

## A BenchCouncil view on benchmarking emerging and future computing

Jianfeng Zhan

Research Center for Advanced Computer Systems, Institute of Computing Technology, Chinese Academy of Sciences, China

### ARTICLE INFO

#### Keywords:

Benchmark science and engineering  
Benchmarking challenges  
Extrinsic property  
Process entanglement  
Instantiation bias  
Unified benchmark definition  
Conceptual framework  
Benchmarking methodology  
Traceability  
Supervised learning  
Emerging computing  
Future computing  
BenchCouncil Plan

### ABSTRACT

The measurable properties of the artifacts or objects in the computer, management, or finance disciplines are extrinsic, not inherent — dependent on their problem definitions and solution instantiations. The processes of problem definition, solution instantiation, and measurement are entangled. Only after the instantiation can the solutions to the problem be measured. Definition, instantiation, and measurement have complex mutual influences. Meanwhile, the technology inertia brings instantiation bias — trapped into a subspace or even a point at a high-dimension solution space. These daunting challenges, which emerging computing aggravates, make metrology cannot work for benchmark communities. It is pressing to establish independent benchmark science and engineering.

This article presents a unifying benchmark definition, a conceptual framework, and a traceable and supervised learning-based benchmarking methodology, laying the foundation for benchmark science and engineering. I also discuss BenchCouncil's plans for emerging and future computing. The ongoing projects include defining the challenges of intelligence, instinct, quantum computers, Metaverse, planet-scale computers, and reformulating data centers, artificial intelligence for science, and CPU benchmark suites. Also, BenchCouncil will collaborate with ComputerCouncil on open-source computer systems for planet-scale computing, AI for science systems, and Metaverse.

### 1. Introduction

Benchmarking is widely practiced in different disciplines without a consensus on a consistent definition. For example, in the computer science discipline, the community uses a set of workload implementations to measure CPU (processor) performances [1,2]. In machine learning, standardized data sets labeled with ground truths are used to define a data science problem [3,4]. In the management discipline, the industry best practices are searched and compared against different products, services, and processes [5,6]. All are called benchmarks or benchmarking. In the previous work, I concluded five categories of benchmarks [6]: measurement standards, standardized data sets with defined properties, representative workloads, representative data sets, and industry best practices.

The inconsistency or chaos results from the following fact. Per JCGM 200 definition, metrology is the science of measurement and its application [7,8]. Metrology measures intrinsic properties independent of an observer, like length, time, and power. There is a true quantity for each inherent property, where a probability could state the coverage interval containing the true value [7,8]. However, the measurable properties of the artifacts or objects in the computer, management, or finance disciplines are extrinsic, not inherent — dependent on their problem definitions and solution instantiations. Unlike the processes in metrology that are linear and static, the processes of a benchmark have

complex mutual influence. The problem definition, solution instantiation, and measurement processes are entangled and indivisible, and only after the instantiation can the solutions to the problem be measured, which I call process entanglement. Users adhere to existing products, tools, platforms, and services, called technology inertia [9]. The technology inertia traps the solution to a problem into a specific exploration path — a subspace or even a point at a high-dimension solution space. The instantiation bias impacts the measurement of the extrinsic properties.

Our society increasingly relies upon information infrastructure with daunting complexity that dwarfs the previous systems, making it difficult to trace the problem definition. Instead, the biased instantiation of solutions becomes the proxy of the problems, missing the forest for the trees. As shown in Fig. 1, these daunting challenges: extrinsic properties, process entanglement, and instantiation bias, result in the benchmark community's inability to reuse the metrology knowledge and the de facto isolation of benchmark communities, like computers, management, and finance, developing different methodologies, tools, and practices. It is pressing to establish independent benchmark science and engineering.

Echoing my past call [6], this article further builds up benchmark science and engineering. I define the benchmark from the perspective of problems and solutions. A benchmark is an explicit or implicit

E-mail address: [zhanjianfeng@ict.ac.cn](mailto:zhanjianfeng@ict.ac.cn).

URL: <https://www.benchcouncil.org/zjf.html>.

<https://doi.org/10.1016/j.tbench.2022.100064>

Available online 25 May 2022

2772-4859/© 2022 The Authors. Publishing services by Elsevier B.V. on behalf of KeAi Communications Co. Ltd. This is an open access article under the CC BY-NC-ND license (<http://creativecommons.org/licenses/by-nc-nd/4.0/>).

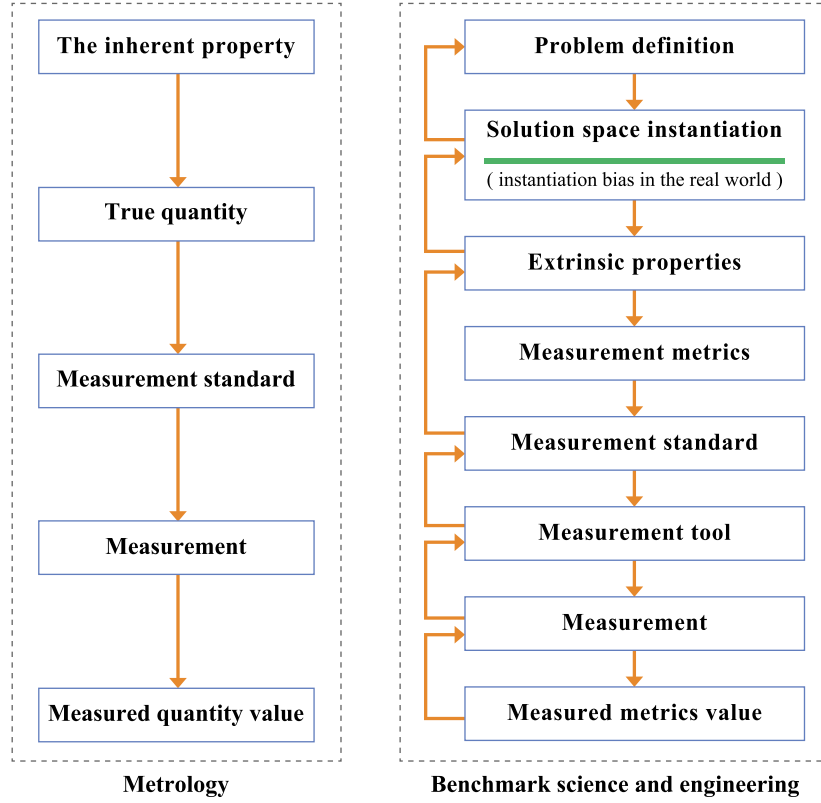


Fig. 1. With respect to metrology, the benchmarking challenges – the extrinsic properties, process entanglement, and instantiation bias – explain why metrology cannot work for the benchmark community. First, the property of a benchmark is not inherent but dependent on its problem definition and solution instantiation. Second, only after the instantiation can the solutions to the problem be measured. The processes of problem definition, solution instantiation, and measurement are entangled, and they have complex mutual influences. Third, instantiation introduces many biases.

definition of a problem, an instantiation of a problem, an instantiation of state-of-the-practice solutions as the proxy to the problem, or a measurement standard that quantitatively measures the solution space. I propose a concise conceptual framework for the benchmark science and engineering, at the core of which is the extrinsic properties. The extrinsic property is a benchmark property that is not inherent but dependent on a problem definition and solution instantiation. I propose a traceable and supervised learning-based methodology to tackle the challenges of extrinsic property, process entanglement, and instantiation bias. The essence of the methodology has two integrated parts: manage the traceability of the processes from the problem definition and solution instantiation to measurement; search for the best solution through supervised learning with reference to a thoroughly-understood process from the problem definition, solution instantiation to measurement.

Also, I discuss BenchCouncil’s plan for emerging and future challenges. The ongoing projects include defining the challenges of intelligence, instinct, quantum computers, Metaverse, planet-scale computers, and reformulating data centers, artificial intelligence for science, and CPU benchmark suites. Also, BenchCouncil will collaborate with ComputerCouncil [10] on the open-source computer systems for Planet-scale computing [11], AI for science [12], and Metaverse [13].

The organization of this article is as follows. Section Two presents the background and challenges and explains why metrology cannot work for the benchmark community. Section Three describes why emerging computing aggravates the benchmark challenges. Section

Four lays the foundation for benchmark science and engineering, including the unifying definition of benchmarks, the conceptual framework, and the benchmarking methodology. Section Five details BenchCouncil’s plan. Section Six concludes.

## 2. Background and challenge: why metrology cannot be directly reused for benchmark science and engineering

In this section, I first introduce the metrology concepts as background and then present the benchmarking challenges and explain why metrology cannot work in the benchmark community.

### 2.1. Background: metrology concepts

As shown in Fig. 2, I provide a simple but systematic metrology concept framework to clarify why metrology cannot be directly reused for establishing benchmark science and engineering. I present and modify most of those concepts from [7,8]. But I define some concepts to emphasize why metrology cannot work for the benchmark community, e.g., inherent properties. To keep concise, I only stay with necessary metrology concepts.

The inherent property is a property of a phenomenon, body, or substance that is independent of an observer, e.g., length and energy [7]. The inherent property can have various magnitudes. True quantity is the magnitude of an inherent property of an individual phenomenon, body, or substance that is independent of an observer, e.g., the radius of a given circle, the kinetic energy of an identified particle in a given system [7,8]. Unit of measurement [8] is a definition and its physical realization, used as a reference to assign a value to a true

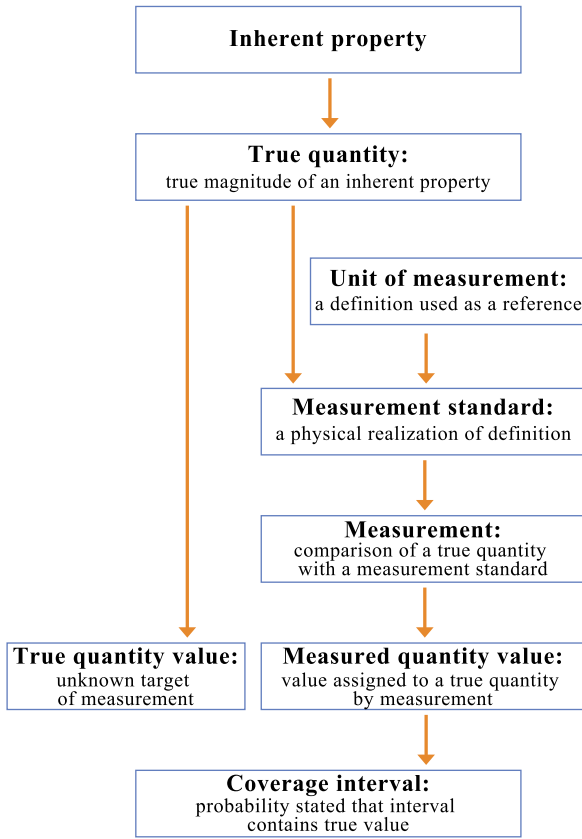


Fig. 2. A simple but systematic metrology conceptual framework is used to clarify why metrology cannot be directly reused for benchmark science and engineering. Some concepts are defined by myself, while the other concepts are reused or modified from [7,8]. Only necessary metrology concepts are reserved to keep concise.

quantity. Measurement standard [8] is a physical realization of a unit of measurement, with a stated quantity value and associated measurement uncertainty.

Measurement [8] is a process of comparing a true quantity with a measurement standard to assign the true quantity one or more quantity values that are traceable to a unit of measurement. A quantity value obtained by the measurement is referred to as a measured quantity value [7]. True quantity value [7] is a quantity value consistent with the definition of a quantity, which is an unknown measurement target [7]. A coverage probability [7] is a probability that a specified coverage interval contains the true quantity value.

## 2.2. The benchmarking challenges: extrinsic properties, process entanglement, and instantiation bias

In the previous work [6], I have noticed the differences in properties of the artifacts or objects in the computer, management, or finance disciplines from those classical ones, like length, time, and power. The properties of the artifacts or objects in the computer, management, or finance disciplines are extrinsic, dependent on their problem definitions and solution instantiations. Instead, the classical properties like time and length are inherent, independent of the observers. From a concept perspective, it is easy to say there are three essential processes: problem definition, solution instantiation, and measurement. However, a problem definition is abstract; only after the instantiation can the solutions to the problem be measured. Moreover, the problem definition, solution instantiation, and measurement processes are entangled and indivisible, which I call process entanglement. Only by fully understanding the side effect of the extrinsic properties and process entanglement can we avoid many traps. I elaborate on this viewpoint from different perspectives.

Before proposing the conceptual framework for benchmark science and engineering (I defer it to Section 4.2), I stay with the metrology concepts in Section 2.1 to depict the challenges.

A subtle change in the definitions of a problem may lead to wildly varied solutions and significantly different measured quantity value. I take the classical matrix multiplication problem [14,15] as an example. Blalock et al. [15] reformulate the classical matrix multiplication problem as follows. The following reformulation is cited from [15]. A and B are two matrices. A is  $R^{N \times D}$  and B is  $R^{D \times M}$ ,  $N \gg D \gg M$ . Given a computation time budget  $\tau$ , the task constructs three functions  $g(\cdot)$ ,  $h(\cdot)$ , and  $f(\cdot)$ , along with constants  $\alpha$  and  $\beta$ , such that

$$\|\alpha f(g(A), h(B)) + \beta - AB\|_F < \epsilon(\tau) \|AB\|_F \quad (1)$$

for as small an error  $\epsilon(\tau)$  possible. For this reformulated problem, they introduce a learning-based algorithm that greatly outperforms existing methods [15]. This is a typical example of a subtle change in the definitions of a problem leading to wildly varied solutions and significantly different measured quantity values.

Furthermore, the solutions instantiations at different levels also interplay with each other and finally impact measured quantity values. The obvious example is deep learning. Algorithms and neural network architectures play a significant role. The hardware implementations, like different precision, e.g., single-precision, double-precision, or mixed precision, impact the learning dynamics. Even for the same system with different scales, the interactions among system size and mini-batch size significantly impact the measured quantity values like time-to-quality – the training time to achieve the state-of-the-art quality [16–20].

The processes in metrology are linear and static. However, for a benchmark, as shown in Fig. 1, the processes of problem definition, solution instantiation, and measurement are entangled, having complex mutual influences. The subtle difference in a problem definition will lead to a wildly varied solution, and its instantiation finally significantly impacts the measured quantity value. The solution instantiation provides the basis for measurement tools, and the latter uses state-of-the-practice instantiations that update frequently, which affects the measured quantity value. Also, the measured quantity values provide hints on searching for the best instantiation in the solution space.

Moreover, the instantiation introduces many biases, which I call instantiation bias. For example, in the computer system and architecture disciplines, Wang et al. [21,22] found that merely conducting microarchitecture-dependent or microarchitecture-independent, or ISA-independent (ISA is short for instruction set architecture) workload characterization (a form of measurement) will lead to misleading or erroneous conclusions. These significant differences in measured quantity values resulted from the solution instantiations at different levels themselves. Before performing microarchitecture-dependent or microarchitecture-independent, or ISA-independent workload characterization, the necessary step is instantiating a computer workload benchmark on a specific microarchitecture, a particular instruction set architecture, or an intermediate representation (very close to the source code), respectively. The community opts for the widely used ISA, IR (intermediate representation) for instantiation.

Matsuoka et al. also found the implementation of biases and complexity traps [23] in the instantiation process: on the one hand, any implementation of a computer workload benchmark entails multiple implicit biases towards algorithms, programming languages, data layouts, and parallelization approaches; on the other hand, the benchmarks, abstracted from large or legacy scientific codes and tuned for previous computer architectures, trap the co-design participants into considering only similar architectures.

Other observations are from the data sets, which many communities like machine learning use to explicitly or implicitly define a problem. It is prohibitively costly to build a representative and fidelity data set that can capture real-world characteristics. Hence, in reality, the goal often degrades to a workable data set. For example, for ImageNet [3],

it is easy to collect familiar animal and plant pictures, while the rare ones are difficult to obtain. Considering the data set is the cornerstone of many challenges like auto-driving and automatic medical diagnosis, this far-fetched methodology has many hidden flaws and risks.

### 3. Emerging computing aggravates the challenges<sup>1</sup>

Modern society is digitized, increasingly relying upon information infrastructure. The information infrastructure consists of massive Internet of Things (IoT), edge devices, data centers, and high-performance computers. Those systems collaborate to handle big data, train AI models and provide Internet services augmented by AI inference for huge end-users with guaranteed quality of services. From a benchmarking perspective, emerging computing like Big Data, AI, and Internet Services are significantly different from the traditional workloads characterized by SPEC CPU (desktop workloads) [1], TPC-C [25], TPC-Web (Traditional web services) [26], and HPL (high-performance computing) [27] benchmarks, raising serious challenges.

The first challenge is fragmentation. There are substantial fragmented application scenarios, a marked departure from the past [24]. For example, hundreds or even thousands of ad-hoc big data solutions, termed NoSQL or NewSQL, are proposed to handle different application scenarios. For AI, the same observation holds. There are tens or even hundreds of organizations that are developing AI training or inference chips to tackle their challenges in different application scenarios, respectively [19,28].

The second challenge is de facto isolation. Internet service provider giants own and treat real-world data sets and workloads or even AI models as first-class confidential issues. The treasure is hidden in data centers and isolated between academia and industry, or even among different providers [29]. The dire situation poses a huge obstacle for our communities towards developing an open and mature research field [29].

The third challenge is the complexity of collaboration: HPC systems, data centers, edge, and IoT devices collaboratively handle the challenges; In the collaborations, different distributions of data sets, workloads, machine learning, or AI models may substantially affect the system's behaviors; the interaction patterns among IoT, edge, and data centers changes fast, embodying different architecture.

The fourth challenge originates from service-based architecture. On the one hand, the software-as-a-service (SaaS) development and deploy model makes the workloads change very fast (so-called workload churn) [30], and it is not scalable or even impossible to create a new benchmark or proxy for every possible workload [31]. On the other hand, modern Internet services adopt a microservice-based architecture, often consisting of various modules with long and complex execution paths across different data centers. As the worst-case performance (tail latency) [32] does matter, the micro-service-based architecture also poses a severe challenge to benchmarking [29,33].

The final but not least challenge is the stochastic nature of AI. AI techniques are widely used to augment modern products or Internet services. The nature of AI is stochastic, allowing multiple different but equally valid solutions [19]. Many factors manifest the uncertainties of AI, e.g., the adverse effect of lower-precision optimization on the quality of the final model, the impact of scaling training on time-to-quality, and run-to-run variation in terms of epochs-to-quality [19]. However, the measurement process mandates being repeatable (the same team) and reproducible (different teams). This conflict raises serious challenges.

Emerging computing aggravates the benchmarking challenges discussed in Section 2.2. First, it is difficult to trace the original problem definition, that is, the target to be achieved. Second, taking the instantiation of solutions as the proxy for the problem aggravates the instantiation bias and makes the community further trapped in the specific solutions.

<sup>1</sup> This section is written based on an unpublished technique report [24], of which I am the lead author.

## 4. Building up benchmark science and engineering

This section proposes the unifying definitions of benchmarks, the conceptual framework, and the benchmarking methodology, which lays the foundation for benchmark science and engineering.

### 4.1. The unifying definition of benchmarks

Previously, I concluded five categories of benchmarks [6]: measurement standards, standardized data sets with defined properties, representative workloads, representative data sets, and industry best practices. In this section, I give a simple and unifying definition to cover five categories of benchmarks and reveal their essence. A benchmark is an explicit or implicit definition of a problem, an instantiation of a problem, an instantiation of state-of-the-practice solutions as the proxy to the problem, or a measurement standard that quantitatively measures the solution space.

A benchmark has three essential processes, some of which often be omitted or implicitly stated in practice: definition, instantiation, and measurement. I explain the process of definition and instantiation from various perspectives in the rest paragraphs of Section 4.1. I leave the discussion of the process of measurement in Sections 4.2, 4.3.

#### 4.1.1. Definition

The first process is the definition. Defining a problem explicitly or implicitly is the fundamental role that a benchmark could play in almost all disciplines. Only after clearly defining a problem can we figure out the solutions and compare them against the others. For example, Alan Turing, in 1950 [34] formulated the problem of what intelligence is as an imitation game: the game tests whether an interrogator can distinguish a machine's ability from a human. Turing's problem definition inspires several generations to explore the solutions to achieve intelligence.

There are many ways to define a problem, e.g., using a natural language or mathematics. From an accuracy perspective, mathematically defining the problem is a better choice. Unfortunately, many problems cannot be accurately depicted in this way.

The NAS parallel benchmarks [35] claimed that the common requirements should be specified in a paper-and-pencil approach [24]. A paper-and-pencil approach is a vague description — It can be mathematical, textual, or even visually. In the computer science discipline, this approach is well-practiced in the database community but not adopted in the computer architecture community.

Shun et al. [36] advocated a methodology to build benchmarks using problem definitions, and they created the problem-based benchmark suite (PBBS). PBBS is a set of benchmarks for comparing parallel algorithmic approaches, parallel programming language styles, and machine architectures across a broad set of problems. Specifically, a problem-based benchmark mandates a problem specification and a set of input distributions while not detailing the requirements in terms of algorithmic approach, programming language, or machine architecture [36].

#### 4.1.2. Instantiation

The second process is instantiating a problem or instantiating state-of-the-practice solutions as the proxy to the problem or challenge. As a replacement or complement, these are two different ways. First, an instantiation of the problem is used. For example, a data set is often used to instantiate a problem in the machine learning community. Li et al. [12] further classified the problem definitions into problem class, problem settings, and problem cases. Second, an instantiation of state-of-the-practice solutions is used as the proxy for the problem. For example, the computer architecture community provides state-of-the-practice implementations of a group of computer workloads like SPEC CPU [2,37]. SPEC CPU is a proxy to the problems.



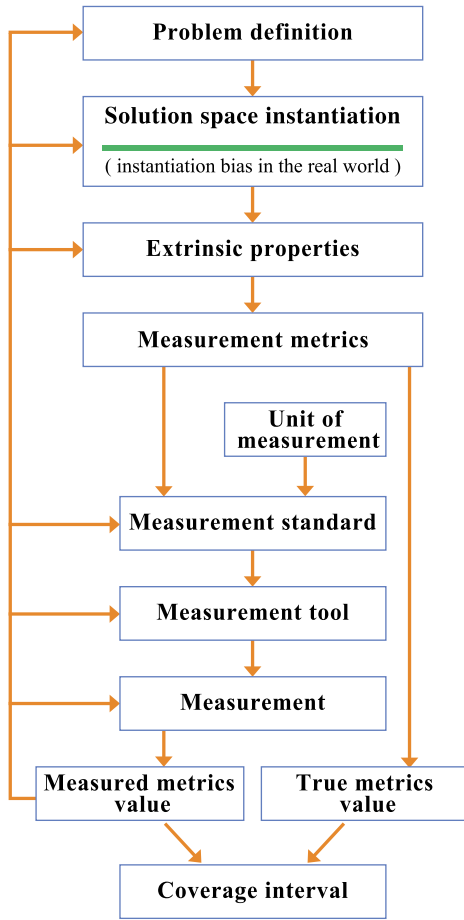


Fig. 3. The conceptual framework of benchmark science and engineering.

There are two reasons for this replacement or complement. First, as a replacement, it serves as the proxy for the problem that is too difficult to define. Second, as a complement, the instantiation brings enriched and necessary details that set more specific problem settings. Each instantiation is a subspace or point – which is often state-of-the-practice – in the solution space to the problem, e.g., using source code or binary code, which brings instantiation bias.

#### 4.2. The conceptual framework of benchmark science and engineering

As shown in Fig. 3, I propose the conceptual framework of benchmark science and engineering. The extrinsic property is a benchmark property that depends on a problem definition and its solution instantiation. The extrinsic property can have various magnitudes. Measurement metrics are the magnitude of a benchmark's extrinsic property, which depends on a problem definition and solution instantiations. Unit of measurement [8] is a definition and its realization, used as a reference to assign a value to a measurement metric. Measurement standard [8] is a realization of a unit of measurement, with a stated metric value, associated measurement uncertainty, and a repeatable (the same team) and reproducible (different teams) measurement methodology. The measurement tool implements a measurement standard that can be calibrated and traceable. Traceability [7] is a property of a measurement result whereby the result can be related to a reference through a documented unbroken chain of calibrations, each contributing to the measurement uncertainty. The measurement tools should be open-sourced and can be replicated by different teams.

Measurement [8] is a process of comparing a measurement metric with a measurement standard to assign one or more measured values

to a measurement metric that are traceable to a unit of measurement. A value obtained by the measurement is referred to as a measured metric value [7]. True metric value [7] is a value consistent with the definition of a measurement metric that is specific to the extrinsic properties of a concrete problem definition and solution instantiation. The true metric value is an unknown measurement target [7]. A coverage probability [7] is a probability that the true metric value is contained within a specified coverage interval.

#### 4.3. The traceable and supervised-learning based benchmarking methodology

A benchmark has no inherent properties, and its extrinsic property is dependent on its problem definition and solution instantiation. Meanwhile, the processes of definition, instantiation, and measurement are entangled, and they have complex mutual influences.

I propose a traceable methodology to tackle the above challenge, at the core of which is to manage the traceability of the processes from the problem definition and solution instantiation to measurement. Fig. 4 shows that problem definition, solution instantiation, extrinsic properties, measurement standard, measurement tool, and measured metrics value have complex mutual influence. The problem definition is the origin of this relationship. No other below entities, like solution instantiation, can impact the problem definition directly. Still, solution instantiations may provide clues for the subtle change of the problem definition, affecting the other entities significantly. At the top level, I suggest a formal definition of problems and tracing the relationship among the different subtle definitions of problems. For many state-of-the-practice benchmarks, the definition process is omitted. It should regularly keep an eye on revisiting the process from the problem definition to solution instantiation, or else the outdated instantiation will be a trap. The solution instantiation provides the basis for measurement tools. It is mandatory to search for state-of-the-art or state-of-the-practice solutions and implement them in the measurement tool.

There is an explosion from the problem definition to solution instantiations. Put in other words, the lower level has more state space [6]. For example, there is increasing state space in a computer workload benchmark, from a mathematical problem definition, an algorithm, an intermediate representation, an ISA-specific representation, to a micro-architectural representation. The technology inertia traps the solution into a specific exploration path — a subspace or even a point at a high-dimension solution space, called instantiation bias. The instantiation bias impacts the measurement of the extrinsic properties. Also, an unguided exploration may drift away from optimized solutions.

I propose the supervised learning-based methodology to tackle the challenge of instantiation bias. Supervised learning is a machine learning branch that trains a predictive model using labeled data with known outcomes. Fig. 4 shows the thoroughly-understood process from the problem definition, solution space instantiation, extrinsic properties, measurement standard, measurement tool, to measured metrics value, standing as a ground truth. From this ground truth, it is easy to learn how the change of the top entities impacts the below. For example, if the problem is reformulated, the solution instantiation changes accordingly. Finally, the measured metrics values are significantly affected. The benchmark plays the role of connecting the problem with its solution space. By exploring the solution space and observing the effect of its change on the measured metrics values, it is possible to search for the best solution. This search process could leverage state-of-the-art deep learning techniques. Of course, this learning dynamic will be very complex. Fig. 5 shows an example on how to use this methodology in computer architecture.

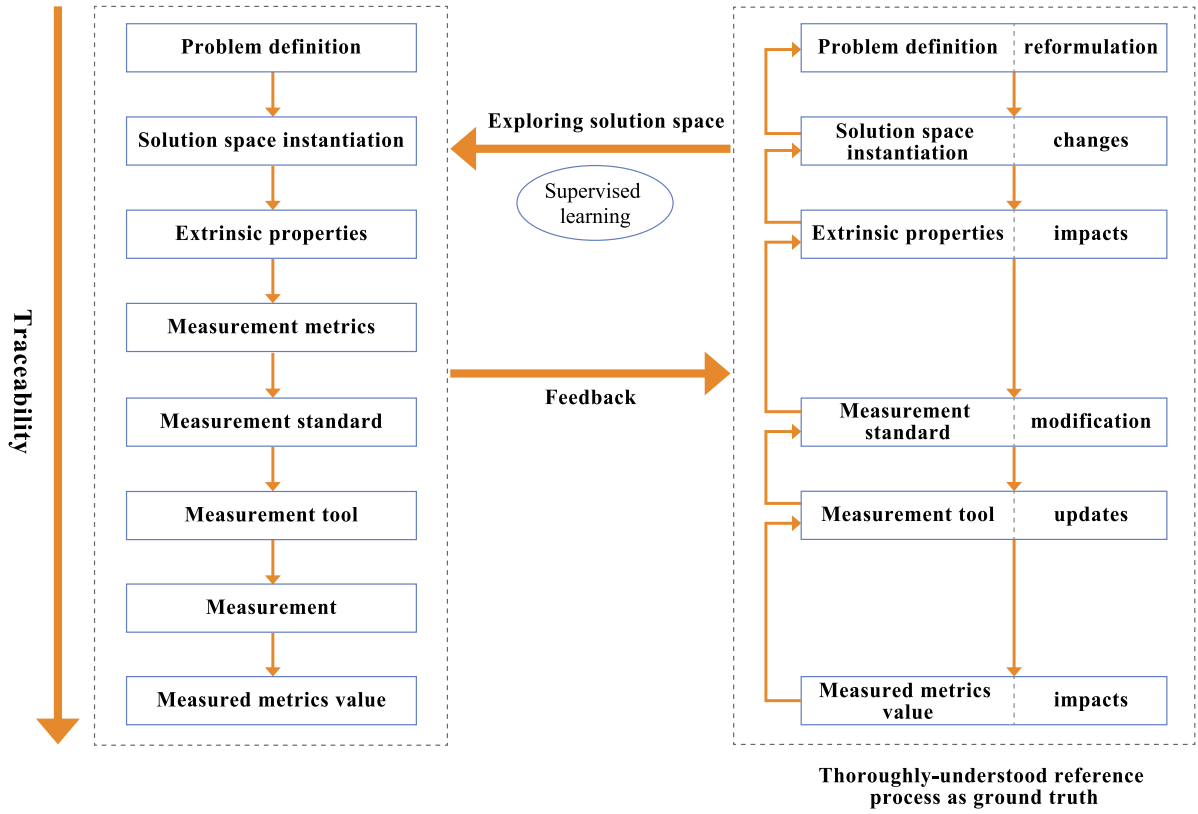


Fig. 4. A traceable and supervised learning-based benchmarking methodology to tackle the challenges of extrinsic property, process entanglement, and instantiation bias.

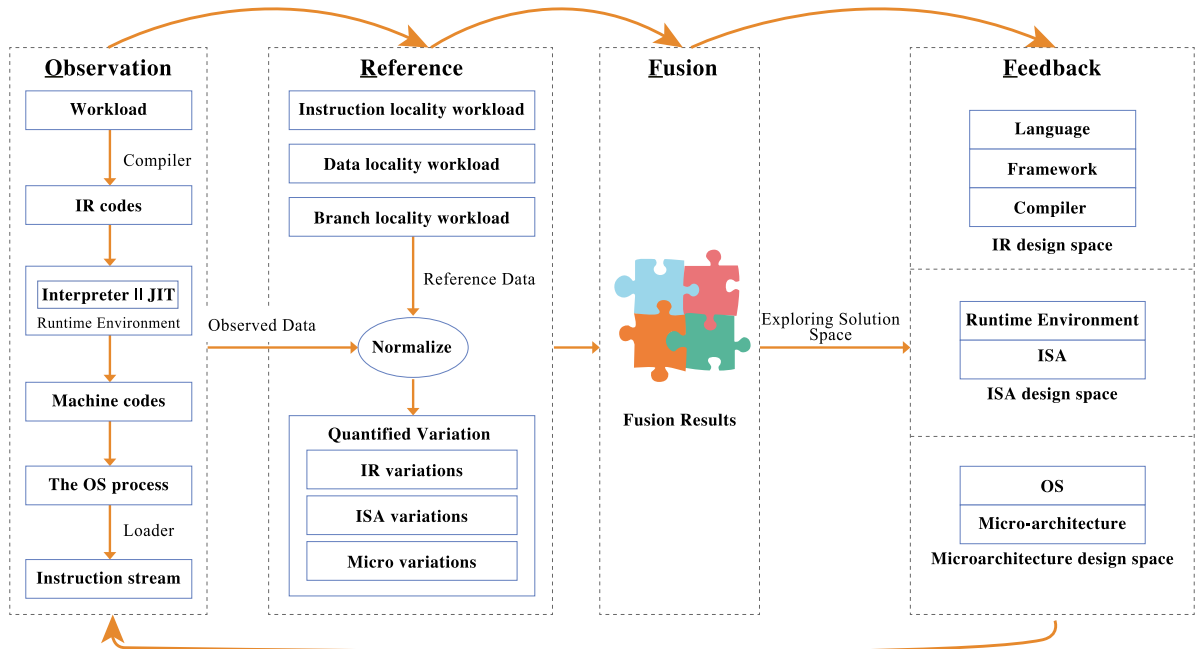


Fig. 5. The application of the traceable and supervision learning-based benchmarking methodology in computer architecture. This figure is cited from [22] with the permission of the authors.

#### 4.4. Re-interpret five categories of benchmarks

I use the benchmark definition proposed in 4.1 to re-interpret five categories of benchmarks defined in [6].

The first category of the benchmark is a measurement standard used to measure the solution space to the problem. I use the Linpack benchmark – an example from high-performance computing – to explain this

category of benchmarks. The Linpack benchmark [38] is widely used to report the performance of a high-performance computer. The problem definition of Linpack is a linear system of equations of an order  $n$ :  $Ax = b$ . The solution uses the LU factorization with partial pivoting. The measurement metrics are the floating-point operations count of the solving algorithm, which is  $(2 * n^3/3 + 2 * n^2)$  operations, and the execution time of running the benchmark. HPL is one of the reference

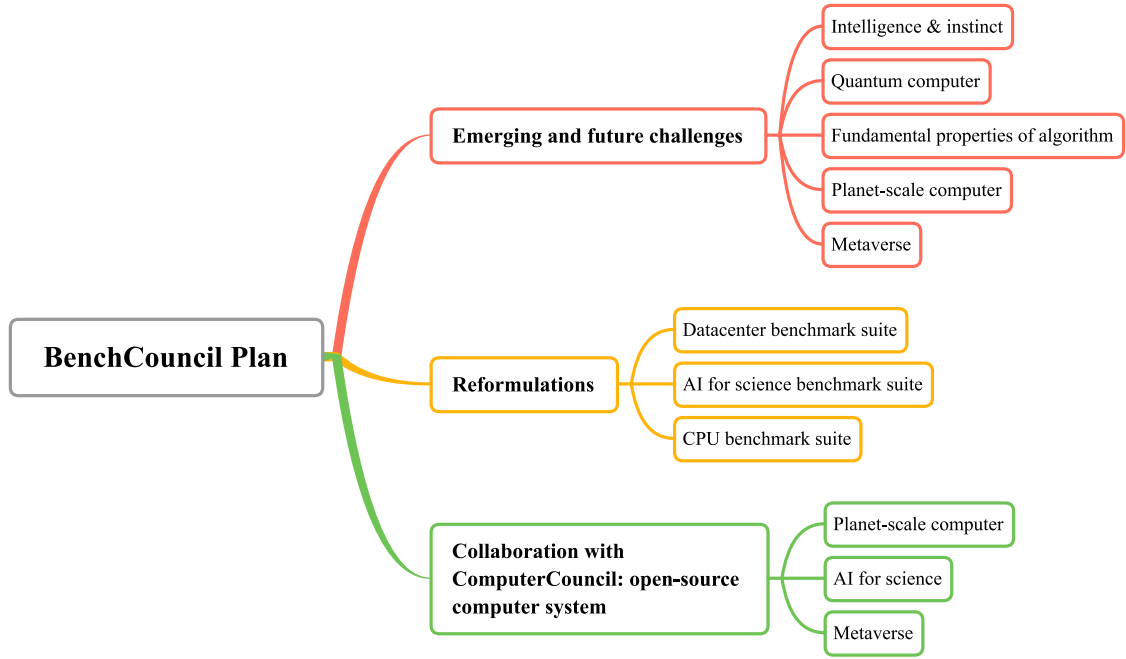


Fig. 6. BenchCouncil's plan on defining the challenges of emerging and future computing and the collaboration with ComputerCouncil on the open-source computer systems.

implementations of the measurement tool used to evaluate against different high-performance computer implementations (solutions). The measurement standard also details the reproducible and repeatable measurement methodology to compare against other solutions: users must report a residual for the accuracy of the solution with  $\|Ax - b\|/(\|A\|\|x\|)$ . The TOP500 list reports the measured metrics values. The measurement metrics highly depends on its problem definition and solution instantiations.

The second category of benchmarks is the representative workloads that run on the systems under measurement [6]. The representative workloads are the definition of the problem or an instantiation of state-of-the-practice solutions as the proxy to the problem. The problem-based benchmark suite (PBBS) [36], TPC-C [25], TPC-Web (Traditional web services) [26] are typical problem definition examples. They also provide the instantiation of state-of-the-practice solutions as the measurement tool. Without explicit problem definition, SPEC CPU (desktop workloads) [1], BigDataBench [39], BigBench [40], AIBench [20,33] and MLPerf [19] are the instantiations of state-of-the-practice solutions and they serves as the proxy to the problem.

The third category of the benchmarks is the implicit definition of the problem using a standardized data set. The standardized data set represents a real-world data science problem with defined properties, some of which have ground truth [6,41]. ImageNet [3] (deep learning benchmark) and MIMIC-III [4] (critical care benchmark) are typical examples.

The fourth category of benchmarks is a representative data set, used as a Ref. [6]. This category of benchmarks is an instantiation of a problem. For example, an index (statistical measure) calculated from a representative set of underlying data and used as a reference for financial instruments or contracts [42] is a benchmark in finance. The London Interbank Offered Rate (Libor) and the Euro Interbank Offered Rate are well-known financial benchmarks [6,42].

The fifth category of benchmarks is the industry best practices in diverse domains [6]. Benchmarking is continuously searching the industry best practices with superior performance and measuring products, services, and processes against them [5,6]. The industry best practices are instantiations of the state-of-the-practice solutions to the problem or grand challenge.

## 5. BenchCouncil's plan on emerging and future computing

Fig. 6 presents BenchCouncil's plan for defining the challenges of emerging and future computing and collaborating with ComputerCouncil on the open-source computer systems. First, I introduce BenchCouncil's plan for emerging and future challenges.

First, what is intelligence? What is instinct? What is the distinction between intelligence and instinct? The pre-trained language models, such as BERT and GPT-3, seem to outperform the capability beyond the Turing test [43]. Many previous works have reformulated the problem of what intelligence is [44,45]. It is necessary to revisit the processes from the intelligence problem definition, solution instantiations to measurement. For example, there are many ways to solve these challenges to somewhat extent, including traditional machine learning, deep learning, and brain-inspired computing. Letting them compete in the same arena is essential. According to [46], instinct is an inborn impulse or motivation to action typically performed in response to specific external stimuli. But how do we distinguish intelligence from instinct? What are the differences between the octopus, birds, apes, and ants' behaviors? Are they intelligence or instinct?

Second, quantum computers emerge as a new computational paradigm with unprecedented capability [47]; what are the problems or grand challenges which the quantum computers do the best? How do state-of-the-practice computers compete against quantum computers in handling different or overlapping domains of problems or grand challenges? The community must ponder this fundamental issue before delving into different levels of instantiations of solutions.

Third, computer algorithms almost govern the running of our society. It is pressing to think, specify, verify and test what fundamental properties an algorithm must have before being embedded in our society. Think about Twitter and Facebook's impact on the election in many vote-based democratic societies. It is vital to propose benchmarks against those algorithms before putting them into practice.

Fourth, information infrastructure becomes the cornerstone of our society [10], and many fundamental applications like medical emergency management and smart cities applications rely upon planet-scale distributed systems consisting of massive Internet of Things (IoT) devices, edges and data centers, which I call planet-scale computers [11]. Different distributions of data sets, workloads, or AI models may substantially affect the system's behaviors, and the system architectures are

undergoing fast evolution regarding the interactions among IoT, edge, and data centers [24]. How can the community propose benchmarks for those ultra-scale emerging and future applications [33]?

Metaverse is an umbrella term. It is predicted to be a brand-new way for people to immersively access the Internet, interact with each other or digital avatars in the cyberworld, and manage digital assets. Though many industry giants are pushing towards those goals, the process itself is in a Cambrian explosion of different things in the forms of concepts, prototypes, products, or services. It is pressing to propose a benchmark suite to define the Metaverse problem or challenge, explore and evaluate state-of-the-art and state-of-the-practice solutions [13].

Many old problems need reformulation. For example, the Berkeley multidisciplinary groups proposed to use thirteen “Dwarfs” [48] – A dwarf is an algorithmic method that captures a pattern of computation and communication – to design and evaluate parallel programming models and architectures. When AI has seen as a new dawn in the traditional and emerging scientific area, how to reformulate those problems [12]?

Datacenters have become the fundamental infrastructure of modern society. There are substantial fragmented application scenarios in big data, AI, and Internet service areas, a marked departure from the past [24]. Virtualization technologies like containers are widely used as resource management and performance isolation facilities. However, the current BenchCouncil benchmark suites like BigDataBench [39] and AIBench [20,33] are fragmented without providing a full-picture definition of the problems or challenges in data centers. Moreover, a lack of simple but elegant abstractions prevents achieving both efficiency and general purpose [24]. For example, hundreds or even thousands of ad-hoc NoSQL or NewSQL solutions are proposed to handle different big data application scenarios [24]. Contrasted, the relation algebra is generalized for the database theory and practice, and any complex query can be written using five primitives like select, project, product, union, and difference [49]. Though domain-specific software and hardware co-design is promising [50], the lack of simple but unified abstractions has two side effects [24]: it is prohibitively costly to build an ad-hoc solution; single-purpose systems and architectures are structural obstacles to resource sharing. Proposing simple but elegant abstractions is an integrated part of managing the traceability of the process from the problem definition to solution instantiation.

The CPU benchmark suite like SPECCPU [2,37] advanced the evolution of different processor architectures. However, the SPEC CPU is an instantiation of state-of-the-practice solutions as the proxy to the problem, severely biased towards market-dominant CPU architecture, high-performance languages like C, and high-performance computing and desktop workloads. The BENCHCPU project [51] will propose a new CPU benchmark suite.

### 5.1. The collaboration with ComputerCouncil

As a non-profit international organization, the International Open-source Computer Council (ComputerCouncil) mission is to unite the science and technology community to tackle the challenges of information technology decoupling [52]. ComputerCouncil initiates the open-source computer system (OSCS) initiative to tackle the challenges of IT decoupling.

ComputerCouncil will choose three emerging areas: planet-scale computers — planet-scale distributed systems and applications built on IoTs, edges, and datacenters [11], AI for science [12], and Metaverse [13] as the initial targets of the OSCS initiative. BenchCouncil will cooperate with ComputerCouncil: the former focuses on the benchmarks, while the latter concentrates on the open-source computer systems for the three emerging areas.

## 6. Conclusion

This article concluded benchmarking challenges as extrinsic properties, process entanglement, and instantiation bias. The measurable properties of a benchmark are not inherent but dependent on their problem definitions and solution instantiations. The processes of problem definition, solution instantiation, and measurement are entangled and have complex mutual influences. The technology inertia leads to a specific exploration path – a subspace or even a point at a high-dimension design space. Those challenges make metrology cannot work for benchmark communities and call for independent benchmark science and engineering.

I proposed a unified benchmark definition, a conceptual framework, and a traceable and supervised learning-based benchmarking methodology, laying the foundation for benchmark science and engineering. A benchmark is an explicit or implicit definition of a problem, an instantiation of a problem, an instantiation of state-of-the-practice solutions as the proxy to the problem, or a measurement standard that quantitatively measures the solution space. At the core of the conceptual framework, the extrinsic property is a benchmark property that depends on a problem definition and its solution instantiation. The essence of the proposed benchmarking methodology has two integrated parts: manage the traceability of the processes from the problem definition and solution instantiation to measurement; search for the best solution through supervised learning with reference to the thoroughly-understood processes from the problem definitions and solution instantiations to measurements. Also, I elaborated BenchCouncil’s plan to define emerging and future computing challenges and collaborate with ComputerCouncil on open-source computer systems.

## Acknowledgments

I am very grateful to Mr. Shaopeng Dai for compiling the references, Mr. Shaopeng Dai and Mr. Qian He for drawing Figs. 1, 2, 3, 4, and 6, and Dr. Lei Wang for discussing and contributing significantly to the presentations of Figs. 1, 4, 6 and proofreading throughout this article. Fig. 5 is cited from [22] as an example with the permission of the authors. Section 3 and a part of Section 5 are based on the unpublished technical report [24], of which I am the lead author. The technical report [24] was based on my presentation at a BoF of SC 2019, and the web link is <https://www.benchcouncil.org/file/BenchCouncil-SC-BoF.pdf>. After the presentation, I drafted the technical report [24] as the first author. I am very grateful to the other authors for their discussions and contributions: Dr. Lei Wang, Dr. Wanling Gao, and Dr. Rui Ren.

## References

- [1] SPEC, SPEC CPU 2017 Benchmark, 2017, <https://www.spec.org/cpu2017/>.
- [2] R. Panda, S. Song, J. Dean, L.K. John, Wait of a decade: Did SPEC CPU 2017 broaden the performance horizon? in: 2018 IEEE International Symposium on High Performance Computer Architecture, HPCA, IEEE, 2018, pp. 271–282.
- [3] J. Deng, W. Dong, R. Socher, L.J. Li, K. Li, L. Fei-Fei, Imagenet: A large-scale hierarchical image database, in: 2009 IEEE Conference on Computer Vision and Pattern Recognition, IEEE, 2009, pp. 248–255.
- [4] A.E. Johnson, T.J. Pollard, L. Shen, L.W.H. Lehman, M. Feng, M. Ghassemi, B. Moody, P. Szolovits, L. Anthony Celi, R.G. Mark, MIMIC-III, a freely accessible critical care database, *Sci. Data* 3 (1) (2016) 1–9.
- [5] M. Zairi, P. Leonard, Origins of benchmarking and its meaning, in: *Practical Benchmarking: The Complete Guide*, Springer, 1996, pp. 22–27.
- [6] J. Zhan, Call for establishing benchmark science and engineering, *BenchCouncil Trans. Benchmarks Stand. Eval.* 1 (1) (2021) 100012.
- [7] I. BiPM, I. IFCC, I. IUPAC, O. ISO, The international vocabulary of metrology—basic and general concepts and associated terms (VIM), 2012, p. 2012, JCGM 200.
- [8] R.N. Kacker, On quantity, value, unit, and other terms in the JCGM international vocabulary of metrology, *Meas. Sci. Technol.* 32 (12) (2021) 125015.
- [9] J. Zhan, Three laws of technology rise or fall, *BenchCouncil Trans. Benchmarks Stand. Eval.* 2 (1) (2022) 100034.



- [10] J. Zhan, Open-source computer systems initiative: The motivation, essence, challenges, and methodology, *BenchCouncil Trans. Benchmarks Stand. Eval.* 2 (1) (2022) 100038.
- [11] ComputerCouncil, The IoTs, edges, datacenter and networks as a computer: Building open-source planet-scale computers (PSC) for emerging and future computing, 2022, <https://www.computercouncil.org/PSC>.
- [12] Y. Li, J. Zhan, SAI-Bench: Benchmarking AI for science, *BenchCouncil Trans. Benchmarks Stand. Eval.* (2022).
- [13] ComputerCouncil, MetaverseBench: Instantiating and quantifying metaverse problems, benchmarks, and challenges, 2022, <https://www.computercouncil.org/MetaverseBench>.
- [14] V.V. Williams, Multiplying matrices faster than coppersmith-winograd, in: *Proceedings of the Forty-Fourth Annual ACM Symposium on Theory of Computing*, 2012, pp. 887–898.
- [15] D. Blalock, J. Guttat, Multiplying matrices without multiplying, in: *International Conference on Machine Learning*, PMLR, 2021, pp. 992–1004.
- [16] Z. Jiang, L. Wang, X. Xiong, W. Gao, C. Luo, F. Tang, C. Lan, H. Li, J. Zhan, Hpc ai500: The methodology, tools, roofline performance models, and metrics for benchmarking hpc ai systems, 2020, *arXiv preprint arXiv:2007.00279*.
- [17] P. Goyal, P. Dollár, R. Girshick, P. Noordhuis, L. Wesolowski, A. Kyrola, A. Tulloch, Y. Jia, K. He, Accurate, large minibatch sgd: Training imagenet in 1 hour, 2017, *arXiv preprint arXiv:1706.02677*.
- [18] Y. You, Z. Zhang, C.J. Hsieh, J. Demmel, K. Keutzer, Imagenet training in minutes, in: *Proceedings of the 47th International Conference on Parallel Processing*, 2018, pp. 1–10.
- [19] P. Mattson, C. Cheng, G. Diamos, C. Coleman, P. Micikevicius, D. Patterson, H. Tang, G.Y. Wei, P. Bailis, V. Bittorf, et al., Mlperf training benchmark, *Proc. Mach. Learn. Syst.* 2 (2020) 336–349.
- [20] F. Tang, W. Gao, J. Zhan, C. Lan, X. Wen, L. Wang, C. Luo, Z. Cao, X. Xiong, Z. Jiang, et al., Aibench training: balanced industry-standard ai training benchmarking, in: *2021 IEEE International Symposium on Performance Analysis of Systems and Software, ISPASS, IEEE*, 2021, pp. 24–35.
- [21] L. Wang, X. Xiong, J. Zhan, W. Gao, X. Wen, G. Kang, F. Tang, Wpc: Whole-picture workload characterization across intermediate representation, isa, and microarchitecture, *IEEE Comput. Archit. Lett.* 20 (2) (2021) 86–89.
- [22] L. Wang, K. Yang, W. Gao, C. Luo, C. Wang, Z. Ge, L. Zhang, F. Zhang, J. Zhan, WPC: Whole-picture Workload Characterization, Technical Report, Institute of Computing Technology, Chinese Academy of Sciences, 2022.
- [23] S. Matsuoka, J. Domke, M. Wahib, A. Drozd, R. Bair, A.A. Chien, J.S. Vetter, J. Shalf, Preparing for the future—rethinking proxy apps, 2022, *arXiv preprint arXiv:2204.07336*.
- [24] J. Zhan, L. Wang, W. Gao, R. Ren, BenchCouncil's view on benchmarking ai and other emerging workloads, 2019, *arXiv preprint arXiv:1912.00572*.
- [25] TPC, TPC-C Benchmarks, <http://www.tpc.org/tpcc/>.
- [26] TPC, TPC-W Benchmarks, <http://www.tpc.org/tpcw/>.
- [27] A. Petitet, HPL—a portable implementation of the high-performance linpack benchmark for distributed-memory computers, 2004, <http://www.netlib.org/benchmark/hpl/>.
- [28] V.J. Reddi, C. Cheng, D. Kanter, P. Mattson, G. Schmuelling, C.J. Wu, B. Anderson, M. Breughe, M. Charlebois, W. Chou, et al., Mlperf inference benchmark, in: *2020 ACM/IEEE 47th Annual International Symposium on Computer Architecture, ISCA, IEEE*, 2020, pp. 446–459.
- [29] W. Gao, F. Tang, L. Wang, J. Zhan, C. Lan, C. Luo, Y. Huang, C. Zheng, J. Dai, Z. Cao, et al., AIBench: an industry standard internet service AI benchmark suite, 2019, *arXiv preprint arXiv:1908.08998*.
- [30] L.A. Barroso, U. Höglze, The datacenter as a computer: An introduction to the design of warehouse-scale machines, *Synth. Lect. Comput. Archit.* 4 (1) (2009) 1–108.
- [31] W. Gao, J. Zhan, L. Wang, C. Luo, D. Zheng, F. Tang, B. Xie, C. Zheng, X. Wen, X. He, et al., Data motifs: A lens towards fully understanding big data and ai workloads, in: *Proceedings of the 27th International Conference on Parallel Architectures and Compilation Techniques*, 2018, pp. 1–14.
- [32] J. Dean, L.A. Barroso, The tail at scale, *Commun. ACM* 56 (2) (2013) 74–80.
- [33] W. Gao, F. Tang, J. Zhan, X. Wen, L. Wang, Z. Cao, C. Lan, C. Luo, X. Liu, Z. Jiang, Aibench scenario: Scenario-distilling ai benchmarking, in: *2021 30th International Conference on Parallel Architectures and Compilation Techniques, PACT, IEEE*, 2021, pp. 142–158.
- [34] A.M. Turing, I.—COMPUTING MACHINERY AND INTELLIGENCE, *Mind* LIX (236) (1950) 433–460.
- [35] D.H. Bailey, Nas parallel benchmarks, in: *Encyclopedia of Parallel Computing*, 2011, pp. 1254–1259.
- [36] J. Shun, G.E. Blelloch, J.T. Fineman, P.B. Gibbons, A. Kyrola, H.V. Simhadri, K. Tangwongsan, Brief announcement: The problem based benchmark suite, in: *Proceedings of the Twenty-Fourth Annual ACM Symposium on Parallelism in Algorithms and Architectures*, 2012, pp. 68–70.
- [37] J.L. Henning, SPEC CPU2000: Measuring CPU performance in the new millennium, *Computer* 33 (7) (2000) 28–35.
- [38] J.J. Dongarra, P. Luszczek, A. Petitet, The LINPACK benchmark: past, present and future, *Concurr. Comput.: Pract. Exper.* 15 (9) (2003) 803–820.
- [39] L. Wang, J. Zhan, C. Luo, Y. Zhu, Q. Yang, Y. He, W. Gao, Z. Jia, Y. Shi, S. Zhang, et al., Bigdatabench: A big data benchmark suite from internet services, in: *2014 IEEE 20th International Symposium on High Performance Computer Architecture, HPCA, IEEE*, 2014, pp. 488–499.
- [40] A. Ghazal, T. Rabl, M. Hu, F. Raab, M. Poess, A. Crolotte, H.A. Jacobsen, Bigbench: Towards an industry standard benchmark for big data analytics, in: *Proceedings of the 2013 ACM SIGMOD International Conference on Management of Data*, 2013, pp. 1197–1208.
- [41] MIT, AutoML benchmark datasets, 2021, [EB/OL]. [https://openml.github.io/automlbenchmark/benchmark\\_datasets.html](https://openml.github.io/automlbenchmark/benchmark_datasets.html). (Accessed 2 December 2021).
- [42] IOSCO, Financial Benchmarks, Technical Report, IOSCO, 2013.
- [43] C.D. Manning, Human language understanding & reasoning, *Daedalus* 151 (2) (2022) 127–138.
- [44] R.M. French, The Turing Test: the first 50 years, *Trends Cogn. Sci.* 4 (3) (2000) 115–122.
- [45] C.H. Hoffmann, Is AI intelligent? An assessment of artificial intelligence, 70 years after turing, *Technol. Soc.* (2022) 101893.
- [46] Britannica, instinct, 2022, <https://www.britannica.com/topic/instinct>.
- [47] T. Tomesh, P. Gokhale, V. Omole, G.S. Ravi, K.N. Smith, J. Viszlai, X.C. Wu, N. Hardavellas, M.R. Martonosi, F.T. Chong, Supermarq: A scalable quantum benchmark suite, 2022, *arXiv preprint arXiv:2202.11045*.
- [48] K. Asanovic, R. Bodik, B.C. Catanzaro, J.J. Gebis, P. Husbands, K. Keutzer, D.A. Patterson, W.L. Plishker, J. Shalf, S.W. Williams, et al., The landscape of parallel computing research: A view from berkeley, 2006.
- [49] E.F. Codd, A relational model of data for large shared data banks, in: *Software Pioneers*, Springer, 2002, pp. 263–294.
- [50] J. Hennessy, D. Patterson, A new golden age for computer architecture: Domain-specific hardware/software co-design, enhanced, in: *ACM/IEEE 45th Annual International Symposium on Computer Architecture, ISCA*, 2018.
- [51] BenchCouncil, BENCHCPU, 2019, <https://www.benchcouncil.org/benchcpu/index.html>.
- [52] ComputerCouncil, International open-source computer council, 2022, <https://www.computercouncil.org/>.



**Dr. Jianfeng Zhan** is a Full Professor at Institute of Computing Technology (ICT), Chinese Academy of Sciences (CAS), and University of Chinese Academy of Sciences (UCAS), the director of Research Center for Advanced Computer Systems, ICT, CAS. He received his B.E. in Civil Engineering and MSc in Solid Mechanics from Southwest Jiaotong University in 1996 and 1999, and his Ph.D. in Computer Science from Institute of Software, CAS, and UCAS in 2002. His research areas span from Chips, Systems to Benchmarks. A common thread is benchmarking, designing, implementing, and optimizing a diversity of systems. He has made substantial and effective efforts to transfer his academic research into advanced technology to impact general-purpose production systems. Several technical innovations and research results, including 35 patents, from his team, have been adopted in benchmarks, operating systems, and cluster and cloud system software with direct contributions to advancing the parallel and distributed systems in China or even in the world. He has supervised over ninety graduate students, post-doctors, and engineers in the past two decades. Dr. Jianfeng Zhan founds and chairs BenchCouncil and serves as the Co-EIC of TBench with Prof. Tony Hey. He has served as IEEE TPDS Associate Editor since 2018. He received the second-class Chinese National Technology Promotion Prize in 2006, the Distinguished Achievement Award of the Chinese Academy of Sciences in 2005, and the IISWC Best paper award in 2013, respectively.