

# TRANSFORMATIVE JOURNEY OF ARTIFICIAL INTELLIGENCE

Garima Joshi, Renu Vig

University Institute of Engineering & Technology, Panjab University, Chandigarh

## Abstract

Artificial intelligence has come a long way in last seventy five years. This article looks at how AI has changed over time, explaining advances in logic, probability, systems theory and neural computation that have helped build intelligent devices. It highlights major milestones and the contributions of key pioneers who connected theoretical principles with practical engineering. It also highlights landmark achievements and innovations that have impacted society and industry. The article illustrates how insights from physics and neuroscience have shaped the learning and training processes of computational models. In addition to tracing the technical evolution, it examines AI's broader societal influence across healthcare, education, communication, industry, and arts. While the technology is already transforming these domains, its long-term consequences will become evident only over time. We need to ensure appropriate policies exist so that society can benefit from these advancements.

**Keywords:** Artificial Intelligence (AI), Machine Learning (ML), Neural Networks, Deep Learning

## 1. Introduction

Artificial intelligence (AI) and machine learning (ML) came into existence through a combination of several disciplines like computer science, mathematics, physics, and neuroscience. Initially, AI was all about following rules and logic, but today algorithms can retrain and generalize, using structures like neural networks and deep learning (Ball, 2019). These new approaches are capable of handling complicated jobs involving vision, understanding language, and making decisions.

Ideas from physics and neuroscience have been helpful in creating models for learning and cognition. Neurobiological discoveries have driven AI research beyond rule-based approaches toward the development of systems inspired by the structure and functioning of the human brain (Barbierato, 2025). For instance, artificial neural networks (ANNs) mimic interconnected "neurons" and synapses. They can store and recall

information much like our own memories, using principles from physics and neurobiology that describe how systems settle into stable states (Boden, 2018).

Neuroscience has given us a better understanding of how our brains process information efficiently. Findings from studying individual neurons and neuroimaging have supported hypotheses like efficient coding, which states that early sensory systems minimize redundant neural firing while maximizing information transmission. Learning methods inspired by biology, such as Hebbian learning (which suggests that when neurons fire together, their connection gets stronger), are key to many models for memory, recognizing patterns, and making associations (Cath, 2017). The emphasis on neurobiological plausibility has encouraged researchers to develop bottom-up models that aim to capture not only cognitive

functions themselves, but also the ways in which these functions could arise from interacting networks of neurons (Domingos, 2015).

Statistical physics and theory of dynamical systems have contributed foundational principles to understanding learning and cognition. The concept of spin glass models, long studied in physics, helped shape early theories about how neural networks store many patterns and deal with noisy or partial inputs (Durrani, 2024). Hopfield’s models and later Boltzmann machines took inspiration from physics using energy functions and concept of free energy to describe network behavior (Dwivedi, 2023). Another significant direction is the development of physics-informed neural networks (PINNs), which incorporate physical laws, often written as differential equations, directly into the training process to ensure that the resulting models satisfy known constraints (Taddeo and Floridi, 2018).

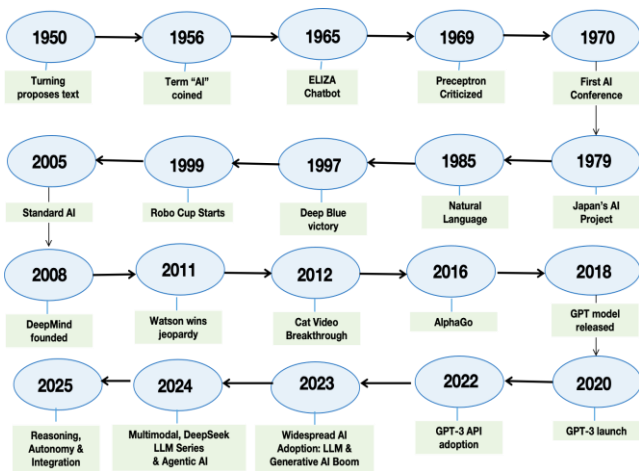
(Fu, 2024).

Figure 1 shows development of AI through distinct stages, with each phase bringing together theoretical insights and practical engineering (Gerdon, 2022). The evolution of AI is therefore more than a record of algorithms but a story of scientific curiosity, interdisciplinary collaboration and transformative potential that continues to redefine the limits of machine capability (Grzybowski, 2024). The 2024 Nobel Prize in Physics honored neural network research, underscoring deep connection between physical theory and computational intelligence (Grace, 2018). The article also addresses current challenges and future directions, emphasizing the need to balance innovation with ethics, transparency and human-centered design.

## 2. Tracing the Arc of AI

### 2.1 Laying the Foundations and Symbolic AI (1950s-1970s)

Back in 1950, Alan Turing surprised everyone through a paper called "Computing Machinery and Intelligence," asking that big, enduring question: "Can machines think?" He came up with this idea for an "imitation game," which is popularly known as the Turing Test, provided a criterion for assessing machine intelligence. Even more remarkably, Turing envisioned machines that could learn as they went along, basically setting up the whole philosophical and computing basis for AI research (Alufaisan, 2021). Then, in 1956, the Dartmouth Conference, spearheaded by John McCarthy, Marvin Minsky *et al.*, officially made AI as a field of study. This era was all about symbolic AI, expert systems, and creating programming languages like LISP (List Processing). Arthur Samuel's checkers program also showed off the idea of algorithms that could

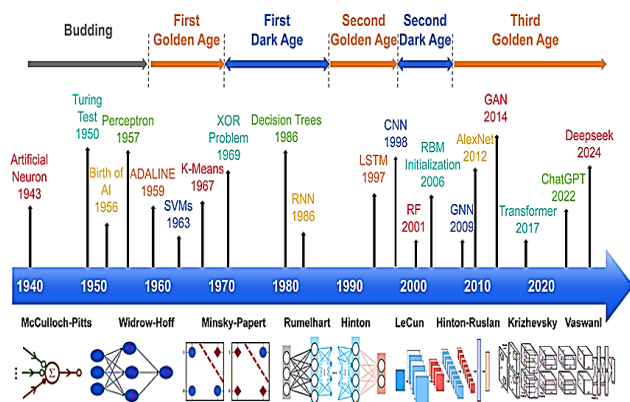


**Fig 1: Evolution of AI**

Recent research has also drawn on concepts from statistical physics and phase transitions to better understand what makes neural networks generalize effectively, how their internal representations evolve, and how learning systems balance trade-off between stability and flexibility

learn themselves, even coming up with the term "machine learning" and primarily laying early foundations for data-driven approaches (Jiang and Wang, 2018).

Then, in 1958, Frank Rosenblatt created the Perceptron, which was the very first ANN that could do simple pattern recognition. Although limited in scope, it demonstrated potential of biologically inspired computation (Jordan and Mitchell, 2015). Concurrent developments in heuristic programming and expert systems started pushing AI into real-world uses like medical diagnoses and engineering, giving us the first real glimpses of machine intelligence in action (Kaplan and Haenlein, 2019).



**Fig 2. Key Algorithms and Techniques in AI and ML (Yu, 2025)**

## 2.2 The AI Winter and Revival (1970s-1990s)

Things got off to a hopeful start with AI in 1950s and 60s, but by the mid-70s, it hit its first "winter." Basically, researchers were very excited about what the technology could do, but enthusiasm outpaced what the technology could deliver. Problems with symbolic AI systems, especially neural networks like perceptrons, became clear. For instance, what single-layer networks could actually compute was pretty limited. Hence, both funding and interest took a nosedive (Lavazza, 2025). Then, in early 80s,

there was a revival led by expert systems, which used rule-based logic to encode expert knowledge for specific domains such as medicine and configuration system. These used logic rules to capture what experts knew about specific areas like medicine or setting up systems. They proved that AI could actually be useful, even if only for very specific tasks.

However, by late 80s and into 90s, expert systems started showing their weaknesses. They were often fragile, costly to keep running, and difficult to expand. Simultaneously, advances in algorithms and increased computing power revived neural network research. The development of the backpropagation algorithm by Hinton, Rumelhart, Williams *et al.* made it much easier to train neural networks with multiple layers. On top of that, machine learning techniques such as decision trees and support vector machines became more feasible and widely used. This combination of better ideas, more data, and robust hardware paved way for major advancements in the deep learning era that came next (Grzybowski, 2024).

## 2.3 Rise of Machine Learning and Neural Networks (1990s-2010s)

During 1990s into 2000s, machine learning and neural network research resurged and began transforming many fields. Systems could now automatically learn useful features from data, rather than relying entirely on handcrafted rules. A landmark event illustrating growing capabilities was IBM's Deep Blue defeating Garry Kasparov in chess in 1997, becoming the first computer to win a full match against a reigning world champion. Another milestone came in 2016, when Google DeepMind's AlphaGo beat Lee Sedol in ancient game of Go, a game considered far more complex for computers

due to its large search space and intuitive reasoning requirements. At the same time, increases in computational power, larger datasets (“big data”), and advances in cloud computing infrastructures provided the tools needed to train deeper models and scale them practically. These combined developments laid foundation for deep learning revolution that dominated AI research in 2010s (Jiang and Wang, 2018).

Deep learning reshaped AI research and applications. Geoffrey Hinton, Yann LeCun, and Yoshua Bengio are designated as the “Godfathers of Deep Learning” for developing methods for training large-scale neural networks. LeCun’s convolutional neural networks (CNNs) transformed image recognition, Bengio advanced deep generative models, and Hinton contributed innovations in representation learning. Parallel efforts by Andrew Ng and others popularized ML education and established large-scale research initiatives such as Google Brain, accelerating global adoption of AI techniques (Lavazza, 2025).

#### **2.4 Modern AI and Deep Learning (2010-2024)**

Since the 2010s, AI has taken huge steps forward due to deep learning and massive neural networks, which have changed how machines can understand and perform tasks. Large Language Models (LLMs) like BERT and GPT-3, which are built on transformers, have dramatically improved how well computers can grasp human language and create text, making conversational AI much smoother, more aware of context, and genuinely helpful. Examples of other LLMs include T5, RoBERTa, and more recent models like Claude, Gemini, and those from the Llama family.

Reinforcement learning has also seen significant progress, enabling AI to get really good at tricky

jobs like playing games, controlling robots, and making choices in changing situations. At the same time, fields like healthcare, finance, and transportation have been using newer AI tools widely. For instance, AI is now being used for treatments tailored to individuals, spotting fraud, self-driving cars, and making supply chains run better. Key breakthroughs include AI-driven scientific research like protein structure prediction and climate modeling, improved medical diagnoses and treatments, more sophisticated image and video generation, and highly capable robots. All these advancements together have led to AI era, bringing both exciting possibilities and fresh challenges. The way AI brings together different fields really shines through in the work of John J. Hopfield and Geoffrey E. Hinton. Hopfield’s associative memory networks showed how ideas from statistical physics could be used to model human brain. Hinton deployed concepts from statistical mechanics in machine learning to make these algorithms more efficient. Their groundbreaking research on neural networks, which allows machines to recognize patterns and learn adaptively, earned them the 2024 Nobel Prize in Physics (The Nobel Prize, 2024).

#### **2.5 AI Advancement**

AI systems are designed for specific tasks like playing chess or recognizing faces. However, Artificial General Intelligence (AGI) would be a versatile “generalist” that can transfer learning across different domains and adapt to new situations. There is a strong drive in AI to create AGI systems that can take smart decisions. These systems often borrow ideas from how working of human brain and include safety features that are tested appropriately. Scientists are putting together huge systems with designs that mimic

the brain or are a mix of different approaches, adding in basic safety measures that are tested in real situations (Le Cun et al., 2015). A cool idea in this field is "sentinel agents." These are special AI agents that keep an eye on how other agents in a system are acting, looking for any bad behavior and making sure things are reliable and trustworthy. For example, a system could use sentinel agents to watch over how agents talk to each other, spot possible dangers, enforce privacy and access rules, and keep detailed logs of everything. Working alongside these sentinel agents are coordinator agents. They make sure policies are followed and manage agents, adjusting those policies based on warnings from the sentinel agents to keep the whole group of agents working properly (Ooi, 2025).

Evaluating the performance of autonomous AI systems requires assessment criteria that extend beyond conventional metrics such as task accuracy and execution speed (Schmidhuber, 2015). A comprehensive evaluation framework should encompass five key dimensions: functional capability and performance, robustness and adaptability, safety and ethical compliance, human-AI interaction, and economic viability. To support assessment across these dimensions, novel evaluation metrics proposed (Siebert, 2021). Evidence from real-world industrial deployments indicates that, although AI systems can substantially increase productivity, critical factors such as fairness, trustworthiness, and long-term sustainability are often overlooked. This underscores the need for multidimensional evaluation approaches that integrate automated performance metrics with human-centered assessments and financial reviews to ensure the responsible and sustainable use of AI systems (Brewka, 1996).

### 3. Current Challenges and Debates

As AI systems become advanced, it is imperative to address issues such as algorithmic bias, transparency in decision-making processes, and societal impacts of AI deployment to ensure that these technologies align with human values (Taddeo and Floridi, 2018). Lately, conversation around AI has shifted to include more technical, ethical, and political concerns (Tai, 2020).

With AI getting stronger, there has been a tug-of-war between better power like faster models, enhanced automation, better performance and accurate AI behaviour and safety (Taddeo and Floridi, 2021). Many experts believe alignment and safety are not in pace with speed of capability development. For example, there is increasing attention on risks that appear after AI systems are deployed, factors like self-improvement, multiple agents working together, or unexpected behaviors that emerge over time (Kaur and Sharma, 2024). Economically, forecasts suggest that while AI may bring big productivity gains, it could also displace many tasks and jobs, so it is crucial to consider how to retrain people and share the gains equitably (Voigts *et al.*, 2024). Generative AI, specifically, has sparked debates about fairness, copyright, data privacy, model transparency and accuracy in systems behaviour which do not always match statistical fairness metrics (Xie *et al.*, 2025).

As a response, policy reports and scholars are advocating for better structural and governance tools. For instance, India's Reserve Bank has proposed a framework for responsible AI in its financial sector, called FREE-AI, which emphasizes risk management, indigenous model development, audit frameworks, and oversight (Mathew *et al.*, 2020). Worldwide, reports like the AI Governance International Evaluation Index

(AGILE Index, 2024-2025) are measuring how countries are doing in AI governance highlighting gaps in regulation, standards, certification and comparing how prepared different countries are (Zuboff, 2015). Governments and organizations are increasingly calling for “governance by design”: building in ethics, human oversight, transparency, and safety constraints into AI systems from the start rather than as afterthoughts. Still progress remains uneven: many AI projects lack strong oversight, accountability mechanisms, or real capacity for enforcing ethical or safety measures.

#### 4. Impact of AI in Various Sectors

As we know, AI has impacted every field, making systems efficient, easier, and tailored. AI is changing every aspect of our lives, bringing both exciting possibilities and tricky challenges.

- In medicine, AI helps doctors spot diseases early, come up with patient specific personalised treatments, and handle patient records smoothly. AI is facilitating better treatment plans. AI programs help with learning and helping medical students practice with virtual scenarios and learn at their own pace.
- For schools, AI tools can create learning experiences that fit each student, give tests that adjust to their level, and automate a lot of the administrative work, making education more welcoming and effective for everyone. AI is helping to tailor lessons to each student, give instant feedback, and teach kids about AI itself.
- The way we communicate has also improved because of AI through enhanced language translation, sentiment analysis, and automated content creation, which helps

break down language and culture barriers.

- Industries leverage AI for predictive maintenance, supply chain optimization, and data-driven decision-making, pushing them to be more innovative and competitive. It is also transforming human-technology communication methods, notably through virtual helpers and chatbots. Industry 5.0 shifts the focus from Industry 4.0’s efficiency-driven automation to a value-driven approach centered on human-machine collaboration, prioritizing the well-being of the worker and leveraging human creativity alongside technology
- Even in arts, AI is helping people create new content, music composition, and visual art generation, opening up new creative avenues and challenging traditional artistic boundaries. Research shows AI is shaking things up in how art is made, how we interact with tools, and how different fields come together. It’s opening up new ways to do visual art and using data to come up with fresh ideas in design and art studies.

All these developments illustrate how important AI is in creating a world which is better connected, runs smoothly, and is full of new ideas. Though benefits are substantial, there are recurring concerns: possible displacement of jobs, bias in algorithms, ethical and privacy risks, and need to ensure fairness and human well-being are part of deployment decisions.

#### 5. AI Transformation of Panjab University

Panjab University (PU) has embarked on a profound strategic initiative, to reposition itself for research and development in Artificial Intelligence (AI). This concentrated effort is

driven by a commitment to establishing robust high-performance computing (HPC) infrastructure, for advanced skill development, and fostering translational DeepTech innovation. The goal is to provide students and researchers with the necessary resources for groundbreaking work in these rapidly evolving fields.

The cornerstone of this initiative is the recent inauguration of a cutting-edge AI Data Centre, established through a substantial Corporate Social Responsibility (CSR) contribution, alongside activation of the Intel Unnati Lab. This multidimensional infrastructure approach immediately positions Panjab University at the forefront of AI based education in this region. The translation of academic research into

commercially viable products is managed through established institutional frameworks as shown in Table 1.

Panjab University's approach integrates infrastructure with human capital development, particularly through faculty enablement programs and AI based curriculum delivery. Institutional frameworks such as the Technology Enabling Centre (TEC) and CIIPP facilitate technology transfer and commercialization, supporting over 200 industries and launching DeepTech startups. Strategic CSR funding further enhances AI-driven materials analysis and prototyping, aligning Panjab Universities's capabilities with national industrial goals.

**Table 1 : Strategic Alignment: PU Initiatives and Industry/Research Outcomes**

<b>Initiative/Facility</b>	<b>Primary Purpose</b>	<b>Measurable Outcome/Impact</b>
Design Innovation Centre	Design-centered synergy catalyzed for a variety of technology-driven projects around a plethora of ideas	Established 6 different Labs and Developed more than 110 prototypes, 16 patents and 10 startups
AI Data Centre	Large Model Training, Deep Learning	Provides critical high-performance computing capabilities, enabling industry-relevant research and custom model training with approximately 64.8 TFLOPS of compute performance.
Intel Unnati Lab	Faculty/Student Skill Enhancement	Standardized curriculum delivery; focus on edge AI optimization and industry skill gap reduction.
SAIF/CIL Integration	Materials Analysis & Prototyping	AI-driven quality control; accelerates new material development and industrial applications.
PU Incubation Centre (CIIPP/TEC)	Commercialization of DeepTech	Successful launch of high-impact, industry-specific startups (e.g., Dach Biotec, Terafac)

## 6. Conclusion

At the pinnacle of this technological revolution, a profound responsibility emerges to navigate the persistent, complex challenges of interpretability, bias, and sustainability while upholding core ethical principles. The "black box" nature of many computational models raises significant concerns about trust and accountability, particularly when AI-driven decisions impact human lives in high-stake domains like healthcare, finance, and criminal justice. Furthermore, without deliberate intervention, existing societal biases can become ingrained in training data, perpetuating and amplifying unfair outcomes that exacerbate inequalities. This must be coupled with addressing the substantial energy consumption and resource demands of large-scale AI, a critical sustainability challenge. The path forward demands an unwavering commitment to transparency, ensuring stakeholders can understand and audit how and why AI systems reach their conclusions. It necessitates robust accountability frameworks that define responsibility for the consequences of AI's actions. Ultimately, the goal is to align future AI advancements with human-centered values, guaranteeing that technology augments rather than undermines human dignity, fairness, and overall well-being. This is not merely a technical task, but a societal imperative to build a future where AI serves as a responsible and trustworthy partner to humanity.

Panjab University's strategic investment in AI infrastructure, skill development, and translational research has yielded a robust ecosystem for DeepTech innovation. The success of startups underscores the university's capacity to convert scientific research into impactful industrial applications. Looking ahead, PU aims

\*Corresponding Author: Garima Joshi: gjoshi5@pu.ac.in

Received: 2<sup>nd</sup> November, 2025 Accepted 31<sup>st</sup> December 2025

to scale its incubation model by expanding shared compute access and prioritizing support for ventures developing complex AI systems, including domain-specific LLMs and robotics platforms. This forward-looking plan positions PU not only as a leader in academic AI research but also as a catalyst for industrial transformation, contributing meaningfully to India's vision of Viksit Bharat@2047.

## References

- Alufaisan, Y., L. R. Marusich, J.Z. Bakdash, Y. Zhou and M. Kantarcioglu. 2021. Does explainable artificial intelligence improve human decision-making? Proceedings of the AAAI Conference on Artificial Intelligence. 35 (8) : 6618–6626.
- Ball, P. 2019. Using artificial intelligence to accelerate materials development. MRS Bulletin. 44 (5) : 335–344.
- Barbierato, E., A. Gatti, A. Incremona, A. Pozzi and D. Toti. 2025. Breaking away from AI: The ontological and ethical evolution of machine learning. IEEE Access. 13 : 55627–55647.
- Boden, M. A. *et al.* 2018. The growing ubiquity of algorithms in society: Implications, impacts and innovations. Philosophical Transactions of the Royal Society A: Mathematical, Physical and Engineering Sciences. 376 (2133) : 1–16.
- Brewka, G. 1996. Artificial intelligence—a modern approach by Stuart Russell and Peter Norvig. Prentice Hall, Englewood Cliffs, NJ. The Knowledge Engineering Review. 11 (1) : 78–79.
- Cath, C. *et al.* 2017. Artificial intelligence and the 'good society': The US, EU, and UK

- approach. *Contemporary Security Policy*. 38 (3) : 305–322.
- Domingos, P. 2015. *The master algorithm: How the quest for the ultimate learning machine will remake our world*. Basic Books, New York, USA.
- Durrani, U .K., A. H. Khan and M. A. Khan. 2024. A decade of progress: A systematic literature review on the integration of AI in software engineering phases and activities (2013–2023). *Proceedings of the IEEE International Conference on Artificial Intelligence Applications*. : 1–6.
- Dwivedi, Y. K. *et al.* 2023. Evolution of artificial intelligence research in Technological Forecasting and Social Change: Research themes, trends and trajectories. *Technological Forecasting and Social Change*. 192 : 1–27.
- Fu, Y. *et al.* 2024. Unleashing the potential: AI empowered advanced metasurface research. *Nanophotonics*. 11 (20) : 1–20.
- Gerdon, F. *et al.* 2022. Social impacts of algorithmic decision-making: A research agenda for the social sciences. *Big Data and Society*. 9 (1) : 1–15.
- Grace, K. *et al.* 2018. When will AI exceed human performance? Evidence from AI experts. *Journal of Artificial Intelligence Research*. 62 : 729–754.
- Grzybowski, A. 2024. A history of artificial intelligence. *Clinical Dermatology*. 42 (3) : 345–350.
- Jiang, C. and J. Wang. 2018. Intelligence originating from human beings and expanding in industry – A view on the development of artificial intelligence. *Chinese Journal of Engineering Science*. 20 (6) : 93–100.
- Jordan, M. and T. Mitchell. 2015. *Machine learning: Trends, perspectives, and prospects*. *Science*. 349 (6245) : 255–260.
- Kaplan, A. and M. Haenlein. 2019. Siri, Siri, in my hand: Who’s the fairest in the land? On the interpretations, illustrations, and implications of artificial intelligence. *Business Horizons*. 62 (1) : 15–25.
- Lavazza, A. 2025. Two ways of considering the ethics of artificial intelligence. *Digital Development*. : 15–36.
- LeCun, Y., Y. Bengio and G. Hinton. 2015. Deep learning. *Nature*. 521 (7553) : 436–444.
- Mathew, A., P. Amudha and S. Sivakumari. 2020. Deep learning techniques: An overview. *Advances in Intelligent Systems and Computing*. 1141 : 599–608.
- Ooi, K. B., *et al.* (2025). The Potential of Generative Artificial Intelligence Across Disciplines: Perspectives and Future Directions. *Journal of Computer Information Systems*, 65(1), 76–107.
- Schmidhuber, J. 2015. Deep learning in neural networks: An overview. *Neural Networks*. 61 : 85–117.
- Siebert, M. *et al.* 2021. Comparing the emergence of technical and social sciences research in artificial intelligence. *Frontiers in Computer Science*. 3 : 1–13.
- Taddeo, M. and L. Floridi. 2018. How AI can be a force for good. *Science*. 361 (6404) : 751–752.
- Taddeo, M. and L. Floridi. 2021. How AI can be a force for good – An ethical framework to

- harness the potential of AI while keeping humans in control. Ethics, governance, and policies in artificial intelligence: 91–96.
- Tai, M. C. T. 2020. The impact of artificial intelligence on human society and bioethics. *Taiwan Journal of Obstetrics and Gynecology*. 59 (5) : 669–671.
- The Nobel Prize. 2024. The Nobel Prize in Physics 2024 – Press release. NobelPrize.org. (Accessed: 30 November 2025).
- Kaur, R. and D. Sharma. 2024. Application and significance of mathematics in artificial intelligence. In: *Proceedings of the International Conference on Soft Computing and Signal Processing*. Springer. : 673–684.
- Voigts, E., D. Elflein and J. Röhnert. 2024. Artificial intelligence – intelligent art? An introduction. *Artificial intelligence – intelligent art?* Transcript Verlag, Bielefeld, Germany. : 9–34.
- Xie, Y., Y. Zhai and G. Lu. 2025. Evolution of artificial intelligence in healthcare: A 30-year bibliometric study. *Frontiers in Medicine*. 11: 1505692.
- Yu, L. *et al.* 2025. Artificial intelligence in traditional Chinese medicine: Advances in multi-metabolite multi-target interaction modeling. *Frontiers in Pharmacology*. 16: 1–8.
- Zuboff, S. 2015. Big other: Surveillance capitalism and the prospects of an information civilization. *Journal of Information Technology*. 30 (1) : 75–89.