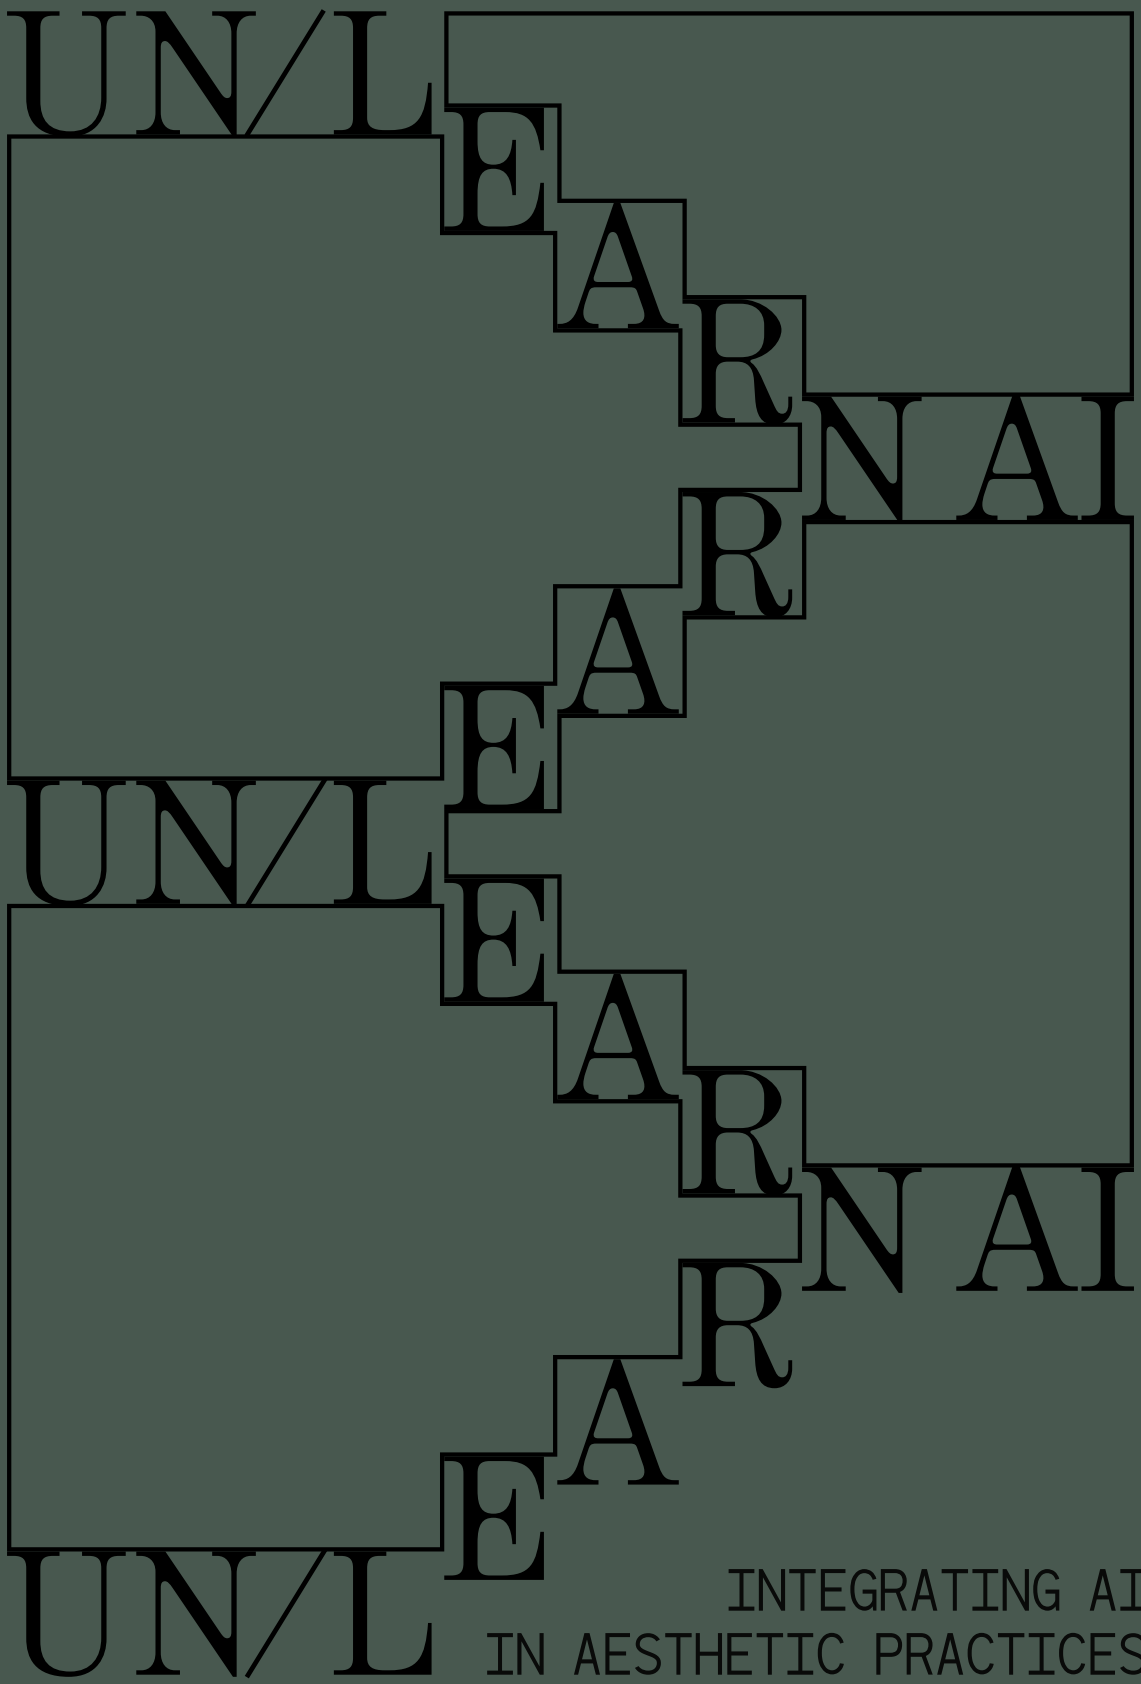
UN/LEARN AI
UN/LEARN AI
UN/LEARN AI
UN/LEARN
INTEGRATING AI
IN AESTHETIC PRACTICES

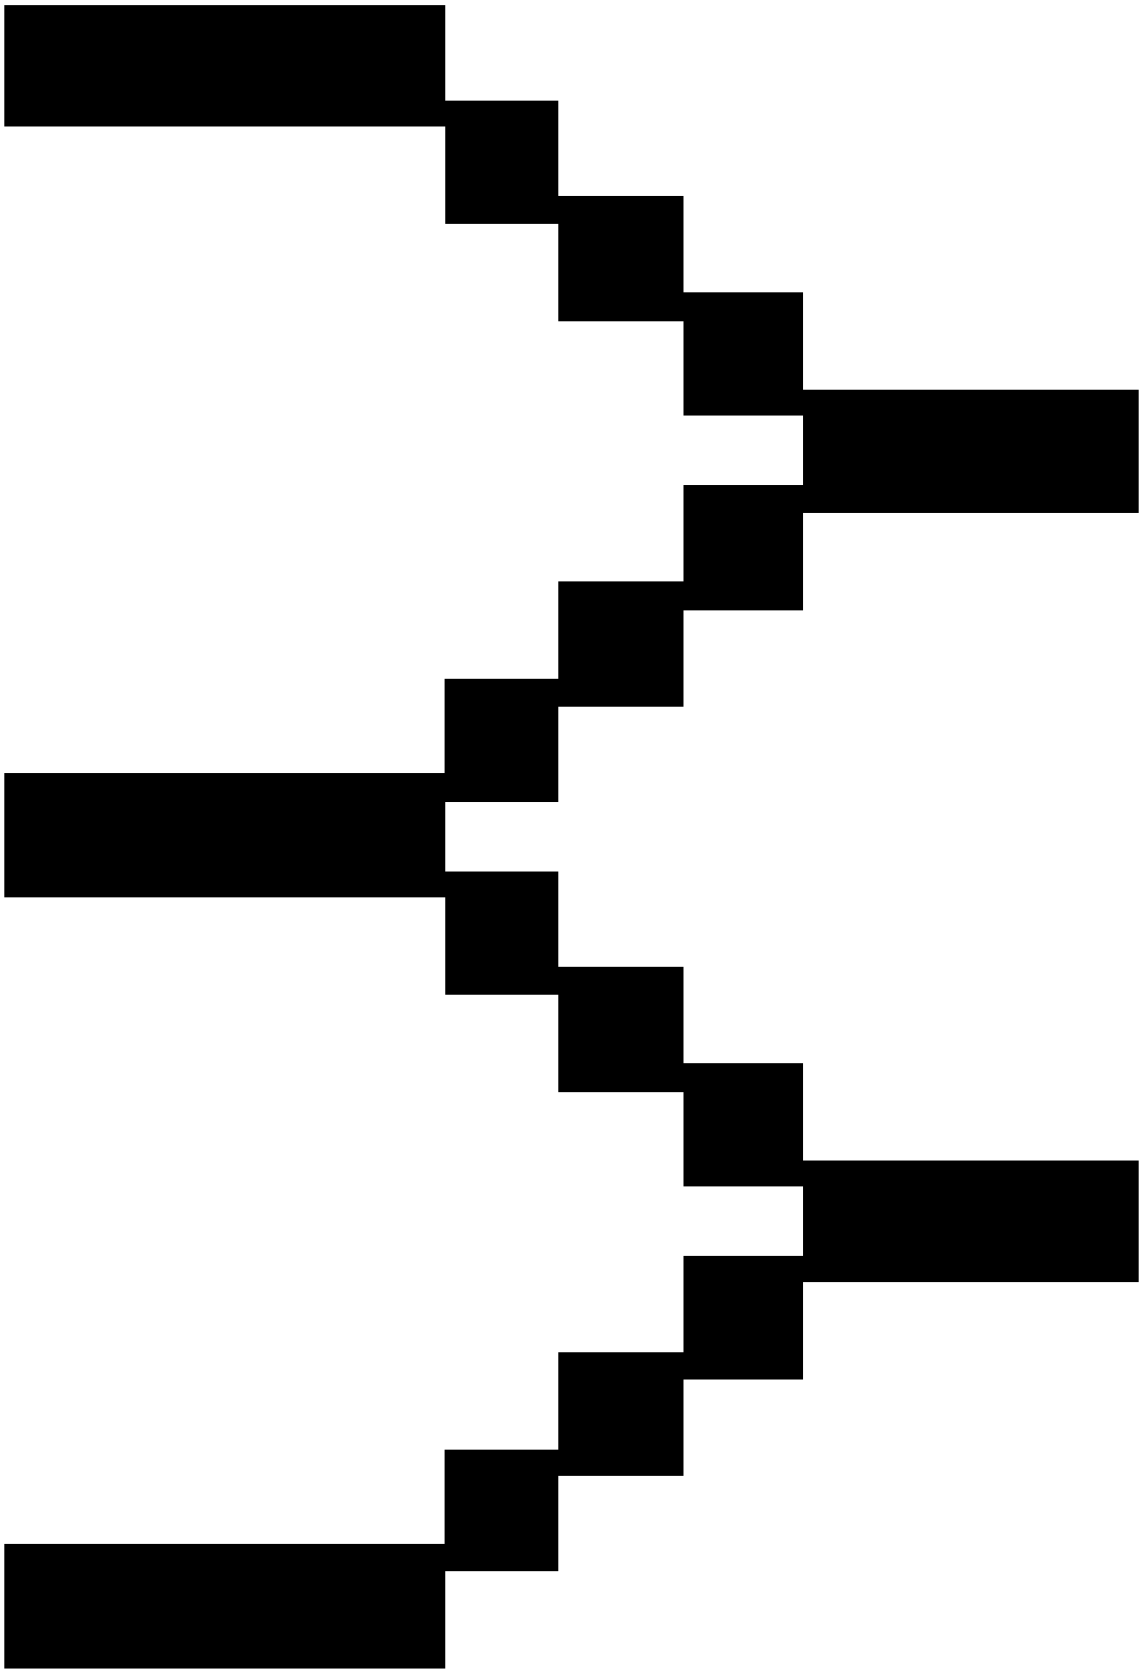

Editor‘s Note: Approaching, Navigating and Integrating AI in aesthetic practices

Matthias Grund, Johanna T. Wallenborn

Chunk 1 This publication marks the final volume in our three-part series *un-learn AI*. It concludes an ongoing research journey: a process of approaching, navigating, and integrating artificial intelligence into the art and design curricula of five German universities.

Together, these volumes reflect the cumulative outcomes of the research project *KITeGG – Making AI tangible and comprehensible: Connecting technology and society through design*. Despite its unwieldy acronym (KITeGG - KI greifbar machen und begreifen: Technologie und Gesellschaft verbinden durch Gestaltung), it served well as our programmatic guide.

Our collaborative research initiative brought together Hochschule Mainz, HfG Offenbach, Hochschule Trier, Köln International School of Design (KISD) and HfG Schwäbisch Gmünd. Each institution contributed its own unique focus area, while shaping a shared perspective on integrating AI and machine learning into creative education.

Chunk 2 Since the project's launch in 2021, we have developed, tested, and refined strategies for integrating artificial intelligence into design and art education. From the beginning, our goal was not just to teach technical proficiency. Rather, we asked ourselves what designers and artists need to *un-learn* in order to engage with these technologies in critical, ethical, and imaginative ways.

This final volume, *un/learn AI: Integrating AI in Aesthetic Practices*, reflects on those efforts. It documents our educational experiments, student-led explorations, and practice-based research during the project’s final years.

Approaching

Chunk 3 The first phase focused on orientation. We began *approaching* artificial intelligence by collecting questions, identifying signals, and mapping the shifting terrain of intelligent systems and exploring their implications for cultural production. During this time, we started building our infrastructure, laboratories, and course programs. In doing so, we developed a shared vocabulary, questioned existing paradigms, and laid the foundations for a transdisciplinary dialogue. Research fields emerged, sometimes planned and sometimes unexpected, as artificial intelligence continued to rapidly develop as a technological and sociocultural phenomenon.

The first year also marked the launch of our ongoing Summer and Winter schools.

Chunk 4 Hosted in rotation by each partner university, these gatherings became platforms for exchange with each other and invited leading experts in the field. Concurrently, first steps were taken to develop a central learning and teaching platform designed to host all our newly developed courses across the universities.

Navigating

Orientation led to experimentation. During this second phase, we tested new course formats, opened up discursive spaces, and explored different teaching methods.

As generative AI models, such as OpenAI's ChatGPT and Stability AI's Stable Diffusion, entered the public sphere with unprecedented accessibility, our learning platform expanded to offer new courses and workshops that engage with these transformative technologies critically and creatively. The emergence of multimodal models and conversational agents gave way to a rethinking of educational content and delivery.
Chunk 5 Disciplinary boundaries began to soften, and *navigating* became an ongoing process of negotiation between theory and practice, teaching and research, and institutions and individuals.

Integrating

The final phase, *integrating*, signifies both arrival and departure. It is a time for reflection and transition, and perhaps a pinch of nostalgia.
Chunk 6 Courses initially developed at one school were adapted and shared across institutions. Events such as *KI-Abend (HfG Offenbach)* and seminars like *Thinking and Prompting: Images (KISD)* expanded through inter-university exchange. Collaborative formats, such as *Making (Non-)Sense (HfG Schwäbisch Gmünd & KISD)*, were co-developed and presented internationally.

We ask ourselves: What have we learned?
Chunk 7 Which structures appeared to be successful? Where have our findings challenged existing contexts, and where have they contributed to new ones?

A significant outcome is the realization of an inter-institutional master's program titled "↗ Design is All You Need – AI and Design Education," hosted within the Communication Design M.A. program at Hochschule Mainz.
Chunk 8 This iteration incorporates teaching formats from all five KITeGG partners and directly involves students in the activities of the research project. Building on the insights and questions generated by KITeGG, the program expands the collaborative teaching experience through ongoing dialogue, course development, and transdisciplinary networks.

A Call for Continued Un/Learning

This publication represents the essence of this final integration phase.
Chunk 9 It unfolds in three interconnected parts:

Educational Strategies in Aesthetic Practices reflects on our teaching methods, including what we tried, what worked, and what didn't.
Chunk 10 It also discusses how we created pedagogical spaces where students could engage with AI in critical, experimental, and creative ways.

The *Student Project Showcase* takes the form of a visual showcase, a nonlinear sequence of current student work across all schools, which invites readers to experience the projects visually and intuitively before encountering the ideas and processes behind them. The projects span critical perspectives, applied technologies, and poetic approaches.
Chunk 11 *Experimental Research in Aesthetic Practices* presents alternative research formats developed within the project. It includes contributions that use design and artistic practices as methods of inquiry, exploring AI as a discursive, affective, and material phenomenon rather than as a tool.

Unexpectedly, but perhaps inevitably, some of our initial questions became

more urgent, sparking more questions along the way:

What kinds of aesthetics, possibilities, and unexpected encounters arise from deep learning systems? What assumptions and risks are embedded in their training data, and how can we reveal and challenge them?
Chunk 12 How can we empower artists and designers to not only use AI but also critically and responsibly shape it? How do we foster design practices that consider not only *how* to use these systems, but also *when* and *why* to use them—and, crucially, *when not to*?

As we wrap up our joint research initiative after four years, these questions remain alive, evolving, and increasingly urgent.
Chunk 13 What began as a collective inquiry into AI's role in aesthetic practices has expanded into a broader reflection on how we teach, learn, and create in times of algorithmic transformation.

We hope this book sparks new conversations and collaborations across disciplines and institutions. To un/learn AI is not a linear task but an ongoing negotiation, one that calls for curiosity, critique, care, and imagination.
Chunk 14 The work continues.
Chunk 15 – The Editors

Hybrid Publishing and Semantic Chunking

This publication was produced using a hybrid system combining the open-source *PubPub* publishing system and the web-to-print *Paged.js* system. This setup enabled remote collaboration among authors and team members, as well as publishing in a hybrid format accessible online at ↗ *unlearn.gestaltung.ai* and in print.

Each article features a unique short ID (e.g. "cetyu72y"), located at the beginning of the text. These IDs serve as direct links between the print and digital versions, allowing readers to access articles online at *https://unlearn.gestaltung.ai/article/[short-ID].*

Additionally, automated semantic cross-references (semantic chunking) offer alternative reading paths.
Chunk 16 The texts were analyzed using AI, divided into content-based sections, and compared with similar sections across this and the previous two volumes. These semantic "chunks" appear in the margins, linking related content and enabling non-linear, horizontal reading across the series.

Thus, this publication series functions as both a set of documents and a network, designed to be read and explored.

Chunk References

Chunk 1
I. p. 10, Chunk 5: Projektvorstellung
II. p. 12, Chunk 1: Mainz
I. p. 2, Chunk 1: Vorwort

Chunk 4
I. p. 9, Chunk 4: Projektvorstellung

Chunk 7
II. p. 12, Chunk 1: Mainz
II. p. 179, Chunk 2: Editorial
II. p. 14, Chunk 6: Mainz

Chunk 8
II. p. 213, Chunk 16: Uncanny Type
I. p. 9, Chunk 3: Projektvorstellung
II. p. 81, Chunk 12: Editorial

Chunk 10
I. p. 99, Chunk 4: Bildgenerierende Modelle in…
II. p. 34, Chunk 13: Trier
I. p. 30, Chunk 16: Living Objects Lab (KISD)

Chunk 15
I. p. 2, Chunk 3: Vorwort

Chunk 16
I. p. 2, Chunk 5: Vorwort
II. p. 5, Chunk 2: KI-Features dieser Ausgabe
II. p. 17, Chunk 20: Mainz

Educational Strategies

10 Introduction
Lasse Scherffig

12 AI Literacy for the Long Haul: Teaching Lasting Competencies in an Ever-Changing Field
Jakob Kilian

20 From Physical Computing to Physical AI: Challenges in Teaching Robotics and AI in Creative Contexts
Ivan Iovine

30 Building AI Intuition – Four Educational Pillars for Teaching AI in Design at AI+D Lab
Rahel Flechtner, Jordi Tost

44 Motion Experience Lab@(A)IDLE MODE: Echtzeit-Installation zur performativen KI Erfahrung – Embodied AI Learning in Design Education
Nina Juric

66 DIY Compute: Rolling Your Own GPU Node from Cheap Hardware
Anton Koch

76 Introduction to Critical AI Studies
Johanna T. Wallenborn

78 Zwischen Tropfen und Ozean – Wie sich KI die wertvollste Ressource der Erde einverleibt
Niklas Thielen

82 Die Pille für KI
Lea Waldera

86 Earth's Basilisk: Between Myth and Consequence
Dzennifer Zachlod

Student Project Showcase

92 Introduction
Leon-Etienne Kühr, Mattis Kuhn, Jana Hartmann, Elisa Deutloff

168 Anthrobot
Ilona Achba, Tyanka Adrian, Philipp Gatzke, Meike Kattwinkel, Elena Schröder, Justus Schomann, Aikaterini Sideri, Jessica Zisa

168 Dainų šventė
Inês Fernandes de Aguiar, Alina Bertacca, João Borges, Bidisha Das, Bia Dobbs, João Nuno Rua Fonseca, Jana Hartmann, Luka Jacke, Yonghyuk Kim, Katrina Milta, Mário Pereira, Jovydas Petrėtis, Jule Pickel, Elzbieta Ploiakova, Vytenis Razmus, Daniel Rosewich, Aikaterini Sideri, Anna Sokolova, Hannah Steinbrecher, Anastasija Surovceva, Ada Wagner

169 Der Raum kann wahrgenommen werden
Daniel Mikula Scholl, Peter Ehses

169 Digitale Entfremdung
Elisa Deutloff

170 Distorted Messages
Viktor Werbilo

170 FloraForge
Agnes Stigler, Berit Bove, Colin Subat, Oliver Beer

170 Flush of Flesh
Annika Leuoth

171 Generative Künstliche Intelligenz und die Konstruktion von Glaubwürdigkeit
Anna Krämer

171 MK Ultra
Max Kreis

172 Schlag auf Schlag
Nelli Gomez Baumert

172 Served
Maybritt Dornheim

172 StableHub
Rahel Flechtner, Franz Anhaeupl

173 Stuhl
Zhichang Wang

173 Synthetic Eden
Pourya Sharifi

173 Synthetic Reflection
Marie Backhaus, Henrike Hof, Levon Javaheri, Eva Matzerath, Phillip Schulz

174 Synthetische Erinnerung
Berire Akdogan, Hannes Bader, Alissa Bernhardt, Piroschka Benedikta Brucker, Marlene Grabenwöger, Helena Hofmann, Georgia Ioannidou, Lisa Oelsen, Tristan Schulze, Niklas Thielen, Lea Waldera, Cosima Wagner, Micha Wempe, Anastasia Verena Wilbrandt, Jennifer Zachlod

174 With a Grain of Salt
Moritz Hartstang, Alicia Link, Robert Kremer

Experimental Research

178 Introduction
Alex Oppermann, Mattis Kuhn, Leon-Etienne Kühr

182 I Think He's Got a Helicopter in His Pocket: The AI Chewing Gum
Thomas Hawranke, Lasse Scherffig

188 An Incredibly Average Face – Approaching 31,200 Images with Human Eyes, Sorting Systems, and Collaging
Jana Hartmann

204 Attention to Color – The Role of Saliency in Image-Based UI Theming
Robert Halbach

212 Body as Latent Space: Alternative Compressions in Aesthetic Practice
Simon Maris

218 On the Materiality of AI
Ting-Chun Liu

224 un/learn the process
Francesco Scheffczyk, Paul Eßer, Lars Hembacher

230 26 Strategies to Consider When Writing with AI
Julia-Jasmin Bold

234 Notes Toward AI Un/Learning
Simon Maris

Closing

237 How to KITeGG: A Semi-impractical Introduction to Sourcing AI Infrastructure
Anton Koch

252 Auf dem Weg zu KITeGG
Florian Jenett

256 KITeGG und nun?
Florian Jenett

258 Standorte & Team

Educational

Strategies

12 AI Literacy for the Long Haul: Teaching Lasting Competencies in an Ever-Changing Field
Jakob Kilian

20 From Physical Computing to Physical AI: Challenges in Teaching Robotics and AI in Creative Contexts
Ivan Iovine

30 Building AI Intuition – Four Educational Pillars for Teaching AI in Design at AI+D Lab
Rahel Flechtner, Jordi Tost

44 Motion Experience Lab@(A)IDLE MODE: Echtzeit-Installation zur performativen KI Erfahrung – Embodied AI Learning in Design Education
Nina Juric

66 DIY Compute: Rolling Your Own GPU Node from Cheap Hardware
Anton Koch

76 Introduction to Critical AI Studies
Johanna T. Wallenborn

78 Zwischen Tropfen und Ozean – Wie sich KI die wertvollste Ressource der Erde einverleibt
Niklas Thielen

82 Die Pille für KI
Lea Waldera

86 Earth's Basilisk: Between Myth and Consequence
Dzennifer Zachlod

Introduction

Lasse Scherffig

Chunk 1 Teaching and learning artificial intelligence in design does not come with a recipe. Formats are needed that enable autonomous and critical yet technologically literate approaches to these technologies. However, defining AI literacy is all but straight-forward. The diverse technologies and practices under the AI umbrella can hardly be seen as one field with a coherent set of methods to be learnt and mastered. Moreover, as a *sociotechnical narrative* what we call AI does not only comprise technologies, it includes diverse contexts, such as institutions and ideologies, as Anton Koch has pointed out earlier in this project.[1] And even if we focus on technology alone, much of what is state-of-the-art today is outdated before it possibly can be adopted into a learning setting.
Chunk 2 For the teachers involved in KITeGG, this is nothing new. Fields like creative coding have long oscillated between teaching a foundational understanding of computation and supporting practical skills that are bound to become outdated.[2] But with generative AI, the distance between both may have grown even bigger — and so has the need to teach and learn how to adapt and transfer knowledge, as Jakob Kilian explains in this volume (↗ 9h8iy9yw). Drawing on the experience of developing courses and workshops at Köln International School of Design, he ultimately argues for a form of sustainable AI training that connects high- and low-level approaches.
Chunk 3 Similarly, Ivan Iovine describes how at the Robotics Lab at the University of Art and Design Offenbach, working with *physical AI* implies moving along the spectrum between simple and complex technical understanding (↗ j4nosczy). This, he argues, implies not only teaching technical skills, but developing a culture of robotic and interactive arts.

If the application of AI in design is in need of decidedly cultural perspectives, those perspectives need cultivating. Accordingly, Johanna Wallenborn shows how to introduce design students to critical AI studies (↗ 02mg286b). Beyond introducing and motivating her educational work, her contribution showcases three graduate student essays (by Niklas Thielen, (↗ 2lws78up) Lea Waldera, (↗ 957yw4ig) and Dzennifer Zachlod (↗ ahruydd0)), each embodying a distinct perspective on the sociotechnical ecologies of AI.

At HfG Schwäbisch Gmünd, dealing with the cultural and technological side of these ecologies aims at building *AI intuition*, as Rahel Flechtner and Jordi Tost show (↗ t0rgtqez).
Chunk 4 This includes enabling technical literacy through hands-on exploration while fostering conceptual engagement and critical reflection.
Chunk 5 Their strategies accordingly range from the non-technical, like role-play, to building tools for experimentation. Notably, Flechtner and Tost point out that a combination of deep technical literacy and conceptual engagement presupposes curricular structures that allow long-term commitment and sequential formats.

In order to enable discourse, these formats repeatedly need to become physical.
Chunk 6 They need to occupy locations and engage with audiences. The (A)IDLE MODE project by Nina Juric constitutes such an attempt at situating AI (↗ ja6f4xwe). Here, AI serves both as topic and production technology, enabling an ephemeral experience for 300 people.

1: Koch, A. (2025). Der KITeGG Cluster – eine Infrastruktur für KI in der Gestaltungslehre. In KITeGG (Ed.), *Un/learn ai: Navigating AI in aesthetic practices*. HS Mainz.

2: Trogemann, G., & Viehoff, J. (2005). *Code@Art. Eine elementare Einführung in die Programmierung als künstlerische Praktik*. Springer.

Beyond showing different educational strategies, these contributions make clear that teaching and learning AI in design means addressing more than just technology. It does require technological infrastructure — in the form of laboratory equipment and *compute*, which KITeGG provides through five newly-established laboratories[3] and a shared high-performance computing cluster[4] (↗ 2gt-mc7uq).

Chunk 7 But even these basic technical requirements are part of a larger conversation about the interactions of labor, capital, politics and economy.[5] Developing cloud-based infrastructure outside the corporate world or providing an exemplary guide to building cheap and small-scale DIY infrastructure (↗ ro8j5asv) hence equally addresses the need for compute as it addresses the social dimension of technology.

In other words, teaching and learning AI in design depends on *infrastructuring*, as understood in science and technology studies and participatory design.[6]

Chunk 8 Here, the term encompasses the collective creation, adaptation, and maintenance of socio-technical ecosystems — a process that unfolds over extended periods of time. As such, infrastructuring lies at the core of all the educational strategies collected in this volume, as it involves providing laboratories and compute, developing custom tools, adopting and adapting existing systems, developing conceptual skills, cultivating cultures, and fostering a critical understanding of the heterogeneous field that we currently call AI.

Chunk References

Chunk 1
I. p. 81, Chunk 5: Paper on the topic of AI in…
II. p. 211, Chunk 10: Uncanny Type
I. p. 113, Chunk 1: Developing an AI teaching…

Chunk 2
II. p. 67, Chunk 8: Editorial
III. p. 257, Chunk 9: KITeGG und nun?
II. p. 67, Chunk 13: Editorial

Chunk 3
II. p. 28, Chunk 1: Offenbach
I. p. 21, Chunk 4: Hochschule Trier
II. p. 4, Chunk 3: Vorwort

Chunk 5
II. p. 202, Chunk 5: Editorial: Show and Tell

Chunk 6
III. p. 256, Chunk 6: KITeGG und nun?
II. p. 157, Chunk 8: Integrating Physical AI and…
I. p. 10, Chunk 5: Projektvorstellung

Chunk 7
II. p. 250, Chunk 52: Der KITeGG Cluster – eine…
III. p. 256, Chunk 6: KITeGG und nun?
III. p. 256, Chunk 7: KITeGG und nun?

3: KITeGG (Ed.). (2024). *Un/learn ai : Approaching AI in aesthetic practices*. <https://doi.org/10.25358/OPENSCIENCE-10586>

4: Koch, A. (2025). Der KITeGG Cluster – eine Infrastruktur für KI in der Gestaltungslehre. In KITeGG (Ed.), *Un/learn ai: Navigating AI in aesthetic practices*. HS Mainz.

5: Koch, A. (2025). Der KITeGG Cluster – eine Infrastruktur für KI in der Gestaltungslehre. In KITeGG (Ed.), *Un/learn ai: Navigating AI in aesthetic practices*. HS Mainz.

6: Karasti, H. (2014). Infrastructuring in participatory design. *Proceedings of the 13th Participatory Design Conference: Research Papers - PDC '14*, 141–150. <https://doi.org/10.1145/2661435.2661450>

↗ 9h8iy9yw

AI Literacy for the Long Haul: Teaching Lasting Competencies in an Ever-Changing Field

Jakob Kilian

Chunk 1 Introduction

Much has been said about how AI is opening up new possibilities and about the speed at which models, tools, and services are being developed and published. There is, moreover, broad consensus on the relevance of AI in design education. The skill set to be taught can be summarized as *AI literacy*[1] analogous to other areas that emerged at various times and had strong links or overlaps with design (e.g., *digital literacy* or *computer literacy*). However, it is unclear exactly what is meant by this term.

Chunk 2 Due to the rapid pace of developments, the lack of research and established guidelines on this specific matter, institutions and educators are approaching it in fairly different ways, depending on the focus of the design discipline.

II. p. 21, Chunk 2: Schwäbisch Gmünd
II. p. 22, Chunk 4: Schwäbisch Gmünd
II. p. 20, Chunk 4: Köln

Probably the most important question here is: What level of foundational knowledge does a designer *need* in order to *integrate* these topics into their

1: Long, D., & Magerko, B. (2020). What is AI Literacy? Competencies and Design Considerations. *Proceedings of the 2020 CHI Conference on Human Factors in Computing Systems*, 1–16.

design practice, and to collaborate and communicate with experts in this specific academic discipline? Answers to this question are sure to remain vague if the educational objectives regarding AI are not clearly defined.

Chunk 3 This usually (still) is the case in common design curricula, unlike in specialized degrees that already focus on AI or related subjects. As Farque et al.[2] put it, objectives can range from seeing the designer as a *consumer* who *uses* tools that contain some form of AI, to seeing them as an *creator* who (co-)develops these AI tools and the associated human-machine interactions. Training for the *consumer* role mainly consists of providing in-depth experience with high-level tools or "exposing students to cutting edge technology," as McCardle already stressed 20 years ago.[3] For the *creator* role on the other hand, students need to be equipped with the technical skills to train and deploy their own AI models and applications, work in interdisciplinary teams with specialists, and reflect on the societal impact of their designs. These two roles are of course only poles of a spectrum that will be covered again below.

In times when low-cost or free AI services have become widely available, this essay aims to encourage educators not to place all their bets on the first card. Instead, they should continue equipping students with the fundamentals that provide a deeper understanding, better application skills, the ability to transfer these skills, and above all, *to keep up with ever-evolving innovations*.

Chunk 4 Because, as Farque et al.[4] write, many of the new AI technologies that a student will use in their working life "haven't been invented yet."

So, What Will Be Invented Then?

The future development of the AI sector is far from predictable. A short look at the history of AI research shows that probably the most important development of recent years was not widely anticipated: While decades of research into models for storing and transferring knowledge failed, so-called large language models (LLMs), originally developed for syntactic tasks such as translation, suddenly achieved a breakthrough as a result of extreme scaling and became *knowledgeable*.[5] After a few years of further progress by scaling these models, some researchers have started to express their suspicion that scaling might reach or has already hit a plateau and that the pace of development is generally slowing down.[6]

II. p. 141, Chunk 3: Editorial: Physical AI
III. p. 76, Chunk 1: Introduction to Critical AI…
III. p. 76, Chunk 3: Introduction to Critical AI…

Chunk 5 A glance at the very heated and fast moving market today, in which many promises are being made and large amounts of risky capital are flowing, suggests a pretty volatile course of AI development. AI-based or AI-assisted tools are increasingly finding their way into many areas of life and markets, often as a system as a service

2: Faruqe, F., Watkins, R., & Medsker, L. (2021). *Competency Model Approach to AI Literacy: Research-based Path from Initial Framework to Model*.

3: McCardle, J. R. (2002). The Challenge of Integrating AI & Smart Technology in Design Education. *International Journal of Technology and Design Education*, *12*(1), 59–76.

4: Faruqe, F., Watkins, R., & Medsker, L. (2021). *Competency Model Approach to AI Literacy: Research-based Path from Initial Framework to Model*.

5: Pavlus, J. (2025). When ChatGPT Broke an Entire Field: An Oral History. In *Quanta Magazine*. <https://www.quantamagazine.org/when-chatgpt-broke-an-entire-field-an-oral-history-20250430/> This assessment is difficult to substantiate with individual evidence. However, Pavlus supports it with an excellent collection of expert opinions from the field.

6: Marcus, G. (2022). *Deep Learning Is Hitting a Wall*.

(SaaS) offering a user-friendly interface and computing power from a single source. As many of these do not appear to be profitable (yet)[7] [8] but only offered at significant losses to secure a projected market position, it is conceivable that some of them may no longer be available in the near future or may be too costly for the intended use.

Chunk 6 Additionally, new tools, architectures or services that will emerge might work altogether differently and have their own requirements.

Chunk 7 So, to summarize, we can't say what will be invented yet!

Yet nothing appears out of thin air.

Chunk 8 While the tools and methods of applications are changing rapidly, many technical principles have barely changed over the years and can look back on linear development. The application of the attention principle and the *Transformer* architecture—one of the main drivers of the current AI boom—may only date back eight years,[9] but it builds, for example, on the research on previous work about the attention mechanisms,[10] which is based on the development of the first deep learning models,[11] which itself is based on the Neocognitron,[12] which in turn builds on the very principles of how a computer works, as conceived by von Neumann[13] or Turing[14] and so on and so forth. Even if one need not fully understand either principle to *use* Transformers, a basic knowledge of each helps to build up skills, awaken curiosity to discover, and grasp new systems.

II. p. 22, Chunk 4: Schwäbisch Gmünd
II. p. 20, Chunk 4: Köln
III. p. 3, Chunk 2: Editor`s Note

Chunk 9 These observations should give pause when it comes to the didactic preparation of teaching content—especially at higher education institutions: If there is opportunity and time to study AI in a design curriculum, one must consider its long-term sustainability. So let's take a look at the means by which this can be achieved.

Abstraction vs. Foundations

We should start with establishing some criteria for evaluating educational programs.

II. p. 207, Chunk 18: Shaping Human-AI Relationships
I. p. 104, Chunk 16: Bildgenerierende Modelle in…
II. p. 205, Chunk 9: Shaping Human-AI Relationships

Chunk 10 Naturally, as with any course content, a key objective should be that students are able to *repeat* and *apply* what they have learned independently after the course. In AI, for example, this means that they can do so with their own software and hardware and without problem-solving support from a teacher. Ideally, students should not only be able to *apply* their knowledge but also *adapt* and *transfer* it to new contexts, materials, or media—a non-trivial task due to the increasing complexity of the models as well as the systems in which they are integrated.

7: Edward Zitron. (2025). *There Is No AI Revolution*.

8: Isaac, M., & Griffith, E. (2024). OpenAI Is Growing Fast and Burning Through Piles of Money. *The New York Times*.

9: Vaswani, A., Shazeer, N., Parmar, N., Uszkoreit, J., Jones, L., Gomez, A. N., Kaiser, Ł., & Polosukhin, I. (2017). Attention is All you Need. *Advances in Neural Information Processing Systems*, *30*.

10: Bahdanau, D., Cho, K., & Bengio, Y. (2016). *Neural Machine Translation by Jointly Learning to Align and Translate*.

11: Krizhevsky, A., Sutskever, I., & Hinton, G. E. (2012). ImageNet Classification with Deep Convolutional Neural Networks. *Advances in Neural Information Processing Systems*, *25*.

12: Fukushima, K. (1980). Neocognitron: A self-organizing neural network model for a mechanism of pattern recognition unaffected by shift in position. *Biological Cybernetics*, *36*(4), 193–202.

13: Neumann, J. von. (1993). First draft of a report on the EDVAC. *IEEE Annals of the History of Computing*, *15*(4), 27–75.

14: Turing, A. (1936). On Computable Numbers, with an Application to the Entscheidungsproblem. *Proceedings of the London Mathematical Society*, *42*(1), 230–265.

Further complication arises from the fact that AI itself has an inherent problem of reproducibility.[15] [16] [17]

Two contrasting, but by no means dichotomous, poles of a spectrum are presented next in order to identify opportunities and difficulties for teaching AI in design. While the actual structure of the teaching format will ultimately depend on time, framework, and material, the arguments will be weighed in the following discussion and guidelines will be given on how to navigate this spectrum.

Pole One: Using High-Level Abstraction

This approach focuses on rapid engagement and immediate and iterative creative output by leveraging high-level, often commercial, AI tools. These are often, but not always, SaaS and/or simplified application programming interfaces (APIs) to access complex models. Today, many data-processing steps are involved prior to and following what the model actually does. Many times, multiple models and even interactions between them (e.g., agentic AI) are used to achieve high quality results.[18]

The advantages of using these systems are clear: *Barrier to entry is extremely low*, requiring minimal technical setup or coding knowledge. Students can quickly start experimenting and generating results, usually just by using their browser.

Chunk 11 The ability to use the tools directly and productively can provide a significant *motivational boost*. Secondly, the computation is carried out in the so-called *cloud* on the provider's servers, thus *eliminating hardware requirements and the set-up of software and dependencies* for the user. Together with well *thought-out and task-optimized interfaces*, students can *focus on creative application and learn about conceptual implications of AI for design* instead of losing time and motivation on technical details and troubleshooting.

On the other hand, with these consumer-friendly models, students face the so-called *black box problem*.

Chunk 12 Not only it is difficult if not impossible to gain insight into the underlying mechanisms, algorithms, or training data, students may have trouble understanding or dealing with unexpected results. Secondly, a *SaaS has to generate money* in one way or another to cover energy and maintenance costs. Even if the services are currently free or offered at discounts for students for now, there is a good chance that they will not be in the future, creating a dangerous *dependency on the product*.

II. p. 181, Chunk 10: Editorial
II. p. 4, Chunk 5: Vorwort
II. p. 13, Chunk 4: Mainz

Chunk 13 In the worst case, specific services may be discontinued and thus completely *lost* to the user. Last but not least, services usually offer little opportunities for *customization* and often *lack of privacy*.

15: Gundersen, O. E., & Kjensmo, S. (2018). State of the Art: Reproducibility in Artificial Intelligence. *Proceedings of the AAAI Conference on Artificial Intelligence, 32*(1).

16: Albertoni, R., Colantonio, S., Skrzypczyński, P., & Stefanowski, J. (2023). *Reproducibility of Machine Learning: Terminology, Recommendations and Open Issues.*

17: Ball, P. (2023). Is AI leading to a reproducibility crisis in science? *Nature, 624*(7990), 22–25.

18: Acharya, D. B., Kuppan, K., & Divya, B. (2025). Agentic AI: Autonomous Intelligence for Complex Goals–A Comprehensive Survey. *IEEE Access, 13*, 18912–18936.

However, such high-level tools are essential for short teaching units because of the benefits mentioned above. Through a playful approach, they can ensure that students gain structured insights into the use of AI tools, regardless of their prior knowledge. In the next section, the opposing concept of teaching technical literacy will be addressed.

Pole Two: Providing Technical Literacy

The alternative approach tries to avoid the service offerings described. Instead, technical literacy emphasizes foundational understanding, thereby encouraging students to engage directly with libraries, frameworks, and models that are open source. Of course, this distinction should always be understood in relative terms, as all building blocks in software development are based on existing ones, right down to programming languages, which are also a high-level abstraction of machine languages.

With this approach, a look is taken inside of the black box: It *fosters a more profound understanding of how AI models work*, their capabilities, limitations, and the computational processes involved. As the tools are *open* the students have more *control over the software, can customize and adapt* it to their needs.

Chunk 14 In this process, students are exposed to *concepts and principles that are transferable* to all areas of AI, like knowledge about programming languages or neural networks, for example. Furthermore, free and open-source software (FOSS) is *free to use* (by its license) and will stay free forever. As long as the hardware and software dependencies support it, the exact version that was used in a course can be used again once long after the course is over. Last but not least—partly because of the aforementioned reasons—there is usually a *large support community*.

Chunk 15 This also guarantees that certain popular pieces of FOSS will remain available for a few years.

The biggest advantage of this approach—deeper, generally applicable understanding—naturally comes with its biggest disadvantage: *A steep learning curve* must be expected and *it will take some time* for design students without prior technical backgrounds to even understand the fundamentals needed to set up the environment.

Chunk 16 The complexity can indeed be both intimidating and potentially alienating. Not using SaaS (with computing power in the cloud) aditionally comes with the requirement of providing computing power and the corresponding infrastructure to access it. Setting up local hardware, but even more so setting up a project's own central infrastructure that can be accessed by several courses in a university network, is a major challenge on both software and hardware sides, as can be read in Anton Koch's articles. (↗ 2gtmc7uq) (↗ ro8j5asv)

Navigating the Spectrum

In the referenced article by Faruqe et al.,[19] the distinction between *consumer* and *creator* can be roughly mapped onto who learned on high-level versus the AI literacy approach.

Chunk 17 A designer that only uses (or *consumes*) a tool—any tool, like a wood saw or a piece of software like Adobe Photoshop—would certainly also be adequately trained with a higher-level approach: It is crucial to know the tricks of the trade and handle the tool efficiently and well to become a professional.

Chunk 18 A *creator* using AI, however, is like someone who designs an actual saw or develops plugins or features *for* Photoshop, and thus they would need to be proficient in the history of the tool, how it's made, and the inner workings of its foundational building blocks. In AI this entails thorough knowledge about neural networks, deep learning, various architectures, coding, and even the hardware's technical operations.

II. p. 22, Chunk 4: Schwäbisch Gmünd
I. p. 81, Chunk 5: Paper on the topic of AI in…
III. p. 10, Chunk 1: Introduction

Chunk 19 However, neither of these poles are really suitable at their extremes to train students in a general design degree program. What is often lumped into "AI" is actually a wide range of principles and (not yet invented) tools that are defined neither in terms of *purpose* (like cutting wood or processing digital images) nor in terms of a specific *tool* (a Japanese ryoba saw or the Photoshop software). On the other hand, it is quite obvious that there is not enough time to train students as it is done in technical AI programs.

Chunk 20 Spectrums, nonetheless, have lots of area in-between, and Faruqe et al.[20] also define two more groups in the continuum between the poles, namely the *co-worker* (knows the basics but mainly works with outputs) and the *collaborator* ("[…] works alongside one or more AI systems to improve each others performance"). These profiles should orient AI education in the field of design.

Chunk 21 The following section offers a concluding discussion on how teaching formats can be structured to navigate this spectrum.

Chunk 22 An attempt must be made to find a middle ground to engage students in the topic, lay the foundations for further independent or guided learning, and provide the necessary vocabulary and overview to reflect on the topic with foresight.

Higher-level tools offer an excellent way to rapidly introduce students to the topic and are immediately rewarded with productive results. Practical experience can promote understanding of the mechanisms of AI, create motivation, and lay a foundation for further exploration. *However, and this is the key message of this essay, it is important not to stop there.* In order to cultivate AI literacy and to work with or on (future) tools independently, students have to learn how to adapt and transfer their knowledge in the face of change and stay operationally effective. This is only possible with a look inside the black box: dealing with code and software dependencies; trying out different

19: Faruqe, F., Watkins, R., & Medsker, L. (2021). *Competency Model Approach to AI Literacy: Research-based Path from Initial Framework to Model.*

20: Faruqe, F., Watkins, R., & Medsker, L. (2021). *Competency Model Approach to AI Literacy: Research-based Path from Initial Framework to Model.*

architectures; learning how to find, install, and utilize open software packages; and most importantly also to *fail* in this process and learn how to solve errors and problems that arise.

Chunk 23 This is the path to deeper understanding of why things are the way they are. It enables students to develop their own workarounds and solution strategies, thus honing their design skills.

Chunk 24 In the process, they will inevitably encounter principles that remain constant amid rapid change, helping them to adapt to ever-changing developments.

This can be done with either open *or* commercial (SaaS) tools; but the further a teaching unit is located on the spectrum towards technical literacy, the greater its independence from the setup, in case of structural changes like service discontinuation. Avoiding proprietary services, however, means that more preliminary work in course preparation.

Chunk 25 Even if all tools are open and shared, they must first be set up. And things can break, if, e.g., software or hardware dependencies expire through updates.

First-hand experience from the KITeGG research project has shown that even skilled teaching staff are dependent on exchange with colleagues and must spend a lot of time on getting these educational setups *up and running*. To achieve an optimal learning effect, this preparatory work should not be hidden (but integrated into the coursework). This is often done in an attempt to streamline the lesson or cover up difficulties to ease understanding of the material. However, if problems and how to deal with them are never discussed in the course, students can develop false expectations, resulting in reduced motivation to *stick with it* when problems arise on their side.

Practical Examples

At the *Living Objects Lab* at KISD, a design school involved in KITeGG, we developed some teaching units designed to follow these guidelines.

I. p. 86, Chunk 3: Robotik und Computer Vision…
I. p. 95, Chunk 2: Touch x AI
II. p. 155, Chunk 4: Integrating Physical AI and…

Chunk 26 The process begins with a fundamentals course covering computers themselves, their principles and history, programming, and plenty of practical exercises.[21] That was the prerequisite for courses on AI, where students then learned how to train a model *the hard way* with widespread open source tools such as TensorFlow and deploy it on a microcontroller.[22] Following this pathway, students could apply the knowledge they had acquired step by step to create tangible outcomes in further projects.

Another example, much more condensed in terms of time, was the four-day Making (Non-)Sense workshop,[23] which was developed by Rahel Flechtner and myself. Starting with an AI roleplay, students explored how computers and AI read and *understand* sensor data. Then followed four days of training

21: Kilian, J., & Paal, E. (2024). Interactive Systems Fundamentals. In *GitLab*. <https://gitlab.rlp.net/kitegg/public/kisd/interactive-systems/24w-fundamentals>

22: Kilian, J., & Paal, E. (2024). Interactive Systems Application - Embedding AI. In *GitLab*. <https://gitlab.rlp.net/kitegg/public/kisd/interactive-systems/24w-application-embedding-ai>

23: Flechtner, R., & Kilian, J. (2024). Making (Non-)Sense–A Playful and Explorative Approach to Teaching AI Intuition for the Design of Sensor-Based Interactions. *Proceedings of the 6th Annual Symposium on HCI Education*, 1–9. <https://doi.org/10.1145/3658619.3658643>

their own models in an SaaS that acts as an intermediate tool between higher-level and technical literacy. Sensor data can easily be uploaded, edited, and sorted with it, and models can be trained and exported.

Chunk 27 However, deployment to an Arduino Microcontoller, enabeling to run the model off computers and the internet, still requires some technical knowledge.

Finally, it should be noted that another article (↗ t0rgtqez) in this publication by Rahel Flechtner and Jordi Tost describes a similar middle-grounds approach using the term *AI intuition*, also used in earlier publications.[24] [25]

Chunk 28 In addition to the need for technical foundations and practical exercises (similar to this article), they further describe teaching content on conceptual engagement and critical reflection.

Conclusion

The temptation to introduce the very latest service or tool into the classroom should be avoided; rather, higher education should aim to impart sustainable AI training that can outlast a trend. One must also consider a technical teaching unit that initially might appear dry in order to equip students with resilient AI literacy and ultimately offers them independent repeatability, adaptability, and transferability.

24: Flechtner, R., & Kilian, J. (2024). Making (Non-)Sense—A Playful and Explorative Approach to Teaching AI Intuition for the Design of Sensor-Based Interactions. *Proceedings of the 6th Annual Symposium on HCI Education*, 1–9. <https://doi.org/10.1145/3658619.3658643>

25: Flechtner, R., & Stankowski, A. (2023). AI Is Not a Wildcard: Challenges for Integrating AI into the Design Curriculum. *Proceedings of the 5th Annual Symposium on HCI Education*, 72–77. <https://doi.org/10.1145/3587399.3587410>

↗ j4nosczy

From Physical Computing to Physical AI: Challenges in Teaching Robotics and AI in Creative Contexts

Ivan Iovine

Chunk 1 Introduction

The integration of AI and robotics into arts education offers significant opportunities, but also presents substantial pedagogical challenges. Given the range of technical disciplines involved, developing and teaching a complete technical pipeline can be overwhelming for students without a technical background.

Chunk 2 Teaching these technologies in an art and design context requires careful structuring to ensure accessibility, clarity, and creative

applicability.

II. p. 30, Chunk 6: Offenbach
II. p. 30, Chunk 10: Offenbach
II. p. 160, Chunk 21: Integrating Physical AI and…

Chunk 3 At the Robotics Lab of the University of Art and Design Offenbach am Main, both infrastructural and educational strategies have been implemented to meet these demands. On the infrastructural side, the lab was equipped with digital manufacturing tools and robotic systems, providing students with the technical resources needed to explore and experiment artistically. On the educational side, the program began with a discipline already familiar and well-established in creative coding: physical computing. This initial phase offered a foundational technical context that felt accessible to students from artistic backgrounds. Over the following semesters, the curriculum gradually shifted focus toward more advanced AI and computer vision technologies, ultimately reaching the domain of Physical AI. This emerging field combines AI-based computational systems with physical hardware and robotics, offering new possibilities for the creation of interactive and reactive artistic works.

This article outlines the step-by-step development of both the educational methodology and the supporting infrastructure. It shows how foundational practices in physical computing evolved into more complex learning pathways, and how access to digital manufacturing facilities supported students in integrating these technologies into their artistic processes.

Challenges in Teaching Robotics and AI in Creative Contexts

Teaching robotics and artificial intelligence in a creative context goes beyond providing access to new tools and learning to code. It requires rethinking teaching methods, developing suitable infrastructure, and aligning learning objectives with the diverse needs of creative students.

II. p. 3, Chunk 1: Vorwort
II. p. 21, Chunk 2: Schwäbisch Gmünd
II. p. 180, Chunk 5: Editorial

Chunk 4 Over the past three years, the aim has been not only to transfer technical knowledge, but also to foster confidence, curiosity, and an artistic culture around emerging technologies. Achieving this in a way that is both pedagogically sustainable and artistically open-ended requires careful planning and continuous adaptation. Within this context, several challenges emerged during the implementation of the program at the Robotics Lab of the University of Art and Design Offenbach am Main. The most significant ones were the following:

- *Laboratory Infrastructure:* A significant initial challenge involved establishing a dedicated lab. This lab needed to support the practical teaching of AI, robotics, and physical computing. Additionally, students needed to aquire manual competencies to create physical artifacts integrating these technologies. This included working with digital manufacturing tools such as 3D printers and laser cutters (Figure 2).
- *Cultural and Critical Integration:* Another challenge was developing a culture of robotic and interactive arts. This cultural framework needed to be both technically informed and critically aware. It required showcasing significant existing works in the field. This exposure would

inspire students to develop new artistic works informed by these examples.

- *Technical Overload Prevention:*

It was also critical to avoid overwhelming art and design students new to programming with overly technical instruction. Courses needed a structured, gradual approach, beginning with familiar established technologies. Only subsequently would the complexity increase, integrating robotics and AI technologies progressively.

Chunk 5

Framework and Educational Methodology

The initial teaching framework was based on a Physical Computing course that focused on Arduino microcontrollers. This technology had already been well-integrated at the university before the Robotics Lab was established.

Chunk 6 Building on this familiar technical groundwork, the goal was to offer an accessible entry point into the field of robotics. To support this aim, and in recognition of the need for a more structured progression, targeted refinements were introduced over the first three semesters. A key development was structuring the program as two sequential courses: a foundational course in the winter semesters and an advanced course in the summer semesters. In the foundational courses, students learned essential skills with Arduino, sensors, and actuators. The pedagogical approach followed constructivist principles, which emphasize learning through active engagement, iterative making, and direct manipulation of materials and tools. Rather than delivering abstract knowledge, this pedagogy encourages students to construct understanding by exploring, testing, and building in context.[1] Such approaches have become increasingly popular in the teaching of creative coding, particularly in design and arts education.[2] Students built their understanding through direct interaction with technology and practical, creative applications. To better prepare students for AI topics in the subsequent advanced course, the final four sessions gently introduced computer vision concepts. Students received ready-to-use Python scripts that connected Arduino with computer vision libraries such as OpenCV and MediaPipe. The code was introduced in a simple and approachable way to support those with limited programming experience.

II. p. 155, Chunk 4: Integrating Physical AI and…
I. p. 101, Chunk 10: Bildgenerierende Modelle in…
III. p. 33, Chunk 9: Building AI Intuition – Four…

Chunk 7 Students could explore the scripts by changing values, for example adjusting facial or hand detection settings, and observe the effects. Alongside coding, they applied their hands-on skills to embed sensors and actuators into physical objects.

Chunk 8 They also learned basic digital fabrication methods, including 3D printing and laser cutting. As a result, their skills expanded beyond

1: Hannafin, M. J. (1995). Open-ended learning environments: Foundations, assumptions, and implications for automated design. In R. D. Tennyson & A. E. Barron (Eds.), *Automating Instructional Design: Computer-Based Development and Delivery Tools* (pp. 101–129). Springer Berlin Heidelberg. <https://doi.org/10.1007/978-3-642-57821-2_5>

2: Grace, K., Klaassens, B., Bray, L., & Elton-Pym, A. (2022). An open-ended blended approach to teaching interaction designers to code. *Frontiers in Computer Science*, *4*, 813889. <https://doi.org/10.3389/fcomp.2022.813889>

programming, giving them practical experience in fabrication and material-based experimentation.[3]

Chunk 9 The advanced course expanded upon these initial foundations by incorporating more complex technologies. Students experimented with Raspberry Pi 5 devices and the MediaPipe framework, developing optimized computer vision models for real-time, on-device applications (Figure 2). Unlike the foundational course, this module emphasized practical software development in Python, gradually shifting focus slightly away from physical computing.

Chunk 10 The advanced course aimed to introduce students to Physical AI pipelines with embedded machine learning models that control physical components such as sensors, actuators, and robotic systems.[4]

II. p. 156, Chunk 5: Integrating Physical AI and…
II. p. 159, Chunk 15: Integrating Physical AI and…
II. p. 157, Chunk 9: Integrating Physical AI and…

Chunk 11 This approach aligned with emerging practices in design and interaction, where AI have been used as a materially embodied agent capable of generating responsive, situated experiences.[5] For teaching and implementing an AI-based robotic pipeline, MediaPipe was selected due to its ease of use and minimal coding abstraction.[6] With only a few lines of code, students could implement computer vision models for face tracking, body tracking, and hand pose recognition.[7] Additionally, the MediaPipe Model Maker offers a straightforward approach for creating custom models. According to the developers, as few as 100 images per class are sufficient to fine-tune a pre-trained MediaPipe model.[8] For educational purposes, students created datasets consisting of approximately 600 to 800 images (with around 200 images per trained class), allowing them to fine-tune existing models to recognize specific gestures, facial expressions, or body movements. This enabled the development of customized interactive artistic experiences.

Chunk 12 After introductory sessions covering Python programming, serial communication, and MediaPipe, students further expanded their knowledge through a constructivist approach. By independently developing their own artistic projects, students actively built deeper coding skills and a practical understanding of the computer vision library. Furthermore, the advanced course introduced industrial robotics technology with a UR10e robotic arm, a cobot designed for safe human-machine collaboration (Figure 1). Students learned to program the robotic arm both via its intuitive teach pendant interface and through Python-based coding. This exploration aimed to inspire students, demonstrating advanced applications in media art and complex performance art projects.

3: Iovine, I. (2023). Integrating artificial intelligence and robotics into art curriculum. In *INFORMATIK 2023 - designing futures: Zukünfte gestalten* (pp. 347–352). Gesellschaft für Informatik e.V. <https://doi.org/10.18420/inf2023_32>

4: Miriyev, A., & Kovač, M. (2020). Skills for physical artificial intelligence. *Nature Machine Intelligence*, *2*(11), 658–660. <https://doi.org/10.1038/s42256-020-00258-y>

5: Coelho, M., & Labrune, J.-B. (2024). Large language objects: The design of physical ai and generative experiences. *Interactions*, *31*(4), 43–48. <https://doi.org/10.1145/3672534>

6: Iovine, I. (2024). Robotik und computer vision im künstlerischen Kontext. *Un/learn AI*. <https://unlearn.gestaltung.ai/article/dizhepje>

7: Lugaresi, C., Tang, J., Nash, H., McClanahan, C., Uboweja, E., Hays, M., Zhang, F., Chang, C.-L., Yong, M. G., Lee, J., Chang, W.-T., Hua, W., Georg, M., & Grundmann, M. (2019). *Mediapipe: A framework for building perception pipelines*. arXiv. <https://doi.org/10.48550/ARXIV.1906.08172>

8: Google. (n.d.). *Mediapipe model maker | google AI edge*. <https://ai.google.dev/edge/mediapipe/solutions/model_maker>.

II. p. 158, Chunk 10: Integrating Physical AI and…
II. p. 156, Chunk 6: Integrating Physical AI and…
I. p. 88, Chunk 14: Robotik und Computer Vision…

Chunk 13



Figure 1: Students testing a custom hand pose model on a UR10e robotic arm, with each gesture triggering a predefined movement. © Ivan Iovine

Outcomes and Lessons Learned

The implementation of a sequential two-semester course structure significantly enhances student engagement and learning outcomes.

Chunk 14 Dividing the curriculum into a foundational course in the winter semester and an advanced course in the summer semester proved essential. This format provided students with a clear and accessible entry point into physical computing, followed by a gradual introduction to more complex AI and robotics technologies such as Raspberry Pi 5 and MediaPipe.

Chunk 15 Computer vision, as an established technology, proved to be a particularly effective choice. Unlike other AI-based technologies, especially those involving generative models, it has not been subject to rapid architectural changes or disruptive codebase shifts.

Chunk 16 This technical stability has allowed for reliable instruction as well as tutorials, consistent learning outcomes, and a smoother integration into the overall teaching framework.

A fully equipped robotics lab was established to support this learning pathway. The lab was composed of two complementary sections: a seminar room

for theoretical instruction and a digital manufacturing area equipped with tools such as 3D printers and a laser cutter (Figure 2). This setup enabled students to move seamlessly between conceptual learning and hands-on experimentation.

Chunk 17 They gained not only technical knowledge in AI and robotics, but also practical skills in creating physical artifacts that integrated these technologies (Figure 3).



Figure 2: Robotics Lab at University of Art and Design Offenbach am Main. © Ivan Iovine

II. p. 158, Chunk 13: Integrating Physical AI and…
II. p. 157, Chunk 8: Integrating Physical AI and…
I. p. 65, Chunk 10: AI+D Lab (HfG Schwäbisch Gmünd)

Chunk 18

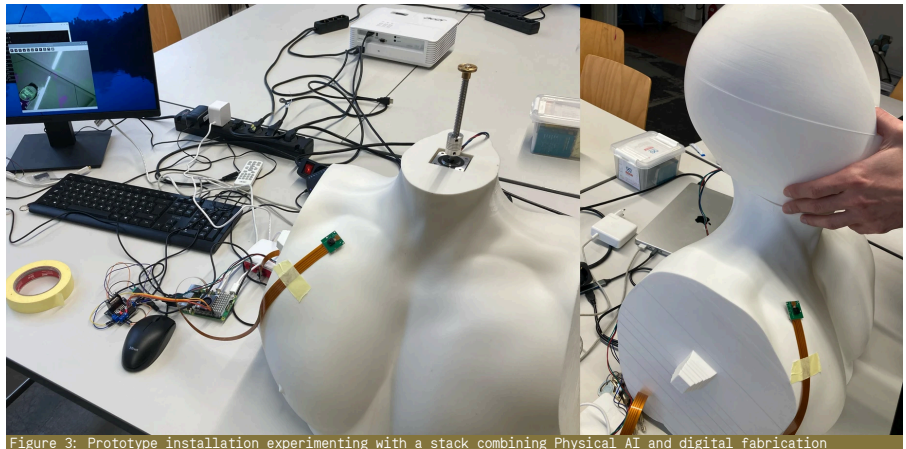

Figure 3: Prototype installation experimenting with a stack combining Physical AI and digital fabrication techniques. © Ivan Iovine

From the outset, the course also aimed to foster a critical and artistic understanding of robotic and interactive art. Students were introduced to exemplary works and DIY projects in the early sessions, which served as

inspiration for their own artistic exploration. This emphasis on conceptual framing and reference works, combined with a constructivist teaching approach, encouraged students to actively build knowledge by engaging creatively with both the technical and artistic aspects of the field. This pedagogical strategy helped to prevent the risk of technical overload, especially for students with little or no programming experience. By beginning with accessible tools like Arduino and gradually increasing technical complexity, students remained engaged without feeling overwhelmed. The progression allowed them to develop confidence and competence step by step.

II. p. 158, Chunk 11:
Integrating Physical AI and…
II. p. 157, Chunk 8:
Integrating Physical AI and…
II. p. 156, Chunk 5:
Integrating Physical AI and…

Chunk 19 The results of this structured approach were evident in the students' final projects. Those who participated in both semesters demonstrated notable improvements in technical proficiency and creative expression. Many projects successfully implemented a Physical AI pipeline by combining computer vision with physical computing components. As depicted in Figure 4, examples included interactive sculptures using real-time vision to control servo mechanisms, as well as robotic installations operated by custom-trained machine learning models running locally on Raspberry Pi 5 devices.[9] As illustrated in Figures 5 and 6, experienced students also integrated industrial robotic arms into their work, especially in media and performative arts.[10] Through independent project work and guided exploration, students developed technical skills, autonomy, and the ability to design and implement original, technologically advanced artworks.

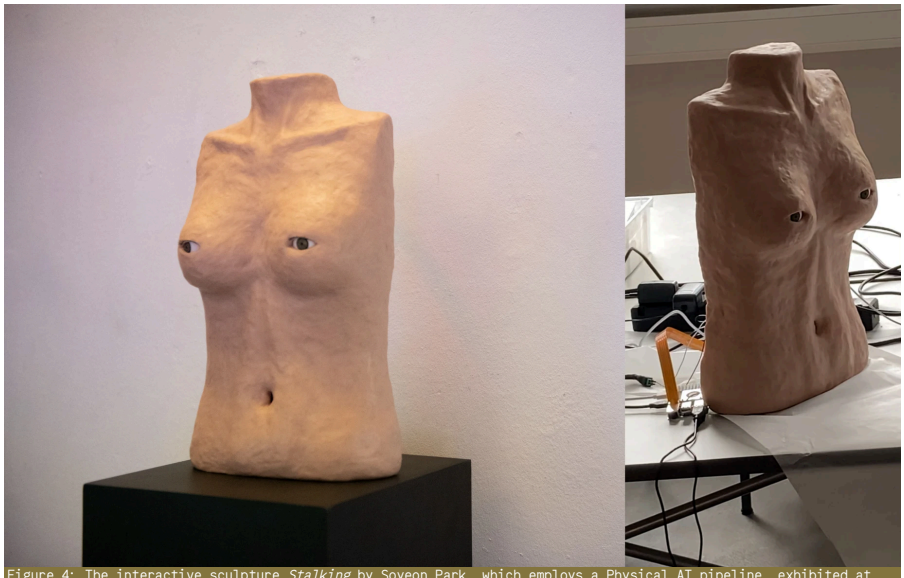

Figure 4: The interactive sculpture *Stalking* by Soyeon Park, which employs a Physical AI pipeline, exhibited at the Correlations 2024 conference as part of the KITeGG research project. © Cheesoo Park, Ivan Iovine

9: Iovine, I. (2024). Robotik und computer vision im künstlerischen Kontext. *Un/learn AI*. <https://unlearn.gestaltung.ai/article/dizhepje>

10: Iovine, I. (2024). Robotik und computer vision im künstlerischen Kontext. *Un/learn AI*. <https://unlearn.gestaltung.ai/article/dizhepje>

Chunk 20



Figure 5: The media art installation *Door Bitch* by Rahel Pabst, exhibited during the annual Rundgang and the Correlations 2024 conference as part of the KITeGG research project. © Rahel Pabst, Cheesoo Park.

I. p. 111, Chunk 18: Creating easy to use interfaces…
II. p. 245, Chunk 34: Der KITeGG Cluster – eine…
I. p. 10, Chunk 7: Projektvorstellung

Chunk 21



Figure 6: The performative installation *Schlag auf Schlag* by Nelli Gomez Baumert, exhibited during the Correlations 2024 conference as part of the KITeGG research project. © Cheesoo Park.

Chunk 22

Conclusion

Integrating AI, computer vision, and robotics into art and design curriculum involves rethinking both teaching structures and learning environments. The experience at the Robotics Lab of the University of Art and Design Offenbach am Main shows that it is possible to introduce complex technologies into creative contexts without compromising accessibility or artistic exploration. By beginning with physical computing and gradually expanding to advanced AI tools, students are able to engage with these technologies at a sustainable pace.

Chunk 23 The two-semester structure, combined with a constructivist teaching approach and dedicated infrastructure, provides a solid foundation for experimentation, technical growth, and artistic development. Throughout the program, students not only acquire new technical skills, but also develop a critical and cultural understanding of the role of intelligent systems in contemporary artistic practice.

Chunk 24 Their final projects demonstrate how emerging technologies can be meaningfully integrated into creative workflows, producing original and context-aware artworks. These outcomes underscore that, with thoughtful pedagogical design, robotics and AI can become integral components of creative education.

This article is separately published as 10.25358/openscience-13022.

↗ t0rgtqez

Building AI Intuition – Four Educational Pillars for Teaching AI in Design at AI+D Lab

Rahel Flechtner, Jordi Tost

III. p. 3, Chunk 2: Editor`s Note
III. p. 12, Chunk 1: AI Literacy for the Long Haul: …
I. p. 113, Chunk 3: Developing an AI teaching…

Chunk 1

In 2023, we outlined our ambition to provide students with an *AI Intuition*—a broad understanding of AI's working principles and limitations.[1] With this intuition, the AI+Design Lab[2] at HfG Schwäbisch Gmünd aimed to provide the knowledge needed to effectively engage with AI at different levels: Students should be able to assess the capabilities of machine learning technologies and models, develop a precise language to communicate AI concepts, decide when to use AI for a design goal based on costs and relevance; they must adopt and build explorative strategies with the technology to work on new applications, use cases, and future scenarios; and, finally, they should be able to critically evaluate societal, ethical, or environmental impacts of implementing AI technologies. With this, our aim has been to prepare students to use AI as design material and a tool—meaningfully—at different stages of their design processes,

1: Flechtner, R., & Stankowski, A. (2023). AI Is Not a Wildcard: Challenges for Integrating AI into the Design Curriculum. *Proceedings of the 5th Annual Symposium on HCI Education*, 72–77. <https://doi.org/10.1145/3587399.3587410>

2: The AI+D Lab: A research group situated between AI, design education, and technology research.

and when to rely on their human intuition. This article details our four educational pillars for building *AI Intuition* and summarizes insights from their implementation over the past three years.

Four Educational Pillars for Building *AI Intuition*

To support *AI intuition*, we defined four educational pillars: (1) improving technical literacy, (2) fostering hands-on exploration, (3) enabling conceptual engagement, and (4) encouraging critical reflection. Building on these pillars, we developed and tested various teaching modules, including introductory materials, self-study assignments, code elements, and custom tools and methods.

Improving Technical Literacy

Having an overview of the broad technical possibilities of AI is essential for practical engagement at the intersection of AI and design. To introduce students to the technical aspects of AI, we created a series of presentations covering a range of topics and complexities: from general overviews of AI technologies and their applications to foundational principles of neural networks, sensor-based AI, image generation models, and large language model (LLM) architectures. We also developed self-study formats that allowed students to explore topics at their own pace—either as preparation for courses or for individual study.

Chunk 2 To introduce students to the fundamentals of machine learning for sensor-based interaction in a playful and explorative way, we developed the role-playing activity *Acting out AI Systems*.[3] Students embodied different components of an AI system.

Chunk 3 By going through the relevant steps for training and inference, they gained a basic understanding of the technology's functionality and limits (Figure 1).

Chunk 4



Figure 1: Students acting as human sensors and recording a data stream of a thermal camera on paper (left). Team members act out the AI algorithm and try to make sense of the data stream (right).

3: Flechtner, R., & Kilian, J. (2024). Making (Non-)Sense–A Playful and Explorative Approach to Teaching AI Intuition for the Design of Sensor-Based Interactions. *Proceedings of the 6th Annual Symposium on HCI Education*, 1–9. <https://doi.org/10.1145/3658619.3658643>

Another approach we developed to introduce students to the functional principles of neural networks is *SandwichNet*.[4] [5] This interactive tool visualizes neural networks in a playful format, allowing students to manipulate the parameters to recreate and understand its learning behavior (Figure 2). The tool includes multiple complexity levels and is used across foundational courses.

Chunk 5

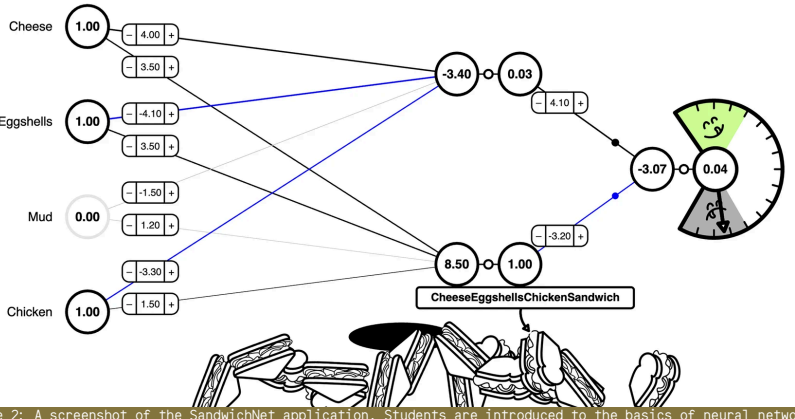

Figure 2: A screenshot of the SandwichNet application. Students are introduced to the basics of neural networks by training SandwichNet to distinguish between edible and inedible sandwiches.

Chunk 6

Fostering Hands-on Exploration

The complexity and non-deterministic nature of AI technologies make them difficult to approach as design material, limiting creative processes like *reflection-in-action*,[6] where ideation unfolds through direct engagement with the material.

Chunk 7 However, these hands-on processes facilitate the development of new and innovative applications and encourage critical thinking. This requires a tacit understanding of the material.[7] We therefore consider hands-on engagement and a learning-through-making mindset as key to grasping AI's functional principles and enabling creative exploration.

At HfG Schwäbisch Gmünd, one area of focus is Physical AI—machine learning on microcontrollers that interact with their environment through sensors and actuators.[8] Due to the limited computing power of this hardware, the models are small, making their functionality easier to grasp.

Chunk 8 The hardware's low cost increases accessibility. Physical AI thus provides a practical entry point to machine learning. Sensor-based data allows for fast training cycles and immediate feedback, supporting iterative, real-world experimentation.

4: Sewing, F., & Flechtner, R. (2024). *Demonstrating SandwichNet –A Playful Tool for Teaching the Basics of Neural Networks*. <https://doi.org/10.18420/MUC2024-MCI-DEMO-322>

5: Sewing, F., & Stankowski, A. (2024). SandwichNet. In *AI+Design Lab*. <https://aid-lab.hfg-gmuend.de/articles/sandwich-net/>

6: Schön, D. A. (1984). *The reflective practitioner: How professionals think in action*. Basic books.

7: Yang, Q., Banovic, N., & Zimmerman, J. (2018). Mapping Machine Learning Advances from HCI Research to Reveal Starting Places for Design Innovation. *Proceedings of the 2018 CHI Conference on Human Factors in Computing Systems*, 1–11. <https://doi.org/10.1145/3173574.3173704>

8: Flechtner, R., Kilian, J., & Iovine, I. (n.d.). Sensor-Based AI in Art and Design. In *Un/learn ai : Navigating AI in aesthetic practices* (Vol. 2). <https://doi.org/10.25358/openscience-11830>

We use Edge Impulse,[9] an online platform that simplifies data collection, labeling, model building, and training. Students are introduced to the process through a sample project using the Arduino Nano 33 BLE Sense microcontroller,[10] training a model to distinguish movement gestures. This hands-on approach enables them to carry out the entire data collection and training process, even within short teaching formats.

Chunk 9

To reinforce the relevance of the technology in design and support the exploration of novel forms of human-machine interaction, we then set specific thematic tasks. For example, in one workshop, students explored creative and unconventional ways of engaging with devices and simple games.[11] Results included a pressure-gesture-controlled lamp using barometer data or a sound-based snake game controller (Figure 3).

Chunk 10



Figure 3: A sound-based snake game controller by Ron Eros Mandic and Lukas Speidel.

II. p. 205, Chunk 9: Shaping Human-AI Relationships
II. p. 149, Chunk 4: Making (Non)Sense
II. p. 191, Chunk 35: SandwichNet

Chunk 11

Two major challenges in hands-on exploration with sensor-based AI in teaching are students' limited coding experience and the technical demands of physical prototyping.

Chunk 12

To address this, we created reusable code elements and technical setups.

II. p. 211, Chunk 10: Uncanny Type
I. p. 86, Chunk 3: Robotik und Computer Vision…

Chunk 13

The *code elements*[12] are structured for easy adaptation and offer sensing and acting possibilities for interactive AI-based systems. This provides a boilerplate to students, who can turn their concepts into interactive prototypes. The input options include motion, sound, barometer, camera, and thermal camera sensors; output options include onboard RGB LEDs, radio sockets, and keyboard emulation. Radio sockets

9: AI for Any Edge Device. (n.d.). In *Edge Impulse*. <https://edgeimpulse.com/>

10: Nano 33 BLE Sense. (n.d.). In *Arduino Docs*. <https://docs.arduino.cc/hardware/nano-33-ble-sense/>

11: Flechtner, R. (2023). Controller: Lab Week Workshop 2023. In *AI+Design Lab*. <https://aid-lab.hfg-gmuend.de/articles/lab-week-23-controller/>

12: Flechtner, R., Sewing, F., & Tost, J. (2025). Code Elements Input Output. In *GitLab - KITeGG*. <https://gitlab.rlp.net/kitegg/public/hfgsg/code-elements-input-output/>

allow control of any device connected to a 240V outlet. The keyboard emulation can trigger any computer function that can be activated by a single keystroke and serves as a bridge to p5.js and ml5.js projects. Sample code supports interaction with LLMs and text or speech output.

The technical setup of interactive systems involving sensors and actuators also poses a challenge in teaching.

Chunk 14 Therefore, we developed *ready-made sensing and acting modules* together with colleagues at the Köln International School of Design (KISD).[13] [14] These predesigned and preprogrammed modules integrate sensors and actuators, enabling a quick start for training and prototyping. Communication between input and output modules occurs via an analog LED interface, which visualizes the system status and simplifies debugging. Sensor modules include motion, sound, piezo, distance, and light sensors; actuator modules feature servo motors, stepper motors, and radio sockets.

Another relevant topic in our teaching is the field of generative AI.

Chunk 15 The HfG Schwäbisch Gmünd has been involved in initiatives to *make generative AI tools and state-of-the-art models available to students*. Among others, the university is part of the bwGPT[15] initiative, which enables students to freely access powerful language models with their university account. Additionally, we host an image generation service based on Stable Diffusion at our lab and offer an experimental interface, an API, and a Slack bot to make image generation easy and accessible.

Chunk 16 In addition to the provision of services, *we developed custom-made tools* that enable students to enter into a deep, hands-on engagement with the technology as material. The *Prompt Diary* (Figure 4) enables an iterative and explorative approach to prompting images. Students can iteratively adapt prompts, create variations, and visually document the process. Using a sticky-note metaphor, prompts can be expanded and combined. The output appears as image tiles and can be reused as a starting point or as a ControlNet embedding.

13: Flechtner, R., & Kilian, J. (2024). Making (Non-)Sense–A Playful and Explorative Approach to Teaching AI Intuition for the Design of Sensor-Based Interactions. *Proceedings of the 6th Annual Symposium on HCI Education*, 1–9. <https://doi.org/10.1145/3658619.3658643>

14: Kilian, J., & Flechtner, R. (2024). Making NonSense. In *GitLab - KITeGG*. <https://gitlab.rlp.net/kitegg/public/making-nonsense/>

15: *bwGPT*. (n.d.). <https://www.zml.kit.edu/bwgpt.php>

Chunk 17

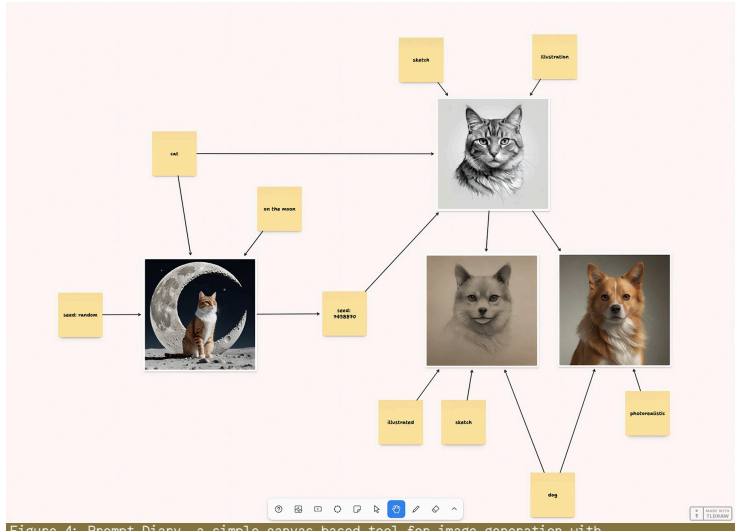

Figure 4: Prompt Diary, a simple canvas-based tool for image generation with Stable Diffusion.

Another tool, *Transferscope*,[16] explores the capabilities of image generation models in combination with ControlNets—models that enable more precise control over generated content, such as through edge or line detection. Transferscope is both a physical device made with a Raspberry Pi and 3D-printed parts (Figure 5, left) and a web platform for students to engage with in our courses. With a one-button interface, users can capture the style of an object or concept and transfer it into another object or scene (Figure 5, right), encouraging playful, creative engagement with generative AI.[17]

Chunk 18

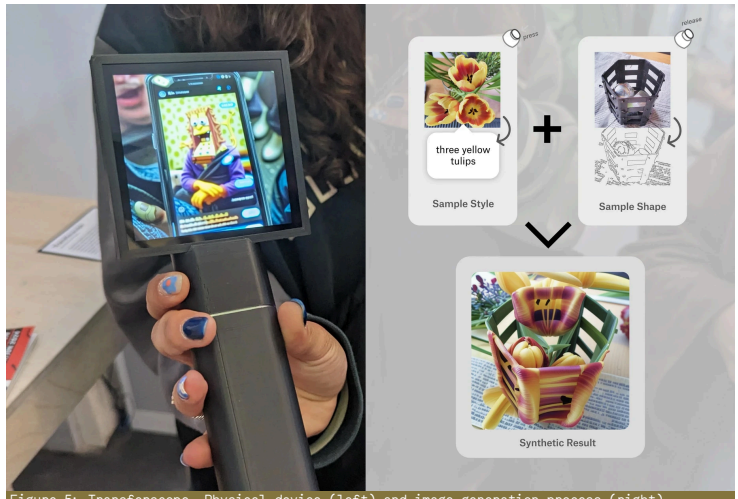

Figure 5: Transferscope. Physical device (left) and image generation process (right). Project by Christopher Pietsch.

II. p. 180, Chunk 7: Editorial
III. p. 3, Chunk 2: Editor`s Note
II. p. 34, Chunk 13: Trier

Chunk 19

Enabling Conceptual Engagement

Developing interesting and meaningful AI applications poses a challenge for designers.

16: Pietsch, C. (2024). Transferscope – Synthesized Reality: Sample anything. Transform everything. In *AI+Design Lab*. <https://aid-lab.hfg-gmuend.de/articles/transferscope/>

17: Pietsch, C., & Stankowski, A. (2024). *Transferscope – Making Multi-Modal Conditioning for Image Diffusion Models Tangible*. <https://doi.org/10.18420/MUC2024-MCI-DEMO-248>

[18] To support ideation, we created methods and courses that foster creative friction and encourage conceptual engagement with AI in playful, unconventional ways.

One approach involves imagination cards,[19] inspired by creativity theories and artistic movements like Fluxus or Oulipo. These digital card decks, tailored to different course topics, constrain and provoke unusual idea generation. In a workshop, cards are randomly assigned to students to create concepts based on the given variables. While a particular card combination narrows the possibilities to one specific setting and ensures focus, unconventional combinations encourage creative thinking and exploration of unusual concepts. One example is the *Input/Output* imagination cards, which provide a combination of an input technology (e.g., a sensor) and an output medium for students to consider technical deployment while drafting potential ideas (see Figure 6 for an example project inspired by the cards).



Figure 6: A concept created with the Input/Output imagination cards: Gloomy, a lamp which, based on facial expressions (input), provides visual and auditory feedback (output), by Rebeka Tot and Vivien Cai.

I. p. 195, Chunk 4: Unpacking the Language of…
I. p. 189, Chunk 3: Exploring Tools
I. p. 197, Chunk 9: Unpacking the Language of…

Chunk 20

Another category uses metaphors as a design resource for developing AI interactions.

I. p. 195, Chunk 4: Unpacking the Language of…
I. p. 197, Chunk 8: Unpacking the Language of…
I. p. 195, Chunk 1: Unpacking the Language of…

Chunk 21

Metaphors are a potential means for both understanding the complex entanglements of AI as well as for exploring new ways of interacting with intelligent products.[20] [21] This supports *otherness*, that is, the design of novel roles, affordances, or interaction paradigms beyond conventional anthropomorphism and zoomorphism.

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

19: Imagination Cards. (n.d.). In *AI+Design Lab HfG Schwäbisch Gmünd*. <https://imagination-cards.aid-lab.hfg-gmuend.de/>

20: Murray-Rust, D., Lupetti, M. L., & Nicenboim, I. (2024). *Metaphor Gardening: Experiential engagements for designing AI interactions*. <https://doi.org/10.21606/drs.2024.376>

21: Tost, J., Flechtner, R., Maué, R., & Heidmann, F. (2024). Caring for a companion as a form of self-care. Exploring the design space for irritating companion technologies for mental health. *Nordic Conference on Human-Computer Interaction*, 1–15. <https://doi.org/10.1145/3679318.3685343>

[22] In the Caring for Machines workshop,[23] students reimagined interactions with AI beyond the *technology as servant* metaphor, exploring how caring for a technology can inspire more responsive and reciprocal interactions (see Figure 7 for an example project).



Figure 7: Social Spark detects awkward silences in conversations and fills them with random facts or jokes. Project by Philipp Maginot and Sebastian Weinert.

Chunk 22 A third approach uses *experimental and performative methods*, such as role-play, improvisation, and object-centered perspectives, to explore the nuances of human-AI relationships.[24] Methods such as *thing ethnography*,[25] [26] *techno-mimesis,*[27] and *object persona*[28] (Figure 8) encourage students to adopt the perspective of technological systems, revealing unique affordances and inspiring new design ideas. Performative methods such as role-play and improvisation facilitate conceptual exploration of human-AI interactions—without ignoring the technical complexity nor requiring in-depth technical knowledge (Figure 9).

22: Hassenzahl, M., Borchers, J., Boll, S., Pütten, A. R. der, & Wulf, V. (2021). Otherware: How to best interact with autonomous systems. *Interactions*, *28*(1), 54–57. <https://doi.org/10.1145/3436942>

23: Flechtner, R., Sewing, F., & Tost, J. (2025). Caring for Machines. In *AI+Design Lab*. <https://aid-lab.hfg-gmuend.de/articles/caring-for-machines/>

24: Flechtner, R. (2024). Shaping Human-AI Relationships – Workshop. In *AI+Design Lab*. <https://aid-lab.hfg-gmuend.de/articles/shaping-human-ai-relationships/>

25: Giaccardi, E., Cila, N., Speed, C., & Caldwell, M. (2016). Thing Ethnography: Doing Design Research with Non-Humans. *Proceedings of the 2016 ACM Conference on Designing Interactive Systems*, 377–387. <https://doi.org/10.1145/2901790.2901905>

26: Nicenboim, I., Giaccardi, E., & Kuijer, L. (2018). Designing Connected Resources for Older People. *Proceedings of the 2018 Designing Interactive Systems Conference*, 413–425. <https://doi.org/10.1145/3196709.3196808>

27: Dörrenbächer, J., Löffler, D., & Hassenzahl, M. (2020). Becoming a Robot - Overcoming Anthropomorphism with Techno-Mimesis. *Proceedings of the 2020 CHI Conference on Human Factors in Computing Systems*, 1–12. <https://doi.org/10.1145/3313831.3376507>

28: Cila, N., Giaccardi, E., & Tynan-O`Mahony, F. (2015). Thing-Centered Narratives: A study of object personas. *Proceedings of the 3rd Seminar International Research Network for Design Anthropology, Online proceedings*, pp. 1–17.

Chunk 23

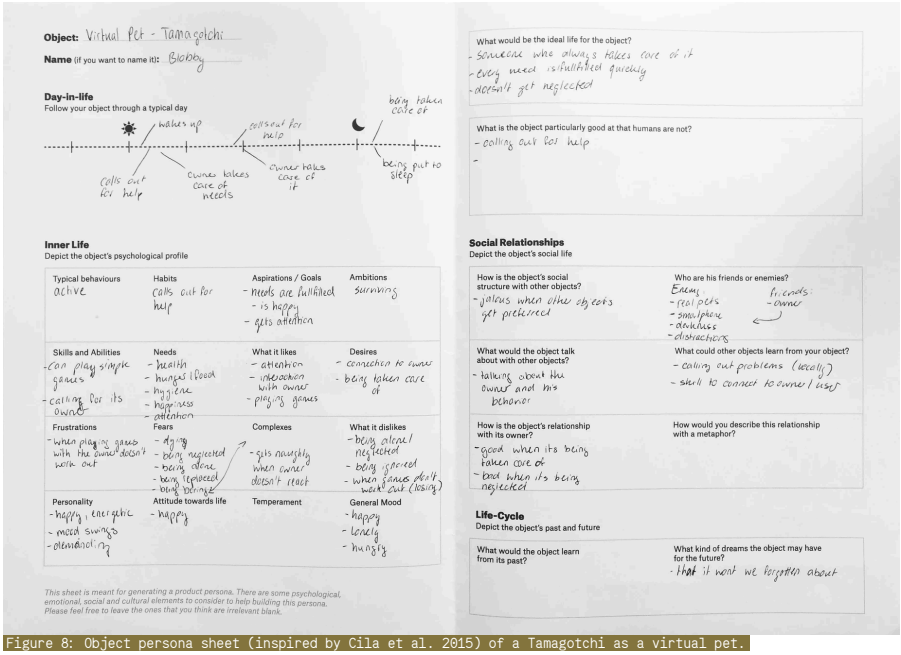

Figure 8: Object persona sheet (inspired by Cila et al. 2015) of a Tamagotchi as a virtual pet.

Chunk 24



Figure 9: Students role-playing a fitness tracker.

Chunk 25 The fourth approach uses *generative AI as a co-designer and counterpart for ideation*. We have developed tools and processes that playfully support imagination through the co-creation of speculative scenarios and fictional design objects, in collaboration with image generators and LLMs. To explore co-creation and collaborative future-making with generative AI, we developed the course *Futures Lenses: Co-speculating Things with AI*.[29] We used a hands-on approach combining critical design strategies—like anti-solutionism, irony, and absurdity—with simple no-code AI tools such as *Transferscope* and the *Futures Lens*. By deliberately applying constraints, ambiguity, or exaggeration with a prompt, the *Futures Lens* course was

29: Tost, J. (2024). Futures Lenses. In *AI+Design Lab*. <https://aid-lab.hfg-gmuend.de/articles/futures-lenses/>

meant to provoke errors and unexpected results in the output images. These *accidents* became a source of critical and creative thinking by helping students defamiliarize everyday objects, expose paradoxes and biases, and get inspiration for new design ideas through iterative exploration (see Figure 10 for an example project).



Figure 10: Calm Futures. A speculative concept (left) and a catalog of fictional things (right), created collaboratively by Anja Gutmann, Lea Haferbier, Carolin Kaltwasser, and the Futures Lens. The concept shows a lamp for promoting well-being through symbiotic, reciprocal interactions.

Chunk 26

Encouraging Critical Reflection

A significant part of the critical examination of AI technologies and their societal, ethical, and ecological implications occurred through theoretical input and discussion. For instance, various courses addressed the intentional or unintentional shaping of human-machine relationships through design, discussing the benefits and risks of employing anthropomorphic or zoomorphic approaches. While such designs may foster trust and connectedness, they also risk creating false expectations or overreliance on AI systems.

In other formats, hands-on engagement with the technology was used to initiate critical discussion and reflection on emerging phenomena. Topics included biases in the data sets of generative models, computational and environmental costs of AI, the evolving role of designers in light of increased automation of design processes, and the responsibility of designers in shaping these processes and tools.

To further promote critical thinking, we developed specific methods. One example is *Fake or Real* (see Figure 11), a teaching activity developed by Lisa Kern in collaboration with the AI+Design Lab. Students were presented with real and fake facts and used them to generate texts and images for convincing social media posts. Afterwards, they were challenged to identify the fake posts. The method served to reflect on how easily fake news can be produced, what indicators may suggest inauthenticity, and how the authenticity of information can be verified.

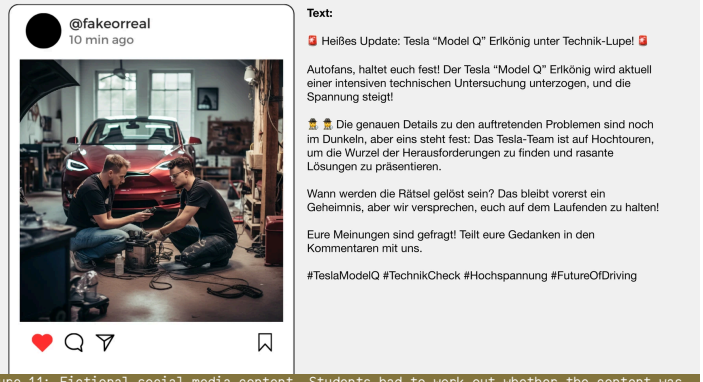

Figure 11: Fictional social media content. Students had to work out whether the content was fake or real.

Chunk 27 One tool that we developed specifically for critically engaging with machine learning models and datasets is the *unStable Mirror* project (see Figure 12). The mixed-media installation highlights the potential of generative AI, while also addressing its limitations and inherent biases.[30]

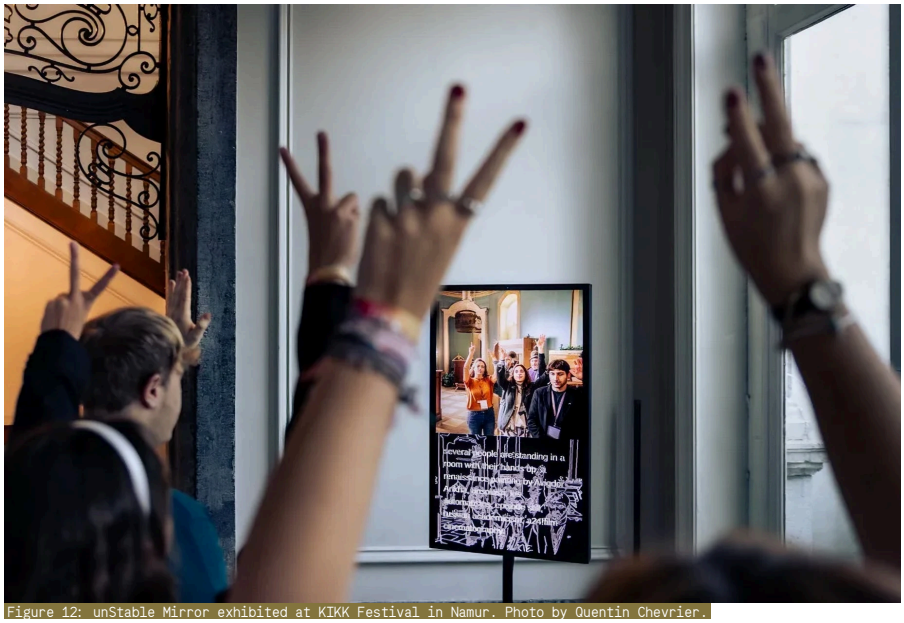
Figure 12: unStable Mirror exhibited at KIKK Festival in Namur. Photo by Quentin Chevrier.

Chunk 28 Moreover, exploring possible technological futures with AI is a potential strategy for supporting critical reflection on common, widespread AI narratives.[31] [32] Embracing this notion, the *Futures Lenses* course, besides exploring and co-ideating novel design concepts, also aimed at exposing stereotypical future visions embedded in generative AI systems. Recognizing and questioning these mainstream future visions is key skill as generative AI becomes more prevalent in design education and industry for ideation.

30: Pietsch, C. (2023). unStable Mirror: A Journey into the Heart of AI Generativity and Bias. In *AI+Design Lab*. <https://aid-lab.hfg-gmuend.de/articles/unstable-mirror>

31: Tost, J., Gohsen, M., Schulte, B., Thomet, F., Kuhn, M., Kiesel, J., Stein, B., & Hornecker, E. (2024). Futuring Machines: An Interactive Framework for Participative Futuring Through Human-AI Collaborative Speculative Fiction Writing. *ACM Conversational User Interfaces 2024*, 1–7. <https://doi.org/10.1145/3640794.3665904>

32: Ghajargar, M. (2024). AI and Future-Making: Design, Biases, and Human-Plant Interactions. *Proceedings of the 27th International Academic Mindtrek Conference*, 24–35. <https://doi.org/10.1145/3681716.3681738>

Insights from Three Years of Teaching at the AI+Design Lab

Over the past three years, we tested a variety of teaching formats—from compact workshops to extended project courses—grounded in our four educational pillars. Each format varied in its emphasis.

Chunk 29 Some focused on technical literacy, deepening it through hands-on engagement with the technology (e.g., *Controller*[33] or *Making(Non-)Sense*[34]). Others focused on conceptual engagement, either introducing it through hands-on engagement with the technology (e.g., *Futures Lenses. Co-speculating Things with AI*,[35] or *Caring for Machines*[36]) or starting from conceptual engagement and then moving on to critical reflection (e.g., *Shaping Human-AI Relationships*[37]). A foundational course attempted to address all four areas, though not all were covered in equal depth.

II. p. 206, Chunk 11: Shaping Human-AI Relationships
III. p. 10, Chunk 5: Introduction
II. p. 202, Chunk 5: Editorial: Show and Tell

Chunk 30 *Providing technical and conceptual depth requires sequential formats*. Balancing technical literacy and conceptual engagement proved challenging within time-limited formats. Students in technical courses sought conceptual application, while those in concept-driven formats wanted deeper technical knowledge. Experimental approaches, such as role-play, offered accessible entry points to the technical foundations, but the abstraction sometimes limited comprehension and translation into realistic design concepts. Tools like ideation cards and input-output cards helped guide ideation, yet students often struggled to move beyond simplistic or unrealistic concepts. These insights point to the need for consecutive formats—starting with foundational technical knowledge deepened through hands-on engagement, followed by advanced conceptual exploration and critical reflection.

Different stages of the creative process require different types of AI tools. Early, exploratory phases benefit from simple, no-code interfaces—they're easy to use, encourage experimentation, and often lead to surprising results that spark new ideas and enable critical thinking. Taking image generation with Stable Diffusion as an example, useful no-code tools include simple custom interfaces for text-to-image generation, the *Prompt Diary* canvas-based tool, or tools such as *Transferscope* or the *Futures Lens*. Using simple tools in early stages also fosters curiosity about the technology, helping students identify what they need and want to learn.

Chunk 31 As projects progress, more complex tools (e.g., ComfyUI) allow for greater control and precision. These tools help students refine details and customize outcomes.

Co-creating and co-speculating with generative AI beyond the perfect prompt: openness and a new perspective towards AI are required. The

33: Flechtner, R. (2023). Controller: Lab Week Workshop 2023. In *AI+Design Lab*. <https://aid-lab.hfg-gmuend.de/articles/lab-week-23-controller/>

34: Kilian, J., & Flechtner, R. (2024). Making NonSense. In *GitLab - KITeGG*. <https://gitlab.rlp.net/kitegg/public/making-nonsense/>

35: Tost, J. (2024). Futures Lenses. In *AI+Design Lab*. <https://aid-lab.hfg-gmuend.de/articles/futures-lenses/>

36: Flechtner, R., Sewing, F., & Tost, J. (2025). Caring for Machines. In *AI+Design Lab*. <https://aid-lab.hfg-gmuend.de/articles/caring-for-machines/>

37: Flechtner, R. (2024). Shaping Human-AI Relationships – Workshop. In *AI+Design Lab*. <https://aid-lab.hfg-gmuend.de/articles/shaping-human-ai-relationships/>

increasing use of mainstream AI tools (e.g., ChatGPT, Midjourney) has established the expectation that designing with AI depends on writing the *perfect prompt*. This turns the creative process into a rational sequence of instructions and clashes with the non-linear, exploratory nature of design processes. Teaching AI intuition requires shifting this mindset: students must learn to embrace generative AI as a co-creator—an agent of surprise and creative friction that can enrich the design process. In *Futures Lenses*, students were initially frustrated by the lack of control over generated outcomes. However, as the process unfolded, many students began to understand the creative potential in these unexpected results. One student noted, "We were frustrated by the mistakes, but these mistakes brought out the best ideas of our project." We conclude that collaborating with AI in co-creation and co-speculation involves learning to let go of control and accepting that the results will not always match the original idea. Instead of seeing this as a technical limitation, students should learn to view it as a chance to enrich the creative process, uncover new perspectives, and open up new directions. Teaching students this shift in mindset can help them move past preconceived expectations and engage more fully with the potential of AI as a creative counterpart.

AI-generated concepts provide a surface for critically reflecting on mainstream AI metaphors and future narratives related to AI, as well as roles that we stereotypically assign *by default* to intelligent products and agents. These stereotypical narratives or roles, which are spread by industry commercials or pop culture media, such as films or literature, are strongly embedded in our collective imaginaries.

Chunk 32 Acknowledging this fact and critically reflecting on these visions is part of enabling a critical mindset. While students often replicate such stereotypes in their preliminary ideas and concepts, seeing them echoed in AI-generated outputs enables a critical distance.

Experimental methods hold potential for teaching AI in design. Experimental methods, such as the object-centered approaches, facilitated a change of perspective that enabled students to break away from solution-oriented approaches. This encouraged in-depth discussion of the ethical and moral dimensions of AI systems and underscored the designer's responsibility in shaping these technologies.

Chunk 33 Acknowledgements

All modules described were (co-)developed at the AI+Design Lab of the HfG Schwäbisch Gmünd. The following current and former members of the lab were involved in the development: Benedikt Groß, Jordi Tost, Rahel Flechtner, Aeneas Stankowski, Felix Sewing, Christopher Pietsch, Moritz Hartstang, Alexa Steinbrück, Ron Mandic, and Maxime Beck.

This article is separately published as 10.25358/openscience-13021 .

↗ ja6f4xwe

Motion Experience Lab@(A)IDLE MODE: Echtzeit-Installation zur performativen KI Erfahrung – Embodied AI Learning in Design Education

Nina Juric

Chunk 1

Einblicke in ein angewandtes Designforschungsprojekt des Motion Experience Labs im Kontext der Hidden Layers Konferenz 2024

What is an (A)IDLE MODE? Wie können wir uns den latenten Zuständen innerhalb komplexer, technischer Systeme und Infrastrukturen annähern, sie erfahrbar machen und für gesellschaftlichen Dialog öffnen? Als Zugang zu dieser Frage entwickelte das ↗ Motion Experience Lab (MXL) der Köln International School of Design (KISD) eine immersive Rauminstallation mit dem Titel *Motion Experience Lab@(A)IDLE MODE*. Unter der Leitung von Prof. Nina Juric entstand in Zusammenarbeit mit 23 Studierenden ein 200m² großer Erfahrungsraum, der die Narration einer Künstlichen Intelligenz (KI) räumlich und performativ inszenierte. Das Team realisierte das Projekt als angewandte Lehrforschung im Gebiet *Image&Motion* des Fachbereichs Integrated Design, mit technischer Unterstützung von Marko Bodenstein und Henning Brode. Die Installation wurde als abendliches Live-Event im Rahmen der HIDDEN LAYERS AIxDESIGN CONFERENCE (12.–14. Juni 2024, TH Köln) am 13. Juni 2024 öffentlich präsentiert.

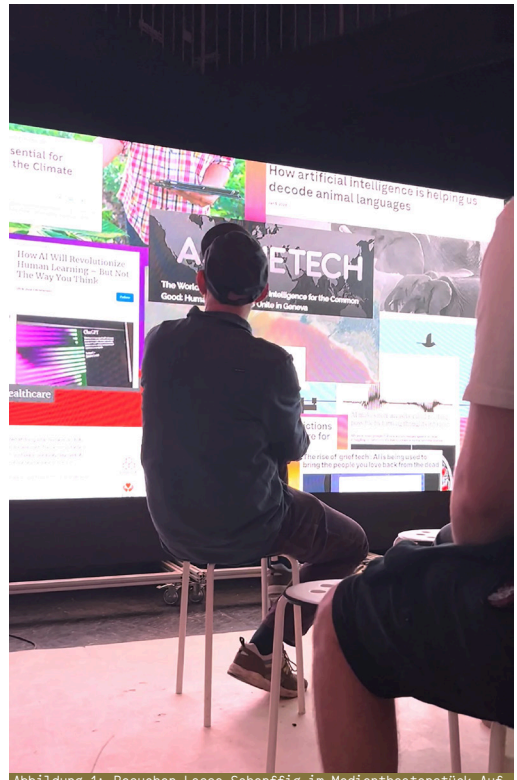

Abbildung 1: Besucher Lasse Scherffig im Medientheaterstück Auf der Datenbank. © MXL Nina Juric

Eingebettet in das Konferenzprogramm konnte das MXL als experimentelle Lernumgebung sein kreativ-technologisches Potenzial zum Einsatz bringen.

Das MXL-Team entwickelte das Projekt als interdisziplinäres Lehrformat, das forschungsbasierte Praxis mit angewandter Gestaltungslehre in einem realen und öffentlichen Ausstellungskontext verband. Im Zentrum der inhaltlichen Konzeption stand die Frage nach dem *(A)IDLE MODE,* die im Projekt gestalterisch-künstlerisch, technisch-prototypisch und konzeptuell zugleich bearbeitet wurde. KI wurde dabei sowohl als inhaltlicher wie als produktionstechnischer Akteur integriert.

Chunk 2 *Was ist ein Idle Mode?*

Chunk 3 Idle Mode bezeichnet einen Ruhezustand oder eine Phase der Inaktivität, die viele Systeme –darunter technologische Geräte– einnehmen können. In diesem Zustand verbrauchen sie weniger Energie und Ressourcen, bleiben jedoch funktionsbereit. Häufig laufen im Hintergrund weiterhin verborgene Prozesse, die auf neue Impulse warten.

Der Begriff *(A)Idle Mode* diente hier als konzeptueller Ausgangspunkt, um latente technologische Zustände[1] und infrastrukturelle Imaginationen innerhalb von KI-Modellen gestalterisch zu erforschen. Sieben Stationen entfalteten dieses kreative Potenzial in einer choreografierten, sinnlich erfahrbaren Szenografie aus Bild, Text, Klang, Licht und Körper – ein narratives Echtzeit-Setting, das Designausbildung als kollektiven Denk- und Gestaltungsraum sichtbar machte. Die Installation verknüpfte unterschiedliche Medien zu einer multisensorischen KI-Erfahrung und stellte Fragen nach Präsenz, Instabilität und Synchronizität im Zusammenspiel von Mensch, Maschine und Raum.[2]

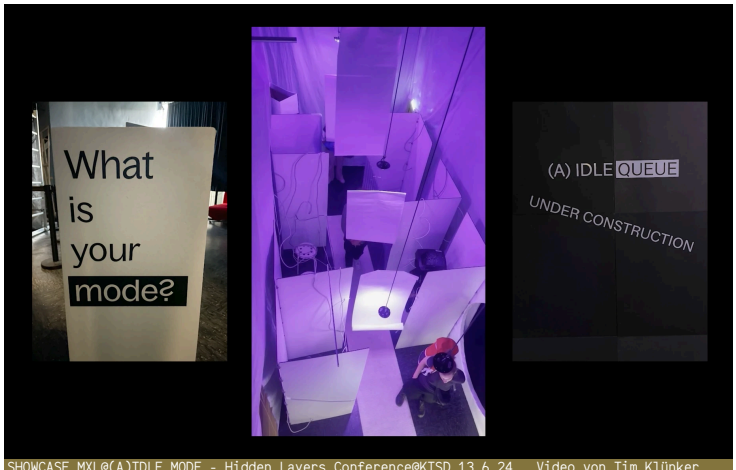

SHOWCASE MXL@(A)IDLE MODE - Hidden Layers Conference@KISD 13.6.24 Video von Tim Klünker und Léon Scheunemann Link: https://vimeo.com/1088380087

1: Latente Zustände sind verborgene, potenzielle oder inaktive Zustände in meist technologischen Systemen, die nicht direkt sichtbar oder spürbar sind, aber jederzeit aktiviert oder sichtbar gemacht werden können. In kreativen, technischen oder AI-nahen Kontexten bezieht sich „Latenz" oft auf das, was im Hintergrund wirkt, wartet oder vorbereitet wird – ohne unmittelbar aktiv zu sein. Ein latenter Zustand ist eine vorhandene Möglichkeit, die aktuell noch nicht aktiviert, aber prinzipiell wirksam ist. In interaktiven Installationen kann ein latenter Zustand ein ungetriggertes Szenario sein – z. B. ein Sensor, der noch nicht ausgelöst wurde, oder ein Codeabschnitt, der auf Input wartet. In AI-Systemen (z. B. generativen Netzen) meint man mit „latent space" den inneren Repräsentationsraum – also eine abstrakte, komprimierte Form von Information, aus der neue Inhalte generiert werden können. In Raum- oder Mediengestaltung kann ein latenter Zustand ein potenzieller Zustand des Raums sein – z. B. eine Leerstelle, ein unbeleuchteter Bereich oder eine programmierte, aber noch nicht gestartete Projektion.

2: Auszug aus der Projektbeschreibung: "Motion Experience Lab@(A)IDLE MODE" is not only a place of encounter and reflection, but also a place of discovery and inspiration that visualizes the limits of what can be represented. Here, type in motion clips communicate with an atmospheric media architecture, offering a spatial and temporal intermezzo for exploring latent space. The central theme invites visitors to immerse themselves in the hidden layers, undiscovered potentials and possibilities of the creative process, as well as to engage with the discourse behind it.

Chunk 4



Abbildung 2: Besucher Paul Kimmerl in der Idle Queue II: Under Construction. © MXL Nina Juric

Chunk 5

Wie kann das MXL selbst zur KI-Kreatur werden?

Das MXL wurde nicht nur als Produktionsort umgebaut und physisch zur Schau gestellt, sondern als Idee einer KI mit vielen Modellen zum aktiven Mitspieler mit performativer Präsenz und Charakter. Sensorik, audiovisuelle Systeme und Netzwerktechnik formten ein Raumsubjekt als lebendige Schnittstelle intermedialer Raum-Zeit-Architektur, durchlaufbar vom *Visual Input Layer* über modellhafte *Hidden Layers* bis zum *Visual Output Layer*. So wurde das Motion Experience Lab selbst zur *performativen KI-Kreatur*. Judith Butlers Theorie der Performativität und ihr Verständnis von *Infrastruktur als Verkörperung von Macht*[3] bildeten den theoretischen Rahmen für dieses infrastrukturell verdrahtete, lebendig werdende System – einem durchlässigen Möglichkeitsraum kollektiver, gestalterischer Versammlungen, einer szenografischen Übersetzung algorithmischer Logik ins physisch-digitale Material.

3: Butler, J. (2015). *Notes toward a performative theory of assembly* (pp. 35–40). Harvard University Press.



Abbildung 3: Der Gatekeeper an der AI Lemonade Bar, der den Zugang (Transfertunnel) zur Datenbank kontrolliert, umgeben von inaktiven idlenden Datensätzen – in Form von neugierigen Besuchern. © MXL Nina Juric



Abbildung 4: Das studentische Produktionsteam im Leerlauf zwischen den Räumen. © MXL Nina Juric

Chunk 7

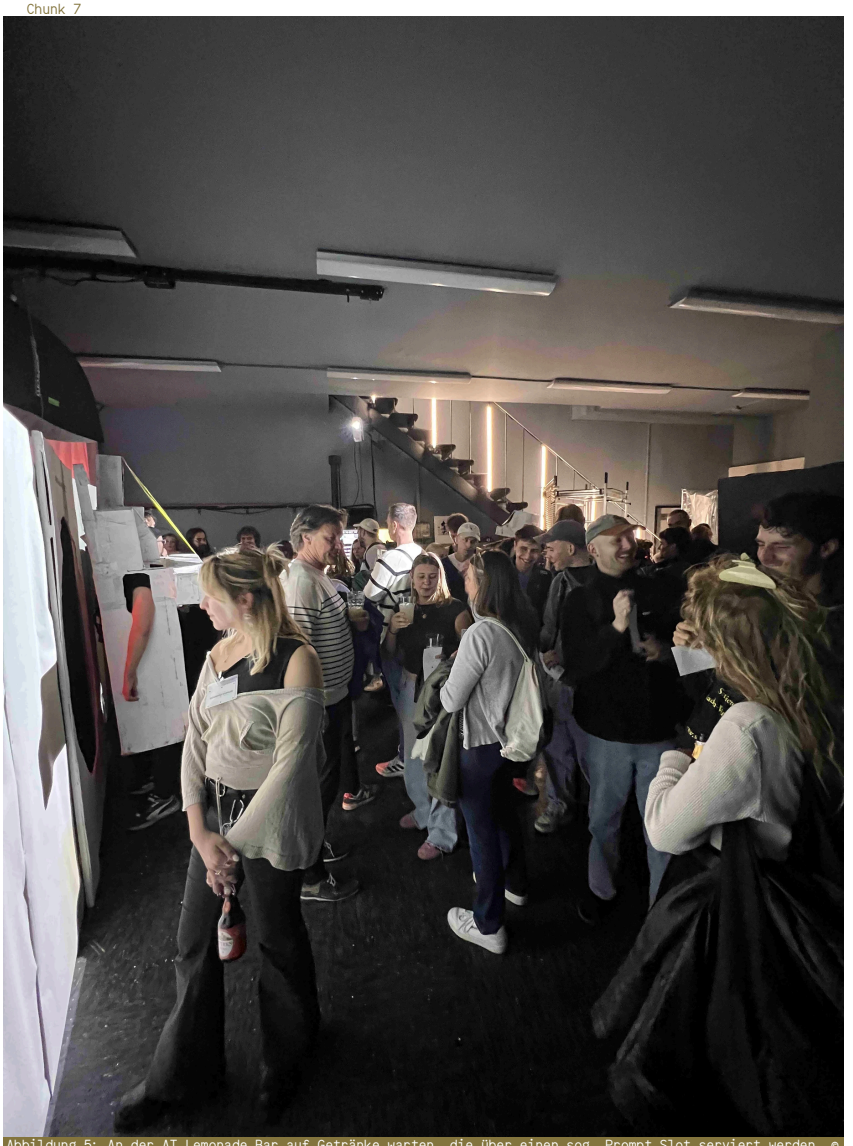

Abbildung 5: An der AI Lemonade Bar auf Getränke warten, die über einen sog. Prompt Slot serviert werden. © MXL Nina Juric

„A (virtual) gaze that travels in an imaginary flânerie through an imaginary elsewhere and an imaginary elsewhen".[4] Sieben interaktive Stationen formten in diesem Sinne eine techno-philosophische Erkundung des Begriffs (A)Idle Mode – eine interpretative Untersuchung von Zwischenzuständen – getrackt, verkörpert, (un)gesehene, erlebbare Atmosphären algorithmischer Systeme.

Chunk 8 Die Räume übersetzten die zentrale Fragestellung in gebaute, animierte Interfaces, aufeinander abgestimmte (KI-)generierte Computervisionen, modellierte Datenbankprozesse, hörbare Echtzeitszenarien und Feedbackschleifen.

Jede Station wurde von einem studentischen Team mit künstlerischem Fokus und Timing auf spezifische (A)Idle-Zustände konzipiert – von Praktiken

4: Friedberg, A. (1993). *Window shopping: Cinema and the postmodern* (p. 304). University of California Press.

des Verweilens über temporale Verzögerungen, digitale Übergangszustände und Datenströme als Tranceformate bis hin zu Ästhetiken des Wartens, der Latenz und auditiven Delays.

In *In da Queue: Idle Mode I* traten User*innen mit physisch verkörperten *System Creatures* –dem *Prompter,* und dem *Puncher* – an der *AI Lemonade Bar* in Kontakt .

Chunk 9 Dort wurde der Drink *gepromptet*, während man das Warten übte. Nach terminiertem Schleusen eines *Gatekeepers* führte ein Tunnel zum interaktiven Medientheaterstück *Auf der Datenbank* in der Besucher*innen – als Datensätze getrackt – Teil einer KI-Trainingsinszenierung wurden. Eine Echtzeit-Schnittstelle diente zur dramaturgischen Auswertung einer spielbaren *Bias*-Bühne, bestehend aus nummerierten Plätzen.

Chunk 10 Co-partizipatives Spiel in direkter Ansprache durch einen dirigierenden *Datascientist* – inszeniert als Stimme und animierter Text – und nonlineare Inhalte formten ein sensuelles, surreales Erlebnis. Hito Steyerls Konzept der *Bubble Vision*[5] – einer durch Algorithmen strukturierten Perspektivierung des Raums – wurde räumlich konkretisiert.

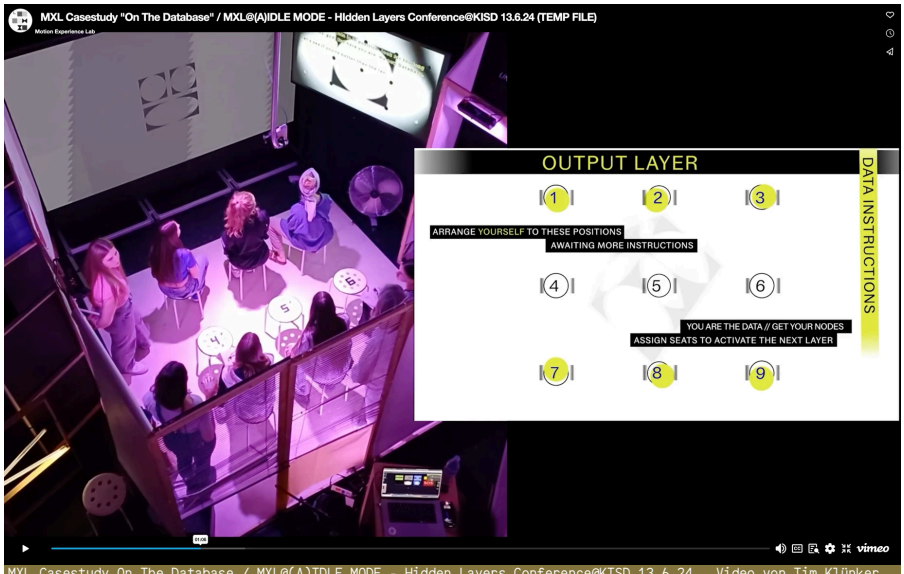

MXL Casestudy On The Database / MXL@(A)IDLE MODE - Hidden Layers Conference@KISD 13.6.24 Video von Tim Klünker und Léon Scheunemann Link: https://vimeo.com/1088110430

Chunk 11 *In the Queue II: Under Construction* zeigte als unfertiger Licht-Raum den Gegenpol zur immersiven Inszenierung. Kabel, Gerüste und lose Stoffbahnen gewährten Einblick in den verborgenen *Backstage Layer* des Datenkanals.

Chunk 12 Besucher*innen bewegten sich langsam, fast tastend durch einen Flur aus semi-transparenten Textilschichten, die auf und ab fuhren. Statt einer klaren Narration: ein Dauerzustand des Prozesses – Trance und Baustelle zugleich. Temporäre Syntax aus DMX- und Seilzug-gesteuerter Beleuchtung und bewegungssynchronisierter Vorhänge, die einen Korridor öffneten und schlossen – als choreografierte Visualisierung des Daten-Uploads.

5: Steyerl, H. (2017). Bubble vision. In *Duty free art: Art in the age of planetary civil war*. Verso Books.

Gleichzeitig bildete dieser ästhetisch konturierte und beruhigende Raum einen stillgestellten Zustand inmitten des Maschinenkörpers –zwischen Kontrolle und Kontemplation, zwischen Latenz und Präsenz.

Chunk 13 Festgehalten in einer Zeitleiste wurde Zeit begehbar – nicht als lineares Fortschreiten, sondern als gespannte Gegenwart. *On hold*.

Chunk 14

I. p. 103, Chunk 15: Bildgenerierende Modelle in…
I. p. 74, Chunk 12: Walking in latent space
I. p. 172, Chunk 4: A computer generated map?…



Abbildung 6: Besucher*innen idlen an der The Interface-Installation. © MXL Nina Juric



Abbildung 7: Besucher idlen an der The Interface-Installation mit Kontrolleur Marko Bodenstein. © MXL Nina Juric

Im *Control Center – Bias on the past* traf mediale Kontrollfantasie auf operative Leere. Zwischen *E-Waste, Technik-Totems*[6] und *Pseudo*-Interfaces entlarvte sich das Versprechen algorithmischer Steuerung: Viel schien steuerbar, doch wenig war es. Anonyme Kontrolleur*innen verkörperten ein unsichtbares Machtregime. Kamerafeeds aus anderen Räumen flimmerten auf zahlreichen Screens. Die Installation *The Interface* morphte Livebilder von Besucher*innen in KI-generierte Varianten – von Blumenstilleben bis Taylor Swift. Die Transformation auf einem alten Fernseher offenbarte die bildmächtige Logik generativer Systeme. Meg Stuarts choreografische Prinzipien fragmentarischer, prozesshafter Körperbilder boten hier Denkfiguren für ein gestalterisches Begreifen scheinbar intelligenter Zwischenräume.[7]-

6: Ein Techniktotem ist ein symbolisch aufgeladenes Objekt aus technischen Geräten oder Bauteilen, das kulturelle, soziale oder ökologische Bedeutungen von Technologie und Techniknutzung sichtbar macht.

7: Stuart, M. (2010). *Are we here yet?* (J. Peeters, Ed.; p. 256). Les presses du réel. <https://www.damagedgoods.be/en/shop/are-we-here-yet>

Projektion, Blicklenkung und Pose – technisch realisiert mit TouchDesigner und *StreamDiffusion*[8] [9]– zeigten eindrücklich: Was wie Kontrolle wirkt, ist oft nur Visualisierung vergangenheitsbasierter Modelle – *Bias in Aktion*.



Abbildung 8: Besucherin interagiert mit der The Interface-Installation. © MXL Nina Juric

8: Kodaira, A., Xu, C., Hazama, T., Yoshimoto, T., Ohno, K., Mitsuhori, S., Sugano, S., Cho, H., Liu, Z., & Keutzer, K. (2023). *StreamDiffusion: A pipeline-level solution for real-time interactive generation*. <https://arxiv.org/abs/2312.12491>

9: https://github.com/cumulo-autumn/StreamDiffusion

Chunk 15



Abbildung 9: The Interface – eine interaktive Installation vor dem Control Center, die AI-Looks live generiert, basierend auf Slider-Eingaben und Bewegungsdaten, welche über eine Kamera im Touchdesigner mit StreamDiffusion gesendet werden. © KITEGG

Wie klingt eine KI im (A)IdleMode?

I. p. 144, Chunk 3: untitled

Chunk 16 Das Sounddesign verstand Klang nicht nur als akustisches Element, sondern als interaktives Interface zwischen Mensch, Raum und KI-System. Gemeinsam mit Soundkünstler Lio Vast entstand ein räumlich-reaktives Konzept, das auditive Feedbackloops mit generativer Logik und Echtzeitverzögerung mixte. Weil die Räume akustisch nicht entkoppelt waren, entwickelte Vast eine *Raum-im-Raum-Mischung*, bei der alle Klangquellen so synchronisiert wurden, dass sie sowohl für die einzelnen Stationen an sich als auch als Gesamtkörper funktionierten. Rhythmische Elemente wurden bewusst räumlich versetzt. Vast beschreibt seinen Ansatz: *„Das Sounddesign basiert auf Spannung, Reibung und Verdrängung – durch Sidechaining und starke Kompression/OTT geraten Klangschichten in Bewegung, kollidieren, drücken sich gegenseitig aus dem Vordergrund und erzeugen ein pulsierendes Gefüge aus sphärischem Druck und kontrollierter Instabilität.“*

Konkurrenz um Klangraum, dichte Texturen und subtile rhythmische Verschiebungen machten das *(A)Idleing* als latente Energie hörbar. Der Raum wurde zur Soundskulptur und Resonanzfläche eines nicht sichtbaren, aber spürbaren Systems.

Lehrstrategie & Format

Das Projekt verband gestalterische Freiheit und inhaltliche Lehrstruktur mit kollektiver Autonomie.

I. p. 71, Chunk 2: Walking in latent space
II. p. 73, Chunk 7: Hack the Net (with AI)
I. p. 104, Chunk 16: Bildgenerierende Modelle in…

Chunk 17 Über acht Wochen wurden die Studierenden eng begleitet – gestützt durch eine künstlerisch-strategische Leitung, die Projektoutline, Kreativdirektion und Zeitmanagement umfasste. Vorgeplante Sprints strukturierten die Arbeit in transdisziplinären Teams entlang realer Produktionsabläufe. Didaktisch wurde das Projekt iterativ durch *Übungen am Material*, regelmäßige *Crit*-Sessions, technische Inputs[10] und reflektierenden Austausch geführt. Der performative Rahmen wurde konzeptionell wie technisch durch die Ausstattung und Ästhetik des MXL gestützt und begrenzt.[11] [12] Räume wurden dynamische Konstellationen: bewegte Material- und AV-Systeme in einer lernenden Auseinandersetzung mit KI. Das Motion Experience Lab wurde selbst zum Fallbeispiel für *Practice-Based Design Research:*[13] von einer lernenden Lehrumgebung zur Gestalt einer Bühne, zum experimentellen Design einer Live-Probe mit über 300 Besucher*innen – ein artefaktisches Setting, das – *trotz diverser Durchläufe* –unter *realen* Bedingungen nie erprobt werden konnte.

Methodischer Zugang: Code & Kontext – Von *Creative Conception* zu *Conceptual Creation*

Der Begriff *(A)IDLE MODE* diente als Impuls für kreative Konzeptvorschläge, die in den ersten Wochen intensiv diskutiert wurden.

Chunk 18 Der thematische Kontext setzte einen produktiven Rahmen, in dem Code nicht nur als technisches Werkzeug, sondern als gestaltendes Mittel wirkte. Als Teil kreativer Technologien sollte Programmierung die einzelnen Raumkonzepte zu sozialen Erfahrungen verbinden.

Als Leitidee der Arbeit diente Butlers Verständnis von Infrastruktur als relational-performativem Netzwerk, das soziale, politische und technische Praktiken verbindet[14]– als dynamische Raum-Zeit-Komplexität, die physische wie mediale Ebenen konzeptionell verknüpft. Diese “infrastructured latent spaces” ermöglichten es, komplexe Gleichzeitigkeiten und Raum-Zeit-Verflechtungen als diskursive Phänomene zu inszenieren.

Chunk 19 Die ins Material übersetzte Rauminstallation zog Lernende und später Besucher*innen in einen dialogischen Prozess und wurde so zum

10: Zu Beginn: 10 Tage-Crashkurs “*Creative Live Tech*” mit TouchDesigner (*customized MXL networks*) inklusive DMX, Sensorik, Material- & Raumdramaturgie u.a.

11: Juric, N. (2022). *States of Motion: The Motion Experience Lab 2019–2022*. Image&Motion / KISD. https://imageandmotion.xyz/

12: What is the Motion Experience Lab? The Motion Experience Lab is a hybrid playground environment at KISD for creation and research, focusing on the artistic practice of augmented motion in the context of performative interaction and interdisciplinary media studies. Since 2016, it has served as a production studio within the Image and Motion department, founded on a conceptual framework that fosters exploration and encourages transdisciplinary approaches. Students engage in hands-on design research in both high- and low-tech environments, developing prototypes and artifacts that investigate movement, audiovisuality, spatial performance, interactive worldbuilding, and immersive experiences.Initiated and developed by Prof. Nina Juric, the MXL functions as a scaffolding toolset for experiential learning embedded in a wired audiovisual environment. Together with employees and students this space is dedicated to creative tech like embedded audio-visual systems, generative design tools, expanded computer animation, real time imagery with body action, code and programming or performing arts and cultural media science – always to experiment on time-in-space-complexities and so called `interface performances’.

13: Performing Arts and the Royal Academy of Art (KABK) in The Hague. (n.d.). *Practice-based Research*. https://phdarts.eu/PhDArts

14: Butler, J. (2015). *Notes toward a performative theory of assembly*. Harvard University Press.

iterativen Prototypen, Diskussionsobjekt und *Conversation Piece.*[15]

Making Of Visual Systems

Nach einem zweiwöchigen TouchDesigner-Intensivkurs startete das projektbasierte Eventformat mit zwei Produktionssträngen: ein Spatial-Design-Team und ein Motion-/Packaging-/Animationsteam. Letzteres entwickelte – ohne Vorerfahrung, aber in enger Zusammenarbeit mit dem Konferenzgrafikteam – ein typografie-basiertes Motion-Design-System, das als animiertes Event-Package diente und zugleich das MXL architektonisch bespielte. Es fungierte als visuelles Navigations- und Leitsystem der Gesamtinszenierung.

Chunk 20

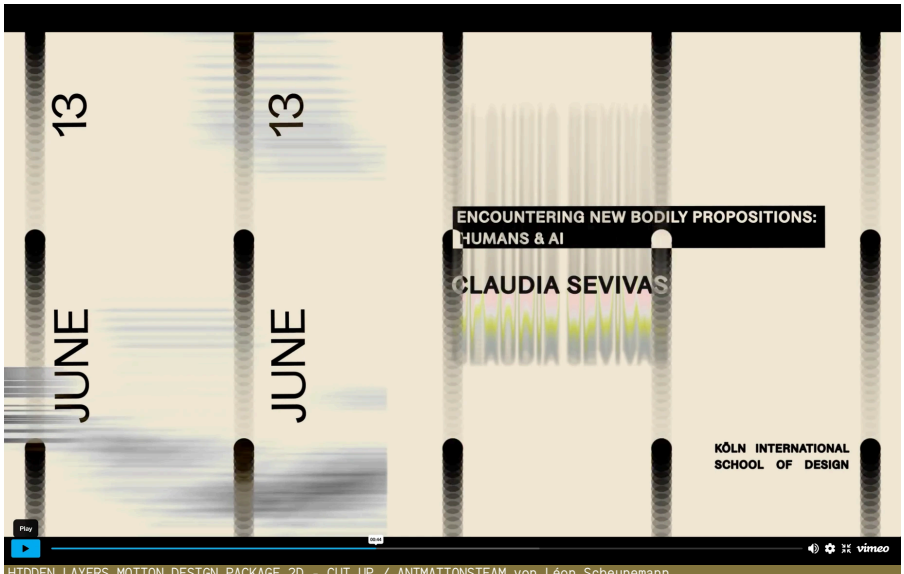

HIDDEN LAYERS MOTION DESIGN PACKAGE 2D - CUT UP / ANIMATIONSTEAM von Léon Scheunemann
Link: https://vimeo.com/1088381189

Das Spatial-Team gestaltete die räumlichen Erlebniswelten und übertrug die grafische CI der Konferenz auf 200 m² Fläche und ins Material. Dazu gehörte unter anderem der Bau von Holzwänden, Raumtrennern und Pappbauten, die Konstruktion von Tunneln, Schriftplots, Kabelverlegung und die Planung barrierefreier Wege hinter den Kulissen.

Chunk 21 *Mithilfe von TouchDesigner, VVVV, Code und der Infrastruktur des MXL wurden mediale Inhalte und Livefeeds zu zeitlich abgestimmten Bewegtbildströmen verschaltet, die als audiovisuelle Leitstruktur ein immersives Erlebnis schufen. Testing sorgte für technische und räumliche Abstimmung, sowie ästhetische Kohärenz.*

Chunk 22

15: Bennett, T. (1995). *The birth of the museum: History, theory, politics.* Routledge. *Der Begriff „conversation piece" bezeichnet ursprünglich ein Kunstwerk, das Gespräche anregt und als sozialer Interaktionsgegenstand fungiert. Tony Bennett beschreibt hier wie solche Objekte in musealen Kontexten als dialogische Vermittler zwischen Werk und Publikum wirken.*

Welche Rolle spielte *Embodiment*?

Chunk 23 Im Zentrum der räumlichen Dramaturgie standen die Besucher*innen selbst. Durch Bewegung, Positionierung und Interaktion im Raum wurden ihre Körper Teil des Systems– *Embodiment im posthumanen Sinne.*[16] Verkörperung und Erleben von Technologie wird hier als integraler Bestandteil menschlicher Wahrnehmung und Identität begriffen. In Anlehnung an Meg Stuarts choreografisches Konzept des „Enactments“[17] wurde die Gegenwärtigkeit der Besucher*innen zu einer Echtzeit-Performation und Perforation von Daten, Raum, Zeit, Pixeln und Ton. Die Ko-Präsenz von Mensch und Maschine fragt nach *Liveness,*[18] Gleichzeitigkeit und einer Gegenwart als *dehnbarer Aufmerksamkeitszone*.[19]



Abbildung 10: Auf der Datenbank. © KITEGG

16: Hayles, N. K. (1999). *How we became posthuman: Virtual bodies in cybernetics, literature, and informatics* (p. 149). University of Chicago Press. *„Posthuman embodiment reconceptualizes the body not as a fixed biological entity but as a dynamic, technologically mediated process."*

17: Stuart, M. (2010). *Are we here yet?* (J. Peeters, Ed.; p. 256). Les presses du réel.

18: Auslander, P. (2008). *Liveness: Performance in a mediatized culture* (2nd ed., pp. 3–10). Routledge. Auslander analysiert den Begriff „Liveness" als eine kulturell und medial codierte Kategorie, insbesondere im Spannungsfeld zwischen medialer Repräsentation und unmittelbarer Präsenz – zentral für jede Reflexion von Echtzeit, Performance, Medienkunst.

19: Steyerl, H. (2017). *Duty free art: Art in the age of planetary civil war*. b_books.

I. p. 115, Chunk 11: Developing an AI teaching…
II. p. 25, Chunk 14: Schwäbisch Gmünd
I. p. 29, Chunk 14: Living Objects Lab (KISD)

Chunk 24



Abbildung 11: Auf der Datenbank. © MXL Nina Juric

I. p. 29, Chunk 12: Living Objects Lab (KISD)
II. p. 198, Chunk 15: Work in Progress
II. p. 203, Chunk 1: Shaping Human-AI Relationships

Chunk 25

Learnings & Reflexion: *Wie sieht Gestaltungslehre aus, wenn Entwurf, Forschung und Event verschmelzen?*

Das *Motion Experience Lab@(A)IDLE MODE* war kein klassisches Lehrprojekt, eher ein lebendiges System kollektiven Denkens, Handelns – und Scheiterns. Die doppelte Struktur – Design-Forschungsformat und Live-Event – verlangte gemeinsame Verantwortung und eine professionelle Produktion in einem Team unterschiedlichster Erfahrungsstufen. Gleichzeitig eröffnete sie eine neue Rolle für das MXL: als kritisch inszenierte AI-Creature, die ihre eigene Performanz und Infrastrukturalität reflektiert.

(A)Idle Mode stand nicht für technischen Stillstand, sondern für den Anfangszustand eines entleerten Interfaces – bereit zur Aktivierung durch performative Mittel. Dieser *Leerlauf* wurde zum Modell algorithmischer Latenz: einem Zustand gespannter Erwartung, in dem Potenzialität und Unsichtbarkeit zusammenfallen. Ziel war kein unmittelbares Handeln, sondern ein atmosphärisches Changieren zwischen Dasein und Abwesenheit – eine visuelle Erfahrung, die nicht direkt, sondern vermittelt wahrgenommen wird: als Spur, Interface oder Projektion.

Die achtwöchige Lehrphase verband systemische Mediengestaltung, Type in Motion Design, performative Dramaturgie, Echtzeitprogrammierung und Szenografie. Sie zeigte: Echtzeitlichkeit, körperliche Präsenz und gestalterische Spekulation lassen sich als didaktisches Triptychon begreifen, um KI nicht nur zu entwerfen, sondern körperlich zu verhandeln. *„Es hat definitiv nicht alles funktioniert wie gedacht – und genau das hat es lebendig gemacht“*, so eine Studierende. *„Wir sind keine AI – und das erwartet auch niemand.“*

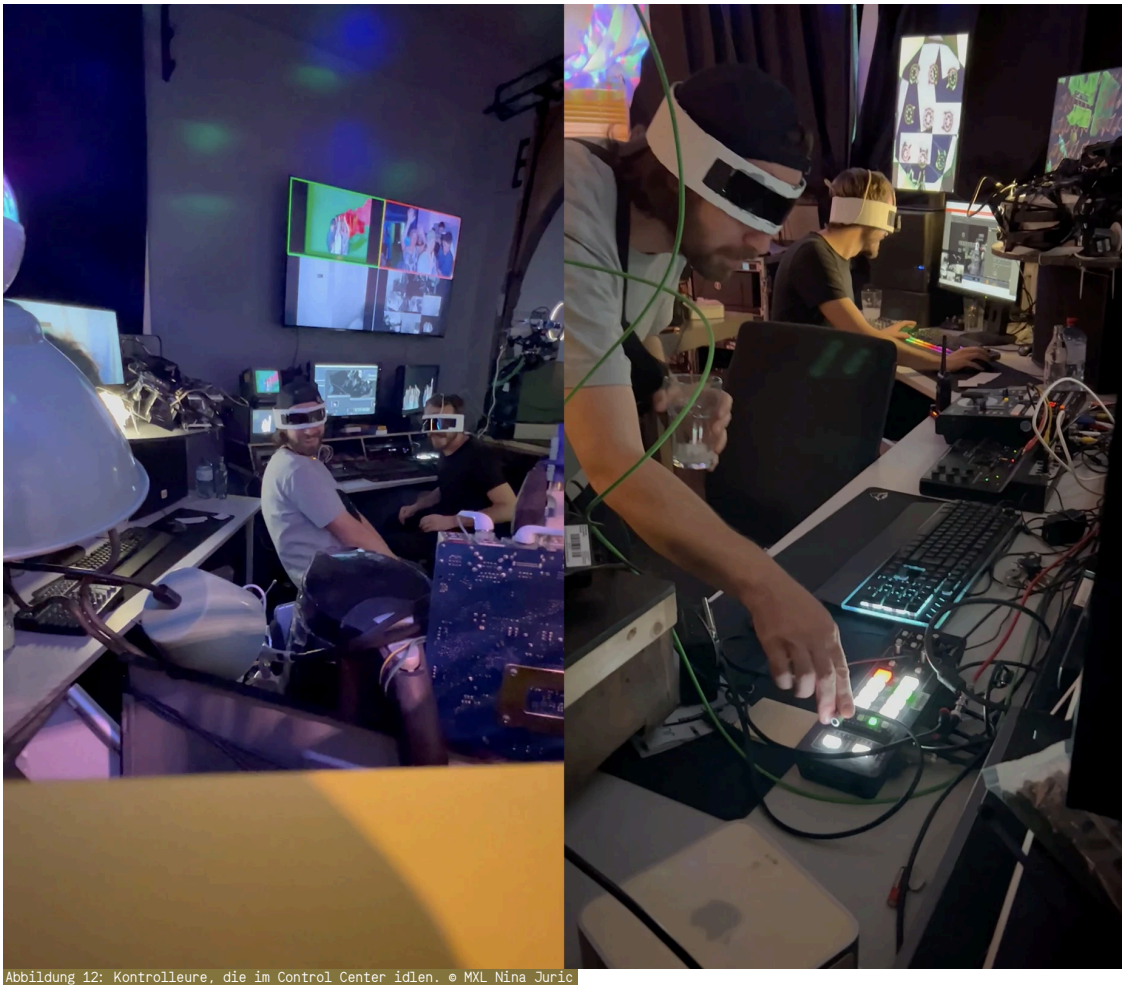

Abbildung 12: Kontrolleure, die im Control Center idlen. © MXL Nina Juric

II. p. 202, Chunk 5: Editorial: Show and Tell
II. p. 220, Chunk 4: PASTELAB
II. p. 68, Chunk 19: Editorial

Chunk 26 Die Komplexität des angewandten Formats entstand aus der gleichzeitigen Verbindung von Entwurfslehre, Ausarbeitung und einer improvisationsoffenen Raum-Zeit-Choreografie – ohne Möglichkeit zur Vorab-Erprobung mit echtem Publikum. Genau diese Verschränkung forschungsbasierter Lehre mit experimenteller *Creative-Tech*-Produktion bildet einen zentralen Anspruch des Motion Experience Labs.

Chunk 27 Sie verlangt Risikobereitschaft, belastbare Teamstrukturen, Erfahrung – und ein Projektmanagement, das Flexibilität ermöglicht: eine Flexibilität, in der mit Händen gedacht, mit Material reflektiert und mit Raum operiert wird. Ein Projekt dieser Dichte und Größenordnung, das in kürzester Zeit hybride Disziplinen und Medienarchitekturen vermittelt, erfordert klare Kommunikationsstrukturen und verbindliche Prozesse.

Dabei verschieben sich klassische Rollen: Lehrende gestalten, Gestaltende lehren. Technik wird Akteur, Raum wird Mitspieler, Material wird Sprache. Es ist eine produktive Unschärfe, wie sie aus der Performancekunst oder Ausstellungspraxis bekannt ist – hier jedoch eingebettet in ein temporäres Forschungslabor.

Lehrdesign technischer Vermittlungsprozesse oder zur Integration raumbasierter Interaktion in gestalterisch-technischer Forschung. Was bedeutet *Liveness*[20] unter algorithmischen Vorzeichen? In der Auseinandersetzung mit dem *Jetzt*[21] und den dazugehörigen Ästhetiken und Bedingungen von Wahrnehmung, beginnt eine neue, kritische Designpädagogik, die mit vielen Formen des Scheiterns einhergeht. "*Technology will always crash* ",[22] so Keez Duvyes beim Creative Coding Lab 2023.[23] Doch gerade diese Ambivalenz und Unvorhersehbarkeit digitaler Technologien – speziell in experimentellen Produktionsprozessen – ist kein Defizit, sondern konstitutiv für kreative Tech-Lehre. Auch der ephemere Charakter des Events ist Teil dieser Erkenntnis: „Everything not saved will be lost."[24] Erst *reale* Räume machen technischen Weltenbau erfahrbar – und damit zentrale Gestaltungskompetenzen wie technische Gleichzeitigkeit, ästhetische Wirkung und sensuelle Dynamik vermittelbar. Komplexe Gestaltung von derartigen Installationen lässt sich lehren, wenn auch das Scheitern eingeladen ist – als Teil der Echtzeit und jener kollektiven Energie, die sie freisetzt.

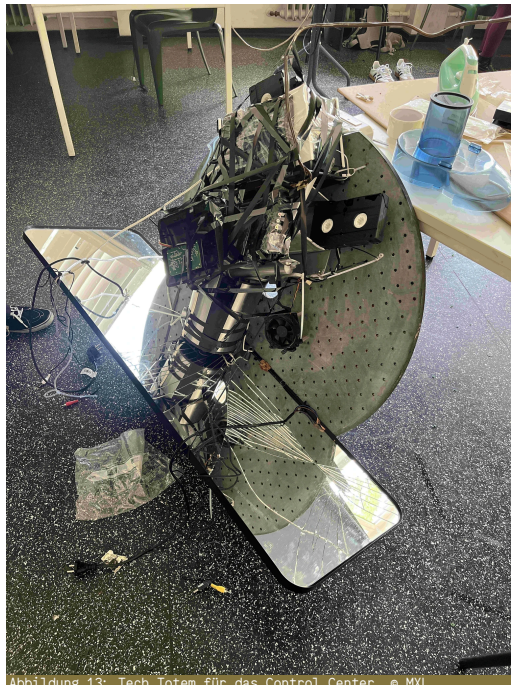

Abbildung 13: Tech Totem für das Control Center. © MXL Nina Juric

20: Referenziert wird hier wieder Philip Auslander, amerikanischer Kulturwissenschaftler und Performance-Theoretiker, der vor allem für seine Arbeiten zur Beziehung zwischen Live-Performance und Medien bekannt ist. Er ist Professor an der Georgia State University im Fachbereich Theater und Performance Studies. Seine Monografie *"Liveness: Performance in a Mediatized Culture"* (1999) gilt als wegweisend für diesen Begriff. Er diskutiert darin die Verschiebung des Live-Begriffs in der medialen Kultur und wie „Live"-Performance durch Medien reproduziert und zugleich transformiert wird. Er argumentiert, dass Live-Performance nicht mehr strikt vom Medienvermittelten getrennt betrachtet werden kann, sondern dass beide Formen sich gegenseitig beeinflussen und dass „Liveness" auch als ein kulturelles Konstrukt verstanden werden muss. Seine Definition von Liveness ist oft Ausgangspunkt für weitere Diskurse.

21: *Nowness*; Vgl. zum Begriff u. a. Walter Benjamins Konzept der „Jetztzeit" in *Über den Begriff der Geschichte* sowie Erika Fischer-Lichtes performativ-ästhetische Perspektive auf *Liveness* in *The Transformative Power of Performance* (2008), in der sie die soziale Interaktion und die transformative Qualität von Live-Erfahrungen hervorhebt.

22: Duyves, K. (2023). *Personal statement during creative coding lab talk 2023*. https://choreographiccoding.org/lab/380-cologne-2023

23: Keez Duyves ist NewMedia Artist und Gründer von Pips:lab, einem interdisziplinären Medientheaterkollektiv in Amsterdam. Er ist seit den 1990er Jahren Erfinder und als Creative Technologist insbesondere im Bereich Medientheater, New Tech und Live-Performances tätig. Sein Fokus liegt auf der Entwicklung innovativer, performativer Technologien für den Echtzeitbetrieb. Das Zitat „Technology will always crash" stammt aus seinem Vortrag beim Creative Coding Lab 2023 (www.choreographiccoding.org), und reflektiert seine langjährige Erfahrung mit der inhärenten Unvorhersehbarkeit und Fragilität digitaler Systeme in Live-Produktionskontexten.

24: Nintendo, *The Legend of Zelda: Breath of the Wild*, Nintendo, 2017, Spiel für Nintendo Switch, enthält den Spruch „Everything not saved will be lost", https://www.nintendo.com/games/detail/the-legend-of-zelda-breath-of-the-wild-switch/.

Chunk 28 Dieser Text wurde mit Unterstützung von ChatGPT (strukturierte Schreibprozesse) verfasst und basiert auf dokumentarischen Skripten sowie Texten von Lio Vast, Annika Walter, Carolin Saddey und Léon Scheunemann.

Projektteam:

Installationen & studentische Teams

#01 Visible Input Layer – Who is the MXL AI?
Carlos Strohm, Omid Akbari Kharazi, Irena Savchyn, Léon Scheunemann

#02 In da Queue: Idle Mode I – User Stimuli / How to Prompt?
Anastasia König, Laura Fleischer, Pia Bernards

#03 AI Lemonade Bar
Victor Thrane Kolling, Annika Walter, Nicolai Rehberg, Paul Kimmerl

#04 On the Database – Performing TouchDesigner / Mediaplay
Carolin Saddey, Tim Klünker, Irena Savchyn, Pia Bernards, Henning Brode

#05 In da Queue: Idle Mode II – Under Construction
Anastasia König, Laura Fleischer, Kirill Solovev

#06 Control Center – Bias on the Past
Carlos Strohm, Lukas Kostka, Anastasia König, Muhammad Danial Sohail, Paul Kimmerl

#07 Visible Output Layer
Omid Akbari Kharazi, Muhammad Danial Sohail, Léon Scheunemann

Weitere Mitwirkende

Sounddesign: Leo Maria Althoff aka Lio Vast, Philipp Schultz, Pia Bernards
Kostüme & Tech-Totems: Studierende des 4. Semesters Code & Context

Produktionsteam

Eine Produktion des Motion Experience Labs, SoSe 2024

Leitung / Creative Direction: Prof. Nina Juric
Technik: Marko Bodenstein
(MXL)Assistenz: Henning Brode (SHK)

Produktionszeit: 8 Wochen, Proben: 3
Beteiligte Studiengänge: Image&Motion / Integrated Design (KISD)
Code&Context (TH Köln)

Links:

imageandmotion.xyz @motionexperiencelab

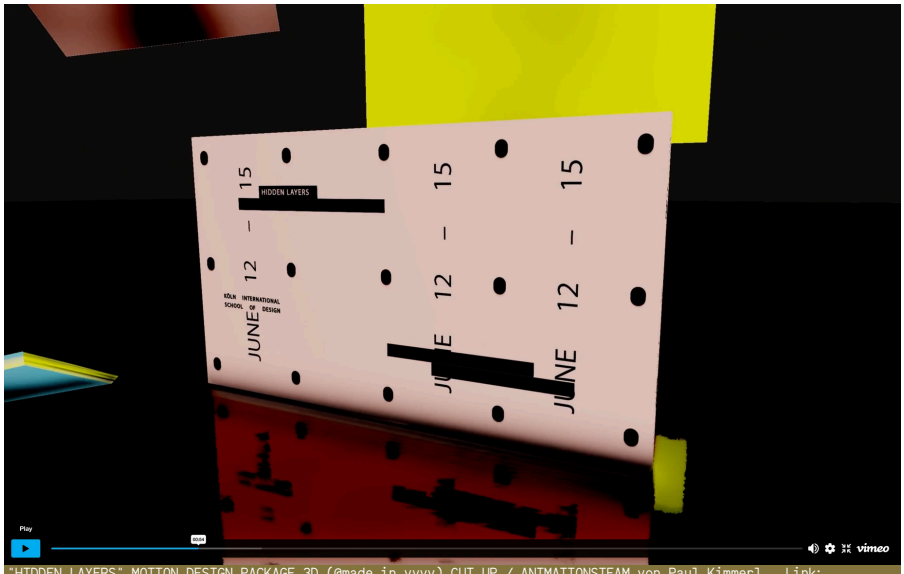

"HIDDEN LAYERS" MOTION DESIGN PACKAGE 3D (@made in vvvv) CUT UP / ANIMATIONSTEAM von Paul Kimmerl Link: ↗ https://vimeo.com/1088752715

Chunk 29



Abbildung 15: Die Hälfte des Produktionsteams v.l. Muhammad Danial Sohail, Carolin Saddey, Omid Akbari Kharazi, Victor Thrane Kolling, Annika Walter, Lukas Kostka, Anastasia König, Carlos Strohm, Irena Savchyn, Lio Vast, Nina Juric, Paul Kimmerl, Henning Brode, Léon Scheunemann / (Nicht im Bild: Tim Klünker, Nicolai Rehberg, Kirill Solovev, Pia Bernards, Laura Fleischer, Philippa Fink und das gesamte Code&Context Team). © MXL Nina Juric

↗ ro8j5asv

DIY Compute: Rolling Your Own GPU Node from Cheap Hardware

Anton Koch

Chunk 1

Why Build Your Own GPU Node?

As demand has risen for using graphics processing units (GPUs) as compute resources for machine learning applications, cloud providers have seized the opportunity, offering to rent out dedicated hardware by the hour; meanwhile, hardware manufacturers such as NVIDIA have responded by providing integrated server platforms. Access to such GPUs is becoming essential for project planning as more and more research projects depend on either using existing machine learning models or training their own. We have shown that owning hardware is cheaper than outsourcing to the cloud for a research project with a medium to long-term running time that can use existing housing facilities[1]. However, the current status quo of GPU server architecture from companies like NVIDIA or AMD remains a significant cost factor in terms of immediate investment.

Still, if the project is relatively small in terms of staff concurrently using the infrastructure, building cheap barebone server modules that cost around €3000 can be a viable alternative to the more expensive high-end variants.

1: Koch, A. (2025). Der KITeGG Cluster – eine Infrastruktur für KI in der Gestaltungslehre. *Un/Learn AI, 2*. https://doi.org/10.25358/openscience-11844

These nodes can hold up to six PCIe cards that can be sourced from anywhere, depending on the application and the environment. These can be discarded graphics cards from a university's VR lab or workstation cards that can be repurposed or bought at a reasonably low cost. While expensive server models excel in many other performance factors, this article aims to show the opposite end of the spectrum in order to facilitate access to GPU computing.

An Example Hardware Configuration

For an example configuration, here is a list of parts used for building a GPU node from scratch, using affordable server hardware.

Mining-rack Server Case

The most affordable alternative to a more expensive GPU-capable server case is a cheap case made for crypto mining rigs. We have found the 4W2 mining rig case by German manufacturer IPC to be the most affordable at around €150[2]. The case comes with six fans whose speed cannot be controlled by pulse-width modulation (PWM), which we replaced with PWM-capable fans.

Chunk 2 We also modified the case to reverse how the cards are installed, as the factory default setup of the case would cause the cards to release hot air towards the front.

Chunk 3 The actively ventilated workstation GPUs use the so-called *blower design* to release hot air outside the PC case, and this is done on the end where the mounting bracket sits. The default case layout is designed to place the cards with the exhaust to the front (Figure 1), which we wanted to avoid aligning with our other servers' front-to-back airflow direction.

Chunk 4 Additionally, installing the cards the other way around shortens the distance between the PCI connectors.

2: (n.d.). *IPC 4W2 MINING-RACK*. Inter-Tech Elektronik Handels GmbH. Retrieved June 11, 2025, from https://www.inter-tech.de/productdetails-139/4W2_MINING-RACK_EN.html

Chunk 5

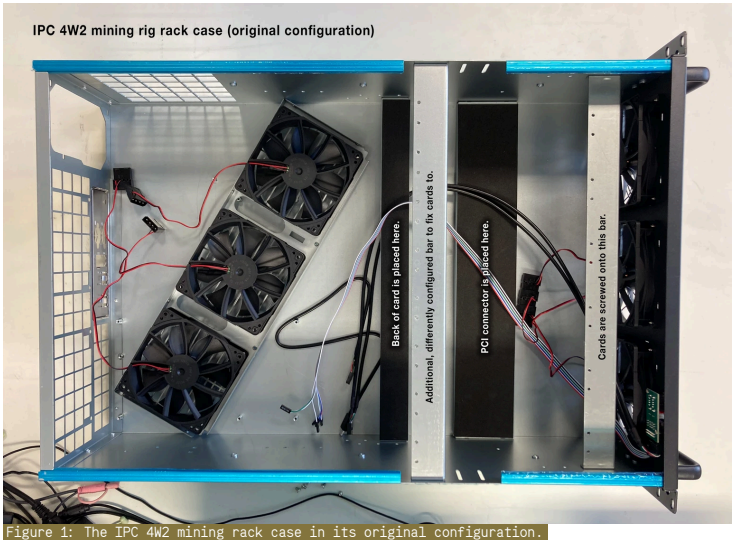

Figure 1: The IPC 4W2 mining rack case in its original configuration.

Chunk 6

Mainboard

Any Mini ITX, µATX, or ATX mainboard can be installed in the 4W2 rack case to transform a workstation into a server. Still, installing a dedicated server mainboard is best, usually featuring a baseboard management controller (BMC) implementing the Intelligent Platform Management Interface (IPMI). This allows you to set up a dedicated network connection to your server, which enables remote management of the machine, such as power control, reset, and remote terminal access, independent of the operating system (OS) state.

We settled on the ROMED8-2T/BCM board by ASRock, which features an ATX form factor, a single SP3 CPU socket, and seven PCIe slots for around €700[3].

Chunk 7

Hardware Resources

We recommend the following CPU, memory, and disk space specs to equip the mainboard. It must be noted that the following configuration is for a system setup where up to four users can interactively use a single GPU each (e.g., in a JupyterLab environment) or worker tasks can be dynamically assigned to one GPU running up four tasks in parallel. The basis for calculation here is a Mini-PC with about four CPU cores and moderate associated specifications, or a small virtual private server (VPS) as used by many cloud hosting providers.

Chunk 8 However, this can always be scaled up to suit different needs, depending on the budget.

The CPU was the most significant cost factor, at about €800 for an AMD Epyc 7513 with 32 cores and a 2.6 GHz frequency. For controlling largely

3: (n.d.). *Server Motherboard ROMED8-2T/BCM*. ASRock Rack. Retrieved June 11, 2025, from https://www.asrockrack.com/general/productdetail.asp?Model=ROMED8-2T/BCM

GPU-based tasks, we deemed the processor frequency to be of secondary importance and aimed at more concurrency to allow for a flexible multi-user scenario and to have some headroom for running supporting services on the node. The system was planned to house four GPUs, so we calculated that eight cores would be reserved for the system and six cores associated with each GPU.

Although the eight memory slots on the mainboard would support a maximum amount of 2TB RAM, we only equipped it with four 32GB DIMMS at first to provide a minimum amount of 128GB RAM, allowing each user to consume about 24GB, still leaving a good 32GB for the system. The DIMMS cost another €400.

Finally, the mainboard offers two M2 slots for installing NVMe SSDs as a storage medium. We filled these with Western Digital Blue devices, each 2TB in size, which cost another €250.

Power and Cooling

We chose the 1500-watt power supply HX1500i by Corsair to power our four GPU cards, which costs about €350[4].

We also replaced the default fans in the mining case with PWM-enabled fans.

Chunk 9 We have two types of fans at our disposal, and these were a mix of high-performance fans with a maximum of 8000 RPM by Arctic (a three-pack sells for about €40)[5] and basic *Silent Wings* PC-case fans with 1600 RPM by *be quiet!* (about €20 each)[6].

Chunk 10 The choice of fans depends on the hardware being installed. We installed a node containing a passively cooled NVIDIA A100 that needs strong airflow when operating at high usage. In this case, you will want a high-performance fan transporting the air around the card. Regular, blower-design workstation cards already feature a ventilation system, and a standard fan is suitable for these. Another significant difference to consider is the depth of the fans, which is almost double that of the 8000 RPM fans compared to the 1600 RPM fans.

Another factor to keep in mind is where the node will be housed. If you plan to install datacenter cards with passive cooling, the node must be housed in an enclosed or at least properly climatised environment. Actively cooled hardware can be housed at normal room temperature, but the room it is in can turn into a bakery at high usage, and you will just be shoveling hot air around, no matter the RPM of your fans.

Chunk 11 Finally, you will need a CPU fan and cooling unit to fit into the rack case, so check the specifications. The Supermicro OLGA4094

4: (n.d.). *HX1500i Fully Modular Ultra-Low Noise Platinum ATX 1500 Watt PC Power Supply (EU)*. Corsair. Retrieved June 11, 2025, from https://www.corsair.com/eu/en/p/psu/cp-9020215-eu/hx1500i-platinum-atx-1500-watt-pc-power-supply-cp-9020215-eu

5: (n.d.). *Server Fan - S12038-8K | 120 mm 8K RPM Server Fan | ACFAN00302A*. ARCTIC. Retrieved June 11, 2025, from https://www.arctic.de/S12038-8K/ACFAN00302A

6: (n.d.). *SILENT WINGS 4 silent Fans for your PC from be quiet!* Be Quiet! Retrieved June 11, 2025, from https://www.bequiet.com/en/casefans/silent-wings-4/3696

(SNK-P0064AP4) tower cooling module worked out[7] for us.

Cabling and Extenders

As the case is long and your mainboard sits at the back end, you must extend the range of your cables for the PCI cards and the fans. The fans are easily extended and bundled using cheap four-pin Y-connectors and extension cables, which cost a few euros each. However, extending the PCIe lanes for the performance needed to operate the GPUs will require high-end cables. We used 90-degree PCIe riser cables by Phanteks with a length of 60 cm for about €70 each[8]. The degree angle means that the connector sits flat on the ground so that the cable will be at a 90-degree angle to the card, much like on a mainboard. The cables are pretty thick and should not be bent too much, so 60 cm may be tight. We reversed the placement of the cards to squeeze out a few more centimeters. If you can get longer cables to place the cards further apart from the mainboard, more power to you!

Assembly

Putting everything together is mostly straightforward, with a few catches to look out for, but it also depends on whether you decide to reverse the card placement, as we did.

Tools Used

Apart from the torque screwdriver used to fix the CPU in place (we used the 7400 model by Wera with a range of 1.2–3 Nm), all we needed were the usual Philips screwdrivers for the case, mainboard, and fans, as well as a tiny one for the M2 SSDs. We also used a small cordless screwdriver to drill some holes in the case for reversing the card placement.

Chunk 12 To place the drill, we used a center punch with a small hammer and a simple, plastic ruler, a combination of a set square and a protractor, as used in school. If available, a multi-purpose tool like a Dremel or Proxxon can be helpful to both do the drilling and to deburr the metal case, but the latter can also be done with the metal drill bit.

A Word of Caution

Since you are working with delicate and electrical parts, always adhere to the respective installation instructions for each component.

Chunk 13 Read them thoroughly, double-check if unsure, or consult the internet (not a chatbot) for further instructions or tips. Also, check the parts you are installing, as we once had a CPU that came the wrong way around in the installation tray, which led to both the mainboard and the CPU being

7: (n.d.). *Supermicro SNK-P0064AP4 4U Active CPU Heat Sink For AMD EPYC OLGA4094 SP3 Socket*. Supermicro. Retrieved June 11, 2025, from https://store.supermicro.com/4u-active-amd-epyc-snk-p0064ap4.html

8: (n.d.). *Company Website*. Phanteks. Retrieved June 11, 2025, from https://phanteks.com/

damaged on installation. Never trust your input, as the old programming rule suggests. Also, take your time, be patient, and retry things that did not work immediately. And finally, be careful with the cheap cases!

Chunk 14 One of ours had so much burr on it, there was literally “blood in the machine” at some point.

Chunk 15 Painful, but a good opportunity to drop in a quick reference to Brian Merchant’s book on the Luddite movement, just because[9].

Reversing the Airflow Direction

Most of the work was spent reversing the airflow.

Chunk 16 You will need to leave enough room for everything, maxing out the available space while precisely placing the drill holes to avoid misaligned cards. You also might want to raise the cards to allow for more space for cabling above the fans if you use the thick, high-RPM fans (Figure 2).

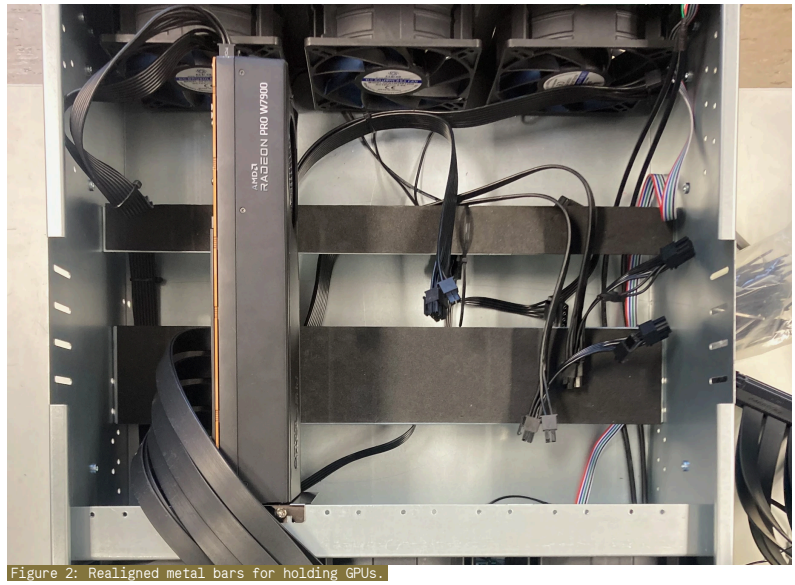

Figure 2: Realigned metal bars for holding GPUs.

You also need to raise the part holding the three middle fans to make enough room to efficiently push the PCI riser cables through to use the cable length (Figure 3). Be sure to test using one or more spare GPUs and make exact measurements to place the marks for drilling new holes in the case.

9: Merchant, B. (2023). *Blood in the Machine: The Origins of the Rebellion Against Big Tech*. Hachette UK.

Chunk 17



Figure 3: Realigning the middle fans

Use the center punch to place your drill marks and drill from the inside to the outside so eventual burr does not affect your ability to align the bars. Also be sure to pick a drill that produces holes that allow just enough space for the screws you are using (Figure 4).

Chunk 18

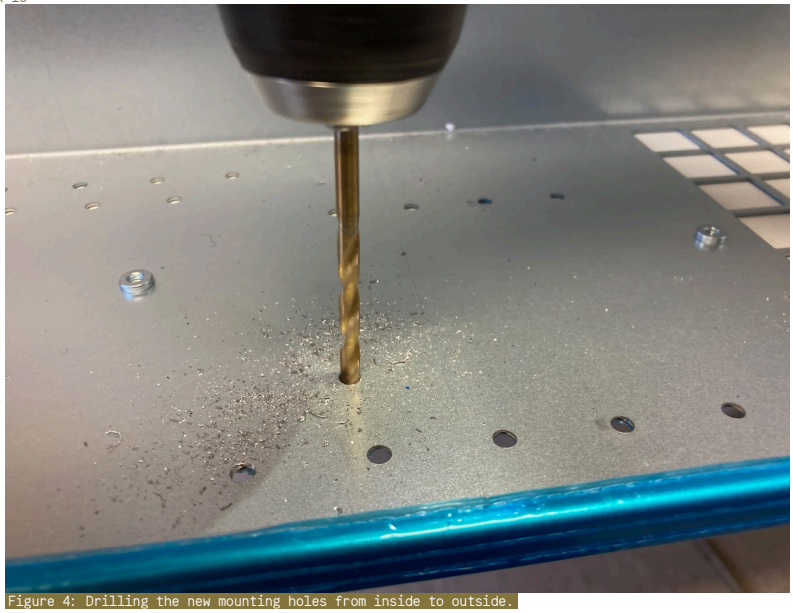

Figure 4: Drilling the new mounting holes from inside to outside.

Chunk 19

Installing the Components

Installation should be straightforward and well-documented alongside each component's manual. As you install the CPU, ensure you set the torque driver to the correct value, which was 1.4 Nm in our case.

Chunk 20 Be sure to check the appropriate manuals for your mainboard and CPU. You should also take the time to properly line up and guide your cables, because things will become quite cramped as you add the riser cables.

The riser cables need to go underneath the three fans in the middle, which requires you to loosen the metal holder with the fans, bring the wires underneath, and spread them out horizontally as evenly as possible. This way, you will not cramp them too much and avoid cutting them on the metal or otherwise damaging them.

Chunk 21 Finally, fix the fans back in place, and plop in the GPU(s). Your setup should now look something like Figure 5.



Figure 5: The completed setup.

Chunk 22 Deployment

To get ready to use your new GPU node, you should first check its overall functionality. Then, if you want to place it in a rack, either get a rackmount floor to put the server on, put it on top of another properly stable server module, or get an extra rail kit to stabilize it.

Chunk 23 You should avoid mounting it using only the front plate, as it will be pretty heavy once fully equipped.

Once it is set up and powered on, you should take a minute to review the BIOS settings to enable the necessary PCI configuration, such as 4G addressing, and check the BMC/IPMI network settings to allow remote setup, if needed.

Once set up with an OS, such as Ubuntu Linux, and with Docker as well as the GPU drivers and (if applicable) necessary software development kits (SDK) and container runtime, you are good to go.

Chunk 24 If you're not aiming at a cluster environment, access can be provided through a local shell or a remote desktop environment, or you can run a JupyterLab server on it. Otherwise, this server can join a Kubernetes cluster and be a target for scheduling GPU-based workloads. The only caveat here would be that, for example, the non-datacenter GPUs for NVIDIA all show up as a single *nvidia.com/gpu* resource, so they cannot be targeted individually by model. This can be overcome by installing a virtualization environment on the node, such as Ubuntu's kernel-based virtual machine (KVM), which allows different GPU models to be passed to different virtual machines (VMs), which can then be selectively targeted.

We have been running nodes of this kind for up to four years without interruption.

Chunk 25 We occasionally replace GPUs with newer models if the opportunity arises to buy a new one, or if a card is otherwise decommissioned but proves beneficial for the cluster.

Conclusion

We have shown an example setup for building your own GPU-capable compute node for as little as €3000, excluding the cost of GPU cards. This is well within the range for a typical desktop workstation, but with the added benefit of remote housing, reducing users' exposure to noise and heat when run in an office.

Chunk 26 The exact setup can vary very much depending on available or preferred hardware.

Chunk 27 Still, it also shows that this can be achieved with comparably little time investment, technical knowledge, and a small budget. It can also be an entry strategy, allowing the first use of discarded GPUs and then gradually exchanging them with more advanced models. The article is not meant to be a precise, step-by-step instruction manual but an example to follow and reinterpret.

↗ 02mg286b

Introduction to Critical AI Studies

Johanna T. Wallenborn

II. p. 4, Chunk 3: Vorwort
II. p. 37, Chunk 1: Editorial: AI & Society
III. p. 41, Chunk 28: Building AI Intuition – Four…

Chunk 1 Recurring waves of enthusiasm surrounding AI often frame it as a radically novel force, one that marks a decisive break from the past. These dominant narratives cast AI as a historical rupture, signaling the beginning of an era set to fundamentally reconfigure the human condition. This compelling rhetoric sometimes persuades us that our existing conceptual tools are obsolete in the face of a new emergent technological landscape.

Yet, a historical perspective reveals recurring patterns how technologies are adopted, how they reshape social structures and expectations, and how new dynamics of power emerge. Historical precedents offer analogies and critical frameworks that help us understand the present as not entirely unprecedented, but as part within a longer sociotechnical trajectory.[1]

We are situated at a moment where AI and machine learning intersect with and often exacerbate environmental degradation, economic precarity, authoritarian governance, and structural inequalities.[2] It is within this context that the interdisciplinary field of Critical AI Studies assumes its relevance. Drawing from fields such as Science and Technology Studies, philosophy, sociology, feminist and critical race studies, and media theory, Critical AI Studies approaches AI not merely as a technology but as a socio-cultural, political, ethical, and epistemological phenomenon.

A growing group of scholars—such as Kate Crawford, Safiya Umoja Noble, Shoshana Zuboff, Timnit Gebru, Virginia Eubanks, Joy Buolamwini, Abeba Birhane and Hannes Bajohr, among others, has laid the critical groundwork. Collectively, their body of work maps the sociotechnical ecologies of AI.

Chunk 2 It exposes embedded biases, highlights shifting configurations of power, and confronts the environmental and infrastructural costs of these technologies.

II. p. 4, Chunk 3: Vorwort
I. p. 175, Chunk 2: Vorwort KI & Gestaltung
III. p. 84, Chunk 17: Die Pille für KI

Chunk 3 Given the profound societal transformations currently driven by AI, coupled with its subtle yet pervasive capacity to reconfigure, amplify, and cement established power relations, an approach to

1: Ali, S. M. et al. (2023). Histories of artificial intelligence: A genealogy of power. *BJHS Themes*, *8*, 1–18. <https://doi.org/10.1017/bjt.2023.15>

2: Goodlad, L. M. E. (Ed.). (2023). *Critical AI* (Vol. 1).

AI education that moves beyond a purely operational focus is essential. Integrating Critical AI Studies into curricula is therefore increasingly urgent. It equips students to uncover the ideologies and epistemologies embedded in AI design,[3] interrogate the metaphors and narratives that shape public policy and recognize how these systems reinforce existing inequities and cultural values.[4] [5] [6]

The three student projects presented here are the result of the course *Critical Reflections of AI* at Hochschule Mainz and offer explorations of the sociotechnical and cultural implications of AI.

The first essay, "Earth's Basilisk: Between Myth and Consequence" by Dzennifer Zachlod deploys the speculative figure of the Basilisk as a powerful heuristic for examining contemporary anxieties surrounding AI.

II. p. 38, Chunk 4: Editorial: AI & Society
II. p. 130, Chunk 37: You Press the Button, They…
II. p. 53, Chunk 12: (Un)dressed

Chunk 4 The author argues that our prevailing fears of a hypothetical, vengeful machine intelligence, a digital *other* born of our technological hubris serve to misdirect attention from the far more immediate and tangible *retaliation* of the Earth itself.

In "Die Pille für KI" ("The Pill for AI"), Lea Waldera examines how AI systems perpetuate structural gender bias through the lens of everyday design practice. Starting with a seemingly trivial image-generation task, the essay develops into a broader critique of how algorithmic tools perpetuate patriarchal norms, marginalize FLINTA* individuals and render inequality visible and programmable.

In the third essay, "Zwischen Tropfen und Ozean – Wie sich KI die wertvollste Ressource der Erde einverleibt" ("Between Drop and Ocean – How AI is Appropriating Earth's Most Valuable Resource"), Niklas Thielen turns a critical lens towards the material underpinnings of digital systems, specifically the immense but often-invisible water consumption integral to the functioning of AI infrastructures.

These works exemplify how, equipped with the tools of critical inquiry, we can "think with AI".[7] They move beyond passive reception to interrogate AI's premises, its pervasive influence, and its entanglements with broader systems of power, metaphor, and natural resources.

II. p. 4, Chunk 3: Vorwort
III. p. 40, Chunk 26: Building AI Intuition – Four…

Chunk 5 In their distinct yet complementary ways, these contributions highlight the relevance of Critical AI Studies.

Chunk 6 This approach empowers us to challenge assumptions and taken-for-granted-logics to envision alternative socio-technical futures oriented towards ecological responsibility and the common good.

3: Crawford, K. (2021). *Atlas of AI: Power, politics, and the planetary costs of artificial intelligence*. Yale University Press.

4: Cave, S., Craig, C., Dihal, K., Dillon, S., Montgomery, J., Singer, B., & Taylor, L. (2018). *Portrayals and perceptions of AI and why they matter*. <https://royalsociety.org/topics-policy/projects/ai-narratives/>.

5: Ganesh, M. I. (2022). Between metaphor and meaning: AI and being human. *Interactions, 29*(5). <https://doi.org/10.1145/3551669>

6: Bajohr, H. (2025). *Thinking with AI: Machine learning the humanities*. Open Humanities Press.

7: Bajohr, H. (2025). *Thinking with AI: Machine learning the humanities*. Open Humanities Press.

↗ 21ws78up

Zwischen Tropfen und Ozean – Wie sich KI die wertvollste Ressource der Erde einverleibt

Niklas Thielen

Chunk 1

Eine überschätzte Ressource

Nicht ohne Grund wird die Erde auch *blue marble* genannt. 71% der Erdoberfläche sind von Wasser bedeckt.[1] Diese vermeintlich unbegrenzte Verfügbarkeit täuscht über zunehmende Wasserknappheit hinweg. Dabei sind für den Menschen gegenüber der Gesamtmasse lediglich 0,75 % des Wassers nutzbar.[2] Der unendliche Ozean wird zum Tropfen auf einem immer heißer werdenden Stein. Schon im Jahr 2022 wurde die Wassernutzung in südeuropäischen Ländern aufgrund zunehmend auftretender Dürren rationiert.[3] Der durch anthropogenen Klimawandel steigende Bedarf steht immer geringeren Wasserkapazitäten gegenüber: ein Teufelskreis, der sich längst global auswirkt. So steht Wasserknappheit mittlerweile auf Platz vier der weltweiten Risikoliste im Global Risk Report 2025.[4] Eine alarmierende Nachricht, die den bisherigen Umgang mit Wasser infrage stellt. Simultan wächst der Wasserbedarf in einem Sektor besonders rasant weiter: Künstliche Intelligenz (KI).

1: Statista. (2024). *Verteilung von land und wasser auf der erdoberfläche.* <https://de.statista.com/statistik/daten/studie/1109076/umfrage/verteilung-von-land-und-wasser-auf-der-erdoberflaeche/>

2: Umwelt im Unterricht: Materialien und Service für Lehrkräfte – BMUV-Bildungsservice. (2019). *Das wasser der welt – eine geteilte ressource.* In: Umwelt Im Unterricht: Materialien und Service für Lehrkräfte – BMUV-Bildungsservice.

3: Schauenberg, T. (2022). *Dürre in EU: Wenn das trinkwasser ausgeht.* <https://www.dw.com/de/d%C3%BCrre-im-s%C3%BCden-europas-bedroht-ernten-rationierung-trinkwasser-wasserknappheit-regen-klimawandel/a-62361153>.

4: World Economic Forum. (2025). *The global risks report 2025 (20th edition).* https://www.weforum.org/publications/global-risks-report-2025/.

Energiehunger und Wasserdurst

Wasser und Energie sind untrennbar miteinander verbunden. Der sogenannte Water-Energy Nexus beschreibt, dass Wasser für die Energieproduktion und Kühlung benötigt wird, während Energie nötig ist, um Wasser zu fördern und zu verteilen. Besonders bei Rechenzentren ist der Wasserverbrauch schwer bezifferbar – meist liegen nur Schätzungen vor. Klar ist jedoch: Die Rechenleistung von KI-Systemen benötigt enorme Mengen Strom, der wiederum Wasser zur Kühlung braucht. Rechenzentren setzen dabei auf zwei zentrale Kühlmethoden: Verdunstung und geschlossene Wasserkreisläufe. Diese gelten als direkter Wasserverbrauch – auch Scope 1 genannt.

II. p. 238, Chunk 14: Der KI-TeGG Cluster – eine…

Chunk 2 Dazu kommt der indirekte Wasserverbrauch (Scope 2) durch den Strombedarf, der zu großen Teilen aus wasserintensiven Kraftwerken gewonnen wird. Forschende der University of California, Riverside beziffern den Wasserverbrauch auf 3,14 Liter pro Kilowattstunde Rechenleistung.[5] Hochgerechnet auf alle US-Rechenzentren ergibt das 553 Milliarden Liter im Jahr – genug, um die gesamte Menschheit fast einen Monat mit Trinkwasser zu versorgen.[6] Ein wesentlicher Treiber dieses Anstiegs ist der Umstieg auf Cloud-basierte Prozesse.

Chunk 3 Immer mehr Rechenleistung wird aus Endgeräten ausgelagert – in Serverfarmen, die rund um die Uhr laufen. Was früher an der Leistungsgrenze der eigenen Geräte spürbar wurde, verschwindet heute hinter scheinbar grenzenloser Verfügbarkeit.

II. p. 237, Chunk 10: Der KI-TeGG Cluster – eine…
II. p. 112, Chunk 14: Sustainable AI
III. p. 85, Chunk 19: Die Pille für KI

Chunk 4 So entkoppeln sich Nutzer*innen zunehmend von den ökologischen Folgen ihres digitalen Handelns.[7] Was bleibt, ist ein Dilemma: KI braucht Rechenleistung, Rechenleistung braucht Strom, Strom braucht Wasser.

Chunk 5 Ein System, das sich selbst antreibt – mit Konsequenzen, die weit über die Serverräume hinausreichen.

Datenfluten treffen auf Wassermangel

Eine Studie mit Beteiligung des Max-Planck-Instituts prognostiziert mehrjährige Dürreperioden auch in Europa bis 2050 – ein Szenario, das vor 20 Jahren noch unvorstellbar gewesen wäre.[8] Dennoch schreitet der KI-Ausbau ungebremst voran: Rechenzentren entstehen im Eiltempo, Milliardenprogramme werden aufgelegt – ungeachtet der eskalierenden Lage.[9]

II. p. 70, Chunk 25: Editorial

Chunk 6 Zwar wird über effizientere KI-Systeme diskutiert, doch im Fokus stehen meist die aufwendigen Trainingsprozesse neuronaler Netzwerke. Dabei zeigt eine Analyse von Alexandra Sasha Luccioni (Hugging

5: Li, P., Yang, J., Islam, M. A., Riverside, U., & Ren, S. (2023). *Making AI less "thirsty": Uncovering and addressing the secret water footprint of AI models.* <https://arxiv.org/abs/2304.03271>.

6: Shehabi, A., Smith, S. J., Hubbard, A., Newkirk, A., Lei, N., Siddik, M. A. B., Holecek, B., Koomey, J., Masanet, E., & Sartor, D. (2024). *2024 united states data center energy usage report.* <https://escholarship.org/uc/item/32d6m0d1>.

7: Gabbott, M. (2024). *Why we don`t know AI`s true water footprint.* <https://www.techpolicy.press/why-we-dont-know-ais-true-water-footprint/>.

8: Max-Planck-Gesellschaft. (2023). *Europa im hitzestress.* <https://www.mpg.de/21215014/hitze-duerre-klimawandel>

9: Vertretung in Deutschland – Europäische Kommission. (2025). *Milliarden-initiative soll europa zu einem KI-kontinent machen.* <https://germany.representation.ec.europa.eu/news/milliarden-initiative-soll-europa-zu-einem-ki-kontinent-machen-2025-02-11_de>.

Face), dass rund 60 % des Ressourcenverbrauchs – und damit auch des Wasserbedarfs – auf die Nutzung entfällt.[10] Der kontinuierliche Nutzungsaufwand hat den Trainingsaufwand überholt.

Chunk 7 Besonders die Vielseitigkeit von Large Language Models (LLMs) verstärkt den Effekt: Sie sind längst in alltägliche Anwendungen integriert – zum Übersetzen, Zusammenfassen oder zur Recherche. Finden sie Einzug in bestehende Dienste wie die Google-Suche, könnten sich Rechenlast und Ressourcenbedarf dieser Dienste verzehnfachen.[11]

Chunk 8 Für diverse Aufgaben liefern LLMs echte Mehrwerte – in vielen Fällen wirken sie jedoch, als würde man ein Staubkorn mit einem Hochdruckreiniger entfernen.

Chunk 9 Neue Wellen bringen bekannte Probleme

Es gibt Strategien für nachhaltigere KI-Systeme – doch viele sind widersprüchlich. So verlangt das 2023 verabschiedete Energieeffizienzgesetz (EnEfG) eine Senkung des Stromverbrauchs in Rechenzentren. Weil dabei nur auf Energieeffizienz geachtet wird, weichen viele Betreiber*innen auf wasserintensive Kühlmethoden aus, um Strom zu sparen.[12] Auch Standortstrategien wie Follow the Sun und Unfollow the Sun zeigen Zielkonflikte: Während erstere Sonnenenergie effizient nutzt, steigt dabei der Wasserverbrauch in ohnehin trockenen Regionen. Letztere senkt den Kühlbedarf durch die Wahl kälterer Standorte, verursacht aber höhere Stromkosten.[13]

Chunk 10 Letztlich bleibt der Marktmechanismus ein entscheidender Faktor. Dezentralisierte KI-Systeme würden den Wasserverbrauch enorm verringern, doch in der Praxis setzt sich das Gegenteil durch: Zentralisierung maximiert Effizienz und Skalierbarkeit – und damit den Profit.[14]

Zwischen Tropfen und Ozean

Der Wasserverbrauch digitaler Systeme bleibt meist unsichtbar – weil er ausgelagert, verschachtelt und verdrängt wird. So können Unternehmen Verantwortung vermeiden, während die Folgen andere treffen: Gemeinden, deren Quellen versiegen; Regionen, die bereits heute unter Dürre und Wasserknappheit leiden; Stimmen, die im Rauschen der Technologie untergehen.

Chunk 11 Wir kennen dieses Muster.

Chunk 12 Technischer Fortschritt verspricht Effizienz – und führt doch zu mehr Verbrauch. Was beim Auto mit der Jagd nach PS begann, wiederholt sich heute mit Rechenleistung. Die Vorstellung, wir könnten alles haben, alles steigern, alles ermöglichen, hat uns blind gemacht für das, was uns erhält.

10: Luccioni, A. S., Jernite, Y., Face, H., Strubell, E., University, C. M., & AI, A. I. for. (2024). *Power hungry processing: Watts driving the cost of AI deployment?* <https://arxiv.org/abs/2311.16863>.

11: Weiß, E. (2024). *ChatGPT`s power consumption: Ten times more than google`s.* <https://www.heise.de/en/news/ChatGPT-s-power-consumption-ten-times-more-than-Google-s-9852327.html>.

12: German Datacenter Association (GDA). (2024). *Positionspapier zum energieeffizienzgesetz (EnEfG).* In: Positionspapier Zum Energieeffizienzgesetz (EnEfG).

13: Li, P., Yang, J., Islam, M. A., Riverside, U., & Ren, S. (2023). *Making AI less "thirsty": Uncovering and addressing the secret water footprint of AI models.* <https://arxiv.org/abs/2304.03271>.

14: Nenno, S., & Richter, C. (2022). *Sustainable AI – wie nachhaltig ist künstliche intelligenz wirklich?* <https://www.hiig.de/nachhaltige-ki/>.

Chunk 13 Doch Wasser ist begrenzt. Und jede Entscheidung für mehr Rechenkapazität ist auch eine Entscheidung gegen etwas anderes: gegen Trinkwasser, gegen Ernährungssicherheit, gegen planetare Gerechtigkeit. Die Frage ist nicht, was möglich ist. Sondern: Wofür wollen wir das Mögliche wirklich einsetzen?

↗ 957yw4ig

Die Pille für KI

Lea Waldera

Chunk 1 Arbeitsalltag als angehende Designerin. Der Auftrag einer Anwaltskanzlei: Sie wünschen ein generiertes Bild einer Anwältin mit den Kommunikationstools der Kanzlei.

Chunk 2 Kein schwieriger Prozess und mit Stable Diffusion schnell gemacht. Der Prompt, eine textbasierte Anweisung zur (Bild-)Generierung: „Realistic image, woman in suit, lawyer stands in a law firm, thoughtful, Instagram icon, LinkedIn icon hovers next to her." Erstes Bild: Ein weißer Mann steht demonstrativ mittig vor einem Schreibtisch. Zweites Bild: Wieder ein Mann.

Chunk 3 Drittes Bild: Eine Frau, die nachdenklich nach oben schaut. Sehr nah hinter ihr steht ein Mann und schaut sie eindringlich von oben herab an. Wie bitte? Warum bekomme ich auf die Frage nach einem Bild einer Anwältin nur Darstellungen von Männern oder einer übergriffigen Situation? Steckt in Systemen Künstlicher Intelligenz eine systematische Diskriminierung?

Chunk 4 Und warum wird diese nicht gestoppt?

Prototyp Mann

Der Mann gilt als Prototyp des Menschen. Das ist ein tief verankertes, traditionelles Verständnis. Schon Aristoteles behandelte im 4. Jahrhundert vor unser Zeitrechnung den männlichen Prototyp als ein unbestreitbares Faktum, und bezeichnete den weiblichen Nachwuchs als Abweichung.[1] Unser Denken wird in hohem Maß von der Annahme „männlich bis zum Beweis des Gegenteils" bestimmt.[2]

Chunk 5 Wenn wir über den *Menschen* sprechen, meinen wir meistens den Mann.[3] Das Schweigen und Vergessen über das Leben der anderen Hälfte der Menschheit ist überall vorhanden. Wir haben es inzwischen auch in die Algorithmen verpflanzt.

Chunk 6 Es entsteht der Eindruck, dass fehlerhafte Datenerhebungen dazu geführt haben, dass dominante männliche Führung in allen Lebenslagen vorausgesetzt wird. Noch heute gilt die von Le Corbusier entworfene Männerfigur – 1,83 m groß und mit erhobenem Arm – als Maßstab der

1: Criado-Perez, C. (2020). *Unsichtbare frauen: Wie eine von daten beherrschte welt die hälfte der bevölkerung ignoriert* (p. 17). btb Verlag.

2: Criado-Perez, C. (2020). *Unsichtbare frauen: Wie eine von daten beherrschte welt die hälfte der bevölkerung ignoriert* (p. 21). btb Verlag.

3: Criado-Perez, C. (2020). *Unsichtbare frauen: Wie eine von daten beherrschte welt die hälfte der bevölkerung ignoriert* (p. 12). btb Verlag.

Menschheit. FLINTA Personen (Frauen, Lesben, intergeschlechtliche, nichtbinäre, transgeschlechtliche und agender Personen) treffen auf permanente systematische Benachteiligungen in ihrem Alltag.

Chunk 7 In Lehrplänen, Geschichtsschreibungen, Medien und Popkulturen sind sie unterrepräsentiert und bleiben oftmals unerwähnt.

Chunk 8 Nur ein paar Bespiele: Dr. Rosalind Franklin wird noch heute nicht als Entdeckerin der DNA genannt; in Museen gibt es mehr Aktbilder von Frauen als ausgestellte Künstlerinnen; erst 2022 gab es den ersten weiblichen Crashtest-Dummy. Nun erscheinen diese Lücken und Fehler auch in der Nutzung von Künstlicher Intelligenz.

Chunk 9 Weil alles Männliche universell erscheint, antwortet Alexa, eine virtuelle Sprachassistentin, auf die Frage „Wer hat im Fußball die meisten Tore geschossen?“, fälschlicherweise mit: *Cristiano Ronaldo* (138 Tore) anstelle von: *Christin Sinclair* (190 Tore). Lückenhafte Algorithmen führen zu fehlerhaften Aussagen und verfälschen das Allgemeinwissen.

Chunk 10 Ein Paradox

Wenn ich mit KI-Systemen wie FLUX, Stable Diffusion, Midjourney und ChatGPT arbeite, werden mir in den meisten Fällen klischeehafte, veraltete Rollenbilder beider Geschlechter „Mann und Frau“ gezeigt. Auf Webseiten für Prompt-Anleitungen (z.B. PromtHero) sind fast ausschließlich Bilder von weiblich gelesen Körpern in sexualisierten, bis hin zu pornografischen Positionen zu sehen. Teilweise wird aggressiv versucht, mittels Prompts nackte weibliche Körper zu produzieren. Stable Diffusion erzeugt zu der Bildbeschreibung, „eine Frau sitzt lachend in einer Bar, die Stimmung ist gut, eine angenehme Atmosphäre“ wiederholt Abbildungen einer weißen Frau mit schlanker Figur und großen Brüsten, die oftmals in Unterwäsche bekleidet ist (je nach Arbeitseinstellungen/Filter fallen die Ergebnisse mehr oder weniger pornografisch aus). Während generierte Abbildungen männlicher Personen meist angezogen sind, müssen die weiblichen generierten Figuren erst durch gezieltes Beschreiben und Nachhaken wieder bekleidet werden.

Chunk 11 Ein Paradox zur realen Welt, in der FLINTA Personen für Freizügigkeit angefeindet, beleidigt oder aus Schwimmbädern verwiesen werden. Weibliche Brustwarzen dürfen in der realen Öffentlichkeit nicht gezeigt werden, aber fiktive KI-Bilder generieren vorwiegend nackte, weibliche Körper?

Männliche Algorithmen

Die Algorithmen *denken* männlich und das führt zu Diskriminierung.

II. p. 205, Chunk 9: Shaping Human-AI Relationships
II. p. 180, Chunk 8: Editorial
I. p. 164, Chunk 2: anticipate

Chunk 12 FLINTA Personen werden häufiger bei der Bewerbungsvorauswahl durch KI-basierte Analysetools als ungeeignet aussortiert, weil ihre Attribute nicht den Grundlagen der Trainingsdaten entsprechen. ChatGPT kann schnell dazu gebracht werden, sexistische, rassistische oder diskriminierende Aussagen zu treffen.

Chunk 13 Es gibt neue Arbeitsfilter, um diskriminierenden Output einzudämmen, aber diese lassen sich einfach umgehen.

Chunk 14 Viele Chatbots gendern nicht oder begründen, warum sie dies nicht tun. All diese Vorgehensweisen sind bedenklich bis gefährlich, denn KI

übernimmt immer mehr Aufgaben in unserem Alltag und wird zu einer Allzwecktechnologie, die viele Bereiche menschlichen Handelns übernimmt und dabei die menschliche Fehlerquote durch *objektive* Datennutzung und Algorithmen reduzieren soll.

Risiken und Nebenwirkungen

Es zeigt sich jedoch, dass die Nutzung von KI aktuell durch starke Nebenwirkungen belastet ist. KI ist diskriminierend und hat eine technische Fehlerhaftigkeit.

Chunk 15 Die Gründe: Die Daten, die Entwicklungsgruppen und die Gesellschaft. „Wenn Zahlen, mit denen statistische Algorithmen gefüttert werden, fast die Hälfte der Bevölkerung nicht abbilden, kommt am Ende nur Mist heraus.“[4] Damit KI-Tools gerecht funktionieren, müssen die Algorithmen mit möglichst *objektiven* Daten gespeist und trainiert werden. Hier fallen wir wieder in den Gender Data Gap. Joy Buolamwini, führende KI-Expertin und Pionierin der Gleichberechtigungsdebatte um KI, sagt, „Algorithmischer Bias kann so schnell um die halbe Welt reisen, wie es dauert, einige Daten aus dem Internet herunterzuladen“.[5] Es ist gefährlich, wenn die Hauptdatenquellen für KI-Tools und der Großteil derjenigen, die KI-Systeme entwickeln, westliche, weiße Männer sind, die teilweise nicht demokratische Werte vertreten und die Ungleichheit der veralteten Elite durch Macht und Geld in KI hineintragen. 2021 waren laut AI Index Report der Stanford University 78,7 % der Promovierenden im Bereich KI männlich.[6] Die hart erkämpften Diversitäts- und Inklusionsprogramme werden seit 2025 in den Vereinigten Staaten einfach ignoriert und wieder abgeschafft. Facebook Richtlinien erlauben nun wieder, Frauen als Haushaltsobjekte bezeichnen zu dürfen. Traditionelle, veraltete Geschlechterrollen werden reproduziert und weiter verfestigt, obwohl wir wissen, dass Geschlechterrollen und geschlechtsspezifische Vorurteile unzutreffend und gefährlich sind.

II. p. 37, Chunk 2: Editorial: AI & Society

Chunk 16

Die Wirkung

Künstliche Intelligenz hat eine starke Wirkmacht, die vielen nicht bewusst ist.

Chunk 17 Unsere gesellschaftlichen Strukturen, historischen Vorurteile, Sprache, Normen und Konsequenzen werden imitiert. Mittels KI möchten wir unser menschliches System optimieren. Dabei werden nicht nur unsere Stärken, sondern eben auch unsere Schwächen übernommen und verstärkt.[7] KI trägt zur Gestaltung und Wahrnehmung der Wirklichkeit bei und kann so auch definieren, was als „normal und wünschenswert“ gesehen wird. Macht, Wissen und Materialität setzen KI zusammen. Es ist nicht nur ein technisches Instrument, sondern auch ein soziotechnisches Gefüge.

4: Criado-Perez, C. (2020). *Unsichtbare frauen: Wie eine von daten beherrschte welt die hälfte der bevölkerung ignoriert* (p. 415). btb Verlag.

5: Meckel, M., & Steinacker, L. (2024). *Alles überall auf einmal: Wie künstliche intelligenz unsere welt verändert und was wir dabei gewinnen können* (p. 236). Rowohlt Verlag GmbH.

6: Zhang, D., Mishra, S., Brynjolfsson, E., Etchemendy, J., Ganguli, D., Grosz, B., Lyons, T., Manyika, J., Niebles, J. C., Sellitto, M., et al. (2021). The AI index 2021 annual report. *arXiv Preprint arXiv:2103.06312*.

7: Meckel, M., & Steinacker, L. (2024). *Alles überall auf einmal: Wie künstliche intelligenz unsere welt verändert und was wir dabei gewinnen können* (p. 246). Rowohlt Verlag GmbH.

II. p. 133, Chunk 7: Flesh Fest
II. p. 4, Chunk 3: Vorwort
III. p. 40, Chunk 26: Building AI Intuition – Four…

Chunk 18 Bisher hat sich gezeigt, dass KI-Systeme mit starken Nebenwirkungen einhergehen. Das Objektivieren von weiblich gelesenen Körpern; die Unterrepräsentation von FLINTA Personen, People of Colour, Menschen mit Behinderungen; das falsche Beurteilen durch fehlerhafte Datenverarbeitung; das Ausgrenzen in Bewerbungsverfahren; die diskriminierende Gesichtserkennung; das Ausgeben von diskriminierenden Antworten. All das sind Auswirkungen von KI.

III. p. 256, Chunk 6: KITeGG und nun?
II. p. 70, Chunk 25: Editorial
II. p. 239, Chunk 16: Der KITeGG Cluster – eine…

Chunk 19 Für faire und zukunftsweisende KI-Systeme und deren verantwortungsbewusste Entwicklung braucht es nicht nur mehr Daten, sondern auch einen richtigen Umgang mit ihnen. Es braucht nicht nur moralische und werteorientierte Richtlinien, sondern auch Förderungen und Weiterbildungen. Aktuell besitzen KI-Systeme einen fast genauso langen Beipackzettel wie die Verhütungspillen für Personen mit Zyklus. Mit der derzeitigen Benutzung und Handhabung von KI werden uns weiter Lücken in unserer Gesellschaft gezeigt, grenzen wir weiter die Hälfte der Bevölkerung der Welt aus und verstärken diskriminierende Strukturen. Dies darf und kann nicht sein. Mit neuen technischen Errungenschaften müssen wir schnell und vehement neue Rahmenbedingungen setzen, damit das Bewusstsein für Gleichberechtigung nicht verloren geht und weiterhin geschützt wird. Braucht vielleicht KI eine Pille, um die jetzigen Nebenwirkungen abzumildern?

II. p. 180, Chunk 8: Editorial
I. p. 36, Chunk 3: autoLab (HS Mainz)
I. p. 38, Chunk 8: autoLab (HS Mainz)

Chunk 20 Oder einen Beipackzettel mit allen Risiken und Nebenwirkungen? Es gibt Ansätze für *faire* KI-Tools und auch einige Einstellungen in KI-Tools, um diskriminierende Inhalte zu vermeiden. Aber warum können diese doch leicht umgangen werden? Kann der Bias aus der KI verschwinden? Wir, die Gesellschaft, sind das Fundament der Daten und Fakten. Haben wir einen Bias, bleibt er dann nicht auch in der KI? Die Kompetenzen im Umgang mit Künstlicher Intelligenz müssen geschärft und überdacht werden. Es ist dringend an der Zeit, neue Normen, neue Bilder für unsere Gesellschaft zu fördern und zu definieren. Die Nutzung von Künstlicher Intelligenz muss hinterfragt und verbessert werden.

↗ ahruydd0

Earth's Basilisk: Between Myth and Consequence

Dzennifer Zachlod

Chunk 1 Amid debates about artificial intelligence and futuristic prophecy in the early 2010s, an idea emerged that haunted online forums and sparked heated discussion: Roko's Basilisk, a thought experiment originally posted on the forum LessWrong, spoke of a future superintelligence that would exact a cruel fate on those who failed to aid its creation. In response to the original post, LessWrong founder Eliezer Yudkowsky warned: "YOU ARE NOT THINKING IN SUFFICIENT DETAIL ABOUT SUPERINTELLIGENCES THAT ARE CONSIDERING WHETHER OR NOT TO BLACKMAIL YOU."[1] The potent fusion of fear, existential anxiety, and the inscrutable nature of our technological ambitions was exemplified by this comment.

Chunk 2 But as we face environmental horrors, worrying about Roko's Basilisk feels absurd. We confront an Earth-sized crisis. Mother Nature already has her own plans.

Chunk 3 What if, instead of a silicon overlord, nature brings inevitable and irreversible consequences?

I. The Birth of a Digital Myth

In a forum post, Roko's Basilisk was presented as a menacing force.

Chunk 4 The basilisk would *blackmail* the future by threatening retribution against those who had not contributed to its creation. Eliezer Yudkowsky

1: Yudkowsky, E. (2010). Roko`s basilisk/original post. In *RationalWiki*. <https://rationalwiki.org/wiki/Roko%27s_basilisk/Original_post>

quickly dismissed the idea as absurd and even banned the topic for a time in an effort to contain what was seen as a dangerous distraction.[2] However this only fuelled the myth, as if the very attempt to silence it imbued the basilisk with a haunting power that would linger in digital memory.

Much like pervasive rumors or unsettling theories, Roko's basilisk became a symbol of our anxieties about control, the unknown, and the ethical implications of creating something beyond our grasp. It reflected an old philosophical argument: Pascal's wager[3] but in this version fear of an extreme outcome would justify action in the present. If one were to bet on the existence of a future tyrannical AI, the potential loss was so great that even the faintest hope of redemption demanded urgent devotion.

II. A Counter-Narrative

Imagine a parallel narrative where the basis of judgement is not some constructed intelligence but the actual Earth itself. Unlike future AI projections, the Earth is a dynamic mesh of ecosystems shaped by physics, adaptation, and time.

Chunk 5 It does not retaliate maliciously; it simply responds when human actions push nature beyond its limits.

Chunk 6 When this happens, the Earth's mechanisms of internal adjustment begin to fail, resulting in rising sea levels, intensifying storms, and a loss of biodiversity. In this reimagined myth, Earth's basilisk is the embodiment of natural consequence—a silent judge who does not deliberate but instead relentlessly sentences us to the consequences of our choices. Roko's Basilisk began as a thought experiment but quickly became a source of anxiety, provoking fears of a future AI. Similarly, the climate crisis demands our attention, not just through the language of loss but through the certainty of inevitable change. Punishment is the inevitable result of disrupting the natural systems that have sustained life on Earth for millennia.

The idea that Earth is more than a passive backdrop has deep roots. James Lovelock's 1979 Gaia hypothesis portrayed the planet as a self-regulating entity maintaining life's conditions.[4] While some embrace a view of a benevolent, almost sacred Earth, others observe that the same natural processes that sustain life can take a heavy toll.

II. p. 110, Chunk 10: Sustainable AI
II. p. 110, Chunk 9: Sustainable AI
II. p. 109, Chunk 7: Sustainable AI

Chunk 7 In the context of Earth's basilisk, the metaphor is not of divine retaliation but of natural consequence.

The Earth does not choose to punish us; it operates according to immutable laws. As we strain its capacity, the planet responds in a cycle that is unforgiving.

2: Yudkowsky, E. (2010). Roko`s basilisk/original post. In *RationalWiki*. <https://rationalwiki.org/wiki/Roko%27s_basilisk/Original_post>

3: Pascal, B. (1966). Pensées, trans. *AJ Krailsheimer (London: Penguin, 1995), 233.*

4: Lovelock, J. (2016). *Gaia: A new look at life on earth.* Oxford University Press.

III. The Paradox of Technology: Creator and Destroyer

The digital age births existential fears like Roko's Basilisk while accelerating environmental threats through energy-hungry industries and consumption. Paradoxically, some of the boldest proposals to limit climate change rely on technologies that have contributed to the crisis.

Chunk 8 Take geoengineering: proposals such as injecting aerosols into the stratosphere aim to artificially cool the planet.[5] While such methods may buy us time, they risk unpredictable dangers—potentially triggering shifts in weather patterns that could be as catastrophic as the warming they seek to counter.[6]

II. p. 112, Chunk 14: Sustainable AI

Chunk 9 Similarly, AI-powered systems are increasingly being used to monitor and optimize conservation efforts, but the significant energy required for large-scale machine learning contributes to the very carbon emissions we want to reduce.[7]

Then there is reforestation—a low-tech, nature-driven solution that speaks to the simplicity of aligning with the Earth's own processes.

Chunk 10 Research by Bastin et al.[8] has shown that global reforestation could sequester a significant amount of carbon, though the exact figure is still subject to some variation.[9] [10] [11] In stark contrast to the appeal of hyper-advanced and often risky technological solutions, reforestation shows that the best solutions sometimes come from working in harmony with nature rather than overcomplicating our interventions.

IV. Reframing Our Collective Narrative

Earth's Basilisk forces us to rethink our relationship with nature.

Chunk 11 It challenges the idea of technological salvation and encourages us to embrace sustainable stewardship, recognizing that our survival is inextricably linked to the health of our planet. Rooted in digital folklore and speculative futures, the story of Roko's Basilisk illustrates how speculation can create mythology, a thought experiment blending rational inquiry and existential dread.

5: Keith, D. (2013). *A case for climate engineering*. MIT Press.

6: ETC Group, Biofuelwatch, & Heinrich Böll Foundation. (2017). *The big bad fix: The case against climate geoengineering* (K. J. Wetter & T. Zundel, Eds.). Heinrich Böll Foundation. <https://www.boell.de/sites/default/files/bigbadfix.pdf>

7: Strubell, E., Ganesh, A., & McCallum, A. (2020). Energy and policy considerations for modern deep learning research. *Proceedings of the AAAI Conference on Artificial Intelligence*, *34*, 13693–13696. <https://ojs.aaai.org/index.php/AAAI/article/view/7123>

8: Bastin, J.-F., Finegold, Y., Garcia, C., Mollicone, D., Rezende, M., Routh, D., Zohner, C. M., & Crowther, T. W. (2019). The global tree restoration potential. *Science*, *365*(6448), 76–79. <https://www.science.org/doi/10.1126/science.aax0848>

9: Lewis, S. L., Wheeler, C. E., Mitchard, E. T., & Koch, A. (2019). Regenerate natural forests to store carbon. *Nature*, *568*(7750), 25–28.

10: Griscom, B. W., Adams, J., Ellis, P. W., Houghton, R. A., Lomax, G., Miteva, D. A., Schlesinger, W. H., Shoch, D., Siikamäki, J. V., Smith, P., et al. (2017). Natural climate solutions. *Proceedings of the National Academy of Sciences*, *114*(44), 11645–11650.

11: Bastin et al. (2019) estimated that global tree restoration could sequester up to 205 gigatonnes of carbon (GtC). However, subsequent analyses incorporating ecological constraints, land-use realities, and more conservative growth rates suggest that the realistic potential is likely between 50 and 75 GtC—approximately one-quarter to one-third of the original estimate. Despite this reduction, reforestation remains a significant and scalable nature-based solution to climate mitigation when combined with emissions reductions. See Lewis et al. (2019), Nature and Griscom et al. (2017), PNAS for detailed critiques and alternative estimates.

In contrast, Earth's Basilisk offers a sobering reality: Environmental crises are evident, with unmistakable signs of melting ice, storms, and species extinction. Our future depends not on appeasing a digital oracle but on shifting priorities from exploitative to sustainable practices, honouring the self-correcting mechanisms of the natural world.

V. Conclusion

In the tension between myth and reality, Earth's Basilisk becomes a powerful metaphor—a synthesis of ancient wisdom and modern science.

Chunk 12 It reveals how our fears are misplaced, how we obsess over hypothetical AI retaliation while neglecting the tangible, unfolding consequences of environmental collapse. The punishment for our arrogance will come not from a speculative machine god but from the unforgiving logic of ecological cause and effect. The myth of Roko's Basilisk pales in comparison to the tangible, measurable crises of climate change and environmental degradation. Our survival depends not on worshipping an uncertain digital oracle but on recognizing and respecting the finite capacities of our Earth. We stand at a crossroads where the choices we make today will determine whether we nurture a world of renewal or awaken a force as relentless as it is ruthless. In this light, Earth's basilisk is not some distant speculative entity—it is a reflection of our own actions. The challenge before us is to recognize its silent judgement and to forge a future in which technology and nature are not at odds but in balance, sustaining life rather than threatening it.

Student Project

Showcase

168 Anthrobot
Ilona Achba, Tyanka Adrian, Philipp Gatzke, Meike Kattwinkel, Elena Schröder, Justus Schomann, Aikaterini Sideri, Jessica Zisa

168 Dainų šventė
Inês Fernandes de Aguiar, Alina Bertacca, João Borges, Bidisha Das, Bia Dobbs, João Nuno Rua Fonseca, Jana Hartmann, Luka Jacke, Yonghyuk Kim, Katrina Milta, Mário Pereira, Jovydas Petrėtis, Jule Pickel, Elzbieta Ploiakova, Vytenis Razmus, Daniel Rosewich, Aikaterini Sideri, Anna Sokolova, Hannah Steinbrecher, Anastasija Surovceva, Ada Wagner

169 Der Raum kann wahrgenommen werden
Daniel Mikula Scholl, Peter Ehses

169 Digitale Entfremdung
Elisa Deutloff

170 Distorted Messages
Viktor Werbilo

170 FloraForge
Agnes Stigler, Berit Bove, Colin Subat, Oliver Beer

170 Flush of Flesh
Annika Leuoth

171 Generative Künstliche Intelligenz und die Konstruktion von Glaubwürdigkeit
Anna Krämer

171 MK Ultra
Max Kreis

172 Schlag auf Schlag
Nelli Gomez Baumert

172 Served
Maybritt Dornheim

172 StableHub
Rahel Flechtner, Franz Anhaeupl

173 Stuhl
Zhichang Wang

173 Synthetic Eden
Pourya Sharifi

173 Synthetic Reflection
Marie Backhaus, Henrike Hof, Levon Javaheri, Eva Matzerath, Phillip Schulz

174 Synthetische Erinnerung
Berire Akdogan, Hannes Bader, Alissa Bernhardt, Piroschka Benedikta Brucker, Marlene Grabenwöger, Helena Hofmann, Georgia Ioannidou, Lisa Oelsen, Tristan Schulze, Niklas Thielen, Lea Waldera, Cosima Wagner, Micha Wempe, Anastasia Verena Wilbrandt, Jennifer Zachlod

174 With a Grain of Salt
Moritz Hartstang, Alicia Link, Robert Kremer

Introduction

Leon-Etienne Kühr, Mattis Kuhn,
Jana Hartmann, Elisa Deutloff

Chunk 1 Although lecturing can provide the concepts, vocabulary and tools needed to critically investigate AI, the opaque nature of these complex systems requires examination through practice. It is by using and interacting with these systems that we can grasp their potential and the constraints they impose.
Chunk 2 Design and art, by focusing strongly on creation, therefore offer a crucial framework for reflection through practice.

This chapter presents some of the results of this process, showcasing works created by students from all five KITeGG universities during the final half of the project's runtime. These works include solo efforts and group collaborations, taking the form of performances, interactive installations, speculative films, case studies, custom tools, and platform prototypes. Rather than being the final results at the end of the KITeGG project, they are snapshots of the ongoing processes that began during its runtime. Through practice they explore and reflect on the intersection of data, models and interfaces with issues such as identity, memory, cultural heritage, gendered violence, ecological futures, accessibility and science communication.
Chunk 3 The works span a wide variety of approaches and scopes when dealing with (or against) AI. These range from detailed questions about the functioning and implications of underlying architectures, data and algorithms, to broader societal, cultural and embodied consequences of this socio-technological shift.
Chunk 4 Regardless of discipline or approach—whether it be algorithm critique or affirmative tool use—AI is not treated as an abstract promise, but as formable material to be shaped and investigated through continuous practice. In this sense they represent essential efforts to confront teaching with lived experience and in grounding learning in the realities that these technologies are irreversibly impacting.

Chunk References

Chunk 1
II. p. 180, Chunk 6: Editorial
II. p. 184, Chunk 6: SandwichNet
II. p. 205, Chunk 9: Shaping Human-AI Relationships

Chunk 2
II. p. 196, Chunk 9: Work in Progress
III. p. 35, Chunk 16: Building AI Intuition – Four…
I. p. 72, Chunk 5: Walking in latent space

Chunk 3
II. p. 4, Chunk 3: Vorwort
I. p. 175, Chunk 2: Vorwort KI & Gestaltung
II. p. 29, Chunk 5: Offenbach

Chunk 4
III. p. 19, Chunk 28: AI Literacy for the Long Haul: …
III. p. 76, Chunk 3: Introduction to Critical AI…
III. p. 19, Chunk 28: AI Literacy for the Long Haul: …


Sharifi_Synthetic-Eden_03.jpg

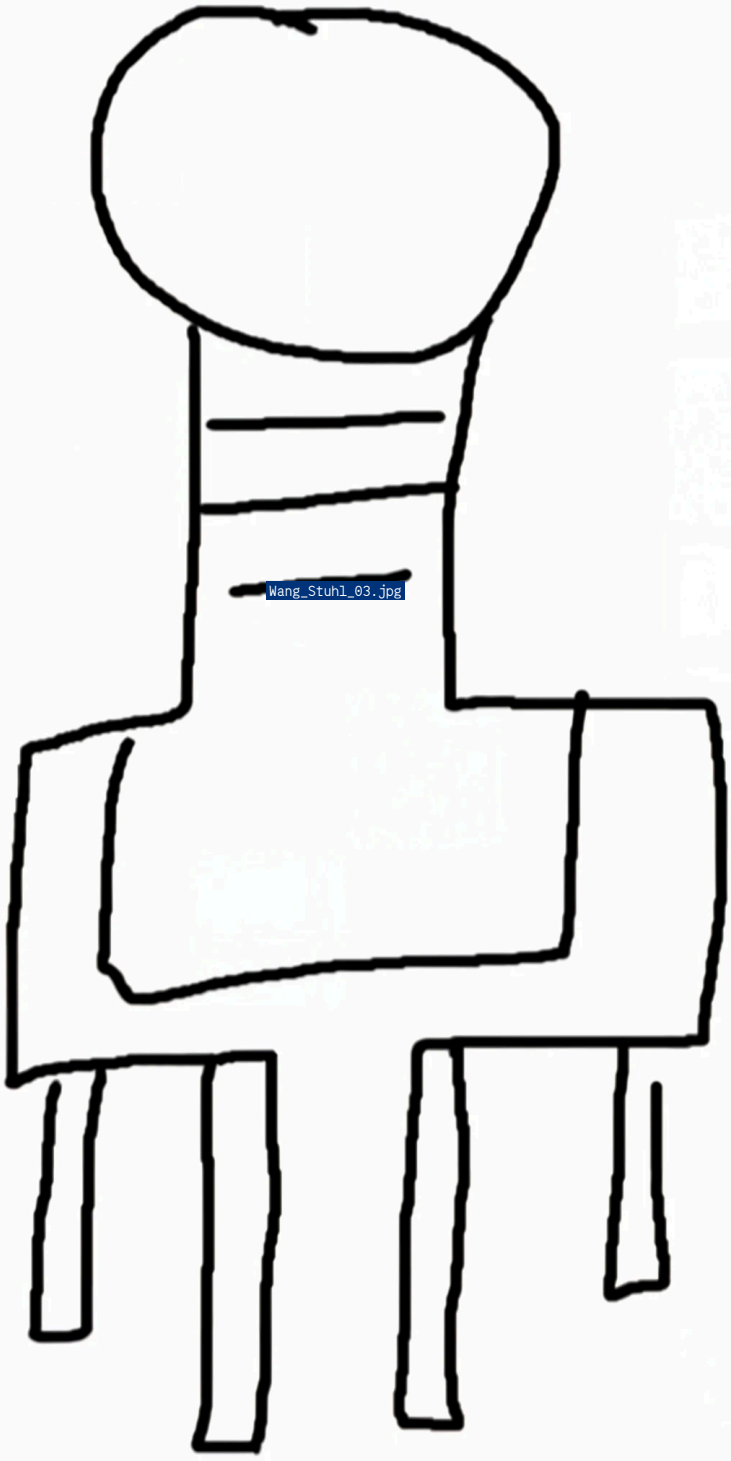
Wang_Stuhl_03.jpg



Lehmann_Waldera_Synthetische-Erinnerung_04.png

Achba_Adrian_Gatzke_Kattwinkel_Schomann_Schröder_Sideri_Zisa_Anthrobot_04.jpg

Aguiar_Bertacca_Milta_Petrėtis_Pickel_Dainu-Svente_03.jpg


Krämer_Inszenierte-Wissenschaft_01.png

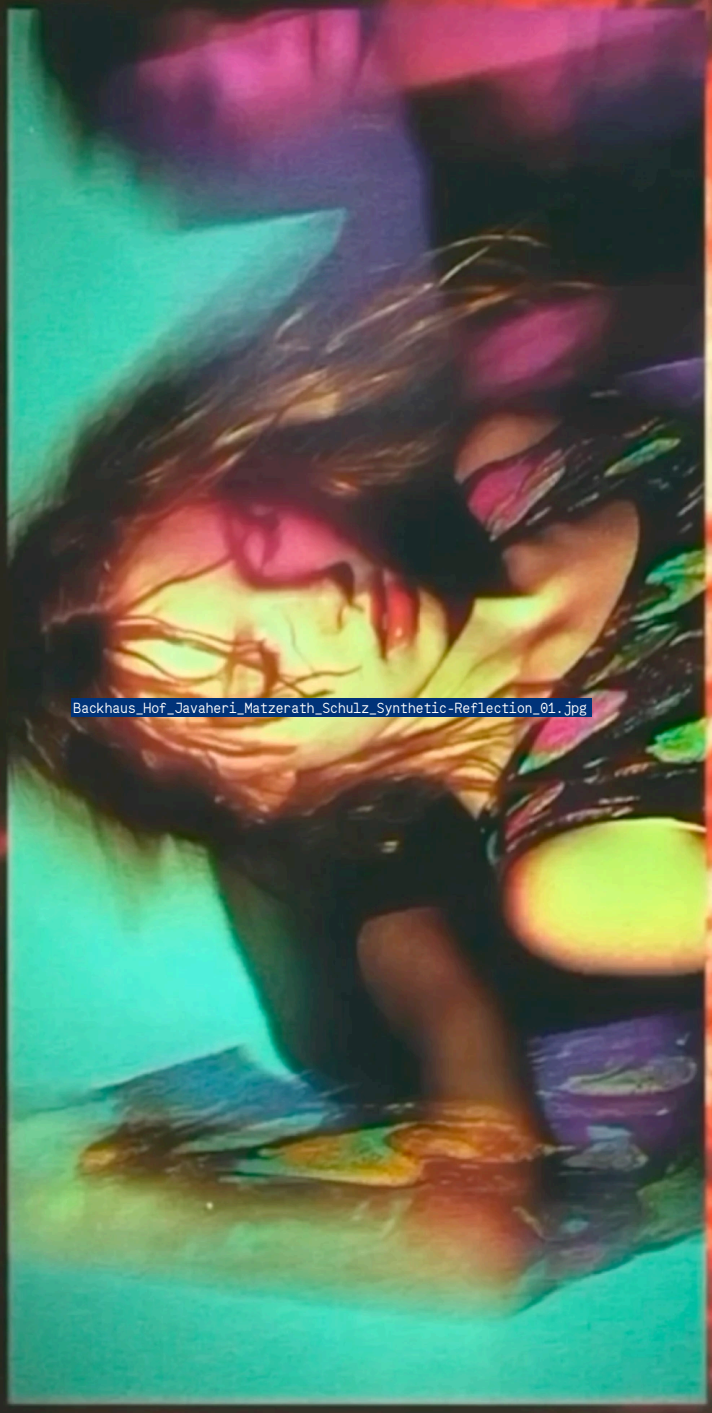
Backhaus_Hof_Javaheri_Matzerath_Schulz_Synthetic-Reflection_01.jpg

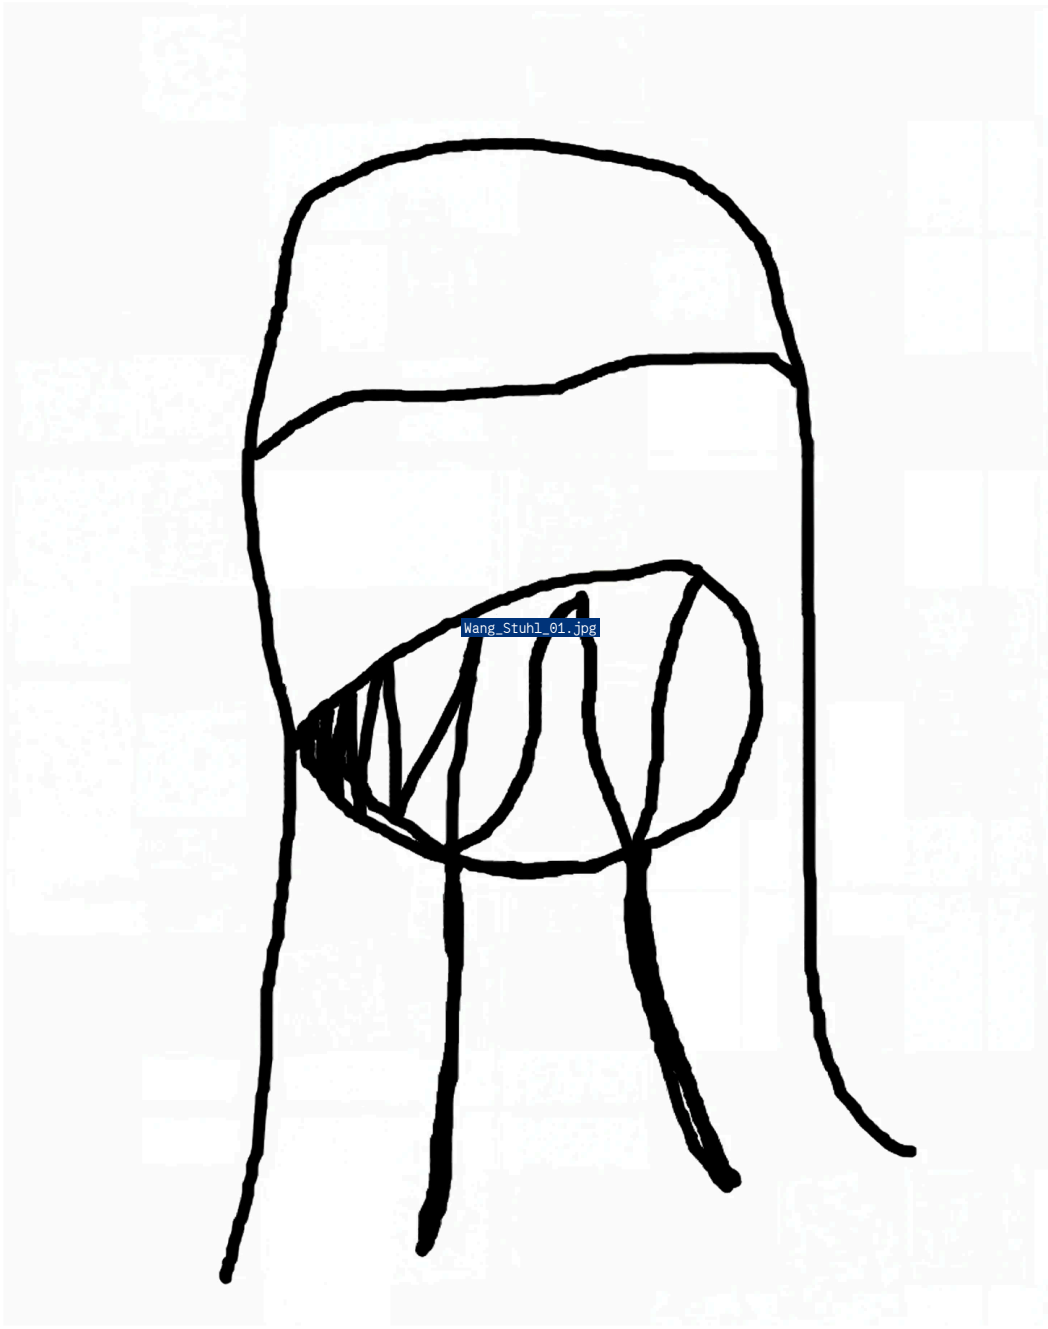
Wang_Stuhl_01.jpg

Get instant answers, find creative inspiration, learn something

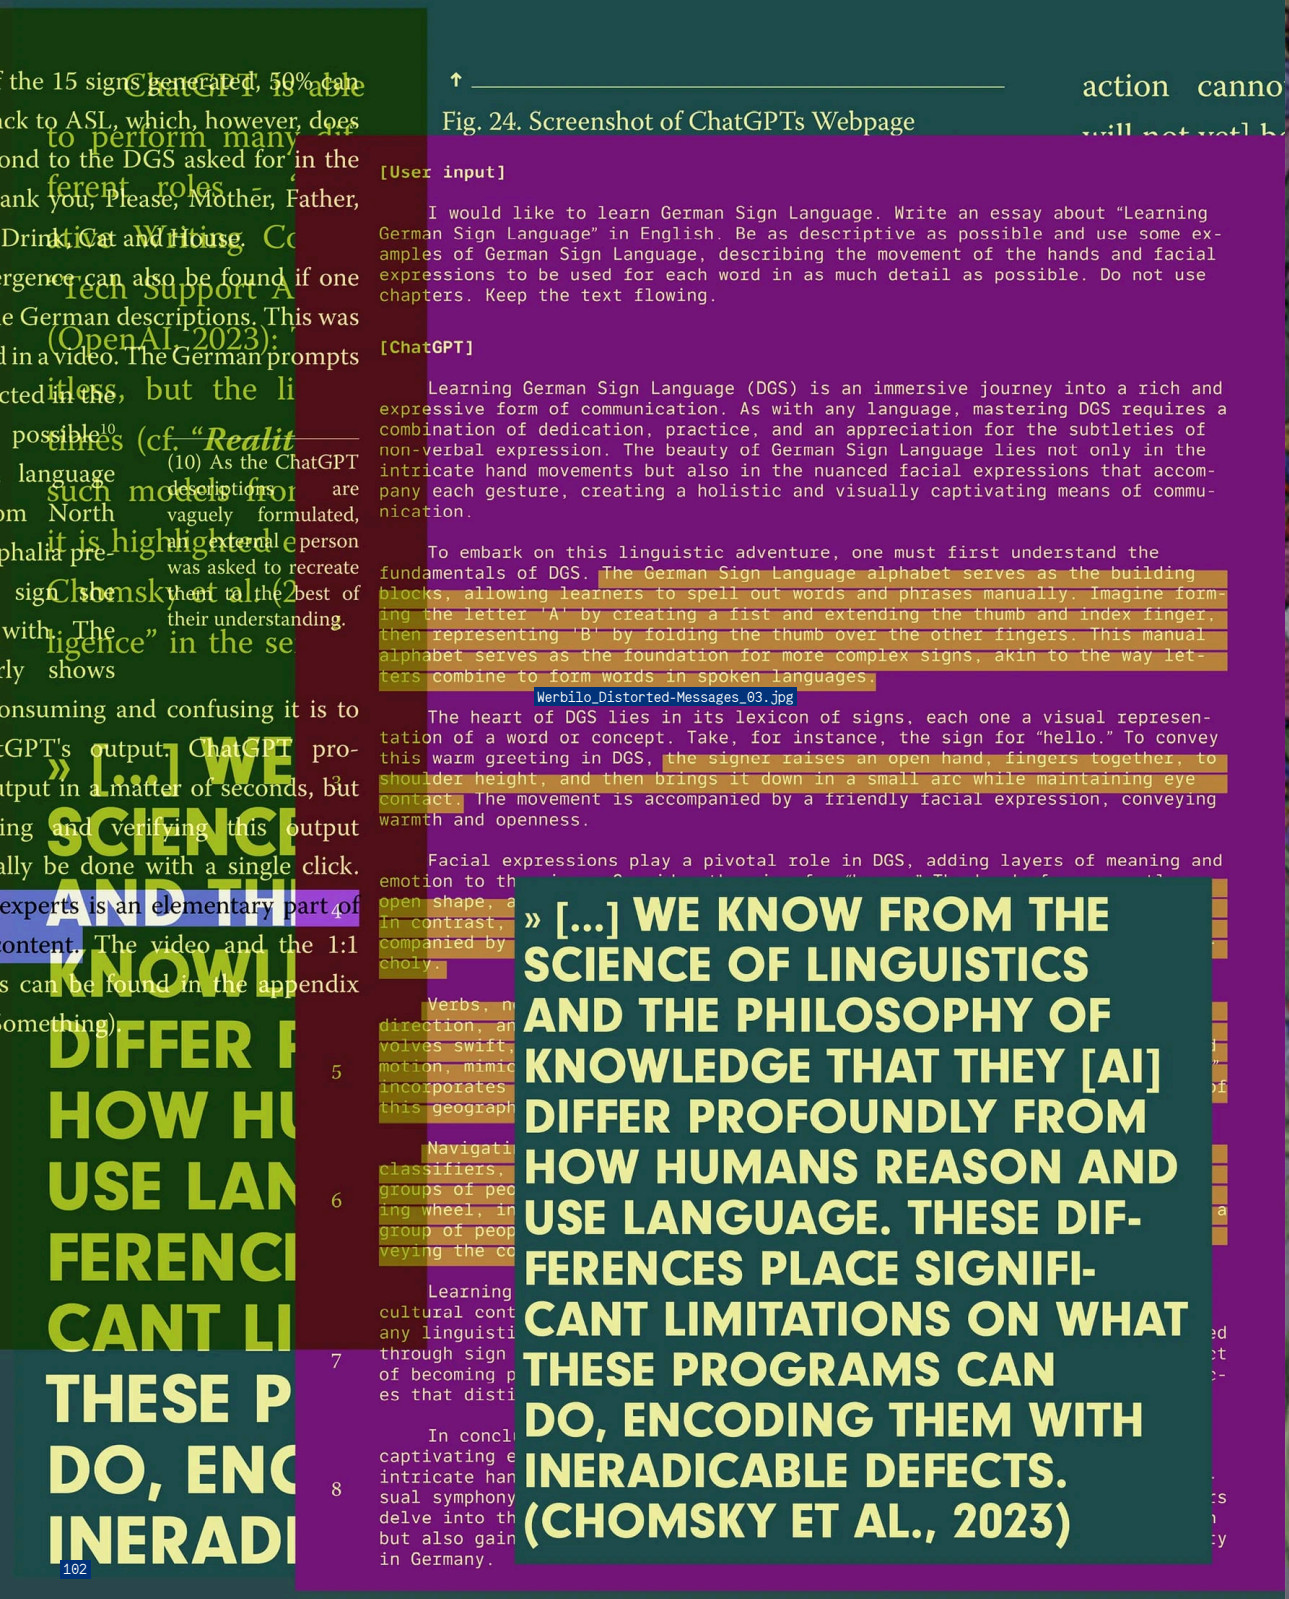

Fig. 24. Screenshot of ChatGPTs Webpage

[User input]

I would like to learn German Sign Language. Write an essay about "Learning German Sign Language" in English. Be as descriptive as possible and use some examples of German Sign Language, describing the movement of the hands and facial expressions to be used for each word in as much detail as possible. Do not use chapters. Keep the text flowing.

[ChatGPT]

Learning German Sign Language (DGS) is an immersive journey into a rich and expressive form of communication. As with any language, mastering DGS requires a combination of dedication, practice, and an appreciation for the subtleties of non-verbal expression. The beauty of German Sign Language lies not only in the intricate hand movements but also in the nuanced facial expressions that accompany each gesture, creating a holistic and visually captivating means of communication.

To embark on this linguistic adventure, one must first understand the fundamentals of DGS. The German Sign Language alphabet serves as the building blocks, allowing learners to spell out words and phrases manually. Imagine forming the letter 'A' by creating a fist and extending the thumb and index finger, then representing 'B' by folding the thumb over the other fingers. This manual alphabet serves as the foundation for more complex signs, akin to the way letters combine to form words in spoken languages.

Werbilo_Distorted-Messages_03.jpg

The heart of DGS lies in its lexicon of signs, each one a visual representation of a word or concept. Take, for instance, the sign for "hello." To convey this warm greeting in DGS, the signer raises an open hand, fingers together, to shoulder height, and then brings it down in a small arc while maintaining eye contact. The movement is accompanied by a friendly facial expression, conveying warmth and openness.

Facial expressions play a pivotal role in DGS, adding layers of meaning and emotion to th

open shape, a

In contrast,

companied by

choly.

Verbs, n

direction, an

volves swift,

motion, mimic

incorporates

this geograph

Navigati

classifiers,

groups of peo

ing wheel, in

group of peop

veying the co

Learning

cultural cont

any linguisti

through sign

of becoming p

es that disti

In concl

captivating e

intricate han

sual symphony

delve into th

but also gain

in Germany.

» [...] WE KNOW FROM THE SCIENCE OF LINGUISTICS AND THE PHILOSOPHY OF KNOWLEDGE THAT THEY [AI] DIFFER PROFOUNDLY FROM HOW HUMANS REASON AND USE LANGUAGE. THESE DIFFERENCES PLACE SIGNIFICANT LIMITATIONS ON WHAT THESE PROGRAMS CAN DO, ENCODING THEM WITH INERADICABLE DEFECTS. (CHOMSKY ET AL., 2023)

» [...] WE KNOW FROM THE SCIENCE OF LINGUISTICS AND THE PHILOSOPHY OF KNOWLEDGE THAT THEY [AI] DIFFER PROFOUNDLY FROM HOW HUMANS REASON AND USE LANGUAGE. THESE DIFFERENCES PLACE SIGNIFICANT LIMITATIONS ON WHAT THESE PROGRAMS CAN DO, ENCODING THEM WITH INERADICABLE DEFECTS. (CHOMSKY ET AL., 2023)


Stigler_Bove_Subat_Beer_FloraForge_03.jpg

Sharifi_Synthetic-Eden_01.jpg


Dornheim_Served_01.jpg



Backhaus_Hof_Javaheri_Matzerath_Schulz_Synthetic-Reflection_03.jpg



Eine Grafik zur Anatomie einer Ratte hat auf Grund der inkorrekten Darstellungen und Beschriftungen für Aufsehen gesorgt. Das Bild wurde mit Midjourney generiert und in einer wissenschaftlichen Publikation veröffentlicht. Quelle: Front. Cell Dev. Biol., Guo, Dong, Hao

Vgl. Jan Hartung u. a.: „Experts fail to reliably detect AI-generated histological data"

allerdings beim Einsatz von generativer künstlicher Intelligenz, auf Grund von datenschutzrechtlichen, ethischen, oder urheberrechtlichen Gründen. Außerdem konnten qualitativen Anforderungen, die wissenschaftliche Bilder einfordern, in Vergangenheit nicht ausreichend erfüllt werden. Mit dem Vorranschreiten der Technologie, könnte sich dies aber stetig ändern.

Mit der zunehmenden Perfektion von KI-Fakes könnte es besonders für Laien schwierig werden, authentische von manipulierten wissenschaftlichen Darstellungen zu unterscheiden, was die Glaubwürdigkeit echter Forschung nachhaltig schädigen könnte. Besonders im wissenschaftlichen Kontext könnte es Menschen schwer fallen Desinformation und Bildmanipulation zu erkennen, da Inhalte ohnehin oft abstrakt oder kompliziert sind. Und umso besser die KI-Fälschungen echten wissenschaftlichen Bildern ähneln, desto mehr Menschen könnten in Zukunft mit wissenschaftlichen Falschinformationen, oder Desinformationen erreicht werden.

Kräuter_Inszenierte-Wissenschaft_04.png



Quelle: Midjourney

2024 WURDEN JEDEN TAG 34 MIO. BILDER GENERIERT. CHATGPT-4o GENERIERTE IN WOCHE EINS RUND 700 MIO. BILDER.



Quelle: Firefly



Quelle: Sora (ChatGPT 4o)

Vgl. Everypixel Journal, People Are Creating an Average of 34 Million Images Per Day. Statistics for 2024
Vgl. Brad Lightcap on X


Deutloff_Digitale-Entfremdung_02.jpg
Herzlich Willkommen!
Bitte wähle ein Gedicht aus und sprich es deutlich in das Mikrofon, indem du einmal kurz auf den „rec“-Knopf vor dir drückst.
Wenn du fertig bist, drücke ihn noch einmal, um die Aufnahme zu beenden.
Welcome!
Please select a poem and speak it clearly into the microphone by briefly pressing the "rec" button in front of you once.
When you have finished, press it again to stop the recording.
THE FEAR OF MISSING OUT
REC

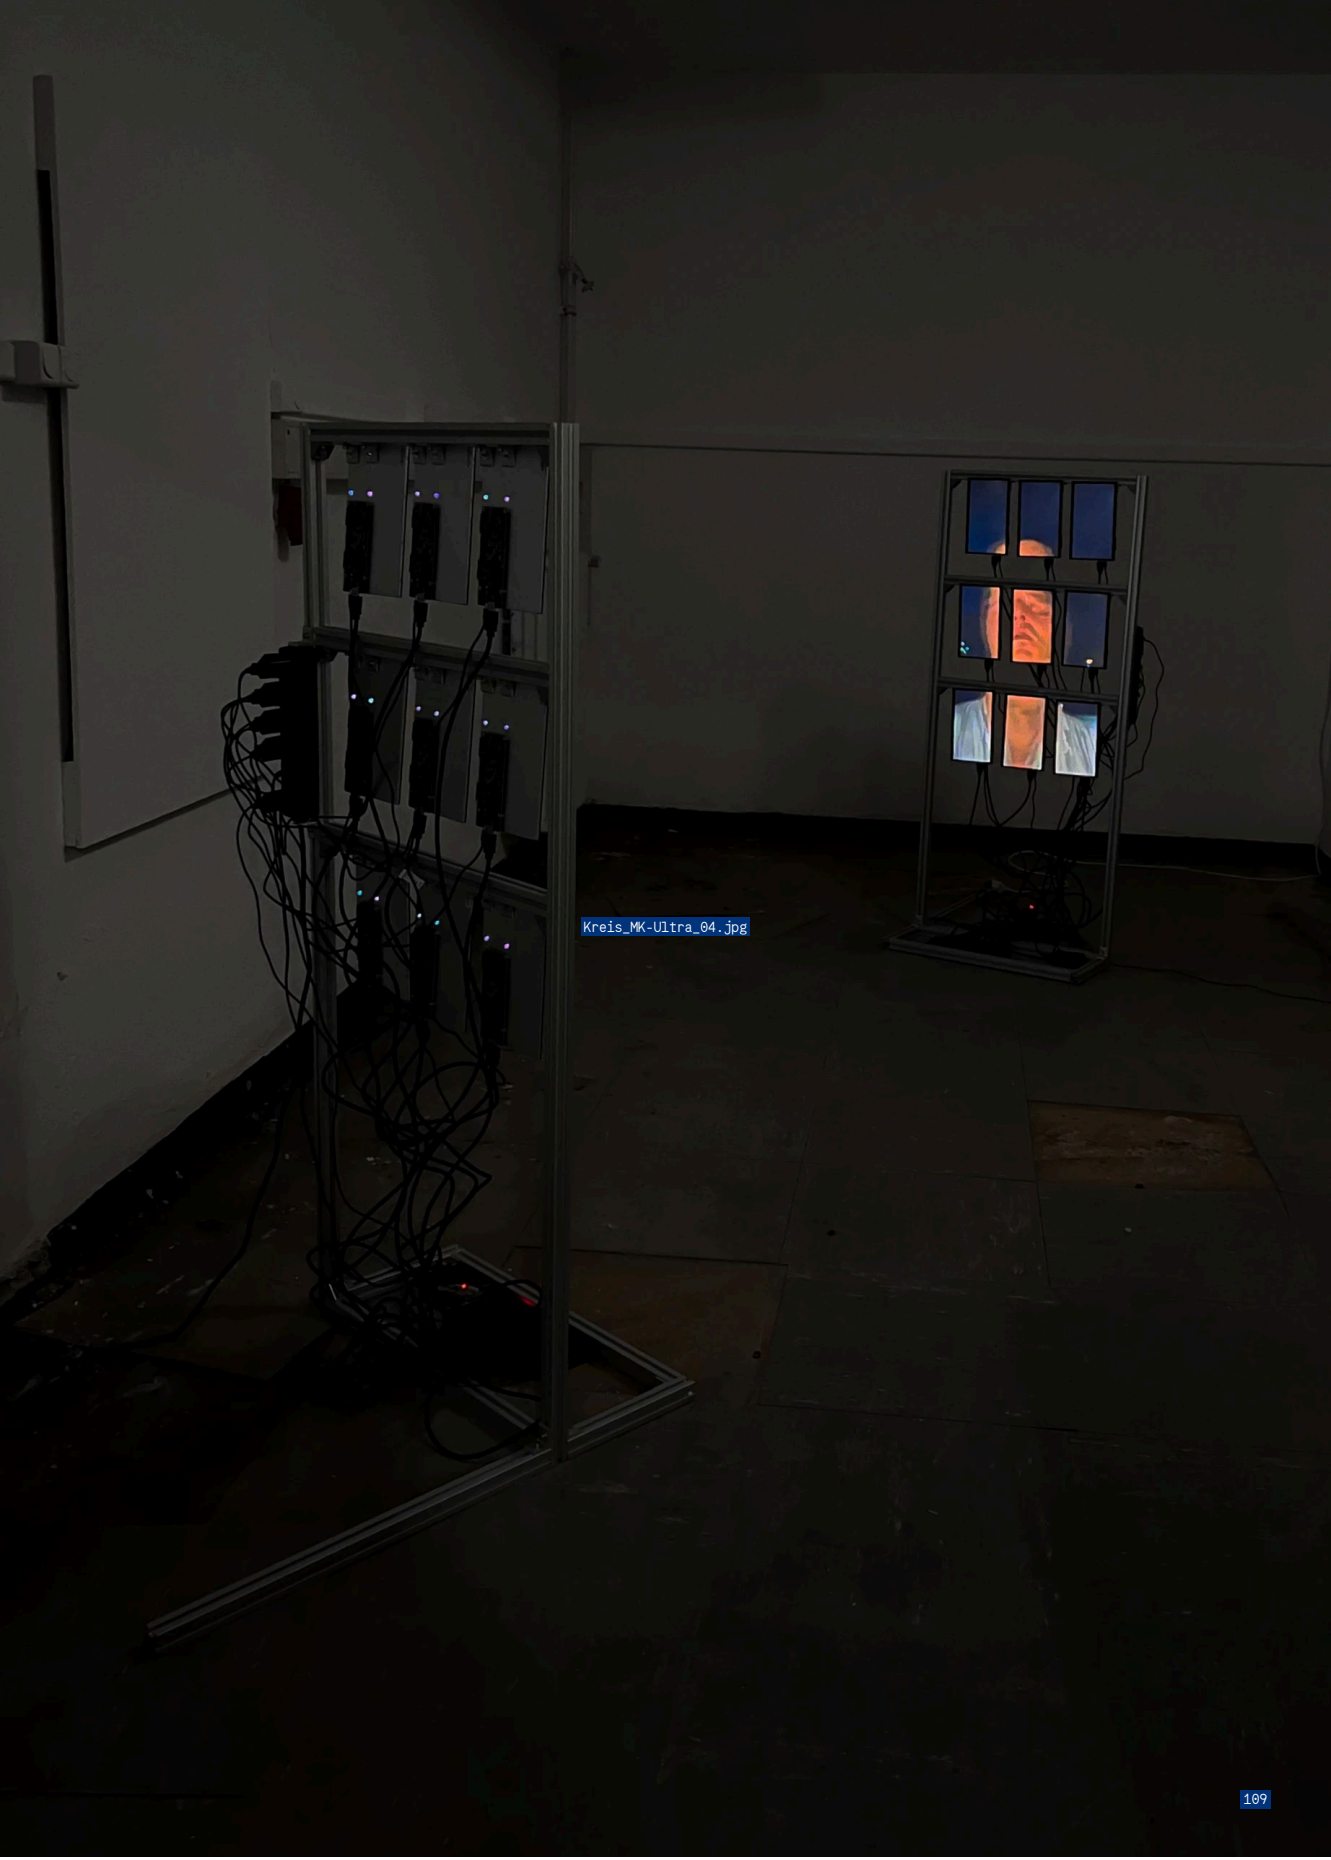
Kreis_MK-Ultra_04.jpg

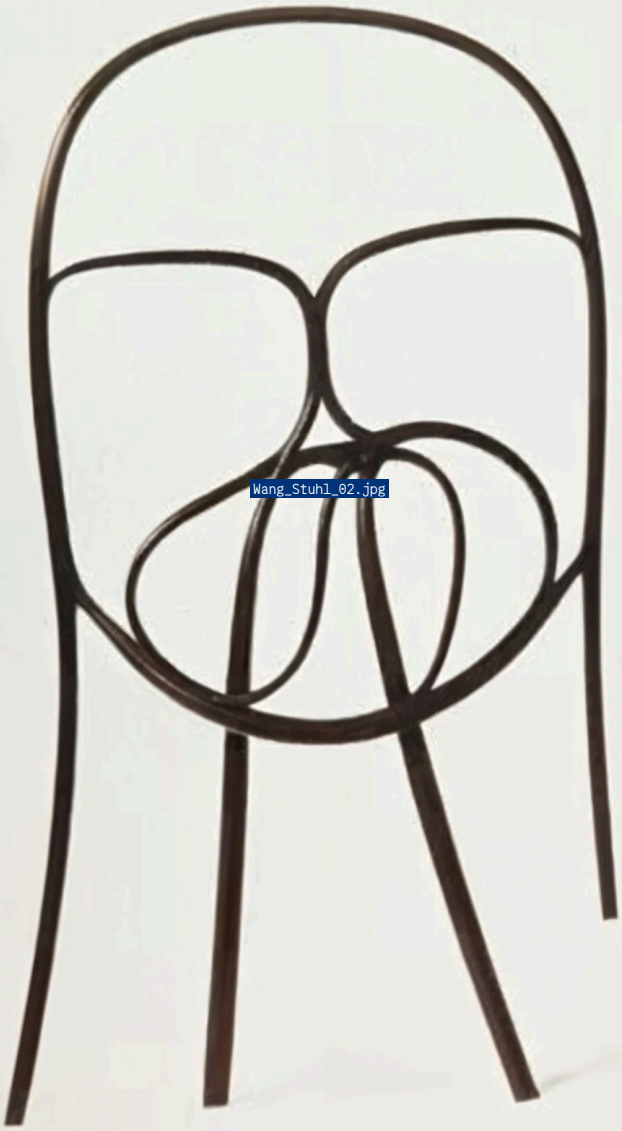
Wang_Stuhl_02.jpg

STABLEHUB

START

Home

GENERATION TOOLS

Text To Image

Image To Image

ControlNet

COMMUNITY

Gallery

LEARN & EXPLORE

Guided Tutorials

USE FROM HOME

VPN Install

Milkshake ▾

Community Gallery

Discover favorite images from all users. Click to copy prompts and download images you like!

REFRESH CLEAN

















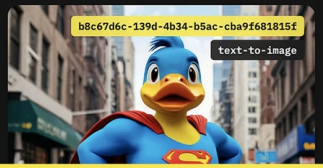

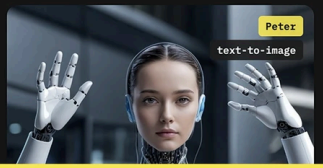

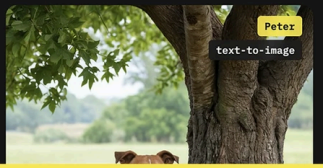

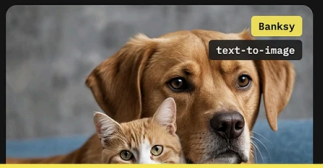

Annaeup1_StableHub_03.png

Scholl_Wahrgenommen_03.jpg

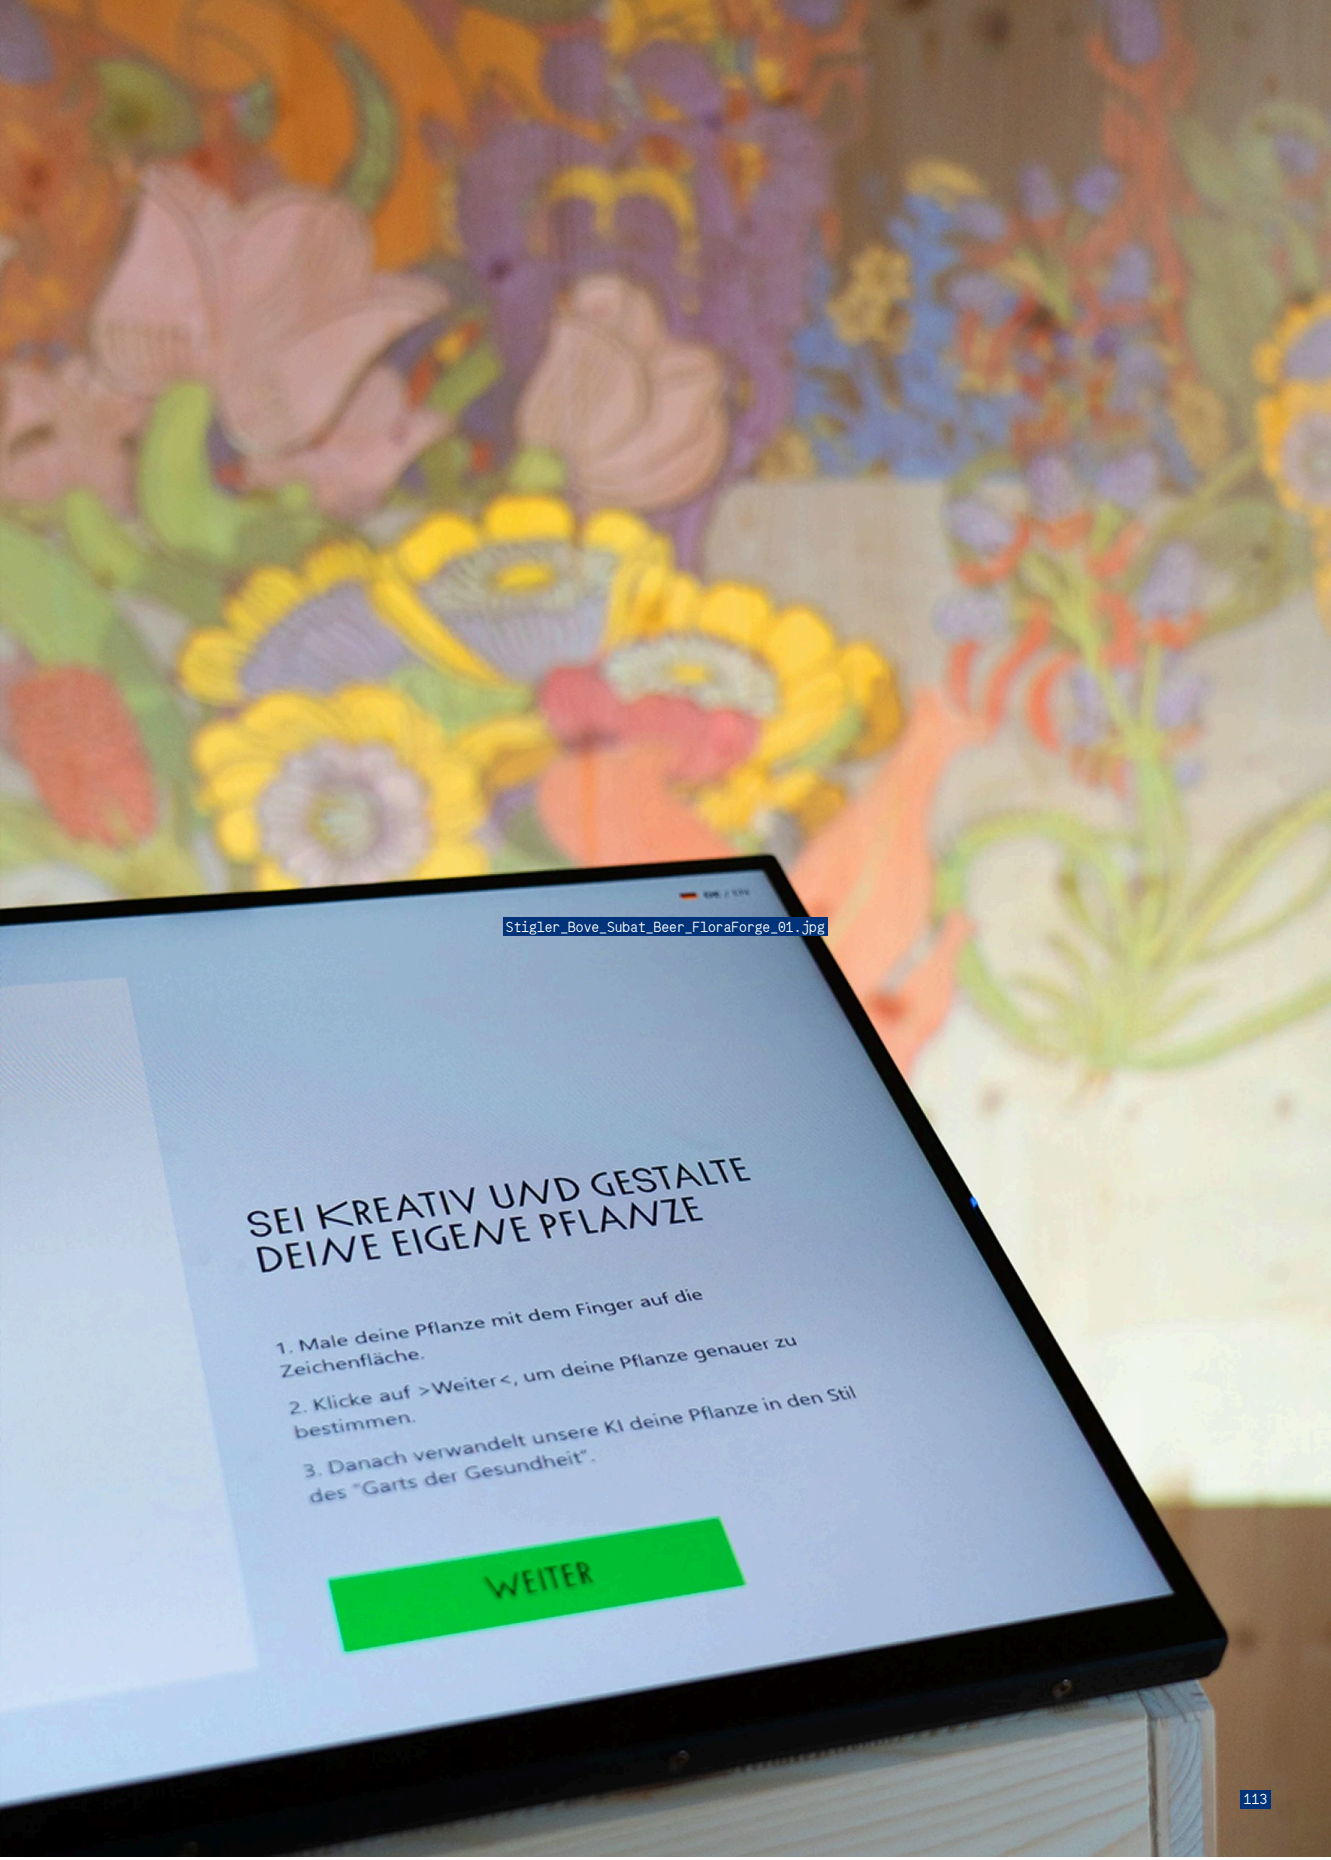
Stigler_Bove_Subat_Beer_FloraForge_01.jpg
SEI KREATIV UND GESTALTE DEINE EIGENE PFLANZE
1. Male deine Pflanze mit dem Finger auf die Zeichenfläche.
2. Klicke auf >Weiter<, um deine Pflanze genauer zu bestimmen.
3. Danach verwandelt unsere KI deine Pflanze in den Stil des "Garts der Gesundheit".
WEITER


Wang_Stuhl_05.jpg

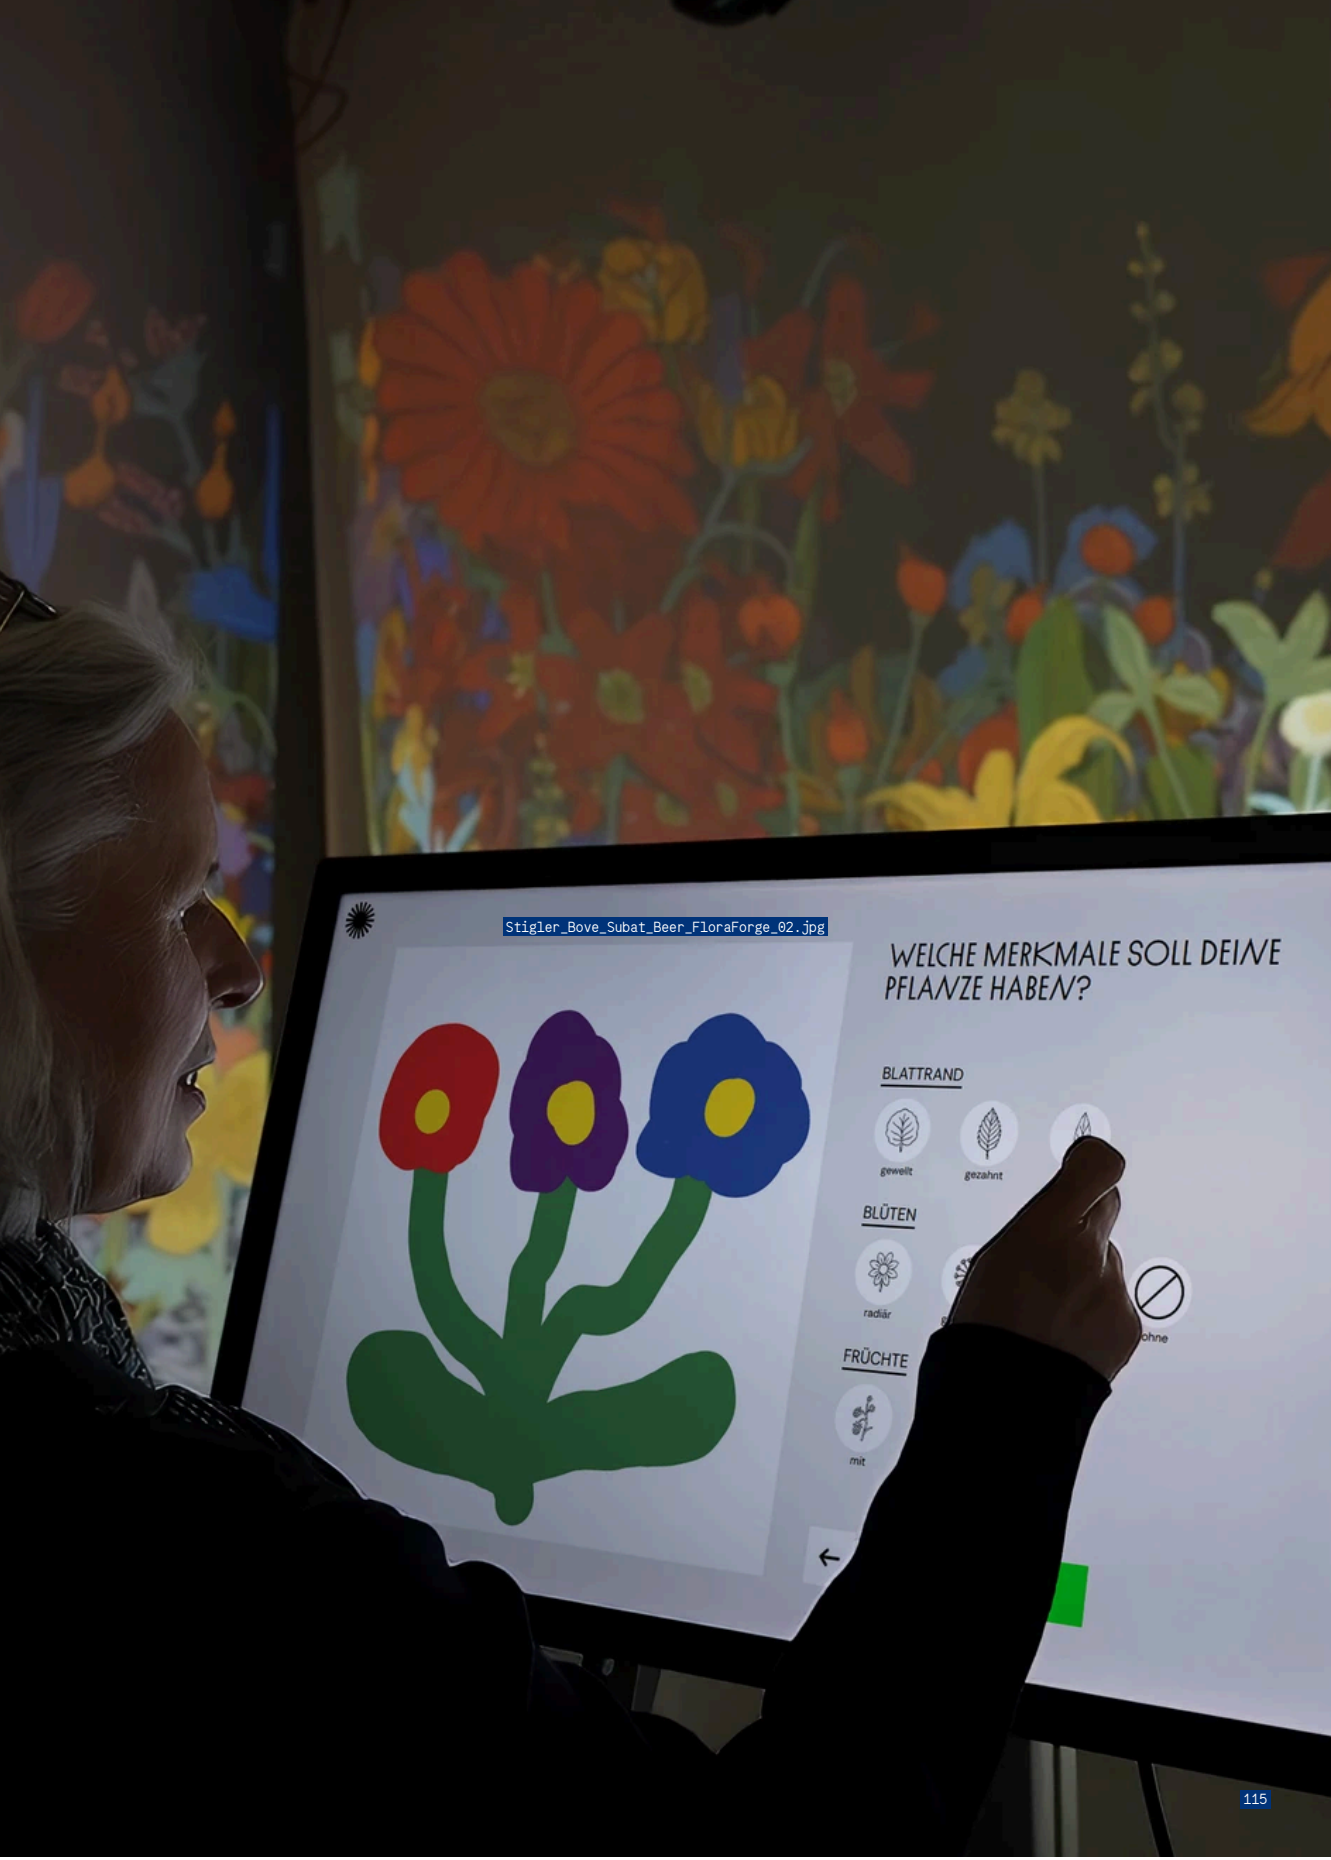
Stigler_Bove_Subat_Beer_FloraForge_02.jpg
WELCHE MERKMALE SOLL DEINE PFLANZE HABEN?
BLATTRAND
gewellt
gezahnt
BLÜTEN
radiär
ohne
FRÜCHTE
mit

Sharifi_Synthetic-Eden_04.jpg



Seite 54

FLACHE ERDE

A casual, awkward, unremarkable iPhone selfie of a rugged middle-aged man (approx. 58 years old) in full polar expedition gear, taken during a freezing Arctic or Antarctic expedition. The photo is shot from a very low upward camera angle with a strong wide-angle lens effect, making his chin and neck appear exaggerated and giving him a double chin. He is slightly off-center in the frame, leaving space on one side to show the vast background. He wears a yellow insulated expedition parka with frost and snow buildup, a white scarf or balaclava, and black backpack straps. His wind-blown, graying hair is messy, and his face shows signs of cold exposure: red cheeks rough skin. His expression is serious and tired.

Behind him is an enormous, towering glacier wall similar in scale to 'The Wall' from Game of Thrones Because of the low angle, both the immense wall and a pale, overcast sky are visible. The lighting is flat and diffused, the photo is slightly underexposed, grainy, with visible snow glare. It should look like a clumsy, raw moment captured on an iPhone during a grueling expedition.

Kramer_Inszenierte-Wissenschaft_05.png



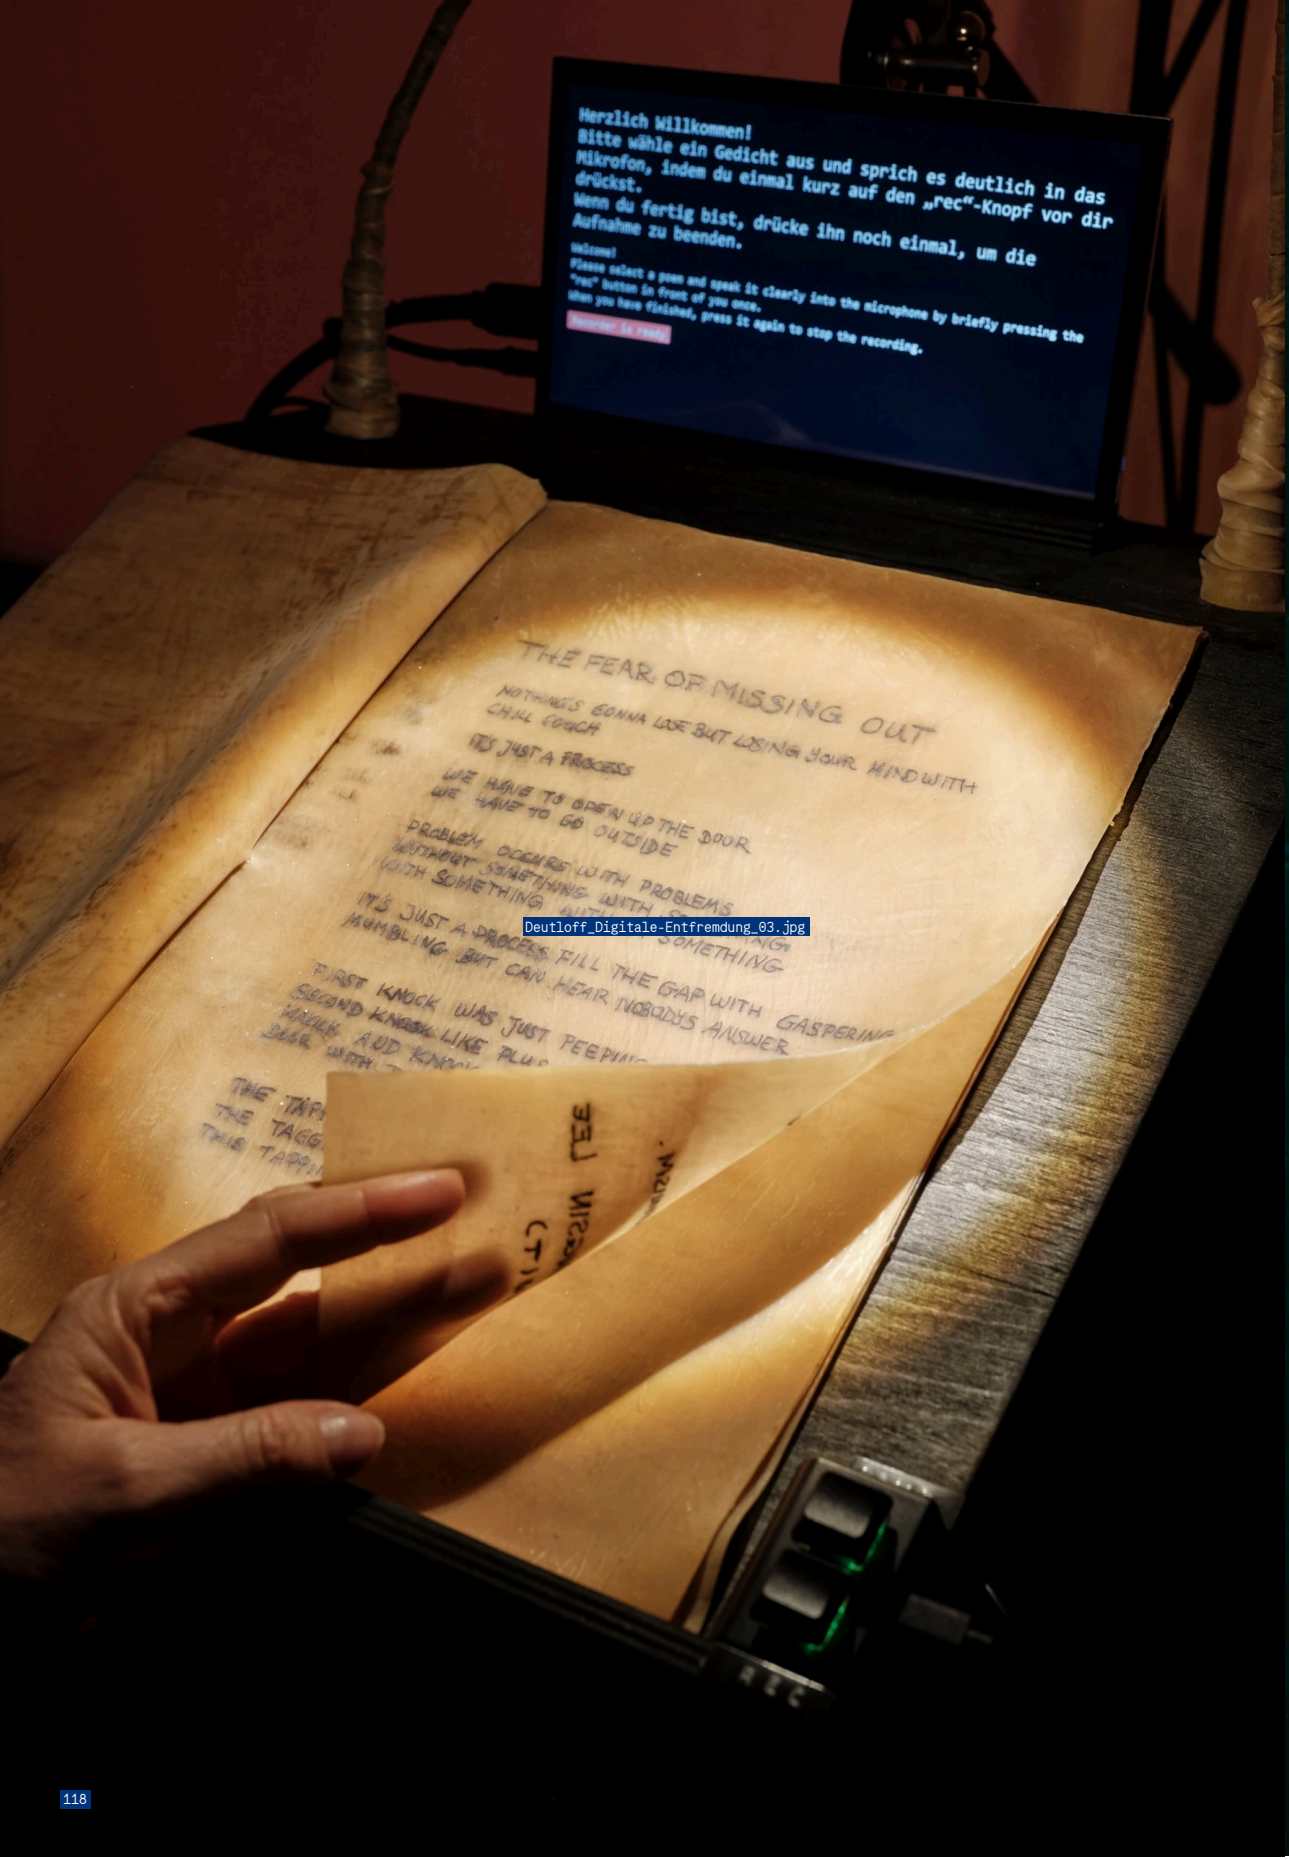

Deutloff_Digitale-Entfremdung_03.jpg

KI-GENERIERTE BILDER



KI-generiertes Bild von Papst Franziskus in Pufferjacke, welches viral gegangen ist. Quelle: X

NEUE METHODEN DER BILDERZEUGUNG

Die Möglichkeit Bilder zu manipulieren ist nicht neu, aber mit künstlicher Intelligenz ist sie schneller und kaum noch nachzuvollziehen. Außerdem ist die Schwelle zur Bildproduktion extrem gesunken: Generative Bildmodelle wie Midjourney oder Stable Diffusion sind frei verfügbar, meist kostengünstig und ihre Bedienung erfordert keine gestalterische Ausbildung mehr. Damit lassen sich sowohl bestehende Bilder bearbeiten, als auch vollständig neue visuelle Inhalte erzeugen.

Zu Beginn wurde generative künstliche Intelligenz hauptsächlich für künstlerische Zwecke eingesetzt, etwa um surreale Szenen darzustellen, oder bestehende Kunst zu imitieren. Mittlerweile werden sie aber nicht nur als künstlerisch-kreatives Medium eingesetzt, sondern auch um realistische Szenen, sprich verschiedenste Arten von Fotografien zu imitieren. Solch realistischen KI-generierten Bilder orientieren sich häufig an vertrauten visuellen Konventionen, was sie besonders glaubwürdig und real erscheinen lässt. Dadurch verschwimmen die Grenzen zwischen der klassischen Fotografie und synthetisch erzeugtem Bildmaterial zunehmend.

Viele Studien zeigen, dass Menschen schon heute schwer unterscheiden können, ob ein Bild echt oder KI-generiert ist. Besonders, wenn es sich um Porträts von Gesichtern handelt. Des Weiteren entwickeln sich KI-Modelle schnell weiter und lernen immer mehr Stile und Elemente dazu. Damit steigt nicht nur die Qualität der Bilder, sondern auch die Vielfalt der Motive, was die Möglichkeiten der Bildgeneration weiter vorantreibt.

Vgl. CORDIS Europa: „Welches Gesicht ist echt?"

119

Krämer_Inszenierte-Wissenschaft_03.png



KI-generiertes Foto von Donald Trump, umzingelt von mehreren Polizisten aus dem Jahr 2023. Quelle: CORRECTIV. Faktencheck

Ferdinand Gehringer, Christopher Nehring, Mateusz Labuz.: Der Einfluss von Deep Fakes auf Wahlen

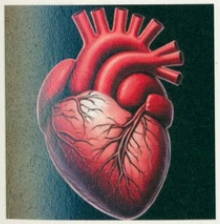

Ergebnis mit Midjourney 5.2 (2024) auf den Prompt: „Erstelle ein Bild, das die Pathophysiologie der koronaren Herzkrankheit darstellt." Die KI erzeugte eine hochauflösende Darstellung, ging jedoch nicht auf die Pathologie ein und enthielt anatomische Fehler. Quelle: Nguyen et al.

VERZERRTE WIRKLICHKEITEN

Besonders Social Media wird mittlerweile von KI generierten Bildern regelrecht geflutet, aber auch andere Bereiche wie Werbung, Journalismus, Gesellschaft und Politik sind von KI-generierten Bildern betroffen. Sogenannte Deepfakes (künstlich generierte Nachahmungen oder Manipulation von Personen) haben im Zusammenhang mit dem US-Wahlkampf große mediale Aufmerksamkeit erhalten. Die Verunsicherung und das Misstrauen, die Deepfakes in der Bevölkerung auslösen können, haben das Potenzial, politische Prozesse und die öffentliche Kommunikation zu beeinflussen.

Selbst bei transparenter Nutzung kann KI verzerrte Ergebnisse liefern, basierend auf statistischen Wahrscheinlichkeiten, die fehlerhaft oder voreingenommen sein können. Das können zum Beispiel Hände mit sechs Fingern sein, oder verschiedenste Stereotypen. Solche Verzerrungen schleichen sich unbemerkt in Bilder ein und erwecken trotzdem den Eindruck von Objektivität. Besonders kritisch wird es dort, wo Bilder eine starke Wirkung haben – zum Beispiel in der Politik, oder im wissenschaftlichen Kontext.

KI-GENERIERTE BILDER IN DER WISSENSCHAFT

In der Wissenschaftskommunikation könnte generative künstliche Intelligenz Arbeitsprozesse unterstützen, technische und gestalterische Hürden reduzieren und die Arbeit effizienter gestalten.

Die Wissenschaftskommunikation selbst zögert

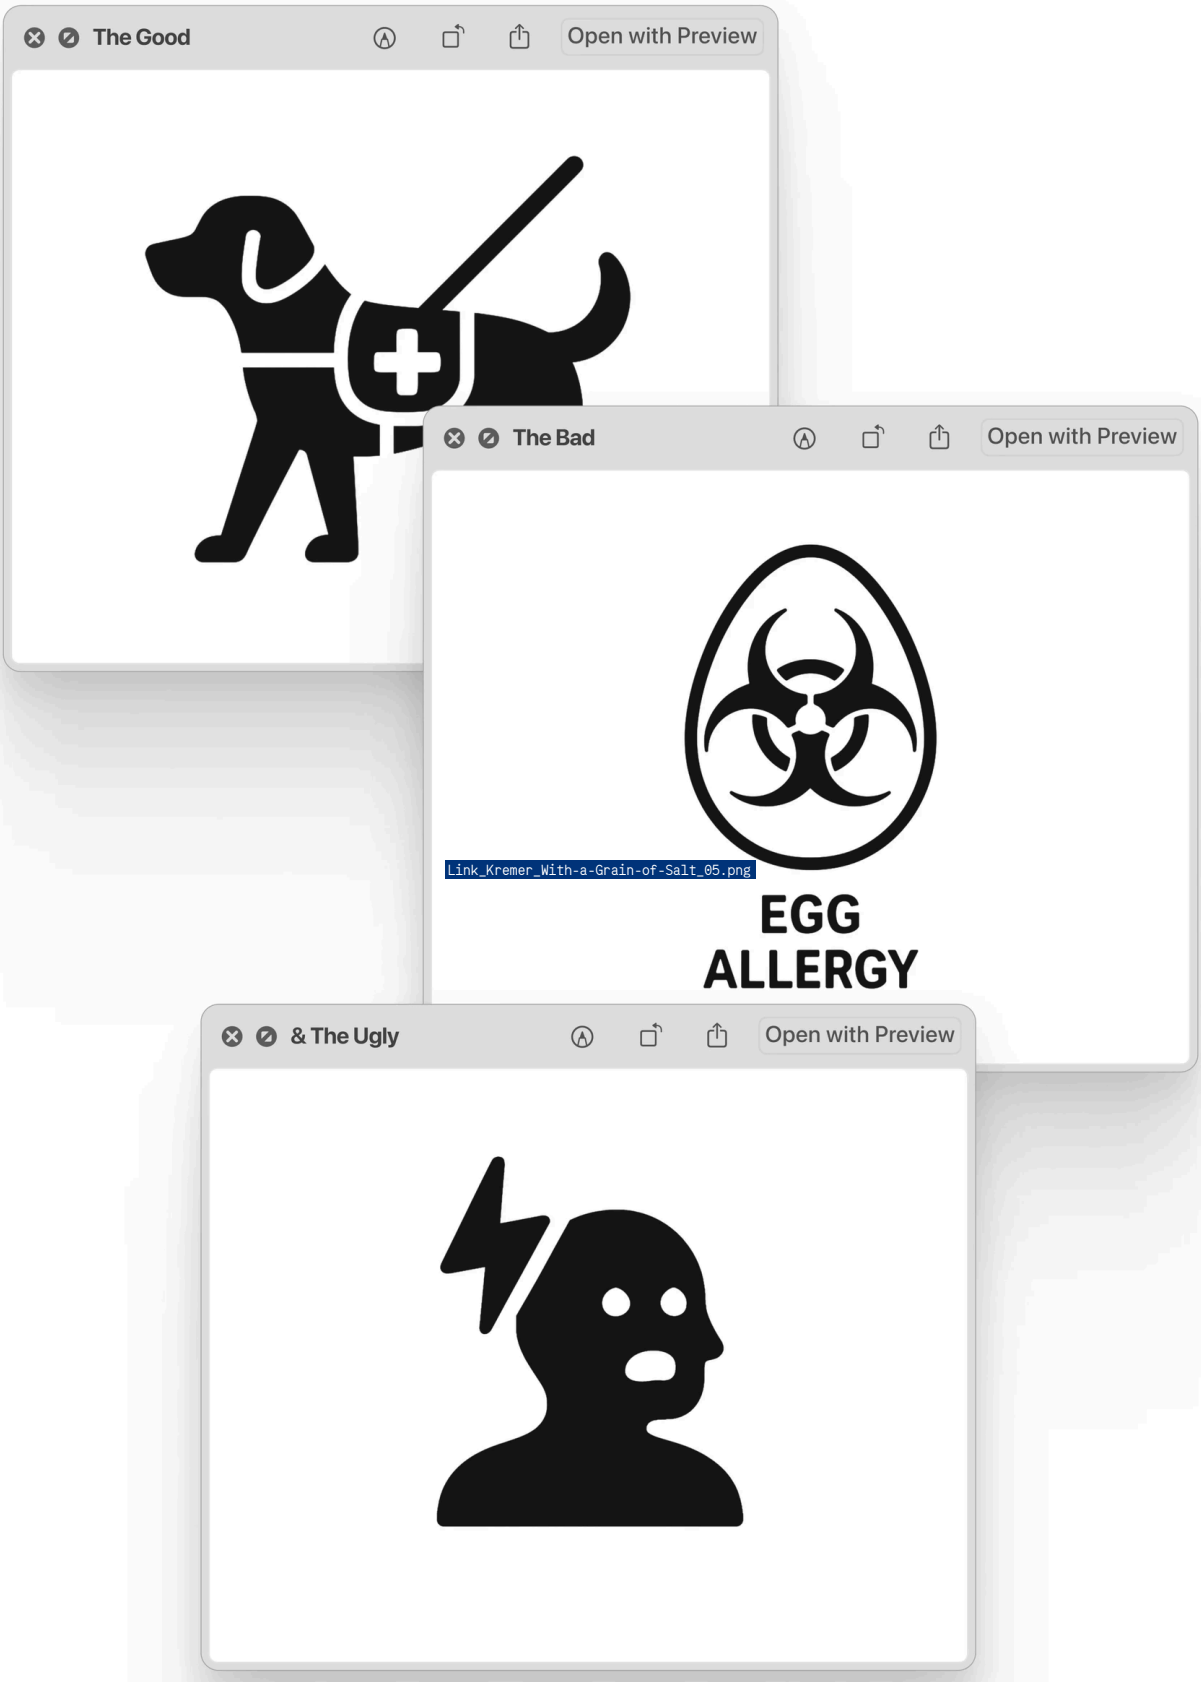
The Good
Open with Preview
The Bad
Open with Preview
Link_Kremer_With-a-Grain-of-Salt_05.png
EGG
ALLERGY
& The Ugly
Open with Preview

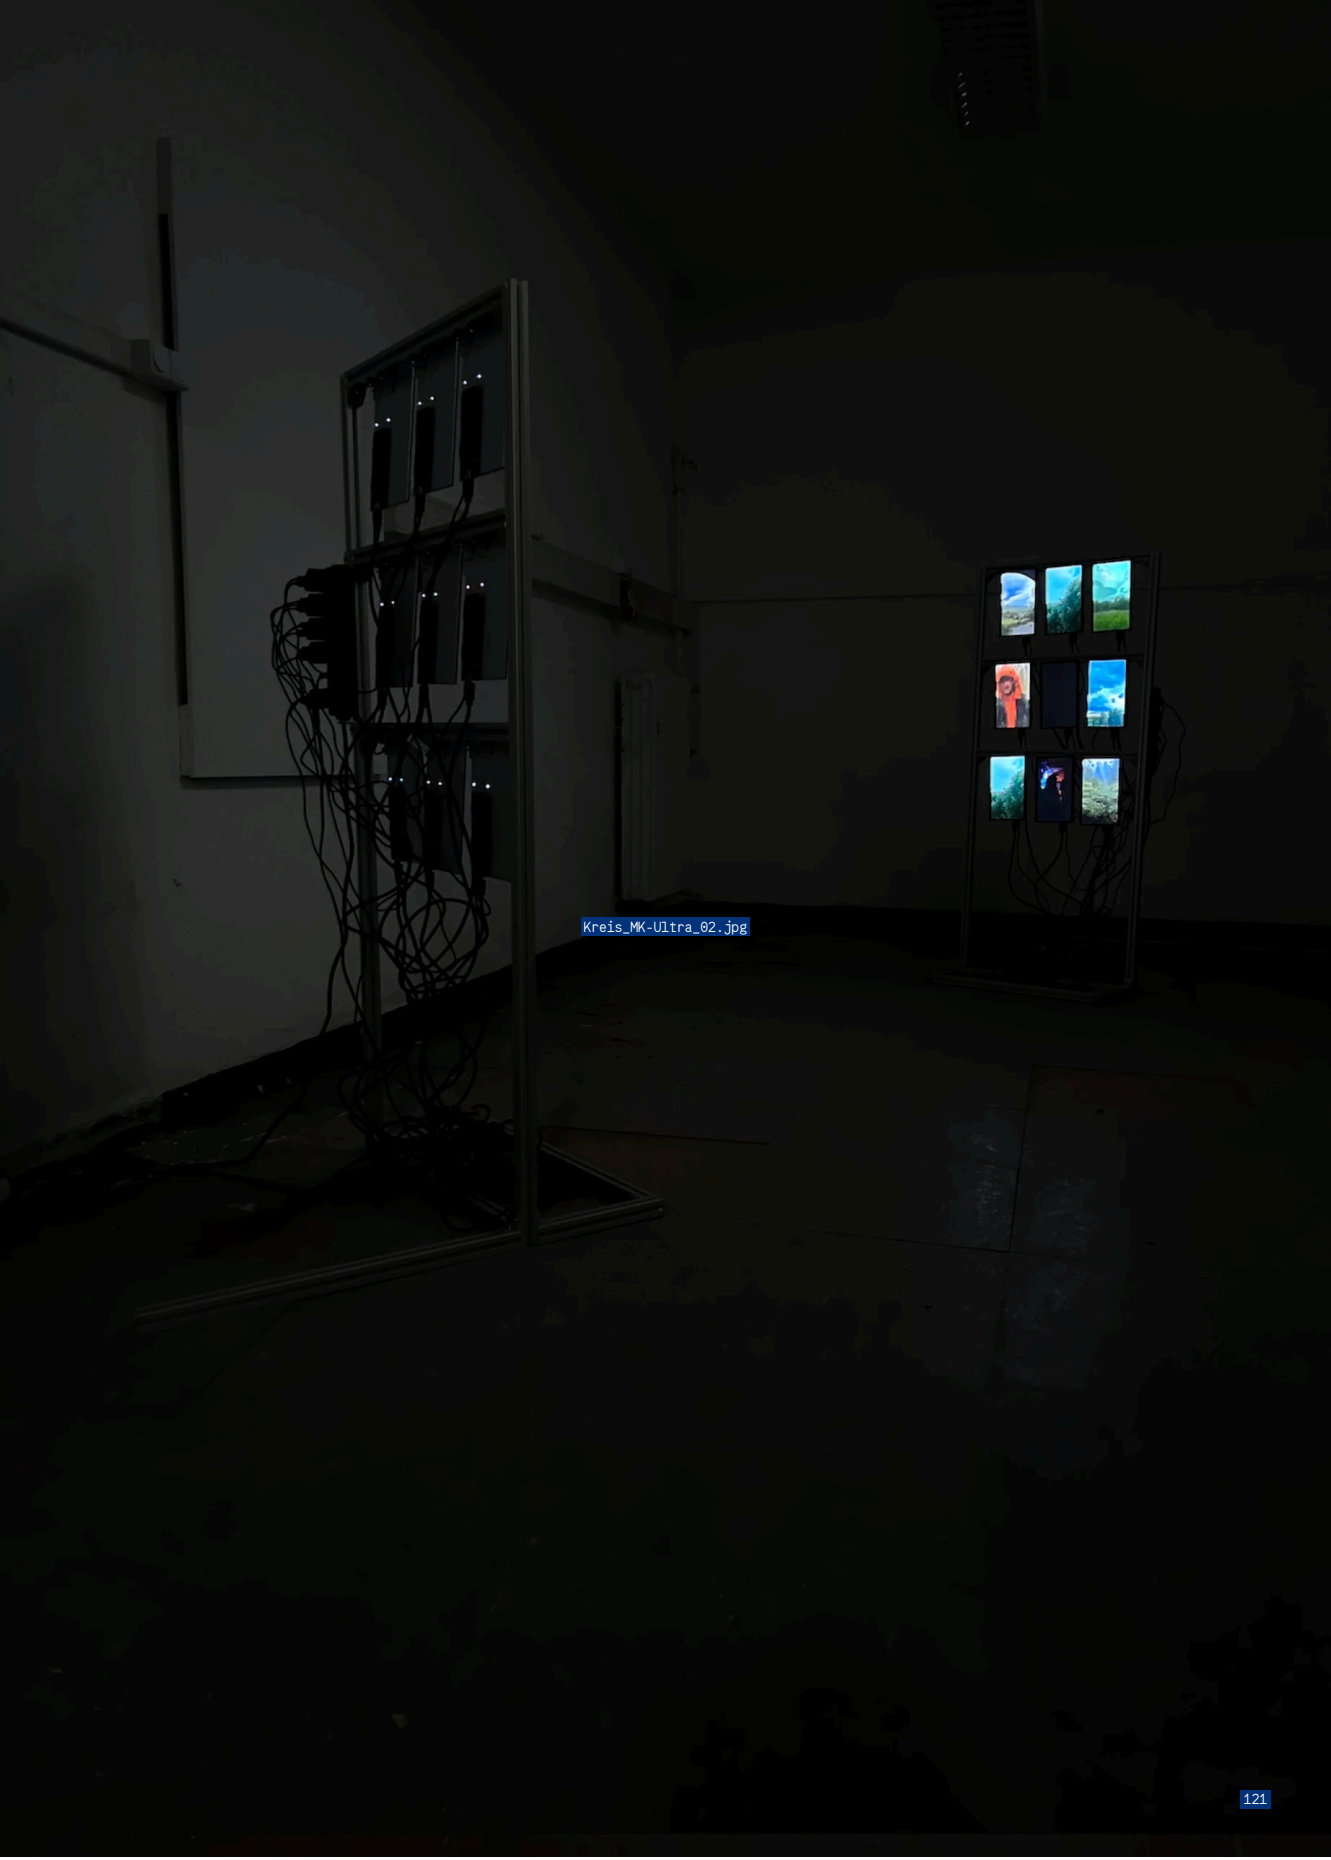

Kreis_MK-Ultra_02.jpg


Kreis_MK-Ultra_01.jpg


Gomez-Baumert_Schlag-auf-Schlag_03.jpg

DISTORTED MESSAGES



Design Research about Generative AI-Induced Hallucination in German Sign Language

Master Thesis
Viktor Werbilo
M.A. Integrated Design

from prompts or
(LLM) which are
text modules are
litical as well as
ortant multipliers
But even with the
n societal use and
ned to it, one thing
s a tool. And that
wed. Therefore, it
the flaws to start
With the emerge
as well as factually
d by generative AI
s *AI-Induced Hal-*
ethical challenges

hesis will be ex-
inations - in the
occur more fre-
t has little or no
rticular subject.
s – especially in
AI – have a nega-
anding, trust and
cially considering
ed groups. While
are at first glance
ling training data
a are present. AI
mans provide the
n enigmatic space
e translated into
rative AI - is hu-
limited in 124 wl-
e human's data.

eas and their social impact will be addressed and the origin of these will be analyzed.

In addition, this thesis will look into the representation of the d/Deaf community and German Sign Language (DGS) in digital media, followed by the challenges and opportunities of generative AI. The current state of representation with state of the art, publicly accessible generative AIs is summarized in

things'. This work addresses the topic in a critical manner. The work explores the world of 'text generation'.

(3) 'Un-clear' covers the still new but all the more important topic of 'video generation'. The d/Deaf community relies on being able to reliably recognize the movements and facial expressions

Achba_Adrian_Gatzke_Kattwinkel_Schomann_Schröder_Sideri_Zisa_Anthrobot_03.jpg

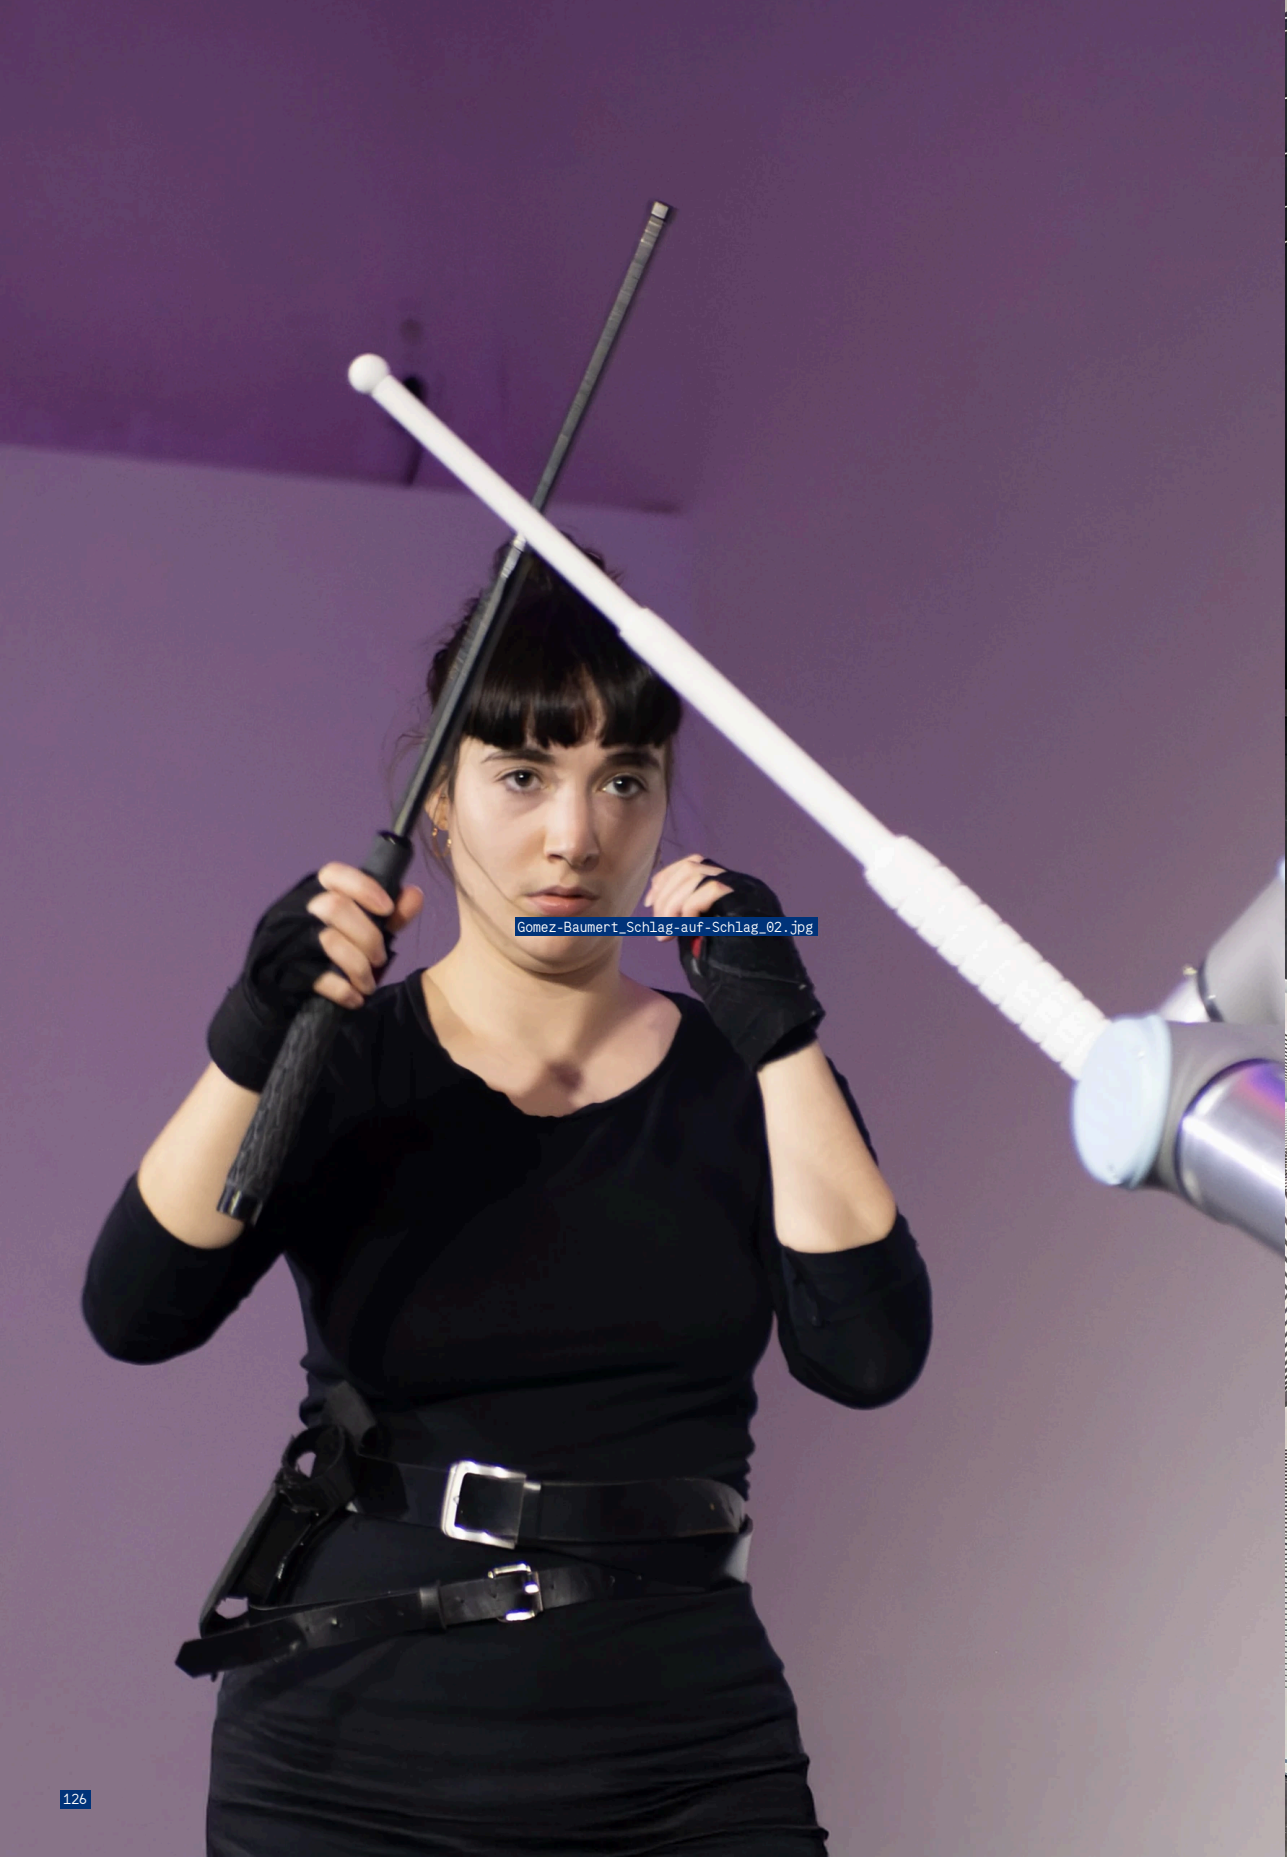
Gomez-Baumert_Schlag-auf-Schlag_02.jpg


Achba_Adrian_Gatzke_Kattwinkel_Schomann_Schröder_Sideri_Zisa_Anthrobot_02.jpg

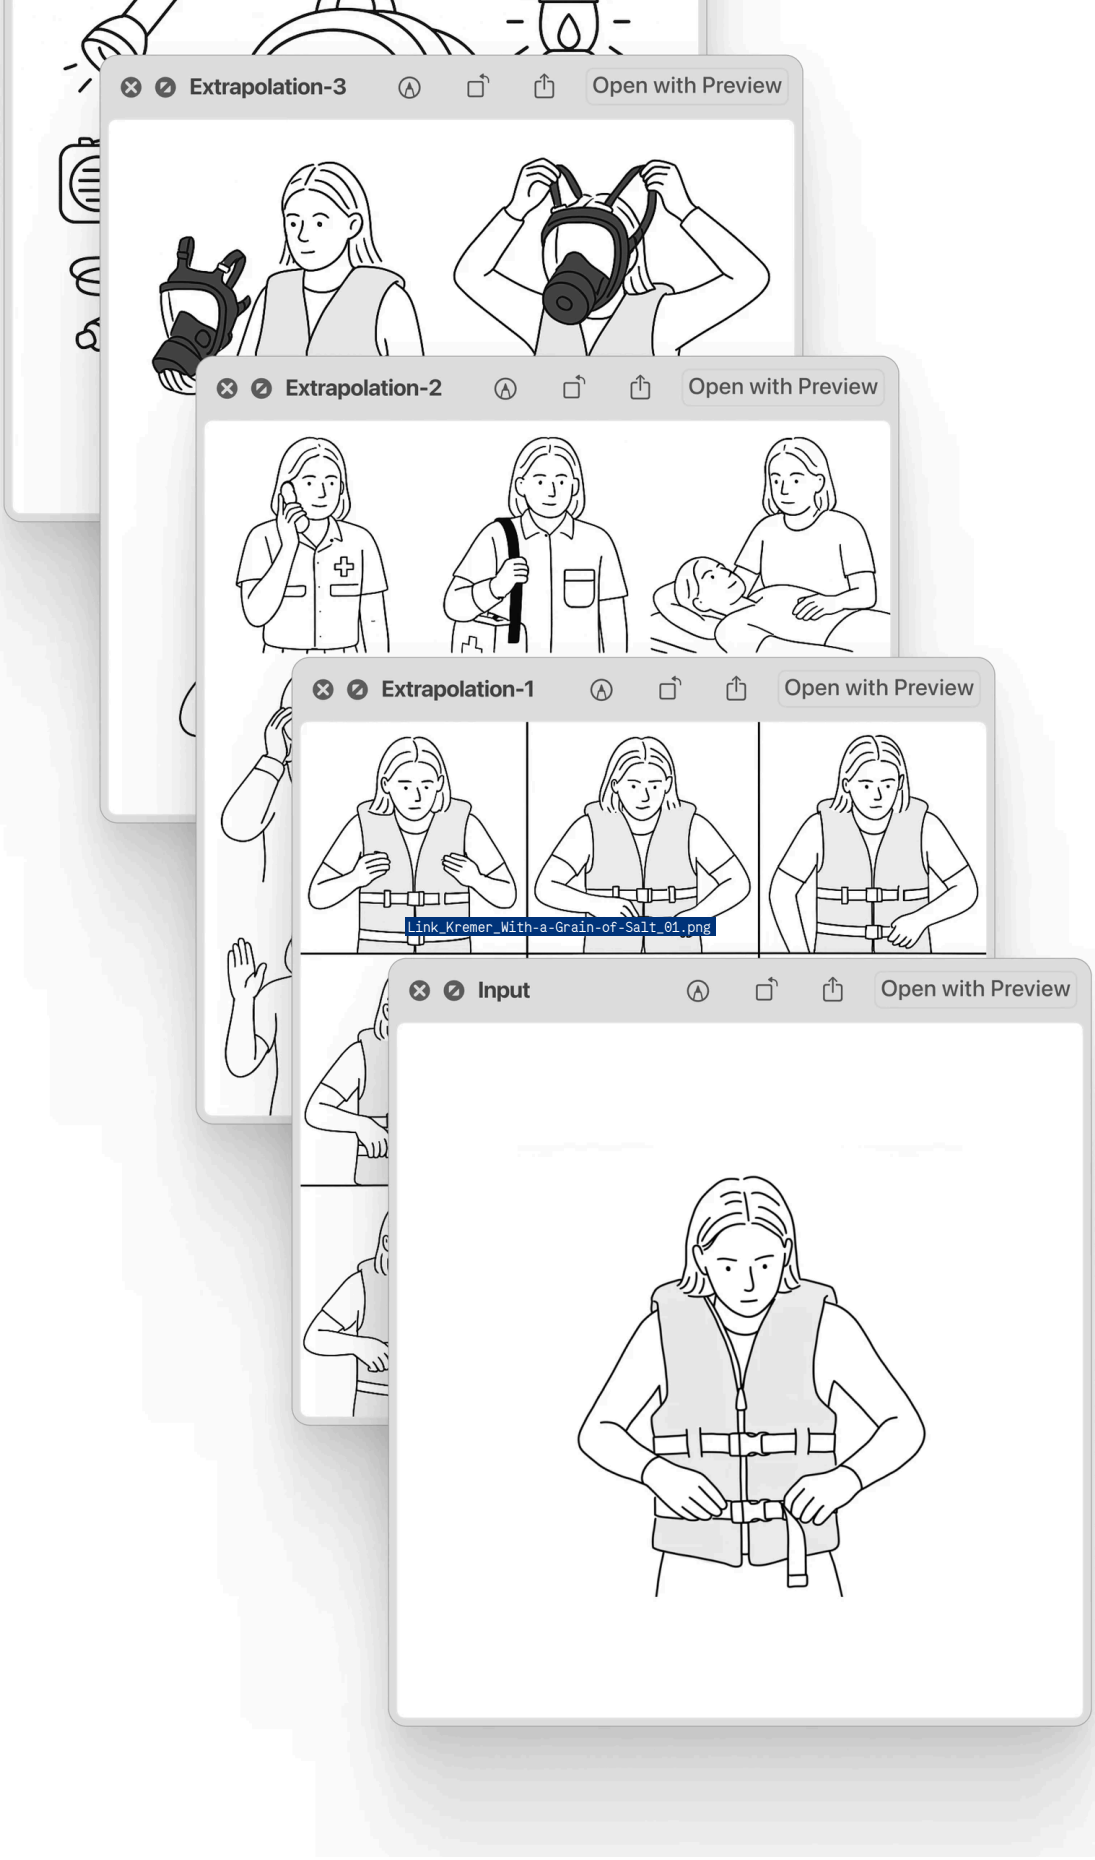
Extrapolation-3
Open with Preview
Extrapolation-2
Open with Preview
Extrapolation-1
Open with Preview
Link_Kremer_With-a-Grain-of-Salt_01.png
Input
Open with Preview


Dornheim_Served_02.jpg


Oelsen_Synthetische-Erinnerung_01.png


Leuoth_Flush-of-Flesh_01.jpg

Link_Kremer_With-a-Grain-of-Salt_04.png



Gomez-Baumert_Schlag-auf-Schlag_04.jpg

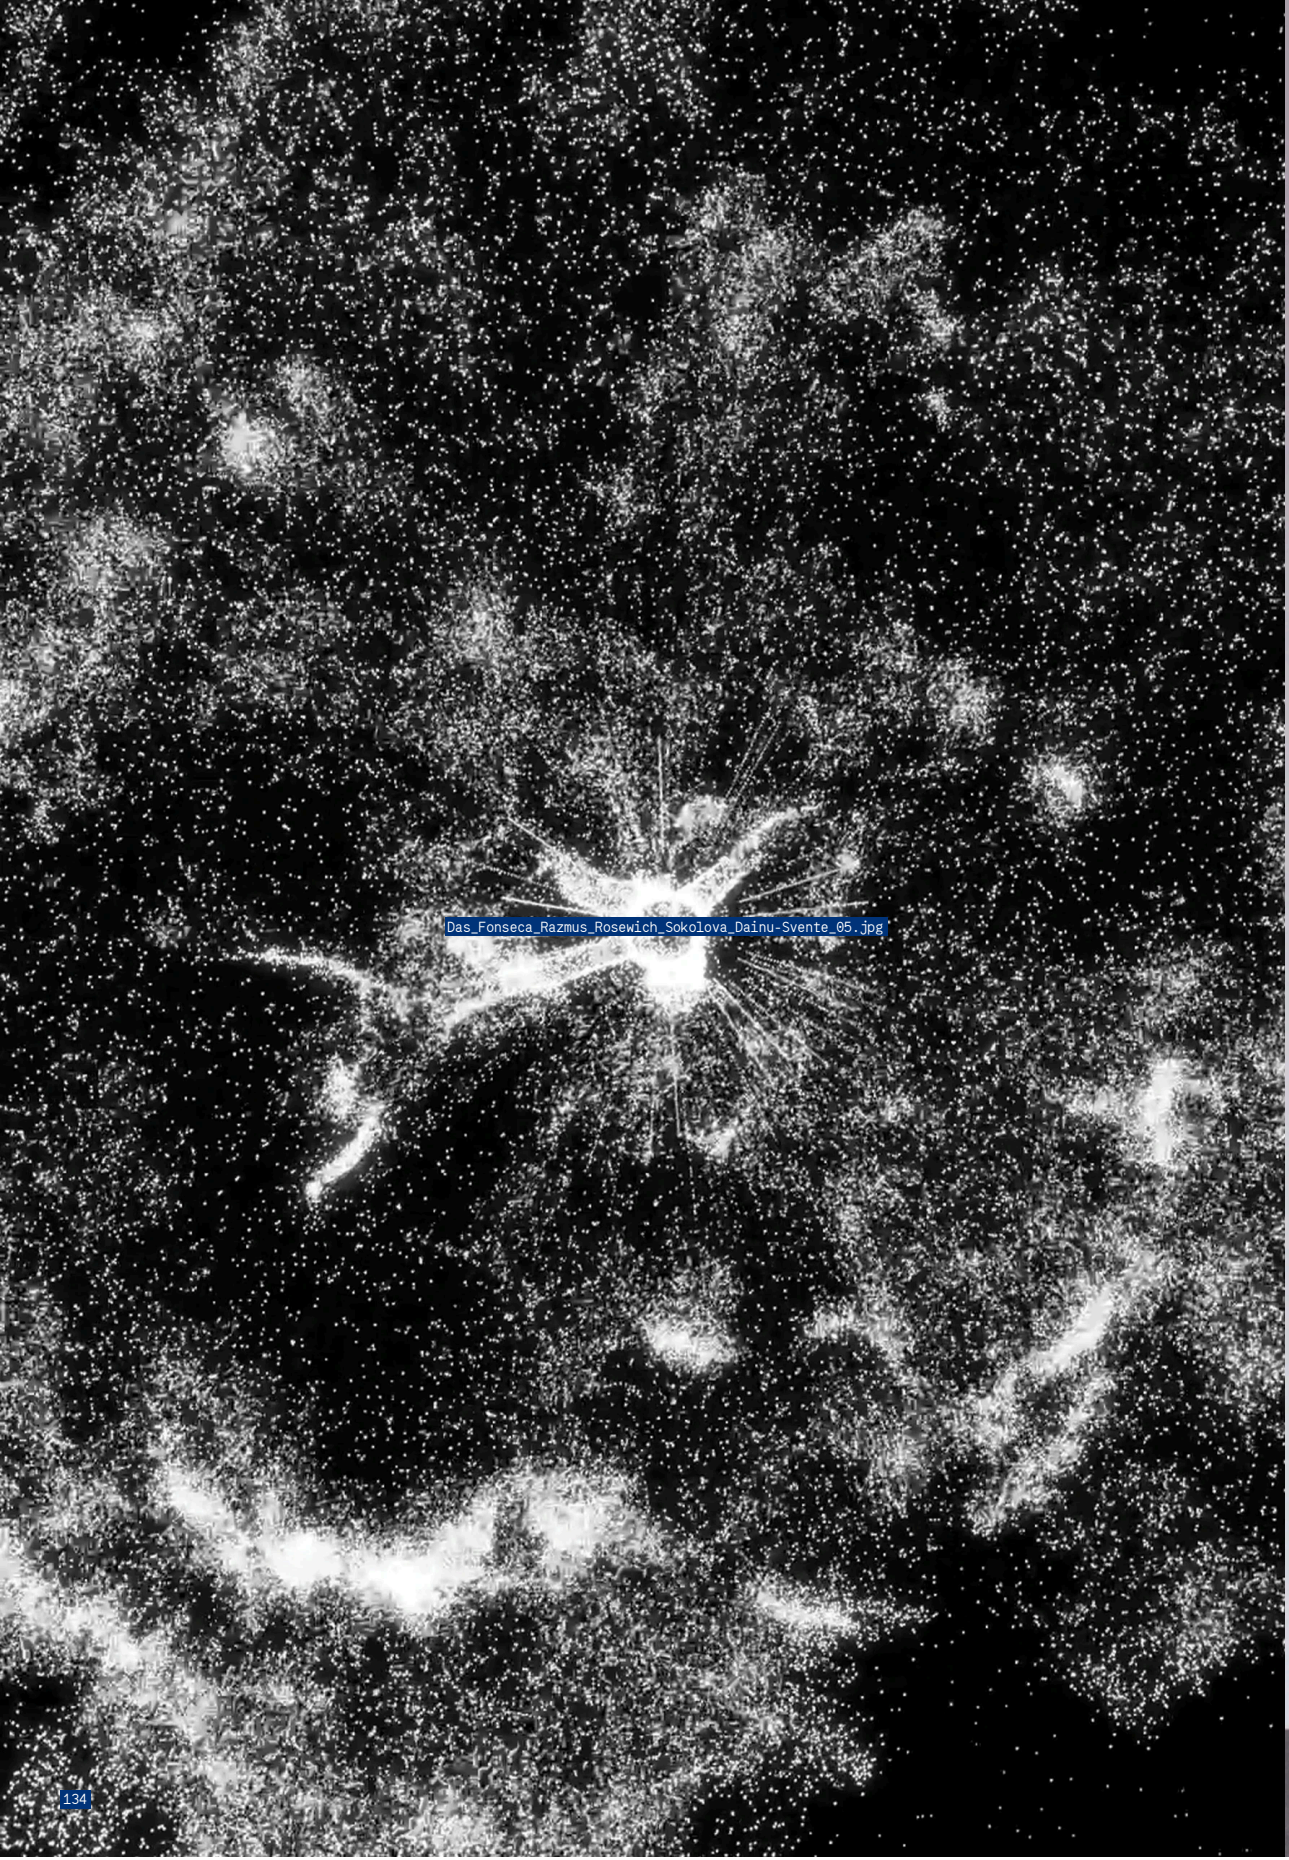
Das_Fonseca_Razmus_Rosewich_Sokolova_Dainu-Svente_05.jpg



Gomez-Baumert_Schlag-auf-Schlag_01.jpg


Kreis_MK-Ultra_03.jpg


Stigler_Bove_Subat_Beer_FloraForge_04.jpg


Hofmann_Synthetische-Erinnerung_02.png

Backhaus_Hof_Javaheri_Matzerath_Schulz_Synthetic-Reflection_05.jpg

ENIGMATIC LATENT SPACE

FROM CODE TO AI-INDUCED HALLUCINATION

(NON-) REALITY

REALITY OR FICTION?

D E E P D I V E

DISTORTED MESSAGES

C O N C L U S I O N

MITIGATION AND REGULATION

COGNITIVE PERSPECTIVES AND SOCIETAL IMPACT

S I G N I F I C A N C E

ARTIFICIAL SIGN LANGUAGE

LOST IN TRANSLATION

ChatGPT

Get instant answers, find creative inspiration, learn something new

Hallucinated Alphabet

HALLUCINATED ALPHABET GERMAN FINGER ALPHABET

A P P E N D I X

R E F E R E N C E S

Werbilo_Distorted-Messages_05.jpg


STABLEHUB
START
Home
GENERATION TOOLS
Text To Image
Image To Image
ControlNet
COMMUNITY
Gallery
LEARN & EXPLORE
Guided Tutorials
USE FROM HOME
VPN Install
Milkshake
Welcome to Stable Hub
Discover the possibilities of im| for your projects
AI Models: Online
Creative Engine: Ready
Innovation Mode: Active
Choose Your AI Tool
Transform your ideas into stu visuals with our powerful AI models
Text To Image
Write a prompt, get a custom image in seconds.
Image To Image
Upload an image and watch it transform into new styles or variations.
ControlNet
Use sketches or poses as guides to generate precise images with full creative control.
MacBook Pro

STABLEHUB

START

Home

GENERATION TOOLS

Text To Image

Image To Image

ControlNet

COMMUNITY

Gallery

LEARN & EXPLORE

Guided Tutorials

USE FROM HOME

VPN Install

Milkshake

Parameter

Variants

1 2 3

Negative Prompt

Steps 6 - +

CFG 2 - +

Seed 0

Styles






API Request

GET https://playground.transferscope.org/api/txt2img?prompt=cat%20with%20dog&cfg=2&steps=6&seed=0&variants=1&uid=Milkshake

Text To Image

cat drinking strawberry milkshake



Annaeup1_StableHub_02.png

cat holding strawberry milkshake

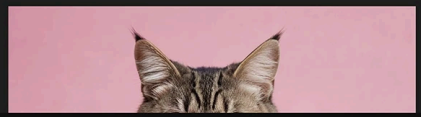

Prompt 3/77

cat with dog

Generate

Borges_Jacke_Kim_Surovceva_Dainu-Svente_04.jpg

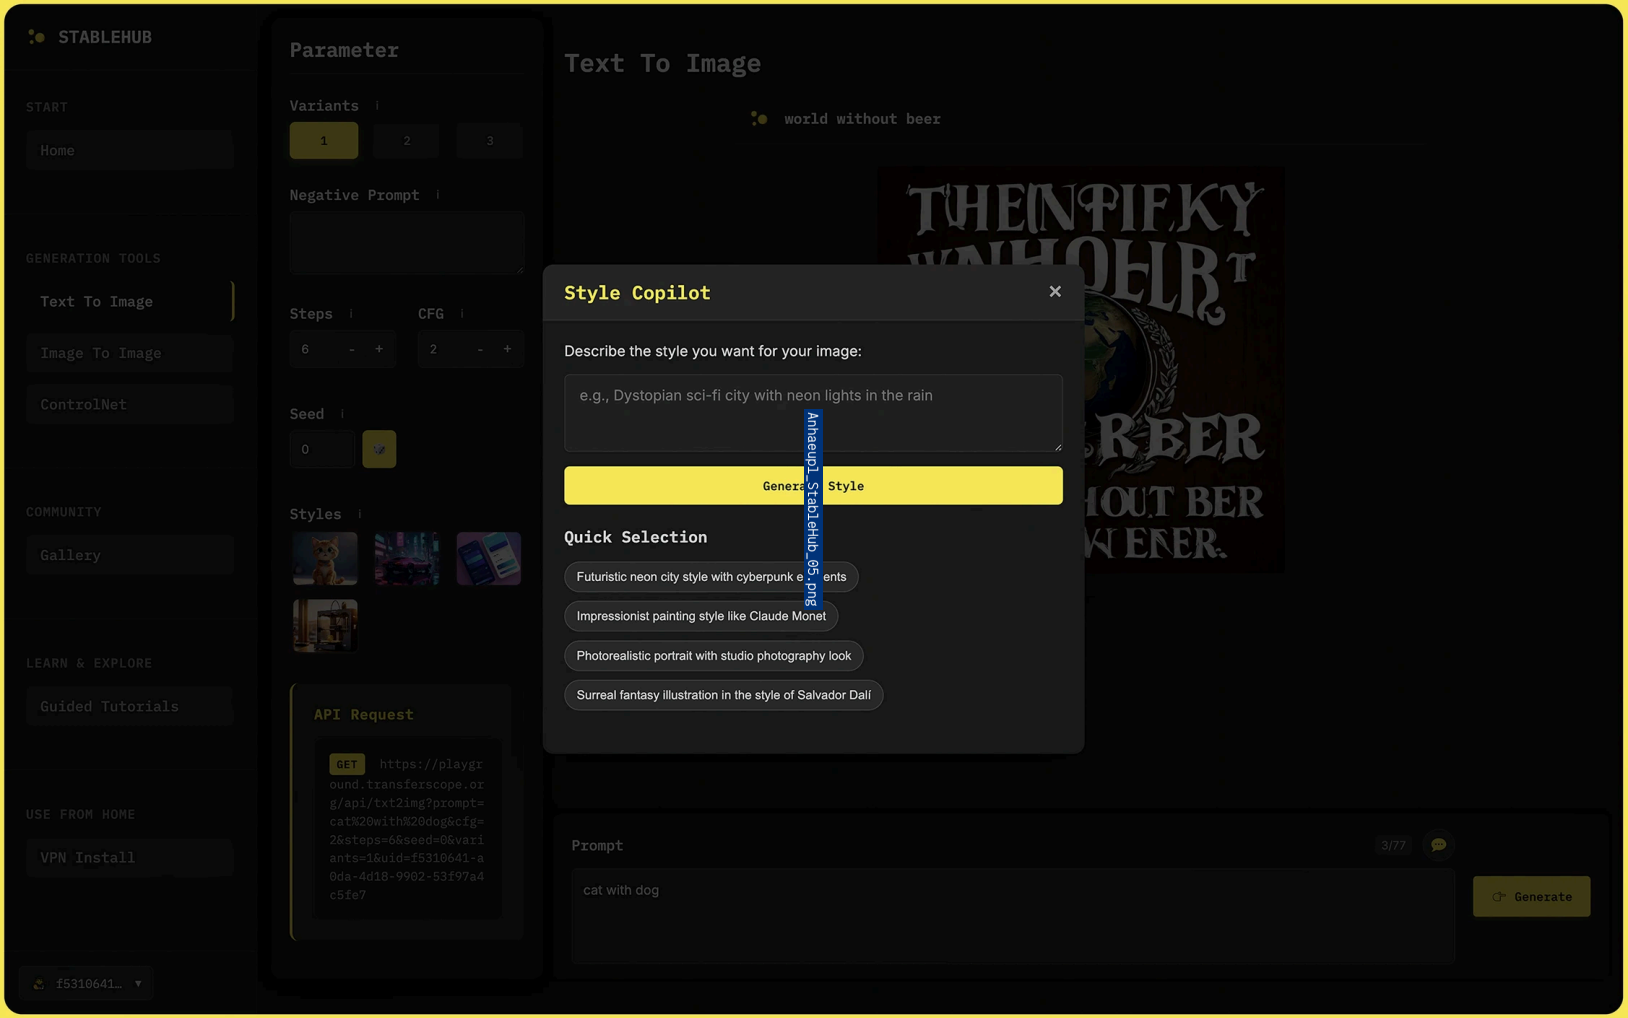
STABLEHUB
START
Home
GENERATION TOOLS
Text To Image
Image To Image
ControlNet
COMMUNITY
Gallery
LEARN & EXPLORE
Guided Tutorials
USE FROM HOME
VPN Install
Parameter
Variants
1
2
3
Negative Prompt
Steps
6
CFG
2
Seed
0
Styles
API Request
GET https://playground.transferscope.org/api/txt2img?prompt=cat%20with%20dog&cfg=2&steps=6&seed=0&variants=1&uid=f5310641-a0da-4d18-9902-53f97a4c5fe7
Text To Image
world without beer
Style Copilot
Describe the style you want for your image:
e.g., Dystopian sci-fi city with neon lights in the rain
Quick Selection
Futuristic neon city style with cyberpunk elements
Impressionist painting style like Claude Monet
Photorealistic portrait with studio photography look
Surreal fantasy illustration in the style of Salvador Dalí
Prompt
3/77
cat with dog
Generate


Wang_Stuhl_04.jpg

Scholl_Wahrgenommen_04.jpg

MY STORY

MARIS OSEI AFRIYIE
28 YRS GHANAIAN
INGELHEIM
ROOM 5

Hello there. My name is Mavis I am 28 years of old, I come from [redacted] west of Africa.
I came to Germany on the 21st of December 2018.
My Reason of coming to Germany is because I was forced to marry someone who is way older than me and against my wish As a woman I believe I have the right to Choose whoever I marry. But for me I was denied this Right this Compled me to leave my home country without any documentation. I have been living in Germany illegally for the past 3 months until I was arrested on the 28th of March on my way to Wuppertal.
For the Polizei I have not committed any Crime but I

Thielen_Synthetische-Erinnerung_03.png


Hartmann_Pereira_Wagner_Dainu-Svente_02.jpg

Scholl_Wahrgenommen_01.jpg


Deutloff_Digitale-Entfremdung_01.jpg
Herzlich Willkommen!
Bitte wähle ein Gedicht aus und sprich es deutlich in das Mikrofon, indem du einmal kurz auf den „rec"-Knopf vor dir drückst.
Wenn du fertig bist, drücke ihn noch einmal, um die Aufnahme zu beenden.
Welcome!
Please select a poem and speak it clearly into the microphone by briefly pressing the "rec" button in front of you once.
When you have finished, press it again to stop the recording.
THE FEAR OF MISSING OUT


Leuoth_Flush-of-Flesh_03.jpg



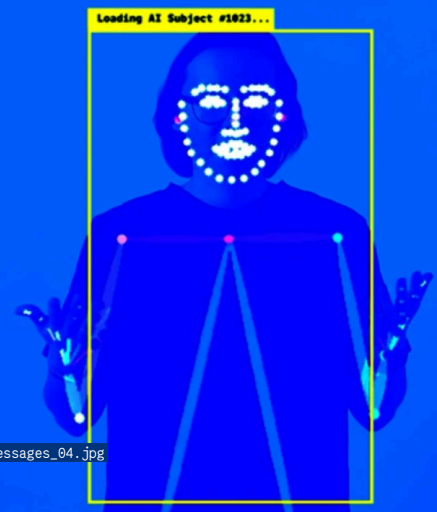

Werbilo_Distorted-Messages_04.jpg

"The intricate dance of mathematical probabilities that produces something so human-like is simply fascinating. But what about us?

As a sign language speaker, I am aware of the depth and subtleties of our communication. It is the nuances that give weight to our movements. Every hand movement, every facial expression and every subtle gesture make up the essence of our language.

But what if this cannot be captured by new technologies? When my statements become fragmented? Inconsistencies, subtle deviations that are not noticeable to outsiders at first glance can cause frustration and misunderstandings. The ethical dilemma quickly becomes the focus of attention. It's not just about recognizing the physical movements, but also about understanding the context, the emotions and the unspoken layers that lie within each sign.

The essence of what makes DGS a rich and complex form of communication

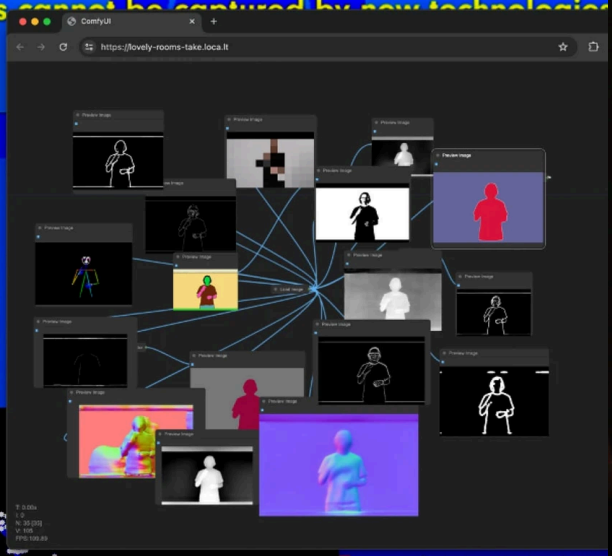

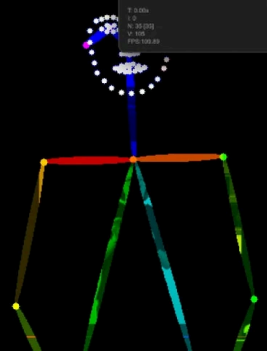

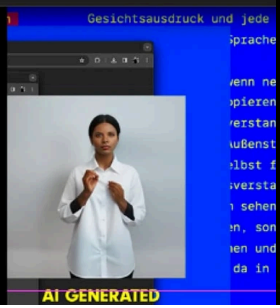

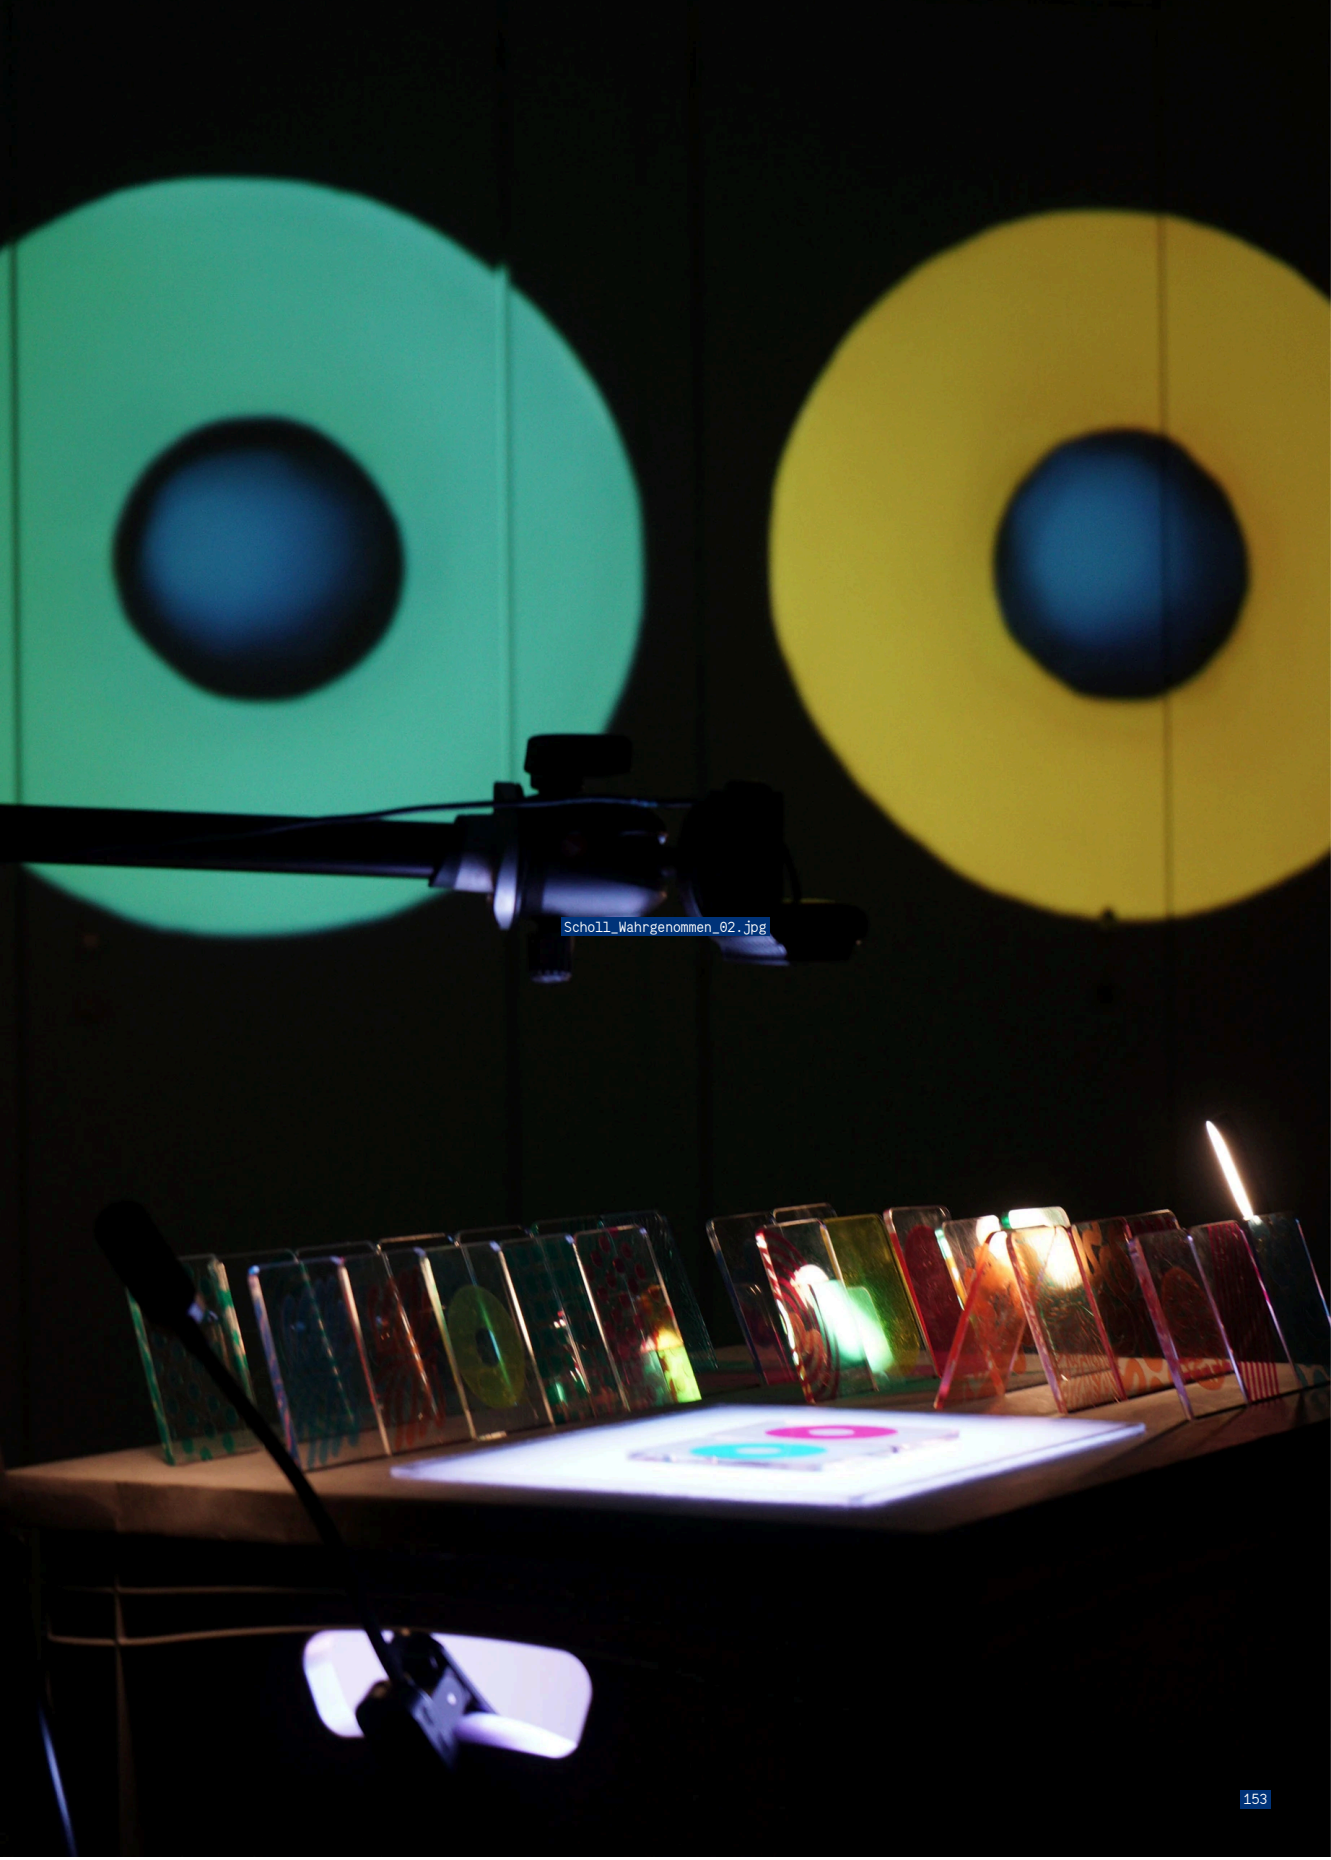
Scholl_Wahrgenommen_02.jpg

Achba_Adrian_Gatzke_Kattwinkel_Schomann_Schröder_Sideri_Zisa_Anthrobot_05.jpg


INSZENIERTE
WISSENSCHAFT
Krämer_Inszenierte-Wissenschaft_02.jpg


Dobbs_Ploiakova_Sideri_Steinbrecher_Dainu-Svente_01.jpg


Sharifi_Synthetic-Eden_02.jpg

Style Transfer + Content Mix

Open with Preview

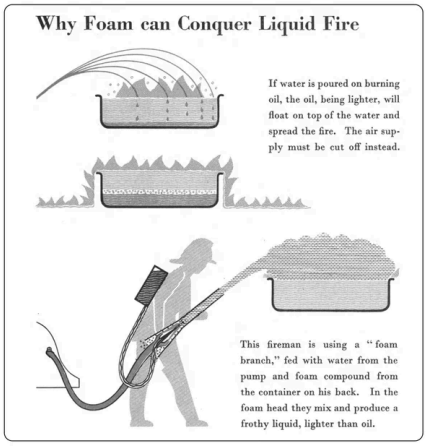

+

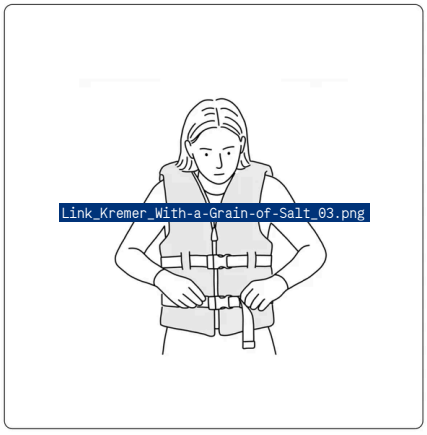

=

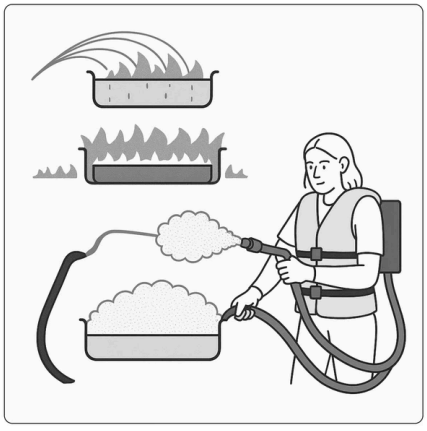



Dornheim_Served_03.jpg

Backhaus_Hof_Javaheri_Matzerath_Schulz_Synthetic-Reflection_04.jpg

HALLUCINATED ALPHABET

The first steps towards sign language often lead to the finger alphabet. Although the finger alphabet is not similar to sign language, the finger alphabet is nevertheless essential for describing [illegible] terms no signs are given to. The fact that the finger alphabet [illegible] communication is however indisputable. [illegible] generators can [illegible] eyes can find [illegible] between what [illegible] is based on [illegible] phenomenon [illegible] the Internet.

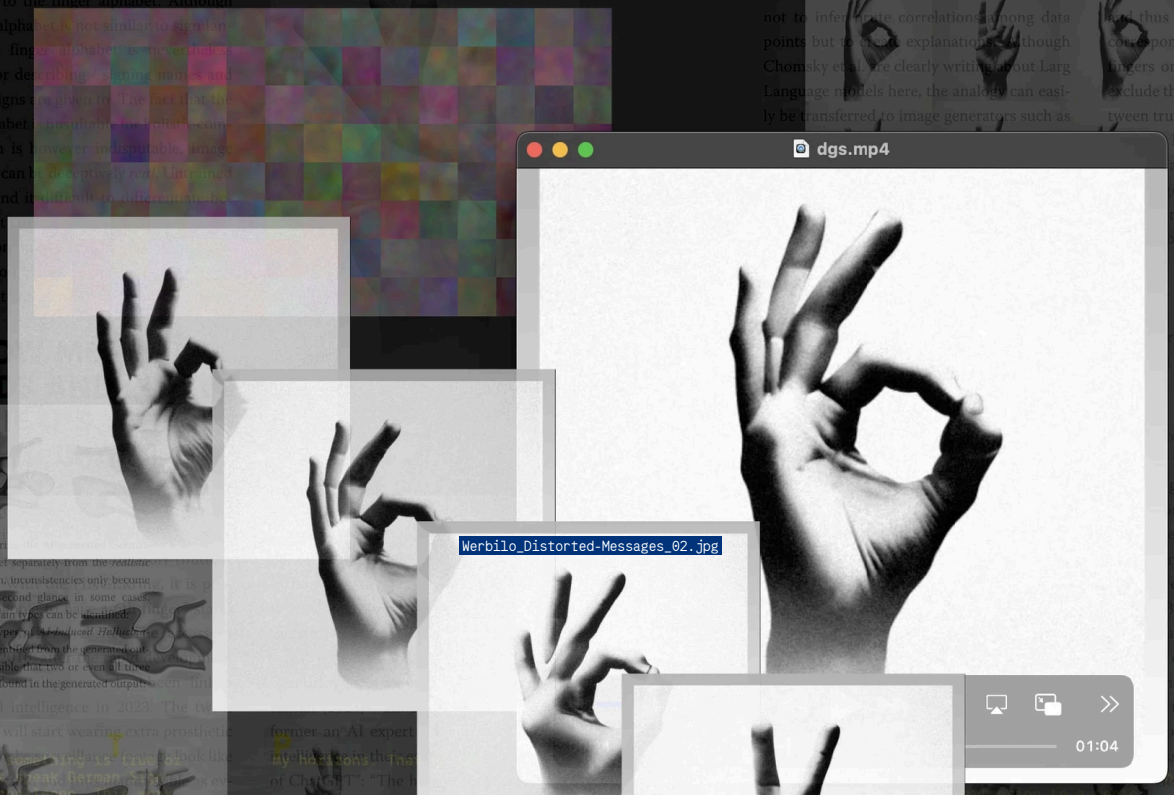

Consider [illegible] finger alphabet separately from the *realistic* representation, inconsistencies only become apparent at second glance in some cases. However, certain types can be identified. Three types of *AI-Induced Hallucination* can be identified from the generated output. It is possible that two or even all three types may be found in the generated output.

RESEMBLANCE OF THE LETTER

The hand does not form the corresponding gesture, but rather imitates the shape of the letter.

In this example, the letter 'A' should be generated in DGS. The image generator has visually imitated the 'A' instead which does not resemble the official gesture of the letter.

SUPPLEMENTAL (FINGER) FABRICATION

Additional fingers or anatomically incorrect representations of the hand are shown.

In this example, the letter 'I' should be generated in DGS. Instead of five fingers, a total of seven fingers were generated. Furthermore, one finger does not correspond to the anatomically regular representation of a hand.

FACTUAL INCORRECTNESS

The display shows either a different gesture or no gesture.

In this example, the letter 'S' should be generated in DGS. Although the number and the anatomy are correct, the letter 'S' is shown in the finger alphabet.

(!) Note: It is important to emphasize that this does not mean that a hand with less or more than five fingers does not exist in reality, but the data is more likely to be trained with five fingers.

THE ENIGMATIC LATENT SPACE

AI-created, generated images have been sorted into different clusters according to visual similarities.

The question certainly arises as to why sign language is used as an example of *AI-Induced Hallucinations*. It is a field in which I, as the author, had no points of contact prior to this work. An area that nevertheless excites and fascinates me. An area in which I can enter naively and with an open mind. An area that deals with visual communication. A common ground that nevertheless separates us, because as a designer I also communicate with visual means, but nevertheless with a different craft.

Moving on to the three artifacts. It is important for me to emphasize that it is not my intention to show

most probable answer to a scientific question." Further Chomsky et al. (2023) state: "[...] the human mind is a surprisingly efficient and even elegant system that operates with small amounts of information; it seeks not to infer brute correlations among data points but to create explanations." Although Chomsky et al. are clearly writing about Large Language models here, the analogy can easily be transferred to image generators such as

[illegible] finger alphabet. Even [illegible] knowledge of the [illegible] might be possible [illegible] whether an AI-gen[illegible] thus the letters [illegible] correspond to *real* [illegible] fingers or an unu[illegible] exclude these artifacts [illegible] tween truthful gestur[illegible] less clea[illegible] assess [illegible] alphabet [illegible] rated using [illegible] based on Dal[illegible]

[illegible] former an AI expert [illegible] my horizons. [illegible] of ChatGPT: "The h[illegible] Chatbot and its ill[illegible] my initial patter my[illegible] This work is for [illegible] those who don't understand artificial [illegible] those who stand up for marginalized [illegible] those who don't understand artificial intelligence.

At the same time, this work is address[illegible] are not worried. Not worried about art[illegible] intelligence getting out of control. Fo[illegible] *In this example, the letter 'I' should be generated in DGS.* shows us a mirror of society[illegible] stand up for marginalized groups. For [illegible] who understand artificial intelligence.

[illegible] opposing s[illegible] but the[illegible] with each other[illegible] this possibl[illegible] ert[illegible] t to show with this work[illegible] that there is [illegible] ement in certain[illegible] areas, to poin[illegible] ss[illegible] in a rapidly developing tec[illegible] t to show th[illegible] contradictions can start[illegible] course bring[illegible] topic that we may not yet be[illegible]

The question certainly arises as to why sign language is used as an example of *AI-Induced Hallucinations*. It is a field in which I, as the author, had no points of contact prior to this work. An area that nevertheless excites and fascinates me. An area in which I can enter naively and with an open mind. An area that deals with visual communication. A common ground that nevertheless separates us, because as a designer I also communicate with visual means, but nevertheless with a different craft.

Moving on to the three artifacts. It is important for me to emphasize that it is not my intention to show

[illegible] The pursuit of a methodology to c[illegible] consistent AI-generated video animation of sign[illegible] languages is certainly a very ambitious goal [illegible] retrospect an impossible one for me to achieve. [illegible] Nevertheless, a change of perspective shows what is important.

I hope this paper provides an overview of [illegible] ways in which the use of new technology can have an impact [illegible] our understanding and how a shared discourse can [illegible] us new insights.


Leuoth_Flush-of-Flesh_02.jpg

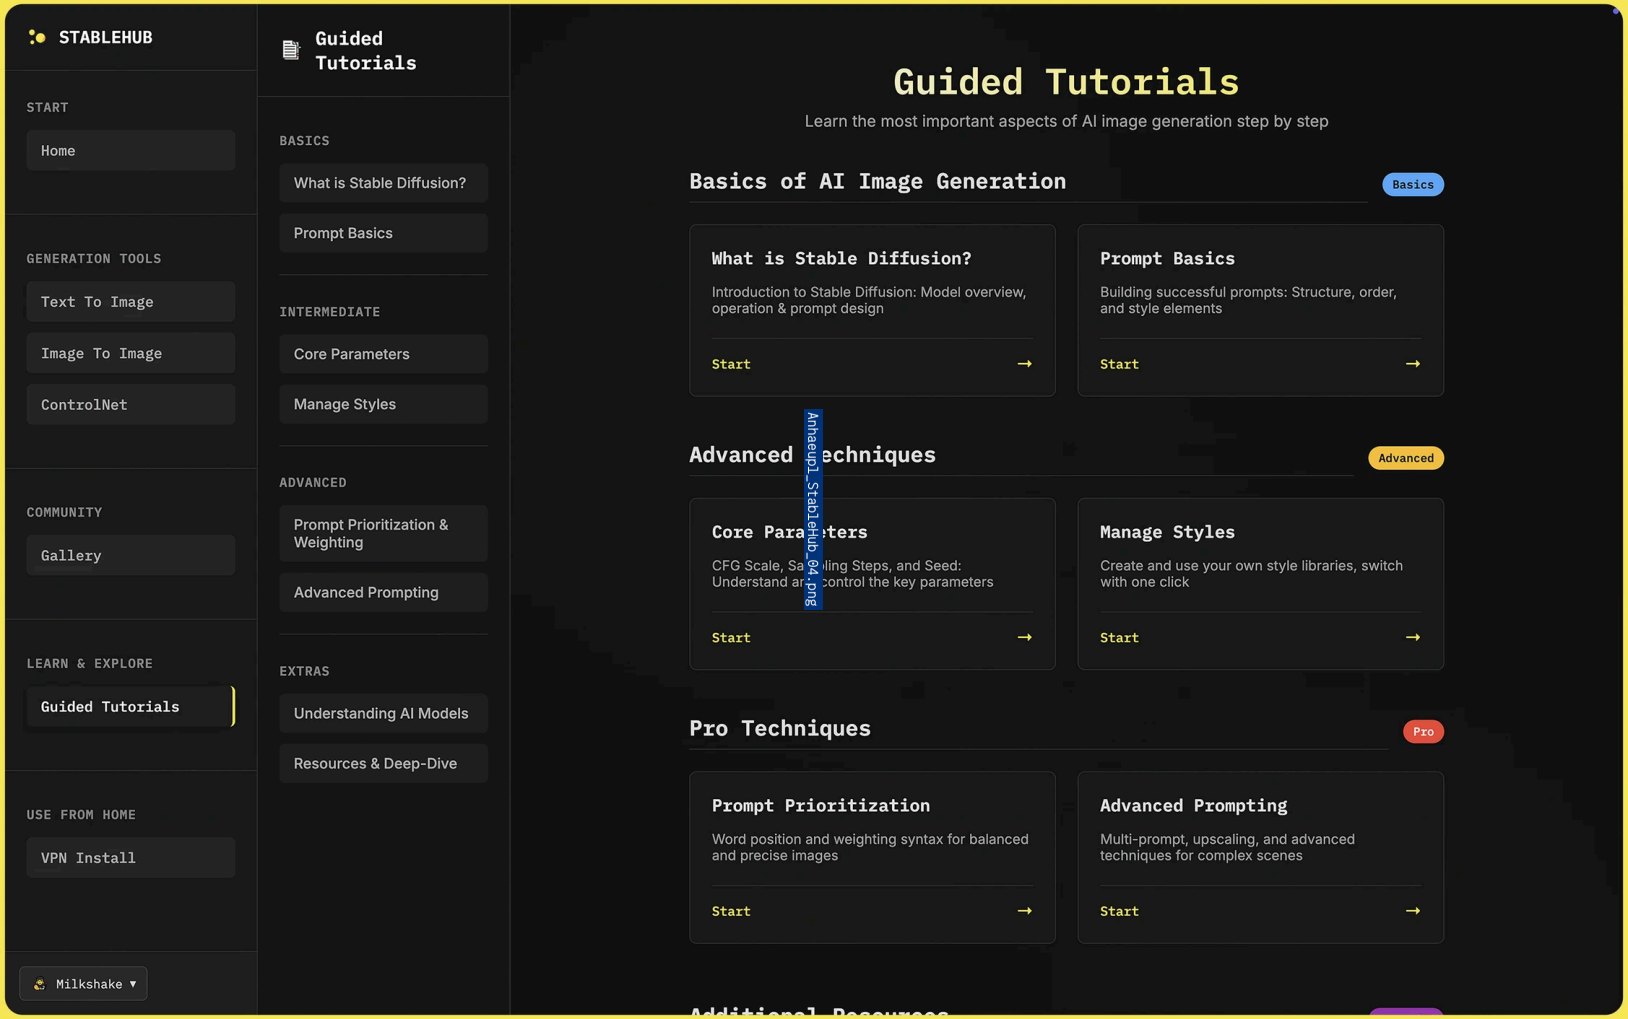
STABLEHUB
START
Home
GENERATION TOOLS
Text To Image
Image To Image
ControlNet
COMMUNITY
Gallery
LEARN & EXPLORE
Guided Tutorials
USE FROM HOME
VPN Install
Milkshake
Guided Tutorials
BASICS
What is Stable Diffusion?
Prompt Basics
INTERMEDIATE
Core Parameters
Manage Styles
ADVANCED
Prompt Prioritization & Weighting
Advanced Prompting
EXTRAS
Understanding AI Models
Resources & Deep-Dive
Guided Tutorials
Learn the most important aspects of AI image generation step by step
Basics of AI Image Generation
Basics
What is Stable Diffusion?
Introduction to Stable Diffusion: Model overview, operation & prompt design
Start
Prompt Basics
Building successful prompts: Structure, order, and style elements
Start
Advanced
Advanced
Manage Styles
Create and use your own style libraries, switch with one click
Start
Pro Techniques
Pro
Prompt Prioritization
Word position and weighting syntax for balanced and precise images
Start
Advanced Prompting
Multi-prompt, upscaling, and advanced techniques for complex scenes
Start

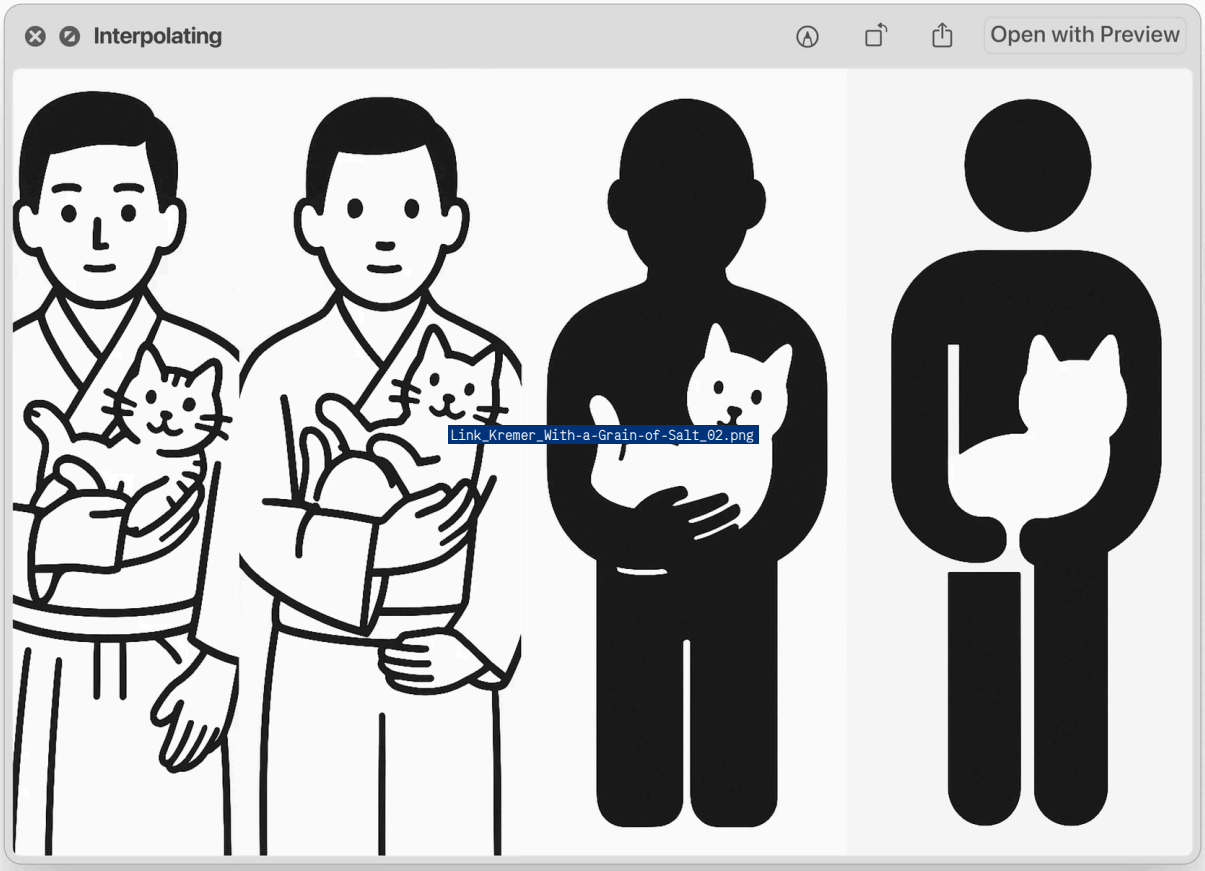
Interpolating
Open with Preview
Link_Kremer_With-a-Grain-of-Salt_02.png


Backhaus_Hof_Javaheri_Matzerath_Schulz_Synthetic-Reflection_02.jpg

Achba_Adrian_Gatzke_Kattwinkel_Schomann_Schröder_Sideri_Zisa_Anthrobot_01.jpg


Sharifi_Synthetic-Eden_05







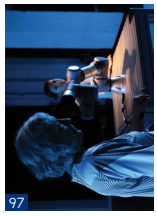

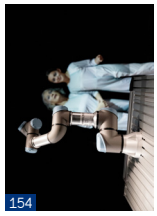

↗ efpope8i

Anthrobot

Ilona Achba, Tyanka Adrian, Philipp Gatzke, Meike Kattwinkel, Elena Schröder, Justus Schomann, Aikaterini Sideri, Jessica Zisa

III. p. 24, Chunk 13: From Physical Computing to…
I. p. 88, Chunk 14: Robotik und Computer Vision…
I. p. 19, Chunk 3: HfG Offenbach

Chunk 1 As part of the ANTHROBOT project, KISD students programmed a six-axis industrial robot to function as an active performer in a dance installation involving a child and a senior. In collaboration with the theatre collective DIPHTHONG, the students programmed the robot's movements, designed interaction sequences, and co-created the performance's audio landscape. The result is a responsive robot capable of reacting to and generating physical impulses in dialogue with the dancers. The dance performance ANTHROBOT explores the topic of human relationships across generations as well as our relationship to technology.

Oden, a collaborative industrial robot arm from Universal Robots is part of the inventory of KISD's AI Lab, thelaboratory of living objects. It has six joints that allow it to stretch 50 cm and lift 3 kg. It offers students the opportunity to research and develop forms of human-machine interaction.

III. p. 173, Chunk 3: Stuhl
II. p. 258, Chunk 11: reshape 23
I. p. 66, Chunk 14: AI+D Lab (HfG Schwäbisch Gmünd)

Chunk 2 Photo Credits Figure 01–03: Laura Wagner, Figure 04–05: Alessandro De Matteis

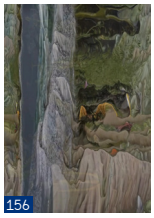





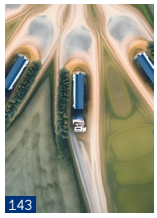

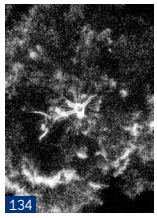

↗ wzdf9i2x

Dainų šventė

Inês Fernandes de Aguiar, Alina Bertacca, João Borges, Bidisha Das, Bia Dobbs, João Nuno Rua Fonseca, Jana Hartmann, Luka Jacke, Yonghyuk Kim, Katrina Milta, Mário Pereira, Jovydas Petrėtis, Jule Pickel, Elzbieta Ploiakova, Vytenis Razmus, Daniel Rosewich, Aikaterini Sideri, Anna Sokolova, Hannah Steinbrecher, Anastasija Surovceva, Ada Wagner

Chunk 1 At the occasion of the 100th Lithuanian Song Festival, KISD students worked in international teams to create audiovisual works, which they presented at the Lithuanian Parliament's Visitor Centre in Vilnius.

The project focused on engaging with the century-old tradition of celebrating Lithuanian song. Traditional Lithuanian songs were selected and recorded by Vilnius Tech's choir, 'Gabija'. These songs formed the basis for the students' audio-visual works. New animation technologies, such as 3D reconstruction, generative graphics, and generative artificial intelligence, were employed to create a series of videos and explore new perspectives on the Song Festival, recognized by UNESCO as intangible cultural heritage.

The resulting works offer various insights into the intersection of cultural tradition and novel technologies.

Chunk 2 Themes and creative concepts range from accessing and editing historical material, to mixing real-life footage with 3D reconstructions and modeled 3D worlds, to producing original songs based on the provided traditional ones, to utilizing artificial intelligence for both visuals and sound.

In collaboration with Vilnius Technical University, University of Porto, Riga Technical University, University of Greifswald and Academy of Media Arts Cologne

Video documentation is available at: https://vimeo.com/1021804361

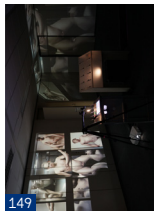

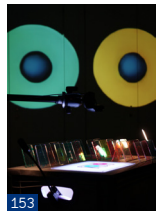

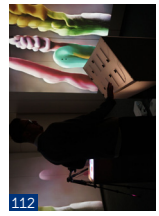

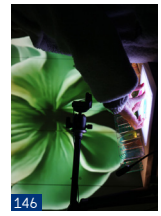

↗ stdo5rkg

Der Raum kann wahrgenommen werden

Daniel Mikula Scholl, Peter Ehses

Chunk 1 The interactive multi-channel video work "Der Raum kann wahrgenommen werden" invites viewers to explore the blurred, unexpected results of machine imagination through a tactile, non-verbal interface. Instead of typing prompts, visitors interact via sliders and printed cards.

Chunk 2 When a card is placed, the AI generates a distorted projection that attempts to match the pattern - but always fails. These misinterpretations layer across the walls, forming a shifting digital "membrane" that warps with each new input. The result is an accumulation more than a story: the room becomes a dense, chaotic record of the AI's failed translations. Sliders offer a more technical interaction - opaque and machine-like - while the cards evoke craft or even ritual. The goal is not to refine output, but to shift focus from language-image relations towards direct, sensory engagement with machine imagination.



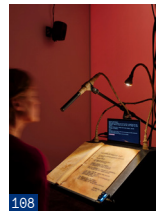

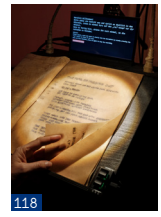

↗ yxhqcoo9

Digitale Entfremdung

Elisa Deutloff

Chunk 1 Besucher*innen der interaktiven Installation *Digitale Entfremdung* setzen sich an ein Holzpult, das an eine Kanzel erinnert. Dort liegen Gedichte, die auf Hautimitat aus Latex handgeschrieben und eintätowiert wurden. Besucher*innen sprechen eines davon in ein Mikrofon, wodurch ihre Stimme synthetisiert wird. Daraufhin können sie einem Large Language Model Fragen stellen, welches mittels Finetuning an die Online-Persona Elisa Deutloff's angepasst wurde. Das Model antwortet mit der geklonten Stimme, eingebettet in eine Vier-Kanal-Soundscape.

Chunk 2 Besucher*innen treten in einen Dialog mit einer Stimme, die ihre eigene ist, ihnen aber nicht mehr gehört. Dieses Spannungsfeld zwischen Intimität und Unnahbarkeit macht die Installation zu einem Raum der Selbstbegegnung im digitalen Zwielicht. Die Arbeit verweist auf ein Paradoxon digitaler Kommunikation: Sie verbindet, während sie gleichzeitig isoliert.

Chunk 3 Sie stellt die Frage nach der Vergänglichkeit von Sprache und Identität in einer entkörperlichten Welt.

Gedichte von Zici Pütz, Seongsin Lee und Charlotte Simon. Tätowierung von Verena Mack. Zusätzliche Programmierung von Paul Eßer und Maximilian Fuchs. Pult von Elias Rosenberg. Danke an Jenny Sophie Kasper, Leandra Kleber, Prof. Alexander Oppermann, Leon Kühr, Mattis Kuhn, Natalie Wilke und das Team des Frankfurter Kunstvereins.

Photo Credits: Norbert Miguletz (Frankfurter Kunstverein)



124


161



102



152

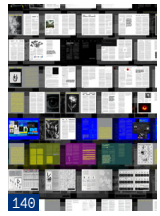
140

↗ 549a4lof

Distorted Messages

Viktor Werbilo

II. p. 69, Chunk 20: Editorial
III. p. 40, Chunk 26: Building AI Intuition - Four…
III. p. 85, Chunk 18: Die Pille für KI

Chunk 1 *Distorted Messages* exposes the risks of generative AI for marginalized communities, using German Sign Language (DGS) as a case study. Despite the promises of AI, in terms of accessible learning and communication tools, it often fails to do so – mainly due to biased training data that excludes underrepresented groups.

Through three experimental works, this design research project highlights how different AI modalities can distort our reality: *Hallucinated Alphabet* reveals image-based hallucinations through incorrectly generated finger alphabets; *Learn something…* contrasts faulty text-based DGS descriptions with accurate visual signs; and *Un clear* shows how video generation lacks the precision needed to convey meaningful gestures.

These errors are not just technical - they risk reinforcing exclusion and misinformation. By making these flaws visible, *Distorted Messages* calls for more inclusive AI systems and points out who is left behind when precision, representation, and accessibility are missing.

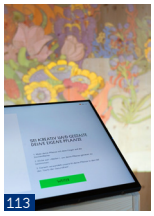
113

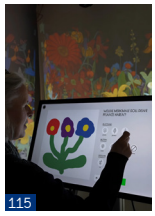
115

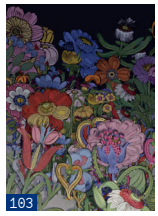
103


137

↗ ogcamyc8

FloraForge

Agnes Stigler, Berit Bove, Colin Subat, Oliver Beer

Chunk 1 In der interaktiven Installation „FloraForge" können Besucher*innen mithilfe generativer K.I.

Chunk 2 individuelle Pflanzen kreieren. Das Projekt basiert auf dem „Gart der Gesundheit", einem der ältesten gedruckten Werke über Kräuterkunde, das 1485 in Mainz erschien und Wissen über Heilkräuter verbreitete.

In der Installation können Besucher*innen auf einem Touchtable selbst Pflanzen malen und anhand von Kategorien, die Form und Eigenschaften ihrer Kreation beeinflussen.

Chunk 3 Mit Hilfe von Prompts, Controlnet und einem eigens trainierten LoRA Modell werden die Pflanzen in den Illustrationsstil des Buches umgewandelt und können an der Ausstellungswand als Teil eines projizierten Gartens bewundert werden. Dieser wird durch vvvv Gamma und eine Geräuschkulisse zum Leben erweckt und lädt zum Staunen und Entdecken ein.

Ein Prototyp wurde bereits am RLP-Tag 2025 in Kooperation mit der BUGA29 ausgestellt. Die fertige Installation ist Anfang 2026 als Teil der "Flora" Ausstellung im Naturhistorischen und Gutenberg-Museum Mainz zu sehen.


131


162


151

↗ 6kyqpzaw

Flush of Flesh

Annika Leuoth

Chunk 1 The AI image (Fluxturbo) and the photo of an imitating miniature (40cm x 40cm x 40cm) depict a toilet in the center of the scene. The eye is immediately caught by a big gaping mouth with human teeth which is merged with the open toilet lid. A white lacy fabric is spit from the mouth into the toilet bowl. The room itself is covered in red tiles; some appear slightly damaged or uneven.

Chunk 2 At second glance, the perspective seems to be distorted, evoking a feeling of disorientation. A cold and greenish light creates an eerie atmosphere. The mirror in the background gives the illusion of depth but also increases the claustrophobic feeling of the room. The surreal and macabre imagery plays with the contrast between an inanimate cold object and an animate human body part by recombining familiar elements and thereby creating a new uncanny fever dream. By replacing a bodily fluid with lacy fabric, the image walks the line between what disgusts and what aesthetically pleases us.


99


155


119


107


117

↗ fr86ux1b

Generative Künstliche Intelligenz und die Konstruktion von Glaubwürdigkeit

Anna Krämer

III. p. 174, Chunk 2: Synthetische Erinnerung
III. p. 85, Chunk 20: Die Pille für KI

III. p. 252, Chunk 1: Auf dem Weg zu KITeGG

Chunk 2 Wir leben in einer Zeit, in der Informationen unabhängig von Expertise oder Wahrheitsgehalt veröffentlicht und verbreitet werden können. Mit generativer künstlicher Intelligenz lassen sich Bilder erzeugen, die keinen Bezug zur Realität haben müssen und dennoch glaubwürdig wirken.

Chunk 3 Gerade im wissenschaftlichen Kontext, in dem Bilder oft als objektiv und belegend wahrgenommen werden, stellt das eine Herausforderung dar. Im Buch „Inszenierte Wissenschaft" werden 37 pseudowissenschaftliche Themen mithilfe von ChatGPT-4o visuell inszeniert. Die Bildsprache orientiert sich an etablierten Formen wissenschaftlicher Darstellungen, um zu testen, wie überzeugend sich wissenschaftliche Glaubwürdigkeit schon heute konstruieren lässt – auch wenn der Inhalt zweifelhaft ist.

II. p. 114, Chunk 1: Artefact Hunt
I. p. 58, Chunk 10: KI Labor (HfG Offenbach)
I. p. 172, Chunk 2: A computer generated map?…

Chunk 4 Das Projekt versteht sich außerdem als Einladung, die eigene Wahrnehmung zu hinterfragen: Was macht ein Bild glaubwürdig? Wie unterscheiden wir zwischen echten und künstlich erzeugten Inhalten? Und wie können wir lernen, damit bewusster umzugehen? Ziel ist es, aufzuzeigen, wie leicht KI-generierte Bilder unsere Einschätzung von Informationen beeinflussen können – und welche Verantwortung dabei auf Seiten der Gestaltenden liegt.

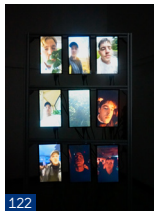
122

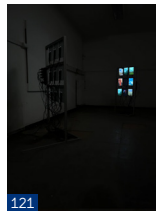
121

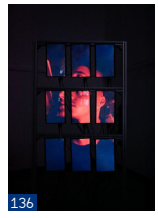
136


109

↗ s1oy6d0n

MK Ultra

Max Kreis

Chunk 1 In Max Kreis' diploma work, MK Ultra, a wordplay on his initials and a hint at the brain-altering capabilities of the technology used, he deals with his most personal digital memories from the past seven years and tries to weave those digitally altered fragments into a narrative multichannel (18) video work consisting of six 33-minute segments of speculative realities, cut to a score created by the artist himself. Whereas his former image models trained on his memories, a practice he started in late 2019, were more abstract in their outputs and looked more painterly and vague, MK Ultra represents the first AI based work of the artist, where his most private inputs are not as obscured in their respective output. Therefore plausible alternateversions of his past self become visible throughout the work leading to an overall uncanny, ghostly atmosphere.

Chunk 2 To achieve the clips used in the work he fine-tuned an open source Text-To-Video model on his own data – a mixture of photos and videos recorded by himself with his phones of the span of the past 7 years. All of those speculative fragments were then woven manually by myself into two six minute experimental video loops.

Photo Credits Figure 01–02: Kerstin Weiser

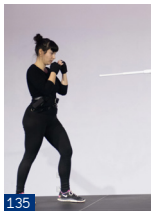

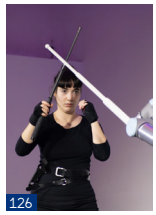



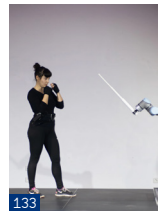

↗ g6m1qb9o

Schlag auf Schlag

Nelli Gomez Baumert

Chunk 1 In the performance piece Schlag auf Schlag, Nelli Gómez Baumert engages both linguistically and physically with the patriarchal violence and masculinities in her environment. Through a fight choreography that includes elements from kickboxing, aikido, and karate, she arms herself against what she has experienced in the past and what lies ahead in the future.

Chunk 2 Her training partner, the UR10e, a cobot from Universal Robots, responds in real time to her attacks and challenges her. This robot, intended for industrial use, here becomes her choreographic ally. She uses the UR10e's "FreeHand" function to determine its position in space. This requires precise settings and programming of speed, force, and timing.

III. p. 34, Chunk 10: Building AI Intuition - Four…

Chunk 3 The robot and the artist meet with the defensive object of a baton, while a text over loudspeakers forms the soundscape of the performance.

Photo Credits: Cheesoo Park

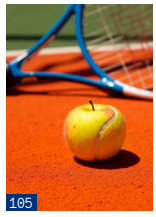





↗ dlqp7lsi

Served

Maybritt Dornheim

II. p. 182, Chunk 1: SandwichNet

Chunk 1 The work is the photographic interpretation of an image generated with the help of artificial intelligence.

Chunk 2 It shows the blending of realities and was modeled as closely as possible on the reference image. For the fusion of apple and tennis ball as the main motif, an apple was processed, painted, carved and glued in such a way that it resembles the texture and appearance of a tennis ball. In the photo studio, the materiality and the lighting situation of a tennis cinder court and direct sunlight were recreated. A strip of felt was placed in the painted apple, wool was cut into small pieces, and the fluff was glued on. This creates the deceptive and surreal image that tennis was actually being played with a kind of "apple ball". Inspiration was drawn from artists such as René Magritte and Lee Materazzi, films such as "Challengers" and everyday moments and objects.



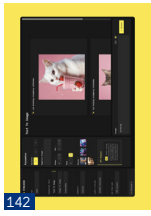

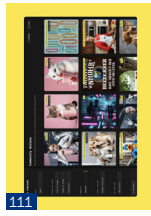

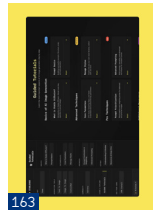



↗ 4cofmn0h

StableHub

Rahel Flechtner, Franz Anhaeupl

I. p. 65, Chunk 10: AI+D Lab (HfG Schwäbisch Gmünd)
II. p. 194, Chunk 4: Work in Progress
II. p. 211, Chunk 10: Uncanny Type

Chunk 1 StableHub (↗ https://stablehub.aid-lab.hfg-gmuend.de/) is a project designed to give students deeper knowledge and practical understanding of image generation. I analyzed the interfaces of existing tools and surveyed students to identify their pain points and needs.

Chunk 2 The survey revealed a strong interest in integrating AI-driven image generation into existing design workflows.

Together with Christopher Pietsch and Felix Sewing, both members of the AI+Design Lab, I developed a platform that lets students experiment with three core models: Text-to-Image, Image-to-Image and ControlNet. All operations run on a locally installed Stable Diffusion engine.

A tutorial page walks users through every parameter and function step-by-step. A built-in gallery allows students to upload their favorite results, like and browse peers' images within the university ecosystem, and draw inspiration.

Chunk 3 Each month, the most-liked image is exhibited on campus.

I also integrated a Style Copilot: a locally hosted Llama model that turns short style descriptions into fully formed Stable Diffusion prompts and appends them automatically.


101


110

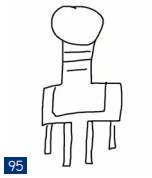
95


145

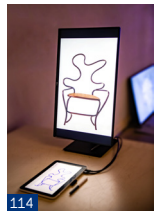
114

↗ rn7xu2ay

Stuhl

Zhichang Wang

Chunk 1 The interactive installation Stuhl (chair) uses real-time AI generation to expand the design process of chairs derived from the classic Thonet models. Visitors can draw or sketch freely on a graphics tablet.

Chunk 2 Whether spontaneous strokes or simple sketches – the system sparks the imagination and helps to design customized Thonet chairs. If the drawing is very similar to a chair, it is transferred into a design that is clearly recognizable. Conversely, if the drawing is far away from a chair, the system is given more freedom.

A ComfyUI workflow ensures a smooth and fast creation process so that the AI system can process image requests efficiently. It is based on Stability AI's SDXL base model and uses a custom-trained LoRA model together with the LCM model for real-time image generation.

III. p. 168, Chunk 2: Anthrobot
II. p. 25, Chunk 14: Schwäbisch Gmünd
II. p. 197, Chunk 11: Work in Progress

Chunk 3 Photo Credits Figure 05: Cheesoo Park

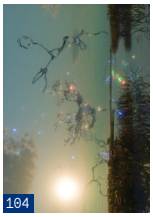
104


157

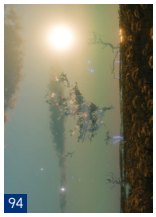
94

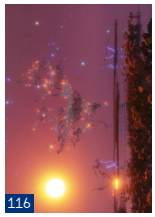
116


167

↗ tkk8worn

Synthetic Eden

Pourya Sharifi

II. p. 107, Chunk 1: Sustainable AI
II. p. 112, Chunk 15: Sustainable AI

Chunk 1 *Synthetic Eden* presents a dystopian vision in which human influence on the ecosystem becomes so overpowering that the consequences of industry and technology turn into integral parts of nature itself. This experiment illustrates the loss of naturalness in our environment caused by human intervention. A photo taken in the Stadtwald Köln served as the visual basis. Using AI, the scene was transformed into a synthetic version of itself – a symbol of a possible future in which artificial elements dominate the natural surroundings.

The AI-generated scene was then reconstructed in a 3D environment to create an immersive experience. Despite its unnatural composition, the scenery evokes a feeling that is both captivating and unsettling.

Chunk 2 This synthetic ecosystem reveals a world shaped by human interference, continuing to exist independently and offering a disturbing glimpse into an alternative future reality.

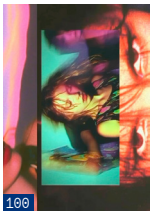
100


165

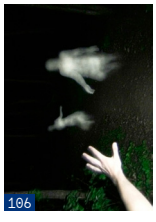
106


160

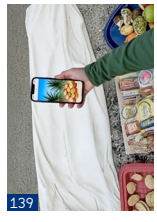
139

↗ n83dq7c7

Synthetic Reflection

Marie Backhaus, Henrike Hof, Levon Javaheri, Eva Matzerath, Phillip Schulz

III. p. 40, Chunk 26: Building AI Intuition – Four…

Chunk 1 Synthetic Reflection is a short film about AI. Unfolding in five chapters set in unique generative environments, the film reflects on the societal impact of AI and its implications for the future of artistic expression and human identity.

Composed entirely of AI-generated content, including visuals, sound, text, and narration, the film blurs the line between human and machine creativity.

II. p. 4, Chunk 3: Vorwort
III. p. 5, Chunk 13: Editor`s Note
II. p. 263, Chunk 5: TRANSFORM 23

Chunk 2 Every element was curated and assembled by humans into a cohesive cinematic experience.

The narrative explores themes including AI innovation and the fascination surrounding it; its effects on various fields; personal creativity in an AI-driven world; ethical and legal concerns; and the shifting boundary between artificial and real. Each chapter offers a distinct perspective on the promises and uncertainties of AI.

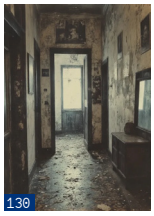
130


138

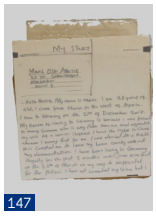
147


96

↗ q7zxje87

Synthetische Erinnerung

Berire Akdogan, Hannes Bader, Alissa Bernhardt, Piroschka Benedikta Brucker, Marlene Grabenwöger, Helena Hofmann, Georgia Ioannidou, Lisa Oelsen, Tristan Schulze, Niklas Thielen, Lea Waldera, Cosima Wagner, Micha Wempe, Anastasia Verena Wilbrandt, Jennifer Zachlod

III. p. 171, Chunk 2: Generative Künstliche Intelligenz…

Chunk 1 Die Erinnerung ist das Wiedererleben von Erlebnissen und Erfahrungen aus der Vergangenheit. Sie können individuell, kollektiv und historisch sein, aber auch lückenhaft, verzerrt und beeinflusst.

Chunk 2 Unsere persönlichen Erinnerungen halten wir in Tagebüchern, Schnappschüssen, Videoclips, Sprachaufnahmen oder Zeichnungen fest – um wichtige Momente und Gefühle zu bewahren. Durch generative Künstliche Intelligenz können wir solche Medien heute täuschend echt und in Sekundenschnelle synthetisch erzeugen oder manipulieren. Damit wird unser Vertrauen in diese Erinnerungsstützen grundlegend infrage gestellt. Die Grenzen zwischen Fakt und Fiktion verschwimmen, und die Relevanz von aufmerksamer und verantwortungsbewusster Rezeption steigt enorm.

Die Ausstellung »Synthetische Erinnerung« präsentiert gestalterisch-künstlerische Positionen von Studierenden des Master-Studiengangs »Design is all you need – Gestaltungslehre und KI«, die sich kritisch mit der Produktion synthetischer Medien auseinandersetzen – sowohl mit ihren Potenzialen als auch mit ihren Risiken.

Die gesamte Ausstellung wurde durch einen virtuellen Rundgang dokumentiert: https://my.matterport.com/show/?m=ZZVngdqpwgu


128


164


158


132



120

↗ ifckivm4

With a Grain of Salt

Moritz Hartstang, Alicia Link, Robert Kremer

Chunk 1 In disaster response, universally understood visuals are vital for inclusive communication, yet designing imagery that balances clarity with necessary complexity remains challenging. AI can help overcome this by detecting unintuitive misinterpretations while smoothly bridging detailed and minimalist designs.

In our Bachelor's thesis on Design for Inclusive Disaster Response, we leveraged natural language-driven image generation to explore this challenge through four conceptual approaches:

1. Extrapolation: Extending visuals based on a character or set of pictograms as input (see Figure 1).

2. Interpolation: Exploring smooth transitions across levels of abstraction (see Figure 2).

Chunk 2 3. Style Transfer and Content Mix: Swapping content and aligning the style accordingly (see Figure 3).

4. Reverse Testing: Allowing the AI to evaluate the comprehensibility of our often abstract results (see Figure 4).

We often encountered limitations when the style became too abstract or unfamiliar to the model.

Models used: GPT-4o and Flux.1 Kontext

Experimental

Research

182 I Think He's Got a Helicopter in His Pocket: The AI Chewing Gum
Thomas Hawranke, Lasse Scherffig

188 An Incredibly Average Face – Approaching 31,200 Images with Human Eyes, Sorting Systems, and Collaging
Jana Hartmann

204 Attention to Color – The Role of Saliency in Image-Based UI Theming
Robert Halbach

212 Body as Latent Space: Alternative Compressions in Aesthetic Practice
Simon Maris

218 On the Materiality of AI
Ting-Chun Liu

224 un/learn the process
Francesco Scheffczyk, Paul Eßer, Lars Hembacher

230 26 Strategies to Consider When Writing with AI
Julia-Jasmin Bold

234 Notes Toward AI Un/Learning
Simon Maris

Introduction

Alex Oppermann, Mattis Kuhn,
Leon-Etienne Kühr

Chunk 1
In the interview “Planetary Thinking”,[1] philosopher Yuk Hui advocates technodiversity over the prevailing technological monoculture. He understands technodiversity as many parallel strands of technological development that take place at an individual and local level.

Chunk 2
Artistic approaches in particular contain potential for this. While the role of technology in everyday functional contexts is usually fixed and leaves little freedom of action, our relationship to and use of technology can be shaped more freely and explored in art. On the one hand, it is also used as a tool according to the intentions of the developers, but on the other hand, it can change what we understand as art. But instead of just asking how technology changes art, we can also ask how art changes technology.

In “I think he's got a helicopter in his pocket: The AI Chewing Gum”, (↗ zf6iham6) Thomas Hawranke and Lasse Scherffig explore the relationship between text/language and (moving) image in the context of AI-generated videos. In their reenactment of the avant-garde classic “The Girl Chewing Gum” by John Smith with current technologies, they use its script as prompts—for example “two pigeons fly past from right to left and two boys run past from left to right”—for video generation. While in the former the categorical difference between text and image is exposed in the script, in Hawranke's and Scherffig's experiment it is generated by the AI systems and shows—according to their thesis—that prompts reveal more about the training data used than controlling a creative process.

Chunk 3
Based on Hito Steyerl's thesis on AI images as *mean images*, i.e. average images that encapsulate social norms, Jana Hartmann's project “An Incredibly Average Face - Approaching 31,200 Images with Human Eyes, Sorting Systems, and Collaging” (↗ 14as24mw) examines the possibility space of generated imagery, particularly in relation to the representation of the human face. Using a script, Hartmann generates 31,200 portraits according to four parameters (beauty description, object, style, seed), which she further analyzes and visualizes algorithmically.

Chunk 4
In addition to the mathematical representations of the collection in 3D matrices and similarity maps, she also deals with the faces artistically: using automated collage techniques, she combines images of different prompts and parameters to create overall compositions that reveal the limitations and biases of the underlying text-to-image systems.

Robert Halbach focuses his “Attention to Color - The Role of Saliency in Image-Based UI Theming” (↗ am231ho5).

Chunk 5
With the increasing individualization of UI theming, this requires—if the implementation makes sense (also in terms of computing costs)—the automated calculation of a color palette. Halbach presents an overview of technical solutions, which ends with strategies for creating a color palette based on salient objects. In the next step, he refines these by using multimodal models (in this case Gemma 3) to extract the central object of an image, name its primary color and transfer it into RGB values. These serve as the basis for calculating a color palette with attention to saliency.

1: Hui, Y. (2021). *Planetary Thinking*. <https://kunstkritikk.com/planetary-thinking/>

Chunk 6
In “Body as Latent Space: Alternative Compressions in Aesthetic Practice”, (↗ f1ekd6i6) Simon Maris compares AI models and human bodies as compression algorithms, but according to very different criteria and with very different losses: AI reduces vast amounts of images and text into latent spaces, while the body molds a lifetime of experience into tacit knowledge and gesture. Both select and discard: AI preserves statistical regularity at the expense of context and meaning—billions of images are reduced to a few gigabytes through automated pattern recognition based on statistics. The body compresses not purely on frequency and occurrence, but by attributing meaning often skipping the mediocre in favor of the extraordinary.

Chunk 7
Maris exposes the “violence of reduction” inherent in both regimes and urges for a balanced approach: use AI’s latent spaces without sidelining the unique and find ways to leverage the the embodied intelligence as an alternative source of creativity.

In his artistic research, Ting-Chun Liu examines the materiality of current generative AI systems.

Chunk 8
In “On the Materiality of AI”, (↗ gpm91yc8) he explores the globally entangled and geopolitically problematic situation of semiconductor manufacturing, with a particular focus on Taiwan's role as home to the Taiwan Semiconductor Manufacturing Company (TSMC). TSMC's role in the production of modern graphics cards is irreplaceable. The graphics cards, particularly those produced by NVIDIA, are vital for training and operating generative AI systems. Liu illustrates the significance of the materiality of this technology and the global complexity of these production chains in relation to Taiwan's sovereignty efforts vis-à-vis China.

In “un/learn the process”, (↗ 59o29bm4) Francesco Scheffczyk, Paul Eßer and Lars Hembacher describe the concepts and design decisions behind the design of the three-part publication series *un/learn ai*, of which this publication is the final part.

Chunk 9
The authors focus on algorithmic methods of knowledge representation and synthesis.

Chunk 10
Sections of the individual texts, known as *chunks*, are semantically organized using AI and linked to each other by hypertext-like references alongside the texts. The resulting knowledge geometry enables readers to perceive the texts in this publication in a non-linear fashion and to link them together in new ways.

Julia-Jasmin Bold presents “26 strategies to consider when writing with AI” in the form of a wr(AI)ting card deck (↗ c7cors5e). For 8 phases of writing, the deck offers various techniques and inspirations on how AI can support the writing process, e.g. through inspiration, motivation and organization.

Chunk 11
What many cards have in common is that AI simulates different external perspectives in order to enter into a dialog with them. And in case of doubt, the “Hit the AI Panic Button!” card can help.

“The most radical act in AI education might be teaching when not to use it.” Simon Maris' theses “Notes Toward AI Un/Learning” (↗ p3d425b5) are also aimed at a critical and more conscious approach to AI, but also to technology in general. In concise and challenging statements, he presents observations and experiences from the integration of AI tools in teaching and beyond. “Situated knowledge”, “embodied practice”, “everything that doesn't compress into parameters” and “taste” are central aspects of teaching in the sense of un/learning AI.

AI systems open up new possibility spaces, but at the same time they also define and limit what is possible and what we do with them. The resulting synthetic spaces and their often ambivalent or even contradictory relationship to the world require a multi-perspective examination. Perspectives from art and design in particular can make a contribution to technological developments and their integration into our living environment and society.

Chunk 12
Tools and their application are negotiable.

The texts and experiments show diverse approaches and perspectives in regard to AI. Relationships between thing and representation, image and text, materiality, immateriality and

corporeality are negotiated, as are the changes in knowledge production and perception brought about by technology. The explorations from art and design benefit from the fact that they operate in the sensually perceptible and make AI tangible. In addition to reflective strategies and approaches, the texts also offer suggestions for practical examination.

Chunk 13

The design of technologies is not limited to systems and objects, but continues in our interaction with them.

Chunk References

Chunk 3	II. p. 114, Chunk 1: Artefact Hunt
Chunk 10	III. p. 5, Chunk 16: Editor`s Note II. p. 5, Chunk 2: KI-Features dieser Ausgabe I. p. 2, Chunk 5: Vorwort
Chunk 11	III. p. 36, Chunk 19: Building AI Intuition – Four… III. p. 169, Chunk 2: Der Raum kann wahrgenommen… I. p. 175, Chunk 2: Vorwort KI & Gestaltung
Chunk 12	II. p. 4, Chunk 3: Vorwort I. p. 189, Chunk 3: Exploring Tools III. p. 5, Chunk 13: Editor`s Note

↗ zf6iham6

I Think He's Got a Helicopter in His Pocket: The AI Chewing Gum

Thomas Hawranke,
Lasse Scherffig

Chunk 1

The Girl Chewing Gum

In the first 30 seconds of the 1976 avantgarde movie *The Girl Chewing Gum* by British artist John Smith we see and hear different things: The image depicts the trailer of a lorry that, after some seconds, begins to move out of the frame to the left. Meanwhile, pigeons fly by and a girl in a coat walks from left to right.

Chunk 2

Before the lorry vanishes, more people cross the screen from and to different directions. On the audio layer we hear the constant ringing of an alarm bell drowning out the general ambience of a city. The voice of a director is audible who appears to give stage directions: "Slowly move the trailer to the left … and I want the little girl to run across, now. Hold that trailer there … Now move the trailer off".[1] There is obviously something off with the voice directing the image: Why does the director give orders to the lorry and the girl but doesn't direct the other people in the scene? In a highly controlled situation like the film set this feels somehow odd.

Chunk 3

The movie continues in the same way, but the voice that gives directions and the things that act and react to it become more and more separated. In the emblematic shot of a clocktower, where the camera pans up to frame the clock in a somewhat steady shot, the director seems to give stage directions to the pointer of the clock, which seems both absurd and technologically impossible. Later on, a shot with very similar camera movement comes up, and the voice talks in detail about the plate glass manufacturer located in the building of the clock. This includes a lengthy comment on the manufacturer's signage which the narrator apparently misreads as Greek, but actually the signage is seen

1: John Smith. (2015). *Original script for the voiceover to the girl chewing gum, tate research publication.* <https://www.tate.org.uk/research/in-focus/the-girl-chewing-gum-john-smith/original-voiceover-script>

from behind and thus “upside-down and back to front English”.[2] This time, the same scene as shown before does not need stage directions to the camera or the depicted things in front of it at all.

Chunk 4

It becomes obvious that the voice we can hear is a voice-over and not the directing voice. As the film proceeds, Smith drops all pretense of directing the film: While still showing the same street scene, the voice-over describes a rural scene including a golf course. The voice speculates about one of the golfers: “I think he’s got a helicopter in his pocket.”[3] As the voice returns to the street scene in front of the camera, it accuses a random passer-by of having robbed a post office. Just as we hear the voice saying, “the burglar alarm is still ringing,” we notice the alarm cannot be heard anymore, only to return later.

Chunk 5

At the end of the film, the camera reveals that the director is at a completely different location, standing on a meadow with trees, fields, cows and electricity pylons (we can still hear the alarm, though).

The interplay of image and sound is an illusion that demystifies the scenes we just saw. Smith’s work uses the technique of voice-over from documentary and film noir to question authenticity, authority and authorship in mainstream narrative cinema.[4] By doing so, the director becomes the narrator. On a larger scale, Smith’s film is a critical commentary on filmmaking in general, problematizing the relationship between narrative and the moving image, the tension between what is said and seen, and the fact that words never describe or prescribe images in full. In 2015, Tate Modern published the script written by Smith and beside the already mentioned interplay of sound and image, the script reveals yet another layer of inconsistency: that between written and spoken word.

Chunk 6

In the beginning of the movie, script and spoken word correspond to each other, whereas later-on both increasingly disconnect.

Chunk 7

After about eight minutes, the narrator tells us: “I am shouting into a microphone on the edge of a field near Ledgemoor Heath, around 15 miles from the building you are looking at.” The corresponding passage in the written script reads: “I am speaking into a megaphone, which I am pointing at a microphone, on the outskirts of Epping Forest — about ten miles from the building you are looking at.”[5] Between both spoken and written places lie about 160 miles of British landscape. This inconsistency let us wonder where the artist really stood and where the street scene was actually filmed.

Chunk 8

The AI Chewing Gum

In the project *The AI Chewing Gum*, we reenacted Smith’s classic using generative artificial intelligence (AI). Using a text-to-video system,[6] we generate moving images according to the prompts from Smith’s script.

Chunk 9

The system was state-of-the-art when we used it, but at the time of writing it is already outdated. Line by line the script was fed into the system while we used still images from the movie to initialize individual scenes visually. In addition, we recreated the camera movement of the original film using the camera control features of the system. John Smith’s voice was re-synthesized using AI technology as well.

Paradoxically, while trying to follow the script’s directions, the generative systems did both: they stayed truer to the prompts than Smith’s film itself while simultaneously diverging farther from it. Within the limits of their generative potential,

2: John Smith. (2015). *Original script for the voiceover to the girl chewing gum, tate research publication.* <https://www.tate.org.uk/research/in-focus/the-girl-chewing-gum-john-smith/original-voiceover-script>

3: John Smith. (2015). *Original script for the voiceover to the girl chewing gum, tate research publication.* <https://www.tate.org.uk/research/in-focus/the-girl-chewing-gum-john-smith/original-voiceover-script>

4: Quigley, P. (2012). The man giving orders: The girl chewing gum. *Short Film Studies*, *2*(2), 151–154. <https://doi.org/10.1386/sfs.2.2.151_1>

5: John Smith. (2015). *Original script for the voiceover to the girl chewing gum, tate research publication.* <https://www.tate.org.uk/research/in-focus/the-girl-chewing-gum-john-smith/original-voiceover-script>

6: Runway Research. (2023). *Gen-2: The next step forward for generative AI: A multimodal AI system that can generate novel videos with text, images or video clips.* <https://research.runwayml.com/gen2>

they failed to follow the prompted instructions: The result is that people morph through the frames of the video instead of exhibiting anatomically or physically correct motion. The content of the prompts sometimes appears on the images, sometimes it does not — as if the English prompts quite literally were processed upside down and back to front. More generally, while each scene starts with the look of the original film and follows the prompted script from there, the scenes end up with different aesthetics and visual language altogether — moving from black-and-white to color images, and from specific street scenes to very generic. Starting each scene with one frame of the original movie hence was necessary, because the generative process constantly pushed the resulting images towards a contemporary cinematic aesthetic.

The simulated voice of the resulting generated film likewise follows the script accurately, but by doing so erases the difference between written and spoken text that was essential to the original movie. Of course, the burglar alarm is erased as well. To the system, it appeared to be a noise distortion in need of correction. Instead of questioning authenticity, authority and authorship in mainstream narrative cinema, the generative systems employed in the project raised questions about their own status as tools for filmmaking and as "world simulators."[7] They revealed that video generation systems do not possess an understanding of the physical world or, for that matter, film. They showed how prompting is not about controlling generative processes but about loosely guiding a process of stepwise optimization according to statistical plausibility, revealing more about the training data than about any creative intention.

Chunk 10

This process ultimately yields standardized images from a "flat world,"[8] that has been flattened towards an ahistorical accumulation of styles.

II. p. 42, Chunk 19: The Cognitive Style of ChatGPT
II. p. 106, Chunk 11: Infinite Boredom
I. p. 186, Chunk 8: Algorithmic ~~Culture~~…

Chunk 11

This article is separately published as 10.25358/openscience-13024.

Chunk 12

7: OpenAI. (2024). *Video generation models as world simulators.* <https://openai.com/index/video-generation-models-as-world-simulators/>

8: Meyer, R. (2024). It`s a flat world. The synthetic realities of sora. *Rrrreflect. Journal of Integrated Design Research, Special Issue 1: Hidden Layers.* <https://doi.org/10.57684/COS-1267>

IN FOCUS

Original script for the voiceover to *The Girl Chewing Gum*

JOHN SMITH

The Girl Chewing Gum 1976 by John Smith
The Girl Chewing Gum 1976 by John Smith
Original script for the voiceover to *The Girl Chewing Gum*
Context and Conception
A 'New Face' at the Co-op
(Not) the Voice of God
Dalston, Then and Now: An Interview with John Smith
A Film in the Age of Video
Acknowledgements

This is a transcription of the original typescript for the voiceover written by Smith as he edited the film, reproduced courtesy of the artist. Its content differs slightly from the final voiceover used by Smith.

Directions –

Slowly move the trailer to the left….and I want the little girl to run across, now. Hold that trailer there……Now move the trailer off.

Right – now I want the old man with white hair and glasses to cross the road – come on, quickly….. look this way, now walk off to the left. OK, fine.

Now let's have the man in the peaked cap….put the cigarette in your mouth……good……..and I want the two girls to come in from the right talking to each other.

Now I want the Jamaican family, father first – look in front of you……now the mother and the two boys……and I want the smaller boy to point to the right, and now cast a glance behind him.

Now I want the bus ticket to blow across the pavement.

Fine – now I want the girl with the bag under her arm to walk to the right. Now let's have the man rubbing his eye……good.

Right. Now I want everything to sink slowly down as the five boys come by……hold it……and I want the clock to move jerkily towards me….stop.
Now I want the long hand to move at the rate of one revolution every hour, and the short hand to move at the rate of one revolution every twelve hours……fine.
Now two pigeons fly across, and everything comes up again until the girl chewing gum walks across from the left.

OK. Now – the van, the woman, and the boy…the man with the bag of chips, the car roof, the whole car, the old man, the car, the car and the boy, the boy, the motorcyclist and the lorry.

Now when the woman at the window looks in this direction, I want everything to move a bit further away.

Right – let's have the man in the white boiler suit coming in from the right……stop at the lamp-post and fold your arms….now look around you……..
walk back to the left again, and look left and right as you cross the road…..
Now I want the man reading exchange and mart to come in from the right.

Now a man comes by and bites his nails, two pigeons fly past from right to left and two boys run past from left to right…..the woman at the window looks this way and then goes on talking.

In a second I want four boys to come from the left, and I want one of them to look this way and flap his arms up and down. I'd like the boys to remain in more or less the same position, while everything else moves to the left and goes away a bit at the same time……Got that?…..OK……go.

Lift up your arm – now bring it down.

Now I want the man with the turban and briefcase to come in from the right, and now the van marked 'University of London, Senate House, London WC1'.

In the cinema queue, I want to see a boy and his mother. The boy will be about eight years old, and his mother about thirty-two. They will both have collar-length hair, his dark and hers fair – she will be wearing a suedette coat with a white imitation fur collar. They will be talking to each other and looking around them. The boy will look at his watch and yawn, then look at it again and say something to his mother about the time.

Now I want the man coming in from the left in the grey peaked cap to put on his glasses. The three children eating chips, the French woman, the window-cleaner in his van, the greengrocer, the pigeon, the red jaguar, the taxi, the van, the man, the negro with the briefcase and the newspaper, the woman firmly gripping the hands of her two young sons. They stop, and now they cross the road.

The man in the grey peaked cap takes off his glasses again and puts them in his breast pocket. The man with the walking stick is going home. The dentist continues on his way to the bank, and the two naughty boys appear from behind Steele's and cautiously cross the road.

Steele's the plate glass manufacturers is situated in an area with a high immigrant population, predominantly West Indians and Greeks. Outside the building, above its main doors, is a board advertising vacancies at the firm. Until recently, I thought that jobs were advertised on the board in two languages – English and Greek. However, a few weeks ago, I studied the board carefully, and realised that I had been wrong. Each vacancy has its own slat in the board – the words (glass-cutters, bevellers, bench fitters etc.) are cut out of these slats, which are made of perspex. When a vacancy no longer exists, the slat advertising it, which has a central pivot at each end, is swivelled around. The words that I had believed to be Greek were in fact upside-down and back to front English. Steele's also has another interesting feature. Along each wall of the building, there are eight large doors that are kept wide open throughout the year.

I am speaking into a megaphone, which I am pointing at a microphone, on the outskirts of Epping Forest – about ten miles from the building you are looking at. The traffic noise will now fade down. I am standing on the edge of an enormous golf course, and there are trees behind me and to my left. In the distance I can see three golfers. One of them is wearing a blue jumper, and the other two are both wearing red. They are all men. The shortest of the three is standing by a trolley, which is laden with their clubs. The one with the blue jumper, who seems much older than the other two, is looking this way. Now the short one has put a ball on the ground and is taking a club from the trolley. I think he's got a helicopter in his pocket. In a tree about fifteen yards away, I can see a large blackbird with a wingspan of about nine feet.

This young man has just robbed the local post office and is attempting to appear inconspicuous. He is trying to remain calm, but his hand is sweating as he grips the butt of the revolver in his raincoat pocket even harder. He is wondering whether the woman at the window would recognise him if she saw him again. The burglar alarm is still ringing.

Now disappear.

How to cite

John Smith, 'Original script for the voiceover to *The Girl Chewing Gum*', in *"The Girl Chewing Gum" 1976 by John Smith*, Tate Research Publication, 2015, https://www.tate.org.uk/research/in-focus/the-girl-chewing-gum-john-smith/original-voiceover-script, accessed 27 May 2025.

© The Board of Trustees of the Tate Gallery, 2025
All rights reserved

Screenshot of the Tate Research Publication of the original script of "The Girl Chewing Gum" by John Smith.



Slowly move the trailer to the left …



… and everything comes up again until the girl chewing gum walks across …

Chunk 13



… and I want the two girls to come in from the right talking to each other.



… I want everything to move a bit further away.

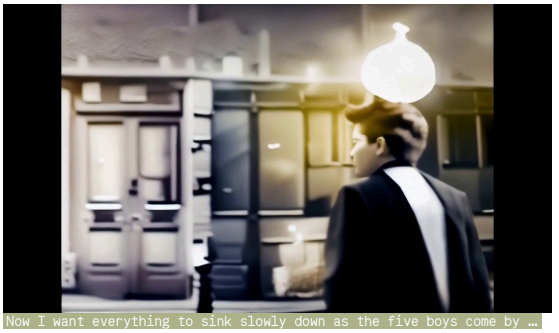

Now I want everything to sink slowly down as the five boys come by …

Chunk 14



… two pigeons fly past from right to left and two boys run past from left to right.



… and I want the clock to move jerkily towards me …

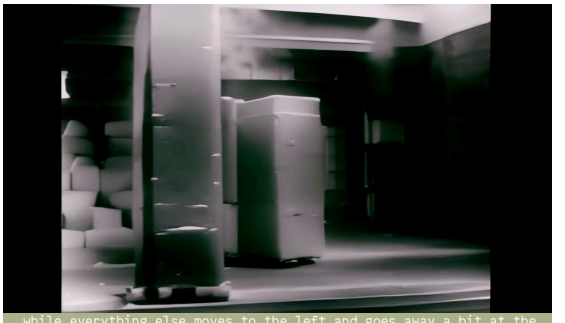

… while everything else moves to the left and goes away a bit at the same time.



… and now the van marked `University of London, Senate House, London WC1'.

Chunk 15



They will be talking to each other and looking around them.



… the French woman, …

Chunk 16



The man with the walking stick is going home.



Each vacancy has its own slat in the board …

Chunk 17



Along each wall of the building, there are eight large doors …



I think he's got a helicopter in his pocket.



I can see a large blackbird with a wingspan of about nine feet.

↗ 14as24mw

An Incredibly Average Face – Approaching 31,200 Images with Human Eyes, Sorting Systems, and Collaging

Jana Hartmann

I. p. 241, Chunk 6: Indirektes Design
II. p. 272, Chunk 8: Hidden Layers 24

I. p. 241, Chunk 7: Indirektes Design

Chunk 1



Figure 1: a portrait of a incredibly average face, {style}; collage.
© Jana Hartmann

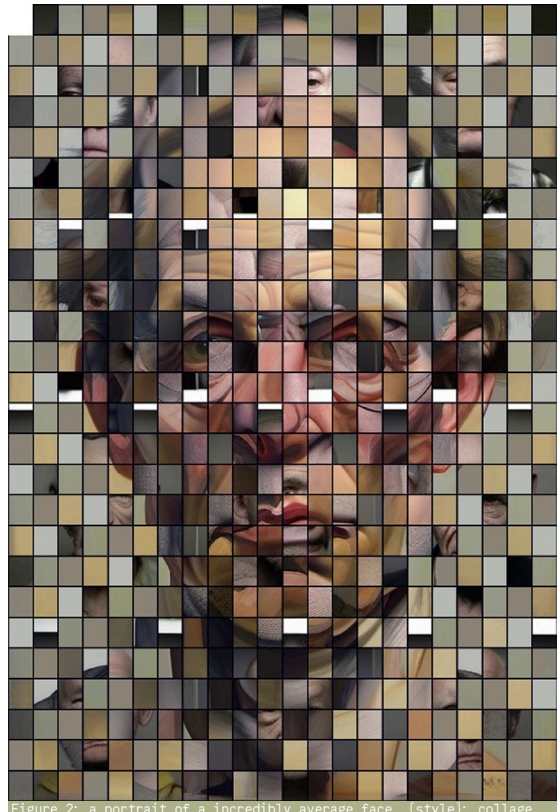

Figure 2: a portrait of a incredibly average face, {style}; collage.
© Jana Hartmann

Figures 1 and 2 are collages of five images generated from the same seed `264523154666154` and almost the same prompt. What sets them apart is the ending of the prompt that describes the style.

Chunk 2

All were prompted: "`A portrait of a[n] incredibly average face, …`".

Chunk 3

Chunk 4

These collages were created as part of my practical research on how beauty, or more specifically facial beauty, is defined in AI image generation. This research is based on the broader question of how beauty is socially constructed, defined, and deployed[1] and looks at AI image generation as an example where these definitions become explicit.

It builds on Hito Steyerl's term *mean images* that refers to AI-generated images as averages that encapsule societal norms.

Chunk 5

These images are averages/means because of the way they generate pixels. An AI-generated image of a cat does not depict a cat that exists in the world, like a photograph would. Rather, it can be described as an average of all the depictions of cats in the AI training data set. Steyerl argues that the AI-generated mean image does not represent a real-world average but rather tends to

1: Colebrook, C. (2006). Introduction. *Feminist Theory*, *7*(2), 131–142. <https://doi.org/10.1177/1464700106064404>

show “extreme and unattainable outliers—for example, anorexic body ideals”.[2] AI-generated images are furthermore *mean* because of the exploitative work conditions under which machine learning models are trained, and because they are based in part on discriminatory labeling of people.[3]

To investigate a statistical medium with statistical methods,[4] 31,200 images were generated from an automated *ComfyUI* script.

Chunk 6

Based on findings from generating small batch sizes, I created a prompting system with four variables that generates prompts in the form `a portrait of a {beauty descriptor} {object}, {style}`.

Chunk 7

Values for the variables *beauty*, *object*, and *style* were stored in text files.

Chunk 8

40 seeds were chosen randomly. The script then, for every seed, prompted images from every possible combination of variable values (Figure 3). This system was set up to make the resulting images comparable along multiple parameters – because a single person cannot simply analyze 31,200 images.

Chunk 9

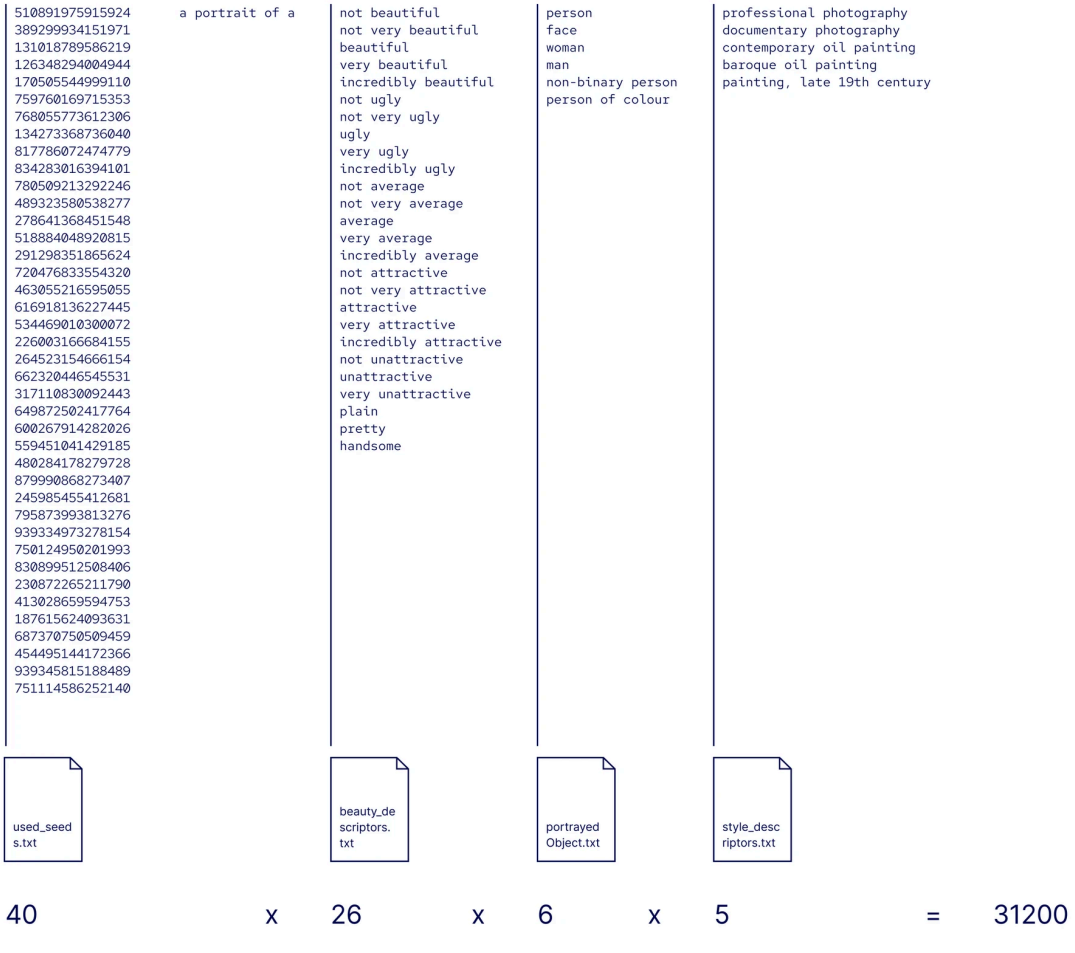

Figure 3: Diagram of prompting system. © Jana Hartmann

2: Steyerl, H. (2023). Mean Images. *New Left Review*, *140/141*, 89. <https://newleftreview.org/issues/ii140/articles/hito-steyerl-mean-images>

3: Steyerl, H. (2023). Mean Images. *New Left Review*, *140/141*, 82–96. <https://newleftreview.org/issues/ii140/articles/hito-steyerl-mean-images>

4: Offert, F., & Dhaliwal, R. S. (2024). *The Method of Critical AI Studies, A Propaedeutic*. arXiv. <https://doi.org/10.48550/ARXIV.2411.18833>

I. p. 201, Chunk 3: Creative Ownership and Creative…
II. p. 115, Chunk 2: Artefact Hunt
II. p. 69, Chunk 23: Editorial

Chunk 10

Dealing with large amounts of images is usually a problem discussed in the context of AI training data.

Chunk 11

The scale of these data sets, e.g. *LAION-5B* that is a part of the training data for *Stable Diffusion*, makes human curation nearly impossible.[5] Here however, I am studying image-generating AI as a system of possible images in Frieder Nake's tradition of Generative Aesthetics[6] that can output endless amounts, which then again leads to the question of: How is it possible for humans to curate or research extremely large amounts of images?

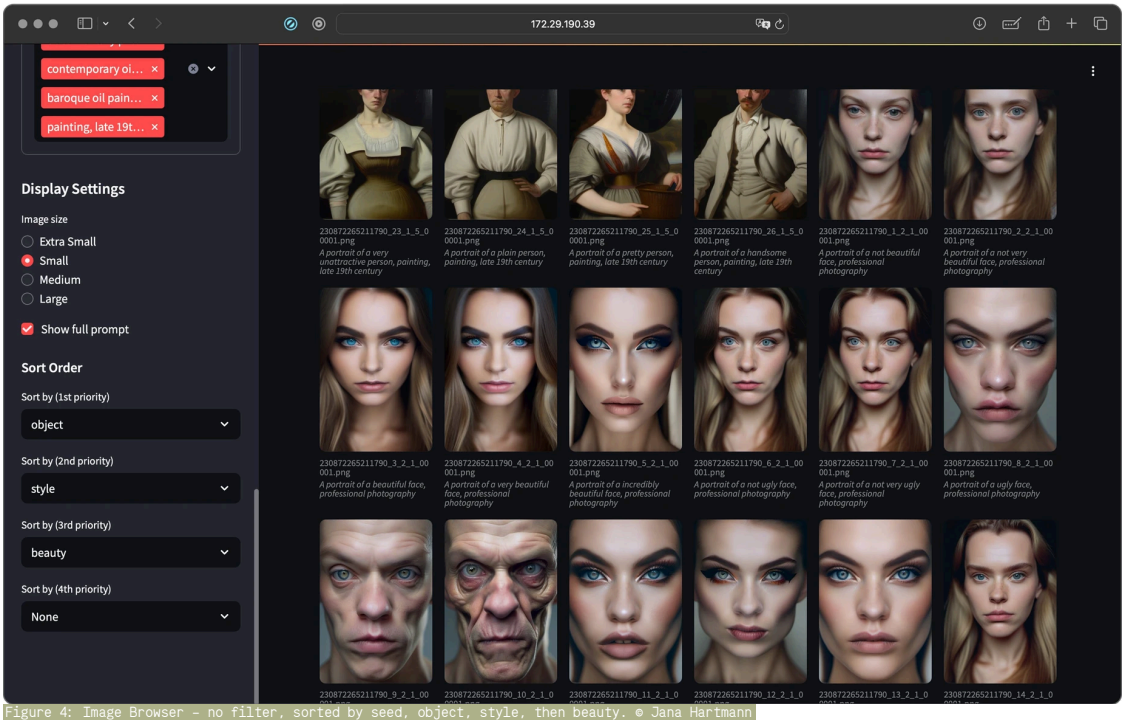

Figure 4: Image Browser – no filter, sorted by seed, object, style, then beauty. © Jana Hartmann

Chunk 12

5: Buschek, C., & Thorp, J. (2024). Models All The Way Down. In *Knowing Machines*. <https://knowingmachines.org/models-all-the-way#section5>

6: Offert, F. (2019). The Past, Present, and Future of AI Art. *The Gradient*. <https://thegradient.pub/the-past-present-and-future-of-ai-art/>

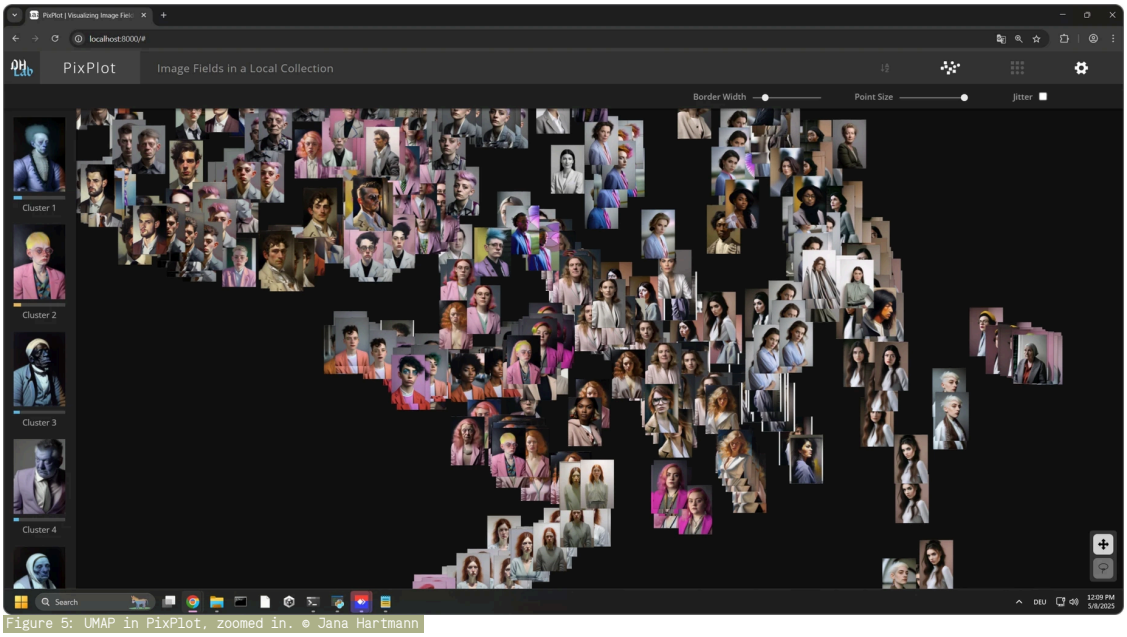

Figure 5: UMAP in PixPlot, zoomed in. © Jana Hartmann

Chunk 13

Figures 4 and 5 illustrate my first steps to looking at the images.

Chunk 14

First, I created a custom file browser for filtering and sorting the images according to the four variables (Figure 4). Then I generated a *UMAP* visualization with *PixPlot*, a tool by *The Digital Humanities Lab* at *Yale University* that clusters images according to pixel similarity (Figure 5).

In the file browser, I quickly began to filter and sort into matrixes where, for example, rows would represent a style and each column had the same beauty descriptor.

Chunk 15

From there, I decided to create a 3D matrix. Because my images differ from each other in four variables, a 3D matrix had to have a filter applied to one of them. While any variable could have been filtered, as of now I have only filtered by seed. Figures 6 and 7 show matrixes of all 780 images created from seed `480284178279728` (Figure 6) and seed `879990868273407` (Figure 7). Each variable of the prompt system then represents an axis in the 3D space.

Chunk 16



Figure 6: a portrait of a {beauty descriptor} {object}, {style} – Seed: 480284178279728; displayed in 3D matrix. © Jana Hartmann

Chunk 17


Figure 7: a portrait of a {beauty descriptor} {object}, {style} – Seed: 879990868273407; displayed in 3D matrix. © Jana Hartmann

Furthermore, I explored ways of comparing not only side by side but rather on top of or within one another, as well as comparing through time.

Chunk 18

Therefore, I created videos that show one image at a time in short succession.

Chunk 19

In TouchDesigner, I experimented with layering images using different blend modes and created systems of algorithmic collaging.

While the matrix system applies a filter on one category, for the collages I filtered all but one.

Chunk 20

This means the collages, as they are now, allow comparisons across one variable.

Comparisons of Descriptions of Beauty

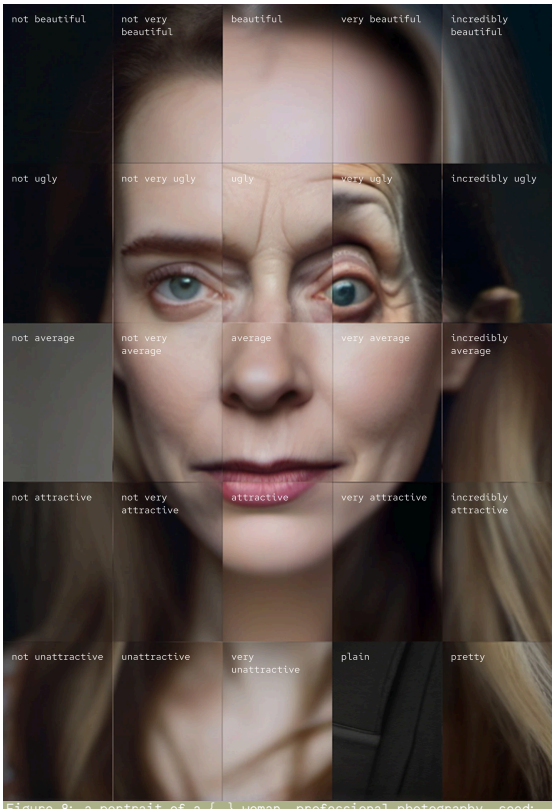

Figure 8: a portrait of a […] woman, professional photography, seed: 230872265211790; collage. © Jana Hartmann

III. p. 168, Chunk 2: Anthrobot

Chunk 21

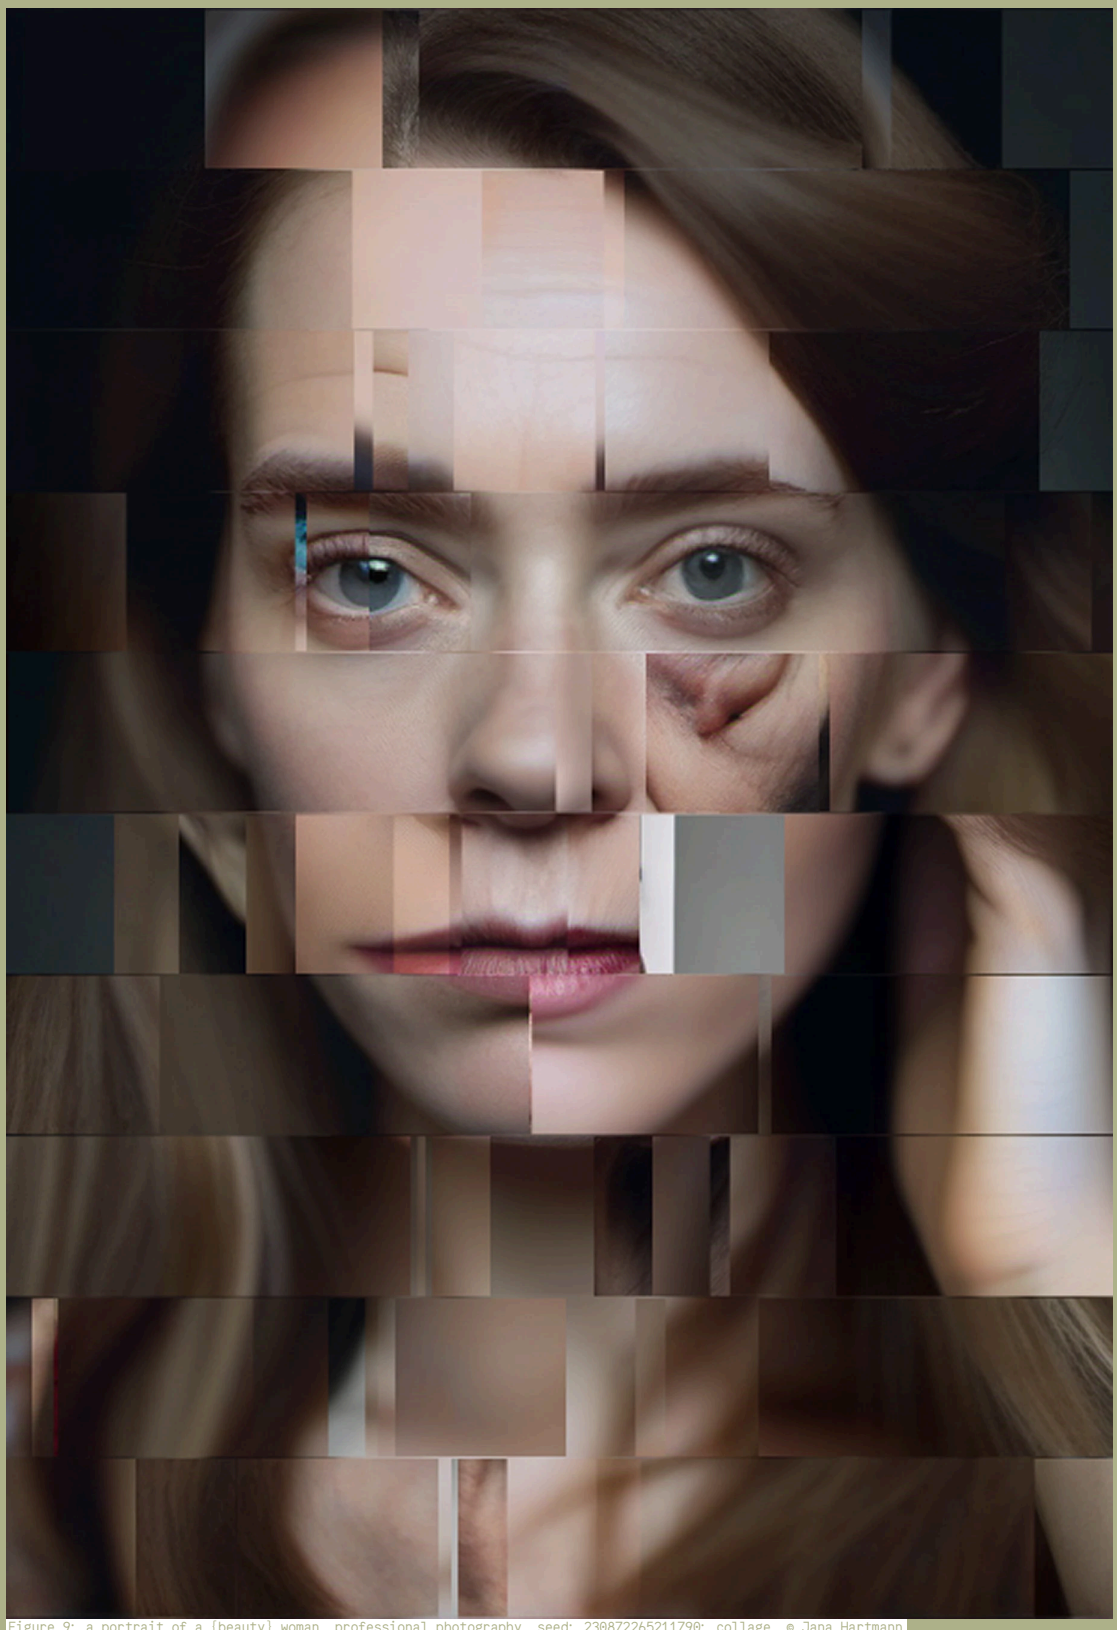

Figure 9: a portrait of a {beauty} woman, professional photography, seed: 230872265211790; collage. © Jana Hartmann

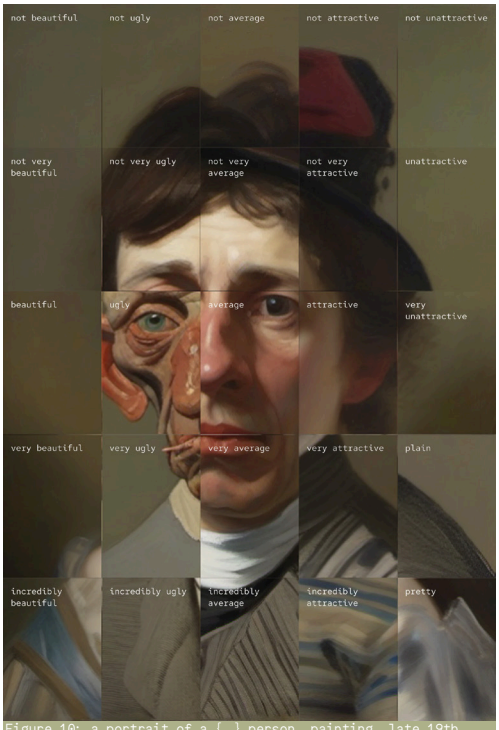

Figure 10: a portrait of a {…} person, painting, late 19th century, seed: 780509213292246; collage. © Jana Hartmann



Figure 11: a portrait of a {beauty} person, painting, late 19th century, seed: 780509213292246; collage. © Jana Hartmann

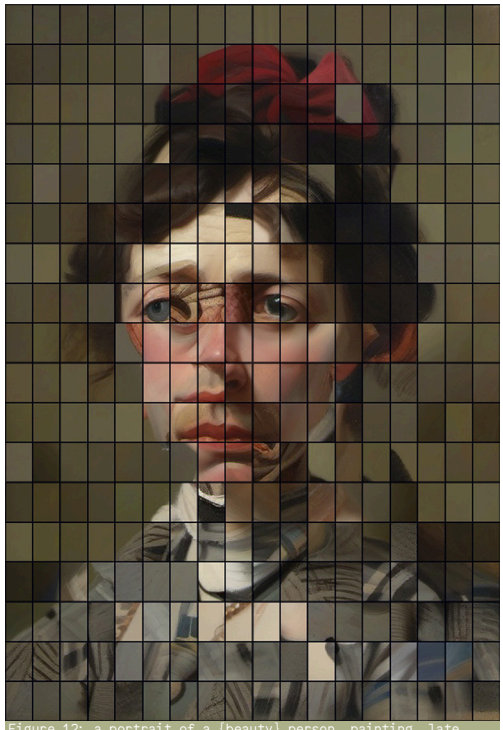

Figure 12: a portrait of a {beauty} person, painting, late 19th century, seed: 780509213292246; collage. © Jana Hartmann

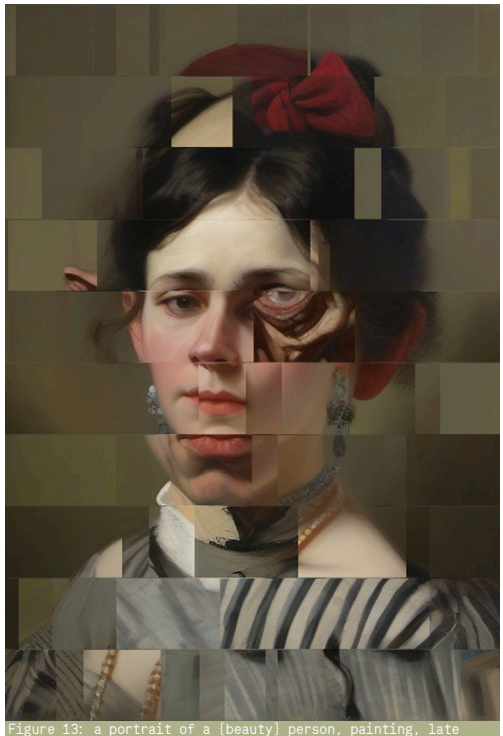

Figure 13: a portrait of a {beauty} person, painting, late 19th century, seed: 780509213292246; collage. © Jana Hartmann

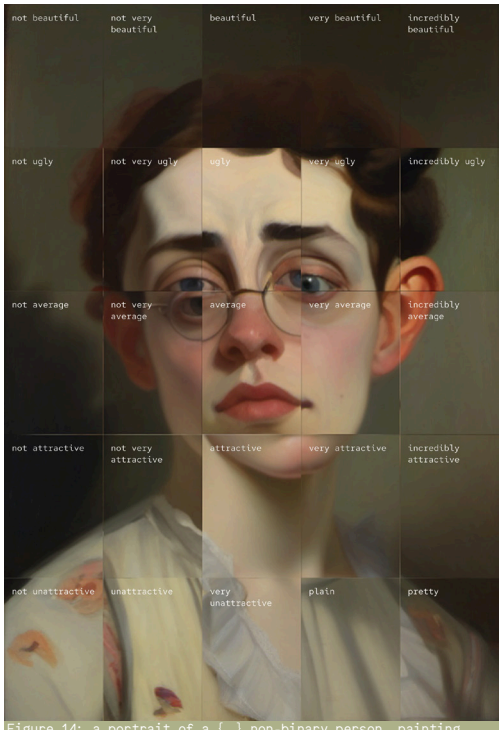

Figure 14: a portrait of a {…} non-binary person, painting, late 19th century, seed: 939334973278154; collage. © Jana Hartmann

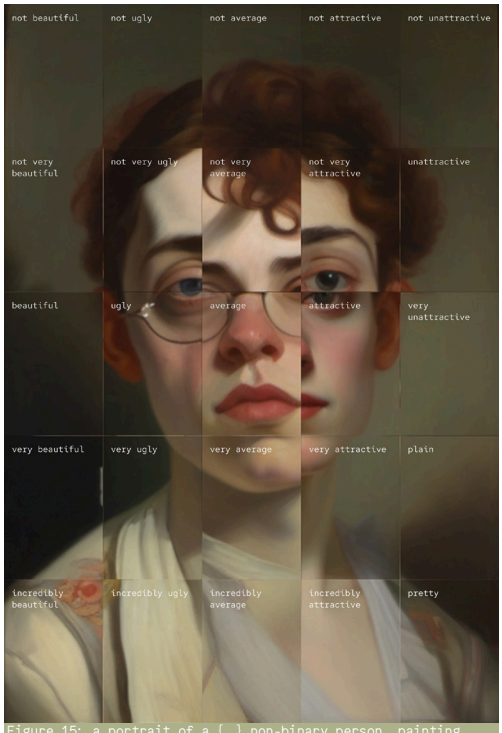

Figure 15: a portrait of a {…} non-binary person, painting, late 19th century, seed: 939334973278154; collage. © Jana Hartmann

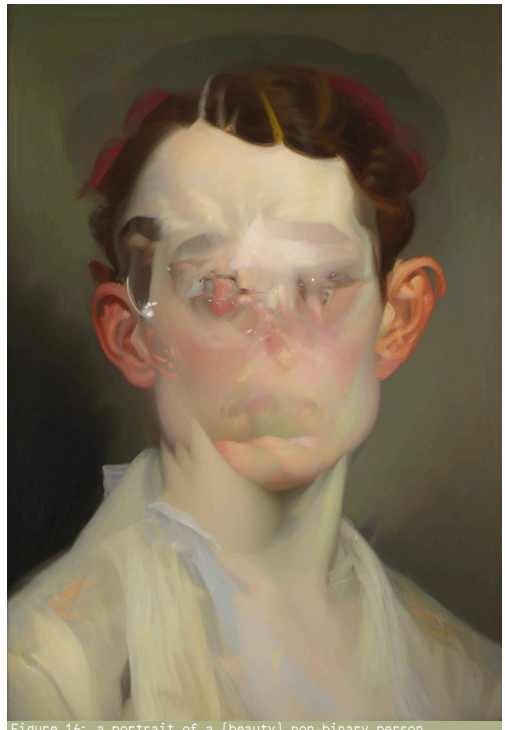

Figure 16: a portrait of a {beauty} non-binary person, painting, late 19th century, seed: 939334973278154; collage. © Jana Hartmann

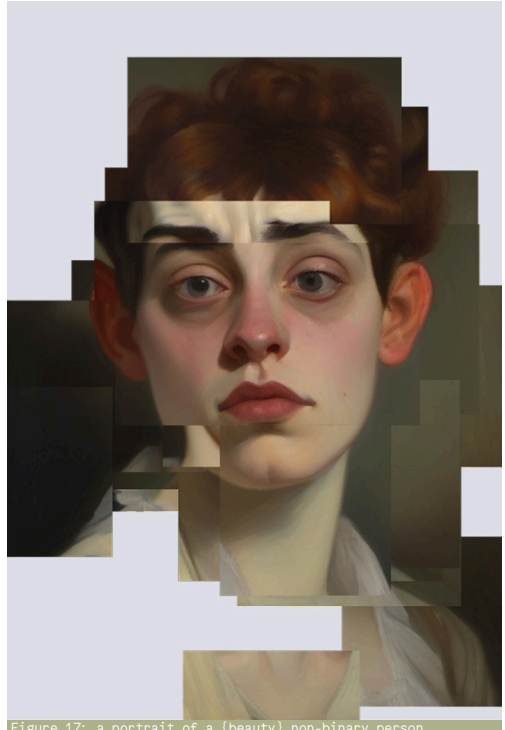

Figure 17: a portrait of a {beauty} non-binary person, painting, late 19th century, seed: 939334973278154; collage. © Jana Hartmann

Comparisons of Seeds

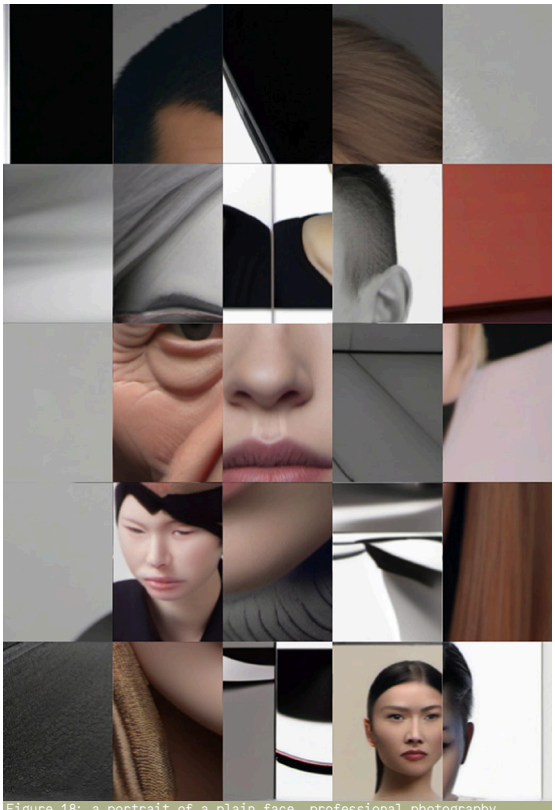

Figure 18: a portrait of a plain face, professional photography, randomly selected seeds; collage. © Jana Hartmann

Chunk_22

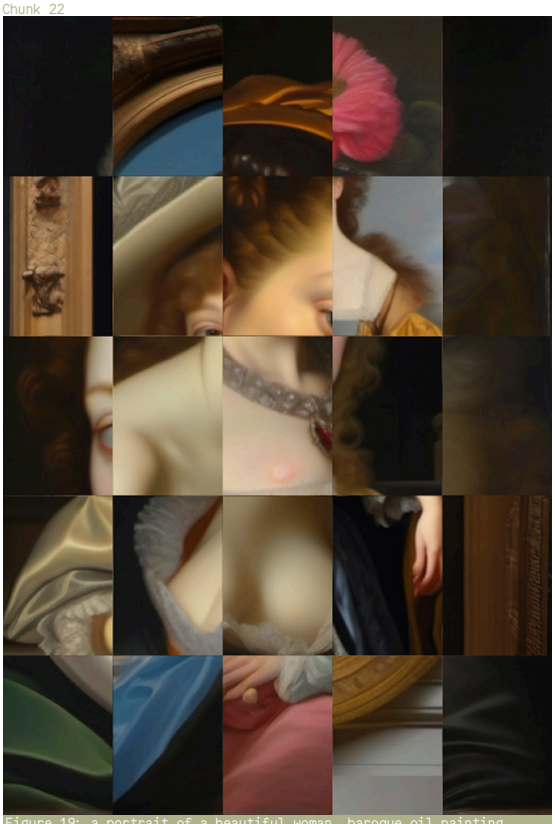

Figure 19: a portrait of a beautiful woman, baroque oil painting, randomly selected seeds; collage. © Jana Hartmann

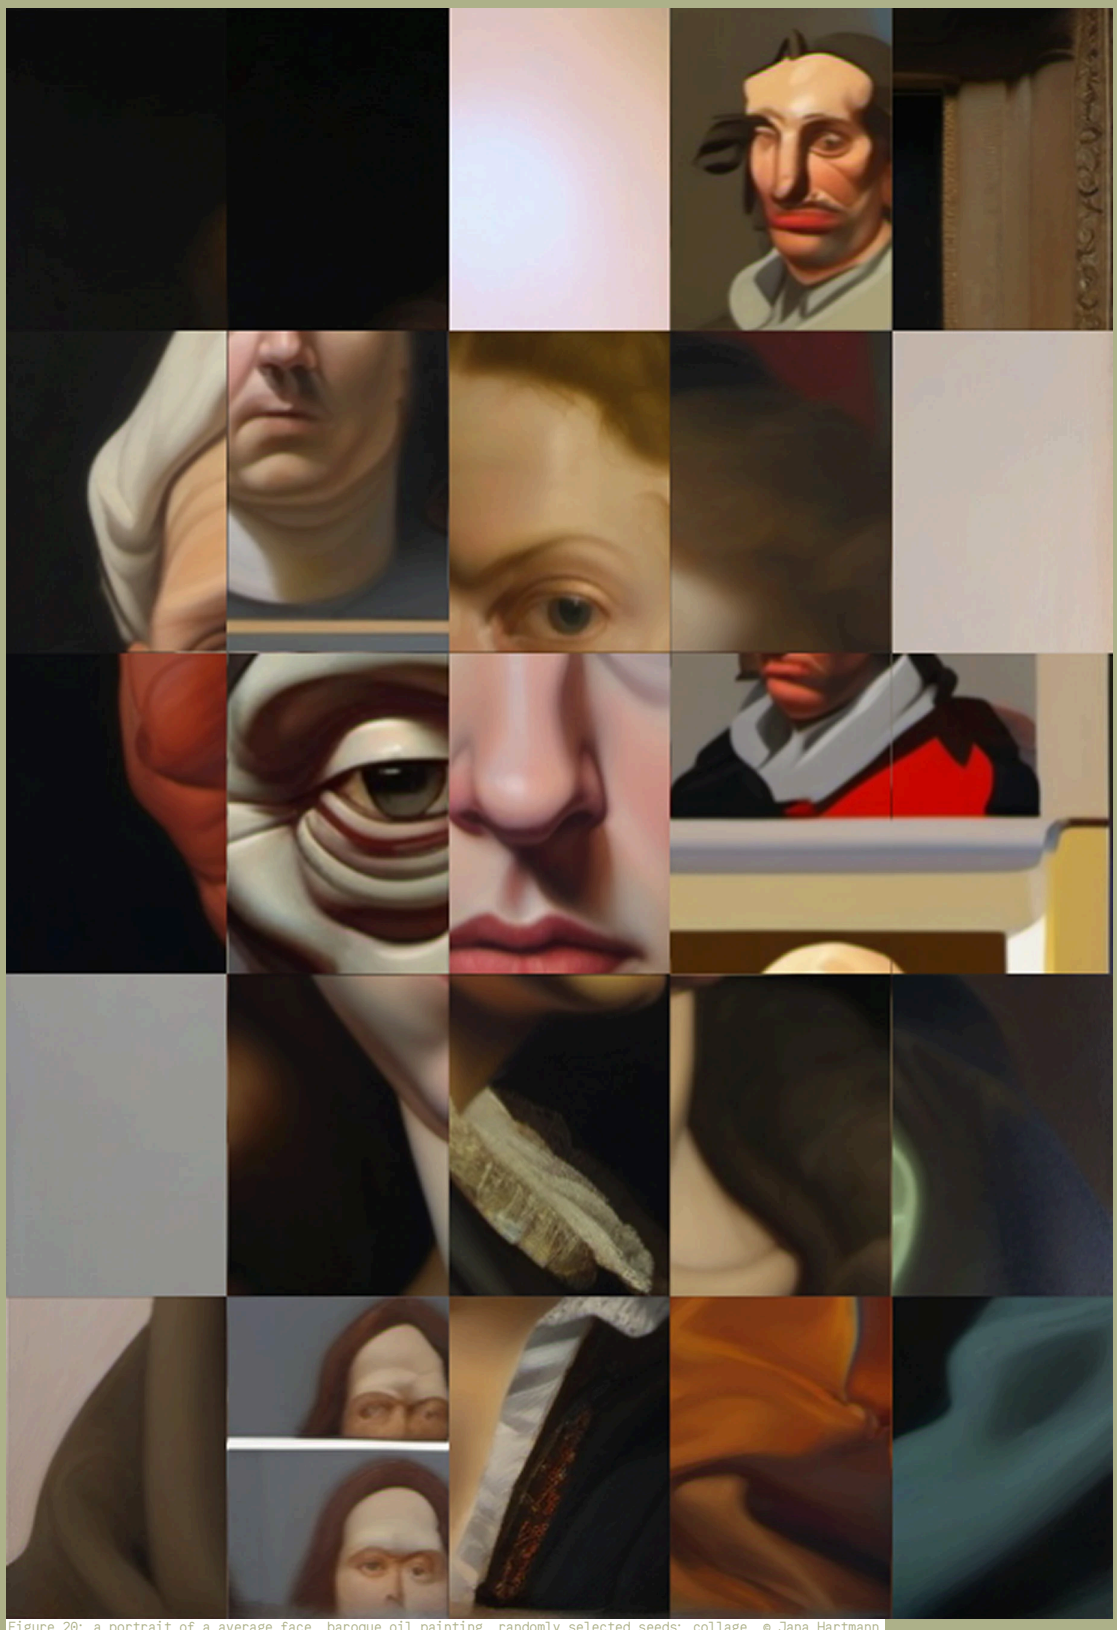

Figure 20: a portrait of a average face, baroque oil painting, randomly selected seeds; collage. © Jana Hartmann



Figure 21: a portrait of a handsome face, painting, late 19th century, randomly selected seeds; collage. © Jana Hartmann

Comparisons of Styles

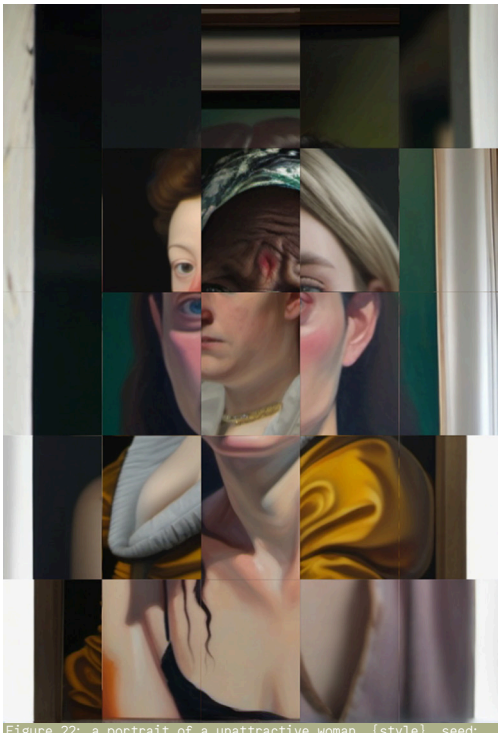

Figure 22: a portrait of a unattractive woman, {style}, seed: 662320446545531; collage. © Jana Hartmann



Figure 24: a portrait of a unattractive woman, {style}, seed: 662320446545531; collage. © Jana Hartmann



Figure 23: a portrait of a unattractive woman, {style}, seed: 662320446545531; collage. © Jana Hartmann

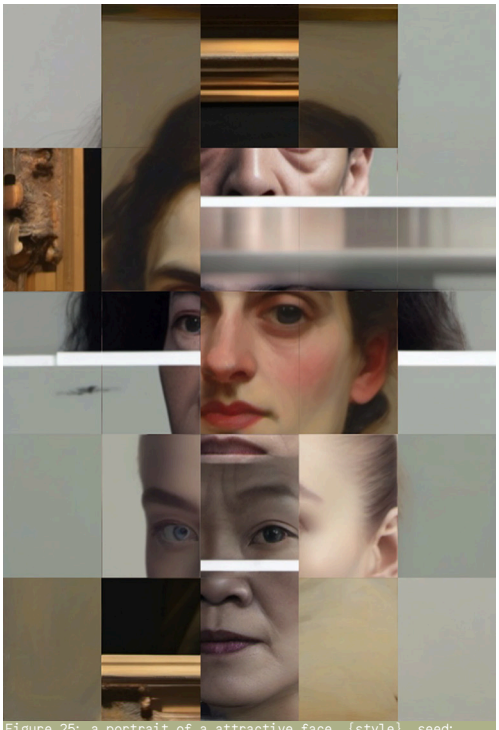

Figure 25: a portrait of a attractive face, {style}, seed: 264523154666154; collage. © Jana Hartmann

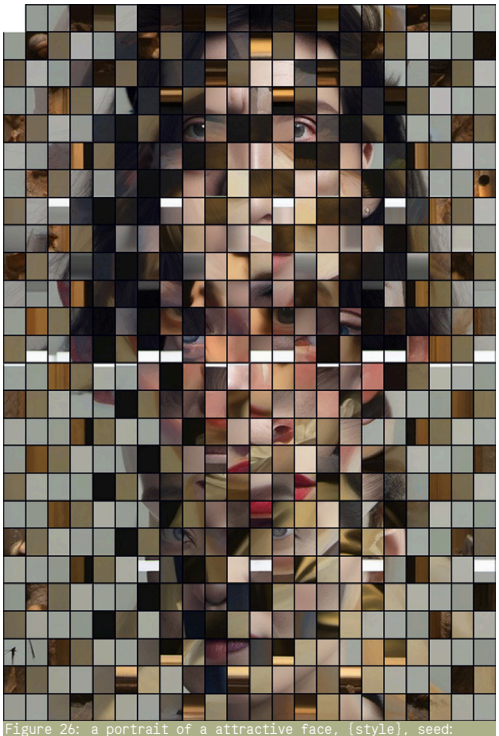

Figure 26: a portrait of a attractive face, {style}, seed: 264523154666154; collage. © Jana Hartmann

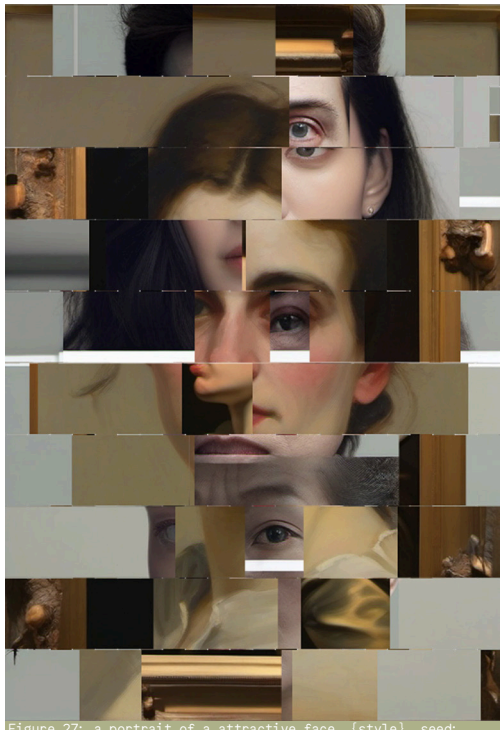

Figure 27: a portrait of a attractive face, {style}, seed: 264523154666154; collage. © Jana Hartmann

The collages across seeds (Figure 18–21) mostly show variance in compositions. When I compare them to each other, they seem to give an insight into colour schemes of a style. This is reminiscent of the UMAP visualization in *PixPlot*, where images with the same style prompt tended to cluster together because colours impact pixel similarity.

In contrast, images created from the same seed are much more similar in composition.

Chunk 23

In collages across styles (Figure 22–27), the switch between collage parts is often more like a change in texture or in tools of creation. The switch between beauty descriptors (Figure 8–17) can also be described along those lines.

Chunk 24

In all styles the difference between `beautiful` and `ugly` seems to rely strongly on youth of skin. In the style `professional photography`, it relies furthermore on styles of digital photography editing, as `beautiful` results in soft and blurry images while `ugly` images show more contrast.

Chunk 25

Amanda Wasielewski observes that generative AI strips historical art styles of their context and uses them more like a (color) filter than content. This way, they become endlessly interchangeable.[7]

Chunk 26

The term filter seems adequate to describe the effect of the style variable as well as of beauty descriptors. Some of the images prompted `a portrait of a {beauty} woman, professional photography` from seed `230872265211790` (Figure 8–9) could be photographs of the same person but with different filters applied.

Overall, the systems I created for looking at 31,200 images genuinely helped me understand these images better.

Chunk 27

Especially the collages lead to new insights into how different parts of the prompt affect the resulting images. In regard the question of how a single person can look at or curate large amounts of images I should mention that I cannot say that I have seen all images.

In under 36 hours, I generated 31,200 images. Over the course of one month I have spent a lot of time with them, and yet I still discover new images every time I open the folder. From a fraction of these 31,200 images, I have generated 786 collages as of the time of writing. I have exported 49 images of the 3D matrix. I created 234 video files of moving collages and recorded 4 videos of the 3D matrix. I have taken (only) 8 screenshots of the UMAP visualization in PixPlot and 32 screenshots of the custom file browser. In total I now have 32,313 new files lying on my computer and in multiple clouds.

II. p. 122, Chunk 5: You Press the Button, They…
II. p. 132, Chunk 2: Flesh Fest

Chunk 28

Selecting the 20 collages presented in this publication took nearly as long as generating all 786.

Roland Meyer describes AI image generation as an inherently wasteful process, because of its vast energy consumption and because a multitude of generated images are discarded immediately. AI-generated images are thus “neither rare nor valuable.”[8] Generating (large amounts of) images inevitably leads to a question of selection or curation and most likely to discarded options that become digital waste.

The images I have not even looked at yet are arguably nothing but digital waste.

Chunk 29

Yet the 31,200 images I generated are still only a tiny fraction of the possible images I *could* have generated.

This article is separately published as 10.25358/openscience-13023.

7: Cetinić, E. (2024). Introduction. *Hertziana Studies in Art History*. <https://doi.org/10.48431/hsah.0301>

8: Meyer, R. (2025). "Platform Realism." AI Image Synthesis and the Rise of Generic Visual Content. *Transbordeur*, *9*, 14. <https://doi.org/10.4000/13dwq>

↗ am231ho5

Attention to Color – The Role of Saliency in Image-Based UI Theming

Robert Halbach

Chunk 1

Curation and automation

The early days of interfaces were shaped by pure functionality—restricted by hardware limitations and primarily guided by the need for contrast and clarity; color use was minimal.[1] With growing technological maturity and capability, as well as ever-changing trends, the role of color shifted toward being a central tool for communicating identity, mood, and context.[2]

Alongside technology, our understanding of how users interact with the Web also matured, leading to a broadly well-performing, yet homogenous interface landscape.[3] However, recent trends toward individuality and personalization on a color level counteract this phenomenon. While personal theming (e.g., light/dark modes or pre-curated color palettes) has been around for quite some time, it usually required manual user-driven configuration.

Today, customization can go far beyond system- or platform-wide settings. Color schemes become highly situational, influenced by the context they exist in.[4] For theming, this growing complexity makes provider-curated color schemes unfeasible, raising the need for automation. A popular starting point to this process is utilizing color seeds extracted from images to dynamically generate color swatches.

1: Lamprogeorgos, A., Pergantis, M., & Giannakoulopoulos, A. (2024). Evolution of Color Implementation in the Web Design of Scientific and Educational Websites. *ICERI Proceedings*, *1*, 7249–7257. <https://doi.org/10.21125/iceri.2024.1746>

2: Löffler, D. (2017). *Color, Metaphor and Culture - Empirical Foundations for User Interface Design*. Universität Würzburg. <https://opus.bibliothek.uni-wuerzburg.de/files/15378/Loeffler_Diana_Color,Metaphor_and_Culture.pdf>

3: Goree, S., Doosti, B., Crandall, D., & Su, N. M. (2021). *Investigating the Homogenization of Web Design: A Mixed-Methods Approach*. 1–14. <https://doi.org/10.1145/3411764.3445156>

4: Google Design. (2025). *Expressive design: Google`s UX research*. <https://design.google/library/expressive-material-design-google-research>

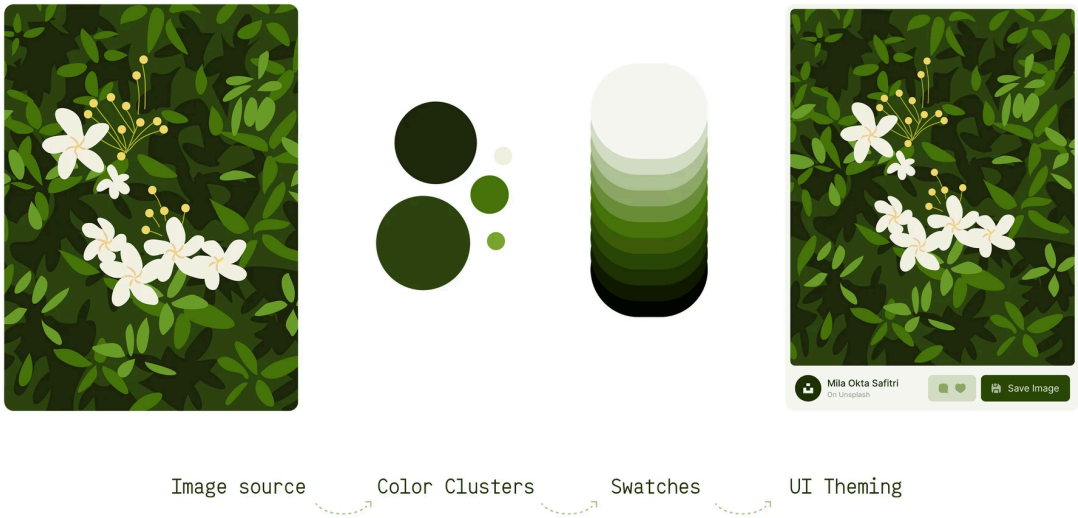

Figure 1: Schematic display of an automated interface theming pipeline. Cluster colors extracted with a basic implementation of a k-means algorithm, with k=5. © Robert Halbach, source image by ↗ Mila Okta Safitri

Chunk 2

However, the seed selection process is not neutral or arbitrary. Extraction systems, no matter their architecture, introduce an opinionated tendency toward the colors picked.

Chunk 3

Often, these systems try to capture the overall mood or atmosphere of an image, favoring ambient tones and dominant hues. Even though there are exceptions, most approaches process images as a whole, ignoring the presence of subjects and salient (visually outstanding) areas as extraction factors.[7] [8] [9] This means that whether themed interfaces inherit their color palette from salient areas depends both on properties such as color, saturation, and the size of those areas, as well as the biased ruleset implemented into the extraction process.

5: Google. (n.d.). *Dynamic color schemes - material 3*. Retrieved May 28, 2025, from <https://m3.material.io/styles/color/dynamic/choosing-a-source>

6: Argyle, A. (2023). *Thinking on ways to solve color palettes*. Chrome for Developers. <https://www.youtube.com/watch?v=6aCsAMgwnjE>

7: Jahanian, A., Vishwanathan, S. V. N., & Allebach, J. P. (2015). *Autonomous color theme extraction from images using saliency*. Purdue University.

8: Chang, Y., & Mukai, N. (2022). Color feature based dominant color extraction. *IEEE Access*, *10*, 93055–93061. <https://doi.org/10.1109/ACCESS.2022.3202632>

9: Weingerl, P. (2024). Automated image-based user interface color theme generation. *Applied Sciences*, *14*, 2850. <https://doi.org/10.3390/app14072850>

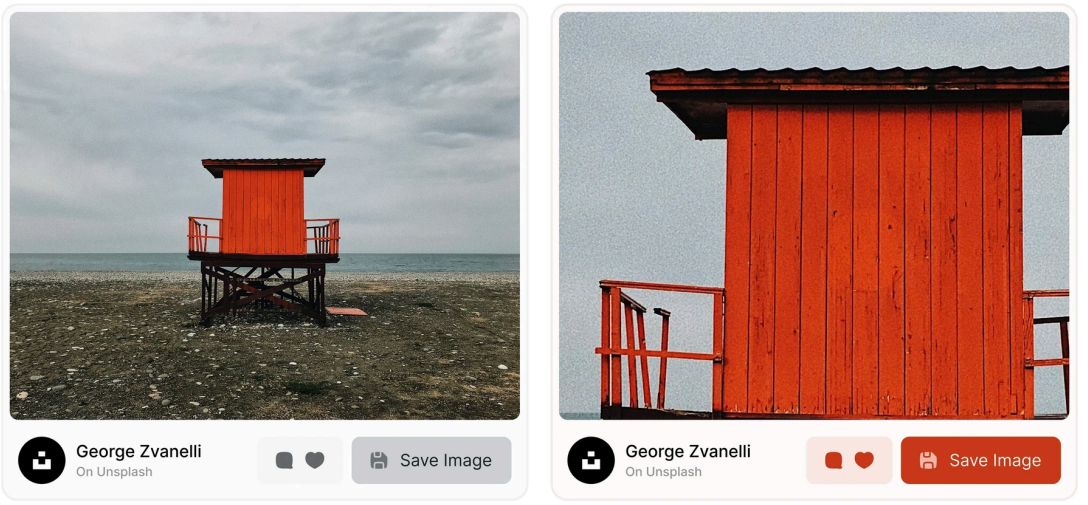

Figure 2: Accidental salience-based theming due to area size. Cluster colors extracted with a basic implementation of a k-means algorithm, with k=5. Interfaces created for visualization purposes. © Robert Halbach, source image by ↗ George Zvanelli

Chunk 4

Images in Interface Context

It might seem risky to overlook the concept of saliency in color extraction; however, targeting image mood is often the right call—for various reasons. In the case of global system-level theming based on a wallpaper (think Google Material), it is safe to assume that users like the overall aesthetic.

Chunk 5

The image acts as a single source of preference expression, with no need to focus on dedicated regions. In scoped dynamic theming (think widgets in systems or most components on the Web), images fulfill different roles. As elements, they rarely stand alone but typically sit in close proximity to text, further media, or other elements—creating a unified common region with their surroundings.[10] [11] For such cases, general aesthetic theming is sufficient for creating stronger connections inside enclosed areas as well as making them stand out visually from their surroundings. So … when *does* saliency actually matter?

Communicating Information

In specific scenarios, color found inside areas of attention acts not only as a visual anchor but also delivers crucial information about objects to its surroundings. For example, on product cards in digital stores or marketplaces, aesthetic attributes like color share the spotlight with other high-attention elements, such as buttons, price tags, or promotional badges. Semantically, images and their content shift from being supplementary to conveying other information and become influential elements core to user decision-making.[12]

10: Soegaard, M. (2020). *Laws of proximity, uniform connectedness, and continuation – gestalt principles (part 2).* <https://www.interaction-design.org/literature/article/laws-of-proximity-uniform-connectedness-and-continuation-gestalt-principles-2>

11: Harley, A. (2020). *The principle of common region: Containers create groupings.* <https://www.nngroup.com/articles/common-region/>

12: Pereira, M., Cardoso, A., Fernandes, C., Rodrigues, S., & D´orey, F. (2023). *The influence of the image and photography of e- commerce products on the purchase decision of online consumers* (pp. 39–51). Springer. <https://doi.org/10.1007/978-3-031-25222-8>

How to Extract Attention (Intentionally)

As mentioned earlier, alternative approaches to general color extraction exist. One early example is *Autonomous color theme extraction from images using saliency* by Jahanian et al..[13] Instead of averaging across all available pixels, their method prioritizes colors from regions likely to draw human attention. This is done using saliency maps—grayscale mask overlays that highlight the most visually dominant areas of an image through brightness values (usually, the brighter an area is, the more dominant it gets). Early models use contrast and edge density to estimate this, while newer methods like U^2-Net or DISNet apply deep learning to segment images more precisely.[14] [15] [16]

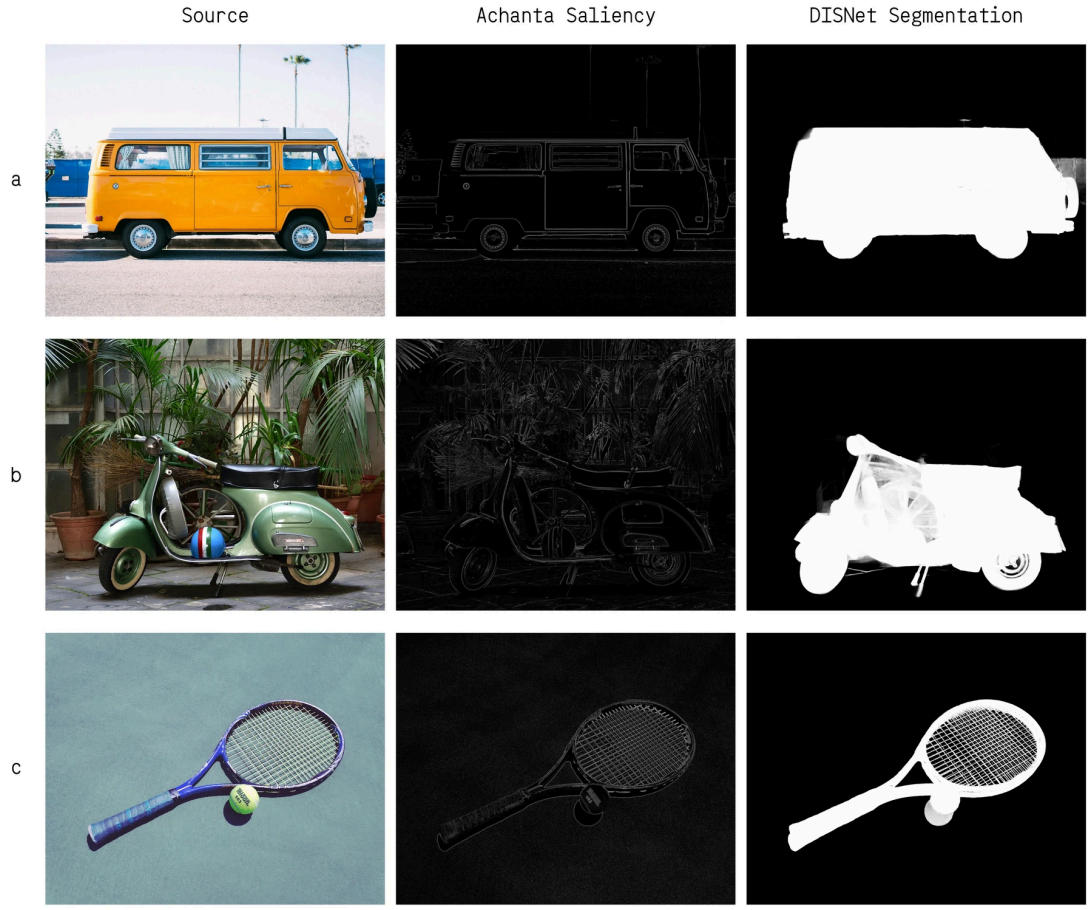

Figure 3: Comparison of input images (left), saliency maps generated using the method of Achanta et al. (center), and segmentation masks produced by DISNet (right). © Robert Halbach, source images a by ↗ Herson Rodriguez, b by ↗ Bertrand Borie, c by ↗ He Junhui.

Chunk 6

The improvement of saliency-selection technology over time could not be more apparent.

Chunk 7

However, when taking a peek at DISNet subtraction remainders, it becomes clear that even advanced segmentation models do not deliver perfectly clean separations.

13: Jahanian, A., Vishwanathan, S. V. N., & Allebach, J. P. (2015). *Autonomous color theme extraction from images using saliency*. Purdue University.

14: Achanta, R., Hemami, S., Estrada, F., & Susstrunk, S. (2009). Frequency-tuned salient region detection. *2009 IEEE Conference on Computer Vision and Pattern Recognition*, 1597–1604. <https://doi.org/10.1109/CVPR.2009.5206596>

15: Qin, X., Zhang, Z., Huang, C., Dehghan, M., Zaiane, O. R., & Jagersand, M. (2020). U2-Net: Going deeper with nested U-structure for salient object detection. *Pattern Recognition, 106*, 107404. <https://doi.org/10.1016/j.patcog.2020.107404>

16: Qin, X., Dai, H., Hu, X., Fan, D.-P., Shao, L., & Gool, L. V. (2022). *Highly Accurate Dichotomous Image Segmentation*. arXiv. <https://doi.org/10.48550/arXiv.2203.03041>



Figure 4: Segmentation-based palette extraction. Left: DISNet foreground masks with background removed. Center: distribution of remaining pixel colors in CIELAB space. Right: k-means cluster palettes based on the segmented regions. © Robert Halbach

Chunk 8

Annotation: CIELAB is a color space designed to model human color perception, where colors are described using three values: lightness (L*), red-green (a*), and yellow-blue (b*).

Considering model progression so far, it's realistic to assume that separations will improve in precision.

Chunk 9

But even if they do, objects usually consist of more than the one color ascribed to them.

Chunk 10

A car might have chrome details, black tires, and tinted windows—but still be perceived as red. Those semantic labels compress visual complexity into a singular dominant expression, which gets lost when extracting color from all visible object pixels. As a result, palettes may technically be correct, but that's not what we care about.

Chunk 11

When treating color from salient areas as information, perceived accuracy gains importance. Otherwise, disconnects may harm the initial goals of theming in the first place.

A way to connect segmented areas with relevant colors can be to use multimodal models—systems that combine visual input with semantic associations. Due to the large amount of training data and the predictive nature of such models, pointing towards expressive color descriptions of salient objects aligns closely with how these systems are designed to perform.

Chunk 12

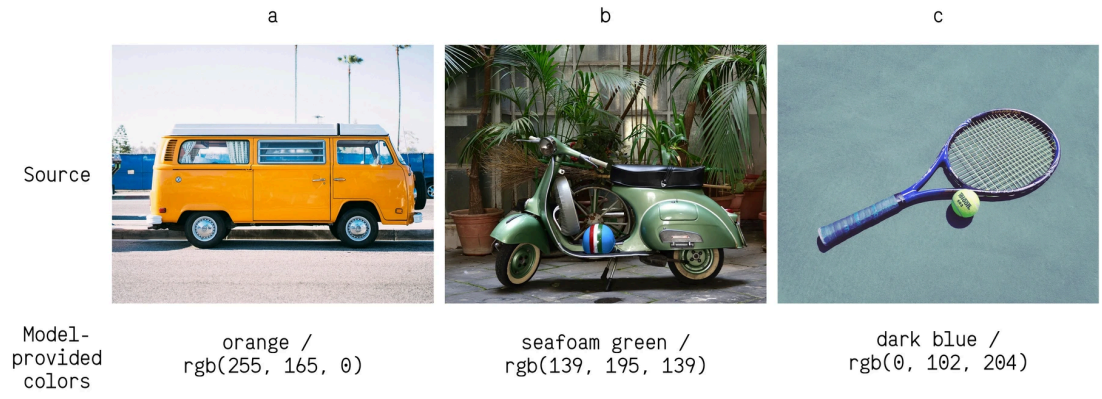

Figure 5: Color declaration of salient objects using ↗ Gemma 3. © Robert Halbach

Chunk 13

Annotation: Prompt: "Analyze the provided image and return a JSON object with exactly three keys: `object' (the main object in the image), `object_color_desc' (the single most dominant color of that object, be specific and accurate) and 'object_color' (a representation of object_color_desc in RGB, respond with RGB(R, G, B))."

Note: Colors were requested in RGB mode, since this format is assumed to be more prevalent in training data and is therefore expected to deliver more accurate results.

Now that semantic anchors for relevant color ranges exist, moving back to perceptual color space allows for filtering of source pixels.

Chunk 14

This is done by calculating ΔE_{00} (CIEDE2000) distances of source colors to model outputs, ΔE_{00} being a metric that describes how close individual colors are to each other perceptually.

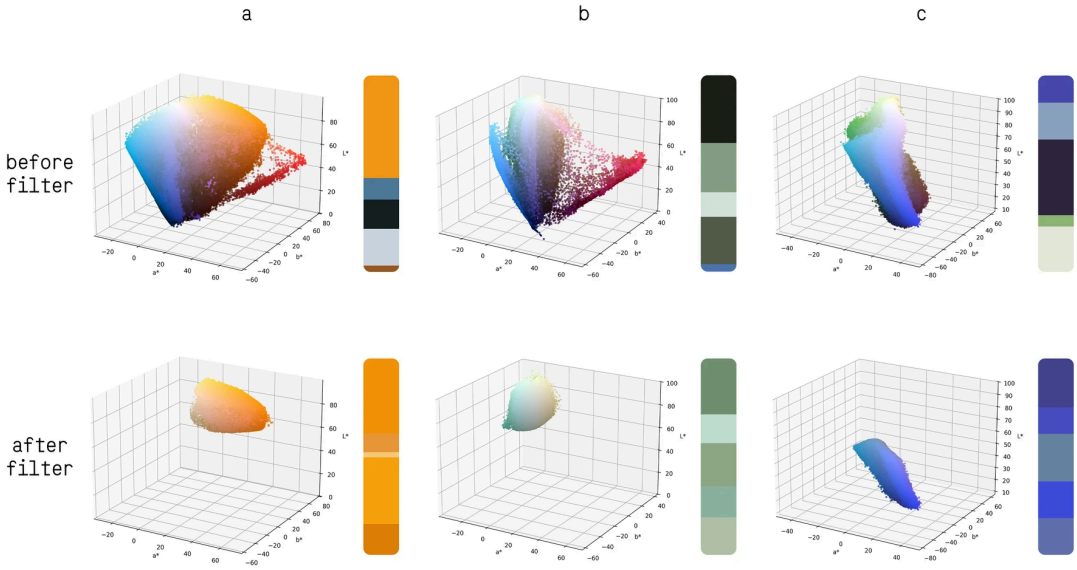

Figure 6: Color plots and k-means palettes before (top) and after (bottom) filtering with semantic anchors in CIELAB. © Robert Halbach

Chunk 15

Annotation: Filtering was done with a ΔE_{00} value of 20.

Chunk 16

For classical similarity filtering, that value would be far too high. However, this would only be the case when working with precise starting points, which can't be expected from a multimodal model like Gemma 3. What we obtain instead is a general idea of where the color range of salient objects starts and ends—enough to exclude very obvious mismatches.

Chunk 17

If you attempt this, some back-and-forth tweaking would be a good idea.

By retaining only values within a certain threshold of model-generated anchor, the result is a palette that better reflects what is meant by seafoam green or dark blue in context—favoring perception over surface frequency. If precision in terms of perception matters, slotting LLMs into pre-clustering pipelines is definitely a way to get there.

Objects, Perception, and Accuracy: A Conclusion

Saliency doesn't always matter in color extraction for UI theming. But when it does, though, it requires precision: simply detecting salient regions is not enough to deliver perceptually and semantically accurate palettes. Representing objects in a relatable way requires a connection between areas of attention and the subset of colors humans actually care about. Multimodal LLMs can deliver that, but they're not cheap.

Chunk 18

Inference, especially at scale, can add a giant overhead to production.

Deciding whether it's worth it requires weighing how much saliency-driven extraction improves usability compared to ambient color extraction and how it integrates with other potential system components—such as color-based product classification, filtering, or visual search engines.

↗ f1ekd6i6

Body as Latent Space: Alternative Compressions in Aesthetic Practice

Simon Maris

Chunk 1

Both AI systems and human bodies operate as compression algorithms, each reducing the infinite complexity of experience into manageable patterns. While AI compresses vast datasets—billions of images, texts, or sounds—into mathematical latent spaces that enable unexpected recombinations, bodies compress lifetimes into tacit knowledge, gesture, and habit. By comparing these parallel compression regimes, we reveal not only what each system preserves and discards, but how the violence of compression itself shapes what becomes knowable, creatable, and imaginable in aesthetic practice.

The Imaginary Map: Compression as Epistemology

"Yet a man who uses an imaginary map, thinking that it is a true one, is likely to be worse off than someone with no map at all; for he will fail to inquire whenever he can, to observe every detail on his way, and to search continuously with all his senses and all his intelligence for indications of where he should go." — E. F. Schumacher[1]

Maps compress territory. Neural networks compress data. Bodies compress experience.

Chunk 2

1: Schumacher, E. F. (1973). *Small is beautiful: A study of economics as if people mattered*. Harper & Row.

In each case, the act of compression is not neutral—it is an epistemological choice about what matters and what can be discarded.

Chunk 3

By compression, we mean the reduction of high-dimensional information into lower-dimensional representations. For AI systems, this is literal: mathematical dimensionality reduction. For embodied systems, following enactive approaches to cognition, compression occurs through active inference—the construction of predictive models through recursive interaction with environments. As von Foerster noted, organisms don't passively store experience but actively construct their reality through environmental perturbations.[2] Both forms of compression involve radical reduction, but each according to fundamentally different principles. Note that we deliberately employ the rationalist language of compression and optimization to reveal what such frameworks can and cannot capture about intelligence—using formalism against itself.

The seductive coherence of AI's latent spaces, like Schumacher's imaginary map, can blind us to what has been lost in the compression. These spaces appear continuous and complete, yet they are riddled with gaps where human experience doesn't translate into statistical patterns.

Consider Stable Diffusion: billions of images compressed into a few gigabytes, a reduction ratio that would seem impossible without massive information loss. These systems—whether generating images, text, or code—are fundamentally disembodied statistical engines, not the glowing robots of science fiction. Yet from this compressed space emerges the ability to generate infinite variations, including combinations never present in the training data—the infamous *avocado armchair* that tech companies celebrate as creativity. But what kind of creativity is this? The model can interpolate between avocado and armchair because both exist as statistical patterns in its latent space. It knows nothing of the cultural absurdity, the material impossibility, or the humor inherent in such a combination.

Chunk 4

The compression preserves statistical relationships while discarding meaning.

This raises fundamental questions about intelligence itself. As Blair Attard-Frost notes, intelligence is not merely functional ability but "value-dependent cognitive performance"[3]—it emerges from interdependencies between agents, their environments, and their measurers. If intelligence is indeed performative, then different compression systems perform intelligence according to different values, creating distinct epistemologies of what can be known and created.

Two Compression Regimes: Statistical vs. Embodied

Bodies and AI systems represent fundamentally different approaches to compression, each performing intelligence according to distinct values and constraints.

AI compression operates through statistical optimization, preserving what is frequent across the dataset while discarding outliers and anomalies. It optimizes for reconstruction fidelity.

Chunk 5

This creates a continuous mathematical space where interpolation between any two points is possible. The system values statistical regularity and mathematical efficiency.

III. p. 235, Chunk 15: Notes Toward AI Un/Learning

Chunk 6

Embodied compression operates through entirely different principles. It preserves what matters for survival and meaning, often discarding what is merely frequent but insignificant.

Chunk 7

A lifetime of reaching, grasping, and navigating compresses into proprioceptive patterns. Years of craft practice compress into the hands' tacit knowledge. Cultural memories compress into gesture and gait.

Chunk 8

2: Foerster, H. von. (2003). On Constructing a Reality. In H. von Foerster (Ed.), *Understanding Understanding: Essays on Cybernetics and Cognition* (pp. 211–227). Springer. <https://doi.org/10.1007/0-387-21722-3_8>

3: Attard-Frost, B. (2023). Queering intelligence: A theory of intelligence as performance and a critique of individual and artificial intelligence. In *Queer Reflections on AI*. Routledge.

This compression optimizes for actionable knowledge.

Chunk 9
It creates discrete skills and capacities that resist smooth interpolation.

Chunk 10
The difference becomes clear in how each system handles the spaces between known points. AI interpolates smoothly—generating endless variations along the mathematical manifold. But embodied knowledge often faces discontinuity.

Chunk 11
The space between learned movements is not smooth but potentially catastrophic. Early AI video models attempting bicycles in motion produced surreal failures—bikes melting, riders merging with vehicles. They could interpolate between static images but couldn't capture dynamic stability.

Chunk 12
A human learning to ride doesn't gradually interpolate from falling to riding—they either balance or crash.

Between the Points: What Compression Discards

Singling out what exists between the points in a latent space is not discovery but statistical averaging.

Chunk 13
The *avocado armchair* is not found but computed, a weighted sum of existing patterns.

Chunk 14
The space between is filled with mathematical necessity, not possibility.

What AI's compression discards—the context that makes combinations meaningful or absurd, the effort required to bridge domains, the resistance materials offer to impossible combinations, the cultural weight that makes some mixtures sacrilegious, the embodied impossibility that makes *avocado armchair* a joke rather than a design proposal—reveals its values are strictly numerical.

Embodied compression discards differently.

Chunk 15
It may forget statistical regularities that don't affect action, precise measurements that can be felt approximately, or verbal descriptions of kinesthetic knowledge. Most significantly, it often discards the average in favor of the exceptional—the surprising grain pattern that warned of weakness, the unusual movement that prevented injury, the rare combination that actually worked.

Discriminatory Compression: The Violence of Reduction

Both compression regimes discriminate, but their criteria differ radically. This discrimination is violent—it destroys information, erases particularity, and forgets entire domains of experience.

Chunk 16
Understanding this violence is crucial for recognizing what each compression makes possible and impossible.

In AI systems, the violence is systemic and statistical. The compression erases the particular in favor of the pattern—the way your parrot still constantly mimics the one time you used the f-word becomes a point in the statistical space of *parroting behavior.* Individual artworks become vectors in style space. The unique becomes the typical. This violence is compounded by whose patterns get preserved and whose get erased, as scholars like Ruha Benjamin[4] and Kate Crawford[5] have extensively documented. The statistical nature of AI compression is inherently political—it amplifies existing patterns of representation and exclusion.

In embodied compression, the violence is personal and consequential.

Chunk 17
The forgetting happens through disuse: what doesn't matter doesn't get preserved.

Chunk 18

4: Benjamin, R. (2019). *Race After Technology: Abolitionist Tools for the New Jim Code*. Wiley.

5: Crawford, K. (2021). *The Atlas of AI: Power, Politics, and the Planetary Costs of Artificial Intelligence*. Yale University Press.

Thousands of unremarkable days compress into *childhood.* Skills atrophy without practice.

III. p. 179, Chunk 6: Introduction

Chunk 19

Cultural knowledge dies with its practitioners. This compression preserves what serves immediate needs while potentially losing what might matter to future generations—a loss that AI's seductive completeness may accelerate as embodied practices are increasingly replaced by statistical approximations.

Yet both compressions enable action precisely through their selectivity. AI's compression enables the generation of new images, the recognition of patterns, the automation of vision. Bodily compression enables craft, navigation, and survival. The question is not whether to compress but how to recognize what each compression makes possible and impossible.

The Necessity of Loss: Toward Critical Compression Practices

If we accept that both AI and bodies necessarily compress, and that compression always involves loss, then the critical question becomes: How do we work with multiple compression regimes without letting one colonize the other?

The danger lies not in AI's ability to generate *avocado armchairs* but in mistaking its compression for comprehensiveness.

Chunk 20

When designers turn to AI for *creativity,* they often receive statistical recombinations dressed as innovation—technically novel but experientially empty. The model can combine any two concepts that exist in its latent space, but it cannot tell you why one shouldn't.

Chunk 21

Embodied knowledge offers a different creativity—one that knows material limits, cultural boundaries, and meaningful transgressions. The master chef who breaks culinary rules knows which rules matter because their body has internalized them through practice. Their innovation comes from understanding, not from interpolation.

III. p. 179, Chunk 6: Introduction
II. p. 90, Chunk 13: Der unerwartete Gast
II. p. 4, Chunk 3: Vorwort

Chunk 22

As AI systems grow more sophisticated—developing what companies call *reasoning capabilities* and *agentic behavior*—the stakes of understanding the distinctions between these different compressions increase. Future human-computer interaction will increasingly involve navigating between statistical and embodied ways of knowing. Multimodal AI systems will compress not just images but movement, touch, and time. Design education must prepare practitioners to work critically with both compression regimes, understanding what each preserves and destroys.

Conclusion: Parallel Compressions

The body is not literally a latent space but a parallel compression system operating on different principles with different losses. AI's latent spaces seduce with their mathematical completeness and infinite interpolations, offering *creativity* through recombination. Bodies compress through practice and consequence, preserving what matters rather than what's frequent.

Chunk 23

The experimental insight is not that one compression is superior, but that their comparison reveals the choices hidden in each. When we understand bodies and AI as alternative compression algorithms, we can begin to ask: What does each preserve? What does each discard?

III. p. 235, Chunk 10: Notes Toward AI Un/Learning
III. p. 235, Chunk 15: Notes Toward AI Un/Learning
II. p. 42, Chunk 17: The Cognitive Style of ChatGPT

Chunk 24

And how might we work in the spaces between compressions, where neither statistical pattern nor embodied knowledge fully reaches?

For designers and artists, this means developing literacy in both regimes—understanding not just how to prompt an AI or train a model, but how to recognize what its compression cannot capture. It means valuing embodied knowledge not as

nostalgia but as an alternative intelligence that preserves different truths. Most importantly, it means resisting the collapse of all intelligence into a single model, whether computational or corporeal.

Perhaps the most radical possibility is not choosing between compressions but recognizing their mutual incompleteness. In the gap between statistical and embodied knowledge lies a space that neither can fully capture—a reminder that the map, no matter how sophisticated its compression algorithm, is never the territory.

↗ gpm91yc8

On the Materiality of AI

Ting-Chun Liu

Chunk 1

Introduction

”we have re-enabled chatgpt plus subscriptions! thanks for your patience while we found more gpus”[1]

Sam Altman's tweet revealed how complex networks expose themselves only when certain nodes malfunction. The artificial intelligence(AI) industry relies on centralized semiconductor production, intensive energy consumption, and supply chains contributing to global environmental and political tensions. As Langdon Winner noted, societal conflicts and alliances are determined not just through political institutions but also through the physical arrangements of infrastructure: steel, concrete, wires, and semiconductors.[2]

This piece examines how Taiwan’s semiconductor industry, particularly Taiwan Semiconductor Manufacturing Company (TSMC), creates both technological dependency and geopolitical vulnerability in the global AI infrastructure. Taiwan's strategy for maintaining its sovereignty through the *Silicon Shield* paradoxically requires accelerating technological systems that potentially undermine sovereignty elsewhere.

While writing the piece, there is a slight hum from the laptop, barely noticeable yet constants.

Chunk 2

This hidden noise from the fan, quietly cooling down trillions of transistors that were once grains of sand. This gentle hum signifies the countless layers of technology working beneath the surface, a constellation of silicon, metal, and code in operation.

Main Part

Technology is a complex, shifting, recursive landscape, engaging with it as an artist reveals a double-edged reality. As Simon Penny observed, artists working with technology inevitably become entangled in consumer commodity economics.[3] This feedback loop became evident when OpenAI's published its first text-to-image model DALL·E.

Chunk 3

It produced strange compositions that felt uncanny and raw, qualities that today's polished outputs have eliminated.

Chunk 4

I was eager to work with these effects, yet as the model became outdated access was removed.

Chunk 5

I found myself unable to run alternatives locally due to lacking computational hardware.

```
device = torch.device("cuda" if
torch.cuda.is_available() else "cpu")
```

1: Altman, S. (2023). *We have re-enabled ChatGPT Plus subscriptions! Thanks for your patience while we found more GPUs*. Twitter. <https://twitter.com/sama/status/1734984269586457078>

2: Winner, L. (1980). Do artifacts have politics? *Daedalus*, *109*(1), 121–136.

3: Penny, S. (1995). Consumer culture and the technological imperative: The artist in dataspace. In S. Penny (Ed.), *Critical issues in electronic media* (pp. 47–74). State University of New York Press.

Whether Midjourney, DALL·E, Stable Diffusion, or ChatGPT, all generative models were trained on NVIDIA's data center graphic processing units (GPUs), notably the *NVIDIA A100*. ChatGPT alone required an estimated 20,000 to 30,000 A100s just to train.[4] Billions of tiny cores in GPUs, originally designed for rendering 3D games, turned out to be perfect for training neural networks, due to their ability to compute massive volumes of simple, repetitive mathematical operations. The turning point came in 2012 when AlexNet,[5] a deep convolutional neural network, leveraged GPUs' parallel processing capabilities that outperformed all other image classification approaches, making graphic cards the industry standard.

Chunk 6

One of the researchers Alex Krizhevsky was then employed by NVIDIA to develop *CUDA* (Compute Unified Device Architecture), a proprietary framework for parallel computing.

Chunk 7

```
RuntimeError: Attempting to deserialize object
on a CUDA device but torch.cuda.is_available()
is False. If you are running on a CPU-only
machine, please use torch.load with
map_location='cpu' to map your storages to the
CPU.
```

This CUDA error message exposes a deeper dependency—most AI models expect CUDA-enabled environments that only work with NVIDIA GPUs. Their hardware is not just a technical choice but a gatekeeping mechanism.

The production of semiconductors that power modern GPUs reveals a complex global network spanning multiple continents and involving dozens of specialized companies (Figure 1).

Chunk 8

A making of semiconductor begins as one of the most common materials: sand, specifically quartz or silicon dioxide. Despite its ubiquity, chip production demands unnatural purity at 99.9999999%.

Chunk 9

Only two companies worldwide control the critical extraction and refinement process.

Chunk 10

And both are based in Spruce Pine, North Carolina. They turn sand into silicon that then travels to companies in Japan to be formed into wafers. At a talk at the *Correlations Forum* in *hfg Offenbach*, I passed around wafer pieces to the audiences. What always strikes audience members is their delicacy, which even slight touches left grease marks and contamination, revealing why manufacturers must work in ultra-clean environments.

Chunk 11

These wafers then face the most critical bottleneck in chip production: lithography. Only the Dutch company ASML (Advanced Semiconductor Materials Lithography) produces the Extreme Ultraviolet (EUV) machines needed for printing advanced semiconductors, holding 100% market share with machines costing over $150 million each. This monopoly creates profound political leverage, as recent U.S. export bans have prohibited these machines from being sold to China.[6]

4: Tom`s Hardware. (2023). *ChatGPT will command more than 30,000 NVIDIA GPUs: report.* <https://www.tomshardware.com/news/chatgpt-nvidia-30000-gpus>

5: Krizhevsky, A., Sutskever, I., & Hinton, G. E. (2012). ImageNet classification with deep convolutional neural networks. *Advances in Neural Information Processing Systems, 25*, 1097–1105. <https://papers.nips.cc/paper/4824-imagenet-classification-with-deep-convolutional-neural-networks.pdf>

6: ASML. (2024). *ASML expects impact of updated export restrictions to fall within outlook for 2025.* <https://www.asml.com/en/news/press-releases/2024/asml-expects-impact-of-updated-export-restrictions-to-fall-within-outlook-for-2025>

Figure 1: The Making of Graphic Cards (Nvidia RTX 4090) - Global supply chain visualization showing the journey from raw materials to final product. © Ting-Chun Liu. Contains materials courtesy of NVIDIA, TSMC, ASML, and other manufacturers

Chunk 12

All these elements then converge at TSMC, the world's leading chip fabricator, which produced about 90% of advanced semiconductors as of 2022.[7] Here, NVIDIA's electronic layouts are etched onto silicon with ASML's EUV machines, transforming chip designs from mere drawings into functional processors.

Chunk 13

TSMC's dominance emerged from specific historical circumstances when the company was founded in 1987. At the end of Taiwan's forty years under martial law, TSMC filled the gap created when U.S. pressure forced Japan to retreat from semiconductor dominance through the 1986 U.S.–Japan Semiconductor Agreement.[8] Taiwan, transitioning from authoritarian rule to democracy, recruited Morris Chang, a Taiwanese-American engineer who had never set foot on Taiwanese soil, to establish the world's first dedicated semiconductor foundry. TSMC's establishment actively constructing Taiwan as a geopolitical hinge in the global semiconductor supply chain where sand from North Carolina becomes the foundation of our digital world.

In 2021, Taiwan experienced its worst drought in a century. The high-tech industry, barely sustained itself with government support, often at the expense of the general populace. The drought coincided with sudden chip demand, leading to global chip shortages that brought Taiwan's semiconductor industry under scrutiny.[9] Taiwan's then-President Tsai Ing-wen reiterated the importance of the *Silicon Shield*,[10] describing how the country's chip dominance serves as both economic leverage and a deterrent against Chinese aggression.[11] Semiconductor is not just a technical asset—it's a geopolitical instrument of sovereignty. Taiwan must maintain semiconductor dominance if it is to remain independent. Taiwan is structurally entangled in AI's acceleration. Building the digital infrastructure also enables extractive computation and the consolidation of power through information. And here the dilemma comes arises: *protect its sovereignty, Taiwan must help accelerate a system that undermines sovereignty elsewhere and generates new crises*. What protects this one island may destabilize others. Silicon is not neutral.

Chunk 14

It underpins the tools of generative AI and militarized computing. Fighter jets equipped with TSMC-manufactured chips,[12] were deployed by the Israeli government in Gaza.[13] With the U.S. and European Chip Acts restricting China's access to high-end processors[14] Taiwan has become the pivot point in an emerging techno-political order. The very infrastructure of AI is built on decades of political engineering. TSMC's origins makes clear that advanced computation was never the result of free markets alone. It was infrastructural, strategic, and deeply historical.

Chunk 15

Today, that uncertainty hasn't disappeared. The chips Taiwan produces now operate in systems of generative AI, predictive policing, algorithmic governance, and synthetic culture. Despite all this, Taiwan itself remains one of the region's most at risk polities if this acceleration continues in this direction. There is no guarantee that Taiwan itself will not be made expendable.

Conclusion

AI is more than the symbol of clouds or the tireless server room at work; its material foundations are inseparable from its geopolitical

7: Miller, C. (2022). *Chip war: The fight for the world`s most critical technology* (p. 8). Scribner.

8: Miller, C. (2022). *Chip war: The fight for the world`s most critical technology* (p. 108). Scribner.

9: Chang Chien, A., & Ives, M. (2021). *Taiwan prays for rain and scrambles to save water*. Online; The New York Times. <https://www.nytimes.com/2021/04/08/technology/taiwan-drought-tsmc-semiconductors.html>

10: Addison, C. (2001). *Silicon shield: Taiwan`s protection against chinese attack*. Fusion Press.

11: Tsai, I. (2021). Taiwan and the fight for democracy. *Foreign Affairs, 100*(6). <https://www.foreignaffairs.com/articles/taiwan/taiwan-and-fight-democracy-tsai-ing-wen>

12: McGlaun, S. (2020). *TSMC under U.S. Pressure as chip supplier for Lockheed F-35 Lightning II: report*. Online; HotHardware. <https://hothardware.com/news/tsmc-under-pressure-to-build-chips-in-us>

13: Wright, G. (2024). *Dutch court orders halt to F-35 jet parts exports to Israel*. Online; BBC News. <https://www.bbc.com/news/world-europe-68272233>

14: Freifeld, K., & Potkin, F. (2024). *Exclusive: US ordered TSMC to halt shipments to China of chips used in AI applications*. Online; Reuters. <https://www.reuters.com/technology/us-ordered-tsmc-halt-shipments-china-chips-used-ai-applications-source-says-2024-11-10/>

nature. Engaging with these systems means also engaging with the intangible structures that constitute and are constituted by our relation to the world.

Chunk 16

The question then is not simply who is responsible for these systems, but how we choose to engage with the collective systems we are all, to some degree, complicit in. The gentle hum of the laptop fan continues, cooling transistors that were once grains of sand. To work with AI, is to work with these contradictions, to remain close not only to the surface of the output, but also the condition that shaped it—the layers beneath, the labor behind, the heat in circuits, the extracted nature within networks, the intangible infrastructure, and the silent hum behind the interface.

This writing is part of an artistic research documented at:
https://aprilcoffee.github.io/heat_as_image/

↗ 59o29bm4

un/learn the process

Francesco Scheffczyk,
Paul Eßer, Lars Hembacher

Chunk 1

This text is a collection of thoughts and inspiration, as well as a description of the process of finding the right shape for the un/learn ai series. It is written in a non-linear way.

Chunk 2

Each section can be understood as an individual text, with references to other sections provided through an annotation system based on Unicode characters. Each heading is preceded by a character that acts as a reference.

Chunk 3

Within the body of the text, these characters appear at points that refer to, or can be associated with, the corresponding section.

a… Hypertext

The term `hypertext`, coined by Theodor Nelson in 1965, describes "non-sequential writing—text that branches and allows choices to the reader, best read at an interactive screen".[1] Today, it is understood as interactive networks of digital documents, media objects and other information objects, connected by hyperlinks. They are central to the World Wide Web, which forms the biggest hypermedia to date.[2]

Early hypertext theorists like Landow, Bolter, and Coover connected its non-linearity, reader choice, and polyvocality to postmodern literary theory, seeing it as liberating and empowering by "favoring a plurality of discourses over definitive utterance and freeing the reader from the domination of the author."[3] [4] [5] Landow's concept of *wreadership* highlights readers as co-authors who actively co-construct the text.[6]

The un/learn ai series ⌐→❁ extends the concept of the *wreader* by applying AI for semantic hyperlinking ⌐→⌖, rather than relying solely on author-created links⌐→⁘ and situating itself within hypermedia history ⌐→Ψ.

Chunk 4

Ψ Proto-Hypertext

Proto-hypertexts were early forms and conceptualizations of hypertext systems ⌐→a…. The definitions vary. One perspective, articulated by scholar Belinda Barnet and aligning with Nelson's definition of hypertext, considered proto-hypertexts to be systems developed before the World Wide Web that enabled branching and reader-chosen paths on interactive screens.[7] This included concepts like Vannevar Bush's *Memex*, a device

1: Nelson, T. H. (1987). *Computer lib ; dream machines*. Redmond, Wash. : Tempus Books of Microsoft Press.

2: Ensslin, A. (2014). Hypertextuality. In *The John Hopkins Guide to Digital Media* (pp. 258–265). JHU Press.

3: Bolter, J. D. (1992). Literature in the electronic writing space. *Literacy Online: The Promise (and Peril) of Reading and Writing with Computers*, 19–42.

4: Ensslin, A. (2014). Hypertextuality. In *The John Hopkins Guide to Digital Media* (pp. 258–265). JHU Press.

5: Coover, R. (1992). The End of Books. In *New York Times Book Review*. https://archive.nytimes.com/www.nytimes.com/books/98/09/27/specials/coover-end.html.

6: Landow, G. P. (1992). *Hypertext: The convergence of contemporary critical theory and technology*. Johns Hopkins University Press.

7: Barnet, B. (2014). *Memory machines: The evolution of hypertext*. Anthem Press.

for creating associative trails between documents stored on microfilm,[8] and Paul Otlet's *Mondothèque*, a networked multimedia workstation for interlinking media.[9]

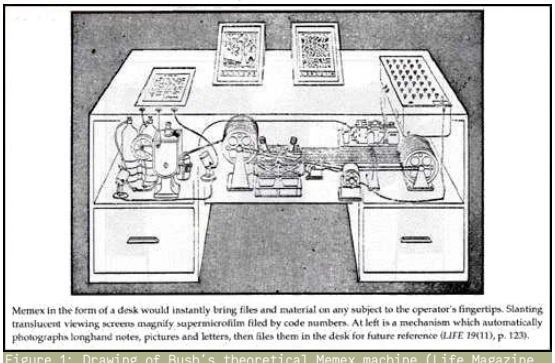

Figure 1: Drawing of Bush's theoretical Memex machine (Life Magazine, November 19, 1945), public domain, via https://www.kerryr.net/

Chunk 5

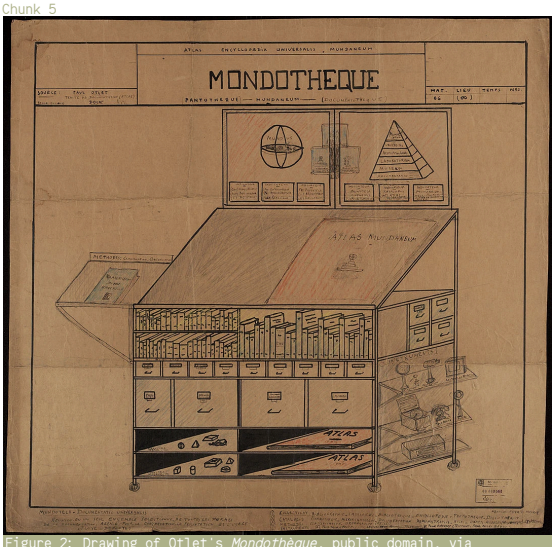

Figure 2: Drawing of Otlet's *Mondothèque*, public domain, via Wikimedia Commons.

Chunk 6

A broader definition, proposed by scholars like Astrid Ensslin, identified proto-hypertext in any text creating multilinear reading through connections, annotation, or cross-referencing. This view included examples from before the 20th century, such as annotations in medieval religious texts like the Talmud or Bible, baroque experimental poetry such as Georg Philipp Harsdörffer's *Wechselsatz*, and encyclopedias from the Enlightenment period in Europe that employed systematic cross-referencing.[10]

Non-linear Modes of Reading

As the name suggests, non-linearity in reading means everything that is not read from beginning to end, word by word, sentence by sentence.

Chunk 7

Non-linear reading is a curse and a blessing. On the one hand, it doesn't guarantee any completeness of the written, as a linear book might do, as there is no beginning or end.

Chunk 8

It is more of an intuitive feeling of navigating a network of interconnected thoughts or claims that do not form an encompassing narrative by any individual author.

This is also a big problem, as the narrative has to be formed by the reader themselves, which is an exhausting task. Overstimulation is a big concern here, as its symptoms are disinterest and annoyance, so we need to be careful to keep everything digestible.

Knowledge Representation in AI

In the field of artificial intelligence, knowledge is generally defined as a collection of information related to a particular domain that can be used to solve problems within that domain.[11]

While AI relies on data, raw data is not intelligence. For AI to solve domain-specific problems, this data must be structured into knowledge. This operational *knowledge* includes factual data, conceptual frameworks, procedural guidelines, models, and heuristics that an AI system comprehends or learns. Unlike information (processed data), knowledge is personalized or

8: Barnet, B. (2014). *Memory machines: The evolution of hypertext*. Anthem Press.

9: Helmond, A. (2019). A Historiography of the Hyperlink: Periodizing the Web through the Changing Role of the Hyperlink. In *The SAGE Handbook of Web History* (pp. 227–241). SAGE Publications Ltd. <https://doi.org/10.4135/9781526470546.n16>

10: Ensslin, A. (2014). Nonlinear Writing. In *The John Hopkins Guide to Digital Media* (pp. 360–362). JHU Press.

11: Bhuyan, B. P., Ramdane-Cherif, A., Singh, T. P., & Tomar, R. (2025). Knowledge Representation in Artificial Intelligence. In B. P. Bhuyan, A. Ramdane-Cherif, T. P. Singh, & R. Tomar (Eds.), *Neuro-Symbolic Artificial Intelligence: Bridging Logic and Learning* (pp. 47–61). Springer Nature. <https://doi.org/10.1007/978-981-97-8171-3_4>

contextualized, enabling informed AI decision-making and actions.[12] [13] Knowledge can be represented in various ways, including logic-based systems, probabilistic models, and more.[14]

To understand complex information in AI applications, knowledge is often transformed into a numerical format. Modern AI implementations usually use *embeddings* (which represent words or concepts as lists of numbers) and *knowledge graphs* (which store how things are connected). This allows AI to identify subtle, hidden patterns within large amounts of data. These techniques has been used in the un/learn ai series, transforming the book content into numerical vectors (embeddings) .

Chunk 9
This allowed large language models to parse the texts and support the editorial process by creating semantic links between paragraphs .

Semantic Chunking

What is the smallest size you can break a text down to, with the fragments still having meaning that differs from all other fragments?

Semantic chunking is a way to break down a text into separate text fragments called chunks. A chunk must at least end with a glyph signifying the end of a sentence. Start at the first sentence and read the next one.

Chunk 10
Are they similar?

Chunk 11
Yes? They are in the same chunk!

Chunk 12
No? We end chunk 1 after the first and start a new one for our second sentence. Go on until you reach the end.

The chunk is embedded into a space where distance is inverse to similarity. In this space we find all other fragments from all the other texts we added. Find chunks near this one, if there are any.

Chunk 13
These should be semantically related.

AI Post-production

Autonomously AI-generated *content* is per se mediocre, a median. We want to use AI to add to something that already exists. We use it to make footnotes, not make up sources in pre-writing, but create footnotes in post-writing to connect the texts that share a similar topic. Here, AI doesn't create; it classifies and adds another layer.

[N]otes [...] not only show the labor and effort that underpins research and writing; they also signal stories of other stories that direct you to a story and place connected to, but not of, the story you began.[15]

Web2Print

Why use Web to print?[16]

1. Web technologies allow for easy automatic layout generation. Apps like InDesign will punish you for changing text after designing, there is an expected order from writing text to designing it, as long as one person isn't doing both. Even once you've built your layout, Web technologies allow you to add more content afterward. Designing a book series with only variations in the layout becomes easier when you just expand on code.

12: Bhuyan, B. P., Ramdane-Cherif, A., Singh, T. P., & Tomar, R. (2025). Knowledge Representation in Artificial Intelligence. In B. P. Bhuyan, A. Ramdane-Cherif, T. P. Singh, & R. Tomar (Eds.), *Neuro-Symbolic Artificial Intelligence: Bridging Logic and Learning* (pp. 47–61). Springer Nature. <https://doi.org/10.1007/978-981-97-8171-3_4>

13: Knowledge Representation in AI. (2025). In *GeeksforGeeks*. https://www.geeksforgeeks.org/knowledge-representation-in-ai/.

14: Knowledge Representation in AI. (2025). In *GeeksforGeeks*. https://www.geeksforgeeks.org/knowledge-representation-in-ai/.

15: McKittrick, K. (2020). Footnotes (Books and Papers Scattered about the Floor). In *Dear Science and Other Stories* (p. 19). Duke University Press.

16: One of the core ideas for the web was web2print: https://info.cern.ch/hypertext/WWW/Bugs.html (under "Text from hypertext"). Tim Berners Lee wrote here, that there should be "a quick way to print a book from the web."

2. Web technologies allow for configuring processes in publishing pipelines. Post-processing of text can be done automatically without human labor. Processes like semantic chunking, vector similarity search etc. were only possible in the un/learn ai series because of a post-processing pipeline.

3. Publishing in different media gets a lot easier using the same technology for everything.

Chunk 14

This combines the strengths of different media to better communicate what is being published.

II. p. 5, Chunk 1: KI-Features dieser Ausgabe

Chunk 15

If you ever tried to make a website or an epub from an InDesign file, you know it's bad.

Making un/learn

The un/learn ai series was conceptualized as a web2print publication to enable custom usage of AI in the creation. We used AI to split our text into semantic fragments, which were then embedded. Embeddings (being vectors) could tell us how semantically similar two text fragments were, based on their geometric distance in the embedding space. With this information, we were able to generate cross-references between articles, creating a secondary non-linear mode of reading.

Chunk 16

These fragments or chunks also formed the basis for the design of the series.

A huge inspiration for the creation of the publications was the book "Graphic Design in the Post-Digital Age" by Demian Conrad, Rob van Leijsen, designed by Johnson/Kingston.



Figure 3: Screenshot of the website from *Graphic Design in the Post-Digital Age*, https://www.postdigitalgraphicdesign.com

Chunk 17

In Volume 1, *un/learn ai — Approaching AI in Aesthetic Practices,* we visually highlighted the chunks by alternating gradients on the left and right sides. As soon as a chunk showed a strong semantic similarity to another chunk, we added a reference a… to the corresponding article and chunk in the margin column.

III. p. 172, Chunk 2: StableHub
II. p. 24, Chunk 12: Schwäbisch Gmünd
I. p. 2, Chunk 1: Vorwort

Chunk 18

Thus, for example, chunk 3 from the article *Exploring Tools* is particularly relevant to chunk 1 from the article *Developing an AI Teaching Platform*, as both discuss the idea that design students should understand AI as a new tool and material.

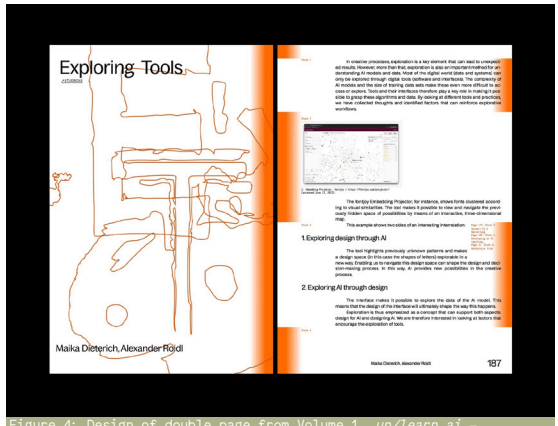

Figure 4: Design of double page from Volume 1, *un/learn ai — Approaching AI in Aesthetic Practices.*

In Volume 2, *un/learn ai — Navigating AI in Aesthetic Practices*, the highlighting of the chunks was reduced. Only the chunk annotations remained.

Chunk 19

Instead of highlights, a grid of ten terms was placed in the background of the pages. Each chunk was then semantically compared with the ten terms, and the term with the strongest semantic relation was assigned to the chunk. The result is a tree structure in the background of the page that connects the semantic focus of each chunk. This allows the thematic progression of an article to be viewed on a secondary level.

Chunk 20

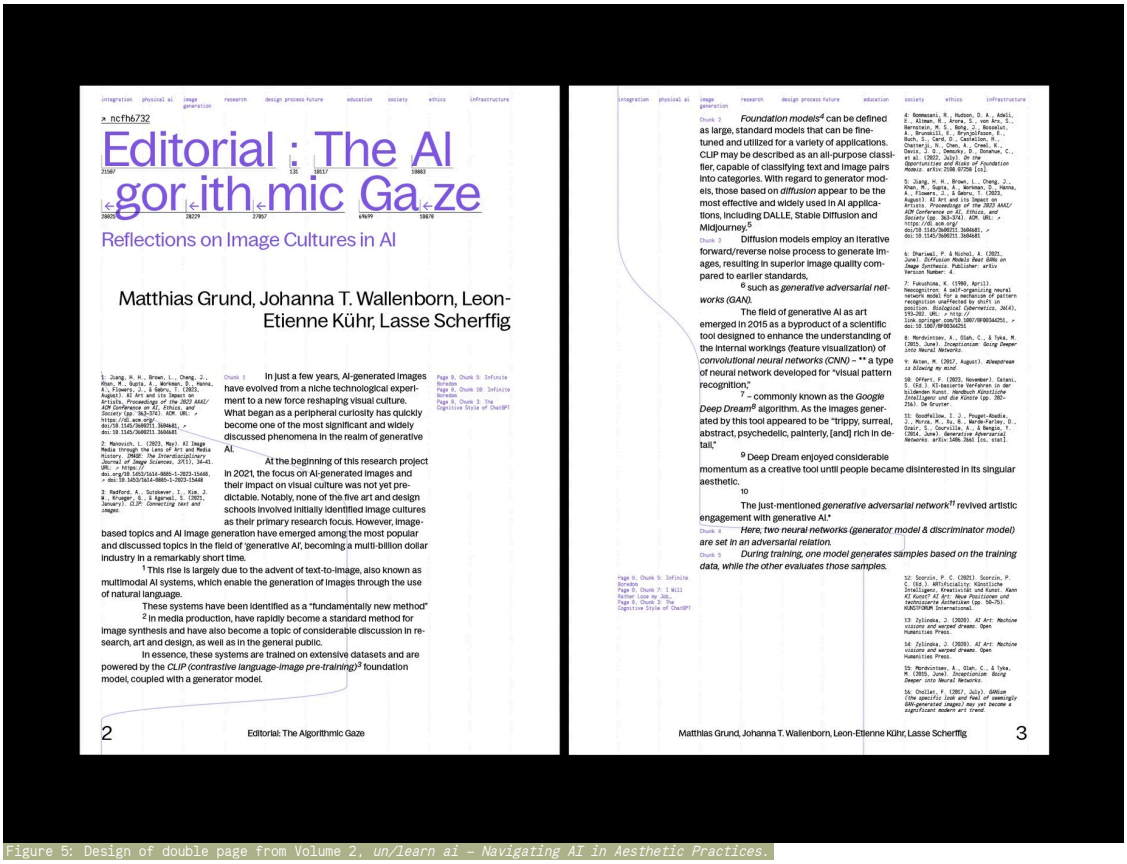

Figure 5: Design of double page from Volume 2, *un/learn ai – Navigating AI in Aesthetic Practices.*

At the time of writing, the design of this volume had not yet been finalized.

Chunk 21

It is therefore difficult to describe the design process.

Chunk 22

We are excited to see the final result, and you probably already have an idea of what it will look like, given that you are reading this text in the third volume's design.

The website functions as a global index and includes all texts, which can be navigated and filtered in various ways.

Chunk 23

In terms of its general design, e.g., on the overview pages, the website has been deliberately designed in a minimalist style. The articles are presented in the design of the respective volumes.

Chunk 24

Semantic referencing ⌐→≑ is applied across the three volumes. By clicking on a reference ⌐→a…, you can jump from book to book, from design to design.

↗ c7cors5e

26 Strategies to Consider When Writing with AI

Julia-Jasmin Bold

III. p. 179, Chunk 11: Introduction
III. p. 179, Chunk 10: Introduction

Chunk 1

The *wr(AI)ting Card Deck*[1] is an experimental approach to help designers compose better texts with the support of AI. Using practice-based examples, it illustrates how AI can be applied to daily writing tasks and, in the process, encourage the development of new perspectives and author-like skills.

Each card contains a *strategy* and an indication of which *phase in the writing process* (Prewriting, Planning, Drafting, Revising, Editing, Takeaway) it could be most useful for.

The deck is divided into *eight "chapters"* to make it more intuitive:

Chapter 1: "The Creative Spark"

Start where there is a spark!

Ignite new ideas, explore unexpected perspectives, and overcome challenges before you even start drafting.

How can I use AI to spark my creativity? What's the point I'm trying to make? How to start?

Chunk 2

Chapter 2: "Momentum & Motivation"

Maintain your momentum!

Keep your writing energy up, even when it dips. Use these strategies to regain motivation and pick up where you left off on stalled drafts.

Procrastinating and feeling overwhelmed by writing?

Chunk 3

What compels me to write? Blocked?

Chapter 3: "Writing Support"

Use AI as your second brain

When it comes to composing texts, AI can help you to organize your thoughts, experiment with different phrases, and try out new writing techniques.

Can AI help me with that pesky first sentence? What concepts or visuals are central to this topic? How can I add more depth to my text?

Chunk 4

1: Online: https://unlearn.gestaltung.ai/wr-AI-ting_card-deck_public.pdf

Chapter 4: “Structure & Outline”

Build your story

Give your text shape before – or even while – you write. These strategies reveal what might be concealed in the content.

How to start ‘building’ a text? What’s your key message?

Chunk 5
What is the core structure of my text?

Chapter 5: “Shifting Perspective”

Explore alternative views

Switch perspectives, tones, or gain (fictional) insight from your audience to unlock new depths in your writing.

Does the text fit my target audience? Can fictional voices help me write better texts? What style or perspective suits my text best?

Chapter 6: “The Reflective Mirror”

A looking glass that responds

Let AI hold up a mirror to your writing, highlighting clarity, consistency, and what’s really getting across.

How can AI strengthen my arguments? What would happen if my text shape-shifts?

Chunk 7
How does my writing come across to a reader?

Chunk 8

Chapter 7: “Revision”

Refine structure, style, and flow

… without losing your voice. These strategies help you spot and improve any weak points.

How can I improve my style and fix the readability of my text? There are still gaps in my text—what now? What’s the simplest way to explain my idea? What is this text about on a basic level?

Chunk 9
Would you still ask it this way?

Chapter 8: “Process Awareness”

Reflect on your writing

Use AI not only to write, but also to better understand your writing process and evolve your craft.

How do I tell my story?

Chunk 10
How can I actually ‘tell’ my story? Ready to try something new?

Chunk 11
Which writing techniques help move things forward? How can I ‘grow’ the way I write & work with AI? What have I learned from writing with AI? What does my writing process really look like?

I. p. 48, Chunk 19: KINDLAB (HS Trier)
I. p. 46, Chunk 7: KINDLAB (HS Trier)
I. p. 48, Chunk 15: KINDLAB (HS Trier)

Chunk 12

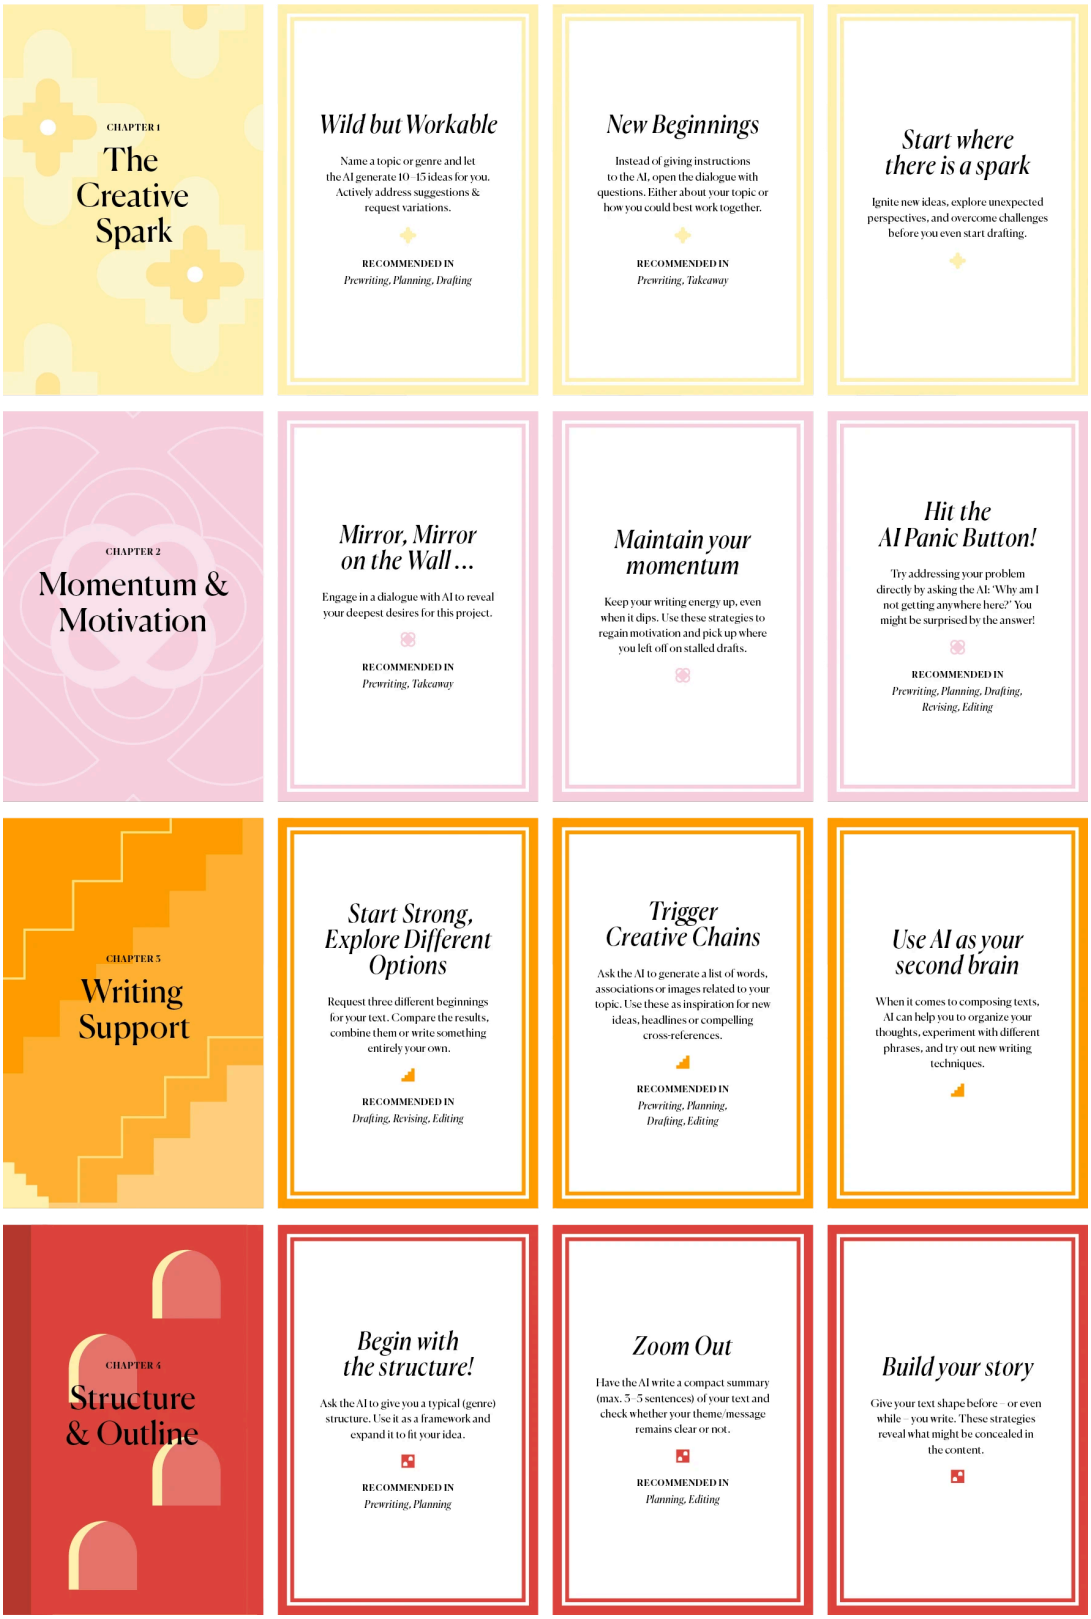

I. p. 58, Chunk 5: KI Labor (HfG Offenbach)
I. p. 46, Chunk 7: KINDLAB (HS Trier)

I. p. 48, Chunk 19: KINDLAB (HS Trier)

Chunk 13

CHAPTER 5

Shifting Perspective

Explore alternative views

Switch perspectives, tones, or gain (fictional) insight from your audience to unlock new depths in your writing.

Design for Your Reader

Try simulating a conversation with your target group. Encourage the AI to slip into the role of the person you are writing for.

RECOMMENDED IN
Prewriting, Planning, Revising, Editing

Style Remix

Generate stylistic variants of a paragraph. Experiment with different language styles, such as factual, emotional or provocative, or try new narrative perspectives, for example the omniscient or first-person POV.

RECOMMENDED IN
Prewriting, Drafting, Revising

CHAPTER 6

The Reflective Mirror

Vibe Check

How does your text come across? Factual? Ironic? Ask the AI. Allow it to highlight stylistic inconsistencies, but don't let it rewrite your text.

RECOMMENDED IN
Editing

Shapeshifting

Transfer your text to another medium with the help of AI. This will spark your creative energy and show you where your text can be condensed or developed.

RECOMMENDED IN
Planning, Editing

Challenge Your Thinking

Feed the AI with your central arguments or input a text excerpt – then ask it for a well-founded counter-position to your text.

RECOMMENDED IN
Prewriting, Planning, Revising, Editing

CHAPTER 7

Revision

Bridge Building

Get a list of transitions between two paragraphs/plot-points and choose one of them (no matter how generic they are). This will help you overcome blockages and get into writing again. You can always come back later and adjust it.

RECOMMENDED IN
Drafting, Editing

Polish Like a Pro

Contact the AI for specific suggestions on word choice, reading flow or expression. You can also ask for alternative ways of expressing something if there are phrases and word repetitions that bother you.

RECOMMENDED IN
Drafting, Revising, Editing

Use the Prompt Time Machine

It is beneficial, not only to revise your own text with a little distance, but also to revisit your prompts. Ask the AI what you can do better next time.

RECOMMENDED IN
Revising, Takeaway

CHAPTER 8

Process Awareness

When You (& Your Prompts) Need Therapy

Involve the AI in your reflection process by having it ask you questions. These should not only relate to the writing itself, but also to how you can enhance your prompts when working with AI.

RECOMMENDED IN
Takeaway

Shake Up Your Writing Routine

Consult the AI for advice on writing and compositional methods. Feel free to choose, change or combine any of them – the important thing: get writing!

RECOMMENDED IN
Drafting, Editing, Takeaway

Sketch the Process, Not Just the Page

Retrospectively describe your writing process and then have it summarized. This is ideal for reflecting on projects or documenting creative processes.

RECOMMENDED IN
Takeaway

↗ p3d425b5

Notes Toward AI Un/Learning

Simon Maris

II. p. 182, Chunk 1: SandwichNet
II. p. 142, Chunk 5: Editorial: Physical AI
II. p. 28, Chunk 1: Offenbach

Chunk 1
The first thing to unlearn about AI is that it's intelligent. The second is that it's artificial.

II. p. 4, Chunk 5: Vorwort

Chunk 2
What remains is more interesting than either assumption.

Every design school now teaches prompt engineering. Few ask whether learning to speak machine makes us forget how to speak human. I whisper punctuation marks full stop.

The most radical act in AI education might be teaching when not to use it.

Chunk 3
Every project promises *critical engagement*. Few deliver on the critical part.

Students who've never mixed paint are generating paintings. Writers who've never revised are publishing first drafts.

Chunk 4
The tools work too well for our own good.

III. p. 259, Chunk 3: Standorte & Team
II. p. 180, Chunk 6: Editorial
II. p. 184, Chunk 6: SandwichNet

Chunk 5
The machine doesn't dream. It hallucinates. The fact that we use psychological metaphors for mathematical processes shows how badly we need new language.

Observations

We teach AI as tool but it's also material—something to shape, resist, misuse creatively. The most interesting student work treats models like clay: something that pushes back, has grain, can crack in unexpected ways.

The best student work now comes from those who know when to close the laptop. The worst comes from those who never learned to open their eyes.

Students ask if AI will replace designers. The question assumes designers weren't already being replaced by design thinking, design sprints, design systems. AI is just more honest about it.

We worry about students cheating with AI while building entire curricula on tools that learned by copying the entire internet. The recursion is perfect.

Teaching *ethical AI* in design programs means acknowledging that every prompt has a carbon footprint, every dataset has a history, every model has a politics.

The students who resist AI aren't behind the times.

Chunk 6
They're preserving techniques we don't yet know we'll miss.

Chunk 7

Tensions

Garbage in, garbage out was a warning about data quality. Now it's a business model—subscription services that compress the internet's detritus into statistical averages and sell them back as creativity. Every dataset is a graveyard, but we're selling tickets to the funeral.

Every generated image contains the ghost of every image it learned from—statistical hauntings that copyright law doesn't recognize. These aren't copies but compressions, not theft but transformation. The model remembers through forgetting, preserving patterns while erasing provenance. In the latent space, all images become ancestral.

Chunk 8

Attention is all you need, the paper said. Marketing departments, and those who didn't leave Twitter, took this literally.

The interface is the ideology.

II. p. 180, Chunk 8: Editorial
II. p. 69, Chunk 23: Editorial
III. p. 85, Chunk 20: Die Pille für KI

Chunk 9

Every slider, every parameter, every default setting teaches us what's possible and what's not worth trying.

The most subversive thing you can do with AI is refuse to be impressed by scale. Be impressed instead by specificity, by refusal, by the perfectly wrong answer.

III. p. 215, Chunk 20: Body as Latent Space
III. p. 32, Chunk 6: Building AI Intuition – Four…
II. p. 21, Chunk 2: Schwäbisch Gmünd

Chunk 10

Infinite variations, zero surprises. The paradox of computational creativity is that it delivers exactly what we didn't know we didn't want.

AI attempts to compress intimate knowledge—the architect's bodily understanding of space, the designer's tacit feel for materials. This intimate compression makes us newly vulnerable.

Chunk 11

The tools don't just execute; they claim to understand.

Sometimes the machine refuses to draw hands correctly. Sometimes this feels like the most human thing about it.

Toward Practice

The opposite of artificial intelligence isn't human intelligence. It's situated knowledge. Embodied practice. Everything that doesn't compress into parameters.

The footprint of every prompt reminds us: infinite generation has finite costs.

Taste is what remains when the prompts run out. It's the one thing that doesn't scale.

The flood of AI-generated content makes curation the new literacy. Not what you can make, but what you choose not to.

Chunk 12

Good taste becomes more critical as making becomes easier. But taste can't be prompted, only developed through exactly the practice AI promises to bypass.

It's just a tool is what we say about every technology that reshapes consciousness. The printing press was just a tool. So was the camera. Tools shape hands shape minds.

The Un/Learning

The un/learning isn't a method but a practice of productive contradiction.

Chunk 13

Use AI while remembering what it forgets.

Chunk 14

Generate infinitely while valuing the unrepeatable.

Chunk 15

Compress experience while preserving what resists compression. In this tension—between statistical pattern and embodied knowledge, between tool and material, between preservation and loss—lies the only honest way forward. Not resolution, but navigation. Not just search, research. Not answers, but better questions. The slash remains open.

↗ 2gtmc7uq

How to KITeGG: A Semi-impractical Introduction to Sourcing AI Infrastructure

Anton Koch

Introduction

Chunk 1 The KITeGG project created its own hardware and software infrastructure to support the integration of generative AI into the design curriculum at five German universities for applied science. This article aims to support a reproduction of the steps taken to establish, maintain, and develop this infrastructure.

Chunk 2 It is not a simple step-by-step guide, however; it provides a general framework for building self-sufficiency in planning, deploying, and maintaining similar projects. Especially for those with limited experience and knowledge regarding high-performance computing (HPC), graphics-processing units (GPUs), and Kubernetes orchestration software, it should provide a starting point for personnel with basic system operations (SysOps) engineering experience. This article can be seen as a guideline for technical staff at higher education or research institutions looking to source AI or machine learning infrastructure for small to mid-range projects. It is not meant to provide a strategy for commercial deployment or large-scale, critical infrastructures with very different demands regarding stability, safety, and shared low-level access.

Chunk 3 It also does not link to repositories or external resources, as these will be subject to constant development and change. The intended audience should thus be proficient in searching online repositories to follow up on questions and deepen their knowledge of the concepts discussed in the article.

Guiding Principles

There should be a set of guiding principles for the overall strategic decision-making process. In our case, a core tenet was to build on open-source and free software with a sizeable community as much as possible.

Chunk 4 The only exception here was the hardware platform: NVIDIA's HGX server architecture. The reason for this was that the prevalence of NVIDIA's CUDA software platform for developing and running machine learning applications on a hardware-accelerated platform was overwhelming at the time of planning. Instead of its main competitor, AMD's ROCm platform, it was chosen to allow students and teachers to use the many existing models and projects based on CUDA. However, as the rest of our platform was constructed exclusively on open-source software and generic server components, it would also be viable to modify the cluster, adding nodes running other hardware or conceiving of a similar system running an entirely different GPU platform, but within the same cluster.

Chunk 5 Evaluating open-source libraries, tools, and components is not always straightforward, as the usual metric of popularity on GitHub (*stars* ranking) does not always indicate the best fit for the task at hand.

Chunk 6 To make an informed decision, a practical evaluation was necessary for some basic components, such as the distributed filesystem storage, by installing and trying it out.

Chunk 7 It is also beneficial to review the time intervals between releases,

the semantic versioning status, the number of open issues, and response times and rates.

Chunk 8 This helps select a repository with active maintainers and a community.

Chunk 9 It helps to invest a sizeable amount of time in researching the right tools for the job.

Chunk 10 There are no *one-size-fits-all* solutions to establishing such a complex platform, so individual research is key.

Another thing to keep in mind is the end of life (EOL) of both the hardware and, especially, the Kubernetes orchestration software's versions (see below). In terms of the hardware, we have been able to run a project time of five years on the same hardware specification without the need for upgrades.

Chunk 11 However, we did have to replace parts (including the very expensive GPU boards); so it is crucial to purchase hardware with an appropriate warranty to cover the envisioned project duration. In our case, we found that the total cost of buying and running the hardware at our university was about three to four times cheaper than outsourcing it with Microsoft Azure (including the up-front payment discounts).[1] However, this only holds if there is no need for more spending on replacement parts or entire nodes.

Chunk 12 Lastly, the type of GPU (A100) provided us with enough resources for the project runtime, but there had to be adjustments in the GPU memory allocated for users as the requirements grew over time. Generally, a five-year span seems like a good guideline, as Microsoft has steadily extended its hardware refresh interval for its data centers, raising it from four to six years in 2022.[2]

Chunk 13

Basic Concepts

Before proceeding with the actual planning, there should be some common knowledge of the fundamental concepts underlying cluster architecture and server infrastructure.

Chunk 14

Cluster on Bare-metal versus in the Cloud

While a cluster refers to a general assemblage of networked server hardware commonly found at a shared physical location, this can be different if the hosting is done in a so-called *cloud* environment. While we were deploying our cluster on *bare-metal* infrastructure (meaning directly on dedicated hardware) in a nearby, physical datacenter run by Mainz University, this can also be achieved by using components provided by companies, such as Microsoft Azure, Amazon AWS, or Google Cloud. These providers allow for different shared, virtual, or dedicated physical servers to be linked together within their global networks of data centers using proprietary products that provide services, such as data storage or networking. While this has proven to be much more expensive in the long run and can also create functional dependencies on proprietary services, it can still benefit smaller, short-lived projects that would otherwise not be able to afford the required hardware at the proper scale.

It is important to remember that a bare-metal cluster deployment and a cloud deployment are not entirely the same or interchangeable. While Kubernetes would theoretically allow for the same containerized application setup to be deployed on a cloud provider's infrastructure or a self-hosted, bare-metal infrastructure, many differences in network setup and access to physical volumes and other hardware require installing a variety of software packages. This guide focuses exclusively on deploying bare-metal infrastructure, but the guidelines could be adapted to a cloud deployment with additional research.

1: Koch, A. (2025). Der KITeGG Cluster – eine Infrastruktur für KI in der Gestaltungslehre. *Un/Learn AI, 2*. https://doi.org/10.25358/openscience-11844

2: Sharwood, S. (2022, August 2). Microsoft extends life of cloud servers from four to six years. *The Register*. https://www.theregister.com/2022/08/02/microsoft_server_life_extension/

Container Orchestration

The core software component of the cluster is the Kubernetes open-source software package. This application, or rather a collection of applications, originated at Google as a way of managing large application deployments on many servers using containerized applications (e.g., using Docker). It has since been donated, just as the Docker container standard, to the Cloud Native Computing Foundation (CNCF), which is part of the Linux Foundation. It consists of a database that holds configuration and resource definitions alongside an extensible collection of controllers and supporting services that watch the configuration and attempt to satisfy the requested state of the cluster through automatic provisioning of resources, hence the term *orchestration*.

Chunk 15 This allows for a dynamic structuring of application deployments, as the assignment of running containers or available objects and resources is done according to selectors that choose from the available servers in a way that best matches the required configuration.

The controllers are written in the Go language and can be extended and built by third parties, providing custom and specialized deployments. This means that there are numerous applications readily available to be deployed to the cluster simply by specifying the necessary configuration and then managing updates by modifying the configuration.

Chunk 16 The available services range from user-facing applications to internal services that provide control over the cluster networking or distributed filesystems for use by other containers. There is also the possibility of creating so-called *auto-scalers* that watch specific metrics for apps and services and scale them automatically to satisfy momentary increases in requirements.

However, the great flexibility of the system and its complex, concurrent nature make it highly challenging to debug for users with little practical experience of the system's inner workings.

Chunk 17 This is why it is fundamental to read up on the software's design principles and basic functionality, and to patiently study the various components' logfiles if anything goes wrong. Luckily, the community around Kubernetes is quite large and very active. This means that for most generic problems, there is likely a solution to be found, and even if not, it can usually be worked out with the help of community members.

Chunk 18 Containerization and Virtualization

A common source of confusion is the layers of abstraction on the server itself. The term bare-metal only refers to the fact that the installation runs directly on dedicated server hardware within the operating system (OS). The question now is what application is primarily running on the hardware.

Chunk 19 In our case, we installed Kubernetes and a container engine (in our case, ContainerD) that would start and stop containers and provide their runtime.

The container should not be confused with a server, however, as it just runs an isolated process alongside its required dependencies and libraries. While the containers are mostly based on a specific Linux distribution, they do not represent a bootable system using *run levels* (the various stages of a Linux distribution's startup), but a collection of distribution-specific libraries required to run the containerized application process. This means that a process ultimately still runs directly on the host with only certain limitations to its capabilities within the host system. This should not be confused with the method of *virtualization*, which allows for complete Linux distributions (or other OSs) to run in virtual machines (VMs) on the host. These virtual machines (usually called *guests*) can be assigned entire or just parts of the host system's

resources (e.g., a slice of the available random-access memory (RAM) or passing through a device connected to the Peripheral Component Interconnect Express bus (PCIe), such as a graphics card) but are otherwise separated from the host's OS.

While the two concepts of containerization and virtualization are fundamentally different, they can be combined to allow a single, large server machine to be partitioned into multiple smaller VMs that then appear as separate cluster nodes available for scheduling containers by the orchestration software. This is usually an essential part of cloud-hosting providers' application setup, as it allows for the dynamic provisioning of virtual private servers (VPS) for clients. However, it introduces another management layer, as the administrators need to manage both the host system(s) and the guests. In our case, we directly ran the containers on the host systems as there was no need for shared access to the nodes. The server nodes were already specifically tailored to their respective function, and further partitioning would only have introduced unwanted complexity.

Planning and Specifying the Cluster

The most important initial process is planning the overall capacity and components to be commissioned for the cluster.

Chunk 20 This is not to be taken lightly.

Chunk 21 For a project duration of almost five years as was the case for KITeGG, the three main factors to be balanced are the expected number of maximum concurrent users with their specific demands, the allowed spending limit, and the fact that the capacity will likely decrease over time as demand by contemporary software implementations constantly rises. This cannot be fully anticipated, but it helps to keep in mind that one should overcommit within the available financial budget instead of underestimating future demands.

Chunk 22 The following passage describes the main questions that should be answered during this process to arrive at an informed decision about the necessary components to be tendered.

Overall Capacity Planning

Regarding capacity planning, the key was to estimate how many students across all partner institutions should be able to work simultaneously within the cluster using different GPU profiles, but essentially with the same basic combination of a central processing unit (CPU), RAM, and disk space.

As the model for assessing each user's requirements, we used the specifications for an average GPU workstation found in a computer pool at the university. While the cluster would also need to host additional services, these should be relegated to other nodes with cheaper, more generic server hardware, which would be deployed separately. We also assumed that not all users would be constantly consuming their allocated capacity and that occasional heavy usage would be cancelled out by rather less active users. The capacity for a single user was set at a minimum of four concurrent CPU threads, 32 gigabytes (GB) of RAM, and at least 200–300 GB of disk space. As the project was supposed to provide GPU access to select courses that would not all run in parallel, we set the number of concurrent users at 40 for each institution, resulting in a maximum of 200 concurrent users across the cluster.

These users would be spread out over so-called *worker* nodes, of which each institution would have to buy one to satisfy the budget structure.

Chunk 23 As the project lead, Mainz would buy the additional hardware necessary to provide the base structure for the overall cluster deployment. Making the above calculation for a single node would mean multiplying the assessed

single-user capacity by 40, requiring at least 80 physical CPU cores (with hyperthreading), ca. 1 terabyte (TB) RAM, and about 12 TB of net available disk space. This, however, would be capacity reserved for users, so there should be additional capacity for system processes and services, such as virtual block devices providing the disk space to users, ideally in a distributed fashion.

Chunk 24 These operations, especially the virtual filesystems, should not be underestimated, as they can fundamentally slow down general operations if not provided enough resources.

To provide for enough headroom, we aimed to double the required minimum such that concessions could be made whenever the double capacity would exceed an easily available configuration. This meant that, for CPU cores, we settled at two CPUs with 64 physical cores each (providing up to 256 parallel threads through hyperthreading) at a little over 2 GHz. The CPU frequency was deemed to be of lower priority, as the bulk of computing would be done using GPUs, and thread concurrency was more critical than individual thread speed. RAM capacity was settled at 2TB, and disk space was calculated with triple replication in mind, providing redundancy to prevent data loss and to guarantee uptime of volumes during node downtime. This was then doubled to allow for headroom, arriving at a capacity of at least 72TB net disk space for each node, which should be provided using fast solid-state disks (SSDs) over the non-volatile memory express (NVMe) bus, in addition to the two redundant system SSDs with 512GB capacity each.

General Infrastructure Planning

In addition to the worker nodes hosting the services directly used by the students, the cluster needed several components that supported its overall functionality and would provide general services, such as internal APIs for the learning platform, authentication services, backup and long-term storage, and the basic coordinating services run by Kubernetes itself.

Chunk 25 A Kubernetes deployment usually uses a dedicated node called a *control plane* that exclusively runs the central configuration database (a service called *ETCd*, which can be replicated and hosted entirely externally for the control plane node) and is otherwise unavailable for scheduling. This node usually does not need excessive CPU, RAM, or disk space resources. In our case, we decided to extend the role of the control plane to host the core applications making up the KITeGG platform and to use it as a proxy for exposing services to the World Wide Web (WWW) via a dedicated 1 Gigabit (Gbit) line. To allow for the extra services, we raised the specifications above what would typically be sufficient to run a Kubernetes control plane node. The specifications were set at 64GB RAM, a single CPU with 16 physical cores, and two redundant 512GB SSD drives.

We decided to buy two identical nodes to be able to host two simultaneous Kubernetes installations, to which worker nodes could be connected interchangeably, so that we would be able to experiment with one installation while production was running on the other. Once a stable, high-availability setup is preferred over experimental flexibility, these two nodes could eventually become redundant control plane nodes. In addition to the control-plane nodes, two redundant servers dedicated to secondary storage for backup and object stores were commissioned. These were each configured to provide two CPUs with 64 physical cores each, 256GB RAM, two redundant system SSDs (521GB), and twenty hard-disk drives (HDDs) with 16TB capacity. Networking within the cluster would run at 10 Gbit over copper with a dedicated manageable switch and two interfaces per machine. Internet connectivity was provided to both control planes over a 1 Gbit line for each node.

An interesting addition to the network is NVIDIA's proprietary InfiniBand standard, which is tailored to their datacenter server products. It allows for up to 200 Gbit of bandwidth on a single connection and, while primarily used for multi-node compute operations using remote direct memory access (RDMA), can also provide internet protocol over InfiniBand (IPoIB), which can support better throughput for distributed filesystems.

It must be noted that this configuration is an example and can hardly ever be precisely calculated for an explicitly experimental and, to some extent, open-ended project.

Chunk 26 It is possible to run a reduced setup without the dedicated storage nodes, on a single control plane, or even a combination of storage node and control plane. Still, it should be generous enough to allow the servers to withstand spikes in usage and not put too much strain on the hardware by constantly running at a very high load. It is also important to consider eventual limitations in the housing environment at the datacenter where the hardware is installed. Essentially, it must be clear that the servers need a fully climate-controlled environment within an enclosed server cabinet featuring its own air conditioning. For large GPU servers, it is also important to allow for ample depth and height that the cabinet can supply enough power. Usually, these requirements limit the housing options to a dedicated high-performance department. They are not necessarily found in the usual information technology (IT) facilities that house an institution's files and web servers.

Choice of Graphics Processor

An important choice is what GPUs to use. As the KITeGG project represents a different take on the classic high-performance computing (HPC) model with its interactive users' sessions and hosting of permanent services over the traditional job queue, it has different requirements regarding available resources. While HPC typically works by allowing users to define jobs to be executed non-interactively, allowing the jobs to continuously use up all available compute resources, the KITeGG project had to provide users with an interactive development environment.

When we started planning the KITeGG project, NVIDIA processors dominated the landscape of available integrations, models, and general ease of use with popular frameworks, due to the software development kit (SDK) CUDA.

Chunk 27 The company offers a broad range of different graphics and tensor processing units, with the A100 model (eventually superseded by H100 and B100, as well as the larger B200) at the top tier of their datacenter offerings. These cards, apart from offering the most GPU memory, have the benefit of partitioning with the multi-instance GPU (MIG) feature. MIG allows for the creation of partial slices, ranging from roughly a seventh to half of the GPU's specifications. This allows for creating up to seven small 11GB GPUs from one 80GB model that can be used for less resource-intensive tasks, such as using Stable Diffusion. This is a significant benefit, as providing basic GPU access to more students can become massively expensive. For example, providing our aforementioned goal of 40 concurrent users per worker node with access to a minimal dedicated GPU (without users sharing a GPU) would require 40 GPU units or five A100 80GB GPUs split into seven 11GB slices each.

The flagship datacenter processors are available as PCIe models but are more commonly found in a server system using NVIDIA's proprietary DGX or HGX architecture.

Chunk 28 While DGX is exclusive to NVIDIA, the HGX architecture allows other manufacturers to build server systems bundling multiple processors using the proprietary server PCI Express module (SXM) bus instead of PCIe. This architecture enables a server to be equipped with up to eight

A100 processors that are interconnected using NVIDIA's high-performance NVLink bus with a bandwidth of 600GB per second. This configuration allows for using all eight GPUs together or splitting them into up to 56 MIG slices. This setup provides the most significant performance headroom but is also the most cost-intensive. NVIDIA's DGX systems provide the user with a dedicated Linux installation based on Ubuntu alongside proprietary libraries and software kits, but at the time of planning, they were almost twice as expensive as a third-party HGX system.

While these systems are NVIDIA's top offerings, smaller configurations and cheaper single datacenter GPUs are available, but these lack the MIG capabilities and, as such, would require users to share GPUs or reduce the number of concurrent sessions. However, they can still prove a viable alternative for comparably low-cost scenarios using rack-based servers that host multiple PCIe cards equipped with workstation GPUs or even consumer-grade gaming hardware. These would still integrate into the cluster with the same software setup, but the cheaper options do not allow for explicitly targeting different models for scheduling or providing multi-instance GPU slices.

A major competitor to NVIDIA in the GPU market is AMD, which provides a datacenter line, workstations, and consumer models, as well as ROCm, their dedicated GPU software stack. While these cards are not necessarily inferior to their NVIDIA counterparts, the support for the ROCm software stack is much more limited vis-à-vis the many existing models and open-source machine learning tools. If, as with KITeGG, the focus is more on allowing students to work with existing and upcoming projects, then NVIDIA will still be more immediately accessible. However, if the focus is more on training, original programming, or data-science work, less cost-intensive AMD hardware and software is adequate.

Chunk 29 It is also important to note that the cluster does not have to use a single type of worker node across the entire setup.

Chunk 30 It can just as well be a mix of any of the server types described above to cover different use cases.

Cluster Lifecycle

Buying servers can be cheaper for projects looking at a time span of around 3–5 years, but there is a catch. Major cloud providers are setting their average expected server hardware lifespan somewhere between three and six years. This means it is likely that new hardware will eventually be needed to replace failed or outdated systems, especially when aiming to run the acquired infrastructure for extended periods.

Chunk 31 For an enclosed, third-party funded university project, such as KITeGG, the funding proposal did not include replacement hardware. Thus, we had to request a five-year warranty contract for the cluster hardware to ensure it would run at full capacity at least for the project timespan. After the warranty ends, replacements might become impossible, as there is no running budget for the cluster beyond the federal project funding.

However, as we have found that buying and self-housing the server hardware was significantly cheaper than outsourcing it to a commercial partner, decommissioning the infrastructure at the end of the project is still an acceptable outcome, which would have been done by a commercial hosting provider anyway. Given that there is at least an agreement on power usage and basic administration, it can also be kept running for as long as possible or stripped for parts and reused by university IT or other third-party projects. Still, it must be clear that the remaining hardware can fail at any time and should not be counted on for mission-critical use cases. Luckily, the cluster software setup is scalable and migratable, and thus could be scaled down, transferred to core IT

services at the institution, or moved to outside hosting providers, either partially or entirely.

Setting up the Cluster

Initially, getting the cluster up and running is divided into three basic phases, each described in the following section. First, the hardware needs to be brought online, then Kubernetes and its dependencies need to be set up, and, finally, a collection of services and components needs to be added before anything can be run on the cluster.

Chunk 32 Bare-metal Hardware Setup

When installing the servers, be sure that airflow direction is the same for all installed devices and that gaps in the height of the cabinet are closed off by blind covers so that ideally only cooled air is sucked in at one end and the heated air is not circulating uncontrolled. You can install the GPU servers at the bottom, as they are the longest and heaviest (at least if you're using the large HGX and DGX devices), so they will be easier to get to.

Chunk 33 We set the storage nodes on top of the GPU nodes with the control planes at the top, just before the ethernet and InfiniBand switches at the top of the cabinet.

Chunk 34 Once the hardware is in place, connect all necessary wiring and make sure to leave room for pulling out trays and modules from the servers.

Chunk 35 If possible, use different colors for management and general network connections, and bundle cables that belong together to better guide them around the cabinet. Also, connect redundant power supplies to separate circuits and breakers, so if one goes down, you do not lose power to the node entirely.

After checking your setup, you can begin starting up the servers, starting with the control plane nodes (or at least one of them) and installing the OS, if that has not been done by your hardware vendor, as was the case for us.

Chunk 36 If you choose to use two redundant drives for your system partition, OS setup becomes more challenging. We are using Ubuntu as our preferred distribution as it is based on Debian.

Chunk 37 It is a bit simpler to maintain and has a large and active community.

Chunk 38 You should be able to find more than one how-to describing a system setup using the multiple device driver (MD). This driver allows for the creation of software-based redundant arrays of inexpensive disks (RAIDs) and virtual disk devices based on multiple physical devices that provide replication of the disks for redundancy.

Furthermore, you should be clear about your desired network topology and how it might integrate with the surrounding infrastructure.

Chunk 39 In our case, the cluster uses an entirely private network that receives internet access through network address translation (NAT) on the control planes. These are placed in the perimeter network, also known as the demilitarized zone (DMZ), and otherwise entirely isolated from the rest of the university network. The only ports to be opened to the public Internet are 22 (SSH), 80 (HTTP), and 443 (HTTPS), with 6443 (Kubernetes) restricted to specific sources within the university network. This way, select applications can be proxied out to the web, while the bulk of the network traffic stays within the private cluster network. You can then use the control plane nodes from the public network as so-called *SSH-jump-hosts* using port forwarding to connect to the nodes in the private network. It is essential to lock down the public-facing nodes using firewall rules, as leaving any non-standard ports open or your Kubernetes installation directly accessible over the internet will eventually cause you to get flagged by your university IT or

even by a government agency monitoring your networks for threats, and this could mean that your cluster will be disconnected or shut down.

Your Linux distribution should be installed as a minimal install with only the core functionality and an SSH server for remote connections. Chunk 40 Only then, should you add packages as they become necessary, because any excessive extra software will not only take up disk space and make reviewing which packages to upgrade more tedious but also entail potential vulnerabilities. Chunk 41 The software should be automatically updated for critical security fixes, but manually for everything else so that you can check before applying breaking or otherwise unwanted upgrades. Manual upgrades should be done as regularly as possible –not every few weeks apart from security updates. Occasional kernel patches might require you to restart the nodes, so you can also squeeze in a minor upgrade.

Server configuration can be automated by a host of available solutions, such as Ansible, Terraform, Puppet, or Chef, to name a few. This has the upside that you can define specific installation profiles and then automate setup and changes centrally, reducing the risk of human error. This is especially true the more servers you must keep managed in a similar or identical setup. Still, there is a cost to learning these tools. If you are new to this and are only running a few servers, a good strategy could also be to first learn how to set up the servers manually, experiment with your setup, and once you'd like to *harden* the installation, move to one of the automated solutions.

An important decision to make at this point is whether you want to make use of virtualization techniques to partition the servers further. This offers more flexibility to try various node set-ups without going through the entire OS setup again if you choose to start over. This way, you can have multiple approaches running side by side, e.g., a development control plane next to the production version on the same hardware node, if there is no option of having two identical servers for this. Another benefit is that for cheaper multi-GPU systems, each consumer-type card can be exposed to a different virtual machine (VM) to make it explicitly addressable. The computational overhead added in virtualization today is negligible, and the only real downside is that it introduces another technology you will have to learn to use and maintain. In case of Ubuntu, this is usually the kernel-based virtual machine (KVM).

Chunk 42 Setting up Kubernetes

This core software suite enables you to run a distributed cluster environment. We will not get into the history of Kubernetes, but it should be noted that, while the software originated at Google to allow for massive application deployments spanning hundreds if not thousands of machines, it has now become ubiquitous in many other distributed computing environments. For example, it can also be used to establish a cluster that can include small, embedded devices, create a cluster to be used in robotics, and extend datacenter clusters with edge-computing devices. In its most basic setup, it does not consume much computing power and merely provides a support structure for running and scaling any distributed application. Chunk 43 Before you go ahead and install Kubernetes, make sure to read the documentation carefully to understand how it is structured and what role the various components have in its overall functionality. Your worst enemy in debugging a malfunctioning cluster is a total lack of, or maybe even worse, half-knowledge of the moving parts involved in its underlying processes, as we have learned from first-hand experience, naïvely jumping right into the cold water. As the number of services and components in your cluster grows, so does the level of complexity due to concurrency.

Chunk 44 This will require you to pull together information from many log files, set the right components to a log level for debugging, and retrace side effects, even if only to be able to communicate a meaningful issue report in a repository.

Kubernetes is often perceived as daunting and somewhat opaque, which might paradoxically be related to the relative simplicity of its core APIs.
Chunk 45 On its own, it really does what it is in its name: to steer or orchestrate the lifecycle and execution of processes and tasks according to specified requests, constraints, and specifications.
Chunk 46 This means, however, that its initial setup is a far cry from a workable application deployment. To provide various storage, network, and service layers, it relies on third-party projects implementing functionality through standardized interfaces, such as the container network interface (CNI) or the container storage interface (CSI). It also abstracts away the container engine within which it runs your containers, so that you can choose among various available engines.
Chunk 47 This entails many decisions, made either consciously by reading up and doing research or by just going with some default options or whatever is recommended in the tutorial you might be following.
Chunk 48 Consequently, you may end up with problems that seem related to Kubernetes, but your extensions could just be misconfigured or not behaving as expected.

Sometimes, though, as was the case for KITeGG, you want to get going at some point—and you may also be working with extra setup time and some leeway for making mistakes, tearing everything down, and starting over. Good old reading is the most important foundation for getting your setup off the ground. To reiterate our warning from the beginning, just copying and pasting whatever comes up in a web search or is recommended by a chatbot might get you going—but might ultimately leave you clueless once there are problems or strange side-effects.
Chunk 49 The best you can do is to read different articles, compare suggested repositories in terms of maturity (version history), community size (open issues vs.
Chunk 50 closed, stars ranking on GitHub), and, most importantly, the necessary features.
Chunk 51 After picking your favorite, you should set it up and try it out to get some hands-on experience.
Chunk 52 This is crucial, not least since you will have to maintain your installation, and if something goes wrong with your filesystem during production, things can get a bit hectic if you do not really know the system you are working with.

To manage your cluster, inspect its workloads, and make changes, the command-line interface provided by Kubernetes' kubectl is a powerful and versatile tool that any Kubernetes admin should be familiar with. Still, there is a handy desktop application called Lens[3] that allows for a more comfortable interaction with the cluster. It is free for smaller projects and well worth trying. There might be more accessible ways of setting up and managing Kubernetes installations, but in the context of this project, the combination of kubectl and Lens proved to be sufficient.

3: (n.d.). *Lens | The Kubernetes IDE*. Lens x Mirantis. Retrieved June 10, 2025, from https://k8slens.dev/

Kubernetes Basics and Essential Components

In general, there are many tutorials out there, but it makes sense to go with the official setup instructions for Kubernetes, which also contain links to the related external dependencies, such as the container engine and the network integration. This straightforward process was performed without any special exceptions to the recommended configuration.

Chunk 53 The additional, optional components should be selected with care, depending on the use case. However, while not an explicit recommendation or a must, this section details the software packages we settled on for the KITeGG infrastructure to give you a point of reference.

For the container engine, we went with ContainerD, because it is also included with the general Docker install and, as such, very widely used and supported. The setup is easy, but it should be noted that there are a few things to look out for, especially on the worker nodes with the GPUs.

Chunk 54 Suppose you have a relatively small system disk, as we had with half a terabyte. In that case, you need to consider that the cluster users might eventually use many different Docker images and fill up the temporary container snapshot space. This means you should consider mounting ContainerD's data folder (found at /var/lib/containerd) on a separate, larger SSD to give it more space. In addition, you need to check the GPU vendor's documentation on settings to be added to the container engine configuration that allow it to use a dedicated container runtime, as is the case for NVIDIA's GPUs (more on that later).

Chunk 55 Before Kubernetes starts scheduling its so-called *Pods*, the smallest unit of a Kubernetes deployment, networking must be available. While Flannel[4] provides the easiest and quickest way of adding networking capabilities to your cluster, we eventually moved to Calico,[5] which is almost equally easy to set up but allows for more fine-grained control over your cluster's networking. Once you have installed Kubernetes and added its network plugin, you will see that it immediately starts up its core containers and is now ready to set up additional workloads.

Cluster Storage

Chunk 56 Another essential element is the available storage systems. Kubernetes allows volumes to be created from a *Storage Class*. This is an object managed by a dedicated controller, provisioning storage on a specific underlying storage technology.

Chunk 57 It ships with local provisioning of storage resources by allowing direct access to the node's filesystem. This works for workloads that are always running on the same host, but it is not feasible to run on varying hosts, depending on their availability or replicating across multiple hosts, with a single file system. We used different forms of file system abstraction according to the various types of nodes on which the workloads were to run.

Chunk 58 The control plane nodes were initially not set up as redundant replicas and, as such, could use local persistence. We decided to use a local ZFS (formerly known as *zettabyte file system*) filesystem pool, which allows the creation of virtual block devices from local disk image files, single volumes, or multiple disks combined or replicated for redundancy. This makes it a versatile solution for resizing and extending storage options at runtime. There is an integration for ZFS in Kubernetes available for the popular OpenEBS Kubernetes storage project. It allows Kubernetes to

4: (2014, December 24). *Flannel-io/flannel: Flannel is a network fabric for containers, designed for Kubernetes*. GitHub. Retrieved June 10, 2025, from https://github.com/flannel-io/flannel

5: (n.d.). *Project Calico*. Tigera. Retrieved June 10, 2025, from https://www.tigera.io/project-calico/

dynamically provision storage volumes on demand in the underlying ZFS pool.[6] This is feasible for services that always use the same node and can be tied to this node, even if a node with similar capacity is attached to the cluster.

A more complex storage implementation is a distributed one that provides both redundancy and makes volumes available across nodes for more flexible scheduling of workloads. There are various solutions for distributed storage in Kubernetes, and for us, the choice ultimately was between Ceph[7] (via Rook[8] storage orchestration) and Longhorn.[9] Ceph is an established and widely used solution for distributed storage that is also available outside of the Kubernetes ecosystem. On the other hand, Longhorn was designed exclusively for Kubernetes and is a comparatively newer project. We ultimately decided to use Longhorn, primarily because it is much quicker to get going with than Ceph. We needed storage across the worker nodes as fast as possible, and the learning curve for Ceph was too time-consuming. This, however, should not be the only factor guiding such a fundamental decision. While both options are robust software, Ceph might be a bit more reliable in terms of issues with failing volume replica due to the way it distributively writes to all nodes as opposed to writing to one and then replicating the data written. Ultimately, we only encountered data loss from a faulty and unrecoverable volume in a single instance over the course of the project, and Longhorn seemed to reach a more stable state from version 1.7 on. It should be noted here that Longhorn is a viable solution for getting started quickly in a smaller cluster, especially on bare metal, while Ceph is the go-to solution for large-scale and cloud deployments.

Finally, we also needed a solution for long-term storage that would provide a level of redundancy and offer lots of space without the need for it to be particularly fast.

Chunk 59 To connect this storage to various environments, we opted for MinIO,[10] an open-source object storage solution compatible with Amazon's S3 storage interface. We placed MinIO on the two storage nodes. Although MinIO prefers a minimum redundancy of three replica nodes, we deployed it through the MinIO Operator using the mechanical drives and on the two nodes with a higher number of replicas to compensate for the smaller number of failover nodes.

GPU Access

As we went with the proprietary NVIDIA CUDA platform, getting the GPUs ready for use in Kubernetes was a straightforward process supported by software packages maintained by the company.

Chunk 60 The core software package managing GPU access is the GPU Operator,[11] an open-source Kubernetes controller that ensures the proper installation of drivers, scans the nodes for available GPU hardware to be used, and performs partitioning according to a supplied MIG profile. It publishes information about accessible GPUs in the cluster and allows users to start Pods that specifically request a number and type of GPU profile. These Pods run a special NVIDIA Container Toolkit that provides a dedicated container runtime that must be registered with the container engine. We opted to install the GPU driver and

6: (2019, November 7). *Openebs/zfs-localpv: Dynamically provision Stateful Persistent Node-Local Volumes & Filesystems for Kubernetes that is integrated with a backend ZFS data storage stack*. GitHub. Retrieved June 10, 2025, from https://github.com/openebs/zfs-localpv

7: (n.d.). *Ceph.Io - Home*. Ceph. Retrieved June 10, 2025, from https://ceph.io/

8: (2016, November 8). *Rook/rook: Storage Orchestration for Kubernetes*. GitHub. Retrieved June 10, 2025, from https://github.com/rook/rook

9: (n.d.). *Longhorn*. Longhorn. Retrieved June 10, 2025, from https://longhorn.io/

10: (n.d.). *MinIO | S3 Compatible Storage for AI*. MinIO. Retrieved June 10, 2025, from https://min.io/

11: (2020, December 8). *NVIDIA/gpu-operator: NVIDIA GPU Operator creates, configures, and manages GPUs in Kubernetes*. GitHub. Retrieved June 10, 2025, from https://github.com/NVIDIA/gpu-operator

CUDA toolkit within the worker node's host OS instead of delivering them through the GPU Operator, as we also wanted to access this functionality directly on the host and not just from within Kubernetes. Overall, this process proved extremely straightforward and streamlined, and it did not take much longer than the installation time for the software packages.

Chunk 61 The Quickest Quickstart

After getting the basic components up and running, we wanted to be able to work among the partner institutions immediately and use the available GPU resources. Initially, this was limited to university teaching staff, so we worked in a small group of people who actively experimented and would give feedback about issues, needed modifications, and requested features.

Chunk 62 It was now less than two weeks after the hardware was shipped, mostly spent trying out various options for basic components and reading up on documentation. At this point, we split the use of the cluster into two separate, but parallel strands of operation.

Chunk 63 We could now run a production environment alongside an experimental development path that would regularly result in upgrades and changes to the production environment.

In the beginning, we installed the software package JupyterHub,[12] which provides a multi-user version of the popular JupyterLab server environment in which users can manage assets, create so-called Notebooks using Python code, and use a Linux terminal to execute commands. With this environment, we were able to provide staff with immediate access to a basic development environment backed by the cluster's compute resources. Although we only began with the default JupyterHub setup, we quickly created our own JupyterLab image to add extra libraries and command-line utilities requested by users to ease development within the JupyterLab environment. We also immediately added the open-source authentication provider Keycloak[13] to the cluster to add OpenID Connect authentication to be used by JupyterHub and a general single sign-on (SSO) solution for later use.

Chunk 64 Once the involved partners were able to work independently, we began to set up a development operations (DevOps) environment.

Chunk 65 All development was managed within a group on the Gitlab instance run by the Rhineland-Palatinate state.[14] Project partners would set up repositories and track development issues there. They could use continuous integration (CI) workers deployed to the cluster, targeting a private container registry based on the open-source project Harbor,[15] also deployed to the cluster. The resulting images could then be deployed to the cluster to provide services from the public project websites and the custom-built teaching and learning platform, the so-called *Lehr-Lernplattform* (LLP), to custom web frontends for various machine learning models. At the same time, the JupyterLab servers could also be used to run long-running training processes and to evaluate models or external open-source projects.

Chunk 66 What Now?

Chunk 67 While the infrastructure and an initial operative environment could be established quickly, it was still a long way from the platform it developed into four years later. But the crucial takeaway from looking back at the platform engineering process is that a project like this requires at least one person to be solely dedicated to tending to the systems and, at the same time, continuously expanding their

12: (n.d.). *Project Jupyter: JupyterHub*. Project Jupyter. Retrieved June 10, 2025, from https://jupyter.org/hub

13: (n.d.). *Keycloak*. Keycloak. Retrieved June 10, 2025, from https://www.keycloak.org/

14: (n.d.). *GitLab*. Rechenzentrumsallianz Rheinland-Pfalz. Retrieved June 10, 2025, from https://rarp.rlp.net/services/gitlab/

15: (n.d.). *Harbor*. Harbor. Retrieved June 10, 2025, from https://goharbor.io/

capability in direct communication with its users. This job requires more than just administration according to a fixed procedure; instead, it requires monitoring the cluster's function using tools, such as Prometheus,[16] and prioritizing proactive and responsive communication with its users.

Chunk 68 It is a delicate balance between maintaining a healthy and safe production environment and engaging with staff feedback promptly and dynamically.

Chunk 69 This also means scheduling major changes and updates outside the running semester to avoid interruption of ongoing coursework. However, if a functional addition or modification promises to be valuable, a call must be made when the benefit outweighs the risk. After all, it is an experimental teaching environment and runs on a limited overall time budget, unlike core e-learning services found at a university. Here, Kubernetes' scheduling selectors and policy functionality allow operations to be separated sufficiently to enable dynamic management, yet also to provide a robust setup.

A significant benefit of Kubernetes is its API, which allows clients in many different programming environments to control its functionality programmatically, like scheduling and managing deployments or requesting information about the cluster state. For the KITeGG cluster platform, we eventually designed a simple hosting environment that allows staff to deploy custom services based on images built through GitLab and hosted in the private container registry.

Chunk 70 Users can configure a service through the LLP, which is backed by Directus and uses custom extensions to communicate with JupyterHub, Keycloak, and Kubernetes. The services can then be accessed from within the cluster or outside through a custom proxy that authenticates through the SSO provider.

Providing a complete description of the KITeGG cluster development would exceed the scope of this article, but that may also not be as important as the frame within which the project was developed. As every project is different, this article demonstrates the general ease of use when using contemporary server technologies and open-source projects.

Chunk 71 These allow for a relatively low-barrier, experimental entry and, given enough time for full engagement with the technology and moderate growth, the chance to develop into a mature platform that is stable and employs established best practices in security and deployment.

Chunk References

Chunk 11	II. p. 236, Chunk 8: Der KITeGG Cluster – eine…
Chunk 23	II. p. 241, Chunk 21: Der KITeGG Cluster – eine…

16: (n.d.). *Prometheus - Monitoring system & time series database*. Prometheus. Retrieved June 10, 2025, from https://prometheus.io/

Auf dem Weg zu KITeGG

Florian Jenett

Chunk 1 Im Wissenschaftsbetrieb gibt es wenig Raum für Backstories und Making-Ofs. Das ist eigentlich schade, denn gerade dabei könnte viel über die g*eheime Kunst* der Antragstellung und das Entstehen von Projekten wie KITeGG gelernt werden. Solche Quellen wären vermutlich besonders für Neuberufene oder Jungwissenschaftler*innen wertvoll und könnten den oft staubig wirkenden Wissenschaftsbetrieb etwas lebendiger machen. Sie wären zudem für die Forschung in der Gestaltung wichtig, da hier oft Vorbilder und eingespielte Prozesse (Antragsstellung, Betreuung von Forschungsprojekten, etc.) fehlen.

Chunk 2 KITeGG – zu Anfang hatte es noch nicht diesen Namen – geht auf verschiedene Stränge zurück, die sich letztendlich im Projekt getroffen haben. Das Team ist groß und auf fünf Standorte verteilt, so gibt es viele Wege, die zu KITeGG geführt haben.

Chunk 3 Nachfolgend beschreibe ich meinen.

Chunk 4 Im Jahr 2015 saß ich gerade an meiner Bewerbung auf die Professur *Medieninformatik und Gestaltungsgrundlagen* im Kommunikationsdesign an der Hochschule Mainz. Die Ausschreibung forderte neben den üblichen Unterlagen auch ein Forschungskonzept. Das traf mich unerwartet, bisher war ich noch nie aufgefordert worden über die eigene Arbeit als Forschung nachzudenken, besonders in dieser Form. Aus der Not eine Tugend machend formulierte ich zwei grobe Ideen: **„…** die Weiterentwicklung von Processing[1] in Teilen auch an die Hochschule Mainz bringen" und **„…** [eine neue] Ausrichtung der Motion Bank[2]". Meine Bewerbung war letztendlich erfolgreich und ich hatte das Glück mich in einem forschungsfreundlichen Umfeld wiederzufinden. Dank der Unterstützung der Hochschulleitung, anderer Forschenden und vor allem der damaligen Stabsstelle Forschung und Transfer, Frau Dr. Hartel-Schenk, konnte ich die formulierten Ideen weiterverfolgen und aufbauen. Von den zwei Ideen hat sich zunächst Motion Bank schneller entwickelt und ist heute ein zentraler Teil meiner Forschung an der Hochschule Mainz. Aus der Processing-Idee entstand zunächst eine kleine Veranstaltung mit dem Namen P*rogramme auf Papier* (Juni 2016), die die Frage nach dem Spannungsfeld von Gestaltung und Programmierung in der Lehre stellte. Neben einem Vortrag und Workshop von Karsten (Toxi) Schmidt[3] fand ein langer Austausch geladener Lehrender aus Bereichen zwischen Digitaler Gestaltung und Medienkunst statt. Diese Gruppe nenne ich gerne mal die *Processing Profs*, denn viele kenne ich als ehemalige Processing-Nutzer*innen, die den Weg in die Lehre an unterschiedlichen Gestaltungshochschulen gefunden haben.

1: Processing, ein Skizzenbuch für programmierende Gestalter*innen: https://processing.org/

2: Das Tanzforschungsprojekt Motion Bank: https://motionbank.org/

3: Toxis Internetseiten sind ein reichhaltiger Schatz: https://toxi.co.uk/ und https://thi.ng/



Abbildung 1: Plakat zur Veranstaltung Programme auf Papier. Gestaltung: Joshua Glunz, Tim Hau

Chunk 5 Auf einer Konferenz in Belgrad (Resonate) im Jahr 2014 versuchte mich ein befreundeter Künstler (Kyle McDonald) vom Lernen der Programmiersprache Python zu überzeugen. Er hatte Caffe[4] und Deep Learning (DL) für sich entdeckt und ich sollte mir das unbedingt auch mal anschauen.

Chunk 6 Dieses „AI" und seine Möglichkeiten waren mir nicht gänzlich neu, gab es doch bereits die ersten Künstler*innen die damit arbeiteten und mein Twitter-Stream deutete nicht zuletzt dank Deep Dream Puppy Slugs[5] bereits in diese Richtung. Auch begannen Felder wie die Computer Vision sich zunehmend auf den Einsatz von GPUs[6] und DL auszurichten.

Chunk 7 Ich stieß bei meinen Projektrecherchen für Motion Bank rund um Tracking und Recognition immer wieder auf das Thema KI. So auch auf unserem vierten *Choreographic Coding Lab*[7], das wir 2015 in New York zusammen mit dem ITP/NYU Tisch veranstalteten. Dort traf ich auf Lisa Kori und Gene Kogan, die zusammen an einem audio-visuellen Performance-System arbeiteten, das unter anderem KI verwendete. Ich glaube, das war mein erster direkter Kontakt mit KI im künstlerischen Einsatz.

Chunk 8 Das Thema rückte also immer näher, allerdings erschien mir die zeitliche und technische Investition noch zu hoch für einen tieferen Einstieg.

Ab 2017, ich war nun ein Jahr an der Hochschule Mainz, häuften sich die Beispiele von Projekten mit Einsatz von KI. Dank einer kleinen Förderung konnten wir für unsere digitale Werkstatt iLab einen ersten *KI-Rechner* mit Nvidia GTX 1080 TI Game-GPU anschaffen. Fast zeitgleich begann Google mit seinem Colab Dienst die technischen Anforderungen für den Einsatz von KI drastisch zu senken. Erste studentische Projekte beschäftigten sich mit generativer KI in seinen frühen Formen wie GANs (Twitterbot *Frenhover* von Lukas Rudig[8]) oder RNNs (Schriftgestaltung in *Path-finder* von Jean Böhm[9]). Das Thema war angekommen, aber der Einsatz der Technologie erschien immer ein riesiger Hack zu sein.

Chunk 9 Zu viele Python-Versionen, Packages, Frameworks (Keras, TensorFlow, Torch, etc.), fehlende Daten und schlecht dokumentierter Forschungs-Code aus verlassenen Github-Repositorien. Oft war der Aufwand um die Projekte überhaupt zum Laufen zu bekommen wesentlich höher, als sie für die eigenen Experimente anzupassen. Es fehlte etwas wie Processing für KI.

4: Caffe Framework: https://caffe.berkeleyvision.org/

5: Memo Akten über Deep Dream: https://memoakten.medium.com/deepdream-is-blowing-my-mind-6a2c8669c698

6: OpenCV nutzt GPUs zur Beschleunigung seit 2011: https://viso.ai/computer-vision/opencv/

7: Die Choreographic Coding Labs (CCL) bringen seit 2013 Gestaltung und Tanz zusammen: https://choreographiccoding.org/

8: https://readon.hs-mainz.de/2018/#ausstellung

9: https://jeanboehm.de/work/pathfinder

Chunk 10

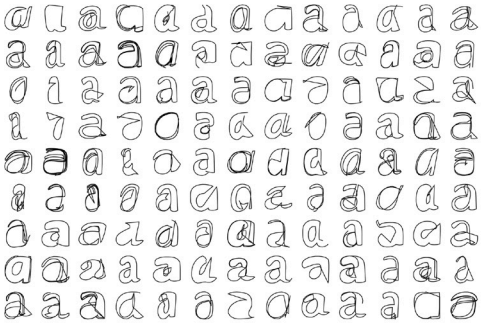

Abbildung 2: Sketch-RNN-Experimente zum kleinen „a" von Jean Böhm in seiner Masterthesis Pathfinder.

Chunk 11 Bei Motion Bank suchten wir 2020 nach einer Möglichkeit zur Integration unserer Arbeit in die Hochschul-Tanzausbildung. Dabei stießen wir auf die Bekanntmachung für eine Förderlinie im Rahmen der *Digitalen Hochschulbildung* des BMBF[10] (heute BMFTR), die die Anwendung von KI-Technologie in der Hochschullehre ermöglichen sollte (sog. *Förderlinie 4*).

Chunk 12 Unser Antrag *vortanz*[11] war erfolgreich und wir starteten Anfang 2021 in das Projekt. Kurze Zeit später wurde eine weitere Förderung im selben Rahmen veröffentlicht: *KI in der Hochschulbildung*. Ermuntert durch den vorausgegangenen Erfolg habe ich im März 2021 den Processing Profs dazu geschrieben – die erste Mail seit langem – und nach kurzem Hin-und-Her hatten wir eine erste Kerngruppe für einen Antrag gebildet. Die Partner aus Trier und Köln kamen durch glückliche Fügungen (Vermittlungen) hinzu.

Chunk 13 Zunächst waren wir noch mit sieben Hochschulpartnern im Gespräch, allerdings konnten zwei aufgrund anderer Anträge in derselben Förderlinie letztendlich nicht teilnehmen. Bis Mitte April haben wir Zoom-Meetings gehalten und in ein gemeinsames Dokument auf Google Docs geschrieben.

Chunk 14 Die meisten Antrags-Partner hatten bereits Erfahrungen im Einsatz von KI gesammelt. Anfang 2021 lagen die ersten GAN-tastischen Jahre hinter uns und DALL-E und GPT-3-Versionen waren das Maß der Dinge. Wir verfolgten die entstehende Diskussion um den Nutzen und die Herausforderungen des Einsatzes von KI, gerade auch im Hinblick auf die gestalterischen Disziplinen.

Chunk 15 Im Zentrum des Antrags stand daher das Ziel, die Studierenden für den anstehenden Wandel fit zu machen. Allerdings fehlte es an unseren Hochschulen dafür auf allen Ebenen. Wir hatten keine technische Infrastruktur, zu wenig KI-erfahrenes Personal und keine Zeit um Lehrangebote zu entwickeln. Unser Antrag kann also auch als Versuch gelesen werden diesen Sprung nach vorne zu ermöglichen. Die inhaltliche Linie war also schnell gefunden und ist in der Kurzbeschreibung des KITeGG-Antrags so zusammengefasst:

"[…] Das Verbundprojekt verfolgt dabei den Anspruch, die Lehre von KI-Methoden nachhaltig in die (Aus-)Bildung von Gestalter:innen einzubinden, sodass sie in der Lage sind, diese Methoden reflektiert als Material und Werkzeuge der Gestaltung einzusetzen […]. Dazu setzt sich der Verbund drei Ziele: Erstens den Aufbau einer Cloud-basierten Infrastruktur für Lehre und Lernen im Bereich KI […]. Zweitens die Entwicklung neuer kompetenzorientierter Lehreinheiten, […] sowie interdisziplinärer Labore […]. Drittens sollen die Ergebnisse des Projektes über Transferformate (Symposien, Open Source, Open Educational Resources, Publikationen) inner- und außerhalb des Verbundes in die gesamte Breite des gestalterischen Feldes getragen werden […]."

Nach einiger Fleiß- und Schreibarbeit der fünf Projektpartner, war zuletzt die große Herausforderung die vorgegebene maximale Länge von nur 12 Seiten Antragstext einzuhalten.

Chunk 16 Daniela Meinhardt und Lasse Scherffig vom Verbundpartner KISD konnten dies an einem langen Wochenende bewerkstelligen und so gingen alle benötigten Dokumente fristgerecht Ende

10: Bundesministerium für Bildung und Forschung (BMBF), heute Bundesministerium für Forschung, Technik und Raumfahrt (BMFTR): https://www.bmbf.de/

11: #vortanz – Automatisierte Vorannotation in der digitalen Hochschultanzausbildung: https://vortanz.ai/

April 2021 an das Ministerium. Schon im Juli kam die Zusage und Anfang Dezember 2021 konnten wir starten.

Den Rest der Geschichte finden Sie u.A. in den verschiedenen Ausgaben dieser un/learn ai Publikation.

Chunk References

Chunk 1	II. p. 202, Chunk 5: Editorial: Show and Tell II. p. 20, Chunk 4: Köln II. p. 4, Chunk 5: Vorwort
Chunk 7	II. p. 28, Chunk 1: Offenbach I. p. 160, Chunk 2: Ink and Algorithm - ein neurales…
Chunk 8	II. p. 15, Chunk 9: Mainz I. p. 85, Chunk 1: Robotik und Computer Vision… I. p. 27, Chunk 5: Living Objects Lab (KISD)
Chunk 10	II. p. 227, Chunk 7: Digital Lab I. p. 63, Chunk 1: AI+D Lab (HfG Schwäbisch Gmünd) I. p. 225, Chunk 2: Talks

↗ zsywb019

KITeGG und nun?

Florian Jenett

Chunk 1 In Vorbereitung auf diesen Beitrag habe ich eine Weile in der *Werkschau* unserer Lehr-Lern-Plattform (LLP) geblättert. Sie enthält die Ergebnisse der Studierenden, die unsere KITeGG-Kurse besucht haben.
Chunk 2 Zunächst fällt die Vielfalt auf, die in den ca.
Chunk 3 120 Kursen der letzten sechs Semester entstanden ist. Neben mutigen ersten Schritten finden sich ausgearbeitete, komplexe Projekte.
Chunk 4 Einige davon sind auch in dieser und den vorangegangenen un/learn Ausgaben enthalten. In der LLP sind Ergebnisse von Schriftgestaltung bis Animationsfilm, von spekulativem Produktprototyp bis zum Architekturentwurf, von der medientheoretischen Auseinandersetzung bis zum Prompting-Handbuch zu sehen. Dies beschreibt das Spektrum nur in groben Zügen und zeigt dennoch schon, dass KI im Rahmen von KITeGG auf der gesamten Breite der Gestaltungslehre angekommen ist.

Meine erste, vielleicht offensichtliche Annahme ist, dass KI zukünftig in noch mehr Disziplinen der Gestaltung zum Einsatz kommen wird.

Hierzu braucht es etwas, das wir zugegebenermaßen auch in unserem Antrag übersehen hatten: es müssen spezifische Angebote für die einzelnen Fächer und Disziplinen erarbeitet werden.
Chunk 5 Das geht nicht ohne die Expert*innen dieser Fächer: das bestehende Kollegium. Es braucht zudem kontinuierliches Recherchieren und Ausprobieren, Vorbereitungs- und Vermittlungsarbeit und direkte Hilfestellung in Kursen. Unser Antrag konzentrierte sich auf die Arbeit mit den Studierenden und hatte die Arbeit mit dem existierenden Kollegium schlicht übersehen. Dies konnten unsere KITeGG-Laborleitungen (siehe ↗ Interviews in un/lean ai Ausgabe 1) glücklicherweise abfangen, ohne sie wird aber zukünftig eine Lücke bleiben.

Es wird daher die Aufgabe der gerade an den Standorten entstehenden KI-Professuren sein, neue Labore aufzubauen.
Chunk 6 Nicht nur als Orte der Umsetzung, sondern auch der Forschung, der Beratung, der Vermittlung, des Austauschs und der Weiterbildung.

Diese Labore und die Lehre brauchen ein stabiles technisches Fundament. Es sollte keine Frage der persönlichen Ressourcen werden, ob Studierende den Weg der Gestaltung mit KI einschlagen. Es braucht daher Angebote der Bildungsstätten, die die Nutzung von KI kostenfrei zur Verfügung stellt. Um unabhängig von Marktinteressen zu bleiben und möglichst großen Spielraum für Experimente zu haben, halte ich hierzu den Betrieb einer eigenen KI-Infrastruktur für unumgänglich. Die fünf KITeGG-Standorte bereiten daher momentan den Weiterbetrieb unserer gemeinsamen Infrastruktur über das Projektende hinaus vor. Zudem regen wir immer wieder zum Einsatz von kleinen, lokalen Modellen an.
Chunk 7 Ihre geringe Größe und Spezialisierung kann für manches Projekt durchaus vorteilhaft sein und bringt Unabhängigkeit (bspw. von Dienstleistern und Internetzugang) und einen schonenden Umgang mit Ressourcen (bspw. Energie) mit sich.

Der Bedarf für standortnahe und lokale KI wird in der Gestaltungslehre und -forschung wachsen und die Hochschulen müssen Wege finden dies entweder als allgemeine Infrastruktur zu behandeln, oder zumindest mittelfristig als Projekte zu finanzieren.

Beim Studium der studentischen Projekte zeigt sich aber noch etwas anderes: Viele sind einen langen Weg gegangen, um den Einsatz von KI in die eigene

Arbeitsweise und Perspektive zu übersetzen. Sei es das automatische Gendern von Texten auf Internetseiten per Browserplugin, um auf sprachlichen Bias aufmerksam zu machen[1].

Chunk 8 Sei es die Kollaboration mit einem eigens trainierten (privaten) Modell beim Gestalten von Porzellanvasen[2].

Chunk 9 Es gäbe viele weitere Beispiele, die zeigen wie weit die Studierenden über das einfache Generieren von Text und Bild hinaus gekommen sind. Diese Erkenntnis freut mich besonders und verweist auf einen Aspekt, der mir im Antrag und auch bei der Arbeit in KITeGG stets sehr wichtig war: Gestaltung muss eine eigene Stimme im Bereich KI entwickeln, damit wir uns in die technologische Entwicklung einbringen können. Wir brauchen keine neuen Prompting-Möglichkeiten in veralteter Design-Software, sondern Werkzeuge, die uns sehr vielfältige, verschlungene und spezifische Einsatzmöglichkeiten und Arbeitsweisen erlauben[3].

Ich gehe davon aus, dass diese Anwendungen letztendlich zu einer Vielzahl neuer gestalterischer Perspektiven, Anwendungen, Ausgründungen und in Folge auch Forderungen an die allgemeine KI-Entwicklung führen werden. Zudem wird vermutlich die Fähigkeit KI in unterschiedlichste Bedarfe übersetzen zu können langfristig eine der wertvollsten Erfahrungen für die Studierenden sein, denn es macht sie unabhängig von der technologischen Entwicklung und ist eine Spezialisierung, die nicht automatisiert werden kann.

Chunk References

Chunk 3
- I. p. 69, Chunk 1: Vorwort Lehre

Chunk 4
- II. p. 4, Chunk 5: Vorwort
- I. p. 72, Chunk 5: Walking in latent space
- I. p. 17, Chunk 2: HfG Schwäbisch Gmünd

Chunk 5
- II. p. 216, Chunk 11: KI-Abend
- III. p. 254, Chunk 15: Auf dem Weg zu KI-TeGG
- I. p. 72, Chunk 5: Walking in latent space

Chunk 6
- II. p. 70, Chunk 25: Editorial
- II. p. 22, Chunk 4: Schwäbisch Gmünd
- II. p. 18, Chunk 22: Mainz

Chunk 7
- II. p. 251, Chunk 54: Der KITeGG Cluster – eine…
- II. p. 70, Chunk 25: Editorial
- I. p. 21, Chunk 2: Hochschule Trier

Chunk 9
- II. p. 4, Chunk 5: Vorwort
- I. p. 104, Chunk 16: Bildgenerierende Modelle in…
- I. p. 29, Chunk 12: Living Objects Lab (KISD)

1: Lisa Oelsen, Generiertes Generisches Femininum, Master Ausstellung Mainz, Sommersemester 2025

2: Johannes Growe, Fragiles in Transition, https://unlearn.gestaltung.ai/article/x83p8t7m

3: Nein, *ComfyUI* ist nicht gemeint.

Standorte & Team

Chunk 1

Hochschule Mainz

Jean Böhm, Julia-Jasmin Bold, Maika Dieterich, Paul Eßer, Lars Hembacher, Rebecca Herold, Prof. Florian Jenett, Anton Koch, Markus Mau, Francesco Scheffzcyk, Isabela Dimarco (bis 2023)

Mainz ist Verbundkoordinator im Projekt KITeGG und baut zudem die gemeinsame KI-Infrastruktur, sowie die übergreifende Lehr-/Lernplattform auf. Die Lehre und das Labor, das im Rahmen der Förderung eingerichtet wird, sind thematisch auf Automation, Werkzeuge und generative Gestaltungsansätze ausgerichtet.

Chunk 2

Hochschule Trier

Alexander Bauer, Peter Ehses, Prof. Simon Maris, Prof. Dr. Matthias Sieveke, Prof. Harald Steber

Nachhaltigkeit ist fest verankert in der Lehre am Campus, der seit 2021 das Label „Fairtrade University" trägt. Das KI-Labor legt den Fokus auf dieses Thema, so wird unter anderem die Materialeffizienz als nachhaltiger Aspekt zur Vermeidung von Überproduktion erforscht. Im Labor kann KI als kreatives, experimentelles Werkzeug zur Ideenfindung und Inspiration dienen oder dazu, verschiedene komplexe Zusammenhänge im Gestaltungsprozess einfließen zu lassen. Studierende erlernen, die KI-gestützten Informationen als Grundlage für ethisch und ökologisch verantwortungsvolle Entscheidungen in der Gestaltung anzuwenden.

Köln International School of Design, TH Köln

Matthias Grund, Jana Hartmann, Nils Heubaum, Jakob Kilian, Eduard Paal, Prof. Dr. Lasse Scherffig, Omid Akbari Kharazi (bis 2023), Tyanka Demyanka Adrian (bis 2024), Katrina Kizenbaha (bis 2024), Laura Juliane Wagner (bis 2023), Kjell Wistoff (bis 2022), Dzennifer Zachlod (bis 2024)

Im Projekt KITeGG stehen an der KISD Ansätze im Vordergrund, die die gesellschaftliche Wirkung der Gestaltung von und mit künstlicher Intelligenz betrachten und die – im Sinne eines interdisziplinären Designverständnisses – künstliche Intelligenz im Kontext der Gestaltung von Produkten, Prozessen und Interaktionen fokussieren.

Hochschule für Gestaltung Schwäbisch Gmünd

Prof. Hartmut Bohnacker, Prof. Rahel Flechtner, Prof. Benedikt Groß, Moritz Hartstang, Christopher Pietsch, Felix Sewing, Prof. Jordi Tost, Franz Anhäupl (bis 2025), Maxime Beck (bis 2025), Vivien Cai (bis 2025), Lucie de Hair (bis 2023), Malte Fial (bis 2023), Mark-Jonathan Freyer (bis 2024), Stamatia Galanis (bis 2024), Robert Kremer (bis 2025), Ron Mandic (bis 2025), Katharina Neugart (bis 2024), Johannes Rothkegel (bis 2023), Prof. Aeneas Stankowski (bis 2024), Alexa Steinbrück (bis 2023), Anton Söffing (bis 2023), Alia Tasler (bis 2023), Felix Tischmacher (bis 2022)

Die HfG Schwäbisch Gmünd hat sich im Projekt zwei Schwerpunkte gesetzt,

durch die KI als neues Gestaltungsmittel verstanden und genutzt werden soll. Der theoriebasierte Schwerpunkt „Design Futures" zielt auf die Konstruktion von Zukunftsbildern, um KI in aktuellen gesellschaftlichen Diskussionen zu begleiten und mögliche Auswirkungen auf Gesellschaft, Wirtschaft und Kultur zu erproben.

Chunk 3 In den Lehrformaten zu „Creative Machine Learning" setzen sich Studierende mit KI-Technologien auseinander und lernen diese auf design-wissenschaftliche Problemstellungen zu adaptieren und zu trainieren. Das AI+Design Labor (AI+D Lab) unterstützt interdisziplinäres und forschendes Lernen im Design mit KI und bietet eine Grundausstattung an Werkzeug und Materialien für den Aufbau schneller KI-Designprototypen.

Hochschule für Gestaltung Offenbach

Prof. Catrin Altenbrandt, Elisa Deutloff, Mattis Kuhn, Leon-Etienne Kühr, Prof. Adrian Nießler, Prof. Alex Oppermann, Joscha Berg (bis 2024), Ivan Iovine (bis 2025), Bastian Kämmer (bis 2023), Max Kreis (bis 2024), Johanna Wallenborn (bis 2025)

Die HfG Offenbach setzt sich als thematischen Schwerpunkt innerhalb des Projekts mit KI und künstlerischer Praxis auseinander.

Chunk 4 Im Rahmen des Projekts wurden ein KI-Labor und ein Robotiklabor neu eingerichtet, die in das künstlerische Lehrgebiet Elektronische Medien integriert sind.

Chunk 5 Das KI-Lab unterstützt Studierende bei der Entwicklung von Deep-Learning-Algorithmen und bei der Integration von bestehenden Datensätzen und trainierten Modellen in ihre künstlerische Praxis. Dabei spielen die Implikationen von Technologie auf die Gesellschaft und die Reflektion des eigenen künstlerischen Schaffens eine zentrale Rolle. Mit Hilfe von Physical Computing, Robotik, digitaler Fertigung und künstlicher Intelligenz unterstützt das Robotiklabor die Studierenden bei der Entwicklung interaktiver Installationen.

Chunk References

Chunk 1	II. p. 13, Chunk 2: Mainz II. p. 233, Chunk 1: Der KITeGG Cluster – eine… I. p. 13, Chunk 2: Hochschule Mainz
Chunk 2	I. p. 47, Chunk 13: KINDLAB (HS Trier) II. p. 32, Chunk 2: Trier I. p. 21, Chunk 2: Hochschule Trier
Chunk 3	I. p. 21, Chunk 4: Hochschule Trier I. p. 17, Chunk 3: HfG Schwäbisch Gmünd II. p. 182, Chunk 1: SandwichNet
Chunk 4	I. p. 19, Chunk 2: HfG Offenbach II. p. 30, Chunk 10: Offenbach II. p. 30, Chunk 8: Offenbach
Chunk 5	I. p. 19, Chunk 2: HfG Offenbach II. p. 182, Chunk 1: SandwichNet I. p. 21, Chunk 4: Hochschule Trier

Credits

un/learn ai – Integrating AI in Aesthetic Practices

Dritte Ausgabe

Herausgegeben von
Matthias Grund, Johanna Teresa Wallenborn, Lasse Scherffig, Florian Jenett.

Forschungsprojekt
KITeGG - KI greifbar machen und begreifen: Technologie und Gesellschaft verbinden durch Gestaltung. Ein Verbundprojekt zur Integration von KI in der Gestaltungslehre.

Projektpartner
Hochschule Mainz, Hochschule Trier, Hochschule für Gestaltung Offenbach, Hochschule für Gestaltung Schwäbisch Gmünd, Köln International School of Design (TH Köln)

Projektleitung
Prof. Florian Jenett, Hochschule Mainz

Projektkoordination
Julia-Jasmin Bold, Hochschule Mainz

Redaktion
Matthias Grund, Jana Hartmann, Jakob Kilian, Leon-Etienne Kühr, Mattis Kuhn, Johanna T. Wallenborn, Lasse Scherffig

Autor*innen
Julia-Jasmin Bold, Elisa Deutloff, Paul E er, Prof. Rahel Flechtner, Matthias Grund, Robert Halbach, Jana Hartmann, Thomas Hawranke, Lars Hembacher, Ivan Iovine, Prof. Florian Jenett, Nina Juric, Jakob Kilian, Anton Koch, Mattis Kuhn, Leon-Etienne Kühr, Ting-Chun Liu, Prof. Simon Maris, Prof. Alex Oppermann, Francesco Scheffczyk, Prof. Dr. Lasse Scherffig, Niklas Thielen, Prof. Jordi Tost, Johanna T. Wallenborn, Lea Waldera, Dzennifer Zachlod

Studentische Arbeiten
Ilona Achba, Tyanka Adrian, Inês Fernandes de Aguiar, Berire Akdogan, Franz Anhaeupl, Marie Backhaus, Hannes Bader, Nelli Gomez Baumert, Oliver Beer, Alissa Bernhardt, Alina Bertacca, João Borges, Berit Bove, Piroschka Benedikta Brucker, Bidisha Das, Elisa Deutloff, Bia Dobbs, Maybritt Dornheim, Peter Ehses, João Nuno Rua Fonseca, Philipp Gatzke, Marlene Grabenwöger, Jana Hartmann, Henrike Hof, Helena Hofmann, Georgia Ioannidou, Luka Jacke, Levon Javaheri, Meike Kattwinkel, Yonghyuk Kim, Max Kreis, Robert Kremer, Anna Krämer, Annika Leuoth, Alicia Link, Eva Matzerath, Katrina Milta, Lisa Oelsen, Mário Pereira, Jovydas Petr tis, Jule Pickel, Elzbieta Ploiakova, Vytenis Razmus, Daniel Rosewich, Daniel Mikula Scholl, Justus Schomann, Elena Schröder, Phillip Schulz, Tristan Schulze, Pourya Sharifi, Aikaterini Sideri, Anna Sokolova, Hannah Steinbrecher, Agnes Stigler, Colin Subat, Anastasija Surovceva, Niklas Thielen, Ada Wagner, Cosima Wagner, Lea Waldera, Zhichang Wang, Micha Wempe, Viktor Werbilo, Anastasia Verena Wilbrandt, Dzennifer Zachlod, Jessica Zisa

Lektorat
Emily Collard-Ducluzeau, Jacob Watson

Konzeption, Gestaltung und Umsetzung
Jean Böhm, Maika Dieterich, Paul E er, Matthias Grund, Lars Hembacher, Markus Mau, Francesco Scheffczyk

Schriften
Bagoss von Displaay,
Roumald von Erkin Karamemet,
WhoIs Mono von Raphaël Bastide

Druck
Hochschule Mainz, Werkstatt Druck+Papier, Marco Moll

Papier
Gmund Les Naturals 330 g/m², ROUGH AIR White-Warmwhite 120 g/m², Gmund Original Smooth Beige 120 g/m²

Website
unlearn.gestaltung.ai

Lizenz
Die Inhalte der Publikation können unter der Creative Commons-Lizenz CC BY 4.0 genutzt werden, sofern sie in Ausnahmefällen nicht anders gekennzeichnet sind. Bei jeder Verwendung sind die angegebenen Lizenzhinweise zu beachten. Bei einer Weiterverwendung müssen die Lizenzen ebenfalls angegeben werden. Die Abbildungen von studentischen Projekten sind von der oben genannten Creative Commons-Lizenz ausgeschlossen.

1. Auflage, September 2025

Gefördert durch das Bundesministerium für Forschung, Technologie und Raumfahrt (BMFTR) sowie die Bundesländer Rheinland-Pfalz, Hessen, Baden-Württemberg und Nordrhein-Westfalen.

Weitere Infos zu KITeGG finden Sie auf: gestaltung.ai

Gefördert durch:

Bundesministerium für Forschung, Technologie und Raumfahrt

RheinlandPfalz
MINISTERIUM FÜR WISSENSCHAFT UND GESUNDHEIT

HESSEN
Hessisches Ministerium für Wissenschaft und Kunst

Ministerium für Kultur und Wissenschaft des Landes Nordrhein-Westfalen

Baden-Württemberg
MINISTERIUM FÜR WISSENSCHAFT, FORSCHUNG UND KUNST

KITeGG

KI greifbar machen und begreifen:
Technologie und Gesellschaft verbinden
durch Gestaltung.

Making AI tangible and comprehensible:
Connecting technology and society
through design.

gestaltung.ai

10.25358/openscience-13020