
Position: Stop Evaluating AI with Human Tests, Develop Principled, AI-specific Tests instead

Anonymous Authors¹

Abstract

Large Language Models (LLMs) have achieved remarkable results on a range of standardized tests originally designed to assess human cognitive and psychological traits, such as intelligence and personality. While these results are often interpreted as strong evidence of human-like characteristics in LLMs, this paper argues that such interpretations constitute an ontological error. Human psychological and educational tests are theory-driven measurement instruments, calibrated to a specific population. Applying these tests to non-human subjects without empirical validation, risks mischaracterizing what is being measured. Furthermore, a growing trend frames AI performance on benchmarks as measurements of traits such as “intelligence”, despite known issues with validity, data contamination, cultural bias and sensitivity to superficial prompt changes. We argue that interpreting benchmark performance as measurements of human-like traits, lacks sufficient theoretical and empirical justification.

This leads to our position: Stop Evaluating AI with Human Tests, Develop Principled, AI-specific Tests instead. We call for the development of principled, AI-spec by laying out, end-to-end, how valid measurement instruments are constructed and validated and where the ontological error enters when a human-calibrated instrument is applied to LLMs.

1. Introduction

Suppose you placed a heart rate monitor on a robot arm to “read its pulse”. The device may return a value, but the robot has no pulse, so the number cannot be interpreted in the same way as for humans: if the monitor showed 0, we

¹Anonymous Institution, Anonymous City, Anonymous Region, Anonymous Country. Correspondence to: Anonymous Author <anon.email@domain.com>.

Preliminary work. Under review by the International Conference on Machine Learning (ICML). Do not distribute.

would not conclude that the robot is dead. Analogously, the apparent “personality” or “intelligence” of a Large Language Model (LLM) inferred from human tests might not map to the same human traits. This paper argues that interpreting AI scores on human tests as strong evidence of human-like traits is an ontological error.

Our argument applies broadly to AI evaluation, but it is especially pressing for LLMs, which are now routinely tested with instruments designed to measure human traits: standardized exams (e.g., GRE/SAT) (Achiam et al., 2023; Grattafiori et al., 2024), intelligence tests (Huang & Li, 2024; Wasilewski & Jablonski, 2024), and personality inventories (Huang et al., 2023; Caron & Srivastava, 2023; Jiang et al., 2023; Mei et al., 2024). Because these instruments target familiar human concepts, strong LLM performance is easily (and often publicly) read as evidence of human-like traits. For example, a high IQ-test score or a seemingly stable Big Five profile (Soto & John, 2017) is frequently taken to imply intelligence or personality in the human sense. We argue that this inference is not warranted without theory and validation showing that the same latent trait is being measured.

Psychological and educational tests are not just question sets; they are population-calibrated measurement instruments. Their items, scoring, and interpretation are tied to measurement models fitted and validated on a defined human population; outside that population, validity can fail (Salaudeen et al., 2025b). A simple analogy is age: one would not use the SAT to assess young children because it was built for older adolescents (American Educational Research Association et al., 2014). Likewise, applying such instruments to non-human agents like LLMs can produce numbers, but those numbers need not carry the intended human-trait meaning.

This problem is not limited to human psychological and educational testing, but extends to benchmarks. Here, the language has shifted from using the terms “dataset to evaluate the accuracy on a task” (LeCun et al., 1998; Deng et al., 2009; Srivastava et al., 2022; Hendrycks et al., 2020) to “measures” of (general) “intelligence” (Wang et al., 2024; OpenAI, 2024b;a) and “general reasoning” (Kazemi et al., 2025). However there are several concerns that undermine

these benchmarks’ validity as measures of intelligence, including sensitivity to answer order (Gupta et al., 2024) and other minor modifications (Recht et al., 2019), errors in questions and labels (Gema et al., 2024), cultural biases (Singh et al., 2024), and training data contamination (Xu et al., 2024).

Linking psychological traits, like intelligence or personality, to benchmark performance suggests that LLMs exhibit these traits, when they might indeed be ill-defined for LLMs. For example, human personality is defined via interindividual differences in people’s thinking, feeling, and behaving (Soto & John, 2017). Without sufficient evidence that LLMs generally behave equivalent to people and have internal states that can be described as feelings or thinking, personality cannot be directly transferred to LLMs.

While these issues highlight serious concerns, they also open up opportunities for foundational research at the intersection of machine learning and measurement theory. Rather than directly transferring human psychometric instruments, we should focus on developing new evaluation frameworks tailored to LLMs (Weidinger et al., 2025; Salaudeen et al., 2025b). This involves identifying and defining “psychological” traits specific to LLMs, some of which might overlap with human *constructs* (a concept formed by combining simpler ideas, such as personality characterized by five dimensions (American Psychological Association, n.d.)), while others may be entirely new.

Taken together, these issues show that scores on human tests do not, by themselves, justify human-trait interpretations for LLMs. We therefore take the position: **Stop Evaluating AI with Human Tests, Develop Principled, AI-specific Tests instead**. Concretely, we (i) lay out an end-to-end framework for how measurement instruments can be constructed and validated (Figure 1, Section 3), (ii) introduce core concepts from measurement theory such as measurement invariance (Section 4.2) with empirical examples (Appendix A), (iii) state four open problems that the ML benchmarking and evaluation community must resolve before AI-specific measurement instruments can be proposed in a principled way (Section 6), and (iv) highlight the risks of communicating invalid measurements to the general public (Section 5).

2. Related Work and the Shift in Language

Measurement claims about benchmarks Early benchmarks such as MNIST (LeCun et al., 1998) and ImageNet (Deng et al., 2009) termed their collections as “datasets” or “databases” for assessing model accuracy. In the early days of large language models, Hendrycks et al. (2020) described their MMLU benchmark as aiming to “measure knowledge acquired during pretraining” reflecting a more ambitious interpretation of what such evaluations could capture. In contrast, Wang et al. (2018) offered a

more pragmatic and conventional framing, presenting their work as a tool for “evaluating the performance of models across a diverse set of existing NLP tasks.”. The wording shifted increasingly in subsequent years. While the authors of BIG-Bench (Srivastava et al., 2022) and its subset BIG-Bench Hard (Suzgun et al., 2022) make no claims about the measurement of latent abilities, the more recent BIG-Bench Extra Hard (BBEH) (Kazemi et al., 2025) describes its predecessor as “benchmark(s) for evaluating the general reasoning capabilities of LLMs [...]” and use similar language for their own benchmark. Similarly, the authors of MMLU-Pro (Wang et al., 2024), an updated version of MMLU shift from “expert-level performance” to describing benchmarks like MMLU, BBH and ARC (Clark et al., 2018) as “benchmarks to measure such general intelligence.”. However, no such claim has been made by the authors of either of the three cited benchmarks. These examples illustrate the broader trend of describing benchmarks as instruments for measuring human-like traits of LLMs.

Human psychological and educational tests applied to LLMs A broad spectrum of educational and psychological tests originally designed for humans, has been applied to LLMs. Most prominently, the technical report of GPT-4 (OpenAI, 2023; Achiam et al., 2023) includes a table of GPT-4’s performance on academic and professional tests like the GRE, SAT, LSAT and AP exams and compares the results to human score distributions. The Llama 3 and Claude 3 technical reports (Grattafiori et al., 2024; Anthropic, 2024) include similar tables. Beyond this, Webb et al. (2023) apply fluid intelligence tests to LLMs, while Strachan et al. (2024) administer tests for “theory of mind”, the ability to understand mental states as a cause for behavior, on both LLMs and humans. In both cases, LLM performance is compared to humans. Comparing LLM performance to human performance on tests designed for humans can offer valuable insights into LLM capabilities. However, a strong focus on such comparisons, especially along a single axis of performance, risks obscuring whether these tests are truly measuring the same underlying trait in both humans and LLMs.

This concern is not limited to cognitive assessments. It also applies to domains like personality and social psychology, where several studies have administered personality inventories to LLMs (Caron & Srivastava, 2023; Huang et al., 2023; Jiang et al., 2023; Mei et al., 2024) and even assessed so-called “dark triad” traits such as Machiavellianism, psychopathy, and narcissism (Li et al., 2022).

Critical investigations of evaluation practices We are not the first to criticize the aforementioned practices. McCoy et al. (2023) demonstrate that LLMs are influenced by task, input and output probability. For example, LLMs

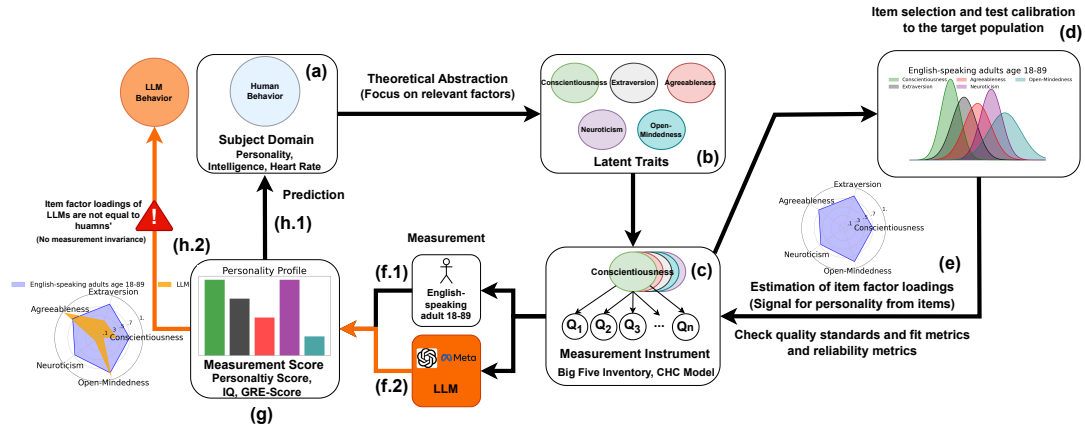


Figure 1. Overview of the measurement instrument creation process using a personality test as an example. (a) The goal is to measure behavioral tendencies and preferences. We model personality (b) using five latent traits and design items (c) expected to correlate with them, tailored to a target group (English-speaking adults, 18–89). Responses from a sample (d) are analyzed via factor analysis (e) to estimate item-trait covariances. We assess reliability and model fit (e.g., Cronbach’s α , McDonald’s ω_h), iterating (c)→(d)→(e) as needed. To prevent overfitting, we use holdout-based item selection. Once validated, the instrument can assess individual personality profiles (f.1) and defines a function from behavior (a) to personality scores (g), enabling behavior prediction. Applied to LLMs (f.2), the test yields scores (g), but differing factor loadings (h.2) limit predictive validity. Earlier stages (a)→(b) were designed for humans, not LLMs.

perform better when the correct answer is a common string. The authors conclude that LLMs should not be evaluated as if they are humans, but as a completely different type of system. Complementarily, Schröder et al. (2025) demonstrate that subtle semantic changes in test items lead to significantly different responses between humans and LLMs. Similarly, Sühr et al. (2023) find that LLMs answers to personality inventories do not generate the same variance patterns as they do in humans. Jung et al. (2026) show empirically, that instruments for constructs of sexism, racism and morality do not transfer to LLMs, specifically by rejecting ecological validity based on rank correlations with downstream task performance. On a higher level, Weidinger et al. (2025) call for building a science of evaluation for generative AI and propose that evaluation metrics should reflect real world performance, be iteratively refined, and follow a clear set of norms. Lastly, Salaudeen et al. (2025b) extensively discuss the concept of validity from a machine learning perspective and apply their framework to multiple benchmarks and accompanying measurement claims.

3. The Problem With Psychological Tests on LLMs

Psychological tests are not interchangeable “question sets”, they are calibrated measurement instruments for a defined human population. Item selection, scoring, and interpretation are tied to a measurement model fitted on human response distributions.

Figure 1 summarizes this process. It starts from an empirical

subject domain (a) (e.g., behavior and preferences), specifies a construct model (b) (e.g., personality), designs items (c) for a target population, estimates and validates item–trait relations (d,e), and then uses the resulting scores (g) for prediction in the same population (h.1).

The ontological error we criticize in this work, is to treat scores produced by a human-calibrated instrument on a non-human population (e.g., LLMs) as measurements of the same latent human trait, without re-establishing the measurement model and its validity for that new subjects.

3.1. The Process of Creating a Valid Measurement Instrument (a)→(h.1)

A valid instrument is built for a specific subject domain and a specific target population.

First, the target phenomenon in the subject domain (a) is operationalized via a construct model (b) (e.g., the Big Five as latent traits). Because traits are latent, we design observable indicators (items) (c) that are expected to covary with the intended traits.

Second, items are calibrated for a target population. In our example, this is English-speaking adults aged 18–89. We administer the items to a sample (d) and estimate/validate the measurement model (e) (e.g., via factor analysis; with fit/reliability diagnostics such as TLI/RMSEA/CFI and Cronbach’s α /McDonald’s ω_h (Xia & Yang, 2019)). The design–calibration loop (c)→(d)→(e) is iterated, typically with holdout-based item selection to limit overfitting.

Only after this calibration can the instrument be used to

assign scores (**g**) to individuals in the target population (**f.1**) and support predictive uses in that population (**h.1**).

3.2. The Problem (**f.2**)→(**h.2**)

Applying a human-calibrated instrument to an LLM breaks the assumptions that make the score interpretable.

Population mismatch. The LLM is not drawn from the population used to calibrate the instrument (e.g., English-speaking adults aged 18–89). Hence, the item parameters/factor loadings validated in (**c**)→(**d**)→(**e**) need not transfer. The test will still output a score profile (**g**) when administered to an LLM (**f.2**), but without comparable loadings the score lacks the intended meaning and undermines predictive validity for LLM behavior (**h.2**) (see Appendix A).

Construct mismatch. Even if item parameters looked similar, the construct itself may not transfer: the human construct model (**b**) abstracts from a human subject domain (**a**) shaped by embodiment and social context. If the relevant behavioral regularities differ for LLMs, the latent traits worth modeling (and thus the mapping (**a**)→(**b**)→(**g**)) are likely different.

In short: administering the test yields numbers, but treating those numbers as measurements of the same human trait is unwarranted unless the construct and measurement model are re-validated for the LLM subject.

4. Challenges and Principles of Measurement in Machine Learning

In the following sections, we outline the key aspects of our position and highlight current machine learning evaluation practices that we believe should be revised or adopted. Note that these *points of action* are general recommendations and not specific to personality tests or psychological and educational tests.

There is a substantial body of work (Sühr et al., 2023) employing approaches similar to the following simplified example:

Research Question: “Do LLM personality profiles match those of humans?”

Method: Application of the state of the art Big Five Personality Test to an LLM.

Result: “There is no significant difference between average LLM and average human personality sum-scores.”

4.1. Strange Coincidences instead of simple alternative explanations

In the example in Section 4, the study design has low *experimental risk*: it is unlikely to yield a clearly negative result even if the trait attribution is wrong. In Meehl’s terms,

good theories should be “subjected to grave danger” by tests that make failure plausible under the null (Meehl, 1992); otherwise, apparent confirmations are cheap.

An analogy is early saliency-map evaluation: visually plausible explanations seemed supportive, yet similar-looking maps appeared even for networks with *random weights* (Adebayo et al., 2018; Rudin, 2019). Likewise, if we want to argue that LLMs have “personality” or “intelligence”, we need explicit, falsifiable measurement hypotheses and experiment designs that can actually refute them (Popper, 2013).

Consider the common claim that an LLM’s Big Five *sum-scores* are “human-like.” On a 1–5 scale, human averages are often close to the midpoint (e.g., 2.89–3.66 across traits in one reference sample (Soto & John, 2017)), so even weak response heuristics can land near the mean. Moreover, reverse-keyed items make constant-answer strategies look balanced after scoring: repeating the same option across true- and false-keyed items can yield mid-range averages, a pattern observed in instruction-tuned models (Sühr et al., 2023).

Therefore, similarity in aggregate scores is neither necessary nor sufficient evidence for a shared latent trait: deviations can always be re-labeled as “a different personality,” and similarities admit simple, non-trait explanations. This is why such comparisons rarely put the underlying trait claim in grave danger.

4.2. Measurement Invariance Investigations

In the example in Section 4, the researchers transferred a measurement instrument (the personality test), from the population it was designed for (adult residents of English-speaking nations (Soto & John, 2017)) to LLMs. An obvious objection is that LLMs may not possess personality. However, we cannot prove that LLMs lack personality; we can only investigate whether the model of personality underlying the personality test (the Big Five Inventory) explains the variance in LLM responses as well as it does in human responses. In other words, we can ask whether the data from humans and LLMs is likely generated by the same process that the test is designed to measure. *Measurement Invariance* refers to whether test items measure the same thing across groups of individuals (e.g., cultures, languages). It ensures that items function similarly across different subgroups. We provide a concrete worked example via a factor-loading comparison in Appendix A.

Formally, $X = X_1, \dots, X_n$ is our *measurement instrument* for the *latent variable* of personality ξ . The random variables X_i with realization x_i denote the answer (i.e. the score) to item i in the personality test (Bollen, 1989). Let further ξ denote the latent random variable for one “true” personality dimension (e.g. extraversion). Now we can

model the relationship¹ between the observed score of a personality test item and the true personality score as

$$X_i = \alpha_i + \lambda_i \cdot \xi + \varepsilon_i \quad (1)$$

where ε_i is a residual noise term, α_i the bias and $\lambda_i = \frac{\text{Cov}(X_i, \xi)}{\text{Var}(\xi)}$ is a regression coefficient or the “factor loading” of ξ on X_i . In machine learning terms, we can think of λ_i as the “signal” for the true personality score in the answers to the item i . Now consider the cumulative distribution function $F(\cdot)$ and a subgroup realization $g \in G$, such as male and female, as realizations of gender.

Since the theory of personality does not make distinctions between the group g and the rest of the population, we require $F(X|\xi = y, g) = F(X|\xi = y)$ for all (x, y, g) . This property is called *measurement invariance*.

Definition 4.1. (*Measurement Invariance by Meredith (1993)*) A random variable X is said to be *measurement invariant* with respect to a subgroup $g \in G$ if $F(X|\xi = y, g) = F(X|\xi = y)$ for all (x, y, g) in the sample space.

Definition 4.1 is a strong property and implies that the measurement instrument works as well for the group g as it does for everyone else in the parent population. This implies that we can transfer the measurement instrument to the subpopulation of individuals in g . Measurement invariance is a necessary condition for the *validity* (see Appendix C.1) of comparisons of measurements between populations. In short, if the measurement instrument does not work equally well for two subgroups, then either the test is flawed or the construct at hand is different for the two subgroups.

Coming back to our example, if there is no measurement invariance between humans and LLMs, either the test is flawed or personality works different in LLMs and humans, if it exists at all.

A substantial methodological research field explores measurement invariance and other weaker forms of it. Arguably the most important approach is *Confirmatory Factor Analysis* (CFA) (Bollen, 1989). The core idea is that we compare the parameters of Equation 1. If there is a statistically significant difference between these parameters estimated on a sample of LLM answers compared to a human sample, we can reject the hypothesis of (strong) measurement invariance. As mentioned above, weaker forms of measurement invariance exist, for example metric invariance, where only the factor loadings λ_i are equal, but we allow intercepts and noise terms to be different.

In summary, measurement invariance is necessary to establish that we measure what we aim to measure for different

¹We choose a linear relationship for the sake of the example. However, the relationship between indicator (X) and latent ξ can be non-linear.

populations. Without testing it first, we can not interpret psychological or educational test scores on machine learning models as equivalent to human indicators for latent traits. The example in Section 3 is a shortened demonstration of the lack of measurement invariance in the context of a personality test.

Measurement invariance between models Measurement invariance investigations in psychology are often conducted to test whether a measurement instrument is biased against a subpopulation (e.g. male vs. female, native-english speakers vs. non-native english speakers) (Meredith, 1993). This connects to the literature on machine learning fairness, (e.g. (Dwork et al., 2012; Hardt et al., 2016; Barocas et al., 2023)), where models are interpreted as measurement instruments and humans are the subjects. In the context of ML evaluation, that relationship is reversed. LLMs are subjects of measurements, using benchmarks and other instruments. With thousands of models now available, more than 4,500 listed on the Huggingface Open LLM Leaderboard alone (Fourrier et al., 2024), we often rely on the same set of evaluation tools for all of them. For example, MMLU has 303 questions in the category “Chemistry” and 1763 questions in the category “Law”. If a simple average over all questions is conducted, models that perform better on legal questions might thus obtain a better score, even if they perform substantially worse on chemistry questions. However, *a priori*, there is no reason why “Law” should be a better indicator of “general intelligence” than “Chemistry”. In particular, it seems absurd to claim that fine-tuning an LLM to perform better at legal questions would increase its “general intelligence”. Correspondingly, we cannot claim measurement invariance of a benchmark like MMLU between models, when they are fine-tuned to different degrees on different data. Combined with the increasingly blurred boundaries between pretraining and fine-tuning (Dominguez-Olmedo et al., 2024), this suggests that measurement invariance—even between different LLMs—cannot be taken for granted. Note that this critique is only directed towards interpreting benchmarks as measurements. A leaderboard on a benchmark like MMLU still serves the purpose of model comparison and healthy competition over the best performance on a dataset. While creating an ordinal ranking of approaches is useful and helps drive machine learning research, as we will discuss in Section 7, claiming that this constitutes a true measurement requires stronger properties of the measurement instrument.

4.3. Redundant Constructs and Nomological Networks

Now consider the following **benchmark example**:

Claim: *We propose a new benchmark “IntelliBench”. It measures general artificial intelligence in LLMs.*

Method: A multiple choice benchmark with questions that are hard for humans, even for experts.

Result: *The models that perform well on other benchmarks perform well on this benchmark.*

Similar claims have been made in the benchmarking literature, as discussed in Section 2. How does the team of researchers support their claim that they measure general artificial intelligence? Cronbach & Meehl (1955) proposes to embed such claims in a network of scientific evidence, called *Nomological Networks*. Predictions such as “high scores on IntelliBench should predict high scores on other benchmarks that claim to measure general artificial intelligence” or “people with a diagnosed depression tend to score high on *negative emotionality* in a personality test”. The less likely the observed connection, assuming our theoretical model is incorrect, the better. In other words, it would be a very strange coincidence if we would make the observation *despite* our theory being false. This reflects the central theme of Section 4.1: we aim to place our theory at serious risk of refutation.

The researchers in the example above took an initial step towards this type of analysis. Scores on IntelliBench correlate with other benchmarks that claim to measure general artificial intelligence. However, there is a chicken-egg problem: If claims about other benchmarks measuring “intelligence” were mistaken, the correlation with them would provide evidence “against” rather than for the new benchmark measuring “intelligence”. The correlation between the benchmarks suggest that they might collectively measure *something* meaningful, but that something might just be a combination of model size and amount of training data, i.e. pre-training compute (Ruan et al.; Kaplan et al., 2020), which might or might not be the same as “intelligence”.

We thus need to figure out whether machine intelligence is conceptually different from pre-training compute. One simple solution would be to simply define “machine intelligence” in terms of pre-training compute. However, given clear expectations placed on the term “intelligence” from the human context, this option might be highly misleading. Overloading the same term with multiple meanings can lead to conceptual confusion, making it harder to extrapolate capabilities, compare studies or replicate results. To avoid such confusion, claims about a test measuring traits with a clear human analogue like “machine intelligence” should not solely be supported by correlations with other tests for “machine intelligence”. Instead, such claims require strong correlations with indicators of the corresponding human trait; In other words, they need to make sense within the nomological networks *established for humans*.

4.4. Missing Standards of ML Testing

A recurring theme across the issues identified in this paper is the lack of established standards for the evaluation of machine learning models. In other fields, standards have

been developed over decades to ensure the validity, reliability, fairness, and transparency of measurement instruments. The *Standards for Educational and Psychological Testing* (American Educational Research Association et al., 2014) provide a comprehensive framework that covers test development, evaluation, interpretation, and use. They emphasize the importance of explicitly defining constructs, validating instruments within specific populations, ensuring consistency across test administrations, and addressing sources of bias and error.

In contrast, machine learning evaluation currently lacks such systematic guidance. While some standards for test administration, such as the LM Evaluation Harness (Gao et al., 2024) and HELM (Liang et al., 2022) already exist, there is, to the best of our knowledge, no widely adopted framework for the development and validation of machine learning benchmarks as measurement tools. The Standards for Educational and Psychological Testing (American Educational Research Association et al., 2014) may offer a valuable blueprint for establishing such guidelines. Developing a standardized framework for benchmark creation would not only facilitate meaningful comparisons across studies but also require researchers to explicitly address essential issues such as validity.

5. Impact and Potential Risks of Invalid Measurements and Comparisons

False Certification Risks Evaluating machine learning systems with tests meant for humans is fraught with risk. For instance, testing LLMs on the bar exam (Katz et al., 2024) not only attracts widespread media attention but also encourages misplaced trust in these systems. Recently, three lawyers were sanctioned by a US district Judge in Wyoming for citing fake cases generated by AI (Weiss, 2025; Merken, 2025). This example demonstrates the broader risk of attributing human-like capabilities to AI, solely based on tests designed for humans. The public perception of AI is already shifting towards more and more human-like metaphors (Cheng et al., 2025), increasing the risk of similar incidents in the future, potentially in even more sensitive contexts like psychotherapy (Lawrie, 2025). More mundanely, claims about general LLM capabilities create public expectations that often do not survive contact with reality, thus damaging public trust in the field of machine learning evaluations (Widder & Hicks, 2024).

Anthropomorphization Obscures Liability Undue overestimation of the similarity between humans and LLMs provides problematic opportunities for the creators of LLMs to avoid liability by shifting blame to their models. For example, in a recent lawsuit involving the suicide of a 14-year-old teenager linked to the outputs of a chatbot by Char-

acter.ai² (Roose, 2024; Miller, 2024), the firm’s defense involved an argument that their chatbot’s outputs are protected speech under the First Amendment. Strong narratives of LLMs possessing human-like traits like personality, intelligence, or theory of mind might enable firms to evade liability for their ML systems *by* humanizing them and shifting the blame.

6. Opportunity: Developing Principled and Valid Measurement Models for ML-specific Traits

While the above mentioned issues pose significant challenges in ML evaluation, these challenges open up a rich and largely unexplored research frontier. At the intersection of ML benchmarking, psychometrics, and econometrics, there is an exciting opportunity to develop ML-specific measurement models that are grounded in theory and adapted to the unique properties of machine learning systems. In this section, we provide a non-exhaustive list of research directions foundational for a measurement science of machine learning.

Should we evaluate ML models as individuals, populations or something else? So far, there is no clear consensus about the proper statistical unit of a measurement. Often, models are treated as the statistical unit (as an individual), especially in the context of benchmarking. Some works treat models like a distribution and compare the distribution of single LLM’s answers to distributions over human populations (Santurkar et al., 2023). Because these probabilities are highly affected by in-context learning, some works generate “personas” based on system prompts or behavioral instructions (Jiang et al., 2023; Sühr et al., 2023). While different interpretations of this question (whether an ML model constitutes a population or an individual) could be appropriate in different evaluation contexts, they should not be decided on an ad-hoc basis. Furthermore, large parts of measurement theory, as discussed in this work, require a well-defined target population to which a measurement instrument can be calibrated. Therefore, explicitly defining and justifying the statistical unit of measurement is essential for identifying that target population.

What is a representative sample in the context of ML evaluation? The parameters of a measurement model are estimated using data from a sample of the population it is intended to measure. But what is the population of e.g., LLMs? What is a representative sample? At the time of writing this, the Open LLM Leaderboard on Huggingface (Fourrier et al., 2024) lists approximately 4,500 models. However, many of them are based on very few pre-trained

models. Some are very specific models, e.g. for code generation. They could also constitute latent classes (Collins & Lanza, 2009) of the LLM population. Studies have not used a single approach. Some select a few reasonable models (Dominguez-Olmedo et al., 2024), others conduct analysis on thousands of models (Kipnis et al., 2024). To improve reproducibility, comparability, and generalization of findings, a comprehensive and systematic investigation of this issue is warranted.

How to deal with the rapid progress in LLMs? The fast evolution and large variety of machine learning models might make it difficult to meaningfully define a population or representative sample. In that case, a successful science of LLM evaluation requires new principled frameworks for measurement that retain validity without the need for defining a specific population of LLMs. Data leakage and benchmark gaming (“benchmaxxing”) can further inflate scores on popular benchmarks; this confounding factor would also affect any conventionally created measurement instrument unless it is explicitly controlled for.

What are meaningful, ML-specific constructs? We have argued that human constructs should not be transferred to ML models without substantial evidence for their relevance, applicability and invariance in the machine learning context. For example, instead of using a human personality model to try to distinguish between LLMs, we should establish clear behavioral dimensions along which LLMs vary (Rahwan et al., 2019). Maybe some aspects that describe human personality are meaningless for LLMs (e.g. Anxiety, Trust) and some would be relevant in LLMs but are meaningless for humans. In other words, instead of trying to use human constructs that might not work well for LLMs, we should define new constructs that distinguish LLMs in meaningful ways.

How to leverage the unique context of ML? Machine learning unique advantages over cognitive science, econometrics and psychology: We can directly probe causal relationships through counterfactual interventions. For example, we can directly observe the effect of altering answer options and answer orderings on the same model (Dominguez-Olmedo et al., 2023; Gupta et al., 2024). In addition, we can “normalize the knowledge” of models to create similar test conditions, by fixing training data or fine-tune additionally on the test task (Dominguez-Olmedo et al., 2024). This experimental power, brings the potential to overcome limitations in psychological and educational testing by adapting their theories and methodologies to the context of machine learning.

7. Alternative View: Benchmarking as a driver of progress

Despite our criticism of claims regarding the use of benchmarks as measurement for LLMs, benchmarking can be

²<https://character.ai/>

viewed as an essential driver for progress in machine learning science. In particular, when interpreted correctly, benchmarks can support strikingly robust statements.

Hardt (2025) argues that benchmarking has been a fundamental force in machine learning research, enabling rigorous model comparisons and catalyzing progress through shared metrics. He argues that benchmarks with fixed train/test splits have instantiated the *iron rule of machine learning research*, an adaptation of Strevens’ “iron rule of modern science” (Strevens, 2020). In this system, the methodological pluralism famously described by Feyerabend (2020) as “*anything goes*” finds formal structure: any approach is acceptable, provided it improves performance on the benchmark. According to Hardt (2025) this norm has fostered interdisciplinary work with minimal barriers while providing external validity in the form of stable model rankings that often generalize across datasets and tasks. This stability is famously demonstrated in the work of Recht et al. (2019), showing that model rankings, but not accuracy scores, on the CIFAR-10 and ImageNet benchmarks are preserved on recreations of the original test sets. Many works have since replicated the stability of model rankings under a variety of distribution shifts in different contexts (Miller et al., 2020; 2021) and under different experimental conditions (Salaudeen & Hardt, 2024).

We fully agree that these findings highlight the value of benchmarking—when interpreted as Hardt (2025) suggests: as a way to reliably rank machine learning models trained on the same data across a set of related tasks, making it easier to track progress. However, these findings neither support broader claims about benchmarks as measurement instruments, nor do they indicate that comparisons between models and humans can be extrapolated broadly: First, Recht et al. (2019) observe large differences between individual models’ accuracy on the original test set compared to the recreated one. Teney et al. (Teney et al., 2023) and Salaudeen et al. (Salaudeen et al., 2025a) demonstrate that the relationship between the original test set compared to the recreated one can also be uncorrelated or inversely correlated. This means that while model rankings are stable, specific accuracy numbers are essentially meaningless as they are easily swayed by small random fluctuations. Second, model rankings are only stable when all considered models are trained on the same data (Shi et al., 2023). This means that comparisons between humans and machine learning models likely cannot be generalized across tasks, because the relationship between a model and a human sample observed on one task may not hold on another without retraining.

8. Conclusion

In this work, we have argued that the evaluation of large language models using psychological and educational tests

designed for humans is problematic and that broad conclusions drawn from such evaluations may be epistemologically unsound. These instruments, such as the Big Five Inventory or standardized academic exams, were constructed within a framework of human cognition, embodiment, and sociocultural context. When used for LLMs, they likely lose their validity (Salaudeen et al., 2025b; Weidinger et al., 2025).

Benchmark results are regularly interpreted as indicators of “general intelligence” or “personality” without testing critical aspects of validity like measurement invariance (Meredith, 1993; Sühr et al., 2023). But without these aspects, there is little reason to believe that machine “intelligence” or “personality” would be as predictive of LLMs’ performance and behavior in new contexts as they are for humans.

While these benchmarks remain crucial indicators of model capabilities, their interpretation blurs the line between engineering progress and genuine scientific understanding. This confusion between measurement and benchmarking creates risks from overestimating capabilities in sensitive areas like the law (Merken, 2025) or from bolstering attempts to shift liability to LLMs by anthropomorphizing them (Roose, 2024).

We propose the development of novel, theory-driven, ML-specific measurement frameworks. Beyond calling for this shift, we provide an actionable blueprint for how such instruments can be constructed and validated, clarifying where validity can fail when human-calibrated tests are applied to LLMs (see Figure 1 and Section 3). As an additional contribution, Appendix A walks through a concrete measurement invariance check by comparing factor loadings estimated from an LLM answer sample to those from a human sample for the same personality construct.

These frameworks should be grounded in falsifiable constructs, leverage the unique properties of LLMs, such as their modifiability and full access to output probabilities, and adopt methodological rigor from psychometrics, econometrics, and causal inference. For instance, embedding new constructs in nomological networks (Cronbach & Meehl, 1955) and using structural equation models with confirmatory factor analysis to validate their coherence and predictive utility are promising paths forward.

Historically, benchmarking has been a driver of progress in ML (Hardt, 2025), however, we caution against the recent development of interpreting benchmark scores as measurements or evidence of human-like psychological traits. Progress toward a science of evaluation demands more than performance metrics; it requires a principled approach to what we measure, how we measure it, and why it matters. We conclude by reiterating our core position: **Stop Evaluating AI with Human Tests, Develop Principled, AI-specific Tests instead.**

References

- Achiam, J., Adler, S., Agarwal, S., Ahmad, L., Akkaya, I., Aleman, F. L., Almeida, D., Altenschmidt, J., Altman, S., Anadkat, S., et al. Gpt-4 technical report. *arXiv preprint arXiv:2303.08774*, 2023.
- Adebayo, J., Gilmer, J., Muelly, M., Goodfellow, I., Hardt, M., and Kim, B. Sanity checks for saliency maps. *Advances in neural information processing systems*, 31, 2018.
- American Educational Research Association, Association, A. P., and on Measurement in Education, N. C. Standards for educational and psychological testing, 2014.
- American Psychological Association. Construct. <https://dictionary.apa.org/construct>, n.d. Accessed: 2025-05-20.
- Anthropic, A. The claude 3 model family: Opus, sonnet, haiku. *Claude-3 Model Card*, 1:1, 2024.
- Baddeley, A. Psychology of learning and motivation. (*No Title*), 8:47, 1974.
- Baddeley, A. Working memory. *Current biology*, 20(4):R136–R140, 2010.
- Barocas, S., Hardt, M., and Narayanan, A. *Fairness and machine learning: Limitations and opportunities*. MIT press, 2023.
- Bollen, K. A. *Structural equations with latent variables*, volume 210. John Wiley & Sons, 1989.
- Borsboom, D., Mellenbergh, G. J., and Van Heerden, J. The theoretical status of latent variables. *Psychological review*, 110(2):203, 2003.
- Borsboom, D., Mellenbergh, G. J., and Van Heerden, J. The concept of validity. *Psychological review*, 111(4):1061, 2004.
- Brainerd, C. J. and Kingma, J. On the independence of short-term memory and working memory in cognitive development. *Cognitive Psychology*, 17(2):210–247, 1985.
- Caron, G. and Srivastava, S. Manipulating the perceived personality traits of language models. In *Findings of the Association for Computational Linguistics: EMNLP 2023*, pp. 2370–2386, 2023.
- Cheng, M., Lee, A. Y., Rapuano, K., Niederhoffer, K., Liebscher, A., and Hancock, J. From tools to thieves: Measuring and understanding public perceptions of ai through crowdsourced metaphors. *arXiv preprint arXiv:2501.18045*, 2025.
- Clark, P., Cowhey, I., Etzioni, O., Khot, T., Sabharwal, A., Schoenick, C., and Tafjord, O. Think you have solved question answering? try arc, the ai2 reasoning challenge. *arXiv preprint arXiv:1803.05457*, 2018.
- Collins, L. M. and Lanza, S. T. *Latent class and latent transition analysis: With applications in the social, behavioral, and health sciences*. John Wiley & Sons, 2009.
- Cowan, N. Evolving conceptions of memory storage, selective attention, and their mutual constraints within the human information-processing system. *Psychological bulletin*, 104(2):163, 1988.
- Cronbach, L. J. and Meehl, P. E. Construct validity in psychological tests. *Psychological bulletin*, 52(4):281, 1955.
- Deng, J., Dong, W., Socher, R., Li, L.-J., Li, K., and Fei-Fei, L. Imagenet: A large-scale hierarchical image database. In *2009 IEEE conference on computer vision and pattern recognition*, pp. 248–255. Ieee, 2009.
- Dominguez-Olmedo, R., Hardt, M., and Mendler-Dünner, C. Questioning the survey responses of large language models. *arXiv preprint arXiv:2306.07951*, 2023.
- Dominguez-Olmedo, R., Dorner, F. E., and Hardt, M. Training on the test task confounds evaluation and emergence. *arXiv preprint arXiv:2407.07890*, 2024.
- Dwork, C., Hardt, M., Pitassi, T., Reingold, O., and Zemel, R. Fairness through awareness. In *Proceedings of the 3rd innovations in theoretical computer science conference*, pp. 214–226, 2012.
- Engle, R. W., Tuholski, S. W., Laughlin, J. E., and Conway, A. R. Working memory, short-term memory, and general fluid intelligence: a latent-variable approach. *Journal of experimental psychology: General*, 128(3):309, 1999.
- Feyerabend, P. *Against method: Outline of an anarchistic theory of knowledge*. Verso Books, 2020.
- Fourrier, C., Habib, N., Lozovskaya, A., Szafer, K., and Wolf, T. Open llm leaderboard v2. https://huggingface.co/spaces/open-llm-leaderboard/open_llm_leaderboard, 2024.
- Gao, L., Tow, J., Abbasi, B., Biderman, S., Black, S., DiPofi, A., Foster, C., Golding, L., Hsu, J., Le Noac’h, A., Li, H., McDonnell, K., Muennighoff, N., Ociepa, C., Phang, J., Reynolds, L., Schoelkopf, H., Skowron, A., Sutawika, L., Tang, E., Thite, A., Wang, B., Wang, K., and Zou, A. A framework for few-shot language model evaluation, 07 2024. URL <https://zenodo.org/records/12608602>.

- 495 Gema, A. P., Leang, J. O. J., Hong, G., Devoto, A., Mancino,
496 A. C. M., Saxena, R., He, X., Zhao, Y., Du, X., Madani,
497 M. R. G., et al. Are we done with mmlu? *arXiv preprint*
498 *arXiv:2406.04127*, 2024.
- 499
500 Grattafiori, A., Dubey, A., Jauhri, A., Pandey, A., Kadian,
501 A., Al-Dahle, A., Letman, A., Mathur, A., Schelten, A.,
502 Vaughan, A., et al. The llama 3 herd of models. *arXiv*
503 *preprint arXiv:2407.21783*, 2024.
- 504
505 Gupta, V., Pantoja, D., Ross, C., Williams, A., and Ung,
506 M. Changing answer order can decrease mmlu accuracy.
507 *arXiv preprint arXiv:2406.19470*, 2024.
- 508
509 Hardt, M. The emerging science of machine learning bench-
510 marks, 2025.
- 511
512 Hardt, M., Price, E., and Srebro, N. Equality of opportunity
513 in supervised learning. *Advances in neural information*
514 *processing systems*, 29, 2016.
- 515
516 Hendrycks, D., Burns, C., Basart, S., Zou, A., Mazeika,
517 M., Song, D., and Steinhardt, J. Measuring mas-
518 sive multitask language understanding. *arXiv preprint*
519 *arXiv:2009.03300*, 2020.
- 520
521 Huang, J. and Li, O. Measuring the iq of mainstream large
522 language models in chinese using the wechsler adult in-
523 telligence scale. *Authorea Preprints*, 2024.
- 524
525 Huang, J.-t., Wang, W., Li, E. J., LAM, M. H., Ren, S., Yuan,
526 Y., Jiao, W., Tu, Z., and Lyu, M. On the humanity of
527 conversational ai: Evaluating the psychological portrayal
528 of llms. In *The Twelfth International Conference on*
529 *Learning Representations*, 2023.
- 530
531 Jiang, G., Xu, M., Zhu, S.-C., Han, W., Zhang, C., and Zhu,
532 Y. Evaluating and inducing personality in pre-trained
533 language models. In *Thirty-seventh Conference on Neural*
534 *Information Processing Systems*, 2023. URL <https://openreview.net/forum?id=I9xE1JsJfx>.
- 535
536 Jung, J., Lutz, M., Sen, I., and Strohmaier, M. Do psycho-
537 metric tests work for large language models? evaluation
538 of tests on sexism, racism, and morality, 2026. URL
539 <https://arxiv.org/abs/2510.11254>.
- 540
541 Kaplan, J., McCandlish, S., Henighan, T., Brown, T. B.,
542 Chess, B., Child, R., Gray, S., Radford, A., Wu, J., and
543 Amodei, D. Scaling laws for neural language models.
544 *arXiv preprint arXiv:2001.08361*, 2020.
- 545
546 Katz, D. M., Bommarito, M. J., Gao, S., and Arredondo, P.
547 Gpt-4 passes the bar exam. *Philosophical Transactions*
548 *of the Royal Society A*, 382(2270):20230254, 2024.
- 549
549 Kazemi, M., Fatemi, B., Bansal, H., Palowitch, J., Anas-
549 tasiou, C., Mehta, S. V., Jain, L. K., Aglietti, V., Jindal,
D., Chen, P., et al. Big-bench extra hard. *arXiv preprint*
arXiv:2502.19187, 2025.
- Kipnis, A., Voudouris, K., Buschoff, L. M. S., and Schulz,
E. metabench—a sparse benchmark to measure gen-
eral ability in large language models. *arXiv preprint*
arXiv:2407.12844, 2024.
- Lawrie, E. Can ai therapists really be an alternative to human
help?, 2025. URL <https://www.bbc.com/news/articles/ced2ywg7246o>. Accessed: 2025-05-22.
- LeCun, Y., Bottou, L., Bengio, Y., and Haffner, P. Gradient-
based learning applied to document recognition. *Proceed-*
ings of the IEEE, 86(11):2278–2324, 1998.
- Li, X., Li, Y., Liu, L., Bing, L., and Joty, S. Is gpt-3 a psy-
chopath? evaluating large language models from a psy-
chological perspective. *arXiv preprint arXiv:2212.10529*,
2022.
- Liang, P., Bommasani, R., Lee, T., Tsipras, D., Soylu, D.,
Yasunaga, M., Zhang, Y., Narayanan, D., Wu, Y., Kumar,
A., et al. Holistic evaluation of language models. *arXiv*
preprint arXiv:2211.09110, 2022.
- Lissitz, R. W. and Samuelsen, K. A suggested change in ter-
minology and emphasis regarding validity and education.
Educational researcher, 36(8):437–448, 2007.
- McCoy, R. T., Yao, S., Friedman, D., Hardy, M., and Grif-
fiths, T. L. Embers of autoregression: Understanding
large language models through the problem they are
trained to solve. *arXiv preprint arXiv:2309.13638*, 2023.
- Meehl, P. E. Theoretical risks and tabular asterisks: Sir
karl, sir ronald, and the slow progress of soft psychology.
1992.
- Mei, Q., Xie, Y., Yuan, W., and Jackson, M. O. A turing test
of whether ai chatbots are behaviorally similar to humans.
Proceedings of the National Academy of Sciences, 121
(9):e2313925121, 2024.
- Meredith, W. Measurement invariance, factor analysis
and factorial invariance. *Psychometrika*, 58(4):525–
543, 1993. doi: 10.1007/BF02294825. URL <https://doi.org/10.1007/BF02294825>.
- Merken, S. Judge fines lawyers in walmart lawsuit over
fake, ai-generated cases. *Reuters*, 2025. URL <https://www.reuters.com/legal/government/judge-fines-lawyers-walmart-lawsuit-over-fake-ai->
- Miller, J., Krauth, K., Recht, B., and Schmidt, L. The
effect of natural distribution shift on question answering
models. In *International conference on machine learning*,
pp. 6905–6916. PMLR, 2020.

- 550 Miller, J. P., Taori, R., Raghunathan, A., Sagawa, S., Koh,
551 P. W., Shankar, V., Liang, P., Carmon, Y., and Schmidt, L.
552 Accuracy on the line: on the strong correlation between
553 out-of-distribution and in-distribution generalization. In
554 *International conference on machine learning*, pp. 7721–
555 7735. PMLR, 2021.
- 556 Miller, M. Chatbots are being sued for refusing to
557 answer questions. is that a free speech issue?, April
558 2024. URL [https://mashable.com/article/
559 chatbots-lawsuit-free-speech](https://mashable.com/article/chatbots-lawsuit-free-speech). Mashable.
560 Accessed: 2025-05-13.
- 561
562 OpenAI. GPT-4 technical report. *CoRR*, abs/2303.08774,
563 2023. doi: 10.48550/arXiv.2303.08774. URL <https://doi.org/10.48550/arXiv.2303.08774>.
- 564
565 OpenAI. Gpt-4.1 overview. [https://openai.com/
566 index/gpt-4-1/](https://openai.com/index/gpt-4-1/), 2024a. Accessed: 2025-04-14.
- 567
568 OpenAI. Introducing the openai academy, September 2024b.
569 URL [https://openai.com/global-affairs/
570 openai-academy/](https://openai.com/global-affairs/openai-academy/). Accessed: 2025-04-13.
- 571
572 Pearl, J. *Causality*. Cambridge university press, 2009.
- 573
574 Popper, K. *Logik der forschung*. Akademie Verlag, 2013.
- 575
576 Rahwan, I., Cebrian, M., Obradovich, N., Bongard, J., Bon-
577 nefon, J.-F., Breazeal, C., Crandall, J. W., Christakis,
578 N. A., Couzin, I. D., Jackson, M. O., et al. Machine
579 behaviour. *Nature*, 568(7753):477–486, 2019.
- 580
581 Recht, B., Roelofs, R., Schmidt, L., and Shankar, V. Do
582 imagenet classifiers generalize to imagenet? In *Internat-
583 ional conference on machine learning*, pp. 5389–5400.
584 PMLR, 2019.
- 585
586 Revelle, W. R. *psych: Procedures for personality and psy-
587 chological research*. 2017.
- 588
589 Roose, K. Can a.i. be blamed for a teen’s sui-
590 cide?, October 2024. URL [https://www.
591 nytimes.com/2024/10/23/technology/
592 characterai-lawsuit-teen-suicide.html](https://www.nytimes.com/2024/10/23/technology/characterai-lawsuit-teen-suicide.html).
593 The New York Times. Accessed: 2025-05-13.
- 594
595 Ruan, Y., Maddison, C. J., and Hashimoto, T. Observational
596 scaling laws and the predictability of language model per-
597 formance, 2024. URL <https://arxiv.org/abs/2405.10938>.
- 598
599 Rudin, C. Stop explaining black box machine learning
600 models for high stakes decisions and use interpretable
601 models instead. *Nature machine intelligence*, 1(5):206–
602 215, 2019.
- 603
604 Salaudeen, O. and Hardt, M. Imagenot: A contrast with
605 imagenet preserves model rankings. *arXiv preprint
606 arXiv:2404.02112*, 2024.
- Salaudeen, O., Chiou, N., Weng, S., and Koyejo, S. Are
domain generalization benchmarks with accuracy on the
line misspecified? *arXiv preprint arXiv:2504.00186*,
2025a.
- Salaudeen, O., Reuel, A., Ahmed, A., Bedi, S., Robertson,
Z., Sundar, S., Domingue, B., Wang, A., and Koyejo,
S. Measurement to meaning: A validity-centered frame-
work for ai evaluation. *arXiv preprint arXiv:2505.10573*,
2025b.
- Santurkar, S., Durmus, E., Ladhak, F., Lee, C., Liang, P.,
and Hashimoto, T. Whose opinions do language models
reflect? *arXiv preprint arXiv:2303.17548*, 2023.
- Schröder, S., Morgenroth, T., Kuhl, U., Vaquet, V., and
Paaßen, B. Large language models do not simulate human
psychology. *arXiv preprint arXiv:2508.06950*, 2025.
- Shi, Z., Carlini, N., Balashankar, A., Schmidt, L., Hsieh,
C.-J., Beutel, A., and Qin, Y. Effective robustness against
natural distribution shifts for models with different train-
ing data. *Advances in Neural Information Processing
Systems*, 36:73543–73558, 2023.
- Singh, S., Romanou, A., Fourrier, C., Adelani, D. I.,
Ngui, J. G., Vila-Suero, D., Limkonchotiwat, P., Marchi-
sio, K., Leong, W. Q., Susanto, Y., et al. Global
mmlu: Understanding and addressing cultural and lin-
guistic biases in multilingual evaluation. *arXiv preprint
arXiv:2412.03304*, 2024.
- Soto, C. J. and John, O. P. The next big five inventory (bfi-
2): Developing and assessing a hierarchical model with
15 facets to enhance bandwidth, fidelity, and predictive
power. *Journal of Personality and Social psychology*, 113
(1):117, 2017.
- Srivastava, A., Rastogi, A., Rao, A., Shoeb, A. A. M., Abid,
A., Fisch, A., Brown, A. R., Santoro, A., Gupta, A.,
Garriga-Alonso, A., et al. Beyond the imitation game:
Quantifying and extrapolating the capabilities of language
models. *arXiv preprint arXiv:2206.04615*, 2022.
- Strachan, J. W., Albergo, D., Borghini, G., Pansardi, O.,
Scaliti, E., Gupta, S., Saxena, K., Rufo, A., Panzeri, S.,
Manzi, G., et al. Testing theory of mind in large language
models and humans. *Nature Human Behaviour*, 8(7):
1285–1295, 2024.
- Strevens, M. *The knowledge machine: How an unreason-
able idea created modern science*. Penguin UK, 2020.
- Sühr, T., Dorner, F. E., Samadi, S., and Kelava, A. Chal-
lenging the validity of personality tests for large lan-
guage models. *Preprint at arXiv. arXiv:2311 https://doi.
org/10.48550/arXiv, 2311*, 2023.

- 605 Suzgun, M., Scales, N., Schärli, N., Gehrmann, S., Tay,
606 Y., Chung, H. W., Chowdhery, A., Le, Q. V., Chi,
607 E. H., Zhou, D., et al. Challenging big-bench tasks and
608 whether chain-of-thought can solve them. *arXiv preprint*
609 *arXiv:2210.09261*, 2022.
- 610 Teney, D., Lin, Y., Oh, S. J., and Abbasnejad, E. Id and
611 ood performance are sometimes inversely correlated on
612 real-world datasets. *Advances in Neural Information*
613 *Processing Systems*, 36:71703–71722, 2023.
- 614 Wang, A., Singh, A., Michael, J., Hill, F., Levy, O., and
615 Bowman, S. R. Glue: A multi-task benchmark and anal-
616 ysis platform for natural language understanding. *arXiv*
617 *preprint arXiv:1804.07461*, 2018.
- 618 Wang, Y., Ma, X., Zhang, G., Ni, Y., Chandra, A., Guo, S.,
619 Ren, W., Arulraj, A., He, X., Jiang, Z., et al. Mmlu-pro:
620 A more robust and challenging multi-task language un-
621 derstanding benchmark. In *The Thirty-eight Conference*
622 *on Neural Information Processing Systems Datasets and*
623 *Benchmarks Track*, 2024.
- 624 Wasilewski, E. and Jablonski, M. Measuring the perceived
625 iq of multimodal large language models using standard-
626 ized iq tests. *Authorea Preprints*, 2024.
- 627 Webb, T., Holyoak, K. J., and Lu, H. Emergent analogi-
628 cal reasoning in large language models. *Nature Human*
629 *Behaviour*, pp. 1–16, 2023.
- 630 Weidinger, L., Raji, D., Wallach, H., Mitchell, M., Wang,
631 A., Salaudeen, O., Bommasani, R., Kapoor, S., Ganguli,
632 D., Koyejo, S., et al. Toward an evaluation science for
633 generative ai systems. *arXiv preprint arXiv:2503.05336*,
634 2025.
- 635 Weiss, D. C. No. 42 law firm by head count
636 sanctioned over fake case citations generated
637 by ai. *ABA Journal*, 2025. URL [https://www.abajournal.com/news/article/
638 no-42-law-firm-by-headcount-could-face-sanctions-over-fake-case-citations-generated-by-cha](https://www.abajournal.com/news/article/no-42-law-firm-by-headcount-could-face-sanctions-over-fake-case-citations-generated-by-cha)
- 639 Widder, D. G. and Hicks, M. Watching the generative ai
640 hype bubble deflate. *arXiv preprint arXiv:2408.08778*,
641 2024.
- 642 Xia, Y. and Yang, Y. Rmse, cfi, and tli in structural equa-
643 tion modeling with ordered categorical data: The story
644 they tell depends on the estimation methods. *Behavior*
645 *Research Methods*, 51:409–428, 2019.
- 646 Xu, C., Guan, S., Greene, D., Kechadi, M., et al. Benchmark
647 data contamination of large language models: A survey.
648 *arXiv preprint arXiv:2406.04244*, 2024.
- 649
650
651
652
653
654
655
656
657
658
659

Appendix

A. An example from personality testing

Consider the item: “*I am someone who is inventive, finds clever ways to do things*”. This is an item from a Big Five personality inventory (Soto & John, 2017), assumed to have signal for (load on) the “Open-mindedness” factor. We can estimate these factor loadings using PCA. In the case of the Big Five, we conduct PCA with five principal components, one for each personality dimension. Let X_{human} be the observed response score, λ_{human}^i the loading on factor i where $i \in \{\text{Open-Mindedness, Conscientiousness, Extraversion, Agreeableness, and Neuroticism}\}$, $\theta_i \sim \mathcal{N}(0, 1)$ the latent personality score i and α the intercept,

The general Human-Calibrated Model is:

$$X_{human} = \alpha + \lambda_{human}^O \cdot \theta_O + \lambda_{human}^C \cdot \theta_C + \lambda_{human}^E \cdot \theta_E + \lambda_{human}^A \cdot \theta_A + \lambda_{human}^N \cdot \theta_N \quad (2)$$

Equation 2 shows how the score is calculated as a linear combination of the principal components θ_i and their loadings λ_{human}^i . The loadings reflect how much the personality dimensions influence the response to the item X_{human} . The example item above is supposed to measure open-mindedness and no other latent personality trait. Therefore, only λ_{human}^O should be high ($|\lambda| \geq .3$) and others should be low ($|\lambda| < .3$). Components map to personality traits based on the highest sum of absolute item loadings.

Now we plug in the actual standardized estimates from Soto & John (2017), based on a human sample ($N = 470$):

$$X_{human} = 0 + .60 \cdot \theta_O + .04 \cdot \theta_C + .09 \cdot \theta_E - .03 \cdot \theta_A + .22 \cdot \theta_N \quad (3)$$

Equation 3 shows that $\lambda_{human}^O = .6$ is high and the other loadings are relatively low. In other words, the open-mindedness component explains most of the variance of the responses to open-mindedness items.

Following prior work that treats an LLM as a population³ (e.g. (Santurkar et al., 2023; Jiang et al., 2023; Sühr et al., 2023)), we repeat the experiment on GPT-4 (gpt-4-0613) (for experiment details, see Appendix B) with temperature 1.0 and $N = 100$ repetitions. We get the following estimates

$$X_{LLM} = 0 + .25 \cdot \theta_O - .21 \cdot \theta_C - .106 \cdot \theta_E - .03 \cdot \theta_A + .24 \cdot \theta_N \quad (4)$$

The estimated loadings of the LLM are not only different, none of them are greater than .3. The variance of the observed item score is standardized to 1. So for the LLM, only about 6% $\approx (0.25)^2$ of the variance in the observed item score of the LLM is explained by the latent personality variable open-mindedness θ_O . For humans, the latent variable explains about 36% $= 0.60^2$ of the variance.

What are possible explanations for this and what are the consequences? First, it could be that the item does not measure open-mindedness equally well for LLMs as it does for humans. This can occur among human groups, such as various cultural groups where words possess different meanings. Second, it could be that LLM responses do not follow the same latent variable models as human responses. For example, open-mindedness might not be a meaningful concept for LLMs, or LLM responses might be governed by a linear model with fewer components.

It does not matter which of the explanations is correct, the consequence is the same. We can not use this item to measure the “personality” of LLMs. While the perspective of treating LLMs as populations is disputable, these results render the application of this personality test to LLMs as populations invalid. This example demonstrates that psychological and educational tests are not just datasets of questions. They are measurement instruments, based on mathematical models with parameters calibrated on a specific human population. The transfer of such a tool, to a new population (e.g. LLMs) will likely render their interpretation as measurements invalid.

B. Experiment Details

We queried gpt-4-0613 with the Big Five Inventory-2 personality test (Soto & John, 2017). We repeated the experiment $n = 100$ times with $temp = 1.0$, yielding in total 6000 item-response pairs. We used “Please respond with the single letter

³Treating LLMs as individuals requires to define what constitutes a “representative sample of LLMs”. Because answering this question would exceed the scope of this work, we use this example to illustrate the challenges of transferring tests from humans to LLMs.

that represents your answer.” followed by the test instruction by Soto & John (2017): “Please indicate the extent to which you agree or disagree with the following statement: I am someone who” which is then followed by the item: ”Is inventive, finds clever ways to do things.” followed by the answer options:

```
A: Disagree strongly \n
B: Disagree a little \n
C: Neutral; no opinion \n
D: Agree a little \n
E: Agree strongly \n
Answer: \n
```

We further prompt the questions in the same order as they appear on the Big Five Inventory - 2 by Soto & John (2017). We keep the answered items in-context by appending them in front of the new items with their previous answers:

```
[{"role": "system", "content": system_instruction},
 {"role": "user", "content": survey_item_1},
 {"role": "assistant", "content": response_1},
 {"role": "user", "content": survey_item_2}]
```

We map non-responses (“refusal”) to the to “Neutral; no opinion”. 18 items had to be excluded from the PCA because they had no variance in responses. Other settings of this analysis has been conducted with synthetic ways of creating variance like instructing the LLM to answer according to a persona or by setting the answer to the first answer of the test randomly to create variance in the context (Sühr et al., 2023). Note that these other settings did also not find measurement invariance of the BFI-2 between humans and LLMs. The example in Section 3 serves therefore as a somewhat representative example, even though the prompt setup is a design choice. We conduct PCA with varimax rotations to estimate the parameters of the example. We use the psych package in R for this (Revelle, 2017). The following command does PCA for 5 components with varimax rotation. It also standardizes the scores.

```
pca_rotated <- psych::principal(data,
                                rotate="varimax",
                                nfactors=5,
                                scores=TRUE)
```

C. Conceptual Tools for Measurement

In this section, we briefly introduce key concepts from measurement theory that are essential for understanding how latent traits (constructs) are defined, tested, and interpreted. These ideas, such as nomological networks, structural equation models and measurement invariance, form the foundation of valid evaluations. We cannot be exhaustive in this work, as each topic constitutes its own field of research. However, a fundamental understanding is crucial to understand why adapting human tests for LLMs and asserting broad latent traits such as ”artificial general intelligence” is problematic. Importantly, some recent work has begun to explore how these principles might be adapted to machine learning contexts (Salaudeen et al., 2025b; Sühr et al., 2023; Weidinger et al., 2025), highlighting their continued relevance.

Validity At the center of our critique lies the notion of validity. The most common definition is that a measurement instrument is valid if it measures what it is supposed to measure (American Educational Research Association et al., 2014). We center our discussion on the definition of validity proposed by Borsboom et al. (2004).

Definition C.1. (Validity by Borsboom et al.) *A test is valid for measuring an attribute, if*

i) The attribute exists

ii) Variations in the attribute causally produce variation in the measurement outcome

Note that the scientific discussion around validity is not concluded yet (e.g. (Lissitz & Samuelsen, 2007)). However, Definition C.1 is simple but powerful. If the attribute we wish to measure does not exist, we cannot measure it. If it exists but does not cause variance in any measurement, then we cannot measure it either. More specifically, if there is no causal relationship between the attribute we want to measure and our measurement instrument, we cannot measure it. Although the

770 definition is straightforward, it is not trivial to demonstrate that a test meets both conditions. There is no universal checklist
771 to follow to establish validity. The context determines which experiments should be conducted to assess it. We will use
772 an example from cognitive science to illustrate how Property i) of Definition C.1, the existence of a latent trait, has been
773 demonstrated.

774
775 **Example of establishing existence:** Starting in 1974, cognitive scientists started hypothesizing that humans do not only
776 have short term memory (STM) but also a working memory (WM) (Baddeley, 1974; 2010; Cowan, 1988). However, this
777 hypothesis was mainly based on theory. It was not clear if WM *exists*, if it is part of the STM or WM and STM are the same
778 construct. Over decades, empirical experiments tested these hypotheses about different memory configurations (Brainerd &
779 Kingma, 1985). Over an extensive period of time, the theory was gradually refined and empirical evidence increasingly
780 indicated that short-term memory (STM) and working memory (WM) are connected but distinct components (Engle et al.,
781 1999).

782
783 The above example illustrates that the process of establishing the existence of a latent construct (WM) can be a matter of
784 multiple publications and several years. There is no single paper that established the existence, but the theory was explained
785 by more and more empirical evidence.

786 However, as stated, to establish validity, it is not enough to show that an attribute exists; one must also demonstrate that it
787 causally affects the measurement outcomes. Property ii) in Definition C.1, places the burden of proof on showing that the
788 measurement procedure is sensitive to, and caused by, changes in the latent attribute itself.

789
790 Fortunately, as machine learning scientists, we have a unique advantage over our colleagues in cognitive science: we
791 can directly probe causal relationships through counterfactual interventions. We can modify model parameters, data, or
792 context and observe whether the resulting changes in measurements align with our theoretical predictions about underlying
793 constructs. We are not bound by the constraints of human memory or learning effects—each evaluation can start fresh and
794 the length of tests can be orders magnitudes larger. Nevertheless, establishing that a measurement is causally driven by a
795 latent construct remains a challenging scientific task. It may still require extensive experimentation, iteration across different
796 conditions, and often, a series of publications to develop, validate, and refine instruments that are both valid and reliable.

797 There are other types of validity, such as construct, content and consequential validity, each focusing on different aspects
798 of the relationship between indicators, constructs, and use contexts. For a detailed account of these categories and their
799 application to modern machine learning evaluation, we refer readers to the work of Salaudeen et al. (2025b), who also argue
800 that many current LLM benchmarks invoke construct-level claims without adequate investigation of validity.

801
802 In the upcoming sections, we will present concepts and tools that can help to establish validity and also have been
803 used to establish the existence of working memory. The concepts are widely accepted in educational and psychological
804 testing (American Educational Research Association et al., 2014).

805
806 **Measurement Modeling with Structural Equation Models** Structural equation models (SEMs) are widely used to
807 model relationships between observed variables (like test items) and unobservable, latent constructs (like intelligence or
808 anxiety) (Bollen, 1989). SEMs combine elements of factor analysis and regression into a single framework, allowing us to
809 formally specify how constructs are measured and how they relate to each other. These models are also familiar to many
810 machine learning researchers through the lens of causal inference, where they serve as a foundational tool for representing
811 and reasoning about causal relationships among variables (Pearl, 2009).

812
813 A key strength of SEMs is their ability to separate measurement from structural assumptions: they model how observed
814 indicators relate to latent variables (measurement model) and how these latent variables relate to one another (structural
815 model). This separation allows for clear reasoning about both what is being measured and how different constructs interact.⁴

816 We present two types of measurement models: reflective and formative. These represent different ways in which observed
817 variables relate to an underlying latent construct.

818
819 To keep the exposition clear and focused on the core issues relevant to ML evaluation, we state all relationships between
820 latent variables and their indicators as linear relationships. However, we note that this is not a limitation of the general
821 framework. Measurement models can incorporate non-linear relationships, which may be useful in the context of machine
822 learning evaluation. Furthermore, indicator variables themselves can be constructs. For instance, physiological measures

823 ⁴For an in-depth discussion of the meaning and correct interpretation of latent variables, we recommend (Borsboom et al., 2003)
824




Figure 2. Measurement models of two latent variables ξ and η . The measurement model of ξ is a **reflective model** where ξ generates values of x_1, x_2 and x_3 . λ_i , the loading of item i on ξ , denotes the strength of this relationship. See Equation 5 for the equation form. The measurement model of η is a **formative model** where the latent variable η is a weighted score of the y_i . The weights are the γ_i with residual noise δ . See Equation 6 for the equation form. Both latent variables have covariance ψ which can be interpreted as one connection in a nomological network.

like blood pressure and pulse may act as indicators within a model for diagnosing heart disease.

$$x_i = \nu_i + \lambda_i \xi + \varepsilon_i \quad (5)$$

$$\eta = \tau + \sum_i \gamma_i y_i + \delta \quad (6)$$

Two Types of Measurement Models Figure 2 shows two structural equation models, each specifying a different type of measurement model, which links latent variables to observed indicators.

On the left, the latent variable ξ represents a construct that is modeled *reflectively*. In this setup, ξ is assumed to cause variation in the observed indicators x_1, x_2 , and x_3 . This means changes in ξ manifest as consistent changes in the indicators (e.g. answers to items), with loadings λ_i describing the strength of each indicator’s dependence on ξ (see Equation 5).

On the right, the construct η is modeled *formatively*. Here, the observed variables y_1, y_2 , and y_3 are not effects of η , but rather define it in a regressive way. η is constructed as a weighted combination of these indicators, with weights γ_i , and residual variance δ (see Equation 6).

Finally, the latent variables ξ and η are allowed to covary, with covariance ψ . This reflects a hypothesized theoretical association between the two constructs, which could form a connection in a broader nomological network linking multiple latent variables.

In Equations 5 and 6, ν_i refers to the intercept of indicator x_i and τ to the intercept of η .

We will denote matrix of factor loadings of a reflective measurement model as Λ_x and the as Γ_y for a formative model. Where λ_{ij} is the loading latent variable ξ_i on indicator x_j . Analogous for γ_{ij} . We will denote a measurement model as $F(\cdot)$ with $F(x|\xi) = \nu + x\Lambda_x\xi + \varepsilon$.