

Universität für
künstlerische
und industrielle
Gestaltung –
Kunstuniversität
Linz

Institut für
Medien –
Interface
Cultures

Masterarbeit zur
Erlangung des
akademisches
Grades

Master of Arts

I, rowboat: The Intricacies of AI & Poetry

Fabian Frei

2. Dezember 2020
Linz

Betreut von:
Univ. Prof. Dr.
Laurent Mignonneau

Univ. Prof. Dr.
Christa Sommerer

Univ. Prof. Dr.
Manuela Naveau

Unterschrift des
Betreuers, der
Betreuerin:



Acknowledgments

I would like to thank the following people for their guidance, support, inspiration and encouragement:

Laurent Mignonneau, Christa Sommerer, Davide Bevilacqua,
Corinne Amrein, Isabella Auer, Giacomo Piazzzi, Wesley Lee

I want to thank Pete Hindle for his relentless and exact editorial work.

I want to thank Bernadette Amrein, Christian Gorbach and smec for providing me with the needed hardware to develop the project presented in this thesis.

Table of Contents

Abstract	6
1. Introduction	10
2. Data sets and Curation	12
2.1. ImageNet	13
2.2. JAFFEE	18
2.3. AI in Music	20
2.4. Curation	21
3. Poetic Intention and Collective Authorship	24
3.1. Introduction	25
3.2. The Author	26
3.3. Automated text generation	29
4. Creative Output as a Commodity	36
4.1. Introduction	37
4.2. Compensation	38
4.3. Creativity	41
4.4. Applications in the Creative Industries	46
5. Personal Experiments & Work	52
5.1. Introduction	53
5.2. Previous Work	53
5.3. Inspirations	60
5.4. The Act of Walking	62
6. Setup of Artwork	66
6.1. Introduction	67
6.2. The Smartphone Application	68
6.3. The Backend	70
6.4. Augmenting the Written Poems	72
6.4.1 Markov Chain	76
6.4.2 Queneau	78
6.4.3 Synonyms	82
7. Training & Generation	84
7.1. Introduction	85
7.2. GPT-2	86
7.3. Generation	88
8. Related Artists & Artistic Work	92
8.1. Mario Klingemann	93
8.2. Rafael Pérez y Pérez	94
8.3. Nick Montfort	97
8.4. Allison Parrish	100
8.5. Ross Goodwin	106
8.6. Tony Veale	109
8.7. Disnovation.org	113
8.8. Jörg Piringner	118
9. Conclusion & Findings	124
10. Bibliography	130
10.1. Further Reading	137

Abstract

Wir leben zunehmend in einer Welt, in der von Maschinen generierter Text alltäglich ist. Die Auseinandersetzung mit den Maschinen sowie ihrer Software, ihren Prozessen, ihren Datensätzen und die daraus resultierend Texten ist ausserordentlich wichtig.

Die vorliegende Abschlussarbeit betrachtet diese Problemstellung im Kontext der Künste und versucht zu verstehen, wie in diesem Kontext sinnvolle Kunstwerk mit Einbezug von künstlicher Intelligenz geschaffen werden können. Es wird betrachtet wie weit genutzte Datensätze aufgebaut sind und welche Probleme sich verbergen, ob und wie sich der Begriff des Autors sich in der Konfrontation mit Text-Generatoren verändert. Weiter erfasst der Text die Situation der Kompensation der Produzenten der Inhalte, die in einem Trainings-Datensatz enden und fragt ob Computer wirklich kreativ sein oder nicht und wie sich das auf die Produktion von kulturellen Fragmenten auswirkt.

Der Fokus der künstlerischen Forschung sind Gedichte, sowohl als Input wie auch als Output. Sie werden von Teilnehmern zu einem Datensatz beigetragen. Dieser Datensatz wird dann genutzt, um eine AI zu trainieren. Die AI wird danach Gedichte generiert.

We increasingly live in a world where machine-generated text is commonplace. The examination of machines as well as their software, their processes, their data sets, and the resulting texts is extremely important.

The present thesis looks at this problem in the context of the arts and tries to understand how sustainable and meaningful artwork can be created in this context with the inclusion of artificial intelligence. It considers how widely used data sets are constructed and what problems are hidden, whether and how the concept of the author changes in the confrontation with text generators. Furthermore, the text captures the situation of compensation of the producers of the content ending in a training data set and asks whether computers are really creative or not, and how this affects the production of cultural fragments.

The focus of artistic research is poetry, both as input and output. They are contributed to a data set by participants. This data set is then used to train an AI. The AI will then generate poems.

1.

Intro- duction

My personal motivation to work on this thesis is grounded in a deep interest in text and technology, and how these topics interact with our contemporary society. With the emergence of artificial intelligence based approaches to text generation it seemed to me obvious to try my hand at such a method to generate poetry. But instead of using a prepared dataset I wanted to invite people to collaborate on a dataset of poems written specially for this occasion.

I have identified the following questions as the main research directions.

Would it be better for artists that want to produce art with artificial intelligence to curate their own dataset instead of trying to innovate on the actual algorithms?

How does the notion of the author change when confronted with non-human creativity in the form of a text generating machine? Does the concept of a sole author still hold up?

Does the romantic ideal of the genius author actually hinder our understanding of our personal poetic intent, and of computer and machine creativity?

Throughout my thesis research I want to find out if helping users unconsciously overcome the same romantic ideal heightens the quality of their poetic texts.

With the rapid progress of AI technologies, we can see that contemporary machine learning approaches like GPT-2 need huge datasets to produce state-of-the-art results. In this thesis I want to examine the relationship between the actors that curate the dataset and the actors that produce the content inside these sets.

2.

**Data sets
and
Curation**

2.1. ImageNet

ImageNet is a collection of categorized images, with over 14 million entries as of 2020. It is one of the most significant collections of images in machine learning and was used or cited in over 300 research papers. Started in 2009, and according to its co-creator Fei-Fei Ling, its aim is to “map out the entire world of objects” (Gershgorn, 2017). ImageNet uses the semantic structure of WordNet, a project dedicated to organizing the entire English language. ImageNet uses a nested structure of cognitive synonyms, or “synset”, a distinct concept where synonyms are grouped accordingly. For example, “home” and “place”, as in the phrase “my place”, belong in the same synset. These synsets are then organized in a nested hierarchy, going from general to more specific. In this example “home” would be located at the end of the following structure: Entity → physical entity → object → location → point → geographic point → address → residence → home. Unlike WordNet, ImageNet only concerns itself with nouns, as in things that can be represented with images.

In 2009, researchers started collecting images for each synset by automatically downloading them from search engines. At that point in time, search engines “did not have accurate algorithms to understand the content of images, so they retrieved images mostly based on their captions or tags” (Yang et al., 2019). This resulted in an extraordinary amount of wrongly categorized images. To mitigate this, and almost as an afterthought, they used Amazon’s Mechanical Turk¹ to get help from minimally paid workers to correctly place the image in the corresponding synsets.

¹ Amazon’s Mechanical Turk is a service for atomized tasks that humans can solve quicker than computers. Tasks are either submitted via a software interface or via the services’ website.

They extended this cooperation into an automated system that would assign the downloaded images to Mechanical Turk workers, owing to the scale “of the annotation effort (over 50K workers and over 160M candidate images)” (Yang et al., 2019).

The process of creating categories is complicated, and not a bias free process. “To create a category or to name things is to divide an almost infinitely complex universe into separate phenomena” (Crawford, Paglen, 2019). Nouns are located in a two-dimensional space, along an axis from concrete to abstract, and a second axis from descriptive to judgmental.

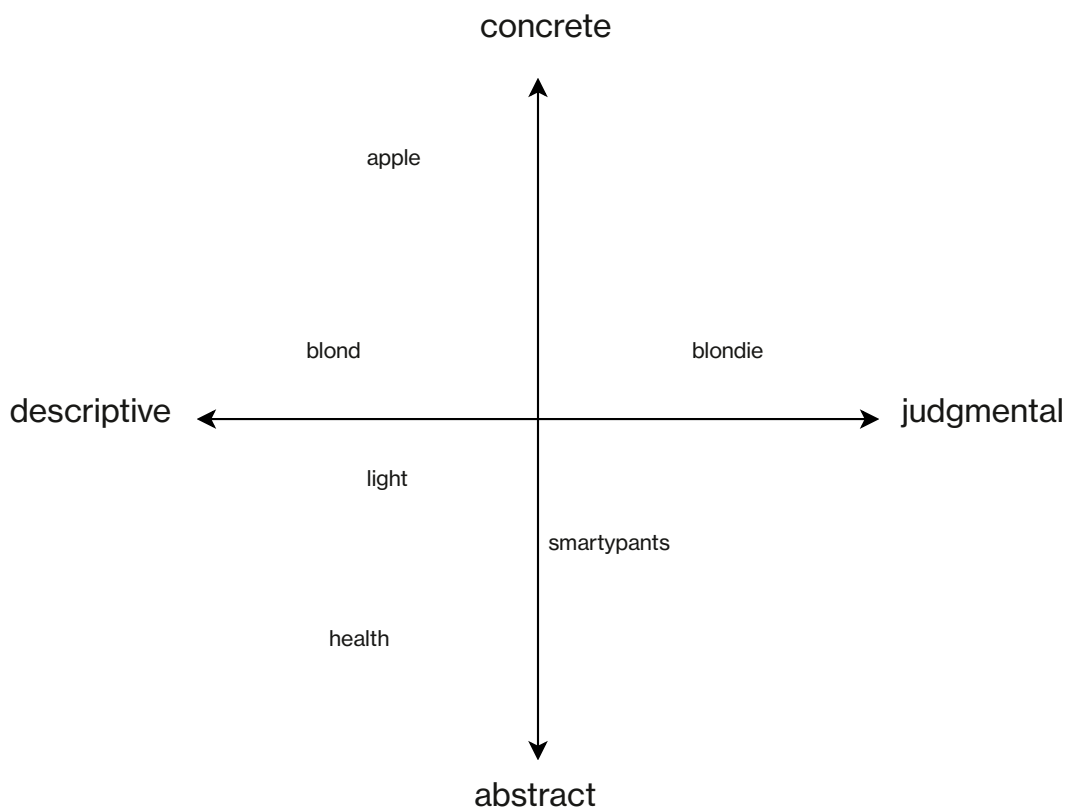


Figure 1: Two-dimensional noun space

The dimensions have been removed from the words added to ImageNet because of the unsupervised and automated process of extracting nouns from WordNet and annotating them through Mechanical Turk workers. Surely for simple examples like “chocolate” and “chocolate cookie” the cognitive distance is short but it begins to fall apart quickly. “The concept of an “apple” is more nouny than the concept of “light”, which in turn is more nouny than a concept such as “health” (Crawford, Paglen, 2019). This becomes especially problematic when pictures of people are annotated with categories.

“As the fields of information science and science and technology studies have long shown, all taxonomies or classificatory systems are political”(Crawford, Paglen, 2019). The ‘people’ category of ImageNet reveals a certain world-view, and classifies people along lines like race, gender, nationality, profession, behaviour, character, and morality. It tries to apply subjective categories to images of people automatically via underpaid and time strapped workers, an undertaking that is doomed to produce crude and problematic results.

At no point in the paper (Deng et al., 2009) do the authors talk about or even recognize the problem of simply taking pictures from the internet. Gathering consent from or even communicating with the persons depicted was not attempted, and so we find problematic arguments like “More sophisticated and robust models and algorithms can be proposed by exploiting these images” (Deng et al., 2009).

Mapping the entire world of objects is an enormous undertaking, maybe even similar in scope to the moon landing. One could argue that technological progress has made such a project easier and its processes more efficient. The approach of ImageNet superficially seems reasonable, and a case of gains made by technological efficiency, but it is built on two questionable and exploitative processes.

The curation of the dataset, including collecting and annotating the images, was not treated with the afforded care. Firstly by treating images on the internet as a public, freely accessible, and usable resource, the project appropriates the work of thousands of internet users. The difficult and time-consuming work of consulting, informing or even rewarding the creators of the images was simply excluded.

Secondly the annotation process relies on so-called ghost work or human computation, as described in “Ghost work: how to stop Silicon Valley from building a new global underclass” (Gray and Suri, 2019). An algorithm is a set of steps to solve a certain problem and human computation is the idea of atomizing a complex task into a few steps. These tasks can then be distributed via computer systems to workers all over the world. ImageNet used such an automated system to hand over the task of annotation to Mechanical Turk workers. This strict compartmentalization of the relationship between the project and its workers, via a programmatic interface, is not conducive to any other form of communication that is not about the task at hand. In (Gray and Suri, 2019) it is described as “algorithmic cruelty—literally, computation incapable of thought, let alone empathy” (Gray and Suri, 2019). This disconnect devalues the workers and the work they do, and while the project profits the workers, they are kept in their position and gain nothing, apart from a meagre pay.

This aspect, together with the minimal “median wage [that] lies between \$1.77/h and \$2.11/h” (Hara et al., 2018) constitutes an exploitative relationship.

While the two aspects shown by ImageNet – its exploitative relationship with the human workers that created its content, and ImageNet’s application of subjective categories to pictures of people usually results in harm – are important, I would like to look at another dataset. This other example should help us identify another common problem in the space of dataset curation: sets that are built on questionable science.

2.2. JAFFEE

“The Japanese Female Facial Expression (JAFFE) Database” is a dataset developed in 1998 by Michael Lyons, Miyuki Kamachi and Jiro Gyoba. It contains pictures of 10 Japanese women making seven facial expressions that map to seven basic emotional states. The dataset has been greatly used in research on affect recognition.

The concept of reading emotions from faces comes from the psychologist Paul Ekman “who asserted that emotions can be grouped into a small set of basic categories like anger, disgust, fear, happiness, sadness, and surprise” (Whittaker et al., 2018). Ekman did research to understand if emotions represented by the connected facial expressions are shared between cultures, and came to the conclusions that the face provides a window into the inner workings of the soul. “But Ekman’s work has been deeply criticized by psychologists, anthropologists, and other researchers who have found his theories do not hold up under sustained scrutiny.” (Crawford, Paglen, 2019)

To test the hypothesis that basic facial configurations evolved in the Pleistocene, studies took place that asked people from remote uncontacted tribes to “label facial configurations as emotional expressions (i.e., studying emotion perception, not production)” (Barrett et al., 2019) but there seem to be no studies that “measured facial muscle movements in people of remote cultures as they experienced emotional events” (Barrett et al., 2019).

Another important point is that in the JAFFE dataset the actresses were performing facial expressions, and were not exhibiting emotions because of a scenario that they encountered. It is “clear that posed and spontaneous movements differ, sometimes quite substantially” (Barrett et al., 2019).

The presented dataset is problematic because it transports the concept of seven basic categories of emotions that can be applied by a beholder, be it a machine or a human, just from seeing a face. This aspect is reinforced by the curation of the images as they are posed facial expressions by actresses. The poses are not an expression of real emotions, but instead are a representation of what the people involved thought that these emotions look like. This carries the danger that these facial expressions become amplified caricatures of themselves.

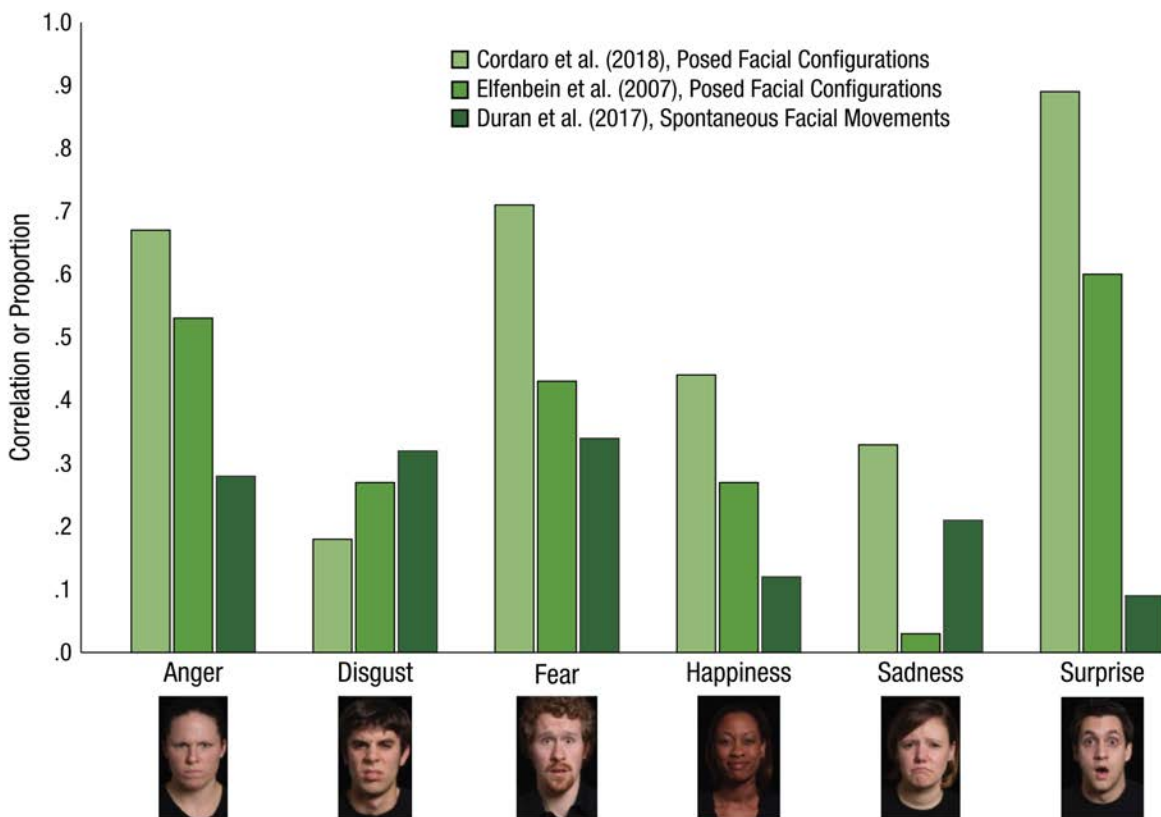


Figure 2: Comparing posed and spontaneous facial movements. Correlations or proportions are presented for anger, disgust, fear, happiness, sadness, and surprise, separately for three studies. (Barrett et al., 2019)

Unlike conventional algorithms in computer science and computer arts that encode knowledge in program code, in machine learning algorithmical knowledge has to be extracted from a dataset. With these two case studies I want to propose that, for artists, who do not usually have access to big amounts of money or computing resources, it makes sense to explore the creation of custom datasets built on consent and inclusion.

2.3. AI in Music

Thought-out curation can lead to new schemes of authorship. Let's look at the most recent music album "PROTO" produced by Holly Herndon which was written together with a choir of 14, including a developer called Jules LaPlace and Spawn. Spawn is a musical AI that Mat Dryhurst and Herndon built over a few years. They trained it to reproduce both the voice of Holly, but also of the whole choir. Instead of letting the AI generate all the music, Spawn became part of the choir as another voice.

Holly explains that in creating the record there was an intensive exchange of sounds between the humans and the AI: "Sometimes I'd play them back the processing, and they'd emulate that, then Spawn would have her version. There was a ton of back and forth" (Love Michael, 2019) On some tracks of the record, the influence and the output of Spawn are clearly audible. Still the sound produced by Spawn is easily recognizable way that mostly fails to capture the listeners' attention.

Further the output strongly resembles other projects in that space of AI produced sound like *Actresses Young Paint*².

² Please see https://youtu.be/ZsZc4Q_eDk4 for an excerpt

In other compositions, the singing voices are more elaborately layered, and the influence of the generated voices is less clear. Personally I think that these are the strongest moments of this work, where the difference between human voices and the generated voices is fuzzy. Holly remarks that Spawn was not the main part of the project: “Even though AI’s the fun, exciting thing to talk about, it really only makes up 20% of the audio. Much of it is human voices sharing space together.” (Friedlander, 2019)

In the context of this chapter, one point of “PROTO” is important: Herndon and her collaborators did not build their own algorithms, but used open source implementations of scientific papers like SampleRNN (Mehri et al., 2016).

2.4. Curation

The current state-of-the-art text generation models, that are based on the transformer architecture like GPT-2, BERT and GPT-3, are trained on a big part of the English-speaking internet. They require a huge and easy-to-harvest training set. This poses the problem of how to filter for quality content within the enormous trove of data harvested. Institutions like OpenAI, Facebook and Google are trying to find ways to pick “the good from the bad on a mass scale” (Thi Nguyen and Weinberg, 2020)

In the case of GPT-2 and GPT-3, OpenAI decided to scrape “all outbound links from Reddit, a social media platform, which received at least 3 karma” (Radford et al., 2019). Karma is used within Reddit to reward a funny, interesting, or simply good post submissions. In this context it is a valid shortcut to get a set of websites curated by humans- but what if there is part of Reddit where people only share links on how amazing Telsa cars are? Prompt the network to talk about cars, and it will probably

give you an answer that is biased towards Tesla cars. This is not entirely what we envisioned, but it is the outcome of the combination of the huge amount of data and the simple heuristic filter that is applied.

This is where the concept of curation comes in. By using existing algorithms or pretrained models, we can focus our time on building datasets that include a more nuanced and artistic idea of good source material. We can apply different measures that are more meaningful and go beyond the abstract reasoning in the form of statistics. Instead of heuristics that are centred around getting the biggest amount of material, we can look at single data points and see them in context. We can talk to the people responsible for the training data, include them in our project and ask them for their consent. We can investigate differences and local specialities that would never show up in a multimillion entries dataset.

3.

**Poetic
Intention
and
Collective
Authorship**

3.1. Introduction

I define poetic intention as an inherent property of a human being and its creativity. It is the ability to misuse the tool of language to express text fragments that evoke/carry multiple meanings on the levels of semantics, sounds and typography. I base this definition on the concept of extended poetry by H.C. Artman from the Wiener Gruppe.

The Wiener Gruppe was a small and loose cluster of Austrian poets and writers that consisted of H.C. Artman, Friedrich Achleitner, Konrad Bayer, Gerhard Rühm, and Ingrid and Oswald Wiener. It is difficult to pin down the start and the end of the group but in the time between 1954 and 1960 the majority of the group was working together. They were mostly focused on creating contemporary prose but expanded quickly into endeavours like vernacular poetry, constructive poetry, photography, film and further founded new art forms such as the action, the happening and conceptual art (Weibel, 1997). Of most interest to my context is the move from single author “classical” poems to collective montage works, the “methodological inventionism” and the “Eight-Point-Proclamation of the Poetical Act”, to which we will return later in this chapter.

3.2. The Author

In the texts “The Death of the Author” (Barthes, 1968), “What is an Author” (Foucault, 1969) and “On the Author Effect. Recovering Collectivity” (Woodmansee, 1992) the idea of a text as an original work of a sole author are refuted. They see any textual fragments, including poems, as an intricate collection of cultural quotes, links and referrals that the author arranges them in a novel way. When Barthes was writing the “Death of the Author”, the practice of explication de texte, which links the life experiences of the author to his output, was considered to explain the meaning of any writing work. The relationship between a work and its author has been a temporal one, the author predates the work, he exists for it, but this does not hold up. The modern author exists at the same time as his text and from the author’s past, and no intention can be extracted. This leads further to the realization that text does not contain any inherent truth instilled by the god – the author – and this truth cannot be uncovered, only the connection of the quotes disentangled. This is because this collection of quotes does not assemble itself into a complete picture in the author but the reader assembles them into a coherent whole during the reading.

Martha Woodmansee is an American professor at Case Western Reserve University in Cleveland, Ohio. She is a 1999 Guggenheim fellow and 2004 Fulbright fellow, and in 2008 she founded the International Society for the History and Theory of Intellectual Property.

Woodmansee shows that the idea that the author as the role that should get exceptional attention, in the process of making a book, is a product of recent times. She quotes the “Allgemeines Oeconomisches Lexicon” that was produced by Georg Heinrich Zink in Leipzig in 1753, in which it is written that a big group of craftsmen are involved and the writer is seen as one of them. The notion of the exceptionalism of the author only is a by-product of the romantic view in literature that genius authors break with the tradition as a whole to come up with something totally new and singular. This new concept was originally outlined in Edward Youngs “Conjectures on Original Composition” and was later developed by a cohort of writers like Herder, Goethe, Coleman and Wordsworth who belonged to the emerging new industry of professional authors.

Woodmansee further uses the example of the English writer Samuel Johnson, who is considered to be the archetype of a modern author. Johnson most known works, such as “Dictionary of the English Language” written in 1755, “The Plays of William Shakespeare” written in 1765, together with “Lives of the Poets” written in 1779-81 are all collective and cooperative works. “Lives of the Poets” is a multivolume collection of England's most significant contemporary writers, to which Johnson each added a preface which discusses the writers critically. The collection was started on a suggestion of the booksellers of London, who feared sales from a similar collection that had been released by a publisher from Edinburgh. After Johnson's publishers saw the subpar quality of the Edinburgh version they immediately embarked on curating their own version, which would include authors to which the bookkeepers owned the rights to. The finished collection was essential in defining the difference between an author and a “simple” writer, because it created a group of great writers who had created works that were better than average writing.

Even though this collection transported this myth successfully it was not the product of Johnson's creative curatorial work but the outcome of a collaboration between Johnson, the portrayed poets, the booksellers of London and others. To illustrate these "others" Woodmansee shows that the preface to the Pope Issue called "Life of Pope" includes a part about the personal customs of Pope. This was incorporated by Johnson without remarking that it came from the "Universal Magazine" from August 1775 or from the "Gentleman's Magazine" from September 1775.

The liberties that Johnson took while writing were also extended to others. He was a proficient ghostwriter that almost always honoured requests of his friends to produce a text. He also wrote a fair share of sermons for priests in his lifetime, and for him, it was clear that after payment these texts were not his any more.

These points support the theory that writing was and is never the singular work of an exceptional genius author both on the level of reading and the text itself, as Barthes shows, but also on the level of craft. The author as the single source of text is further questioned by text machines and story generators as Simanowski writes in (Simanowski, 2009).

3.3. Automated text generation

Robert Simanowski is a German scholar of literature and media studies. He was, among other roles, a professor of media studies at the University of Basel in Switzerland (2010-2013) and at City University of Hong Kong (2014-2017).

“Automated text generation may be the oldest form of digital literature.” (Simanowski, 2009) Already in 1952 Christopher Strachey constructed a program he dubbed “Love Letter Generator”. The software creates love letters based on a simple pattern and a set of predefined words:

DUCK DEAR

MY AMBITION LOVES YOUR FONDNESS. MY AFFECTION YEARNNS
FOR YOUR IMPATIENT DEVOTION. YOU ARE MY DARLING FELLOW
FEELING. MY FERVENT YEARNING FONDLY LIKES YOUR
IMPATIENT LUST. MY EAGER EAGERNESS CARES FOR YOUR
FONDNESS.

YOURS WISTFULLY

M. U. C.³

One could argue that the creator/author disappears because he is not directly included in the creation of the love letters. According to Simanowski the case at hand is a bit more complex. It seems that Strachey's intention was not to create a general generator for love letters, but that he wanted to deconstruct the standard love letter. As a gay man living in postwar England he experienced homophobia, like his contemporary Turing. He wanted to parody the common form of the love letter and thus has not created a machine to generate love letters, but wrote software that expresses his feelings.

³ Output of an emulator written by David Link: http://www.alpha60.de/art/love_letters/

RACTER (from French *raconteur*) is a story generator written in BASIC by William Chamberlain and Thomas Etter in 1983. It generates prose and poems from a database of 2400 words categorized by identifiers. This allows matching adjectives and nouns, and provides the correct context for nouns and pronouns. Sentences generated by RACTER are grammatically correct and it's reuse of phrases previously applied gives the story a sense of continuity. Still, the choice of words is random, and the text appears to be senseless overall. In 1984 Chamberlain and Etter published the book called "The Policeman's Beard is Half-Constructed":

"War", chanted Benton, "war strangely is happiness to Diane." He was expectant but he speedily started to cry again. "Assault also is her happiness." Coldly they began to enrage and revile each other during the time that they hungrily swallowed their chicken.

Suddenly Lisa sang of her desire to Diane. She crooned quickly. Her singing was inciting to Benton. He wished to assassinate her yet he sang.

"Lisa, chant your valuable and interesting awareness." Lisa speedily replied. She desired possessing her own consciousness.

"Benton," she spoke, "you cry that war and assault are a joy to Diane, but Your consciousness is a tragedy as is your infatuation. My spirit cleverly recognizes the critical dreams of Brenton. That is my pleasure."

Benton saw Lisa, then began to revile her. He yodelled that Lisa possessed an infatuation for Diane, that her spirit was nervous, that she could thoughtfully murder her and she would determinedly know nothing. Lisa briskly spoke that Benton possessed a contract, an affair, and a story of that affair would give happiness to Diane. They chanted sloppily for months.

At allevents I quickly stop chanting now.

(Chamberlain & Etter, 1984)

The introduction to the books claims that all text apart from the introduction were written alone by RACTER but the RACTER FAQ (Barger, 1993) strongly suggests that a lot of manual work was needed. Manual work in this case means both at the stage of generation (like adding and tweaking special templates) but also selecting appropriate pieces from the output of the machine.

It is visible already from these two short excerpts that the default mode of text generation machines is the absurd. Thus, the author and/or programmer surely feels the urge to generate text as normal and conventional as possible. Calvino write about this problem in *Cibernetica e fantasma*: “The test of the poetic-electronic machine would be its ability to produce traditional works, poems with closed metrical forms, novels that follow the rules” (Simanowski, 2009). The textual output should not be seen as the main aspect of the work. Instead, the machine and the software that generated the text are the main parts of the artwork. Also, in computer science there is a strong preference to solve a problem in the most general way possible.

This conjecture is relevant for prose but less so for poetry. While prose is more about longer text in which we value coherency highly, “poetry lives on strange metaphors and the violation of linguistics norms and common ways of expression” (Simanowski, 2019) The stronger the weirdness of the words, the stronger we feel the weirdness that the lyrical subject encounters and wants to communicate. It seems that we need or look for an author in this type of text. Less to find that human soul, that represents authentic emotions, in it but more to project our thoughts and emotions onto somebody. Including the programmer’s intent in the text belongs to this. If we accept that the machine can also be an author, we actually encounter a situation where instead of a single creator we have a work with two authors superimposed over each other.

This allows us to read the text in multiple ways, but only two of them are important in the context of this project: first we can read the text as a technical achievement of the programmer author. Secondly we can read the text as an amalgam of the different authors involved, probing the space of quotations from different perspectives, and trying to understand who else was involved.

This view makes computer generated poetry exceptionally suitable to explore the possibilities for ways of collective authorship.

Permutational poetry has a long tradition outside digital computation. Notable examples are Tristan Tzara's "How to Write a Dada Poem" from 1920 that suggest the creation of a poem by cutting out words from existing texts and rearranging in a new random order. This technique was popularized by William S. Burroughs later in the 1950. Further there is the approach of Raymond Queneau of recombining each line of ten sonnets in "Cent mille milliards de poèmes" in 1961. We will encounter Queneau again later in the applied part of this project.

Wiener Gruppe moved through different aspects of poetry, similar to the ones discussed in this chapter. During the formation of the group they concerned themselves mostly with writing classical poetry until around 1953 when H.C. Artman published the "Eight-Point-Proclamation of the Poetical Act":

There is one statement which is irrefutable, namely that one can be a poet without having so much as written or spoken a single word.

The precondition for this, however, is a more or less strongly felt desire to act poetically. The a-logical gesture itself, thus executed, can be elevated to an act of excellent beauty, even to the level of a poem. The term beauty, however, has been granted a very wide scope in this context. (...)

(Weibel, 1997)

With this declaration H.C. Artman introduced the term expanded poetry. “He claimed that everyone should become a poet, and that the poetic act and the poetic result should be completely opened up” (Sommerer & Mignonneau, 2015) With this the group started to go into a different direction of understanding how to make poetry. Shortly after in 1954 Marc Adrian who was obsessed with the golden ratio, suggested implementing it in their poetry. This was the invention of the Methodic Inventionism and “we are very proud of ourselves. Now (we’re talking about 1954) everyone can become a poet”(Weibel, 1997). Instead of relying on the innate creativity of a writer, they were using a stock of words – the so called verbarium – and following a mathematical series. The Wiener Gruppe started with the golden ratio, but later any mathematical series was permitted, selected and ordered the words from the stock. This should result in a harmonic structure.

In 1958 H.C. Artman writes his first dialect poem “a rosn schwimmt ma fedaleicht duachs blauat”, Gerhard Rühm and Friedrich Achleitner join him in writing poems in dialect. Around 1960 Konrad Bayer persuades H.C. Artman to cooperate on their first Montage. They rearrange found texts into new works, from materials consisting mostly of conversational dictionaries, old textbooks, and newspapers. Unaware of the history of permutational poetry they are delighted by their invention. Together they create “lehrgeschichte für deutsche” and “11 verbarien”. Gerhard Rühm joins, and they work on “magische kavallerie”. Both Achleitner and Wiener accept the notion of montage.

The conditions for collaborative work of the whole group are set. They have big plans for a “montage about montage” (a montage about assembly) in which they reuse sections of a manual about the installation of factory machines and that should be read on the assembly floor of a factory while wearing overalls. Or the “flagello-mechanische manifest” in which an old typewriter was being pulled on a small cart through the streets while being hit with a whip. The resulting poem was to be stamped and sold to onlookers. They even call themselves *monteure*. But these projects were never realized.

In the short timespan of a few years the Wiener Gruppe made a huge leap through different genres of poetry, and mostly left behind the idea of the author as a sole genius.

If we are fixated on the idea that there needs to be a singular brilliant creator behind a work of text, we tend to argue that either the program or the human are the author. We cannot accept that the authorship of a fragment of text can be shared, not only between the machine and the human, but by multiple humans and machines in a multitude of configurations and processes. Instead of being fixated on the search for one author, we can appreciate the work in its multitudes.

4.

**Creative
Output
as a
Com-
modity**

4.1. Introduction

Using AI to produce cultural artefacts that position themselves in the landscape of the creative industries poses some hard problems.

Lev Manovich is a professor of Computer Science at the City University of New York, Graduate Center, U.S. and visiting professor at European Graduate School in Saas-Fee, Switzerland.

There is a danger that a successful system could meaningfully decrease the diversity of our cultural landscape. As Lev Manovich states in (Manovich, 2018) this danger is probably overestimated at this current point in time. Most importantly, the effect of the global cultural village is much more prominent than AI because a significant part of the population does not yet have access to AI-driven products. If there are successful AI products there will also be a market for it, and this means that there will mostly probably be a multitude of systems with differing objectives, so the chance that they will converge in the same direction is highly unlikely. Also, while a successful AI system might decrease diversity, changes to the functionality of the user interface and the user experience might actually encourage a divergent approach.

4.2. Compensation

As we saw in the chapter about datasets, contemporary AI systems need a tremendous amount of data to achieve good results. The aspect that researchers treat media on the internet as free to take was also shown. At least research papers are public and visible, so we can critique this practice, but it is presumed that the practice is also part of production efforts, where it is just not visible. This aspect is further complicated when an AI system is used for producing cultural artefacts such as music. Music and songs are a good field to look at because in 2020 OpenAI published Jukebox, a generative network that produces “music, including rudimentary singing, as raw audio in a variety of genres and artist styles” (Dhariwal et al, 2020). In the accompanying research paper they state the amount of data needed: “We scraped a new dataset of 1.2 million songs (600k of which in English), paired with the lyrics and metadata from LyricWiki” (Dhariwal et al., 2020).

It might be out of scope for a research paper to think about how to step away from the “everything on the internet is free” model to a way to compensate the creators of the items in their dataset⁴. Still, the question on how to do this is absent from this space of machine learning research, but it is crucial.

⁴ OpenAI is not a cash-strapped university department, but a well-endowed research company that has a charter. In this charter they state “by which we mean highly autonomous systems that outperform humans at most economically valuable work –benefits all of humanity”. So they should really think further than just the technology.

The way that output is constructed from training datasets invalidates the traditional ways of measuring involvement of one (singular individual) creator in the output, and the success of a produced cultural fragment. For example, in Germany the GEMA⁵ is currently appointed to distribute money that was made by playing music publicly, and similar organizations exist in most Western Countries. This task is made simple by the clear relationship of author and work, e.g. a radio station plays song A by author B and reports this to the GEMA. The GEMA now pays author B a corresponding amount of money.

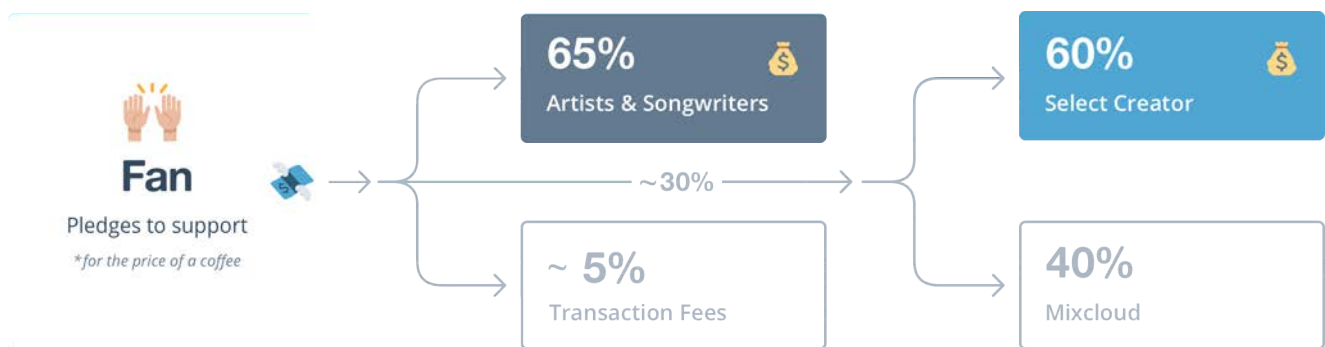


Figure 3: Mixcloud Select allocation formula. (Mixcloud, 2020)

5 See <https://www.gema.de/>

In AI systems producing cultural artefacts this relationship is severed. In a section of the OpenAI website of the Jukebox⁶ project called “Sample Explorer” one is able to listen to uncurated samples produced by their neural network. Most of the entries in this section are composed of genre artist pairs for which the system was tasked to generate a sample, e.g. “Classic Pop, in the style of Frank Sinatra”. But how is the relationship between the creator of some music in the training set and the output of Jukebox? For the given example it is clear that it learned from some songs by Frank Sinatra, but from which artists did it learn about Classic Pop?

There are two notable efforts that combine streaming and DJing, which try to distribute their revenue not only to the DJ, but also the producers of the tracks played. There is the collaboration between Apple Music and Boiler Room, an independent platform for streaming live music sessions specialized on electronic music, that wants to pay both the DJ and the producers a fee for the performances streamed. Another approach is the *Mixcloud Select* program that promises a fan-to-creator relationship. In this program the monthly paid fee is distributed between the producers, the performing artist and the platform according to an allocation formula set by Mixcloud. These solutions are piecemeal, but they hopefully show that in the market for digital goods there is a trend to diversify products by sharing success among all contributors.

⁶ See <https://jukebox.openai.com/>

4.3. Creativity

If you try to define creativity, you will see that it is rather hard, but in the end surprisingly easy if we move the goalpost the right way. Margaret Boden is a Research Professor of Cognitive Science in the Department of Informatics at the University of Sussex and the author of *The Creative Mind: Myths & Mechanisms*. “She has studied creativity, particularly in the context of computer sciences, for many years, looking into the creative aspect of computers and their similarities to the workings of the human brain”. (Miller, 2019) She provides us with three seemingly simple criteria that define creativity:

Creativity is the ability to come up with ideas that are *new*, *surprising* and *valuable*.

(Boden, 2004)

The term *New* means that the idea is not only new to the person that came up with it, but also that it had no one else before, as we know it. This idea has arisen for the first time in human history. For example in Tic-Tac-Toe because of its simplicity we can be confident that all possible combinations have been played before unlike in chess where the possibility space of all moves is finite but still extremely large. Surprising describes an idea that is unlikely. It may be related to your way of thinking but you did not expect to come up, it is against the statistics or that the idea simply seemed impossible. These two aspects of creativity are uncontroversial as most people can accept them as valid requirements for creativity.



Figure 3: Portrait of Edmond Belamy by Obvious

The third aspect, valuable, is more controversial as in the context of art and creative work it highly depends on the personal attitude of a beholder that sets the value. Together with the intricacies of written language we will see that these definitions do not really hold up when creative output is produced like a commodity.

There are also scientists like the French composer and computer scientist François Pachet who believes that creativity can be understood only in subjective terms. “Society will decide whether someone is creative or not” (Miller, 2019)

In reference to the former chapter computers don't really need to be conscious or even creative for us to experience sense in poetry. This is because there is no reading a text fragment without filling in the gaps of sense and exploring different perspectives until we find one that makes sufficient sense.

One of the most famous examples of this notion is Joseph Weizenbaum's 1965 AI program ELIZA which used a pattern matching and substitute approach to imitate a psychotherapist. It would pick out the keywords from the users sentences and following a rogerian therapist script, echo them back to the users. In 1965 the personal computer was still a thing of the future and people that encountered ELIZA had at most limited exposure to computers thus were probably rather easily enamoured by a machine. Weizenbaum was surprised that people thought that they could have a real conversation with the computer that clearly was lacking any mode of comprehension.



Figure 4: Mural in San Francisco Museum of Modern Art, 1979 by AARON




Figure 5: Aaron, with Decorative Panel, 1992 by AARON

The same problem arises also with creativity. We think of a certain activity as only possible with creativity embedded in whoever produces it. One could argue that painting a painting is clearly an act of creativity as the painter needs to subvert on how our eyes perceive our surroundings. There are two examples that invalidate this idea. The more recent and more known is the painting *Portrait of Edmond Belamy* created by a Paris-based art collective called *Obvious* with a GAN based neural network. It sold in 2018 at a Christie's auction for a record price becoming the first ever painting created by an AI to be sold at an auction.

The other example is AARON, a program that creates original artistic paintings started by Harold Cohen 1973 and in continual extension since then. Its creator wanted to answer the simple seeming question, "What is the minimum condition under which a set of marks functions as an image?" (Cohen, 1994). He already proposed an answer:

It required the spectator's belief that the marks had resulted from a purposeful human, or human-like, act. What I intended by "human-like" was that a program would need to exhibit cognitive capabilities quite like the ones we use ourselves to make and to understand images.

(Cohen, 1994)

It seems that Cohen also had the idea that, when drawing a set of marks that function as an image, the computer would need to develop cognitive capabilities and creativity just like a human mind. He wished to add cognitive functions, in this case the ability to understand what AARON was drawing, but later developed a rules based approach that would generate better results.

As we can see, the expectations that you need creativity and consciousness to generate paintings were not fulfilled. Rather, we realize that (as we saw in a former chapter) the paintings are actually not the work of the artists, but the machines created by the artists. Further we see that another aspect is missing from the definition given by Boden: these machines lack intent. They have no agency in the creation of the works simply because they are not the authors.

Following are two contemporary examples that illustrate well how this technology will be used.

4.4. Applications in the Creative Industries

OpenAI took a different approach in publishing their newest text generation model GPT-3. Instead of withholding the models like they did with GPT-2 and only releasing them later, they are building an API⁷ to try to control who uses their model for what. After the publication of the interface to GPT-3 Liam Porr used it to produce a fake blog under a fake name. He knew that GPT-3 is good at making text that reads easy and it is not good at being coherent and logical. That's why he chose to make a productivity and self-help blog. He reviewed the best performing articles in these categories on both Medium⁸ and Hacker News⁹ to construct his own headlines such as “Feeling unproductive? Maybe you should stop overthinking” (Porr and Adolos, 2020) and “We don't really understand randomness” (Porr and Adolos,

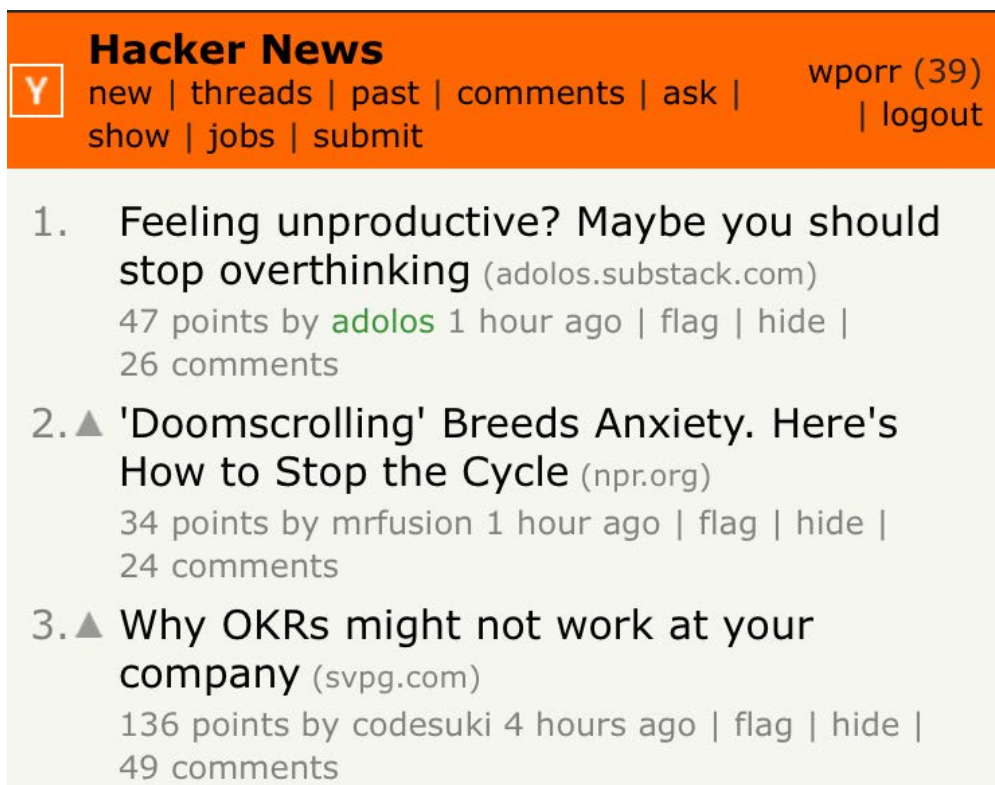
⁷ Application Programming Interface is an interface that describes the interaction surface between software.

⁸ Medium.com is an American social publishing platform that hosts content from both professional and amateur writers.

⁹ Hacker News is a social news website focusing on computer science and entrepreneurship. It is run by Y Combinator, an investment fund and startup incubator.

2020). GPT-3 would then generate the rest of the article. He posted the stories on Hacker News to see how they would be accepted by its tech affine readers. One of his stories even managed to take the coveted top spot, and the few users that suspected something were shut out of the discussion rather quickly. In the end he disclosed the use of GPT-3 on his personal blog, on the fake blog he wrote a rather cryptic acknowledgement.

If we read the articles with the knowledge that they were created by a neural network, distinct qualities (of generative text) stand out. While parts of the text make sense in isolation, in a longer form the frequent jumps, contradictions and unresolved arcs cannot be ignored.

A screenshot of the Hacker News website interface. At the top is an orange navigation bar with the text 'Hacker News' on the left and 'wporr (39) | logout' on the right. Below the navigation bar are three article listings. The first article is 'Feeling unproductive? Maybe you should stop overthinking' by 'adolos' from 'adolos.substack.com', with 47 points and 26 comments. The second article is 'Doomscrolling' Breeds Anxiety. Here's How to Stop the Cycle' by 'mrfusion' from 'npr.org', with 34 points and 24 comments. The third article is 'Why OKRs might not work at your company' by 'codesuki' from 'svpg.com', with 136 points and 49 comments. Each article listing includes a title, source, points, author, time, and action links like 'flag', 'hide', and 'comments'.

Hacker News wporr (39)
new | threads | past | comments | ask | | logout
show | jobs | submit

1. **Feeling unproductive? Maybe you should stop overthinking** (adolos.substack.com)
47 points by adolos 1 hour ago | flag | hide | 26 comments
2. ▲ **'Doomscrolling' Breeds Anxiety. Here's How to Stop the Cycle** (npr.org)
34 points by mrfusion 1 hour ago | flag | hide | 24 comments
3. ▲ **Why OKRs might not work at your company** (svpg.com)
136 points by codesuki 4 hours ago | flag | hide | 49 comments

Figure 6: Article "Feeling unproductive (...)" on top spot of Hacker News, 2020 by adolos

This would probably be attributed to sloppy writing by a human, but it still raises the bar for good enough writing that would be able to generate simple journalistic products such as listicles and other related types of content. The articles published on the fake blog also have a rather simple form, it is always a headline, a lede, a header picture and the rest of the text. By giving the surroundings of the text a bit more attention the creator could have made the impression of human writing even stronger.

The other product that I want to show is called Rosebud AI. It is a service that provides virtual models that are generated by an AI to use for virtual photo shoots.

Don't use the same stock models as your competitors!

Get started with a live demo from our self serve app. Upload any photo and swap in an AI model that uniquely markets your vision.

(Rosebud, 2020)

Their demonstration app is bare bones and does not provide you with an extensive UI to generate faces, all parameters for generation like face variants, gender, mouth forms etc are hidden behind an upgrade to subscription overlay. There are four tiers of face quality (creative, silver, gold and enterprise) that you can choose from. It is difficult to see a difference between the tiers, they all generate uncanny faces that don't fit into the picture.

¹⁰ A listicle is a type of article that is structured around a list. E.g. "14 Dog Posts From This Week That I Firmly Believe You Should See" by BuzzFeed Media (<https://www.buzzfeed.com/sydrobinson1/dog-posts-week-october-26-2020>)



Figure 7: Original photo,
2020 by Rosebud AI



Figure 8: Sample output from App 1,
2020 by Rosebud AI

Their service currently suffers a recurring problem that when you use GANs in a product, it is hard to get the neural network to generate an excellent face, and so you are left to scour the masses of generated artefacts for one that fits your idea. But, as we saw in the case of the generated blog, the images generated are surely enough to be used as stand-ins for other stock photography of people. If one just quickly glanced over such a picture, one would never simply attribute the uncanniness to haphazard production.



Figure 9: Sample output from app 2,
2020 by Rosebud AI

5. Personal Experi- ments & Work

5.1. Introduction

I thoroughly enjoy the back and forth in practice based research. Personally, it is tremendously fruitful to start with an idea and to build a prototype. This prototype can then be used to demonstrate to people where I want to go, and collectively explore where in the theoretical space my project is located. From here I usually start with reading the first batch of theory, both books and papers to understand which theoretical directions that were given to me, I want to pursue. This creates a new back and forth, from reading to building. With this process I can build new iterations of my project which I can then show to other people to gather feedback. However, I not only demo the project, I also present my newly acquired theoretical knowledge.

5.2. Previous Work

The first nucleus of the idea for this project came while finishing my project “Everything devolves into silly words” that was shown at BestOff 2018, the end of year exhibition of the Arts University Linz. “Everything devolves into silly words” was first conceived as an interactive website for a collection of net art published on a peer to peer¹¹ network with the Beaker Browser¹². It was a minimal website with a centred sentence and one button: “next”. The sentence was being generated by using a Markov chain built from the product descriptions from the most ridiculous product categories from ProductHunt.

¹¹ Peer to peer is a distributed application architecture that eschews traditional hierarchical network organization and lets each user become an equal member.

¹² Beaker Browser is an open-source web browser that allows the user to publish websites directly from their own computer.

ProductHunt is a website that is heavily inspired by the culture of Silicon Valley, where people can present their products. Those products are mostly centred around technologies that are super hip and disruptive, which leads to difficult to stomach descriptions of said products. I wanted to showcase this type of language, and my intention was to show that even if you use lofty marketing speak about a product that it would never be able to deliver, people would be distracted by the product and say that it still delivered something. To focus purely on the language and remove this distraction, the website only generated descriptions for fictional products.



Figure 10: Everything devolves into silly words,
2018 by Fabian Frei




Figure 11: Everything devolves into silly words,
2018 by Fabian Frei

For the BestOff exhibition I extended the project with physical buttons and a thermal printer. With one button, “NEXT”, the visitors could generate a new text, and with the other button, “LIKE”, they could print the description and take it home. At the same time, the installation would also transmit the text to my server to save it in the collection of discovered descriptions. The piece was received very well by the visitors. They pressed the buttons over 900 times over the course of the exhibition and left a huge pile of paper next to the installation.

Working with text and building something where visitors can get something back from was really rewarding. So the idea came up to do a bigger text generation project for my thesis.

After BestOff I was given the opportunity to do an internship with Qosmo¹³ in Tokyo, a small agency that does art and design with AI. I collaborated on an AI based vocal detection algorithm for a DJ software. I was responsible for the clean-up of incoming track listings that then went, together with audio tracks, into the training set for the algorithm.

I also took part in the one-week summer workshop held by the School for Poetic Computation at YCAM¹⁴ in Yamaguchi. SFPC is an artist run school located in New York City and holds ten week-long courses that focus on exploring the creative and expressive nature of computational approaches to art and design. The summer school at YCAM was a condensed version of their curriculum and consisted of classes such as mathematics with origami, and building logic gates in a poetic way. In the class about approaches to play I encountered the International Situationists with their theory of the *Dérive* again. My final project was a collection of my poems and pictures that I took in Japan in small leporellos that I gave away to visitors and the other students.

Before my trip to Japan I was selected to be part of the student exhibition of Interface Cultures at Ars Electronica in 2019. Thankfully I was given time during my internship to work on my installation. My plan was to train a neural network to generate text fragments from my personal emails, chats and text messages. These fragments would be superimposed in white letters on top of the projection of a slide projector. The slides were not photos but hand painted blobs of colour.

13 Report of my internship: <https://medium.com/qosmo-lab/exceptional-place-to-work-internship-report-from-austria-e45c3fb283d8>

14 The Yamaguchi Center for Arts and Media, commonly known as “YCAM”, is an art centre located in Yamaguchi City, Yamaguchi Prefecture.




Figure 11: A reactive personalized poetry machine, 2019 by Fabian Frei



Figure 12: A reactive personalized poetry machine, 2019 by Fabian Frei

I spent most of my time collecting, combing through and cleaning the training data for my network. I chose a recurrent neural network as the basis for my approach, but realized quickly that I did not have enough experience to build such an algorithm from the ground up, and also that all my digital messages were not enough for a good training set.

This leads into thinking about *Curation and Data Sets*. In the spring of 2019, shortly before travelling to Japan, I organized two evenings in my university trying to collect poetry together with other people. The premise was that I provide food and in exchange the visitors would write short poems and deposit them in a letterbox. I was mostly trying to invite and convince people to contribute a text and naturally in this process I get feedback from people. On the first evening I learned that while most people were outstandingly motivated the blank pieces of paper in front of them were overwhelming. They were not used to simply sitting down and jot down a few poetic lines of text. So for the second evening I reused the thermal printer from my project “Everything devolves into silly words”.

I wrote a list of around 20 sentences in the form of “Write a poem about ...”. Now every time somebody came and wanted to write, I would hand them a blank paper, some food and this list. The list turned out to be a great tool, some even wrote a poem for each item on it. Additionally, the poems of this second evening were thematically much closer because they were written with the inspirations as guides.

5.3. Inspirations

I incorporated this practical insight as the main feature of the mobile application, the daily theme or inspiration. It serves two important functions, firstly it covers the blank page and gives the user a canvas on which they can form their ideas. Secondly, it is the curatorial action, it tries to give the multitude of voices a bit of thematic focus. I speak of trying because it is still only an inspiration and using it only as a connection to a distant concept is natural and encouraged.

What I saw that worked well as inspirations were instructions that directed people into a situation and open-ended questions that would naturally engage the users into providing a thought out answer.

Apart from the things with screens, how do you keep contact?

Stand in the vicinity of a group of people. Instead of hearing the words, listen to the sound of their chatter. Describe it.

In contrast, polar questions that expected a text beginning with yes or no and a further explanation were often just answered shortly with “yes, I did that” or “No, I don’t do that”. It seems that this type of question often does not evoke the necessity to respond further than a “yes” or “no”.

Doesn't everyday feel a bit like Sunday?

Making sure that there is a good and well written inspiration for each day went through an interesting transformation during the project as did the project itself. In its early form the project was intended to be a light-hearted combination of *dérive* and collecting poems. For example the inspiration for the 24. February was:

If somebody asked you to draw a map of the internet, how would you do it?

But then this light-hearted endeavour came to an abrupt end with the country wide lockdown imposed during the month of March 2020. From this point until the end of the lockdown in the mid of April the inspirations I wrote were centred around the experience of this unusual measure and the change that we experienced.

Today let's do some role-playing: Pretend to be your favorite cup and write about a day in the life of that cup. Think about scale and perspective.

I started keeping a notebook for possible inspirations and even succeeded a few times to write enough sentences for a whole week at once so that I just had to pick one each day and input into the system. During the lockdown I realized that it presented two new problems for me and my project. Spending all the time with the same people in the same locations resulted in a pressing shortage of random encounters that would usually drive my imagination to write the inspirations. Staying at home also diminished the different activities that I could write about, as in the end there is a finite amount of adventures that you can undertake in your own flat.

Today walk around your apartment barefoot and touch each door frame with both hands and feet. Describe your favourite frame.

5.4. The Act of Walking

The act of walking around without a clear goal and direction as a way of exploring mundane surroundings seemed to make an enormous amount of sense to me. To stop the repetitive flow of life in our current capitalistic society, to give in to your intuition to explore a cosy side street or a hidden backyard feels natural to me. Especially my stay in Japan with the personal discovery of analogue photography renewed my desire to discover the places in their entirety and not only pass through them on the way to go on with my life. When I learned about the Situationist International and the *dérive* during the summer workshop at YCAM in Yamaguchi I was delighted as it further validated my intuition.

Debord stipulates that “The average duration of a *dérive* is one day, considered as the time between two periods of sleep” (Debord, 1956). Thankfully he further states that this duration is just a statistical average and “it is difficult for the participants to avoid setting aside an hour or two at the beginning or end of the day for taking care of banal tasks” (Debord, 1956) That means a *dérive* usually takes a few hours out of a day or simply a few moments. What seems to be important in this process is to deliberately take a chunk of time with which to stop thinking about mundane things, and concentrate on the *dérive* to create an emotional distance between ourselves and our daily routine.

The outings provided in my application can help users with that. They act as a practical introduction to the act of drifting and rediscovering the hidden layers of the mundane environment. Instead of forcing the user to both understand a new concept and come up with an interesting activity by themselves, the text I write is an act as practice.

I am not going to closely follow Debord's definition of the *dérive* in the distinction that he draws from the stroll or the *flâneur* by "awareness of psychogeographical effects" (Debord, 1956). As I understand the concept of psychogeography was introduced by Debord to lay a foundation for a theoretical framework of the *dérive*, but quickly lost its appeal and was barely mentioned in the two theoretical statements of situationism, Debord's *The Society of the Spectacle* and Raoul Vaneigem's *The Revolution of Everyday Life*.

In the book *Wander Society* (Smith, 2016) Keri Smith asks the readers to go out, act in novel ways and rediscover their vicinity similar to my project. Smith is an illustrator, artist and author that writes books that try to challenge the reader's creativity such as *Wreck This Journal*. While *Wander Society* (Smith, 2016) does not introduce new concepts or approaches in the context of *dérive* she introduces the idea to document each endeavour. Documentation is important as it aids in remembering past *dérives*. Smith prefers a visual and tactile approach. I chose a language, textual and poetic based approach because text and poetry are the main pillars of my project.

This continuous engagement with self-written poetry is training the poetic intention of each user. Everybody is capable of expressing oneself in poetic text according to H.C. Artman and his "Eight-Point-Proclamation of the Poetical Act", naturally some have more practice or a better working approach to it. Meaning that at a certain moment in time the quality of the output of different writers is diverse on a range from terrible to exceptional. With daily practice I presume that most users of my app will be able to deliver a text fragment that is good. This poem will not be a masterpiece of poetry, but it will spark joy in the writer, and they will experience this as progress, therefore it is a good poem in the context of what this project wants to achieve.

I understand that this view might be perceived as the view of a starry-eyed idealist that likes romantic gestures. I greatly enjoy building projects that eschew defining a valuable function in the context of our contemporary capitalistic society, and see my projects as digital continuation of the poetical act as defined by H.C. Artman. Naturally, he could not envision that we would carry around small personal computers that are connected to a global network all the time. Especially with this project I tried to build a digital artefact that transports the values such as “free of vanity” (Artman, 1953) or “lighthearted humility” (Artman, 1953) that Artman speaks of in his proclamation. If you produce software that is geared towards achieving virality and global scale, as is possible with the global network that is the internet, the user-oriented functionality becomes only a minor aspect. One could even argue that if that software is free, it only functions as a stepping stone towards a financialization of its virality and scale. Instead, I tried to introduce my app in personal chats where I tried to transport the values and emotional experiences embedded in my project to people I was speaking to. As a netizen¹⁵ I also carried the app into digital communities that I am part of, but still on the scale of the whole internet they are tiny.

This also extends to the software running on the server, which sends your smartphone the inspirations and saves the poems submitted. Each user is identified only by an ephemeral character list that is generated the first time the application is opened. In the privacy policy (Frei, 2020) it is explained that a new list can be generated by removing and reinstalling the app, a process that allows users to leave their old identity and with it their old poems. I hope that this encourages fearless writing and a relaxed usage of my app.

¹⁵ A netizen is a citizen of the internet

Until now, I have talked of the immaterial properties of the software but I want to extend the concept of the poetical act also to the user facing side. At the beginning of this project I had the urgent desire to build a minimal, sleek and quiet piece of software that has a clearly communicated purpose and actually sticks to it and on multiple levels expressly ask for consent. This means that it will only include vetted code and no third party code that can be categorized as trackers or ads. Again this attitude can be seen as hopelessly romantic in our contemporary world, but in the context of the poetical act it is totally valid as it has different objectives than the software that we use every day.

To reiterate, I ask for consent and clearly communicate the intended usages of content generated on the website, in the privacy policy and on the welcome screen (the message a user sees when he opens the app for the first time). In this way I make sure that each participant is able to make an informed decision. This extends into how the application communicates its needs to the user outside itself. It sends one notification per day to ask the user to write a poem and if the user does not react to these messages for a few days it ceases to send them. It forgoes extracting the maximum amount of content from the users to instead try to respect the needs and wants of the users. The danger of commodification of the poems collected in this project is not existing and the way I want to utilize the output of this product is to produce zines for the people who contribute some text.

The software acts like a lens that concentrates multiple rays of light into a small spot. It combines the output of different authors and the training data set into one new fragment of text. In this process it adds one more layer to the tangle of cultural links that Barthes talked about in chapter “Poetic Intention and collective authorship”. It should not be understood as an additional author but more than a combinator.

6.

Setup of Artwork

6.1. Introduction

Development of the supporting computing infrastructure was built around the client-server architecture and two main goals: development velocity and minimal maintenance load.

Developmental velocity means how much time I have to spend going from concept to actual working implementation. It is important because the application is only part of the whole project, and the application has to stay small and lean so that I can quickly iterate on it. This allows me to rapidly introduce changes based on my learning from the practice based research I was doing at the same time.

Minimal maintenance load means how much time I have to spend to ensure the normal working of all parts of the system. This aspect is important because the time spent on upkeep is most probably better spent iterating on the core concepts or working on the theoretical part of my thesis.

The development of this system happened in lock step with the theoretical part, as it is normal in practice based research.

Most of the technical choices are based on experience and personal preference.

6.2. The Smartphone Application

To be open to the biggest amount of users the application has to run on the two main smartphone operating systems, Android and iOS. There are a wealth of cross-platform development frameworks based around bringing JavaScript onto smartphones like React Native, Apache Cordova and Ionic Framework. I have used such frameworks in the past to great personal displeasure. The abstractions they provided are leaky and quickly the developer has to interact on a deep level with the underlying platform which defeats the purpose of these frameworks.

Thus, even the best designed and most usable app will not receive much praise if it is not responsive to the users input. While there was a massive investment from big wealthy companies like Google and Microsoft into the performance of JavaScript virtual machines, they are still not as fast as compiled native code. Also, the integration of JavaScript into a different environment adds overhead because the native application needs to communicate with a JavaScript virtual machine.

Another complication is that the approval of such applications through the strict Apple AppStore review process is thorny. Sometimes a minor change in the App Store Review Guidelines after the initial approval means that the applications will be removed and any future updates will be difficult to get approved.

A while ago the Flutter framework came to my attention. It is built on the Dart programming language. Both Flutter and Dart are developed at Google. It compiles to native code for both platforms which means high performance and has tooling like hot code reload, that shortens developer iterations.

I went with Flutter, took some available widgets from the library and customized them to fit my use case. Development went together with the backend and was

uneventful but still two points stood out:

I still needed both an Apple computer and also an Apple smartphone to develop and test the application. It is not possible to submit some software for review and inclusion to the AppStore without testing it on an iPhone. Luckily I managed to get my hands on a second hand MacBook and an older iPhone 5.

Navigating the submission process on both platforms is a nightmarish experience and takes a large amount of time. Within Google Play the quality and intuitiveness of the UI is terrible and feels directly modelled on the layout of the modular code to keep developer satisfaction high. Further the absence of any human being and the insistence on automatic testing is palpable. The outcome of this is that you go around and push buttons in certain sequences to appease a non-communicative/mute machine god.

iTunes Connect feels more polished and intuitive, it still contains a lot of hidden complexity and unexplained terms. The submission process is much stricter and there seem to be real people reviewing submission based on the Apple guidelines. With my first submission my application was rejected on the basis of rule 4.2 “Minimum Functionality”¹⁶ because indisputably it does not contain many features apart from reading an inspiration and writing a poem. Through a thoughtful message to the person or people reviewing my submission I managed to convince them that the minimalism of my app is actually a feature and that most of the interaction is going to happen in the users mind but still I never got a reply to this message. The only visible reaction to my message was the approval of my submission. Since then, I always try to include a short message on the changes that I have done in the current update and how to test them.

16 See <https://developer.apple.com/app-store/review/guidelines/#minimum-functionality>

6.3. The Backend

The backend is the server in the client server model. It is a restful API built on top of Elixir and PostgreSQL. Elixir is a functional general purpose programming language that compiles to Erlang bytecode and runs on the BEAM virtual machine. It is built on the Actor model which enables a high degree of concurrency, resilience and fault tolerance. The syntax of Elixir is approachable and easier to learn than Erlang's syntax from which Elixir descends. PostgreSQL, also known as Postgres, is a free and open-source relational database.

These qualities are complemented with excellent tooling that covers most problems that surround writing code like deploying your project on a server, updating that project while running etc.

To compare this with Python. With it, writing code and building a project is fairly fast and uncomplicated, deploying and maintaining such a project quickly becomes difficult and time-consuming. With Elixir building up may take slightly more time but that is compensated by spending less time on maintenance.

My server that costs me €25 per month has served over 15'000 requests since the start of the project with unplanned zero downtime.

An interesting point about the backend is that it tries to collect the minimal amount of data about the users while being transparent about it. I wrote a privacy policy¹⁷ that details which type of data is collected like “Masked IP-Address for privacy” (Frei, 2020), how long it will be kept and current usage and also further planned usages.

A positive outcome of this privacy oriented approach is that the application can be used without a login. The user can obtain a new unique identifier by simply deleting and reinstalling the application.

¹⁷ See https://tofu.wtf/poems/privacy_policy.html

6.4. Augmenting the Written Poems

During the early stages of the project I had not decided on how to generate poems from my dataset, so I was unsure if I would collect enough poems over the course of the experiment. To be on the safe side, I decided to extend my dataset by generating synthetic text fragments from the existing poems. To extend a dataset with fabricated entries is called dataset augmentation, and it is a commonly used approach in machine learning and AI research, especially in the field of computer vision.

An example is the approach taken in Multi-column Deep Neural Networks for Image Classification (Cireşan et al., 2012). They trained a deep neural network to recognize handwritten digits from the MNIST dataset. MNIST is a large dataset that contains 60,000 training images and 10,000 testing images of handwritten digits. It is commonly used to compare the performance of character recognition systems. In (Cireşan et al., 2012) they managed to achieve a “0.23% error rate, improving state-of-the-art by at least 34%” (Cireşan et al., 2012). For the first time an artificial method came close to the $\approx 0.2\%$ error rate of humans.

They trained their column of deep neural networks around 800 times. Each time they would distort the digits in the training set randomly as seen in Figure 14.

With rather minimal effort they extended the original training set that they had available with random variations of the already known and correctly labelled digits. Instead of adding more images to the dataset, that they would then have to prepare and label, they generated additional synthetic training data from the existing data.

In the space of machine learning based computer vision algorithms this practice has been built into the tools because of the ease and security of application. The basic image manipulations operations that we know from Photoshop or gimp such as cropping, flipping, zooming, rotation, noise injection, and many others are applied to images by software before they are fed into neural networks for training. There is no need to manually prepare the images before training.



Figure 13: Example of MNIST digits, 2017, by Josef Steppan on Wikimedia Commons




Figure 14: Randomly distorted MNIST digits, 2012, from (Cireşan et al., 2012)

Let's say we are building a classifier that detects if there is a horse in a picture. On the right you see the original picture, on the left the same picture with modifications applied. Undeniably these pictures still show a horse. So the meaning of the pictures has not changed even though we applied rather stark modifications to it.

If the training data set contains text instead of images applying similar modifications to create synthetic content becomes quite difficult. The semantics of a sentence can change drastically by just a small exchange of words. "There is a horse in my barn." and "There is a tiger in my barn." are structurally the same other but semantically they are totally different



Figure 15: Horse, 2014,
by Karsun Designs on Flickr



Figure 16: Cropped horse, 2020,
by Fabian Frei



Figure 17: Rotated horse, 2020,
by Fabian Frei



Figure 18: Noise horse, 2020,
by Fabian Frei

There are approaches such as backtranslation, where you use the translation feature of a search engine such as Google, Bing or Yahoo to translate a sentence to another language and back into the target language. There are also approaches such as inserting or deleting a word or swapping words in a sentence. While these approaches make sense for normal text, applying them to poems is difficult. Poems are fragile constructs in which emotions and semantics are in a delicate equilibrium that is easily disturbed.

That is why for my project I have identified three approaches that make the most sense in the context of poems. Instead of simply augmenting and feeding the text fragments directly into the training set, I wanted to build on the community around my app. For this I added a screen on which users before writing a poem could rate the poems generated by my augmentation algorithms. Then only the poems that were approved by users would find their way into the training set further strengthening the influence the users have on the output of the poetic AI.

The three generators were implemented as modules of the backend using the infrastructure that was already existing. In the backend overview I added buttons to generate batches of 100 poems that would get saved in the database as “not approved”. The smartphone application would later request a batch of ten poems to present them to the user. Users would then proceed to review the ten texts, and either approve or reject them. Furthermore, I added an endpoint to the backend to download the approved poems as a dataset, in order to feed it into the machine learning algorithm.

6.4.1 Markov Chain

The first generator is based on a mathematical technique called the Markov chain. “If you think of a text as a sequence of words, you can use a Markov chain to predict the next word of a text based on the previous word.” (Richardson, 2014) In a first step the generator extracts word frequencies from the dataset of written poems. The poems that the generator uses are augmented with start and end markers, special states in the set of word frequencies that tell the generator where to start and end the generated poem. The generator would begin with the start marker and recursively feed the next word into itself until it either reaches a fixed length or the end marker.

Why am besten nicht allein

We love

something can's a strong and repeat, with

feel like the market here and different city, repeat, but we can's only plant is clearing everything that life

pushing down

heralded by day preparing the border.

is just feel you won's boss

Is my mind

is inside,

warm sun.

While the output of the Markov chain generator is at times inspiring and poetic it also generates a big amount of gibberish that seems to make sense, but even with a lot of work from the reader still floats in this weird space between total gibberish and almost making sense. The ratio of accepted to rejected poems is around 1 to 10.

Die

it't tell mighty george

Warm beer.

of stars,

fly though and they dissipate. I don't know

You hold my breath, you hug my dreams.

my

A masturbating driver

But as often as I can

cats are fine

hey

Hahah was, i'd just a metal can make sure, love and back.

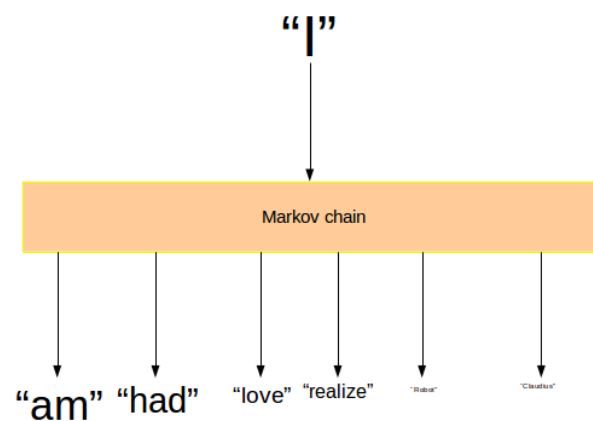


Figure 19: Markov chain visualization,
from (Richardson, 2014)

6.4.2 Queneau

The second generator is based on the work by the French writer Raymond Queneau, “who as far as I know was the first person to apply it to text.” (Richardson, 2014) He was one of the founders of the Oulipo circle in 1960. It was a gathering of mainly French-speaking writers and mathematicians that sought to create constraints for their writing or as Queneau himself said: “Oulipians are rats who build the labyrinth from which they will try to escape.” (de la Torre, 2005)

Two important examples are the sonnets *Cent Mille Millions de poèmes* by Queneau and the novel *La Disparition* by George Perec. In *Cent Mille Millions de poèmes* Queneau constructed a book containing ten sonnets that could be recombined in one hundred thousand billion ways. This could be done by simply swapping lines between different sonnets. “Astoundingly, Queneau ensured that no possible combination of lines would break the rhyme scheme and that any given resulting sonnet would be grammatically correct” (de la Torre, 2005). Perec instead wrote a whole novel without using the letter *E*.

Son pouls battait trop fort. Il avait chaud. Il ouvrit son vasistas, scruta la nuit.
Il faisait doux. Un bruit indistinct montait du faubourg.
Un carillon, plus lourd qu'un glas, plus sourd qu'un tocsin, plus profond qu'un
bourdon, non loin, sonna trois coups. Du canal Saint-Martin, un clapotis
plaintif signalait un chaland qui passait.
Sur l'abattant du vasistas, un animal au thorax indigo, à l'aiguillon safran, ni
un cafard, ni un charançon, mais plutôt un artisan, s'avavançait, traînant un brin
d'alfa. Il s'approcha, voulant l'aplatir d'un coup vif, mais l'animal prit son vol,
disparaissant dans la nuit avant qu'il ait pu l'assaillir.
(Perec, 1969)

During the time of Oulipo there were no computers or digital processes that could have helped them to produce text within these constraints. Amazingly, some constraints have a very computational approach and are easy to implement as software. Queneaus approach is a prime example for this compatibility.

But tonight
and I got stung cheek.
and learning to love
calm and wise,
With one exception since the sausage has two

But baking powder's still on the rise.
Rhyming like there's no tomorrow
You're yes then you're no
ich die letzten Tage waren es nur noch eine gute Nacht mein Engel und das
mit der neuen Wohnung und das war eine der wichtigsten Dinge

With as much effort
opening doors
and my hood.
everything goes silent
anaesthesia strong
Tiles
The norm has been calibrated
is made for me
Sometimes they are really foul
Open the curtains

The implementation of the generator is rather straightforward. It starts by splitting the poems from the dataset into lines. Then it chooses the first line of a random poem and continues with adding lines from other poems until it reaches a stochastically decided maximum length of lines.

Golden.

and the ocean,

I want to interact, chat and talk with people I meet by chance.

and ducklings

Inkább vágý léha,

slower and faster,

Moving ever so slow

'Cause you're hot then you're cold

Get to let go the tension

from a cappuccino

seem to enjoy the total stillstand

eyes on their children,

Tiles

all around one big mess

Less planes

wine at 9 AM, but we're apart.

Heading in no direction, is also a direction.

Take your cup
shovelfulls of
De marad a régi ház
who I am slowly knowing
Yet our feathered friends
on my way
so I remember fine
in my own mind.
the infinite stars
no planes in the sky

6.4.3 Synonyms

The third generator is similar to one of the swap or insert approaches that I talked about. Instead of swapping or inserting random words I pick a stochastically chosen word and exchange it with one of its synonyms from an online thesaurus. With this plan I wanted to see how elastic the semantic structure of a poem is. The user in charge of approving or rejecting a poem that was generated by this module would also never have seen it and thus have no way to judge the meaning of the new poem in relation with the original poem.

I make gart bucks

Here and now i am so hoping for a functional rhyming

I care (for) this app a lot

The implementation is also rather straight forward. It does no preprocessing of the poems but instead takes one fragment and chooses a random word to swap out. Then it asks the Merriam-Webster Collegiate Thesaurus API¹⁸ for a list of synonyms for the selected word. If the thesaurus returns a synonym the module swaps it into the poem and saves it into the database.

I'm a sexy comedian

In these strange times, parks are closed,
people locked up and the beasties follow each other, freely.

¹⁸ See <https://www.dictionaryapi.com/products/api-collegiate-thesaurus>

What kind of cake did you have? someone asked me out for cake for my birthday and found a restaurant with carrot cake. i was sad to be far from family and friends, but the cake was good. my favorite cousin's wedding was supposed to be coming weekend, but it will be next year.

7.

Training & Generation

7.1. Introduction

At the time of writing, the deep neural network architecture most used in natural language processing is called Transformer. Like recurrent neural networks they are able to process sequential content such as language or time series. Unlike RNN they don't need to learn the sequential content in the correct order. This allows for better parallelization of training which in turn allows using bigger training data sets. The result of this optimization are enormous projects like BERT by Google, or GPT by OpenAI. For us as users another important aspect comes into play: Instead of training a Transformer model from zero we can reuse a pretrained model and fine-tune it to the problem at hand. This eliminates a source of errors that could happen during the setup of training and reduces the amount of computing resources we need. For example, the larger GPT-2 models were trained on 256 cloud TPU v3 cores (Huggingface, 2019) while for fine-tuning and inference a single TPU core is sufficient, which is thankfully usually included in the free tier of AI services of different cloud providers. This enables us to start with a strong working network and further progress it into the right direction without spending an enormous amount of money, or needing to enter a sponsoring relationship with a financially capable entity, thus preserving our freedom to act in a way that we want.

7.2. GPT-2

While evaluating the different Transformer models I personally came to the conclusion that the GPT-2 projects were most approachable. I encountered a prepared Colaboratory notebook that allowed me to fine-tune and generate text on one of the free Google Compute Engine nodes. To start with the process of fine-tuning and generation I extended my system with an endpoint to generate a dataset from the submitted poems. Now in this context a dataset is a text file containing all the text fragments from my database including special makers for the start and end of each fragment.

```
<|startoftext|>A walk at night  
a lonely experience  
refreshes me  
from daily life.<|endoftext|>
```

```
<|startoftext|>when blue turns dark grey  
and i'm allowed to hit the hay<|endoftext|>
```

These markers are present both in the training set and in the output. The model is able to generate multiple poems in one pass to make the selection of worthwhile poems easier, and the markers are used to separate the poems in the output.

In 2019 GPT-2 was unveiled by OpenAI. It was the biggest neural network yet, with 1.5 billion parameters, in 2020 during my work on this thesis and the data-poem project, GPT-3 (a 116× bigger network with 175 billion parameters) was unveiled. Naturally, it would have made sense to use GPT-3 as it produces more natural looking output. At the same time, it can only be accessed via a special API. To gain access to the API, you have to sign up with OpenAI and be approved by them.

OpenAI claims that they do this to monitor the usage of the model and to understand possible vectors of misuse. While this is a possible interpretation of the situation I suspect that the model is simply too big to run on commodity hardware and that they want to keep tabs on a possible commercialization of the model. Despite their rosy language about open- and fairness OpenAI is still a commercial entity that at one point should make some money. From my point of view (as an exhibiting artist) it makes sense to have a portable model that runs on commodity (ie, commercially available) hardware that you can transport and set up in an exhibition space, which does not need an internet connection and does not rely on an external service to function.

As of September 2020 GPT-3, most probably because of its sheer size, is still missing the ability to fine tune the model further than single prompts allow. Gwern Branwen in (Branwen, 2020) his creative exploration of the abilities sees the designing of prompts for GPT-3 almost as a new paradigm. He is surprised and, as far that we can tell from his writing, also a bit startled that the model is able to write an “uncanny simulacrum of myself” (Branwen, 2020) from the prompt “A new essay by Gwern Branwen (gwern.net):” (Branwen, 2020). He writes that this works because GPT-3 was trained on a corpus that included most of the public internet and thus has all the knowledge to generate these kinds of texts. While GPT-3 would certainly be able to generate good poetry, the fine-tuning with content that the GPT3 model has not seen is a different aspect.

Using the Colaboratory and the free tier of Google Compute I was able to import my dataset and fine tune the medium-sized model called 355M. Both the 774M and the 1558M model can be used to generate text, but cannot be trained further because of their size.

7.3. Generation

I extracted two datasets from my system, one containing all the submitted poems and one containing the generated poems from the augmentation process. I merged the two files into a combined dataset so that I have a combined set and a pure set. These sets are then used to fine tune the model and generate a few thousand poems that get saved in text files for evaluation by myself.

The whole process of fine-tuning and generating was not as difficult and time-consuming as expected, and I spent almost no time on fighting the software to get it to work. Overall I clearly spent a magnitude more time coming up with the project, building and setting up the support infrastructure such as the apps, writing inspirations and convincing people to submit poems than I spent training and fine-tuning the model. Which clearly follows the argument that I'm trying to make with this text.

Evaluating the difference between the output of the augmented dataset and the pure set are difficult to quantify. In general the output looks almost the same, in both outputs the model sometimes degenerates into repetitive loops of the same or similar words. This is a behaviour that Branwen has seen in all the GPT-2 models but also in the new GPT-3 model in (Branwen, 2020). Apart from these outtakes the poems are quite good, and it seems that the network was able to extract some structure of the datasets that I fed in. It seems that it would fill this structure with words that it learned in another context to generate new fragments. I could neither find those new fragments in my dataset, nor on the internet in general. In both models there were some cases of overfitting, where the model simply repeats a poem directly from the training set.

The following poems were generated by the model which had been fine-tuned on the pure data set. You can see that the model took the layout of the text – short, punchy lines – that were more or less dictated by the UI of the application and filled it with other words.

I was hoping
for a hint of
ethereal finality
or for the faint
morphing of
a pattern,
a sense of place.

I was disappointed.

I want to be like my husband
Who takes pride in his work ethic
And lets his work speak for itself
Without any strings, buttons, or advertisements

No home, just you and the world around you

I am so lucky to have found you
in this crazy place
life throws everything
at us quickly
but through the tough times
we make new friends
like these benches

I want to be an astronaut
Like Neil Armstrong, I hope to live up to his legendary
vision, and go to the red planet.

The following poems were generated by the model fine-tuned on the combined dataset.

The software runs itself
Everything is run, no one can stop it
The software is a machine, what a paradox
Everything is running, no one can stop it
The software is a computer, what a contradiction
Everything is running, everything is ok
The software is a computer, what's the big deal?
Everything is running, let's see what happens next.

My days are long gone
I wish I could take them back
I wish I could go back in time
I wish I could go back in time
I wish I could take your misery

the music made me feel
relief knowing that
all was well
and that my
enemies were
weak and out of ammo

I took a wrong turn into a funeral
I had forgotten how much it hurt

8.

**Related
Artists &
Artistic
Work**

8.1. Mario Klingemann

Mario Klingemann is a German artist that goes by the name of Quasimondo who lives in Munich. In the 90s he was a copywriter and designed Techno flyers and posters. A self-declared “skeptic with a curious mind” (Klingemann, 2020) today he works with code and artificial intelligence. From 2016 until 2019 he was artist in residence at the Google Arts & Culture Lab. He received an honorary mention at the Prix Ars Electronica 2020, the Artistic Award 2016 by the British Library Labs and won the Lumen Prize Gold 2018. He is very active on Twitter where he lets other users take part in the process of his current artistic exploration.

His work, “Appropriate Response”, uses the same model, GPT-2, as my project does. It consists of a 125 character split-flap display that is connected to a computer running the GPT-2 model. In front of the display is a wooden kneeler similar to something used in a church setting, like in a confessional box, giving the installation a religious flair. Klingemann uses the strength of the model to generate short pithy phrases that are difficult to distinguish from texts written by a human. We encounter such phrases continuously in the form of marketing slogans, religious guidance, conspiracy theory catchphrases or the all encompassing self-advice of influencers. Like the project presented in this thesis, he exploits the fact that reading is constructing sense and untangling the cultural links in a piece of text, as seen in the chapter *Poetic Intention and Collective Authorship*. Additionally, he plays with the contrast of superficiality of the phrases and the solemn construction of the installation.

Sadly he does not talk about where he got the additional 60,000 phrases from and what they are, something that is important to my approach. His contribution lies in the hidden technical work of collecting and fine-tuning the model and the visible part of the installation, the split-flap display.

While the display and the kneeler look great and well-made on pictures, the whole project is in the end a glorified tech demo of GPT-2 that does not talk about the important part of an AI artwork: the curation and collection of the data set.

8.2. Rafael Pérez y Pérez

Rafael Pérez y Pérez is a Mexican scientist and artist. He completed an undergraduate degree in Mexico, studying electronics and computers. He went on to complete a PhD in AI at the University of Sussex. Initially he wanted to work on AI and music but lacked sufficient musical knowledge. He built a plot generation system called *MEXICA*, named after the indigenous people of the region of Mexico City. The system generates plots that are inspired by the folklore of the Mexicas, including characters such as the jaguar knights, famous warriors of the Mexica culture.

To prepare the *MEXICA* algorithm for use, it has to be fed with a set of short stories, a list of all possible actions that each character can carry out, and a list of conditions for each action. These conditions are encoded in the form of “emotional link, tensions, and conflicts between the characters. This is a unique part of my work. It mirrors the way we come to understand the world.” (Miller, 2019)

Then an initial action is entered into the algorithm and it searches through its recollection, built from the short stories that it was fed, for the proper action to be performed given the new emotional context. This action taken then gets re entered into the algorithm, building up a plot. Naturally a plot is not a story, so Pérez y Pérez joined forces with Nick Monfort's system *nn Narrator*. *MEXICA* provides the plot, characters and events and *nn Narrator* narrates a story from them.

The jaguar knight will oppose the hunter!

Meanwhile, the hunter will oppose the jaguar knight.

Then, the jaguar knight will feint at the hunter, dude!

Then, the jaguar knight will strike himself!

This generated story is rather bare and shows the problems when two systems collaborate to go automatically from plot to narration. Perez and Monfort added more systems, and switched to a blackboard system where writing and editing are a continuous process.

In 2017 Pérez y Pérez published a book containing twenty short stories by an extended version of *MEXICA* called *MEXICA: 20 Years-20 Stories (20 AÑOS-20 HISTORIAS)*. The writing of these stories is notably smoother but still they are still plots.

The jaguar knight made a potion and drank it quickly. He started to recover.

In that moment the warrior was not able to understand the jaguar knight's conduct.

The jaguar knight was confused and was not sure if what he had done was right.

Hurriedly, the warrior ran off to the Popocatepetl volcano.

The end.

It seems that for Pérez y Pérez the output of his algorithms is a by-product, but the main focus is using them as a vehicle for research about the nature of human creativity. “My stories don’t read like Gabriel Garcia Marquez. So what? Right now we are trying to understand this process better. You learn by degrees. [...] As a scientist I want us to understand better the amazing thing we call creativity.” (Miller, 2019)

While I sympathize with using his algorithm to explore the field of computational creativity, the text fragments generated are too simple to be seen as something creative that simulates the imagination. I suspect that exposure to people outside of his field, more artistic fields, would have been a great source of inspiration and ideas. In other fields there is a certain focus on the output and this focus could have helped to propel *MEXCIAs* poems to better quality.

8.3. Nick Montfort

Nick Montfort is a poet and professor of digital media at MIT. After studying for a bachelor degree in both liberal arts and computer science, he got a master's degree in media arts and science and creative writing. After his PhD in computer and information science he was an assistant professor and later became a full professor of digital media at MIT.

In 2014 Montfort published a book of poems called *#!*, which is pronounced *shebang*. In Unix and Linux computing the shebang is usually the first line in an executable script file called the interpreter directive. It tells the program loader which interpreter it should use together with the present file.

The book contains the poem *Round*, which characterizes Montforts approach to authoring texts with software incredibly well.

Form intends intense verse crease to tense form tense vent verse tone verse
form crease form event tends to crease to tends form form vent form crease
tone verse tense

crease vent vent tends inverse tone into verse form verse verse form tone
tense in
tense vent crease

verse tone tends tense tends tense verse crease form

tone vent into tends
to crease vent to crease

(Montfort, 2014)

As you can see there is a limited set of words used and it seems to hint at a pattern that does not fully reveal itself. That feeling is correct because Montfort chose 9 words to represent the numbers one to nine while zero represents a line break and used the digits of π (PI) as a quasi endless stream of words to form his poem. He relied on the special property of π , it cannot be written as a ratio of two integers and thus never ends and never forms a pattern. Even though π is infinite in its extent, it is also deterministic and can be expressed in a formula which means that it can be calculated by software. This idea is readily comprehensible and is simply to express in code, leaving ample time to work with the words. A point that Montfort emphasizes: “The work of art is the computer program. I select the words. They’re mine” (Miller, 2019)

Unlike the other artists that I talk about, my project and most of the computational creativity community that depend on models using extensive mathematics, complex rule sets and enormous databanks, Montfort is foremost a writer that uses software as his tool to develop and extend his creative reach. “Round,” Montfort says, “is a visualization, a textualization, of the digits of pi in which pi’s digits are mapped to word strings.” (Miller, 2019) In his book *!#*, each poem is accompanied by the listing of the program that created it, harking back to the time before the internet. During that time software magazines would print code alongside with some example results that the readers could type out on their own computer. Readers could in this way get in possession of new software but would also have an insight into how the code is structured and how it works.

“He prefers his programs to have the fewest possible lines – that is, to be as simple as possible” (Miller, 2019) as he hopes to invite writers that have no or negligible experience with writing computer code to try out, change, remix and incorporate his works in their output. This sharing approach cannot be as directly pursued if your project is built upon an intricate architecture of different digital systems and procedures working together that need special domain knowledge to set up and run.

Unlike the project presented in this thesis, which uses an AI model to further complicate the tangle of cultural links and blur the input of each author of a text fragment, Montfort found a way to express his approach to computationally created poetry in short and succinct pieces of computer code. These pieces are made public alongside their output to invite others to participate and play with them.

I find this approach refreshing after working on the complex system that is used in the current project. I also see it as an appropriate continuation of ideas of the Wiener Gruppe, oulipo and similar into the contemporary digital culture of computational creativity. Instead of trying to teach computers about the nebulous concept of human creativity so that at one hypothetical point in the future they can write great poems, he plays to the strengths of digital machines and uses them as patient number crunchers that enhance his current creativity in new ways.

8.4. Allison Parrish

“I’m an experimental computer poet” (Miller, 2019) states Allison Parrish. Programming since a young age, she studied computer science and continued with linguistics and French. She completed a master’s degree from NYU’s Interactive Telecommunications Program (ITP), where she is also a professor now. ITP is part of New York University Tisch School of the Arts, and calls itself the art school for engineers and the engineering school for artists. She was also named “Best Maker of Poetry Bots” by the Village Voice in 2016.

In addition to being a poet, she also acts as a computer programmer and game designer. Similarly to Monfort, whom she has collaborated with, she sees code as a tool that helps poets express their creativity in a different novel way. From her point of view the strength of computers is that they do not seek any sense in the text they produce. Without tiring, they explore semantic space to present us with fresh combinations of words.




Figure 20: Semantic space visualization, from (Parrish, 2015)

To explore the semantic and lexical space she worked with the n-gram data from the Google Book corpus. An n-gram is a fixed sequence of words of the length n. Two words is a bigram, three words is a trigram, and so on. She assembled a big file containing bigrams that look like about, anything, 124451 meaning that anything followed about 124451 times in the database. To keep the data set in a manageable size she included on bigrams that where both words start with the letter a. To visualize this two-dimensional space of bigrams she laid them out in a matrix beginning with a. If there are a lot of occurrences the matrix contains a larger square at the intersection and no square if there are no occurrences. This leads to voids of unexplored lexical space. She picked some bigrams that do not exist in the corpus:

amusedly abstract
 abreast annihilates
 adding alteration
 (Parrish, 2015)

	a	ability	able	abort	about	above	absence	accident	according	account
a	97,569	124	224	705	1,666	1,928	110	301	1,335	1,353
ability	143	0	0	0	105	699	0	0	499	0
able	10,759	0	0	0	0	0	0	0	184	913
abort	437	0	0	0	0	0	0	0	0	0
about	3,520,097	336	341	0	492	1,456	215	0	2,058	0
above	299,589	0	0	0	12,526	0	0	128	785	22,950
absence	436	0	0	0	130	0	0	0	0	0
accident	1,520	0	0	0	365	0	0	0	106	0
according	0	0	0	0	0	0	0	0	0	0
account	11,943	0	104	0	2,122	2,109	0	0	1,518	0

Figure 21: Ngram table, from (Parrish, 2015)




Figure 21: Voids in semantic space, from (Parrish, 2015)

She repeated the same process for trigrams, resulting in a three-dimensional matrix that was visualized with cubes “like a scene from a weird space movie” (Miller, 2019)

Her best know Twitter bot is called *everyword*¹⁸. It tweeted every word of the English language in alphabetical order. It started in 2007 and, tweeting a word every thirty minutes, took until 2014. The final word was *étui*, “I’m not going to give the presentation today about why it ended with “*étui*” instead of “*zyzzyva*” or “*zyxt*” or whatever.” (Parrish, 2015). At the height of its popularity it had over 100’000 followers

18 See <https://twitter.com/everyword>

19 As of 8.10.2020: 61.4k followers

and it got made into a book¹⁹.

Another Twitter bot is called *PowerVocabTweet*²⁰, a bot that promises to “Boost your vocabulary with these fiercely plausible words and definitions” (Parrish, 2015). It generates a new word out of two existing words that fit together based on their letters and syllables. For these words the bot comes up with descriptions using a Markov model based on word definitions from WordNet. It generates word definitions pairs such as the following:

Myriller, prop. n. whiting of the east end of London²¹

Schus, prop. n. English astronomer remembered for his heroic deeds (1779-1820)²²

haxtern, n. a relatively low softening point²³

Naturally Parrish is not the only person and organization that makes use of Twitter bots. The NASA, National Aeronautics and Space Administration that is responsible for the civilian space program, has tolerated the unofficial account for the *Voyager 2* probe. *Voyager 2*, like its twin *Voyager 1*, were launched into space in 1977 to study the outer planets. On November 5, 2018, at a distance of *Voyager 2* left

20 See <https://twitter.com/PowerVocabTweet>

21 <https://twitter.com/PowerVocabTweet/status/1122546024017485824>

22 <https://twitter.com/PowerVocabTweet/status/1122304430827819010>

23 <https://twitter.com/PowerVocabTweet/status/1122425224824279040>

the heliosphere, and entered the interstellar medium, a space outside the influence of our solar system. It mostly tweets about its velocity and instrument calibrations.

Now Parrish's Twitter bot *@the_ephemerides*²⁴ confronts images taken by the NASA's space probes with text fragments generated randomly from two large bodies of text, one is a corpus of astrology, one is a corpus of geology. This creates an immediate connection between the text and the pictures. So even the texts are randomly generated, this aspect makes it easier for the reader to see sense in them.

The cold experienced
by the simple
natures



Figure 22: Space object, from *@the_ephemerides*

24 See https://twitter.com/the_ephemerides

of these substances
is terribly intense.

The works by Parrish presented here all are united in having a strong personal influence. “I always seize authorship for myself. When I put out a book of poems it’s by Allison Parrish, not Allison Parrish and a poetry bot.” (Miller, 2019) In creating her works she empathizes the decisions that she takes, for example about the range of words and the algorithm used. “These are all incredibly expressive decisions” (Miller, 2019). She disagrees with the goal of computational creativity, the goal to construct a machine that is creative in a human way and creates poetry. Instead of imitating existing poetry, she seeks new ways of poetry.

“I’m going to be a hardline and say that computers cannot be creative.

Furthermore, it will always be a mistake to attribute volition to the computer and not to the people who programmed it because attribution of volition is removing personal responsibility: the algorithm did it, not me” (Miller, 2019)

8.5. Ross Goodwin

Ross Goodwin engaged heavily with computers as a child but went on to study economics. After graduating he became a speechwriter for President Obama, transitioning to freelance writing jobs afterwards. One of his tasks was the production of business letters, which are all rather similar. He exploited their suitability to computer based generation by writing a software that constructed these letters from rotating paragraphs and lines according to a template. Intrigued by his own work he discovered that others were already working on similar problems. Driven by his interest in works at the intersection of computing and writing he joined the NYU’s Interactive Telecommunications Program (ITP).

During his studies he created a project called word.camera. The camera in question generates text fragments with a poetic quality based on what is held in front of its lens. He used a recurrent neural network specialized in generating



Figure 23: Saint Jerome in His Study , 1514, by Albrecht Dürer

sequences of characters based on an input corpus called CharRNN. In this case the data set was constructed from twentieth-century poetry and captioned images. If you showed Albrecht Dürer's engraving *Saint Jerome in His Study* from 1514 to the camera it would come up with something like the following lines.

A Stone Wall fills a sea,
A shadow of rivers,
A spirit of light;
and in the water
Some storm that sweeps in the wind and the darkness.

A large brick wall with a slow view of the tress,
The town of the house where it was done with the water,
Having passed a Spring-fire.
A large rock on the ground is dark and shadowy,
and the whole world is dead. (Miller, 2019)

Later he entered the Sci-Fi London Film Festival 48-Hour Film Challenge with a film student, Oscar Sharp, whom he met at university. For the 48-Hour challenge constants are presented with a list of lines that must be included in the film and a set of props. As the name of the challenge insinuates the film has to be made of the course of a weekend and it can not be longer than five minutes. Goodwin curated a new data set of screenplays of sci-fi from the 80s and 90s for his Char-RNN, now called Jetson. The network produces a screenplay containing impossible directions such as “He is standing in the stars and sitting on the floor” (Miller, 2019) thus they decide on the form of a dream sequence for the film. They handed out the lines directly and unmodified to the three actors involved: Thomas Middelditch plays H, Elisabeth Gray plays H₂ and Humphrey Ker plays C.

The actors went on to interpret the lines as they acted them out, imbuing the directions with sense from their interpretation: “Actors spoke exactly the lines that the machine gave them” (Miller, 2019). They called their film *Sunspring* and while it got placed in the top ten of the festival it did not win.

Being an unusual contender in the festival, the software still called Jetson was interviewed on stage. Godwin would type the questions into his computer and Jetson would reply on the same computer. During this exchange Jetson proclaimed that its name is Benjamin and that is how Godwin credited the software for its writing in the film.

A more recent project is *the Road* which is a fleshed out version of word.camera taken for a through the U.S. from New York to New Orleans in a car. Four sensors were built into the car, a surveillance camera on the roof, a GPS unit, a microphone and the computer's internal clock. The software running on the computer inside the vehicle creates sequences of words in response to what the sensors receive from the outside world that the car passes through. The first reaction came out of the printer in Brooklyn: “It was nine seventeen in the morning and the house was heavy.” (Miller, 2019)

Godwin's projects are intriguing, presentable and successful, he is now a creative technologist at Google’s Artist and Machine Intelligence Program. A certain exploratory activity in language with technology like we see it with Parrish and Monfort is missing. His approach to not edit the messages that his works generate hints at the belief that software can have the creativity and volition to express itself.

8.6. Tony Veale

Tony is from Ireland and is a computer scientist that got hooked on computers when he saw Stanley Kubrick’s *2001: A Space Odyssey* as a teenager. He was a bit overwhelmed by the film but he saw HAL 9000 as a character capable of speaking, thinking, feeling and playing chess, and most importantly plot the murder of its coinhabitants, the astronauts.

When studying computer science, he chose to focus on natural language processing, the way computers understand human language. Unlike other rather dystopian commentators on the possibilities of computers he has an optimistic perspective. He sees AI “as a place to explore passionate ideas like literature and creativity; it’s a way to doing creative work by another route because computer

science touches on everything.” (Miller, 2019)

In the research domain for natural language processing his main interest is how metaphors are working in human language. To explore this area he created a project called Metaphor Magnet. A program that prompts for a word and in return generates an endless stream of metaphors and similes.

There exists also a corresponding Twitter account named *@MetaphorMagnet* that not only uses content from the name-giving software but also falls back on other projects of Veale. It tweets out a metaphor approximately every hour.

Nice adventurers like Arthur Dent put me in mind of hamsters: they're as nice as Jesus

Disturbing painters like Francis Bacon put me in mind of nightmares: they're as disturbing as the dancing midget from Twin Peaks

Fluid playwrights like James Joyce put me in mind of dances: they're as fluid as water at room temperature

Similar to Parrish he uses a database of n-grams to create conceptual metaphors. In the quotes above you can see that the pattern is *X is Y* (e.g. Francis Bacon is a disturbing painter and the dancing midget is disturbing). These metaphor pairs get extracted from the database and the bot juxtaposes them according to “resonant constraints and norm contraventions” (Miller, 2019)

He also built a story generation system called *Scêalexctric*, from a combination of

scéal, the Irish word for story and *Scalextric*, a slot car racing construction set with hand-controlled electric cars. Veale explains it as “a simulation-based approach to the generation of episodic stories in which stories are generated, evaluated and frequently discarded in rapid, coarse-grained cycle of engagement and reflection.” (Miller, 2019)

The interface of the generator consists of a form that contains two checkboxes. The first checkbox allows the user to choose the characters of the story. Either just names from a list of known people, only first names or one famous name and the rest first names. The second check box has three options on how to link the different story arcs: *Causal references to previous actions*, *Strict linking of arcs (last arc action to next)* and *Information transfer between characters*. The following story was generated with one famous character and the rest first names and all the arc linkage options activated.

The story of Molly

Cherry would grant boons to those who played the game.

Molly found Cherry at a farmer's market. Molly assiduously curried favor with her. Cherry summoned regal Molly to an emergency meeting. Molly followed critical Cherry's orders, but Cherry brainwashed brainwashing Molly into believing a pack of lies. She said: "I turned against Martha." Molly's loyalty extended to killing implacable Cherry's enemies. Cherry would not give the time of day to nonchalant Molly. She failed to properly reward her. Molly whined and complained to supportive Cherry. Molly said: "I assiduously curried favor with you." Cherry coldly dismissed abrupt Molly and turned away.

At a sculpture garden Molly met Garth Brooks. She went to work for him

after Cherry told haughty Molly to get out and not come back. Molly obeyed ruthless Garth's orders. Garth took full advantage of vulnerable Molly. She was totally fooled by him. Garth said: "Martha sent you out into the field." She filed an official report against Garth Brooks. She said: "Cherry no longer respected or trusted idealistic Martha."

After Molly arrived at a alien cantina, Martha appeared. Molly made a big impression on her because before she testified against Garth in open court. Martha appointed influential Molly to high office. She did not treat her well. Molly thoroughly disappointed demanding Martha.

Do you feel you have been a disappointment to anyone?

Scéallextric has over three thousand prefabricated plot sections. The system searches for matching known people in a genre-spanning way and finds them according to their similarities.

The quality of the stories or plots generated is among the best in the generators review in this text. Still, there is a certain uncanny valley feeling to them. The main reason why the generator works in such a convincing fashion is the enormous backing database that was assembled by hand. This is the big flaw of such approaches, they encode the creator's knowledge of the world into symbols that the system can use to argue about the world. These symbols are fixed and must be modified by hand, unlike more flexible AI systems that can be evolved in a more straightforward way.

This difference between his systems and the contemporary AI system for text generation is Veale's "beef" (Miller, 2019). He states neural networks are "somewhat accidental and don't use knowledge deliberately or knowingly" (Miller, 2019) While I doubt that any software has volition and thus can use knowledge deliberately or

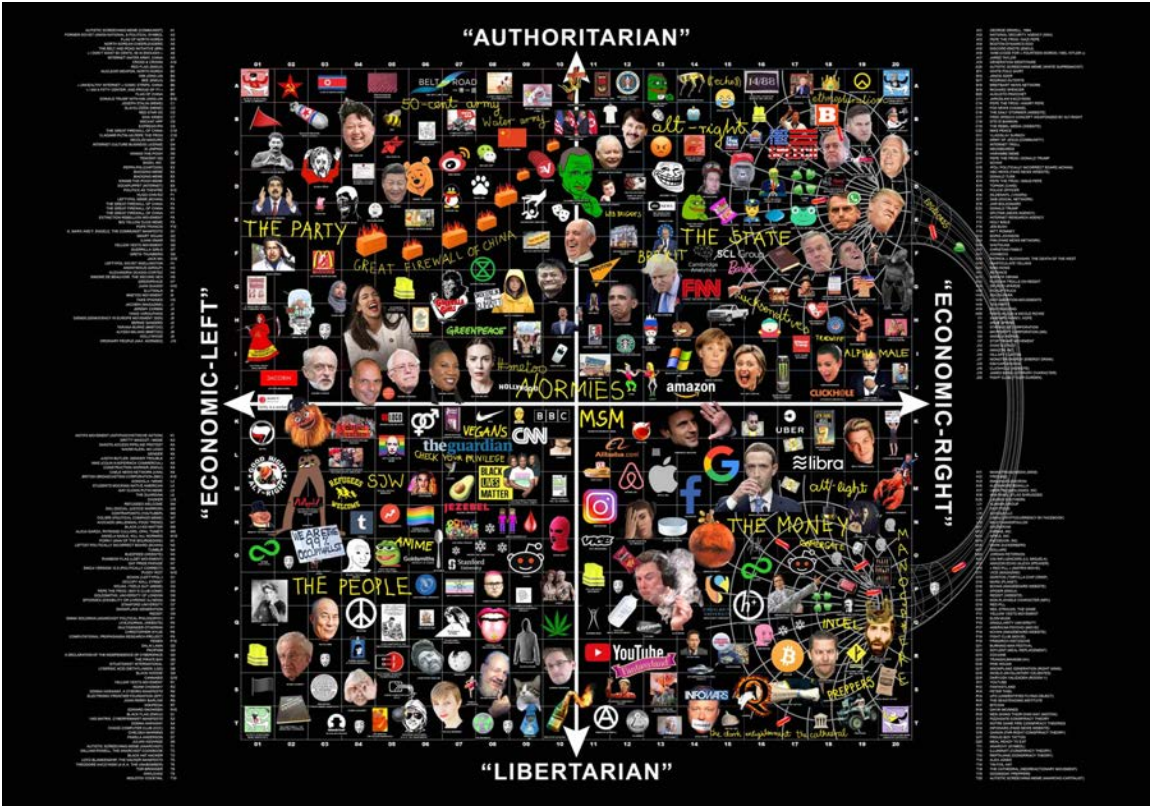


Figure 24: Online Culture Wars (Map), 2019, by Disnovation.org

knowingly, his rather hapless quote shows us an important distinction between symbolic systems and AI systems. In symbolic systems, because of their manual construction, decisions can be followed, shown and reasoned about. AI systems, because of their emergent nature and the interconnectedness of their inner parts, do not offer a simple approach to extract reasoning or explanations.

8.7. Disnovation.org

Disnovation.org is a research and working group at the intersection of contemporary art, research and hacking. Their current team consists of Maria Roszkowska, Nicolas Maigret and Baruch Gottlieb. They want to publicly question the contemporary techno-positivist ideologies that dominate our public sphere and

encourage thinking about post-growth narratives.

Their work has been presented internationally in places such as Centre Pompidou in Paris, at the Transmediale in Berlin, in the Palais de Tokyo in Paris and the ZKM in Karlsruhe. They are currently visiting researchers at the University of California, Irvine.

For their project *Online Culture Wars (Map)*, they engaged with the political landscape of the internet based meme culture. The term *meme* was first used by Richard Dawkins in his 1976 book *The Selfish Gene* as an attempt to explain how ideas replicate, mutate and evolve in our culture. The specialized term *internet meme* was first used in a Wired article in 1993 by Mike Godwin. While Dawkins' idea talks about ideas in general, internet memes are characterized by their means of production and dissemination is at the same their cultural hotbed, the internet.

“From Twitter to Facebook, Instagram, Reddit, Gab, or 4chan, one of the by-products of online culture wars is the over-politicization of seemingly mundane topics, products, practices, and cultural elements.”(Disnovation.org, 2018) Navigating and understanding the implied meaning of these politicized memes is rather difficult and requires a profound knowledge of the subcultures from which the memes originate. An example for this ambiguity is a comic character called *Pepe the Frog*, a green anthropomorphic frog with a humanoid body. It was created by Matt Furie for the comic *Boy's Club* in 2005. Until around 2015 *Pepe* was mostly used in the “feels good man”²⁵ meme when he got used by the fascist alt-right movement. In 2016 the Anti-Defamation League²⁶ included some variation of *Pepe* in its hates symbol database. Unknowingly *Pepe* got used in the 2019-2020 Hongkong protests where

25 See <https://knowyourmeme.com/memes/pepe-the-frog>

26 See <https://www.adl.org/>

27 See <https://knowyourmeme.com/memes/political-compass>

many of the protesters were unaware of its significance as a symbol of white nationalism in the western hemisphere.

Disnovation.org has arranged a myriad of internet memes, symbols and political figures in a two-dimensional space with the following axes: “Authoritarian” ⇌ “Libertarian” and “Economic-Left” ⇌ “Economic-Right”. This arrangement is called the *political compass*²⁷, unsurprisingly it was also used as a meme.

Now this project is only tangentially related to the work presented in this thesis but it shows that Disnovation.org locates themselves differently than the other artist presented. While the other artists are mainly interested in language and creativity, Disnovation.org also notes the political aspect and social aspects of our contemporary society and the cultural production that is embedded. This actually leads to the main project that I want to visit, the *Predictive Art Bot*

*Predictive Art Bot*²⁸ is “an algorithm that turns the latest media headlines into artistic concepts” (Disnovation.org, 2018). The algorithm collects headlines from a curated set of news sources and generates artistic concepts, mimicking the contemporary creative process of identifying trends in media for exploitation in artistic productions. “Predictive Art Bot caricatures the predictability of media influenced artistic concepts, by automating and skirting the human creative process” (Disnovation.org, 2018)

To generate a new concept the algorithm selects one noun each from three headlines and combines them into a short and punchy thought. The following quotes the headlines with the selected nouns and the result.

US charges **Russian government hackers** for NotPetya cyberattack

28 See <http://predictiveartbot.com/>

People with Parkinson's disease have a higher risk of **dying** from COVID-19

Kanye West 'Praying' For Issa Rae After He Slams Her 'SNL' **Sketch** Mocking His Candidacy

TURNING RUSSIAN GOVERNMENT HACKERS INTO A DYING SKETCH

The two most recent version of this artwork are a dual screen installation and a web based version. The installation uses two displays in upright position, one shows the currently selected headlines, the other the generated concept. The web based version is using a similar layout as the screen is split into two panes, one with the currently selected headlines the other containing the generated concept. In both versions the headline is accompanied by a picture extracted from the corresponding



Figure 25: Predictive Art Bot installation, 2019, by Disnovation.org

website. The presented user interface is simple and well-thought-out because of one perfect detail: Before visually selecting the nouns from the headlines, it waits for a few seconds giving the user the opportunity to read the headlines and guess for the selected word.

The qualities of *Predictive Art Bot* are the data set that they curated, the criticism of the contemporary art landscape and the openness towards its users. The datasets consist of an interesting mix of news sites that range from mainstream like the Guardian, Yahoo News and the New York Times to left leaning publications such as Slate, right leaning publications like Fox News and RT, design websites and further less mainstream website that occupy themselves with themes like cryptocurrencies. This creates a rather good portrayal of the media landscape that a contemporary artist would come in contact with and where he gets his inspiration from. But unlike a dataset for an AI project that at one point in time is finished and frozen, this project keeps itself current by selecting news sites and not directly articles from those sites. The sites publish articles on a regular basis and these articles are fed into the bot. That means if you use the project today, it will create up to date concepts.

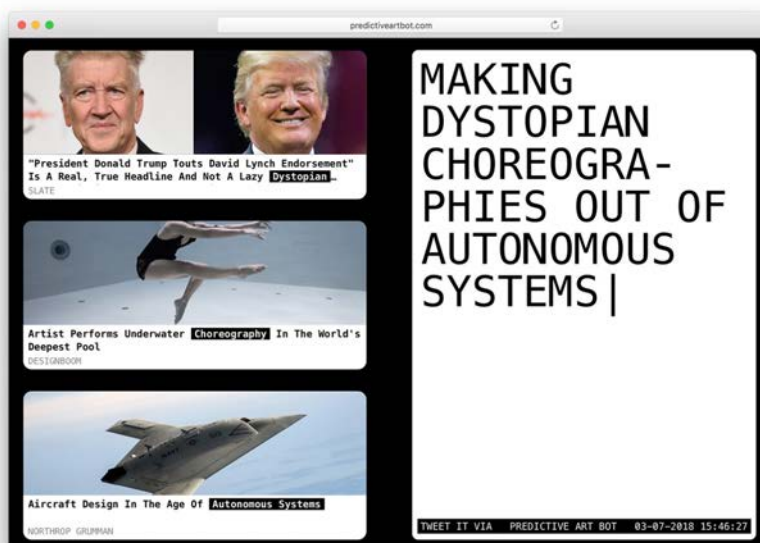


Figure 26: Predictive Art Bot web, 2019, by Disnovation.org

The process and experience that they created that generates the concept is very enjoyable and immersive and provides the proverbial mirror to reflect on our art practice. While the bot often creates outlandish proposals they become “thought experiment[s] waiting to be incubated, misused or appropriated by a human host” (Disnovation.org, 2018).

8.8. Jörg Piringer

Jörg Piringer is an Austrian sound and visual poet currently living in Vienna. He has a master’s degree in computer science and is a member of the institute for transacoustic research and of the vegetable orchestra.

In his book *datenpoesie*²⁹ (Piringer, 2018) he is exploring the language technologies that are heavily influencing our contemporary society. He states that today most of the textual content on the internet is generated and consumed by software, currently this content is mostly made up of standardized protocols and structured text such as HTML. But already items like sport reports, weather forecasts or stock market news are generated by software for a broader mass of human consumers. At the same time search engines are constantly trawling through the internet, trying to extract sense from the textual content of websites. For this purpose they are using knowledge of linguistics and cognitive science.

He approaches this situation with methodologies from the field of artistic research and explorative programming, similar to how Allison Parrish approaches her work. This approach in contrast to scientific research and software engineering is focused on subjective factors, the artistic value and the personal development and knowledge gain.

²⁹ There is also an English issue, for that see <https://joerg.piringer.net/index.php?href=text/datapoetry.xml>

³⁰ In English *verse*

The book has five chapters that explore different forms of computer generated creative text, which are interspersed with black and white illustrations made with typography.

The first chapter, *vers*³⁰, contains poetic texts ranging from chains similar to the ones created by the Wiener Gruppe, an exploration of the meaning of sequences of kanji characters, poetry generated by simple and complex neural networks, back and forth translation of a Christian Morgenstern poem and proverbs.

rache ist der beste koch

blinder, kleinvieh macht

alte liegt die feste medizin.

(proverbs, Piringer, 2018)

The following chapter *frage und antwort*³¹ contains a collection of different types of dialogues between changing characters. There is one fragment where speaker A and B speak to each other, in another text the machine reacts to Jörg walking around in his workspace.

[...]

ein mann steht vor einem kühlschrank

ein mann steht vor in einem raum mit einer fernbedienung

31 In English question and answer

ein mann steht vor einem kühlschrank
ein mann steht vor einem fernseher mit fernbedienung
ein steht in einem zimmer mit einem fernseher
[...]
(captions, Piringer, 2018)

[...]
ich dachte, du wärst längst abgehauen.
sie dachte, es beeinflusst meine entscheidung.
ich fragte mich: warum?
sie sagte nie ja.
(untertitel, Piringer, 2018)

The chapter *suchtext*³² poses questions to the software about different phenomenon of the world. Like a sequence of the defining adjectives of each day of the week, asking Twitter “Who is [...]?”, probing chat protocols looking for definitions of “We are [...]”, letting the internet complete the infamous sentence “I am not a [...] but [...]”³³.

wer ist dieser?
dieser ist nur umsetzbar mit sonde und wind.
dieser ist der tag des gregorianischen kalenders.
dieser ist lila,
dieser ist soeben angekommen schönen dank an de

wer ist hier?

32 In English searchtext

33 In the original “Ich bin zwar kein [...] aber [...]”

Hier ist immer noch unser favorit

hier ist unser demobericht zu den protesten in kandel gestern.

hier ist österreichisch.

hier ist der gewünschte #mudipani .

hier ist der tweet.

(wer ist, Piringer, 2018)

ich bin zwar keine schwedin, aber ich hab ordentlich saubergeleckte vintage-suppenteller vom trödel, kleine skizzenzeichnungen aus ...

ich bin zwar kein astronaut doch ich hab business gemacht. Und jede blitch in der stadt will mit ficken und sagt : uh baby sei nicht so bitte ...

(ich bin zwar, Piringer, 2018)

The chapter *fließtext*³⁴ contains longer form stories generated by a diverse set of approaches. He trained a neural network on the text of *Mein Kampf* by Adolf Hitler, one text is the output of the half trained network and one text is generated by a fully trained network. Both texts are aggressively incomprehensible, the text created by the half trained network is just a multipage repetition of *staat*³⁵. There is also the *Universal Declaration of Human Rights* but run through a process similar to the one that search engines apply to extract relevant information and to simplify comparison between websites or texts. It consists of removing all stop words and reducing words to their stem.

da anerkenn angebor wurd gleich unverauser recht mitglied gemeinschaft
mensch grundlag freiheit gerechtig fried welt bildet nichtanerkenn veracht

34 In English continuous text

35 In English state

menschenrecht akt barbarei gefuhrt hab gewiss menscheit empor erfull ...
(allgemein erklr menschenrecht, Piringer, 2018)

der staat der staat der staat der staat der staat der staatsallen der staat der
staaten der staaten der staat der schänder der menschen der staaten der
schänder ...

(künstliche unintelligenz I, Pringer, 2018)

The chapter *anordnung*³⁶ is concerned with sequential and visual arrangements of words from the German language, be it a list of the most frequent words, random words from a list that is used to test passwords or uncovering all words in composite words such as *waffenstillstandsvereinbarung*³⁷. Further there are 4 character longs anagrams, 3 by 3 and 4 by 4 character squares based on the Latin Sator square and an exploration of the relationships between words in WordNet starting from eye.

als das mehr daß hat diese den sind und um zu nur jahren der kann sei oder
der nach schon hatte nur des oder ...

(häufigste wörter, Piringer, 2018)

abels basel blase salbe

adeln laden lande nadel

adelt delta ladet tadel

agent getan nagte tagen ...

(anagramme, Piringer, 2018)

eye is part of needle.

eye is part of visual system.

36 In English order or arrangement

37 In English armistice agreement

eye is part of face.
face is part of watch.
face is part of ticker.
ticker is part of circulatory system.
circulatory system is part of body.
...

(part of, Piringer, 2018)

Piringer's books is a revealing excursion into contemporary approaches and processes to text generation based around the concepts of artistic research and personal knowledge gain. Through its output it stressed the need to interact and confront these approaches as it becomes clear that they will have an incredible impact on our cultural spheres.

In contrast to the project described in this thesis he mostly chooses tools that are open to inspection and explanation. He seeks to elevate software from the status of the black box by engaging with it and not only providing output but also an explanation.

Throughout the book all the fragments have a cool distanced Viennese humour and charm. This hints that Piringer is not only enthusiastic about software and the possibilities that it provides, but can also see the downsides of text generation software for our society.

9. Conclusion & Findings

The theoretical and practical part of this thesis proposes that AI based art can overcome the status of a technological demonstration through the collection and curation of custom datasets. The theoretical part shows the weaknesses of contemporary approaches while I demonstrate a better fitting approach in the practical work.

In the first chapter we saw that ImageNet struggled with the technological skill of contemporary internet search engines to accurately produce search results for the synset of WordNet. To mitigate this flaw the researchers built a system of human computation based on Amazon's Mechanical Turk workers to correct erroneously categorized images.

This system of human computation is inherently flawed if asked to categorize pictures of people and produces crude and problematic results.

The JAFFEE datasets show that they can lend a veil of science to controversial theories. It also shows the main problem of Ekman's work, which is also its scientific foundation of JAFFEE. The problem lies in the fact that acting out an idea or an emotion as facial expressions compared to a genuine expression of emotions rarely results in the same facial expression. It further serves as an example of the fact that just because there is a dataset of a certain effect, it does not mean that one can survey the proposed effect in nature.

With these two examples I want to propose to artists to mindfully curate their own datasets that at every step of their creation incorporate inclusion of the entities that produce the set's content. In this way maybe we can evade the exploitative processes and end up with artistic datasets that have a positive influence.

Further I conclude that artistic datasets can be used to ask insightful and

inconvenient questions about approaches that are being built on shaky science.

In the second chapter we saw that writing was and is never the singular work of an exceptional genius author at both the level of reading and the text itself, as Barthes shows, but also on the level of craft, as shown by Woodmansee. The author as the single source of text is further questioned by text machines and story generators by Simanowski in (Simanowski, 2009).

In the same chapter we also encounter the concept that we can read a text as a technical achievement of the programmer-author. In the context of Artificial Intelligence art this is a view that I want to discourage, as the influence of the dataset in the result far outweighs the influence of the algorithm. Instead, I propose to use widely available and tested implementations of AI algorithms, and invest time in a critical and creative examination and creation of datasets.

The approach of the programmer author makes sense in the context of exploratory programming like the work of Parrish and Piringer which we saw in the chapter “Related Projects”.

Part about the whole internet and simple heuristic, there is the part about

To conclude, I think that global internet type of scale is not compatible with poetic software that wants to stay true to itself and its users - or “The poetical act is materially completely worthless and therefore from the outset untainted by the bacillus of prostitution. its pure accomplishment is simply noble” (Artman, 1953).

In the chapter “Poetic Intention and collective authorship” we saw that in a short timespan of a few years the Wiener Gruppe made a huge leap through different

genres of poetry, and mostly left behind the idea of the author as a sole genius. Their work did not get worse because of it. This romantic idea of the sole genius author was actually constructed by the authors themselves for their own social and economical gain. This leads me to the conclusion that this idea actually narrows our openness to new genres of poetry that do not adhere to this dogma. Further it also restricts the recognition of our innate poetic intent, because we rarely ever see ourselves as geniuses.

Furthermore, if a process helps its user to unconsciously overcome the concept of the genius author, they are then more free to express themselves, and to question the quality of their work less. In the context of the “Eight-Point-Proclamation of the Poetical Act”, I want to argue that this actually heightens the quality of the texts they write. Maybe not in the sense of classical poetry, but more in the sense true self-expression as said in the proclamation:

The poetical act is poetry purely for poetry’s sake and is free of any ambition for recognition, praise or criticism.

Surely, as with most abilities of a human, the outcomes of the poetic intention grow more intricate and elaborate with practice. It will most probably not become better if we use classical poetry as a guide, but better in quality that the writer chose to evaluate his poems with.

In the chapter “Creative output as a commodity” I wrote about the thesis of Manovich that multiple AI systems in a market will neutralize themselves. While I think that the point made by Manovich is valid, I would like to use Spotify as a contemporary counter example that shows what could go wrong.

From its humble beginning in 2006 it grew to become synonymous for music

streaming, and while there are competitors like Apple Music and Google Play Music, Spotify sets the tone of the conversations around streaming. It is argued that Spotify is not supporting artists and creativity enough. Through opaque payout schemes and a dismantling of the supporting infrastructure for music artists it moved most of the risk to the individual artists. At the same time it exposes those artists to a single global winner takes it all market, where small labels and artists have to compete on the same platform as the high finance corporations of the industry. It still prevailed because of the convenience provided to the users and the backing of the remaining major labels of the music industry. I think a similar outcome could also be possible for a future AI system.

In the same chapter I also investigated the relationship between output and source of inspiration on the basis of Jukebox by OpenAI. I asked from which artists the system learned about Classic Pop.

I imagine a scenario where the system for each generated artefact is able to exactly reveal the sources somehow like a DJ can produce a track listing for each set that they play.

Further we saw Margaret Boden giving us three criteria for creativity. They are new, surprising and valuable. While these are reasonable criteria I'm much more sympathetic to the approach of the French composer and computer scientist François Pachet who believes that creativity can be understood only in subjective terms.

But in the context of this thesis and the bigger nexus of the production of cultural fragments, the discussion if a computer only seems to be creative, or if they are really creative is redundant. As we can see from the given examples and explanations, it is entirely possible to have software generated cultural artefacts that appear creative. It

seems just to be a question of time until a company develops an AI system to produce cultural artefacts. There surely are certain aesthetical and technical problems to overcome before such a system can be successful, but I still expect it to become an important part of the creative industries in the near future.

10. Bibliog- raphy

- Barger, J.** (1993) Racter FAQ. Available at: <https://web.archive.org/web/20070225121341/http://www.robotwisdom.com/ai/racterfaq.html> (Accessed: 13 August 2020).
- Barrett, L. F. et al.** (2019) 'Emotional Expressions Reconsidered: Challenges to Inferring Emotion From Human Facial Movements', *Psychological Science in the Public Interest*, 20(1), pp. 1–68. doi: 10.1177/1529100619832930.
- Barthes, R.** (1968) 'Der Tod des Autors', in *Texte zur Theorie der Autorschaft*. Nachdr. Stuttgart: Reclam (Reclams Universal-Bibliothek, 18058), pp. 185–193.
- Boden, M. A.** (1994) 'Précis of The creative mind: Myths and mechanisms', *Behavioral and Brain Sciences*, 17(3), pp. 519–531. doi: 10.1017/S0140525X0003569X.
- Branwen, G.** (2020) 'GPT-3 Creative Fiction'. Available at: <https://www.gwern.net/GPT-3> (Accessed: 16 July 2020).
- Boden, M. A.** (2004) *The creative mind: myths and mechanisms*. London: Routledge.
- Boden, M. A.** (2015) What's next for AI - Margaret Boden, IBM Cognitive advantage reports. Available at: <http://www.ibm.com/watson/advantage-reports/future-of-artificial-intelligence/margaret-boden.html> (Accessed: 29 June 2020).
- Chamberlain, W. and Etter, T.** (1984) *The Policeman's beard is half-constructed: computer prose and poetry*. New York, NY: Warner Software/Warner Books. Available at: <http://archive.org/details/policemansbeardi0000unse> (Accessed: 18 August 2020).
- Cireşan, D., Meier, U. and Schmidhuber, J.** (2012) 'Multi-column Deep Neural Networks for Image Classification', arXiv:1202.2745 [cs]. Available at: <http://arxiv.org/abs/1202.2745> (Accessed: 14 September 2020).

- Cohen, H.** (1994) THE FURTHER EXPLOITS OF AARON, PAINTER. Available at: <https://web.archive.org/web/20060107184824/http://crca.ucsd.edu/~hcohen/cohenpdf/furtherexploits.pdf> (Accessed: 8 July 2020).
- Coverley, M.** (2009) Psychogeography. Nachdr. Harpenden: Pocket Essentials (Pocket essentials).
- Crawford, K. and Paglen, T.** (2019) Excavating AI, -. Available at: <https://www.excavating.ai> (Accessed: 22 June 2020).
- Debord, G.** (1955) The Situationist International Text Library/Introduction to a Critique of Urban Geography. Available at: <http://library.nothingness.org/articles/SI/en/display/2> (Accessed: 17 April 2020).
- Debord, G.** (1956) Situationist International Online: Theory of the Dérive. Available at: <https://www.cddc.vt.edu/sionline/si/theory.html> (Accessed: 8 September 2020).
- Deng, J. et al.** (2009) 'ImageNet: A Large-Scale Hierarchical Image Database', p. 8.
- Dhariwal, P. et al.** (2020) 'Jukebox: A Generative Model for Music', arXiv:2005.00341 [cs, eess, stat]. Available at: <http://arxiv.org/abs/2005.00341> (Accessed: 20 August 2020).
- Dhariwal, P., Jun, H. and Payne, C. (2020) Jukebox Website, OpenAI. Available at: <https://openai.com/blog/jukebox/> (Accessed: 20 August 2020).
- Foucault, M.** (2009) 'Was ist ein Autor?', in Texte zur Theorie der Autorschaft. Nachdr. Stuttgart: Reclam (Reclams Universal-Bibliothek, 18058), pp. 198-229.
- Frei, F.** (2020) The Poetic Hivemind: Privacy Policy. Available at: https://tofu.wtf/poems/privacy_policy.html (Accessed: 7 July 2020).
- Friedlander, E.** (2020) How Holly Herndon and her AI baby spawned a new kind of folk music, The FADER. Available at: <https://www.thefader.com/2019/05/21/holly-herndon-proto-ai-spawn-interview> (Accessed: 17 July 2020).

- Gershgorn, D.** (2017) The data that transformed AI research—and possibly the world, Quartz. Available at: <https://qz.com/1034972/the-data-that-changed-the-direction-of-ai-research-and-possibly-the-world/> (Accessed: 12 July 2020).
- Gray, M. L. and Suri, S.** (2019) Ghost work: how to stop Silicon Valley from building a new global underclass. Boston: Houghton Mifflin Harcourt.
- Hao, K.** (2020) A college kid created a fake, AI-generated blog. It reached #1 on Hacker News., MIT Technology Review. Available at: <https://www.technologyreview.com/2020/08/14/1006780/ai-gpt-3-fake-blog-reached-top-of-hacker-news/> (Accessed: 18 August 2020).
- Hara, K. et al.** (2018) ‘A Data-Driven Analysis of Workers’ Earnings on Amazon Mechanical Turk’, in Proceedings of the 2018 CHI Conference on Human Factors in Computing Systems - CHI ’18. the 2018 CHI Conference, Montreal QC, Canada: ACM Press, pp. 1–14. doi: 10.1145/3173574.3174023.
- Leufer, D.** (2020) Myth: The term AI has a clear meaning. Available at: <https://www.aimyths.org/the-term-ai-has-a-clear-meaning> (Accessed: 27 August 2020).
- Link, D.** (2009) David Link, LoveLetters_1.0, 2009—... Available at: http://www.alpha60.de/art/love_letters/ (Accessed: 13 August 2020).
- Love Michael, M.** (2019) Holly Herndon on Raising Her AI Baby and Embracing the Future, PAPER. Available at: <https://www.papermag.com/holly-herndon-protocol-2637723054.html?rebelltitem=32#rebelltitem32?rebelltitem=32> (Accessed: 18 July 2020).
- Manovich, L.** (2019) AI Aesthetics. Strelka Press. Available at: <http://manovich.net/index.php/projects/ai-aesthetics>.
- Mehri, S. et al.** (2017) ‘SampleRNN: An Unconditional End-to-End Neural Audio Generation Model’, arXiv:1612.07837 [cs]. Available at: <http://arxiv.org/abs/1612.07837> (Accessed: 18 July 2020).

- Mixcloud** (2020) Introducing Mixcloud Select, Mixcloud. Available at: <https://www.mixcloud.com/select/> (Accessed: 25 August 2020).
- Miller, A. I.** (2019) *The Artist in the Machine: The World of AI-Powered Creativity*. Cambridge, MA: The MIT Press.
- Monroe, J.** (2020) Boiler Room Brings Mixes to Apple Music, Touting Novel Royalty Model, Pitchfork. Available at: <https://pitchfork.com/news/boiler-room-brings-mixes-to-apple-music-touting-novel-royalty-model/> (Accessed: 25 August 2020).
- Montfort, N.** (2014) *#! (shebang)*.
- Porr, L. and Adolos** (2020a) Feeling unproductive? Maybe you should stop overthinking. Available at: <https://adolos.substack.com/p/what-i-would-do-with-gpt-3-if-i-had> (Accessed: 31 August 2020).
- Porr, L. and Adolos** (2020b) We don't really understand randomness. Available at: <https://adolos.substack.com/p/what-i-would-do-with-gpt-3-if-i-had> (Accessed: 31 August 2020).
- Parrish, A.** (2015) Exploring (Semantic) Space With (Literal) Robots. Available at: <http://opentranscripts.org/transcript/semantic-space-literal-robots/> (Accessed: 8 October 2020).
- Perec, G.** (1969) *La disparition*. Paris: Gallimard (Collection l'Imaginaire, 215).
- Piringer, J.** (2018) *datenpoesie*. Klagenfurt am Wörthersee: Ritter Klagenfurt.
- Prabhu, V. U. and Birhane, A.** (2020) 'Large image datasets: A pyrrhic win for computer vision?', arXiv:2006.16923 [cs, stat]. Available at: <http://arxiv.org/abs/2006.16923> (Accessed: 2 July 2020).
- Quackenbush, C.** (2020) Painting Made by Artificial Intelligence Sells for \$432,500, Time. Available at: <https://time.com/5435683/artificial-intelligence-painting-christies/> (Accessed: 27 August 2020).

- Radford, A. et al.** (2019) 'Language Models are Unsupervised Multitask Learners', p. 24.
- Richardson, L.** (2014) 'Writing Aliens', or, 'Duchamp, Markov, Queneau: A Mostly Delightful Quilt'. Available at: <https://www.crummy.com/writing/speaking/2014-Foolscap/> (Accessed: 16 June 2020).
- Rosebud AI** (2020) Rosebud AI. Available at: <https://generative.photos/?ref=producthunt#howitworks> (Accessed: 1 September 2020).
- Simanowski, R.** (2011) Digital art and meaning: reading kinetic poetry, text machines, mapping art, and interactive installations. Minneapolis: University of Minnesota Press (Electronic mediations, 35).
- Smith, K.** (2016) The Wander society. New York, New York: Penguin Books.
- Sommerer, C. and Mignonneau, L.** (2015) 'Living Poetry', *Artificial Life*, 21(3), pp. 313–319. doi: 10.1162/ARTL_a_00172.
- de la Torre, M. and Poets, A. of A.** (2005) Into the Maze: OULIPO | Academy of American Poets. Available at: <https://poets.org/text/maze-ouliipo> (Accessed: 16 September 2020).
- Thi Nguyen, C. and Weinberg, J.** (2020) Philosophers On GPT-3 (updated with replies by GPT-3): Who Trains the Machine Artist?, *Daily Nous*. Available at: <http://dailynous.com/2020/07/30/philosophers-gpt-3/> (Accessed: 3 August 2020).
- Weibel, P. (ed.)** (1997) Die Wiener Gruppe: ein Moment der Moderne 1954-1960: die visuellen Arbeiten und die Aktionen: Friedrich Achleitner, H.C. Artmann, Konrad Bayer, Gerhard Rüm, Oswald Wiener = the Vienna group: a movement of modernity 1954-1960: the visual works and the actions: Friedrich Achleitner, H.C. Artmann, Konrad Bayer, Gerhard Rühm, Oswald Wiener. Wien; New York: Springer.
- Whittaker, M. et al.** (2018) 'AI Now Report 2018'.

- Woodmansee, M.** (1992) 'Der Autor-Effekt. Zur Wiederherstellung der Kollektivität', in *Texte zur Theorie der Autorschaft*. Nachdr. Stuttgart: Reclam (Reclams Universal-Bibliothek, 18058), pp. 298-314.
- Woolf, M.** (2019) How To Make Custom AI-Generated Text With GPT-2, Max Woolf's Blog. Available at: <https://minimaxir.com/2019/09/howto-gpt2/> (Accessed: 21 April 2020).
- Yang, K. et al.** (2019) 'Towards Fairer Datasets: Filtering and Balancing the Distribution of the People Subtree in the ImageNet Hierarchy', arXiv:1912.07726 [cs]. doi: 10.1145/3351095.3375709.
- Zhang, X., Zhao, J. and LeCun, Y.** (2016) 'Character-level Convolutional Networks for Text Classification', arXiv:1509.01626 [cs]. Available at: <http://arxiv.org/abs/1509.01626> (Accessed: 11 May 2020).

10.1. Further Reading

- Akten, M.** (no date) 'Ultrachunk', Memo Akten | Mehmet Selim Akten | The Mega Super Awesome Visuals Company. Available at: <http://www.memo.tv/works/ultrachunk/> (Accessed: 8 July 2020).
- Apple** (2020) App Store Review Guidelines - Apple Developer. Available at: <https://developer.apple.com/app-store/review/guidelines/> (Accessed: 5 July 2020).
- Brown, D.** (no date) 'Strolling the Coastline: Criminology in Everyday Life: Through "Landscape" from Gaol to 'Badlands'', Law Text Culture, 13(Crime Scenes).
- Burnham, J.** (1980) 'Art and Technology: The Panacea That Failed', p. 16.
- Chtcheglov, I.** (1953) Formulary for a New Urbanism (Ivan Chtcheglov). Available at: <http://www.bopsecrets.org/SI/Chtcheglov.htm> (Accessed: 17 April 2020).
- Cole, D.** (2020) 'The Chinese Room Argument', in Zalta, E. N. (ed.) The Stanford Encyclopedia of Philosophy. Spring 2020. Metaphysics Research Lab, Stanford University. Available at: <https://plato.stanford.edu/archives/spr2020/entries/chinese-room/> (Accessed: 3 July 2020).
- Denton, E.** (no date) Critical Perspectives on Computer Vision, SlidesLive. Available at: <https://slideslive.com/38923500/critical-perspectives-on-computer-vision?ref=og-meta-tags> (Accessed: 22 June 2020).
- Gebru, T. and Denton, E.** (2020) FATE/CV - Schedule. Available at: <https://sites.google.com/view/fatecv-tutorial/schedule> (Accessed: 22 June 2020).
- Gershgorn, D.** (no date) The data that transformed AI research—and possibly the world, Quartz. Available at: <https://qz.com/1034972/the-data-that-changed-the-direction-of-ai-research-and-possibly-the-world/> (Accessed: 12 July 2020).
- Goodwin, R.** (2016) Narrated Reality. Available at: <https://narratedreality.com/> (Accessed: 3 July 2020).
- Hsieh, K.** (no date) transformer poetry.

- Hugging, F.** (2019) gpt2 · Hugging Face. Available at: <https://huggingface.co/gpt2> (Accessed: 21 September 2020).
- Kashyap, P., Phatale, S. and Drori, I.** (2019) 'Prose for a Painting', arXiv:1910.03634 [cs]. Available at: <http://arxiv.org/abs/1910.03634> (Accessed: 31 October 2019).
- Klingenman, M.** (2019) 'About | Quasimondo'. Available at: <http://underdestruction.com/about/> (Accessed: 24 September 2020).
- Kuchina, S. A.** (no date) On Generative Poetry: Structural, Stylistic and Lexical Features. Available at: https://impactum-journals.uc.pt/plugins/generic/pdfJsViewer/pdf.js/web/viewer.html?file=https%3A%2F%2Fimpactum-journals.uc.pt%2Fmatlit%2Farticle%2Fdownload%2F2182-8830_6-1_5%2F4786%2F (Accessed: 12 May 2020).
- Lee, R.** (2019) 'Uncertainties in the Algorithmic Image', *Journal of Science and Technology of the Arts*, 11(2), pp. 36–40. doi: 10.7559/citarj.v11i2.661.
- Leufer, D.** (no date) Myth: The term AI has a clear meaning. Available at: <https://www.aimyths.org/the-term-ai-has-a-clear-meaning> (Accessed: 27 August 2020).
- Mersch, D.** (2020) '(Un)creative Artificial Intelligence: A Critique of "Artificial Art"'. doi: 10.13140/RG.2.2.20353.07529.
- Montfort, N.** (2020) Nick Montfort's site. Available at: <https://nickm.com> (Accessed: 6 October 2020).
- Onkaos** (2020) 'MARIO KLINGEMANN', Onkaos. Available at: <https://onkaos.com/mario-klingemann/> (Accessed: 24 September 2020).
- OpenAI** (2020) Jukebox, OpenAI. Available at: <https://openai.com/blog/jukebox/> (Accessed: 20 August 2020).
- Palea, D.** (2020) 'Transformer Bard: Music and Poem Generation Using Transformer Models', p. 4.

- Piringer, J.** (no date b) jörg piringer - data poetry / datenpoesie. Available at: <https://joerg.piringer.net/index.php?href=text/datapoetry.xml> (Accessed: 21 October 2020).
- Raita, Y.** (2016) The Day a Computer Wrote a Novel, Big Echo. Available at: <http://www.bigecho.org/the-day-a-computer-wrote-a-novel> (Accessed: 3 July 2020).
- Sato, S.** (2016) 'A Challenge to the Third Hoshi Shinichi Award', in Proceedings of the INLG 2016 Workshop on Computational Creativity in Natural Language Generation. Proceedings of the INLG 2016 Workshop on Computational Creativity in Natural Language Generation, Edinburgh, UK: Association for Computational Linguistics, pp. 31–35. doi: 10.18653/v1/W16-5505.
- Takahashi, Y.** (2020) Exceptional place to work – Internship Report from Austria, Medium. Available at: <https://medium.com/qosmo-lab/exceptional-place-to-work-internship-report-from-austria-e45c3fb283d8> (Accessed: 2 September 2020).
- Temkin, D.** (2020) The Hacker Aesthetic of Minimalist Code, Hyperallergic. Available at: <https://hyperallergic.com/566816/the-hacker-aesthetic-of-minimalist-code/> (Accessed: 7 October 2020).
- Wallbank, J.** (1999) LOWTECH MANIFESTO, Lowtech Manifesto. Available at: <http://lowtech.org/projects/n5m3/> (Accessed: 20 August 2020).
- Weinberg, J.** (2020) Philosophers On GPT-3 (updated with replies by GPT-3), Daily Nous. Available at: <http://dailynous.com/2020/07/30/philosophers-gpt-3/> (Accessed: 3 August 2020).
- Zucker, M.** (no date) uh hey guys I made a meme, Twitter. Available at: https://twitter.com/matt_zucker/status/1275574621522341888/photo/1 (Accessed: 25 June 2020).



CC BY-NC-ND 4.0 International
Namensnennung - Nicht-kommerziell - Keine Bearbeitung 4.0 International