

Generative Artificial Intelligence

Pallab Debnath, Supriya Shaw, Dip Sarkar, Sanjana Maity, Debashis Sarkar, Dr. Ranjan Kumar Mondal

BCA Students, School of Computer Science, Swami
Vivekananda University, Kolkata, WB
Assistant Professor, Department of CSE, Swami
Vivekananda University, Kolkata, WB
ranjankm@svu.ac.in

Abstract

Computational methods that can produce original, meaningful content—like writing, pictures, or audio—from training data are called generative AI. The broad use of this technology, shown by-products like GPT-4, Copilot, and Dall-E 2, is now transforming how we collaborate and do business. In addition to being utilized creatively to produce fresh text that mimics authors or new visuals that resemble artists, generative AI systems can and will help people by acting as intelligent question-answering systems. Applications in this context include IT help desks, where generative AI assists with routine chores like food preparation, medical advice, and transitional knowledge job duties. According to industry projections, generative AI can potentially replace 300 million knowledge worker positions and increase the global gross domestic product (GDP) by 7% (Goldman Sachs, 2023). This significantly impacts the Business & Information Systems Engineering (BISE) community. Not only will we face revolutionary opportunities, but we will also face risks and challenges that we must manage to steer the technology and its use sustainably and responsibly.

Introduction

2.1 Generative AI's Mathematical Foundations

The foundation of generative AI is generative modeling, which differs mathematically from discriminative modeling and is often used in data-driven decision support. Generally speaking, discriminative modeling (e.g., in classification problems using $Y \in \{0, 1\}$) aims to learn decision boundaries between data points X to partition them into distinct classes Y . Generative modeling, on the other hand, seeks to infer some accurate data distribution. Two examples are the joint probability distribution $P(X, Y)$ of the inputs and outputs or $P(Y)$. In most cases, Y is from a high-dimensional space. In doing so, a generative model can create new synthetic samples (i.e., new observations X given a target value Y or new observation-target pairs (X, Y)).

Extending what was previously said, generative AI models are generative modeling instances that are instantiated with machine learning architectures (such as deep neural networks) and may, therefore, generate new data samples based on learning

Additionally, a generative AI system includes the

the whole infrastructure includes the user interface, data processing, and model. The model is the fundamental element of the system that enables interaction and application in the broader environment. Finally, “generative AI applications” describe these systems’ real-world applications and use cases, including article creation for search engine optimization (SEO) or code generation that addresses pressing issues and fosters creativity in various fields. Figure 1 presents a systematization of generative AI from the model-, system-, and application-level viewpoints (explained in detail in the next section) and across a few chosen data modalities (text, picture, and audio).

Take note of Fig. 1’s modalities. One is incomplete and unique, and further information may be provided. Additionally, various exceptional use cases, such as modeling the functional features of proteins, may be represented using another modality, like text.

2.2 An Overview of the Model, System, and Application Levels of Artificial Intelligence

A generative AI model is a machine learning architecture that builds on the patterns and correlations in the training data to generate new data instances using AI algorithms. It is vitally important but not yet complete, as systems and applications need to be used to refine it further for specific purposes.

Because deep neural networks can be designed using various architectures to model different data types, such as sequential data like human language or spatial data like images, they are exceptionally well suited for data generation (Janiesch et al., 2021; Kraus et al., 2020). An overview of typical generative AI model designs and underlying principles is provided in Table 1. Examples include diffusion probabilistic models for text-to-image generation, transformer architectures, and (large) language models (LLMs) for text creation. For instance, the famous family of LLMs known as GPT (short for generative pre-trained transformer) is used, for example, in the conversational agent ChatGPT to generate text.

An overview of generative AI models and many chosen data modalities is shown in Figure 1. These models have been trained on enormous volumes of data. Please note how the models are organized in Fig. 1 by their output mode, such as X-to-image or X-to-text. For instance, OpenAI's well-known conversational agent ChatGPT (OpenAI, 2023a) uses the most current generative AI model, GPT-4, which can produce text outputs from picture and text inputs. Likewise, Midjourney allows for the generation of pictures using both modalities. In light of this, unimodal and multimodal generative AI models may also be distinguished from one another. The input type for unimodal models is the same as the output type (text, for example). On the other hand, Multimodal models can receive input from several sources and provide output in multiple formats. There are multimodal models for many data modalities, such as text, picture, and audio. Some notable examples are the text-to-image generation technique called Stable Diffusion, the text-to-music generation technique called MusicLM, the text-to-code generation technique called Codex and AlphaCode, and the image-to-text and text-to-text generation technique called GPT-4.

Different generative AI models have very different underlying training processes. In order to eventually get the distribution of synthetic samples close to the distribution of training samples, generative adversarial networks (GANs), for instance, are trained through two competing objectives. The first objective is to create new synthetic samples, while the second aims to detect synthetic samples from the actual training samples. On the other hand, systems that employ conversational models, like ChatGPT, incorporate reinforcement learning from human feedback (RLHF). The three steps of RLHF as used by ChatGPT are as follows: first, demonstration data for prompts are created; next, users are asked to rate the quality of various outputs for a prompt; and last, a policy that produces desirable output via reinforcement learning is learned, ensuring that the output would score highly during ranking.

2.2.2 View at the System Level

Every system is made up of many components that are linked to one another and exchange information. In the case of generative AI systems, this includes the previously described generative AI model, the supporting framework, user-facing elements and their mode, and the associated data processing (for example, for prompts). One example would be incorporating deep learning models, such as Codex, into a more feature-rich and interactive system, such as GitHub Copilot, to enable more productive coding by its users. Similarly, users may utilize Discord bots to engage with an unnamed X-to-image generation model to produce pictures using Midjourney's image-generating system. In order to provide an interface for user interaction, generative AI systems integrate the functionality of the underlying mathematical model. By improving the model-specific capabilities, this stage makes the model more valuable and practicable for a broader range of real-world use scenarios.

Scalability (e.g., distributed computing resources), deployment (e.g., in several contexts and for multiple devices), and usability (e.g., a user-friendly interface and intent detection) are often the primary issues when integrating deep learning models in generative AI systems. The availability of trained open-source alternatives to closed-source, proprietary models

to their users—whether businesses or individuals—becomes increasingly crucial as these models continue to be produced. Unexpected declines in model performance over time for both closed-source and open-source models emphasize the need for ongoing model monitoring. Even though robust text-generating models were available before the ChatGPT system's debut in November 2022, one of the main reasons for its rapid global popularity was its simplicity of use, even for non-expert users.

Furthermore, a generative AI system's many components may be coupled or integrated at the system level with other systems, external databases that provide domain-specific information, or platforms. For instance, typical restrictions

Many generative AI models suffer from two main flaws: either they were trained on historical data with a set cut-off date and do not retain information beyond that point, or they undergo information compression, which causes the generative AI models to lose some of the information they observed during training. The model's accuracy and utility may be significantly increased by adding real-time information retrieval capability, which can help offset both drawbacks. Similarly, online language modeling solves the issue of out-of-date models in text creation by continually training them on

current data.² As a result, these models may be informed of current happenings that their static counterparts would not be aware of because of their training deadlines.

2.2.3 View at the Application Level

Applications of generative AI are generative AI systems installed in businesses to provide value by resolving specific business issues and meeting stakeholder demands. They may be considered information or human-task-technology systems that boost human capabilities to execute specific tasks via generative AI technology. This level of generative AI covers numerous real-world applications: they include creating SEO content artificial intelligence (AI) music, movies, and software development based on natural language.

Applications of generative AI will lead to new forms of technology-enabled labor. Users' perceptions of these innovative apps will change with time, influencing their trust, distrust, and use. Applications will probably evolve from routine jobs like composing basic letters and making dinner reservations to more delicate duties like requesting legal or medical guidance. They will need increasingly significant choices, maybe even including moral assessment (Kru'gel et al., 2023). Because generative AI applications are becoming increasingly widespread, there is an urgent need to outline guidelines and best practices for dependable and trustworthy designs and carefully consider how these designs will affect the user to balance attributes like trust properly. In order to fully understand the innovation potential and affordances of generative AI artifacts, end users and organizations must consider several critical socio-technical concerns while using and adopting such applications.

2.3 Generative AI from a Socio-Technical Perspective

The definition and scope of artificial intelligence (AI) constantly evolve as technology progresses, while the benchmark of human intellect remains relatively stable.

With an emphasis on the relationship between humans and AI, the idea of delegation has been explored so far for analytic AI to create a decision-making hierarchy. In generative AI, a human user interacts with an AI system via prompts to produce content; the AI then deciphers the user's goals and feeds back to needing more prompts. This also seems to follow a delegation pattern at first look. However, the next step doesn't, as the AI's output might indirectly or directly influence the other and shape their further engagement. As a result, the creation process more closely resembles co-creation, which works together in various positions to align and provide a range of perspectives to direct a design process. Initiation is not exclusive to humans when seen through the prism of agentic AI creations.

The interactions above influence how we conceptualize hybrid intelligence, the fusion of AI and humans using their unique abilities. Hybrid intelligence aims to overcome

the shortcomings of each form of intelligence by fusing human intuition, creativity, and empathy with the precision, scalability, and processing capacity of AI systems to improve decision-making and problem-solving skills. With the advent of generative AI and its capacity for co-creation, our conception of this collective intelligence has started to change. Therefore, to guarantee that AI agency and reach are included and to allow effective and efficient usage in AI applications, unique models and patterns of human-AI interaction may become essential to explain and steer the behavior of people and AI systems.

Three Drawbacks of Existing Generative AI

The following highlights four critical generative AI limits that, in our opinion, represent significant restrictions for practical use. The following restrictions are technical as they speak to how existing generative AI models conclude; as such, they are model-specific. This means that long-term limits are likely to exist, with repercussions at the system and application levels.

Inaccurate results. Models using generative AI may provide incorrect results. This might be attributed to the fundamental characteristics of machine learning models, which depend on probabilistic methods to conclude. Generative AI models, for instance, provide the most likely answer to a query, but they are not always the right one. Because of this, there are difficulties since outputs may now be mistaken for natural material and provide false information or trick people.

The quality of training data and the corresponding learning process significantly impact how accurate generative AI models are. Correctness checks may be used by generative AI systems and applications to prevent specific outputs. However, since modern AI models are opaque (Rai, 2020), users' confidence in the systems' ability to provide accurate results is crucial to their use. This reality is exacerbated by the closed-source nature of commercial, off-the-shelf generative AI systems, which prevents further model adjustment and retraining. Using generative AI to provide justifications or references that users may verify is one way to alleviate the downstream effects of inaccurate outputs. Though these

justifications are also probabilistic and prone to inaccuracies, they might nonetheless assist users in determining whether and not to accept generative AI results.

The academic literature is paying more attention to bias and fairness in AI, but it is still an open subject that needs further investigation. For instance, “probing and understanding the limitations and biases of generative models” is a crucial topic of study highlighted by the creators of Stable Diffusion. According to some researchers, models can self-correct morally, which might reduce worries about ingrained biases and increase fairness. Additionally, mitigating methods may be used at the system and application levels to remove biases ingrained in the deep learning models and provide more varied outputs (e.g., Dall-E 2 changing the prompts “under the hood” to boost the outputs’ demographic diversity). However, additional investigation is required to approach the concept of equitable AI.

Infringement of copyright. Because generative AI models, systems, and apps may generate outputs that mimic or even replicate existing works without the consent of or payment to the original authors, they may violate copyright laws. There are two common possible infringement hazards here. Generative AI can potentially infringe upon authors’ rights to reproduction by producing illicit copies of works. This may occur, among other things, if a generative AI is educated on copyright-protected original material and then creates duplicates. Therefore, a joint inference is that copyrights must not be included in the training data used to create generative AI systems. Importantly, copyright infringement may still occur even in cases where the generative AI has never seen a copyrighted work before. An example of this would be if AI created a trademarked logo that resembles Adidas’ emblem without ever having seen it before. However, generative AI has the potential to produce derivative works, infringing upon artists’ rights to transformation. As a result, the balance between originality and creativity in generative AI systems gives rise to legal concerns. In keeping with this, legal issues also surface around who owns the intellectual property for works (including patents) generated by generative artificial intelligence.

Environmental issues. Last, the development and application of generative AI systems raise severe environmental issues since they are usually based on large-scale neural networks, which result in high electricity consumption and a significant carbon footprint during development and operation. For instance, it was calculated that the carbon emissions used to train generative AI models, such as GPT-3, generated 552 t CO₂, equal to the yearly carbon emissions of several dozen families. Because of this, attempts are being made in AI research to build and implement AI algorithms that are less harmful to the environment by using more effective training algorithms, smaller neural network topologies, and better hardware.

4 . Consequences and Prospects for the BISE

Collective

As an application-oriented, socio-technical research field, the BISE community may directly benefit from the implications and future research paths we outline in this section. On the other hand,

Nonetheless, their multidisciplinary nature provides a wealth of study possibilities, particularly for BISE researchers. We arrange our thoughts based on the many departments of the BISE journal (for a summary of sample research issues, refer to Table 2).

4.1 Management of Business Processes

Generative AI will significantly impact Business Process Management (BPM) since it can help with process automation, employee and customer satisfaction, and process innovation opportunities, particularly in creative processes. Different stages of the BPM lifecycle model may be linked to specific consequences and future areas for study. For instance, generative AI models might be used to create process descriptions in the context of process discovery, which can assist companies in recognizing and comprehending the many steps of a process. Generative process models have the potential to facilitate idea creation and creative process (re-)design activities from the standpoint of business process

improvement. Generative AI has much promise in supporting both exploratory and exploitative BPM design methodologies. Furthermore, generative AI may apply for BPM-related natural language processing tasks like process extraction from text, even without requiring further prompt fine-tuning. Similarly, other stages may benefit from generative AI's capacity to understand intricate and non-linear linkages in dynamic business processes, which can be used, among other things, for simulation and predictive process monitoring.

The short-term benefits of robotic process automation include the ability to allow new forms of automation by retrofitting and intelligentizing existing software, in addition to replacing previously handwritten processing rules. We also see great potential in the long term to help the typical BPM phase of business process execution. In particular, we expect to create a fresh wave of process-guiding systems. Large enterprise-wide trained language models provide the foundation for more dynamic and adaptable systems, while conventional system designs rely on static and manually crafted knowledge bases. These systems have the potential to enhance knowledge retrieval tasks from a multitude of diverse sources, such as job descriptions, wikis, emails, manuals, and handbooks. This creates new research opportunities for integrating dispersed and unstructured organizational information into intelligent process guidance systems.

4.2 Data science and decision analytics

Even with the tremendous advancements in recent years, several analytical and technological issues regarding the development of generative AI remain unresolved. One unanswered topic is how generative AI may be successfully tailored for domain-specific applications and raise contextualization levels to enhance performance. For instance, customizing generative AI-based conversational bots for finance or medicine requires innovative and scalable methods. This will be essential in addressing specific BISE-related tasks where customization might result in significant performance advantages. Innovative customization strategies need to be created with the security of confidential information in mind and keep it hidden from prying eyes.

Furthermore, new frameworks created with the user in mind and supporting interpretability and usability are required for quick engineering.

Enhancing generative AI systems' dependability is a significant area of study. For instance, algorithmic approaches are required to identify and reduce hallucinations using generative AI. Further work is required to provide user-centered solutions or ways that users may lessen the likelihood of falling for wrong results, in addition to algorithmic solutions, for example, by producing improved methods for verifying outcomes (e.g., providing more details or references).

Lastly, concerns are raised on how generative AI can bridge the divide between domain expertise and modeling specialists to natively assist data science and decision analytics initiatives. For example, it is well known that many AI models used in business analytics are complex for non-experts to grasp. To address this issue, generative AI might provide explanations that clarify the reasoning behind business analytics models, improving the decision logic's comprehensibility. For instance, generative AI could produce user-friendly descriptions of inherently interpretable models or translate post hoc explanations from techniques like SHAP or LIME into more comprehensible textual descriptions.

4.3 Digital Management of Businesses and Digital

Headship

Generative AI has significant promise for enhancing various value-generation processes, such as autonomous agency, task augmentation, and knowledge development. However, to successfully guide the technical potential of generative AI in the right direction, the required organizational capacities and circumstances are also needed; additional study is needed to investigate these components in more detail.

Conclusion

A kind of artificial intelligence known as "generative AI" can produce fresh writings, photos, or audio files that are becoming increasingly similar to works of human

craftsmanship. Because of this, fields and sectors that depend on creativity, innovation, and knowledge processing stand to gain significantly from generative AI. Specifically, it opens up new applications like individualized education and services, digital art, and realistic virtual assistants that were previously unfeasible or impracticable for automation. Therefore, as an interdisciplinary research community, BISE practitioners and researchers stand to benefit significantly from generative AI. In our Catchword essay, we conceptualized the fundamentals of generative AI, describing the limits of existing generative AI from the perspectives of the social-technical, system, and application levels. Finally, using the BISE discipline as a lens, we highlighted the many affordances that generative AI provides by offering an influential research agenda for the BISE community.

References

- Brown, T. B., Mann, B., Ryder, N., Subbiah, M., Kaplan, J., Dhariwal, P., ... Amodei, D. (2020). Language models are few-shot learners. *Advances in Neural Information Processing Systems (NeurIPS)*, 33, 1877–1901.
- Goodfellow, I., Pouget-Abadie, J., Mirza, M., Xu, B., Warde-Farley, D., Ozair, S., ... Bengio, Y. (2014). Generative adversarial nets. *Advances in Neural Information Processing Systems*, 27, 2672–2680.
- Vaswani, A., Shazeer, N., Parmar, N., Uszkoreit, J., Jones, L., Gomez, A. N., ... Polosukhin, I. (2017). Attention is all you need. *Advances in Neural Information Processing Systems*, 30, 5998–6008.
- Radford, A., Narasimhan, K., Salimans, T., & Sutskever, I. (2018). Improving language understanding by generative pre-training. *OpenAI Technical Report*.
- Radford, A., Wu, J., Child, R., Luan, D., Amodei, D., & Sutskever, I. (2019). Language models are unsupervised multitask learners. *OpenAI Technical Report*.
- OpenAI. (2023). GPT-4 technical report. *arXiv preprint arXiv:2303.08774*.

Kingma, D. P., & Welling, M. (2014). Auto-encoding variational Bayes. *International Conference on Learning Representations (ICLR)*.

Ho, J., Jain, A., & Abbeel, P. (2020). Denoising diffusion probabilistic models. *Advances in Neural Information Processing Systems*, 33, 6840–6851.

Ramesh, A., Pavlov, M., Goh, G., Gray, S., Voss, C., Radford, A., ... Sutskever, I. (2021). Zero-shot text-to-image generation. *International Conference on Machine Learning (ICML)*.

Ramesh, A., Dhariwal, P., Nichol, A., Chu, C., & Chen, M. (2022). Hierarchical text-conditional image generation with CLIP latents. *arXiv preprint arXiv:2204.06125*.

Bommasani, R., Hudson, D. A., Adeli, E., Altman, R., Arora, S., von Arx, S., ... Liang, P. (2021). On the opportunities and risks of foundation models. *arXiv preprint arXiv:2108.07258*.

Bender, E. M., Gebru, T., McMillan-Major, A., & Shmitchell, S. (2021). On the dangers of stochastic parrots: Can language models be too big? *Proceedings of the ACM Conference on Fairness, Accountability, and Transparency*, 610–623.

Floridi, L., & Chiriatti, M. (2020). GPT-3: Its nature, scope, limits, and consequences. *Minds and Machines*, 30(4), 681–694.

Russell, S., & Norvig, P. (2021). *Artificial Intelligence: A Modern Approach* (4th ed.). Pearson Education.

Goodfellow, I., Bengio, Y., & Courville, A. (2016). *Deep Learning*. MIT Press.

Kietzmann, J., Paschen, J., & Treen, E. (2018). Artificial intelligence in advertising: How marketers can leverage AI. *Journal of Advertising Research*, 58(3), 263–267.

McKinsey Global Institute. (2023). *The economic potential of generative AI: The next productivity frontier*. McKinsey & Company.

Goldman Sachs. (2023). *The potentially large effects of artificial intelligence on economic growth*. Goldman Sachs Global Investment Research.

UNESCO. (2023). *Guidance on generative AI in education and research*. UNESCO Publishing.

Rai, A., Constantinides, P., & Sarker, S. (2019). Next-generation digital platforms: Toward human-AI hybrids. MIS Quarterly, 43(1), iii-ix.