Engineering Education Review

EDITORIAL

From Narrow Intelligence to General Intelligence: Technological Evolution, Industrial Logic, and Educational Transformation of Artificial Intelligence

Bo Zhang[1], Juan Yang[2,*]

[1]Department of Computer Science and Technology, Tsinghua University, Beijing 100084, China
[2]Higher Education Research Institute, China Agricultural University, Beijing 100083, China

ABSTRACT

Current mainstream artificial intelligence (AI) does not simply replicate the internal mechanisms of human intelligence; rather, it primarily simulates human intelligent behaviors—such as thinking, perception, action, and emotion—thereby charting a development path for "silicon-based intelligence" that remains distinct from human intelligence. This study systematically maps the trajectory of AI across three dimensions: technological paradigms, industrialization logic, and educational transformation. First-generation AI, centered on knowledge-driven approaches and symbolic reasoning, possesses strong explainability but suffers from high knowledge-acquisition costs and narrow application domains, making large-scale industrialization difficult. Second-generation AI, powered by big data and deep learning, achieved breakthroughs in perception, recognition, and prediction, yet confronts challenges such as black-box opacity, unreliable execution, and ambiguous liability attribution. Third-generation AI, represented by Large Language Models (LLMs), exhibits emergent abilities through the scaling of parameters and data; this shifts machines from specific task executors to open-domain language interactors and content generators, while evolving further toward general intelligence systems capable of operating stably in complex environments. The key to AI industrialization lies not in the mere pursuit of model performance, but in the concurrent maturation of domain adaptation, explainability, reliability, engineering implementation, hardware-software co-design, and business models. Confronting the labor restructuring brought about by AI developments, educational systems must synchronously reform their foundational paradigms, instructional methods, learning modes, governance structures, and evaluation mechanisms, while embedding AI literacy throughout basic education, higher education, and lifelong learning.

Keywords: artificial intelligence, large language models, industrialization, artificial general intelligence, educational transformation, AI literacy

ARTIFICIAL INTELLIGENCE: FROM IMITATING INTELLIGENCE TO SIMULATING INTELLIGENT BEHAVIOR

Public understanding of artificial intelligence frequently suffers from a fundamental misconception: equating AI with a direct replication of the human brain or human cognitive mechanisms. In fact, the primary objective of mainstream AI technologies is not to replicate the internal mechanisms underlying human intelligence, but to achieve the functional realization of intelligent behavior in machines. Because a unified, operational defini-

***Corresponding Author:**
Juan Yang, Higher Education Research Institute, China Agricultural University, Beijing 100083, China. Email: yangjuan@cau.edu.cn; https://orcid.org/0000-0003-1049-502X

Received: 14 September 2026; Revised: 9 October 2026; Accepted: 19 October 2026
https://doi.org/10.67229/eer.26100026

This is an open access article distributed under the terms of the Creative Commons Attribution-NonCommercial-ShareAlike 4.0 International License, which allows others to copy and redistribute the material in any medium or format non-commercially, as long as the author is credited and the new creations are licensed under the identical terms.

Table 1: Technological Paradigms and Industrialization Characteristics Across Three Generations of AI

Generation & Paradigm	Core Mechanism	Primary Capabilities	Prominent Advantages	Major Boundaries & Limitations
First-Generation: Knowledge-Driven	Knowledge bases, rules, & symbolic reasoning	Reasoning, diagnosis, planning, scheduling	Explainable & traceable	Narrow domains, high knowledge-acquisition costs
Second-Generation: Data-Driven	Big data, deep learning, & artificial neural networks	Recognition, perception, prediction, screening	Strong learning capacity, high recognition precision	Black-box opacity, heavy data dependence, liability dilemmas
Third-Generation: Model-Driven	Foundation models, scale expansion, & capability emergence	Open dialogue, generation, transfer, & tool utilization	Universal interface, multimodal, highly interactive	Hallucinations, high cost, safety concerns, domain reliability issues

tion of intelligence remains elusive, what can be observed and measured are primarily its external behavioral manifestations. Consequently, evaluating the efficacy of an AI system rarely depends on whether its internal mechanism is identical to that of the human brain, but on whether its output achieves human-like functional effects.

Accordingly, two distinct technological pathways can be delineated:

1. **The Behavioral Simulation Pathway.** The mainstream approach oriented toward behavioral simulation, which utilizes computational models, data, and engineering systems to enable machines to demonstrate capabilities in thinking, perception, action, and interaction. This route generates silicon-based intelligence, which fundamentally differs from human intelligence.

2. **The Internal Mechanism Pathway**. The brain-like computing approach oriented toward simulating internal mechanisms, which seeks to understand and reproduce the brain's information processing and cognitive architecture. While the latter holds immense scientific value, human understanding of brain mechanisms remains limited; hence, it has yet to establish an industrial foundation comparable to that of mainstream AI. The major achievements of the current AI industry are almost entirely built upon the former pathway.

From a functional perspective, human intelligent behavior can be summarized into four domains: thinking, perception, action, and emotion. *Thinking* encompasses reasoning, decision-making, diagnosis, design, planning, and creation; *perception* includes visual, auditory, tactile, and other environmental information acquisition; *action* involves manual manipulation, physical mobility, and intervention in the physical environment; and *emotion* is expressed through language, voice, facial expressions, and body language. In real-world settings, these four dimensions of intelligent behavior are closely interconnected, forming a closed-loop process of perception, reasoning, decision-making, action, and feedback. The recent rise of AI agents essentially reflects the organization of these capabilities into a system that continuously interacts with and acts upon its environment.

This perspective provides a critical prerequisite for analyzing the AI industry: technological breakthroughs do not automatically convert into industrial success. Whether AI can enter real-world production and daily life depends not only on algorithmic sophistication, but also on whether outputs are interpretable, systems are reliable, hardware is adapted, liability is clearly defined, costs are acceptable, and users and their social environments are adequately prepared. Consequently, the industrialization of AI is far more complex than that of conventional information technology, requiring the synchronized evolution of technology, institutions, and human capabilities.

GENERATIONAL EVOLUTION OF AI AND ITS INDUSTRIAL BOUNDARIES

The development of artificial intelligence has not been a linear progression, but a continuous shifting among different technological paradigms. Based on core resources and primary capabilities, this development can be summarized into three stages: knowledge-driven, data-driven, and model-driven (Table 1). These three stages do not represent simple chronological replacements; while newer paradigms expand the boundaries of machine capability, earlier technologies retain unique value in scheduling, rule-based reasoning, and high-explainability tasks. Industrial practice requires combining technologies according to task demands rather than dismissing the old for the new.

First-generation AI relied on human knowledge and experience as its core resources, utilizing symbolic representation, knowledge bases, and reasoning engines to simulate expert thinking processes. Expert systems for medical diagnosis represent a typical application: professional medical knowledge and diagnostic logic were encoded into a knowledge base, enabling the system to perform rule-based inference. The primary advantage of such systems lies in their explainability and traceability—

users can understand the underlying rationale behind a conclusion, making these systems well-suited for high-stakes decision-support tasks. However, their limitations are equally pronounced: the cost of acquiring and encoding knowledge is exceptionally high, and a system is typically confined to a specific domain or task type, making cross-domain transfer difficult. While highly valuable for enterprise scheduling, computer-aided diagnosis, and planning, first-generation AI struggles to scale into universal products for broader consumer markets.

Second-generation AI treats data and learning as its primary resources. By training deep neural networks on large-scale datasets, capabilities in image recognition, speech processing, anomaly detection, and predictive modeling have improved dramatically. Unlike first-generation AI, which relied on explicit knowledge, second-generation AI can extract features from data, reflecting the underlying technological premise that learning constitutes a fundamental source of intelligence. Its applications span facial and speech recognition, preliminary screening of medical images, cybersecurity, financial risk management, sales forecasting, and public opinion analysis, contributing to a marked acceleration in the development of the AI industry around 2010.

Nevertheless, data-driven models usually lack sufficient explainability. In high-risk domains such as medical diagnosis, judicial adjudication, and autonomous driving, performance metrics cannot substitute for normative liability judgments. In medical imaging, for instance, even if recognition accuracy reaches impressive thresholds, the output can typically serve only as a preliminary screening tool, with final diagnostic accountability remaining with certified professionals. Similarly, autonomous driving and service robotics remain constrained by round-the-clock environmental perception and system-wide reliability requirements. A technology's ability to perform a particular task under experimental conditions does not necessarily imply that it can operate reliably in open, complex, and unpredictable real-world environments.

Third-generation AI is represented by Large Language Models and multimodal foundation models. Since 2022, LLMs have demonstrated capabilities in open-domain dialogue, content generation, contextual interaction, and cross-task transfer, transforming AI from dedicated single-task tools into general-purpose interfaces capable of handling diverse cognitive assignments. The core shift here is not that machines have achieved human-like cognition, but that the expansion of parameter scales, training data, and compute investments has triggered emergent capabilities. As machines evolved from merely classifying and predicting to generating, explaining, and interacting, the AI industry entered a new phase of rapid growth.

The creativity and non-determinism of large language models stem from the same underlying mechanism. Open-ended generation necessitates maintaining output diversity; however, while this diversity enables the generation of novel expressions, it also introduces factual inaccuracies and so-called 'hallucinations.' Governance of large models should not naively pursue absolute zero-error outputs. Instead, generative freedom should be calibrated against task risk: permitting higher divergence during creative design and exploratory brainstorming, while strictly enforcing Retrieval-Augmented Generation (RAG), rule constraints, fact verification, and human-in-the-loop oversight in healthcare, finance, and mission-critical operations. Tiered governance tailored to specific application scenarios offers a pragmatic path for balancing innovation with safety.

FROM TECHNOLOGICAL CAPABILITY TO INDUSTRIAL CAPABILITY: THE REALIZATION LOGIC OF AI INDUSTRIALIZATION

Task-Centric Application Design

The foundational principle of AI industrialization is placing real-world tasks at the center. A major reason for the slow industrial expansion of first- and second-generation AI was their hyper-specialization, which restricted systems to single organizations or isolated scenarios. Although foundation models provide a broad, general interface, enterprises must still embed domain knowledge, operational workflows, and organizational rules into these systems to deliver deployable products. Consequently, Retrieval-Augmented Generation, domain-specific knowledge bases, tool orchestration, and agentic workflows are assuming greater importance, collectively converting general language capabilities into specialized task competencies.

Absolute Reliability as the Industrial Baseline

Reliability serves as the industrial baseline for service robotics, autonomous vehicles, and mission-critical enterprise systems. When service robots interact with vulnerable populations—such as the elderly, children, or medical patients—a single operational failure can cause severe harm. Autonomous vehicles operating under adverse conditions (wind, rain, snow, and mixed traffic) must similarly handle long-tail events that laboratories cannot exhaustively simulate. Thus, evaluating an AI product requires looking beyond demonstration effects and average accuracy rates to scrutinize failure modes, edge-case resilience, safety redundancies, auditability, and continuous maintenance. The closer a technology gets to physical safety and human life, the higher its reliability requirements become.

Architectural Logic for Embodied AI

Embodied intelligence should follow the principle of

maximizing software generalizability while tailoring hardware configurations to diverse application scenarios. Different physical tasks demand distinct levels of force, precision, speed, structure, and sensory modalities. Enforcing complex humanoid forms across all domains purely for aesthetic universality leads to inflated costs, degraded reliability, and mismatched market demand. A more viable industrial path is pairing universal foundation models and control software with diverse, task-specific hardware, enabling domain-tailored optimization across industrial manufacturing, healthcare, household care, agriculture, and specialized field operations.

Engineering Capacity as a Bridge to Scale

Engineering capabilities can translate technological potential into advantages of scale. While model performance is crucial, cost optimization, inference/training efficiency, supply chain organization, system integration, and rapid product iteration play decisive roles in market success. China's comparative advantages in the artificial intelligence industry lie primarily in its capacity for rapid engineering and deployment, cost reduction, and large-scale application, while its capabilities in original theoretical research and breakthrough innovation still need to be strengthened. Future competitiveness cannot permanently rely on cost-optimizing existing technical trajectories; it demands the capacity to independently propose original models, architectures, and technological paradigms.

High-Risk Investment Mechanisms

AI entrepreneurship requires financing mechanisms tailored to high-risk innovation. Research and development in AI demand massive capital investments accompanied by high failure rates. If this financial risk is borne entirely by researchers or individual entrepreneurs, innovative activity will be severely suppressed. A more effective mechanism leverages venture capital to pool and diversify risk, allowing professional investment institutions to absorb high uncertainties while enforcing market-driven governance over project directions. This must be coupled with robust intellectual property protection, technology transfer frameworks, and viable exit mechanisms. Crucially, venture funding cannot replace demand validation: an entrepreneurial initiative retains long-term value only when its technology solves genuine problems and generates sustained user willingness to pay.

Co-evolution of Technology, Institutions, and Human Behavior

Technological deployment must adjust synchronously with social rules and human behavior. The challenges faced by autonomous driving in mixed-traffic environments demonstrate that even the most advanced technology cannot operate in isolation from societal structures. Regulatory awareness, public infrastructure, industry standards, ethical frameworks, and societal trust form integral parts of the industrial ecosystem. Governance of AI, therefore, involves more than constraining algorithms; it encompasses remodeling application environments, clarifying organizational liabilities, and raising user literacy. Industrial policy should shift from supporting isolated technologies toward building an integrated ecosystem of "technology, institutions, scenarios, and talent."

AI-DRIVEN EDUCATIONAL TRANSFORMATION AND TALENT RECONSTRUCTION

The impact of AI on the labor market concentrates primarily on white-collar occupations centered around information processing, expression, judgment, and strategy formulation. Fields such as law, finance, social services, education, and academic research face profound restructuring. Crucially, this shift is not a wholesale substitution of humans by machines, but a reallocation of task labor between humans and machines. Individuals who understand AI, formulate high-quality prompts, evaluate output quality, and take ultimate responsibility will progressively replace those who lack these competencies. The ability to collaborate effectively with AI is becoming a foundational professional skill.

Accordingly, the educational system must undergo five systemic transformations:

1. Paradigm Shift from Knowledge Acquisition to Value-Added Competencies. As foundation models can instantly retrieve vast bodies of knowledge, education can no longer treat memorization and repetition as primary objectives. Instead, it must prioritize problem definition, evidence evaluation, interdisciplinary understanding, creative expression, ethical responsibility, and action in complex scenarios. Students must move beyond merely "knowing how to use a tool" to understanding its boundaries—identifying hallucinations, structural biases, and uncertainties while accepting accountability for final outcomes.

2. Instructional Transition from Uniform Lecturing to Human-AI Collaboration. Educators can leverage AI to generate preliminary instructional materials, design differentiated tasks, provide real-time feedback, and assist in diagnostic learning assessments—thereby reallocating their energy toward high-order mentoring, emotional support, and value orientation. Concurrently, AI-generated content should not enter classrooms without rigorous verification. Teachers must establish clear pedagogical guidelines specifying allowable usage, disclosure protocols, fact-checking procedures, and student data privacy safeguards.

3. Learning Mode Shift from Standardized Answers to Iterative Co-Creation. Proficiency with large models does not mean retrieving answers in a single prompt; it involves an iterative cycle of questioning, evaluating drafts, adding supporting evidence, refining logic, and continuously interacting with the model to elevate output quality. The higher a user's domain knowledge, the better equipped they are to spot flaws in model outputs and create works with personal insight and professional depth. Thus, AI does not render learning obsolete; it amplifies the importance of foundational knowledge, domain expertise, and metacognitive capabilities.

4. Evaluation Reform from Outcome Grading to Process-Based Evidence. Traditional written assignments are easily completed by generative AI, making it difficult to gauge a student's true competence from final texts alone. Educational evaluation must shift toward assessing problem formulation, literature selection, prompt design, fact-checking workflows, revision histories, oral defenses, and real-world application. Project-Based Learning (PBL) should be leveraged to introduce technological frontiers, industry needs, and authentic societal issues into the curriculum. Using complex assignments as vehicles allows educators to assess both a student's ability to use AI tools and their capacity to think critically beyond them.

5. Universalization of AI Literacy Across All Educational Stages. AI literacy should not be restricted to computer science majors; it must serve as a foundational public competency akin to digital literacy. At the basic education level, students should be equipped with foundational concepts, appropriate habits of AI use, safety awareness, and ethical judgment. At the higher education level, the deep integration of AI across disciplines should be promoted to develop curricula and practical training that integrate AI with disciplinary learning. Continuing education, meanwhile, should provide working professionals with support for career transitions. The objective is not to turn every individual into an algorithm engineer, but to empower everyone to learn, work, and act responsibly in an AI-infused society.

International engineering education offers a vital perspective on these reforms. Talent development in AI must align with domestic industrial demand while adhering to international engineering standards, professional accreditations, and global engineer qualification frameworks. Curricula should bridge technical capabilities with project management, teamwork, engineering ethics, cross-cultural understanding, and social responsibility. In international engineering projects, many non-technical barriers relate directly to compliance standards, governance rules, communication, and sign-off authority. Advancing professional accreditation, qualification reciprocity, and international exchange will elevate global competency among students and bolster the international influence of engineering education.

CONCLUSION AND OUTLOOK

The evolution of artificial intelligence reflects a clear paradigm shift: moving from knowledge-driven to data-driven approaches, and now to model-driven architectures. Namely, knowledge-driven plus data-driven. It also mirrors an extension of machine capabilities from single, narrow tasks to open interaction, and from digital environments into the physical world. The explainability of first-generation AI, the perceptual learning capacity of second-generation AI, and the generative, transfer capabilities of third-generation AI each possess distinct advantages. Future intelligence systems will likely integrate diverse models, structured knowledge bases, and specialized software tools, rather than being dominated by any single, monolithic technology.

From an industrial perspective, the true value of AI lies not in producing impressive technical demonstrations, but in solving real-world problems stably, reliably, and cost-effectively. Scenario demands, domain expertise, engineering execution, hardware-software integration, institutional environments, and venture capital mechanisms collectively determine whether a technology can mature into a viable industry. Enterprises and policymakers must capitalize on the window of opportunity offered by foundation models while remaining cautious of hype detached from practical demand, strategically directing resources toward products and services capable of yielding sustainable value.

From an educational perspective, AI is driving a systemic restructuring of talent development goals and operational paradigms. The core task of future education is not to compete with machines in memory capacity or text generation speed, but to cultivate human beings capable of directing tools, auditing outputs, and taking responsibility for their actions. To this end, educational institutions must advance AI literacy across all academic levels, rebuild curricula and evaluation frameworks around authentic projects, empower educators to navigate pedagogical transitions, and embed professional accreditation and global competency into talent cultivation pipelines. Only through the synchronized advancement of technological progress, institutional innovation, and human development can artificial intelligence truly be transformed into a force that drives industrial upgrading and societal advancement.

DECLARATIONS

Conflict of Interest

The authors declare that there is no conflict of interest.