Beyond AutoML: Mindful and Actionable AI and AutoAI With Mind and Action

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

Automated machine learning (AutoML), in particular, neural architecture search (NAS) for deep learning, has ignited the fast-paced development of automating data science (AutoDS) and artificial intelligence. However, in the existing literature and practice, AutoML, AutoDS, and autonomous AI (AutoAI) are highly interchangeable and primarily centered on the automation engineering of data-driven analytics and learning pipelines. This challenges the realization of the full spectrum of AI paradigms and human-like to human-level intelligent and autonomous systems. Going beyond the state-of-the-art paradigm of AutoML and their automation engineering, there is an expectation that the new age of AI and autonomous AI (or AutoAI⁺) will incorporate mind-to-action intelligence and integrate them with autonomy. We pave the way for this new AI and AutoAI integrating **mindful AI and AutoAI** ***with AI mind and mindfulness and*** **actionable AI and AutoAI** ***with AI actions and actionability and translating AI mind to AI action for autonomous, all-around AI systems.***

Automation has been increasingly incorporated into machine learning (ML), driving the emergence of automated machine learning (AutoML).[1,2,3,4] This paradigm is also reflected in the evolution of data science (DS) and artificial intelligence (AI), driving automated data science (AutoDS) and autonomous artificial intelligence (AutoAI) for applications. While ML, DS, and AI share overlaps but also significant conceptual, disciplinary, and technical differences, the conceptions and functionality of AutoML, AutoDS, and AutoAI in the literature and vendor solutions are rather duplicated, interchangeable, confusing, or even misleading. Data-driven analytical and learning tasks, workflows, pipelines, and functions define AutoML, AutoDS, and AutoAI in the literature. On the other hand, autonomy and autonomous systems have a long research history grounded in various disciplines, including cybernetics, control systems, and AI areas such as multiagent systems. This article revisits the conceptual systems and landscapes of AutoML, AutoDS, and AutoAI in terms of the ML, DS, and AI nature and autonomous systems. Going beyond the existing conceptual and technical scope of AutoML, AutoDS, and AutoAI, we propose *AI mind* and *AI mindfulness* for *mindful AI* engaging mind[5] and mindfulness, and *AI actions* and *AI actionability* for *actionable AI* enabling action and actionability.[6,7] We further discuss a new perspective or generation of AutoAI: autonomous AI (AutoAI⁺, or interchangeably AutoAI for its new conception) with mind and action, forming mindful AutoAI and actionable AutoAI and differentiating them from the existing automation-centered conceptions. Our motivation is to align AutoAI with the broad spectrum of AI visions and diverse paradigms of intelligences[7] and enable AI and AutoAI systems with human-like to human-level and all-round mind, thinking, action, and autonomy.

1541-1672
Digital Object Identifier 10.1109/MIS.2022.3207860
Date of current version 24 October 2022.

AUTOML AND AUTODS

Given a dataset and a corresponding learning task, AutoML[1] automates their data processing, feature engineering, ML model (algorithm) selection, process pipelining, hyperparameter selection, or performance evaluation. A configuration space of features, models (algorithms), parameters, hyperparameters, structures or architectures, and pipelines is selected for the best or best-first performance of the learning task and mostly suits the data. AutoML involves domain knowledge, metadata, and metaknowledge. Figure 1 summarizes the landscape of AutoML and AutoDS research and their iterative optimization processes. In the following, we introduce the typical research agenda and

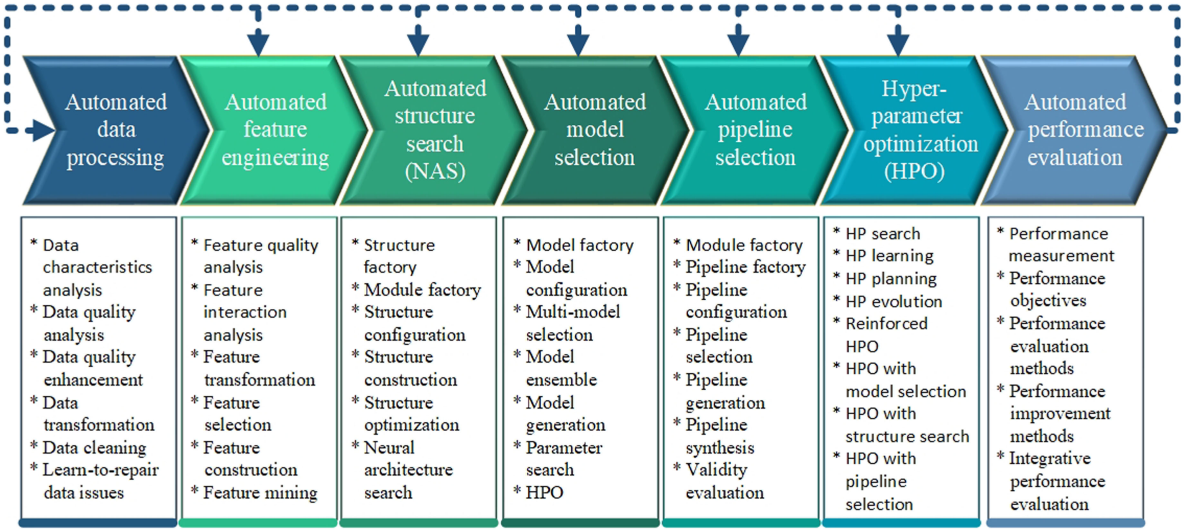

FIGURE 1. AutoML and AutoDS landscape: the automation engineering of analytics and learning.

tasks of AutoML and AutoDS in terms of each automation stage and their evolution.

AutoML and AutoDS Landscape

The landscape of AutoML and AutoDS covers the respective ML and DS stages,[8] including automating data processing, feature engineering, structure search, model selection, pipeline selection, hyperparameter optimization (HPO), and performance evaluation.

Automated data processing: The existing research focuses on automatic data cleaning rather than the full spectrum of automated data processing including *automated data quality processing*. *Automatic data cleaning* removes data errors, such as missing, redundant, inconsistent, irregular, incorrect, invalid, or outlying values and their corresponding objects. Typical methods rely on human–machine interaction and cooperation,[9] where domain experts, data modelers, or data scientists predefine data error detection and repairing functions, rules, and operators. They are often then hard-coded into AutoML and AutoDS systems for automatic error detection, scaling, normalization, deletion, or imputation. Exploratory data analysis, metafeatures, metadata, metaknowledge, and domain knowledge are often involved in guiding the design and automation of error detection and mitigation. More advanced data cleaning detection and mitigation apply data representation, embedding, and encoding for data characteristics analysis, data quality analysis, and data quality enhancement. However, the full-spectrum data quality analysis, data processing, learning-to-fixing data quality issues, and data-agnostic and domain-agnostic general data processing automation are rarely involved in the existing research.

Automated feature engineering (AFE): AFE is another major area of AutoML and AutoDS. AFE consists of tasks for automated feature extraction, transformation, selection, mining, construction, or generation. The main focus of AFE has been on automated feature transformation and selection. *Automated feature transformation* creates techniques, such as

1) detecting and filtering noisy, irrelevant, redundant, erroneous, or misleading features using feature quality analysis or feature interaction analysis;
2) transforming features to usable formats using transformation methods, such as mapping, compression, scaling, and normalization; and
3) representing features in new formats using representation learning automatically or in a human–machine cooperative manner.

Typical methods for *automated feature selection* include selecting features based on measuring feature importance, interactions, relevance, correlations, mutual information, redundancy, and stability. Predefined functions, rules, and operators are designed by data engineers for feature transformation and selection. Wrapper or embedded methods then select features by extracting subspaces or during learning processes. *Automated feature construction* (or automated feature generation) applies techniques, including mathematical transformation, such as discretization, randomization, masking, or derivation; arithmetic operations, such as feature interaction analysis, feature aggregation and mapping, and feature representation and embedding learning to create, derive, infer, and represent new features from an existing feature set. *Automated feature mining* applies data mining techniques on a feature set to discover feature subsets (subspaces) satisfying mining objectives. Limited research has been conducted on automated or human–machine-cooperative feature transformation, construction, and discovery in complex data, feature interactions, and learning contexts.

Automated structure search and configuration: Suitable structures are searched from a structured factory or configured from an ML/DS module factory, where different candidate ML/DS building blocks are configurable to fulfill an ML/DS structure, architecture, process, or pipeline. Various ML/DS modules with general and specific functions are available for different stages of the learning, including data processing, learning objectives, tasks, models, and evaluation. Module configuration and a pipeline structure search are undertaken by techniques, which involve, for example, domain knowledge, hierarchical planning, structure construction (e.g., network, graph, tree, sequential, or parallel structures[10]), and structure optimization (e.g., using a Monte Carlo tree search and genetic programming).

In deep AutoML, *neural architecture search* (NAS)[11] retrieves, configures, and generates the architectures of deep neural networks (DNNs) from a factory of configurable DNN building blocks. The *DNN factory* produces DNN building blocks, such as neural nodes (e.g., input neurons, activation function, gate function, and output neurons), functional layers (e.g., convolutional layer, attention layer, and classification layer), neural operators (e.g., convolution, dropout, pooling, encoding, and decoding), microlevel network connectors (recurrency, bidirection, chain, stacking, and multibranching), macrolevel network architectural frameworks (e.g., convolutional, recurrent, adversarial, autoencoding, and attentive frameworks), and performance evaluators or learning controllers (e.g., learning rate, early stopping, and batch normalization). The building blocks are then configured into DNNs following certain architectural frameworks and configuration specifications. Alternatively, well-performing architectures are searched from a space of configurations with candidate structures, mechanisms, and strategies, forming a DNN with various possible parameter combinations, thus generating different network architectures. NAS applies search strategies to retrieve configurations of DNN building blocks to form well-performing network architectures in the search space evaluated by performance estimation strategies.[11] Typical NAS search strategies include converting an NSA problem to a random search, evolutionary search, Bayesian optimization, Monte Carlo tree search, gradient-based search, or reinforced search. The search performance is then estimated using performance estimation or evaluation strategies, such as the learning performance, computational costs, and search complexity of the identified or configured neural architectures. Automated structure search and NAS face the challenges of speeding up search algorithms, optimizing search strategies, and handling complex search settings, such as unsupervised learning systems, complex learning tasks, multitask or multiobjective learning, and large-scale learning systems.

Automated model selection or *automated algorithm selection:* Typically, they select or configure the best-performing or first-best model (or algorithm) from

1) set of different models;
2) configuration space of models with similar modeling principles, schematic processes, backbones, or frameworks;
3) model candidates with different parameters; or
4) models with different hyperparameter settings.

The first scenario involves multimodel selection or ensembles[12] in terms of learning objectives and tasks. For ensemble models, base learners are combined per techniques, such as bagging, boosting, stacking, and cascade generation. The second scenario generates and selects the best-performing models from a set of configurable model candidates. The third one conducts a parameter search or optimization within the domain of the parameters. The last one selects the best-performing model by automating its HPO. Accordingly, techniques for model configuration, comparison, integration, ensemble, and HPO are incorporated into ML/DS systems for an automated model search or generation. These also involve the techniques for model evaluation.

Automated pipeline selection: An ML/DS pipeline is composed of building blocks to fulfill the different functions required at different stages of ML/DS. First, the ML/DS factory produces various building blocks for stages, such as data processing, feature selection, model selection, HPO, and performance evaluation. Second, automated pipeline selection configures, generates, or integrates configurable pipeline candidates. This involves pipeline configuration specifications, learning tasks, domain knowledge, pipeline prototypes, structure configurations, etc. Finally, automated strategies are designed to search, compare, evaluate and filter pipeline candidates, and best-performing pipelines. Accordingly, automated pipeline selection identifies suitable structures, modules, and workflows for ML/DS. The resultant workflow or pipeline synthesis searches for a space of optimal or acceptable configurations consisting of a set of candidate models (algorithms). Each model also searches for its most suitable parameters and hyperparameters for a possible best match with the given data. Pipeline performance evaluation verifies the validity, learning performance, and computational efficiency of the constructed pipelines.

Automated HPO: HPO was the first and most intensively studied area in AutoML and other related areas, such as mathematical optimization, Bayesian optimization, planning, evolutionary computing, and reinforcement learning.[13] Techniques and methods for AutoML HPO are mainly inspired by these areas or apply their developments. Typical HPO methods include grid search, random search, tree search, nonlinear optimization with gradient descent, Bayesian optimization, Monte Carlo hierarchical planning and tree search, genetic algorithm, particle swarm optimization, fuzzy set, and multiarmed bandit optimization. Early work on HPO relies on domain knowledge and user-defined hyperparameters. More recent studies combine HPO with other AutoML tasks, such as structure search and pipeline configuration. More research is expected on the data-driven parameter search in an end-to-end manner during the pipeline selection and optimization and fine-tuning the hyperparameters of pretrained DNNs.

Automated performance evaluation: This involves:

1) performance measurement corresponding to AutoML objectives and tasks, such as learning accuracy, computational efficiency, and scalability;
2) performance targets to be optimized, such as the accuracy of automated data quality repairing; and
3) improvement methods for optimizing performance, such as speeding up optimization using the one-shot supermodel in NAS.

These functions are incorporated into ML/DS processes and systems.

Evolving AutoML and AutoDS Research

With the paradigm migration from specific ML to broad-reaching DS, the area of AutoML has evolved significantly in the recent two to three decades. This evolution from preliminary AutoML to advanced AutoDS reflects the following evolving perspectives and paradigm shifts.

- *From hyperparameter to model:* transforming HPO to automated model selection.
- *From hyperparameter to system, workflow, process, and pipeline:* transforming hyperparameter selection to the configuration of complete ML/DS systems and pipelines consisting of data processing, feature engineering, model selection, and their hyperparameter tuning to form ML/DS scenarios, pipelines, and ecosystems.
- *From module to structure and architecture:* transforming modules, such as parameters of models to model/algorithm structures and architectures, describing the combination of parameters, synopsis of modules, pipeline structures, and modeling logic and processes;
- *From simple to complex settings:* transforming simple settings, such as HPO and static automation to dynamic, adaptive, online, active, or open (e.g., open-domain) automation.
- *From specific to general settings:* transforming the task-, objective-, data-, domain- or model-specific settings to task-, objective-, data-, domain- or model-agnostic and general optimization and automation.
- *From individual to hybrid objectives and methods:* transforming individual tasks, objectives, objects, and methods to multiple, hybrid, ensemble, hierarchical, or integrative tasks, objectives, objects, and methods for optimization and automation, e.g., combining algorithm selection with its HPO.
- *From search and selection to configuration, generation, and optimization:* transforming search and selection-oriented automation to higher level, more complex, and advanced automation objectives, operations, or settings through the configuration and generation of optimal ML/DS parameters, structures, models/algorithms, pipelines, and ecosystems.
- *From search to learning:* transforming search-focused methods and strategies, such as grid and random search, Monte Carlo tree search or guided search of the best parameters or parameter combinations to the hierarchical planning, automated learning, and reinforced learning of parameters, parameter combinations, process pipelining, and system configuration.
- *From small scale to large scale:* transforming a small-scale, low-dimensional hand-crafted search and selection of parameter combinations and a configuration space of structures and pipelines to a large-scale and high-dimensional automated search, selection, learning, or optimization.
- *From shallow to deep automation:* transforming handcrafted, predefined, and *ad hoc* automation

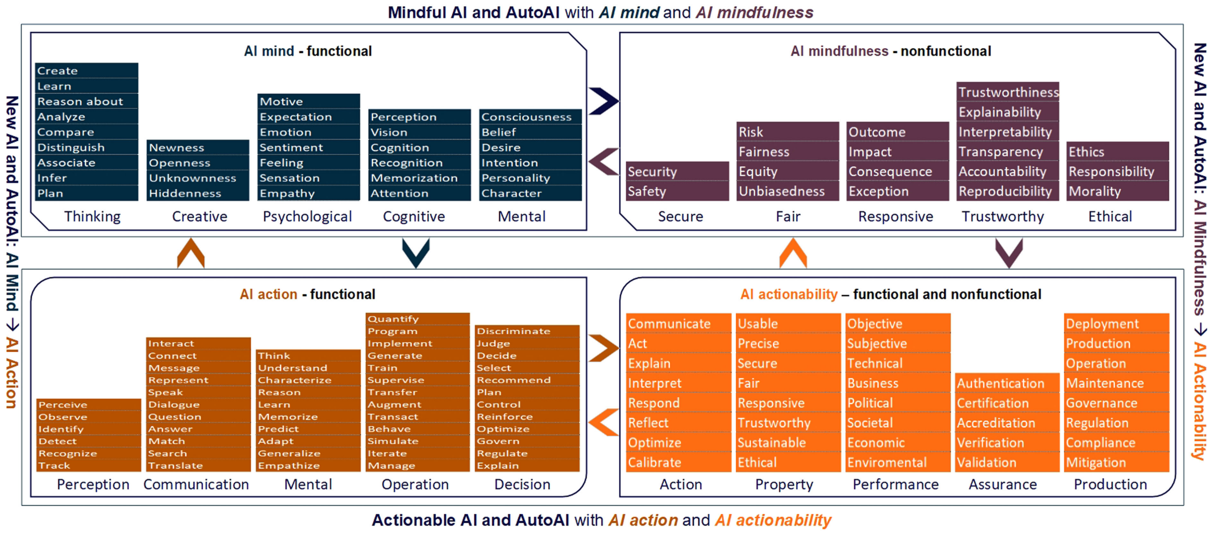

FIGURE 2. Landscape of mind-to-action AI and AutoAI: synergizing mindful AI with AI mind and AI mindfulness, and actionable AI with AI action and AI actionability.

to algorithm-driven, end-to-end, or ecosystem-oriented automation.

MISCONCEPTIONS OF AUTOML AND AUTOAI

There are many references available online and in the body of literature addressing the concepts and issues in relation to AutoML and AutoAI. However, there are overwhelming misconceptions, myths, pitfalls, and fallacies relating to the concepts of autonomous techniques and systems, AutoML, or AutoAI. Here, we briefly discuss some.

Are AutoML, AutoDS, and AutoAI Interchangeable?

In the literature, it is common to see such arguments that

- AutoAI is equivalent to AutoML;
- AutoAI is also called AutoML;
- AutoAI is interchangeable with AutoML;
- AutoAI is a special area of AutoML;
- AutoAI is a variation of AutoML; or
- AutoAI is an extension of AutoML.

The abovementioned arguments can be more commonly seen in discussing the relations between AutoDS and AutoML in the literature and commercial solution manuals.[2] Typically, AutoDS and AutoML are interchangeably used in the literature and vendor solutions.

In addition, it is also common to see many vendors treat AutoML, AutoDS, and AutoAI interchangeably in introducing their commercial solutions.[2] The mixture of AutoML, AutoDS, and AutoAI is also widely reflected in open source and commercial tools, in particular, end-to-end, data-driven, metalevel, and analytical/learning algorithm-centric platforms and toolsets. Such systems focus on automating the workflow from data cleaning, feature engineering, model selection, and HPO to an end-to-end, metalevel, black box learning manner. The learning box is "black" without human involvement and intervention, and metalevel learning and optimization replace the usual expert choice of models and hyperparameters for the learning box.

These misconceptions of AutoML, AutoDS, and AutoAI highlight their commonality but overlook or muddle their distinct disciplinary objectives, scopes, concepts, tasks, processes, or agendas. Arguably and technically, AutoML is a fundamental subset of both AutoDS and AutoAI,[a] while AutoDS and AutoAI intersect but complement each other, as shown in Figure 2 and further discussed in the following section on AutoAI.

- AutoML is centered on automating analytics and learning toward autonomously discovering data

[a]Cao[7] summarizes the 70 years of AI evolution with different AI paradigms and shows the broad AI disciplinary spectrum; while Cao's work[8] shows the cross-disciplinary nature of DS. In both AI and DS, ML forms a core area and technique for their broad disciplinary systems.

intelligence by exploring analytics, learning, and their algorithmic intelligence;

- AutoDS fuses the broad disciplinary power and agendas of statistics, computing, and informatics with a much broader spectrum and depth of data-driven scientific vision, mission, capacity, and capabilities, where "autonomous data-driven discovery" is only a core mission, going beyond automating analytics and learning;
- AutoAI carries forward its broad-reaching vision and mission of simulating and mimicking human-level, natural, and social intelligence, where data, analytics/learning, and algorithm intelligence form a small (core) part of the broad AI spectrum.

Unfortunately, there is overwhelming confusion and several misconceptions relating to ML, DS, and AI in the literature.[b] It is also very common to see similar confusion and misconceptions in widespread undergraduate and postgraduate syllabus and curricula globally.[15] In such courses, ML subjects are often interchangeably labeled DS and AI courses or dominate the disciplinary body of knowledge of DS and AI courses. It is also not uncommon to hear that DS and AI are outdated and subjects on either classic ML or deep learning become the only or major disciplinary focus in training undergraduate or postgraduate DS and AI specialists.

Myths and Pitfalls of AutoML, AutoDS, and AutoAI

Over the development journey of AutoML, AutoDS, and AutoAI, there have been different understandings, debates, and widespread misconceptions about their definitions and research scope and methods, etc. in the literature and in commercialization. The misconceptions, myths, and pitfalls involve the goals, methodologies, techniques, processes, and tools for AutoML, AutoDS, AutoAI, and autonomous systems both in general and specifically.

The first category of misconceptions, myths, and pitfalls is about the core concept "autonomy" and its implementations—the automation of AutoML and AutoAI systems. They involve the argument about whether the goal of AutoML, AutoDS, and AutoAI are to allow technical layman and domain experts to use the systems "without understanding anything"[1]; whether the autonomy of AutoML, AutoDS, and AutoAI should be sole, discrete, or full; and whether more or higher autonomy would be better.

- *Discrete autonomy:* AutoML and AutoAI systems can be implemented as dashboards or widgets without the need for human involvement or interactions with the environments.
- *Sole autonomy:* AutoML and AutoAI systems are unidimensional without other dimensions and an AI capacity and capability set, hence they are self-directed and self-sufficient.
- *Full autonomy:* AutoML and AutoAI systems have unbounded and full autonomy without the involvement, attention, and governance of humans and other agents.
- *More autonomy:* More autonomy would make AutoML and AutoAI systems better, smarter, and more capable, etc.
- *Higher autonomy:* A higher level of autonomy would make AutoML and AutoAI systems better, smarter, more capable, etc.

Accordingly, AutoML and AutoAI systems could be overautomated, fully automated, partially automated, semiautomated, or underautomated. However, it would be naive to believe AutoML and AutoAI systems should take on automation independent of or without teamwork, constraints, interactions, coordination, and collaboration. This is related to teamwork and human involvement in automation, inspiring the following category.

The second category is about the relationships between machines and humans in the automation of AutoML and AutoAI. This involves questions, such as: whether automation should be machine-centric or learning algorithm-centric; whether automation should be supported in teamwork; and whether humans should be part of autonomous systems.

- *Hyperparameter-centric automation:* The automation of AutoML and AutoAI systems focuses on HPO, which searches for a hyperparameter configuration that minimizes the objective loss or combines the search for algorithms and their hyperparameters.
- *Machine-centric automation:* The automation of AutoML and AutoAI systems is machine-only, machine-centric, unmanned without human involvement, and human-out-of-the-loop.
- *Algorithm-centric automation:* The automation of AutoML and AutoAI systems is solely dependent on algorithms to implement automation without human involvement.

[b]For example, Cao[14] discusses various pitfalls and misconceptions of DS.

- *Teamwork automation:* The automation of AutoML and AutoAI systems is implemented through team interactions and collaborations and individual autonomy during teamwork.
- *Human–machine teaming automation:* The automation of AutoML and AutoAI systems is implemented through human–machine teaming, human-in-the-loop, and automation with human involvement.

The level of human involvement and the degree of human–machine teaming determine whether an AutoML system or an AutoAI system is human-out-of-the-loop, human-in-the-loop, or human-centric where automated machines are fully operated by humans.

In either machine-centric or human–machine-teaming AutoML and AutoAI, there are different aspects of autonomy and automation properties that determine whether the autonomy and automation in AutoML and AutoAI are constrained, bounded, or governed; and whether AutoML and AutoAI are transparent, explainable, moral, controllable, accountable, or trustworthy.

- *Bounded automation:* The automation of AutoML and AutoAI systems is bounded or constrained with boundaries, limitations, or constraints.[16]
- *Governed automation:* The methodologies, design, implementation, operation, or exception handling, etc. of AutoML and AutoAI systems are governed and regulated by humans and authorities following governance specifications.
- *Transparent automation:* The automation of AutoML and AutoAI systems, with their designs, operations, decision making, consequences, etc. is transparent to end users and authorities.
- *Explainable automation:* The automation of AutoML and AutoAI systems is explainable,[17] with their designs, operations, decision-making, outcomes explained, understandable, and accountable.
- *Controllable automation:* The automation of AutoML and AutoAI systems is controllable, with the controllability of their operations, decision-making, and consequences, etc. directed, influenced, regulated, or managed by humans, machines, or human–machine cooperation.
- *Accountable automation:* The automation of AutoML and AutoAI systems is accountable, with accountability justifying their designs and operations, and being responsible for their outcomes, etc.
- *Trustworthy automation:* The automation of AutoML and AutoAI systems is truthful and dependable, where the end-user has trust or confidence in their designs, operations, decision-making, consequences, etc.

Accordingly, different AutoML and AutoAI systems may have different levels or degrees of boundness, governance, transparency, explainability, controllability, accountability, and trust in implementing automation and achieving autonomy. This layered thinking turns AutoML and AutoAI systems into black, gray, or white ML/DS/AI boxes for stakeholders with different backgrounds, expertise, purposes, and implementation modes.

- *Black-box automation:* AutoML and AutoAI systems are designed as black learning boxes for ML/AI laymen, with settings, such as minimizing user interactions with ML/AI system inputs and outputs, generating user-friendly outputs, having zero access to internal ML/AI rationale and end-to-end design, implementation, and operation, and making no response to exceptions.
- *Gray-box automation:* AutoML and AutoAI systems are designed as gray boxes for ML/AI-savvy decision-makers, with settings, such as identifiable and explainable input–output mapping, essential access to internal systems and information, and routine response to exceptions.
- *White-box automation:* AutoML and AutoAI systems are designed as white boxes for subject experts, with settings, such as full access to the system rationale, design, implementation, operation, and the management of risk and consequences.

It is noted that the abovementioned misconceptions, myths, and pitfalls may be further reflected in the overwhelming industry and economy digitalization and smartness movement, including in driverless cars, drones, military robots, smart cities, smart homes, digital health, and digital finance. The misconceptions may not only distort the design and operation of these autonomous techniques and systems but also incur new problems in particular vulnerability, safety, security, machine ethics, algorithm ethics, data ethics, and machine morality. They thus must receive sufficient attention, and be clarified and discussed in their contexts, and more strategic considerations of the above issues should be taken in AutoML and AutoAI design, execution, and their management philosophies, methodologies, and techniques.

WHAT MAKES AUTOAI?

The abovementioned misconceptions, myths, and pitfalls do not simply surround the debate on strong versus weak AI (or alternatively, broad versus narrow AI, hard versus soft AI), i.e., whether an AI machine should encompass the full range of human intelligence, including self-awareness, consciousness, imaginary thinking, and afflatus (inspiration). Migrating from automated AI to autonomous AI is essential for human-like to human-level AI systems. Autonomous AI should be built on the broad-reaching human, natural, social to machine intelligence paradigms and cross-disciplinary, all-round AI foundations.[7]

AutoAI$^+$: Autonomous AI Systems

In general, *automated systems* are operated by predefined rules and configurations under predictable environments, aiming to automate repeatable tasks and operations. The existing AutoML and AutoAI systems follow this path. Although it has not been fully developed and there are various misconceptions as discussed previously, the existing AutoAI systems are deemed automated with predefined AI tasks, designs, rules, and algorithms for intelligent computing and informatics.

In contrast, *autonomous systems* have been increasingly studied in the recent three decades, aiming for autonomous decision-making, behavior change, and adaptation to unanticipated events and environments during their operation. This suggests a different perspective and a more intrinsic conception of AutoAI: *autonomous AI*. We categorize such autonomous AI systems as *AutoAI*$^+$ for the convenience of distinguishing them from the misconceptions of "automated AI" and the narrowly defined AutoAI in the literature. For simplicity and to refresh the conception of AutoAI for its alignment with the broad-reaching AI vision and mission, we interchangeably use AutoAI and *AutoAI*$^+$ in this article except when special clarification is required.

AutoAI$^+$ has to be built on the broad-reaching AI and AutoAI landscape. It expands *AI engineering automation* to monitoring, optimizing, and adjusting AI system structures, behaviors, and functions to

- fulfill the broad-reaching AI aims and tasks (beyond learning and data-driven tasks);
- realize various paradigms of intelligence (beyond data, algorithmic, and learning intelligence);
- incorporate cognitive skills into AI algorithms and systems (beyond predefined symbols, logic, and rules);
- engage a wide range of AI techniques (beyond analytics and learning);
- operate on various AI objects (beyond data, models, and algorithms);
- learn to handle unanticipated and uncertain settings (beyond unambiguous conditions);
- learn to behave in uncertain environments (beyond static and stationary contexts);
- generate alternatives responding to uncertainty (untested, unforeseen, unpredictable, and unmanaged tasks, inputs, environments, exceptions, or consequences); and
- learn to optimize operations, decisions, and actions (beyond selection and search in given domains).

The resultant AutoAI systems will go beyond automated AI software engineering and be autonomous and actionable.

AutoAI$^+$ Landscape

Here, we discuss the research landscape of AutoAI$^+$, which is illustrated in Figure 2. This AutoAI landscape is motivated by and consistent with the broad-reaching, evolving AI field over its 70 years of history.[c]

Consistent with the broad AI landscape,[7] AutoAI is empowered by broad AI research aims, tasks, representative paradigms of intelligence, and their areas and objects.

- *AutoAI aims and tasks:* AutoAI systems may fulfill general or specific aims and tasks of AI. These may include AutoAI systems and functions to understand, simulate, perceive, reason about, match, plan, learn, discover, search, recognize, transfer, act, reconstruct, reinforce, communicate, optimize, decide, reflect, explain, think, and govern the underlying artificial, human-involved, natural, or social systems (represented as intelligent agents) and their environments.
- *AutoAI intelligence paradigms:* AutoAI systems may be implemented by specific intelligence paradigms, simulate some paradigms of intelligence, or aim to achieve such intelligence. Typical intelligence paradigms include object intelligence, symbolic intelligence, connectionist intelligence, learning intelligence, behavioral intelligence, natural intelligence, networking intelligence, data intelligence, social intelligence, algorithmic intelligence, emotional intelligence, ethical intelligence, human intelligence, and system intelligence.

[c]As shown by the broad AI landscape in Figure 1 in Cao's[7] work.

- *AutoAI areas and techniques:* This may cover any areas of classic and modern AI research. They include but are not limited to symbolic reasoning, expert systems, knowledge engineering, probabilistic reasoning, pattern recognition, ML, knowledge discovery, computer vision, natural language processing, robotics, multiagent systems, behavior informatics, evolutionary computing, DS, deep learning, and their applications, such as outlier detection, recommender systems, and ethical AI.
- *AutoAI objects:* AutoAI systems may study various objects and target a wide spectrum of work products in the AI system space. Examples are mind, hypothesis, logic, observation, knowledge, uncertainty, connectivity, behavior, vision, sensing, conversation, experiment, memorization, emotion, data, program, algorithm, architecture, experience, feedback, attention, intention, governance, and ethics.

Accordingly, there are many aspects of AutoAI worthy of noting: task phases, functionalities, stakeholders, roles, properties, and application scenarios. In the following, we briefly discuss these.

First, AutoAI involves different phases of AI tasking to automate AI aims, tasks, and intelligence paradigms. Typical stages include problem definition, and hypothesis making; data acquisition, pipelining, processing, and governance; model selection, search, and evaluation; system configuration, and structure or architecture search; hyperparameter and parameter search, selection, and optimization; system testing, and performance evaluation; system monitoring, deployment, and governance; system updating and optimization; feedback, reflection, and improvement; and system output generation, presentation, and communication.

Second, AutoAI may be expected to hold different functionalities to implement or achieve various paradigms of intelligence. Typical AutoAI systems may perform functions including enabling cognitive skills, commonsense modeling, context-aware interpretation, interaction within teams and with the environment, human-in-the-loop (if required), self, active and online learning, adaptation, socio-technical skills, reflection, introspection, metacognition, understanding data quality issues, action with alerts, decision policies, exception mitigation with strategies, trust, and explainability.

Third, different or specific stakeholders and roles may exist for an AutoAI system. They could be technical layman, junior experts, newcomers, subject experts, business, industrial or management personnel, or policymakers. AutoAI systems are then required to customize the system's usability for their specific stakeholders and configure their role-friendly functions.

Finally, different scenarios of AutoAI applications may require the automation of their respective target objects, the customization of the corresponding objectives, tasks, and functions, and the delivery of application-specific AutoAI designs and systems. What is to be optimized may go beyond models, hyperparameters, results, and performance. In real-world settings, AutoAI may be associated with the dynamic, online, adaptive automation of generating, searching, evaluating, optimizing, monitoring, and adjusting their system structures and architectures, and tasking processes and pipelines in complex contexts. Complex application contexts may be associated with very large, high-dimensional, and poor-quality data, evolving problems and boundaries, complicated interactions, changing AI objectives and tasks, and uncertain and open[18] contexts and environments.

PUTTING MIND AND ACTION INTO AUTOAI

The implementation of the all-round AutoAI$^+$ landscape requires the engagement of mind and action into AI and autonomous AI systems. Inspired by the human brain and mind-like and human-level functionalities and thinking, we propose the concepts of the "AI mind" and "mindful AI" with AI mindfulness. Inspired by real-world system engineering and decision-making, we introduce "AI action" and "actionable AI" with AI actionability. This *mind-to-action AI* synthesizes all-round AI aims, tasks, intelligence, techniques, and objects. It aims to enable broad-reaching AI objectives and tasks, going beyond the highly programmed automatic algorithms or prespecified pipelines.

Mindful AI and AutoAI With Mind

AI and AutoAI systems should be mindful and thoughtful, forming *mindful AI and AutoAI* (or *thoughtful AI and AutoAI*). Mindful AI goes beyond classic symbolic AI such as rules and rule-based systems, reactive AI, and predictive AI to amalgamate mind and mindfulness with AI machines. *Mindful AI and AutoAI* aim to empower AI machines with *AI mind* and *AI mindfulness* for human-like intelligence and mind-like autonomous systems. Accordingly, mindful AutoAI requires various functional AI actions and nonfunctional AI mindfulnesses in autonomous AI systems.

Functionally, mindful AI aims to simulate, be inspired, incorporate, or generate a human-like *AI mind*, fulfilling thinking mechanisms,[19,20] cognitive and psychological traits and behaviors, such as human sensations, thoughts, reasoning, representation, communication, imagination, self-knowledge, presentation, and critical analysis. *AI mind* may also process abstract and conceptual symbols, such as mental imagery like commonsense and feelings, imaginary symbols in music and artwork, and visual patterns in design. These functional mind-like capabilities and properties of mindful AI can be categorized into:

- *thinking functions and behaviors* to create, understand, learn, reason about, analyze, compare, distinguish, associate, infer, derive, and plan knowledge and intelligence;
- *creativeness and openness* to new concepts, open systems and domains, and unknown futures;
- *psychological and emotional merits, traits, and activities* to understand, process and achieve motives, expectation, emotion, sentiment, feelings, sensation, empathy, and mindfulness;
- *cognitive abilities* of perception, vision, cognition, recognition, memorization, and attention;
- *mental functions and activities* to generate, understand and manage consciousness, belief, desire, intention, personality, and character.

The functional capabilities and properties of mindful AI and AutoAI support the engagement and implementation of mental, psychological, and emotional traits into AI tasks, processes, behaviors, and systems for autonomous human-like AI systems. However, this is nontrivial and challenging. One possibility is to build human-AutoAI symbiosis for mindful AutoAI, i.e., human-in-the-loop autonomous AI systems, supporting mind and machine teamwork. Furthermore, enabling open-form AutoAI is challenging, as human interactions and guidance by humans-in-the-loop are essential when self-, auto-, and meta-learning cannot be achieved by AI systems and without human knowledge and expert assistance. New concepts, functions, and behaviors may then be introduced or inferred to mindful AI by developing novel means such as the constructive induction, propositionalization, and reformulation of newness, openness, and unknownness in emerging tasks, classes, knowledge, data characteristics, exceptions, or impacts.

On the nonfunctional side, mindful AI, and AutoAI check, identify and ensure the *mindfulness* of AI and AutoAI systems, judgment, choice-making,[19] behaviors, and consequences. Mindfulness enables AI and AutoAI systems to be heedful and vigilant of the operations of AI systems, tasks, models (algorithms), processes, and behaviors, and their occurring and potential problems, consequences, and impacts. Specifically, such nonfunctional mindfulness will enable AI and AutoAI with capabilities and properties being:

- mindful of AI security and safety;
- mindful of AI risk, fairness, equity, and unbiasedness;
- mindful of AI consequences, and impact;
- mindful of AI trustworthiness, explainability, interpretability, transparency, and accountability;
- mindful of AI machines, data and model ethics, responsibility, and morality.

These properties of AI mindfulnesses should be built into AI methodologies, designs, systems, processes, behaviors, performance, evaluation, and management.

The aforementioned functional and nonfunctional capabilities and properties of mindful AI and AutoAI are interrelated and complementary. They ensure both the mind and mindfulness of AI and autonomous AI systems.

Actionable AI and AutoAI With Action

Mindful AI and AutoAI embed mind and mindfulness into AI systems, which is critical for generating human-like AI and AutoAI and enabling AI with mind-like mechanisms for thinking, cognition, mental and psychological activities, and thoughtful and mindful functions and behaviors. However, AI and AutoAI with only mind and mindfulness are incomplete and may be nonactionable, inspiring actionable AI and AutoAI.

Actionable AI and AutoAI aim to incorporate *AI actions* and *AI actionability* into AI and AutoAI systems and make AI systems actionable in production. It also bridges the gap between AI design and AI practice and between AI designers (researchers and scientists) and AI practitioners (e.g., domain experts). Accordingly, actionable AI and AutoAI implement various functional AI actions and ensure the functional and nonfunctional AI actionability of AI and autonomous AI systems.

On the functional side, AI and AutoAI aim to involve and support various AI actions that convert the AI mind and thinking to AI reality through AI machines and operations. Such *AI actions* are comprehensive and systematic, such as to perceive, reason, plan, learn, recognize, respond, reinforce, communicate, optimize, decide, reflect, explain, and govern. We further categorize the functional AI actions for actionable AI into:

- *perception activities* such as perceive, observe, identify, detect, recognize, and track;
- *communication activities* such as interact, connect, message, represent, speak, dialogue,

question, answer, match, search, retrieve, and translate;
- *mental activities* such as think, understand, characterize, reason (induce, deduce), infer, learn, memorize, predict (forecast), adapt, generalize, and empathize;
- *operation activities* such as quantify, program, implement, generate, construct, train, supervise, transfer, augment, transact, behave, simulate, imitate, iterate, and manage;
- *decision actions* such as discriminate, judge, decide, select, recommend, plan, control, reinforce, optimize, govern, regulate, and explain.

In both theory and practice, the aforementioned AI actions may be chosen, categorized, incorporated or integrated into AI modules and machines on demand. Not all of them are intelligent and intellectual, which means they have to be incorporated into an intelligent translation framework to transform, create, optimize or upgrade their process, knowledge, or intelligence. Any AI and AutoAI systems have to implement some of the actions following their corresponding intelligent translation frameworks and processes, which differentiate such intelligent systems from unintelligent ones.

However, AI actions do not essentially result in the actionability of AI systems. In fact, many AI systems are nonactionable not because of their lack of appropriately implemented AI actions but for other reasons. Correspondingly, *AI actionability* consists of functional actions and functional and nonfunctional properties, performance (evaluation), assurance, and production to make AI systems operable and deployable and their decision and production actionable. These build on, integrate and implement the AI mind and mindfulness but also expand them to take decisive actions and implement these into operation, and production. This also requires AutoAI^{+} to not only be autonomous but also adaptive, active, self-to-lifelong learning-enabled, responsible, pathetical, ethical, and reliable.

First, functional AI actionability supports and implements AI-driven decisions and decision actions. Such actions include:

- *communicate* AI insights and findings to stakeholders in a friendly manner;
- *act* on the recommendations, selection, and control made by AI;
- *explain* AI rationale, designs, and outcomes;
- *interpret* the relations between causes and effects of AI results;
- *respond* to the issues, exceptions, consequences, and impacts of AI deliverables;
- *reflect* on the performance, outcomes, and issues of AI systems;
- *optimize* the structures, parameters, pipelines, behaviors, and responses of AI systems;
- *improve* AI designs, processes, behaviors, and decisions with issues resolved and solutions optimized;
- *calibrate* the quality, performance, error, and impact against standards, benchmarks, or ground truth of AI deliverables.

Second, actionable AI and an autonomous AI system are expected to have various properties which enable actionability. They make AI systems:

- *rational* in relation to respecting truths and facts and obedient to rules and orders;
- *usable* for nontechnical and technical stakeholders and users;
- *precise* for accurate computation, inference, and results;
- *secure* and *safe* for deployment, operation, and application;
- *fair* and *unbiased* for outcomes, decisions, and recommendations;
- *responsive* to consequences, impact, risk, and exceptions;
- *trustworthy* in relation to operations, results, consequences, and decisions;
- *sustainable* for operations and reproduction;
- *ethical* and *moral* for deployment, operation, decisive action-taking, and impact.

Third, the performance of actionable AI and autonomous AI systems may be evaluated using multifaceted measures, such as

- *objective* measures of the objectivity, impact, and benefit of AI systems, such as truth, pure nature, reality, unbiasedness, and impersonal beliefs;
- *subjective* measures of the subjectivity, impact, and benefit of AI systems, such as personal judgment, opinion, feelings, and beliefs;
- *technical* measures related to specific techniques, subjects, or sciences for developing AI systems;
- *business* measures about the significance, impact, influence, and benefit, etc. of AI systems on a sector or from a commercial perspective;
- *political* measures relating to the significance, impact, influence, and benefit, etc. of AI systems from organization, government, policy, and public affair perspectives;

- *societal* measures relating to the impact and benefit etc. of AI systems related to human society or community;
- *economic* measures relating to the significance, impact, influence, and benefit, etc. of AI systems from trade, industrial, financial, or monetary perspectives;
- *environmental* measures relating to the significance, impact, influence, and benefit, etc. of AI systems from natural, ecological, and living perspectives.

Fourth, the qualification and assurance of actionable AI and autonomous AI systems are ensured by the corresponding mechanisms and activities, such as

- *authentication* for the genuineness, authority, and truth of AI modeling, processes, decisions, and quality;
- *certification* for the guarantee, attestation, and testimony of AI modeling, processes, decisions, and quality;
- *accreditation* for the official authority, approval or resulting statuses of AI modeling, processes, decisions, and quality;
- *verification* for the proof, authenticity, or validity of AI design, processing, results, and decisions;
- *validation* for the truths, facts, rules, or orders of AI design, processing, results, and decisions.

Finally, the actionability of actionable AI and autonomous AI systems is embodied through their production processes, including:

- *deployment* of moving AI systems into the real world or reality;
- *production* of converting AI systems to products or into manufacturing;
- *operation* of functioning AI systems in working conditions or places;
- *maintenance* of caring, upkeeping, or supporting AI systems after deployment and in production;
- *governance* of sustaining, enforcing, managing, or controlling the specifications, rules, standards, and norms of AI systems;
- *regulation* of the restrictions, obligations, or authority of AI systems;
- *compliance* of the requirements, rules, regulations, and standards of AI systems;
- *mitigation* of issues, exceptions, negative consequences, impacts, and outcomes of AI systems.

CHALLENGES AHEAD

Implementing the landscape of AutoAI and the potential of mindful AI and actionable AI is highly challenging. Several of these challenges are as follows.

- Challenge of implementing a multifaceted functional and nonfunctional AI mind.
- Challenge of implementing multiaspect functional and nonfunctional AI actionability.
- Challenge of connecting and transforming an AI mind to AI action in one system.
- Challenge of jointly supporting all-round autonomous functions, such as autonomous task formulation, autonomous functional modularization, autonomous module configuration, autonomous data processing, autonomous structure formation, autonomous modeling, autonomous parameterization, autonomous evaluation formulation, autonomous role-function matching and configuration, and autonomous system calibration.
- Challenge of jointly supporting all-round stakeholders and bridging the bottlenecks between their different backgrounds, expertise, functionality requirements, and expectations.
- Challenge of jointly supporting engineering automation and run-time autonomous adaptive problem-solving.
- Challenge of implementing autonomous AI workflow formation, pipelining, and selection in dynamic and uncertain contexts.
- Challenge of developing business friendly AutoAI with active and online optimization for evolving and uncertain contexts.

CONCLUDING REMARKS

Automating ML, DS, and AI is a trending research area and challenge, accelerated by the emergent NAS toward end-to-end automation engineering of analytics and learning. Going beyond the engineering nature of existing AutoML and the confusion that AutoML is interchangeable with AutoDS and AutoAI, we argue for the transformation to all-round autonomous AI systems with an AI mind and mindfulness, and AI action and actionability. This requires an ambitious perspective and generation of AI, and AutoAI, i.e., AutoAI^{+}, enabling all-round AI visions, mechanisms, paradigms, and applications. The original and systematic research on mindful AI and actionable AI is thus essential for full-spectrum, human-like to human-level, and real-world AI and autonomous AI systems.

ACKNOWLEDGMENTS

This work was supported in part by the http://dx.doi.org/10.13039/100000163ARC Discovery under Grant DP190101079 and Future Fellow under Grant FT190100734.

REFERENCES

1. F. Hutter and L. J. Kotthoff Vanschoren Eds., *Automated Machine Learning: Methods, Systems, Challenges*. Berlin, Germany: Springer, 2019.
2. S. Santu, M. Hassan, M. Smith, L. Xu, C. Zhai, and K. Veeramachaneni, "AutoML to date and beyond: Challenges and opportunities," *ACM Comput. Surv.*, vol. 54, no. 8, pp. 175:1–175:36, 2022.
3. M. Zoller and M. Huber, "Benchmark and survey of automated machine learning frameworks," *J. Artif. Intell. Res.*, vol. 70, pp. 409–472, 2021.
4. M. Bahri, F. Salutari, A. Putina, and M. Sozio, "AutoML: State of the art with a focus on anomaly detection, challenges, and research directions," *Int. J. Data Sci. Anal.*, vol. 14, no. 2, pp. 113–126, 2022.
5. S. Pinker, *How the Mind Works*. Baltimore, MD, USA: Penguin Press, 2015.
6. L. Cao, "AI science and engineering: A new field," *IEEE Intell. Syst.*, vol. 37, no. 1, pp. 3–13, Jan./Feb. 2022.
7. L. Cao, "A new age of AI: Features and futures," *IEEE Intell. Syst.*, vol. 37, no. 1, pp. 25–37, Jan./Feb. 2022.
8. L. Cao, *Data Science Thinking: The Next Scientific, Technological and Economic Revolution*. Berlin, Germany: Springer, 2018.
9. P. Daugherty and H. Wilson, *Human Machine: Reimagining Work in the Age of AI*. Boston, MA, USA: Harvard Bus. Rev. Press, 2018.
10. S. Islam et al., "Automatic parallelization of representative-based clustering algorithms for multicore cluster systems," *Int. J. Data Sci. Anal.*, vol. 10, no. 2, pp. 135–159, 2020.
11. T. Elsken, J. Metzen, and F. Hutter, "Neural architecture search: A survey," *J. Mach. Learn. Res.*, vol. 20, no. 55, pp. 1–21, 2019.
12. N. Reunanen, T. Raty, and T. Lintonen, "Automatic optimization of outlier detection ensembles using a limited number of outlier examples," *Int. J. Data Sci. Anal.*, vol. 10, no. 4, pp. 377–394, 2020.
13. P. Wurman, P. Stone, and M. Spranger, "Challenges and opportunities of applying reinforcement learning to autonomous racing," *IEEE Intell. Syst.*, vol. 37, no. 3, pp. 20–23, May/Jun. 2022.
14. L. Cao, "Data science: Nature and pitfalls," *IEEE Intell. Syst.*, vol. 31, no. 5, pp. 66–75, Sep./Oct. 2016.
15. L. Cao, "Data science: Profession and education," *IEEE Intell. Syst.*, vol. 34, no. 5, pp. 35–44, Sep./Oct. 2019.
16. S. Gardner et al., "Constrained multi-objective optimization for automated machine learning," in *Proc. IEEE Int. Conf. Data Sci. Adv. Anal.*, 2019, pp. 364–373.
17. L. Gilpin, V. Penubarthi, and L Kagal, "Explaining multimodal errors in autonomous vehicles," in *Proc. IEEE Int. Conf. Data Sci. Adv. Anal.*, 2021, pp. 1–10.
18. S. Alnegheimish et al., "Cardea: An open automated machine learning framework for electronic health records," in *Proc. IEEE 7th Int. Conf. Data Sci. Adv. Anal.*, 2020, pp. 536–545.
19. D. Kahneman, *Thinking, Fast and Slow*. New York, NY, USA: Farrar, Straus and Giroux, 2013.
20. K. Dutton, *Black and White Thinking: The Burden of a Binary Brain in a Complex World*. London, U.K.: Bantam Press, 2020.

LONGBING CAO is a professor and an ARC future fellow (professorial level) with the University of Technology Sydney, Sydney, NSW, 2007, Australia. His research interests include AI, data science, machine learning, behavior informatics, and their enterprise applications. Contact him at longbing.cao@uts.edu.au.