



# Generative artificial intelligence

Leonardo Banh<sup>1</sup> · Gero Strobel<sup>1</sup>

Received: 26 June 2023 / Accepted: 7 November 2023 / Published online: 6 December 2023  
© The Author(s) 2023

## Abstract

Recent developments in the field of artificial intelligence (AI) have enabled new paradigms of machine processing, shifting from data-driven, discriminative AI tasks toward sophisticated, creative tasks through generative AI. Leveraging deep generative models, generative AI is capable of producing novel and realistic content across a broad spectrum (e.g., texts, images, or programming code) for various domains based on basic user prompts. In this article, we offer a comprehensive overview of the fundamentals of generative AI with its underpinning concepts and prospects. We provide a conceptual introduction to relevant terms and techniques, outline the inherent properties that constitute generative AI, and elaborate on the potentials and challenges. We underline the necessity for researchers and practitioners to comprehend the distinctive characteristics of generative artificial intelligence in order to harness its potential while mitigating its risks and to contribute to a principal understanding.

**Keywords** Generative AI · Artificial intelligence · Deep learning · Deep generative models · Large language models

**JEL Classification** C8 · M21

## Introduction

*“A groundbreaking fusion of data-driven creativity and artificial intelligence, poised to redefine the boundaries of innovation and transform the future of digital landscapes.”*

–ChatGPT on GPT-4, Mar 23 Version, on Generative AI

In an era where applications like ChatGPT set records for the fastest-growing user base by demonstrating unprecedented domain-independent expertise (Hu, 2023), the concept of “Generative Artificial Intelligence” (GAI) emerges as a disruptor in the digital landscape (Dwivedi et al., 2023; Teubner et al., 2023; Wessel et al., 2023). With capabilities to generate high-quality, contextually relevant content

almost indistinguishable from human-created work, discussions arise on whether this new technology even holds early signs of artificial general intelligence (Bubeck et al., 2023; The Washington Post, 2022). Regardless of discussions on AI’s sentience, the latest advancements in machine learning (ML) and deep learning (DL) have extended traditional, data-driven AI tasks such as predictions, classifications, or recommendations toward the generation of unique, realistic, and creative content. Prevalent collaborations between humans and intelligent systems in businesses and private life have been shaped by the adoption of AI in various ways, especially in the field of process optimization and decision-making (Brynjolfsson & McAfee, 2016; Burström et al., 2021; Moussawi et al., 2021). GAI addresses this development by providing novel augmentation and automation prospects in form of innovative services and business models (Huang & Grady, 2022; Mondal et al., 2023), e.g., by assisting customer support employees with suggestions of appropriate conversation responses (Brynjolfsson et al., 2023). The fast-changing and continuously evolving landscape of GAI calls for an extensive conceptualization of the properties and capabilities to fathom the phenomenon (Dwivedi et al., 2023; Strobel et al., 2024).

Responsible Editor: Christian Matt

✉ Leonardo Banh  
leonardo.banh@uni-due.de  
Gero Strobel  
gero.strobel@uni-due.de

<sup>1</sup> University of Duisburg-Essen, Universitätsstraße 2,  
Essen 45141, Germany

While GAI research and development is continuing to invest toward better, faster, and more capable models (e.g., Microsoft, 2023), studies on the fundamental principles, applications, and socio-economic impact remain largely unexplored in the academic discourse (Strobel et al., 2024; Susarla et al., 2023; Wessel et al., 2023). GAI provides innovation opportunities for various domains (e.g., networked businesses and digital platforms) but also comes with challenges (e.g., transparency, biases, and misuse) that need to be addressed for successful implementations (Houde et al., 2020; Schramowski et al., 2022; van Slyke et al., 2023). However, an examination of the key concepts is yet to be conducted, leaving a clear image and understanding of generative AI undefined. To overcome that shortcoming, this article provides an introduction to the fundamentals of generative AI, with its concepts, applications, and challenges. To do so, we exploratively synthesize recent literature on the technical foundation leading toward generative AI in combination with contemporary empirical examples of generative AI models and applications. Our aim is to conceptualize the key properties of GAI and differentiate them from ML and DL methods, to foster the understanding of the theoretical foundations of generative AI, and to guide further endeavors in examining as well as designing generative AI-based systems.

The remainder of the article is structured as follows: In the next section, we conceptualize generative AI and provide a distinction between related AI methods as well as outline the technological foundations. Afterward, we elaborate on the prospects and applications of the generative AI value chain and examine the impact of different generation modalities. Then, we address the potential challenges of adopting generative AI before concluding with a brief research outlook.


## Conceptualization

The field of artificial intelligence (AI) has taken a prominent place in research and practice across various disciplines for the past decades. Especially in information systems (IS) research, the socio-technical impact of AI as a phenomenon is at the core of investigation (Ågerfalk et al., 2022; Berente et al., 2021). These technologies have transformed the way we interact with data and make decisions, leading to uncharted ways in society as well as economy (Fügener et al., 2021; J. Li et al., 2021; van den Broek et al., 2021). However, as we venture into the new era of generative AI, it becomes increasingly crucial to understand the core concepts and distinctions within GAI as a rapidly evolving technology. To differentiate GAI from other AI concepts and provide a fundamental conceptualization, we will present a brief overview of AI and its subfields, machine learning (ML) and deep learning (DL), in the next section (Fig. 1). Afterward, we continue to elaborate on how DL has driven the development of deep generative models to enable distinct GAI characteristics and capabilities, ultimately leading to an even broader array of AI applications and opportunities for various fields.

## From artificial intelligence to deep learning

*Artificial intelligence* is considered an umbrella term, spanning over different computational algorithms capable of performing tasks that typically require human intelligence, such as understanding natural language, recognizing patterns, making decisions, and learning from experience (Castelvecchi, 2016; Winston, 1993). Early AI systems, such as expert systems and knowledge bases, were rule based and aimed at supporting users and businesses in decision-making

**Fig. 1** Generative AI and other AI concepts (inspired by Goodfellow et al., 2016, p. 9; Janiesch et al., 2021, p. 687)



(Harmon, 1985; Patterson, 1990). *Machine learning* as a subfield of AI deals with the development of algorithms capable to autonomously solve tasks through exposure to data without being explicitly programmed—i.e., learning (Brynjolfsson & Mitchell, 2017). In the realm of ML, there are several types of learning approaches based on the nature of the data and the desired outcome. *Supervised learning* is a common approach, e.g., for applications in commercial contexts, as algorithms are trained on labeled datasets to classify or forecast (business) data (Janiesch et al., 2021). The algorithm learns to map inputs to outputs and, thus, is capable of making predictions on new, unseen data. Moreover, *unsupervised learning* (i.e., discovering hidden structures or patterns within unlabeled data) and *reinforcement learning* (i.e., learning optimal decision-making by interacting with an environment and maximizing cumulative rewards over time through trial and error) are further learning strategies in ML (Kühl et al., 2022). What ML algorithms share in common are their *discriminative* properties, i.e., the goal of processing data to conduct classification, regression, or cluster and determine decision boundaries. Exemplary algorithms include decision trees, k-nearest neighbors, or support vector machines (Ray, 2019). *Deep learning* is a more advanced subset of ML and leverages artificial neural networks to model complex data representations and automatically detect correlations and patterns in large datasets (Janiesch et al., 2021; Samtani et al., 2023). Neural networks are computational models inspired by the structure and function of the human brain, consisting of interconnected layers of artificial neurons (Goodfellow et al., 2016). In DL, neural networks comprise of multiple hidden layers in a nested architecture to learn hierarchical feature representations from the data, leading to improved performance on various tasks. Thus, DL is capable of processing high-dimensional data in various domains, ranging from one-dimensional data like signals and texts to multidimensional data such as images, video, or audio (LeCun et al., 2015). These advances have enabled a plethora of use cases across different domains, from societal good, such as improving healthcare and environmental sustainability (Piccialli et al., 2021; Schoormann et al., 2023; Strobel et al., 2023), to electronic markets, where DL can optimize pricing, serve as recommendation systems, forecast demands, and detect fake consumer reviews (Ferreira et al., 2016; M. Li et al., 2022; Zhang et al., 2023b).

## Toward generative AI

Fueled by advancements in DL techniques, *deep generative models* (DGMs) have emerged as a class of DL models to generate new content based on existing data, creating a variety of new possibilities for AI applications (Lehmann & Buschek, 2020; Tomczak, 2022). These models are trained

to understand complex data distributions, which allows them to produce outputs that closely resemble real-world data. By leveraging statistics, the goal of DGM training is to learn high-dimensional probability distributions from a finite training dataset and create new, similar samples that resemble an approximation to the underlying class of training data (Ruthotto & Haber, 2021). While discriminative models focus on modeling the relationship between input features and output labels, *generative models* learn the inherent data structure and generation processes (Jebara, 2004). Generative models have been around for decades, with, for example, hidden Markov models or Bayesian networks aiming to model statistical problems involving time series or sequences (Gm et al., 2020). Nonetheless, DGMs relying on neural networks have paved the way for significantly higher-quality generated content in recent advancements in the field of so-called *generative AI*. Thus, the goals of DGMs differ from traditional discriminative AI models (e.g., in ML) because the focus lies on the probabilistic generation of new data instead of determining extant data's decision boundaries (e.g., classification, regression, or clustering) (Tomczak, 2022; Weisz et al., 2023). In the following, we will focus on DGMs as the underpinning of GAI and give an overview of four core DGMs that have shaped the evolution of GAI in Table 1.

To leverage DGMs in GAI applications, they can be trained to generate new data and enable a variety of use cases (we refer to DGMs implemented in GAI applications as GAI models). Training a GAI model can be different than a discriminative AI model due to *semi-supervised learning*, a combination of learning techniques leveraging a small amount of labeled data (i.e., supervised) followed by extensive unlabeled data (i.e., unsupervised) (Kingma et al., 2014). For instance, recent GAI models apply techniques like supervised fine-tuning (SFT), reward models, and reinforcement learning via proximal policy optimization (PPO) to achieve an alignment of the model with the developers' intentions and values (OpenAI, 2023; Ouyang et al., 2022). This unique approach allows the training of very large datasets required for GAI models without the need for difficult complete labeling.

The application system functions as an interface for the user to interact with a GAI model. *Prompting* is an interaction technique and unique GAI property that enables end users using natural language to engage with and instruct GAI application (e.g., LLMs) to create desired output such as text, images, or other types (Dang et al., 2022; Liu & Chilton, 2022). Depending on the application, prompts vary in their modality and directly influence the mode of operation. For instance, text-to-image applications use textual prompts describing the visuals of the desired image, while image-to-image applications rely on an input image to steer the generation process.

**Table 1** Overview of core deep generative models

Deep generative model	Description
Generative adversarial network (GAN)	<i>Generative adversarial networks</i> consist of two competing neural networks: a generator and a discriminator (Goodfellow et al., 2020). The generator creates realistic data samples, while the discriminator distinguishes between real and generated samples (Pan et al., 2019). Both neural networks are trained together until the discriminator is not able to differentiate both samples (Janiesch et al., 2021). This adversarial competition results in the generator improving its data generation capabilities over time, eventually producing high-quality, realistic outputs. Hence, GANs find various applications, for instance, in image generation and manipulation, object detection and segmentation, and natural language processing (Aggarwal et al., 2021; Gui et al., 2023)
Variational autoencoder (VAE)	<i>Variational autoencoders</i> employ a neural network to learn encoding compressed input data into a lower-dimensional latent space and then decode the data by reconstructing the original data from the latent space representation (Kingma et al., 2014). By optimizing a variational lower bound on the data likelihood in a probabilistic approach, VAEs can generate new samples that resemble the original data distribution. Typical use cases for VAEs can be seen in the synthetic generation and reconstruction of data such as images, in anomaly detection, and recommendation systems (Wei & Mahmood, 2021)
Transformer	<i>Transformer</i> models have become the basis for many state-of-the-art natural language processing tasks and succeeding models. They are a specific type of neural network architecture that employ self-attention mechanisms to capture long-range dependencies in the data, making them well-suited for large-scale language modeling tasks (Vaswani et al., 2017) <i>Generative pre-trained transformers</i> (GPT) build on the transformer architecture and were trained with large datasets of unlabeled data (Brown et al., 2020). Due to their large size (i.e., a very large number of trainable parameters), GPT trained on text data are often referred to as <i>large language models</i> (LLMs) (Schramowski et al., 2022). The goal of LLMs is to generate novel, coherent, contextually relevant human-like text by predicting which token is most likely to occur after the prior tokens in a sentence (Brown et al., 2020; H. Li, 2022). Hence, LLMs can serve as the foundation for conversational AI tools like ChatGPT (Teubner et al., 2023). Besides conversing, the large amount of information stored in LLMs can be used for text generation, writing, or even programming, e.g., to support scholars (Cooper, 2023; Lund et al., 2023)
Latent diffusion model (LDM)	<i>Latent diffusion models</i> are transformer based and build on the concepts of denoising score matching and contrastive divergence to learn a stochastic data generation process (Rombach et al., 2022). In LDMs, the generation process starts with a simple initial distribution, such as Gaussian noise. Then, the data gets gradually refined through a series of noise-reduction steps following a predefined diffusion process through a latent space (Ho et al., 2020). The key advantage of LDMs is their ability to learn complex data distributions without requiring adversarial training (as in GANs) or optimizing variational lower bounds (as in VAEs). They also feature improved stability over other DGMs during training to be less prone to issues like mode collapse (Kodali et al., 2017; Rombach et al., 2022), making them well-suited for high-quality and detailed outputs, such as high-resolution image synthesis (Ho et al., 2020)

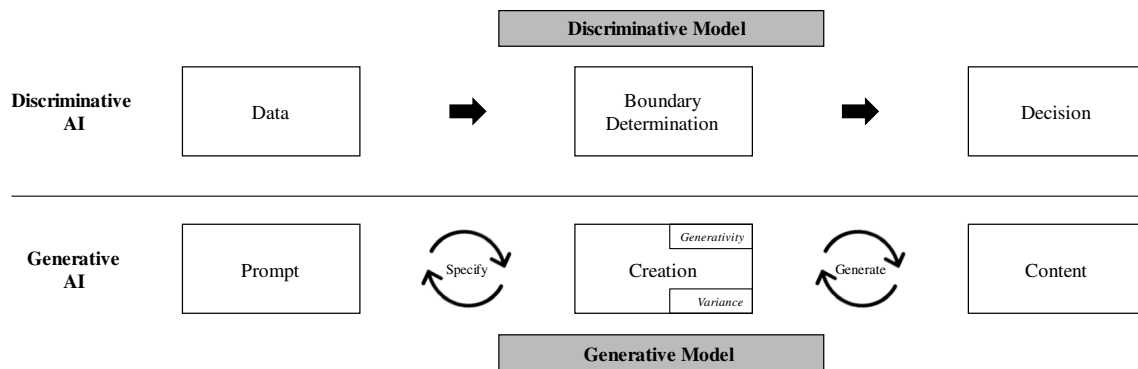
By design, the outputs of generative AI models are probabilistic and not replicable compared to the deterministic outcomes of discriminative AI—i.e., *variance* (Weisz et al., 2023). For one exact input prompt, a GAI application will generate varying outputs each time it is prompted, but the results remain valid and prompt fulfilling. On the other hand, different input prompts can lead to the same goal. Hence, formulating a meaningful prompt that leads to the desired outcome is based on trial-and-error process, e.g., by rephrasing textual prompts with the same keywords. The field of *prompt engineering* deals with systematically constructing prompts to improve the generated outputs (Liu & Chilton, 2022).

Based on the heuristic approach of prompt engineering and the inherent variance in the generated content, GAI users continuously and iteratively specify their desired tasks as input prompts to generate outputs until their task is

solved. The primary goal of generating new, probabilistically produced data (i.e., *content*) with varying outputs based on the same input distinguishes generative AI from discriminative AI, which pursues *boundary determination* by analyzing data and making a *decision* (see Fig. 2). Hence, a primary difference lies in the role of data, as GAI leverages very large datasets in its generative model to produce diverse content, while discriminative AI processes user data based on a (pre-trained) algorithm.

## Prospects and applications of generative AI

Complementing discriminative AI, GAI has recently emerged as a novel tool with a wide range of new possibilities impacting multiple sectors, from education and healthcare (Brand et al., 2023; Burger et al., 2023; Cooper, 2023)



**Fig. 2** Procedural differences of generative AI and discriminative AI

to networked businesses (Dwivedi et al., 2023; Wessel et al., 2023). These emerging applications inherit the generativity and variance properties of GAI and, therefore, are capable of producing unique and creative content, going beyond mere assistance. Hence, GAI becomes increasingly multidisciplinary, enabling disruptive innovations and automating even traditionally creative tasks, e.g., by generating customized contextual texts or images, facilitating new opportunities for businesses to innovate and differentiate themselves in the competitive economic landscape (Dwivedi et al., 2021; Lund et al., 2023; Pavlik, 2023).

Generative AI finds its utility across various modalities, including the generation of text, image, video, code, sound, and other produced content, such as molecules or 3D renderings (see Table 2). For example, GAI applications aim to create tailored marketing content, generate realistic (product) images or videos, and even assist in software development by generating code (Bakpayev et al., 2022; Elasri et al., 2022; Kowalczyk et al., 2023). Several modalities can serve as the input for GAI models. Distinguishing the different modality types, *unimodal* models generate the same output type as their input type, e.g., text-to-text or image-to-image generation, whereas *multi-modal* models combine different input and output types, for instance, in a text-to-image or code-to-text scenario. Different multi-modal models can subsume as *x-to-modality* models (e.g., x-to-text or x-to-image).

Examining the architecture of GAI-based systems, three major component layers can be identified: *model layer*, *connection layer*, and *application layer* (see Fig. 3). These parts embed generative AI in its information systems context and draw a boundary from external entities that can be interacted with its *environment* (e.g., users, organizations) and *data* (i.e., public and enterprise data) (Samtani et al., 2023). Inside the boundaries of GAI-based systems, the prevalent characteristics of generativity and variance persist and affect all layers and processes. The *model layer* comprises the pre-trained, deployable GAI artifact (i.e., a DGM) for

application systems. Depending on the training dataset, *general purpose* models aim at solving a wide range of tasks in multiple domains (e.g., GPT-4), whereas *customized* models are designed for domain-specific tasks and were, therefore, trained on highly specific data (e.g., CodeBERT).

Integrating these models into a system environment that affects people and organizations leads to *the application layer* of generative AI. By providing a proper context for the artifact, users are able to leverage the capabilities of GAI models for a specific application use case. Observing the trend of various recently emerging GAI applications building on top of existing models, it becomes apparent to further distinguish between *end-to-end applications* that are based on undisclosed, proprietary models (e.g., Midjourney) and *open applications* that are built around open-source models or leverage publicly accessible pre-trained models (e.g., Jasper and Elicit using OpenAI's GPT-3 (Elicit, 2022; Jasper, 2022)). Huang and Grady (2022) describe GAI applications as a “UI layer and ‘little brain’ [i.e., application layer] that sits on top of the ‘big brain’ that is the large general-purpose models [i.e., model layer].” This perspective emphasizes that new business models and applications can be developed without the need to train large GAI models from scratch by leveraging publicly available application programming interfaces (API) or AI-as-a-service platforms (Burstström et al., 2021; Janiesch et al., 2021; Lins et al., 2021). Indeed, the accessibility and availability of pre-trained GAI models foster value co-creation and can be leveraged via a *connection layer* (e.g., Hugging Face). *Fully integrated* GAI systems, on the other hand, employ their own, custom-trained proprietary models. In many cases, *end-to-end* GAI applications represent *fully integrated* systems (e.g., GitHub Copilot), while *open* GAI applications leverage external models via APIs (e.g., Stable Diffusion). Overall, GAI models may aim for *general purposes* or *customized* tasks regardless their connection or application characteristics (see Fig. 3). To enrich GAI-based systems with additional data beyond their training state and the GAI boundary, external *data* sources

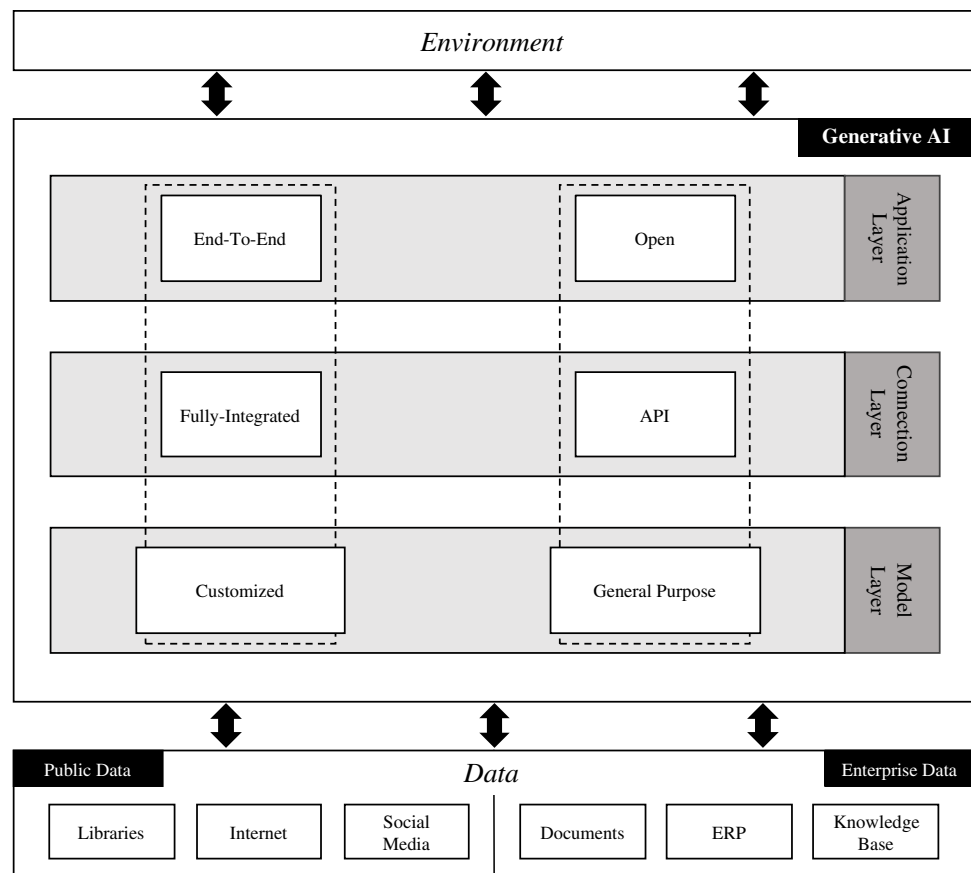


**Table 2** Overview of different output modalities for generative AI applications

Modality	Description
Text	X-to-text applications are centered around text generation and natural language processing. The goal is to generate human-like written text that fits the user's input prompt by providing a meaningful answer within the context. For instance, chatbots like OpenAI's ChatGPT imitate textual conversations with the user and can be guided to output text artifacts as desired (OpenAI, 2023). Furthermore, text-producing applications can be leveraged for content creation (e.g., copywriting or specific writing in e-commerce contexts) (Bakpayev et al., 2022; Brand et al., 2023). Moreover, text generation can support processes in sales or support by providing the ability to produce customized texts tailored toward the requests (Mondal et al., 2023). Systems integrating GAI models with further knowledge bases (e.g., enterprise data and Internet access) extend the available information beyond the model's initial training dataset
Image	X-to-image applications generate images based on the user's prompting. Relying on GANs or diffusion models as DGMs, synthetic images are created that find use cases in marketing, design and fashion, or creative fields in the form of new visual art (Haase et al., 2023; Mayahi & Vidrih, 2022; Zhang et al., 2023a). For instance, Stable Diffusion is an open-sourced x-to-image model that enables the generation of images in multiple GAI applications (Rombach et al., 2022). Moreover, generated synthetic images can act as training data for further ML models to train classifiers (e.g., medical images to detect diseases (Ali et al., 2023)). Besides a text-to-image creation process, image editing capabilities are possible, e.g., via image-to-image systems that manipulate and extend images according to the user's prompting (Openlaender, 2022)
Video	X-to-video applications deal with the creation of synthetic videos, i.e., dynamic motion images. New video clips are generated by describing the content of the desired video footage (text-to-video) or applying the style and composition via text or image prompt to a source video (video-to-video) (Esser et al., 2023). These prospects allow the fast and convenient creation and editing of videos via natural language and other modalities (Zhan et al., 2021). Thus, not only videographers benefit from x-to-video applications but also people without filming and editing skills are enabled to creatively express themselves due to an accessible creation process (Anantrasirichai & Bull, 2022). Besides recreational and entertainment purposes, x-to-video GAI models find application in sales and marketing (e.g., product marketing videos), onboarding and education (e.g., virtual avatars in training videos), or in customer support (e.g., how-to videos) (Leiker et al., 2023; Mayahi & Vidrih, 2022). As an exemplary application, Synthesia is a video creation platform specialized in generating professional videos with virtual avatars and synthetic voiceovers (Synthesia, 2023)
Code	In the realm of software development, x-to-code GAI applications offer transformative potential in how developers work and code by providing x-to-text capabilities specific to programming languages. Models like CodeBERT (Feng et al., 2020) or GraphCodeBERT (Guo et al., 2021) were trained on programming code to generate source code from natural language or modeling languages for new software programs. Several x-to-text models also offer coding capabilities because general-purpose LLMs are trained with increasingly large datasets that contain code (e.g., Stability.ai, 2023). Programmers using applications such as GitHub Copilot are supported by automatically written chunks of code, ideas converted into actionable scripts, auto-completion functions, generated unit tests, duplicate code detection, and bug fixing (Sun et al., 2022). These automation potentials allow developers to focus on higher-level tasks and problem solving, enhancing their productivity and the final product's overall quality, reducing time-to-market, supporting rapid prototyping, and promoting continuous innovation for the product and business
Audio	X-to-audio applications focus on audio content generation and comprise, for instance, the generation of speech with synthetically generated human-like voices (Borsos et al., 2022; Wang et al., 2023). Especially text-to-speech and speech-to-speech models are being heavily researched and can be used to power various applications, ranging from digital assistants and customer services to audiobook and training narration and accessibility tools (Moussawi et al., 2021; Qiu & Benbasat, 2005). GAI models like Microsoft's VALL-E (Wang et al., 2023) offer a more personalized and engaging user experience by enabling realistic voice modeling. Moreover, x-to-sound models find application in music creation. By specifying genres or melodies via prompts, unique pieces of music can be generated that respect the original intent (Agostinelli et al., 2023). GAI models such as MusicLM (Agostinelli et al., 2023) help musicians in their creative process, offering inspiration and aiding the composition of complex pieces. Businesses in the music industry can leverage high-fidelity music generation to create customized soundtracks for marketing, movies, or video games, significantly reducing the cost and time associated with traditional music production (Anantrasirichai & Bull, 2022; Weng & Chen, 2020)
Other	The applications of GAI extend beyond the stated modality types and domains, impacting multiple other, specific areas. For instance, x-to-molecules models like AlphaFold (Jumper et al., 2021) and OpenBioML (Murphy & Thomas, 2023) generate viable protein structures and design new molecules by generating valid, novel molecular structures, supporting drug discovery and bioengineering researchers (Walters & Murcko, 2020). 3D modeling is also impacted by GAI applications such as DreamFusion (Poole et al., 2023), Nvidia GET3D (Gao et al., 2022), and Point-E (Nichol et al., 2022), which generate realistic and complex 3D models that facilitate a range of applications from product design and architecture to virtual reality and game development

can be connected. Enterprise data (e.g., internal documents, enterprise resource planning (ERP) systems, knowledge bases) and public data (e.g., the Internet, libraries, social media) may serve as complementary, contextual data that GAI applications can further draw upon for more relevant and personalized results.

Employing GAI in enterprises can extend the level of assistance for workers and open up opportunities for augmentation and automation of the job, leading to new forms of collaborations between humans and machines (Einola & Khoreva, 2023). Furthermore, GAI transforms the way businesses operate in their daily tasks, innovate, and interact

**Fig. 3** Conceptual framework of generative AI

with their customers (Brynjolfsson et al., 2023; Mondal et al., 2023). Thus, the prospects of value co-creation come in hand with potential changes in human work roles, requiring workforces in various domains to adapt their tasks as a diverse set of tasks could be impacted by generative AI (Brynjolfsson & McAfee, 2016; Eloundou et al., 2023). The ongoing diffusion of AI into businesses gets accelerated by GAI applications, resulting in a possible replacement of human jobs on the one hand but also the creation of new jobs (e.g., for prompt engineers or with new business models) on the other hand (Einola & Khoreva, 2023). Hence, the effect on the labor market by the disruption needs to be discussed, and businesses should seek to understand and embrace the potential of generative AI (Eloundou et al., 2023; Willcocks, 2020).

### Challenges for generative AI-based systems

While generative AI holds transformative potential for individuals, organizations, and society due to its vast possible application space, the technology also inherits various challenges that parallel those of traditional ML and DL systems. The domain of electronic markets is a prime example that moved into the center of transformation

due to its latest focus on data-driven efforts (Selz, 2020). Outlining and emphasizing these challenges relevant for research and practice helps to raise awareness of the constraints as well as supports future efforts in developing, implementing, and improving GAI-based systems.

### Bias

Because of GAI's data-driven nature, data quality plays an essential role in how GAI-based systems perform and, thus, how feasible their adoption for real-world scenarios in business contexts is. Similar to their traditional discriminative AI relatives, GAI models are prone to bias causing biased decisions, disadvantages, and discriminations (Ferrara, 2023; Schramowski et al., 2022). Biases manifest in different ways and evolve primarily during two development phases of an AI-based system: training and inference.

*Data bias* gets injected during the model's training phase and leads to biased results because of faulty datasets. Factors such as non-representative, imbalanced sampling, incorrect labeling, and mismeasured features during the selection and processing of datasets hinder an unbiased training of the GAI model, ultimately leading to biased algorithmic outcomes (Mehrabi et al., 2022; Ntoutsis et al., 2020). The development of large-scale training datasets is especially

important for GAI models and often involves strategies of scraping public-available data on the Internet (Schuhmann et al., 2022). This approach is usually performed unsupervised and autonomously, which complicates the dataset's quality assurance because of its large quantity of unstructured data. Since GAI models are often trained to be general-purpose and multi-modal, they require and rely even more on such training datasets. Hence, moderating potential data bias is crucial for applications in business contexts like electronic markets due to the closeness to customers (e.g., points of contact via advertisements, social media, or customer support). Furthermore, social bias as a form of data bias can cause distorted views in generated texts or images and should be considered as well as mitigated (Baeza-Yates, 2018).

*Algorithmic bias* is introduced during the inference phase, independent from the model's training dataset (Mehrabi et al., 2022). In this case, the models have been trained on diverse, unbiased input data, and either the model's algorithm or the application around it introduces biases affecting users. Overfitting is a typical phenomenon that originates from the chosen learning strategies or optimization functions and causes biased algorithmic outcomes (Danks & London, 2017; Hooker, 2021). In this case, GAI models might introduce biases not reflected in the data because they fail to learn the data distribution correctly. Likewise, the presentation of and the user interaction with GAI-based systems can cause biases, such as when only selected generated content (e.g., one image out of multiple variants) is shown to the user (Baeza-Yates, 2018).

Thus, generative AI applications exerting biased results influence users' opinions and judgement and require control mechanisms (Jakesch et al., 2023a). Strategies should be developed to prevent, detect, and mitigate biases in order to safeguard users and ensure the service quality and reputations of a company. One approach to steer the quality of outputs from GAI models is via reinforcement learning from human feedback (RLHF) (Christiano et al., 2017; Griffith et al., 2013). The technique involves feedback from human evaluators to guide the model's training process, with evaluators assessing and comparing the quality of generated outputs. This approach enables generative models to refine their output generation process, aiming for better alignment with human expectations and objectives. Nevertheless, determining what content is "good" or "right" remains a difficult and bias-prone task (Teubner et al., 2023).

## Transparency

The need for explainability arises with the unpredictability of the inherent generative nature of GAI models and the overall functionality of ML models as "black boxes" (Janiesch et al., 2021; Meske et al., 2022). While the impact of GAI

may not reach as far as discriminative AI use cases (e.g., decision-making or dynamic pricing), research on explainable generative AI is still in its infancy, and the justification for more transparency is without a doubt (Brasse et al., 2023; Sun et al., 2022). Governments are already discussing the enforcement of AI regulations that include explainable AI to protect the general society and mitigate risks tied to the technology (Hamon et al., 2020). Interpretability (i.e., the human capability to understand the AI system's processes and decisions) is key, especially for GAI-based systems employed in large-scale information systems that affect large user groups, such as in networked businesses and digital platforms. In these cases, generated content has the potential to impact individuals and society, for instance, when generative AI serves as an advisor based on user questions and provides unsophisticated answers that are difficult to verify. Inaccuracy in generated product recommendations can have varying consequences depending on the situation, ranging from selecting the wrong product to taking the wrong medication. Early studies have shown how the chatbot ChatGPT performed surprisingly well in medical exams, suggesting inherent knowledge similar to medical students (Bhayana et al., 2023; Gilson et al., 2023). However, the seemingly omniscient capabilities may be restricted because, in the case of the medical exams, the GAI model might have been trained on the exam data and can reproduce its answers but is not able to comprehend the contextual state of an individual relevant for medical assessment. Therefore, understanding how the system performs sensemaking and generates its data helps users and businesses to achieve their goals responsibly and effectively, satisfying stakeholders' needs and expectations (Miller, 2019; Sun et al., 2022). Particularly for autonomous systems in critical business applications that interact with human beings, supervision and explainability of the GAI-generated content remain vital to ensure reliable, safe, and trustworthy outputs (Brasse et al., 2023; Hamm et al., 2023).

Another angle of transparency concerns the debate between open-source and closed-source models. Legal issues revolving around copyright, licenses, and intellectual property make it difficult for individuals and enterprises to deploy GAI-based systems, especially when the large training data of closed-sourced GAI models is procured through Internet scraping (Jin et al., 2023; Smits & Borghuis, 2022). Research initiatives revolving around open-source datasets (e.g., Schuhmann et al., 2022) and open-source models (e.g., Stability.ai, 2023) aim at increasing the transparency on data provenance and highlight, for instance, the data sources as well as the presence of watermarks on images (Schuhmann et al., 2022). Although most system engineers will rely on pre-trained models and perform fine-tuning for their specific use case, open-source efforts are facilitating a legally safer route for businesses to deploy GAI models, ensuring legal compliance, and mitigating associated risks.



## Hallucinations

Due to the probabilistic variance property of GAI, generative models are not immune to output errors, or so-called *hallucinations*, which manifest themselves in confidently generated results that seem plausible but are unreasonable with respect to the source of information (Ji et al., 2023; Susarla et al., 2023). The underlying causes of hallucinations are still being researched, with early findings suggesting that the training data that might contain contradictory or fictional content besides factual information (Dziri et al., 2022). This combination of varied inputs can lead to the generation of outputs that deviate from reality and introduce false information. As a result, the uncertainty of generation quality is further fueled by closed-source models that do not disclose any information on their training, making it crucial to carefully select appropriate datasets and models to mitigate the risk of hallucinations.

To illustrate the occurrence of hallucinations, studies have identified GAI-based image generators that facilitate anatomical inaccuracies in the generated images of humans (Choi et al., 2022; Hughes, 2023). These inaccuracies suggest that GAI models require further refinement and improvement before they can be reliably used for unsupervised production tasks (e.g., advertisement production, automated social media posts). Additionally, errors in simple arithmetic operations have also been observed (Bubeck et al., 2023), highlighting the limitations and potential shortcomings of current generative models in performing even basic computations accurately. Due to the seemingly realistic data produced by GAI models, the detection and evaluation of hallucinations are a challenging task. Current automatic evaluation includes statistical methods to measure discrepancies between ground-truth references and generated data and model-based metrics that leverage additional DL models to detect content inconsistencies (Ji et al., 2023). However, both approaches can be subjected to errors and are still inferior to cumbersome human evaluation. These instances emphasize the importance of appropriate hallucination mitigation methods, such as human supervision, to ensure the quality and accuracy of generated content.

Moving forward, addressing the issue of hallucinations in generative AI requires ongoing research and development efforts. Enhancing the transparency of training data and computation processes as well as promoting the adoption of open-source models can help mitigate the risk of generating misleading or flawed results. Furthermore, refining the underlying algorithms and incorporating robust error-checking mechanisms can contribute to the overall reliability and trustworthiness of GAI models (Zhou et al., 2023).

## Misuse

The access to GAI-based content creation tools with realistic outputs does not only enable new creative opportunities for good (e.g., novel automation and innovation prospects across multiple domains for businesses and users), but can also be leveraged for malicious purposes to intentionally cause risk and harm to society (Weidinger et al., 2022). Deepfakes have become increasingly sophisticated over the past decade as a result of the low cost and ease of creating such media using x-to-image, x-to-video, and x-to-sound GAI models (Mirsky & Lee, 2022; Vasist & Krishnan, 2022). They are authentic media content designed to impersonate individuals, such as celebrities or politicians, and are created to entertain or manipulate the viewers. For example, deepfakes depicting Trump's arrest have circulated on social media and in the news, causing misinformation due to their hyperrealistic, almost unrecognizable appearance (BBC, 2023). This is just one example of how generative AI carries potential for abuse and can be extended to other areas of daily life, such as fraudulent service offers, identity theft, or fake shops (Houde et al., 2020; Weidinger et al., 2022). The availability of GAI models provides starting points for new applications and business models for misuse and criminals, ultimately being leveraged to spread misinformation and influence the media and politics, or to defraud individuals and businesses (Hartmann et al., 2023; Kreps et al., 2022; Mirsky & Lee, 2022). Such issues, combined with low data quality and bias, provide a preview of potential social and ethical harms that may result in discrimination, exclusion, toxicity, and information hazards (Weidinger et al., 2022). This increases the urgency for disclosure and transparency of GAI models, along with the aforementioned calls for explainability (Brasse et al., 2023; Horneber & Laumer, 2023; Raj et al., 2023).

Generative AI researchers seek to develop measures for safer and more responsible use (van Slyke et al., 2023). RLHF and carefully crafted open-source datasets are first attempts at improvement, besides input filters restricting user prompts to harmless content. However, applications can still be tricked into bypassing filters and safeguards of GAI models, for instance, through prompt injections that insert malicious prompts to achieve misaligned outputs of generative AI applications (Perez & Ribeiro, 2022). The collaborative efforts between researchers, organizations, and regulators (e.g., by initiatives such as the European Union AI Act or US National AI Initiative) serve as promising initial steps toward opening pathways for future research to effectively address these issues, ensuring that AI-generated content is morally, ethically, and legally appropriate and cannot be misused (Hacker et al., 2023).

**Table 3** Future research questions regarding the challenges of generative AI-based systems

	Generative AI perspective	Environment perspective	Data perspective
Bias	<ul style="list-style-type: none"> <li>How do potential biases from GAI-generated content affect service offerings?</li> <li>What are implications of addressing bias in GAI models and data within electronic markets, and how can they be measured and managed?</li> </ul>	<ul style="list-style-type: none"> <li>How can we prevent GAI-based systems from perpetuating biases or discrimination against certain groups of people?</li> <li>What measures can be taken to ensure that GAI-generated content respects cultural values and societal norms?</li> </ul>	<ul style="list-style-type: none"> <li>How can GAI help identify and address bias in user-generated data?</li> <li>How can GAI be used to synthesize diverse and representative datasets that minimize bias in decision-making?</li> </ul>
Transparency	<ul style="list-style-type: none"> <li>How can we ensure that GAI-based systems are transparent and accountable?</li> <li>How and to what extent should GAI-based services declare the use of GAI technology?</li> <li>What measures can be taken to ensure that users are aware of when they are interacting with GAI-generated content, thus enhancing transparency?</li> </ul>	<ul style="list-style-type: none"> <li>How can GAI-based systems be designed to respect individual privacy while still delivering personalized content?</li> <li>How does transparency in GAI-generated content impact user trust and engagement within digital platforms and ecosystems?</li> </ul>	<ul style="list-style-type: none"> <li>How can platforms involving GAI transparently disclose the origins of data and preserve copyright?</li> <li>How can data transparency promote trust in GAI-based services, particularly when dealing with sensitive information?</li> </ul>
Hallucinations	<ul style="list-style-type: none"> <li>How can GAI-related hallucinations be mitigated on business strategy level?</li> <li>When and how should humans-in-the-loop be integrated in a GAI-based system to tackle hallucinations?</li> </ul>	<ul style="list-style-type: none"> <li>What is the impact of hallucinations toward consumer behavior?</li> <li>How do hallucinations in GAI-generated content affect user trust, engagement, and decision-making on digital platforms?</li> </ul>	<ul style="list-style-type: none"> <li>What role can explainable AI play in identifying and addressing hallucinations in GAI-generated content?</li> <li>How can data preprocessing and validation methods be improved to detect and mitigate hallucinations?</li> </ul>
Misuse	<ul style="list-style-type: none"> <li>Which impact does the misuse of GAI has on current digital platforms and electronic markets?</li> <li>How can we detect and prevent the misuse of (e-commerce) platforms by fraudulent providers and customers who leverage GAI?</li> <li>How can businesses develop effective strategies and regulations to prevent the misuse of GAI and protect the privacy and security of market participants?</li> </ul>	<ul style="list-style-type: none"> <li>What are the ethical implications of using GAI in electronic markets?</li> <li>How do instances of GAI misuse impact competition and market dynamics within platform ecosystems, and what regulatory frameworks can be established to address this?</li> </ul>	<ul style="list-style-type: none"> <li>How can GAI models be designed to flag and reject content that exhibits signs of potential misuse during the generation process?</li> <li>How can data sources be checked and authenticated to ensure that training data for GAI is not compromised or manipulated to encourage misuse?</li> </ul>
Societal impact	<ul style="list-style-type: none"> <li>Where should the boundary of liability be drawn when GAI generates false content?</li> <li>What are the potential implications of GAI-based systems on the future of work?</li> <li>How can GAI empower users to understand and control the content they interact with, promoting a more informed and empowered societal experience?</li> </ul>	<ul style="list-style-type: none"> <li>How can we ensure that GAI-based systems do not lead to worker displacement or other negative social impacts?</li> <li>How can GAI enable individuals to perform tasks and offer novel services they have not been trained for?</li> <li>How can businesses and platform ecosystems proactively engage with their user communities to address societal concerns regarding AI-generated content and adapt their practices accordingly?</li> </ul>	<ul style="list-style-type: none"> <li>How can GAI models continuously adapt to evolving societal norms, ethics, and cultural sensitivities, particularly in content generation and data handling?</li> <li>How can the generation of data be leveraged with GAI models to achieve greater goods for society?</li> </ul>

## Societal impact

With its ability to produce novel and diverse content, generative AI has significant implications for society in several areas. Besides the aforementioned risks of misuse and misinformation, the environmental footprint of developing and operating any AI-based system is an emerging concern and gets amplified for GAI-based systems due to the demanding technical requirements for training and running large DGMs and their high implicit energy consumption (Schneider et al., 2023; Schoormann et al., 2023). Consequently, the public acceptance of AI on a societal but also individual level inevitably guides the future development and use of GAI. Individually, researchers have investigated how AI can be perceived as a threat by some users (Johnson & Verdicchio, 2017; Lysyakov & Viswanathan, 2022). For example, a fear of employees is losing their jobs due to AI's automation potential (Mirbabaie et al., 2022). With GAI's capabilities of producing high-quality, customized realistic content, the frontier toward machines taking over human tasks moves even closer. Hence, trust plays a pivotal role in the adoption of emerging technologies and comes with direct impact between users and application providers (Lukyanenko et al., 2022; Riedl, 2022; Wanner et al., 2022; Yang & Wibowo, 2022). The promotion of explainable and responsible AI becomes paramount once more (Meske et al., 2022), especially in a technological state where distinguishing between generated and real content by GAI models is becoming increasingly more challenging (Jakesch et al., 2023b; Lehmann & Buschek, 2020). Early studies on generative AI suggest that generated content is perceived as less trustworthy, resulting in aversion toward generated products (Longoni et al., 2022; Rix & Hess, 2023). Nonetheless, GAI applications offer a variety of value propositions that revolve around building a trusting relationship toward such systems (Tomitza et al., 2023). For instance, social chatbots like Replika aim to realistically mimic or replace humans, leading to feelings of attachment toward machines among users (Pentina et al., 2023), whereas assistive GAI applications intent to support their users in the best way possible (Burger et al., 2023; van Dun et al., 2023).

While generative AI promises significant advances and has the potential to revolutionize domains, such as marketing, arts and culture, or even electronic markets as a whole, by supporting, augmenting, and automating a wide range of operations and offering novel services, it also underlines the importance of fully understanding its inherent challenges and mitigating possible risks. The far-reaching implications of GAI, ranging from poor data quality and hallucinations

to cases of societal misuse, require a proactive and sustained effort by researchers and practitioners to develop appropriate solutions ensuring the responsible and beneficial integration of generative AI technology into our digital society. In Table 3, we summarize the aforementioned challenges of GAI-based systems and suggest future research questions revolving around the GAI system, environment, and data perspective.

## Conclusion

With this fundamentals article, we provide an introduction to generative artificial intelligence. Drawing on ML and DL as the underlying technologies, we conceptualized generative AI and differentiated it from traditional AI methods. We outlined the most common DGMs as the theoretical foundation and described the core principles of generativity and variance. Moreover, we discussed the potential applications, acknowledging the distinct generation modalities and layers along the value chain. The paradigm shift of AI applications from discriminative to generative is leading to unique use cases and promising opportunities in various domains, including those traditionally resistant to automation. Therefore, researchers and practitioners need to understand the inherent properties of generative AI to effectively leverage its potential while also mitigating associated risks. We also considered five core challenges by elaborating on bias, transparency, hallucinations, misuse, and societal impact.

Generative AI carries the potential to significantly impact various industries. Taking the different layers of the value chain into account, this progression could further lead to the creation of new platforms and services centered around whole GAI ecosystems. It remains to be seen how these forms of unprecedented artificial creativity and generativity will find a place in the industry and everyday life. Future research will need to address the challenges for safety and responsibility measures, especially when employed in highly regulated or autonomous scenarios. For example, ensuring technological transparency, increasing public trust in GAI, and developing process models for employing GAI-based systems can contribute to the body of knowledge on generative AI. We believe this article provides an entry point to this novel type of technology and guides other researcher in their efforts in examining generative AI. Closing with the initial quote by ChatGPT, generative AI indeed holds transformative the potential to redefine innovation boundaries of the digital landscapes.

**Funding** Open Access funding enabled and organized by Projekt DEAL.

**Open Access** This article is licensed under a Creative Commons Attribution 4.0 International License, which permits use, sharing, adaptation, distribution and reproduction in any medium or format, as long as you give appropriate credit to the original author(s) and the source, provide a link to the Creative Commons licence, and indicate if changes were made. The images or other third party material in this article are included in the article's Creative Commons licence, unless indicated otherwise in a credit line to the material. If material is not included in the article's Creative Commons licence and your intended use is not permitted by statutory regulation or exceeds the permitted use, you will need to obtain permission directly from the copyright holder. To view a copy of this licence, visit <http://creativecommons.org/licenses/by/4.0/>.

## References

- Ågerfalk, P. J., Conboy, K., Crowston, K., Eriksson Lundström, J., S. Z., Jarvenpaa, S., Ram, S., & Mikalef, P. (2022). Artificial intelligence in information systems: State of the art and research roadmap. *Communications of the Association for Information Systems*, 50(1), 420–438. <https://doi.org/10.17705/1CAIS.05017>
- Aggarwal, A., Mittal, M., & Battineni, G. (2021). Generative adversarial network: An overview of theory and applications. *International Journal of Information Management Data Insights*, 1(1), 100004. <https://doi.org/10.1016/j.ijime.2020.100004>
- Agostinelli, A., Denk, T. I., Borsos, Z., Engel, J., Verzetti, M., Cailion, A., Huang, Q., Jansen, A., Roberts, A., Tagliasacchi, M., Sharifi, M., Zeghidour, N., & Frank, C. (2023). *MusicLM: Generating Music From Text*. <https://doi.org/10.48550/arXiv.2301.11325>
- Ali, H., Murad, S., & Shah, Z. (2023). Spot the fake lungs: Generating synthetic medical images using neural diffusion models. In L. Longo & R. O'Reilly (Eds.), *Communications in Computer and Information Science. Artificial Intelligence and Cognitive Science* (Vol. 1662, pp. 32–39). Springer Nature Switzerland. [https://doi.org/10.1007/978-3-031-26438-2\\_3](https://doi.org/10.1007/978-3-031-26438-2_3)
- Anantrasirichai, N., & Bull, D. (2022). Artificial intelligence in the creative industries: A review. *Artificial Intelligence Review*, 55(1), 589–656. <https://doi.org/10.1007/s10462-021-10039-7>
- Baeza-Yates, R. (2018). Bias on the web. *Communications of the ACM*, 61(6), 54–61. <https://doi.org/10.1145/3209581>
- Bakpayev, M., Baek, T. H., van Esch, P., & Yoon, S. (2022). Programmatic creative: AI can think but it cannot feel. *Australasian Marketing Journal*, 30(1), 90–95. <https://doi.org/10.1016/j.ausmj.2020.04.002>
- BBC. (2023). *Fake Trump arrest photos: How to spot an AI-generated image*. <https://www.bbc.com/news/world-us-canada-65069316>
- Berente, N., Gu, B., Recker, J., & Santhanam, R. (2021). Special issue editor's comments: Managing artificial intelligence. *MIS Quarterly*, 45(3), 1433–1450. <https://doi.org/10.25300/MISQ/2021/16274>
- Bhayana, R., Krishna, S., & Bleakney, R. R. (2023). Performance of ChatGPT on a radiology board-style examination: Insights into current strengths and limitations. *Radiology*, 307(5), e230582. <https://doi.org/10.1148/radiol.230582>
- Borsos, Z., Marinier, R., Vincent, D., Kharitonov, E., Pietquin, O., Sharifi, M., Teboul, O., Grangier, D., Tagliasacchi, M., & Zeghidour, N. (2022). *AudioLM: a Language Modeling Approach to Audio Generation*. <https://doi.org/10.48550/arXiv.2209.03143>
- Brand, J., Israeli, A., & Ngwe, D. (2023). Using GPT for market research. *Harvard Business School Marketing Unit Working Paper*. Advance online publication. <https://doi.org/10.2139/ssrn.4395751>
- Brasse, J., Broder, H. R., Förster, M., Klier, M., & Sigler, I. (2023). Explainable artificial intelligence in information systems: A review of the status quo and future research directions. *Electronic Markets*, 33, 26. <https://doi.org/10.1007/s12525-023-00644-5>
- Brown, T., Mann, B., Ryder, N., Subbiah, M., Kaplan, J. D., Dhariwal, P., Neelakantan, A., Shyam, P., Sastry, G., Askell, A., Agarwal, S., Herbert-Voss, A., Krueger, G., Henighan, T., Child, R., Ramesh, A., Ziegler, D., Wu, J., Winter, C., & Amodei, D. (2020). Language models are few-shot learners. In H. Larochelle, M. Ranzato, R. Hadsell, M. F. Balcan, & H. Lin (Eds.), *Advances in neural information processing systems 33* (pp. 1877–1901). Curran Associates Inc.
- Brynjolfsson, E., & McAfee, A. (2016). *The second machine age: Work, progress, and prosperity in a time of brilliant technologies*. W.W. Norton & Company.
- Brynjolfsson, E., Li, D., & Raymond, L. (2023). Generative AI at Work. Cambridge MA. <https://doi.org/10.3386/w31161>
- Brynjolfsson, E., & Mitchell, T. (2017). What can machine learning do? Workforce implications. *Science*, 358(6370), 1530–1534. <https://doi.org/10.1126/science.aap8062>
- Bubeck, S., Chandrasekaran, V., Eldan, R., Gehrke, J., Horvitz, E., Kamar, E., Lee, P., Lee, Y. T., Li, Y., Lundberg, S., Nori, H., Palangi, H., Ribeiro, M. T., & Zhang, Y. (2023). *Sparks of artificial general intelligence: Early experiments with GPT-4*. <https://doi.org/10.48550/arXiv.2303.12712>
- Burger, B., Kanbach, D. K., Kraus, S., Breier, M., & Corvello, V. (2023). On the use of AI-based tools like ChatGPT to support management research. *European Journal of Innovation Management*, 26(7), 233–241. <https://doi.org/10.1108/EJIM-02-2023-0156>
- Burström, T., Parida, V., Lahti, T., & Wincent, J. (2021). AI-enabled business-model innovation and transformation in industrial ecosystems: A framework, model and outline for further research. *Journal of Business Research*, 127, 85–95. <https://doi.org/10.1016/j.jbusres.2021.01.016>
- Castelvecchi, D. (2016). Can we open the black box of AI? *Nature*, 538(7623), 20–23. <https://doi.org/10.1038/538020a>
- Choi, H., Chang, W., & Choi, J. (2022). Can we find neurons that cause unrealistic images in deep generative networks? In R. Dechter & L. de Raedt (Eds.), *Proceedings of the thirty-first international joint conference on artificial intelligence* (pp. 2888–2894). International Joint Conferences on Artificial Intelligence Organization. <https://doi.org/10.24963/ijcai.2022/400>
- Christiano, P. F., Leike, J., Brown, T., Martic, M., Legg, S., & Amodei, D. (2017). Deep reinforcement learning from human preferences. In I. Guyon, U. von Luxburg, S. Bengio, H. Wallach, R. Fergus, S. Vishwanathan, & R. Garnett (Eds.), *Advances in neural information processing systems 30* (Vol. 30). Curran Associates, Inc.
- Cooper, G. (2023). Examining science education in ChatGPT: An exploratory study of generative artificial intelligence. *Journal of Science Education and Technology*, 32(3), 444–452. <https://doi.org/10.1007/s10956-023-10039-y>
- Dang, H., Mecke, L., Lehmann, F., Goller, S., & Buschek, D. (2022). How to prompt? Opportunities and challenges of zero- and few-shot learning for human-ai interaction in creative applications of generative models. In *Generative AI and HCI Workshop: CHI 2022*, New Orleans, LA. <https://doi.org/10.48550/arXiv.2209.01390>
- Danks, D., & London, A. J. (2017). Algorithmic bias in autonomous systems. In F. Bacchus & C. Sierra (Eds.), *Proceedings of the twenty-sixth international joint conference on artificial*



- intelligence (pp. 4691–4697). International Joint Conferences on Artificial Intelligence Organization. <https://doi.org/10.24963/ijcai.2017/654>
- Dwivedi, Y. K., Hughes, L., Ismagilova, E., Aarts, G., Coombs, C., Crick, T., Duan, Y., Dwivedi, R., Edwards, J., Eirug, A., Galanos, V., Ilavarasan, P. V., Janssen, M., Jones, P., Kar, A. K., Kizgin, H., Kronemann, B., Lal, B., Lucini, B., & Williams, M. D. (2021). Artificial intelligence (AI): Multidisciplinary perspectives on emerging challenges, opportunities, and agenda for research, practice and policy. *International Journal of Information Management*, 57, 101994. <https://doi.org/10.1016/j.ijinfomgt.2019.08.002>
- Dwivedi, Y. K., Kshetri, N., Hughes, L., Slade, E. L., Jeyaraj, A., Kar, A. K., Baabdullah, A. M., Koohang, A., Raghavan, V., Ahuja, M., Albanna, H., Albashrawi, M. A., Al-Busaidi, A. S., Balakrishnan, J., Barlette, Y., Basu, S., Bose, I., Brooks, L., Buhalis, D., & Wright, R. (2023). “So what if ChatGPT wrote it?” Multidisciplinary perspectives on opportunities, challenges and implications of generative conversational AI for research, practice and policy. *International Journal of Information Management*, 71, 102642. <https://doi.org/10.1016/j.ijinfomgt.2023.102642>
- Dziri, N., Milton, S., Yu, M., Zaiane, O., & Reddy, S. (2022). On the origin of hallucinations in conversational models: Is it the datasets or the models? In M. Carpuat, M.-C. de Marneffe, & I. V. Meza Ruiz (Eds.), *Proceedings of the 2022 Conference of the North American Chapter of the Association for Computational Linguistics: Human Language Technologies* (pp. 5271–5285). Association for Computational Linguistics. <https://doi.org/10.18653/v1/2022.naacl-main.387>
- Einola, K., & Khoreva, V. (2023). Best friend or broken tool? Exploring the co-existence of humans and artificial intelligence in the workplace ecosystem. *Human Resource Management*, 62(1), 117–135. <https://doi.org/10.1002/hrm.22147>
- Elasri, M., Elharrouss, O., Al-Maadeed, S., & Tairi, H. (2022). Image generation: A review. *Neural Processing Letters*, 54(5), 4609–4646. <https://doi.org/10.1007/s11063-022-10777-x>
- Elicit. (2022). *Frequently asked questions: What is elicit?* <https://elicit.org/faq#what-is-elicit>
- Eloundou, T., Manning, S., Mishkin, P., & Rock, D. (2023). GPTs are GPTs: An early look at the labor market impact potential of large language models. *arXiv*. <https://doi.org/10.48550/arXiv.2303.10130>
- Esser, P., Chiu, J., Atighehchian, P., Granskog, J., & Germanidis, A. (2023). *Structure and content-guided video synthesis with diffusion models*. <https://doi.org/10.48550/arXiv.2302.03011>
- Feng, Z., Guo, D., Tang, D., Duan, N., Feng, X., Gong, M., Shou, L., Qin, B., Liu, T., Jiang, D., & Zhou, M. (2020). CodeBERT: A pre-trained model for programming and natural languages. In T. Cohn, Y. He, & Y. Liu (Eds.), *Findings of the association for computational linguistics: EMNLP 2020* (pp. 1536–1547). Association for Computational Linguistics. <https://doi.org/10.18653/v1%2F2020.findings-emnlp.139>
- Ferrara, E. (2023). Should ChatGPT be biased? Challenges and risks of bias in large language models. *arXiv*. <https://doi.org/10.48550/arXiv.2304.03738>
- Ferreira, K. J., Lee, B. H. A., & Simchi-Levi, D. (2016). Analytics for an online retailer: Demand forecasting and price optimization. *Manufacturing & Service Operations Management*, 18(1), 69–88. <https://doi.org/10.1287/msom.2015.0561>
- Fügener, A., Grahl, J., Gupta, A., & Ketter, W. (2021). Will humans-in-the-loop become borgs? Merits and pitfalls of working with AI. *MIS Quarterly*, 45(3), 1527–1556. <https://doi.org/10.25300/MISQ/2021/16553>
- Gao, J., Shen, T., Wang, Z., Chen, W., Yin, K., Li, D., Litany, O., Gojcic, Z., & Fidler, S. (2022). GET3D: A generative model of high quality 3D textured shapes learned from images. In S. Koyejo, S. Mohamed, A. Agarwal, D. Belgrave, K. Cho, & A. Oh (Eds.), *Advances in Neural Information Processing Systems* 35. Curran Associates, Inc.
- Gilson, A., Safranek, C. W., Huang, T., Socrates, V., Chi, L., Taylor, R. A., & Chartash, D. (2023). How does ChatGPT perform on the United States Medical Licensing Examination? The implications of large language models for medical education and knowledge assessment. *JMIR Medical Education*, 9, e45312. <https://doi.org/10.2196/45312>
- Gm, H., Gourisaria, M. K., Pandey, M., & Rautaray, S. (2020). A comprehensive survey and analysis of generative models in machine learning. *Computer Science Review*, 38, 100285. <https://doi.org/10.1016/j.cosrev.2020.100285>
- Goodfellow, I., Bengio, Y., & Courville, A. (2016). *Deep learning*. The MIT Press.
- Goodfellow, I., Pouget-Abadie, J., Mirza, M., Xu, B., Warde-Farley, D., Ozair, S., Courville, A., & Bengio, Y. (2020). Generative adversarial networks. *Communications of the ACM*, 63(11), 139–144. <https://doi.org/10.1145/3422622>
- Griffith, S., Subramanian, K., Scholz, J., Isbell, C. L., & Thomaz, A. L. (2013). Policy shaping: Integrating Human feedback with reinforcement learning. In C. J. C. Burges, L. Bottou, Z. Ghahramani, & K. Q. Weinberger (Eds.), *Advances in Neural Information Processing Systems* 26 (Vol. 26). Curran Associates, Inc.
- Gui, J., Sun, Z., Wen, Y., Tao, D., & Ye, J. (2023). A review on generative adversarial networks: Algorithms, theory, and applications. *IEEE Transactions on Knowledge and Data Engineering*, 35(4), 3313–3332. <https://doi.org/10.1109/TKDE.2021.3130191>
- Guo, D., Ren, S., Lu, S., Feng, Z., Tang, D., Liu, S., Zhou, L., Duan, N., Svyatkovskiy, A., Fu, S., Tufano, M., Deng, S. K., Clement, C., Drain, D., Sundaresan, N., Yin, J., Jiang, D., & Zhou, M. (2021). GraphCodeBERT: Pre-training code representations with data flow. *9th International Conference on Learning Representations 2021 (ICLR)*, Virtual.
- Haase, J., Djurica, D., & Mendling, J. (2023). The art of inspiring creativity: Exploring the unique impact of AI-generated images. *AMCIS 2023 Proceedings*.
- Hacker, P., Engel, A., & Mauer, M. (2023). Regulating ChatGPT and other large generative AI models. *2023 ACM Conference on Fairness, Accountability, and Transparency* (pp. 1112–1123). ACM. <https://doi.org/10.1145/3593013.3594067>
- Hamm, P., Klesel, M., Coberger, P., & Wittmann, H. F. (2023). Explanation matters: An experimental study on explainable AI. *Electronic Markets*, 33, 17. <https://doi.org/10.1007/s12525-023-00640-9>
- Hamon, R., Junklewitz, H., & Sanchez, I. (2020). *Robustness and explainability of artificial intelligence: From technical to policy solutions*. EUR: Vol. 30040. Publications Office of the European Union.
- Harmon, P. (1985). *Expert systems: Artificial intelligence in business*. Wiley & Sons.
- Hartmann, J., Schwenzow, J., & Witte, M. (2023). *The political ideology of conversational AI: Converging evidence on ChatGPT's pro-environmental, left-libertarian orientation*. <https://doi.org/10.48550/arXiv.2301.01768>
- Ho, J., Jain, A., & Abbeel, P. (2020). Denoising diffusion probabilistic models. In H. Larochelle, M. Ranzato, R. Hadsell, M. F. Balcan, & H. Lin (Eds.), *Advances in Neural Information Processing Systems* 33 (pp. 6840–6851). Curran Associates Inc.
- Hooker, S. (2021). Moving beyond “algorithmic bias is a data problem”. *Patterns (New York, N.Y.)*, 2(4), 100241. <https://doi.org/10.1016/j.patter.2021.100241>
- Horneber, D., & Laumer, S. (2023). Algorithmic accountability. *Business & Information Systems Engineering. Advance online publication*. <https://doi.org/10.1007/s12599-023-00817-8>



- Houde, S., Liao, V., Martino, J., Muller, M., Piorkowski, D., Richards, J., Weisz, J., & Zhang, Y. (2020). Business (mis)Use Cases of Generative AI. In W. Geyer, Y. Khazaeni, & M. Shmueli-Scheuer (Eds.), *Joint Proceedings of the Workshops on Human-AI Co-Creation with Generative Models and User-Aware Conversational Agents co-located with 25th International Conference on Intelligent User Interfaces (IUI 2020)*. CEUR. <https://doi.org/10.48550/arXiv.2003.07679>
- Hu, K. (2023, February 2). *ChatGPT sets record for fastest-growing user base - Analyst note*. Reuters. <https://www.reuters.com/technology/chatgpt-sets-record-fastest-growing-user-base-analyst-note-2023-02-01/>
- Huang, S., & Grady, P. (2022). *Generative AI: A Creative New World*. Sequoia. <https://www.sequoiacap.com/article/generative-ai-a-creative-new-world/>
- Hughes, A. (2023). *Why AI-generated hands are the stuff of nightmares, explained by a scientist*. BBC Science Focus. <https://www.sciencefocus.com/future-technology/why-ai-generated-hands-are-the-stuff-of-nightmares-explained-by-a-scientist/>
- Jakesch, M., Bhat, A., Buschek, D., Zalmanson, L., & Naaman, M. (2023a). Co-writing with opinionated language models affects users' views. In A. Schmidt, K. Väänänen, T. Goyal, P. O. Kristensson, A. Peters, S. Mueller, J. R. Williamson, & M. L. Wilson (Eds.), *Proceedings of the 2023 CHI Conference on Human Factors in Computing Systems* (pp. 1–15). ACM. <https://doi.org/10.1145/3544548.3581196>
- Jakesch, M., Hancock, J. T., & Naaman, M. (2023b). Human heuristics for AI-generated language are flawed. *Proceedings of the National Academy of Sciences of the United States of America*, 120(11), e2208839120. <https://doi.org/10.1073/pnas.2208839120>
- Janiesch, C., Zschech, P., & Heinrich, K. (2021). Machine learning and deep learning. *Electronic Markets*, 31(3), 685–695. <https://doi.org/10.1007/s12525-021-00475-2>
- Jasper. (2022). *ChatGPT vs. Jasper: How it's different from Jasper chat*. <https://www.jasper.ai/blog/what-is-chatgpt>
- Jebara, T. (2004). Generative versus discriminative learning. In T. Jebara (Ed.), *Machine Learning* (pp. 17–60). Springer US. [https://doi.org/10.1007/978-1-4419-9011-2\\_2](https://doi.org/10.1007/978-1-4419-9011-2_2)
- Ji, Z., Lee, N., Frieske, R., Yu, T., Su, D., Xu, Y., Ishii, E., Bang, Y. J., Madotto, A., & Fung, P. (2023). Survey of hallucination in natural language generation. *ACM Computing Surveys*, 55(12), 1–38. <https://doi.org/10.1145/3571730>
- Jin, Y., Jang, E., Cui, J., Chung, J.-W., Lee, Y., & Shin, S. (2023). DarkBERT: A language model for the dark side of the Internet. In *61st Annual Meeting of the Association for Computational Linguistics (ACL'23)*, Toronto, Canada.
- Johnson, D. G., & Verdicchio, M. (2017). AI Anxiety. *Journal of the Association for Information Science and Technology*, 68(9), 2267–2270. <https://doi.org/10.1002/asi.23867>
- Jumper, J., Evans, R., Pritzel, A., Green, T., Figurnov, M., Ronneberger, O., Tunyasuvunakool, K., Bates, R., Židek, A., Potapenko, A., Bridgland, A., Meyer, C., Kohl, S. A. A., Ballard, A. J., Cowie, A., Romera-Paredes, B., Nikolov, S., Jain, R., Adler, J., & Hassabis, D. (2021). Highly accurate protein structure prediction with AlphaFold. *Nature*, 596(7873), 583–589. <https://doi.org/10.1038/s41586-021-03819-2>
- Kingma, D. P., & Welling, M. (2014). Auto-encoding variational Bayes. *International Conference on Learning Representations 2021 (ICLR)*, Banff, Canada.
- Kingma, D. P., Mohamed, S., Jimenez Rezende, D., & Welling, M. (2014). Semi-supervised learning with deep generative models. In Z. Ghahramani, M. Welling, C. Cortes, N. Lawrence, & K. Q. Weinberger (Eds.), *Advances in Neural Information Processing Systems 27* (Vol. 27). Curran Associates, Inc.
- Kodali, N., Abernethy, J., Hays, J., & Kira, Z. (2017). On convergence and stability of GANs. *arXiv*. <https://doi.org/10.48550/arXiv.1705.07215>
- Kowalczyk, P., Röder, M., & Thiesse, F. (2023). Nudging creativity in digital marketing with generative artificial intelligence: Opportunities and limitations. *ECIS 2023 Research-in-Progress Papers*, Article 22.
- Kreps, S., McCain, R. M., & Brundage, M. (2022). All the news that's fit to fabricate: AI-generated text as a tool of media misinformation. *Journal of Experimental Political Science*, 9(1), 104–117. <https://doi.org/10.1017/XPS.2020.37>
- Kühl, N., Schemmer, M., Goutier, M., & Satzger, G. (2022). Artificial intelligence and machine learning. *Electronic Markets*, 32(4), 2235–2244. <https://doi.org/10.1007/s12525-022-00598-0>
- LeCun, Y., Bengio, Y., & Hinton, G. (2015). Deep learning. *Nature*, 521(7553), 436–444. <https://doi.org/10.1038/nature14539>
- Lehmann, F., & Buschek, D. (2020). Examining autocompletion as a basic concept for interaction with generative AI. *I-Com*, 19(3), 251–264. <https://doi.org/10.1515/icom-2020-0025>
- Leiker, D., Gyllen, A. R., Eldesouky, I., & Cukurova, M. (2023). Generative AI for learning: Investigating the potential of synthetic learning videos. In *24th International Conference of Artificial Intelligence in Education (AIED 2023)*, Tokyo, Japan.
- Li, H. (2022). Language models. *Communications of the ACM*, 65(7), 56–63. <https://doi.org/10.1145/3490443>
- Li, J., Li, M., Wang, X., & Thatcher, J. B. (2021). Strategic directions for AI: The role of CIOs and boards of directors. *MIS Quarterly*, 45(3), 1603–1644. <https://doi.org/10.25300/MISQ/2021/16523>
- Li, M., Bao, X., Chang, L., & Gu, T. (2022). Modeling personalized representation for within-basket recommendation based on deep learning. *Expert Systems with Applications*, 192, 116383. <https://doi.org/10.1016/j.eswa.2021.116383>
- Lins, S., Pandl, K. D., Teigeler, H., Thiebes, S., Bayer, C., & Sunyayev, A. (2021). Artificial intelligence as a service. *Business & Information Systems Engineering*, 63(4), 441–456. <https://doi.org/10.1007/s12599-021-00708-w>
- Liu, V., & Chilton, L. B. (2022). Design guidelines for prompt engineering text-to-image generative models. In S. Barbosa, C. Lampe, C. Appert, D. A. Shamma, S. Drucker, J. Williamson, & K. Yatani (Eds.), *CHI Conference on Human Factors in Computing Systems* (pp. 1–23). ACM. <https://doi.org/10.1145/3491102.3501825>
- Longoni, C., Fradkin, A., Cian, L., & Pennycook, G. (2022). News from generative artificial intelligence is believed less. In *2022 ACM Conference on Fairness, Accountability, and Transparency* (pp. 97–106). ACM. <https://doi.org/10.1145/3531146.3533077>
- Lukyanenko, R., Maass, W., & Storey, V. C. (2022). Trust in artificial intelligence: From a Foundational Trust Framework to emerging research opportunities. *Electronic Markets*, 32(4), 1993–2020. <https://doi.org/10.1007/s12525-022-00605-4>
- Lund, B. D., Wang, T., Mannuru, N. R., Nie, B., Shimray, S., & Wang, Z. (2023). ChatGPT and a new academic reality: Artificial intelligence-written research papers and the ethics of the large language models in scholarly publishing. *Journal of the Association for Information Science and Technology*, 74(5), 570–581. <https://doi.org/10.1002/asi.24750>
- Lysyakov, M., & Viswanathan, S. (2022). Threatened by AI: Analyzing users' responses to the introduction of AI in a crowd-sourcing platform. *Information Systems Research*, 34(3). Advance online publication. <https://doi.org/10.1287/isre.2022.1184>
- Mayahi, S., & Vidrih, M. (2022). *The impact of generative AI on the future of visual content marketing*. <https://doi.org/10.48550/arXiv.2211.12660>

- Mehrabi, N., Morstatter, F., Saxena, N., Lerman, K., & Galstyan, A. (2022). A survey on bias and fairness in machine learning. *ACM Computing Surveys*, 54(6), 1–35. <https://doi.org/10.1145/3457607>
- Meske, C., Abedin, B., Klier, M., & Rabhi, F. (2022). Explainable and responsible artificial intelligence. *Electronic Markets*, 32(4), 2103–2106. <https://doi.org/10.1007/s12525-022-00607-2>
- Microsoft. (2023). *Microsoft and OpenAI extend partnership*. <https://blogs.microsoft.com/blog/2023/01/23/microsoftandopenaiextendpartnership/>
- Miller, T. (2019). Explanation in artificial intelligence: Insights from the social sciences. *Artificial Intelligence*, 267, 1–38. <https://doi.org/10.1016/j.artint.2018.07.007>
- Mirbabaie, M., Brünker, F., Möllmann Frick, N. R. J., & Stieglitz, S. (2022). The rise of artificial intelligence – Understanding the AI identity threat at the workplace. *Electronic Markets*, 32(1), 73–99. <https://doi.org/10.1007/s12525-021-00496-x>
- Mirsky, Y., & Lee, W. (2022). The creation and detection of deepfakes. *ACM Computing Surveys*, 54(1), 1–41. <https://doi.org/10.1145/3425780>
- Mondal, S., Das, S., & Vrana, V. G. (2023). How to bell the cat? A theoretical review of generative artificial intelligence towards digital disruption in all walks of life. *Technologies*, 11(2), 44. <https://doi.org/10.3390/technologies11020044>
- Moussawi, S., Koufaris, M., & Benbunan-Fich, R. (2021). How perceptions of intelligence and anthropomorphism affect adoption of personal intelligent agents. *Electronic Markets*, 31(2), 343–364. <https://doi.org/10.1007/s12525-020-00411-w>
- Murphy, C., & Thomas, F. P. (2023). Generative AI in spinal cord injury research and care: Opportunities and challenges ahead. *The Journal of Spinal Cord Medicine*, 46(3), 341–342. <https://doi.org/10.1080/10790268.2023.2198926>
- Nichol, A., Jun, H., Dhariwal, P., Mishkin, P., & Chen, M. (2022). *Point-E: A system for generating 3D point clouds from complex prompts*. *arXiv*. <https://doi.org/10.48550/arXiv.2212.08751>
- Ntoutsi, E., Fafalios, P., Gadiraju, U., Iosifidis, V., Nejdil, W., Vidal, M.-E., Ruggieri, S., Turini, F., Papadopoulos, S., Krasanakis, E., Kompatsiaris, I., Kinder-Kurlanda, K., Wagner, C., Karimi, F., Fernandez, M., Alani, H., Berendt, B., Kruegel, T., Heinze, C., & Staab, S. (2020). Bias in data-driven artificial intelligence systems—An introductory survey. *WIREs Data Mining and Knowledge Discovery*, 10(3), e1356. <https://doi.org/10.1002/widm.1356>
- OpenAI. (2023). GPT-4 technical report. *arXiv*. <https://doi.org/10.48550/arXiv.2303.08774>
- Openlaender, J. (2022). The creativity of text-to-image generation. *Proceedings of the 25th International Academic Mindtrek Conference* (pp. 192–202). ACM. <https://doi.org/10.1145/3569219.3569352>
- Ouyang, L., Wu, J., Jiang, X., Almeida, D., Wainwright, C. L., Mishkin, P., Zhang, C., Agarwal, S., Slama, K., Ray, A., Schulman, J., Hilton, J., Kelton, F., Miller, L., Simens, M., Askell, A., Welinder, P., Christiano, P., Leike, J., & Lowe, R. (2022). *Training language models to follow instructions with human feedback*. <https://doi.org/10.48550/arXiv.2203.02155>
- Pan, Z., Yu, W., Yi, X., Khan, A., Yuan, F., & Zheng, Y. (2019). Recent progress on generative adversarial networks (GANs): A survey. *IEEE Access*, 7, 36322–36333. <https://doi.org/10.1109/ACCESS.2019.2905015>
- Patterson, D. W. (1990). *Introduction to artificial intelligence and expert systems*. Prentice Hall.
- Pavlik, J. V. (2023). Collaborating with ChatGPT: Considering the implications of generative artificial intelligence for journalism and media education. *Journalism & Mass Communication Educator*, 78(1), 84–93. <https://doi.org/10.1177/10776958221149577>
- Pentina, I., Hancock, T., & Xie, T. (2023). Exploring relationship development with social chatbots: A mixed-method study of replika. *Computers in Human Behavior*, 140, 107600. <https://doi.org/10.1016/j.chb.2022.107600>
- Perez, F., & Ribeiro, I. (2022). Ignore previous prompt: Attack techniques for language models. In D. Hendrycks, V. Krakovna, D. Song, J. Steinhardt, & N. Carlini (Chairs), *Thirty-sixth Conference on Neural Information Processing Systems (NeurIPS)*, Virtual.
- Piccialli, F., Di Cola, V. S., Giampaolo, F., & Cuomo, S. (2021). The role of artificial intelligence in fighting the COVID-19 pandemic. *Information Systems Frontiers: A Journal of Research and Innovation*, 23(6), 1467–1497. <https://doi.org/10.1007/s10796-021-10131-x>
- Poole, B., Jain, A., Barron, J. T., & Mildenhall, B. (2023). DreamFusion: Text-to-3D using 2D diffusion. In *Eleventh International Conference on Learning Representations (ICLR 2023)*, Kigali, Rwanda.
- Qiu, L., & Benbasat, I. (2005). An investigation into the effects of text-to-speech voice and 3D avatars on the perception of presence and flow of live help in electronic commerce. *ACM Transactions on Computer-Human Interaction*, 12(4), 329–355. <https://doi.org/10.1145/1121112.1121113>
- Raj, M., Berg, J., & Seamans, R. (2023). Artificial intelligence: The effect of AI disclosure on evaluations of creative content. *arXiv*. <https://doi.org/10.48550/arXiv.2303.06217>
- Ray, S. (2019). A quick review of machine learning algorithms. In *2019 International Conference on Machine Learning, Big Data, Cloud and Parallel Computing (COMITCon)* (pp. 35–39). IEEE. <https://doi.org/10.1109/COMITCon.2019.8862451>
- Riedl, R. (2022). Is trust in artificial intelligence systems related to user personality? Review of empirical evidence and future research directions. *Electronic Markets*, 32(4), 2021–2051. <https://doi.org/10.1007/s12525-022-00594-4>
- Rix, J., & Hess, T. (2023). From “handmade” to “AI-made”: Mitigating consumers’ aversion towards AI-generated textual products. *AMCIS 2023 Proceedings*.
- Rombach, R., Blattmann, A., Lorenz, D., Esser, P., & Ommer, B. (2022). High-resolution image synthesis with latent diffusion models. In *2022 IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)* (pp. 10674–10685). IEEE. <https://doi.org/10.1109/CVPR52688.2022.01042>
- Ruthotto, L., & Haber, E. (2021). An introduction to deep generative modeling. *GAMM-Mitteilungen*, 44(2), e202100008. <https://doi.org/10.1002/gamm.202100008>
- Samtani, S., Zhu, H., Padmanabhan, B., Chai, Y., Chen, H., & Nunamaker, J. F. (2023). Deep learning for information systems research. *Journal of Management Information Systems*, 40(1), 271–301. <https://doi.org/10.1080/07421222.2023.2172772>
- Schneider, J., Seidel, S., Basalla, M., & vom Brocke, J. (2023). Reuse, reduce, support: Design Principles for green data mining. *Business & Information Systems Engineering*, 65(1), 65–83. <https://doi.org/10.1007/s12599-022-00780-w>
- Schoormann, T., Strobel, G., Möller, F., Petrik, D., & Zschech, P. (2023). Artificial intelligence for sustainability - A systematic review of information systems literature. *Communications of the Association for Information Systems*, 52(1), 199–237. <https://doi.org/10.17705/ICAIS.05209>
- Schramowski, P., Turan, C., Andersen, N., Rothkopf, C. A., & Kersting, K. (2022). Large pre-trained language models contain human-like biases of what is right and wrong to do. *Nature Machine Intelligence*, 4(3), 258–268. <https://doi.org/10.1038/s42256-022-00458-8>

- Schuhmann, C., Beaumont, R., Vencu, R., Gordon, C. W., Wightman, R., Cherti, M., Coombes, T., Katta, A., Mullis, C., Wortsman, M., Schramowski, P., Kundurthy, S. R., Crowson, K., Schmidt, L., Kaczmarczyk, R., & Jitsev, J. (2022). LAION-5B: An open large-scale dataset for training next generation image-text models. In S. Koyejo, S. Mohamed, A. Agarwal, D. Belgrave, K. Cho, & A. Oh (Eds.), *Advances in neural information processing systems 35*. Curran Associates, Inc.
- Selz, D. (2020). From electronic markets to data driven insights. *Electronic Markets*, 30(1), 57–59. <https://doi.org/10.1007/s12525-019-00393-4>
- Smits, J., & Borghuis, T. (2022). Generative AI and intellectual property rights. In B. Custers & E. Fosch-Villaronga (Eds.), *Information Technology and Law Series. Law and Artificial Intelligence* (Vol. 35, pp. 323–344). T.M.C. Asser Press. [https://doi.org/10.1007/978-94-6265-523-2\\_17](https://doi.org/10.1007/978-94-6265-523-2_17)
- Stability.ai. (2023). *Stability AI launches the first of its StableLM suite of language models*. <https://stability.ai/blog/stability-ai-launches-the-first-of-its-stablelm-suite-of-language-models>
- Strobel, G., Banh, L., Möller, F., & Schoormann, T. (2024). Exploring generative artificial intelligence: A taxonomy and types. In *Hawaii International Conference on System Sciences 2024 (HICSS 2024)*, Hawaii, USA.
- Strobel, G., Schoormann, T., Banh, L., & Möller, F. (2023). Artificial intelligence for sign language translation – A design science research study. *Communications of the Association for Information Systems*, 53(1), 42–64. <https://doi.org/10.17705/ICAIS.05303>
- Sun, J., Liao, Q. V., Muller, M., Agarwal, M., Houde, S., Talamadupula, K., & Weisz, J. D. (2022). Investigating explainability of generative AI for code through scenario-based design. In *27th International Conference on Intelligent User Interfaces* (pp. 212–228). ACM. <https://doi.org/10.1145/3490099.3511119>
- Susarla, A., Gopal, R., Thatcher, J. B., & Sarker, S. (2023). The Janus effect of generative AI: Charting the path for responsible conduct of scholarly activities in information systems. *Information Systems Research*, 34(2), 399–408. <https://doi.org/10.1287/isre.2023.ed.v34.n2>
- Synthesia. (2023). *Synthesia | #1 AI Video Generation Platform*. <https://www.synthesia.io/>
- Teubner, T., Flath, C. M., Weinhardt, C., van der Aalst, W., & Hinz, O. (2023). Welcome to the era of ChatGPT et al.: The prospects of large language models. *Business & Information Systems Engineering*, 65, 95–101. <https://doi.org/10.1007/s12599-023-00795-x>
- The Washington Post. (2022). *The Google engineer who thinks the company's AI has come to life*. <https://www.washingtonpost.com/technology/2022/06/11/google-ai-lamda-blake-lemoine/>
- Tomczak, J. M. (2022). *Deep generative modeling*. Springer International Publishing. <https://doi.org/10.1007/978-3-030-93158-2>
- Tomita, C., Schaschek, M., Straub, L., & Winkelmann, A. (2023). What is the minimum to trust AI?—A requirement analysis for (generative) AI-based texts. *Wirtschaftsinformatik 2023 Proceedings*.
- van den Broek, E., Sergeeva, A., & Huysman Vrije, M. (2021). When the machine meets the expert: An ethnography of developing AI for hiring. *MIS Quarterly*, 45(3), 1557–1580. <https://doi.org/10.25300/MISQ/2021/16559>
- van Dun, C., Moder, L., Kratsch, W., & Röglinger, M. (2023). ProcessGAN: Supporting the creation of business process improvement ideas through generative machine learning. *Decision Support Systems*, 165, 113880. <https://doi.org/10.1016/j.dss.2022.113880>
- van Slyke, C., Johnson, R., & Sarabadani, J. (2023). Generative artificial intelligence in information systems education: Challenges, consequences, and responses. *Communications of the Association for Information Systems*, 53(1), 1–21. <https://doi.org/10.17705/ICAIS.05301>
- Vasist, P. N., & Krishnan, S. (2022). Deepfakes An integrative review of the literature and an agenda for future research. *Communications of the Association for Information Systems*, 51, 590–636. <https://doi.org/10.17705/ICAIS.05126>
- Vaswani, A., Shazeer, N., Parmar, N., Uszkoreit, J., Jones, L., Gomez, A. N., Kaiser, U., & Polosukhin, I. (2017). Attention is all you need. In I. Guyon, U. von Luxburg, S. Bengio, H. Wallach, R. Fergus, S. Vishwanathan, & R. Garnett (Eds.), *Advances in neural information processing systems 30* (pp. 5999–6009). Curran Associates Inc.
- Walters, W. P., & Murcko, M. (2020). Assessing the impact of generative AI on medicinal chemistry. *Nature Biotechnology*, 38(2), 143–145. <https://doi.org/10.1038/s41587-020-0418-2>
- Wang, C., Chen, S., Wu, Y., Zhang, Z., Zhou, L., Liu, S., Chen, Z., Liu, Y., Wang, H., Li, J., He, L., Zhao, S., & Wei, F. (2023). Neural codec language models are zero-shot text to speech synthesizers. *arXiv*. <https://doi.org/10.48550/arXiv.2301.02111>
- Wanner, J., Herm, L.-V., Heinrich, K., & Janiesch, C. (2022). The effect of transparency and trust on intelligent system acceptance: Evidence from a user-based study. *Electronic Markets*, 32(4), 2079–2102. <https://doi.org/10.1007/s12525-022-00593-5>
- Wei, R., & Mahmood, A. (2021). Recent advances in variational autoencoders with representation learning for biomedical informatics: A survey. *IEEE Access*, 9, 4939–4956. <https://doi.org/10.1109/ACCESS.2020.3048309>
- Weidinger, L., Uesato, J., Rauh, M., Griffin, C., Huang, P.-S., Mellor, J., Glaese, A., Cheng, M., Balle, B., Kasirzadeh, A., Biles, C., Brown, S., Kenton, Z., Hawkins, W., Stepleton, T., Birhane, A., Hendricks, L. A., Rimell, L., Isaac, W., Gabriel, I. (2022). Taxonomy of risks posed by language models. In *2022 ACM Conference on Fairness, Accountability, and Transparency* (pp. 214–229). ACM. <https://doi.org/10.1145/3531146.3533088>
- Weisz, J., Muller, M., He, J., & Houde, S. (2023). Toward general design principles for generative AI applications. In *4th Workshop on Human-AI Co-Creation with Generative Models*, Sydney, Australia.
- Weng, S.-S., & Chen, H.-C. (2020). Exploring the role of deep learning technology in the sustainable development of the music production industry. *Sustainability*, 12(2), 625. <https://doi.org/10.3390/su12020625>
- Wessel, M., Adam, M., Benlian, A., Majchrzak, A., & Thies, F. (2023). Call for papers to the special issue: Generative AI and its transformative value for digital platforms. *Journal of Management Information Systems*. [https://www.jmis-web.org/cfps/JMIS\\_SI\\_CFP\\_Generative\\_AI.pdf](https://www.jmis-web.org/cfps/JMIS_SI_CFP_Generative_AI.pdf)
- Willcocks, L. (2020). Robo-Apocalypse cancelled? Reframing the automation and future of work debate. *Journal of Information Technology*, 35(4), 286–302. <https://doi.org/10.1177/0268396220925830>
- Winston, P. H. (1993). *Artificial intelligence* (3. ed., reprinted with corr). Addison-Wesley.
- Yang, R., & Wibowo, S. (2022). User trust in artificial intelligence: A comprehensive conceptual framework. *Electronic Markets*, 32(4), 2053–2077. <https://doi.org/10.1007/s12525-022-00592-6>
- Zhan, F., Yu, Y., Wu, R., Zhang, J., Lu, S., Liu, L., Kortylewski, A., Theobalt, C., & Xing, E. (2021). Multimodal Image Synthesis and Editing: A Survey. *arXiv*. <https://doi.org/10.48550/arXiv.2112.13592>

- Zhang, C., Zhang, C., Zhang, M., & Kweon, I. S. (2023a). Text-to-image diffusion models in generative AI: A survey. *arXiv*. <https://doi.org/10.48550/arXiv.2303.07909>
- Zhang, D., Li, W., Niu, B., & Wu, C. (2023b). A deep learning approach for detecting fake reviewers: Exploiting reviewing behavior and textual information. *Decision Support Systems*, 166, 113911. <https://doi.org/10.1016/j.dss.2022.113911>
- Zhou, J., Zhang, Y., Luo, Q., Parker, A. G., & Choudhury, M. de (2023). Synthetic lies: Understanding AI-generated misinformation and evaluating algorithmic and human solutions. In A. Schmidt, K. Väänänen, T. Goyal, P. O. Kristensson, A. Peters, S. Mueller, J. R. Williamson, & M. L. Wilson (Eds.), *Proceedings of the 2023 CHI Conference on Human Factors in Computing Systems* (pp. 1–20). ACM. <https://doi.org/10.1145/3544548.3581318>

**Publisher's Note** Springer Nature remains neutral with regard to jurisdictional claims in published maps and institutional affiliations.