

Research article

Evaluatology: The science and engineering of evaluation

Jianfeng Zhan^{a,b,c,*}, Lei Wang^{b,a,c}, Wanling Gao^{b,a,c}, Hongxiao Li^{b,c}, Chenxi Wang^{b,c},
Yunyou Huang^{d,a}, Yatao Li^f, Zhengxin Yang^{b,c}, Guoxin Kang^{b,a,c}, Chunjie Luo^{b,a,c},
Hainan Ye^{g,b,c}, Shaopeng Dai^c, Zhifei Zhang^{e,a}

^a The International Open Benchmark Council, Delaware, USA^b ICT, Chinese Academy of Sciences, Beijing, China^c University of Chinese Academy of Sciences, Beijing, China^d Guangxi Normal University, Guilin, Guangxi, China^e Capital Medical University, Beijing, China^f Microsoft Research Asia, Beijing, China^g Hong Kong International Evaluation and Benchmark Research, Hong Kong, China

ARTICLE INFO

Keywords:

Evaluation

Benchmark

Scale

Index

Evaluation condition

Evaluation model

Evaluation system

Evaluation standard

Equivalent evaluation condition

Least equivalent evaluation condition

Evaluatology

Benchmarkology

ABSTRACT

Evaluation is a crucial aspect of human existence and plays a vital role in each field. However, it is often approached in an empirical and ad-hoc manner, lacking consensus on universal concepts, terminologies, theories, and methodologies. This lack of agreement has significant consequences. This article aims to formally introduce the discipline of evaluatology, which encompasses the science and engineering of evaluation. We propose a universal framework for evaluation, encompassing concepts, terminologies, theories, and methodologies that can be applied across various disciplines, if not all disciplines.

Our research reveals that the essence of evaluation lies in conducting experiments that intentionally apply a well-defined evaluation condition to individuals or systems under scrutiny, which we refer to as the *subjects*. This process allows for the creation of an evaluation system or model. By measuring and/or testing this evaluation system or model, we can infer the impact of different subjects. Derived from the essence of evaluation, we propose five axioms focusing on key aspects of evaluation outcomes as the foundational evaluation theory. These axioms serve as the bedrock upon which we build universal evaluation theories and methodologies. When evaluating a single subject, it is crucial to create evaluation conditions with different levels of equivalency. By applying these conditions to diverse subjects, we can establish reference evaluation models. These models allow us to alter a single independent variable at a time while keeping all other variables as controls. When evaluating complex scenarios, the key lies in establishing a series of evaluation models that maintain transitivity. Building upon the science of evaluation, we propose a formal definition of a benchmark as a simplified and sampled evaluation condition that guarantees different levels of equivalency. This concept serves as the cornerstone for a universal benchmark-based engineering approach to evaluation across various disciplines, which we refer to as benchmarkology.

1. Introduction

Evaluation, a fundamental and significant undertaking in human existence, possesses a multifaceted nature. It spans a wide spectrum of domains, encompassing the assessment of computer performance, the evaluation of societal interventions to determine their efficacy [1], the ranking of educational institutions, and even the appraisal of political leaders through electoral processes. As a result, evaluation assumes a pivotal role that permeates every discipline. Nevertheless, it is pertinent

to recognize that evaluation practices often adopt ad-hoc and empirical approaches, displaying inherent variations among various disciplines.

Collectively, evaluations within diverse disciplines lack universal concepts, terminologies, theories, and methodologies. In the disciplines of computer science, social sciences, and psychology, the communities develop different methodologies that design experiments to deliberately impose conditions on individuals or systems under scrutiny, which we refer to as the *subject*, to measure and analyze their responses [2]. In the field of computer science, a *benchmark* is utilized as a tool

* Corresponding author.

E-mail address: jianfengzhan.benchcouncil@gmail.com (J. Zhan).URL: <https://www.zhanjianfeng.org> (J. Zhan).<https://doi.org/10.1016/j.tbench.2024.100162>

Received 19 March 2024; Accepted 22 April 2024

Available online 30 April 2024

2772-4859/© 2024 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/>).

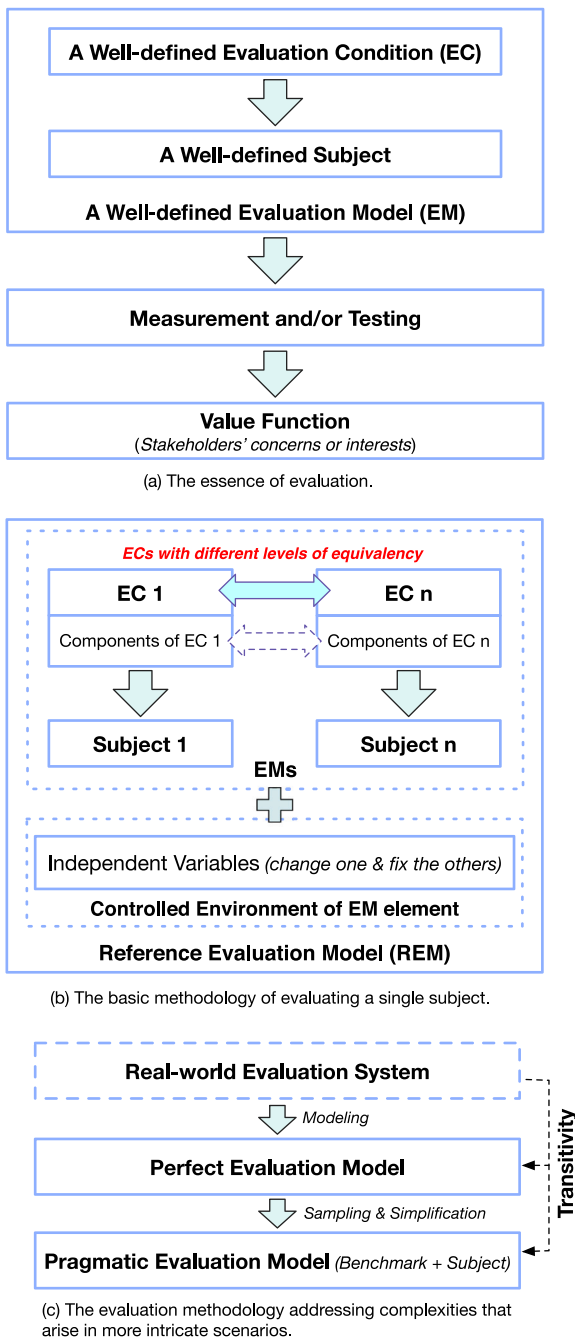


Fig. 1. The universal concepts, theories, and methodologies in evaluatology.

and methodology [3–6] to evaluate the effectiveness and efficiency of system designs and implementations. In the realm of social sciences, evaluation assumes the application of social research methodologies to systematically investigate the effectiveness and efficiency of intervention programs aimed at enhancing societal conditions, as defined by Rossi et al. [1]. Within the psychology domain, social and personality psychologists often employ *scales* such as psychological inventories, tests, or questionnaires [7] to quantify psychometric variables [7].

Conversely, evaluations in the business, finance, and education domains take different observational study methodologies [2]. The field of business embraces the concept of *benchmarking* as a means to identify exemplary practices that can propel companies towards superior performance [8]. In the realms of finance and education, evaluation often utilizes a tool named *index*. Widely employed to gauge

the overall performance of the system, the index is derived through the meticulous calculation of the weighted average, utilizing a select group of representative individuals [9].

Discussions concerning the true essence of evaluation are seldom found, which often results in confusion with measurement and testing and lacks clear differentiation. A critical consequence of this absence is the lack of previous endeavors to establish universally applicable foundational evaluation principles and methodologies that cut across diverse disciplines, ultimately giving rise to significant ramifications.

Even within computer sciences and engineering, it is not uncommon for evaluators to generate greatly divergent evaluation outcomes for the same subject. These discrepancies can range from significant variations to the extent of yielding contradictory qualitative conclusions. An example of this phenomenon can be observed when using multiple widely recognized CPU benchmark suites to assess the performance of the same processor. This often leads to greatly divergent evaluation outcomes that are *incomparable* across different benchmarks. Such circumstances give rise to valid concerns surrounding the reliability, effectiveness, and efficiency of these approaches when appraising the subject that is critical to safety, missions, and businesses. Further details on this issue can be found in Section 5.

To the best of our knowledge, this article, for the first time, formally introduces the discipline of evaluatology, encompassing the science and engineering of evaluation. We present an all-encompassing concept, terminology, theory, and methodology framework for evaluation that can be universally applied across diverse disciplines if not all disciplines.

We highlight that the essence of evaluation lies in conducting deliberate experiments where a well-defined Evaluation Condition (EC) is applied to a well-defined subject. The purpose is to establish a well-defined Evaluation Model (EM). By measuring and/or testing this EM, we can then infer the impacts of the subjects being evaluated. Derived from the core essence of evaluation, we present five axioms as the foundational principles of evaluation theory. The five axioms focus on key aspects of evaluation outcomes, including true quantity (The first and second axioms), traceability of discrepancy (the third axiom), comparability (the fourth axiom), and estimate (the fifth axiom).

Based on the five evaluation axioms, we present the universal evaluation theories and methodologies from two distinct dimensions: evaluating a single subject and complex scenarios.

A well-defined EC serves as a prerequisite for meaningful comparisons and analyses of the subjects. We propose a universal hierarchical definition of an EC and identify five primary components of an EC from the top to the bottom.

In the process of evaluating subjects, it is of utmost importance to prioritize the use of the equivalent ECs (EECs) across diverse subjects. This means that in order to establish two EECs, it is crucial to ensure that the corresponding components within the same layer of the two ECs are equivalent. By maintaining equivalency at each layer, we can ensure fair and unbiased evaluations, enabling meaningful comparisons and assessments between different subjects.

In certain cases, achieving complete equivalence between two ECs at all levels can be a challenging or even unattainable task. In such cases, we propose a minimum requirement of ensuring uniformity in the most essential components of the two ECs, which we refer to as the least equivalent evaluation conditions (LEECs). To establish the LEECs, we identify the most governing component within an EC that must exhibit equivalency. This component, known as the evaluation standard, plays a crucial role in defining the LEECs.

We apply ECs with different levels of equivalency to diverse subjects to constitute EMs. An EM element refers to a specific point within the EM state space, and each EM element may have many independent variables. To eliminate confounding, we propose a new concept named a reference evaluation model (REM). An REM mandates that each element of an EM change only one independent variable at a time while keeping the other independent variables as controls. Subsequently, we utilize the measurement and/or testing to gauge the functioning of the

REM. Finally, from the amassed measurement and testing data of the evaluation systems, we then deduce the cause–effect impacts of the different subjects.

Addressing the complexities that arise in more intricate scenarios, we reveal that the key to effective and efficient evaluations in various complex scenarios lies in the establishment of a series of EMs that maintain transitivity.

In real-world settings, we refer to the entire population of real-world systems that are used to evaluate specific subjects as the *real-world evaluation system (ES)*. Assuming no safety concerns are present, the real-world ES serves as a prime candidate for creating an optimal evaluation environment, enabling the assessment of diverse subjects. However, there are several significant obstacles to consider, i.e., the presence of numerous confounding, the challenges of establishing an REM, prohibitive evaluation costs resulting from the huge state spaces, multiple irrelevant concurrent problems or tasks taking place, and the inclination to exhibit bias towards certain clusters within the EC state space.

We posit the existence of a *perfect EM* that replicates the real-world ES with utmost fidelity. A perfect EM eliminates irrelevant problems or tasks, has the capability to thoroughly explore and comprehend the entire spectrum of possibilities of an EC, and facilitates the establishment of REMs.

However, the perfect EM possesses huge state space, entails a vast number of independent variables, and hence results in prohibitive evaluation costs. To address this challenge, it is crucial to propose a pragmatic EM that simplifies the perfect EM in two ways: reducing the number of independent variables that have negligible effect and sampling from the extensive state space. A pragmatic EM provides a means to estimate the parameters of the real-world ES.

We put forth four fundamental issues in the evaluations and formally formulate the problems mathematically: ensure the transitivity of EMs; perform a cost-efficient evaluation with controlled discrepancies; ensure the evaluation traceability; connect and correlate evaluation standards across every discipline.

Building upon the science of evaluation, we formally define a benchmark as a simplified and sampled EC, specifically a pragmatic EC, that ensures different levels of equivalency. Based on this concept, we propose a benchmark-based universal engineering of evaluation across different disciplines, which we aptly term “benchmarkology”. Fig. 1 presents the universal concepts, theories, and methodologies in Evaluatology.

The article is structured as follows: Section 2 elucidates the background. Section 3 introduces a comprehensive theoretical and methodological framework for evaluatology. Section 4 outlines the principles and methodologies of benchmarkology. Section 5 reviews the state-of-the-art and state-of-the-practice evaluations and benchmarks and expounds upon the imperative to cultivate the science and engineering of evaluation. Ultimately, Section 6 manifests the overarching conclusion.

2. Background

This section provides an overview of the background.

2.1. Basic concepts

This subsection presents several concepts, like individual, systems, populations, samples, variables, models, confounding, and control based on several undergraduate textbooks [2,10,11].

An *individual* can be defined as the object that is described by a given set of data. A *system* is an interacting or interdependent group of individuals, whether of the same or different kinds, forming a unified whole [12,13]. A system could be a recursive structure. That is to say, a high-level system could consist of an interacting or interdependent

group of low-level systems, whether of the same or different kinds, forming a unified whole.

A *population* is the entire group of individuals or systems we wish to study and understand, while a *sample* represents a smaller subset of individuals or systems from the population [2]. A *variable* or *quantity* is any *property* of an individual or system. A *parameter* is a number that describes some property of the population, while a *statistic* is a number that describes some property of a sample. *Inference* is the process of drawing conclusions about a parameter of a population on the basis of the statistic of sample data [2].

According to [11], a function, denoted as f , is a rule that assigns a unique element, referred to as $f(x)$, from a set R to each element in a set D . In this context, the domain, denoted as D , refers to the set of all possible values for which the function is defined. On the other hand, the range of the function, denoted as $f(x)$, consists of all the possible values that $f(x)$ can take as x varies within the domain. The *independent variable* is represented by a symbol that encompasses any arbitrary number within the domain of the function. A *dependent variable*, represented by a symbol, is used to denote a number within the range of the function.

A *model* is a simplified version of a system that would be too complicated to analyze in full detail [10]. A model could be a physical model or a *mathematical model*. A mathematical model is a mathematical description, typically through *functions* or equations, of a system, the purpose of which is to understand the system and to make predictions about its behavior [11]. Throughout the remainder of this article, we will use the terms “system” and “model” interchangeably unless explicitly stated otherwise.

An *observational study* observes individuals or systems and measures variables of interest without any attempt to influence their responses, while an *experiment* is designed to deliberately impose conditions on individuals or systems to measure and analyze their responses [2].

In the realm of understanding cause and effect, it is crucial to rely on experiments rather than observational studies. Even if an observational study is based on a random sample, it still falls short in effectively measuring the impact of changes in one variable on another variable [2]. Experiments, on the other hand, provide us with compelling and conclusive data, making them the only source that truly convinces us of cause-and-effect relationships.

Confounding arises when two independent variables are associated in a manner that makes it challenging to differentiate their specific effects on a dependent variable. In other words, the influence of these independent variables becomes entangled, making it difficult to attribute specific impacts to each one. In such cases, the independent variable responsible for this confounding effect is referred to as a *confounding variable*. *Control* means keeping other independent variables that might affect the response the same [2], and the main purpose of a *control group* is to provide a baseline for comparing the effects of the other treatments.

2.2. Metrology

Metrology is the science of measurement and its applications [14]. In this section, we present a simplified yet systematic framework for understanding metrology concepts based on the works of [14,15]. To maintain conciseness, we focus only on the essential metrology concepts (see Fig. 2).

The essence of metrology lies in quantities and their corresponding measurements. A *quantity* is a property whose instances can be compared by ratio or only by order [14]. Furthermore, Psychologist Stanley Smith Stevens developed a well-known measurement classification with four levels based on empirical operations, mathematical group structure, and permissible statistics (invariant): nominal, ordinal (based on order [14]), interval, and ratio [16].¹

¹ In the original article, Stanley Smith Stevens used the term “levels or scales of measurement”. We have only used “levels” to avoid confusion with the specific meaning of “scales” in psychology

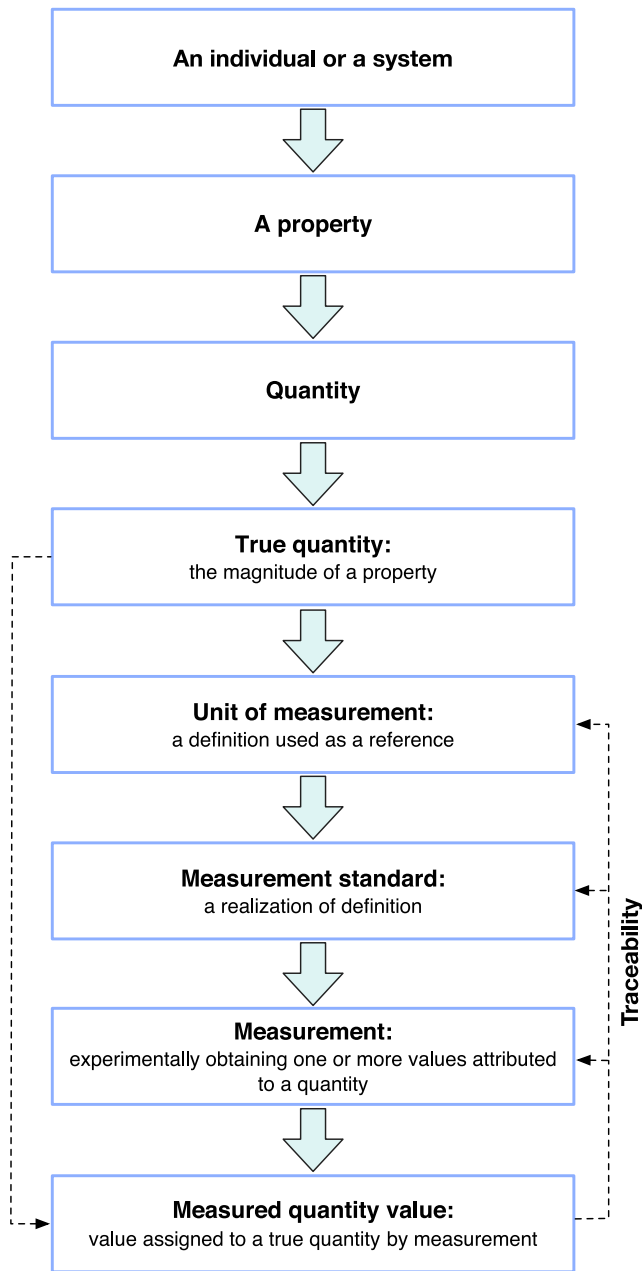


Fig. 2. A simplified yet systematic conceptual framework for metrology [14,15].

A nominal level is the most basic form of measurement, where numbers are used as labels or type numbers to establish an equality relation. An ordinal level, on the other hand, involves ranking the items in a particular order. An interval level exhibits an equality of interval relation, where (1) the choice of a zero point is a matter of convention or convenience, (2) there is rank ordering, and (3) the scale remains invariant when a constant is added to all values, preserving the differences between them. A ratio level allows for all four types of relations: equality, rank-ordering, equality of intervals, and equality of ratios.

The international system of metrology encompasses seven fundamental quantities: time, length, mass, electric current, thermodynamic temperature, amount of substance, and luminous intensity [14].

Consistent with the definition of a quantity, the true value of a quantity represents the magnitude of a property or characteristic of an individual or system, e.g., a phenomenon, body, or substance that is

independent of any observer. For example, it can be a specific circle’s radius or a particular particle’s kinetic energy within a given system [14,15]. For measurement, the true quantity value is an unknown measurement target [14].

In the field of measurement, the *unit of measurement* [15] plays a crucial role. It is a real scalar magnitude that is defined and adopted by convention. Its purpose is to allow for the comparison of quantities of the same kind.

Measurement standard [15] is a realization of the definition of quality. It is characterized by a stated metric value and an associated measurement uncertainty.

To establish a measurement standard, it is important to use a *measurement methodology* that is both repeatable (performed by the same team) and reproducible (performed by different teams). This ensures consistency and reliability in the reference for measurements. Such measurements can be conducted using *measuring instruments* or *measuring systems* [14], providing a reliable foundation for further analysis and comparison.

Measurement is experimentally obtaining one or more values attributed to a quantity and other relevant information [14]. Another widespread definition of measurement in the social sciences is “the assignment of numerals to objects or events according to some rule.” [16], dating back to 1946. Quantity values obtained by the measurement are measured (quantity) values, representing the measurement results [14].

The hierarchy of measurement standards follows a progression from lower to upper levels, with increasing accuracy and cost. This progression starts from national measurement standards and extends to international standards. As a property of a measurement result, *measurement traceability* [14] establishes a connection between the result and a reference (measurement standards, measuring instruments, and measuring systems). This connection is established through a documented, unbroken chain of calibrations, with each calibration contributing to the measurement uncertainty. To ensure accuracy, each level of measurement standards in the hierarchy should be calibrated using a higher standard with greater precision.

2.3. Testing

A *test oracle* is a method used to verify whether an individual or system being tested has performed correctly during a specific execution [17]. Test oracles can include, but are not limited to, specifications, contracts, reference products, previous versions of the same product, and relevant performance or quality of service criteria [18].

As shown in Fig. 3, *testing* is the process of executing an individual or system to determine whether it (1) conforms to the specified behavior defined by the test oracles [19] (the first category) and/or (2) operates correctly within its intended environment as defined by the test oracles (the second category).

In the first category of testing, the test oracle compares the actual output with the specified output to identify incorrect behavior, which is considered a *failure* [19]. Another type of failure is often encountered in the second category of testing, where an individual or system fails to meet environmental constraints or falls outside the specified requirements. Examples of such failures include running out of memory, slow execution, and incompatibility with operating systems [17]. It is important to note that these two types of failures are not isolated incidents. Failures in the second category can lead to failures in the first category, such as running out of memory, which results in incorrect program execution.

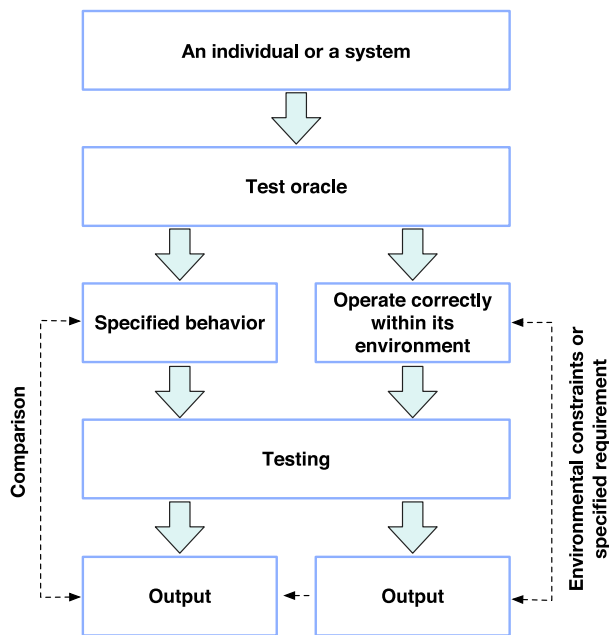


Fig. 3. A simplified yet systematic conceptual framework for testing [17,19].

3. The science of evaluation

Within this section and the forthcoming one, we first propose the universal evaluation concepts and terminologies. Then, we elucidate the theory and methodology governing the realm of evaluation, collectively referred to as evaluatology.

We present the science of evaluation from two distinct perspectives: evaluating a single subject and evaluating complex scenarios.

We postpone the exposition of the principles and methodology pertaining to benchmarkology, the benchmark-based engineering of evaluation, to Section 4.

3.1. Basic evaluation concepts and terminology

In this subsection, we will examine an illustrative case study in evaluation and meticulously analyze the fundamental *components* inherent in an evaluation. For the sake of convenience, henceforth, within this article, each case study shall be assigned a unique case ID for differentiation purposes.

In the first case, which we will call Case One, an organization is in the process of acquiring a computer. To make an informed decision, the organization decides to evaluate various computer options by executing its applications on each one. During this evaluation, the organization will collect extensive data on performance and energy efficiency.

Based on this data, the organization will then formulate an explicit or implicit function to express its preferences for different computers. This function will serve as a way to quantify and articulate the organization's priorities and requirements in terms of performance and energy efficiency. By doing so, the organization can make a well-informed decision and choose the computer that best aligns with its needs and preferences.

In any evaluation process, we refer to an individual or a system under scrutiny as a *subject*. In the context of Case One, the computers that are being considered for evaluation are the subjects.

Another important component of an evaluation is the presence of *stakeholders*. A stakeholder is defined as an entity that holds a stake of responsibility or interest in the subject matter. In this case, the stakeholders involved in the procurement of computers include the organization seeking to acquire the computers, the specific users, the

designers responsible for creating the computer specifications, and the producers who manufacture the computers.

Each subject, in this case, the computer, has its own set of stakeholders who will render judgments based on the data collected from measurements. Metrology provides the foundational principles, methodology, and instruments for measurements within this process. It ensures that the measurements are accurate, reliable, and consistent, enabling the stakeholders to make informed decisions based on the gathered data.

In Case One, the prospect of measuring the mere attributes of a subject, such as its weight and power consumption, possesses a certain degree of utility. Nevertheless, such measurements fall considerably short of meeting stakeholders' evaluation requirements. The stakeholder seeks comprehension of the subject's effectiveness and efficiency when applying a specific condition or intervention to the subject. In this case, it is to execute the stakeholder's primary or forthcoming applications, which we informally label as an *evaluation condition (EC)*. The EC represents the third critical component of an evaluation, which shall be formally expounded upon subsequently.

Within this framework, an important question arises: How can organizations establish a framework to determine the preferences of distinct subjects when they exhibit varying levels of performance across different applications?

In the current state-of-the-practice, a more intuitive approach entails executing applications on computers sequentially. Subsequently, we proceed to measure the computers' performance when operating distinct applications individually. Following each execution, data is collected encompassing factors such as the duration of each application's execution and its corresponding energy consumption.

It is imperative to establish a function that can map the compiled measurement data to one or several composite evaluation metrics capable of capturing the stakeholders' concerns and interests. In the rest of this article, we refer to this function as a *value function*. Once the evaluation outcomes have been obtained, it becomes feasible to define a *reference subject* and its *reference evaluation outcome* against which the evaluation results of alternative subjects can be compared.

3.2. The essence of evaluation

From the aforementioned analysis in Section 3.1, it is evident that the challenge in evaluation arises from the inherent fact that evaluating a subject in isolation falls short of meeting the expectations of stakeholders. Instead, it is crucial to apply a well-defined EC that reflects the stakeholders' concerns or interests. By doing so, evaluation can be viewed as an intentional experiment that deliberately imposes a specific EC on the subject itself.

Based on the definitions provided in Section 2.1, when a subject is equipped with an EC, it forms an *evaluation system (ES)* or an *Evaluation Model (EM)*. An evaluation model (EM) is a simplified version of an ES that would be too complicated to analyze in full detail [10].

Based on the analysis presented earlier, it becomes clear that the *core essence of evaluation lies in conducting deliberate experiments where Equivalent ECs (EECs) are applied to a diverse range of subjects, resulting in the establishment of equivalent EMs. Subsequently, we can effectively evaluate the subjects by measuring the equivalent EMs.*

Therefore, we formally define evaluation as an *experiment that applies EECs to diverse subjects and establishes equivalent EMs, enabling the measurement of these equivalent EMs, the inference of the subjects' impact, and the subsequent judgment of them.*

3.3. Five evaluation axioms

In this section, we present five evaluation axioms that are derived from the core essence of evaluation, serving as the foundational principles of evaluation theory. They focus on key aspects of evaluation outcomes, including true quantity (The first and second axioms),

traceability of discrepancy (the third axiom), comparability (the fourth axiom), and estimate (the fifth axiom).

The First Axiom of Evaluation: The Axiom of the Essence of Composite Evaluation Metrics. This axiom declares that the essence of the composite evaluation metric either carries inherent physical significance or is solely dictated by the value function.

In nature, *a composite evaluation metric refers to a combined quantity that is constructed using base quantities and other quantities that possess physical significance*. If a composite evaluation metric does not carry inherent physical significance, the value function serves as a mechanism that maps base quantities and other quantities carrying physical meaning into a composite evaluation metric. The widespread acceptance of the composite evaluation metric relies on it being embraced by the community of evaluators.

The Second Axiom of Evaluation: The Axiom of True Evaluation Outcomes. This axiom declares that when a well-defined EC is applied to a well-defined subject, its evaluation outcomes, including its quantities and composite evaluation metrics, possess true values.

“The magnitude of a property of an individual phenomenon, body, or substance is associated with a true quantity” [14,15]. Additionally, each testing procedure yields a definitive outcome relative to its respective test oracle. Building upon this inference, it is reasonable to presume that when a well-defined subject is equipped with a well-defined EC, the quantities within the corresponding well-defined EM possess true values.

For a well-defined EM, each composite evaluation metric is derived from measurement and/or testing outcomes, utilizing a definite value function that translates the base quantities and other quantities into a composite evaluation metric. Consequently, the evaluation results are exclusively determined by the measurements and/or testing procedures employed. It is reasonable to assume that for a well-defined EM, its composite evaluation metric possesses a true value.

The first and second axioms are concerned with the true evaluation outcome. The first axiom provides the basis for defining value functions. With a well-defined value function, a well-defined EM possesses true quantities of evaluation outcomes.

The Third Axiom of Evaluation: The Axiom of Evaluation Traceability. This axiom declares that for the same subject, the divergence in the evaluation outcomes can be attributed to disparities in ECs, thereby establishing evaluation traceability. This axiom focuses on the traceability of discrepancies in the evaluation outcomes.

For the same subject, this axiom is deemed rational as disparities in evaluation outcomes can be rationalized as the consequence of variations in the ECs. In the absence of this axiom, the differences observed in evaluation outcomes would be inexplicable, contradicting our scientific and engineering intuitions.

The Fourth Axiom of Evaluation: The Axiom of Comparable Evaluation Outcomes. This axiom declares when each well-defined subject is equipped with EECs, their evaluation outcomes are comparable. It goes without saying this axiom is related to the comparability of the evaluation outcomes.

Only when each EC is well-defined, and two ECs achieve complete equivalence at all levels can we refer to them as EECs. When each well-defined subject is equipped with EECs, its evaluation outcomes possess true values. Additionally, when well-defined subjects are subjected to EECs, their evaluation outcomes accurately reflect the impacts of different subjects under the same conditions, making them comparable.

The Fifth Axiom of Evaluation: The Axiom of Consistent Evaluation Outcomes. This axiom asserts that when a well-defined subject is evaluated using different samples from a population of ECs, their evaluation outcomes consistently converge towards the true evaluation outcomes of the population of ECs. This axiom provides an estimate of the true evaluation outcomes under the population of ECs.

According to the Second Axiom of Evaluation, when a well-defined subject is equipped with a well-defined population of ECs, the resultant EM possesses the true evaluation outcomes. When a sample is taken

from a population of ECs, it serves as an approximation of the entire population. As a result, different samples yield consistent evaluation outcomes that gradually converge towards the evaluation outcomes of the entire population of ECs. This convergence is influenced by the sample's ability, which is determined by the chosen sampling policy, to represent the underlying characteristics of the population accurately.

3.4. Basic evaluation methodology

This section outlines the fundamental methodology for evaluating a single subject. Drawing upon the discussion of the essence of evaluation in Section 3.2, we propose a rigorous evaluation methodology to determine the impacts of the subjects as follows.

We create equivalent ECs (EECs) and apply EECs to diverse subjects to constitute equivalent EMs. An EM element refers to a specific point within the EM state space, and each EM element may have many independent variables. To eliminate confounding, we propose a new concept named a reference evaluation model (REM). An REM mandates that each element of an EM change only one independent variable at a time while keeping the other independent variables as controls. Subsequently, we utilize the measurement to gauge the functioning of the REM. Finally, from the amassed measurement data of the evaluation systems, we then deduce the cause-effect impacts of the different subjects.

In this methodology, we emphasize five essential steps to ensure a comprehensive evaluation, as shown in Fig. 4. These steps are crucial in accurately determining the impacts of the subjects.

The first step is to establish a rigorous definition of an EC. According to The Second and Three Axioms of Evaluation, when a well-defined subject is equipped with a well-defined EC, its evaluation outcomes possess true values; when the same subject is equipped with different ECs, any divergence in the evaluation outcomes can be attributed to disparities in these ECs. Therefore, the key focus in this phase is to clearly present a well-defined EC.

The second step involves the establishment of EECs. As outlined in the Fourth Axiom of Evaluation, when EECs are applied to diverse subjects, their evaluation outcomes become comparable. Since the primary objective of evaluation is to compare different subjects, the establishment of EECs becomes an essential step in the process.

The third step involves the elimination of confounding variables. Given that each element of an EM consists of multiple independent variables, it becomes essential to establish an REM. An REM serves as a controlled evaluation environment where only one independent variable is altered at a time while the other independent variables remain constant as the control. This approach helps in isolating the effects of individual variables and ensures a more accurate evaluation of the subject's performance.

The fourth step is to define the value functions that map the base quantities and other quantities to composite evaluation metrics that represent the stakeholders' primary concerns or interests. According to the first axiom of evaluation, the essence of the composite evaluation metric either carries inherent physical significance or is solely dictated by the value function. So, when defining a value function, it is crucial to make it become the consensus of the community.

Finally, we utilize the measurement to gauge the functioning of the REM. From the amassed measurement data of the REM, we then deduce the cause-effect impacts of the different subjects.

3.5. Basic evaluation theory

This subsection presents the basic evaluation theory, including the hierarchical definition of an EC, universal concepts across different disciplines, the establishment of EECs, LLECs, evaluation standards, and the establishment of an REM.

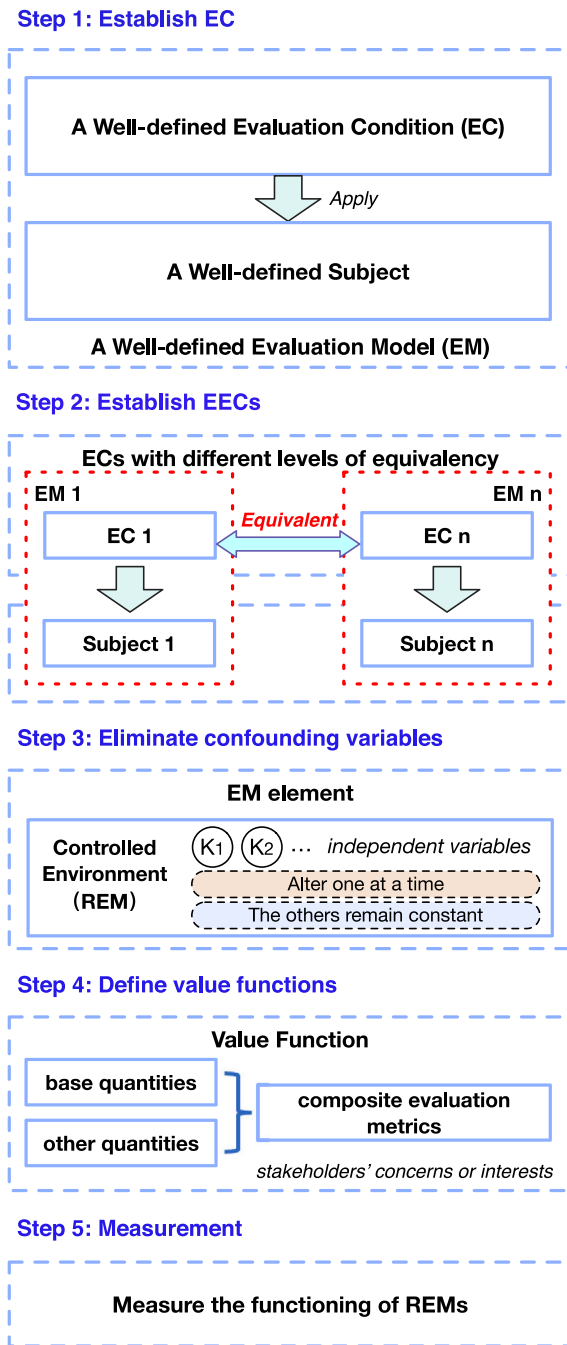


Fig. 4. Basic evaluation methodology.

3.5.1. The hierarchical definition of an EC

In the preceding subsection, we deduced that there exists solely one feasible approach to evaluation: applying the EECs to diverse subjects and establishing an REM. Regrettably, even within a rudimentary evaluation setting such as Case One, an EC has multifarious components.

To address the above challenges, in this subsection, we propose a hierarchy definition of an EC. We start defining an EC from the problems or task spaces that these stakeholders face and need to address with the following two reasons. First, the concerns and interests of the relevant stakeholders are at the core of the evaluation. These concerns and interests are best reflected through the problems or tasks they must face and resolve, which provide a reliable means to define an EC.

Second, utilizing the same problem or task can ensure the comparability of evaluation outcomes.

Taking Case One as an example, we observe that in Case One, an EC encompasses numerous constituents. Notably, we identify four primary components of an EC from the top to the bottom. The first top component is a set of equivalent definitions of problems or tasks. While the problem or task itself serves as the foundation for the evaluation process, it cannot solely serve as the evaluation itself because the problem or task is often abstract and requires further instantiation to determine its specific parameters. The second component is the set of a collective of equivalent problem or task instances, each of which is instantiated from the element of the first component. Different from the first component, an equivalent problem or task instance is specific and could serve as the evaluation directly.

After a problem or task instance is proposed, it is necessary to figure out a solution. The third component consists of the algorithms or algorithm-like mechanisms, each of which provides the solution to a specific problem or task instance. An algorithm-like mechanism refers to a process that operates in a manner similar to an algorithm. This term is proposed because, in numerous disciplines, such as social and biological sciences, it is not currently feasible to formulate mathematical algorithms explicitly. These domains often involve complex and nuanced phenomena that defy precise mathematical modeling.

The fourth component encompasses the implementation of an algorithm or instantiation of an algorithm-like mechanism. Its implementation or instantiation involves understanding the algorithm or the algorithm-like mechanism and implementing it in a specific system. This process ensures that the algorithm or the algorithm-like mechanism can effectively and efficiently solve the intended problem instance or perform the desired task instance within the given context.

In addition to the four components of an EC that we discussed above, other components can be involved in the other complex evaluation scenarios, which we will discuss later.

3.5.2. Universal concepts across different disciplines

In Section 3.1, we conducted an examination of the fundamental constituents of an evaluation, which encompass “subjects” and “evaluation conditions”. Additionally, we put forth definitions for various fundamental concepts, namely “subject”, “stakeholders”, and “value functions”. Through this analysis, we unveiled that the core nature of evaluation is to intentionally apply EECs to diverse subjects and establish an REM to infer the impact of the subjects for judgments.

However, given the multiplicity of evaluation scenarios, two critical questions must be addressed: (1) Do these concepts suffice for diverse scenarios? (2) Can we formulate a comprehensive and universally applicable conceptual framework? In this subsection, we delve deeper into various evaluation cases across diverse disciplines, with the primary aim of enhancing our comprehension of the evaluation process.

Evaluating an AI algorithm

In Case Two, the objective is to evaluate an AI algorithm, specifically focusing on an Image Classification task as a case study. Real-world images are gathered and annotated with accurate labels, such as a cat or dog. A portion of these images is randomly selected to construct the training, validation, and test datasets based on a designated percentage. To assess the image classification algorithm (the subject in this case), it must be implemented on a computer system utilizing a specific programming framework, such as PyTorch or TensorFlow.

During the evaluation process, the test data is provided to the algorithm, which generates an output. This output is then compared to the ground truth labels. In Case Two, the ground truth can be considered as a test oracle. Additionally, measurements are collected for each run of the evaluation process.

The fundamental evaluation process in Case Two bears a resemblance to that of the baseline case. However, there exist three notable distinctions. Firstly, an EC in Case Two differs. The presence of a dataset labeled with ground truth constitutes a vital aspect of an EC.

A dataset labeled with the ground truth represents a specific instance of a problem or task. Regrettably, in Case Two, it is impractical to define a problem or task mathematically. Hence, diverse problem or task instances are devised in an ad-hoc manner, for example, selecting images randomly to form training, validation, and test datasets based on a predetermined percentage. This approach may introduce biases.

Secondly, both the subject and the algorithms must be implemented on a computer, prompting the introduction of another concept called “a support system” to elucidate this facet, which may assume diverse manifestations in alternative evaluation scenarios. The support system represents an additional essential constituent of an EC. Therefore, in Case Two, we introduce two novel concepts: “the support system” and “the subject instantiation”. Fig. 5 shows a hierarchy definition of an EC.

Thirdly, in Case Two, upon feeding the algorithm with the test data, it generates an output that is then compared with the ground truth, also known as the test oracle. Apart from measurements, there exist other forms of activities in the evaluation process, namely testing, as expounded upon in Section 2.3. Hence, in this instance, we modify the essence and definition of evaluation, as delineated in Section 3.2, as follows. The essence of an evaluation lies in “conducting deliberate experiments where EECs are applied to a diverse range of subjects, resulting in the establishment of equivalent EMs. Subsequently, we can effectively evaluate the subjects by measuring and/or testing the equivalent EMs”. We formally define evaluation as “an experiment that applies EECs to diverse subjects and establishes equivalent EMs, enabling the measurement and/or testing of these equivalent EMs, the inference of the subjects’ impact, and the subsequent judgment of them”.

Drug and policy evaluations

The Third Case (Case Three) and the Fourth Case (Case Four) exhibit similarities, as they involve evaluating a drug and evaluating a policy aimed at addressing drug addiction within a community. In these cases, the subject refers to a specific drug or a policy aimed at addressing drug addiction intervention, while the support system encompasses the participants targeted by these interventions.

When comparing Cases Three and Four with Cases One and Two, we have made three observations. Firstly, the specific problem or task instances could be defined in a literal manner. Case Three focuses on the cure of a specific illness, whereas Case Four aims to address drug addiction and improve the overall well-being of individuals within a designated community. Regrettably, we currently lack a mathematical understanding of these problems or tasks, making it challenging to provide a mathematical definition of the abstract problem or task. Nevertheless, even if the problems can only be defined in a literal sense, having detailed and comprehensive definitions of problem or task instances, as well as a profound comprehension of the interrelationships between different problem or task instances (from biological or social perspectives), proves advantageous for the purposes of knowledge reuse and sharing. Furthermore, it is anticipated that in the future, we will strive to gain a deeper understanding of the connections between various diseases or social issues, potentially through mathematical means.

The second observation revolves around the distinctions of the support systems found in Cases Two, Three, and Four, specifically referring to the substantial variability in conditions within the target participants. Evaluations commonly utilize a methodology known as randomized controlled trials (RCT), which serves to eliminate confounding.

In practical application, Randomized Controlled Trials (RCTs) are widely recognized as the gold standard for conducting evaluations. In an RCT, subjects and support systems in the treatment group and control group are randomly assigned. This random assignment helps to minimize confounding variables that may arise from differences in support systems. Additionally, the allocation of participants to either

the treatment or control group is kept concealed from the evaluator and relevant stakeholders.

Lastly, when it comes to Cases Three and Four, the third component of an EC often lacks a mathematical-form algorithm, although it relies on a scientifically valid mechanism that includes biological, social, or psychological interactions. That is the right reason for us proposing the term “algorithm-like mechanism”, which we have explained before. Moreover, Case Four presents a significantly more intricate situation compared to Case Three, as the instantiation of the algorithm-like mechanism encounters challenges in maintaining consistent quality due to diverse factors, such as different attitudes towards interventions and varying levels of communication skills that hold considerable importance.

In summarizing, across various evaluation scenarios in different disciplines, such as Cases One, Two, Three, and Four, it is possible to develop a comprehensive conceptual framework that can be universally applied.

3.5.3. The establishment of EECs

To lay the groundwork for the formalization of an EC, it is essential to establish a clear understanding of some crucial notations. The notations E' , S , and U represent three crucial components. Specifically, E' represents the problem or task space. S represents the support system space. Finally, U represents the subject space. e' , s , u is an element of E' , S , and U respectively. We note $e' \in E'$, $s \in S$, $u \in U$.

In addition to the aforementioned notations, we also define several other fundamental notations. For each problem or task, $e'_i \in E'$, there is a set of problem or task instances noted as E_i . For all problems or tasks in E' , there is a collection of a set of problem or task instances, which we noted as $SE = (e'_i, E_i)$. We use the division notation SE/E' to denote E . E can be defined as the union of all E_i .

We introduce the notation SA' to represent the algorithm-like mechanism space. This space, denoted as SA' , consists of a set of algorithm-like mechanisms that are associated with each problem or task instance. Specifically, for a given problem or task e'_i in the problem or task space E' , and for each instance e_{ij} in the corresponding instance space E_i , we define a set of algorithm-like mechanisms as $SA' = (e'_i, e_{ij}, A'_{ij})$.

To represent the algorithm-like mechanism space SA' in relation to the problem or task space E' and the problem or task instance space E , we use the division notation $SA'/E'/E$ to denote A' . A' is a union of all A'_{ij} .

We introduce the notation SA to represent the instantiations of the algorithm-like mechanism space. This space, denoted as SA , consists of a set of instantiations of algorithm-like mechanisms that are associated with each problem or task instance, algorithm-like mechanism, and support system. Specifically, for a given problem or task e'_i in the problem or task space E' , for each instance e_{ij} in the corresponding instance space E_i , for each algorithm-like mechanism a'_{ijk} in the algorithm-like mechanism space A'_{ij} , and for each support system s_l in the support system space S , we define a set of instantiations of algorithm-like mechanisms as $SA = (e'_i, e_{ij}, a'_{ijk}, s_l, A_{ijkl})$.

To represent the instantiations of the algorithm-like mechanism space SA in relation to the problem or task space E' , the problem or task instance space E , the algorithm-like mechanism space A' , and the support system space S , we use the division notation $SA/E'/E/A'/S$ to denote A . A is a union of all A_{ijkl} .

By introducing these notations, we establish a comprehensive framework that allows us to delineate the various components of an EC and their respective roles. This formalization enhances our understanding of the key components and their relationships within the EC framework.

We can formalize an EC as $C = E' \times E \times A' \times A \times S$.

In the realm of EC spaces, the concept of EECs plays a significant role. Two EC spaces, denoted as C_1 and C_2 , are considered to be EECs if and only if there exists a bijection, denoted as β , between the two spaces: $\beta : C_1 \mapsto C_2; \beta^{-1} : C_2 \mapsto C_1$. This equivalence is denoted as $C_1 \sim C_2$.

Evaluation Condition (EC)	
Component 1	a set of equivalent definitions of problems or tasks (E')
Component 2	the set of a collective of equivalent problem or task instances (E)
Component 3	the algorithms or algorithm-like mechanisms (A')
Component 4	the implementations of algorithms or instantiations of algorithm-like mechanisms (A)
Component 5	support systems (S)

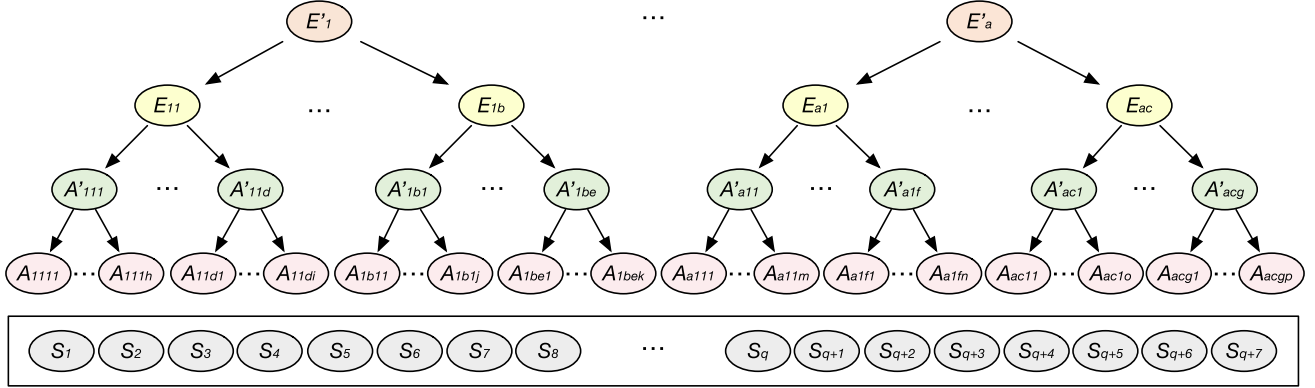


Fig. 5. The hierarchical definition of an EC.

3.5.4. LEECs and evaluation standard

In certain cases, establishing EECs can be a challenging task. It proves to be an arduous or unattainable task to ensure the equivalence of two ECs at all levels. However, it is crucial to ensure the relative comparability of evaluation outcomes.

In this section, we propose the concept of the least equivalent evaluation conditions (LEECs) as the foundation to attack this challenge. In the event where we are unable to guarantee or establish the equivalence of two ECs at all levels in the hierarchy, we propose a minimum requirement of ensuring uniformity in the most essential components of the two ECs, which we refer to as the least equivalent evaluation conditions (LEECs).

We propose the establishment of LEECs at the levels of the first and second top components of ECs. Firstly, the first and second top components of an EC serve as the foundation upon which the other two lower-level components are derived. Therefore, they provide the most primitive components when setting ECs. Secondly, to enable effective comparison of evaluation outcomes, it is crucial to establish equivalence between the first high-level components of two ECs. If the first high-level components are not equivalent, the evaluation outcomes cannot be compared reliably. Thirdly, relying solely on the first component may not provide enough specificity and certainty. To address this, in addition to the equivalence of the first high-level component, it is necessary for two ECs to possess two sets of definite and solvable problem or task instances that are equivalent.

In certain situations, it is possible to relax the requirement of strict equivalency between the second high-level component if the scale of the problem or task instances can be defined. In such cases, we can consider a scenario where two ECs have the same problem or task but differ in the scales of problem or task instance as LEECs.

We formally define LEECs as follows:

For two ECs, denoted as $C_1 = E'_1 \times E_1 \times A'_1 \times A_1 \times S_1$ and $C_2 = E'_2 \times E_2 \times A'_2 \times A_2 \times S_2$, if there is equivalence between their first two subspaces (E' and E), that is, if and only if there is a bijection, denoted

as $\hat{\beta}$, between $E'_1 \times E_1$ and $E'_2 \times E_2$, we can establish that they are LEECs, denoted as $C_1 \approx C_2$. In other words, $\hat{\beta}$ is a function mapping from $E'_1 \times E_1$ to $E'_2 \times E_2$, and its inverse function $\hat{\beta}^{-1}$, maps from $E'_2 \times E_2$ to $E'_1 \times E_1$, denoted as $\hat{\beta} : E'_1 \times E_1 \mapsto E'_2 \times E_2; \hat{\beta}^{-1} : E'_2 \times E_2 \mapsto E'_1 \times E_1$.

To effectively define the least equivalent ECs (LEECs), it is crucial to identify the most governing component of ECs that must exhibit equivalency. This component, known as the *evaluation standard*, plays a crucial role in defining the LEECs. By establishing and adhering to this evaluation standard, we can ensure that the evaluation outcomes are relatively comparable.

An evaluation standard should embody the characteristics that are solvable, definite, and equivalent (abbreviated as SDE). First, it should be amenable to a solvable framework, employing specific mechanisms. These mechanisms could encompass mathematical steps executed in a mechanical fashion [20] or incorporate biological, social, and psychological mechanisms and interactions. It is noteworthy that algorithms can be regarded as specific applications of mathematical steps executed in a mechanical manner. If an evaluation standard does not lend itself to a solvable framework, it becomes essentially meaningless. Second, it should possess definiteness, whereby there exists a unanimous understanding among evaluators without any uncertainty. Third, it should exhibit equivalence across multiple evaluators, ensuring consistency among their assessments.

We propose the establishment of an evaluation standard at the level of the definition of an individual problem or task instance. There are several valid reasons for this approach.

Firstly, the first and second high-level components define LEECs, and they are the cornerstone of the evaluation conditions. These two high-level components serve as the foundation upon which the other two lower-level components are derived.

Secondly, the first high-level component, which pertains to the definition of a problem or task, is not definite and specific, as it may encompass a population of different instances. To be precise, an equivalent, definite, and solvable problem or task instance qualifies as “an

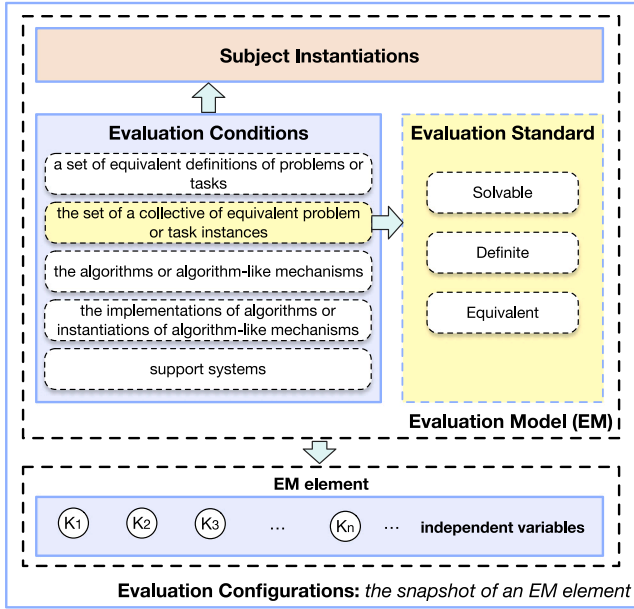


Fig. 6. The relationships among evaluation configurations, evaluation conditions (C), Evaluation models (M), and evaluation standards. Please note that $C = M/U$.

evaluation standard". This definition serves as the basis for conducting evaluations on the subjects. Fig. 6 shows the relationships among the evaluation standard, evaluation conditions, evaluation configuration, and evaluation standard.

We discuss the subtle differences between LEECs and evaluation standards. LEECs are defined on the first and the second high-level components, while the evaluation standard is only defined at the second high-level component, and they are closely related but with distinct implications. Their shared objective is to guarantee the comparability of evaluation outcomes. The aim of LEECs is to ensure the least equivalence between two ECs, while the evaluation standard is to provide the most governing component of an EC that ensures the comparability of the evaluation outcomes. Second, LEECs imply a state space, while an evaluation standard is an evaluation criterion.

3.5.5. The establishment of an REM

Based on the notations defined in Section 3.5.3, we formalize an EM as $M = C \times U = E' \times E \times A' \times A \times S \times U$.

An element of an EM, denoted as $m \in M$, can be expressed as $m = (c, u) = (e', e, a', a, s, u)$. Here, $e' \in E'$ represents a given problem or task, $e \in E$ represents a specific instance of the problem or task, $a' \in A'$ represents a particular algorithm-like mechanism, $a \in A$ represents an instantiation of the algorithm-like mechanism, $s \in S$ represents a support system, and $u \in U$ represents a subject instantiation.

In addition, we will define how to establish an REM. We assume that an EM element, denoted as $m \in M$, is made of n' independent variables, rewritten as $m = (k_1, \dots, k_{n'}) = (e', e, a', a, s, u)$. We note $M = K_1 \times \dots \times K_n = C \times U = E' \times E \times A' \times A \times S \times U$. Please bear in mind that the number of variables n is greater than n' .

For each EM element, represented as $m = (k_1, \dots, k_{n'})$, we follow a specific methodology in which only one independent variable at a time, from the set $k_1, k_2, \dots, k_{n'}$, is allowed to vary while keeping the remaining variables constant. This controlled experimentation approach is referred to as an REM, as defined in Section 3.4.

We define the evaluation cost of an EM or ES as the costs of constructing, traversing, and assessing its corresponding REM, where $\|M\|$ stands for the capacity of an EM or ES and μ stands for a constant coefficient:

$$\text{cost}(M) = \mu \|M\| = \mu \|K_1\| \times \|K_2\| \times \dots \times \|K_n\|$$

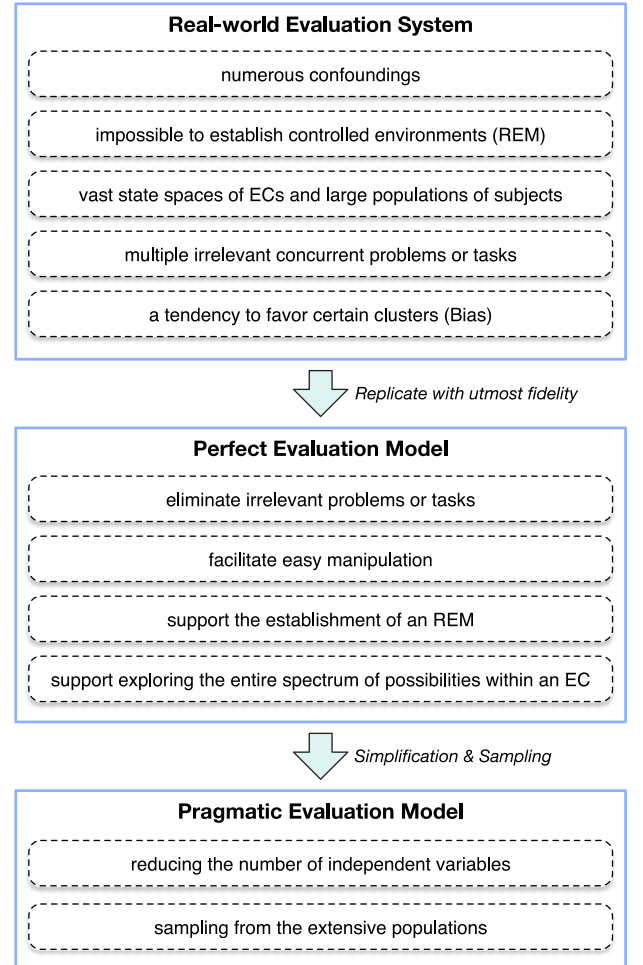


Fig. 7. Universal evaluation methodology in complex scenarios.

3.6. Universal evaluation methodology in complex scenarios

From the revealed essence of the evaluation, it seems that performing an evaluation is straightforward. Unfortunately, in reality, there are evaluation scenarios with different levels of complexity. Fig. 7 presents a universal evaluation methodology in complex scenarios.

We refer to the entire population of real-world systems that are used to evaluate specific subjects as the *real-world ES*. Assuming no safety concerns are present, the real-world ES serves as a prime candidate for the assessment of the subjects of interest. Unfortunately, there are five significant obstacles to consider when assessing diverse subjects within a real-world ES.

Firstly, the presence of numerous confounding in the real-world ES poses a considerable challenge. It is often difficult, if not impossible, to completely eliminate these confounding. They can complicate the evaluation process by introducing variables that make it challenging to isolate the effects of different independent variables.

Secondly, manipulating the real-world ES is a formidable task, making it virtually impossible to establish controlled environments, known as REM, for evaluating subjects. Additionally, the interconnected nature of subjects, support systems, and other components of ECs further complicates the establishment of an REM.

Thirdly, the vast state spaces of ECs and the large populations of subjects result in high evaluation costs. The sheer scale of these systems makes it expensive and time-consuming to thoroughly performing assessment and analysis.

Fourthly, within the real-world ES, multiple irrelevant concurrent problems or tasks may be taking place simultaneously, which may not directly align with the subject being assessed. This further adds complexity to the evaluation process and introduces confounding.

Lastly, it is important to acknowledge that the real-world ES, regardless of the nature of its problems or tasks, tends to exhibit a bias towards certain clusters. This bias can manifest through problem or task instances, algorithm-like mechanisms, and their instantiations. However, this bias towards specific groups can limit our ability to fully explore and understand the entire range of possibilities available to us.

When assessing a specific quantity of interest in an evaluation experiment, it is crucial to closely examine the relationship between the quantity of the EM and the corresponding quantity of the real-world ES. The ratio between these two quantities serves as a key indicator of accuracy. Ideally, a ratio closer to 100% signifies higher accuracy in the EM's modelling of the real-world ES.

In our research, we propose the concept of a "perfect EM" that aims to replicate the real-world ES with the highest level of fidelity, achieving a remarkable ratio of 100%. In theory, a perfect EM would possess several characteristics that enhance the evaluation of subjects within the EC framework.

Firstly, it would eliminate irrelevant problems or tasks that may be used to establish ECs, ensuring that the evaluation focuses on specific and directly applicable contexts. This targeted approach would enhance the relevance and applicability of the evaluation process.

Secondly, a perfect EM would facilitate easy manipulation, allowing for the free and artificial configuration of different evaluation settings. This flexibility would enable researchers to explore various scenarios and assess subjects under a range of conditions, enhancing the depth and breadth of the evaluation process.

Thirdly, a perfect EM would support the establishment of an REM, effectively eliminating confounding. By isolating and controlling variables of interest, researchers could gain more accurate insights into the impact of specific factors on the subjects being evaluated.

Furthermore, a perfect EM would have the capability to thoroughly explore and understand the entire spectrum of possibilities within an EC. This would include problem or task instances, algorithms-like mechanisms, and their instantiations. By encompassing this comprehensive range, researchers could gain deeper insights into the behavior and performance of subjects within the EC framework.

However, it is important to note that achieving a truly "perfect EM" may be challenging, if not impossible. The real-world ES is complex and dynamic, and replicating it with absolute fidelity is a monumental task. While we can strive to create more accurate and representative evaluation environments, it is crucial to recognize the inherent limitations and constraints that exist in the real world.

Nevertheless, by considering the concept of a perfect EM and its accompanying characteristics, we can strive to improve the evaluation of subjects within the EC framework and enhance our understanding of their performance within real-world contexts.

The characteristics of the perfect EM, such as encompassing large populations of problem or task instances, algorithm-like mechanisms, instantiations, and support systems, as well as a vast number of independent variables, can lead to significant evaluation costs. However, to address this challenge, it is important to propose a pragmatic EM that simplifies the perfect EM in two key ways.

Firstly, to reduce the evaluation costs associated with a large number of independent variables, it is crucial to identify and focus on the variables that have a significant impact on the evaluation outcomes. By identifying and prioritizing these variables, researchers can streamline the evaluation process and allocate resources more efficiently. Negligible variables that have a minimal effect can be excluded or controlled for, reducing complexity and costs. It is worth emphasizing that the simplification involved in creating a pragmatic EM will inevitably lead to a decrease in the accuracy of the evaluation model.

Secondly, sampling techniques can be employed to manage the extensive populations of problem or task instances, such as algorithm-like mechanisms, their instantiations, and support systems. Rather than evaluating every single possibility, researchers can select representative samples that capture the diversity and range of the population. This approach allows for a more manageable evaluation process while still maintaining a good level of coverage and representation.

In nature, a pragmatic EM is a subject equipped with a simplified and sampled EC. It can be considered as a sample of a perfect EM without taking into account any potential decrease in its accuracy. In order to measure the extent to which the statistics of a pragmatic EM can infer the parameters of a perfect EM, we employ the criterion of confidence interval and confidence level. The confidence level provides us with the probability that the estimated parameters of a perfect EM fall within a specific range of values. Meanwhile, the confidence intervals establish a range of values within which we can reasonably expect the true parameters of a perfect EM to fall. By utilizing these statistical measures, we can assess the degree of alignment between the statistics of a pragmatic EM and the parameters of a perfect EM. This allows us to gauge the effectiveness and validity of the pragmatic EMs.

By implementing these simplifications in a pragmatic EM, researchers can strike a balance between comprehensiveness and feasibility. The pragmatic EM allows for a more practical and efficient evaluation of subjects within the EC framework, mitigating the challenges posed by evaluation costs and the complexity of the perfect EM.

In representing different ECs, we use specific symbols. The symbol C_r denotes the EC in a real-world ES (a real-world EC), which can be calculated as $E'_r \times E_r \times A'_r \times A_r \times S_r$. Similarly, the EC in a perfect EM (a perfect EC) is denoted by C_p , calculated as $E'_p \times E_p \times A'_p \times A_p \times S_p$. Lastly, the EC in a pragmatic EM (a pragmatic EC) is represented by C_g , calculated as $E'_g \times E_g \times A'_g \times A_g \times S_g$.

Likewise, we use symbols to denote a real-world ES (M_r), a perfect EM (M_p), and a pragmatic EM (M_g). These are represented as:

$$M_r = E'_r \times E_r \times A'_r \times A_r \times S_r \times U_r$$

$$M_p = E'_p \times E_p \times A'_p \times A_p \times S_p \times U_p$$

$$M_g = E'_g \times E_g \times A'_g \times A_g \times S_g \times U_g$$

These symbols help us distinguish and calculate the various components of ECs and EMs in different contexts.

3.7. Fundamental issues in evaluatology

This subsection presents three fundamental issues in Evaluatology.

3.7.1. Ensure transitivity of EMs

The key to the effectiveness and efficiency of evaluations in different scenarios is to establish a series of EMs that ensure the transitivity of the primary characteristics.

In Section 3.6, we have effectively examined and described a real-world ES, along with its corresponding EM, which we call a perfect EM, and their interconnections.

The real-world ES system presents three notable obstacles. Firstly, the presence of numerous confounding creates a challenge as they cannot be completely eliminated. Secondly, the existence of multiple irrelevant problems or tasks adds another layer of complexity. Lastly, regardless of the nature of the problems or tasks involved, there is a tendency for the system to exhibit bias towards certain clusters exhibited by problem or task instances, algorithmic mechanisms, and their instantiations.

To overcome these obstacles, a perfect EM, which possesses several key characteristics, is proposed. Firstly, the model eliminates any irrelevant problems or tasks that may be used to derive evaluation standards for assessing different subjects. Secondly, the model enables the free setting of an REM. Thirdly, the model allows for a comprehensive

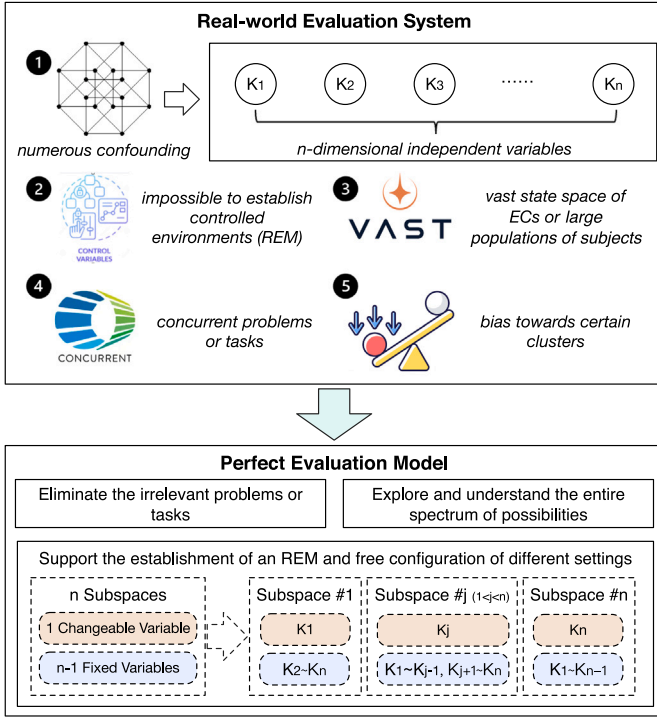


Fig. 8. A perfect EM resembles a real-world ES.

exploration and understanding of the entire spectrum of possibilities in terms of problem or task instances, algorithmic mechanisms, and their instantiations.

Compared to a real-world ES $M_r = E'_r \times E_r \times A'_r \times A_r \times S_r \times U_r$, a perfect EM eliminates its irrelevant problems or tasks, represented by $M_p = E'_p \times E_p \times A'_p \times A_p \times S_p \times U_p$. Transforming a real-world ES into a perfect EM should ensure the transitivity of the following characteristics (see Fig. 8).

$$E'_p \subset E'_r$$

$$SE_p/E'_p \supset SE_r/E'_r$$

$$SA'_p/E'_p/E_p \supset SA'_r/E'_r/E_r$$

$$SA_p/E_p/E_p/A'_p/S_p \supset SA_r/E_r/E_r/A'_r/S_r.$$

The perfect EM encompasses large populations of problem or task instances, algorithm-like mechanisms, their instantiations, and support systems. Also, it entails a vast number of independent variables.

To overcome this difficulty, it becomes essential to propose a pragmatic EM that simplifies the perfect EM in two ways: (1) reducing the number of independent variables and (2) sampling from the extensive populations of support systems, problem or task instances, algorithm-like mechanisms, and their instantiations.

A pragmatic EM adopts a sampling approach on the perfect EM, resulting in a smaller space to work with. To formalize this process, we introduce the notation $s(\cdot)$ to represent the sampling function. Additionally, the pragmatic EM streamlines the independent variables within the perfect EM by excluding those that have minimal impact.

For each element m_g in the sampled space M_g , which is a subset of the perfect EM M_p , we denote the corresponding element in M_p as m_p . When transforming a perfect EM into a pragmatic EM, it is essential to maintain the transitivity of the following characteristics:

$M_g = s(M_p)$: The sampled space M_g is obtained through the application of the sampling function s on the perfect EM M_p .

$M_g \subset M_p$: The sampled space M_g is a subset of the perfect EM M_p .

$m_g = (k_1, \dots, k_{n'}) \in M_g$: Each element m_g in the sampled space M_g consists of a set of independent variables $(k_1, \dots, k_{n'})$.

$m_p \in M_p$ is the matched element in the perfect EM M_p corresponding to m_g . m_p consists of a set of independent variables $(k_1, \dots, k_{n''})$. n'' is greater than n' , ensuring that the corresponding element in the perfect EM includes at least as many independent variables as that of the element in the pragmatic EM.

3.7.2. Perform cost-efficient evaluation with controlled discrepancies

By disregarding the accuracy of an EM, conducting evaluations solely through the establishment of an REM within a perfect EM may indeed result in maximum confidence. However, this approach also comes with a significant drawback — the exorbitant cost it entails. The process of creating an REM within a perfect EM can be prohibitively expensive, making it impractical for many organizations. Therefore, it is crucial to strike a balance between ensuring the discrepancy threshold of the evaluation outcomes and managing the associated costs when implementing evaluation processes.

When creating a pragmatic EM from a perfect EM, the discrepancy threshold ϵ , which is a discrepancy limit that can be tolerated in an evaluation scenario, holds the potential to exert a profound influence on the evaluation results and, in certain instances, it would give rise to grave concerns, particularly in the context of safety-critical tasks where failure could lead to detrimental side effects such as harm, loss of life, or significant environmental damage. So, after thoroughly understanding the stakeholders' evaluation requirements, a risk function $\gamma(\cdot)$ could be predefined. When the stakes are high, and there is a greater risk associated with the evaluation outcomes, it becomes imperative to have a lower discrepancy threshold between the evaluation outcomes of a pragmatic EM and a perfect EM. This is because the potential consequences of making a wrong decision or drawing inaccurate conclusions become more significant.

In Section 3.7.1, we use the notation $s(\cdot)$ to represent the sampling function. In creating a pragmatic EM from a perfect EM, the accuracy of EM decreases. We use the notation $m(\cdot)$ to represent this modeling process. We use the notation $e(\cdot)$ to represent the process of ensuring different equivalency levels of EC.

We introduce a discrepancy function of the evaluation outcomes $\text{disc}(\cdot)$ between M_g and M_p . When the discrepancy is 0, it indicates that M_g and M_p are equivalent.

The discrepancy function of the evaluation outcomes $\text{disc}(\cdot)$ between M_g and M_p is defined as follows. In the formulation, $\rho(\cdot)$ is a measurement function, and $v(\cdot)$ is a value function. Besides, we define the evaluation cost as the product of a constant μ and the space capacity of M_g . This cost factor allows us to incorporate the resource constraints and practical considerations associated with the evaluation process.

$$\begin{cases} \text{discrepancy threshold } \epsilon = \gamma(\cdot), \\ \text{disc}(M_g, M_p) = \text{disc}(v(\rho(e(m(s(M_p)))))), v(\rho((M_p))), \\ \text{cost}(M_g) = \mu \|M_g\|. \end{cases}$$

Based on the above formulation, the evaluation issue of balancing evaluation cost and the discrepancies in the evaluation outcomes can be framed as an optimization problem. The objective is to minimize the evaluation cost, represented by cost while ensuring the discrepancies in the evaluation outcomes, denoted as $\text{disc}(M_g, M_p)$, do not exceed a predefined discrepancy threshold ϵ (see Fig. 9).

The optimization problem can be formulated as follows:

$$\arg \min \text{cost}(M_g) \text{ subject to } (\text{disc}(M_g, M_p) < \epsilon).$$

3.7.3. Ensure evaluation traceability

According to the third axiom of evaluation, for a well-defined subject, the divergence in the evaluation outcomes can be attributed to disparities in ECs, thereby establishing evaluation traceability.

Conceptually, traceability asks for a quantified mapping between the differences in the input and output of the value function $v(\cdot)$ decided by the evaluation community and the measurement function $\rho(\cdot)$

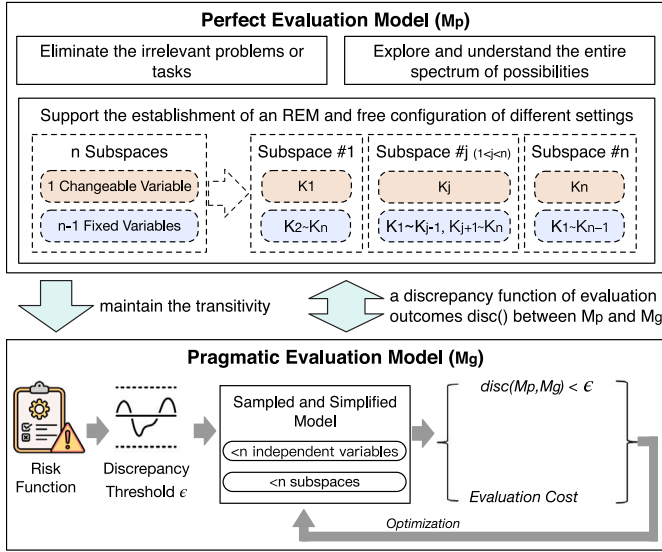


Fig. 9. Proposing a pragmatic EM based on the evaluation risk function.

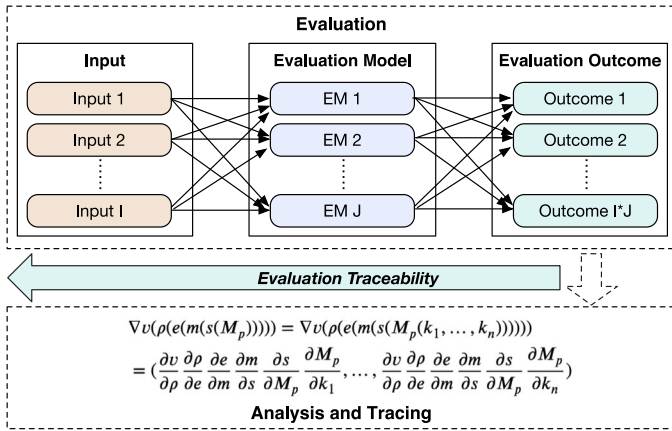


Fig. 10. Evaluation traceability.

through the evaluation process, described by the mathematical model we formulated above, and the model in the following formulation is a perfect EM. We discover that this concept aligns well with the mathematical notation of the gradients of a function, which gives the rate of changes in the output for each input variable. Fig. 10 shows how to ensure the evaluation traceability. In the context of evaluation, the gradient of evaluation outcomes can be written as follows, which is a matrix or tensor:

$$\nabla v(\rho(e(m(s(M_p)))))) = \nabla v(\rho(e(m(s(M_p(k_1, \dots, k_n))))))$$

$$= \left(\frac{\partial v}{\partial \rho} \frac{\partial \rho}{\partial e} \frac{\partial e}{\partial m} \frac{\partial m}{\partial s} \frac{\partial s}{\partial M_p}, \frac{\partial v}{\partial \rho} \frac{\partial \rho}{\partial e} \frac{\partial e}{\partial m} \frac{\partial m}{\partial s} \frac{\partial s}{\partial M_p} \frac{\partial M_p}{\partial k_1}, \dots, \frac{\partial v}{\partial \rho} \frac{\partial \rho}{\partial e} \frac{\partial e}{\partial m} \frac{\partial m}{\partial s} \frac{\partial s}{\partial M_p} \frac{\partial M_p}{\partial k_n} \right)$$

The closed-form mathematical expression is not always available for various EC components in evaluation. Nevertheless, we can follow the method of acquiring gradients in numerical methods by creating perturbations in the ECs for various input variables and observing the differences in the composite evaluation outcomes, thus approximating the gradients.

3.7.4. Connect and correlate evaluation standards across diverse disciplines

While the constituents comprising ECs may differ across distinct evaluation scenarios, a governing fundamental component of ECs

emerges as the evaluation standard. The shared qualities of evaluation standards, as denoted by the SDE characteristics, indicate the potential for establishing a correlation between evaluation standards across diverse disciplines.

The evaluation standard serves as a fundamental pillar within any evaluation model. By establishing connections between evaluation standards across various disciplines, we have the potential to construct a comprehensive framework encompassing evaluation issues in all fields. This holistic framework, known as the grand unified theory of evaluatology, allows for a thorough exploration of evaluation-related matters.

Before defining an evaluation standard, it is essential to understand the nature of the stakeholders' primary problems or tasks. In the rest of this article, a problem or task refers explicitly to a computational problem. Make sure to distinguish between a problem and a problem instance: A problem is an infinite collection of problem instances, each of which is a problem with concrete configurations, which we have elaborated in Section 3.5.1.

Computational complexity theory provides the basis for understanding the nature of primary problems. For example, complexity classes – that are defined by bounding the time or space used by the algorithm – can be used to understand the problem's nature [21]. Computability theory [22] seeks a more general question about all possible algorithms that could be used to solve the same problem. That is to say, the computability theory provides a viable solution to answer whether a problem is solvable. In understanding the problem instance, the theory of analysis of algorithms [23] can be used to analyze the amount of resources needed by a particular algorithm to solve a problem instance.

While the computational complexity and other theories mentioned above lay the foundation, the formulation of evaluation standards introduces novel concerns. The first challenge lies in addressing scenarios where articulating a mathematical model explicitly becomes an insurmountable task. Regrettably, this circumstance is not uncommon, as numerous problems defy expression through a mathematical model at present.

To ensure the definitiveness and equivalency of the evaluation standard, it is imperative to establish a rigorous problem space definition and a problem instance space definition, which provides the quantitative foundation for the comparability and traceability of different EMs.

The representativeness of the evaluation standard is a crucial aspect that warrants discussion. Understanding the composition of problems is crucial in identifying the problem or task that best represents the whole. For instance, across various scientific and engineering disciplines, problems often exhibit a hierarchical structure, where a significant problem can be broken down into several smaller problems commonly referred to as "dwarfs" [24]. This pattern can be considered one of the foundational structures within problem domains. Gaining profound insights into the structural aspects of problems proves immensely valuable when assessing complex and multifaceted subjects.

Fig. 11 shows how to correlate evaluation standards across diverse disciplines. In an optimal scenario, we can identify evaluation standards that embody the SDE characteristics across various disciplines. Ultimately, we can establish connections and correlations between evaluation standards from different fields, giving rise to the grand unified theory of evaluatology. The objective of this theory is to present a hierarchical framework of evaluation standards. Within this framework, we can identify a minimal set of fundamental evaluation standard "dwarfs" along with their respective variations. Complex evaluation standards are formed by combining two or more of these evaluation standard dwarfs and their variants. This hierarchical structure of evaluation standards will greatly facilitate the reuse and sharing of knowledge.

4. Benchmarkology: the engineering of evaluation

This section unveils the core essence of a benchmark and introduces the benchmark-based engineering of evaluation, which we call benchmarkology. Furthermore, we provide guidelines and workflows within the realm of benchmarkology.




Fig. 11. Correlating evaluation standards across diverse disciplines.

4.1. The essence of a benchmark

Benchmarks are extensively employed across various disciplines, albeit lacking a formal definition. Based on the science of evaluation, we propose a precise delineation of a benchmark as a *simplified and sampled EC*, specifically a *pragmatic EC*, that ensures different levels of equivalency, ranging from LEECs to EECs.

Within the framework of this definition, a benchmark comprises three essential constituents. The first constituent is the *stakeholder's evaluation requirements*, which encompass various factors. These include the risk function, which evaluates the potential risks associated with the benchmark. Additionally, the discrepancy threshold of the evaluation outcomes, which determines the acceptable level of deviation in evaluation outcomes, is considered. The evaluation confidence interval plays a crucial role in predicting the parameter of a perfect EM. Lastly, the evaluation cost of EM is taken into account, and the resources required for conducting the evaluation are assessed. By considering these elements, the benchmark can effectively address the evaluation requirements of stakeholders.

The second constituent of the benchmark framework is the *EC configuration and mechanisms*. This includes several elements crucial for the benchmark's effectiveness. Firstly, it involves defining the set of problems or tasks that the stakeholders face when addressing them. Additionally, it encompasses the set of equivalent problem or task instances, which helps ensure specificity in the evaluation process. The benchmark also considers algorithm-like mechanisms and their instantiations, which play a significant role in solving the defined problems or tasks. The support systems, which provide necessary resources and environments, are also taken into account.

Moreover, the benchmark provides the means to eliminate confounding variables that may affect the evaluation outcomes. Also, the benchmark provides the mechanism to ensure varying levels of EC equivalency, determining the extent to which different benchmark instances can be considered equivalent.

By considering these EC configurations and mechanisms, the benchmark can provide a comprehensive and standardized approach to evaluating problems or tasks.

The third constituent is the *metrics and reference*, including the definitions of quantities, the value function, composite evaluation metrics, the reference subject, and the reference evaluation outcomes.

In the subsequent sections of this article, we will refer to these three constituents as the complete constituents of a benchmark. Fig. 12 shows the three essential constituents of a benchmark.

4.2. The goal of benchmarkology

As expounded upon in Section 3, the science and engineering of evaluation, known as evaluatology, aims to apply the EECs to various subjects and establish an REM. A benchmark can be viewed as a simplified and sampled EC, specifically a pragmatic EC, that ensures different levels of equivalency, ranging from LEECs to EECs. In this context, a benchmark-based approach to the evaluation problem is considered a feasible engineering methodology, given the widespread use of benchmarks across various disciplines. Consequently, we propose

the formal definition of *benchmarkology* as an engineering discipline concerned with the quantitative assessment of diverse subjects using benchmarks.

Undoubtedly, the theory of evaluatology serves as the foundation for benchmarkology. Nevertheless, benchmarkology has its unique objective — to furnish guiding principles and engineering evaluation methodologies.

4.3. The principles in building benchmarks

In Section 3, we have extensively discussed the fundamental axioms of evaluatology. Nonetheless, this particular subsection delves deeper into the principles that underpin the creation of benchmarks derived from these four evaluation axioms.

The first principle focuses on the validity of metrics within a benchmark. According to the First Axiom of Evaluation, also known as the Axiom of the Essence of Composite Evaluation Metric, there are three criteria to determine the validity of metrics in benchmarks. If a metric does not meet these criteria, it is considered invalid. The three criteria are as follows:

1. The metric should be a base quantity.
2. The metric should represent another quantity that has inherent physical significance.
3. The metric should be a composite evaluation metric that is explicitly defined by a value function.

The second principle pertains to the comprehensiveness of the configurations of a valid benchmark. According to the Second Axiom of Evaluation, known as the Axiom of True Evaluation Outcomes, when a well-defined subject is equipped with a well-defined EC, its evaluation outcomes possess true values. A well-defined EC should reveal all its well-defined components. Each component within the EC plays a critical role in determining the evaluation outcomes. When the components of the EC are not well-defined, uncertainty is introduced into the evaluation process. Without clear and specific components, the evaluation outcomes become unpredictable and lack reliability.

The evaluation outcome using a benchmark should reveal the complete evaluation configurations. Unfortunately, many contemporary benchmarks, both in terms of state-of-the-art and state-of-the-practice, have failed to disclose their comprehensive evaluation configurations fully.

Some benchmarks, like the widely-used CPU benchmark SPEC CPU, may omit certain constituents or their components, e.g., the support system. In such cases, it becomes essential to clearly define the conditions under which the simplification is made, ensuring that the benchmark can still provide meaningful and valid results. By providing these detailed evaluation configurations, we can ensure that the benchmark remains a reliable tool for evaluation purposes.

The third principle centers around the concept of benchmark traceability. In accordance with the Third Axiom of Evaluation, also known as the Axiom of Evaluation Traceability, it is crucial to establish benchmark traceability to enable the comparison of different benchmarks. This means that it is a top priority to trace the discrepancies

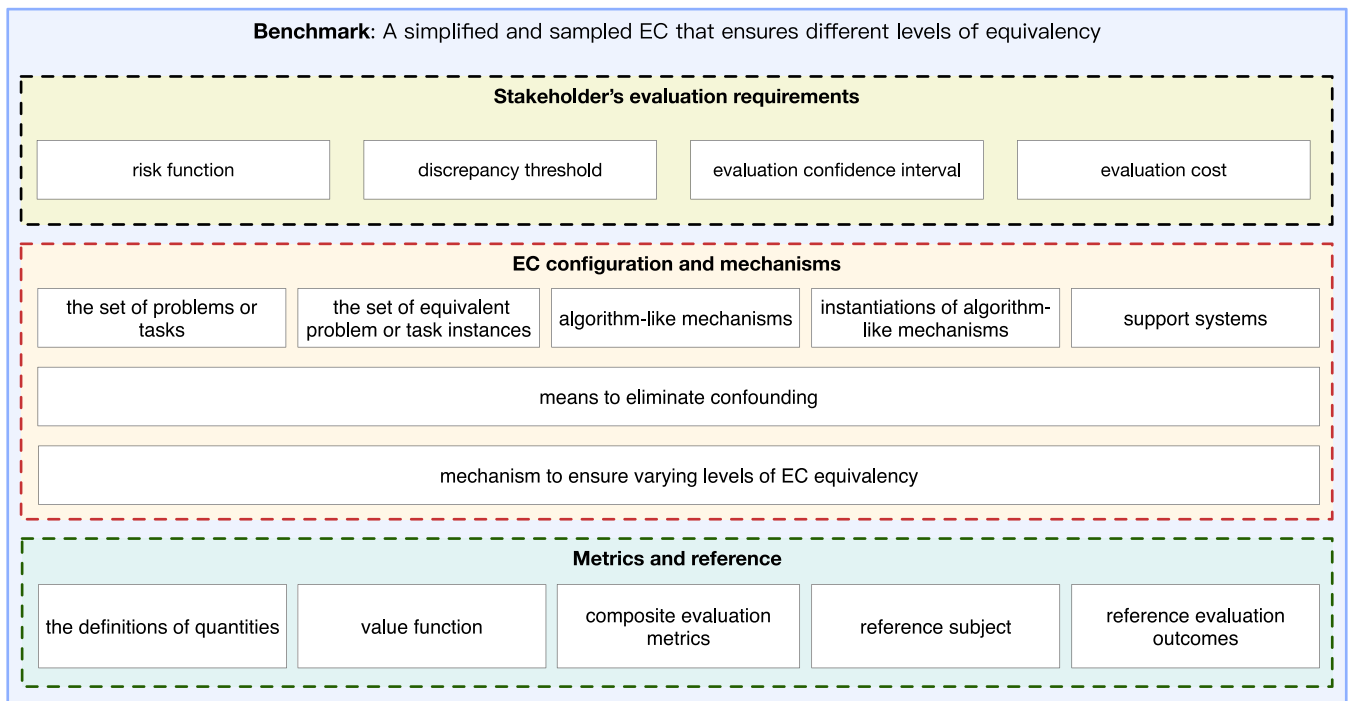


Fig. 12. A benchmark comprises three essential constituents.

in evaluation outcomes back to variations in the configurations of different benchmarks.

The fourth principle encompasses the validity of the speedup of two evaluations obtained using the benchmark. It emphasizes the importance of understanding the implications of the speedup obtained in different pragmatic EMs in relation to those of the real-world ES or perfect EM.

Comparison is a common practice using benchmarks. In practice, the speedup is often a key metric of interest that is obtained in different pragmatic EMs. When we compare two evaluations, we use the benchmarks as a means to infer the true speedup in the real-world ES or perfect EM.

To illustrate this principle, let us consider a specific example where we compare two subjects, Subjects A and B, and obtain a speedup (either greater than or less than 1).

In a simulated CPU scenario (a pragmatic EM), the reported speedup is 1.3, indicating an improvement in performance for Subject A compared to Subject B. However, the ratio between the speedup in the EM and the corresponding speedup in the real-world ES has a 90% confidence interval of [0.7, 1.9]. This means that the ratio could be any value within this interval, including 1.6.

If the ratio of speedup is indeed 1.6, the actual speedup in the real-world ES would be 1.3 divided by 1.6, resulting in a value of 0.8. This indicates a degraded performance for Subject A in comparison to Subject B, reported on the real-world ES, which is contradicted by an improvement in performance reported on a simulated system.

It is crucial to note that relying solely on the reported speedup in the EM without considering its implication in the real-world ES can lead to misleading interpretations and decisions. To ensure the validity of evaluation outcomes, it is essential to take into account the confidence interval associated with the speedup ratio between the EMs and real-world ES.

4.4. The universal methodology in benchmarkology

The aforementioned principles offer valuable insights into the fundamental components of a benchmarkology workflow, as shown in Fig. 13.

4.4.1. Understand stakeholders' evaluation requirements

During the initial phase, it is essential for evaluators to gain a comprehensive understanding of the stakeholders' evaluation requirements, the first constituent of the benchmark, which we discussed in detail in Section 4.1. This crucial step allows for the alignment of these requirements with the overall purpose of the evaluation.

During this phase, a thorough examination of the evaluation risk function, the discrepancy threshold of evaluation outcomes, and the evaluation cost should be conducted. Additionally, for quantities or variables of interest, it is crucial to establish their evaluation confidence interval when using the benchmark. This quantification allows for an assessment of how effectively the benchmark can infer or predict the parameters of a perfect EM.

Unfortunately, in state-of-the-art or state-of-the-practice benchmarks, the importance of this phase is often overlooked. There are two possible reasons for this oversight.

Firstly, in certain evaluation scenarios, the discrepancy in evaluation outcomes, whether intermediate or large, may not have significant consequences. However, it is crucial to note that this is not the case in scenarios involving safety-critical, mission-critical, and business-critical applications. In these situations, even minor deviations can have severe impacts on the overall outcome.

Secondly, the benchmark process itself is an engineering practice that emphasizes iterative and refined operation. As a result, it implicitly incorporates some procedures of this phase.

Therefore, it is imperative to recognize the significance of understanding stakeholders' evaluation requirements, particularly in scenarios where the stakes are high and any discrepancy from expected outcomes can have critical implications.

4.4.2. Design and implement intricate evaluation mechanisms and policies

This phase plays a crucial role, particularly in complex evaluation scenarios, and can be quite costly. Its primary aim is to provide a solid foundation for generating a benchmark.

The real-world ES reflects the complexities and nuances of actual evaluation environments. In this phase, the evaluator takes on the crucial task of building and investigating the real-world ES. The

methodology discussed in Section 3.6 can serve as a helpful guide for this process.

During the investigation, evaluators need to consider several aspects carefully. One aspect involves identifying and eliminating irrelevant problems or tasks that may not be applicable to the assessment of the subjects under evaluation. This ensures that the evaluation focuses on relevant and meaningful aspects of the subject. Another important consideration is that evaluators must recognize any constraints or factors that may impact the evaluation process. This understanding lays the groundwork for creating a perfect EM that explores all the possibilities. A perfect EM serves as an ideal evaluation model, and the design and implementation of this model are also key focus in this phase. We have discussed its main concerns in Section 3.6.

Additionally, in this phase, two key policies and their accompanying procedures are of utmost importance. Firstly, the modeling policy and procedure guide the process of transforming the real-world ES into an EM, striking a balance between accuracy and cost. This involves capturing the essential elements and characteristics of the real-world ES in the model while ensuring that the modeling process is efficient and cost-effective.

Secondly, the sampling policy and procedure play a vital role in transitioning from a perfect EC to a pragmatic EC. This transition aims to save on evaluation costs while still maintaining a high level of confidence in the evaluation results. The sampling policy and procedure guide the selection of representative samples from the perfect EC, ensuring that the pragmatic EC captures the essential aspects and characteristics of the perfect EC while being more practical and resource-efficient.

By following these policies and procedures, the evaluation process is adapted to real-world conditions and constraints. The modeling policy and procedure enable the creation of an EM that accurately represents the real-world ES, while the sampling policy and procedure ensure that the pragmatic EC reflects the essential elements of the perfect EC. This allows for a more effective and efficient evaluation process that balances accuracy, cost, and confidence.

Overall, with the facilitation of the study of real-world ES and the perfect EM, these modeling and sampling policies and procedures are essential in this phase to guide the modeling and sampling processes, ensuring that the evaluation process is well-suited to real-world conditions and constraints.

4.4.3. Decide representative evaluation standards

The third phase is to decide the representative evaluation standards that guarantee the least EECs, ensuring the comparability of the evaluation outcomes. The main objective of this phase is to carefully consider the relevant stakeholders involved and gain a deep understanding of their principal interests and concerns. This phase requires evaluators to identify and comprehend the primary problems or tasks that need to be addressed.

It is important to note that each stakeholder has a unique perspective, leading to subtle differences in the problems or tasks they face. Therefore, it is crucial for evaluators to recognize and take into account these varying perspectives, ensuring that the evaluation standards are comprehensive and reflective of the diverse interests of the stakeholders involved.

In Section 3.5.4, we have explored the concept of evaluation standards and how they are derived from the abstract problem or task at hand. Building upon this understanding, the evaluator's next step is to determine the specific evaluation standards by selecting representative instances of the primary problems or tasks faced by the stakeholders.

It is important to recognize that different stakeholders will have distinct evaluation standards. This is because their perspectives and priorities vary based on their unique roles and interests. As a result, it is essential for evaluators to consider these differences and tailor the evaluation standards accordingly to ensure that they capture the specific needs and concerns of each stakeholder involved.

It is true that previous evaluation and benchmark practices have often lacked consistent discussions on what qualifies as an evaluation standard. However, we believe that our universal definition of evaluation standards has the potential to be applicable across various evaluation scenarios in different disciplines. Our aim is to provide a comprehensive framework that can guide evaluators in establishing effective evaluation standards.

Furthermore, we acknowledge the importance of recognizing that a realistic benchmark can only capture a small sample of huge populations of instances that are derived from a large population of problems or tasks. Despite the challenges that may arise when explicitly stating the problem or task, we remain committed to adopting a systematic approach to our thinking. This allows us to navigate through such complexities and develop meaningful evaluation standards that align with the objectives of the evaluation process.

4.4.4. Design and implement ECs with different levels of equivalency

Based on the outputs from Phases Two and Three, the subsequent stage involves the design and implementation of ECs with different levels of equivalency. This task varies from different subjects. In general, this phase needs to consider algorithm-like mechanisms and their instantiations, as well as the support systems.

Furthermore, in this phase, it is essential to address the levels of EC equivalency. This involves determining the extent to which different benchmark instances can be considered equivalent. It requires careful consideration of which components can be disregarded or simplified in order to streamline the benchmark process while maintaining its validity and reliability.

Additionally, the benchmark needs to establish mechanisms to eliminate confounding that may impact the evaluation outcomes. Confounding variables can introduce biases or distortions into the evaluation results, affecting their accuracy and reliability. One approach to address confounding is by employing our proposed REM methodology. This methodology provides a systematic framework to identify and eliminate confounding variables and ensure that the evaluation outcomes are not influenced by extraneous factors.

4.4.5. Perform measurement and/or testing

The fourth phase encompasses measurements and/or testing, guided by the principles and practices of the metrology and testing theory. The measurement and testing process serves multiple purposes. Firstly, evaluators must determine which properties or quantities to measure, keeping in mind what base quantities and other quantities that carry physical meaning are.

Furthermore, evaluators must also consider the cost of measurement and testing, ensuring that it aligns with budgetary constraints. It becomes imperative for them to make informed decisions on various aspects, such as how, when, and to what extent to perform testing, sampling, and measuring these properties or quantities. By doing so, they can effectively manage resources while still obtaining valuable data for their research.

4.4.6. Perform assessments

The fifth phase entails assessment, wherein the defining of a value function takes precedence. The rationale behind establishing a value function lies in the aim of encapsulating numerous quantities that surpass our capacity for recognition into a singular metric. This value function serves as a proposed function, mapping the target properties or quantities measured during the preceding phase to the evaluation outcomes in order to reflect the concerns or interests of the stakeholders. Given that stakeholders often possess varying concerns or interests, it is common to propose multiple value functions from different perspectives. Different communities may reach a consensus on how to define a value function. Generally, evaluators must engage in consultation with the stakeholders to define a composite evaluation metric in the form of

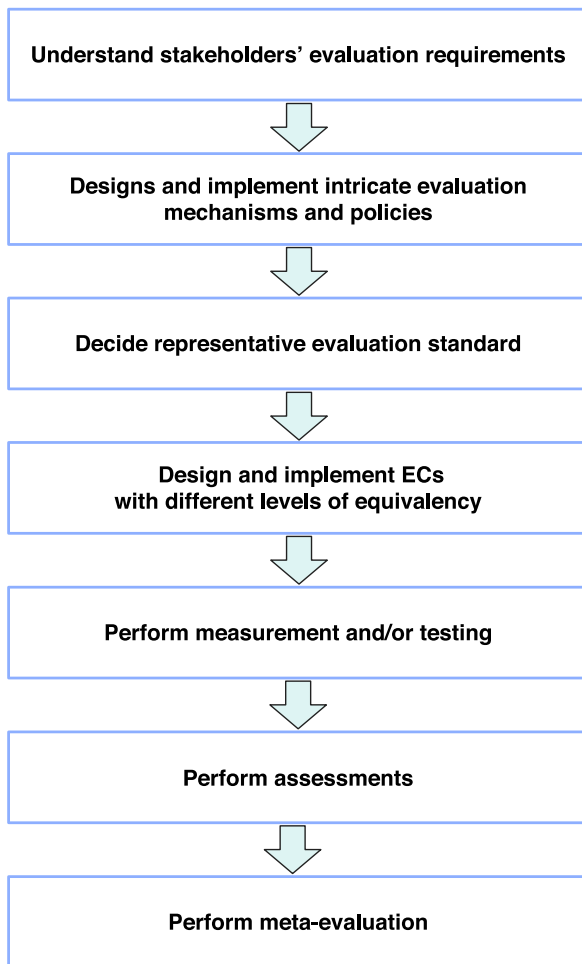


Fig. 13. The universal methodology in benchmarkology.

a value function that truly encapsulates the stakeholders' concerns or interests.

Subsequently, evaluators can compare the obtained evaluation outcomes to the reference evaluation outcomes, ultimately enabling judgment on different subjects, such as performance, value, merit, weaknesses, worth or significance, as well as positive or negative effects.

4.4.7. Perform meta-evaluation

The last phase involves conducting meta-evaluations. The evaluators are tasked with reviewing all the evaluation processes and determining whether the theory, data, or evidence produced can substantiate the conclusions drawn in the evaluations. In this phase of meta-evaluation, various perspectives are taken into account, including the dimensionality of measurement data [7], the reliability of measurement results, the validity of evaluation outcomes, the traceability of benchmarks, as well as the cost and cost-efficiency of the evaluation itself [1].

The "dimensionality" of measurement data refers to the number and nature of variables that are reflected in the assessment [7]. The reliability of measurement results pertains to the extent to which the measured values accurately reflect the true values. The validity of the evaluation outcome denotes the degree to which the statistics of the benchmark can infer the parameter of the real-world evaluation setting or a perfect evaluation model in terms of the metrics of confidence level and confidence interval.

Two approaches can be undertaken to manage benchmark traceability. Firstly, it is essential to develop a comprehensive mathematical

model capable of capturing the influence of the discrepancies of the evaluation configurations on the evaluation results. This model will serve as a foundation for interpreting and comprehending the evaluation outcomes. Additionally, the community must engage in continual benchmark comparisons. Drawing inspiration from the calibration practices in the field of metrology, this consistent comparison and alignment of different benchmarks will ensure consistency and accuracy in evaluation procedures.

Constrained by the limitations of the budget, the aforementioned six phases can be carried out iteratively, employing a trial-and-error methodology.

5. Why it is essential to develop evaluatology?

In this section, we will begin by conducting a comprehensive review of the state-of-the-art and state-of-the-practice evaluation methods and benchmarks. This review will provide us with a solid understanding of the current landscape and help identify areas where advancements are needed.

To illustrate the importance of advancing the science and engineering of evaluation, we will focus on the evaluation of CPU performance as a prime example. Meanwhile, we will critically reflect on the existing state-of-the-art and state-of-the-practice evaluation methods and benchmarks. This reflection will enable us to identify any limitations or gaps that need to be addressed for more accurate and meaningful evaluations.

Furthermore, we will explore the advantages that the field of evaluatology brings to the table. To ensure clarity and understanding, we will provide a concise summary of the distinctions between evaluation, measurement, and testing.

To ensure consistency and alignment with our proposed universal terminology, we will utilize our established terminology framework when discussing state-of-the-art and state-of-the-practice evaluation and benchmark cases.

5.1. Evaluations across different academic fields

This subsection presents a concise overview of the cutting-edge evaluations in a range of academic disciplines, as well as the prevailing evaluation practices.

5.1.1. Observation study methodologies

Observational study methodologies are widely used in the fields of business science, finance, and education. Even based on a random sample, an observational study still falls short of effectively revealing the cause-and-effect relationships.

Evaluations in the field of business science:

Camp [8] defines benchmarking as "the search for those best practices that will lead to the superior performance of a company". Benchmarking consists of two primary steps [8]: (1) establishes operation targets based on industry best practices; (2) "a positive, proactive, structured process leads to changing operations and eventually achieving superior performance and competitive advantage". In the study conducted by Andersen et al. [25], the essence of benchmarking is summarized as the quest for knowledge and learning from others.

Evaluations in the fields of finance and education:

In the fields of finance and education, indices are widely used as benchmarks to assess the overall performance of the individuals or systems under study. These indices are derived by calculating the weighted average of a selected group of individuals or systems [9].

For example, stock market indices are used as benchmarks to assess the stock market's performance in the finance field. These indices are derived by calculating the weighted average of a selected group of representative stocks [9]. Some widely recognized stock market indices include the Dow Jones Industrial Average, the S&P 500, the NASDAQ

Table 1
The base quantity, value function, and the reference machine specified in different CPU benchmark suites in Section 5.

Benchmark	Category	Base quantity	Value function	Reference machine
SPEC CPU2006	SPECSpeed	Execution time	$x_i = \text{time on a reference machine/time on the evaluation machine [27]; } \bar{x} = \sqrt[n]{\prod_i x_i}$	The Ultra Enterprise 2 with 296 MHz UltraSPARC II chips [28]
	SPECrate	Execution time	$x_i = \text{number of copies * (time on a reference machine/time on the evaluation machine) [27]; } \bar{x} = \sqrt[n]{\prod_i x_i}$	
SPEC CPU2017	SPECSpeed	Execution time	$x_i = \text{time on a reference machine/time on the evaluation machine [27]; } \bar{x} = \sqrt[n]{\prod_i x_i}$	The Sun Fire V490 with 2100 MHz UltraSPARC-IV+ chips [27]
	SPECrate	Execution time	$x_i = \text{number of copies * (time on a reference machine/time on the evaluation machine) [27]; } \bar{x} = \sqrt[n]{\prod_i x_i}$	
PARSEC		Execution time	$x_i = \text{time on the evaluation machine}$	

Composite, and the Shanghai Stock Exchange Composite Index. Different indices employ varying calculation methods. The most common approach is the weighted average method, which determines the index value based on the weighted average of the constituent stock prices. Another method is the geometric mean method, which calculates the geometric average of the stock prices and adjusts it using a base period price. Typically, stock market indices are published at the close of each trading day. Some index providers offer real-time index data, enabling investors to stay informed about the latest market conditions.

The Brent benchmark is used to determine the price of Brent crude oil [26]. Brent crude oil is a type of light and low-sulfur crude oil produced from oil fields in the North Sea region. Due to its relatively stable supply and high quality, Brent crude oil has become a significant benchmark in the international oil market. Traders, investors, and industry participants worldwide reference the Brent benchmark to track and evaluate the price of Brent crude oil.

In the finance discipline, indexes or benchmarks serve as reference measurements or evaluation results. However, these practices often prioritize data collection and processing over building a solid evaluation theory framework.

5.1.2. Experimental methodologies

Experimental methodologies are widely used in the fields of social sciences, computer sciences, psychology, and medicine.

Evaluations in the field of social sciences:

According to Rossi et al. [1], at the earliest, Thomas Hobbes and his contemporaries tried to “use numerical measures to assess social conditions and identify the cause of mortality, morbidity, and social disorganization in the discipline of social science”.

Rossi et al. [1] define program evaluation as the process of using social research methods to systematically assess programs aimed at “improving social conditions and our individual and collective well-being”, with the goal of providing answers to the stakeholders. Rossi et al. [1] summarize the five domains of evaluation questions and methods that exhibit strong interplays: (1) the need for the programs, (2) program theory and design, (3) program process, (4) program impacts, and (5) program efficiency.

Evaluations in the Field of Computer Science:

The SPEC CPU benchmark suite, known as SPEC CPU [29], is widely recognized as the most renowned benchmark suite for CPU performance evaluation. Throughout its history, six versions of the SPEC CPU benchmark suite have been released, with the latest version being SPEC CPU2017, which can be found in Table 1. The SPEC CPU workloads cover a broad range of CPU-intensive tasks.

The performance evaluation metric used in SPEC CPU is based on the execution time. The reported score of SPEC CPU represents the ratio of its execution time compared to that of a reference machine. The

specific details of a reference machine can be found in Table 1. To ensure the credibility of the results, the overall metrics are calculated as the geometric mean of each respective ratio. Each ratio is based on the median execution time from three runs or the slower of the two runs.

Dongarra et al. [30] proposed the LINPACK benchmark for evaluating high-performance computing (HPC) systems. The LINPACK Benchmark is designed to solve dense linear systems of equations of order n , represented by the equation $Ax = b$. It originated from the development of the LINPACK software package in the 1970s.

The LINPACK benchmark is commonly used to evaluate HPC systems, and the measurement metric is the number of floating-point operations per second (FLOPS). FLOPS represents the count of floating-point operations (FLOPs) performed by the solving algorithm of the LINPACK benchmark, which is calculated as $(2 * n^3 / 3 + 2 * n^2)$ operations divided by the execution time of the benchmark.

As shown in Fig. 14, ImageNet is a significant benchmark in the field of computer vision, consisting of 14,197,122 high-resolution images manually annotated across 21,841 distinct categories, commonly known as ImageNet-21K [31]. These categories encompass a wide range of objects, animals, and scenes. The ILSVRC (ImageNet Large Scale Visual Recognition Challenge) is an annual computer vision competition that focuses on a subset of ImageNet-21K called ImageNet-1K [32]. It aims to evaluate the performance of deep learning models in tasks such as image classification and object detection, providing specific task configurations and evaluation criteria. ImageNet-1K is primarily used for image classification tasks and consists of 1,281,167 training images, 50,000 validation images, and 100,000 test images. The evaluation metrics commonly used in ILSVRC include Top-1 accuracy, which measures the match between the predicted category and the true category of the image, and Top-5 accuracy, which indicates if the true category of the image is among the top five predicted categories by the model.

Evaluations in the field of medicine:

The evaluation in the field of medicine can be traced back to the early medical eras, although there are no documented records. A rigorous modern medical evaluation methodology and system were established as early as 1938 [34]. Clinical trials, with a history spanning over 250 years, are the primary and widely recognized method for medical evaluation. They are defined as experimental designs to evaluate the potential impact of medical interventions on human subjects [35].

Currently, clinical trials based on experimental designs can be categorized into various types, including randomized trials, double-blind trials, prospective trials, and retrospective trials [36].

As illustrated in Fig. 15, Randomized Controlled Trials (RCTs), considered the gold standard for medical evaluation, possess a rigorous and reliable theoretical framework [37]. However, their high

Table 2

The evaluation outcomes of the Intel Xeon Gold 5120T Processor using different CPU benchmarks show significant discrepancies. SPEC CPU2017 includes two sub-suites: SPECrate and SPECspeed. The results derived from SPECrate and SPECspeed are further categorized into two groups, known as floating-point (FP) and int.

Benchmark	Support platform	Workload	Time (s)	Score	Result	Score meaning	
SPECrate	Unix (AIX, HP-UX, Linux, Mac OS X, Solaris), Windows [27]	503.bwaves_r	1483	379.0	96.9 (FP)	Higher scores mean that more work is done per unit of time [27]	
		508.namd_r	636	83.7			
		510.parest_r	1742	84.1			
		511.povray_r	1128	116.0			
		519.lbm_r	1420	41.6			
		521.wrf_r	1682	74.6			
		526.blender_r	787	108.0			
		527.cam4_r	998	98.2			
		538.imagick_r	1479	94.2			
		544.nab_r	893	106.0			
		549.fotonik3d_r	2001	109.0			
		554.roms_r	1429	62.3			
		500.perlbench_r	926	96.3			84.3 (INT)
		502.gcc_r	758	105.0			
		505.mcf_r	1059	85.5			
		520.omnetpp_r	1217	60.4			
		523.xalancbmk_r	786	75.3			
		525.x264_r	1179	83.2			
531.deepsjeng_r	715	89.8					
541.leela_r	1197	77.5					
548.exchange2_r	1338	110.0					
557.xz_r	824	73.4					
SPECspeed	Unix (AIX, HP-UX, Linux, Mac OS X, Solaris), Windows [27]	603.bwaves_s	224	263.0	48.7 (FP)	Higher scores mean that less time is needed [27]	
		619.lbm_s	182	28.8			
		621.wrf_s	522	25.4			
		627.cam4_s	155	57.2			
		628.pop2_s	532	22.3			
		638.imagick_s	507	28.4			
		644.nab_s	191	91.5			
		649.fotonik3d_s	244	37.3			
		654.roms_s	245	64.2			
		600.perlbench_s	832	2.1			3.2 (INT)
		602.gcc_s	823	4.8			
		605.mcf_s	1369	3.5			
		620.omnetpp_s	815	2.0			
		623.xalancbmk_s	444	3.2			
		625.x264_s	703	2.5			
631.deepsjeng_s	651	2.2					
641.leela_s	999	1.7					
648.exchange2_s	807	3.6					
657.xz_s	492	12.6					
PARSEC3.0	Linux/i386, Linux/AMD64, Linux/Itanium, Solaris/Sparc [33]	blackscholes	133		133		
		bodytrack	346		346		
		canneal	258		258		
		facesim	771		771		
		fluidanimate	974		974		
		freqmine	776		776		
		streamcluster	2037		2037		
		swaptions	424		424		
		x264	144		144		
		dedup	58		58		
		raytrace	245		245		
		vips	179		179		
CINT2006 (speed)	Unix (AIX, HP-UX, Linux, Mac OS X, Solaris), Windows [28]	400.perlbench	742	13.2	19.6	Higher scores mean that less time is needed [27]	
		401.bzip2	603	16.0			
		403.gcc	373	21.6			
		429.mcf	283	32.2			
		445.gobmk	567	18.5			
		456.hmmer	433	21.6			
		458.sjeng	880	13.7			
		462.libquantum	409	50.6			
		464.h264ref	983	22.5			
		471.omnetpp	434	14.4			
473.astar	553	12.7					

(continued on next page)

Table 2 (continued).

CFP2006 (speed)		410.bwaves	406	33.5	23.9	Higher scores mean that less time is needed [27]
		433.milc	549	16.7		
		434.zeusmp	400	22.8		
		435.gromacs	334	21.4		
	Unix	436.cactusADM	385	31.1		
	(AIX,	437.leslie3d	229	41.0		
	HP-UX,	444.namd	485	16.5		
	Linux,	450.soplex	278	30.0		
	Mac OS X,	453.povray	276	19.2		
	Solaris),	454.calculix	963	8.57		
	Windows [28]	459.GemsFDTD	354	30.0		
		465.tonto	497	19.8		
		470.lbm	337	40.8		
		482.sphinx3	665	29.3		

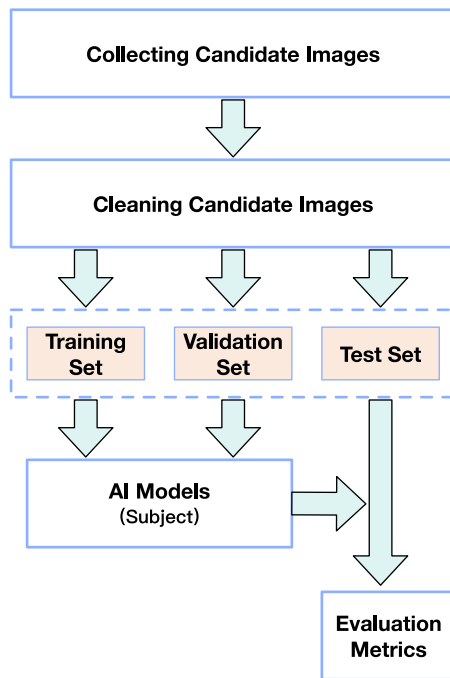


Fig. 14. The ImageNet evaluation working process.

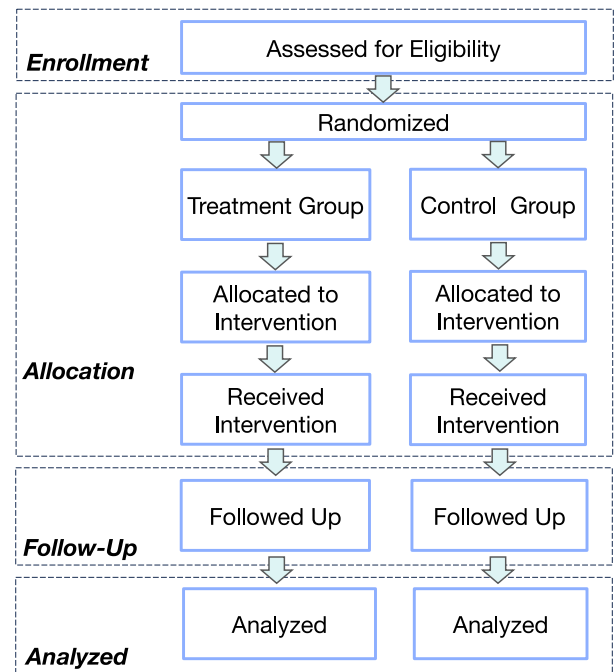


Fig. 15. The randomized controlled trials (RCT) evaluation process [40].

time and financial costs limit their application. To compensate for the shortcomings of RCTs, emerging clinical evaluation methods, such as Real-World Data (RWD) assessment and digital clinical trials, have been proposed [38,39]. These novel medical assessments are still in their early stages and have noticeable deficiencies in their theoretical foundations, such as lacking rigor and reliability.

Evaluations in the field of psychology:

In the field of psychology, social and personality psychologists often rely on scales, such as psychological inventories, tests, or questionnaires [7], to evaluate psychometric variables [7]. These variables include attitudes, traits, self-concept, self-evaluation, beliefs, abilities, motivations, goals, social perceptions, and more [7].

It is important to note that cognitive biases, which are systematic patterns of deviation from norm or rationality in judgment [41], may introduce distortions in self-report style evaluations.

5.2. A case study on CPU benchmarks

Within this scenario, we assess the same CPU utilizing diverse CPU benchmarks, namely SPEC CPU2006 [28], PARSEC 3.0 [33], and SPEC CPU2017 [27], which are proposed by distinct entities employing diverse methodologies.

Employing these benchmark suites, we evaluate the performance of a subject, the Intel Xeon Gold 5120T processor, and proceed to compare the resultant evaluation outcomes. During the experiments, apart from the processor itself, we provide the support system in the following manner: a 384 GB of memory, a 16TB disk, and the utilization of Ubuntu 20.04 as the operating system. To facilitate the compilation process, we employ the GNU Compiler Collection (GCC) version 9.4. We also use the largest data set of each benchmark suite (for SPEC CPU, it is a ‘ref’ data set, and for PARSEC, it is a ‘native’ data set) and run each workload three times for a comprehensive evaluation. For the SPEC CPU2006 benchmark suite, we use the default configuration file of SPECspeed Metric.

The evaluation outcomes are presented in Table 2. The discernible discrepancies observed in the evaluation outcomes can be comprehensively elucidated by taking into account the significant disparities inherent in different benchmark suites. The variations encompass the selection of distinct problem or task instances, the algorithms, the implementation of algorithms, the value functions utilized, the composite metrics employed for evaluation, the reference support system, and the reference subject, which is witnessed by Table 1.

The experimental findings illustrated in Table 2 reveal significant discrepancies in the evaluation outcomes of the same CPU when assessed using different CPU benchmark suites (the SPEC CPU and PARSEC benchmark suites). Furthermore, comparing evaluation outcomes from different evaluators becomes challenging when they employ different benchmarks. There are several reasons as follows. Firstly, they utilize different value functions. Secondly, the SPEC CPU benchmark suite encompasses a reference machine, whereas the PARSEC benchmark suite lacks such a basis for comparison. Moreover, the EECs cannot be ensured even when using the same value function and reference machine. This is because different benchmark suites introduce variations in problem or task instances, algorithms, and the implementation of algorithms.

Moreover, there are significant differences in evaluation outcomes of the same CPU, even when using the same benchmark suite with different versions. It stems from the different implementations of algorithms on varied support systems. Let us take the gcc workload in the SPEC CPU benchmark suite as an example. When evaluated under the CPU2006 benchmark suite, the CPU achieved a score of 21.6 in the 403.gcc workload, while it scored 4.8 in the 602.gcc_s workload under the SPECspeed benchmark of SPEC CPU2017. These scores show a disparity of nearly five-fold. The discrepancies in evaluation outcomes can be mainly attributed to variations in the reference machine used by the respective benchmark suites, as outlined comprehensively in Table 1. Moreover, the two workloads utilize different GCC compiler versions, with the 403.gcc workload utilizing GCC version 3.2 and the 602.gcc_s workload utilizing GCC version 4.5. Although the command flag of both workloads is 'ref,' the input data of the 403.gcc workload consists of nine C-code workloads, while the input data of the 602.gcc_s workload is the preprocessed GCC compiler code. Additionally, 403.gcc is not multi-threaded, while the multi-threaded is permitted for 602.gcc_s.

Furthermore, when being implemented with the same version of GCC compiler and using the same input data, the variances in evaluation outcomes for the gcc workload between the 502.gcc_r workload in SPECrate benchmark suite and 602.gcc_s in SPECspeed benchmark suite are more than twentyfold, which stem from the adoption of disparate value functions and the distinct implementations of the same algorithm. The SPECrate benchmark suite workloads are designed to assess throughput, employing multiple copies of a single-thread implementation during evaluations, while the SPECspeed benchmark suite workloads solely measure execution time, and the utilization of multiple threads is optional throughout the evaluation process. For the evaluation condition, 502.gcc_r workload makes fifty-six copies, running with a single thread, while 602.gcc_s workload has only one copy but runs with fifty-six threads.

The observed variations in evaluation outcomes for a particular CPU across different benchmark suites underscore the necessity of advancing the science and engineering of evaluation. While state-of-the-art CPU benchmarks have made significant progress, they do have certain drawbacks that need to be addressed. Firstly, the lack of comparability among evaluation results from different evaluators is a significant concern. Secondly, the significant discrepancies in evaluation outcomes cannot be traceable. Lastly, state-of-the-art CPU benchmarks often struggle to provide a realistic estimate of the parameters of real-world systems (ES) with a high level of confidence.

According to the comprehensive elements of a benchmark discussed in Section 4.1, many CPU benchmarks, such as the SPEC CPU benchmark suites, primarily focus on the reference implementation of algorithms, metrics, and references while neglecting other essential constituents and components. In the following section, we will carefully analyze and highlight the shortcomings of state-of-the-art and state-of-the-practice evaluation and benchmarks, employing our own terminology.

5.3. The reflections on state-of-the-art and state-of-the-practice benchmarks and evaluation

To further illustrate the limitations of existing evaluation and benchmarking practices, we present Fig. 16, which showcases these shortcomings within the evaluatology framework. By examining this figure, we can gain a clearer understanding of the areas where state-of-the-art and state-of-the-practice evaluation and benchmarks fall short.

It is evident that a lack of consensus exists regarding concepts and terminologies across different areas of study. This lack of consensus often leads to confusion and misinterpretation, especially when the same terms are used in different disciplines with varying meanings.

For example, the term "benchmark" is commonly employed in computer science, finance, and business disciplines but without a formal definition. Moreover, even within these fields, the definition of "benchmark" can be vague and subject to interpretation. In contrast, psychology may use the term "scale" as a concept similar to benchmark, while social science and medicine may not have an analogous concept at all.

Recognizing this challenge, our work has aimed to propose universal concepts and terminologies that can bridge these disciplinary gaps. By establishing clear and standardized definitions, we seek to promote a shared understanding and facilitate effective communication and collaboration across different areas of study.

Few works discuss the essence of evaluation, let alone reaching a consensus on it. Evaluation is often mistakenly equated with measurement or testing without clear differentiation. For instance, in computer science and psychology, evaluation and measurement are often used interchangeably. In the context of testing, where the goal is to determine whether an individual or a system aligns with the expected behavior defined by test oracles, evaluation is often conflated with testing. For instance, according to the SPEC terminology, a benchmark refers to "a test, or set of tests, designed to compare the performance of one computer system against the performance of others" [42,43]. SPEC is a highly influential benchmark organization. Our work has revealed the essence of the evaluation.

The proposed evaluation theories and methodologies are often domain-specific, with a lack of universally applicable foundational principles and evaluation methodologies that transcend diverse disciplines. Different disciplines do not delve into the underlying principles of evaluation. Instead, they adopt a pragmatic approach and prioritize guidelines for conducting evaluations within specific contexts.

For instance, in the medical discipline, the focus is primarily on eliminating confounding variables within the specific groups or cohorts being studied. In the business discipline, efforts are concentrated on searching the state of the practice.

The most rigorous theoretical foundation can be found in the field of clinical trials. For instance, Randomized Controlled Trial (RCT) techniques are employed to rule out the effect of confounding variables. However, there is a lack of universal problem formulations or fundamental solutions that fully consider the intricate interactions among the key components of EMs in diverse scenarios.

There are two serious drawbacks to the RCT methodology and its variants. Firstly, there is a lack of a stringent hierarchical definition of EC and EECs. The variations in ECs can introduce confounding that may affect the results and make meaningful comparisons difficult. Without ensuring EECs, it becomes an illusion to expect comparable evaluation outcomes.

Secondly, when it comes to studying complex systems such as human beings or experimental animals, which we refer to as support systems, the RCT methodology and its variants may struggle to establish an REM. This kind of support system is characterized by a multitude of independent variables, making it difficult to isolate and control all relevant factors in a controlled experimental setting. Consequently, it becomes challenging to eliminate confounding variables and ensure unbiased evaluation outcomes completely.

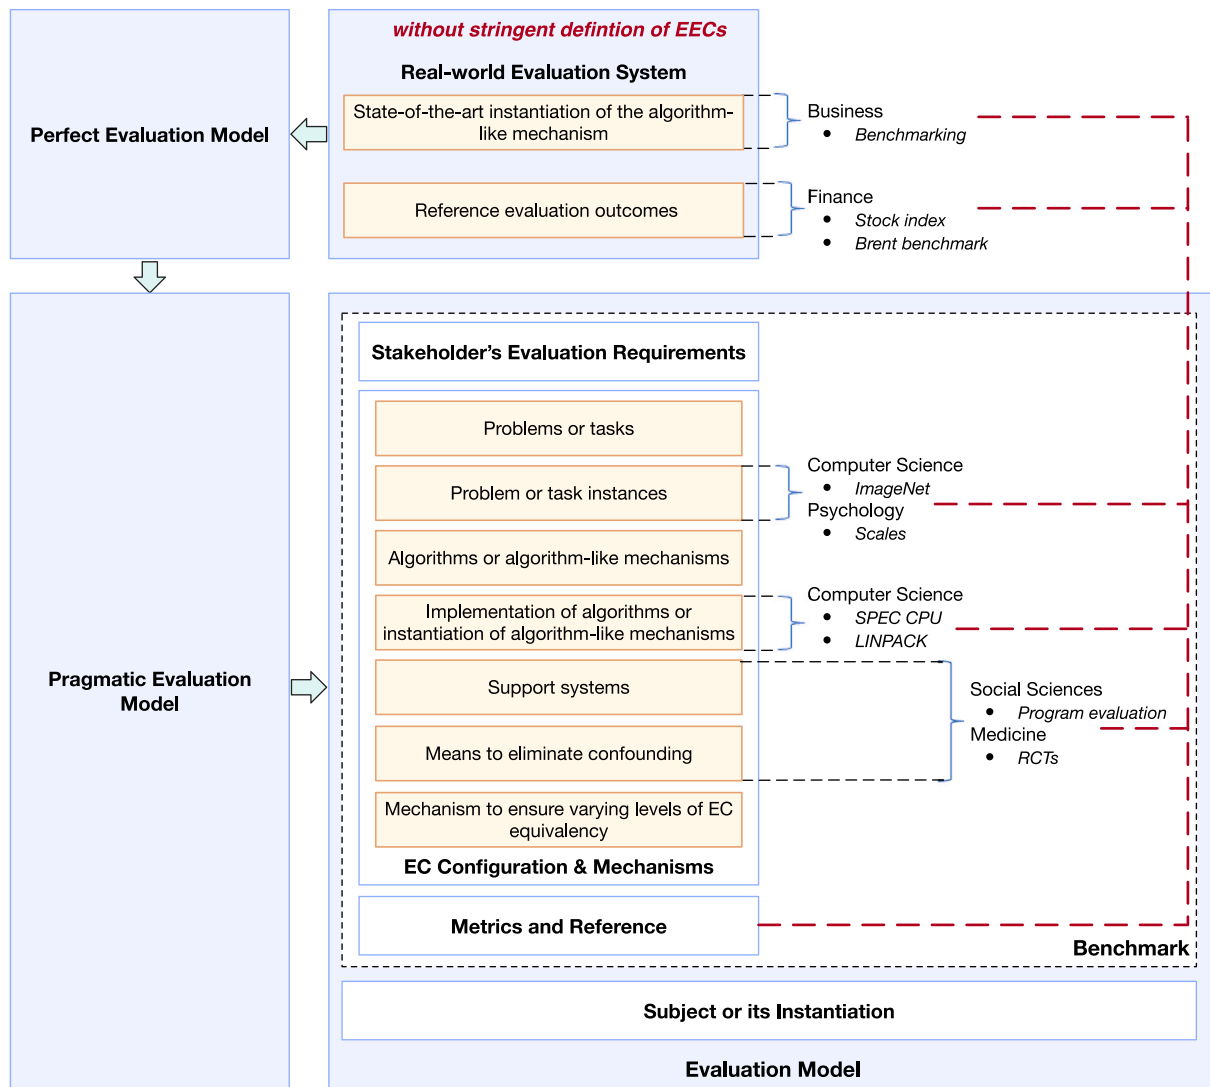


Fig. 16. The reflections on state-of-the-art and state-of-the-practice benchmarks and evaluation are based on the science and engineering of evaluation.

In the realms of business and finance, different observational study methodologies are widely used, and we have revealed their inherent limitations in Section 2.1. An observational study is not even an experiment. Certainly, it cannot eliminate confounding variables and reveal the cause-and-effect relationships. In the business discipline, benchmarking assumes the state-of-the-art instantiation of the algorithm-like mechanism and the reference evaluation outcomes. In finance and education disciplines, benchmarks or indexes assume the role of reference evaluation outcomes in an observational study that measures variables of interest but does not attempt to influence the response [2].

Rossi et al. [1] propose a valuable framework for evaluating methodologies in the field of social science. However, they do not provide a universal theory that can be applied to different disciplines. Their limitations stem from their narrow focus on assessing social programs without developing a generalized theory for evaluating other subjects in complex conditions.

Rossi et al. indeed utilized or developed some approaches to isolate the social programs' impacts, e.g., comparison group designs and randomized controlled trials (RCT), but they failed to explicitly state the underlying principles and methodology for universal science and engineering of evaluation.

Within the computer science field, there are varying viewpoints and perspectives. For example, Hennessy et al. [4] highlight the significance

of benchmarks and define them as programs specifically selected for measuring computer performance. On the other hand, John et al. [3] compile a book on performance evaluation and benchmarking without providing formal definitions for these concepts. Kounev et al. [42] present a formal definition of benchmarks as “tools coupled with methodologies for evaluating and comparing systems or components based on specific characteristics such as performance, reliability, or security”. The ACM SIGMETRICS group [5,6] considers performance evaluation as the generation of data that displays the frequency and execution times of computer system components, with a preceding orderly and well-defined set of analysis and definition steps.

In psychology, social and personality psychologists often utilize scales, such as psychological inventories, tests, or questionnaires, to assess psychometric variables [7,7]. While these tools are commonly used, it is important to recognize that they rely on virtual assessments and self-report-style evaluations, which may introduce potential distortions.

To overcome this limitation, we suggest implementing a physical application of an EC to the subjects, supplemented with a variety of measurement instruments. This approach aims to provide a more objective and accurate assessment of various aspects, including attitudes, traits, self-concept, self-evaluation, beliefs, abilities, motivations, goals,

and social perceptions [7], by incorporating tangible and observable data.

Various disciplines have proposed engineering approaches to evaluations. However, they fail to provide universal benchmark concepts, theories, principles, and methodologies.

For instance, benchmarks are commonly utilized in finance, computer science, and business, albeit with inconsistent meanings and practices. Regrettably, there have been limited discussions in previous works regarding universal benchmark principles and methodologies that can be applied across different disciplines. From a computer science standpoint, Kounev et al. [42] provide a comprehensive foundation for benchmarking, including metrics, statistical techniques, experimental design, and more.

Most state-of-the-art and state-of-the-practice benchmarks overlook an essential aspect: the stakeholders' evaluation requirements. This oversight leads to a failure to consider different and diverse evaluation requirements. For instance, they do not enforce the discrepancy threshold in evaluation outcomes, nor do they consider evaluation confidence, among other crucial factors. As a result, most CPU benchmarks are ill-equipped to meet the evaluation requirements in scenarios involving safety-critical, mission-critical, and business-critical applications.

Another issue is the lack of a stringent definition for similar concepts, such as an EECs or LEECs. For example, most CPU or AI (deep learning) benchmarks, like ImageNet, fail to provide a clear definition of an EECs or LEECs. Instead, they jump directly into the implementation of algorithms or a specific dataset labeled with the ground truth without proper justification. Additionally, the support system, which plays a crucial role in some cases, is omitted without any explanation of the condition of simplifying the benchmarks. Furthermore, most of the methodologies fail to discuss the confounding elimination mechanism. This oversight can potentially introduce bias and inaccuracies in the evaluation outcomes.

Not surprisingly, the intricate evaluation mechanisms and policies introduced in Section 4.4.2 are not explicitly discussed in the design and implementation of most benchmarks. For instance, it fails to address important aspects such as investigating and characterizing real-world ES, the design and implementation of a perfect EM, the modeling policy and procedure from a real-world ES to an EM, and the sampling policy and procedure from a perfect EC to a pragmatic EC. This omission makes it difficult for the benchmark to adapt to intricate evaluation scenarios.

It is crucial to include these mechanisms and policies to ensure the benchmark's applicability and effectiveness in complex evaluation scenarios. Without explicit discussion of the real-world ES, it is difficult to establish an EC that captures the characteristics and requirements of real-world evaluations. Furthermore, exploring different sampling and modeling policies is essential to gain the confidence of the evaluation community in using the benchmark for inferring parameters of real-world ES. By carefully designing these policies, we can strike a balance between achieving high accuracy in evaluation outcomes and managing the associated evaluation costs.

There are many widely used AI (deep learning) benchmarks. Taking the ImageNet dataset as an illustrative example [31], we reveal their limitations. Firstly, a specific AI benchmark like ImageNet cannot be traced back to an explicit formulation of a problem or task and instead manifests itself in the form of a dataset containing ground truth, which may possess certain biases. In other scenarios, we also encounter challenges in identifying a precise mathematical function that accurately models the chemical and biological activities within the human body (Case Three in Section 3.5.2) or the social dynamics within the target population (Case Four in Section 3.5.2). Secondly, the benchmark relies on an unverified assumption that the data distribution within the real world closely aligns with that of the collected dataset to a considerable extent. Thirdly, in real-world applications, we use the statistic of a sample – a specific benchmark – to infer the parameters of the entire population. However, we do not know their confidence levels and intervals.

5.4. What is the benefit of evaluatology?

Evaluatology serves as the foundational theory that encompasses evaluations in various fields of study. It provides a universal framework for optimizing the evaluation process, with four fundamental axioms serving as its basis. Formulating the core evaluation issues mathematically presents opportunities for seeking optimal solutions based on theoretical grounds. As a subdivision of evaluatology, benchmarkology offers a comprehensive engineering approach and methodology for evaluation, applicable across diverse disciplines.

Together, evaluatology and benchmarkology contribute reusable knowledge to different domains, encompassing universal terminology, principles, and methodologies. By sharing this knowledge base, they facilitate advancements in both the state-of-the-art and state-of-the-practice of evaluation in various realms. This unification of communities embarks on a collective journey to address future challenges.

5.5. The differences between evaluation, measurement and testing

Drawing on the preceding analysis, this subsection elucidates the marked disparity between evaluation, measurement, and testing.

First and foremost, it is important to acknowledge that measurement or testing serves as a preliminary constituent within the broader framework of evaluation. In addition to measurement and testing, an evaluation encompasses a series of steps, which we have discussed in Section 3. These steps involve defining and applying evaluation conditions to a diverse range of subjects, which ultimately leads to the creation of an evaluation model or system. Once the evaluation model or system is established, the impacts of different subjects can be inferred through the process of measuring and/or testing.

Furthermore, it is crucial to recognize that the measurement results are of an objective nature, assuming the existence of an inherent truth value for each measured quantity. Similarly, testing results also possess an objective nature as they typically yield either a positive or negative outcome for each test conducted.

Conversely, evaluation results possess a certain degree of subjectivity, such as the formulation of value functions based on the underlying measurement data, which we have discussed in the first evaluation axiom in Section 3.3.

By virtue of the aforementioned reasons, we can assert that metrology or testing serves as but one foundational aspect in the realm of evaluations.

6. Conclusion

This article formally introduces evaluatology, a discipline encompassing both the science and engineering of Evaluations. Our contributions are three-fold.

First, in order to promote consistency and facilitate cross-disciplinary understanding, we propose the adoption of universal evaluation concepts and terminologies centered around evaluation conditions.

Secondly, we reveal the essence of evaluation and propose five evaluation axioms as the foundational evaluation theory. Furthermore, we introduce the universal evaluation theory, principles, and methodology that govern the field of evaluation.

We create evaluation conditions with different levels of equivalency and apply them to diverse subjects to establish reference evaluation models that alter a single independent variable at a time while keeping all other variables as controls. We discover that the key to effective and efficient evaluations in various complex scenarios lies in establishing a series of evaluation models that maintain transitivity.

Third, building upon the science of engineering, we formally define a benchmark as a simplified and sampled evaluation condition that ensures different equivalency levels. We present a benchmark-based universal engineering of evaluation across different disciplines, which we refer to as benchmarkology.

CRedit authorship contribution statement

Jianfeng Zhan: The science and engineering of evaluation, encompassing the universal evaluation concepts, terminologies, theory and methodology, the universal benchmark concepts, principles, theories, and methodologies, and mathematical formulation of fundamental evaluation issues. **Lei Wang:** Whole-session discussion, the axiom of the true evaluation results, the CPU benchmark experiments in Section 5.2, the summary of the related work for finance and HPC benchmarks in Section 5.1, and the mathematical formulation of fundamental evaluation issues. **Wanling Gao:** Presentations of all figures, whole-session discussion, and the discussion of the benchmark inspired the first author to think about what is essential of the benchmark. **Hongxiao Li:** Whole-session discussion, mathematics notations, mathematical formulations of three fundamental issues in Evaluatology, and mathematical formulations of EECs and LEECs. **Chenxi Wang:** Whole-session discussion, design, implementation, analysis, and presentation of experiments in Section 5.2. **Yunyou Huang:** Discussion, and the summary of the related work of drug evaluations in Section 5.1. **Yatao Li:** Discussion, and the mathematical formulations of evaluation traceability in Section 3.7.3. **Zhengxin Yang:** Discussion and presentation of metrology, and real-world evaluation of safety-critical AI systems. **Guoxin Kang:** Discussions, and related work about business benchmarking in Section 5.1. **Chunjie Luo:** Discussions. **Hainan Ye:** Discussions. **Shaopeng Dai:** Presentations of some figures. **Zhifei Zhang:** Discussions, and the summary of related work of drug evaluations.

Declaration of competing interest

The authors declare that they have no known competing financial interests or personal relationships that could have appeared to influence the work reported in this paper.

References

- [1] P.H. Rossi, M.W. Lipsey, G.T. Henry, *Evaluation: A Systematic Approach*, Sage publications, 2018.
- [2] D.S. Starnes, D. Yates, D.S. Moore, *The Practice of Statistics*, Macmillan, 2010.
- [3] L.K. John, L. Eeckhout, *Performance Evaluation and Benchmarking*, CRC Press, 2018.
- [4] J.L. Hennessy, D.A. Patterson, *Computer Architecture: A Quantitative Approach*, Elsevier, 2011.
- [5] J.C. Browne, An analysis of measurement procedures for computer systems, *ACM SIGMETRICS Perform. Eval. Rev.* 4 (1975) 29–32.
- [6] M.E. Knudson, A performance measurement and system evaluation project plan proposal, *ACM SIGMETRICS Perform. Eval. Rev.* 13 (1985) 20–31.
- [7] M. Furr, Scale construction and psychometrics for social and personality psychology, *Scale Constr. Psychometr. Soc. Pers. Psychol.* (2011) 1–160.
- [8] R.C. Camp, *Benchmarking: The Search for Industry Best Practices that Lead to Superior Performance*, Asq Press, 1989.
- [9] L. Fisher, Some new stock-market indexes, *J. Bus.* 39 (1966) 191–225.
- [10] H.D. Young, R.A. Freedman, L.A. Ford, *University physics with modern physics*, 2020.
- [11] J. Stewart, *Single Variable Calculus: Concepts and Contexts*, Cengage Learning, 2018.
- [12] System, 2024, <https://www.merriam-webster.com/dictionary/system>. (Accessed 6 February 2024).
- [13] A. Backlund, The definition of system, *Kybernetes* 29 (2000) 444–451.
- [14] I. BiPM, I. IFCC, I. IUPAC, O. ISO, *The international vocabulary of metrology—basic and general concepts and associated terms (vim)*, Joint Comm. Guides Metrol. 200 (2012) 2012.
- [15] R.N. Kacker, On quantity, value, unit, and other terms in the jcgim international vocabulary of metrology, *Meas. Sci. Technol.* 32 (2021) 125015.
- [16] S.S. Stevens, On the theory of scales of measurement, *Science* 103 (1946) 677–680.
- [17] L. Baresi, M. Young, *Test oracles*, 2001.
- [18] D. Choudhary, V. Kumar, *Software testing*, *J. Comput. Simul. Model.* 1 (2011) 1.
- [19] J.A. Whittaker, What is software testing? And why is it so hard? *IEEE Software* 17 (2000) 70–79.
- [20] P.J. Denning, The science of computing: What is computer science? *Am. Sci.* 73 (1985) 16–19.
- [21] D.S. Johnson, A catalog of complexity classes, in: *Algorithms and Complexity*, Elsevier, 1990, pp. 67–161.
- [22] S.B. Cooper, *Computability Theory*, Chapman and Hall/CRC, 2017.
- [23] R. Sedgewick, P. Flajolet, *An Introduction to the Analysis of Algorithms*, Addison-Wesley Longman Publishing Co. Inc, 1996.
- [24] 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.
- [25] B. Andersen, T. Fagerhaug, S. Randmæl, J. Schuldmaier, J. Prenninger, *Benchmarking supply chain management: finding best practices*, *J. Bus. Ind. Mark.* 14 (1999) 378–389.
- [26] D.P. Scheitrum, C.A. Carter, C. Revoredo-Giha, Wti and brent futures pricing structure, *Energy Econ.* 72 (2018) 462–469.
- [27] SPEC, SPEC CPU2017 benchmark suite, 2017, <https://www.spec.org/cpu2017/>.
- [28] SPEC, SPEC CPU2006 benchmark suite, 2006, <https://www.spec.org/cpu2006>.
- [29] SPEC, SPEC CPU Benchmark Suite. <https://www.spec.org/benchmarks.html#cpu>.
- [30] J.J. Dongarra, P. Luszczek, A. Petitet, The linpack benchmark: past, present and future, *Concurr. Comput. Pract. Exp.* 15 (2010) 803–820.
- [31] 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, pp. 248–255.
- [32] O. Russakovsky, J. Deng, H. Su, J. Krause, S. Satheesh, S. Ma, Z. Huang, A. Karpathy, A. Khosla, M. Bernstein, et al., Imagenet large scale visual recognition challenge, *Int. J. Comput. Vis.* 115 (2015) 211–252.
- [33] C. Bienia, S. Kumar, J.P. Singh, K. Li, The parsec benchmark suite: Characterization and architectural implications, in: *Proceedings of the 17th International Conference on Parallel Architectures and Compilation Techniques*, 2008, pp. 72–81.
- [34] D.F. Cavers, The food, drug, and cosmetic act of 1938: its legislative history and its substantive provisions, *Law Contemp. Probs.* 6 (1939) 2.
- [35] J. Jenkins, S. Hubbard, *History of clinical trials*, in: *Sem. Oncol. Nurs.*, 1991, pp. 228–234.
- [36] J.W. Sellers, C.M. Mihaescu, K. Ayalew, P.D. Kronstein, B. Yu, Y.M. Ning, M. Rodriguez, L. Williams, N.A. Khin, Descriptive analysis of good clinical practice inspection findings from us food and drug administration and european medicines agency, *Ther. Innov. Regul. Sci.* 56 (2022) 753–764.
- [37] S. Monti, V. Grosso, M. Todoerti, R. Caporali, Randomized controlled trials and real-world data: differences and similarities to untangle literature data, *Rheumatology* 57 (2018) vii54–vii58.
- [38] O. Inan, P. Tenaerts, S. Prindiville, H. Reynolds, D. Dizon, K. Cooper-Arnold, M. Turakhia, M. Pletcher, K. Preston, H. Krumholz, et al., Digitizing clinical trials, *NPJ Digit. Med.* 3 (2020) 101.
- [39] S.V. Ramagopalan, A. Simpson, C. Sammon, Can real-world data really replace randomised clinical trials? *BMC Med.* 18 (2020) 1–2.
- [40] D. Moher, K.F. Schulz, D.G. Altman, C. Group*, The CONSORT statement: revised recommendations for improving the quality of reports of parallel-group randomized trials, *Ann. Intern. Med.* 134 (2001) 657–662.
- [41] BIAS, FOC, The evolution of cognitive bias, *Handb. Evol. Psychol.* (2015) 2.
- [42] S. Kounev, K.D. Lange, J.Von. Kistowski, *Systems Benchmarking*, Springer, 2020.
- [43] SPEC, SPEC glossary, 2023, <https://www.spec.org/spec/glossary>.



Dr. Jianfeng Zhan is a Full Professor at the Institute of Computing Technology (ICT), Chinese Academy of Sciences (CAS), and University of Chinese Academy of Sciences (UCAS). He is also the director of the Research Center for Distributed Systems at ICT, CAS. He holds a B.E. in Civil Engineering and an MSc in Solid Mechanics from Southwest Jiaotong University, obtained in 1996 and 1999 respectively. He completed his Ph.D. in Computer Science from the Institute of Software, CAS, and UCAS in 2002.

Dr. Zhan's research encompasses a wide range of areas, including chips, systems, and benchmarks. His work involves benchmarking, designing, implementing, and optimizing diverse systems. He has successfully translated his academic research into advanced technologies that have had a significant impact on general-purpose production systems. His team's innovations and research results, including 35 patents, have been adopted in benchmarks, operating systems, and cluster and

cloud system software, contributing to the advancement of parallel and distributed systems in China and globally.

Dr. Zhan is the founder and chair of the International Open Benchmark Council (BenchCouncil), a prominent organization in the field of benchmark and Evaluation. He also serves as the Co-Editor-in-Chief of TBench alongside Prof. Tony Hey. From 2018 to 2022, Dr. Jianfeng Zhan served as an Associate Editor for the IEEE TPDS journal. He has received several accolades for

his work, including 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.

Throughout his career, Dr. Zhan has supervised numerous graduate students, post-doctoral researchers, and engineers, with a total of over ninety individuals under his guidance.