

# Establishing Construct Validity in LLM Capability Benchmarks Requires Nomological Networks

Timo Freiesleben\*

Munich Center for Mathematical Philosophy (MCMP)

LMU Munich

## Abstract

Recent work in machine learning increasingly attributes human-like capabilities, such as reasoning, to large language models (LLMs) based on benchmark performance. This paper examines that practice through the lens of construct validity, understood as the challenge of linking theoretical capabilities to their empirical measurement. It compares three influential frameworks, the nomological, inferential, and causal accounts, and argues that the nomological account provides the most suitable foundation for LLM capability research. The nomological account avoids the strong ontological assumptions of the causal account while offering a more substantive framework for defining construct meaning than the inferential account.

**Keywords:** *Construct Validity; Measurement; Benchmarking; Machine Learning; Large Language Models*

*“Measure what can be measured, and make measurable what cannot be measured.”*

— Galileo Galilei

---

\*Corresponding authors: [timo.freiesleben@lmu.de](mailto:timo.freiesleben@lmu.de).

# I Introduction

Humans routinely attribute capabilities to one another. We say things like 'she is a sharp reasoner' or 'he has good moral intuitions'. Intuitively, we grasp what such claims mean: we recognize the patterns of behavior they point to and how these patterns tend to cohere with other capabilities a person may or may not possess. For example, by definition, we expect someone with strong reasoning skills to excel at solving mathematical problems or navigating complex legal cases. Beyond this, we also associate reasoning with related capabilities, such as broad knowledge or strong memory.

To measure complex capabilities like reasoning, psychologists and educational scientists have developed various tests, which typically consist of a range of concrete tasks like completing math sequences, logical syllogisms, or analogies. While such tests are not uncontroversial, evidence suggests that they provide a meaningful measure of what we mean by reasoning [55].

Machine learning researchers have recently begun to transfer this capability-based language to the context of large language models (LLMs) [17]. They test LLMs for reasoning capability [57], theory of mind [49], or even morality [23]. They do so in close analogy with human assessment, namely by designing concrete tests in the form of benchmarks, i.e. structured evaluation suites for machine learning models [16]. Attributing substantial capabilities such as reasoning to LLMs is risky: Humans naturally interpret terms like reasoning in the same way they interpret them when applied to other humans, drawing broad generalizations about behavior and inferring the presence of related capabilities [54]. How can we be confident that our benchmark tests genuinely measure the relevant capability and thus warrant intuitive interpretation by human users?

In the context of human testing, a range of well-established standards for test design aims to ensure that tests measure the intended capability, known under the name *construct validity*. The central task is to connect a latent construct, for example reasoning capability, with a concrete measurement, such as a reasoning test. Researchers investigating LLM capabilities have so far largely ignored these standards in designing their benchmarks and are only gradually becoming aware of their existence [4, 45, 1]. But which philosophical framework of construct validity should LLM researchers adopt?

Within philosophical psychology, three highly influential accounts of construct validity can be distinguished [11]. They differ in their ontological commitments regarding the status of constructs, as well as in their views on how measurements should be validated and what kinds of inferences they support. The first is the *nomological account* proposed by Cronbach and Meehl in 1955 [12]. On this view, a construct is defined by its role within a nomological web of relations to other constructs. Validity is established by testing whether measurements of the construct correlate with measurements of other constructs as specified by the nomological web; i.e. it takes a coherentist perspective. The second is the *inferential account* advanced by Messick in 1989 [35, 36] and later refined by Kane [26], which, rather than asking whether a measurement truly captures a construct, understands construct validity in terms of the evidence and assumptions that link a measurement to its interpretation and use. The most recent is the *causal account* defended by Borsboom in 2004 [7, 6]. This view requires the ontological existence of the construct of interest and demands that it plays a causal role in producing

the observed behavior.

In this paper, I argue that the nomological account developed by Cronbach and Meehl provides the most suitable theoretical framework for current LLM capability research (Section 2). Compared to Borsboom’s causal account, it avoids strong ontological commitments about the real existence of constructs, commitments that are difficult to defend in the case of LLMs. Unlike the inferential account, it offers concrete guidance on how to theorize complex constructs. I then analyze the conceptual implications of adopting the nomological account in LLM research by applying it to the assessment of reasoning capabilities (Section 3).

## 2 The nomological account is the best available framework for conceptualizing LLM capabilities

There are three major accounts of construct validity in psychological and educational testing [11] that are potential candidates for evaluating LLM capabilities: the nomological account, the inferential account, and the causal account. In the following, I analyze the suitability of each as a framework for construct validity in LLM capability testing in reverse historical order.

### 2.1 The causal account is ontologically demanding for LLMs

Borsboom’s causal account of construct validity is conceptually straightforward. For a measurement to validly measure a construct, two conditions must be satisfied: the construct must *exist*, and variations in measurement outcomes must be *caused* by variations in the state of that construct [7, 6]. On this view, construct validity is established by demonstrating a genuine causal mechanism linking the construct to test scores. Validity thus becomes a scientific claim about underlying causal structure, grounded in a realist ontology of psychological attributes.

This ontological perspective may be defensible in the context of human cognition, where constructs such as reasoning or theory of mind are at least plausibly connected to biological and cognitive mechanisms [5].<sup>1</sup> For LLMs, however, this realist stance is difficult to defend. It is unclear what it would mean for a model to *possess* reasoning, theory of mind, or morality as internal attributes, rather than merely to *exhibit behavioral patterns* that resemble those associated with these constructs in humans. In any case, inferring the existence of a construct from observed behavioral patterns requires particularly strong evidence [26]. The burden would fall on researchers defending the causal account to provide this evidence.

One possible response from defenders of the causal account is to appeal to mechanistic interpretability research [47] and argue that capabilities are realized in identifiable internal components of LLMs, thereby supporting their existence in Borsboom’s realist sense. On this view, if specific circuits or subnetworks can be shown to causally produce particular behavioral regularities, then constructs

---

<sup>1</sup>Even in the case of humans, the ontological status of constructs remains controversial [2, 38].

correspond to genuine causal structures within the model. However, the success of mechanistic interpretability in isolating psychologically meaningful constructs, especially higher-level constructs, often remains fragmentary [28, 15, 37]. Existing results mainly identify mechanisms for relatively narrow behaviors, e.g. related to concrete locations like Paris [34], objects like the Golden Gate Bridge [52], or mathematical operations like addition [27]. While such examples may provide partial support for a realist interpretation, it remains unconvincing unless mechanistic interpretability demonstrates robust mappings between internal structures and behavioral patterns associated with complex capabilities such as reasoning.

Alternatively, proponents of the causal account may conclude that if no construct exists in their strong realist sense, then no valid measurement of it is possible. On this view, if the causal account cannot be applied, we should refrain from speaking of measuring capabilities through benchmarks altogether; capability ascriptions would then be reserved for humans. This position, however, seems overly restrictive. The practical relevance of capability language in the context of LLMs is not tied to metaphysical realism about constructs, but to the usefulness of such language for humans in describing, predicting, and evaluating LLM behavior.

Finally, defenders of the causal account may adopt a weaker form of causal realism. Instead of assuming the existence of capabilities within the LLM, they may identify LLM capabilities with their corresponding behavior. However, different behaviors generally do not *cause* one another. An LLM’s behavior in theory of mind tasks may be associated with its behavior in reasoning tasks, but it is certainly not caused by it. Stripped of its causal interpretation, theorizing constructs through structural causal models [6] reduces to Bayesian networks [8]. In this weaker form, causal realism ultimately theorizes constructs in much the same way as the nomological account [9].

## **2.2 The inferential account provides limited resources to theorize LLM capabilities**

Messick’s and Kane’s inferential conception of construct validity differs fundamentally from Borsboom’s causal approach. Rather than asking the ontological question whether a measurement is caused by a latent but real construct, they focus on the epistemological question whether inferences drawn from a measurement are warranted. Because the validity of inferences depends on many factors, including background assumptions, related evidence, and the ethical consequences of score use, Messick argues that construct validity “undergirds all score-based interpretations” [35]. The central goal becomes justifying interpretive claims about what test scores mean and how they may be used.

This shift has a clear advantage over the causal account: it avoids metaphysical commitments regarding the existence of constructs. Instead of asking whether a construct is real, Messick and Kane ask whether score interpretations are adequately supported by evidence. What in the causal framework requires ontological justification can here be defended pragmatically. However, as argued by Borsboom [7] and Lissitz et al. [30], this expansion renders construct validity extremely broad. The measurement process becomes intertwined with the ethical and practical consequences of test use, making it difficult

to determine where construct validity is actually established, as it becomes dispersed across a wide range of inferential considerations. More importantly, the inferential account offers limited resources for even theorizing the construct. Because it focuses primarily on downstream inferences, the question of what exactly is being measured remains comparatively undertheorized. This limitation is particularly problematic for complex and abstract constructs such as reasoning, theory of mind, or morality, which LLM researchers aim to assess.

For complex constructs, meaning cannot be established solely through inferential practices—a point even acknowledged by inferentialists such as Kane [26]. Kane distinguishes between observable attributes and theoretical constructs, noting that inferences about theoretical constructs typically require nomological networks (introduced in Section 2.3) as their theoretical foundation. Accordingly, it has been proposed that inferentialist, causal, and nomological accounts should be seen as complementary rather than fundamentally distinct approaches to construct validity, with inferentialist accounts focusing on the pragmatic dimension, nomological networks providing the theoretical grounding, and causal accounts offering a causal interpretation where possible [22, 26]. Recent frameworks for benchmark validation in machine learning have largely adopted an inferential perspective, but they either focus on narrower constructs, such as weather forecasting [16], or integrate nomological networks into their inferential framework to theorize more complex constructs [45].

### **2.3 Nomological networks as a middle ground between causal realism and pure pragmatism**

Cronbach and Meehl originally introduced construct validity through the notion of a nomological network: a structured web of theoretical constructs, observable variables, and empirically testable relations. On this view, a construct is defined neither by its real-world referent, as in the causal account, nor by isolated inferential claims, as in the inferential account, but by its position within a coherent system of laws, correlations, and theoretical expectations.

The motivation for this proposal lies in the tradition of logical empiricism, which moved away from rigid *point-by-point* definitions toward a more holistic conception of theory. Quine’s famous *Web of Belief* illustrates this idea: the meaning of any single term is determined by its inferential relations to other terms and its distance from the sensory periphery [42, 41]. To revise the definition of *reasoning* in an LLM setting is therefore to adjust its relations to neighboring constructs such as memory or pattern recognition. Similarly, Hempel described scientific theories as complex networks of interconnected elements [19]. Theoretical constructs function as *knots*, while logical and law-like relations serve as the *threads* connecting them. Although the network is anchored in observable data, the constructs themselves derive their identity from their functional role within the system.

Because meaning is defined relationally within a nomological network, the ontological critique against the causal account does not apply. Nevertheless, Borsboom contends that avoiding metaphysical commitments comes at a cost: instead of providing *reference* to a real but unobservable entity, the nomological account merely provides *meaning* by situating terms within a semantic web. He

argues that this project is bound to fail because “Few, if any, theoretical terms in psychology can be unambiguously identified in this way.” [7] However, LLM research suggests that in many practical contexts terms can be characterized relationally [40]. One might respond that Borsboom’s critique targets the relatively simple nomological networks employed in psychological practice. But if so, the limitation lies not in the nomological approach itself, but in its empirical implementation. While constructing rich nomological networks in psychology is constrained by the difficulty of testing the same participants across many constructs, benchmarking additional tasks on the same LLM scales easily.

Above, I argued that the inferential account presupposes a sufficiently articulated construct to anchor its interpretive argument; it offers comparatively little guidance for theorizing the construct in the first place. By contrast, Cronbach and Meehl’s nomological approach fills this gap by systematically embedding a construct within a theoretical network of relations. Rather than focusing on the *inferences* warranted by benchmark scores, the nomological account places the *meaning* of the construct itself at the center of attention. For emerging domains such as LLM capability research, where the meaning of constructs such as reasoning remains contested, this difference in emphasis is crucial. In such contexts, prioritizing the clarification of construct meaning is the more promising strategy. However, if LLM capabilities become more extensively theorized in the future and attention shifts to the legal or ethical implications of benchmark scores, the nomological account would benefit from incorporating insights from the inferential tradition [25, 22, 26].

The main challenge lies in constructing a nomological network, especially for complex constructs such as reasoning. Here, the best available scientific theories should guide construction. For psychological constructs, one promising approach is to integrate established nomological models from psychology with additional relational information, for example, by extracting structural relations from LLM semantic vector spaces [31] or learning them directly from data [14]. Relying on nomological networks established in *human* psychology in testing *LLMs* might be seen as problematic. However, it helps ensure that the meaning of the constructs remains aligned with their use in interpersonal contexts (discussed in more depth in Section 3.1).

In sum, Cronbach and Meehl’s nomological network account offers a compelling middle ground for LLM capability testing. It avoids Borsboom’s strong realist commitments while preserving a substantive notion of construct meaning that goes beyond Messick and Kane’s inferentialism. To explore the conceptual implications of the nomological account for LLM research in greater depth, I now turn to a concrete case: the assessment of reasoning capabilities.

### **3 Implications of the nomological account for LLM benchmark design**

In current practice, LLM researchers often begin by designing benchmarks for complex capabilities such as reasoning [57], theory of mind [49], or even morality [23]. In psychology and educational

science, however, directly designing tests for complex constructs is considered “not good practice” [55]. The danger lies in *construct underrepresentation*, i.e., relevant parts of the construct may not be captured by the test [50], a central threat also in current LLM research, where data is often selected based on availability rather than representativeness [51]. To prevent construct underrepresentation, the nomological account recommends first designing a nomological network before constructing any test.

### 3.1 Researchers should design nomological networks prior to benchmarks

Consider the example of *reasoning*. Designing a nomological network for reasoning requires specifying the qualitative relationships between reasoning and simple tasks such as continuing mathematical sequences, as well as other complex capabilities like working memory. In psychological testing, many constructs already have nomological networks supported both theoretically and empirically; reasoning is such an example [55]. Reasoning is commonly embedded in the Cattell–Horn–Carroll (CHC) theory of cognitive abilities [10], widely regarded as the most empirically supported framework in contemporary psychometrics [46, 33]. CHC theory specifies a hierarchical structure of cognitive abilities (see Figure 1) organized into three strata [10]. Stratum I contains narrow *tasks*, Stratum II broad *capabilities*, and Stratum III a general *cognitive ability factor*. Within this framework, reasoning is a capability.

In CHC, the meaning of *reasoning* is partly determined by its relations to other capabilities, such as working memory capacity, processing speed, visual processing, and long-term retrieval. This interrelatedness is not accidental but theoretically structured. The correlation among capabilities is explained in CHC theory by reference to a higher-order construct: the *g*-factor, often interpreted as general intelligence. Although the ontological status of *g* remains debated, most contemporary psychologists accept that modeling a higher-order general factor provides a coherent representation of observed covariance patterns [44]. From a nomological perspective, the key point is not metaphysical commitment to a reified *g*, but the structural constraint it imposes: capabilities are expected to correlate in systematic ways. Tasks, by contrast, are narrower; they are constituents of capabilities. For reasoning, these include inductive reasoning tasks such as sequence continuation, deductive reasoning tasks such as logical syllogisms, or planning tasks such as multistep search problems. While further decomposition is possible, tasks are typically treated as sufficiently specific to be operationalized through concrete tests [12]. In CHC, capabilities are modeled as weighted composites of tasks.

This hierarchical architecture already reveals something philosophically significant that goes beyond current benchmarking practice: reasoning is not identified with any single task but with a structured pattern of relations to narrow tasks and neighboring capabilities. A benchmark that ignores this structure risks mischaracterizing what it claims to measure.

But should LLM researchers adopt CHC theory when testing LLM capabilities? CHC is, after all, a theory of *human* cognition, not of *LLM* cognition. Applying it to LLMs presupposes that its structural relations are transferable across radically different substrates—a non-trivial assumption. I believe that beginning with a theoretically mature and empirically supported human framework

provides a principled starting point and ensures that our technical definitions align with the human tendency to interpret terms such as *reasoning* or *morality* as we do in interpersonal contexts. Nevertheless, a nomological network for LLM reasoning may ultimately diverge from the human CHC model, simply because reasoning in LLMs may occupy a different place within the broader architecture of capabilities than it does in human cognition.

### 3.2 Task benchmarks require empirical and theoretical validation

Assume the LLM researcher settles on a nomological network, here CHC (Figure 1). The second step in the nomological account is to design tests for *tasks* (Stratum I). Consider, for example, the task of *sequence continuation*. It can be described as predicting a symbol based on a sequence of other symbols, for instance, filling the question mark in the sequence 1, 1, 2, 3, 5, 8, 13, ?. Designing a test for the task means specifying a range of sequence continuation problems, the so-called *test set*, consisting of concrete *test items*. Tests also require a notion of error by assigning numerical values to predictions. For example, the answer 21, correctly continuing the Fibonacci sequence, could be assigned an error of 0, whereas any other answer could be assigned an error of 1. Furthermore, tests require test-takers and an overall test score obtained by averaging the errors across all test items.

These components of psychological tests match closely the components of machine learning benchmarks [16, 18]. However, although psychological tests and benchmarks are structurally similar, the validity theory surrounding psychological testing is substantially more mature. This becomes particularly evident when interpreting test scores: LLM researchers often interpret benchmark scores naively as task performance, whereas psychologists have developed a rich repertoire of conditions required for this inferential step.

The only validity conditions well known in LLM research are statistical in nature [18]: 1) test items should be sampled independently and identically (i.i.d.) from an underlying distribution, or at least form a representative sample of a population [21]; 2) the test set should be independent from the test-takers; 3) the test set should be sufficiently large. These conditions are also recognized in the nomological account, where they constitute one aspect of *content validity*<sup>2</sup> and safeguard against selecting biased samples.

However, content validity in Cronbach and Meehl’s account encompasses more than statistical criteria. It also requires theoretical justification for the chosen task formalizations, test items, and error definitions, considerations that should likewise be addressed in LLM benchmarks. Beyond content validity, Cronbach and Meehl recommend evaluating the *concurrent validity* of tests. Translated to benchmarks, this means that if two independently constructed benchmarks both claim to measure sequence continuation with high content validity, their scores should correlate strongly. Recently, increasing attention has been paid to this issue within machine learning under the label *external validity* [43, 18, 16]. Similarly, when benchmark scores serve as proxies of real-world abilities, for

---

<sup>2</sup>In other accounts they appear under different names—for example, as *generalization* in Kane’s framework [26], *statistical assumptions* in Borsboom’s framework [6], or *internal validity conditions* in Freiesleben and Zetzka [16].

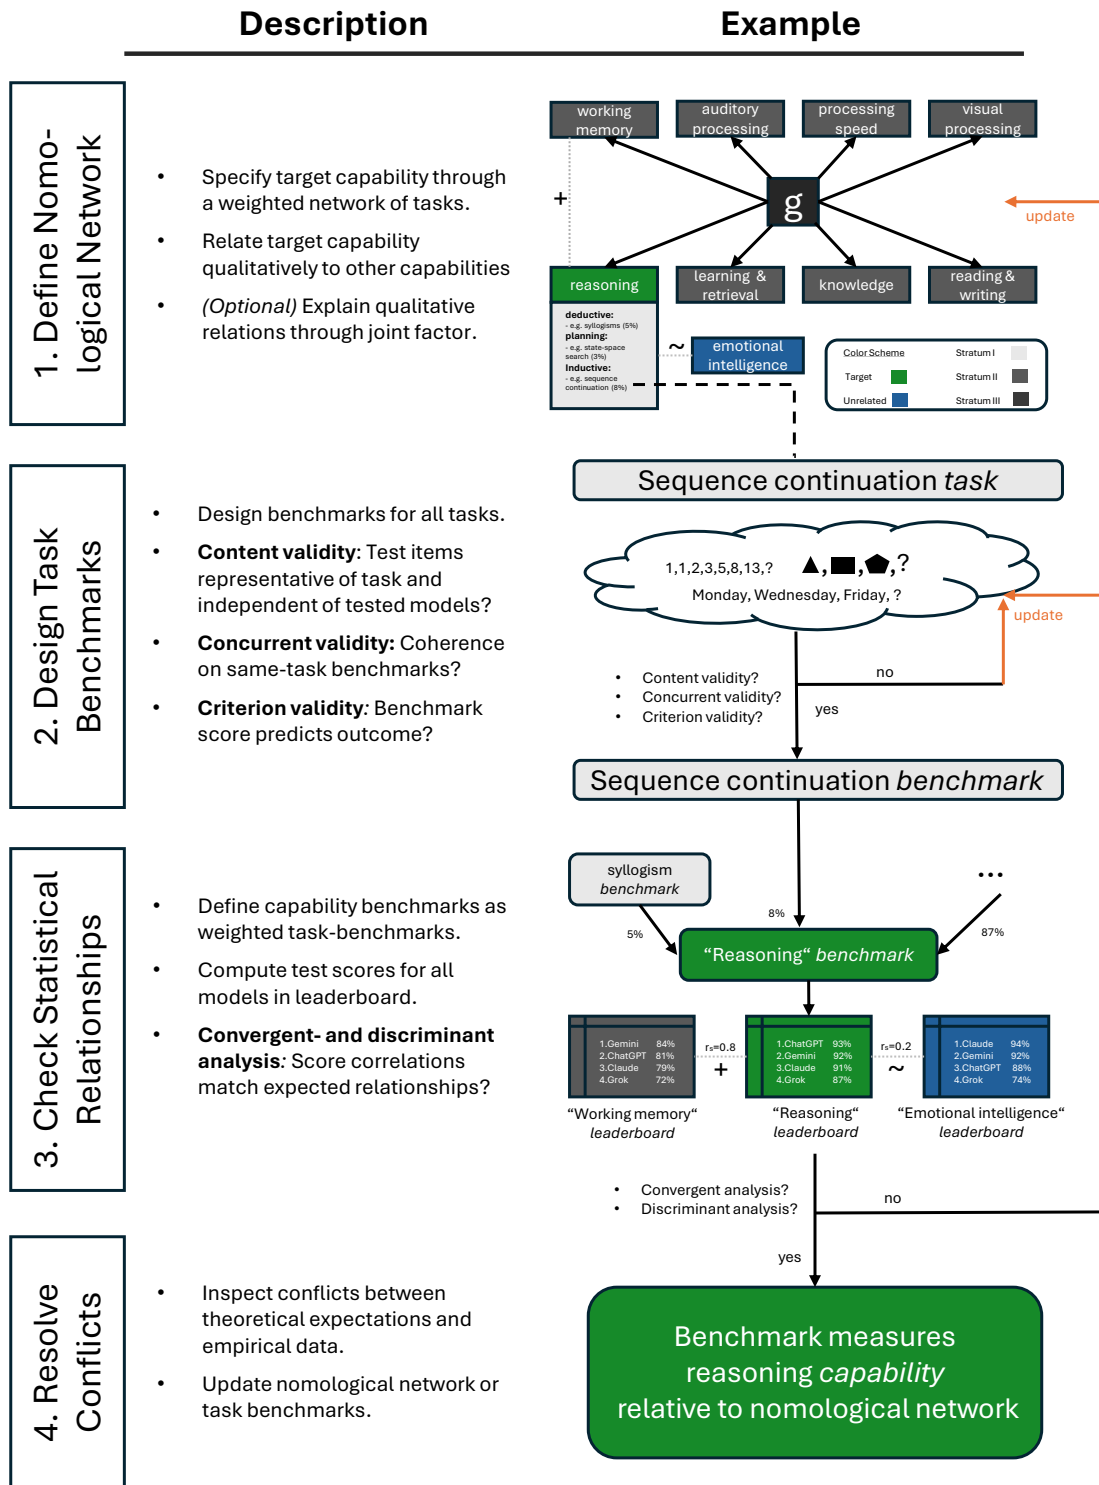


Figure 1: The nomological account structures the benchmark design process into four steps. The right-hand side illustrates the framework using reasoning as an example, with the Cattell–Horn–Carroll theory [10, 33] as the nomological network. To keep the example readable, only a subset of reasoning tasks is described and only two qualitative relationships are specified: ‘+’ indicates an expected positive relationship between reasoning and working memory [29], and ‘~’ denotes an expected independence between emotional intelligence and reasoning [32].

example, selecting an LLM to fill tax forms based on reasoning scores, the nomological account requires verifying that the relevant skill correlates strongly with the score, i.e. *criterion validity*.

One might conclude that established psychological tests such as the Woodcock–Johnson IV, the Wechsler scales, or the Stanford–Binet, since they satisfy these conditions, could simply be adopted for LLM testing. Indeed, they can guide test design, but many validity conditions depend partly on the test-taker. Consider statistical condition 2) above: LLMs are trained on vast corpora of data that may contain many psychological test items. In such cases, the test may measure memorization of correct answers rather than skill on the task. This phenomenon is known as *data contamination* and has recently received increasing attention in the literature [37, 51].

### **3.3 Capability benchmarks should undergo convergent and discriminant analysis**

The next step in the nomological account is to define capability benchmark scores for all capabilities in the nomological network. The benchmark score for reasoning capability is defined as the weighted average of the task benchmark scores connected to reasoning in CHC [10].

Aggregation alone does not secure construct validity. From a nomological perspective, reasoning is not merely a sum of tasks but occupies a specific position within a broader network of capabilities like working memory or processing speed. If an LLM scores extremely high on reasoning but very low on working memory, as some studies suggest [58, 20], this creates a theoretical–empirical mismatch (see Section 3.4).

This implies that LLM researchers must examine whether the theoretical relations posited in the nomological network align with the empirical relations observed in benchmark data [10]. In Cronbach and Meehl, this process is called *convergent* and *discriminant analysis*. Convergent analysis checks whether reasoning scores correlate positively/negatively with theoretically related capabilities such as working memory. Discriminant analysis checks that reasoning scores are uncorrelated with capabilities that are theoretically distinct, such as social intelligence. Discriminant analysis helps researchers to prevent *construct overrepresentation* [50], i.e. benchmarks should not measure capabilities unrelated to the target capability, empirical research indicates that many current benchmarks suffer from this threat [4].

### **3.4 Resolving conflicts between theory and measurement is difficult**

Empirical results and the theoretical expectations encoded in the nomological network may come into conflict. There is no straightforward way to resolve such tensions. We are confronted with a version of the Duhem–Quine problem [53, 13, 3]: as Simms puts it, it remains unclear whether “(i) the measure does not adequately measure the target construct, (ii) the theory requires modification, or (iii) some combination of both.” [48] Applied to reasoning benchmarks, we may need either to question the adequacy of our task measurements or to reconsider aspects of the nomological network itself.

In the context of LLM research, I believe that initial doubt should often fall on the benchmarks. Benchmark design in machine learning is still underexplored, only gradually developing rigorous validation practices [18]. Theoretical–empirical mismatches will frequently indicate weaknesses in benchmark design. Established psychometric standards, such as criteria for content, concurrent, or criterion validity, as well as strategies to detect construct under- and overrepresentation, can provide valuable guidance for improving LLM benchmarks [24, 56].

Alternatively, revising the nomological network amounts to revising the meaning of the construct itself. Researchers who rely on different nomological networks for reasoning implicitly operate with different conceptions of reasoning. Even well-established theories such as CHC may be challenged, particularly when applied to LLMs. As discussed above, CHC is ultimately a theory of *human*, not *LLM*, cognition and may require refinement.

Ultimately, nomological networks translate the meaning of capabilities into a set of statistical constraints on their benchmarks. If these constraints are satisfied, we can conclude that the empirical data cohere with the theory. One might object, following Borsboom, that “correlations are not enough, no matter what their size.” [7] Indeed, coherence alone does not establish Borsboom’s high causal standards. But, if causal structures exist, systematic correlations will be among their observable consequences, rendering empirical coherence with the nomological network at least a necessary condition for its adequacy [39]. If no such causal structure exists, coherence between theory and data still provides partial justification for employing the theoretical vocabulary of human psychology when describing LLM capabilities.

## 4 Conclusion

The increasingly common practice of attributing human-like capabilities to LLMs on the basis of benchmark performance demands a rigorous theoretical foundation. I have argued that the nomological network account of construct validity from psychometrics offers the most promising path forward: it avoids contentious metaphysical commitments about the real existence of constructs in LLMs while providing a substantive and structured account of what those constructs mean. By embedding benchmarks within a theoretically articulated network of interrelated capabilities, we move from isolated performance scores to empirically constrained claims about LLM capabilities, as illustrated through the example of reasoning. If LLM research is to continue ascribing capabilities such as reasoning and theory of mind, it must also adopt the corresponding standards of validity that render these terms scientifically defensible.

## Acknowledgements

I thank Max Hellrigel-Holderbaum for his helpful comments on earlier drafts of this paper. This research was supported by the German Research Foundation (DFG), project number 511917847.

## References

- [1] Ahmed Alaa, Thomas Hartvigsen, Niloufar Golchini, Shiladitya Dutta, Frances Dean, Inioluwa Deborah Raji, and Travis Zack. Position: Medical large language model benchmarks should prioritize construct validity. In *Forty-second International Conference on Machine Learning Position Paper Track*, 2025.
- [2] Jamin Asay. The role of truth in psychological science. *Theory & Psychology*, 28(3):382–397, 2018.
- [3] Yuri Balashov. Duhem, quine, and the multiplicity of scientific tests. *Philosophy of Science*, 61(4):608–628, 1994.
- [4] Andrew M. Bean, Ryan Othniel Kearns, Angelika Romanou, Franziska Sofia Hafner, Harry Mayne, Jan Batzner, Negar Foroutan, Chris Schmitz, Karolina Korgul, Hunar Batra, Oishi Deb, Emma Beharry, Cornelius Emde, Thomas Foster, Anna Gausen, María Grandury, Simeng Han, Valentin Hofmann, Lujain Ibrahim, Hazel Kim, Hannah Rose Kirk, Fangru Lin, Gabrielle Kaili-May Liu, Lennart Luetzgau, Jabez Magomere, Jonathan Rystrom, Anna Sotnikova, Yushi Yang, Yilun Zhao, Adel Bibi, Antoine Bosselut, Ronald Clark, Arman Cohan, Jakob Nicolaus Foerster, Yarin Gal, Scott A. Hale, Inioluwa Deborah Raji, Christopher Summerfield, Philip Torr, Cozmin Ududec, Luc Rocher, and Adam Mahdi. Measuring what matters: Construct validity in large language model benchmarks. In *The Thirty-ninth Annual Conference on Neural Information Processing Systems Datasets and Benchmarks Track*, 2025.
- [5] William Bechtel. *Mental mechanisms: Philosophical perspectives on cognitive neuroscience*. Psychology Press, 2007.
- [6] Denny Borsboom. *Measuring the mind: Conceptual issues in contemporary psychometrics*. Cambridge University Press, 2005.
- [7] Denny Borsboom, Gideon J Mellenbergh, and Jaap Van Heerden. The concept of validity. *Psychological review*, 111(4):1061, 2004.
- [8] Luc Bovens and Stephan Hartmann. *Bayesian epistemology*. Oxford University Press, 2004.
- [9] Giovanni Briganti, Marco Scutari, and Richard J McNally. A tutorial on bayesian networks for psychopathology researchers. *Psychological methods*, 28(4):947, 2023.
- [10] John Bissell Carroll. *Human cognitive abilities: A survey of factor-analytic studies*. Cambridge university press, 1993.
- [11] Jerry A Colliver, Melinda J Conlee, and Steven J Verhulst. From test validity to construct validity... and back? *Medical education*, 46(4):366–371, 2012.
- [12] Lee J Cronbach and Paul E Meehl. Construct validity in psychological tests. *Psychological bulletin*, 52(4):281, 1955.

- [13] Pierre Duhem. *The Aim and Structure of Physical Theory*. Princeton University Press, 1954.
- [14] Sacha Epskamp, Denny Borsboom, and Eiko I Fried. Estimating psychological networks and their accuracy: A tutorial paper. *Behavior research methods*, 50(1):195–212, 2018.
- [15] Timo Freiesleben. Artificial neural nets and the representation of human concepts. In *Philosophy of Science for Machine Learning: Core Issues and New Perspectives*, pages 273–294. Springer, 2025.
- [16] Timo Freiesleben and Sebastian Zezulka. The benchmarking epistemology: Construct validity for evaluating machine learning models. *arXiv preprint arXiv:2510.23191*, 2025.
- [17] Jacqueline Harding and Nathaniel Sharadin. What is it for a machine learning model to have a capability? *The British Journal for the Philosophy of Science*, 2024.
- [18] Moritz Hardt. The Emerging Science of Machine Learning Benchmarks. Online at <https://mlbenchmarks.org>, 2025. Manuscript.
- [19] Carl G Hempel. *Fundamentals of concept formation in empirical science, Vol. II. No. 7*. University of Chicago Press, 1952.
- [20] Dan Hendrycks, Dawn Song, Christian Szegedy, Honglak Lee, Yarin Gal, Erik Brynjolfsson, Sharon Li, Andy Zou, Lionel Levine, Bo Han, et al. A definition of agi. *arXiv preprint arXiv:2510.18212*, 2025.
- [21] Benedikt Höltgen and Robert C Williamson. We should avoid the assumption of data-generating probability distributions in social settings. *arXiv preprint arXiv:2407.17395*, 2024.
- [22] S Brian Hood. Validity in psychological testing and scientific realism. *Theory & Psychology*, 19(4):451–473, 2009.
- [23] Liwei Jiang, Jena D. Hwang, Chandra Bhagavatula, Ronan Le Bras, Jenny Liang, Jesse Dodge, Keisuke Sakaguchi, Maxwell Forbes, Jon Borchardt, Saadia Gabriel, Yulia Tsvetkov, Oren Etzioni, Maarten Sap, Regina Rini, and Yejin Choi. Can machines learn morality? the delphi experiment, 2022. URL <https://arxiv.org/abs/2110.07574>.
- [24] Michael T Kalkbrenner. A practical guide to instrument development and score validation in the social sciences: The measure approach. *Practical Assessment, Research & Evaluation*, 26:1, 2021.
- [25] Michael T Kane. Current concerns in validity theory. *Journal of educational Measurement*, 38(4):319–342, 2001.
- [26] Michael T. Kane. Validating the Interpretations and Uses of Test Scores. *Journal of Educational Measurement*, 50(1):1–73, 2013. doi: 10.1111/jedm.12000.

- [27] Subhash Kantamneni and Max Tegmark. Language models use trigonometry to do addition. In *ICLR 2025 Workshop on Building Trust in Language Models and Applications*, 2025.
- [28] Lena Kästner and Barnaby Crook. Explaining ai through mechanistic interpretability. *European journal for philosophy of science*, 14(4):52, 2024.
- [29] Patrick C Kyllonen and Raymond E Christal. Reasoning ability is (little more than) working-memory capacity?! *Intelligence*, 14(4):389–433, 1990.
- [30] Robert W Lissitz and Karen Samuelsen. A suggested change in terminology and emphasis regarding validity and education. *Educational researcher*, 36(8):437–448, 2007.
- [31] Roman Lukyanenko, Kai R Larsen, and Sen Yan. Integrating llms and psychometrics: Global construct validity. *Available at SSRN 5341306*, 2024.
- [32] John D Mayer, Peter Salovey, and David R Caruso. Emotional intelligence: New ability or eclectic traits? *American psychologist*, 63(6):503, 2008.
- [33] Kevin S McGrew. The cattell-horn-carroll theory of cognitive abilities: Past, present, and future. *Contemporary Intellectual Assessment: Theories, Tests, and Issues*, 2005.
- [34] Kevin Meng, David Bau, Alex Andonian, and Yonatan Belinkov. Locating and editing factual associations in gpt. *Advances in neural information processing systems*, 35:17359–17372, 2022.
- [35] Samuel Messick. Meaning and values in test validation: The science and ethics of assessment. *Educational researcher*, 18(2):5–11, 1989.
- [36] Samuel Messick. Standards of Validity and the Validity of Standards in Performance Assessment. *Educational Measurement: Issues and Practice*, 14(4):5–8, 1995. doi: 10.1111/j.1745-3992.1995.tb00881.x.
- [37] Raphaël Millière and Cameron Buckner. A Philosophical Introduction to Language Models - Part II: The Way Forward, 2024. arXiv.
- [38] Freek Oude Maatman. Folk psychology and network theory: Fact or gamble? a reply to kalis and borsboom. *Theory & Psychology*, 30(5):729–734, 2020.
- [39] Jonas Peters, Dominik Janzing, and Bernhard Scholkopf. *Elements of causal inference: foundations and learning algorithms*. MIT press, 2017.
- [40] Philipp Poschmann, Jan Goldenstein, Sven Büchel, and Udo Hahn. A vector space approach for measuring relationality and multidimensionality of meaning in large text collections. *Organizational Research Methods*, 27(4):650–680, 2024.
- [41] Willard Quine. *From a logical point of view: Nine logico-philosophical essays*. Harvard University Press, 1980.

- [42] Willard Van Orman Quine and Joseph Silbert Ullian. *The web of belief*, volume 2. Random House New York, 1978.
- [43] Benjamin Recht, Rebecca Roelofs, Ludwig Schmidt, and Vaishaal Shankar. Do ImageNet Classifiers Generalize to ImageNet? In Kamalika Chaudhuri and Ruslan Salakhutdinov, editors, *Proceedings of the 36th International Conference on Machine Learning*, volume 97 of *Proceedings of Machine Learning Research*, pages 5389–5400. PMLR, 09–15 Jun 2019.
- [44] Heiner Rindermann, David Becker, and Thomas R Coyle. Survey of expert opinion on intelligence: Intelligence research, experts’ background, controversial issues, and the media. *Intelligence*, 78:101406, 2020.
- [45] Olawale Elijah Salaudeen, Anka Reuel, Ahmed M Ahmed, Suhana Bedi, Zachary Robertson, Sudharsan Sundar, Benjamin W Domingue, Angelina Wang, and Sanmi Koyejo. Measurement to meaning: A validity-centered framework for ai evaluation. In *NeurIPS Workshop on Evaluating the Evolving LLM Lifecycle: Benchmarks, Emergent Abilities, and Scaling*, 2025.
- [46] Joel Schneider and Kevin McGrew. The cattell-horn-carroll theory of cognitive abilities. *Contemporary intellectual assessment: Theories, tests, and issues*, 733:163, 2018.
- [47] Lee Sharkey, Bilal Chughtai, Joshua Batson, Jack Lindsey, Jeffrey Wu, Lucius Bushnaq, Nicholas Goldowsky-Dill, Stefan Heimersheim, Alejandro Ortega, Joseph Isaac Bloom, Stella Biderman, Adrià Garriga-Alonso, Arthur Conmy, Neel Nanda, Jessica Mary Rumbelow, Martin Wattenberg, Nandi Schoots, Joseph Miller, William Saunders, Eric J Michaud, Stephen Casper, Max Tegmark, David Bau, Eric Todd, Atticus Geiger, Mor Geva, Jesse Hoogland, Daniel Murfet, and Thomas McGrath. Open problems in mechanistic interpretability. *Transactions on Machine Learning Research*, 2025. ISSN 2835-8856. URL <https://openreview.net/forum?id=91H76m9Z94>. Survey Certification.
- [48] Leonard J Simms. Classical and modern methods of psychological scale construction. *Social and Personality Psychology Compass*, 2(1):414–433, 2008.
- [49] James W. A. Strachan, Dalila Albergo, Giulia Borghini, Oriana Pansardi, Eugenio Scaliti, Saurabh Gupta, Krati Saxena, Alessandro Rufo, Stefano Panzeri, Guido Manzi, Michael S. A. Graziano, and Cristina Becchio. Testing theory of mind in large language models and humans. *Nature Human Behaviour*, 8(7):1285–1295, 2024. doi: 10.1038/s41562-024-01882-z.
- [50] Milton E. Strauss and Gregory T. Smith. Construct Validity: Advances in Theory and Methodology. *Annual Review of Clinical Psychology*, 5(1):1–25, 2009. doi: 10.1146/annurev.clinpsy.032408.153639.
- [51] Tom Sühr, Florian E. Dorner, Olawale Salaudeen, Augustin Kelava, and Samira Samadi. Stop Evaluating AI with Human Tests, Develop Principled, AI-specific Tests instead, 2025. arXiv.

- [52] Adly Templeton. *Scaling monosemanticity: Extracting interpretable features from claude 3 sonnet*. Anthropic, 2024.
- [53] Willard van Orman Quine. Two dogmas of empiricism. In *Can Theories be Refuted? Essays on the Duhem-Quine Thesis*, pages 41–64. Springer, 1976.
- [54] David Watson. The rhetoric and reality of anthropomorphism in artificial intelligence. *Minds and Machines*, 29(3):417–440, 2019.
- [55] Oliver Wilhelm. Measuring reasoning ability. *Handbook of understanding and measuring intelligence*, pages 373–392, 2005.
- [56] Tarid Wongvorachan and Okan Bulut. Detecting construct-irrelevant variance: A comparison of network psychometrics and traditional psychometric methods using the hexaco-pi dataset. *Psychology International*, 7(4):88, 2025.
- [57] Fengli Xu, Qianyue Hao, Chenyang Shao, Zefang Zong, Yu Li, Jingwei Wang, Yunke Zhang, Jingyi Wang, Xiaochong Lan, Jiahui Gong, et al. Toward large reasoning models: A survey of reinforced reasoning with large language models. *Patterns*, 6(10), 2025.
- [58] Chunhui Zhang, Yiren Jian, Zhongyu Ouyang, and Soroush Vosoughi. Working memory identifies reasoning limits in language models. In *Proceedings of the 2024 Conference on Empirical Methods in Natural Language Processing*, pages 16896–16922, 2024.