Defining *Generative Artificial Intelligence*

An Attempt to Resolve the Confusion about Diffusion.

Raphael Ronge*†1, Markus Maier*‡1, and Benjamin Rathgeber§1

[1] *Department of Philosophy of Nature and Technology, Munich School of Philosophy, Kaulbachstraße 31a, 80539 München, Germany*

Abstract

The concept of *Generative Artificial Intelligence* (GenAI) is ubiquitous in the public and semi-technical domain, yet rarely defined precisely. We clarify main concepts that are usually discussed in connection to GenAI and argue that one ought to distinguish between the technical and the public discourse. In order to show its complex development and associated conceptual ambiguities, we offer a historical-systematic reconstruction of GenAI and explicitly discuss two exemplary cases: the generative status of the Large Language Model BERT and the differences between protein structure predictions from AlphaFold 2 and 3. Our analysis shows that there is no unique and unambiguous definition of GenAI based on a purely technical account of the term. Following this conclusion, we argue that the public discourse is not simply a less complex way of speaking, but instead transcends its technical basis. As a means to structure this newly emerging discussion landscape we introduce a non-exhaustive list of four central aspects of GenAI: *(multi-)modality*, *interaction*, *flexibility*, and *productivity*. These dimensions constitute a first step towards defining GenAI beyond its technical basis.

Keywords— Generative Artificial Intelligence, Generative Models, Definition of GenAI, Conceptual Analysis, History of AI, Philosophy of AI

1 Introduction

Since the second half of 2022 *Generative Artificial Intelligence* (GenAI) is on everyone's lips. But despite the ubiquitous use of this phrase, a precise definition of GenAI is rarely given. Rather, people are usually content with naming a few exemplary systems such as ChatGPT or Stable Diffusion. Is this because there is no need for a precise definition of GenAI or are there other reasons for this apparent conceptual vagueness? Are there any ambiguities in the discourse regarding GenAI and, if so, *what generates the confusion about Diffusion* (and other models)?

On closer inspection it seems as if the concept of GenAI often functions as a collective term, which is mostly used as an external description by laypeople or researchers from other disciplines such as philosophy (see, e.g., Zohny, McMillan, and King 2023), the social sciences (see, e.g., Bail 2024) or in overview articles aimed at a broader audience (see, e.g., Gozalo-Brizuela and Garrido-Merchan 2023). As a matter of fact, the term "Generative Artificial Intelligence" does not appear in any of the relevant technical publications – be it on GANs (Goodfellow, Pouget-Abadie, et al. 2014), VAEs (Kingma and Welling 2022), diffusion (Sohl-Dickstein et al. 2015; Ho, Jain, and Abbeel 2020) or other models. Instead, the researchers refer exclusively to *generative models*, but the exact relationship to GenAI remains unclear. This discrepancy between technical and public way of speaking seems to suggest that the concept of GenAI transcends the purely technical basis of generative models and includes further, functionally relevant characteristics of certain AI-systems.

In order to shed some light on the concept of GenAI and its usage, it is helpful to explain and clarify some common concepts that are often discussed in the context of GenAI and to locate where or rather *by whom* different concepts are used. As depicted in Fig. 1, the conceptual pairs we chose span a "discussion landscape" regarding the understanding and usage of the term GenAI.

*These authors contributed equally.

†raphael.ronge@hfph.de, ORCiD: 0009-0003-6133-1911

‡markus.maier@hfph.de, ORCiD: 0000-0003-1854-9416

§benjamin.rathgeber@hfph.de

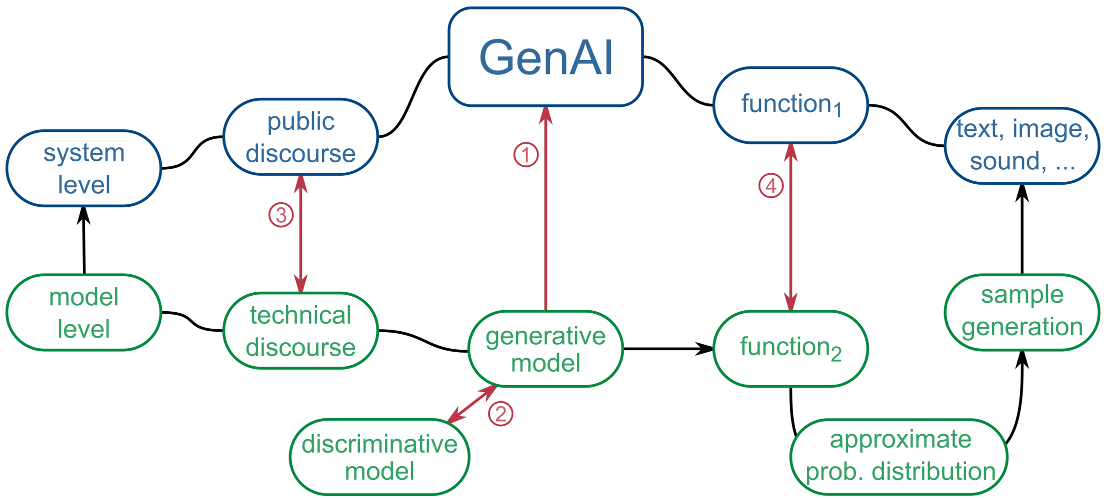

Figure 1: An overview of the current GenAI-discussion landscape. The numerically labeled pairs are discussed in the continuous text.

(1) Generative AI vs. Generative Model

The first distinction we want to point out is the one between GenAI on the one hand and generative models on the other hand. Technically speaking, a generative model is a statistical model that tries to approximate the underlying probability distribution given a set of data samples (observations) and can be used, in the case of GenAI, to create (*generate*) new, random instances of this learned probability distribution (cf. Ng and Jordan 2001). In a semi-technical setting – e.g., technically informed researchers writing for a broader or interdisciplinary audience – GenAI is usually understood to somehow refer to generative models (cf. Hacker, Engel, and Mauer 2023), but this relationship is often not specified further. Feuerriegel and colleagues define GenAI as "generative modeling that is instantiated with a machine learning architecture" and can therefore "create new data samples based on learned patterns." (Feuerriegel et al. 2024, 112)[1]

This definition attempt makes it apparent that generative modeling and GenAI are deeply intertwined, but not identical concepts – generative models are not limited to AI and have been discussed and used long before the concept of GenAI emerged. Still, it seems that GenAI necessarily includes some generative modeling capacity in the sense that the output of such a system has to meet the same criteria as the output of a generative model – namely novelty and usefulness. But, as we want to show in the following, this generative aspect is best understood only as a necessary (and not a sufficient) condition for these systems and furthermore might not even be identical to the technical definition of a generative model. In sum, it is fair to say that every GenAI has to produce new and useful output, but – as we will show in Chapter 2 below – the relationship to generative models is not as straightforward as it might seem.

(2) Generative Model vs. Discriminative Model

The distinction between generative model and discriminative model is a purely mathematical distinction of two different probabilistic models that describe different paradigms of AI research. However, this distinction is neither clear-cut nor exclusive to (generative) AI.

Discriminative models are often used for classification tasks, i.e., assigning data instances into different pre-existing categories. This includes well-known examples such as object recognition or spam detection (cf. Sarker 2021). As the underlying probability distribution for this task is relatively easy to model, AI quickly excelled at discriminative modeling and became as ubiquitous as it is today. Generative models, in contrast, can also be used to synthesise new data instances. Since good performance by machine learning models on generative tasks strongly correlates with larger networks and training data sets,[2] it is considerably more resource intensive than discriminative modeling (cf. Kaplan et al. 2020, 18-19).

In practice, "[m]any AI [models] employ both in tandem. In a generative adversarial network, for example, a generative model generates sample data and a discriminative model determines whether that data is *real* or *fake*. Output from the discriminative model is used to train the generative model until the discriminator can no longer

[1]Note that this is practically identical to the definition of a *Deep Generative Model* (see, e.g., Ruthotto and Haber 2021).
[2]In recent years, more powerful computers, increased dataset and model sizes and refined training has led to a tremendous growth of *Deep Learning* in general (cf. Goodfellow, Bengio, and Courville 2016, 18-26).

discern ‘fake’ generated data.” (IBM 2024) Notice that, in order to generate new data samples, a generative model also requires a probabilistic element: “If our model is merely a fixed calculation, such as taking the average value of each pixel in the dataset, it is not generative because the model produces the same output every time. The model must include a *stochastic* (random) element that influences the individual samples generated by the model.” (Foster 2019, 3)

(3) Public Discourse vs. Technical Discourse

There exists a tremendous difference between the public discourse on GenAI and its technical dimension. As already mentioned above, GenAI, Generative AI or similar terms do not appear in the technical literature, but are only used in semi-technical (e.g., in communication with other disciplines, public talks, lectures, outreach, etc.) or public domains whenever there is the need to communicate something *general* about systems such as ChatGPT, Stable Diffusion, etc. This distinction between a broad, public discourse and a narrow, technical discourse is partly mirrored by the distinction between the model level and the system level of GenAI as introduced in a recent paper by Feuerriegel and colleagues: Whereas the model level refers to the underlying technical implementation (e.g., GPT4), the system level refers to the embedding of the model functionality into a larger infrastructure (e.g., ChatGPT) (cf. Feuerriegel et al. 2024, 112).

According to Google Trends (2024), search volume for the term “generative AI” only really started in the second half of 2022 – around the same time ChatGPT and Stable Diffusion were launched and created a huge public debate about the capabilities of these systems. As we will show in Chapter 2, their technical basis – namely generative models instantiated with a machine learning architecture – has been around long before that. But in contrast to today, there was simply no need for such a term.

(4) Functions of GenAI vs. Functions of Generative Models

As mentioned above, the terms “Generative AI” and “generative models” are closely connected. Both terms describe models that (ultimately) produce new data points of textual, visual, auditory, ... data. However, the *generative* function that both terms include is a very different one. In the public discourse, *generative* in Generative AI refers to the generation of new text, images, sounds, or similar modalities (cf. Gartner 2024; Coursera 2024; Fruhlinger 2023) and GenAI is thus often characterized by means of its system level functionality: “Often, Generative AI models are marked by their wider scope and greater autonomy in extracting patterns within large datasets. In particular, LLMs’ capability for smooth general scalability enables them to generate content by processing a varying range of input from several domains.” (Novelli et al. 2024, 2)

The term *generative* in generative model on the other hand refers to its function of describing “how a dataset is generated, in terms of a probabilistic model. By sampling from this model, we are able to generate new data.” (Foster 2019, 2) The result of a generative model (by sampling) is typically a new data instance which ties it to *generative* of GenAI, but only indirectly and not necessarily (see also point (1) above).

The above description of the current discussion landscape on GenAI points to an important facet that should be kept in mind for the following analysis. There are two mechanisms of definition at work in the debate, referring to the *public* and the *technical* discourse. They are different in various respects, but they have complex interrelationships. In what follows, we will not simply take the position that the technical and public debates about GenAI are not identical, which would be a trivial statement. Rather, we argue that (1) the technical definition of GenAI is less clear and less unambiguous than one might think, and (2) that the public debate is not just an informal or less complex way of speaking, but that it actually transcends its technical basis regarding certain aspects. *It is one goal of this paper to offer a non-exhaustive list of aspects of GenAI that are central to the public debate, yet not captured by a purely technical discourse.*

To clarify this point, consider the following analogy: There is considerable public interest in the theory of general relativity (GR), as it is concerned with the evolution of the universe and our cosmological fate. It is good practice in introductory lectures or public talks to demonstrate certain characteristics of this theory – e.g., the curvature of space-time induced by the stress-energy tensor of an object – by invoking the image of a massive sphere on a rubber sheet. This is an attempt to visualize an important feature of GR and to convey a basic understanding of some of its important properties to laypeople (cf. Regt 2017, 91, 114). Even though this visualisation features a dramatically reduced complexity and some necessary errors and inaccuracies, it is still firmly rooted in the technical discourse. On the other hand, consider cosmological models based on GR implying that the universe is of finite age and started from an incredibly dense, primordial state which expanded subsequently. The theory does not allow one to make any statements about space-time *before* this so-called Big Bang – in fact, this expression in itself is meaningless. While it might be genuinely interesting to many people and even important for personal questions of meaning and worldview to consider a time before the Big Bang, GR (or the appropriate cosmological model) is silent on this topic – it is simply not in reach of the technical discourse. The public debate about GenAI, we argue, is best understood by analogy to the latter type of situation in the example above. While the technical discourse forms the basis of the public discourse, it does not touch on some relevant features of GenAI.

The rest of the paper is organized as follows. Since neither the technical nor the public debate is settled in itself there are changing definitions of terms, blurred boundaries and varying links between technical progress and its public perception. This is addressed in chapter 2, which gives an overview of the technical and conceptual history of GenAI. As the technical debate changes rapidly, partly due to innovations, there are always edge cases, two of which (BERT and AlphaFold) we will explicitly discuss in chapter 2.3. These case studies will support our argument that there are aspects of GenAI that are only insufficiently covered by a purely technical definition. In chapter 3, we propose a non-exhaustive list of four aspects that play an important role in the public debate about GenAI: *(multi-)modality*, *interaction*, *flexibility*, and *productivity*. In chapter 4 we conclude our analysis of GenAI and put it into a coherent picture with regard to current debates about creativity (of and with AI) and Artificial General Intelligence (AGI).

2 Historical-Systematic Reconstruction of GenAI

The history of GenAI can be told in many different ways. The most fruitful and interesting narratives for our intents and purposes are the conceptual and technical origins of GenAI. But even when we focus exclusively on these historical developments, no single common thread can be found. Rather, important turning points can be identified in retrospect among the many approaches, programs and ideas, on the basis of which a more or less coherent story can be told. This heterogeneity is still present today and is part of the problem that this paper attempts to resolve. The two case studies at the end of this chapter will analyse selected GenAI models to exemplify the complexity.

2.1 Technical Reconstruction

If one wants to define the term "Generative Artificial Intelligence" more precisely, it is necessary to analyze the tradition of the AI models that are nowadays mentioned in connection with GenAI. Two popular and impressive examples are Stable Diffusion from Stability AI[3] and the Generative Pre-trained Transformer (GPT) from OpenAI[4]. Although their success and public reception is unprecedented, AI models attempting something similar have been around for a long time. In what follows, we attempt to present a history of GenAI and how this technology has led to the current situation. This narrative is necessarily selective and does not claim to be exhaustive.

In general it can be observed that GenAI is not tied to a specific type of AI. The current success of GenAI can be traced back to Artificial Neural Networks (NNs). However, until the rise of NNs, various other methods were used with the aim of producing similar results.

The first attempts that can be associated with GenAI date back to the late 1960s and early 1970s. The first GenAI algorithms were all expert systems (ES). In other words, they were based on rules explicitly specified by humans. Unlike today's AI algorithms, this eliminates the need for training, which requires large data sets and appropriate learning techniques. On the other hand, ES are very limited in their capabilities. Programs designed to generate text can only generate pre-defined sentences or phrases. Drawing-ESs can only generate human-defined shapes, and ESs built for composition can only generate melodies according to predetermined rules. Between 1964 and 1966 Joseph Weizenbaum developed ELIZA, a chat program that simulated a psychotherapist (cf. Weizenbaum 2020, 14-15). Despite its very limited conversational capabilities, users were enthusiastic about its possibilities. Although, or perhaps because, it was "only" a computer program, users entrusted it with many personal details (cf. ibid., 19). At the same time, Harold Cohen programmed AARON, an ES for drawing images (cf. Cohen 1995). R. Kh. Zaripov devoted himself from 1959 to a program that could generate melodies based on predefined rules (cf. Zaripov and Russell 1969). The limitations of all three ESs are of a different nature. ELIZA can only generate answers for a limited area of knowledge and even then only on the basis of a certain set of rules (cf. Weizenbaum 2020, 15). If ELIZA cannot provide an answer, a question is generated that more or less repeats the input (cf. ibid., 15). For the domains of art and music, both systems had strict aesthetic rules, based on which they could compose elements (e.g., shapes and notes): "[The method] is based on the formalization and programming of certain laws of musical structure and rules of composition" (Zaripov and Russell 1969, 129). Unlike today, it is not possible for all three systems to adapt their output to descriptive inputs (so-called prompts) (e.g., *compose something in the style of Mozart*).

In the years that followed, technical progress in the field of GenAI was made primarily in the area of Natural Language Processing (NLP). Expert Systems were followed by models which have the advantage over Expert Systems that they are not limited to pre-defined phrases. Instead, they can learn the statistical sequence of words (and phrases) in a language. Based on ideas from the early 20th century, Hidden Markov Models (HMM),

[3]stability.ai
[4]openai.com

for example, became increasingly important for these NLP tasks from the 1980s onwards, thanks to computer technology that allowed for their training (cf. 589-590 Jiang 2010). The next great leap forward was achieved with the help of "long short-term memory" (LSTM) (cf. Hochreiter and Schmidhuber 1997). These models were the first NNs capable of generating text of any substantial quality. Although their structure is very different, they are similar to HMMs in that they generate text autoregressively. Both share this functionality with today's popular transformer models. LSTM models were able to solve problems of earlier NNs: Because of their additional "knowledge memory", they can store information from previous words for longer and process more complex relationships.

On the image generation side, the first major breakthrough came in 2014 with Generative Adversarial Networks (GANs) (cf. Goodfellow, Pouget-Abadie, et al. 2014). One reason for this delay in comparison to NLP is the sheer size and complexity of image data. Progress in this area has only been made possible by the computing power needed to train large NNs such as GANs. In the training process of GANs, two neural networks – the so-called generator and discriminator – are cleverly combined. The former network generates images from random noise, while the latter tries to distinguish these generated images from real training examples. Metaphorically speaking, the generator tries to outsmart the discriminator with its images. If you combine the training of both, they improve each other: the generated images become more and more similar to real images, as the discriminator also becomes better at distinguishing them. After training, the generator can now be used to generate new images from noise. A disadvantage of this type of image generation is that the generator has no incentive to be creative. The images that look most like the originals in the training dataset are the ones that are hardest to distinguish from the real thing. So there is no reason to generate something *new* (cf. Creswell et al. 2018, 63). This problem was solved in 2022 by so-called *Diffusion Models* (e.g. Stable Diffusion) (cf. Rombach et al. 2022). They also generate images, but work in a fundamentally different way. Unlike GANs, Diffusion Models do not attempt to generate an image from noise in a single step. Instead, they have been trained to extract (previously superimposed) noise from real images in many small steps. When given pure noise as input, the noise is reduced until an image is produced. This process can be controlled by text prompts, so that not just any image is produced, but, for example, an "astronaut on a horse" (ibid., 10687).

Just before the latest leap in innovation for image generation, there was a similar leap in the field of natural language tasks. In 2017, a team of researchers developed the so-called "Transformer" architecture (cf. Vaswani et al. 2017). The innovation in text processing introduced by Transformers is an "attention mechanism". Unlike its predecessors, HMM and LSTM models, this attention mechanism can map complex relationships within a text (between words and phrases), which is essential for successful text generation. Due to increased computing power, it is now possible to use questions as input, which are answered in real time by a transformer model (e.g. GPT).

2.2 Conceptual Reconstruction

The second historical path we want to follow for a successful analysis of GenAI is that of the concept itself. It is by no means the case that the term GenAI has always been used in its current form – term and meaning have only grown together over the years. This historical convergence still leads to inaccuracies in meaning today. As such, it is instructive to sketch the history of the term GenAI here.

The term "Generative Artificial Intelligence" as a whole only emerged in the mid-2010s. However, similar terms were used in similar domains much earlier. This is mainly due to generative models, which can be used for classification. The idea is as old as Bayes' theorem. Generative models, in contrast to discriminative models, do not directly model the probability of a class based on the observation, but model the joint probability distribution of class and observation, from which the class can then be calculated (using Bayes' theorem) (cf. Jebara 2004, 18-22). Accordingly, terms such as "generative models" and "generative learning" can also be found in connection with AI methods such as hidden Markov models and Bayesian networks, for example in 2001 (Ng and Jordan 2001) and 2004 (Jebara 2004). The connection between generative models and GenAI is not entirely clear. The term "generative model" explicitly describes a statistical learning method that makes no statement as to whether its output can also be used for classification, for example. However, GenAI is based on generative models and attempts to (approximately) model some ground-truth reality (its probability distribution) and not just calculate an observation-label probability distribution, as discriminative models do. This means that the terms operate in a similar space, but are not identical in wording and do not describe the same phenomena exactly.

To our knowledge, it was in his 2010 dissertation that van der Zant (2010, 2) used the term "Generative Artificial Intelligence" explicitly for the first time to refer to a clearly defined phenomenon. However, he describes something that has nothing to do with the current meaning and understanding of the concept. In his work, he describes an approach with which AI can no longer only be optimized in relation to one point, but generates new goals in order to improve itself dynamically (cf. ibid.). His understanding of the term is perhaps along the same lines as the dreams of some developers today that GenAI is a step in the direction of "Artificial General Intelligence". We will come back to the relationship of GenAI and AGI in Chapter 4. For the time being, it suffices to note that van der Zant's idea certainly does not describe what we understand by GenAI today.

Term and phenomenon had their first real point of contact in 2014 thanks to Goodfellow and colleagues with their development of Generative adversarial nets (GANs). The term "Generative AI" does not yet exist at that time, but these generative nets describe an AI approach which is capable of generating images of noteworthy quality for the first time (cf. Goodfellow, Pouget-Abadie, et al. 2014). This has familiarized people beyond the purely technical domain of computer science with the concept of image generation and the adjective *generative*.

A year later, the term "Generative Artificial Intelligence" reappeared in a publication by Tony Veale. The term is used to denote a text-generating AI approach in the context of Twitterbots (cf. Veale 2015). This is the first time the term appears in a context similar to our current understanding of GenAI and how we would use it today. At this point, the problem of why GenAI is so vaguely defined becomes clear. In fact, Veale is not concerned with a precise definition of GenAI, or even a delineation of GenAI as a field of computer science research. He uses GenAI only once, as a catch-all term to roughly define in the abstract why his work is interesting: "As generative AI systems grow in sophistication, so do our expectations of their output" (ibid., 5). His focus is not on computer science at all, but on what machine creativity means for metaphor generation.

The problem of lack of definition and precision persisted, although use was very sporadic at the time. Another example can be found in 2017 in an article on AI applications in astronomy (cf. Castelvecchi 2017, 17). Again, the focus is not on AI, but on astronomy and how image generation methods can support it. And again, GenAI is used as a catch-all term to summarize different approaches and illustrate where they can help astronomy – regardless of their exact design.

The fact that the term "Generative AI" was introduced from outside the field rather than from within it is also indicated by the fact that it does not appear in the most influential AI papers on the subject. The paper on GANs was mentioned above. Even the seminal 2017 paper on transformers, now synonymous with successful text generation, does not even mention GenAI (see Vaswani et al. 2017). Even in 2022, in the development paper of Stable Diffusion, one of the most successful image generation algorithms ever, only the technical basics of generative models are discussed, but the approach is not directly related to GenAI (cf. Rombach et al. 2022).

As we have shown, there is a conceptual history to the term "Generative AI" that is not identical to its technical history. From the beginning it was used in a loose and imprecise manner to denote general functionalities regarding the generation of certain outputs and not in a strictly technical way. One might think that with the already mentioned historical convergence there is now a unique and unambiguous meaning to GenAI, but this conclusion would be premature. If we want to talk about GenAI, it is important to be aware of the technical underlying structure. But it is – depending on the exact context – possibly as important to do justice to the multiple layers and complexities that are connected to GenAI which are not satisfied by a purely technical discourse. In order to refute the idea that the term GenAI has converged to a uniform and well-defined concept, we will in the following analyse two edge cases of AI systems that hint towards a pragmatic or pluralistic understanding of GenAI.

2.3 Exemplary Reconstruction

The genesis of the term makes it clear that GenAI is not congruent with a specific technical system or model, but in its current use represents a cross-section of existing technologies. In order to demonstrate the inadequacy of a purely technical definition of GenAI, we will in the following analyze two different case studies. The analysis of the Large Language Model BERT will demonstrate that it is not as straightforward as one might think to determine whether a specific system includes or constitutes a generative model in the technical sense. The second case study of AlphaFold 2 and 3 on the other hand suggests that even in cases where the technical dimension is unambiguous, there are reasons to extend the narrow, purely technical idea of GenAI in order to accommodate the actual practice and intuition of researchers working with these systems.

Case Study 1: BERT

BERT, which stands for Bidirectional Encoder Representations from Transformers (Devlin et al. 2019), was introduced by Google in 2018 and is one of the classical machine learning models in the context of NLP. After its release it quickly surpassed traditional models such as LSTM models and constituted the state-of-the-art system for a variety of NLP related benchmarks (cf. Rogers, Kovaleva, and Rumshisky 2020, 842). BERT is sometimes considered as an early example of GenAI, but at first glance it does not seem to include a generative component in a technical sense. While it shares some central architectural features with current LLM models such as GPT (e.g., both are transformer-based models), it is – as its name already suggests – *bidirectional*, meaning that it considers both left and right context of a given token. Unlike GPT, which predicts each token unidirectional based on the previous token, it is not autoregressive. Indeed, the vanilla BERT model was not built to *generate* text, but to excel at other NLP related tasks such as classification, understanding or translation.

Despite this fact, researchers quickly noticed that it is still possible to use BERT to generate text without changing its core architecture. This is surprising, since BERT – by design – is a so-called "encoder-only" system, which should not be able to generate text in any straightforward manner (cf. Footnote 4 in Devlin et al. 2019). A possible explanation for this unintentional functionality is that in order to train a bidirectional model such as

BERT, the developers relied on introducing a "Masked Language Model" (MLM) during pre-training. For technical details on MLMs, we refer the interested reader to the relevant literature (see, e.g., Chapter 11 in Jurafsky and Martin 2023). The important point to be made in the context of our argument is that there is an ongoing debate in the literature if – and how – BERT (or MLMs in general) could be represented and understood as generative models in disguise (cf. Wang and Cho 2019; Goyal, Dyer, and Berg-Kirkpatrick 2022; Torroba Hennigen and Y. Kim 2023).

Apparently, the seemingly easy task to decide if a specific system includes or constitutes a generative model is far from trivial. Against this background, it is questionable if it makes much sense to exclusively rely on a narrow and technical definition of GenAI based on generative models. It is unquestionable that the technical dimension constitutes the basis of any GenAI, but if a system has the necessary functionality to produce novel and useful outputs – to generate – it is irrelevant for most people if it includes a generative model in the technical sense. In order to emphasise this point, we will in the following consider the example of AlphaFold 2 and 3.

Case Study 2: AlphaFold

The AlphaFold project revolves around the central problem of predicting the three-dimensional structure of proteins from their amino-acid sequences. The problem of protein folding is by no means new. Since 1994 there exists a biennial competition for protein structure prediction – the Critical Assessment of Structure Prediction (CASP) –, which has been dominated by deep-learning-based approaches in recent years. While AlphaFold is not the only such approach to the prediction of tertiary structure of proteins (see, e.g., Baek et al. 2021; Callaway 2022), it is of particular interest for our purposes: While AlphaFold 2 predicts the three-dimensional structure of proteins from their amino-acid sequences, AlphaFold 3 is "capable of predicting the joint structure of complexes including proteins, nucleic acids, small molecules, ions and modified residues" (Abramson et al. 2024, 493). On a technical level, AlphaFold 3 adds a diffusion model to the underlying structure of AlphaFold 2. This added diffusion model leads AlphaFold 3 and structurally similar models (cf. Lin and AlQuraishi 2023; J. Lee, J. Kim, and P. Kim 2023), as opposed to AlphaFold 2, to be labeled as GenAI (see, e.g., Erdmann, Schumann, and von Lindern 2024; Oldfield 2023; Schwaller 2023).

Unlike with BERT, there is no technical debate about the generative model structures of AlphaFold 2 and AlphaFold 3. However, this does not mean their differences in generative capabilities would be without potential confusion. For the debate among practicing researchers of biology and medicine it is mostly irrelevant if a GenAI model was only added in AlphaFold 3. In personal communication, researchers confirmed this intuition and stated that AlphaFold 2 already *generated novel and useful output*s for their intents and purposes. The upgrade to AlphaFold 3 is not to be downplayed and, as mentioned above, adds additional capabilities that were not present on AlphaFold 2, but it is in the perception of practicing scientists not necessarily a paradigm shift with respect to its generative capabilities but rather an extension of the already existing generated content. The same phenomenon is observable in research papers on the use of AlphaFold 2, where the authors write that AlphaFold 2 "generates models of protein structures" (Thornton, Laskowski, and Borkakoti 2021, 1666) or that AlphaFold 2 "may have [...] success in *novel* [emphasis added] protein design" (Laurents 2022, 4).

This reference to the generative capabilities of models without a component that is usually associated with GenAI points to a confusion in the debate that is independent from technical ambiguities. Above, we have shown that the technical configuration of GenAI plays an important role as the basis for GenAI models and systems. It is the starting point of the current hype and debate surrounding GenAI. However, we have also shown that there is no clear effect relationship between technical basis and (use of the) term. In the following chapter, we attempt to resolve some of this confusion by introducing a set of relevant aspects for the public discourse about GenAI.

3 Aspects of GenAI

As we have shown in the previous chapters, the technical and public discourse about GenAI are not isomorphic. Rather, the public dimension transcends the technical one in certain important aspects. It is our goal in this chapter to analyse this incongruity and to propose a set of four aspects – *(multi-)modality*, *interaction*, *flexibility* and *productivity* that help to structure this debate.

These aspects are a first step towards *defining* GenAI. Importantly, we do not claim that this list is exhaustive or that the four aspects are exclusive. Rather, we propose these different dimensions in order to span the public discussion landscape on GenAI and help to allocate oneself in this abstract space. As such, these aspects pick out important and irreducible characteristics of GenAI and thus help to structure the debate and to demarcate GenAI from other types of AI.

3.1 (Multi-)Modality

The modalities of GenAI significantly determine the debates about it. Looking at the public and semi-technical discourse, three modalities are commonly associated with GenAI: text (language or code), image (or video) and sound (music or speech). In addition, authors usually provide little technical information, preferring to mention a few well-known exemplary models such as GPT or DALL-E (cf. Singleton 2024; Coursera 2024; Eliassi-Rad 2024; Satariano 2024). The article on regulation of LLMs by Hacker, Engel and Mauer (2023, 1113-1114) is an interesting exception, as they try to discuss the technical foundations of GenAI models, but mostly stay at the system level or equate GenAI with generative models. In contrast, what is not mentioned in many articles and videos on the subject is a consideration of how these models work. This is a new twist in the AI debate in general, as discussions of AI have often focused on learning methods (as opposed to conventional algorithms) and the structure of AI models (e.g. NNs). Input and output modalities have been secondary. This emphasis has now shifted.

This phenomenon becomes even clearer when we look at AI models that are not part of the GenAI debate. Think back to our case study of AlphaFold 2 above: We showed that there is a relatively clear technical distinction between the non-generative AlphaFold 2 and the generative AlphaFold 3. However, for all intents and purposes of the practicing scientist, AlphaFold 2 *generates* new data – in this case, the 3D structure of proteins. Nevertheless, in the public debate it has not been considered to be GenAI. We argue that one reason for this fact is precisely because it does not generate text, images or sound. The same applies to other examples such as DeepBlue[5] and AlphaGo[6]. Both generate new data, but it is not text, image or sound data, but new game moves. The modalities that an AI model can process are obviously important in assessing whether it is GenAI, as opposed to technical functionality.

Focusing on input modalities reinforces the point made above. Models typically considered as GenAI are applicable to their own output. It is possible to use their output as input for the next generation cycle. This is necessary for the conversational capabilities of text generation models, but image generators are also capable of adapting their generated images based on new prompts. This iterative self-application is a cornerstone of GenAI models and massively influences their interaction capabilities (see next chapter 3.2).

However, if we take a look at the history (see chapter 2), it becomes clear that text, image or sound output was not an automatic prerequisite for the description as GenAI. On the contrary, there were AI models that produced results similar to GPT or DALL-E long before the concept of GenAI (see ELIZA (cf. Weizenbaum 2020), AARON (cf. Cohen 1995), compositions by Zaripov (cf. Zaripov and Russell 1969)). At this point it could be argued that the difference between historical AI models and today's GenAI is *multi*modality. Today, output (whether speech, code, music, voices, images or video) can be influenced and modified at will by text prompts. Even AIs that were originally limited to text, such as the GPT series, can process visual input with the GPT4 version (cf. OpenAI 2023). However, these boundaries are fluid and quickly become blurred when moving away from large language or image models. The Visual Question Answering (VQA) tool by Malinowski and Fritz is a good example. The user can ask (via text promts) for the position of objects in a scene and gets a text based answer. Although VQA fulfils all the above conditions (images and text questions as input and text answers as output (cf. Malinowski and Fritz 2014)), it would hardly be classified as GenAI. The reason is that its input and output is very rudimentary and based on predefined categories (cf. ibid., 8). Thus, the classification as GenAI seems to not only be based on modalities of an AI. It is also the flexibility within and interaction with these modalities that is important for an AI model to be labeled as GenAI.

3.2 Interaction

An important feature that distinguishes GenAI from other types of AI is its intuitive interaction capabilities. Today, the AI systems that define the term "Generative AI" are particularly accessible to laypeople. Other types of AI are not excluded from intuitive interaction, but models that deal with humanly recognisable text, images and sound are predestined to be integrated into easily accessible systems.

The clearest example of this phenomenon are language models (see e.g. OpenAI 2023; Narang and Chowdhery 2024). There are two processes at work. First, language models are tasked with generating human-readable text as output (either in conversation, or as translations, etc.). Second, while they are not inherently capable of handling human-readable input, the better they can handle diverse, complex input, the more likely they are to be adopted. As mentioned in the previous chapter, this interaction process is supercharged if the AI's input and output are (partly) of the same modality, and the output of one generation cycle can be used as input for the next. This principle can also be seen in GenAI models that deal with sound (see e.g., Radford et al. 2023). Any model that works with speech input/output must, by definition, be able to work with human language. The ease of interaction for the layman is least obligatory for image generation tools (see, e.g., Rombach et al. 2022). The connection between textual input and visual output is not as direct as for language models. However, the second

[5] ibm.com/history/deep-blue
[6] deepmind.google/technologies/alphago

process of language models also applies to image generation: the easier it is to use, the more users there will be. An image generation tool that requires a tabular description of which parts of an image should be filled with which objects loses its appeal and even its usefulness.

These examples show an important circle that leads to easily accessible GenAI. First, as mentioned in the previous chapter, the modalities of AI models that are to be called GenAI are humanly understandable modalities. Secondly, this makes it easy for laypeople to interact with them. Thirdly, easy interaction contributes massively to their popularity. Finally, this popularity forces developers to support ever better interaction capabilities. For example, by broadening the scope of the model and building so-called "Foundation models" (see next chapter 3.3) that can work on a variety of tasks (cf. Bommasani et al. 2021, 3). This in turn increases the humanly accessible modalities (see step one) and their usefulness for everyday problems. While interaction with AI systems is not specific to GenAI, we argue that it is an important and central aspect of it.

3.3 Flexibility

One can observe an explosion in AI coverage – especially in relation to GenAI – in semi-technical and public domains (cf. Singleton 2024; Coursera 2024; Eliassi-Rad 2024; Satariano 2024). As we have seen in the historical overview, technological developments, while accelerating, are a gradual process. Technical progress brings new possibilities, but neither the ideas nor the application are particularly new. What has changed massively is the scope of the currently developed systems, which introduces an unprecedented *flexibility* for the users. New text-generating GenAI models are known as "Foundation Models", which are capable of simulating conversations on any conceivable topic. GPT and similar models can simulate the writing style of any author and are, at the same time, able to traverse complex math problems, play games, etc.

While interaction plays a crucial role in the accessibility of GenAI, earlier models such as ELIZA (cf. Weizenbaum 2020) were also able to simulate conversation in English. However, earlier models were much more limited in their knowledge. ELIZA had to be *retrained* (given a new set of rules) for conversations other than as a therapist. A similar trend can be seen with image-generating models. Although the first models in this area appeared much later, their limitations were not as pronounced as for text generation. Today, however, their flexibility is often even greater. Open source applications such as Stable Diffusion can be fine-tuned to generate images in niche art styles or in the styles of specific artists. This ties into our expectation of creativity with regard to these models. Today, they are capable of producing an unimaginable number of unique images and are no longer limited to generating photorealistic human portraits, as StyleGAN was a few years ago (cf. Karras, Laine, and Aila 2018).

This change is partly due to the increase in computing power that made these extensions possible. In large part, however, this new flexibility was a conscious design decision by the developers and led, among other things, to the creation of GPT: a Generative *Pretrained* Transformer.

3.4 Productivity

There is one fact about GenAI that is not directly related to the models or their inputs and outputs, but to the way people use these new systems for professional and personal tasks: productivity. The public debate is already predicting large productivity gains from the use of GenAI (c.f. Brynjolfsson, Li, and Raymond 2023; Chui et al. 2023). This is not to say that GenAI systems are the only ones capable of improving productivity. On the contrary, specialised AI models have been able to do this for years in their respective domains: predictive maintenance, spam filters, automated investments. In itself, GenAI is no more *productive* than other types of AI. There is no *a priori* difference in the way it produces its outputs or processes its inputs that would make it more productive than other AI models.

What makes GenAI a superior candidate for productivity improvement is a combination of the aspects we have introduced above, as they allow people to use GenAI systems in a variety of settings and to apply them to a variety of problems. Ease of interaction, (multi-)modality and flexibility support the foundational nature of the models and ensure their application in a diverse range of different contexts. As such, GenAI may be able to generate productivity gains (with varying intensity) in a variety of sectors (cf. Brynjolfsson, Li, and Raymond 2023). Thanks to this wide applicability, it is easy to imagine that GenAI systems will lead to productivity gains in almost every field of work "from banking to life sciences" (Chui et al. 2023), which might explain why productivity seems like a natural candidate to include in our set of relevant aspects.

It is interesting to observe that widespread productivity in turn requires greater flexibility and interactive capabilities of GenAI, which links it back to the other aspects we have introduced. In this sense, these four dimensions influence and enhance each other, which further substantiates our claim that they pick out some relevant structural characteristics in the current GenAI discussion landscape.

4 Conclusion

In the previous chapters we have tried to work out central distinctions and differentiations with regard to the concept of GenAI, both historically and systematically. We have identified various, sometimes contradictory, sources and contexts that are repeatedly used in the field and which potentially explain the prevailing *confusion about Diffusion* (and other GenAI models). We showed that there is not a single, unambiguous concept for understanding GenAI: It is not possible to find an exhaustive list of necessary and sufficient conditions for characterizing GenAI, as there are always counterexamples that elude a uniform characterization. A purely technical definition does not do justice to the multiple layers and complexities of the public discourse. This does not mean, however, that there is nothing to be said about this discourse and its different aspects. We identified four main features in order to structure the discussion landscape – *(multi-)modality*, *interaction*, *productivity* and *flexibility* –, which can be used to distinguish GenAI from other forms of AI and thus represents a first step towards *defining* GenAI.

These aspects potentially function as an explanation to the question of why GenAI is so fascinating and to the associated hype in the public domain. GenAI opens up a total field of possibilities that can be utilized both in a private as well as an economic context. As Foster pointed out, due to its generative capabilities, GenAI is not deterministic but randomly open for "individual samples generated by the model" (Foster 2019, 3). Because it is not only focused on one function but contains an unprecedented freedom to use it in different contexts, it approaches a status that seems qualitatively different from more traditional AI applications and hints for the first time towards an Artificial General Intelligence (AGI). Even though the concept of AGI and associated predictions of future technological development are sometimes considered as problematic or unscientific (cf. Altmeyer et al. 2024; Fjelland 2020), it is undeniable that (1) the prospect of such systems spark substantial public interest and that (2) current GenAI is the closest candidate for such systems.

In this field of (multi-)modality and multi-functionality, there is a variety of different applications that opens up the realm of productivity and creativity in an unprecedented manner for a large share of people around the globe. It combines a dynamic, user-oriented process with the effective creation of a useful and new product (cf. Zhou and D. Lee 2024). Importantly, this framing of a convergence of GenAI towards a state of AGI is far from fictional scenarios about dystopian nightmares of AI taking over control – or a utopia where AI solves all of humanities problems for that matter. Instead, it tries to focus the debate on the actual functionalities and potential risks of current GenAI systems and the way they already influence our private and working life.

The four aspects we have worked out can be described as a form of "collaborative creativity" (cf. Kaufman and Sternberg 2019, Part IV) between humans and AI systems. In this realm of possibilities and potentiality, GenAI is a unique and novel mode of operation and clearly detached from earlier reductionist misunderstandings (cf. Schmidhuber 2010). It is not just tailored to a specific process, but rather to a multi-functional range of applications which redefines the interface between humans and AI. It shapes a new intersection between the real and the digital world, a paradigm shift that has already arrived in the public discourse: "While each of these architectures – GANs, autoregressive networks, and diffusion models – sits on a bedrock of mathematical rigor, their real-world applicability extends far beyond sterile equations." (Corbeel 2023)

The aspects of *(multi-)modality*, *interaction*, *productivity* and *flexibility* we have proposed in this article are attempts to map out exactly this interface. As such, they help to clarify what function and relevance GenAI has beyond its technical basis.

References

Abramson, Josh et al. (2024). "Accurate structure prediction of biomolecular interactions with AlphaFold 3". In: *Nature*. DOI: 10.1038/s41586-024-07487-w.

Altmeyer, Patrick et al. (2024). "Position: Stop Making Unscientific AGI Performance Claims". In: *Proceedings of the 41st International Conference on Machine Learning*. Vienna, Austria.

Baek, Minkyung et al. (2021). "Accurate prediction of protein structures and interactions using a three-track neural network". In: *Science* 373.6557, 871–876. DOI: 10.1126/science.abj8754.

Bail, Christopher A. (2024). "Can Generative AI improve social science?" In: *Proceedings of the National Academy of Sciences* 121.21, e2314021121. DOI: 10.1073/pnas.2314021121.

Bommasani, Rishi et al. (2021). *On the Opportunities and Risks of Foundation Models*. arXiv: 2108.07258v3 [cs.LG]. URL: http://arxiv.org/pdf/2108.07258v3.

Brynjolfsson, Erik, Danielle Li, and Lindsey Raymond (2023). "Generative AI at Work". In: DOI: 10.3386/w31161.

Callaway, Ewen (2022). "AlphaFold's new rival? Meta AI predicts shape of 600 million proteins". In: *Nature* 611, 211–212. DOI: 10.1038/d41586-022-03539-1.

Castelvecchi, Davide (2017). "Astronomers explore uses for AI-generated images". In: *Nature* 542.7639, 16–17. DOI: 10.1038/542016a.

Chui, Michael et al. (2023). "The economic potential of generative AI: The next productivity frontier". In: *McKinsey & Company*.

Cohen, Harold (1995). "The further exploits of AARON, painter". In: *Stanford Humanities Review* 4.2, 141–158.

Corbeel, Brecht (2023). *What is the difference between GANs, autoregressive networks, and diffusion models?* Medium. URL: https://medium.com/@brechtcorbeel/what-is-the-difference-between-gans-autoregressive-networks-and-diffusion-models-8494410dc84e (visited on 06/04/2024).

Coursera (2024). *Generative AI Examples and How the Technology Works*. URL: https://www.coursera.org/articles/generative-ai-examples (visited on 04/23/2024).

Creswell, Antonia et al. (2018). "Generative Adversarial Networks: An Overview". In: *IEEE Signal Processing Magazine* 35.1, 53–65. DOI: 10.1109/MSP.2017.2765202.

Devlin, Jacob et al. (2019). *BERT: Pre-training of Deep Bidirectional Transformers for Language Understanding*. arXiv: 1810.04805 [cs.CL].

Eliassi-Rad, Tina (2024). *Just Machine Learning*. Santa Fe Institute. URL: https://santafe.edu/events/Community-Lecture-just-machine-learning (visited on 06/17/2024).

Erdmann, Elena, Florian Schumann, and Jakob von Lindern (2024). *Alphafold: Endlich eine KI, die wirklich hilft*. Zeit. URL: https://www.zeit.de/digital/2024-05/alphafold-google-kuenstliche-intelligenz-molekuele-dna (visited on 06/13/2024).

Feuerriegel, Stefan et al. (2024). "Generative AI". In: *Business & Information Systems Engineering* 66.1, 111–126. DOI: 10.1007/s12599-023-00834-7.

Fjelland, Ragnar (2020). "Why general artificial intelligence will not be realized". In: *Humanities and Social Sciences Communications* 7.10. DOI: 10.1057/s41599-020-0494-4.

Foster, David (2019). *Generative Deep Learning: Teaching Machines to Paint, Write, Compose, and Play*. 1st ed. Sebastopol CA: O'Reilly.

Fruhlinger, Josh (2023). *Was ist Generative AI?* Computerwoche. URL: https://www.computerwoche.de/a/was-ist-generative-ai,3614061 (visited on 04/23/2024).

Gartner (2024). *Generative AI: What Is It, Tools, Models, Applications and Use Cases*. URL: https://www.gartner.com/en/topics/generative-ai (visited on 04/23/2024).

Goodfellow, Ian, Yoshua Bengio, and Aaron Courville (2016). *Deep Learning*. Cambridge, Massachusetts: MIT Press.

Goodfellow, Ian, Jean Pouget-Abadie, et al. (2014). "Generative Adversarial Nets". In: Z. Ghahramani et al., eds. *Advances in Neural Information Processing Systems*. Vol. 27. Curran Associates, Inc.

Google Trends (2024). *Generative AI*. URL: https://trends.google.de/trends/explore?date=today%205-y&q=generative%20AI (visited on 06/05/2024).

Goyal, Kartik, Chris Dyer, and Taylor Berg-Kirkpatrick (2022). "Exposing the Implicit Energy Networks behind Masked Language Models via Metropolis–Hastings". In: *International Conference on Learning Representations*.

Gozalo-Brizuela, Roberto and Eduardo C. Garrido-Merchan (2023). *ChatGPT is not all you need. A State of the Art Review of large Generative AI models.* arXiv: 2301.04655 [cs.LG].

Hacker, Philipp, Andreas Engel, and Marco Mauer (2023). "Regulating ChatGPT and other Large Generative AI Models". In: *2023 ACM Conference on Fairness, Accountability, and Transparency.* New York, NY, USA: ACM, 1112–1123. DOI: 10.1145/3593013.3594067.

Ho, Jonathan, Ajay Jain, and Pieter Abbeel (2020). "Denoising diffusion probabilistic models". In: *Advances in neural information processing systems* 33, 6840–6851.

Hochreiter, Sepp and Jürgen Schmidhuber (1997). "Long short-term memory". In: *Neural computation* 9.8, 1735–1780. DOI: 10.1162/neco.1997.9.8.1735.

IBM (2024). *What Is an AI Model?* URL: https://www.ibm.com/topics/ai-model (visited on 06/04/2024).

Jebara, Tony (2004). *Machine Learning: Discriminative and Generative.* Vol. 755. The Springer International Series in Engineering and Computer Science. Boston, MA: Springer. DOI: 10.1007/978-1-4419-9011-2.

Jiang, Hui (2010). "Discriminative training of HMMs for automatic speech recognition: A survey". In: *Computer Speech & Language* 24.4, 589–608. DOI: 10.1016/j.csl.2009.08.002.

Jurafsky, Daniel and James H. Martin (2023). *Speech and Language Processing. An Introduction to Natural Language Processing, Computational Linguistics, and Speech Recognition.* Stanford University and University of Colorado at Boulder. URL: https://web.stanford.edu/~jurafsky/slp3/ed3bookfeb3_2024.pdf.

Kaplan, Jared et al. (2020). *Scaling Laws for Neural Language Models.* arXiv: 2001.08361 [cs.LG].

Karras, Tero, Samuli Laine, and Timo Aila (2018). *A Style-Based Generator Architecture for Generative Adversarial Networks.* arXiv: 1812.04948 [cs.NE].

Kaufman, James C. and Robert J. Sternberg, eds. (2019). *The Cambridge Handbook of Creativity.* 2nd ed. Cambridge: Cambridge University Press. DOI: 10.1017/9781316979839.

Kingma, Diederik P and Max Welling (2022). *Auto-Encoding Variational Bayes.* arXiv: 1312.6114 [stat.ML].

Laurents, Douglas V. (2022). "AlphaFold 2 and NMR Spectroscopy: Partners to Understand Protein Structure, Dynamics and Function". In: *Frontiers in molecular biosciences* 9, 906437. DOI: 10.3389/fmolb.2022.906437.

Lee, J.S., J. Kim, and P.M. Kim (2023). "Score-based generative modeling for de novo protein design". In: *Nat Comput Sci* 3, 382–392. DOI: 10.1038/s43588-023-00440-3.

Lin, Yeqing and Mohammed AlQuraishi (2023). *Generating Novel, Designable, and Diverse Protein Structures by Equivariantly Diffusing Oriented Residue Clouds.* arXiv: 2301.12485 [q-bio.BM].

Malinowski, Mateusz and Mario Fritz (2014). "A Multi-World Approach to Question Answering about Real-World Scenes based on Uncertain Input". In: Z. Ghahramani et al., eds. *Advances in Neural Information Processing Systems.* Vol. 27. Curran Associates, Inc.

Narang, Sharan and Aakanksha Chowdhery (2024). *Pathways Language Model (PaLM): Scaling to 540 Billion Parameters for Breakthrough Performance.* Google. URL: https://research.google/blog/pathways-language-model-palm-scaling-to-540-billion-parameters-for-breakthrough-performance/ (visited on 04/25/2024).

Ng, Andrew and Michael Jordan (2001). "On Discriminative vs. Generative Classifiers: A comparison of logistic regression and naive Bayes". In: T. Dietterich, S. Becker, and Z. Ghahramani, eds. *Advances in Neural Information Processing Systems.* Vol. 14. MIT Press.

Novelli, Claudio et al. (2024). *Generative AI in EU Law: Liability, Privacy, Intellectual Property, and Cybersecurity.* Working Paper. DOI: 10.2139/ssrn.4694565. URL: https://ssrn.com/abstract=4694565.

Oldfield, Jim (2023). *Researchers use generative AI to design novel proteins.* Science X. URL: https://phys.org/news/2023-05-generative-ai-proteins.html (visited on 05/23/2024).

OpenAI (2023). *GPT-4 Technical Report.* arXiv: 2303.08774 [cs.CL].

Radford, Alec et al. (2023). "Robust Speech Recognition via Large-Scale Weak Supervision". In: Andreas Krause et al., eds. *Proceedings of the 40th International Conference on Machine Learning.* Vol. 202. Proceedings of Machine Learning Research. PMLR, 28492–28518.

Regt, Henk W. de (2017). *Understanding Scientific Understanding.* New York: Oxford University Press.

Rogers, Anna, Olga Kovaleva, and Anna Rumshisky (2020). "A Primer in BERTology: What We Know About How BERT Works". In: *Transactions of the Association for Computational Linguistics* 8, 842–866. DOI: 10.1162/tacl_a_00349.

Rombach, Robin et al. (2022). "High-Resolution Image Synthesis With Latent Diffusion Models". In: *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*, 10684–10695. DOI: 10.1109/CVPR52688.2022.01042.

Ruthotto, Lars and Eldad Haber (2021). "An introduction to deep generative modeling". In: *GAMM-Mitteilungen* 44.2, e202100008. DOI: https://doi.org/10.1002/gamm.202100008.

Sarker, Iqbal H. (2021). "Machine Learning: Algorithms, Real-World Applications and Research Directions". In: *SN Computer Science* 2.160. DOI: 10.1007/s42979-021-00592-x.

Satariano, Adam (2024). *Mistral, a French A.I. Start-Up, Is Valued at $6.2 Billion*. The New York Times. URL: https://www.nytimes.com/2024/06/11/business/mistral-artificial-intelligence-fundraising.html (visited on 06/17/2024).

Schmidhuber, Jürgen (2010). "Formal Theory of Creativity, Fun, and Intrinsic Motivation (1990–2010)". In: *IEEE Transactions on Autonomous Mental Development* 2.3, 230–247. DOI: 10.1109/TAMD.2010.2056368.

Schwaller, Fred (2023). *Generative-AI: Dreaming up proteins*. Deutsche Welle. URL: https://www.dw.com/en/generative-ai-inventing-proteins-is-changing-medicine/a-66356415 (visited on 05/23/2024).

Singleton, Tom (2024). *New AI tools much hyped but not much used, study says*. BBC News. URL: https://www.bbc.com/news/articles/c511x4g7x7jo (visited on 06/17/2024).

Sohl-Dickstein, Jascha et al. (2015). "Deep Unsupervised Learning using Nonequilibrium Thermodynamics". In: Francis Bach and David Blei, eds. *Proceedings of the 32nd International Conference on Machine Learning*. Vol. 37. Proceedings of Machine Learning Research. Lille, France: PMLR, 2256–2265.

Thornton, Janet M., Roman A. Laskowski, and Neera Borkakoti (2021). "AlphaFold heralds a data-driven revolution in biology and medicine". In: *Nature Medicine* 27.10, 1666–1669. DOI: 10.1038/s41591-021-01533-0.

Torroba Hennigen, Lucas and Yoon Kim (2023). "Deriving Language Models from Masked Language Models". In: Anna Rogers, Jordan Boyd-Graber, and Naoaki Okazaki, eds. *Proceedings of the 61st Annual Meeting of the Association for Computational Linguistics (Volume 2: Short Papers)*. Toronto, Canada: Association for Computational Linguistics, 1149–1159. DOI: 10.18653/v1/2023.acl-short.99.

Vaswani, Ashish et al. (2017). "Attention is All you Need". In: Ulrike von Luxburg et al., eds. *Advances in neural information processing systems 30*. Vol. 30. Red Hook, NY: Curran Associates Inc.

Veale, Tony (2015). "Unnatural Selection: Seeing Human Intelligence in Artificial Creations". In: *Journal of Artificial General Intelligence* 6.1, 5–20. DOI: 10.1515/jagi-2015-0002.

Wang, Alex and Kyunghyun Cho (2019). "BERT has a Mouth, and It Must Speak: BERT as a Markov Random Field Language Model". In: Antoine Bosselut et al., eds. *Proceedings of the Workshop on Methods for Optimizing and Evaluating Neural Language Generation*. Minneapolis, Minnesota: Association for Computational Linguistics, 30–36. DOI: 10.18653/v1/W19-2304.

Weizenbaum, Joseph (2020). *Die Macht der Computer und die Ohnmacht der Vernunft*. 15th ed. Frankfurt am Main: Suhrkamp.

Zant, Colin Martijn van der (2010). *Generative AI: A neo-cybernetic analysis*. Groningen: s.n.

Zaripov, R. Kh. and J. G. K. Russell (1969). "Cybernetics and Music". In: *Perspectives of New Music* 7.2, 115. DOI: 10.2307/832298.

Zhou, Eric and Dokyun Lee (2024). "Generative artificial intelligence, human creativity, and art". In: *PNAS Nexus* 3.3, 1–8. DOI: 10.1093/pnasnexus/pgae052.

Zohny, Hazem, John McMillan, and Mike King (2023). "Ethics of generative AI". In: *Journal of Medical Ethics* 49.2, 79–80. DOI: 10.1136/jme-2023-108909.