AI Science and Engineering: A New Field

Longbing Cao, *University of Technology Sydney, Sydney, NSW, 2007, Australia*

Artificial Intelligence (AI) first emerged in the 1950s and over the next few decades, experienced both advances and obstacles.[1] However, there is now a new era of AI. The history of AI has been an acceleration from *object intelligence* (e.g., on symbol, behavior, and agent) to *system intelligence* (e.g., human, nature, and society), and from *individual intelligence* (e.g., learning intelligence) to *metasynthetic intelligence* (hybridizing and synthesizing intelligences). Building on the tumultuous AI evolutions, this new-generation AI is accelerating its pace of innovating, differentiating, transforming, and reshaping the world. The new-generation AI not only enables a smarter and more resilient humanity, well-being, and economy, but also everything else. What lessons can we learn from reviewing these AI advances and challenges? What makes AI science and engineering (AISE, or intelligent science and technology) a solid and comprehensive scientific field in addition to transforming other scientific and engineering disciplines and translating businesses and economy into their smart editions? What are the fundamental questions to be addressed in AISE? What forms the body of knowledge of AISE? What type of profile should AI professionals have to meet the requirements of AISE? What should AI education look like to produce qualified AI professionals? These questions deserve enduring, comprehensive, deep, creative, and critical thinking, ideas, and actions first and foremost to establish the AI field. Here, I share my limited view on AISE as a new discipline and the imperative developments, including the AI profession and AI education, to drive and enable the *intelligent digital era* and *Industry 4.0*. I hope my humble opinions will spur valuable debate and exchange and systematic and strategic developments of AISE in the broad AI community.

AISE BODY OF KNOWLEDGE

Is AISE a new discipline, or is it an independent discipline? What comprises the discipline of AISE? On the one hand, there has been a substantial spectrum of AI developments that sufficiently formulate the disciplinary knowledge map and body of knowledge of AI. This is evidenced by the various research areas, tasks, and achievements made over the 70 years of AI (see Figure 1: the landscape of AI evolution, in Cao's work[2]). On the other hand, AI is a young discipline, still evolving and requiring comprehensive and strategic thinking and development. Figure 1 illustrates the body of knowledge of AISE. In this article, I briefly discuss some of the areas.

1541-1672 © 2022 IEEE
Digital Object Identifier 10.1109/MIS.2022.3151993
Date of current version 11 April 2022.

AI Science

AI science addresses fundamental disciplinary questions in AI. What makes AI a natural science rather than just transforming sciences and delivering techniques for building intelligent machines? What constitutes AI foundations beyond neuroscience?[3] What makes "intelligent science" a major scientific discipline like physics? Or will AI science follow the path of social science, such as finance and management science, by focusing on applying relevant natural sciences to studying intelligent worlds?

Accordingly, a critical disciplinary agenda is to develop and enrich the body of knowledge for AI science. This includes the following: 1) *AI fundamentals* to understand and quantify intelligence and intelligent worlds, such as using neuroscience, human mental models, natural evolution theories, biology, and societal awareness; and understanding the essence of intelligent phenomena, behaviors, and consequences, such as by selection, immunity, evolution, and adaptation; 2) *AI technologies* to acquire, represent, compute, learn, discover, transfer, implement, and optimize "intelligence" from knowledge, behavior, emotion, sensing, vision, conversation, interaction, and response and to develop intelligent machines, such as robotics, driverless cars, and UAVs; and 3) *AI foundations*, such as philosophy, cognitive, and psychological sciences, and probabilistic programming to inspire and formulate the abovementioned fundamentals and technologies.

AI Engineering

Will *AI engineering* become a widely applicable engineering area like electrical engineering and computer

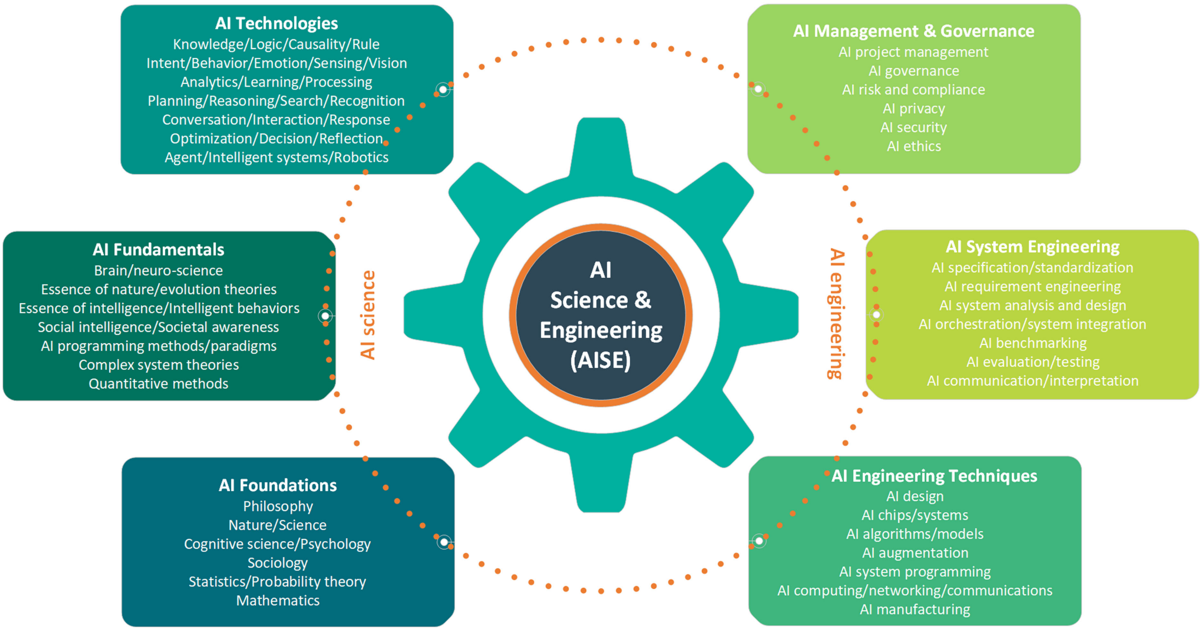

FIGURE 1. AI science and engineering: the body of knowledge.

engineering, or an applied area like financial engineering? What techniques, skills, and processes are required in transforming AI science to intelligent systems and applications? How to define, characterize, and manage intelligent problems, behaviors, and observations? What will ensure the standard, quality, and governance in applying AI technologies to solve real-life problems and applications? How to design, conduct, evaluate, and manage experiments and case studies for better AI engineering and practice? These are essential questions to be addressed in AI engineering.

Accordingly, I categorize AI engineering into following three layers: 1) *AI engineering techniques:* consisting of techniques for AI designs, AI chips, AI systems, AI algorithms/models, and AI computing, etc.; 2) *AI system engineering:* such as AI standardization, AI requirement engineering, AI system analysis and design, AI orchestration, and AI benchmarking; and 3) *AI management and governance:* composed of AI project management, AI risk and compliance, AI privacy and security, and AI ethics. Many of these areas involve interdisciplinary techniques and achievements made in the relevant engineering areas, such as software engineering, information systems, and intelligent services.

AI SCIENTIFIC REVOLUTION

Over the past 70 years, AI has evolved from a fuzzy motivation and ambition to a rich research area, which has become a truly universal field. The ups and downs of the AI evolution have resulted in significant transformations in both AI research and applications.[1]

Like other ever-evolving research areas, the two ages of AI[2] have experienced the following eight major disciplinary transformations and paradigm shifts:

- *From general to specialized aims:* the shift from weak artificial general intelligence to strong artificial specialized intelligence; and from GPSs with agnostic and generalized reasoning, planning, and learning to domain-specific and specialized intelligent behaviors, devices, architectures, and algorithms.
- *From closed to open domains:* the research on AI problems and intelligent applications has expanded from specific and known domains, such as game theory and robotics, to universal, new and open domains, such as translational medicine and Internet finance.
- *From simplified to real-life problems:* including migrating from highly simplified and manipulated scenarios to addressing the reality of real-world problems, systems, behaviors, and data, such as the shift from task-specific individual autonomous agents to dynamic context-aware human-guided multiagent systems.
- *From hypothesis-driven to hypothesis-free settings:* conventional research and methods typically assume certain hypotheses [such as rationality, independent

and identically distributed (IID) assumption, and constraints] on problems, data, and solutions; however, modern AI research focuses on unknown and invisible states and spaces with partial to no hypotheses (limited rationality, non-IID data, and partial constraints), as illustrated by the transition from logical planning to unsupervised exploratory reinforcement.

- *From shallow to deep design:* the past theories and tools were featured by lightweighted, handcrafted, and small-scale designs, while the current ones cope with both lightweight (e.g., for mobile and IoT devices and applications) and massive large-scale (e.g., deep neural networks with millions of parameters) system components and strategies, as illustrated by the transform from classic Monte Carlo reinforcement to deep Q-learning reinforcement.
- *From individual to blended approaches:* the past focus on individual research topics and methods has transferred to hybrid and ensemble methods for modern ecosystems, including transforming from reasoning only to consolidating and hybridizing reasoning with learning and optimization.
- *From small to large scale:* in contrast to traditional research where small samples are targeted, today's AI is simply a game of scale, in the sense of data size, search space, computing infrastructure, and power, AI accelerator and processor, and beneficiary end-users.
- *From narrow to broad communities:* many past specific, logic-centric research groups have evolved to diversified research institutions; the AI communities are multi, cross, and transdisciplinary; AI not only dominates the core ICT areas, but also dives into social science, such as archaeology and finance and other disciplines, such as physics, chemistry, biology, and medicine science.

Table 1 illustrates the above AI transformations w.r.t. transitional highlights in the two ages: 1) *philosophy* in terms of mind, behavior, nature, and universe; 2) *aims* in terms of intelligence and objective; 3) *foundations* in terms of philosophy, statistics and mathematics, neuroscience,[3] and psychology, ICT and engineering, and social science; 4) *research areas* in terms of capability, capacity, approaches, and designs; 5) *applications* in terms of domains, performance, and evaluation; 6) *profession* in terms of research and business; and 7) *education* in terms of higher education, secondary education, and primary education. In the following, we illustrate several aspects of these AI evolutions: AI thinking, intelligence, complexity in AI, uncertainty in AI, AI programming, AI design, AI resources, AI computing, and AI engineering.

AI Paradigmatic Shift

The evolution of AI thinking: The original ambition of creating "thinking machines" that can think, perceive, and act humanly and rationally has substantially expanded. Today, AI thinking highlights: 1) creating intelligence and intelligent machines that can think, perceive, act, and reflect humanly, naturally, socially, and ethically; 2) making intelligent designs, mechanisms, strategies, architectures, and interfaces inspired, enabled by and involving human, nature, society, domain, data, network, computation, and any other human civilization; and 3) pursuing higher level (both general and strong, human-level, hybrid, transdisciplinary, cross-domain, and human–machine-integrated) to super (superhuman, supernatural, and human–nature–machine-integrated) intelligence and performance.

The evolution of intelligence: The concept of intelligence and its scope and content are continually evolving. Its early focus on exploring symbolic intelligence, connectionist intelligence, situated intelligence, and nature-inspired intelligence has shifted to the recent highlight on both advanced human intelligence (HI), such as common sense, emotion,[4] consciousness, learning, imagination, and reflection, and other diversified nonhuman *X-intelligences*, including data intelligence, social intelligence, organizational intelligence, and environmental intelligence, and their metasynthesis.[5]

The evolution of complexity in AI: Complexity is a key challenge in AI. Its definition, type, and degree have evolved substantially, reflected in 1) the scope, aims, resources, and expected outcomes of AI problems, research, and practice; 2) the openness, uncertainty, diversity, scale, hierarchy, heterogeneity, interactions, distributions, and dynamics of target problems and intelligent agents, systems, and devices; 3) the advancements of AI settings, strategies, designs, mechanisms, and architectures; and 4) the adversary, negative to uncertain and unknown consequences, and the impact of AI behaviors, systems, and applications.

The evolution of uncertainty in AI: Uncertainty has been a fundamental concern in AI and its research landscape. Surpassing hypothesis and certainty factors, uncertainty evolves 1) as a research challenge of observations, such as uncertain states, spaces, domains, knowledge, actions, behaviors, processes, interactions, relations, dynamics, and changes; 2) as a reflection of unknownness, such as unknown entities, states, behaviors, relations, structures, distributions, and evolution; 3) as an effect of AI tasks and behaviors, such as positive and negative consequences, risks, and adversaries; and 4) for problem-solving approaches, such as nondeterministic reasoning,

TABLE 1. AI Disciplinary and paradigm shift from the old age to the new age: From a research area to a universal field.

Aspect	Dimension	Old age (before year 2000)	New age (2000 onward)
Philosophy	Mind	Simulate how humans think (mental model)	Develop brain-like cognitive and neural systems
	Behavior	Simulate how humans act and value	Act, communicate, decide and reflect rationally, critically, creatively and ethically
	Nature	Inspired by natural (physical, chemical, and biological) systems	Inspired by and develop natural and supernatural intelligence
	Universe	Human and nature-inspired machine intelligence	Ubiquitous X-intelligences inspired by the universe
Aims	Intelligence	Human-like machine intelligence and nature-inspired computational intelligence	X-intelligences including human, nature, data, machine (network, computing), behavioral, organizational and social intelligences
	Goal	Rational intelligent agents	Intelligent systems that can think and behave humanly, naturally, socially and rationally
Foundation	Philosophy	Reductionism	Reductionism, holism, systematism
	Statistics mathematics	Logic, probability, optimization	Logic, probability, algebra, numerical methods, optimization
	Neuroscience psychology	Mental model, behaviorism	Brain informatics, cognitive science, psychology, human–machine interaction and partnership
	ICT and engineering	Computer engineering, information theory, cybernetics/ control theory, etc.	Computer science, data/information science, machine/deep learning, quantum computing, communications and networking, intelligent device and manufacturing, etc.
	Social science	Economics (decision theory, operations research)	Sociology, economics, finance, management science, ethics
Areas	Capability	Perceive, plan, reason about, act and optimize	Perceive, plan, reason about, learn, decide, act, optimize, communicate, reflect
	Capacity	Point-based, component, simple systems	Process, system, large to giant, complex systems
	Approaches	Knowledge, reasoning, planning	Knowledge, reasoning, planning, interaction, communication, decision-making, reflection
	Designs	Hypothesis test, symbolism, behaviorism, probability	Hypothesis free, data-driven discovery, emotion, intention, experience, feedback
Applications	Domains	Reasoning/search, planning, game, robotics, simulation	Anytime, anywhere and anything, including business, government, society, military, outer-space
	Performance	Human-like	Human-level or super human/natural-level
	Evaluation	Technical measure, such as accuracy, error, and risk statistics	Technical performance and business effect including actionability, reproducibility, accountability, trust, explainability, ethics, and impact
Profession	Research	Individual researchers and groups	General all-level research roles and institutions
	Business	Specific system/application developers	General all-level professionals
Education	Higher	Undergraduate subjects and knowledge modules	Undergraduate to master's and doctoral courses
	Secondary	None to limited knowledge modules	Compulsory to elective subjects
	Primary	None to limited knowledge modules	Compulsory to elective subjects

approximate optimization, probabilistic programming, stochastic inference, adversarial learning, variational inference, and nonoccurring behavior analytics. Uncertainty does not rule out causality, which explores the cause and effect of MI.[6]

The evolution of AI programming: Here, AI programming broadly refers to the approaches and processes of solving AI problems. AI programming has extensively evolved in terms of 1) deterministic to nondeterministic programming; 2) precise to approximate and probabilistic inference; 3) differentiable to nondifferentiable programming; and 4) machine-programmable to human–machine-cooperative programming.

The past AI programming has been quantitative, logical, precise, calculable, differentiable, back-propagatable, tractable, and machine-programmable for intelligent processing, computing, optimization, management, and systems. This has been further inflated by end-to-end deep neural networks and automated machine learning. The challenges and openings remain with nondifferentiable and intractable AI problems and advancing approximate and fair sampling-based, prioritized, and strategical AI programming theories and methods.

The evolution of AI design: AI design thinking, tasks, and techniques have been enhanced in terms of 1) *mechanism:* to support from reactive and situated to interactive, evolutionary, adaptive, proactive, online, and real-time intelligent behaviors, functionalities and systems; 2) *scale:* from nano, micro, and meso to macro scale; from small closed problems and systems to open giant intelligent problems and systems; from point-based and small-scale to process-oriented and large-scale components and systems; from a handful to millions or even billions of parameters in models and systems; 3) *diversity:* from pure and individual to multiple, distributed and heterogeneous aims, tasks, domains, methods, data modalities, and sources; 4) *depth:* from flat and shallow to deep and hierarchical problems, models, architectures, and systems; from observations and know-what to deep insights and explanation of know-how and know-why; and from the identification, recognition, and extraction of explicit phenomena and scenarios to the discovery and explanation of implicit and unknown opportunities.

The evolution of AI resources: The involvement of AI-related data, systems, infrastructures, models and algorithms has extended from selective, manipulated, proprietary, in-house and centralized resources to open, realistic, shared, leased and distributed hardware, software, data and infrastructure. The traditional intelligence dependence on specialized domain knowledge and skills, customized infrastructures and systems, and private data and resources has been augmented and transformed into mobile, cloud, leased, and shared infrastructures, devices, and data. The roles of open data and shared AI infrastructures and algorithms are increasingly critical for enterprise AI, AI economies, and AI for social good.

The evolution of AI computing: AI computing has achieved significant improvement in 1) *computing operations:* from point, vector, matrix, tensor, and array operations to function-specific AI operations; 2) *architectures:* from general-purpose central processing units for main operations to various special-purpose processors and accelerators for specialized tasks, such as graphics processing units for graphics processing; tensor processing units for specific architectures and frameworks for neural processing and learning operations; where accelerating operations and architectures become the trending direction; 3) *computing environment:* from personal and mobile environments to the cloud, edge, and fog computing environment, where shared data, shared models, shared computing (including storage, architectures, programming, and systems) and shared services (including management and maintenance of resources, storage, models, systems, communications, and security) are commonly encouraged and available for enterprise intelligent systems; and from hardware accelerators to in-memory computing and heterogeneous computing environments.

The evolution of AI engineering: The engineering and industrialization of AI emerge as a trending economy, which requires further substantial developments in 1) *AI commercialization:* the formation and prosperity of AI-specialized vendors and service providers, including both large and multinational enterprises and small startups; 2) *AI profession and education:* a critical mass of AI scientists and engineers and the continuous training and accreditation of new-generation AI qualifications and professionals; (3) *AI engineering methodology:* the standards, specifications, frameworks, and metrics to define and govern AI project management and evaluation; and for AI requirement engineering, system analysis, design, manufacturing, and evaluation; 4) *mature AI capabilities and tasks:* including perception, recognition, detection, monitoring, learning, planning, choice-making, judgement, adaption, deployment, evaluation, optimization, and governance; 5) *AI processes:* such as from thinking and requirement analysis to design, manufacturing, service, and reproduction; 6) *AI toolkits:* widely available off-the-shelf, open-source, or shared general and configurable AI factories, software, high-performing general pretrained models, and user-friendly interfaces for scientific, educational, and commercial purposes and AI layman without the relevant qualifications; 7) *AI benchmarking:* benchmarkable and testable designs, systems, use cases, data, evaluation, and best practice; and 8) *AI applications:* both flagship AI application domains and emerging areas and opportunities.

AI Topic Transformation

In the following, we illustrate some of the main transformations of topics of interest in terms of the following dimensions:

- *aim:* from simulating partial human mental activities to modeling human-like, human-level, and super human/natural-intelligent behaviors and systems;
- *scope:* from local and personal to global, enterprise, distributed, and Internet-scale problems, computing, solutions, and usability;
- *originality:* from simulation and replication to creative and critical thinking-driven exploratory innovation;
- *target:* from small samples and predefined settings to bigger data, digitization of intelligent behaviors and systems, the complexities of problems, systems, environment and data, the analytics and learning of knowledge and intelligence, and the involvement of new types of intelligence;
- *task:* from representation, reasoning, detection, recognition and identification to understanding emotion, intention, insight, and reflection;
- *representation:* from representation engineering to representation learning and augmentation;
- *planning:* from predefined planning to learning to plan and optimizing to plan;
- *learning:* from shallow, handcrafted, and structured input–output mapping to deep, exploratory and online estimation, encoding, and discovery;
- *optimization:* from numerical and probabilistic to evolutionary, neural, and hybrid methods;
- *design:* from predefined designs, logic, and mechanisms to online, active, dynamic, real time, and predictive mechanisms and adaptive, personalized, and optimal autonomy and customization;
- *programming:* from logic centered and experimental to data- and learning-driven problem-solving;
- *factor:* from objective, observable, and internal to external and hidden factors, and subjective aspects, such as emotions, personality, and motives;
- *algorithms:* from operational to decisive and intelligent algorithms, and algorithm learning to be more intelligent (smarter in learning, adaptation, optimization over time, and process);
- *intelligence:* the "essence" of intelligence evolves from knowledge and single sources of intelligence to hybrid intelligences from human, nature, and the society and their metasynthesis in open complex intelligent systems (OCISs);
- *knowledge:* from structured to semistructured, unstructured, and new structures of knowledge;
- *autonomy:* from constrained to unconstrained, adaptive autonomy and on-the-fly autonomous decision-making;
- *openness:* from small and closed systems to semiclosed, semiopen, and open complex giant systems[5];
- *relation:* from association, correlation, and dependence to causality, coupling relationships, and interactions;
- *uncertainty:* from object-oriented and stationary uncertainty to process, distributional, structural, outcome-oriented, and nonstationary uncertainty;
- *causality:* from probabilistic causation to counterfactual relation, causal calculus, and other theories to answer why[6];
- *heterogeneity:* from individual and homogeneous to heterogeneous, multimodal, multilabel, multitask, cross-domain, and mixed scenarios;
- *rationality:* from rational to limitedly rational and reality-based (mixing rational and irrational);
- *nonlinearity:* from linear to nonlinear and exponential scalability and capabilities of data volume, complexities, transformations, search, learning, and computing;
- *visibility:* from explicit observations to implicit factors, features, interactions, relations, structures, and outcomes;
- *knownness:* from known, observable, explicit and expected to unknown, hidden, implicit, nonoccurring, and unpredictable awareness and scenarios;
- *environment:* from context-free settings to contextual, interactive, dependent, evolving/dynamic, uncertain (random), unexpected and unknown external objects, spaces, behaviors, processes and their states, influence and dynamics;
- *system:* from process and workflow construction to architectural design; from single components to ensemble, distributed, federated and hybrid systems with communications, collaborations and governance;

- *ethics:* from unethical and ethics-agnostic to ethical intelligence, design, programming, algorithms, and machines.

AI ENGINEERING

AI applications are everywhere in the human body and physical, socioeconomic, and virtual worlds. *AI engineering* standardizes and governs AI practice, formulates its methodologies and processes, and develops tools and systems to support, monitor, and evaluate AI systems, developments, and applications. AI engineering is critical for developing OCISs and provides technologies and practice for AI requirement engineering, AI design, programmable AI, AI orchestration, AI benchmarking, AI governance, and enabling actionable intelligence.

Open Complex Intelligent Systems

Intelligent systems have migrated to a new stage that focus on 1) open intelligent systems: with adaption to open domains and environments, capability for programmable problems, tasks and settings, and adjustability in system architectures, workflows, and parameterization; 2) human–machine-integrated systems: including human–machine-cooperative and human-centered intelligent systems, where humans form a constituent of the system or problem-solver or directly contribute to partial problem-solving in an end-to-end or closed form; 3) ethical enterprise intelligent systems: connecting heterogeneous business domains, multimodal data resources, local and global AI tasks, enterprise privacy and ethics, automated processing and delivery of impactful and explainable insights and results, and justified decision-making actions. For example, automated intelligent devices evolve from goal-specific, static and certain, controlled and constrained, and repetitive tasks to personalized, flexible, dynamic and evolving, uncertain, active, and online perception, tracking, updating, and adaptation. Such devices may work on live, online, or end-to-end mode.

In addressing open complex problems, the metasynthesis of *X*-intelligences generates hybrid intelligences in building OCISs as the problem-solvers.[7] Such metasynthetic AI systems involve domain experts in the process of problem-solving and support an iterative process to progressively understand, quantify, and analyze the problems and their complexities and challenges, call the relevant models and analytical/learning methods to deepen the understanding and quantification. During the problem-solving, domain experts may engage in discussions to build their consensus on the basis of a deepened understanding and quantitative evidence extracted by progressive modeling and learning.

AI Requirement Engineering

Methodologies and experience accumulated in areas, such as large-scale agent-oriented software engineering are useful for AI requirement engineering. Two extra tasks in AI requirement engineering include: 1) experience transfer: by learning general (high- and meta-level) knowledge, principle and experience (rather than specific systems) and transferring them from one domain to another that is highly different or a new domain; 2) reflection of performance, preference, and feedback: by understanding, tracking, and predicting user (agent) preference and feedback from agent-environment and agent-space interactions, including positive, negative, missing, and inaction (nonoccurring) feedback, and differentiating agent preference from decisions influenced or manipulated by external factors (e.g., misinformation, advertising, and recommendation).

AI Design

AI design becomes increasingly critical for AI science, engineering, and the economy. *AI design* develops 1) *AI design thinking:* innovative AI design principles, design patterns, design templates, and the processes and criteria for design checking, quality assurance, and ethical assurance; 2) *Intelligent devices:* general-purpose or specialized hardware to undertake common or specific AI functions and tasks, such as intelligent sensing, storage, computing, networking, communications, and control in a specific environment or for a specific purpose; 3) *AI chips:* specialized semiconductors for efficiently processing or accelerating dedicated AI tasks, such as structured programming, neural computation, and programmable learning; 4) *AI software:* empowering intelligent problem-solving and the operation of intelligent systems, such as purposeful intelligent mobile apps and automated perception, analysis, and warning systems. An increasingly important area is robotics, where intelligent soft bots (soft robots and agents) are integrated into smart and efficient AI chips to fulfill various intelligent tasks, constituting hard robotic machines, such as self-driving cars, humanoid robots, unattended drones, and UAVs. The engineering of such robotics must synthesize AI design, cognitive engineering, material engineering, mechanical engineering, electrical engineering, and computer science.

Programmable AI

The commercialization and promotion of AI engineering demand-programmable AI. Programmable AI enables and converts AI programming (problem-solving approaches) to AI engineering for the analysis and design of interactive and automated AI systems. *AI programming* supports the coding of AI theories and algorithms, such as differentiable programming, probabilistic programming, and approximate inference, into computing machines. AI engineering develops methodologies, languages, and project management tools for the analysis, design, development, management, evaluation, and governance of AI programming. Programmable AI also develops an AI algorithmic library, resource access interfaces, human/brain interfaces, system architectures, AI resources, configuration, and a computing environment.

AI factory lays the foundation of programmable AI. Shareable design patterns, modeling templates, AI algorithms, architectures, workflows, and benchmarks are essential. 1) Shareable pretrained models: built for general or specialized intelligent tasks (e.g., mobile object detection, people reidentification, and multiclass classification), design mechanisms, learning paradigms (e.g., Bayesian networks, tree models, or convolutional neural networks) or for specific business problem-solving (e.g., fraud detection, outlier detection, or customer churning). Such pretrained models can be provided in off-the-shelf or online toolkits. They may also be further customized and retrained to fit user-specific business goals and tasks. 2) Shared data: relevant and benchmarkable quality data are costly to obtain. Creating a shareable data benchmark for specific business problems or goals and usable for checking and evaluating learning designs, tasks, and systems is critical for AI engineering.

AI Orchestration

Future AI will be an orchestration of general and narrow AI, strong and weak AI, and AI over hardware, inmemory, and cloud. As the world is overloaded with complexities, its problem-solver (OCISs) has to synthesize and meta-synthesize MI with other *X*-intelligences on demand and over time. On the one hand, we expect autonomous machines to perform like humans, including replicating human cognitive and psychological functions, such as consciousness, empathy, and imagination, and obeying the rules of law, ethics, and governance. On the other hand, AI will further elevate and reform its artificial super intelligence to perform intellectual tasks that a human cannot perform well.

AI exploitation and exploration: AI has been a manifestation of exploitation and exploration. Future AI will further *exploit* new transformations and breakthroughs in specialized objectives, tasks, architectures, components, systems, designs, mechanisms, and strategies; and *explore* new thinking and far-out ideas in disciplinary, cross-disciplinary, and transdisciplinary blue-sky research and translations. Both AI exploitation and exploration are increasingly sensitive to the scale, performance, and autonomy of the aforementioned AI paradigms, functions, topics, and areas in this new age.

Actionable Intelligence

The overall goal of actionable intelligence is to satisfy the reality and provide smart decision-making actions. *Intelligence is actionable* if it 1) addresses business impact and value; 2) is secure, trustworthy, and selfregulatory to avoid negative impact or adversary consequences; 3) is operable into business; 4) predictive for next-best communications and intervention; and 5) is self-adaptive for optimization, alerting, and safe intervention. In complex intelligent machines, ethical assurance may require high-level authorization of priority tasks or emergency handling.

Actionable intelligence: It translates AI theories to best practice, including in new and more advanced applications, such as translational medicine. Enabling actionable intelligence needs to identify the gaps between AI research and problem reality, between wellcurated case studies and real-life large-scale experiments, and between customized performance evaluation and practical business value and impact. Actionable intelligence applies AI thinking and science to transform other sciences, technologies, industries, governments, the economy, and the society. Delivering actionable intelligence requires the transformation of existing AI thinking, system designs and mechanisms, workflows and processes, the output presentation and evaluation, the input/output interfaces, and user interfaces with devices and equipment.

AI Benchmarking and AI Testing

The current general-purpose or specialized benchmarks of AI-related systems, datasets, pretrained models, and algorithms face various limitations and issues. They include 1) a strong manipulation of problems, hypotheses, designs, systems, and data; 2) the mismatch between AI designs and outputs and the reality of reallife problems, requirements, and data; 3) the biases in benchmark designs, systems, data, ground truth, and performance metrics; 4) the limitations in AI test designs and methods; and 5) the limitations of performance metrics, e.g., the poor actionability of results in addressing realities.

We highlight AI testing, technical benchmarks, data benchmarks, and performance benchmarks for AI benchmarking and AI testing, which are critical for AI designs, algorithms, systems, and their engineering. 1) *AI testing:* including defining the test tasks, objectives, test designs and specifications, test platforms, test data, test metrics, and test analytics for AI hardware, software, systems, or algorithms. Such AI testing may be specified objectively and independently from the to-be-tested AI tasks and systems, but undertaken by exploring the reality of the tasks, data, and context, and the limitations (such as arbitrarity, subjectivity, misunderstanding, manipulation, and exclusion of context) of designs and systems. 2) *Technical benchmarks:* including shareable systems and models that provide a benchmark for a) specific problems (e.g., learning tasks, such as classification problems, time-series forecasting problems); b) AI tasks that address the problems; and by c) specific systems (e.g., infrastructures, such as mobile NPUs or learning architectures, such as convolutional neural networks), designs, and mechanisms (e.g., learning tricks, such as batch normalization for deep neural networks, Gibbs sampling for Bayesian networks, and reward/value functions for reinforcement learning). 3) *Data benchmarks:* including shareable datasets that reflect the reality of target problems, own intrinsic characteristics, and provide unbiased ground truths for evaluation; robust data samplings for unbiased and benchmarkable data samples; having structured specifications and evaluation of specific real-life data characteristics and complexities (e.g., non-IID data, networking data, outlying data, noisy, or missing data); and reproducible performance matrices for specific AI tasks (e.g., planning, outlier detection, and object reidentification). 4) *Performance benchmarks:* including the comprehensive aspects of verification for a proposed AI test; a checklist of suitable, relevant, and unbiased performance measures; result analytics and evaluation of performance; reflection and mitigation strategies to enhance AI design, testing, and data per the evaluation findings.

AI Governance

The potential negative-to-adverse consequences of AI systems and algorithms raise increasing concern. It demands the healthy development and quality assurance of AI engineering. The research on AI governance develops methodologies, principles, workflows, and KPI metrics to identify, evaluate, monitor, and ensure 1) the performance evaluation of executing an AI task, such as statistical and technical significance, and the business impact (costs and benefits) of implementing AI systems, algorithms, and results; 2) the actionability of a system and results in supporting operations and decision-making; 3) quality issues, including integrity, robustness, trust, bias and fairness of AI tasks, algorithms, and results; 4) potential negative-to-adverse consequences of executing AI tasks; 5) machine ethics and machine rights, including transparency, interpretability, accountability, responsibility, and empathy; 6) the security and privacy of AI systems; and 7) risk and compliance with audits and regulate the abovementioned concerns in intelligent machines.

AI PROFESSION AND EDUCATION

The enduring development and prosperity of AI is determined by the availability, qualification, and dissemination of a qualified AI workforce and professionals and quality education programs.

It is encouraging to see that major economies and organizations have built national and state-level AI plans into their strategic agendas.[8–12] Many more AI centers and courses have been established in major research institutions and multinational vendors over the recent five years. AI is regarded as one of the most strategic and paramount areas for building national competitiveness and future advantage.

AI Profession

The history of AI is much longer than the recently emergent data science. These two areas have become increasingly convergent and integrative, while their education and profession are widely divided. Data science plays a fundamental role in propelling this new epoch of AI. In contrast to the new AI hype, in job markets, AI roles are seldom as popular as those in data science. There is also significant imbalance between their job vacancies. Almost every enterprise holds some roles in data science, from upper management like the chief analytics officer, chief data officer, and chief data scientist, to technical titles, such as data engineer, data scientist, and business analyst. AI professionals have not formed such an evident professional niche and hierarchy. Roles like chief AI officer, AI scientist, and AI engineer are rather rare. Systematic AI education is thus imperative to nurture well-qualified and hierarchical AI professionals.

AI Education

AI education aims to ensure AI technologies and professionals are work ready. In China, many universities have established a School of AI to train AI engineers and specialists. No qualifications are accredited by AI professional bodies. There is a significant imbalance between AI and data science and also between AI research and AI education. While many research groups emerged in the AI journey, only recently have universities and R&D vendors established AI research institutions. In contrast, similar movements for data science have already solidified a critical mass of resources, funding, and projects, since its first research initiative in 2011. In curricula, the recent decade has witnessed the growth of on-campus and online courses on data science and analytics in most of educational institutions at all levels, from undergraduate and master's to doctoral programs. In contrast, similar AI courses have only emerged in the recent 2–3 years, most of which simply repackage data science courses. Most data science and AI courses are simply overloaded with well-received subjects related to data analytics, machine learning, pattern recognition, computer vision, NLP, and R/Python programming. In fact, many lecturers often interchange, mix, or confuse the content of these various subjects in their lectures.

To nurture the AI profession and the AI economy, AI education should cover the entire spectrum of education with hierarchical curricula customized for primary schools, secondary schools, tertiary education at undergraduate, graduate and doctoral levels, and professional training. Training in AI thinking, knowledge and technologies, skills and capabilities, and practice should reside in the national educational and innovation agenda. Doctoral students focus on AI science for cutting-edge research and innovation. Undergraduate and postgraduate students build AI thinking, fundamentals, technologies, and engineering skills to be accredited as AI professionals. Primary students may join elective course modules, such as interactive games and robotics, to cultivate their sense and interest in intelligent machines. Secondary students may be equipped with certain knowledge and skills in applying and creating intelligent technologies, devices, and applications. In a nutshell, AI should be the core of STEM education at all levels. It is increasingly imperative to ensure young generations are AI-ready sooner and faster.

ROLE OF IEEE INTELLIGENT SYSTEMS (ISs)

IEEE Intelligent Systems is the first IEEE publication on AI and has been promoting AI theories, technologies, and applications for more than 35 years. With such a unique prestige and in the golden new age of AI, I am thrilled to collaborate with all authors, readers, editors, and reviewers on reporting the most significant advances in AISE, AI thinking, and expert opinion on the future of AI.

IEEE Intelligent Systems will build on its rich history to introduce strategic initiatives to promote AISE. These include several new columns, namely Editor's Perspective, AI Future, AI Expert, AI Focus, and AI Community. They will highlight expert visions and perspectives from senior principal AI scientists and engineers on the latest AI milestones, spotlights, and prospective AI trends and its future. *IEEE Intelligent Systems* will also enhance its engagement with global and regional AI institutions and professional bodies. From 2022, the prestigious "AI's 10 To Watch" will be substantially enhanced to promote next-generation AI leadership and celebrate some of the most selective world-leading accomplishments made by top young-generation AI stars. *IEEE Intelligent Systems* will also launch its annual "Test of Time AI Award" this year to acknowledge the most influential AI papers published over the past 15 years in *IEEE Intelligent Systems*.

Last but not least, I thank V. S. Subrahmanian for his leadership as the editor-in-chief of *IEEE Intelligent Systems*. His vision and direction have significantly contributed to AI innovations and developments in many important areas over his five years of service, including but not limited to AI and economics, big data, cognitive computing, cybersecurity, automated teaming, chatbots, AI and FinTech, self-driving cars, and various applications of machine learning. As the new editor-in-chief, it is my privilege to warmly welcome your most important achievements and contributions and special issue proposals on important and emergent AI areas to *IEEE Intelligent Systems*.

REFERENCES

1. S. Russell and P. Norvig, *Artificial Intelligence: A Modern Approach*, 4th ed. London, U.K.: Pearson, 2021.
2. L. Cao, "A new age of AI: Features and futures," *IEEE Intell. Syst.*, vol. 37, no. 1, 2022.
3. D. Purves, G. Augustine, D. Fitzpatrick, W. Hall, A. LaMantia, R. Mooney, M. Platt, and L. White, Eds., *Neuroscience*, 6th ed. Oxford, U.K.: Oxford Univ. Press, 2018.
4. M. Minsky, *The Emotion Machine: Commonsense Thinking, Artificial Intelligence, and the Future of the Human Mind*. New York, NY, USA: Simon and Schuster, 2007.
5. X. Qian, J. Yu, and R. Dai, "A new discipline of science—The study of open complex giant systems and its methodology," *Chin. J. Syst. Eng. Electron.*, vol. 4, no. 2, pp. 2–12, 1993.
6. J. Pearl and D. McKenzie, *The Book of Why: The New Science of Cause and Effect*. LaVergne, TN, USA: Ingram Publisher Services, 2017.
7. L. Cao, R. Dai, and M. Zhou, "Metasynthesis: M-space, M-interaction and M-computing for open complex giant systems," *IEEE Trans. Syst., Man, Cybern.–Part A*, vol. 39, no. 5, pp. 1007–1021, Sep. 2009.

8. China's State Council, *A Next Generation Artificial Intelligence Development Plan*, China, Jul. 2017. [Online]. Available: http://www.gov.cn/zhengce/content/2017-07/20/content_5211996.htm or [Online]. Available: https://na-production.s3.amazonaws.com/documents/translation-fulltext-8.1.17.pdf
9. U.S. National Security Commission on Artificial Intelligence, NSCAI, Arlington, VA, USA, Final Rep., Mar. 2021. [Online]. Available: https://www.nscai.gov/2021-final-report/
10. European Union, "White Paper on Artificial Intelligence—A European Approach to Excellence and Trust." Brussels, Belgium COM(2020) 65 Final, Feb. 19, 2020.
11. OECD, "The OECD AI Principles," May 2019. [Online]. Available: https://oecd.ai/en/ai-principles
12. The Global Partnership on Artificial Intelligence (GPAI). Accessed: Dec. 26, 2021. [Online]. Available: https://www.gpai.ai/

LONGBING CAO is a Professor and an ARC Future Fellow (professorial level) with the University of Technology Sydney, Ultimo, NSW, Australia. He is the Editor in Chief of *IEEE Intelligent Systems* magazine. Contact him at longbing.cao@uts.edu.au.