
Position: Meaning Is Not A Metric: Using LLMs to make cultural context legible at scale

Anonymous Author(s)

Affiliation

Address

email

Abstract

1 This position paper argues that large language models (LLMs) can make cultural
2 context, and therefore human meaning, legible at an unprecedented scale in AI-
3 based sociotechnical systems. We argue that such systems have previously been
4 unable to represent human meaning because they rely on thin descriptions: numeri-
5 cal representations that enforce standardization and therefore strip human activity
6 of the cultural context that gives it meaning. By contrast, scholars in the humani-
7 ties and qualitative social sciences have developed frameworks for representing
8 meaning through thick description: verbal representations that accommodate het-
9 erogeneity and retain contextual information needed to represent human meaning.
10 While these methods can effectively codify meaning, they are difficult to deploy at
11 scale. However, the verbal capabilities of LLMs now provide a means of (at least
12 partially) automating the generation and processing of thick descriptions, poten-
13 tially overcoming this bottleneck. We argue that the problem of rendering human
14 meaning legible is not just about selecting better metrics, but about developing new
15 representational formats (based on thick description). We frame this as a crucial
16 direction for the application of generative AI and identify five key challenges:
17 preserving context, maintaining interpretive pluralism, integrating perspectives
18 based on lived experience and critical distance, distinguishing qualitative content
19 from quantitative magnitude, and acknowledging meaning as dynamic rather than
20 static. Furthermore, we suggest that thick description has the potential to serve as a
21 unifying framework to address a number of emerging concerns about the difficulties
22 of representing culture in (or using) LLMs. In addressing these challenges, we
23 present a pathway to developing systems that can better represent and support
24 meaningful human experiences across domains including healthcare, education,
25 and sustainability.

26 1 Introduction

27 The premise of this paper is that it is important, but difficult, to develop technologies that can support
28 meaningful human experience. Our position aims to articulate how AI technologies can be used not
29 just to yield superficial increases in convenience, but to offer something deeper: a tool for helping
30 people tap more directly into *what really matters*. This is not an inevitable outcome of technological
31 progress. Dystopian fiction is full of cases in which a hypothetical society has advanced far beyond
32 our own, only to find that their sophisticated technology has had a detrimental effect on human
33 experience. Of course, the easy part is proposing that technology should have positive societal
34 impacts; the difficult part is defining “what really matters.”

35 For example, consider the case of social media. Initially these platforms were presented as a tool for
36 human social connection [10]. Over time, however, they came to reflect the more superficial dynamics
37 of engagement metrics; content was prioritized not to facilitate more profound social interactions, but
38 on the basis of likes, views, and shares [103, 30, 63]. This led to a wide range of negative societal
39 outcomes, such as rampant misinformation [55, 106], increased feelings of loneliness [78, 101],
40 addictive behavior [40], and other psychological and social costs [22].

41 In response, technologists, social scientists, and activists have proposed other methods for prioritizing
42 content on these platforms—ones that would, in one form another, allow us to more directly target
43 what really matters [34, 97, 19]. A representative example is the Time Well Spent movement (now
44 the Center for Humane Technology) led by Tristan Harris [34]. The goal of this movement was to
45 develop technology that supports “meaningful” social connections and “choices we cherish” [72]. In
46 response, at least one major social media platform alleged to have changed their recommendation
47 algorithm to prioritize these kind of “meaningful” interactions [96]. However, it was disputed whether
48 this effort had the intended effect, or simply substituted one superficial metric for another [72, 63].

49 While laudable in theory, proposals about orienting technology around meaningful experience have
50 remained difficult to implement in practice. This is not a new problem, nor is it limited to social media
51 platforms. For instance, in the early twentieth century Bertrand Russell predicted that technological
52 development would soon lead to the widespread adoption of a four-day work week, thus allowing
53 people to spend less time on drudgery and more time on what they actually care about [84]. Despite
54 continued technological advancement, Russell’s prediction has yet to materialize.

55 We contend that in order to develop sociotechnical systems that can increase time well spent or reduce
56 drudgery, we first need a way of systematically codifying what kind of experiences humans find
57 meaningful. We approach this problem—how to design sociotechnical systems that support human
58 meaning—as one of representation. The core of our argument is that this issue has been difficult to
59 address not simply because we have been using the wrong metrics, but because “metrics” (as we have
60 traditionally conceived of them) are the wrong representational format. The formats we typically
61 use for measuring behavior at scale exclude key contextual information necessary for representing
62 human meaning. The primary aim of this paper will be to describe why that is the case and how we
63 can begin to address it.

64 **Our position is that we can use LLMs to make human meaning legible at an unprecedented scale;**
65 **that “thin” metrics alone will not be able to do this, because they cannot adequately represent**
66 **cultural context; and that developing “thick” representational formats might eventually allow**
67 **our society’s most powerful and impactful sociotechnical systems to support the heterogeneous**
68 **range of experiences that really matter to people.** We argue that such systems have previously
69 been unable to represent human meaning because they typically rely on thin descriptions: numerical
70 representations that enforce standardization, and therefore strip human activity of the broader context
71 that gives it meaning [9, 71, 59, 89]). By contrast, scholars in the humanities and qualitative social
72 sciences have developed robust frameworks for representing meaning based on thick descriptions:
73 verbal representations that allow for heterogeneity, and therefore retain crucial contextual information
74 necessary for analyzing cultural and experiential meaning [27, 108, 87, 82]. These methods provide
75 a way of codifying meaning—but they are difficult to deploy at scale.

76 Previously, reliance on thick descriptions would lead to a bottleneck: only human analysts had the
77 necessary verbal capabilities to generate and evaluate such representations. With LLMs, this is no
78 longer the case. They provide a means of (at least partially) automating the generation and evaluation
79 of thick descriptions. Overall, we make the case that meaning is not a metric: thicker representational
80 formats are needed to render human meaning legible at scale.

81 2 The Legibility of Meaning

82 In his classic political science text *Seeing like a State*, James C. Scott (1998) introduces the concept
83 of legibility [89]. This term describes the connection between a sociotechnical system (in Scott’s
84 case, a political state) and its information processing capacities. Political states cannot represent the
85 goals and desires of their constituents directly. Simply asking each person what they want would
86 result in too much heterogeneity of response. Rather, the state must develop standardized metrics
87 that can be applied across the whole population. These help the state understand the overall trends
88 and dynamics of the populace. However, as Scott argues, by enforcing this standardization the state
89 makes some kinds of information legible—while rendering other kinds illegible.

90 Drawing on Scott’s concept, we frame the main challenge of this paper as how to make human
91 meaning *legible* at scale. To clarify: we are not concerned with how to make human meaning legible
92 within AI models, but rather how AI models can be used to render new kinds of information legible
93 within sociotechnical systems. We use the term *sociotechnical system* to refer to any large-scale
94 configuration of people and machines that cannot be understood in terms of human or technological
95 dynamics alone but that depend on a complex entanglement of their interactions [41, 5, 100]. This
96 includes many systems that have a major impact on society, such as political states, social media
97 platforms, corporations, and institutions for healthcare, education, and science. As AI technologies
98 become increasingly integrated into these societally impactful systems, how can we use these
99 technologies to make new kinds of information (such as human meaning) legible at scale?

100 2.1 What Do We Mean By “Meaning”?

101 The term “meaning” has been used by many different people, across many disciplines, to refer
102 to many different (though often overlapping) concepts [27, 11, 4, 33, 23, 36, 24]. While there is
103 undoubtedly heterogeneity in these different usages, there are also recurring themes—such as the
104 relational, contextual, and contested nature of meaning. In this paper, we are primarily concerned
105 with what we call “human meaning,” to distinguish our usage from debates about whether LLMs
106 really “understand” the *semantic* meaning of the text they generate [6, 50]. We take human meaning
107 to have two components: cultural and experiential.

108 *Cultural meaning*. This is the endorsement of certain symbols as significant within the context of
109 a given community. These symbols could include cultural emblems (such as the Union Jack or the
110 Olympic rings), behaviors (such as bowing or raising one’s middle finger), rituals (such as a marriage
111 ceremony or graduation procession), events (such as the Battle of Waterloo or the Moon landing),
112 or works of art (such as James Joyce’s *Ulysses* or Picasso’s *Guernica*). This is how the term is
113 typically used in humanities and social science fields such as literary studies [32], history [17], and
114 cultural anthropology [27]. This usage is consistent with other usages of “culture” in contemporary
115 AI [21, 113, 86, 1, 44]. For our purposes, this is the kind of meaning we can most directly make
116 legible with existing frameworks from the humanities and social sciences.

117 *Experiential meaning*. This is the individual mental state of believing a particular event or experience
118 to be especially worthwhile or to matter in some grander sense. It is how the term is typically used
119 in fields such as psychology and mental health [36, 105]. More informally, this is the sense used
120 by religion, self-help, and other pathways for orienting one’s life toward *what really matters* [24].
121 Often, experiential meaning is framed as a magnitude estimate: whether a given experience is more
122 or less meaningful than another one [38, 37, 94]. However, these magnitude judgments often miss
123 key aspects of the cultural basis of meaning [11, 12, 14, 15, 35, 104].

124 It is not always easy to disambiguate between cultural and experiential meaning. For example,
125 experiential meaning derives from cultural meaning in ways that are crucial but often difficult to
126 systematize [13, 52, 114]. The more we understand about a person’s cultural background, the more
127 likely we are to understand what they find personally meaningful. But while people’s individual
128 sense of meaning is shaped by their cultural milieu, their individual perceptions of meaning can
129 either conform to a given cultural template—or reject it. While the framework we describe below
130 can be used to codify experiential meaning [11, 60, 109, 82], for the purposes of this paper we will
131 focus primarily on cultural (rather than experiential) meaning, since that is traditionally where thick
132 description has been most directly applied. We see this as a crucial first step toward making human
133 meaning, in its broadest sense, legible within sociotechnical systems.

3 Representing Meaning via Thick Description

In this paper, our primary reference point for meaning is Clifford Geertz’s (1973) account of “thick description [27], which provides a framework for the interpretation of symbols and behaviors within their cultural context to uncover their layers of significance. Methodologically, it is a bedrock of qualitative approaches in the social sciences [87, 64, 58]; theoretically, it is a unifying framework for influential schools of thought in disciplines such as literary criticism [32], history [17], political science [108], philosophy [98], and more. It is arguably one of the most influential frameworks in the humanities and social sciences.¹ Thick description offers a conceptual and methodological bridge between cultural and experiential meaning, as well as a guide for how human meaning can be rendered legible—both to human interpreters and computational systems—at scale.

Geertz proposed thick description as a means of clarifying the goal of ethnographic fieldwork in cultural anthropology. The aim is not just to provide a surface-level account of physical behaviors, rituals, stories, or cultural norms (i.e., *thin* description). Rather, it is to provide an interpretation of the symbolic dimensions of a given behavior or practice—to analyze its role in providing an individual with a larger sense of meaning and purpose in the context of their broader community (i.e., *thick* description). Geertz contended that this approach to the study of culture is “not an experimental science in search of law but an interpretive one in search of meaning” [27].

An example of the contrast between thin and thick description can be seen in different ways of describing an economic activity. For instance, an entire country’s economic activity can be encapsulated in the useful, if incomplete, measure of gross domestic product. This thin description can be used to track whether a nation’s economy is trending in a positive or negative direction but does not provide the language to distinguish between, for example, the economic considerations of purchasing a house (a major outlay of monetary resources) versus the experiential considerations of a buying a home (a significant choice about the composition of one’s family, lifestyle, and community).

By contrast, a canonical instance of thick description is Malinowski’s (1922) investigation of the Kula Ring [57, 29]. The Kula Ring refers to an elaborate ritual in which inhabitants of the Trobriand Islands crossed treacherous swathes of water in small boats to trade necklaces and armbands which in themselves had no obvious utility: the necklaces and armbands were never actually worn. Malinowski’s account showed how this ostensibly economic activity could only be understood by an analysis of its symbolic dimensions. That is, it cannot be accounted for by an analysis of the transactions through an externally observable value metric (such as money). Rather, it requires an analysis of the meaning assigned to such transactions internally by the people participating in them. The Kula necklaces and armbands were symbols of achievement and social standing. Comparably, an observer attempting to understand the value of a Super Bowl trophy in terms of how much water or soup it holds will likely miss important aspects of the broader cultural activity it represents.

As a theory of meaning, thick description operationalizes meaning as a relationship mediating a symbol (including artifacts such as a Super Bowl trophy or the Bible; actions such as attending a protest or making a particular hand gesture; and mental states such as belief in the sanctity of marriage or satisfaction with clearing one’s email inbox) and its sociocultural context. Influenced by Wittgenstein (1953), Geertz claimed that meaning is not an intrinsic property of an individual mind but only exists within the broader context of who is interpreting a given symbol, the background or experience they bring, and the time and place in which the interpretation occurs [27, 110].

As a result, the meaning of any symbol is inherently multiplicitous: no singular meaning prevails, but rather multiple meanings can coexist depending on the interpreter and their sociocultural conditions. For example, a hand signal such as a “thumbs up” may signify approval in Western cultures, but in parts of the Middle East and Africa can be perceived as an offensive expression of contempt. The meaning of a symbolic figure can, therefore, only be understood against its appropriate cultural ground: what makes a description “thick” is the extent to which it provides the necessary contextual grounding to locate a particular meaning.

¹According to Google Scholar’s citation count [20 May 2025], Geertz (1973) has >100k citations. For comparison, foundational frameworks in related disciplines published around the same time include: Kahneman & Tversky’s (1979) prospect theory in the behavioral sciences (89k citations) and Granovetter’s (1973) weak ties social network analysis in the quantitative social science (76k citations).

3.1 Thick and Thin Description in Sociotechnical Systems

Thus, one promising conceptual framework for making meaning legible at scale is to make thick descriptions legible at scale. However, the default format of representation within most sociotechnical systems is thin; generally speaking, only numerical signals—canonical thin descriptions such as GDP, engagement metrics, or stock price—are legible within the information processing capacities of such systems [9, 71, 59, 89]. In this section, we argue that, across a wide range of sociotechnical systems, the use of thin descriptions strips away the kind of contextual data needed to encode a signal of human meaning.

The distinction between thin and thick descriptions is related, but not identical, to the distinction between quantitative and qualitative data. Thin descriptions account for facts, figures, or observable behavior without context; thick descriptions account for the context, interpretations, subjective perspectives, and meanings that make sense of those facts. Not all qualitative data conveys a thick description; otherwise, it would not be necessary to have a theory of what makes a given qualitative analysis more or less thick. Conversely, while thick description traditionally refers to qualitative analysis, it is nonetheless possible for a quantitative figure to capture thick elements (e.g., with multiple variables). The distinction is not binary: a representation can be “thinner” or “thicker.”

The key problem is that, historically, the only way to reach legibility at scale is via thin description. Standardization must be applied, such that every individual performance can be assessed in the same generic terms. However, this kind of standardization is antithetical to encoding meaning via thick description. If concepts can only be made legible within a system on the condition of standardization, then that creates a natural tension with the context-rich information necessary for a thick description.

For example, it is often the case that the primary goal of a system is qualitative: to promote health, education, urban livability, social connection, justice, and so on. But what is legible within the system is quantitative (e.g., body mass index, standardized test scores, crime statistics, social media shares, or recidivism rates); this quantitative metric becomes a tangible proxy for the more abstract, multi-faceted goal. While such data-driven optimization is a central aspect to the functioning of modern sociotechnical systems [75, 71], there are well-documented issues with overreliance on such quantitative signals [66, 43]. We argue that relaxing this mandate for quantitative standardization by creating representations based on thick description can help overcome these issues.

Healthcare is a prime example. Many conceptual definitions of health take a thick perspective. For example, the Constitution of the World Health Organization defines health as “a state of complete physical, mental and social well-being and not merely the absence of disease or infirmity” [111]. Likewise, several prominent theories of health and well-being emphasize the functional nature of health—that what constitutes healthiness depends on what is necessary to physically support a meaningful life [99, 90, 70]. This definition of health is highly context-dependent: there is no single standardized template for health, just as there is no standardized template for the life people want to live. However, hospitals and healthcare systems need tangible metrics by which to measure their efficacy, such as length of stay, readmission rates, and mortality rates. Though important, over-reliance on these thin metrics runs the risk of missing crucial qualitative aspects of care such as emotional, familial, occupational, and social facets of human life [74]. Thicker representations are needed to codify this context-dependent, holistic notion of health in a large-scale sociotechnical system [73].

This basic pattern holds across a range of domains. In education: standardized test scores, grade point average, and graduation rates fail to capture student growth, creativity, critical thinking, and character development [80, 51]. In scientific research: journal impact factors and citation counts attempt to quantify research impact but can miss crucial aspects like conceptual innovation, methodological rigor, societal impact, the authors’ capacity for mentorship, and influence on society outside traditional citation patterns [77, 20]. In sustainability: a narrow focus on carbon emissions metrics (though important) can miss deeper notions of sustainability based on cultural and spiritual connections to land, traditional ecological knowledge, and community-specific environmental practices [47]. In employee hiring and evaluation: quantitative metrics like sales numbers, tasks completed, or hours worked fail to capture crucial qualities like mentorship, team morale building, institutional knowledge sharing, and maintaining positive client relationships [76].

In each of these cases, there is a more abstract notion of “what really matters” that gets consolidated into a thin signal which is legible at scale. We contend that employing thick description can help build systems that can represent something closer to the meaningful target notion.

4 Scaling Thick Description: Five Challenges for Measuring Meaning

In the previous section, we argued that robust strategies for representing meaning already exist. They are used by scholars in the humanities and qualitative social sciences [27, 87, 82]. Accordingly, we contend that if it were possible to employ a sufficiently large number of literary critics, cultural anthropologists, historians, or other relevant scholars within large-scale sociotechnical systems, then these systems would have robust capabilities for representing and processing information related to cultural context: in other words, to thickly describe human meaning at scale.

Our proposal is to use LLMs to fulfill at least part of the functionality of human experts in these domains—to act as cultural critics at scale. Working in conjunction with human experts, LLMs can help solve the problem of scaling thick description. By prompting LLMs to produce thick descriptions, this approach would seek to reproduce a form of analysis that has typically required deep, situated human judgment. Crucially, this strategy operates within existing computational architectures, without requiring a fundamental redesign of current AI technologies. However, the feasibility of this approach depends on addressing two key challenges: first, the conditions under which LLMs are empirically capable of producing thick descriptions that approximate human-level interpretations; and second, whether such descriptions can be systematically translated into forms that are actionable within computational systems without erasing their cultural and interpretive complexity. We articulate how progress could be made on this front by describing five key challenges for measuring meaning, and for each one offer initial ideas for paradigms that would address them.

4.1 Meaning Only Exists in Context

The thrust of our proposal is not that qualitative representations are better than quantitative ones. It is that in order to represent meaning, we need to represent context. There are many ways to do this, with qualitative thick description being a canonical one. But for any given implementation strategy, the standard is not whether it is qualitative or quantitative—but how well it retains information about context in which the meaningful experience or action originally occurred.

This context-dependence of meaning holds across a range of domains and definitions of “meaning.” For example, words are meaningful, but only when used in the context of a given language; a swastika is a meaningful symbol, but its meaning changes when found in Western Europe versus a Buddhist country; and marriage ceremonies can be meaningful, but what specific aspects of the ceremony will be perceived as meaningful depends on the culture in which it is taking place—not to mention the individuals getting married. The difficulty is heterogeneity. What kind of context needs to be represented in order to capture meaning is underconstrained. Representing context is not simply a matter of listing proxy variables, like country, language, or ethnicity, nor can it be encapsulated in a feature set of norms, preferences, or taboos [113, 44, 85]. This is why standardization is antithetical to the representation of human meaning.

One means of translating between thinner and thicker representations is via novel methods for qualitative coding, in which subjective interpretations are condensed into a numerical value by a human analyst [87]. Indeed, the question of translating between thinner and thicker descriptions can be reframed in terms of such coding: How do we create domain-general methods for coding qualitative contextual factors into numerical values? Traditionally, qualitative coding schemes are bespoke for a given dataset—a human researcher must develop, test, and train other researchers to use them. However, with LLMs it may be possible both to develop a unique coding scheme for relevant aspects of a given context and then to employ the newly created scheme [49, 88]. This could help avoid the mandate for standardization by allowing for an ad hoc process of evaluation, adaptable to an arbitrarily large number of contexts, which nonetheless yields numerical values.

4.2 There Is No Single Source Of Truth for Meaning

The point of codifying meaning is not that there are “right” and “wrong” answers; it is that there are no *easy* answers and that any oversimplification necessarily omits crucial information about human meaning. The same symbol, artifact, or experience can carry different meanings depending on who is interpreting it and in what context. Systems for representing meaning should maintain this plurality rather than collapsing diverse interpretations into singular representations.

For example, consider the case of a student for whom we have comprehensive educational data, including standardized test scores, demographic information, personal essays, psychological evaluations by third-party experts (think: an augmented college admissions application). One evaluator might interpret an application as providing evidence of technical focus and emerging leadership in STEM fields. Another might view it as the story of a creative, emotionally attuned individual with potential in design or writing. A third might focus on signs of perfectionism or risk-aversion. Each interpretation is grounded in the same underlying data, yet each foregrounds different elements of meaning. This highlights a core difficulty: even when the input is constant, thick descriptions can diverge depending on the interpretive lens. The goal, then, is not to eliminate this divergence but to build systems that can recognize and represent such pluralism. How do we apply interpretive judgment at scale in ways that are consistent, fair, and reflexive?

Recent work shows how LLMs can be used to reckon with this kind of ambiguity, specifically to identify the presence of dog whistles [61, 53]. These are a form of coded communication that uses covert social or ideological signals to convey a generally acceptable meaning to a wider audience while conveying a more controversial or offensive meaning to a targeted audience. This leads to conflicting, audience-dependent layers of meaning. This work shows that LLMs are capable of word-sense disambiguation by analyzing contextual usage, distinguishing between benign and coded meanings, giving a preliminary instance of how LLMs can be used to support interpretive pluralism.

4.3 Both Lived Experience and Critical Distance Are Crucial

Expert judgments from trained humanists and social scientists are useful, but they must be grounded in the lived experience of the individuals they are trying to understand. Geertz (1983) distinguishes between the concepts of *experience-near* (informal usages that reflect a subjective perspective) and *experience-distant* (formal usages that reflect an analytical or scientific perspective) [28]. For example, a gambler and an AI researcher might describe the same situation in different terms (e.g., experience-near: slot machine, hitting it big; experience-distant: n-armed bandit, sparse rewards). Both specialized scholarly knowledge (experience-distant) and lived cultural experience (experience-near) are necessary for robust representations of meaning. Neither should be privileged at the expense of the other, and systems should facilitate integration across these different forms of expertise.

A range of recent work could support this effort. In natural language processing, researchers have begun to explore perspectivist approaches: data sets that are labeled to represent the judgments of annotators with differing, sometimes conflicting, points of view—rather than labeling data according to a single ground truth [25, 2, 112, 93]. Similarly, work in pluralistic alignment has sought to create AI systems that reflect diverse values and perspectives [92, 91, 56]. Efforts like these could specifically target a balance of experience-near and experience-distant perspectives.

4.4 What \neq How Much

A key problem in mapping between cultural and experiential meaning is the difference between content and magnitude. So far we have discussed qualitative methods for assessing meaning [27, 87, 82]. However, psychological studies of meaning typically use quantitative assessments: surveys in which respondents are asked to rate on a numerical scale how meaningful they perceive a given experience, activity, or event to be [94, 36, 38, 37]. These result in fundamentally different questions about meaning. Qualitative approaches ask a question about content: “what is it?” By contrast, quantitative approaches ask a question about magnitude: “how much is there?” Effective systems for representing meaning should maintain this distinction rather than reducing one aspect to the other.

At first glance, creating a mapping between these two facets of meaning may seem like a straightforward empirical problem; however, this is not the case [15, 35]. For instance, it might seem that the solution would simply be to gather enough human-generated data in which people rate experiences as more or less meaningful. With enough features coded into the data, it could be determined which experiences people find meaningful—and those could be supported through technological interventions. However, this kind of approach would have difficulty accounting for the ambiguity in asking participants about meaning versus happiness [69], the radical flexibility of the interpretive process underlying judgments of meaning [52], the difficulties of codifying context as a feature set [113, 44, 85], that the same events can be perceived as more or less meaning depending on their narrative framing [83], the discrepancies between self-report and revealed preference [63], among

many other difficulties in combining cultural and psychological perspectives on meaning [11, 98]. It therefore remains an open question how best to develop mappings from qualitative interpretations about the cultural content of meaning to quantitative estimates of its psychological magnitude.

4.5 Meaning Is Made, Not Found

Humanists and social scientists often talk about “meaning-making,” a term which better captures the dynamic, dialogic process underlying judgments of meaning [26, 3]. Culture and meaning are not static; they evolve continuously through social negotiation and changing contexts [7]. Meaning-making may be usefully framed as a search process [24], but it is one through a space that is undergoing constant reorganization. What is meaningful can quickly become meaningless (or vice versa) as interpretive frames shift. Systems for representing meaning must be capable of adapting to these constant changes rather than treating cultural meaning as fixed.

Recent psychological evidence shows that people’s judgments about meaning can change based on the narrative framing of their experiences [83]: the same events elicit different ratings of meaningfulness when they occur in the context of different stories. This suggests that not only will LLMs be useful for measuring meaning—but they will become participants in the meaning-making process. For example, they may influence people’s perceptions about the meaningfulness of their own experiences through guided journaling [68, 46], personalized personas [18], or psychotherapy [67, 45]. Furthermore, how LLMs categorize people as having a shared contextual frame may have crucial downstream effects on their interactions with others [16]. Overall, this is consistent with what humanists and social scientists call the “reflexivity” of meaning: that the way we measure, codify, or understand meaning has a direct impact on the meaning-making process itself [31, 110, 8].

5 Discussion

In this paper, we have argued that LLMs can be used to make human meaning legible at an unprecedented scale. We contend that the primary obstacle to this effort is that human meaning can only be encoded via thick description—but, historically, the only way to reach legibility at scale is via thin description. We propose that LLMs can be used to make thick descriptions a useful representational strategy in large-scale sociotechnical systems. A crux of this is about adequately representing cultural context. Accordingly, there are useful gradations of thickness between descriptions that are fully thick and fully thin—and therefore many ways to usefully make progress on this problem. These gradations exist through relaxation of the mandate for standardization, mainly by finding ways to accommodate the inherent heterogeneity of socially-situated behavior.

5.1 Better Metrics Are Not Enough

Our proposal shares a premise with other calls to use AI in reshaping our society’s sociotechnical infrastructure (such as recommender systems) to reflect a broader sense of human well-being [97, 95, 42]. However, these efforts are primarily concerned with identifying which proxy metrics will most effectively capture human well-being. Our position is that all metrics will inevitably succumb to proxy failure [43] and what is needed instead is an entirely new representational format (i.e., thick description). To illustrate, this is why scholars who are concerned with meaning in the humanities and social sciences write papers full of long verbal exposition rather than ledgers of numbers [27, 28, 29]. Our proposal is similarly aligned to the nascent field of representation engineering [115], which places the representation of information (rather than its processing) at the core of technical development.

5.2 Without Thick Description, We Cannot Represent Key Aspects of Culture

Several recent papers make compelling cases about the difficulties of representing culture in (or using) generative AI models [113, 50, 1, 44, 86, 54, 102, 21]. While most of these do not focus primarily on “meaning” in the same way that we do, they reach broadly similar conclusions about the difficulties of representing cultural context in a nuanced way. Our proposal—to use thick description as a framework for making human meaning legible at scale—offers a potential unifying framework for these various perspectives. Thick description is a representational strategy designed to accommodate heterogeneity. A commonality among these accounts is that generative AI models depend on [21, 50], fail to account for [1, 44, 112, 86], or are not sufficiently representative of [54, 102] this kind of

heterogeneous cultural information. We argue that thick description gives us a framework for making progress on this wide variety of issues in representing the nuance and diversity of human culture.

5.3 This Is A Crucial Application Of Approaches To Scaling Qualitative Methods

Social scientists have begun exploring a range of efforts to scale qualitative methods (including thick description) using LLMs and other forms of generative AI [79, 23, 88, 49, 65, 107, 81]. Some offer thick description (or other social science approaches) as a method for evaluating AI systems [79, 107, 88]. A notable case is Qadri and colleagues (2025), who argue that “thick evaluations” are necessary to measure whether AI can represent key aspects of minority cultures [79]. While our position is aligned with their work in many ways, we argue that thickness is needed to represent all cultures (not just minorities) and that it should be used as a representational scheme itself and not just for evaluation. Other efforts have used LLMs to conduct qualitative analysis on large corpora of culturally-situated text [23, 81, 49, 65]; building on this work will be foundational toward ensuring the reliability of deploying thick description at scale. Especially promising work in this vein uses qualitative methods in not just in data analysis but in the generation of data as well [62].

5.4 Alternative Views, Risks, and Points of Failure

The following views are not mutually exclusive with the position offered in this paper. Rather, they present crucial risks and points of failure in pursuing this line of research.

This is going to lead to bad outcomes because you will fail at robustly representing meaning; like social media, you will start off trying to measure “what really matters” but ultimately end up with something much more narrow insidiously masquerading as the deeper thing. As implied in our section presenting five challenges, a lot can go wrong when trying to represent meaning. It is easy to overextend a given generalization beyond where it would be useful, accurate, or beneficial. It is difficult to convey the essence of culture in a way that does not trivialize it [113, 44]. One way our proposal could go wrong is by claiming to liberate a sociotechnical system from “enforced standardization”—only to make such standardization harder to discern. The risk is overgeneralization: applying one cultural standard for meaning in a context (or to an individual) where it does not belong.

This is going to lead to bad outcomes because you will succeed at robustly representing meaning; this will create a powerful lever not only for positive actors, but bad ones as well. For example, if a political state has the power to use qualitative analysis of personal data for noble purpose, then it would also have the power to use it for nefarious purposes as well [48]. In 1984, George Orwell imagined a future where “thoughtcrime” was enforced by legions of spies and informants. But the same thing could be accomplished if a large amount of personal data were available online and the state could employ the kind of automated qualitative analysis we describe.

You claim to draw on humanities and social science frameworks. But thick description is only one such framework; others paint a much more complicated picture. For example, Haraway (1988) argues that all knowledge is intrinsically partial, embodied, and dependent on specific historical and social contexts [33]. Thus, when the scientific community assumes that the perspective of its work is neutral and objective, it is simply imposing its own partial, contextually-bound perspective—usually a WEIRD one [39]—on an authoritative body of knowledge without having to acknowledge the imposition. According to Haraway, this paper’s proposal could potentially provide a more powerful mechanism for imposing a singular way of seeing things on a diverse cultural ecosystem. Likewise, Barad (2007) argues that the act of making something legible within a sociotechnical system does not just observe it in a neutral sense, but in fact transforms it [4]. For example, when a recommender system measures user preferences, it is not just recording pre-existing preferences; it is creating the categories that make certain preferences possible, likely, or even inexpressible. According to Barad, this paper’s proposal could potentially exert a corrupting influence on the original cultural phenomenon itself, pressuring it to fit into the categorical box that renders it legible within the sociotechnical system.

This isn’t a technical problem; it’s a political one. Even if you could represent meaning in the way you’re suggesting, there’s no guarantee that powerful sociotechnical systems would employ such representations in a beneficial way. In fact, they probably wouldn’t. This is a significant concern. However, in order to achieve the political solutions we want, we first need to have the technological solutions to support them. Our position is offered in service of this effort.

References

- [1] Muhammad Adilazuarda, Sagnik Mukherjee, Pradhyumna Lavania, Siddhant Singh, Alham Aji, Jacki O'Neill, Ashutosh Modi, and Monojit Choudhury. Towards measuring and modeling "culture" in llms: A survey. In *Proceedings of the 2024 Conference on Empirical Methods in Natural Language Processing*, pages 15763–15784, 2024.
- [2] Paula Akemi Aoyagui, Kelsey Stemmler, Sharon A Ferguson, Young-Ho Kim, and Anastasia Kuzminykh. A matter of perspective (s): Contrasting human and llm argumentation in subjective decision-making on subtle sexism. In *Proceedings of the 2025 CHI Conference on Human Factors in Computing Systems*, pages 1–16, 2025.
- [3] MM Bakhtin. *The dialogic imagination: four essays*. University of Texas Press, 1981.
- [4] Karen Barad. *Meeting the universe halfway: Quantum physics and the entanglement of matter and meaning*. duke university Press, 2007.
- [5] Gordon Baxter and Ian Sommerville. Socio-technical systems: From design methods to systems engineering. *Interacting with computers*, 23(1):4–17, 2011.
- [6] Emily M Bender, Timnit Gebru, Angelina McMillan-Major, and Shmargaret Shmitchell. On the dangers of stochastic parrots: Can language models be too big?. In *Proceedings of the 2021 ACM conference on fairness, accountability, and transparency*, pages 610–623, 2021.
- [7] Peter Berger and Thomas Luckmann. *The social construction of reality: a treatise in the sociology of knowledge*. 1966.
- [8] Geoffrey C Bowker and Susan Leigh Star. *Sorting things out: Classification and its consequences*. MIT press, 2000.
- [9] danah boyd and Kate Crawford. Critical questions for big data: Provocations for a cultural, technological, and scholarly phenomenon. *Information, communication & society*, 15(5):662–679, 2012.
- [10] danah boyd and Nicole B Ellison. Social network sites: Definition, history, and scholarship. *Journal of computer-mediated Communication*, 13(1):210–230, 2007.
- [11] Jerome Bruner. *Acts of meaning: Four lectures on mind and culture*, volume 3. Harvard university press, 1990.
- [12] Jerome Bruner. The narrative construction of reality. *Critical inquiry*, 18(1):1–21, 1991.
- [13] Jerome Bruner. Culture and mind: Their fruitful incommensurability. *Ethos*, 36(1):29–45, 2008.
- [14] Vivien Burr. *An Introduction to Social Constructionism*. Routledge, 3 edition, 2015.
- [15] Jennifer D Campbell, Paul D Trapnell, Steven J Heine, Ilana M Katz, Loraine F Lavallee, and Darrin R Lehman. Self-concept clarity: Measurement, personality correlates, and cultural boundaries. *Journal of personality and social psychology*, 70(1):141, 1996.
- [16] Elijah L Claggett, Robert E Kraut, and Hirokazu Shirado. Relational ai: Facilitating intergroup cooperation with socially aware conversational support. In *Proceedings of the 2025 CHI Conference on Human Factors in Computing Systems*, pages 1–22, 2025.
- [17] Robert Darnton. *The great cat massacre: and other episodes in French cultural history*. Basic Books, 1984.
- [18] Stefano De Paoli. Writing user personas with large language models: Testing phase 6 of a thematic analysis of semi-structured interviews. *arXiv preprint arXiv:2305.18099*, 2023.
- [19] Catherine D'ignazio and Lauren F Klein. *Data feminism*. MIT press, 2023.
- [20] Wendy Nelson Espeland and Michael Sauder. Rankings and reactivity: How public rankings transform organizational behavior. *American Journal of Sociology*, 113(1):1–40, 2007.
- [21] Henry Farrell, Alison Gopnik, Cosma Shalizi, and James Evans. Large ai models are cultural and social technologies. *Science*, 387(6739):1153–1156, 2025.
- [22] Max Fisher. *The chaos machine: The inside story of how social media rewired our minds and our world*. Little, Brown, 2022.

- 492 [23] Samuel Flanders, Melati Nungsari, and Mark Cheong Wing Loong. Big meaning: Qualitative analysis on
493 large bodies of data using ai. *arXiv preprint arXiv:2504.08213*, 2025.
- 494 [24] Viktor E. Frankl. *Man's Search for Meaning*. Beacon Press, 1946.
- 495 [25] Simona Frenda, Gavin Abercrombie, Valerio Basile, Alessandro Pedrani, Raffaella Panizzon, Alessan-
496 dra Teresa Cignarella, Cristina Marco, and Davide Bernardi. Perspectivist approaches to natural language
497 processing: a survey. *Language Resources and Evaluation*, pages 1–28, 2024.
- 498 [26] Hans-Georg Gadamer. *Truth and Method*. A&C Black, 1960.
- 499 [27] Clifford Geertz. *The Interpretation of Cultures*. Basic books, 1973.
- 500 [28] Clifford Geertz. *Local Knowledge: Further essays in interpretive anthropology*. Basic Books, 1983.
- 501 [29] Clifford Geertz. *Works and lives: The anthropologist as author*. Stanford University Press, 1988.
- 502 [30] Carolin Gerlitz and Anne Helmond. The like economy: Social buttons and the data-intensive web. *New
503 media & society*, 15(8):1348–1365, 2013.
- 504 [31] Anthony Giddens. *Modernity and Self-Identity: Self and Society in the Late Modern Age*. Stanford
505 University Press, 1991.
- 506 [32] Stephen Greenblatt. *Renaissance self-fashioning: from More to Shakespeare*. University of Chicago
507 Press, 1980.
- 508 [33] Donna Haraway. Situated knowledges: The science question in feminism and the privilege of partial
509 perspective. *Feminist Studies*, 14(3):575–599, 1988.
- 510 [34] Tristan Harris. Distracted? let's make technology that helps us spend our time well. TEDxBrussels, 2014.
511 Talk presented on December 16, 2014.
- 512 [35] Steven J Heine, Darrin R Lehman, Kaiping Peng, and Joe Greenholtz. What's wrong with cross-cultural
513 comparisons of subjective likert scales?: The reference-group effect. *Journal of personality and social
514 psychology*, 82(6):903, 2002.
- 515 [36] Samantha J Heintzelman and Laura A King. Life is pretty meaningful. *American psychologist*, 69(6):561,
516 2014.
- 517 [37] Samantha J Heintzelman and Laura A King. (the feeling of) meaning-as-information. *Personality and
518 Social Psychology Review*, 18(2):153–167, 2014.
- 519 [38] Samantha J Heintzelman, Jason Trent, and Laura A King. Encounters with objective coherence and the
520 experience of meaning in life. *Psychological science*, 24(6):991–998, 2013.
- 521 [39] Joseph Henrich, Steven J Heine, and Ara Norenzayan. The weirdest people in the world? *Behavioral and
522 brain sciences*, 33(2-3):61–83, 2010.
- 523 [40] Yubo Hou, Dan Xiong, Tonglin Jiang, Lily Song, and Qi Wang. Social media addiction: Its impact,
524 mediation, and intervention. *Cyberpsychology: Journal of psychosocial research on cyberspace*, 13(1),
525 2019.
- 526 [41] Thomas P Hughes et al. The evolution of large technological systems. *The social construction of
527 technological systems: New directions in the sociology and history of technology*, 82:51–82, 1987.
- 528 [42] Chenyan Jia, Michelle S Lam, Minh Chau Mai, Jeffrey T Hancock, and Michael S Bernstein. Embedding
529 democratic values into social media ais via societal objective functions. *Proceedings of the ACM on
530 Human-Computer Interaction*, 8(CSCW1):1–36, 2024.
- 531 [43] Yohan J John, Leigh Caldwell, Dakota E McCoy, and Oliver Braganza. Dead rats, dopamine, performance
532 metrics, and peacock tails: proxy failure is an inherent risk in goal-oriented systems. *Behavioral and
533 Brain Sciences*, 47:e67, 2024.
- 534 [44] Ariba Khan, Stephen Casper, and Dylan Hadfield-Menell. Randomness, not representation: The unrelia-
535 bility of evaluating cultural alignment in llms. *arXiv preprint arXiv:2503.08688*, 2025.
- 536 [45] Taewan Kim, Seolyeong Bae, Hyun Ah Kim, Su-woo Lee, Hwajung Hong, Chanmo Yang, and Young-Ho
537 Kim. Mindfuldiary: Harnessing large language model to support psychiatric patients' journaling. In
538 *Proceedings of the 2024 CHI Conference on Human Factors in Computing Systems*, pages 1–20, 2024.

- [46] Taewan Kim, Donghoon Shin, Young-Ho Kim, and Hwajung Hong. Diarmate: Understanding user perceptions and experience in human-ai collaboration for personal journaling. In *Proceedings of the 2024 CHI Conference on Human Factors in Computing Systems*, pages 1–15, 2024.
- [47] Robin Wall Kimmerer. *Braiding sweetgrass: Indigenous wisdom, scientific knowledge and the teachings of plants*. Milkweed editions, 2013.
- [48] Hannah Rose Kirk, Bertie Vidgen, Paul Röttger, and Scott A Hale. The benefits, risks and bounds of personalizing the alignment of large language models to individuals. *Nature Machine Intelligence*, 6(4):383–392, 2024.
- [49] Elisabeth Kirsten, Annalina Buckmann, Abraham Mhaidli, and Steffen Becker. Decoding complexity: Exploring human-ai concordance in qualitative coding. *arXiv preprint arXiv:2403.06607*, 2024.
- [50] Lauren Klein, Meredith Martin, André Brock, Maria Antoniak, Melanie Walsh, Jessica Marie Johnson, Lauren Tilton, and David Mimno. Provocations from the humanities for generative ai research. *arXiv preprint arXiv:2502.19190*, 2025.
- [51] Alfie Kohn. *The Case Against Standardized Testing: Raising the Scores, Ruining the Schools*. Heinemann, Portsmouth, NH, 2000.
- [52] Cody Kommers and Simon DeDeo. Sense-making, cultural scripts, and the inferential basis of meaningful experience. 2025.
- [53] Julia Kruk, Michela Marchini, Rijul Magu, Caleb Ziems, David Muchlinski, and Diyi Yang. Silent signals, loud impact: Lfms for word-sense disambiguation of coded dog whistles. In *Proceedings of the 62nd Annual Meeting of the Association for Computational Linguistics (Volume 1: Long Papers)*, pages 12493–12509, 2024.
- [54] Maik Larooij and Petter Törnberg. Do large language models solve the problems of agent-based modeling? a critical review of generative social simulations. *arXiv preprint arXiv:2504.03274*, 2025.
- [55] David MJ Lazer, Matthew A Baum, Yochai Benkler, Adam J Berinsky, Kelly M Greenhill, Filippo Menczer, Miriam J Metzger, Brendan Nyhan, Gordon Pennycook, David Rothschild, et al. The science of fake news. *Science*, 359(6380):1094–1096, 2018.
- [56] Seongyun Lee, Sue Hyun Park, Seungone Kim, and Minjoon Seo. Aligning to thousands of preferences via system message generalization. *Advances in Neural Information Processing Systems*, 37:73783–73829, 2024.
- [57] Bronislaw Malinowski. *Argonauts of the western pacific*. 1922.
- [58] Catherine Marshall and Gretchen B. Rossman. *Designing Qualitative Research*. Sage Publications, 2014.
- [59] Viktor Mayer-Schönberger and Kenneth Cukier. *Big data: A revolution that will transform how we live, work, and think*. Houghton Mifflin Harcourt, 2013.
- [60] Dan P McAdams. The psychology of life stories. *Review of general psychology*, 5(2):100–122, 2001.
- [61] Julia Mendelsohn, Ronan Le Bras, Yejin Choi, and Maarten Sap. From dogwhistles to bullhorns: Unveiling coded rhetoric with language models. In *The 61st Annual Meeting Of The Association For Computational Linguistics*, 2023.
- [62] Han Meng, Renwen Zhang, Ganyi Wang, Yitian Yang, Peinuan Qin, Jungup Lee, and Yi-Chieh Lee. Deconstructing depression stigma: Integrating ai-driven data collection and analysis with causal knowledge graphs. In *Proceedings of the 2025 CHI Conference on Human Factors in Computing Systems*, pages 1–21, 2025.
- [63] Smitha Milli, Micah Carroll, Yike Wang, Sashrika Pandey, Sebastian Zhao, and Anca D Dragan. Engagement, user satisfaction, and the amplification of divisive content on social media. *PNAS nexus*, 4(3):pgaf062, 2025.
- [64] Haradhan Kumar Mohajan. Qualitative research methodology in social sciences and related subjects. *Journal of economic development, environment and people*, 7(1):23–48, 2018.
- [65] Anders Giovanni Møller and Luca Maria Aiello. Prompt refinement or fine-tuning? best practices for using lfms in computational social science tasks. *arXiv preprint arXiv:2408.01346*, 2024.
- [66] Jerry Muller. *The tyranny of metrics*. Princeton University Press, 2018.

- [67] Hongbin Na, Yining Hua, Zimu Wang, Tao Shen, Beibei Yu, Lilin Wang, Wei Wang, John Torous, and Ling Chen. A survey of large language models in psychotherapy: Current landscape and future directions. *arXiv preprint arXiv:2502.11095*, 2025.
- [68] Subigya Nepal, Arvind Pillai, William Campbell, Talie Massachi, Michael V Heinz, Ashmita Kunwar, Eunsol Soul Choi, Xuhai Xu, Joanna Kuc, Jeremy F Huckins, et al. Mindscape study: Integrating llm and behavioral sensing for personalized ai-driven journaling experiences. *Proceedings of the ACM on interactive, mobile, wearable and ubiquitous technologies*, 8(4):1–44, 2024.
- [69] Shigehiro Oishi, Hyewon Choi, Minkyung Koo, Iolanda Galinha, Keiko Ishii, Asuka Komiya, Maike Luhmann, Christie Scollon, Ji-eun Shin, Hwaryung Lee, et al. Happiness, meaning, and psychological richness. *Affective Science*, 1:107–115, 2020.
- [70] Michael Oliver and Colin Barnes. *The new politics of disablement*. Bloomsbury Publishing, 2012.
- [71] Cathy O’Neil. *Weapons of math destruction: How big data increases inequality and threatens democracy*. Crown, 2017.
- [72] Will Oremus. Facebook’s new philosophy: Will emphasizing “time well spent” fix the company’s russia problem—or just make it more money? *Slate*, nov 2017.
- [73] Kirsten Ostherr. Artificial intelligence and medical humanities. *Journal of Medical Humanities*, 43(2):211–232, 2022.
- [74] Kirsten Ostherr. The future of translational medical humanities: bridging the data/narrative divide. *Medical Humanities*, 49(4):529–536, 2023.
- [75] Amy I Parkes, Adam J Sobey, and Dominic A Hudson. More consistent learnt relationships in regression neural networks using fit to median error measure. *Engineering Applications of Artificial Intelligence*, 144:110113, 2025.
- [76] Jeffrey Pfeffer and Robert I Sutton. *Hard facts, dangerous half-truths, and total nonsense: Profiting from evidence-based management*. Harvard Business Press, 2006.
- [77] Jason Priem, Dario Taraborelli, Paul Groth, and Cameron Neylon. Altmetrics: A manifesto. <http://altmetrics.org/manifesto/>, 2012.
- [78] Brian A Primack, Ariel Shensa, Jaime E Sidani, Erin O Whaite, Liu yi Lin, Daniel Rosen, Jason B Colditz, Ana Radovic, and Elizabeth Miller. Social media use and perceived social isolation among young adults in the us. *American journal of preventive medicine*, 53(1):1–8, 2017.
- [79] Rida Qadri, Mark Diaz, Ding Wang, and Michael Madaio. The case for" thick evaluations" of cultural representation in ai. *arXiv preprint arXiv:2503.19075*, 2025.
- [80] Diane Ravitch. *The Death and Life of the Great American School System: How Testing and Choice Are Undermining Education*. Basic Books, New York, 2010.
- [81] Kevin Reuter, Lucien Baumgartner, and Pascale Willemsen. Tracing thick and thin concepts through corpora. *Language and Cognition*, 16(2):263–282, 2024.
- [82] Catherine Kohler Riessman. *Narrative Methods for the Human Sciences*. Sage Publications, 2008.
- [83] Benjamin A Rogers, Herrison Chicas, John Michael Kelly, Emily Kubin, Michael S Christian, Frank J Kachanoff, Jonah Berger, Curtis Puryear, Dan P McAdams, and Kurt Gray. Seeing your life story as a hero’s journey increases meaning in life. *Journal of Personality and Social Psychology*, 125(4):752, 2023.
- [84] Bertrand Russell. In praise of idleness: And other essays. 1935.
- [85] Jonathan Rystrom, Hannah Rose Kirk, and Scott Hale. Multilingual!= multicultural: Evaluating gaps between multilingual capabilities and cultural alignment in llms. *arXiv preprint arXiv:2502.16534*, 2025.
- [86] Sougata Saha, Saurabh Kumar Pandey, and Monojit Choudhury. Meta-cultural competence: Climbing the right hill of cultural awareness. *arXiv preprint arXiv:2502.09637*, 2025.
- [87] Maggi Savin-Baden and Claire Major. *Qualitative research: The essential guide to theory and practice*. Routledge, 2023.
- [88] Hope Schroeder, Marianne Aubin Le Quéré, Casey Randazzo, David Mimno, and Sarita Schoenebeck. Large language models in qualitative research: Can we do the data justice? *arXiv e-prints*, pages arXiv–2410, 2024.

- 637 [89] James C Scott. Seeing like a state. 1998.
- 638 [90] Amartya Sen. *Development as Freedom*. Alfred A. Knopf, New York, 1999.
- 639 [91] Taylor Sorensen, Liwei Jiang, Jena D Hwang, Sydney Levine, Valentina Pyatkin, Peter West, Nouha Dziri,
640 Ximing Lu, Kavel Rao, Chandra Bhagavatula, et al. Value kaleidoscope: Engaging ai with pluralistic
641 human values, rights, and duties. In *Proceedings of the AAAI Conference on Artificial Intelligence*,
642 volume 38, pages 19937–19947, 2024.
- 643 [92] Taylor Sorensen, Jared Moore, Jillian Fisher, Mitchell Gordon, Niloofar Mireshghallah, Christo-
644 pher Michael Rytting, Andre Ye, Liwei Jiang, Ximing Lu, Nouha Dziri, et al. A roadmap to pluralistic
645 alignment. *arXiv preprint arXiv:2402.05070*, 2024.
- 646 [93] Taylor Sorensen, Joshua Robinson, Christopher Michael Rytting, Alexander Glenn Shaw, Kyle Jeffrey
647 Rogers, Alexia Pauline Delorey, Mahmoud Khalil, Nancy Fulda, and David Wingate. An information-
648 theoretic approach to prompt engineering without ground truth labels. *arXiv preprint arXiv:2203.11364*,
649 2022.
- 650 [94] Michael F Steger, Patricia Frazier, Shigehiro Oishi, and Matthew Kaler. The meaning in life questionnaire:
651 assessing the presence of and search for meaning in life. *Journal of counseling psychology*, 53(1):80,
652 2006.
- 653 [95] Jonathan Stray. Aligning ai optimization to community well-being. *International Journal of Community*
654 *Well-Being*, 3(4):443–463, 2020.
- 655 [96] Jonathan Stray. Aligning ai to human values means picking the right metrics. *Partnershiponai.org*, 2020.
- 656 [97] Jonathan Stray, Alon Halevy, Parisa Assar, Dylan Hadfield-Menell, Craig Boutilier, Amar Ashar, Chloe
657 Bakalar, Lex Beattie, Michael Ekstrand, Claire Leibowicz, et al. Building human values into recommender
658 systems: An interdisciplinary synthesis. *ACM Transactions on Recommender Systems*, 2(3):1–57, 2024.
- 659 [98] Charles Taylor. Sources of the self: The making of the modern identity. *Harvard UP*, 1989.
- 660 [99] Hayley Thomas, Geoffrey Mitchell, Justin Rich, and Megan Best. Definition of whole person care in
661 general practice in the english language literature: a systematic review. *BMJ open*, 8(12):e023758, 2018.
- 662 [100] Eric Lansdown Trist and Kenneth W Bamforth. Some social and psychological consequences of the
663 longwall method of coal-getting: An examination of the psychological situation and defences of a work
664 group in relation to the social structure and technological content of the work system. *Human relations*,
665 4(1):3–38, 1951.
- 666 [101] Jean M Twenge. *iGen: Why today’s super-connected kids are growing up less rebellious, more tolerant,*
667 *less happy—and completely unprepared for adulthood—and what that means for the rest of us*. Simon and
668 Schuster, 2017.
- 669 [102] Ted Underwood, Laura K Nelson, and Matthew Wilkens. Can language models represent the past without
670 anachronism? *arXiv preprint arXiv:2505.00030*, 2025.
- 671 [103] José Van Dijck and Thomas Poell. Understanding social media logic. *Media and communication*,
672 1(1):2–14, 2013.
- 673 [104] Tyler J VanderWeele and Byron R Johnson. Multidimensional versus unidimensional approaches to
674 well-being. *Nature Human Behaviour*, pages 1–7, 2025.
- 675 [105] Tyler J VanderWeele, Byron R Johnson, Piotr T Bialowolski, Rebecca Bonhag, Matt Bradshaw, Thomas
676 Breedlove, Brendan Case, Ying Chen, Zhuo Job Chen, Victor Counted, et al. The global flourishing study:
677 Study profile and initial results on flourishing. *Nature Mental Health*, pages 1–18, 2025.
- 678 [106] Soroush Vosoughi, Deb Roy, and Sinan Aral. The spread of true and false news online. *Science*,
679 359(6380):1146–1151, 2018.
- 680 [107] Hanna Wallach, Meera Desai, Nicholas Pangakis, A Feder Cooper, Angelina Wang, Solon Barocas,
681 Alexandra Chouldechova, Chad Atalla, Su Lin Blodgett, Emily Corvi, et al. Evaluating generative ai
682 systems is a social science measurement challenge. *arXiv preprint arXiv:2411.10939*, 2024.
- 683 [108] Michael Walzer. *Thick and thin: Moral argument at home and abroad*. University of Notre Dame Press,
684 2019.
- 685 [109] Michael White and David Epston. *Narrative means to therapeutic ends*. WW Norton & Company, 1990.

- 686 [110] Ludwig Wittgenstein. *Philosophical Investigations*. Blackwell Publishing, 1953.
- 687 [111] World Health Organization. Constitution of the World Health Organization, 1946.
- 688 [112] Michael JQ Zhang, Zhilin Wang, Jena D Hwang, Yi Dong, Olivier Delalleau, Yejin Choi, Eunsol Choi,
 689 Xiang Ren, and Valentina Pyatkin. Diverging preferences: When do annotators disagree and do models
 690 know? *arXiv preprint arXiv:2410.14632*, 2024.
- 691 [113] Naitian Zhou, David Bamman, and Isaac L Bleaman. Culture is not trivia: Sociocultural theory for
 692 cultural nlp. *arXiv preprint arXiv:2502.12057*, 2025.
- 693 [114] Tania Zittoun and Alex Gillespie. Internalization: How culture becomes mind. *Culture & Psychology*,
 694 21(4):477–491, 2015.
- 695 [115] Andy Zou, Long Phan, Sarah Chen, James Campbell, Phillip Guo, Richard Ren, Alexander Pan, Xuwang
 696 Yin, Mantas Mazeika, Ann-Kathrin Dombrowski, et al. Representation engineering: A top-down approach
 697 to ai transparency. *CoRR*, 2023.