

AI FOR EXPERIENCE: Designing with Generative Adversarial Networks to evoke climate fascination

Master thesis
Frederik Ueberschär

Design for
Interaction

TU Delft
February 2021

**AI for Experience:**

Designing with Generative Adversarial
Networks to evoke climate fascination

Master thesis

Design for Interaction

Frederik Ueberschär

February 2021

Delft University of Technology

Faculty of Industrial Design Engineering

Committee

Derek Lomas

Alessandro Bozzon

Contact details

frederik-ueberschaer.de

© Frederik Ueberschär 2021

Whereas it is accepted that the present is caused by the past it is also possible to think of it being shaped by the future, by our hopes and dreams for tomorrow.

– Anthony Dunne & Fiona Raby (*Speculative Everything*)

PREFACE

When I started my Design for Interaction Masters at the TU Delft, I was hoping to find answers. What kind of designer do I want to be? What can and should I work on in the future? What even is design? Instead of finding simple, clear answers to my questions, I was left with even more. And to me, this is something incredibly beautiful.

My past two years studying and all the wonderful and talented people I had the fortune to encounter on the way helped me to see design from novel perspectives. Instead of discovering the one thing I want to work on in the future, I found hundreds. Rather than finding the clear definition of what kind of designer I want to be, I now see it as a never-ending process, with the chance to always grow, challenge yourself and embrace the joy and excitement that it brings.

For this self-initiated graduation project, I saw the opportunity to combine two themes that I personally care about and that never fail to engage me: technology and our climate. When combining the two themes, I was happy to see that the enthusiasm was not only my own but answered from multiple sides. This project helped me to grow as a designer and again discover new perspectives, new questions that want to be answered. I hope that you, my dear reader, will also take something out of my journey captured in this report, and if it's only a small spark of inspiration, a novel thought, or plain appreciation of intriguing earth imagery.

When I started the project six months ago, the world was still a different one. The persistent reality of COVID-19 and the uncertainty it brought with it turned this project into a wild ride, and keeping a positive spirit was not always an easy task. At this point, I want to take the opportunity and thank everyone who guided and supported me through this graduation project:

To my supervisors Derek and Alessandro, I want to thank you for your mentorship throughout the project, continuously pushing me to go to new lengths, and giving me the ability to grow. This project would certainly be a different one without you.

To Luke and the team members of Google AI, thank you for allowing me to collaborate with such talented people. Mauro, thank you especially, for our weekly meetings. You helped guide me through this project, always provided me with an encouraging perspective, and motivated me to keep going.

To my family, thank you for always being there for me. I am endlessly grateful for your support and I could not have done my studies and my continuous exploration without you.

To the members of ID StudioLab, with special regards to Aadjan, thank you for welcoming me in your midst and providing me with a space to work.

To the countless other people who supported me throughout this project, by testing my prototypes or providing their expertise: thank you! I hope you found as much joy in our conversations as I have.

Lastly, I want to thank my dear friends and my partner Cristina, who accompanied me throughout my studies and turned this time into the most wonderful experience imaginable. You helped me see structure and paths when I was lost, and always offered an open ear and kind words when I needed them, and ultimately helped me keep my sanity throughout the project.

A handwritten signature in black ink, reading 'F. Ueberschär'.

Frederik Ueberschär

19. Feb 2021, The Hague

EXECUTIVE SUMMARY

Our world climate is an incredibly complex and ever-changing system. The reality of climate change and the consequential pressure to act fast to reduce greenhouse gas emissions requires immediate attention. And yet, despite this unprecedented urgency, active public engagement needed to motivate meaningful change is still missing.

This graduation project explores how to utilize Generative Adversarial Networks (GANs) – a subset of artificial intelligence (AI) – to motivate public engagement through the emotion of fascination and element of play. The final design of this exploration is an interactive exhibition piece that allows its users to create novel, AI-generated landscapes.



In the first phase of the project, hands-on exploration of the use of GANs as a design opportunity was conducted and the experiences of the exploration process documented. Further research through interviews with climate researchers and literature review into the topics of climate change helped to get an understanding of the climate context and map out and narrow down the solution space for the following ideation phase. The context research revealed that the current climate debate is deeply emotional, with mostly negative emotions like fear and hopelessness in its center. Yet, it was also found that fostering positive emotions like fascination and joy is a potential means to increase interest and motivate people to engage with climate initiatives. Based on the context exploration, a series of design prototypes were created and tested with participants to inform the final design.

The final concept LANDSHAPES is an interactive exhibition piece intended for the museum context. It features a circular projection of AI-generated aerial landscapes – images that look real, yet have never been seen by anyone before. The installation enables participants to become absorbed in the morphing landscapes and experience their potential impact on our surroundings in a playful, open-ended manner.



A preliminary installation was evaluated with participants using qualitative methods of observation and interviews. The installation evoked a deep sense of beauty, calmness, and inspiration that contributed to a meditation-like, engaging experience. Even without a predefined narrative, participants automatically associated the morphing images with our changing planet and ascribed meaning to the abstract visuals. The addition of control over the moving landscapes resulted in a mindset shift, as the experience transformed from a passive and meditative experience to an experience with additional cognitive involvement. The preference of experience differed between participants.

The design research conducted in this project strengthens the assumption that experiences with positive emotion in its center can be used to foster initial engagement with the topic of climate change. Furthermore, it showed that GAN technology can be used to evoke an emotional response, which offers interesting opportunities for future work. More specifically, the research suggests the potential for further development of playful experiences to make climate data accessible and engaging to the general public, as well as an opportunity to explore GANs that utilize emotional value as an input to generate targeted media that reinforces or counteracts the viewer's emotional response.



GLOSSARY

Artificial Intelligence (AI, field)

The broad field of computer science focussed on building smart machines.

Artificial Intelligence (AI, noun)

A system with the ability to interpret external data correctly, to learn from such data, and to use those learnings to achieve specific goals and tasks through flexible adaptation (Kaplan & Haenlein, 2019).

Climate Change

A change in climate patterns, both local and global, attributed largely to manmade increase of CO₂ in the earth's atmosphere.

Design Goal

Describes the target for a design work without limiting the solution space.

Empirical Research Through Design (ERDM)

A design methodology that utilizes quick concept, prototyping, and validation cycles to approach the final design.

Generative Adversarial Network (GAN)

A machine learning framework utilizing two competing machine learning systems to generate and detect fake data. By self-assessment, the system learns to output increasingly realistic novel data.

Interaction Vision

A design method that uses an analogy to a commonly known situation or feeling to derive its characteristics and use it to describe an envisioned interaction with a novel design.

Machine Learning (ML)

A sub-type of artificial intelligence that uses vast amounts of data to learn from it without the instructions given by a programmer.

StyleGAN

A popular type of Generative Adversarial Network developed by NVIDIA.

TABLE OF CONTENTS

01	Introduction	15
01.01	Climate change and the need for engagement	16
01.02	Design and technology	17
01.03	Collaboration with Google Research	18
01.04	Project aim	18
01.05	Approach	20
02	Theme exploration: Artificial Intelligence	23
02.01	A short history: What is an AI?	24
02.02	Data and training: how an AI comes to be	25
02.03	Analyze and generate: use cases, capabilities and shortcomings of AI	28
02.04	Current limitations and challenges	29
02.05	The ethical debate surrounding AI	30
02.06	From theory to practice: creating and utilizing an AI System	32
02.07	Exploration journal: AI as a design tool	33
02.08	What is a Generative Adversarial Network?	33
02.09	Initial exploration experience	38
02.10	Difficulties faced on the way when working with GANs	56
02.11	Reflection on working with a generative AI model	59
02.12	Chapter conclusion	60
03	Theme exploration: Climate	63
03.01	Our ever-changing climate	67
03.02	From the past to the present to the future	72
03.03	Acting against climate change	74
03.04	Represented stakeholders	74
03.05	Current and potential climate actions	75
03.06	Public understanding of climate change	80
03.07	Emotional element of climate change	81
03.08	Chapter conclusion	83

04	Ideation & Prototyping	85
04.01	Ideation Basis	88
04.02	Design Goal & Interaction Vision	90
04.03	Target group	91
04.04	Hypothesis and design methodology	92
04.05	Design scope and further restrictions	93
04.06	Design Cycle 1: Divergence	94
04.07	Design Cycle 2: Interaction	98
04.08	Additional design exploration	102
04.09	Ideation conclusion	103
05	Final Design: Landshapes	105
05.01	Concept description	106
05.02	Developing a proof of concept	118
05.03	Building a dataset	120
05.04	Training a StyleGAN network	124
05.05	Concept Validation	126
05.06	Insights & recommendation for the final design	134
06	Project conclusion	137
06.01	Project summary & outcomes	138
06.02	Answering the design goal	139
06.03	Knowledge contributions	140
06.04	Future work	142
06.05	Personal Reflection	144
06.06	Tips for other students	145
	References	146



01 INTRODUCTION

In this introductory section, I will give a general overview of the project and its context. This includes a description of the project background, its aim, relevance, stakeholders, and guiding themes of climate change and artificial intelligence. Furthermore, a visual representation of chosen design methodology corresponding to the resulting reading structure of the report will give an overview of the concluding report structure.

01.01 CLIMATE CHANGE AND THE NEED FOR ENGAGEMENT

Our world climate is incredibly complex and ever-changing. It is a collective issue, too big to be understood and solved by one person, one project, or even one country alone. The reality of climate change and the consequential pressure to act fast to reduce our greenhouse gas emissions and general impact on the world has reached an unprecedented amount of attention from public, political, corporate, and academic communities (Schmidt et al., 2013). Recent years have also shown an increased public interest in climate developments, manifesting in climate protests like Friday for Future (2020), but further and stronger engagement is needed to cause meaningful change.

Currently, vast resources are pooled into researching our climate with modern data collection methods and intelligent algorithms helping scientists in their understanding and predictive efforts of our climate development (EU Funds, 2017). This communication of climate findings commonly characterized by numbers, graphs, and elaborated reports in scientific language. While this form of communication is presumably fitting for their main target group, its complexity and scientific nature make our climate and ongoing development a topic that is hard to grasp and imagine for a novice. The general public as an active shareholder in our climate discussion requires a different form of communication and preparation of climate findings to be able to participate in the discussion.

Furthermore, climate change is not only a factual topic but substantially driven by emotions. The public emotional response to global warming overwhelmingly consists of negative emotions, like worry, anger, sadness, or fear (Smith & Leiserowitz, 2014). However, studies have also found that the use of "positive emotions appear to play an important role in public support for climate policies" (Smith & Leiserowitz, 2014).

01.02 DESIGN AND TECHNOLOGY

The field of design has always been driven by innovation and changes. What to design, who to design for and how to design are in constant shifts that incorporate technological changes as well as changes in societal and individual needs.

Innovation in material science and manufacturing processes allowed for the redesign and improvement of existing products, while developments like the digital revolution even opened up completely new behavior, product and design domains (King & Chang, 2016) and dramatically influenced how design is practiced nowadays.

One of the more recent technical innovations concerns the notion of artificial intelligence (AI). While not being a particularly novel concept in itself, recent developments and improvements in AI technology contributed to its increased utilization for scientific research, enterprise solutions, and lastly consumer products. Nowadays, many industries produce and rely on products with a smart component with prominent examples of smart consumer products including search engines like Google Search, voice assistants like Amazon Alexa, recommendation systems in services like Netflix and Spotify, or AI-curated social media feeds of Facebook and Twitter. The increased popularity of products with an AI component can be seen as an encouragement for designers to familiarize themselves with the technology and explore how they could utilize AI as a practical tool in their workflow and (industrial) design processes.

In recent years, the fields of design – including that of industrial design – continually broadened their scope. Instead of solely providing concrete product solutions to concrete problems, designers nowadays commonly work in interdisciplinary teams, design vast systems rather than actual objects and use their expertise to tackle big societal issues (Tromp & Hekkert, 2018) – like the issue of climate change. AI is already utilized to help tackle climate change, for instance by improving the energy efficiency of industries, optimizing supply chains, or modeling future climate scenarios to inform policymakers (Rolnick et al., 2019), but the potential usage also extends further into the field of climate change communication.

01.03 COLLABORATION WITH GOOGLE AI

One of the driving forces at the forefront of AI development is the company Google. For parts of this project, I had the fortune to work together with a team of Google AI, with members based in Berlin, London, New York, and San Francisco. The team formed around a pitch proposed by our graduation team and sparked the first company-internal excitement for developing an experience that had the potential to showcase Google technology and result in a gamified experience to approach the topic of climate change.

The collaboration focused mainly on the first stages of the project, exploring the general climate space through expert interviews, sharing research insights, and collecting examples of prior art that could be used as a basis for the ideation..

01.04 PROJECT AIM

As stated in the design brief (Appendix 1), this graduation project will take a deep dive into exploring the opportunity of transforming climate data through AI into a playable experience and make it engaging and immersive for stakeholders beyond the scientific community. However, the goal is not to build a finished and financially viable product, but rather to explore the domain of climate research and prototype an engaging form of novel communication. By translating climate data into a playable, visual experience is meant to pave the way for further work that fosters public climate engagement by bringing the elements of AI, climate, and play together.

Hypotheses

The three hypotheses for the design exploration are:

1. **An experience to provoke the positive emotion of fascination is an appropriate approach to engage people with the topic of climate change.**
2. **Generative Adversarial Networks can produce media to provoke an emotional response.**
3. **A playful and interactive component strengthens the engagement with the installation.**

Initial research questions

The described approach, goal and hypothesis results in several research questions to explore further in the scope of the project:

1. **What are the challenges and opportunities when working with a GAN as a design tool?**
2. **What are the current challenges and opportunities of climate communication and engagement?**
3. **Is a visual, playful experience a suitable means to engage people?**
4. **How can and should the playful interaction look like?**
5. **Can a GAN be used as a suitable tool for the playable experience?**
6. **What climate data can be used to motivate engagement?**
7. **What emotional response can and should the experience evoke?**

01.05 APPROACH

In this project, I explore the usage of Generative Adversarial Networks (GANs) – a subset of generative AI technology – a hands-on tool in a design process to support climate engagement. The latter, concrete design goal functions as a facilitator for the exploration process with the goal of knowledge generation. This meta-level exploration (Figure 1), operating simultaneously on a concrete design level as well as a higher design tool and methodology level, as described by Stappers & Visser (2014) also guided my process and the methodology chosen to fulfill goals for both levels.

My process is further heavily inspired by the empirical research through design methodology (ERDM) (Keyson & Bruns, 2009) with various design activities and rapid prototyping iterations as a means to generate new knowledge as well as proceed to fulfill the setout design goal. Based on input from two theme explorations of both the AI and climate space, I formed a hypothesis and designed and tested a variety of artifacts based on that underlying hypothesis. Through testing and iteration of intermediate prototypes, the final iteration was formed and tested in a final experiment set out to answer the initially formed hypothesis. The project approach is reflected in the report structure (Figure 2).




Figure 1: *Moving between different meta levels allows for exploration on the knowledge and design level*

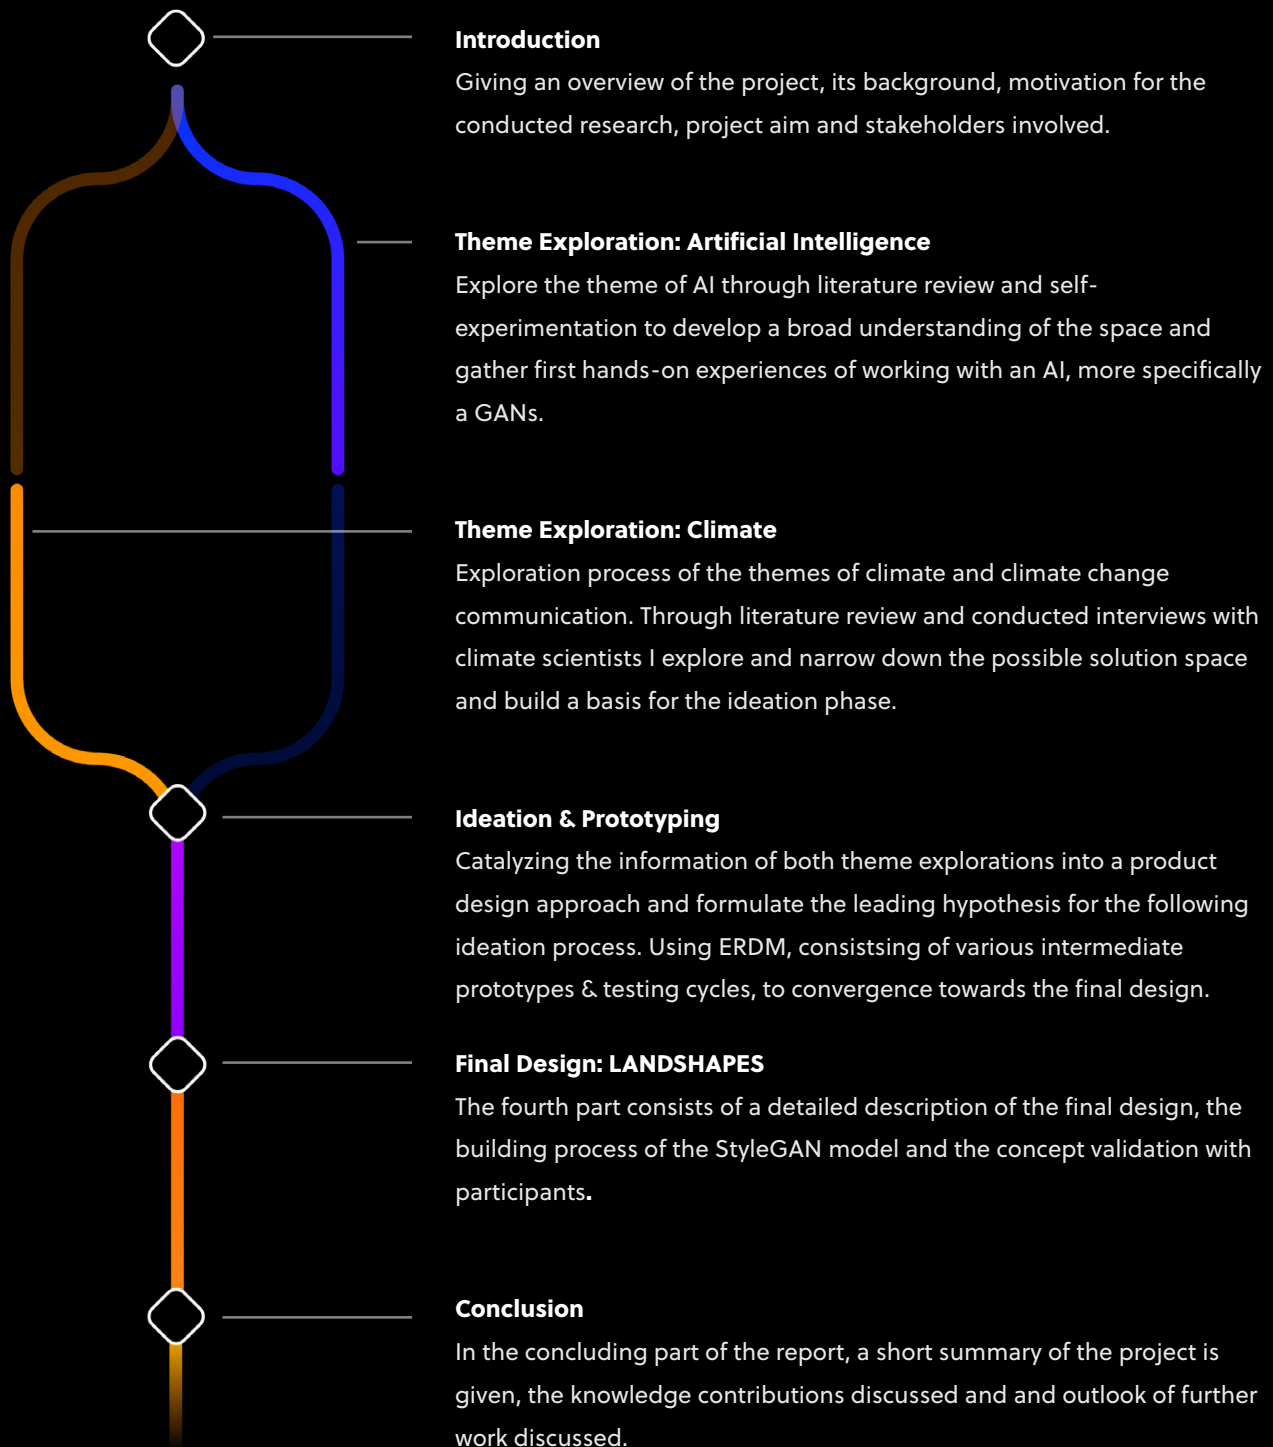


Figure 2: Project process visualized in a flow from top to bottom



02 THEME EXPLORATION: ARTIFICIAL INTELLIGENCE

To know how to use a tool to its fullest, one needs to understand its capabilities as well as limitations. This chapter will give a brief overview of artificial intelligence, how it is utilized in consumer products to date already. This exploration will allow me to decide what kind of AI I can and explore further.

AI is a highly complex computational field. As I am exploring AI from the perspective of a technical novice, I am not striving to understand and explain all the deeply technical elements of an AI, but rather get a sense of the bigger picture. However, as my mentor accurately framed it: "To paint, you don't necessarily need to know everything about how the pigments in the paint are made".

| *Methods used: Literature review and hands-on prototyping exploration.*

02.01 A SHORT HISTORY: WHAT IS AN AI?

The human thought of machine intelligence goes back to the antique and has been present in literature and fiction ever since. What started with stories of Talos, a figure of Greek mythology that was said to be made out of bronze and circle the island of Crete to protect it from invaders is also present in more recent culture with figures like the terminator or HAL 9000 from '2001: A Space Odyssey'.

This fiction started to become more real with the development of the first computers with the term of artificial intelligence (AI) being coined during a workshop in 1956 by computer scientist John McCarthy. Despite its long existence, the terminology of AI is still quite fuzzy with different definitions floating around (Kaplan & Haenlein, 2019). Nowadays, AI is often used as an umbrella term for various types of different advanced algorithms and the computer science sector that specializes in the development of those systems (Jakhar & Kaur, 2020).

As its definition is so broad, there is no archetype of "one AI", but a multitude of different systems for different applications and use cases. A definition given by Kaplan & Haenlein (2019) will be used throughout the report and says:

AI as a system's ability to interpret external data correctly, to learn from such data, and to use those learnings to achieve specific goals and tasks through flexible adaptation.

02.02 DATA AND TRAINING: HOW AN AI COMES TO BE

To be able to interpret data correctly and create a sense of what we might call intelligence, AI systems use machine learning (ML) and an underlying neural network (Figure 3) to analyze big amounts of data, recognize patterns and relations, and transfer them into a statistical model to make informed predictions (Pasquinelli & Joler, 2020).

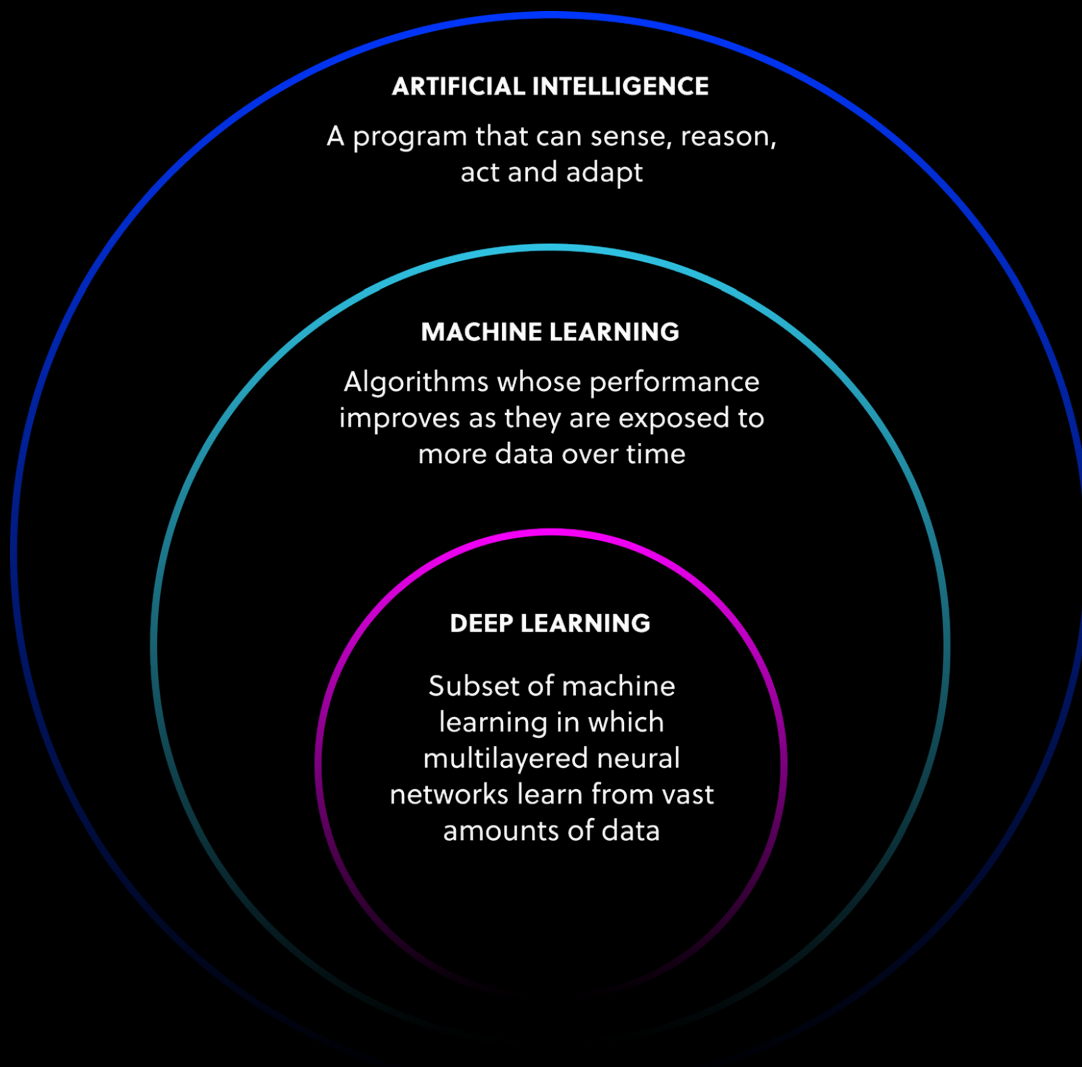


Figure 3: Distinction and relation between AI, ML and DL, adapted from (Singh, 2017)

All AI systems need data to be trained on, but the training process itself can differ depending on the chosen network structure. With this training, the system learns and improves over time. Kaplan & Haenlein (2019) divide the types of learning into three broad categories:

Supervised learning, where the given set of inputs is related to a given set of usually labeled outputs. A system like this could be trained to distinguish between the picture of a dog and a cat (Figure 4).




Figure 4: Supervised learning using labeled input data

Unsupervised learning builds upon labeled inputs but unlabeled outputs. The algorithm needs to find and cluster an underlying structure in the data (Figure 5). The number of possible outputs is not known in advance to the human nor system, which makes the accuracy and trustworthiness output hard to assess. These systems are often used for speech recognition.



Figure 5: Unsupervised learning using unlabeled input data to form clusters

Reinforcement learning gives the system an objective, like achieving a high game score, and a set of possible actions and lets the system try, fail or succeed and learn for the next round to improve itself (Figure 6).



Figure 6: Reinforced learning with the objective to maximise a score

In the end, algorithms turn input data into an ML model (condensed training set), which when prompted can output new data, make predictions or even act upon the prediction made.




Figure 7: Basic structure of a machine learning model, adapted from (Sevarac, 2020)

02.03 ANALYZE AND GENERATE: USE CASES, CAPABILITIES AND SHORTCOMINGS OF AI

Increased computing power, as well as access to vast amounts of data, led to a significant increase in AI applications over the past years (Columbus, 2018) and the notion of it seems to be everywhere. AI systems enable companies and services to work with amounts of data that would not be comprehensible for a single human anymore. The created smart systems have the ability to support or even replace humans in the process of making sense of the data (Wang et al., 2019).

Consumer products

Some already common use cases for consumer products include recommendation systems for online services, autocomplete functionality on smartphones, personal assistants like Alexa or Siri (Figure 8), efficient routing of goods and people (like in Google Maps).

Research

The uses of AI go beyond consumer products and are also present to support the research context. In 2020, the AI-focussed company DeepMind published a case study on their AI-driven solution to protein folding, a 50-year old challenge in biology, which “has the potential to accelerate research in every field of biology” (DeepMind, 2020) and assist researchers in their work, for instance in the development of vaccines. The system was trained on the amino acid sequence and the resulting structure of over 100.000 known proteins. The resulting model can now accurately predict the 3D structure of a protein by having solely the amino acid sequence as a given input, supporting the work of researchers immensely (Figure 9).

Art and creative purposes

“Humans have used tools to extend their creative capabilities since the stone age. But in the last decade, AI and machine learning has transformed our ability to create more than any time in history. These new technologies not only open new possibilities, but also pose important questions about how we’ll interact with technology now and in the future.” - Mission statement (AI Artists, n.d. a)




Figure 8: Smart assistants like Siri (image by Apple, n.d.) are powered by a form of AI




Figure 9: Protein structures predicted by AI (DeepMind, 2020)

The creative and artistic field has been working with AI for many years already. Through generative art pieces (Figure 10), artists explore possibilities and limitations of AI, as well as target questions surrounding the topic, like "Can autonomous machines be creative and truly create art?", "How will AI systems enhance and augment human creativity?" or "What values should we embed in AI?" (AI Artists, n.d. b)

These questions show that there are still many open questions surrounding AI, ranging from philosophical to technical questions.

02.04 CURRENT LIMITATIONS AND CHALLENGES

AI systems to date, as advanced and capable as they may seem, are still far from the imagined capabilities of a futuristic all-knowing, all-capable AI presented in pop culture, also called a "strong" or "general" AI. The systems we have so far are classified as "weak" or "narrow" AI, which means that they have the capability to perform a limited amount of specific tasks well (Kaplan & Haenlein, 2019).

But even these narrow AI systems reached a complexity that often is not understood by their creators anymore and (without the ability for the machine to reason its decision making) lacks explainability. Vladan Joler and Matteo Pasquinelli, authors of "The Noosphere Manifested" (2020) describe it as follows:

"AI is now at the same stage as when the steam engine was invented, before the laws of thermodynamics necessary to explain and control its inner workings, had been discovered."

Without a way of a system to reason its decision-making in a human-comprehensible form, the result is what is often referred to as a black box. This lack of explainability and controllability are two of the main reasons for some experts like Veale et al. (2018) to advocate for the limited use of AI for critical decision making, at least in fully-automated form.



Figure 10: *Example of an art piece utilizing AI by Mario Klingemann (2018)*

02.05 THE ETHICAL DEBATE SURROUNDING AI

Every technological innovation will be surrounded by a debate about its uses, potentials, and shortcomings sooner or later. AI is no exception to this, and recent years have shown an active debate about the ethics of AI-powered systems. To reflect the depth of this ongoing discussion is beyond the scope of this exploration, but a selection of the discussion points addressed by Lepri et al. (2018) target the points of:

1. Fairness

How do we make sure that AI systems make decisions that are free of bias and discrimination, but treat everyone fairly?

2. Transparency

How do we make design an AI that we can understand, which can explain itself and reason its decision making?

3. Accountability

AI makes mistakes just as humans do. Who is responsible for a mistake made by an autonomous system?

4. Agency

How do we remain agency over AI systems and control them?

02.06 FROM THEORY TO PRACTICE: CREATING AND UTILIZING AN AI SYSTEM

The current popularity, ongoing development, and increasing number of use cases for ML and AI applications strongly suggest that the technology will stick around and its usage will only increase in the coming years and decades. Recent years have also shown an increased number of platforms and services that allow researchers, makers, and designers to explore, build and utilize various kinds of ML. Some notable platforms and services, some of which require no or very little coding knowledge are ml5.js, RunwayML, and Google Teachable Machine (Figure 11).

However, it is one thing to explore the topic of AI from a theoretical standpoint, and another to apply it practically – especially without a degree in computer science. With the field of design always being driven by technological progress and means (King & Chang, 2016), the AI progress and increased accessibility to tools present an opportunity for designers to familiarize themselves with the technology in a practical manner.

In this section, I will describe my first attempt at training an AI system and reflect on the experience with the perspective of applying it as a design tool in my process.



Figure 11: Logos of three examples of accessible machine learning platforms, ml5.js, Runway and Google Teachable Machine

02.07 EXPLORATION JOURNAL: AI AS A DESIGN TOOL

For this graduation project, I set out to explore how I can use AI as a practical tool in my design process. Since many different types of AI systems technologies exist, I decided to limit myself to one type of AI, in my case a Generative Adversarial Network (GAN). The coming section will explain the basic functionality of a GAN further and showcase why it might be a suitable tool for design.

Although I possess some basic coding knowledge and background knowledge about the utilization, workings, and capabilities of an AI, I can be described as a novice in both fields. From the first day onwards, I logged my experiences and my collected learnings in a journal. The purpose of the journal is to be able to analyze my progress retrospectively and identify hurdles as well as accelerators I faced on the way.

02.08 WHAT IS A GENERATIVE ADVERSARIAL NETWORK?

The framework of Generative Adversarial Networks was first proposed by Ian Goodfellow and his team back in 2014. The ML structure behind a GAN consists of two competing networks, a so-called generator, and an opposing discriminator (Figure 12).

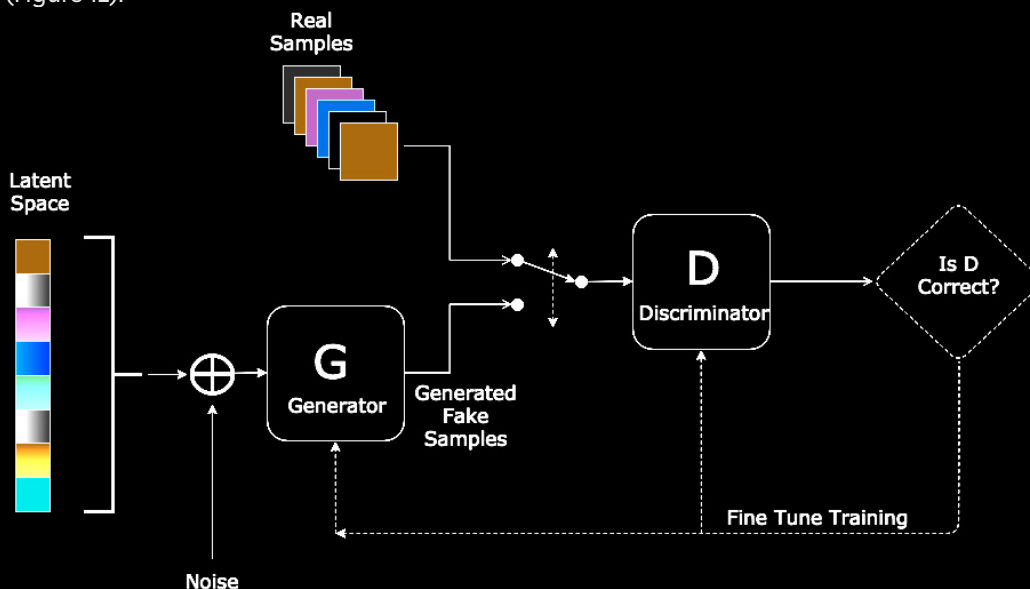


Figure 12: Schematic overview of a GAN shows the relation between data, generator, discriminator and the resulting feedback loop. Image adapted from (Gharakhanian, 2016).

When trained on a dataset, the generator produces new data to mimic the structures identified in the given data. The discriminator's job is to identify if the newly produced images are authentic or artificially generated. Through the process of generating, assessing, and adjusting, the GAN learns to produce new data with increasing fidelity and realism (Goodfellow, 2020).

GANs can be trained on various data, like images, videos, sounds, or text, making it a versatile tool with numerous application opportunities. Prominent examples of GAN application are the generation of faces (Karras et al., 2019), transferring the style of one image to another (Zhu et al., 2017), or elaborate text generation with a network like GPT-3 (Brown et al., 2020).

Over the past years, the scientific attention towards GANs (Figure 13) increased noticeably, with over 3000 mentions on Scopus in 2019 (Salehi et al., 2020). The capabilities of GANs increased equally, as shown in figure 14.




Figure 13: Number of publications mentioning GANs on SCOPUS over the years.

Figure adapted from (Salehi et al., 2020)



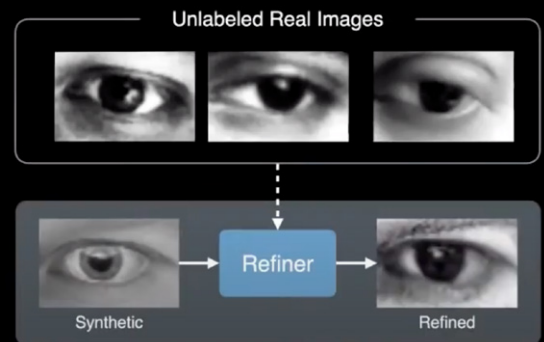
Figure 14: Faces generated with various GAN models show the yearly progress in resolution and fidelity from 2014 to 2018.

Figure adapted from (Salehi et al., 2020)

At a conference in 2020, titled “GANs for Good”, Ian Goodfellow presented some more examples of how this technology can be and is already used in practice (DeepLearning.AI, 2020, 04:09-13:50):

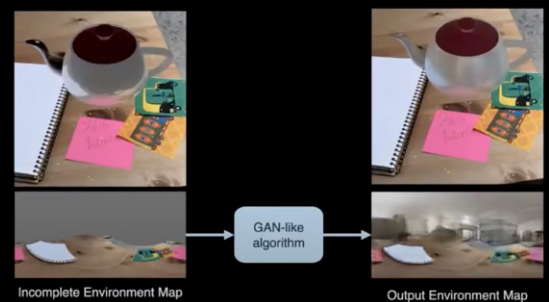
1. Generating medical data

The medical sector is always in need of data for training purposes. As medical information often poses privacy concerns, this data often can not be used to its fullest without further anonymization. GANs can be of help by taking rich medical data as an input and generating new, synthetic but realistic data for training purposes.



2. Apple AR

3D models in augmented reality applications require lighting information to be rendered realistically, however this information is seldom available. GANs can support this process by generating realistic environmental images just by getting an incomplete snippet as an input.



3. Generating dental crowns

Fitting dental crowns to a patient’s mouth can be a labour-intensive process that often requires the use of a temporary crown. GANs can speed up the process by taking images of the jar as an input and generating fitting crowns in return.


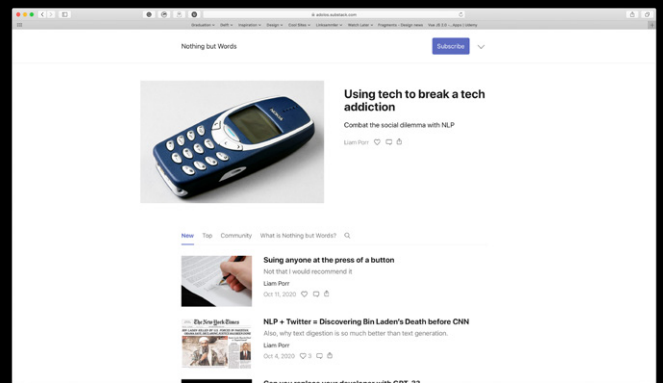


Figure 15: Examples of practical use cases of GANs show the potential of the technology

The list of different GAN architectures and their corresponding examples grows every week. A selection of some recent examples, also with consideration to the potential usage for designers, are:

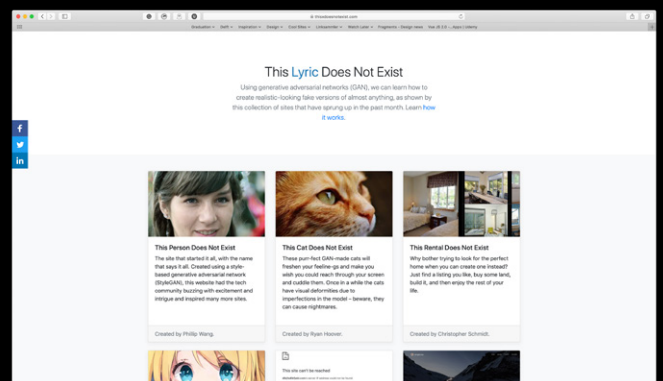
1. Building an AI-generated blog

Using the language model GPT-3, Porrr created a blog whose content was purely AI-generated, and yet succeeded to attract thousands of readers (Porrr, n.d.).



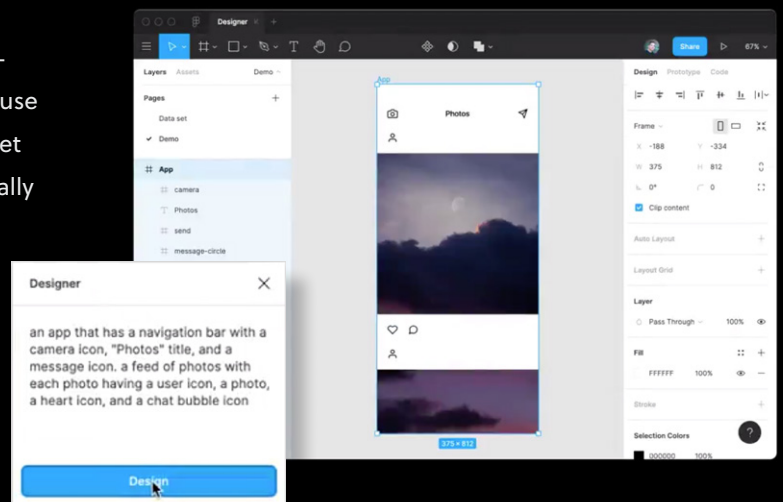
2. This X does not exist

A collection of examples hosted online that use GANs to generate novel content, like human faces, cats, web designs, or vases (Hora, n.d.).



3. Text to design

Modern design tools often work with a text-based file structure. This allows us to again use the language GPT-3 to write an intent and let it create the resulting design file automatically (Singer, 2020).



4. Text to illustrations

In 2020, OpenAI published a novel model called DALL-e that builds upon their GPT-3 model to create thousands of images and illustrations fitting to a text prompt alone (OpenAI, 2021).


TEXT PROMPT an illustration of a baby fox in a suit playing a grand piano

AI-GENERATED
IMAGES



5. Image Style Transfer

One of the first and most prominent examples of GAN capability allows transferring the style of one image or a set of images, like the style of a famous painter, to any other image (Gatys et al., 2015).



02.09 INITIAL EXPLORATION EXPERIENCE

Designers, myself included, often work visually to communicate or even to guide their thought process. The recent advantages in GAN research as well as their often visual nature make it an intriguing technology to explore further from a design perspective.

Throughout my project I logged my weekly learning experience with GANs in a diary format, starting as a novice completely from scratch and trying to work myself into the matter. As the output of GANs can touch every thinkable data format, I restricted myself to visual outputs and guided by the second theme of my graduation – climate change – to earth imagery of some kind, in other words: satellite or landscape images.

In the following, I will present some of the intermediate outcomes, platforms, and tools I used and things I learned on the way.

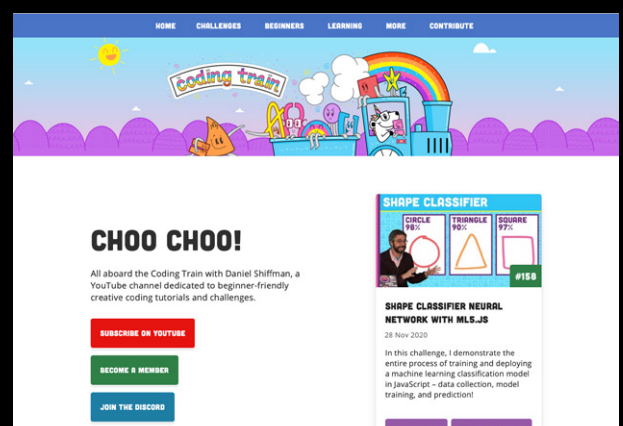
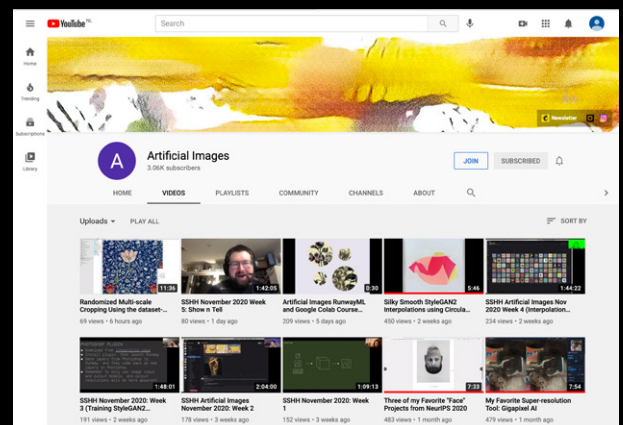
Knowledge resources

The Coding Train

While usually focussed on creative coding tutorials using non-ML tools, The Coding Train's introduction videos to Runway (Shifman, n.d.) are an easy way to get started with ML.

Artificial Images

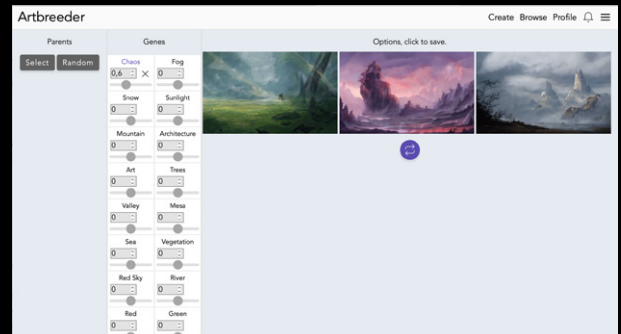
Derrick Schultz is an artist and programmer who continuously creates and publishes content and online code repositories (Schultz, n.d.) specialized in creating visual art with GANs and turned out to be an incredibly useful source in my exploration process.



Platforms and tools used

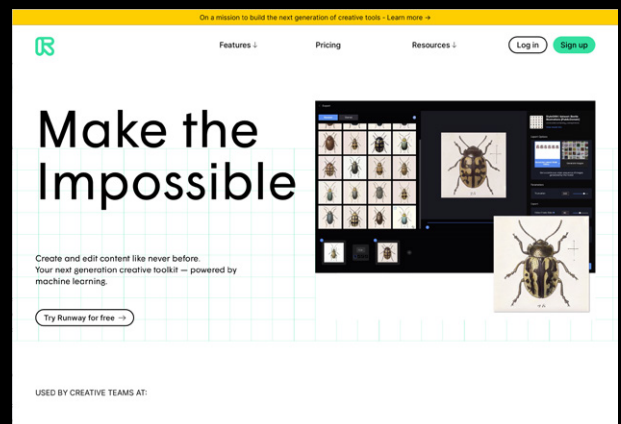
1. Artbreeder

As the name suggests, Artbreeder is an online platform focussed on creating art – through “breeding” various images together. Its hosted GAN models allow for easy still image and video synthesis (Artbreeder, n.d.-a)..



2. Runway

Usually, machine learning requires some coding knowledge. Runway is a platform with the goal of making machine learning accessible for designers & creatives through an easy-to-use interface and connected servers to render content (Runway, n.d.).



3. ml5.js

With the motto “Friendly machine learning for the web!”, ml5.js strives to be an easy-to-use coding library, build upon TensorFlow, to allow creators to use machine learning in their projects in a beginner-friendly manner (ml5.js, n.d.).



LOG 1: A FIRST STYLE TRANSFER

Guided by some video resources by Daniel Shiffman on The Coding Train (Shifman, 2019) and using a pre-trained model (AdaIN-Style-Transfer - Runway, 2019) on the machine learning platform Runway, my first tryout was to transfer the style and characteristics of one image to that of another (Figure 16).

The StyleTransfer model on Runway takes two input images and generates an output image with a merged style. Both images I obtained by taking screenshots of different locations on Google Earth and uploaded them through the Runway Interface.

The needed calculations of the model can be done either locally or remotely server-side on dedicated graphic cards, where only the result is then sent back to your local machine. Higher resolution images required remote execution of the model. The benefit is that I can run intensive graphic calculations even on my somewhat dated computer (MacBook Pro Retina, 2012).

Presumably, to pay for the necessary infrastructure and develop their service, Runway charges a small fee for each run of the model, in the case of my StyleTransfer a few cents per image. This makes it not a free service, but (at the stage of my tryout) Runway offered 15€ worth of free credit when signing up.

Takeaway

It is surprisingly easy to get something out of the StyleTransfer model, although the results might be a bit funny or not particularly realistic looking (Figure 17). From the first experience, Runways service offers a user-friendly way to get a quick result without the need to code and looks promising for further exploration.

However, I was surprised by the computational power needed to perform an image generation, at least at higher resolutions.

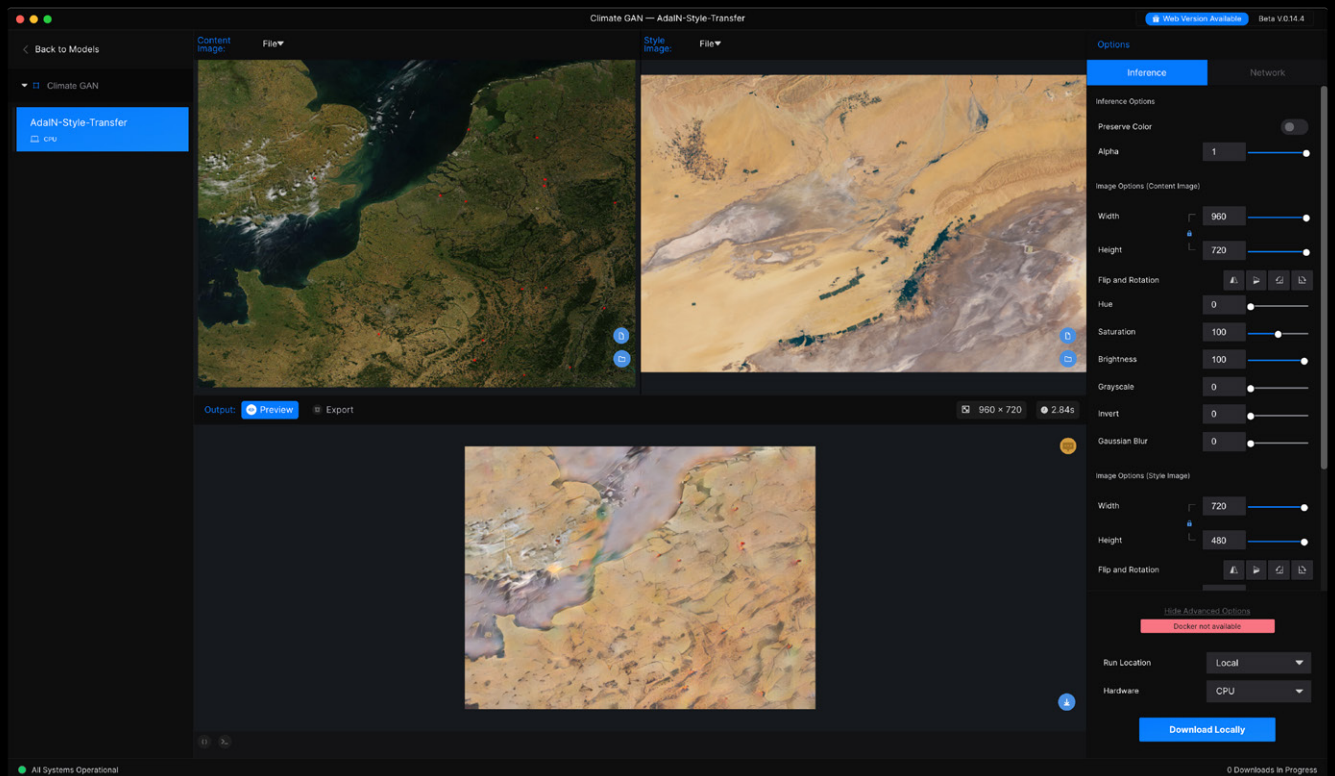


Figure 16: A screenshot of the Runway interface, showing the ease of combining the style of two images with a graphical user interface




Figure 17: The results of a style-transfer can sometimes look hallucinating

LOG 2: DISCOVERING THE LATENT SPACE

For my second experiment, I took a closer look at the Artbreeder, an online service that allows you to look at AI-generated images or generate new images yourself. The free version allows you to upload 3 images per month and download 5 high-resolution images – a higher limit requires a paid subscription (Artbreeder, n.d.-b).

I uploaded an image of the skyline of The Hague with the intention to transform it into something else. Artbreeder (or at least this particular model) can take an uploaded image and embed it into the GAN's latent space (Figure 19). This means that it analyses the uploaded image for its features and tries to find the closed possible alternative that the GAN model can generate. The resulting alternative possesses some resemblance but does not look like the original anymore.

Artbreeder offers the possibility to "cross-breed" different images (Figure 18). For instance, I can take a wintery image and apply the style to my upload. Additional sliders offer some form of additional control, but I found it still challenging to control the output in a meaningful way. Another option is to seamlessly transition from one image to another, a so-called "latent walk" (Figure 20).

Takeaway

Using Artbreeder seems slower and more cumbersome than using Runway. The added abstraction layer makes it difficult to get a sense of what is going on, one element that adds to the challenge of making the output of the models controllable in a meaningful way. However, the morphing of different images using the latent walk transports a fascinating quality that I want to explore further.

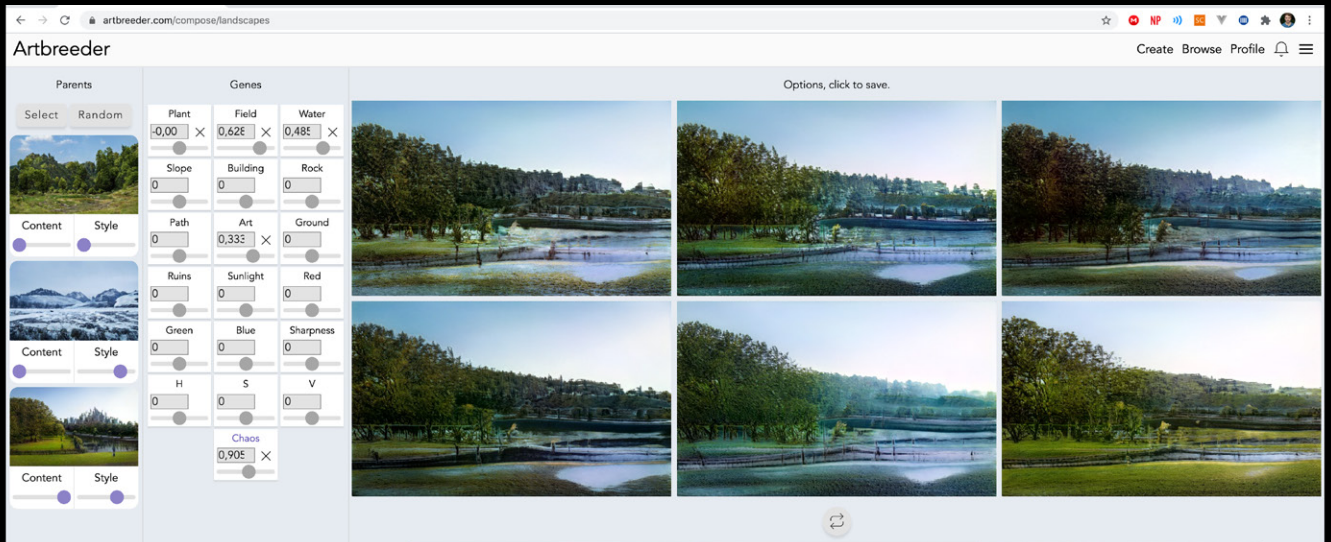


Figure 18: Artbreeder allows the mixing of different images, their content and style through a web interface



Figure 19: The original image (left) gets translated into the most similar looking image in the latent space (right)





Figure 20: A latent walk calculates the intermediate steps and their corresponding images between to target images

LOG 3: THE ISSUE OF LIMITED COMPUTING POWER

The third platform I want to take a look at is ML5.js. Contrary to Runway and Artbreeder, ML5 is not a service, but a coding library that is open-source and accessible to anyone. The added transparency of the codebase might help in the controllability of the output.

Using a provided demo sketch of ML5, I tried to apply the style of a painting to an input image (Figure 21), similar to Log 1 with Runway. However, in this case, the model is trained to always apply the style of that one specific painting to the new input image, and not just combine two completely unrelated images.

I dared a second attempt by running a DCGAN model trained on satellite images to produce new low-resolution satellite images.

Takeaway

In this experiment, I ran into the first hardware-related wall. Even only applying the StyleTransfer model requires too much computing power for my MacBook to provide, which made it very slow to run and often even crash. Training such a model myself for a different image would require even more computing power.

The DCGAN model on the other hand ran surprisingly fast, although the output resembled noise more than anything else (Figure 22). So the level of detail and resolution seems to be a tradeoff to the speed of computation.

For further exploration, the server-side approaches of Artbreeder or Runway seem more suitable for my limited hardware availability.


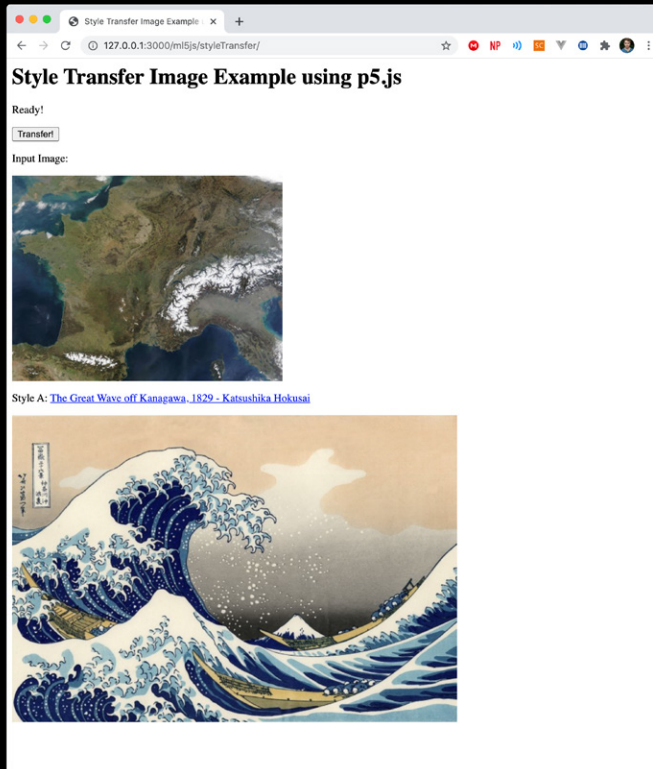


Figure 21: Combining the shape of a custom image (top left) with the style of a pre-trained model (bottom left) to calculate a combination of the two images (right)




Figure 22: The output images of the DCGAN model are quickly calculated, but low in resolution and noisy

LOG 4: GENERATING LANDSCAPES

Returning to Runway as a platform for experimentation, I looked at some other models it had to offer, two of which were created with the StyleGAN architecture mentioned in chapter 02.08. While one StyleGAN was only trained on (assumed by the name of the model) satellite images of Google Earth (GoogleEarth_02 - Runway, 2020), the other model (Runway, 2019) offered the generation of different kinds of images, selectable through choosing a different checkpoint in the interface.

In contrast to the StyleTransfer model explored prior in Log 1, those two models take a random vector input and based on that generates a grid of output pictures, which can then be downloaded. The only two options for further control are a slider for “Truncation” and “Sampling distance”, but no further explanation on what those mean is offered.

Takeaway

It showed again that Runway’s server-side rendering is way faster than anything on my personal computer, so this will be the way to go for now. The user-friendly interface also makes it easy to generate output and explore new models without the need to code, yet this added layer of abstraction also functions as a barrier for me to explore the underlying technology further and explore how to control and manipulate the output. Another barrier is the specific lingo used to control the output. What is the vector distance? What does truncation mean? Those presumably technical terms without further explanation hinder the understanding process of the limited given controls.

Taking a look at the output of both models, both can sometimes develop a more “arty” style, where this is more noticeable in the generated landscapes (Figure 24) in comparison to the top-down satellite images (Figure 23). This type of aesthetic, the need for server infrastructure, and limited controllability might be critical design constraints to keep in mind.



Figure 23: *Generated satellite images using StyleGAN*



Figure 24: *Generated landscape images using StyleGAN*

LOG 5: LATENT WALK VS CONVENTIONAL TRANSITION

In addition to the preview generation, Runway offers a separate section for exporting media. The two options are either to generate a folder of random images or generate a latent walk video, similar to the output of Artbreeder showcased in Log 2.

The image option offers truncation as well as the number of generated images as export parameters, while the latent walk video allows the user to select keyframes from a selection of random images and determine the duration of the transition. Once satisfied with the parameters, Runway starts to export either video or images and notifies the user once the server-side generation is completed.

Using the pre-trained Google Earth model, I created a latent walk video between two images. In addition, I exported the two keyframe images of the video still images and attempted to replicate the morphing effect through the pixel shift method in the video editing tool Adobe AfterEffect (Figure 25).

Takeaway

The transition of one image to the other generated by the GAN and the fluidity of it transports a mesmerizing and aesthetically pleasing quality that might be worth exploring further. The latent walk transition feels like a smoother transition of higher quality compared to the image morph replica attempted in AfterEffects and makes me feel and appreciate the beauty of math and vector interpolation.

The possibility to generate thousands of images at once, if required, is also a potential of the technology that I did not consider like this before. Yet, if using a service of Runway, extensive amounts of video or image generation can get quite expensive quickly with around 5€ per 500 images.




Figure 25: *After Effects Morph (left) and GAN latent walk (right) in comparison show the additional detail of the GAN output*

LOG 6: A VAST UNIVERSE OF GANS

Looking through the offered models on Runway, but also beyond by considering the current academic publications of GANs, it becomes apparent that my explorations so far only touched the surface. A user-curated list on Github (Hindupur, 2018) now lists over 500 different GAN models, each with its niche or novel approach. Which kind of GAN I could use for further exploration depends on the resource accessibility in terms of hosted models and teaching material as well as subjective usefulness for the project.

Let's take a look at some of the more prominent visual GANs and other synthesizing networks:

StyleGAN

StyleGAN (Karras et al., 2019) takes an input vector, meaning an array of 512 numbers, to generate novel images. The landscapes and satellite images from previous logs were generated with this model. It, as well as its successor StyleGAN2 and StyleGAN2-ADA, are trained on a big collection of images and learn to distill features from those images to generate realistic-looking results. The models are quite popular and teaching materials as well as hosted models on Runway available.

Pix2Pix

Pix2Pix (Isola et al., 2017) is a network for image-to-image translation. In contrast to StyleGAN, Pix2Pix is trained with labeled data, and on two paired datasets. The first dataset describes the input, while the second dataset contains the corresponding output. Then, when presented with new input, the network can generate a novel output based on its prior learning. The labeling, as well as visual input, might increase controllability for those networks.

CycleGAN

CycleGAN (Zhu et al., 2017) pursues a similar approach as Pix2Pix and aims to translate one type of image to another. Unlike Pix2Pix, CycleGAN does not require paired images for the training process.

SPADE

SPADE (Park et al., 2019) was developed by researchers at Nvidia and allows to translate a segmentation map into a photorealistic image. The novelty of this approach is that the style of the image can be controlled separately from the semantic structure.

BigGAN

As the name suggests, BigGAN's approach is to go big and increase the training size for the network immensely, "with two to four times as many parameters and eight times the batch size compared to prior art." (Brock et al., 2018, p. 1). They also incorporated some other tricks and learnings from prior GAN developments, like the ability to generate base on an input vector as well as a labeled class.

Takeaway

The issue when working with GANs is certainly not a lack of choice. Each model pursues a sometimes slightly, sometimes more drastically different approach. Understanding those differences and their significance, however, requires some deeper knowledge and understanding of the space. The biggest difference from a design perspective seems to be the varying input, either in form of an image, a class, or a vector.

Focussing on one of them seems like a reasonable decision to not get lost in all the possibilities. Judging from the availability of teaching material and hosted models, StyleGAN seems to be a good candidate for further use.

LOG 7: TRAINING A CUSTOM STYLEGAN

After using pre-trained GANs to generate new output, it was time to take the next step and train my own StyleGAN. Runway offers the option to upload a custom dataset, in the case of image synthesis make a choice between StyleGAN and StyleGAN2 as a training architecture and start the process. Runway requires using the method of transfer learning in contrast to training from scratch. In transfer learning, a prior training checkpoint on another dataset (for instance to generate cars or faces) is used and the network then exposed to the new data to learn the new features and how to generate them. Transfer learning is quicker than ground-up learning and also more stable. In the following I will go through the learning process step by step:

1. Searching for data

Suitable data in the right format and sufficient quantity is a key requirement for StyleGAN training. Finding accessible data was the most challenging part of this training process, as aerial images are often inaccessible without purchasing them, requiring you to download them one by one through an interface or do not come in the right format.

I ended up using the AID dataset (Figure 26), a collection of aerial photos published by Xia et al. (2017), to train the network. The dataset contains 10.000 images with a 600x600 pixel resolution in RGB format, structured in folders. I selected only folders containing ~2.559 images without prominent human presence, like farmland, deserts, and mountains.



Figure 26: Examples from the AID dataset

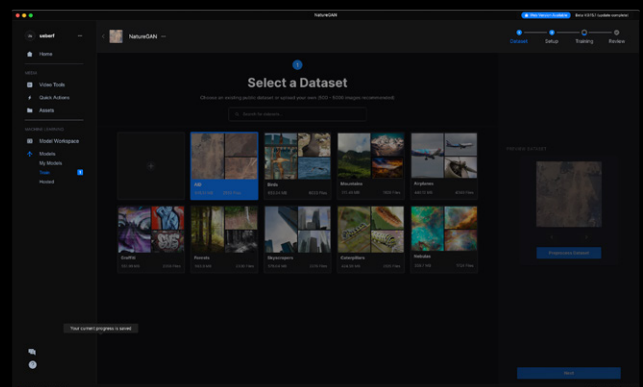


Figure 27: Interface to select a custom dataset

2. Uploading and preparing the data

Making the data accessible for the training process was a matter of uploading them through the Runway interface (Figure 27). GAN training often requires the images to be square and in the same resolution in order for the network to learn the image features correctly. As all AID images already came in a square format with a unified resolution, no additional processing had to be done.

3. Training

For my training, I am using StyleGAN2 with the faces as a starting point for the transfer learning (Figure 28) and chose the default amount of 3000 training steps. After clicking the starting button, the training began.

The training process can be observed by generated samples that Runway automatically provided after each training step (Figure 30). Through those samples, one can see how the network incorporates the new data and how the generated faces slowly transition into generated landscapes with each new step (Figure 31).

4. Generating

After the training process is completed, the model can be used just like every other Runway StyleGAN model to export images or latent walk videos.

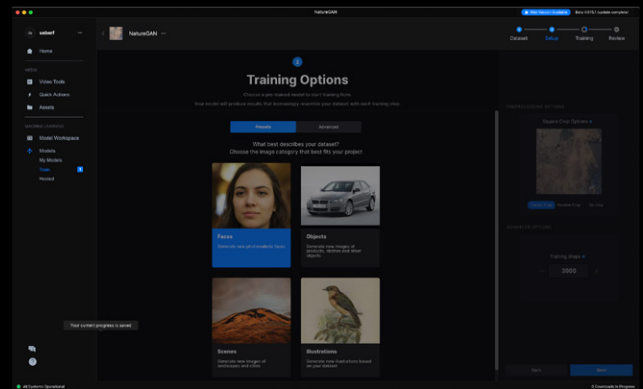


Figure 28: Different pre-trained models provide a starting point

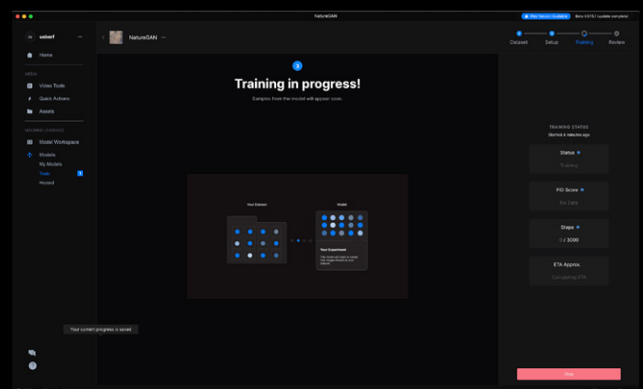


Figure 29: The training process can be observed

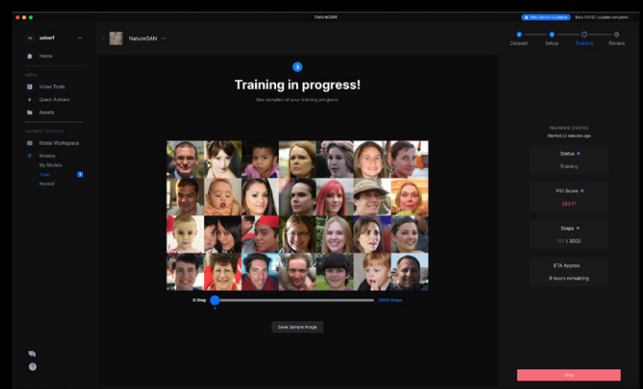


Figure 30: Generated samples during the training process

Takeaway:

Although the training process of a StyleGAN model on Runway was surprisingly easy, I felt strangely proud to get something out of my own trained network. In my case, the training was also surprisingly quick with 3-4 hours to complete the first 3000 steps. The biggest challenge in the process was finding the right data, as GANs usually need over 1000 images to train on in order to perform well.

However, the output was still lacking some of the qualities I hoped for. GANs act like mirrors – what you show them, they will also reflect back in a slightly transformed way. This increases the importance of the right dataset even further to make sure that the generated images fit the output you envision. In my case, the output images (Figure 32) are muted in color, the same as the characteristics of the dataset, and seem to be missing water bodies and beaches. The output has a resolution of 1024x1024 pixels, but as the input data was only 600x600 pixels, it is lacking in visual detail.



Figure 31: *With transfer learning, the GAN slowly picks up on the new dataset and learns to generate landscapes instead of faces.*




Figure 32: *Six examples of the possible output of the newly trained GAN*



02.10 DIFFICULTIES FACED ON THE WAY WHEN WORKING WITH GANS

My exploration process with GANs showed me that it is possible, even with my limited knowledge, to get results out of those networks. However, I also identified some limitations and difficulties that make it difficult to get into the space of AI and potentially use it for a later design.

Controllability

A trained GAN model takes an input, in the case of styleGAN a 512-dimensional vector (a list with 512 numbers between -1 and 1), to generate an output image. Changing numbers in this input vector results in a different image, with small changes resulting in a probably very similar image, and big changes in a probably very different image. The complexity of the input makes it difficult to control for a human and to generate the exact results you envision.

The controllability gets complicated further through so-called feature entanglement. Feature entanglement means that changing one dimension of the vector changes more than one feature of the image. To give a simplified example: Let's say we want to train a GAN on images of dogs with the intent to generate new images of dogs. In our dataset of existing dog images, white dogs happen to always look to the right, while black dogs look to the left. When training and identifying patterns, the GAN does not know that white dogs can also look in other directions – it entangles the feature of viewing direction with the color of the dog. Changing the vectors responsible for the color of the dog might then also result in changing the viewing direction of the dog.

GANs are still a novel technology, and projects that aim to make GANs more understandable and controllable are still rare and under active development. Two examples I found are GANSpace (Härkönen et al., 2020) and TL-GAN (Guan, 2018), which offer the ability to analyze a GAN and add interpretable controls (Figure 33). Yet, those approaches require some advanced (coding) knowledge as well as access to the GAN source file to attempt to add controllability.

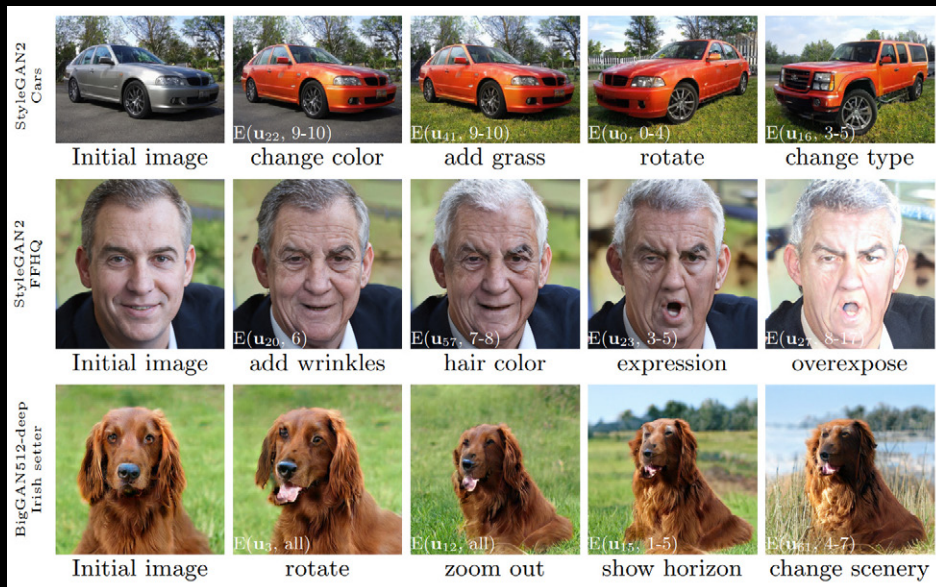


Figure 33: GANSpace allows to detect some feature vectors in GANs and make them somewhat controllable

Resource & time intensity

Generating an image with a pre-trained GAN requires access to hardware resources with decent graphic performance or extended time to render the result. Of course, this will be different for every GAN, but generating high-resolution imagery (even with good graphic cards) in real-time is still an ongoing challenge. Services like Google Colab or Runway offer the possibility to render on their servers equipped with powerful graphic cards, but this requires either paying for it (Runway) or coding knowledge (Google Colab).

Training a GAN is even more compute and time-intensive than generating an output with a pre-trained model. Services like Google Colab and Runway offer the possibility to train on their servers as well, but the required training time can still be hours, sometimes days, or even weeks to get usable results.

Lacking transparency

The lack of transparency, as discussed in section 02.06 already, is also present in GANs. It is hard to understand why a GAN does (or does not do) what it does, which makes it difficult to work with.

Services like Runway and Artbreeder add a layer of abstraction between the technical base and the user controls, which makes it easier to get started, but also in some instances more difficult to understand. Often the user does not have access to the underlying dataset or trained GAN model, which hinders the process of analyzing, understanding, and then potentially controlling it further.

Accessibility

GANs and the ongoing research concerning them are without a doubt a deeply technical field. The technical complexity (and so far necessity) makes it a challenging space to explore from a novice perspective and wrap one's head around.

The exploration and learning process is further exacerbated by specialized lingo used in papers, articles, and software around GANs. The use of specialized vocabulary should not come as a surprise, as novices are also seldom the target group of GAN publications, yet even Runway – targeting creatives with their software – makes use of language that make it sometimes difficult to understand. For example, Runway allows the user to control a StyleGAN by choosing a random input vector, but then only further by setting a Truncation as well as Sampling Distance value (Figure 34). What those settings mean is not explained further and needs active research to understand.

Another barrier when it comes to training a GAN is the access to usable data that the network could be trained on. Precomposed datasets in the right format are often hard to find, and creating a dataset is often challenging and time-consuming with knowledge in coding required to compile hundreds or thousands of needed images.

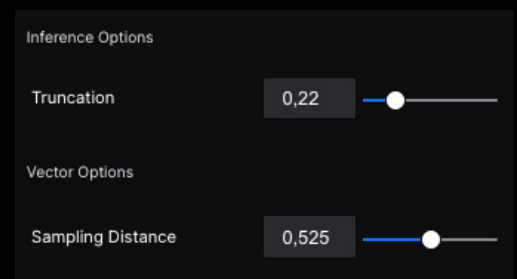


Figure 34: Screenshot of the Runway interface, only showing two values to influence the output

02.11 REFLECTION ON WORKING WITH A GENERATIVE AI MODEL

Despite the difficulties faced along the way, like the lack of transparency, accessibility, or needed time investment, it was surprisingly easy and quick to get some kind of result out of a GAN. The fact that I was able to do some basic exploration without any further technical knowledge is amazing in itself.

A variety of pre-trained models can be utilized to generate visual output using services like Runway or Artbreeder, and once trained, they allow for the speedy output of a multitude of different images. The aesthetic of most GAN output, the possibility to seamlessly morph between images, and the knowledge that it is completely generated transport a strange feeling of fascination.

But what does that mean for the design process? Are GANs already or will be a feasible tool to work with? My intermediate answer is: it depends.

In the design process, I could see how GANs can be used for asset generation or as a tool for inspiration. Controlling a GAN output in a meaningful way remains challenging so far, and once leaving the safe space of utilizing pre-trained models, a steep learning curve awaits anyone who would like to explore more. Yet, the fast progress of GANs over the past years makes it a technology that designers should keep an eye on. When deciding to work with a GAN, it is advisable to either have a decent amount of time to spend on exploration or to work in a multidisciplinary team with computer scientists as a collaboration partner.

02.12 CHAPTER CONCLUSION

In this chapter we took a look at the history, basic workings, use cases, and ongoing development of AI systems, both from a theoretical as well as practical perspective. It is apparent that the capabilities, and as a consequence thereof, the uses of this technology will only increase in the future. This prospect encourages further exploration and familiarization for designers.

With regards to AI as a practical tool in the design process, the initial exploration of GANs for visual output revealed that it is an interesting, yet challenging tool to work with. Barriers of technical complexity, lacking accessibility and controllability, needed hardware and time resources can make it a tedious process, which might not fit every design process, type of designer or warrant the difficulties faced on the way.


My initial exploration allowed me to only scratch the surface, but the successes motivated me to explore GANs further and also explore advanced fields like the curation of a dataset, training as well as debugging of a network. The authors of Noosphere (2020) compare the current state of AI to that of the steam engine before the laws of thermodynamics were discovered and written down. We have something that does something, but not always in a controllable manner, grounded in theory and common understanding.

Equally, I would compare the current stage of AI as a design tool to the early stages of computers. Their potential was perceptible from an early stage, yet the first versions were clunky, limited in capabilities, and needed extensive handbooks and knowledge to be usable. Nowadays, the computer is an essential tool for most design and engineering work, present in most workflows, and the days with rooms full of people drawing blueprints by hand are long gone (Figure 35).

The same potential that was seen in early computers is currently foreseeable for AI systems. It is already possible to use them for the generation of visuals, use them to take over normally tedious work like removing the background of an image or scale-up low-resolution imagery. Yet, AI systems still need extensive research and development work until they might become a truly valuable tool for designers.



Figure 35: *Engineers drawing blueprints before the availability of computers (Vintag, 2018)*



03 THEME EXPLORATION: CLIMATE

Climate is all around us: It shapes the environment we live in, and in recent decades has equally shaped the media landscape. Documentaries about environmental change, news reports about climate catastrophes, or political summits on the topic of environmental actions are all around us. Yet, the topic of climate and climate change can still seem strangely vague and distant, especially in an effort to frame it as something approachable from a concrete design direction.

To get a clearer understanding of the climate context, I conducted a deep dive into literature, guided by keyword search, and carried out interviews with climate researchers from various domains. Although most of the climate research will not be visible in the final design, it was a crucial part of my process that allowed me to check and correct initially made assumptions and build a basis for the following design stages.

The following chapter provides a broad overview of our world climate and summarizes the gathered insights from literature and interviews.

| *Methods used: Literature review and interviews with climate researchers*

We live on a blue planet, orbiting a fireball, with a moon that moves
our oceans, and yet you decided not to believe in miracles?

– Marteria, Welt der Wunder, translation from German

SIDE BLOCK: EXPERT INTERVIEWS WITH CLIMATE RESEARCHERS

When exploring a new context, designers regularly rely on the insights of stakeholders to navigate the unknown space. In my case, I looked for the knowledge and opinions of climate researchers to get a personal perspective on assumptions made in the design brief of the project, fill in gaps and the potential blindspots of my literature review and overall get a clearer understanding of the climate space.

The interviewees were researchers of various climate domains located at universities and research institutions in Germany and The Netherlands. The interviews were conducted remotely via Zoom, and the statements and insights were recorded manually. Insights gathered during the conversations will be mentioned in the following chapter, when relevant.

The main questions targeted in the interviews were:

1. Can you tell me a bit about your work and the process from start to finish?
2. What are the main challenges in the communication of climate findings?
3. What is the goal of the communication?
4. How important is communication to non-scientific stakeholders?
5. Do you know some examples of great communication work?
6. How is machine learning applied in climate science (more specific your work) at the moment?

03.01 OUR EVER-CHANGING CLIMATE

The word “climate” appears regularly in our everyday language use and everyone probably has a rough idea of its general meaning. Outside the strictly scientific context it is even used to describe a “general attitude or feeling”, for instance when we talk about a social or a political situation (Oxford University Press, n.d.). But what does it mean when scientists refer to climate? Is it all about temperature, clouds, and rain?

What we mean when we say “climate” is not the daily changing weather phenomena we see when looking out the window, but the long-term average of the weather at a certain location. The usual period used to create this average is around 30 years, but can also be over thousands of years (Planton, 2013). Each location on earth has its very own climate, but they can be classified into different climate groups (Figure 36), for instance tropical, dry or polar regions, each with more sub-classification (Arnfield, 2020).

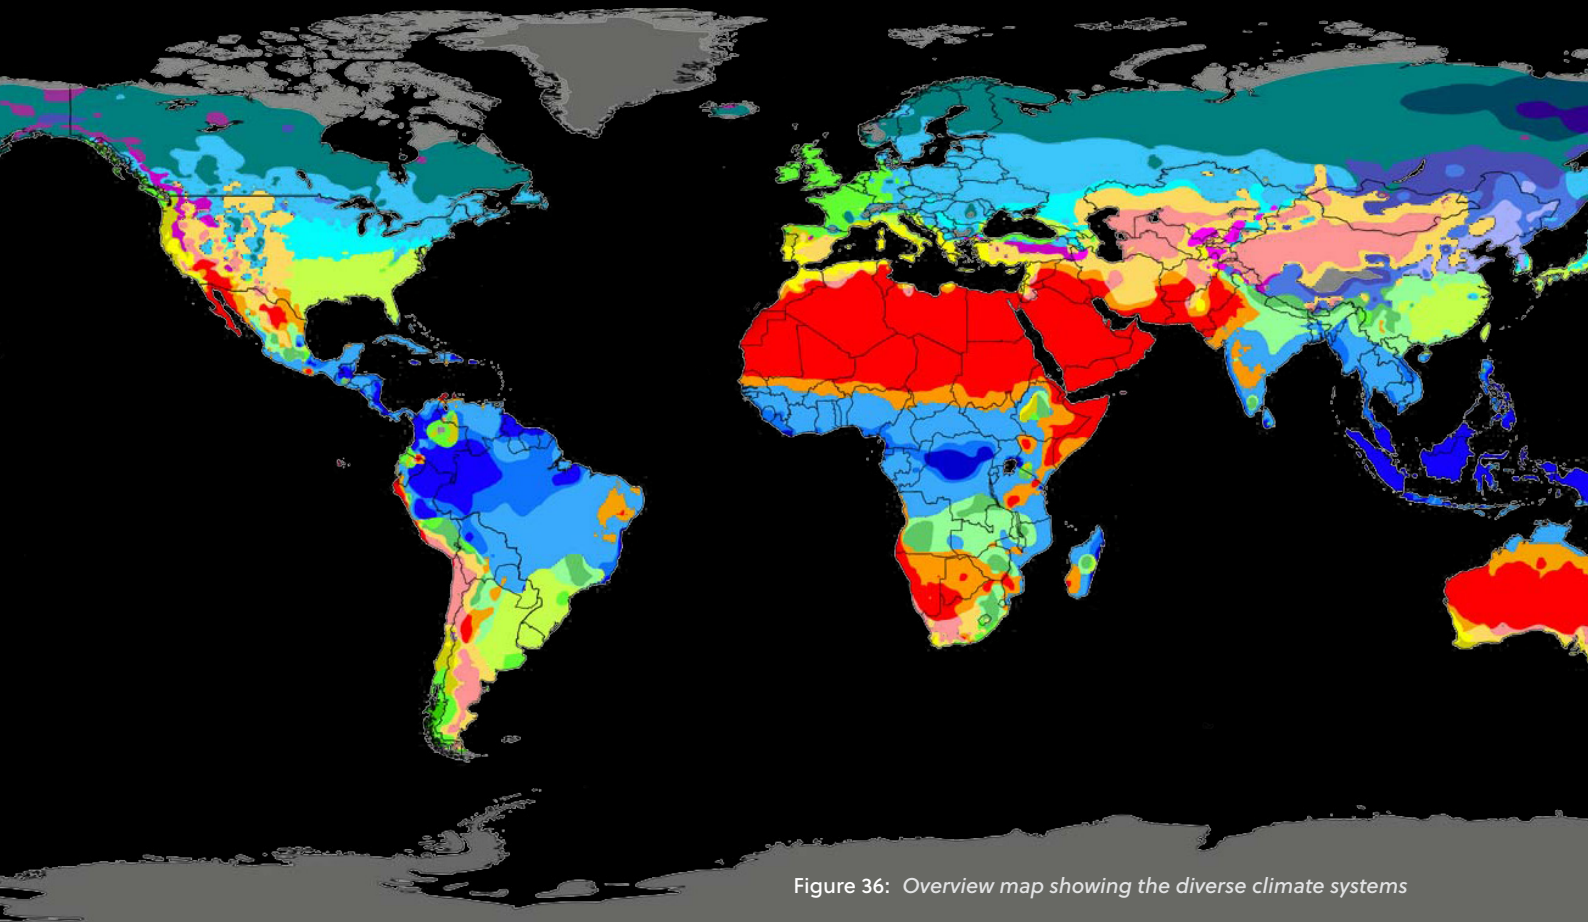


Figure 36: Overview map showing the diverse climate systems

The various subclimates are forming a complex system, also known as our earth's climate system. This ever-changing system consists of and is shaped by the interactions between the subcomponents of the

1. atmosphere (an envelope of gas surrounding the earth),
2. hydrosphere (liquid elements of the earth, like oceans, lakes, rivers, underground reservoirs),
3. cryosphere (ice sheets and frozen grounds),
4. lithosphere (rock and earth layer) and
5. biosphere (all ecosystems of living organisms, on land and in the ocean) (Planton, 2013).

Everything happening on our earth is somehow involved in and influences our climate system. For millions, if not billions of years, did it shape and was shaped by the world we know, and has defined who we are and how we live.

Nowadays it is hard to think of climate without immediately thinking of the notion of climate change, too. But what do we mean when we talk about climate change?

Changes in our earth's climate are driven through so-called "forcings" Natural forcing like solar variability have always influenced the climate system (Lean et al., 1995). An example of a natural forcing could be the eruption of a volcano, like the one of Mount Pinatubo in the Philippines. On June 15, 1991, Mount Pinatubo erupted and tossed millions of tons of material and gases into the higher atmospheres (Figure 37), which caused an increased sunlight reflection and led to a subsequent, measurable cooling of the earth's climate (NASA, 2020; U.S. Geological Survey, n.d.).

When we talk about climate change nowadays we more likely refer to effects caused by non-natural, human-introduced (anthropogenic) forcings, like our emission of greenhouse gases. Greenhouse gases, like carbon dioxide, methane, ozone, but also water vapor exist naturally in our atmosphere (Planton, 2013) and contribute to the earth being a planet inhabitable by biological organisms. By absorbing heat, some coming from the earth itself, but mainly as heat in the form of infrared radiation coming from our sun, they allow life to exist on earth (Planton, 2013).

While organisms also emit carbon dioxide and other greenhouse gases, plants on land and in the sea capture those and bind them while releasing oxygen back into the atmosphere. They function as natural carbon sinks, and once buried underground and with enough time turn into coal, oil, and gas (Bolin, 1970).



Figure 37: Photo showing the scale of the Mount Pinatubo eruption

The beginning of the industrial revolution in 1760 brought an increased need for energy and correspondingly a rapid increase in burning fossil fuels like oil, coal, and gas, to meet the demands (Ritchie, 2017). Burning fossil fuels releases the prior captured greenhouse gases back into the atmosphere, where they contribute to increased absorption of heat radiation, and eventually measurable warming of the That we have an impact on our earth and contribute to its warming by emitting greenhouse gases and removing natural carbon captures like forests are already known for over a century. The first time this particular phenomenon, which later should be known as the "greenhouse effect", was theorized was in 1896 by the Swedish scientist Svante Arrhenius (Spencer Weart & American Institute of Physics, 2017b). Measurements and calculations conducted during the following decades support the theory, and today more than 97% of scientists agree that anthropogenic climate change is real and that its impact on our earth's climate undeniable (Anderegg et al., 2010).




Figure 38: Correlation of CO₂ emission and global temperature increase between 1880 and 2019.

Figure from (Climate Central, 2020).

But is a change in our climate a bad thing? After all, changes in our climate were always present and are only natural. While this might be the case, the issues at hand are not the changes in themselves, but rather the rapid speed they are happening in. These natural changes, if not caused by an unexpected event like the eruption of a volcano or impact of a meteor, normally take thousands, if not millions of years to take place. Their gradual change allows ecosystems and organisms (humans included) to adapt to the changing environment. The rapid surge in the emission of greenhouse gases, the consequential drastically increased warming of our earth, and change in climate will alter our living space forever (Kerr, 2007).

Although these predictions come with some uncertainty, especially when it comes to the exact timing and intensity of these events, scientists current predictions (Spencer Weart & American Institute of Physics, 2017a) of global warming and climate change impact are:

6. The earth will get warmer with some regions suffering from severe heat waves, with the increased temperature also leading to an increased spread of tropical diseases
7. Sea levels will rise significantly due to glacier and ice sheet melting, endangering particularly low lying regions and coastal cities
8. Weather patterns will keep changing and become more unpredictable and extreme
9. Ecosystems will be stressed. The ocean for instance will get more acidic, endangering coral reefs and marine ecosystems
10. Many more unforeseen impacts will happen which we can't predict yet

Some of these impacts, like rising sea levels, increased number and intensity of hurricanes or more wildfires, can be seen and experienced today already (European Commission & European Commission. Climate Action DG., 2018, pp. 8–14).

But how do we know that this will all happen? The next chapter will examine the scientific process from data to future climate predictions

03.02 FROM THE PAST TO THE PRESENT TO THE FUTURE

As shown and discussed in the previous chapter, our climate is an incredibly complex system with many interlinking subcomponents that influence each other. Equally many scientific disciplines are involved to get an understanding of the space and understand all those connections, looking at it from a global scale down to the microscopic level. But how do we know what happened to our climate in the past and what will happen to our climate in the future, at least to a degree of certainty?

The whole process of getting an understanding of our climate system is a long and intense process that will be individual for each specific research project. In one of my context interviews, a climate modeler from Utrecht University talked me through their research process of understanding and predicting the development of arctic ice sheets. Their (now very simplified) process starts with the collection of data, remotely via satellite imagery and on-site by local researchers, to determine the distribution, thickness, and consistency of ice sheets. For this remote assessment, they have to calculate and consider the relative position of the sun and the satellite at the time of capture and have a good understanding of the reflectiveness of snow and ice (Albedo) to analyze the satellite image correctly. The remote data gets extended and checked by local sensor readings, or ice core samples obtained from the depth of a glacier by on-site research teams.

The data, together with outputs and findings of other climate researchers, get analyzed for patterns and correlations between forcings and ice sheet development and then translated into a climate model. A climate model is a mathematical equation that sums up the processes and developments and exists in many different forms and complexities, ranging from local, high resolution, to global low-resolution models. Increased resolution and complexity mean higher computing power to get to a result, so the balancing act is a tradeoff between computing power needed to perform all calculations of the model and the gathering of predictions to be useful for the scientific process.

One way of testing the validity of models is to “hindcast”. Hindcasting is the process of producing “a forecast made for a period in the past using only information

available before the beginning of the forecast” (Planton, 2013). Since the outcome of that period is already known through recordings, it can be used to judge and tweak a model’s capabilities and accuracy. If the predictions of past events are correct, meaning that the relations and interactions between forcings and outcomes are captured accurately, it suggests a high probability to also accurately predict a future climate development based on changing inputs. Lastly, the data, the created model, and its output are then communicated further to other stakeholders, for instance through papers, university press publications, and added to public databases, so that it can be used by other researchers.

No model is perfect and can deliver the ultimate truth of climate developments on a global scale. Therefore, many models are taken to predict future development to see if there is a distinguishable, general trend that can be used for a climate scenario. Climate or climate change scenarios are a usually simplified projection of the future climate development, that takes current emission scenarios, the current state of the climate, and predicted climate developments into account (Planton, 2013). They are used mainly for outside communication (like in IPCC reports), and to give policymakers and other stakeholders a solid direction to base their interventions on (Moss et al., 2010). The creation of a climate scenario is again a complex process of different disciplines working together, checking each other, and basing the projection process on different model outcomes. This whole process is often very energy, labor, and time-intensive, as so many disciplines have to wait for input from other sciences to continue or check their work (Moss et al., 2010).

03.03 ACTING AGAINST CLIMATE CHANGE

The previous sections described the foundational elements of our climate system, as well as the scientific abilities to measure, compute and predict its changes. Now that we have an overview of our climate system, climate change, and its likely impacts, the question arises: what can we do about it? And who should be doing it? Due to the complexity of the climate space, I can only give a broad overview, based on the information found in literature and my conversations with scientists. Nevertheless, even on a rough level, can the system map and its relations be used to get a better understanding of the space and identify potential opportunities for a later design direction.

03.04 REPRESENTED STAKEHOLDERS

Since climate change is happening on a global scale, at some places more severe than in other parts of the world, its impact should and will concern everyone sooner or later. The stakeholder network of climate (in)action is as complex as the climate system itself, starting with universities and other research institutions, which are conducting research about climate change and informing other stakeholders. These stakeholders include politicians and policymakers, private and public companies of all sizes, non-governmental organizations, and of course the general public, who all have the choice to act or not act on scientific insights and advice, and need to balance it with other interests and constraints. Other acting stakeholders are the media, overarching climate organizations like the Intergovernmental Panel on Climate Change (IPCC), and industry lobbyists.

It is important to notice that the interaction between science and other stakeholders is not a one-way street. Research institutions depend on governmental funds and funding from private companies to conduct their work, and a change in governance can have a severe impact on its capabilities. To give an example: after pulling out of the Paris climate accords in 2017, the US administration under Donald Trump continued to exacerbate the work of climate scientists, for instance by quietly defunding the NASA Carbon Monitoring System (Voosen, 2018), a system with the

aim to better track sources, sinks and movements of carbon on our planet.

Likewise is the connection from companies and governments to the public no one-way-street but in democratic systems a constant interaction. Citizens have a direct impact on companies through their daily consumer decisions and an impact on governments through their vote in elections. They also have the opportunity to apply additional pressure through activism and demonstrations, giving their views and opinions a voice. The media functions as an important connector between science, governments, and companies, and the public and helps communicate insights and decisions understandably.

With so many stakeholders involved, it is easy to point fingers and shift the responsibility from oneself to other parties. Often, the distinction made is between individual actions that everyone can be doing daily or systemic changes that influence causes of climate change on a bigger level. While some stakeholders might have a bigger impact and leverage than others, in the end, everyone is affecting climate change through their choices and is affected by made decisions in return.

03.05 CURRENT AND POTENTIAL CLIMATE ACTIONS

Measurements against climate change are already taken today to prevent its main causes or to prepare and shield oneself from climate impacts. The Paris Agreement, signed by over 195 in 2015, aims to coordinate global governmental efforts and encourage every participating country to set ambitious goals to keep global warming under 2°C and limit as well as prepare for the impacts of climate change (United Nations Framework Convention on Climate Change, n.d.). Nonprofits like The Climate Group (2003) and their initiative RE100 (n.d.) aim to coordinate and help with the same transition, informed by the Paris Agreement, for the business sector, while campaigns like “Act Now” by the United Nations (2018) aim to guide people in their personal transition to a more sustainable lifestyle. The big consensus with all those initiatives is that we have to cut global emissions drastically while also starting to recapture the greenhouse gases we emitted into the atmosphere in order to limit the damaging effects of climate change.

Project Drawdown, a nonprofit organization founded in 2014, is one of many sources for climate insights and concrete actions to tackle global warming. Their list of concrete, impactful solutions (Project Drawdown, 2020) is over 80 entries long and divided into three connected areas:

1. Reduce Sources — bringing emissions to zero
2. Support Sinks — uplifting nature's carbon cycle
3. Improve Society — fostering equality for all

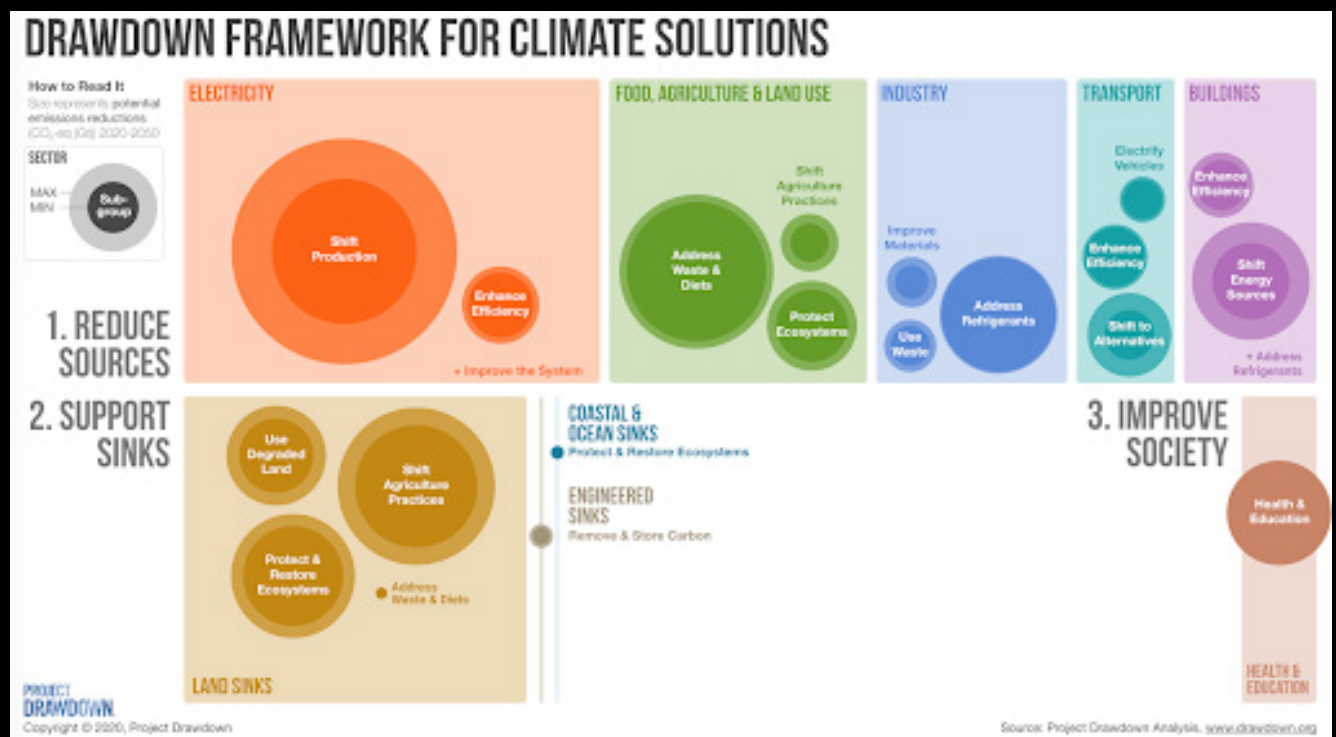


Figure 39: Solution map of Project Drawdown summarizes the three main targets (Project Drawdown, 2020)

Other initiatives, like Climate Action Tracker, analyze governmental actions taken and compare them against the set goals of the Paris Agreement. Their latest report from 2019, while being able to recognize some positive developments, summarizes that there still is a significant gap between set global climate targets and current actions taken to achieve those goals, with “[...] many governments are still failing to meet their often insufficient targets.” (Climate Action Tracker, 2019).

With the urgency being clear and solutions available, it begs the question of why not more is done already to tackle climate change on a global as well as national and individual level? While this question – as it seems so often when it comes to climate change – is too complex to answer in a simple way, two hindering factors are the invisibility of the issue as well as conflicting motivations.

As Weber & Stern (2011) point out, “some fundamental attributes of climate change make it hard to understand” with the “main causes of climate change (greenhouse gases) [being] invisible, its impact [...] geographically and temporarily distant for most [...] and its signals hard to detect.”. Reduced visibility here might lead to a reduced sense of urgency, both in citizens as well as politics. Likewise isn’t climate change the only global issue, but only 1 of 22 listed by the United Nations, and fighting it is not everyone’s main concern. On the contrary, fossil fuel companies continue to spread misinformation about the causes of climate change and lobby against political climate actions that would harm their greenhouse gas emitting businesses (Oreskes, 2018).

And even with the best intentions, acting against climate change is not a straightforward process, but a delicate balancing act of funds, different interests, and prioritization. The tradeoffs don’t need to be clear opposites, but can be for example the decision between investing money now, prevent climate change but lack the funds for social initiatives, or use the money at a potential later point to limit the change’s impacts, or prioritizing reduction of pollution versus the reduction of emission.

The Montreal Protocol is another example that showcases for one the complexity of these decisions, but also the power of acting decisively and commonly to solve a global issue. In 1987, all 197 member countries of the United Nations came together to sign a treaty aimed at reducing the global use and emission of substances damaging the atmosphere's ozone layer (United Nations Environment Programme, n.d.). The ozone layer is an important aspect of our atmosphere, responsible for blocking sunlight radiation in the form of ultraviolet light from reaching the earth's surface (NASA, 1999) and keeping crops and humans safe from cell damage (Molina & Zaelke, 2017). After discovering the damaging effects of so-called chlorofluorocarbons, a compound often used in fridges, to the ozone layer, the treaty was set in place, planning the gradual phase-out and replacement of these harmful substances (European Commission & European Commission. Climate Action DG., 2018, p. 5). The plan was a success, also due to the establishment of a multilateral fund that should help developing countries with the transition, and the ozone layer slowly started to recover. It was only later that scientists discovered that the replacement substances for the ozone-damaging chlorofluorocarbons were also an extremely competent greenhouse gas and therefore contributing to global warming. The planned replacement for these substances was amended to the Montreal Protocol at a later stage, but the gradual adjustment to new scientific insights contributed to the efforts to be called "one of the world's most successful environmental treaties of all time" (United Nations Environment Programme, n.d.).

Environmental actions to tackle climate change can originate from all kinds of different groups and stakeholders, from governmental and non-governmental organizations over private and public corporations to civic initiatives. If they gain momentum like in the example of the Montreal Protocol, they can have massive impacts on our environment. But for most of these changes to happen, the public has to agree to these initiatives to become a success. According to Weber & Stern (2011), as well as the scientists I spoke to, an understanding of our climate and climate change is a crucial first step to involve the public, but an active, civic engagement is the thing we truly need in order to tackle climate change – something I will take a closer look at in the following section.




Figure 40: The progress of the Montreal Protocol in comparison to predicted change can be seen as an inspirational example of international climate collaboration. Image by Balkan Green Energy News (2018)

03.06 PUBLIC UNDERSTANDING OF CLIMATE CHANGE

Nowadays, the topic of climate change is present in many contexts and through various formats, like geography classes in school, news reports, movies, documentaries, books, or podcasts. Advertisement campaigns, climate demonstrations, or art installations have the potential to further strengthen awareness and knowledge about the issue. All the knowledge about climate change is already available, however, our understanding of climate change and its severity is often tinted by our personal experiences (Weber & Stern, 2011). If we don't see the impacts immediately or experience them ourselves, it is easy enough to discard any thoughts in that direction or underestimate the severity.

Common misconceptions about important distinctions can hinder the process of understanding our climate system further (Weber & Stern, 2011), like the confusion of weather versus climate, or pollution versus global warming. When someone would ask me, I would intuitively say that fewer particles in the air (less pollution, so to say) is always better. While this might be true concerning our general respiratory health, it might not be better for our global climate. In one of my conversations with a climate researcher of the Delft University of Technology, I got to know more about a common misconception surrounding clouds and air pollution. To explain the complexity of the topic, the researcher guided me through the process of cloud formation and its importance on the global climate.

A cloud persists of countless little water droplets, with each of those microscopic water droplets forming around an even smaller dust particle floating in the air. More particles generally mean more water droplets, and more clouds as a result. As clouds reflect sunlight back into space, the more particles, the more clouds, the cooler the earth will become. However, not every particle and not every cloud is equal. While clouds closer to the ground persist out of water and reflect solar radiation into space, clouds higher in the atmosphere start to freeze and get thinner. While those thin clouds don't reflect much sunlight anymore, they are still able to reflect heat radiation coming from the earth and therefore contributing to global warming.

On a global scale, rising temperature around the equator causes air to rise upwards, pushing clouds towards the north and south poles. As fewer clouds result in less reflection of sunlight around the equator, this, in turn, amplifies the global warming effect again – a positive feedback loop.

The complexity of the climate space makes it difficult for non-experts to form accurate mental models of the workings of our climate systems. And yet, literature also reports that a pure lack of understanding is actually not the main issue. For a long time, it was understood that more information and knowledge transported to the public will gain their approval when it comes to taking appropriate actions. More recent literature criticizes this “information deficit model” as inaccurate (McDivitt, 2016), and points towards not lack of understanding, but rather a lack of engagement from the public as the main issue (Weber & Stern, 2011).

The researchers I talked to ascribed this also partially to a language gap between science and the average citizen. Communicating in pure facts and numbers, just hearing that the world will heat by 1.4°C might not mean much to everyone. And yet, to quote the cloud researcher from earlier: “Everybody has an underlying concern, something emotional - and this is how you need to communicate with and target them, not just with facts and figures.”

03.07 EMOTIONAL ELEMENT OF CLIMATE CHANGE

The scientific approach to tackle climate change might make it seem like a purely logical, factual issue. However, people seldom make purely logical decisions but base them on emotions and feelings instead (Schwarz, 2000). As climate change is ultimately caused by emotional people, it makes the issue deeply emotional as well.

The public emotional response to global warming overwhelmingly consists of negative emotions, like worry, anger, sadness, or fear, which might lead to a lack of interest, apathy, and public disengagement (Smith & Leiserowitz, 2014). However, studies have also found that the use of “positive emotions appear to play an important role in public support for climate policies” (Smith & Leiserowitz, 2014).

The authors continue by arguing that an active interest, caused by positive emotions, might provoke learning, further information seeking, and awareness of the issue.

Designing for public engagement through positive emotions offers a variety of options to drive the experience. Desmet (2012) gives an overview of positive emotions and their effects. A selection of emotions that could be used for the further design stage include:

1. Fascination

The feeling when you encounter something new and interesting that you do not immediately understand. You feel an urge to explore or investigate to find out more about it.

2. Inspiration

The feeling when you suddenly have a new idea or insight or see the world in a different light. You have an urge to express or actualize this new insight.

3. Pride

The feeling when you possess (or have accomplished) something that exceeds your own expectations, or that others find praiseworthy. You enjoy a sense of self-worth, feel vigorous, and have the desire to show your accomplishments to others.

4. Sensory pleasure

The feeling when something happens that pleases your senses. You feel mesmerized and are motivated to savor the experience.

5. Awe


The feeling when you encounter something that is greater or more powerful than yourself. You feel overwhelmed and need a moment to adjust.

03.08 CHAPTER CONCLUSION

Through expert interviews and explorative literature research into the foundations of climate science, a broader understanding of the climate space, the world climate system, climate change, involved stakeholders, climate action, and the public understanding of climate change was gathered.

The research showed that the main issue is not a lack of possible climate actions, nor a lack of understanding of the nature of climate change, but rather a missing public engagement to motivate meaningful change. Using positive emotions might be a strong driver to encourage that engagement and build support for climate actions.

My exploration and conversation with scientists also showed that the world's climate system is beautifully complex and fascinating. Designing an experience that builds on positive emotions like fascination and awe, that showcases earth's beauty, without trying an immediate agenda, might be a valuable approach to foster public climate engagement.



04 IDEATION & PROTOTYPING

Based on the two theme explorations of climate and AI, I was able to build a basis for the following design phase of ideation and prototyping. After scoping the solution space in a feasible direction and formulating my initial design goal and interaction vision, two cycles of concept ideation, prototyping, and testing were conducted to gather insights and converge towards the final design.

Methods used: Rapid prototyping, empirical research through design, speed dating

The days of designers dreaming on behalf of everyone have passed but designers can still facilitate a dreaming process that unlocks people's imaginations.

– Anthony Dunne & Fiona Raby (*Speculative Everything*)

04.01 IDEATION BASIS

In the initial project brief of this project, I set out to explore the use of AI and play to build an experience to foster the understanding of climate change and make it engaging and immersive for stakeholders beyond the scientific community.

If it is the development of tools for professional climate stakeholders or tools for individual behavior change – the solution space of climate change is nearly endless. My exploration of the climate space supports the initial choice of the public as a target group, as their collective actions and votes have the potential for substantial influence on climate developments. However, my context research also indicates that it is not a lack of understanding or knowledge like initially assumed, but rather a lack of engagement with the topic of climate change. Play and playful experiences, like those found in games, have the strong potential to engage people and seem suitable to explore further. Furthermore, there is also a need for designs and experiences that allows people to engage without pushing an immediate agenda, but rather motivate engagement for engagement's sake.

The current climate debate is already a deeply emotional one, with most current communication being driven either by neutral facts or by negative emotions like fear and hopelessness. Yet, the literature indicates a value in the fostering of positive emotions as a means to increase interest and motivate people to engage with climate initiatives as well. In my conversations with climate scientists, I could see a strong fascination that they felt for their work and our climate system, one that makes them appreciate its beauty and its complexity and motivates them to continue the work, despite the sometimes grim prospects of climate change. As a first design direction, it seems interesting to aim at transporting the feeling of fascination for our climate to non-scientific stakeholders as well.

To achieve this goal, using GANs is a suitable approach. My initial exploration of AI and GANs has shown that these systems can synthesize vast amounts of data and generate new, often intriguing-looking data in return. Those vast amounts of data can also be found in climate research, which suggests that it is a suitable tool to explore and use further in the design process. It also allows me to continue gathering insights for my knowledge goal of how to use AI as a hands-on design tool.

Definition of play

Play can come in many different shapes and forms. Paidia and Ludus are often used to describe the opposite ends of the spectrum, with Ludus describing a formal, rule-driven game, and Paidia being used to describe a rather unconstrained version of play (Lucero & Arrasvuori, 2010). For this project, I will use play and playfulness as a state of mind, following the definition given by Fullerton et al. (2004), and explore concepts rather defined as games as well as experiences of free play and playfulness.

Appendix 2 features a collection of climate education games and games with a climate component that further informs the ideation phase.

04.02 DESIGN GOAL & INTERACTION VISION

Based on the prior research and basis for ideation, I can formulate my first design goal and interaction vision (Pasman et al, 2011) as guidance for my following design exploration.

My design goal is to make users feel an engaging fascination for our world climate system.

The envisioned effects on the user are to enjoy the experience, create a feeling of fascination for our climate and allow them to dream of a positive future again



Interacting with my design should feel like playing with soap bubbles

Characteristics: aesthetically pleasing, simple, playful and short-lived, active, passive, careful and engaged



04.03 TARGET GROUP



The initial target group of my design are **non-scientific visitors of a museum experience.**

Further characteristics of the target group are:

- English speaking, living in the western world, moderately climate-informed
- Personality: Curious, empathic, open
- Values: Accountability, compassion, exploration, learning, fairness, sustainability, self-responsibility

The reason to choose this target group is their position in the middle of the spectrum between climate professionals or activists, and active climate deniers on the other end of the spectrum. This target group makes up a big percentage of our population, which also makes it one that has the potential for meaningful influence on systemic developments through their choices, voices, and votes.

Important to note is that this does not mean that the final design is solely meant for this group, but rather provides the first focus to guide the design exploration process.

04.04 HYPOTHESIS AND DESIGN METHODOLOGY

As stated in the introduction already, I will utilize empirical research through design methodology as the main driver of my design exploration. Through a multitude of designs, rapid prototypes, and testing cycles I will approach the final design iteration.

The three main hypotheses for the design exploration are:

1. An experience to provoke the positive emotion of fascination is an appropriate approach to engage people with the topic of climate change.
2. Generative Adversarial Networks can produce media to provoke an emotional response.
3. A playful and interactive component strengthens the engagement with the installation.

04.05 DESIGN SCOPE AND FURTHER RESTRICTIONS

While the exploration process and possible directions can be potentially endless, I have to be aware of the time and resource constraints of a graduation project. Restrictions caused by the COVID-19 pandemic might limit the exploration process further and forbid access to a concrete context to host the design, while my set choice of using GANs as a main technology behind the design might lead to unforeseen complications that prevent further design iterations.

However, as stated in the design brief already, the goal of this project is not to build a finished and financially viable product, but to explore the domain of climate research and prototype a different form of communication by translating it into a playable, visual experience.

04.06 DESIGN CYCLE 1: DIVERGENCE

Based on my ideation basis, insights from literature, conversations with climate scientists, analysis of prior art (Appendix 2), and further input from Google Research, I was able to create my first design directions. Aided through informal brainstorming sessions with other students from Industrial Design Engineering I conceptualized a variety of game experiences that fit the broad direction of the Google project and roughly aligned with my initial design goal. My approach was to go as wide as possible, explore many different directions and use initial feedback from participants to narrow down the scope to a more targeted direction. I supported my ideation by asking a variety of "what if" questions to break up the space and come up with interesting concepts (Figure 41) that use these questions as a basis.

The ideas were tested through the speed dating method (Zimmerman & Forlizzi, 2017), where first sketches and visualizations together with a description of the concept are run by targeted stakeholders to collect quick and qualitative feedback about their perceived desirability.

A detailed description of the various concept premises, mechanics, as well as gathered feedback can be found in Appendix 3.

Main questions to answer with the prototypes:

1. What concept direction is perceived as desirable by the participants?
2. What concept direction seems most feasible?
3. What concept aligns with the set-out design goal, interaction vision, and knowledge goal?




Figure 41: A selection of concept drawings that were presented to the participants

Main takeaway

While most concepts were perceived with a positive response, participants also voiced their skepticism towards more elaborate game ideas, especially with the knowledge that developing a good game is very challenging. This view aligns with a priorly conducted climate game analysis (Appendix 2). More artful, yet playful concepts on the other hand were seen as more intriguing and interesting.

One of the artful concepts called "Landshapes" was based closely on my exploration of GANs in chapter 02.08, and included a physical art piece with generative landscape visuals that could be influenced and shaped by a web interface (Figure 42). This concept provoked the most responses, some of them pointing out open questions, but also sparked the most observed fascination in the participants. The framing of an art piece and the basis on GAN technology further allow me to explore the topic of AI while providing a low entry barrier, yet high ceiling in contrast to a full-fledged game.

The basis of "Landshapes" was transported into the following cycle to develop and test further.

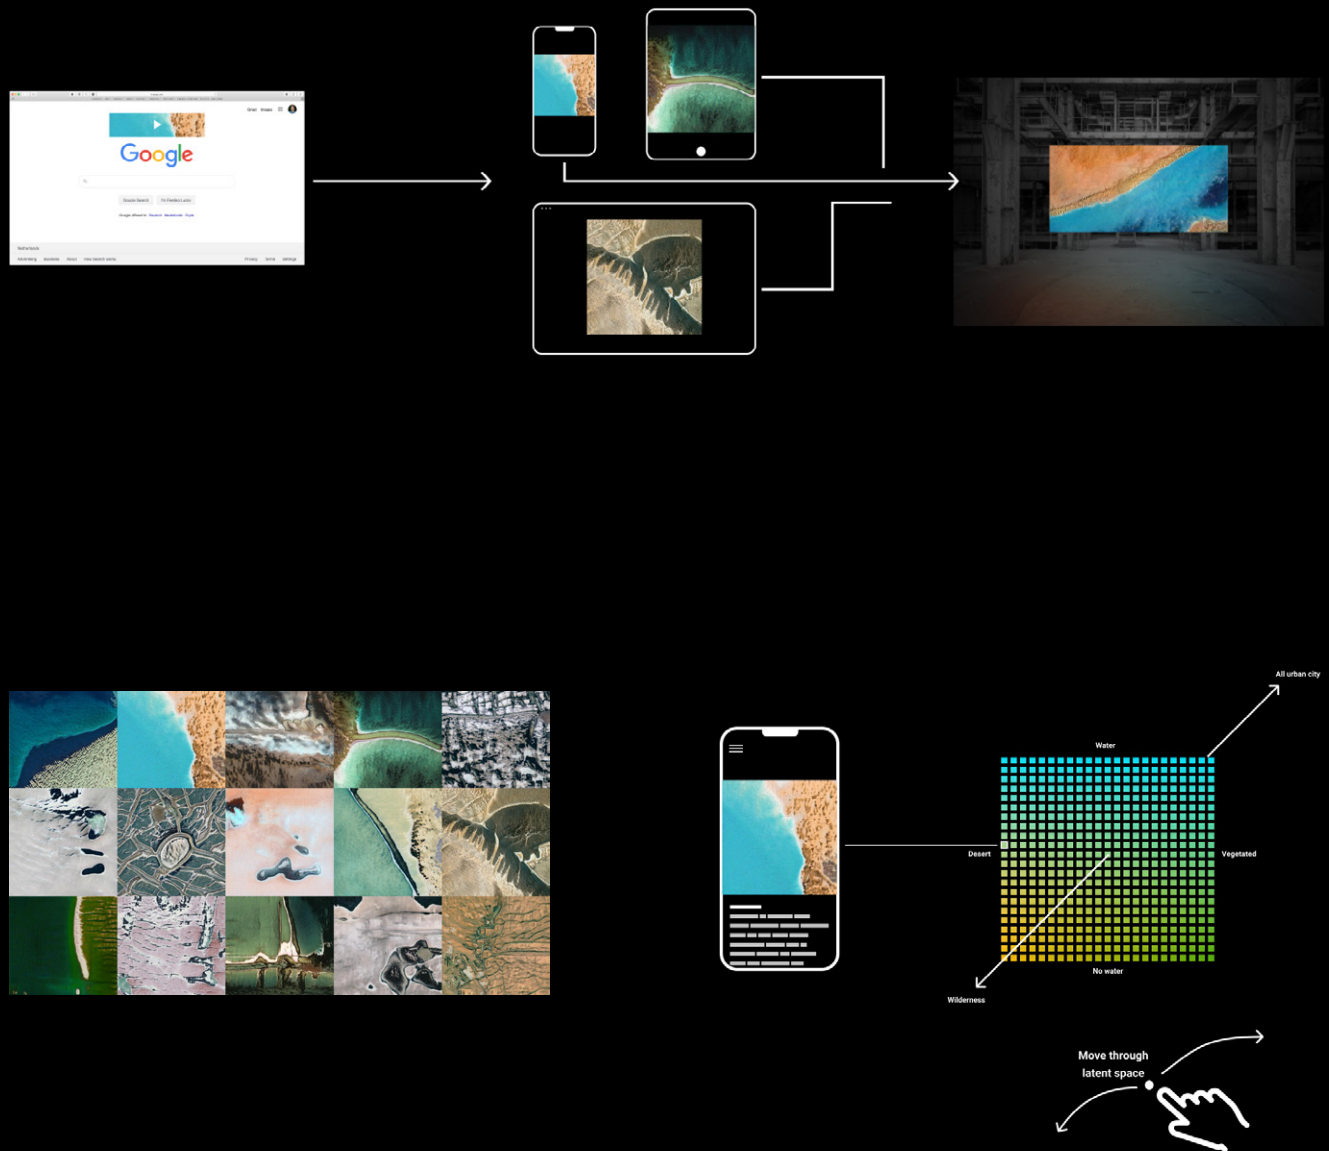


Figure 42: *Landshapes* uses an interactive web component in interplay with a physical art piece to allow participants to shape their own landscape and merge it with the landscapes of other.

04.07 DESIGN CYCLE 2: INTERACTION

Based on the concept direction of Cycle 1, a set of five prototypes was developed that allowed users to choose or influence different landscapes. The goal of this cycle was to gather feedback on an interactive version of the concept and test different interaction methods: Swiping, one-click, choosing, 3D navigation, and drawing (Figure 43).

Each prototype varied in complexity, used visuals, the scope of possible actions, and incorporated different main interactions. A surrounding narrative for the prototypes was purposefully left out to avoid distortion of the feedback on the interaction itself.

The visuals for the prototypes were generated using Runway and the pre-trained Google Earth StyleGAN model discussed previously (see chapter 02.08).

Although the concept incorporates a physical end piece, due to COVID-19 restrictions all prototypes had to be digital in order to test with participants remotely. Yet, the prototypes still allowed me to gather general insights that could inform the later physical design.

The procedure, analysis, and individual responses can be found in Appendix 3.

Main questions to answer with the prototypes:

1. Which aesthetic finds the most appeal?
2. How does each interaction feel?
3. Which interaction fits the set interaction vision?
4. Does the prototype spark fascination?

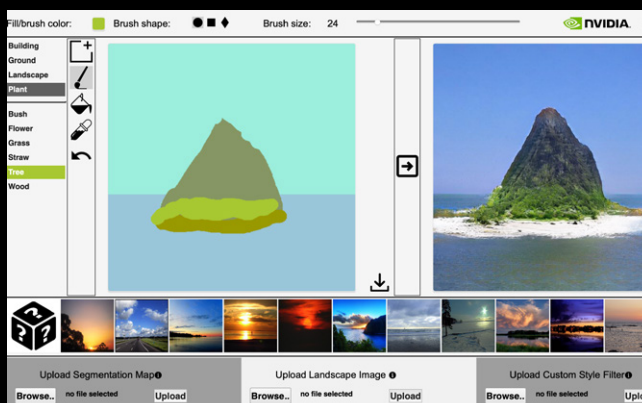





Figure 43: The prototypes use the output of GANs and different interaction methods to let participants shape, explore or select a generated landscape

Main takeaway

Testing the interactive prototypes (Figure 44) showed again that the GAN-generated imagery, especially transitions between images through latent walks, spark fascination and engage the participant. They can be seen as the main driver of the experience, especially when connected to direct interaction. Just clicking a button seems too simple of an interaction, while drawing and exploring a 3D space seem too complex. In contrast, swiping over the screen to influence the landscape seems to fit the visual transition and balance simplicity and interaction freedom.

The testing also revealed the potential to enhance the experience by underlying it with sound and providing high-resolution, less abstract images that can be directly connected to landscapes. Participants also voiced that the final design should have clear indications of a start and end of the experience, with free choice of how much time to spend on the interaction. Furthermore, the scope of possible interactions should be understood immediately, while a clear narrative (which was purposefully avoided in this testing) could provide a beneficial frame for the design.

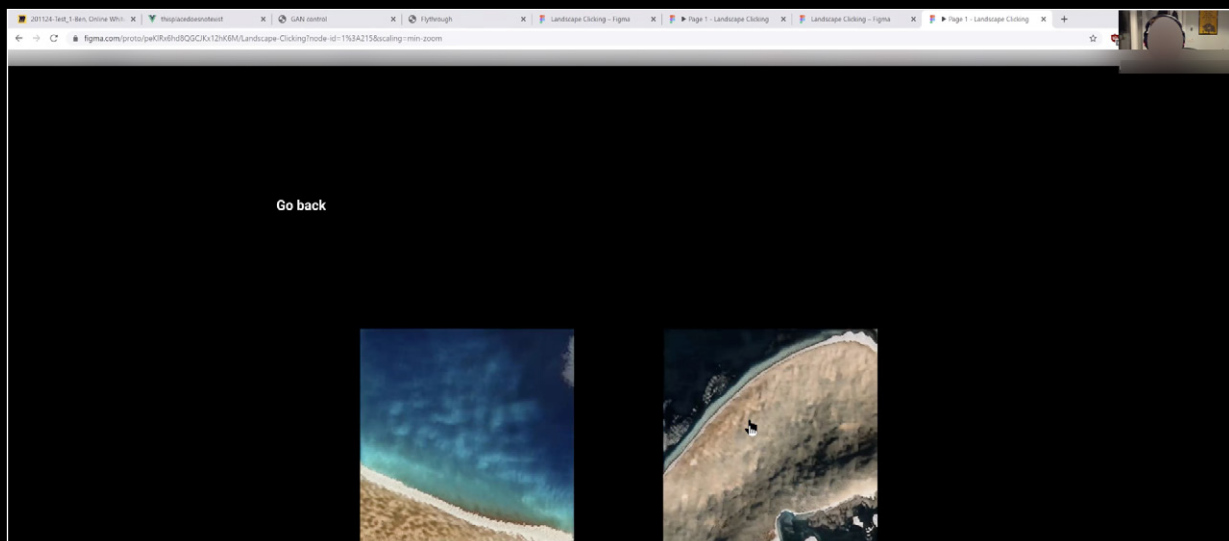

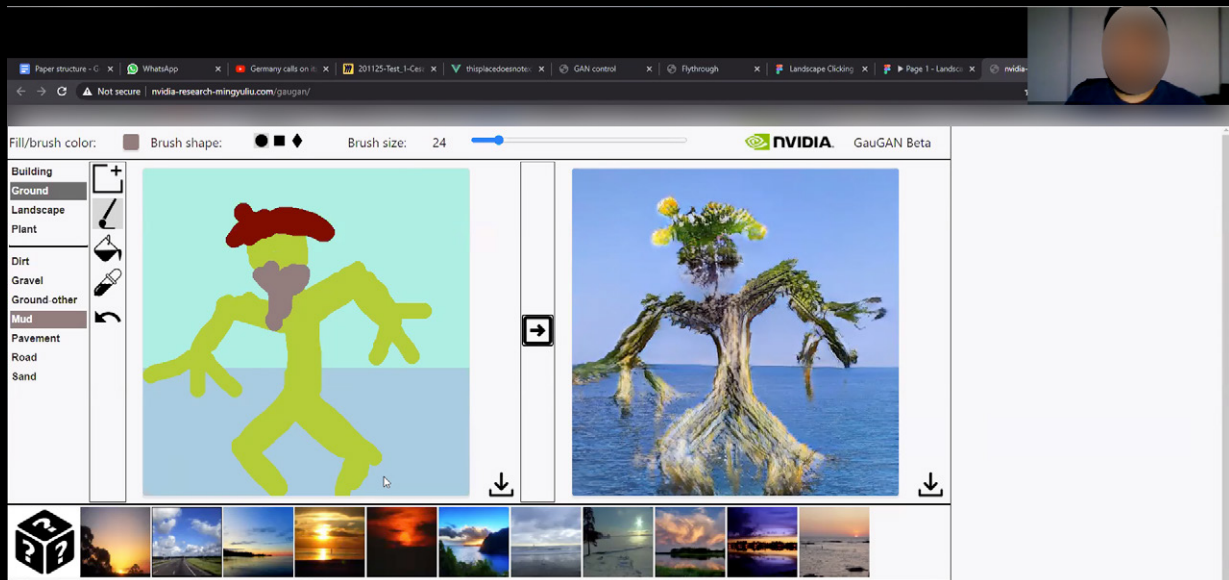


Figure 44: Participants testing the prototypes online, while being interviewed and observed using Zoom.

04.08 ADDITIONAL DESIGN EXPLORATION

The validated concept directions presented in Cycle 1 and Cycle 2 showcase only a part of the conducted design exploration. For instance, one exploration experimented with the manifestation of the GAN visuals, not in 2D, but 3D form. Other approaches included the concept of a website that (next to the morphing landscapes) offered a collection of real-life climate stories I came across during my research, like that of the gravitational force of glaciers, or about the flying river of the Amazon.

While first informal conversations about the directions make me feel confident about their value, their formal validation with participants lay outside the scope of the project and excluded from the final design. Nevertheless, these initial explorations present an opportunity for future work and a next version of the final design concept.

A collection of stories and the web design can be found in Appendix 3.


04.09 IDEATION CONCLUSION

The initial exploration of different design concepts with their corresponding prototype and testing provided a design direction that focuses on a compelling experience rather than a full-fledged game, which might have also gone beyond the scope of this project. Furthermore, the testings generated three main insights that I deem important for the final design:

First, I see a clear value in an open, abstract art experience – inspired by the design goal of creating fascination for our climate system – that allows people to engage freely, develop a sense of wonder, and get lost in their thoughts and feelings. A compelling experience like this could function as a hook to motivate further engagement with and will to explore the topic of climate.

Second, the visuals and transitions generated by GAN are a crucial component of the experience and already successfully sparked a sense of fascination on a small screen. The experience would benefit from high-resolution visuals as well as a physical component to strengthen the art setting and provide a secluded frame for the visitor to properly engage with the design. A physical representation of the swiping interaction seems appropriate to facilitate a potential interaction.

And third, a surrounding, clear narrative surrounding the experience might give more explanation of the design's purpose and functionality but is no unconditional necessity. More value might be found in letting visitors freely discover the possibilities through open play, choose their level and duration of involvement, and interpret the experience themselves to develop a narrative. The involvement of a surrounding narrative to enrich the experience might be explored in future design iterations, but goes beyond the scope of the final design and its validation.



05 FINAL DESIGN: LANDSHAPES

The months of research, prototyping, and testing results in the final concept design of LANDSHAPES. The concept development was guided and shaped by the underlying GAN technology and is based on what I know to be feasible, desirable, and valuable.

In this chapter, the design is described in detail. Furthermore, a brief overview of the building process of the underlying GAN model is given and the validation process of the design and its results described. The chapter concluded with a reflection on the limitations and recommendations for the final design.

| *Methods used: Interview, observation, prototyping*

05.01 CONCEPT DESCRIPTION

The final design outcome of my graduation project is LANDSHAPES, an interactive installation intended for the museum context. The center of the installation is the circular projection of AI-generated aerial landscapes – images that look real, yet have never been seen by anyone before.

The framing handrailing offers people a moment to stay still and become absorbed in the slowly changing landscapes. However, we are no passive beings when it comes to our planet and its climate. A pushable handle, following the shape of the railing, allows the visitor to become an active participant in the installation, to move along with the changing artwork, and experience their potential impact on our surroundings in a playful, open-ended manner.

The installation evokes a new perspective on the world, as it shows it not as static, but as an ever-changing entity, open to be shaped by the vision of any daring hand. The AI – trained on images from all around our world – functions as a neutral mediator of the experience. With the potential to depict any possible landscape, it is not the designer who imposes a vision of a potential future landscape, but the creation being the result of a collaboration between the machine and the guiding visitor.




Figure 45: An exhibition visitor observes the morphing images of the bright landscapes projected on the floor of the dark room

Context and setup

The installation in its current design is meant for the museum context. Framing it as an art piece and placing it in a gallery presents it not as a solution-driven product, but a place to find wonder and get inspired.

The darkness of the room stands in contrast to the bright, projected images on the floor, allowing the landscapes to take up the prominent focus. A circular mirror floating above the installations functions as a complement to the earth-bound projection and reflects the projection towards the room entrance, immediately evoking curiosity for the entering visitor.





Figure 46: *Layout of a fictional museum space with the installation as the centerpoint of the exhibition*




Figure 47: A mirror hanging over the installation reflects the earth-bound visuals towards the entrance to evoke immediate curiosity from a distance

User experience storyboard



The visitor decides to visit a museum. While strolling through the exhibitions, he comes across a dark and mysterious corridor. At the end of the corridor, he can see the warm glow of lights reflecting from the walls.




As he moves closer, he notices abstract music and sounds playing in the background. Turning around a corner, he sees the center piece of the room. Slowly moving and changing visuals are projected on the floor and throw a shimmering reflection on the walls and surrounding floor.



Leaning against the handrail, the visitor watches for a while, just seeing different landscapes morphing from one stage to the next. He thought he might have seen an ice landscape, but it was already gone the next moment.

After some time, he looks around and notices what looks like a wooden knob on the other side of surrounding handrail.



Intrigued, he walks towards the kob and gave it a slight nudge. The handle moves effortlessly, and immediately, the music and the moving landscapes come to a stop. At the same time, a surrounding light inside the handrail lights up ever so slightly.



The visitor grabs the handle, and moves it along the railing. As he moves the handle, the ambient music fades in again and the projected landscapes start transforming again, this time in synchrony with the now walking visitor. He quickly realizes that abrupt movements change the landscapes equally fast, while slow and delicate movements allow him to influence the landscapes in fine steps, giving him time to take in every little detail.



The visitor continues to slowly walk around the installation, always exploring novel perspectives on the shown imagery and guessing where those beautiful landscapes might be. After a while, he stops and lets go of the handle. The movements and music stops.

As he walks towards the exit, he can hear the music fading in, he sees the shimmering reflections move again – the installation getting ready for its next visitor.

Physical design

Nowadays, many of our experiences are accessed through our digital devices. While this cascade of potential perspectives have the benefit of being accessible from all around the world, they also tend to be rather short lived. The next experience, the next story is often only a click away.

The physical design is meant to provide an immersive experience, motivating visitors to take the time to interact and engage with the installation. The physicality of the installation grounds the fluent landscapes in the real world and gives them additional value, as well as allowing for physical movements as an influencing interaction.

The handrail design surrounding the landscape projection is kept intricate and light, allowing visitors to see the imagery at a first glance and inviting them to step closer. The movable wooden handle, being made out of a warm and smooth material, together with its rounded shape afford to be touched and played with.

A lightstrip integrated into the handrail provides a warm glow as soon as a visitor engages with the installation by moving the handle. This subtle cue is supposed to communicate the transition from passive to active participant and reassure that the interaction is welcomed by the installation.





Figure 48: *The handle design is kept smooth and warm to invite the visitor to touch and interact with it*






Figure 50: Wireframe presenting the installation from an aerial perspective




Figure 49: Wireframe presenting the installation from a sideview.

System diagram

The installation consists of a digital projector, oriented towards the ground, and a connected computer with a powerful graphics card to provide the projected imagery. Sensors embedded into the handrail capture handle movements, directed by the visitor, and pass them onto the underlying digital system (Figure 51).

The core of the generative system is a generative StyleGAN model. The model constantly generates novel landscapes. While generating new imagery, it simultaneously forgets prior landscapes, making each visit a truly novel experience. When left untouched, the landscapes autonomously, slowly morph from one state to the other. Moving the handle gives the visitor the opportunity to slow down, stop or even reverse time (Figure 52) to look at new and old landscapes and their transitions in detail, or by moving quickly cause substantial and immediate change (Figure 53).

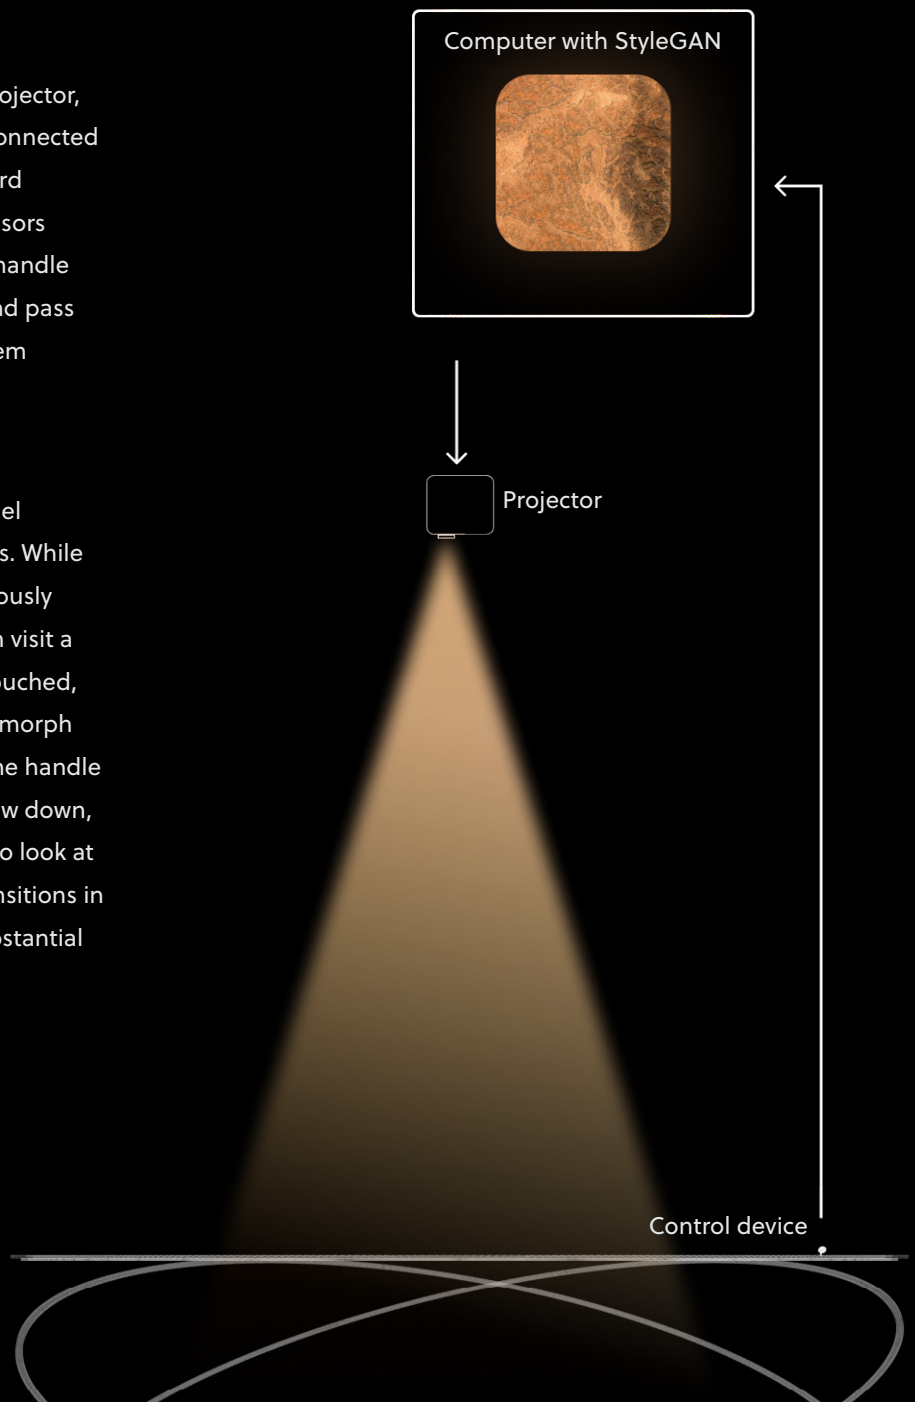


Figure 51: The control device can be used to further influence the StyleGAN model, which generates visuals for the projection.




Figure 52: *Moving the handle slowly results in gradual and gentle transitions between landscapes*




Figure 53: *Moving the handle fast and abrupt results in more substantial change*

Used imagery

The visuals displayed in the installation are based on a curated dataset of high-resolution satellite imagery from all around the world (Figure 54). Earth imagery taken from space possesses an unmistakable beauty and shows a view most people will never experience in their lives. The intricate shapes and colors invite us to explore the details of a landscape and appreciate its appeal. At the same time, those images – taken from a far distance – only suggest details of the landscape and provokes the viewer to fill in those missing pieces with their imagination. This perspective allows oneself to be detached from reality, if only for a brief moment, and quite literally puts the emphasis on the bigger picture.

A single image can only capture and represent the snapshot of a brief moment in time. Especially on the grand scale of our planet and its history, it would be false to assume that the shown views and landscapes will always stay the same. Presenting an interpolation between landscapes generated by the underlying GAN (Figure 55) shows our environment not as static, but as an ever-evolving entity. The morphing images possess an intriguing and mesmerizing quality that evokes fascination and desire for further exploration in the viewer. Yet, the presented changes do not possess any judgment. They can be interpreted as both positive and negative, and the narrative is shaped by the visitors themselves.

Lastly, the used images to train the GAN are open-source data of the public domain, created by the SENTINEL-2 program of the European Space Agency (2015). Having access to big amounts of data is a crucial requirement to allow a GAN to generate realistic and appealing new images. Sourcing the training data from locations all around the world ensures that every location, every nation has the same chance of being represented. Besides the practical benefit, the choice of data also contributes to the installation's narrative. When the underlying images belong to all of us, every visitor seeing the installation already has a part of themselves in the installation and contributed to its existence.




Figure 54: *Real images from the SENTINEL-2 dataset*




Figure 55: *Generated images from the StyleGAN model*

05.02 DEVELOPING A PROOF OF CONCEPT

The final concept of LANDSHAPES is one possible result of the design explorations conducted in prior iterations, documented in chapter 3 of this report. Building a fully functional physical experience, especially when limited not only by time, but also persistent COVID-19 restrictions met in the period of this project, would go beyond the scope of this graduation. Nevertheless, this restriction should not prevent further validation of the concept (Figure 56).

As the main drivers of the envisioned experience are the projected visuals of generated landscapes and the ability to interact with them playfully, this is where the validation focus was set. In the following sections, I will summarize my process of building a dataset and training a StyleGAN model on the collected images, as well as the validation process and results of a simplified LANDSHAPES experience.




Figure 56: *Photograph showing a person with the wall-projected images at the final validation*

05.03 BUILDING A DATASET

In my initial exploration of GANs, described in chapter 02.09 of the report, I was able to gather first experiences with the training process of a custom StyleGAN model. This experience suggested that obtaining a fitting dataset would be the most crucial and also potentially trickiest part of the development process.

My first exploration and training attempts fortunately also clarified which kind of data I was looking for. Ideally, the dataset would consist of:

1. satellite images in RGB JPEG format,
2. 1024x1024px in resolution,
3. diverse scenery,
4. with ideally over 2000 images

The publicly funded organizations NASA and ESA provide daily updated remote sensing data, including satellite imagery with channels of the RGB spectrum suitable for my needs. When accessing the data through web portals hosted by NASA and ESA directly, the images can be downloaded in their original GeoTIFF format. While the GeoTIFF format provides the full resolution and all available channels of the satellite capture, they often reach sizes of 200Mb per image. For a dataset of 2000 images, the individual files would add up to 400Gb of data that would require further processing to achieve the targeted data format and resolution.

Another access point to the data is Google Earth Engine (2010), a platform where a diverse set of geoscience-related datasets are hosted and openly accessible for further analysis. Most importantly in my case, the platform also includes the ability to define an area and export it as an image preview in JPEG or PNG format.

To collect my dataset of diverse satellite images, I build an export pipeline (shown right) that exports atmospherically corrected images from the SENTINEL-2 program (Sentinel-2 MSI: MultiSpectral Instrument, Level-2A, 2015). The pipeline utilizes the open-source Geographic Information System QGIS (2009), Google Earth Engine, and Google Colaboratory (2014) to run the export process. All code used to build the dataset can be found in Appendix 4.



1. Generate locations

Using QGIS and a shapelayer of earth's coastline, thousands of coordinates were generated and exported as a table.


QGIS

Longitude	Latitude
48.3325104751309	80.7655050089461
25.227826874662	35.405652432138
80.103522776003	9.66528817970192
136.202091011706	36.418449544462
8.74773892124874	41.6649534702982
...	...

2. Search SENTINEL-2

A script running on Colab iterates over the table, requests the satellite image at that location from the SENTINEL-2 dataset and calculates a square export region laying within the image.

CO + Google Earth Engine



3. Generate the export link

Using the calculated export region, a request to Google Earth Engine returns a link with the desired region as a JPEG.


Google Earth Engine

<https://earthengine.googleapis.com/v1alpha/projects/earthengine-legacy/thumbnails/df5066652ffd763ebde48e6d02f30ee1-f73e794ab0168dd0b0f18058b9b0eeb5:getPixels>

4. Save the image and repeat

Another script running on Colab requests the image behind the link and saves it in Google Drive before repeating the whole process for the next location.

CO + Google Earth Engine + 



The exported images were prefiltered by finding and deleting duplicates and the remaining images were imported to Adobe Lightroom (2020) for further processing. The manual process of color and exposure correcting the images, as well as filtering out undesired samples with visible artifacts or distortions, ensured that the data was suitable for the following training process.

After filtering and further processing, the final dataset contained 4.040 high-quality RGB images in 1024x1024px resolution, separated in ~2000 images showing a coastline and the same amount depicting pure landmass images.



Figure 57: *The raw data (left) required further post-processing, like in this case increasing the exposure to come to a usable result (right)*






Figure 58: *The dataset consists of very diverse images, ranging from mountains, over desert and arctic regions to inhabited coastal regions*

05.04 TRAINING A STYLEGAN NETWORK

With the dataset in place, the training process of my custom StyleGAN could begin. Similar to my first GAN training, described in Log 7 of chapter 1, an existing model was used to provide a starting point. Through exposing the existing model to new data, the process of transfer learning allows the GAN to identify and generate patterns it recognizes in the novel data while cutting down on training time and increasing network stability. For this training, I used the latest StyleGAN iteration: StyleGAN2-ADA, which promises faster training times and further network improvements in comparison to older versions (Karras et al., 2020).

The time and graphics-intensive process of training the GAN was executed on Google Colab, a server platform that provides infrastructure and allows code execution for researchers, and an adapted version of the original StyleGAN code repository (StyleGAN2-Ada Custom Training, 2020).

The collected dataset could simply be uploaded to Google Drive and converted into the required tfrecords format, which makes it usable for the underlying TensorFlow infrastructure. I decided to split the training into multiple segments, with a division into four datasets that showed either arctic landscapes, coastlines, landmass imagery, or a combination of landmass and coastline data (Figure 59).

The training ran in intervals over multiple days, for a total of 747 epochs. With each epoch taking ~ 8 minutes to complete, the total training time sums up to approximately ~100 hours. Every eight epochs, the code saves a copy of the current network and exports an image sample that allows for a visual assessment of the training process. Although additional quantitative metrics like the Fréchet Inception Distance (Dowson & Landau, 1982) exist to help judge the training process, they were not used in this training process. While they would add additional insights into the sometimes oblique GAN training, they also extend the overall training time substantially. As the main objective of the GAN is not to reproduce outright realistic output but to create aesthetically pleasing imagery, I decided to rely on qualitative assessment alone.

The results of the training process are four StyleGAN2 models, able to generate either satellite images resembling Antarctica, coastlines, landmass, or a combination of them all.

Utilizing another Google Colab notebook (StyleGAN2-ADA Generate, 2020), I was able to use the provided server infrastructure to test the trained models and successfully generate new landscape images to use for the final concept validation.



Figure 59: *The four checkpoints generate landscapes with different characteristics.*

From left to right: Images of antarctica, images of coastlines, images of landmass, diverse images of all kind

05.05 CONCEPT VALIDATION

The custom training of the StyleGAN model on high-resolution satellite imagery and concomitant ability to generate new visuals showed that LANDSHAPES was feasible from a technical perspective. While building the full installation lays outside the scope of this graduation, a final user testing aimed at providing further validation for the overall concept direction.

Mostly, I was interested to shine light on the the questions of:

1. How are participants responding emotionally to the different visuals projected at full scale?
2. How does the addition of a control mechanism influence the experience in terms of engagement?
3. How would participants interpret the visuals and open narrative of the installation?

Furthermore, the validation should indicate the design compliance with the set-out design goal of making the participant “feel an engaging fascination for our world climate system.”

Test setup

For the final concept validation, I pre-exported three videos with different controllable characteristics: The variety of landscapes depicted in the video, the speed of interpolation between landscapes, and if it was a looping video or not. The videos are linked below the corresponding graphic (Figure 60).

The test was conducted with students and employees of Industrial Design Engineering, TU Delft. The installation was set up in a dark room at the faculty of Industrial Design Engineering, with a FullHD projector projecting the visuals on the wall. Abstract background music played on a speaker was meant to enrich and simulate an atmosphere closer to the envisioned experience.

The control handle envisioned for the final concept was replaced by a rotary encoder with 360° turnability. Connected to an Arduino and further via Serial to Processing (2020) as a video player, it allowed participants to influence the playback position of the looping video. Not interacting with the device for a time of 5 seconds automatically started the playback again. The code for the interactive sketch can be found in Appendix 4.



Calm

<https://youtu.be/G1RZMsRG110>



Abrupt

<https://youtu.be/8Ypel6EW-yM>



Calm and looping

https://youtu.be/_h9-d76JUBE

Figure 60: The exported videos had different characteristics, in this figure symbolized with a line drawing. From left to right: Calm movements, chaotic movements, calm and looping movements

User testing procedure

The invited participants (n = 5) tested the prototype in succession. The research approach and procedure were explained to the participant and time given to answer any potential questions. Furthermore, each participant was asked if they had any underlying condition of epilepsy or complications of flashing lights, as some of the fast-moving images might evoke a physical response. The introduction procedure concluded with the filling-in and signing of the appropriate consent form (Appendix 3).

The three videos were shown in the order Calm, Fast, and Controllable. For each video, the participants were asked to watch the visuals and think out loud, meaning to speak whatever came to their mind. The qualitative data of comments and feedback was captured manually, and the collection process was supported by occasional clarifying or probing questions.

After one loop of the controllable video, the control device was given to the participants with the request to explore it freely, without any further explanation on its working.

Analysis

The statements and feedback of the participants were transferred onto a digital Miro canvas and turned into individual post-its for further analysis. The post-its were given a color code based on the referred video and clustered into groups of overarching themes.

Results

The resulting clusters of the given statements and responses can be broadly categorized into the themes of "Aesthetics", "Narrative", "Emotions" and "Control". Each category consists of multiple sub-clusters, like "Visual interpretation", "Experience Narrative", "Emotional Response", "Negative emotions", etc. The full cluster overview of the clustering shown in figure 61 with the raw data can be found in Appendix 3. In the following, I will summarize the main insight per category.

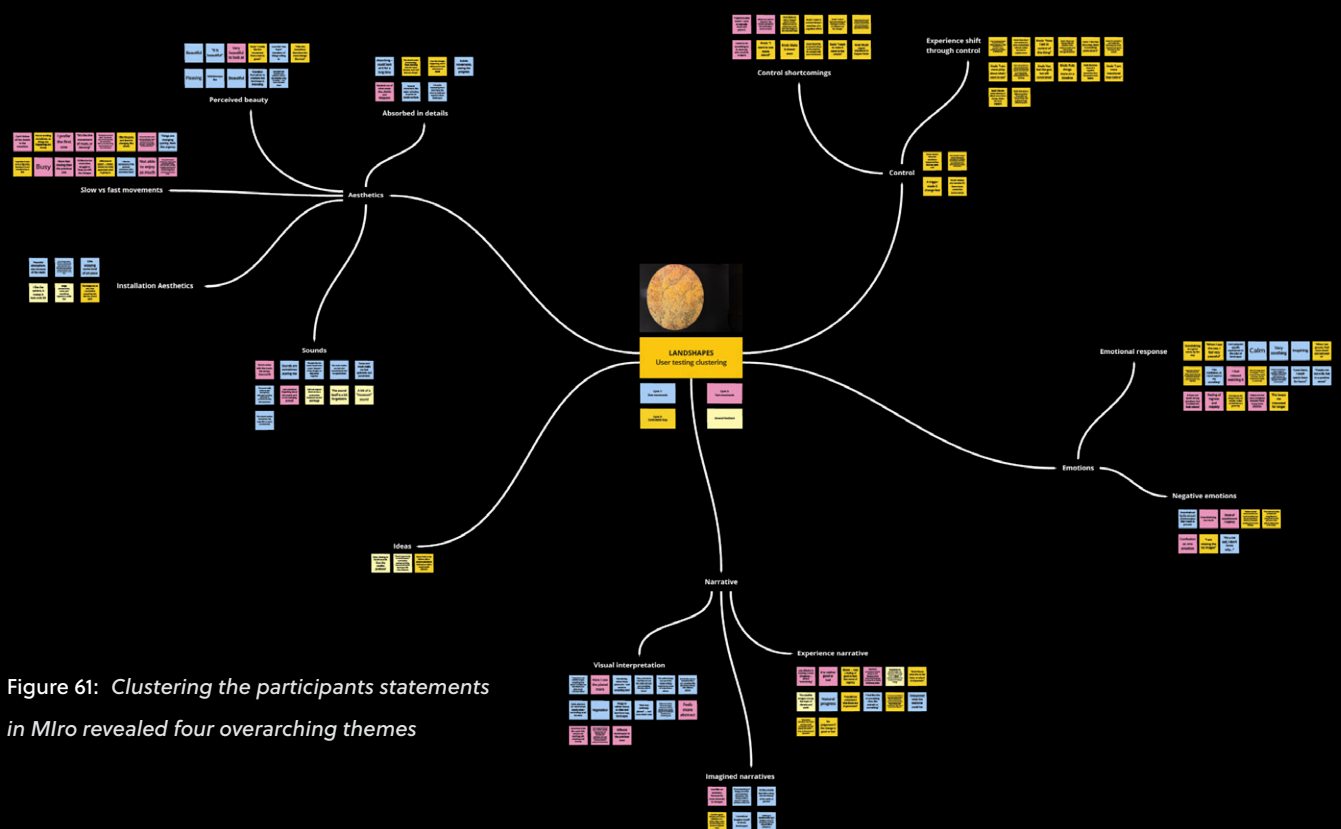


Figure 61: Clustering the participants statements in Miro revealed four overarching themes

Cluster 1: Aesthetics

All participants experienced the projected visuals as “pleasing” or describing them as “beautiful”. The beauty was accredited to the natural movements of the morphing visuals, the richness of details, the circular installation design, and the interplay with accompanying background music.

The speed difference between the slow-moving and the fast-moving video influenced the experience and perceived aesthetics. The fast-moving images were described as “more alive”, yet also as “busy” as well as “tense and a bit confusing”. The majority of participants preferred the slow-moving version, as it allowed them to “follow all the details in a transition” and “appreciate what is going on”.

Cluster 2: Emotions

The experience and contributing visuals were able to evoke a mainly positive emotional response. The slow-moving visuals were described as “calm”, “very soothing” and “inspiring”, with it all forming a meditation-like experience that they “could watch for hours”.

The fast-moving images caused a feeling of being energized in some participants, as part needed to pay more attention. However, some participants voiced that the same visuals caused a feeling of confusion, unsettling and urgency, and being overwhelmed.

Yet, the feeling of overwhelmedness was described as both positive and negative by different participants. Some say that the perceived scale of the landscapes had a “feeling of bigness and majesty” that made them “feel tiny, but in a good way”, while others reported to “feel insignificant in comparison to the pictures of the planet, taken from so far away”.

Interestingly enough, interpreting the still-standing visual alone seems to be able to spark an emotional response as well. One participant voiced to miss the ice images present in the first video, while another said that seeing the sea gives a feeling of peacefulness while seeing images of the ground evokes a feeling of closeness and narrowness.

Cluster 3: Narrative

For this testing, the participants were not given the context of the project or explicit priming of what they might see. Yet, all were able to interpret the visuals as aerial landscapes or satellite images, although some intermediate states of the morphing felt more abstract and resembled other visuals, like tiny cells on a petri dish, the surface of the moon, or in some yellow bits the surface of the sun. The abstract, yet realistic nature of some images evoked a feeling of wonder, as participants try to figure out what they see and where this place might be.

Multiple participants voiced that they imagine and make up their own narrative to explain changes in the landscape or imagine themselves in those landscapes. Some imagine "a story of the planet, and the rise and fall of civilizations", see an evolution in the images, or imagine a meta-narrative of "an experience of change, and change of the environments themselves".

The experience narrative, so not only one found in the visuals, but ascribed to the installation as a whole, differed between participants. Some describe the overall experience as without judgment if the presented change is good or bad and as natural progress, while others said that recognizing the images as satellite images was important, as it conveys the topic of climate and earth, but also "brings in the baggage of what you expect to see in reports of climate change.". On the other hand, the experience was also seen as being "more about enjoying earth and appreciating what is beautiful" without expectations of "what the outcome could be".

The element of using an AI to generate the images was not understood by all participants, or why it would be important.

Cluster 4: Control

For the last testing, participants were given a control device to influence the shown visuals by allowing them to move forward or backward through the looping video. Adding a control device caused a shift in experience for most participants. Overall, it shifts the experience from being a passive and relaxed one to a more cognitive task.

This led to some being “pickier about what I want to see” and being “more intentional than before”. While several participants voiced that they liked the added control, as it “makes [them] want to explore more”, “feel like a god”, or being “more engaged”, the shift in experience also has the potential to take away from the appreciate prior experience, as they “can’t enjoy the landscapes as much [when being] focussed on the control”. The ability to go back and forth shifts the interpretation of videos and “puts things more on a timeline” and the participant in the position of a researcher, turning the control to look at what did happen or what might still be happening.

While some participants voiced the wish to be able to control and pause the movements in the non-interactive videos, the control used in the testing had some shortcomings. Participants asked for a control ability that lets them influence the landscape in very small steps, to see the transitions happening in every detail, while at the same time also having the ability to cause swift and considerable change in the landscapes. The control device should be usable without significant cognitive effort and also fit the presented visuals in scale. One participant voiced this as “The control does not fit the images. The images are big and overwhelming, and I can’t believe that I can change something so big with the small remote”.

Testing limitations

The validation of the final concept was among other things limited by COVID-19 restrictions, available technology, resources, and time.

This restricted the number ($n = 5$) as well as the diversity of testable participants, whose feedback and responses might not be representable for other potential users of the installation. Furthermore, the design and form of interaction of the tested prototype diverged from the final design, which did not allow for testing every element in detail.

While providing a closer representation of the envisioned experience, the added background music also influenced the purely visual assessment and might have altered the participant's assessment of the presented visuals.

Lastly, for this testing, no further indication of the project's context and its link to climate change was given. This context knowledge, as accommodating the final design, might also influence the participant's perception of the experience.

05.06 Insights & recommendation for the final design

The validation of the final prototype indicated that the design has the potential to evoke an emotional response and that participants also managed to build the connection to the climate without a surrounding narrative. The installation, with the generated landscape visuals as the main feature, was described as a beautiful, engaging, and overall pleasant experience.

The comparison of three different video types suggests that the highest engagement and positive emotional response would be provoked from slow-moving and detailed videos. However, fast and drastic changes could also be used to evoke a feeling of urgency and confusion, to trigger a potentially more negative emotional response.

It was interesting to see that the lack of a clear narrative surrounding the installation motivated participants to come up with their own narrative and fill in the gaps. If the lack of surrounding narrative is more beneficial for creating a feeling of engagement in comparison to a given narrative needs further evaluation.

As anticipated, the addition of a control device changed the experience from passive to active engagement. While it added an extra benefit for some participants, it also had the potential to take away from the meditative experience of just looking at the beautiful visuals. This insight supports the design decision to offer an input opportunity to influence the installation, but not place it as a necessity for people to use it to enjoy the experience. The gathered feedback of participants asking for a control bigger in size and with less cognitive investment also falls in line with the envisioned design. The planned interaction of using walking around the installation and physically moving the control device in space can be expected to fit those requirements, yet should be tested in an additional validation.

The initial design exploration and final design were guided by the setout design goal to “make users feel an engaging fascination for our world climate system.”. The final design still builds upon the main element of our climate system, represented by the utilized satellite imagery, although in an artistic and rather nuanced fashion. The design validation showed that it is possible and valuable to use climate data to create an engaging experience that evokes fascination in the user towards the shown imagery.

Prior design explorations conducted in this project experimented with the element of short stories showcasing fascinating interactions taking place in our climate, such as the story of the gravitational force of glaciers, powerful enough to attract mountains of waters around the world’s ice sheets. While outside the scope for this version of the final design, including narrative material surrounding, or within the installation in the form of climate stories has the potential to make a more explicit link between the design experience and our climate system. As design and climate literature see narratives as a powerful tool for engagement (Corner et al., 2018), I feel confident that exploring this more and including direct linkages to the climate system, such as with narratives, can have the potential to enrich the overall experience and answer more explicitly the design goal as currently stated.



06 PROJECT CONCLUSION

While the conclusive parts of the previous chapter discussed the insights and recommendations for the final design of LANDSHAPES, this section will give a summary of the project and its outcomes. The knowledge contributions towards society, design, and the field of AI are discussed, and opportunities for future work are given. A personal reflection on my learnings during this project will complete this report.

06.01 PROJECT SUMMARY & OUTCOMES

This graduation project explored how to utilize Generative Adversarial Networks (GANs) to motivate public engagement through the emotion of fascination and element of play.


Conducted interviews with climate researchers and literature review into the topic of climate change showed that the current climate debate is deeply emotional, with mostly negative emotions like fear and perceived hopelessness in its center. Yet, it was also found that fostering positive emotions like fascination and joy is a potential means to increase interest and motivate people to engage with climate initiatives. A series of design prototypes, based on an initial hands-on exploration of the use of GANs as a design tool, was created and tested with participants to inform the final concept.

The final design of this exploration is an interactive exhibition piece that allows its users to create novel, AI-generated landscapes. The evaluation of the installation showed that it evoked a deep sense of beauty, calmness, and inspiration that contributed to a meditation-like, engaging experience. Participants automatically associated the morphing images with our changing planet and ascribed meaning to the abstract visuals. The addition of control over the moving landscapes resulted in a mindset shift, as the experience transformed from a passive and meditative experience to an experience with additional cognitive involvement.

The design research conducted in this project strengthens the assumption that experiences with positive emotion in its center can be used to foster initial engagement with the topic of climate change. Furthermore, it showed that GAN technology can be used to evoke an emotional response, which offers interesting opportunities for future work.

06.02 ANSWERING THE DESIGN GOAL

The initial design exploration and final design were guided by the setout design goal to “make users feel an engaging fascination for our world climate system.”. The final design still builds upon the main element of our climate system, represented by the utilized satellite imagery, although in an artistic and rather nuanced fashion. The design validation showed that it is possible and valuable to use climate data to create an engaging experience that evokes fascination in the user towards the shown imagery.



06.03 KNOWLEDGE CONTRIBUTIONS

This graduation project conducted a deep dive into the fields of AI and field of climate research to address the societal issue of climate change from a design perspective, while simultaneously exploring the usage of contemporary AI in the form of GANs as a tool in the design process.

The three leading hypotheses for this project were:

- 1.** An experience to provoke the positive emotion of fascination is an appropriate approach to engage people with the topic of climate change.
- 2.** Generative Adversarial Networks can produce media to provoke an emotional response.
- 3.** A playful and interactive component strengthens the engagement with the installation.

The exploration and outcome of the project show the potential for using GANs to design engaging experiences. Supporting the initial hypothesis, it furthermore demonstrates the ability to use GANs, through visual representation and interaction with the model, as a way to trigger an emotional response and inspire people about a societal issue. The playful interaction with the installation strengthened participant's engagement.

Climate change, the particular societal issue targeted in this project, is a wicket issue that requires engagement, novel ideas, and addressing from all sides. The final design of this project reframes the issue by using, contrary to many contemporary approaches, positive emotions like fascination as a first stepping stone to engage people with the topic.

The documented experience of using AI technology as a tool in the design process showed that AI (or GANs, in this example) can be a powerful tool in a designer's toolbox. However, the exploration also revealed difficulties that designers should be aware of when choosing AI as a design tool. These difficulties include the potentially intensive time investment of training an AI, exploring the space – a process hindered by often specialized lingo –, finding the right data, defining the right use cases, and making the output controllable and usable.

The knowledge and resources collected during this project can be used and built upon by other design practitioners interested in exploring AI.

06.04 FUTURE WORK

The project started with grand aspirations and the vision of making climate models playable and accessible for people. The final design was shaped by the met technical limitations and presents only one of many possible forms of use of the underlying technology for motivating climate engagement.

The COVID pandemic and met restrictions as a consequence thereof further hindered the design process and influenced the final outcome. Nevertheless, the positive design validation shows the potential for a variety of further work and exploration.

Opportunity 1: Building upon LANDSHAPES

The design of LANDSHAPES and positive validation presents the clear opportunity to build the concept a fully functional model and be exhibited in a museum context of a post-COVID world. The inclusion of compelling, real-world climate stories (Appendix 3) has the potential to enrich a surrounding exhibition and spark not only fascination but also more targeted interest towards our climate and its workings.

Furthermore, as described before, the chosen manifestation of LANDSHAPES shows one possible interaction and one possible physical design. The general concepts of playful interactions with a GAN trained on aerial satellite imagery can be explored further through different kinds of interactions and media. One imaginable form is the adaptation of the physical experience into a digital web-based experience to allow access to a broader audience. Another kind of interaction, potentially more suitable for kids, is to use drawing as an input method. First explorations in that direction were already conducted during this project, with user tests showing the high potential of drawing for prolonged engagement (Appendix 3), and GAN exploration of a Pix2Pix network showing its feasibility.

Opportunity 2: Playable climate models

The final concept of landscape builds upon a simple, unidirectional interaction that allows visitors to push for change in the shown landscape, yet not influence any specific parameters of the created world. And still, it showed the potential around engaging and aesthetically pleasing climate experiences.

With access to more data, time, and hardware resources, it is not far-fetched to envision a version of LANDSHAPES that truly enables visitors to shape and explore a landscape based on their imagination. A contemporary version of a playable earth, like once envisioned by SimEarth (1990), might be a valuable approach to let users engage with climate data and its beauty in a playful way.

Opportunity 3: Emotional GANs

The evaluation of LANDSHAPES showed the ability of GANs to produce media that evoke an emotional response in the viewer. These first research results present the interesting opportunity to measure these perceived emotions and their intensity and feed them back into the GAN as additional training data.

A hypothesis might be that it is possible to use a GAN, trained on perceived emotions, to produce imagery that evokes or counteracts a specific emotional response in the viewer. It is not suggested here that this research is, by all means, desirable, yet it does provoke interesting ethical questions about our envisioned future relationship of AI and human emotions.

06.05 PERSONAL REFLECTION

In the initial project brief of my graduation, I described my motivation to combine two of my passions, technology, and environmental efforts, as well as my personal ambitions on what to learn. I wanted to improve my prototyping and building skills by building an installation, finally, try to work with data and machine learning in a hands-on fashion and build professional relationships on the way.

Looking back on my graduation project, my process, and all the things that did not make it into this report, it feels somewhat surreal how many experiences and learnings I collected on the way. It was a challenging journey with many moments of struggling, but I can also proudly say that I still achieved most of what I set out to do. True, COVID-19 did not provide the best circumstances to build an installation, but this for instance pushed me to finally learn more about 3D modeling and rendering – looking at the bright side.

I knew that my deep dive into the two domains of climate research and AI, both of which I knew very little about, would be tough and connected to a lot of uncertainty, and walking the line between research and design was not always easy. At the end of this project, I realize that I only scratched the tip of the iceberg, with so much more to explore. Yet, I see this not as a demotivating realization, but the rather exciting prospect of future exploration.

My design process and the outcome were certainly influenced by the technical challenges I faced on the way, as they forced me to readjust or divert energy from further working on the conceptual part. Nevertheless, I can only recommend to everyone to try and build something with data yourself. For future projects, I'd be eager to work in a multidisciplinary team with computer scientists and engineers on board to push the boundaries of my creativity.

When this project taught me one thing, then it is humbling respect for other researchers and engineers working in the two domains of AI and climate research. Not only conceptualizing but actually conducting the research and building working AI systems, still seems like a marvelous task to me.

06.06 TIPS FOR OTHER STUDENTS

There are some things I would have liked to do during my graduation project, but also some small elements that I did do, and which worked out well for me. They might seem obvious, but perhaps some other student reading those words will find them useful, nonetheless.

1. Write things down while you go

Throughout your graduation you will do a lot of different activities in a short time, some of them simultaneously. It is only too easy to get lost in your activities and forgetting to document what you were doing and what you learned on the way.

Continuously writing down insights and references of where or how you found that information will save you a lot of stress in the long run and make reporting your project substantially easier.

2. Keep a diary of your activities

One thing that helped me during my project was to keep a small diary of my activities and write down all my meetings, activities, struggles, and decisions in a weekly summary. When going back through my process, just seeing a keyword about a day's activity helped me immensely to recall the event and my decision-making.

3. Consider carefully if you want to graduate without a company

I self-initiated my graduation and decided to do it in big parts without a company. While this approach certainly gave me more freedom to pursue my interests, this freedom can also be challenging to work with. Especially in COVID times and working remotely, I sometimes wished to have a clearer context.

4. Try to find pride in the small things

Graduation and writing your master thesis will be a challenge to most. Sometimes you will be lost, have no idea what to do, or if anything you are doing even makes remote sense. However, it will help to still see the project as a big learning process and take pride in even the smallest successes.

REFERENCES

- [Photograph of engineers drawing on paper]**. (n.d.). <https://www.vintag.es/2018/08/life-before-autocad.html>
- AdaIN-Style-Transfer - Runway**. (2019). Runway. <https://app.runwayml.com/models/reiinakano/AdaIN-Style-Transfer>
- Adobe Lightroom**. (2020). [Photo editing software]. Adobe. <https://lightroom.adobe.com>
- AI Artists**. (n.d. a). About Us. AIArtists. Retrieved 12 January 2021, from <https://aiartists.org/about>
- AI Artists**. (n.d. b). Unanswered Questions. AIArtists. Retrieved 12 January 2021, from <https://aiartists.org/unanswered-questions>
- Anderegg, W. R., Prall, J. W., Harold, J., & Schneider, S. H.** (2010). Expert credibility in climate change. *Proceedings of the National Academy of Sciences*, 107(27), 12107-12109.
- Arnfield, A. John** (2020, November 11). Köppen climate classification. Encyclopedia Britannica. <https://www.britannica.com/science/Koppen-climate-classification>
- Artbreeder**. (n.d.-a). Artbreeder. Retrieved 21 January 2021, from <https://www.artbreeder.com/>
- Artbreeder**. (n.d.-b). Artbreeder. Retrieved 21 January 2021, from <https://www.artbreeder.com/pricing>
- Balkan Green Energy News**. (2018, December 10). 30 years of Montreal Protocol – environmental success story. <https://balkangreenenergynews.com/30-years-of-montreal-protocol-environmental-success-story/>
- Bolin, B.** (1970). The carbon cycle. *Scientific American*, 223(3), 124-135.

- Brown, T. B., Mann, B., Ryder, N., Subbiah, M., Kaplan, J., Dhariwal, P., ... & Amodei, D.** (2020). Language models are few-shot learners. arXiv preprint arXiv:2005.14165.
- Climate Action Tracker** (2019). Governments still showing little sign of acting on climate crisis.
- Climate Action Tracker.** (2019, December 10). Global update: Governments still showing little sign of acting on climate crisis | Climate Action Tracker [Press release]. <https://climateactiontracker.org/press/global-update-governments-showing-little-sign-of-acting-on-climate-crisis/>
- Climate Central.** (2020, February 19). Global Temperatures and CO2 Concentrations (2020). <https://www.climatecentral.org/gallery/graphics/global-temperatures-and-co2-concentrations-2020>
- Columbus, L.** (2018, January 19). 10 Charts That Will Change Your Perspective On Artificial Intelligence's Growth. Forbes. <https://www.forbes.com/sites/louiscolumbus/2018/01/12/10-charts-that-will-change-your-perspective-on-artificial-intelligences-growth/?sh=1bb6935f4758>
- Corner, A., Shaw, C., & Clarke, J.** (2018). Principles for effective communication and public engagement on climate change: A Handbook for IPCC authors. Climate Outreach: Oxford, UK.
- DeepLearning.AI.** (2020, September 30). GANs for Good- A Virtual Expert Panel by DeepLearning.AI [Video]. YouTube. <https://www.youtube.com/watch?v=9d4jmPmTWmc>
- DeepMind.** (2020, November 30). AlphaFold: A solution to a 50-year-old grand challenge in biology. <https://deepmind.com/blog/article/alphafold-a-solution-to-a-50-year-old-grand-challenge-in-biology>
- Dowson, D. C., and B. V. Landau.** "The Fréchet distance between multivariate normal distributions." *Journal of multivariate analysis* 12.3 (1982): 450-455.

- ESA.** (2015). Sentinel-2 - Missions - Sentinel Online - Sentinel. The European Space Agency. <https://sentinel.esa.int/web/sentinel/missions/sentinel-2>
- EU funds.** (2017, February 16). Climate Action - European Commission. https://ec.europa.eu/clima/policies/adaptation/financing/funds_en
- European Commission & European Commission. Climate Action DG.** (2018). Our Planet, Our Future. UTB.
- Fridays For Future is an international climate movement active in most countries and our website offers information on who we are and what you can do.** (2020, September 1). Fridays For Future. <https://fridaysforfuture.org/>
- Fullerton, T., Swain, C., & Hoffman, S.** (2004). Game design workshop: Designing, prototyping, & playtesting games. CRC Press.
- Gatys, L. A., Ecker, A. S., & Bethge, M.** (2015). A neural algorithm of artistic style. arXiv preprint arXiv:1508.06576.
- Gharakhanian, A.** (2016). GANs: One of the Hottest Topics in Machine Learning. LinkedIn. https://www.linkedin.com/pulse/gans-one-hottest-topics-machine-learning-al-gharakhanian/?trk=pulse_spock-articles
- Goodfellow, I., Pouget-Abadie, J., Mirza, M., Xu, B., Warde-Farley, D., Ozair, S., ... & Bengio, Y.** (2020). Generative adversarial networks. Communications of the ACM, 63(11), 139-144.
- Google Colaboratory.** (2014). [Software]. Google. <http://colab.research.google.com/>
- Google Earth Engine.** (2010). [Software]. Google. <https://earthengine.google.com>
- GoogleEarth_02 - Runway.** (2020). Runway. https://app.runwayml.com/models/j3nsykes/GoogleEarth_02
- Guan, S.** (2018). TL-GAN: transparent latent-space GAN. GitHub. https://github.com/SummitKwan/transparent_latent_gan

- Härkönen, E., Hertzmann, A., Lehtinen, J., & Paris, S.** (2020). GANSpace: Discovering Interpretable GAN Controls. arXiv preprint arXiv:2004.02546.
- Hindupur, A.** (2018, September 30). The Gan Zoo. GitHub. <https://github.com/hindupuravinash/the-gan-zoo/blob/master/gans.tsv>
- Brock, A., Donahue, J., & Simonyan, K.** (2018). Large scale gan training for high fidelity natural image synthesis. arXiv preprint arXiv:1809.11096.
- Hora, K.** (n.d.). This X Does Not Exist. This X Does Not Exist. Retrieved 21 January 2021, from <https://thisxdoesnotexist.com>
- Isola, P., Zhu, J. Y., Zhou, T., & Efros, A. A.** (2017). Image-to-image translation with conditional adversarial networks. In Proceedings of the IEEE conference on computer vision and pattern recognition (pp. 1125-1134).
- Jakhar, D., & Kaur, I.** (2020). Artificial intelligence, machine learning and deep learning: definitions and differences. *Clinical and experimental dermatology*, 45(1), 131-132.
- Kaplan, A., & Haenlein, M.** (2019). Siri, Siri, in my hand: Who's the fairest in the land? On the interpretations, illustrations, and implications of artificial intelligence. *Business Horizons*, 62(1), 15-25.
- Karras, T., Aittala, M., Hellsten, J., Laine, S., Lehtinen, J., & Aila, T.** (2020). Training generative adversarial networks with limited data. arXiv preprint arXiv:2006.06676.
- Karras, T., Laine, S., & Aila, T.** (2019). A style-based generator architecture for generative adversarial networks. In Proceedings of the IEEE conference on computer vision and pattern recognition (pp. 4401-4410).
- King, S., & Chang, K.** (2016). *Understanding Industrial Design: Principles for UX and Interaction Design* (1st ed.). O'Reilly Media. <https://www.oreilly.com/library/view/understanding-industrial-design/9781491920381/ch01.html>

- Klingemann, M.** (2018). Mario Klingemann – Artist Profile (Photos, Videos, Exhibitions). AIArtists.Org. <https://aiartists.org/mario-klingemann>
- Laurans, G., & Desmet, P. M.** (2017). Developing 14 animated characters for non-verbal self-report of categorical emotions. *Journal of Design Research*, 15(3-4), 214-233.
- Lean, J., Beer, J., & Bradley, R.** (1995). Reconstruction of solar irradiance since 1610: Implications for climate change. *Geophysical Research Letters*, 22(23), 3195-3198.
- Lepri, B., Oliver, N., Letouzé, E., Pentland, A., & Vinck, P.** (2018). Fair, transparent, and accountable algorithmic decision-making processes. *Philosophy & Technology*, 31(4), 611-627.
- Lucero, A., & Arrasvuori, J.** (2010, September). PLEX Cards: a source of inspiration when designing for playfulness. In *Proceedings of the 3rd International Conference on Fun and Games* (pp. 28-37).
- McDivitt, P.** (2016). *The Information Deficit Model is Dead. Now What? Evaluating New Strategies for Communicating Anthropogenic Climate Change in the Context of Contemporary American Politics, Economy, and Culture* (Doctoral dissertation, University of Colorado at Boulder).
- ml5.js.** (n.d.). ml5js-Friendly Machine Learning For The Web. Retrieved 21 January 2021, from <https://ml5js.org>
- Molina, M., & Zaelke, D.** (2017). The Montreal Protocol: triumph by treaty. UN Environment. <https://www.unep.org/news-and-stories/story/montreal-protocol-triumph-treaty>
- Moss, R. H., Edmonds, J. A., Hibbard, K. A., Manning, M. R., Rose, S. K., Van Vuuren, D. P., ... & Wilbanks, T. J.** (2010). The next generation of scenarios for climate change research and assessment. *Nature*, 463(7282), 747-756.

NASA. (1999). Ozone. NASA Earth Observatory. https://earthobservatory.nasa.gov/features/Ozone/ozone_2.php

NASA. (2020). Volcanoes and Climate Change | Earthdata. Earthdata. <https://earthdata.nasa.gov/learn/sensing-our-planet/volcanoes-and-climate-change>

National Center for Environmental Information. (n.d.). Temperature Change and Carbon Dioxide Change. Retrieved 18 February 2021, from <https://www.ncdc.noaa.gov/global-warming/temperature-change>

OpenAI. (2021, January 13). DALL-E: Creating Images from Text. <https://openai.com/blog/dall-e/>

Oreskes, N. (2018, February 14). Fossil fuel firms are still bankrolling climate denial lobby groups. The Guardian. <https://www.theguardian.com/environment/2015/mar/25/fossil-fuel-firms-are-still-bankrolling-climate-denial-lobby-groups>

Oxford University Press. (n.d.). climate noun - Definition, pictures, pronunciation and usage notes | Oxford Advanced Learner's Dictionary at OxfordLearnersDictionaries.com. Oxford Learner's Dictionary. <https://www.oxfordlearnersdictionaries.com/definition/english/climate>

Park, T., Liu, M. Y., Wang, T. C., & Zhu, J. Y. (2019). Semantic image synthesis with spatially-adaptive normalization. In Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition (pp. 2337-2346).

Pasquinelli, M., & Joler, V. (2020). The Noosphere manifested: AI as an instrument of knowledge extractivism. *AI & society*, 1-18.

Planton, S. (2013). Annex III Glossary.

Porr, L. (n.d.). Nothing but Words. Nothing but Words. Retrieved 12 January 2021, from <https://ados.substack.com/>

- Processing** (3.5.4). (2020). [Desktop Application]. Processing Foundation
<https://processing.org>
- Project Drawdown**. (2020, August 20). Drawdown Framework | @ProjectDrawdown.
<https://drawdown.org/drawdown-framework>
- QGIS** (3.16.3). (2009). [Desktop Application]. OSGeo. <https://www.qgis.org/>
- RE100**. (n.d.). RE100 - We are accelerating progress towards 100% renewable power.
<https://www.there100.org>
- Ritchie, H.** (2017, October 2). Fossil Fuels. Our World in Data. <https://ourworldindata.org/fossil-fuels>
- Rolnick, D., Donti, P. L., Kaack, L. H., Kochanski, K., Lacoste, A., Sankaran, K., ... & Bengio, Y.** (2019). Tackling climate change with machine learning. arXiv preprint arXiv:1906.05433.
- Runway**. (2019). StyleGAN - Runway. <https://app.runwayml.com/models/runway/StyleGAN>
- Runway**. (n.d.). Runway | Make the Impossible. Retrieved 21 January 2021, from <https://runwayml.com>
- Salehi, P., Chalechale, A., & Taghizadeh, M.** (2020). Generative Adversarial Networks (GANs): An Overview of Theoretical Model, Evaluation Metrics, and Recent Developments. arXiv preprint arXiv:2005.13178.
- Salehi, P., Chalechale, A., & Taghizadeh, M.** (2020). Generative Adversarial Networks (GANs): An Overview of Theoretical Model, Evaluation Metrics, and Recent Developments. arXiv preprint arXiv:2005.13178.
- Schmidt, A., Ivanova, A., & Schäfer, M. S.** (2013). Media attention for climate change around the world: A comparative analysis of newspaper coverage in 27 countries. *Global Environmental Change*, 23(5), 1233-1248.

- Schultz, D.** (n.d.). Artificial Images. YouTube. Retrieved 21 January 2021, from <https://www.youtube.com/user/bustbright/about>
- Schwarz, N.** (2000). Emotion, cognition, and decision making. *Cognition & Emotion*, 14(4), 433-440.
- Sentinel-2 MSI: MultiSpectral Instrument, Level-2A.** (2015). [Dataset]. ESA. https://developers.google.com/earth-engine/datasets/catalog/COPERNICUS_S2_SR
- Sevarac, Z.** (2020, June 5). Machine Learning Tutorial for Java Developers. Deep Netts Blog. <https://www.deepnetts.com/blog/machine-learning-tutorial-for-java-developers.html>
- Shifman, D.** (2019, August 8). Introduction to Runway: Machine Learning for Creators (Part 1). YouTube. <https://www.youtube.com/watch?v=ARnf4ilr9Hc>
- Shifman, D.** (n.d.). The Coding Train. The Coding Train. Retrieved 21 January 2021, from <https://thecodingtrain.com>
- Sim Earth.** (1990). [Video Game]. Maxis.
- Singer, J.** (2020, July 18). Jordan Singer. Twitter. <https://twitter.com/jsngr/status/1284511080715362304>
- Singh, N.** (2017). How to Get Started as a Developer in AI. Intel. <https://software.intel.com/content/www/us/en/develop/articles/how-to-get-started-as-a-developer-in-ai.html>
- Smith, N., & Leiserowitz, A.** (2014). The role of emotion in global warming policy support and opposition. *Risk Analysis*, 34(5), 937-948.
- Spencer Weart & American Institute of Physics.** (2017a). Impacts of global warming | impacts. Impacts of Global Warming. <https://history.aip.org/climate/impacts.htm>
- Spencer Weart & American Institute of Physics.** (2017b). Introduction - Summary. The Discovery of Global Warming. <https://history.aip.org/climate/summary>.

- Stappers, P. J., & Visser, F. S.** (2014). Meta-levels in design research: Resolving some confusions. *DRS*, 16-19.
- StyleGAN2-ada custom training.** (2020). [Google Colab Notebook]. Derrick Schultz. https://github.com/dvschultz/ml-art-colabs/blob/master/Stylegan2_ada_Custom_Training.ipynb
- StyleGAN2-ADA generate.** (2020). [Software]. Derrick Schultz. https://github.com/dvschultz/stylegan2-training/blob/main/SG2_ADA_Generate.ipynb
- Tromp, N., & Hekkert, P.** (2018). *Designing for society: products and services for a better world*. Bloomsbury Publishing
- United Nations Environment Programme.** (n.d.). The Montreal Protocol. UN Environment. <https://www.unenvironment.org/ozonaction/who-we-are/about-montreal-protocol>
- United Nations Framework Convention on Climate Change.** (n.d.). The Paris Agreement. UNFCCC. <https://unfccc.int/process-and-meetings/the-paris-agreement/the-paris-agreement>
- United Nations.** (2018). Act Now. <https://www.un.org/en/actnow>
- USGS.** (n.d.). Volcanoes Can Affect Climate. USGS.Gov | Science for a Changing World. Retrieved 18 February 2021, from <https://www.usgs.gov/natural-hazards/volcano-hazards/volcanoes-can-affect-climate>
- Veale, M., Van Kleek, M., & Binns, R.** (2018, April). Fairness and accountability design needs for algorithmic support in high-stakes public sector decision-making. In *Proceedings of the 2018 chi conference on human factors in computing systems* (pp. 1-14).
- Voosen, P.** (2018, May 18). Trump White House quietly cancels NASA research verifying greenhouse gas cuts. *Science | AAAS*. <https://www.sciencemag.org/news/2018/05/trump-white-house-quietly-cancels-nasa-research-verifying-greenhouse-gas-cuts>

- Wang, D., Weisz, J. D., Muller, M., Ram, P., Geyer, W., Dugan, C., ... & Gray, A.** (2019). Human-AI collaboration in data science: Exploring data scientists' perceptions of automated AI. *Proceedings of the ACM on Human-Computer Interaction*, 3(CSCW), 1-24.
- Weber, E. U., & Stern, P. C.** (2011). Public understanding of climate change in the United States. *American Psychologist*, 66(4), 315.
- Xia, G. S., Hu, J., Hu, F., Shi, B., Bai, X., Zhong, Y., ... & Lu, X.** (2017). AID: A benchmark data set for performance evaluation of aerial scene classification. *IEEE Transactions on Geoscience and Remote Sensing*, 55(7), 3965-3981.
- Zhu, J. Y., Park, T., Isola, P., & Efros, A. A.** (2017). Unpaired image-to-image translation using cycle-consistent adversarial networks. In *Proceedings of the IEEE international conference on computer vision* (pp. 2223-2232).
- Zimmerman, J., & Forlizzi, J.** (2017). Speed dating: providing a menu of possible futures. *She Ji: The Journal of Design, Economics, and Innovation*, 3(1), 30-50.

