VIEW POINT

THE EVOLUTION OF GENERATIVE AI MODELS

Abstract

Generative AI is rapidly blurring the lines between human and machine competencies. With capabilities ranging from content generation to the creation of complete virtual environments, the evolution of artificial intelligence has reshaped our personal and professional spaces in unimaginable ways.

From rule-based systems to complex autoencoders, transformers, and diffusion models, AI has evolved at an unprecedented pace. New, progressive frameworks support futuristic trends such as multimodal content creation, style transfer, data augmentation, and much more. Interestingly, this is just the beginning.

What are the enablers of this fascinating journey, and where are we headed next? Is it possible to control the pitfalls of an open ecosystem and ensure that technology serves the larger good? Here, we trace the evolution of AI from the 1950s to the present and delve briefly into the technical and ethical concerns in widespread AI adoption.

Infosys®
Navigate your next

Before the turn of the millennium, entertainment gadgets that ran our favourite shows or playlists on command, or faceless virtual assistants that managed our schedules and personal finances belonged to science fiction. And yet, in less than three decades, what seemed like magic is an imminent reality.

We have since normalised bantering with a disembodied Siri, issuing commands to an obliging Alexa, and requesting assistance from ChatGPT for academic assignments. A few aspects of our personal, social, and professional lives are untouched by the evolution of artificial intelligence.

We have come a long way from the days of binary logic when gadgets executed mechanical tasks while attributes like higher-order reasoning, creativity, decision-making, and empathy belonged to humans. Generative AI is increasingly blurring the lines between human and machine capabilities, surpassing humans in several domains.

Today, AI transcends mere automation and has evolved to generate original content, write code, train for specific skill sets, creatively solve problems, and create virtual environments in astounding detail. Studies following the advent of ChatGPT in November 2022 show that generative AI usage can add up to $4.4 trillion to the global economy annually.

The global generative AI market size was estimated at $10.79 billion in 2022. It is expected to grow at a CAGR of 27.02% over the decade, reaching $118.06 billion by 2023. Studies reveal that one in four companies in the US use generative AI tools for business. Developing nations, too, are investing heavily in AI self-reliance.

The generative AI evolution underlines the transformative power of technology that silently rewrites the way we live, work, interact, and consume content. Let us have a look at the fascinating journey of AI from its nascent phases to the present.

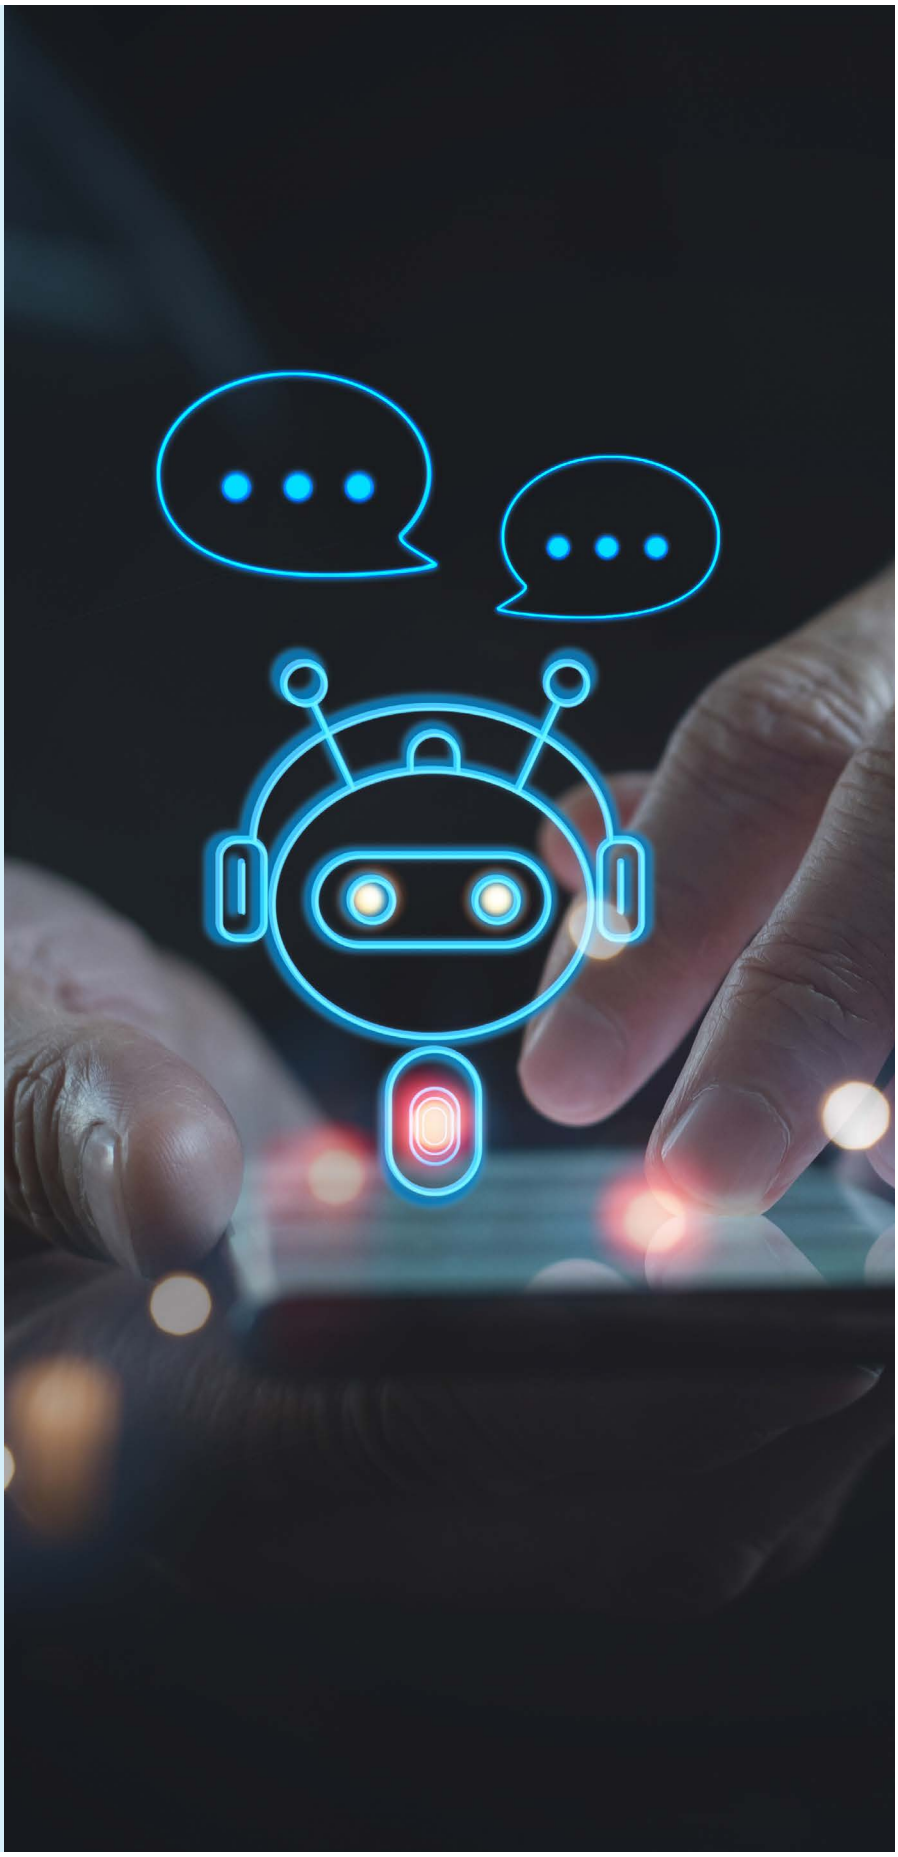

Evolution of generative artificial intelligence: A brief timeline

Generative artificial intelligence, as we know it today, can trace its roots to the mid-20th century. Here is a timeline of the key milestones in the evolution of generative AI:

• 1950s

Scientists first explored the possibilities of machine learning (ML) using algorithms in the 1950s. Early research revolved around rules-based systems that could replicate human reasoning and decision-making.

The Logic Theorist, an automated reasoning programme developed in 1956, was among the pioneering efforts that laid the foundation for the generative AI evolution.

John MacCarthy, an American computer scientist, coined the term 'Artificial Intelligence' in 1956 at the Dartmouth Conference.

• 1960s

The mid-60s ushered in the first chatbot, ELIZA, engineered at MIT. It was a Natural Language Processing (NLP)-powered bot that generated simple text responses to user inputs using pre-programmed phrases and questions.

Interestingly, Joseph Weizenbaum, the creator, named it after Eliza Dolittle, Shaw's Cockney character in The Pygmalion, who is taught to engage with high society with laconic responses that often reflect the speaker's mood and tone.

• 1980s – 90s

The Recurrent Neural Networks (RNNs) of the 80s and 90s laid the foundation for advanced generative models. Trained through backpropagation, RNNs could learn and remember patterns and make probability-based decisions.

1997 saw a significant breakthrough in AI-generated evolution when IBM's Deep Blue, a chess-playing computer, defeated world champion Gary Kasparov.

• 2000s

ML models came of age with support vector clustering, decision trees, random forest, and other complex algorithms that supported supervised learning.

In 2003, Yoshua Bengio engineered the first neural probabilistic language model that predicted the next word in a sentence. This was a groundbreaking achievement in NLP modelling.

• 2010 – 2020

This decade heralded the evolution of generative AI as we know it today. Some pivotal events that defined the gen AI landscape were as follows:

• Siri (2011)

Apple introduced Siri, the first AI/ NLP voice assistant. Siri truly brought generative AI to the masses.

• Release of GAN (2014)

Ian Goodfellow and his team introduced Generative Adversarial Networks (GANs) that could generate realistic images and videos, revolutionising the world of art, entertainment, and education.

• Diffusion model (2015)

Jascha Sohl-Dickstein developed the first diffusion model (DM) at Stanford. DMs have several use cases in applications requiring image search, text-to-video synthesis, and image-to-image translation.

• WaveNet (2016)

Google DeepMind developed WaveNet, a generative model that produces natural-sounding speech.

• Transformers (2017)

Google researchers proposed Transformer architecture, an important landmark in the AI evolution. Unlike RNNs and LSTMs that require sequential data for training, transformers allow for parallel computation, resulting in faster training times.

A versatile technology, transformer models form the base of many pathbreaking content generation and text-to-image applications, such as GPT 3 and DALLE-E, in subsequent years. Transformers garnered the highest AI market revenue share, exceeding 42%, in 2022.

- Progressive GANs (2017)

  Developed by NVIDIA, Progressive GANs enabled extremely high-resolution synthetic images with unprecedented clarity.

  ProGAN networks have found applications in multiple fields, including healthcare, architecture, entertainment, facial recognition technologies, and more, opening the doors to generative AI for business.

- OpenAI's GPT series (2018 – 2020)

  Open AI's GPT (Generative Pre-trained Transformer) series, based on Transformer architecture, is a game-changer in language generation tasks.

- 2021 to the present

The AI-generated evolution continues at a dizzying pace.

ChatGPT 3.5 followed the GPT series in 2022. It garnered one million users within five days of its launch. By making its offerings publicly available, ChatGPT is rapidly popularising open-source AI. The openai.com site attracts approximately 1.5 billion visitors every month. As of November 2023, ChatGPT boasts over 180 million users.

Other accomplishments in recent years include -

- Microsoft's DALLE – E and DALLE – E 2, which can generate original, high-quality images from language prompts
- BARD, Google's conversational chatbot
- Falcon LLM (Large Language Model) with advanced linguistic applications
- MusicGen, a Transformer model that can generate music samples from text or audio prompts
- Meta's Voicebox, a groundbreaking model for text-to-speech tasks that include audio editing, stylising, and sampling
- Stable Diffusion, Stability AI's image generation model
- Jasper AI, a versatile content-generation tool for blog posts, social media posts, marketing copy, and more
- Inworld, a multimodal Character Engine that generates lifelike characters for the metaverse

Drivers of the generative AI evolution

Generative AI owes its recent exponential growth to multiple technological advancements in the new millennium. Here are the key drivers of the millennial explosion of gen AI applications:

- Data availability

Generative AI models require enormous quantities of data from diverse sources for training. As the world becomes increasingly digital, access to high-quality training data gets easier, enabling realistic, diverse, and sustained outputs.

- Computational power

Among the recent advances in hardware, the evolution of powerful graphical processing units (GPUs) stands out the most. GPUs with specialised accelerators provide the computational power for matrix multiplication and other intensive operations required for training sophisticated AI models.

- Cloud capacities

The cloud offers multiple advantages, including open-source distributed storage systems, high-speed networking solutions, and affordable infrastructure, which have contributed significantly to the rapid evolution of AI-driven applications.

- Rise of the open-source model

The open-source framework facilitates knowledge-sharing and collaboration among research teams, IT teams, and other stakeholders. Open-source models have greatly contributed to the AI evolution and the widespread adoption of generative AI tools for business.

- Improved algorithms

Neural network architectures and deep learning frameworks have contributed to the development of new ML algorithms, which are significantly more effective than traditional algorithms in accelerating the gen AI evolution.

- Variational Autoencoders (VAEs), GANs, Diffusion models, LLMs

These networks have taken AI innovation to the next level by enabling multimodalities, cross-modalities, image synthesis, audio synthesis, text-to-speech/video technologies, and much more.

- Application diversity

Generative AI's ability to lend itself to domain-specific applications has motivated industries to facilitate research and develop innovative enterprise solutions. This led to a significant rise in generative AI for business in the past decade.

Looking beyond: The future of the AI evolution

While gen AI's current capabilities are remarkable, they are a tiny fraction of its future potential. In the next decade, the gen AI evolution will ascend to unprecedented heights, powered by its ability to address real-world problems. Here is how generative AI will transform our lives and workspaces in the next decade:

• Automation of creative and knowledge-based work

Gen AI has hitherto automated mundane physical tasks. Its capabilities will now spread to intellectual and creative activities. With technology supporting research, design, original content creation, and realistic art generation, educators, artists, and professionals can soon outsource some of their tasks to AI.

• Generative design

The convergence of data synthesis, data augmentation, NLP, cross-modalities, and fusion architectures will facilitate generative design and revolutionise design-oriented industries such as gaming, advertising, architecture, and robotics.

It will also open up newer possibilities in audio-visual content creation, such as dynamic storytelling, personalised virtual environments, multidimensional films, and interactive art installations.

• Physical applications

The current gen AI applications revolve around digital services and software development. By merging AI with 3D printing technology, gen AI will soon be able to design and manufacture physical products from scratch.

The automobile and aerospace industries have successfully leveraged this approach to produce lightweight and sustainable vehicle components. The technology will soon support a range of customised clothing, footwear, machine parts, and even prosthetic limbs.

This approach can cause a paradigm shift in the manufacturing process by enabling cost-effective, zero-waste, sustainable, and on-demand production.

• Improved healthcare

Healthcare is among the top industries impacted by the generative AI evolution. Gen AI will revolutionise the medical industry further with three pathbreaking tools:

- Healthcare-specific LLMs like MedPaLM, ClinicalBERT, and BioGPT
- CRISPR technology for genome editing
- 3D printing technologies for manufacturing medical equipment and prosthetic limbs

AI will facilitate faster drug discovery, genetic modifications, medical imaging, robotic surgery, and more. Remote diagnosis and remote patient care will soon be available to all.

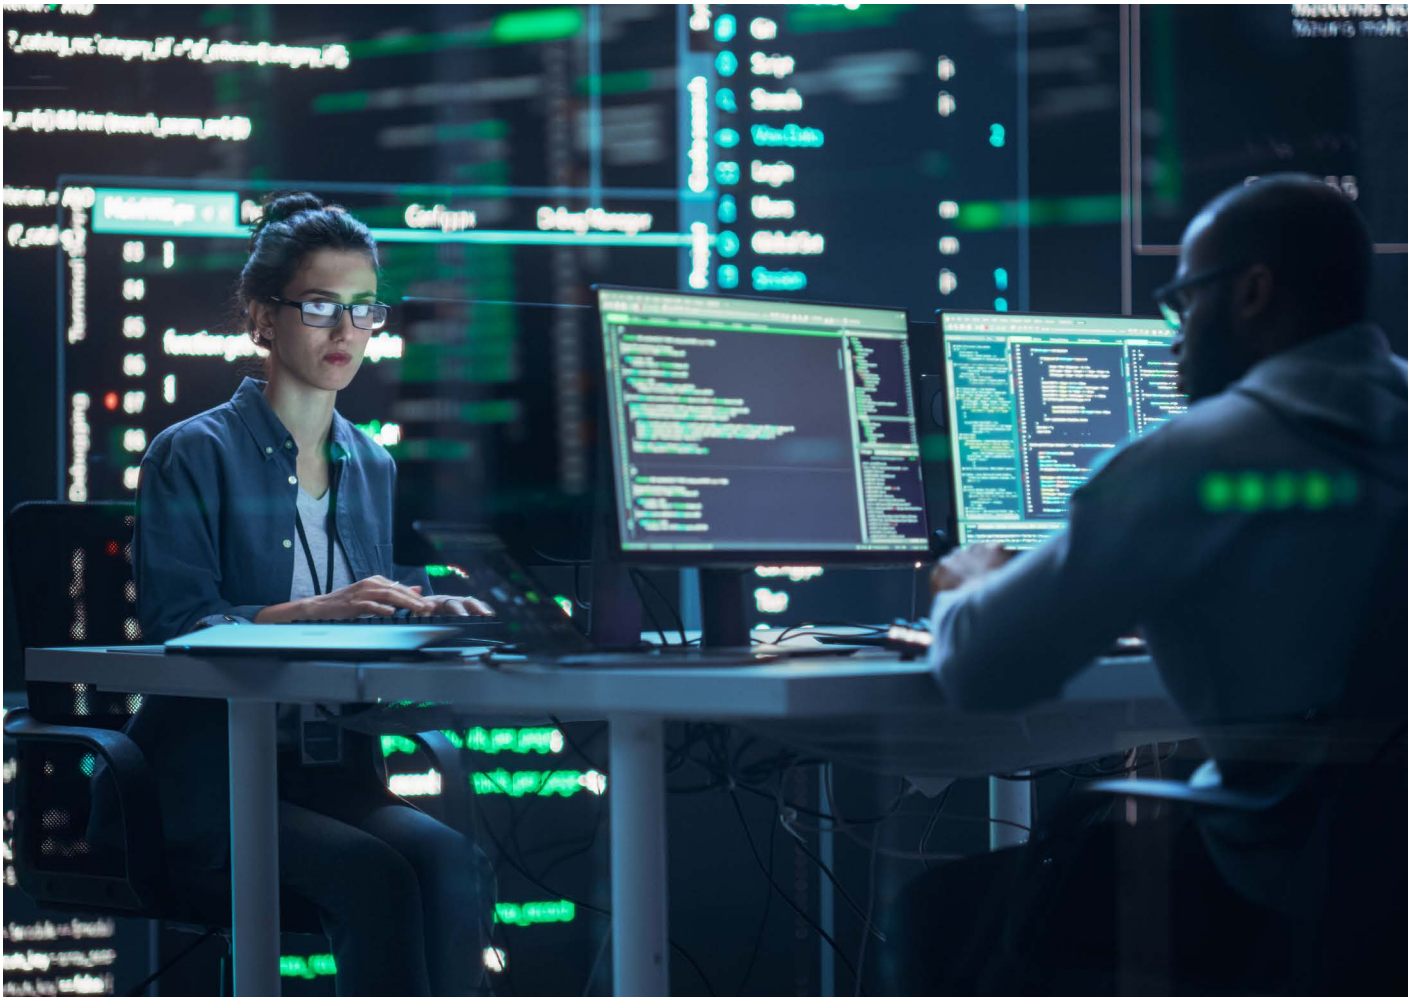

Navigating challenges in generative AI adoption

While generative AI has contributed significantly to the larger good, it is crucial to address concerns around safety, integrity, and ethics before endorsing its large-scale adoption. Below are the potential challenges in the widespread use of gen AI for business and research:

• Data privacy and security concerns

Maintaining data privacy is a key concern while handling enormous amounts of proprietary and open data for model training. Businesses and legal teams must ensure strict compliance with regulatory laws, ethical guidelines, and governance policies to prevent potential data breaches.

• Ethical considerations

Ethical concerns with generative AI usage stem mainly from two sources:

- Misuse of AI-generated content such as malicious deepfakes, doctored videos, or identity theft
- Experimentation with living organisms in ways that may interfere with their autonomy and integrity or beget unintended consequences such as accidental genetic mutations

These are critical concerns that may have a long-term impact on social and environmental harmony. Industries and governments must work together to draw strict guidelines and policies for the ethical use of generative AI.

• Data bias

AI models trained on biased datasets will generate biased outcomes, impacting the business's strategies, decision-making, and trustworthiness. IT and research teams must mitigate bias by standardisation and other techniques to ensure objectivity and inclusivity in outcomes.

• Ownership of content

With the proliferation of open-source systems, there are valid concerns about the ownership of AI-generated content, source attribution, and accountability. We require strong regulatory measures and governance practices to minimise instances of plagiarism and breach of intellectual property rights.

Conclusion

From the traditional rules-based systems of the 50s to complex, self-learning models with near-human attributes, the evolution of AI to the present has been eventful and transformative.

With capabilities that include generating content, conducting in-depth research, creating virtual environments, composing music, and writing code, generative AI models have pushed the boundaries of what seemed possible a few years ago.

The business use of generative AI is already mainstream. In the future, it will have a more profound impact on creative and knowledge-based fields. However, we must acknowledge that generative AI is not a replacement for human abilities but a catalyst that supplements and augments human capabilities in multiple ways.

It is crucial to address the potential pitfalls in this journey and leverage generative AI solutions ethically and responsibly to ensure that they change the world for the better.

For more information, contact infosysbpm@infosys.com

Infosys®
Navigate your next