Epilogue
In search of Trurl's electric bard

Whenever I speak publicly about AI and address a general audience, discussions inevitably arise that go beyond purely technical considerations and touch on philosophical and aesthetic issues as well as ethical and economic questions. In view of the latest advances in the field of generative AI, questions arise about the value, meaning and classification of literary texts, images and musical compositions produced with the help of computer programs.

As far as ethical and economic issues are concerned, I am merely an interested layman. But i really enjoy the philosophical and artistic discussions. After all, it was my interest in generative processes that led me to get into software development in the first place many years ago: A friend, who knew about my interest in mathematics and electronic sounds, had given me a copy of the computer algebra software Mathematica. He said that it was suitable for producing sounds using code. My curiosity was aroused. With almost no prior theoretical knowledge, I started programming and doing algorithmic composition.

Algorithmic composition

I was thus entering a tradition that was much older than I realized at the time. The emergence of the idea of composing music using algorithmic processes goes back much further than the construction of the first digital computers: when the first programmer, Ada Lovelace, formulated the idea of a creative computer in 1842, she too initially thought of music - the Analytical Engine designed by Charles Babbage could "work with things quite different from numbers" and "compose sophisticated and scientifically sound pieces of music of any complexity and length. "[1]

Back in the middle of the 17th century, Athanasius Kircher provided a concrete example of algorithmic composition rules with his "Arca Musarithmica" : He wrote a combinatorial procedure that enabled musical amateurs to compose.

1 Quoted from Jens Schröter: "Bilder weben, Musik komponieren. Ada Lovelace und das Universalmedium Computer" in Krämer 2015, p. 69.

The number of works potentially hidden in his invention runs into the millions and highlights the need for selection - a problem of generative art that is as topical as ever[2].

> When real computing systems became available for non-military purposes for the first time in the post-war period, musical applications came onto the scene shortly afterwards. In the late 1950s, Lejaren Hiller and Leonard Isaacson presented the Illiac Suite, the first string piece composed by a computer.[3]

There is a good reason why the first generative processes were primarily used to compose music: musical structures *are directly numerical* [4] and do not need to be translated before they can be treated algorithmically - with or without a digital computer.

Computer poetry and information aesthetics

An early example of algorithmically generated literature is Raymond Queneau's "Cent Mille Milliards de Poèmes "[5] (One Hundred Thousand Billion Poems), published in 1961. The work contains 10 sonnets of 14 lines each. The author allows these lines to be freely combined and exchanged, so that 10 to the power of 14 different sonnets can be created. Here, too, the above-mentioned selection problem arises.

Also in the 1960s, the philosopher Max Bense, the mathematician and artist Frieder Nake and the electrical engineer Abraham A. Moles developed information aesthetics[6]. This is a field of research that combines aesthetics with information theory, mathematics and computer science. The theoretical and practical results of these efforts are highly remarkable, and it is definitely worth taking a closer look at them. However, they can by no means claim to encompass the entire field of human artistic creation, but only occupy niches in art history and art studies.

2 Fred K. Prieberg: "Musica Ex Machina. Über das Verhältnis von Kunst und Technik", Berlin 1960, p. 106 f.

3 Lejaren Hiller and Leonard Isaacson: "Experimental music; composition with an electronic computer", New York 1959.

4 See section 14.2, "Coding words into numbers for advanced users".

5 Raymond Queneau: "Cent Mille Milliards de Poèmes", Paris 1961.

6 For example in Max Bense: Aesthetica, Baden-Baden 1965, Abraham A. Moles: "Kunst & Computer", Cologne 1973, Frieder Nake: "Ästhetik als Informationsverarbeitung", Vienna 1974.

This is true of all the achievements of early algorithmic art. They were dominated by combinatorial and pseudo-random methods, which incidentally have very little to do with what we understand by AI today. The focus was on top-down methods, such as the formalization of individual aspects of musical and linguistic structures for subsequent use in production rules.

The limits of the calculable

The limitations of early algorithmic art and the corresponding art-scientific efforts are obviously due to the restriction to what could be formalized and algorithmically processed with the theoretical and practical tools of the time. This shortcoming was particularly evident in early computer-generated literary texts, which were based primarily on this formalizable part of language: amusing experiments, which, however, look poor in comparison to great literary works.

Has scientific and technological progress changed anything since the 1960s? Have the boundaries of what can be calculated and formalized shifted?

In fact, the AI processes whose successes are currently causing a sensation do not pursue the goal of formalization at all, but rather follow a statistical approach. Here, it is not the handwritten algorithms that determine the result, but the data sets used for learning. The product of the program is therefore no longer predominantly something immanent to the algorithm, but something added from outside and contained in the training data. The cautious formulation suggests that it is difficult to draw a line here, because every AI project with the associated Preparation of the training data [7] and programming of the evaluation function contains unspoken assumptions and presuppositions.

There is no doubt that the statistical approach has proven to be much more successful than manual formalization. Language models such as ChatGPT, Bard or Jasper, which have been trained using unimaginably large amounts of text, can communicate with users in natural language. They often provide meaningful and correct answers to the questions they are asked.

7 See Section 11.1, "Supervised learning", and Section 13.6, "Summary and outlook".

Where are we now

The large language models not only answer questions on scientific and technical topics, they also provide recipes and ideas for organizing children's birthday parties. They can also compose stories and poems on given topics or write song lyrics. The results are sometimes amusing and astonishing, but only very few readers are likely to take them seriously as literature.

The weaknesses are not only apparent in literary attempts. Although the current language models are able to present *the familiar in an* appropriately structured and largely correct way, when it comes to more specific questions and more complex topics, they all too often deliver the wrong or at least the very ordinary. Surprising perspectives, unexpected ideas, original presentations and successful linguistic turns of phrase remain the exception and tend to give the impression of chance hits. The enumerative character dominates. If you want to discuss philosophical questions with ChatGPT, for example, you will usually receive answers that are so balanced that they may suffice for a school presentation if necessary, but rarely allow for deeper insights and are more likely to lead to cultivated boredom.

Whether literature or philosophy: the limits of current language models are obvious. The aesthetic possibilities of text and image generators can be well understood and described from the perspective of the philosopher Byung-Chul Han[8] . He states that digital capitalism promotes everything that can be processed in digital networks with the least possible resistance. A look at the content on Instagram, TikTok and similar platforms proves that these are short and isolated fragments. Whether it's a self-help phrase, a successful transition in a DJ set, a funny scene from a movie, a sporting trick, an artisanal tip or a political statement: all of this is easily consumable in the so-called social medium. Even grief and pain become a product here when influencers "generate reach" with posts about personal losses. The "other", the "unwieldy", the "strange" that can touch, surprise or disturb us in art does not take place here.

The aesthetics of AI is an aesthetics of fragments. Collage and remix are parts of artistic practice that should not be underestimated. And indeed, the visual creations of transformers like Midjourney are sometimes remarkable. But so far they have exhausted themselves in the remix aesthetic. This is very reminiscent of the above-mentioned first computer poems, which were also only able to depict a small part of linguistic possibilities.

8 See for example in "Transparenzgesellschaft" (Han 2012) and "Im Schwarm" (Han 2013).

In an interview with Deutschlandfunk radio, the literary and media scholar Joseph Vogl remarked in view of the disappointing results of the "production of pictures in the style of Raphael" that the "concept of style is an extremely conservative category under these conditions, just as perhaps this entire project is ultimately testing certain conservative criteria in this respect: Who is the author? How does the author write? How does he put subject and predicate together?". Vogl notes that this "reproduction of conservatism" is ultimately rooted in the AI system, as "these machines deal with probability operations. This means that they have a large memory [...] and they can manage Markov chains very well, i.e. the transition from one letter to another, from one word to another, from one sentence to another ... "[9, 10]

Thin ice

Anyone who tries to make predictions about future technical developments based on the current state of technology is treading on thin ice. One of the reasons for this is that past forecasts have all too often proved to be wrong: What the major language models are capable of today, hardly anyone would have thought foreseeable ten years ago. Despite all the justified criticism of the quality of the results: The fact that it works to mathematically abstract and reproduce characteristics such as "written by Thomas Bernhard" or "in the style of an advertising brochure from the 1970s" is highly astonishing and deserves admiration and attention.

Caution is also advisable because the question of the limits of AI very quickly puts us in the questionable position of defending humans against cold and unfeeling machines. This battle has proven to be a rearguard action, at least in recent decades. So I don't want to take this position, but I would still like to raise the question of the limits.

9 Joseph Vogl and Carsten Hueck: "Wie verändert KI die Literaturkritik" (Interview), Deutschlandfunk, 22.06.2023.

10 See Chapter 2, "Building texts with Markov", and Section 14.1, "An external view of a language model".

A lot helps a lot?

Two camps can be identified in the current debate on improving performance:[11]

Proponents of scaling conclude from experience that increasing the size of models and adding even larger training data sets has brought success in the past that this will also be the case in the future. So if we could train a model that contains as many weights as the human brain contains synapses, then this model could also achieve human intelligence.

Scaling *skeptics*, on the other hand, call for more basic research and assume that current models have fundamental, structural flaws that cannot be remedied by scaling. They criticize the fact that transformers cannot represent objects and do not contain a model of the world in which they operate: The imaginative action we are all familiar with is completely alien to them.

Whoever is right in the end: Further progress can be expected, whether through scaling, new model architectures or the gradual further development of existing models. What if, for example, more curated training data could distinguish between original and less original texts? Is it even possible to mathematically abstract and reproduce originality?

Lem's electric bard

In the short story "Trurl's Electro-Bard", Stanisław Lem recounts the adventure of the designer Trurl, who wants to build a machine "capable of writing flawless poetry". The project proves to be extremely tricky, because "the program that an average poet has in his head was created by the civilization in which he was born".

This ultimately makes it necessary to retrace this development "until the eve of creation, when the bits ... were still completely disorganized in primordial chaos ... were buzzing around". To be able to write poetry, Trurl's machine has to simulate a large part of the history of the universe.[12]

A similar problem, it seems to me, is faced by those who hope for literary works from an AI that are as inconclusively significant as, for example

11 Merkert and Bogartz 2022.

12 Lem 1985, pp. 47-62.

those of Franz Kafka. The author of a work is never its sole originator. Simulating everything that may have influenced the creation of a work may be conceivable *in principle*, but firstly it is *practically* impossible - and secondly, why should we even bother?

So are AI art dealt a bad hand? Is it stuck in collage, in parody, in reproducing surface effects? Not necessarily. It seems remarkable that the terms "art", "intelligence" - and "artificial intelligence" in any case! - similarly elude a precise definition. If we try, we all too quickly find ourselves enumerating technical and social practices. Perhaps this ambiguity may point the way?

From the point of view of art historian and critic Hanno Rauterberg, it is precisely the imperfections that open up aesthetic possibilities. "The saw is meant for sawing. The programmer's device, on the other hand, remains as underdetermined as a painter's canvas. There is no definite purpose [...] Nor is it inscribed in the digital machine what is to take place in it, to whom and how it is to be used."[13]

In fact, imitating and mimicking human abilities is only rarely the purpose of serious AI art. We can expect a large variety of things from art. Our expectations of machines are similarly diverse and changeable. Art and technology condition, shape and overlap each other. And algorithmic art is capable of reflecting and fantasizing about the existence of machines and humans in a media-technically constituted world as well as about the nature of art: "Artificial intelligence in a strong, world-penetrating version that includes social knowledge may still be a long time coming - but here, in the largely undefined field of aesthetics, algorithms are already able to succeed. Here, appearance counts far more than reality, and at least for the apparatuses are able to effortlessly stand in for an apparent illusoryness, an art of second-order simulacrum."[14]

This text follows on from the article "Artificial intelligence - the long road to art from the machine", which I wrote for heise online in 2019. I have used some of the wording from the article for this text.

13 Rauterberg 2021, p. 18 f.
14 Rauterberg 2021, p. 55 f.