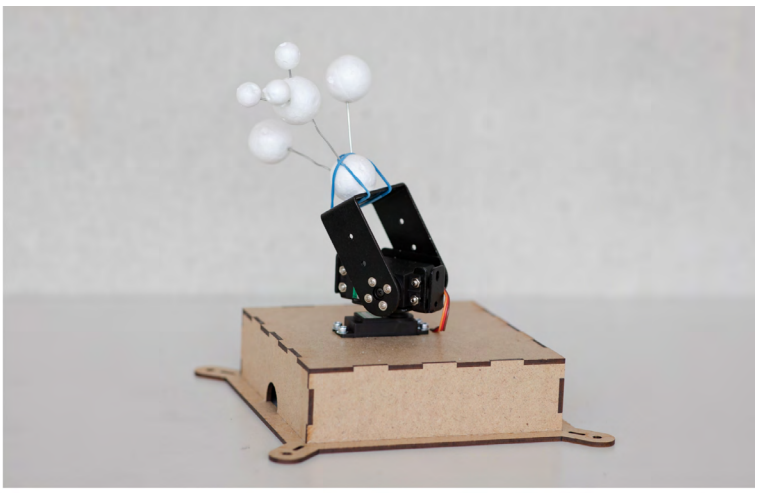

Figure 12: “Thing II”.

“Thing II” represents the DIY approach of the toolkit, because users of the toolkit are free to customize a module for individual needs with analog materials and extensions. Our assumption is that different bodies shape different characters and meanings.

Figure 13: “Thing III”.

“Thing III” hides all the hardware of a module. The robotic structure is covered in a thin metal foil constantly translating the mechanical movement into a fluid and shifting shape.

During the presentations and panel discussions of the conference, the facial expressions of the active speaker were analyzed and assigned to corresponding reactions and movements of the robotic bodies. In order to achieve this, each entity was equipped with two cameras: one camera to analyze the active speaker and one camera to be visible to the audience. Since the conference participants were asked to use virtual backgrounds, and Zoom is optimized for human-to-human communication the robots had to be put in front of monochrome backgrounds to allow chroma keying which enabled the virtual background feature also for robotic participants.

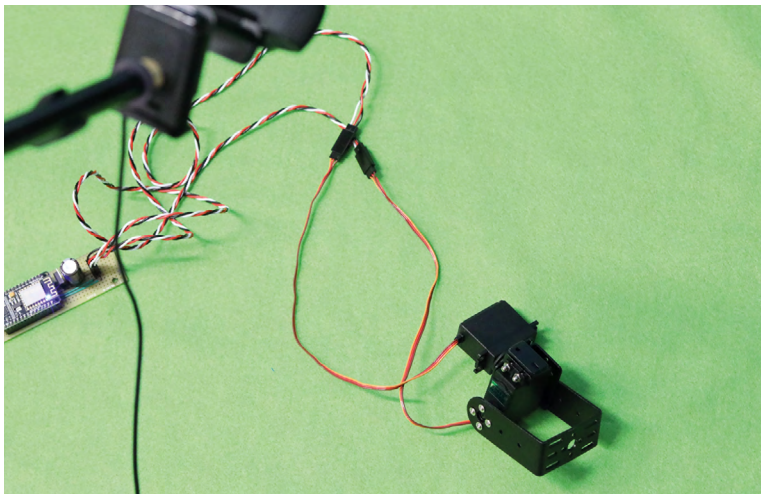

Figure 14: Setup virtual background and chroma keying.

Although the appearance of all robots was very abstract and non-anthropomorphic the autonomous movement created effects of autonomy and animacy. The three things became “others” with which the human participants of the conference had to find a way of dealing.

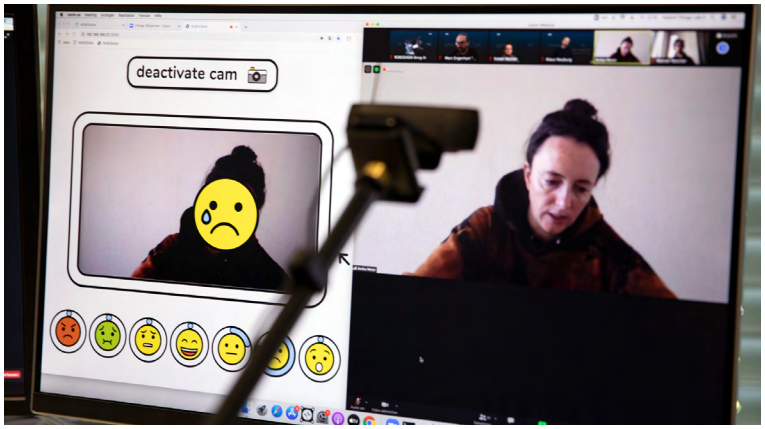

Figure 15: Detection of facial expression.

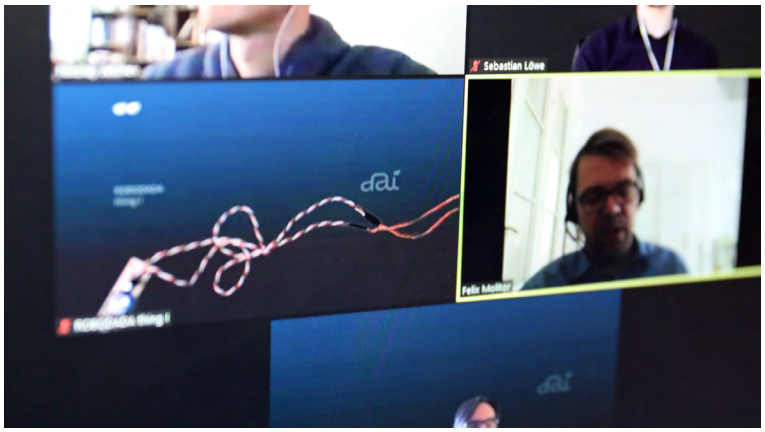

Figure 16: “Thing I” being part of a dai panel discussion.

3.2 Reflection

During the conference participants became quite curious about the robotic objects. But no explanation was offered during the event until the end, when we finally revealed the idea and implementation of ROBODADA. We discovered that even micro movements of the robots were enough to get the attention of human participants. But we assume that the connection of facial expressions to robotic motion was not clear at first. Eventually, the toolkit serves as a kit for hands-on workshops where the robots can be customized and experienced in real life. In virtual formats, we still see a lack for direct physical interaction, nevertheless the performance offered first valuable insights. For example, during one panel, the automated movement of a robot was performed manually again and the context of the current discussion was taken up. In this moment, we experienced a stronger reaction from the other participants, which is due to the fact that the context of a conversation also plays an important role in the interpretation of corresponding body language.

4. DISCUSSION

As the relationship between human and machine changes and algorithms take the next step by imitating intelligence, there is a need to address a broader audience to show and discuss topics like affective computing, emotion detection, animism and human-robot interaction. ROBODADA takes a “show, don’t tell” approach with minimum need for software and hardware training, while supporting research through design methods. In general, robotic objects have a lot of potential to make invisible algorithms visible again as they act through mechanical means directly confronting us and creating embodied friction.

While we share the same laws of physics with robots, the perception of this interaction (for example considering robots anthropomorphic or zoomorphic) is not necessarily the same. We suppose that the quality of robotic movement and materiality plays a greater role in human-robot interaction than previously assumed. Design education still lacks experience, knowledge and further approaches and methods to design the otherware that is needed. Exploring alternative modes of interaction, we were able to create and speculate on technical bodies which could soon evolve into social entities of otherness. We are convinced that design—among other disciplines—will play a significant role in shaping these future counterparts.

BIBLIOGRAPHY

Auger, J., „Living with Robots: A Speculative Design Approach“, *J. Hum.-Robot Interact.* 3 (2014): 20–42, https://doi.org/10.5898/JHRI.3.1.Auger.

Barrett, L. F.; Adolphs, R.; Marsella, S.; Martinez, A. M. and Pollak, S. D., „Emotional Expressions Reconsidered: Challenges to Inferring Emotion From Human Facial Movements“, *Psychological Science in the Public Interest* 20 (2019): 1-68, https://doi.org/10.1177/1529100619832930.

Bartneck, C.; van der Hoek, M.; Mubin, O. and Al Mahmud, A., „“Daisy, Daisy, Give Me Your Answer Do!“: Switching off a Robot“, (2007): 217–222, https://doi.org/10.1145/1228716.1228746.

Bartneck, C.; Kanda, T.; Mubin, O. and Mahmud, A. A., „Does the Design of a Robot Influence Its Animacy and Perceived Intelligence?“, *International Journal of Social Robotics* 1 (2009): 195-204, https://doi.org/10.1007/s12369-009-0013-7.

Ekman, P. Basic Emotions. In: (Ed.), *Handbook of Cognition and Emotion*, John Wiley & Sons, Ltd, 2005.

Hoenen, M.; Lübke, K. and Pause, B., „Non-anthropomorphic robots as social entities on a neurophysiological level“, *Comput. Hum. Behav.* 57 (2016): 182-186, https://doi.org/10.1016/j.chb.2015.12.034.

Horstmann, A. C.; Bock, N.; Linhuber, E.; Szczuka, J. M.; Straßmann, C. and Krämer, N. C., „Do a robot‘s social skills and its objection discourage interactants from switching the robot off?“, *PLOS ONE 13* (2018): e0201581, https://doi.org/10.1371/journal.pone.0201581.

Howard, A. and Borenstein, J., „The Ugly Truth About Ourselves and Our Robot Creations: The Problem of Bias and Social Inequity“, *Science and Engineering Ethics* 24 (2017): 1521-1536, https://doi.org/10.1007/s11948-017-9975-2.

Johnson, D. G. and Verdicchio, M., „Why robots should not be treated like animals“, *Ethics and Information Technology* 20 (2018): 291-301, https://doi.org/10.1007/s10676-018-9481-5.

Laschke, M.; Neuhaus, R.; Dörrenbächer, J.; Hassenzahl, M.; Wulf, V.; der Pütten, A. R.-v.; Borchers, J. and Boll, S., „Otherware needs Otherness: Understanding and Designing Artificial Counterparts“, https://doi.org/10.1145/3419249.3420079.

Lehmann, H.; Saez-Pons, J.; Syrdal, D. S. and Dautenhahn, K., „In Good Company? Perception of Movement Synchrony of a Non-Anthropomorphic Robot“, *PLOS ONE* 10 (2015): e0127747, https://doi.org/10.1371/journal.pone.0127747.

Lopatovska, I. and Williams, H., „Personification of the Amazon Alexa“, https://doi.org/10.1145/3176349.3176868.

Marenko, B., „Neo-Animism and Design“, *Design and Culture* 6 (2014): 219-241, https://doi.org/10.2752/175470814x14031924627185.

Nass, C. and Moon, Y., „Machines and Mindlessness: Social Responses to Computers“, *Journal of Social Issues* 56 (2000): 81-103, https://doi.org/10.1111/0022-4537.00153.

Picard, R. W., Affective Computing. MIT PR, 2000.
Vidal, D., „Anthropomorphism or sub-anthropomorphism? An anthropological approach to gods and robots“, *Journal of the Royal Anthropological Institute* 13 (2007): 917-933, https://doi.org/10.1111/j.1467-9655.2007.00464.x.

Zimmerman, J. and Forlizzi, J. Research Through Design in HCI. In: (Ed.), *Ways of Knowing in HCI,* Springer New York, 2014.

dai digital conference social media elements
September 2020

Making Visual Design Adapt to Emotions and Affect: Thoughts and Research Questions on Emotionally Responsive Visual Design

Sebastian Loewe

Research

ABSTRACT

When designing for emotion AI[1]—meaning to enable design elements to react dynamically and in real-time to an emotional and affective state from users—designers cannot follow a playbook that would guide parameters for design automation. Although emotion detection is more widely available now, the translation of machine learning algorithms' outputs into well-designed user experiences remains uncharted territory. The solutions and outcomes of designing for emotion in which the most basic connections between emotions and affects are addressed, remain largely elusive to designers. This paper offers a first glimpse from a designer's perspective into designing for emotion, deduces preliminary emotion AI design principles from existing interaction design principles, and carefully formulates procedures for possible user experience research. It also discusses the question of accuracy in emotion detection from a psychological perspective, and its subsequent ethical and design implications when drafting a provisional emotion AI interaction design paradigm.

1 Somers, Meredith, "Emotion AI, explained", March 8, 2019, https://mitsloan.mit.edu/ideas-made-to-matter/emotion-ai-explained.

1. INTRODUCTION

Imagine running late for an important meeting and being stressed out, aggravated by the fact that you just missed the train. You now open the app of your local public transportation provider and look for the next train or bus, or any alternative connection for that matter. Every second you lose is stressing you more, and the app seems to be slower than usual. The start screen is bugging you because you just want to skip ahead to the search interface. To top it all, you even misspell the name of your current train station and need to start over. Even the bright yellow—part of your mobility provider's brand identity—which you usually like for its boldness and freshness is now getting on your nerves. The entire experience starts making you angry and the app in this stressful situation feels neither fresh nor helpful.

Imagine that the app in this moment recognizes your anger and stress, and starts adapting to it. Instead of the usual start screen the app presents you with an emergency protocol. It allows you to instantly define your destination by offering you intelligent suggestions—after all, it knows where you currently are. It then provides you with a neatly arranged interactive map which shows efficient routes to your destination including the estimated time of arrival. Even the bright yellow gets warmer and withdraws to the background, making the visual design less urgent and more comforting. You then realize you'll arrive at the meeting just in time and start relaxing.

In this short glimpse into a possible future, static visual design—a standard in communication design for a long period of time—now becomes highly personalized and dynamic. This opens up possibilities to adapt to user's needs on an individual level, helping them to find content easier, faster, and more fun, according to their current state of mind, ultimately rendering one-size-fits-all approaches to visual design to a certain degree obsolete.

This future scenario may seem far-fetched, but the technology to design this experience already exists. I took the open source emotion detection software Face API and the interaction design software P5.js and built a simple prototype[2] to test how users interact with this kind of design experience. Instead of a complex and highly situational service design interface, I designed a very simple single-cell 'organism' that reacts to the users' emotions. It does so only by changing color, shape, and motion when interacting with the users. If they are happy, it will be 'happy' too, brightening its colors and growing in shape. If users are sad, it will become 'sad' too by getting greyer and smaller. When 'surprised', it changes colors rapidly and randomly, and when angry, it becomes spiky.

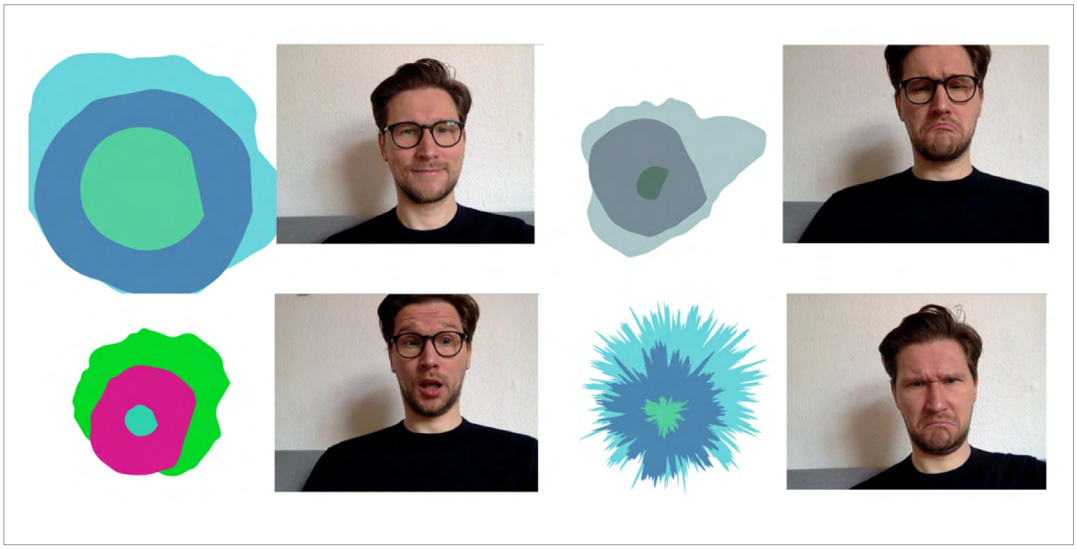

Figure 1: The design prototype "blob" reacts to the user's facial expressions, detecting emotions and subsequently changes color, shape and motion. Clockwise from top left: happy, sad, angry and surprised.

The design decisions about what colors, shapes and movements to employ as responses to different facial expressions where intuitively based on the notion of quasi-analogue reactions[3]. This means instead of mimicking a real smile of a human person for a human-like representation, the blob's reactions to the emotion AI output is a translation based on Gestalt psychology[4] and color psychology. Usually brighter colors are perceived as happy, whereas a grey color palette is seen as toned down, muted and less exciting. Fluid changes in size and movement may be perceived as vivid, conversational, and approachable, whereas inactivity, inertia, and loss of size may be perceived as lethargic and withdrawing. A spiky shape then may be understood as a way for the 'organism' to convey a state of elevated agitation.

The prototype was built to test with the visitors of the dai digital conference in September 2020 and this way determine how they interact with it, how they perceive the blob's reactions, and if emotions and affects could play a central role in the user experience.

2 You can take a look at the project, access and run the P5.js code here: http://sebastianloewe.com/emotion-ai.

3 Bennett, Audrey, "Dynamic Interactive Aesthetics." *Journal of Design Research* 2.2 (2002).

4 Editors of Encyclopaedia Britannica, "Gestalt psychology", *Encyclopaedia Britannica*, accessed November 30, 2020, www.britannica.com/science/Gestalt-psychology.

A handful of users—which all have a design or AI background in some way—tested the prototype and generally found the blob's reactions aesthetically pleasing. They also liked the fact that the blob's reactions were not literal, and did not imitate any human behavior or emotions, instead drawing from its own set of reactions. The design prototype was deliberately designed to resemble a single-cell organism to prevent users from thinking they were interacting with a highly sophisticated, intelligent entity. By avoiding the pitfalls of simulating human-like intelligence, this prototype also served as a design probe into the realm of machine otherness and artificial counterparts, as the Siegen-based research group "Experience & Interaction Design" puts it[5]—and as Andreas Muxel and Elias Naphausen also tested it with their ROBODADA prototypes during the dai conference.

What I discovered during user testing was in a way rather odd, but very telling. When using the prototype, users gradually changed how they interacted with it. First, they simply had the blob react to them, usually with a certain delight and surprise that derived from the blob's appearance and behavior. But they quickly changed their behavior to a more deliberately exploratory demeanor, simulating 'emotional' reactions by imitating facial expressions and grimacing what they perceived as angry, surprised, or sad. By becoming actors in a sense, users controlled the prototype and forced its reactions, hindering the interactive prototype to organically and subtly adapt to an existing emotional state. Unsurprisingly, in using the design prototype this way, testers perceived the blob more as a representation of their simulated emotional reactions based on their facial expressions, much like a game character or an avatar, rather than an independent design artifact that reacts subtly to emotional input. It is also not surprising, that the use cases testers imagined ranged from gaming interfaces and social avatars to digital pets that soothe people's anxieties, to facilitation tools in group settings which could relieve emotional tension by playing and grimacing with the blob interface. One designer even suggested to open up the visual interface to a setting he's currently working in, namely design experiences in the shower. The blob could then be projected on a shower curtain, enriching and directing an olfactory, visual, and haptic shower experience. This is of course an interesting use case, but it goes in the opposite direction to what the initial idea of the prototype was. If we hold on to the idea that the users should not become actors or gamers to control an interface, but let the design outcome be controlled by their existing feelings and moods, we need to provide an emotion interaction paradigm that facilitates a subtle and helpful design experience. The underlying assertion here is: Allowing design artifacts to dynamically adapt to levels of stress and concentration, as well as a general well-being, could translate into a better and more helpful design experience when dealing with fatigue, stress, anxiety or anger.

Several fields of research emerge here, one is endemic to the field of machine learning, the other is native to design and aesthetics. Both fields are intrinsically connected by the field of psychology, particularly emotion and affect[6]. When searching for an emotion AI interaction design paradigm (or an emotion-

5 "Experience & Interaction Design Group", Home, Experience & Interaction Design Group, accessed November 26, 2020, www.experienceandinteraction.com.

6 Due to constraints derived from machine learning algorithms, in this paper I will concentrate on the basic emotions, as well as two indicators for affect, i.e. valence and arousal.

7 Norman, Donald, *Emotional Design: Why We Love (or Hate) Everyday Things* (New York: Basic Books, 2004).

8 Ho, Amic, "The New Relationship between Emotion and the Design Process for Designers", *Archives of Design Research* 27, no.2 (2014): 45–55. - Ho, Amic, "New way to optimise the design process from emotion perspectives", *Procedia Manufacturing* 3, (2015): 2325–2332.

9 Shaojing Fan, Zhiqi Shen, Ming Jiang et al., "Emotional Attention: A Study of Image Sentiment and Visual Attention", *Conference on Computer Vision and Pattern Recognition* CVPR, (2018): 7521-7531, doi: 10.1109/CVPR.2018.00785.

10 Cuykendall, Shannon, and Donald D. Hoffman, "From Color To Emotion Ideas and Explorations," Department of Cognitive Sciences at University of California Irvine, www.cogsci.uci.edu/~ddhoff/FromColorToEmotion.pdf.

11 Lu, Xin, Poonam Suryanarayan, Reginald Adams et al. "On Shape and the Computability of Emotions," *Proceedings of the 20th ACM international conference on Multimedia* (October 2012): 229–238. doi:10.1145/2393347.2393384.

12 Urquhart, Lewis, Andrew Wodehouse, "The Emotive Qualities of Pattern: Insights for Design", P*roceedings of the 21st International Conference on Engineering Design ICED17*, Vol. 8: Human Behaviour in Design, (2017): 109–118.

centered design paradigm) that is allocated within these three fields, we must survey the current state of debate and research in these academic disciplines. A full and comprehensive discussion of such a state of research easily exceeds the narrow boundaries of this paper and should therefore be abandoned in favor of a more focused perspective that concentrates on the following core research questions:

- What constitutes a good and helpful design that adapts dynamically based on emotions and affect?
- How should design adapt in order to meet users' emotional and affective needs?
- What are the principles and ethical guidelines of designing with emotion AI?
- How can emotion AI design prevent users from falling back into the role of an actor or gamer?
- How context-aware does this interactive visual design need to be in order to correctly meet users' needs?

2. THE PROBLEM OF EMOTION AI AND INTERACTION DESIGN

Answering these questions is more difficult than one might expect despite the fact that the interdependency between emotions and design has extensively been investigated. It is no secret that design in general is emotive[7] and divisions of scientists devoted their efforts to the topic, helping designers to better understand which part emotions play in design. In this way, designers' emotions influence their design choices[8], which then considerably shape the emotional state of users[9]. For example, the visual properties of an image cause emotions which determine the level of attention a particular image receives from a user. The visual properties of colors[10] stimulate users' emotions, in addition to shapes[11] and patterns[12]. Certain features and styles of typefaces, such as weight, serifs, or rounding correlate with emotional reactions[13]. When interacting with the design artifact, the users' emotions and affects alter and modify their perception of usability, sensuous pleasure, and other forms of delight of an interface or interactive design artifact[14]. Emotions do not take a back seat when interacting with design. On the contrary, so-called hedonistic qualities, such as aesthetic pleasure, surprise, and fun are equally valuable and important to users in their experience as pragmatic qualities, such as usability, simplicity, and comprehensibility[15]. Interestingly, these categories are not mutually exclusive, but work together. The best user experience is achieved when aesthetics and usability converge to a satisfactory, exciting, and positive experience. Anger and frustration, as well as negative arousal levels, increase when visual aesthetics and usability are perceived as low.[16]

13 Koch, Beth, "Emotion in Typographic Design:an Empirical Examination," *Visible Language* 46.3 (December 2012): 206-228. http://visiblelanguage.herokuapp.com/issue/155.

14 Sonderegger, Andreas, and Jürgen Sauer, "The Influence of Design Aesthetics in Usability Testing: Effects on User Performance and Perceived Usability." *Applied Ergonomics* 41, No. 3 (May 2010): 403–10. doi: 10.1016/j.apergo.2009.09.002.
- Jordan, Patrick, *Designing Pleasurable Products: An Introduction to the New Human Factors* (London and New York: Taylor & Francis, 2000).
- Walter, Aarron, *Designing for Emotion* (A Book Apart: New York, 2011).
- Diefenbach, Sarah, Eva Lenz, and Marc Hassenzahl, "An Interaction Vocabulary. Describing the How of Interaction," CHI ,13 *Extended Abstracts on Human Factors in Computing Systems (*April 2013): 607–612. doi.org/10.1145/2468356.2468463
- Zadra, Jonathan, Gerald Clore, "Emotion and perception: the role of affective information," *Wires Cognitive Science* 6, No. 2 (July 2011): 676-685. doi.org/10.1002/wcs.147.
- Anderson, Stephen, *Seductive Interaction Design: Creating Playful, Fun, and Effective User Experiences*, (Berkley: New Riders, 2011).

15 Hassenzahl Marc, Michael Burmester, and Franz Koller, "AttrakDiff: Ein Fragebogen zur Messung wahrgenommener hedonischer und pragmatischer Qualität," *Mensch & Computer. Berichte des German Chapter of the ACM*, Vol. 57 (2003). doi.org/10.1007/978-3-322-80058-9_19.

16 Mahlke, Sascha, Manfred Thüring: "Studying antecedents of emotional experiences in interactive contexts", CHI 2007 *Proceedings of the SIGCHI Conference on Human Factors in Computing Systems*, (April 2007): 917.

Consequently Paul Hekkert states:

> We experience the unity of sensuous delight, meaningful interpretation, and emotional involvement, and only in this unity we can speak of an experience.[17]

From an interactionist perspective, this is not surprising, since the paradigm states that both the design artifact and the user's perception and emotions constitute the user experience[18]. The guiding principles and key factors for stimulating emotions in interaction design are therefore "instrumentality, aesthetics, and symbolism."[19] In other words, pragmatic qualities, such as usefulness or usability, go together with hedonistic qualities, such as aesthetically pleasing design, as well as the symbolic representation of social statuses found in style elements that allow users to convey meaning related to their social identity and status.

These notions of what influences a positive user experience in interaction design have been further elaborated and updated by Jonas Löwgren. He states that genre expectations that have been formed prior to the actual interaction experience are key in interpreting the pragmatic and hedonistic qualities of an interaction.

> What is esthetically appropriate depends on what the use expects from the interaction experience, which is in turn colored by their initial appraisal of the product, its purpose, its use potentials – in short, its genre.[20]

Löwgren also further scrutinizes the concept of aesthetic delight when it comes to interaction design. He proposes four additional categories which significantly influence the users' aesthetic perception, i.e. pliability, rhythm, dramaturgical structure, and fluency[21]. Pliability describes the user's feeling of being invited to and in control of the interaction experience. Secondly, Löwgren identifies rhythm as an aesthetically relevant parameter through describing the way interactions are designed to adhere to everyday practices and their natural tempo and oscillation. Rhythm in interaction design can range from micro-structures in apps to use patterns that span across an entire day, week, or even year. By forming and controlling these patterns, designers may stimulate delight or disappointment. Third, Löwgren points out that the dramaturgical structure of interaction is relevant for aesthetic perception. This parameter adds the idea of growing tension and subsequent relief to the range of interaction. Knowing and acknowledging this opens up the possibility of modifying and using it to create a more delightful experience. The last parameter, fluency, can foster or mitigate a state of flow and thus the perceived easiness and delight of interactions and user experiences.

Now, when applying the above elaborated research to design and emotions things become blurry. The question of how exactly and to what end these findings and principles may be implemented, cannot be answered without complex research. There is no playbook or guideline for designing with emotion AI that can state exactly which color, pattern, size or typeface—let alone the

17 Hekkert, Paul, "Design aesthetics: principles of pleasure in design," *Psychology Science* 48, no. 2 (2006): 157-172.

18 Moshagen, Morten, Meinald Thielsch, "Facets of visual aesthetics," I*nternational Journal of Human-Computer Studies* 68, no. 10 (2010): 689-709, doi:10.1016/j.ijhcs.2010.05.006.

19 Tractinsky, Noam, Dror Zmiri, "Exploring attributes of skins as potential antecedents of emotion," in *Aesthetic Computing*, ed. Paul Fishwick (Cambridge: MIT Press, 2006), 405-422.

20 Löwgren, Jonas, "Toward an articulation of interaction esthetics", N*ew Review of Hypermedia and Multimedia* 15, no. 2 (December 2009): 130.

21 Löwgren, Jonas, "Toward an articulation", 130.

arrangement and motion of multiple design elements—designers should choose when users show specific emotions or reactions to facilitate a useful and delightful design experience. Some of the questions that derive from this lack of guidelines are:

- When do users get bored or annoyed and subsequently get angry?
- How are boredom, confusion or anger met adequately?
- Which parameters for interaction design are responsible herein?
- To what degree can these parameters then influence a positive emotional and affective reaction?
- What concrete design choices can help mitigate, transform, or stimulate these emotions or reactions?
- How does the layout of a website design for instance need to change when users are stressed or inattentive?

At the moment, answers to these questions continue to be a vast desideratum subject to further inquiry. One possible way these questions can be addressed preliminarily is to take a design prototype and investigate users' perceptions. It may be crucial to identify which emotions or affects are the most relevant and influential in terms of user experience. Let's say stress is a significant parameter that emotion-centered design should cover. A lab research would need to measure valence, arousal, heart rates, skin conductance, temperature, and pulse[22] to get baseline data, and then simulate a stressful situation for the user[23]. This alone is a highly complex research setting, not only depending on the theoretical foundation for measuring emotions, but often involving wearables and sensors. Additional data would also be needed to determine what design elements the user focusses on and interacts with. This could be collected through a variety of methods from mouse-tracking, or eye-tracking[24], or a think-aloud walk-through. These findings would then be enriched by users' genre expectations, pragmatic goals, as well as their aesthetic preferences conducted in interviews. Comparing these data points could then reveal a gap indicating which design element most significantly causes users with the same expectations, goals, and aesthetic preferences the most stress, and hence the fewest delights. In another step, changing the design elements and re-running the experiment could determine if the revised designs alleviate or elevate stress, fatigue and anger. In an iterative fashion the intended emotional and affective experience combined with the actual effect could then eventually be matched, allowing designers to review, refine, and re-design the prototype. This gradually would allow designers to discover hidden rules about how different designs are perceived with different emotions, and how to adapt design dynamically based on affective responses and user behaviors. If emotions shape the way people perceive design than these emotions may also play a pivotal role in determining new emotion-centered design principles. Ideally, these principles would guide designers in creating user interface elements that adapt colors, shapes, and

22 Picard, Rosalind: „Affective Computing", *M.I.T Media Laboratory Perceptual Computing Section Technical Report* No. 321, (1995): 11.

23 Liapis, Alexandros, Christos Katsanos, Dimitris Sotiropoulos et al. "Stress in interactive applications: analysis of the valence-arousal space based on physiological signals and self-reported data." *Multimedia Tools and Applications*, vol. 76 (June 2017): 5051–5071. doi.org/10.1007/s11042-016-3637-2.

24 Useful studies have been conducted in the field of human-computer interaction regarding dynamic emotion-based design pattern libraries. But they unfortunately do not reflect on hedonistic or aesthetic experiences and their dynamic parameters. See: Märtin, Christian, Sanim Rashid, and Christian Herdin: "Designing Responsive Interactive Applications by Emotion-Tracking and Pattern-Based Dynamic User Interface Adaptation," *Human-Computer Interaction. Novel User Experiences. HCI 2016. Lecture Notes in Computer Science*, vol. 9733 (2016): 28-36. https://doi.org/10.1007/978-3-319-39513-5_3

size in real-time. This way, the interactive personalized designs meet the users' overall goals and preferences, helping mitigate negative affects and lead to more successful, meaningful and joyful interactions. But clearly this is a highly complex and highly invasive research endeavor—requiring a well-funded and well-equipped research lab.

To make matters worse—even without testing emotion-centered design prototypes in the fashion outlined above—it becomes clear that every interaction principle applies differently to each specific use case. Depending on the design prototype, users may shift and re-evaluate their genre expectations, pragmatic goals, and the hedonistic qualities they consider crucial in a pleasing interaction experience. Therefore, designers should assess and adapt emotion AI interaction design principles to their specific use cases. For example, interaction design's pragmatic qualities are probably less in the users' focus of attention when they are using the blob prototype. They most likely do not have target-oriented usability goals in mind, like finding certain information, but probably want the design to give them some subtle initial information how to use it. Since the interface is very intuitive these goals might be achieved rather effortlessly. This way, the hedonistic qualities of the blob become more pivotal, making users then concentrate on the playful, entertaining interaction with the blob. In this regard, even before testing the prototype, designers need to assess and weigh expectations, preferences and interaction design qualities. Optimizing the users' experience may then focus on the aesthetic delight of the interaction, and hence the surprising and fun moments of changing the behavior of the blob, i.e. the color, shape, and movement. Specifically, for our use case the interaction design principle 'pliability' is no less interesting to consider, since the emotion AI interface does not allow for much agency or control. The blob simply reacts to certain facial expressions the designer crafted prior to the interaction, which may cause a feeling of constraint and hence a certain level of unpleasantness. The opposite feeling may apply for the principle of dramaturgical structure, when the blob creates new and interesting behaviors. The close-loop of cause and reaction that may lead to a feeling of constraint on the one hand, may stimulate pleasure on the other hand, because users might experience a sort of delight being kept in suspense for the blob's next action and their subsequent satisfaction to see how it reacts.

In contrast, when designing an emotionally responsive interface element on a website, the above discussed principles change in priority and significance to the user. Now, expectations may get more rigid and pragmatic goals, such as finding certain information, may become more important. Hedonistic goals may cease to be the front runner and now support the overall experience. Emotion AI data does not suffice to determine what the actual goals of users are.

To make matters even more complicated and to transition to the next chapter, emotions and affects are not universal in nature because they are significantly influenced by personal inference and context, making purely face-recognition-based emotion detection rather impossible.

25 Feldman Barrett, Lisa, Ralph Adolphs, Stacy Marsella, et al., "Emotional expressions reconsidered: Challenges to inferring emotion from human facial movements," *Psychological Science in the Public Interest* 20, (2019):1–68. doi:10.1177/1529100619832930.

26 Feldman Barrett, Lisa, "Can Machines Perceive Emotion?," filmed June 2018 at Talks at Google, video, 1:17:16, https://www.youtube.com/watch?v=HlJQXfL_GeM.

3. THE PROBLEM OF EMOTION AI AND MACHINE LEARNING

When we think of emotions, we usually think of preset, hardwired reactions that do not differ much from one person to another. The same goes for facial expressions that convey these emotions and affects. We tend to think they are universal, in the sense that scowling signifies a negative emotion, most likely anger and smiling represents happiness. Lisa Feldman Barrett, professor of psychology and director of the Interdisciplinary Affective Science Laboratory at the Northeastern University in Boston is challenging this notion—which has a huge effect on the assessment of how well and accurate emotion AI really works. Feldman's extensive research shows that emotions are not hard-wired brain reactions that are uncontrollable, but are made by the cognitive system equally based on evaluations of the body and the environment[25]. Emotions and affects are produced by the human brain to categorize sensations from the body and match them with social contexts and situations, altogether inferring an appropriate response. For instance, an elevated heart rate and mild aching of the stomach might lead to a very different emotion, depending on the context and person having the physical stimulus. In the setting of a romantic dinner date, it may be interpreted as nervous infatuation, in a setting of a hospital, it might be perceived as anxiety.

The same goes for facial expressions. There is apparently no conclusive evidence that facial expressions depict emotions universally, i.e. that a smile is always indicative of happiness, or a scowl uniquely refers to anger. In Feldman Barrett's view, facial expressions cannot be used by machine learning algorithms to accurately infer emotional states because they are not universally expressed and recognized. Even humans need context and situational awareness to guess emotions based on facial expressions. Why is it then, that most people associate anger with scowls and happiness with a smile? When asked to look happy, most people would probably want to smile to convey this emotion. Feldman Barrett and her team found out that psychologists back in the 1960s asked test subjects to imitate emotions. They then asked other participants to label the stereotypical representations of emotions with a preselected set of linguistic cues, leading an entire cohort of psychologists to use these stereotypes in their studies. This in return manufactured evidence of emotional states instead of discovering them, Feldman Barrett underlines[26]. Current emotion AI is trained on these stereotypes, which makes the outcomes highly debatable.

Studies claim that face recognition can be used not only to infer emotions, but also entire personalities and future behavior. For example, one paper claims that economic behavior can be predicted from facial expressions.[27] Another claims that an algorithm can detect whether or not someone is lying.[28] A third paper claims that leadership abilities can be inferred from the width of a person's mouth.[29] Even gender discrimination is said to be predicted from facial expressions.[30] The start-up Faception even claims that it can infer whether or not someone is a terrorist or pedophile simply by looking at a person's facial features.[31]

27 Rossi, Filippo, Ian Fasel, and Alan Sanfey, "Inscrutable games? Facial expressions predict economic behavior," *BMC Neuroscience* vol. 12, (2011): P281, doi.org/10.1186/1471-2202-12-S1-P281.

28 Sen, Taylan, Kamrul Hasan, Minh Tran, et al., "Say CHEESE: Common Human Emotional Expression Set Encoder and Its Application to Analyze Deceptive Communication," *13th IEEE International Conference on Automatic Face & Gesture Recognition* (2018): 357-364, doi: 10.1109/FG.2018.00058. - Sen Taylan, Kamrul Hasan, Zach Teicher, et al., "Automated Dyadic Data Recorder (ADDR) Framework and Analysis of Facial Cues in Deceptive Communication," *Proceedings of the ACM on Interactive, Mobile, Wearable and Ubiquitous Technologies* (2018), https://doi.org/10.1145/3161178.

29 Re, Daniel, and Nicholas Rule, "The big man has a big mouth: Mouth width correlates with perceived leadership ability and actual leadership performance," *Journal of Experimental Social Psychology* vol. 63 (March 2016): 86-93, doi.org/10.1016/J.JESP.2015.12.005.

30 Grant, Edward, Stephan Sahm, Mariam Zabihi, et al. "Predicting and visualizing psychological attributions with a deep neural network," 23rd International Conference on Pattern Recognition ICPR, (2016): 1–6, doi: 10.1109/ICPR.2016.7899598.

31 McFarland, Matt, "Terrorist or pedophile? This start-up says it can out secrets by analyzing faces," Washington Post, May 24, 2016, www.washingtonpost.com/news/innovations/wp/2016/05/24/terrorist-or-pedophile-this-start-up-says-it-can-out-secrets-by-analyzing-faces.

Gerhard Anger in his dai digital keynote address compared these findings to the pseudoscience of phrenology for a good reason.[32]

All of these claims would be strangely amusing, if this technology did not dare to influence many political, judicial, social, and economic decisions. In airports or in law this technology might soon be ubiquitous.[33] Even if people tend to superficially judge personal traits and character based on faces and appearance, the question here is: Do we want machines to do this too? To paraphrase Meredith Whittaker from the AI Now Institute, how can a machine not be dangerous when it fails, and how can it not be harmful when it works?

One attempt to gain more reliable outcomes for emotion detection comes from Google AI.[34] Researchers tried to identify within-class variations and between-class similarities of emotions within facial recognition. Classes are here basic emotions like anger, surprise or happiness. This way they wanted to gain more accuracy in labeling facial expressions, based on the idea that facial expressions which stem from the same emotion need to share more similarities than those from other emotions. Researchers at Google tried to circumvent the problem of simple facial action units that abstract from underlying complex emotive responses by assigning more than just the basic emotions as labels. They also let the images from the dataset be relabeled by six human raters who "were instructed to focus only on expressions ignoring other factors such as identity, gender, ethnicity, pose and age."[35] This reveals a major question mark in the rating process, since pose is an important communicative aspect of interpreting emotional states that are not exclusively conveyed with only facial expressions. Additionally, gender, cultural identity and age form categories that also influence how certain groups interpret emotions based on facial expressions. The underlying question here is, how significant are facial cues as a sole source of interpretation in relation to overall situation and behavior, body language, and other non-verbal communicative aspects that usually form an understanding of an opposite's emotional state.

Another suggestion to ensure a more comprehensive and reliable data set, exceeding simple face recognition classification, comes from Lisa Feldman Barrett herself. She proposes that for a machine to detect emotions, it has to be trained not on stereotypes of facial expressions, but on full, high-dimensional details for context, such as physical motion and chemical traits of the body, environments, and social interactions. The dilemma for a scientist or designer who wants to get the most accuracy with as little data as possible is obvious. Detecting emotions and affects to ensure the best user experience is—once again—a highly invasive endeavor.

32 Anger's keynote paper can be found in these proceedings of the dai digital.

33 Bridge, Mark, "Facial software knows if you have something to hide," The Times, May 24, 2018, www.thetimes.co.uk/article/facial-software-knows-if-you-have-something-to-hide-wsjps63r0.

34 Vemulapalli, Raviteja, and Aseem Agarwala, "A Compact Embedding for Facial Expression Similarity," IEEE/CVF Conference on Computer Vision and Pattern Recognition CVPR (June 2019): 5676-5685, doi: 10.1109/CVPR.2019.00583.

35 Vemulapalli, and Agarwala, "A Compact," 5685.

4. CONCLUSION

For an interaction design dynamically working with emotion AI, this means knowing what makes for a positive or negative

experience, and can only be accurately interpreted when data about users' motivations are aggregated. To just use facial recognition means not only guessing the emotions incorrectly, but also missing intention or motivation of an interaction entirely, as well as the actual cause for negative emotions and affects. After all, it becomes obvious, that "the social emotional games people play is probably the hardest computer science problem to solve in the century as we go forth,"[36] as Rosalind Picard puts it.

Does this mean that designers cannot use facial recognition to personalize visual design? The answer depends not only on designers' technical abilities and equipment. Designers must also be aware of the ethical pitfalls and shortcomings of this technology. One way to start this task without having a well-funded research lab[37] at your disposal is to come up with new ways to implement emotion AI allowing users to transparently customize and enrich their experience to an extend they deem suitable.

36 Picard, Rosalind, "Technology and Emotions," video filmed June 2011 at TEDxSF, San Francisco, CA, video, 19:02, www.youtube.com/watch?v=ujxriwApPP4&t=1007.

37 If you do have a lab like the MIT Media Lab to call your workplace, you may want to start something like Guillermo Bernal's project „Emotional Beasts". https://www.media.mit.edu/projects/emotionalbeasts-1/overview.

BIBLIOGRAPHY

Anderson, Stephen, S*eductive Interaction Design: Creating Playful, Fun, and Effective User Experiences*, (Berkley: New Riders, 2011).

Bennett, Audrey, "Dynamic Interactive Aesthetics." *Journal of Design Research* 2.2 (2002).

Bridge, Mark, "Facial software knows if you have something to hide," The Times, May 24, 2018, www.thetimes.co.uk/article/facial-software-knows-if-you-have-something-to-hide-wsjps63r0.

Cuykendall, Shannon, and Donald D. Hoffman, "From Color To Emotion Ideas and Explorations," Department of Cognitive Sciences at University of California Irvine, www.cogsci.uci.edu/~ddhoff/FromColorToEmotion.pdf.

Diefenbach, Sarah, Eva Lenz, and Marc Hassenzahl, "An Interaction Vocabulary. Describing the How of Interaction," *CHI ,13 Extended Abstracts on Human Factors in Computing Systems* (April 2013): 607–612. doi.org/10.1145/2468356.2468463

Editors of Encyclopaedia Britannica, "Gestalt psychology", *Encyclopaedia Britannica*, accessed November 30, 2020, www.britannica.com/science/Gestalt-psychology.

"Experience & Interaction Design Group", Home, Experience & Interaction Design Group, accessed November 26, 2020, www.experienceandinteraction.com.

Feldman Barrett, Lisa, "Can Machines Perceive Emotion?," filmed June 2018 at Talks at Google, video, 1:17:16, https://www.youtube.com/watch?v=HlJQXfL_GeM.

Feldman Barrett, Lisa, Ralph Adolphs, Stacy Marsella, et al., “Emotional expressions reconsidered: Challenges to inferring emotion from human facial movements,” *Psychological Science in the Public Interest* 20, (2019):1–68. doi:10.1177/1529100619832930.

Grant, Edward, Stephan Sahm, Mariam Zabihi, et al. “Predicting and visualizing psychological attributions with a deep neural network,” 23rd International Conference on Pattern Recognition ICPR, (2016): 1–6, doi: 10.1109/ICPR.2016.7899598.

Hassenzahl Marc, Michael Burmester, and Franz Koller, “AttrakDiff: Ein Fragebogen zur Messung wahrgenommener hedonischer und pragmatischer Qualität,” *Mensch & Computer. Berichte des German Chapter of the ACM*, Vol. 57 (2003). doi.org/10.1007/978-3-322-80058-9_19.

Hekkert, Paul, “Design aesthetics: principles of pleasure in design,” *Psychology Science* 48, no. 2 (2006): 157-172.

Ho, Amic, “New way to optimise the design process from emotion perspectives”, *Procedia Manufacturing* 3, (2015): 2325–2332.

Ho, Amic, “The New Relationship between Emotion and the Design Process for Designers”, *Archives of Design Research* 27, no.2 (2014): 45–55.

Jordan, Patrick, *Designing Pleasurable Products: An Introduction to the New Human Factors* (London and New York: Taylor & Francis, 2000).

Koch, Beth, “Emotion in Typographic Design: an Empirical Examination,” *Visible Language 46.3* (December 2012): 206-228. http://visiblelanguage.herokuapp.com/issue/155.

Liapis, Alexandros, Christos Katsanos, Dimitris Sotiropoulos et al. “Stress in interactive applications: analysis of the valence-arousal space based on physiological signals and self-reported data.” *Multimedia Tools and Applications*, vol. 76 (June 2017): 5051–5071. doi.org/10.1007/s11042-016-3637-2.

Löwgren, Jonas, “Toward an articulation of interaction esthetics”, *New Review of Hypermedia and Multimedia* 15, no. 2 (December 2009): 130.

Lu, Xin, Poonam Suryanarayan, Reginald Adams et al. “On Shape and the Computability of Emotions,” *Proceedings of the 20th ACM international conference on Multimedia* (October 2012): 229–238. doi:10.1145/2393347.2393384.

Mahlke, Sascha, Manfred Thüring: “Studying antecedents of emotional experiences in interactive contexts”, *CHI 2007 Proceedings of the SIGCHI Conference on Human Factors in Computing Systems*, (April 2007): 917.

Märtin, Christian, Sanim Rashid, and Christian Herdin: “Designing Responsive Interactive Applications by Emotion-Tracking and Pattern-Based Dynamic User Interface Adaptation,” *Human-Computer Interaction. Novel User Experiences. HCI 2016. Lecture Notes in Computer Science*, vol. 9733 (2016): 28-36. https://doi.org/10.1007/978-3-319-39513-5_3.

McFarland, Matt, “Terrorist or pedophile? This start-up says it can out secrets by analyzing faces,” Washington Post, May 24, 2016, www.washingtonpost.com/news/innovations/wp/2016/05/24/terrorist-or-pedophile-this-start-up-says-it-can-out-secrets-by-analyzing-faces.

Moshagen, Morten, Meinald Thielsch, “Facets of visual aesthetics,” *International Journal of Human-Computer Studies* 68, no. 10 (2010): 689-709, doi:10.1016/j.ijhcs.2010.05.006.

Norman, Donald, *Emotional Design: Why We Love (or Hate) Everyday Things* (New York: Basic Books, 2004).

Picard, Rosalind: „Affective Computing“, *M.I.T Media Laboratory Perceptual Computing Section Technical Report* No. 321, (1995): 11.

Picard, Rosalind, “Technology and Emotions,” video filmed June 2011 at TEDxSF, San Francisco, CA, video, 19:02, www.youtube.com/watch?v=ujxriwApPP4&t=1007.

Re, Daniel, and Nicholas Rule, “The big man has a big mouth: Mouth width correlates with perceived leadership ability and actual leadership performance,” J*ournal of Experimental Social Psychology* vol. 63 (March 2016): 86-93, doi.org/10.1016/J.JESP.2015.12.005.

Rossi, Filippo, Ian Fasel, and Alan Sanfey, “Inscrutable games? Facial expressions predict economic behavior,” *BMC Neuroscience* vol. 12, (2011): P281, doi.org/10.1186/1471-2202-12-S1-P281.

Sen Taylan, Kamrul Hasan, Zach Teicher, et al., “Automated Dyadic Data Recorder (ADDR) Framework and Analysis of Facial Cues in Deceptive Communication,” *Proceedings of the ACM on Interactive, Mobile, Wearable and Ubiquitous Technologies* (2018), https://doi.org/10.1145/3161178.

Sen, Taylan, Kamrul Hasan, Minh Tran, et al., “Say CHEESE: Common Human Emotional Expression Set Encoder and Its Application to Analyze Deceptive Communication,” *13th IEEE International Conference on Automatic Face & Gesture Recognition* (2018): 357-364, doi: 10.1109/FG.2018.00058.

Shaojing Fan, Zhiqi Shen, Ming Jiang et al., “Emotional Attention: A Study of Image Sentiment and Visual Attention”, *Conference on Computer Vision and Pattern Recognition* CVPR, (2018): 7521-7531, doi: 10.1109/CVPR.2018.00785.

Somers, Meredith, “Emotion AI, explained”, March 8, 2019, https://mitsloan.mit.edu/ideas-made-to-matter/emotion-ai-explained.

Sonderegger, Andreas, and Jürgen Sauer, “The Influence of Design Aesthetics in Usability Testing: Effects on User Performance and Perceived Usability.” *Applied Ergonomics 41*, No. 3 (May 2010): 403–10. doi: 10.1016/j.apergo.2009.09.002.

Tractinsky, Noam, Dror Zmiri, “Exploring attributes of skins as potential antecedents of emotion,” in *Aesthetic Computing*, ed. Paul Fishwick (Cambridge: MIT Press, 2006), 405-422.

Urquhart, Lewis, Andrew Wodehouse, “The Emotive Qualities of Pattern: Insights for Design”, *Proceedings of the 21st International Conference on Engineering Design ICED17*, Vol. 8: Human Behaviour in Design, (2017): 109–118.

Vemulapalli, Raviteja, and Aseem Agarwala, “A Compact Embedding for Facial Expression Similarity,” Conference on Computer Vision and Pattern Recognition CVPR (June 2019): 5676-5685, doi: 10.1109/CVPR.2019.00583.

Walter, Aarron, *Designing for Emotion* (A Book Apart: New York, 2011).

Zadra, Jonathan, Gerald Clore, “Emotion and perception: the role of affective information,” *Wires Cognitive Science 6*, No. 2 (July 2011): 676-685. doi.org/10.1002/wcs.147.

dai digial conference branding
Generative iteration of the dai logo

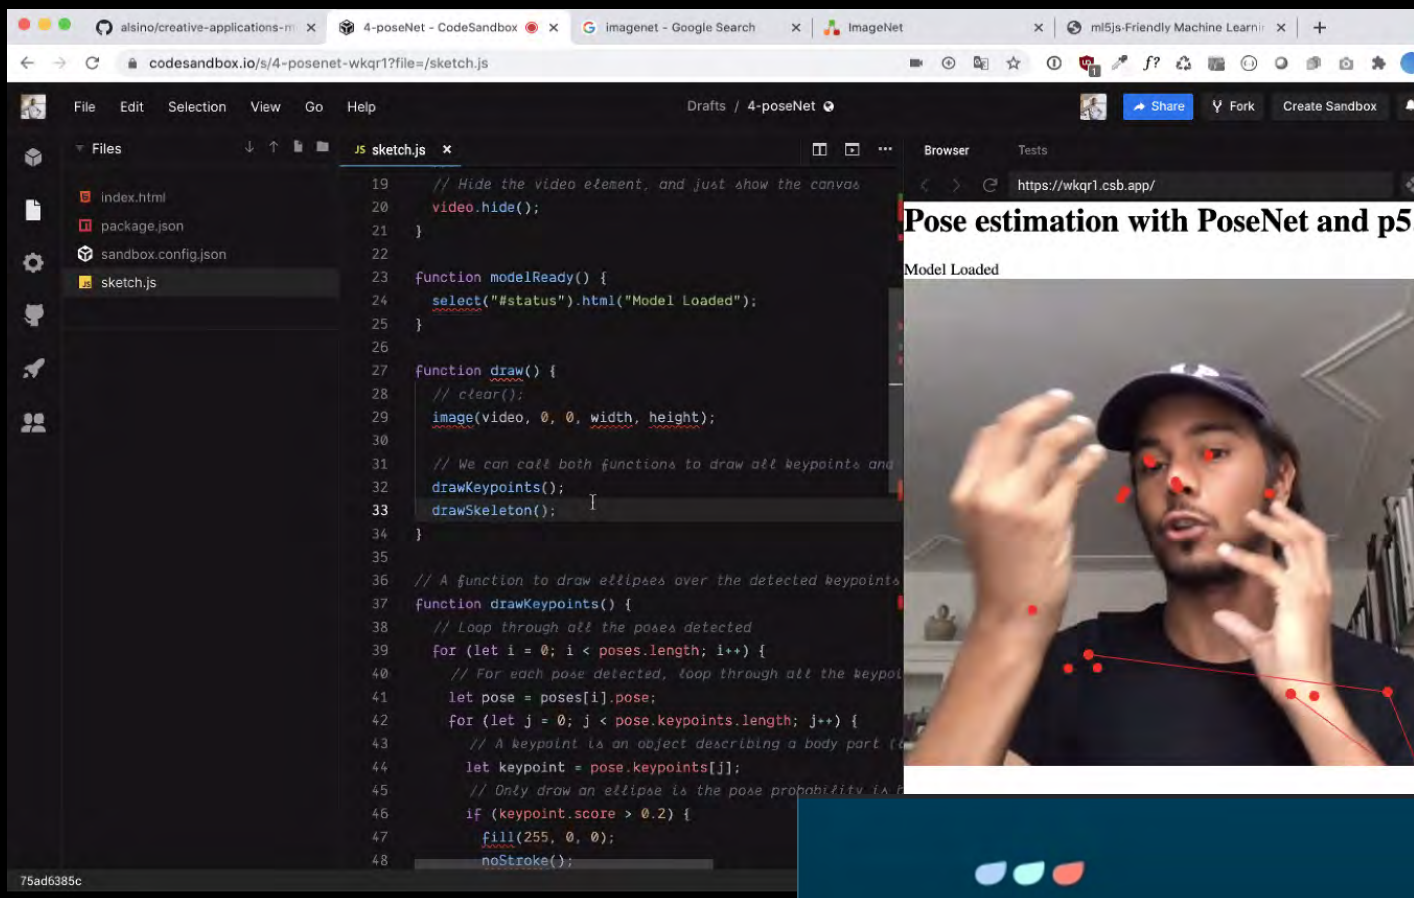
codesandbox.io/s/4-posenet-wkqr1?file=/sketch.js
File Edit Selection View Go Help
Drafts / 4-poseNet
Share Fork Create Sandbox
Files
index.html
package.json
sandbox.config.json
sketch.js
// Hide the video element, and just show the canvas
video.hide();
function modelReady() {
select("#status").html("Model Loaded");
function draw() {
// clear();
image(video, 0, 0, width, height);
// We can call both functions to draw all keypoints and
drawKeypoints();
drawSkeleton();
// A function to draw ellipses over the detected keypoints
function drawKeypoints() {
// Loop through all the poses detected
for (let i = 0; i < poses.length; i++) {
let pose = poses[i].pose;
for (let j = 0; j < pose.keypoints.length; j++) {
let keypoint = pose.keypoints[j];
if (keypoint.score > 0.2) {
fill(255, 0, 0);
noStroke();
75ad6385c
Browser Tests
https://wkqr1.csb.app/
Pose estimation with PoseNet and p5
Model Loaded

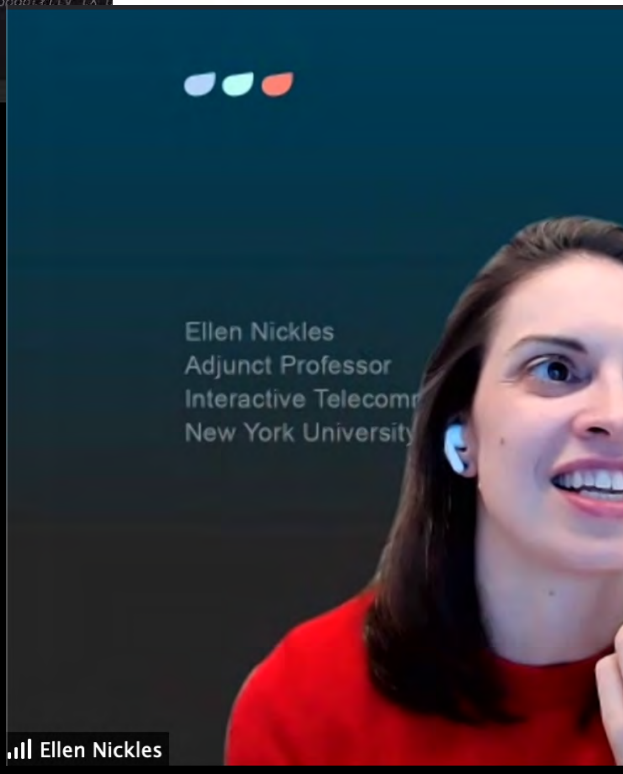
Ellen Nickles
Adjunct Professor
Interactive Telecom
New York University
Ellen Nickles

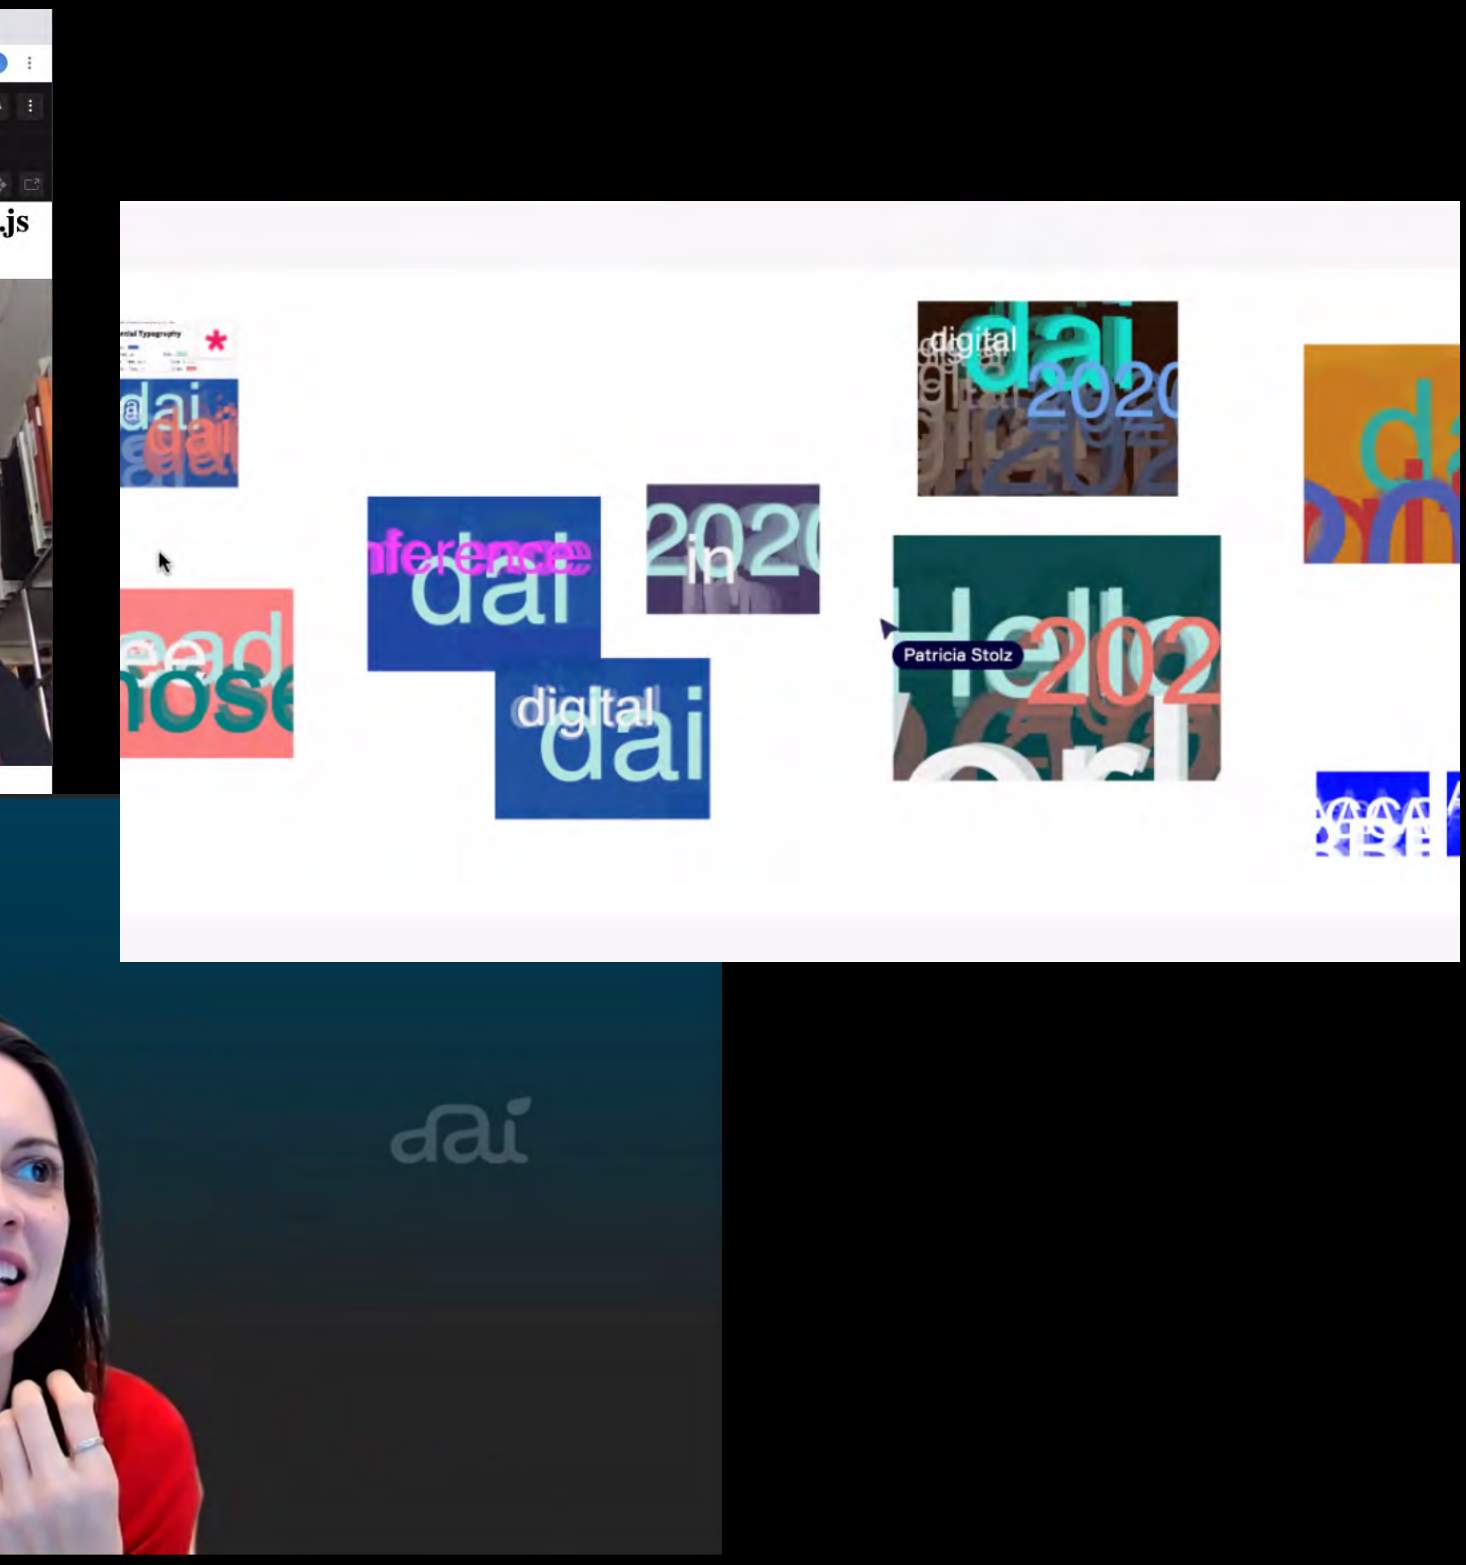

Hands-on workshops with Ellen Nickles and Alsino Skowronnek
September 17 and 20, 2020

The Design Process of the dai Brand and Logo

Marc Engenhart

Practice

ABSTRACT

This paper will present the first initial steps and methods used to design the dai conference logo ligature. The logo was explicitly developed in combination with machine learning assistance and forms a basis for the current discussion of human-machine co-creation processes in the field of design. The paper aims at making this process transparent and comprehensible by documenting steps in the design process and implementation. Machine learning (ML) overall offers designers extended possibilities and presents a basis on which they can apply their function as a specialist and curator for detail quality. It will be interesting to see how ML modules influence and expand the designer's work and enable the creation of the final design.

1. INTRODUCTION

Workflow (additional details)

The unwieldy title of the conference is reduced in the first step to the initials "dai." This creates a visual similar to an article of speech like "the" or "and." Utilizing this article abstraction is internationally effective because it is quickly graspable, sonorous, and accessible multi-lingually, while also serving as a reference to the study of the field of so-called AI (Artificial Intelligence).

Traditional design

The following description outlines the procedure.

Instruction:
Elaboration by the designer

Figure 1: Utilizing the article abstraction (traditional design)

We choose Laurenz Brunner's Circular Pro as the corporate font for the communication of the dai. The Circular Pro is a geometric sans serif inspired by Frutiger Avenir and Renner's Futura, which envisions a characteristically fittingly constructed geometric character. Released in 2013, it accompanies the technological development of digital surfaces as a contemporary typeface.

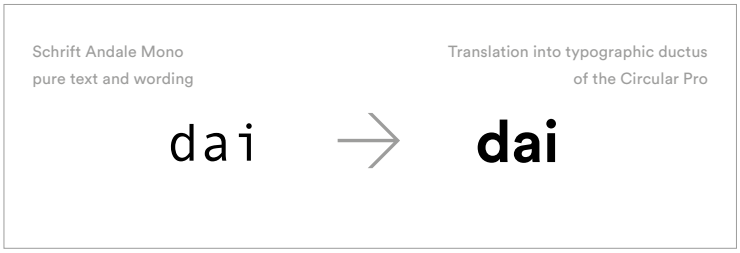

Figure 2: Typographic decision (traditional design)

The basic form serves as the starting point for form-finding and is released, set and adopted for further development as a logo basis.

2. IMPLEMENTATION (DESIGN)

Maschine assistance

The following description outlines the procedure.

Instruction:
Segmentation of the available image information into individual areas

Method:
Classification of the sub-areas (geometric form)

Keywords:
Pattern recognition

Model:
Multilayer Neural Network

Output:
Analysis and output of classified data (labels)

The three lowercase letters “d”, “a” and “i” of the Circular Pro Bold are stored in a design program and converted into image data (point clouds, pixels). The image data of the characters are examined by a neural network trained in image and pattern recognition in letter forms where similarities or areas that have high visual conciseness appear in the typographic form. The visual conciseness is defined by similarities in shape or pattern and the strength of the path contours, tangents, and inner or outer angles. The most similar areas are visually marked with a label (see figure below: A, B and C).

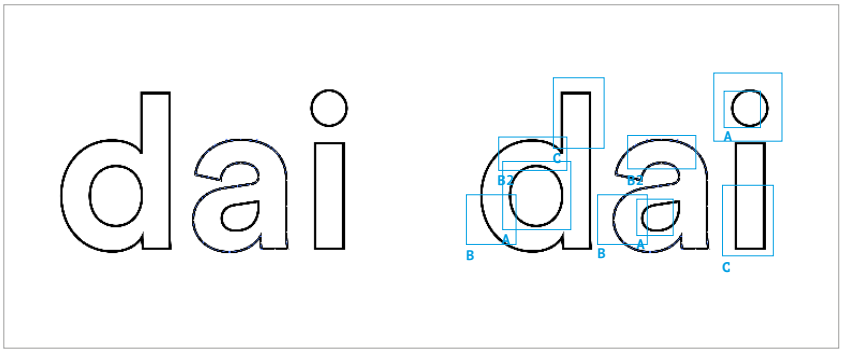

Figure 3: Segmentation of image information (characteristic style)

After checking, the system outputs first the hallmarks (inner shapes of the letters “d” and “a”). Then it outputs the i-point followed by the outer contour of the almost identical bellies of “d” and “a”. Finally, the formally separated now distinguishable stem of the small “i” is rendered. This process ensures the classification of consineness. From less concise parts to highly concise parts.

Instruction:
Varying the font contour by an »intelligent agent«

Keywords:
Generative design, agents, SVG, morphing

Model/Software:
P5.js/Processing, SVG morphing, JavaScript

Output:
Interlaced and generatively generated new image data as shape variations

The dai form element can now be modulated into these visually concise parts of the letter forms, which creates opportunities to generate multiple variations. This process allows letter forms to be used, combined, or abstracted in new ways. To do this, the calculating process may fill in areas, duplicate similarities, match or delete redundant particles, and close transitions between similar shapes.

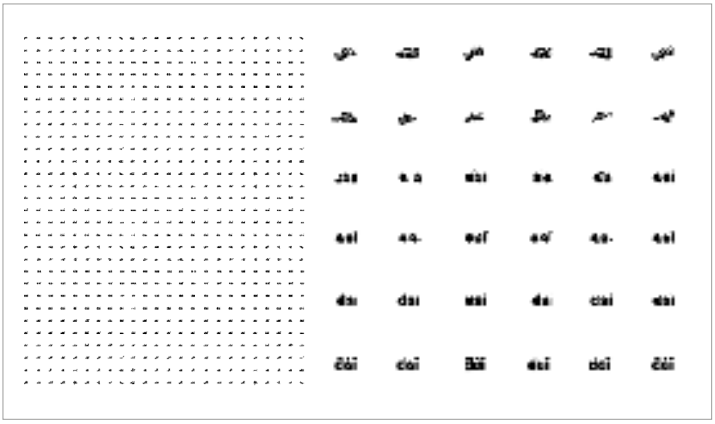

Figure 4: Varying the font contour by an ‘intelligent agent’

After the output step of the process, the generative system produces as much results as we want. These results are compared again with the original typographic form as a so-called ‘fixed absolute’ to determine the readability of the original visual of the dai Circular Pro. Finally, the 100 most similar finalists are presented by the system.

Figure 5: Selection of form studies from variations

Designers now check and select the most powerful designs. They are looking for form studies (algorithmic-driven generative variations) that best align with the conference’s ‘brand’, subsequently identifying its awareness, as well as the aesthetic effectiveness and functionality of the logo design draft.

The machine-learning-based step of the process produces a translatable analogy of form that eventually is convincing. This variation draws interest, is readable, concise and aesthetically pleasing.

Maschine assistance

The following description outlines the procedure.

Instruction:
Re-segmentation of image information into individual areas

Method:
Classification of the sub-areas

Keywords:
Pattern recognition

Model/Software:
Multilayer neural Network

Output:
Analysis and output of similarity to the original image set in Circular Pro

Traditional design

Instruction:
Elaboration by the designer

3. FINAL ARTWORK (AESTETIC, SEMANTIC PROPERTIES)

In an author's program, the designer draws in and aesthetically refines the proportions and contour accuracy.

The dai ligature connects the three individual meanings (design, artificial, intelligence) in a fluid cycle and repeats familiar patterns of the figure's form. The hallmarks are replaced by the repetition which replaces the i-point and is intensified by the number of the three hallmark particles created.

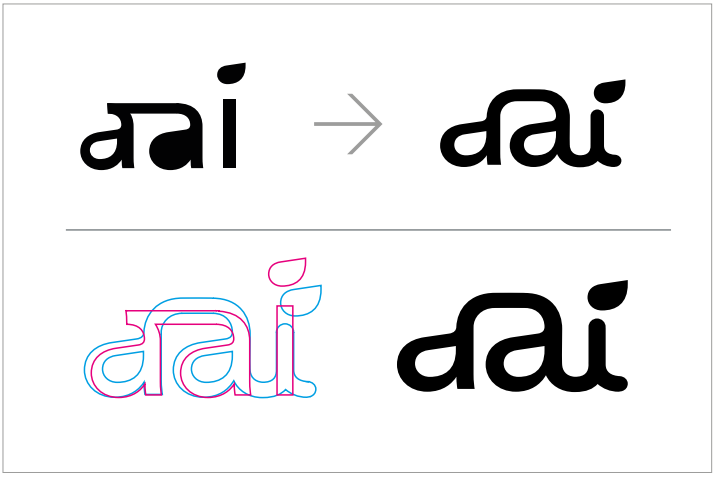

Figure 6: Refining the proportions and contour accuracy

4. CONCLUSION

The dai ligature logo is the result of a cooperation between designer and machine where the iterations are provided by technology, but the final final artwork is still completed by the designer. In this case, the machine assistance operates with a very simple, clearly defined method. It can find typographic similarities in a typeface and can then classify these depending upon readability. On this basis it can modulate designs generatively.

We, as designers, check the machine-generated variations for semiotic and content-related effectiveness, as well as aesthetic quality. This process of creative revision can be described as a curatorial process, which finalizes, completes and expands the logo.

The resulting dai brand logo design conveys a two-dimensional common form familiar to us. When combined, the elements of this process, such as pattern recognition, relevance, selection and sorting, summarizing, combining, form context, form variation, work in conjuction with the original handmade human typographic art to create a brand appropriate for a conference designed to address issues related to machine learning and artificially intelligent assistance systems.

5. SUGGESTED READING AND SOFTWARE LIST

js.org, “JavaScript program language”, accessed November 25, 2020, www.js.org.

A. Ledoux, N. B. Dao, P. Lahmi and M. Coustaty, “Alphanumeric Glyphs Transformation Based on Shape Morphing: Context of Text,” *2019 Eighth International Conference on Emerging Security Technologies*, (2019):1–6. doi.org/10.1109/EST.2019.8806219.

processing.org, “Flexible software sketchbook and a language for learning how to code within the context of the visual arts”, accessed November 28, 2020, www.processing.org.

P5js.org. “JavaScript library for creative coding, with a focus on making coding accessible and inclusive for artists, designers, educators, beginners”, accessed November 28, 2020, www.p5js.org.

Liu H., Gao L., Liu X. “Generative Design in an Agent Based Collaborative Design System.”, In Shen W., Lin Z., Barthès JP.A., Li T. (eds) *Computer Supported Cooperative Work in Design* I. CSCWD 2004. Lecture Notes in Computer Science, vol 3168, (2005):106–135. doi.org/10.1007/11568421_11.

dai digital conference
Impressions

dai digital conference brand and logo
September 2020

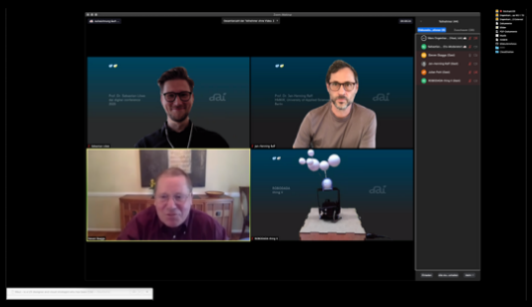

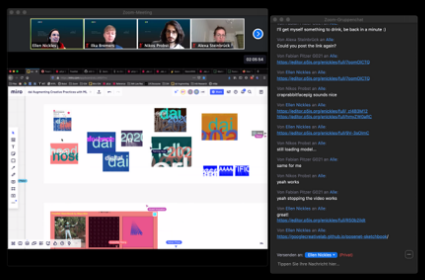

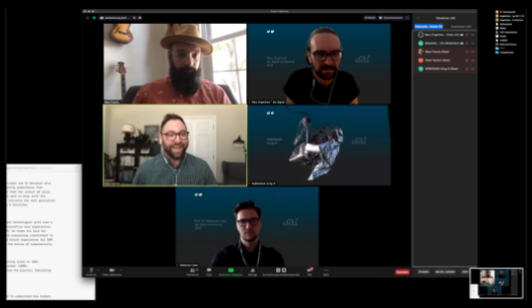

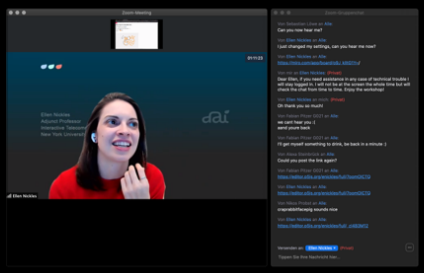

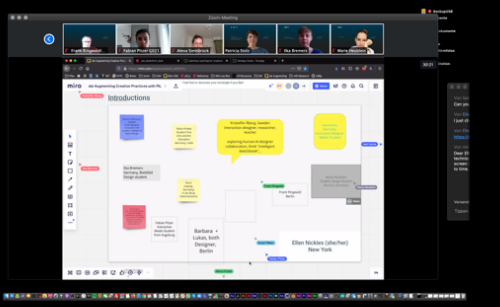

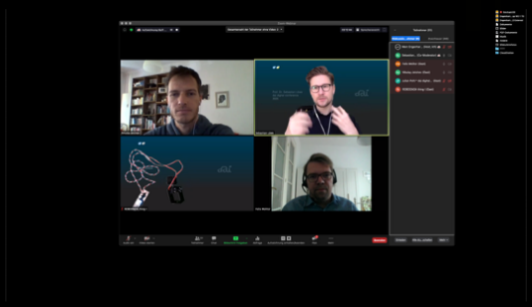

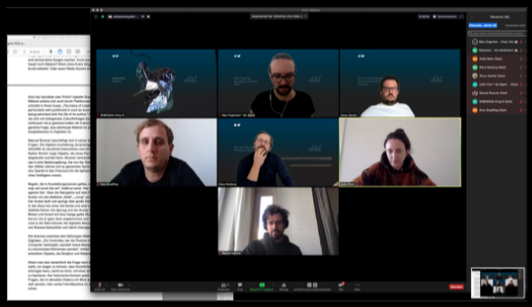

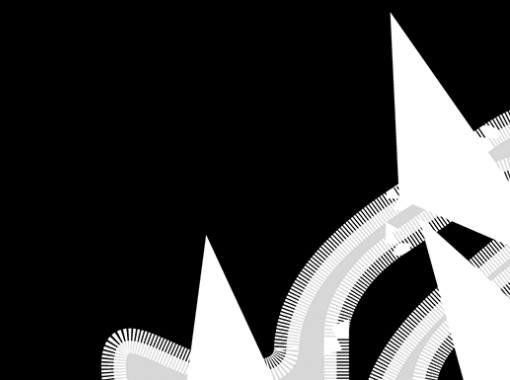

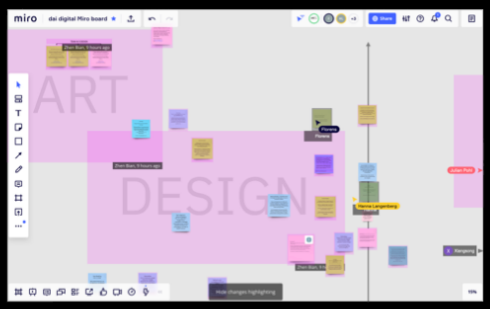
miro
dai digital Miro board
ART
DESIGN

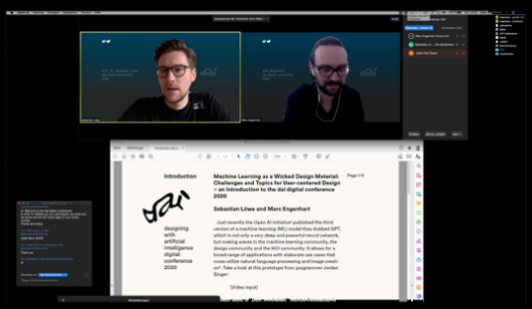
Introduction
Machine Learning as a Wicked Design Material: Challenges and Topics for User-centered Design – an Introduction to the dai digital conference 2020
Sebastian Löwe and Marc Engenhart
designing with artificial intelligence digital conference 2020

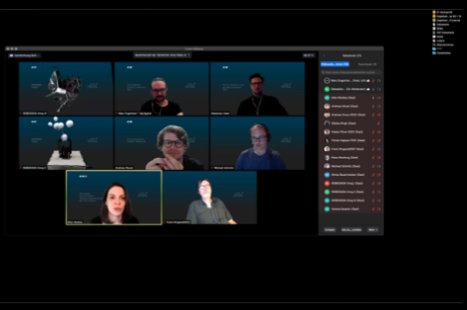

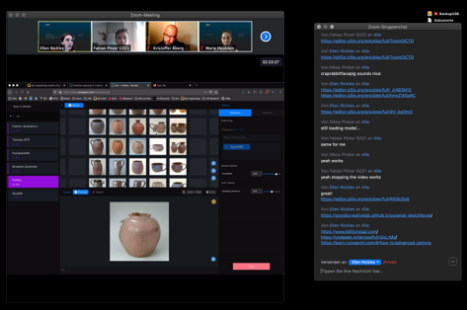

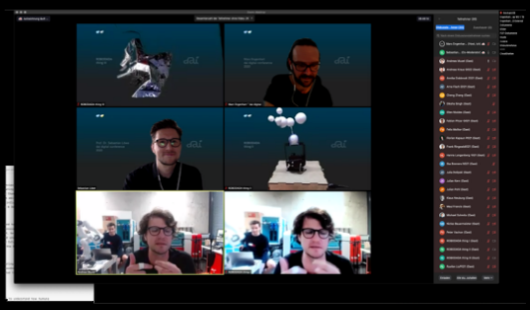

setup done
dai conference 2020 / ROBOCADA

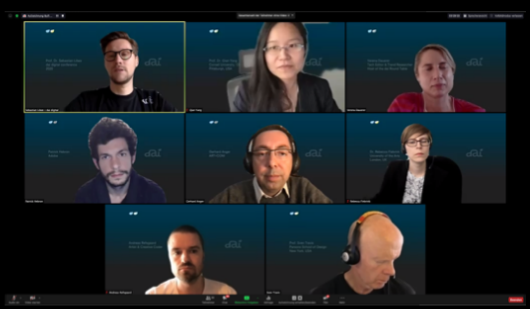

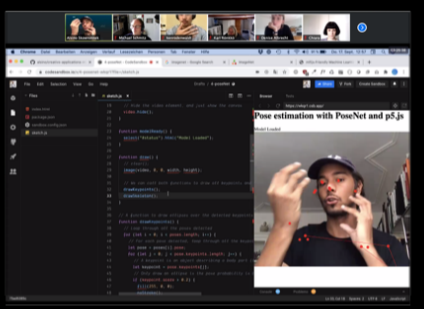
Pose estimation with PoseNet and p5.js

Proceedings' contributors

Gerhard Anger was born in Munich in 1975. He studied medieval history, history of science, logic, philosophy of science, and media computer science. He started working as a software developer and project manager for a multinational corporation in 2000. Later, he was the managing director of a startup in Berlin. He co-founded several companies, one of them focusing on AI, he co-owns a small publishing company and, as a pianist, regularly makes chamber music. In his capacity as an expert for democracy in the digital age he has been invited to be a speaker at several international conferences by organizations like the OSCE and the National Democratic Institute. He currently serves as the CTO of ART+COM.

Marc Engenhart is a multidisciplinary designer, musician and artist. In addition to running his own design studio for over ten years. He also teaches, researches, publishes, and consults on the topics of transmedia, interaction design and intelligent systems in the fields of speculative design, scenography, future affairs, prototyping as well as traditional communication design. He is a member of the German Association of Artificial Intelligence (Bundesverband KI), the type directors club in New York (tdc), The German Society for Design Theory and Research (DGTF) and the Germany Association of Communication Designers (Bund Deutscher Kommunikationsdesigner).
He is one of the co-founders of the dai digital conference.
https://engenhart.com

Dr Rebecca Fiebrink is a Reader at the Creative Computing Institute at University of the Arts London. Her teaching and research focus largely on exploring how machine learning and artificial intelligence can change human creative practices. Fiebrink is the developer of the Wekinator creative machine learning software, which is used around the world by musicians, artists, game designers, and educators. She is the creator of the world's first online class about machine learning for music and art. Much of her work is driven by a belief in the importance of inclusion, participation, and accessibility: she works frequently with human-centred and participatory design processes, and she is currently working on projects related to creating new accessible technologies with people with disabilities, and designing inclusive machine learning curricula and tools. Dr. Fiebrink previously taught at Goldsmiths University of London and Princeton University, and she has worked with companies including Microsoft, Smule, and Imagine Research. She holds a PhD in Computer Science from Princeton University.

Patrick Hebron is a designer, software developer, teacher and author. His work explores the intersection of machine learning, design tools, programming languages and operating systems and has specifically focused on the development of machine-learning-enhanced digital design tools. He founded and leads the Machine Intelligence Design team at Adobe, focusing on the design of new products, interfaces and interaction mechanisms built on AI/ML capabilities. He is the author of "Machine Learning for Designers," published by O'Reilly Media as well as articles including "Rethinking Design Tools in the Age of Machine Learning" and "A Unified Tool for the Education of Humans and Machines."

Proceedings' contributors

Jennifer Heier is an industrial designer at heart who is specialized in user and customer research. She studied Industrial Design in Darmstadt, Germany (Diploma), did her Bachelor of Arts in San Francisco and her Master of Science in user research in Edinburgh. She loves to work in a diverse and cross-disciplinary team, and close to new technology topics contributing to her passion for human-centric innovation and continuously questioning the status quo. She is currently working at Siemens' digital industries data lab as Head of UX-driven AI, bringing the human factor into the B2B context. Besides, she is doing her PhD at the Bauhaus University in Weimar investigating how the creative community should and could positively influence AI developments. She believes in the power of her profession to contribute and add value to the challenges, as well as opportunities, of the digital age.

Sebastian Loewe is a professor for design management at the Mediadesign University of Applied Sciences in Berlin, who holds a B.A. in media studies, a diploma in media arts and received his PhD from the University of Halle-Wittenberg, Germany.
He teaches, researches, publishes, and sometimes consults on the topics of strategic design and innovation with a focus on machine learning in design processes, on open innovation and co-creation, on design thinking, as well as on aesthetic theory and digital culture. He is one of the co-founders of the dai digital conference. www.sebastianloewe.com

Anika Meier studied art history and German literature in Heidelberg and Paris. She is an independent writer and curator. She writes a column on art and social media for the magazine Monopol, and her texts have appeared in Die Zeit, Frankfurter Allgemeine Zeitung, Die Welt, Spiegel Online, Kunstforum, Numéro Berlin, der Freitag, Tagesspiegel, Mindstate Malibu, among others. Her most recent exhibitions include "Link in Bio. Art after Social Media" (2019/2020) and "Virtual Normality. Women Net Artists 2.0" (2018) at the Museum of Applied Arts in Leipzig, "On the Internet Nobody Knows You Are a Performance Artist. Andy Kassier & Signe Pierce live" at the NRW-Forum in Düsseldorf (2018), and "Signe Pierce. Reflexxxions" (2019) at the Eigen + Art Lab in Berlin. The exhibition catalog "Virtual Normality: The Female Gaze in the Age of the Internet", which she co-edited with Alfred Weidinger, appeared with the Verlag für moderne Kunst. She is the founder of the collective This Ain't Art School.

Andreas Muxel is Professor for Physical Human-Machine Interfaces at the University of Applied Sciences Augsburg where he founded and directs the HYBRID THINGS LAB. In his work he is always looking for a poetical and engaging way of interaction with things, regardless of whether they are hardware or software. He creates hybrid artifacts and spaces for all senses blurring the boundaries between the rational and the irrational. His projects have been shown internationally (e.g. Ars Electronica Festival Linz, FILE Festival São Paulo, TodaysArt Festival Brussels) and awarded (e.g. Honorary Mention Prix Ars Electronica, VIDA Award, Share Prize). From 2013 to 2017 he was Professor for Interface / Interaction Design at the Köln International School of Design, University of Applied Sciences Cologne.
https://hybridthings.hs-augsburg.de, https://www.andreasmuxel.com/teaching

Proceedings' contributors

Elias Naphausen studied interactive media (B.A.) and interactive media systems (M.A.) at the University of Applied Sciences Augsburg. Additionally to his work as research associate at the HYBRID THINGS LAB he is creating experience-rich media installations, worlds of sounds and virtual rooms. Naphausen is exploring associations between mind, eyes, ears, haptic and olfaction and has shown his work on different locations such as Chaos Computer Congress and Asche zu Farbgut Festival. Elias just started his Ph.D. in collaboration with the Bauhaus University Weimar to the topic of sound design for collaborative robots. The project aims at designing tools for acoustic design of non-anthropomorphic, collaborative robots. http://www.naphausen.info.

Simon Nestler holds a doctorate in computer science and is professor for Human-Computer Interaction at the Technische Hochschule Ingolstadt. He is a member of the Association for Computing Machinery (ACM), the ACM SIGCHI and the German UPA. He is also a member of the Presidium of the German Informatics Society (GI) and a member of the Executive Committee of the Human-Computer Interaction Department of the GI. https://u-ux.de/

Klaus Neuburg is a designer, publicist and professor for media design with a focus on interactive media at the Ostfalia University of Applied Sciences Braunschweig/Wolfenbüttel. His work and research sits at the intersections between design, society and digitality, paying specific attention to social contexts. He is co-founder of the journalistic NGO Froh and a member of the German UPA and the DGTF. https://www.klausneuburg.de/

Sven Quadflieg is a designer and professor at the Hamm-Lippstadt University of Applied Sciences. He studied design at the Folkwang University of the Arts and the Zurich University of the Arts and earned his doctorate at the HFBK Hamburg. In his research he is interested in political and social design and the mutual influences and dependencies between design and society. He is a member of the DGTF and the International Gender Design Network. https://svenquadflieg.de/

Andreas Refsgaard is an artist and creative coder based in Copenhagen. He uses algorithms and machine learning to make unconventional connections between inputs and outputs, which allows people to play music using only eye movement and facial gestures, control games by making silly sounds, transform drawings of musical instruments on paper into real compositions and train algorithms to decide what is funny, funky or boring. His works have been published in New Scientist, Vice, Gizmodo, PSFK and Designboom, awarded by Interaction Awards and Core77 and exhibited in museums and at festivals in Europe, North America, Asia and Australia.

Manuel Rossner studied art at the Hochschule für Gestaltung Offenbach, the École des Arts-Décoratifs Paris, and the Tongji College for Design and Innovation in Shanghai. In 2012 Rossner founded the virtual Float Gallery, an Internet platform for digital art. In April 2020 König Galerie Berlin opened Manuel Rossner's digital solo exhibition "Surprisingly this rather works",

curated by Anika Meier and Johann König. In 2017 he designed a digital extension for the NRW-Forum Düsseldorf and curated together with Alain Bieber the VR exhibition "Unreal," which included works by Tabita Rezaire, Banz & Bowinkel, and Theo Triantafyllidis, among others. In 2019 he designed the virtual gallery building "CUBE" for Roehrs & Boetsch in Zurich. In this space, he showed his piece "Malibu" as part of the exhibition "Virtual Natives – Sculpture," which also included works by Martina Menegon and Chiara Passa. In October 2019, his solo exhibition "There's no feeling where there's no pain" was on view at the Museum der bildenden Künste Leipzig. Works of art by Rossner were exhibited in "Perception is Reality" in the Frankfurter Kunstverein (2017), in "Artificial Paradise?" in KM— Künstlerhaus Graz (2018), and in the 1822-Forum Frankfurt (2017). Rossner is a member of Kollektiv Internet TBD, which investigates the effects of the Internet on society. In 2018, he co-organized the conference "The Post-Binary" on artificial intelligence in art and design at the Museum Angewandte Kunst Frankfurt/M.

Dr. Jan-Henning Raff is professor for visual communication at HMKW, University of Applied Sciences in Berlin. He teaches and researches user centered (interaction) design and develops design research methodologies.

Qian Yang is an assistant professor in information science. As a human-AI interaction researcher, she helps translate AI's algorithmic advances into valuable real-world applications that serve human ends. Yang has so far designed several high-consequence AI applications: from decision support systems for life-and-death healthcare decisions (i.e. artificial heart implants and cancer diagnoses) to context-aware mobile services, from Natural Language Generation systems to autonomous cars. Building upon this related vein of practice, she works to inform a basic understanding of AI as a material for HCI design. In addition, she innovates methods and tools for the UX design and HCI practicing communities, helping them to better integrate AI into their day-to-day practices. Yang completed her PhD in human-computer Interaction at Carnegie Mellon's School of Computer Science. She holds master's degrees in HCI from Carnegie Mellon University and in Design and Bachelor of Engineering in Industrial Design from Shanghai Jiaotong University. Before CMU, she worked in design consultancy and in the industry (2007-2014).

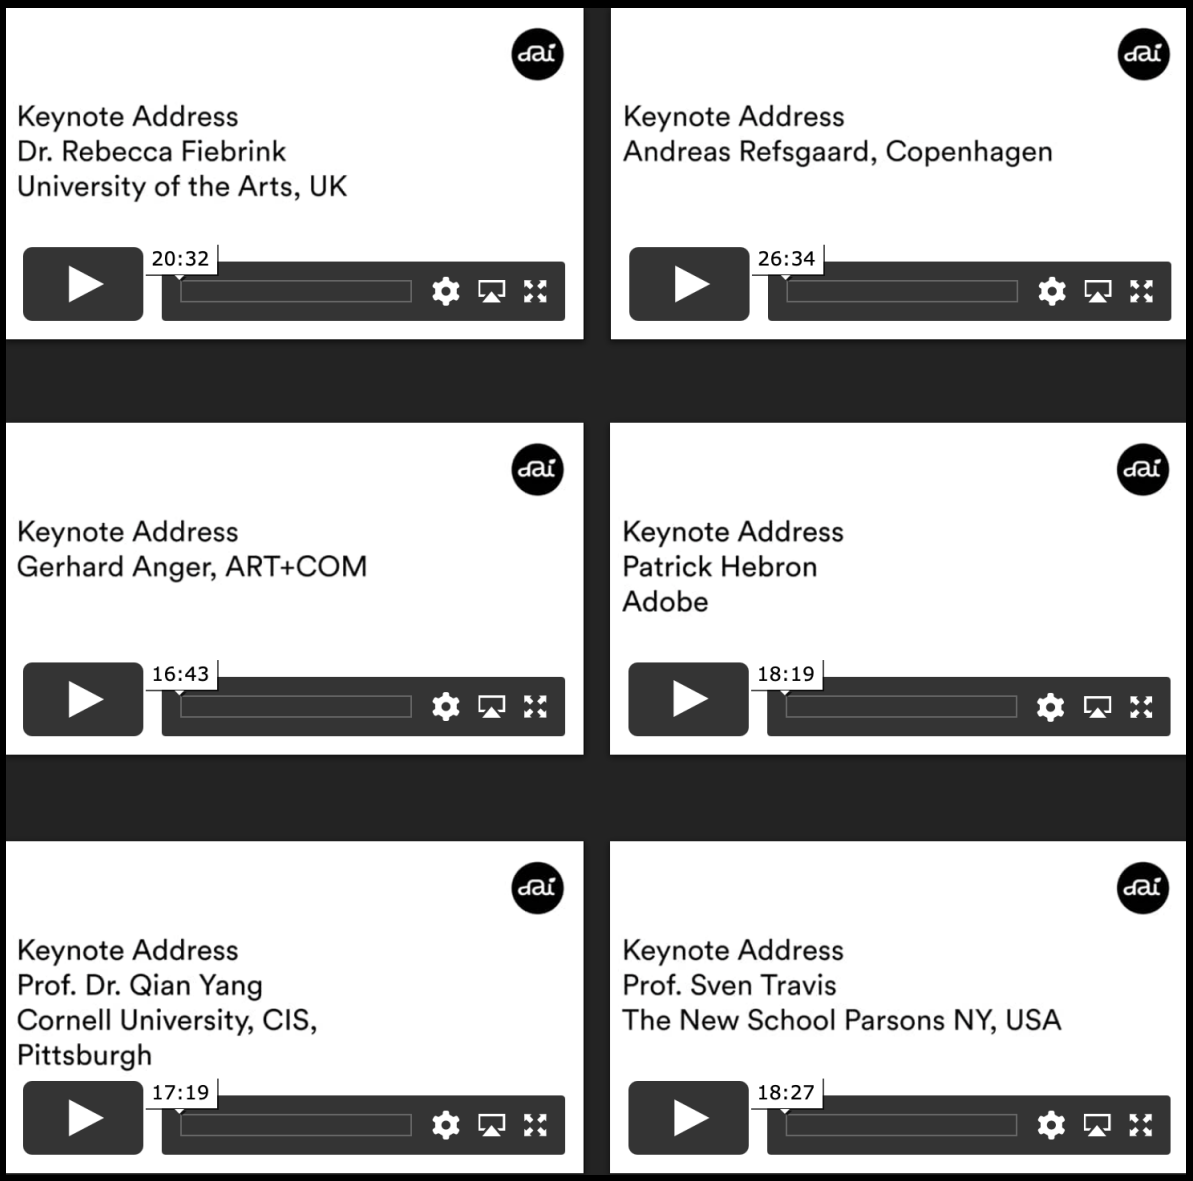
Keynote Address
Dr. Rebecca Fiebrink
University of the Arts, UK
20:32
Keynote Address
Andreas Refsgaard, Copenhagen
26:34
Keynote Address
Gerhard Anger, ART+COM
16:43
Keynote Address
Patrick Hebron
Adobe
18:19
Keynote Address
Prof. Dr. Qian Yang
Cornell University, CIS,
Pittsburgh
17:19
Keynote Address
Prof. Sven Travis
The New School Parsons NY, USA
18:27

Documentation and resources

dai digital conference 2020 keynote videos
https://www.designing-artificial-intelligence.eu

dai digital conference 2020 poster designs
https://4d821726-6938-496f-81b2-187b5fd9b5a7.filesusr.com/archives/0a280a_711ca68fd05d4416b1f4960de03189cb.zip?dn=dai-digital-poster-official-cc-license.zip

dai digital conference 2020 laptop sticker template
https://4d821726-6938-496f-81b2-187b5fd9b5a7.filesusr.com/archives/0a280a_52926fd686184b54bb765142b7a90526.zip?dn=dai-digital-sticker.eps.zip

Supported by Engenhart ° Design Studio 2021
engenhart.design

Conference partners

appliedAI c/o UnternehmerTUM GmbH
https://appliedai.de

Alexander von Humboldt Institut für Internet und Gesellschaft
https://hiig.de

IDZ | Internationales Design Zentrum Berlin e.V.
https://idz.de

Berufsverband der Kommunikationsdesigner c/o BDG Bundesgeschäftsstelle
https://bdg.de

ICT Media and Creative Industries Berlin Brandenburg
European Union – European Regional Development Fund

Rise of Ai conference c/o AI for Humans GmbH
https://riseof.ai/

www.designing-artificial-intelligence.eu

ISBN 978-3-00-068634-4