Supporting and Performing Culture from the Inside

Anonymous ACL submission

Abstract

Anthropology and related disciplines which study the concept of *culture* have found it useful to distinguish their epistemological and methodological approaches as either the insider or emic perspective (how is the world viewed by a member of the culture) or the outsider or etic perspective (the scientific cataloguing of cultural knowledge and practice). We adopt this lens to systematically analyse assumptions in recent research on culture in NLP along the full pipeline of task selection, data collection, system design and evaluation. We expose underlying assumptions and suggest a broader view and a path forward which would lead to more useful, safe and culturally appropriate development of culturally attuned AI systems.

1 Introduction

Studying human culture is complex because it is a complex and fluid phenomenon and because it can hardly be isolated from the culture that a researcher brings to an investigation. The same challenge applies to NLP research with a focus on culture, or cultural alignment, which has seen a surge of attention in recent years (Pawar et al., 2025).

To systematize epistemological and methodological approaches and make them explicit, anthropologists have proposed to distinguish emic and etic approaches to cultural studies (Pike, 1967; Sullivan, 2009). An *etic* perspective adopts the outsider's view, classifies surface forms, is typically scalable and applies generic measures across cultures for a comparative view. The *emic* perspective adopts the insider's view of a member of the target culture, with a focus on the underlying cultural system and it is often specific to a single culture. The concepts of *etic* and *emic* originate in structuralist linguistics, as cultural generalisations of "phonetic" and "phonemic". The term *phonetic* refers to observable surface forms which can be described using a universal classification system (here, the International Phonetic Alphabet). The term *phonemic* refers to a language-internal system (here, sound contrasts that are significant for meaning). Pike (1967) observed the parallel that human behaviours can be observed, described, and classified from outside a culture (etic), yet have a significance that is only accessible to insiders (Headland et al., 1990). A popular representation of the etic/emic distinction is the "cultural iceberg" (Fig. 1).

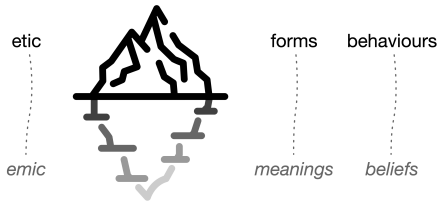

Figure 1: Hall's Cultural Iceberg: forms and behaviours can be observed by outsiders, while corresponding meanings and beliefs are only accessible to insiders (after Hall, 1976)

This paper connects the etic/emic framework with NLP research. It surveys the recent literature on cultural NLP and, perhaps unsurprisingly, identifies a dominance of etic approaches. It then applies the emic/etic framework and (1) uncovers hidden assumptions in how our community designs approaches, data sets and model evaluations; (2) points to gaps in research assuming that the ultimate goal is to develop systems which are valuable to the target culture; and (3) sketches out a path forward towards a more emic research agenda.

2 The Etic/Emic Framework

In this section we will see that the practice of generative AI is strongly etic, in all stages of development. Task definitions are understood as general purpose solutions to externally defined problems, which means that few researchers see any need to engage individual communities about what they need (Sec. 2.1). Primary data about the world is

taken to be text, yet this contains many layers of culture, and may effectively be a blindspot for NLP researchers (Sec. 2.2). Systems are designed to replicate or replace humans without due consideration for how this may impact the meaningfulness of lives and livelihoods on the ground (Sec. 2.3). Finally, deployments are optimised to serve global measures of efficiency, productivity, and profitability, and performance benchmarks are applied across the board, regardless of local cultural norms and expectations (Sec. 2.4).

In addition, we reconsider each of these development stages from an emic perspective. We incorporate findings from a systematic review of the literature which covers all papers published in the Findings or Main streams of major recent conferences (*ACL, EMNLP, 2022–4/2026) containing the word 'culture' or 'cultural' in the title (N=107). Please see Appendix A and the data file accompanying this submission for the full list of papers and labels.

2.1 Task

An emic approach begins locally and seeks local meaning. It values local authenticity and fidelity over ambit claims concerning the efficiency and scalability of universal "solutions". Emic approaches begin with local matters of concern. These include topics of economic participation because minoritised communities are often economically disadvantaged (Bird and Yibarbuk, 2024). A common linguistic challenge is to make sense of outside concepts which have no local translation, so called "untranslatable words" (Helm et al., 2024).

Local matters of concern go further than economic participation to include cultural survival. This is because leaders of local cultures often aspire to pass on their knowledge practices and lore (Abram, 1997; Kimmerer, 2013). A common linguistic challenge is to transmit traditional ecological knowledge – such as collecting and preparing food – with the associated knowledge of the country, its seasons, and various duties to kin.

The tasks tackled in recent NLP papers did not take these considerations into account. The majority of surveyed papers (N=78, 73%) focussed on assessing and/or augmenting cultural knowledge in LLMs, i.e., serve the intrinsic goal of 'trivia-centered' cultural knowledge representations (Zhou et al., 2025; Oh et al., 2025). The remaining papers tackle diverse downstream tasks ranging from culturally appropriate dialogue over culturally attuned translation over culture-aware IR, however tasks are conceived by the researchers who are embedded in a dominant culture (by upbringing and/or by academic affiliation), typically with generalizability to many cultures in mind (etic). We did not encounter any papers which co-developed a task with a cultural community.

2.2 Data

For most of the world's cultures, the textual record comes from ethnographic observations that are subject to the cultural conventions of outsiders (Clifford and Marcus, 2023). The record depends on the observable behaviours of insiders, and what insiders may report about the underlying meanings and beliefs, reports which may be unreliable or even deceptive (Borofsky, 2005). As a case in point, in his book *Orientalism*, Said argues that the construct of the "Orient" – and everything written about it – is a projection of western culture (Said, 1978).

This problem with texts is not removed by eliminating outsider descriptions, and only working with texts produced by insiders. These texts should not be considered a reliable or comprehensive representation of reality, as there are further layers of culture that intervene. First is the culture of writing itself:

> *... writing, wherever it exists, is always only one of several communication channels available to the members of a society. Consequently, the conditions under which it is selected and the purposes to which it is put must be described in relation to those of other channels. (Basso, 1974, p426)*

There is also the nature of observation: that there are no pure facts (Popper, 1935); that differentiating figure from ground is a perceptual act (Schacter et al., 2009); that measurement involves culturally-informed choices (Walter and Andersen, 2013); and that "the world as it is" differs from "the world according to data" (Mitchell and Fleischman, 2020, p145). Even the conversion of speech to text is implicated, being trained on speech transcriptions which are theory-laden (Ochs, 1979; Cucchiarini, 1993). Vector space representations of text have been branded as meaning representations, but this is a misnomer (Bender and Koller, 2020); word and document vectors instantiate a global, outsider, classificatory view, and are therefore etic. In short, texts are separated from "reality" by 3 cultural layers: what is observed, what meaning is attributed,

and how this is rendered into text, ie., a case of "stereotypes all the way down" (Zhou et al., 2025).

A notable fraction of the surveyed literature consulted with members of the community in the data collection process, rather than relying fully on existing texts as a ground truth: 19 papers drew on community input in concept selection, example creation and/or data validation where data was typically contributed by members of the target culture which were based at the respective research institution, or via crowd sourcing. However, we also observe a strong trend of creating 'cultural' benchmarks in an at least partially synthetic way (33 papers), where often concepts and/or context are first generated by LLMs and later filtered by humans in a practice that necessarily distorts the emic, intra-cultural perspective (see discussion above). The practice of translating English benchmarks was widely criticised and occurred only three times. 31 papers based their concept samples or examples on existing data sets which often spanned many languages (news, social media, Wikipedia, . . .) and imply the concerns laid out above.

2.3 System design

The construct of chatbots was founded on the model of a human amateur consulting a human expert (Weizenbaum, 1966; Woods, 1973). Studies of the culture of expertise show that it is not to be understood in terms of information possessed and imparted, but in terms of interaction and ideology. For each domain there is a culture whereby experts are socialised, evaluated, and authorised, and there are particular linguistic performances that produce jargon, structure interactions, and animate evidence (Carr, 2010). Amateur-expert consultations have a *shape* which is particular to the domain, e.g. in medicine with opening, complaint, examination/test, diagnosis, treatment/advice, and closing (Ten Have, 1989). Across many domains, better outcomes are linked to interpersonal rapport (Wilson et al., 2010).

Set against this, we have the project of making domain experts out of chatbots (Metzler et al., 2021), with the redirection of web queries into chatbot interactions, all in the name of retrieving answers rather than documents. However, to make this shift is to "miss the big picture of why people seek information and how that process contains value beyond simply retrieving relevant information" (Shah and Bender, 2022, p221). Moreover, it replaces diverse cultures of expert-shaped learning interactions with a series of prompt-response dyads shaped by the novice.

Most of our surveyed papers focussed on intrinsic evaluation and did not consider a "downstream system" (cf. Sec. 2.1). Those which were more task oriented did not question the design of a potential system, and simply adopted the predictable designs for IR, classification, or dialogue settings as established in the dominant culture and NLP research community.

2.4 Deployment / evaluation

Globally LLMs are predominantly used as chatbots. The chatbot construct was founded on the model of a human amateur consulting a human expert (Weizenbaum, 1966; Woods, 1973). Human-human information seeking culture is, however, grounded in interaction; there is a culture whereby experts are socialised, evaluated and authorised; interactions follow established linguistic performances that structure interactions (Carr, 2010; Ten Have, 1989). Chatbots are presented as domain experts (Metzler et al., 2021), generating (over) confident responses, often formulated to encourage continued engagement. This contrasts with purposeful socially and culturally grounded human-human learning, specifically when considering minority groups such as indigenous cultures where experts shape the process and novices listen unless asked to speak (Harris, 1984; Kimmerer, 2013).

Generative AI is increasingly deployed in agentic systems, designed with the goal of simulating and replacing humans (prominant recent cases including software engineers or academic reviewers). This reflects the culture of efficiency and automation underlying the dominant western view. An alternative take is to design AI agents to be *augmentative* (Bird, 2024), to enhance human expertise rather than replacing it. This approach bodes better with minority cultures' matters of concern (Section 2.3), and may well lead to more socially sustainable AI development more broadly as we discuss in Section 3.

Only a minority of papers in our survey focus on task-oriented cultural NLP systems and none of these papers explicitly questioned the design of a potential system. They adopted the predictable designs for IR, classification, or dialogue settings as established in the dominant culture and NLP research community.

	Task	Data	System Design	Deployment / Evaluation
Etic	generic solns to external problems • generate facts about any culture	data extraction for LLMs • translated or synthesized common sense benchmarks	simulate or improve on humans • text chat; adapting to user language; flexibility in tone/register	efficiency and scalability • correctness and completeness of responses
Emic	local agency in shaping their world, identifying problems, exploring solutions	elicit cultural concepts and logics	build / augment human capacity in sustaining cultural vitality	local decision making, maintaining culture and locally meaningful livelihoods
	• understanding concepts from western culture which have no translation (e.g. "risk", "safety")	• exegesis of western cultural concepts into local languages	• redesign of text interaction to follow local teacher-led pattern	• equipped with actionable information, local people make better-informed choices
	• transmitting traditional knowledge practices (e.g. food collection/preparation)	• videos of people demonstrating knowledge practices, sharing lore	• collection management and navigation, spoken term retrieval	• measure user-reported utility and safety

Table 1: Etic vs Emic Perspectives on Four Stages of AI Development

3 Paths Forward

How can we navigate into the emic? Here we speculate about new paths which adapt mainstream methods to a local culture, or which reproduce mainstream methods under local control through idiomatic interaction designs, or which take seriously diverse local understandings of intelligence. Through this section we see increasingly emic approaches.

Local adaptation. Rather than seeking to contribute cultural data to shape mainstream LLMs, this is a space of adapting language models to local cultures, e.g. in national deployments (Biju and Gayathri, 2025; Raj et al., 2025) or language-specific deployments (Mahelona et al., 2023). Harms that come from outsider views and from homogenisation towards dominant norms (cf. Peterson, 2025) are mitigated by restricting training data to in-country sources (even though such sources may still embed colonial stereotypes). This remains etic because sources are textual (even when they contain emic accounts), and they have the limitations described in Section 2.2. The benefit of training and deploying within a single cultural area is that the cultural assumptions of users are more likely to align with the cultural assumptions embedded in texts.

Local control. Here we ask about approaches to the design of AI which puts the local community's aspirations at the centre (e.g. Bird and Yibarbuk, 2024; Cooper et al., 2024; Markl et al., 2024). There is history in adjacent fields like morphology and speech processing, often applied to use cases like literacy and text entry (Moshagen et al., 2023; Pine et al., 2025).

An emic approach would let data collection be driven by the community: relevant concepts are defined by the community, examples are collected and validated, narratives are compiled (cf. Sultana et al., 2022), data does not leave the community (Mahelona et al., 2023). When it comes to archives of ancestral speech, there are no technological mashups of recordings (technology's role is intermediary not mediator; Latour 2007). The interaction design of a chatbot is reshaped along the lines of local knowledge transmission. Initiative is held by the teacher, not by a learner who has no automatic entitlement to knowledge. Knowledge transmission is relational, where both teacher and learner are human and where language technologies exist to support not replace them. We provide some suggestions in Table 1.

Local diversity. To navigate further into the emic involves cultivating openness to diverse local understandings of intelligence (Bridle, 2022; Lewis et al., 2024), and to approaches to language technology not predicated on extracting behaviour in order to simulate or replace humans (Crawford, 2021; Zuboff, 2022). We embrace cultural hybrids in place of "pure" or "essentialised" notions of culture (Lenard, 2020), embrace parochial technologies in place of one-size-fits-all solutions (Srinivasan, 2017), and embrace the agency of local, bicultural people to shape, leverage, and hybridise technology to their own ends. This will often lead to language processing tasks that respect local self-determination in maintaining culture, language, and life ways.

4 Conclusion

We have presented the etic/emic framework, applying it to notions of culture that arise in contemporary NLP research, drawing on a systemic review of the cultural NLP literature. We have described the theoretical limits to cultural alignment that are intrinsic to etic approaches, given the layers of culture that cannot be removed from the textual record accessible to LLMs. Finally, we have suggested three ways to develop more emic research agendas.

5 Limitations

This position paper was shaped by the authors' own cultural perspectives. While our literature survey is *representative* of research related to culture in NLP, we do not claim that it is *exhaustive* and we might have missed a small number of more emic approaches with our approach. Still, our main conclusions that etic approaches dominate NLP and that a more emic angle would benefit the utiliy and value of cultural NLP work still holds.

References

David Abram. 1997. *The Spell of the Sensuous: Perception and Language in a More-Than-Human World*. Vintage.

Keith H. Basso. 1974. The ethnography of writing. In Richard Bauman and Joel Sherzer, editors, *Explorations in the Ethnography of Speaking*, pages 425–32. Cambridge University Press.

Emily M. Bender and Alexander Koller. 2020. Climbing towards NLU: On meaning, form, and understanding in the Age of Data. In *Proceedings of the 58th Annual Meeting of the Association for Computational Linguistics*, pages 5185–5198. ACL.

P. R. Biju and O. Gayathri. 2025. Indic approach to ethical AI in automated decision making system: implications for social, cultural, and linguistic diversity in native population. *AI and Society*, 40(8):6645–6670.

Steven Bird. 2024. Must NLP be extractive? In *Proceedings of the 62nd Annual Meeting of the Association for Computational Linguistics (Volume 1: Long Papers)*, pages 14915–14929, Bangkok, Thailand. Association for Computational Linguistics.

Steven Bird and Dean Yibarbuk. 2024. Centering the speech community. In *Proceedings of the 18th Conference of the European Association for Computational Linguistics*, pages 826–839. ACL.

Robert Borofsky. 2005. *Yanomami: The fierce controversy and what we can learn from it*. University of California Press.

James Bridle. 2022. *Ways of Being: Animals, Plants, Machines: The Search for a Planetary Intelligence*. Penguin.

E. Summerson Carr. 2010. Enactments of expertise. *Annual Review of Anthropology*, 39:17–32.

James Clifford and George Marcus. 2023. *Writing Culture: The Poetics and Politics of Ethnography*. University of California Press.

Ned Cooper, Courtney Heldreth, and Ben Hutchinson. 2024. "It's how you do things that matters": Attending to process to better serve indigenous communities with language technologies. In *Proceedings of the 18th Conference of the European Chapter of the Association for Computational Linguistics*, pages 204–211. ACL.

Kate Crawford. 2021. *The Atlas of AI: Power, Politics, and the Planetary Costs of Artificial Intelligence*. Yale University Press.

Catia Cucchiarini. 1993. *Phonetic Transcription: A Methodological and Empirical Study*. Ph.D. thesis, Radboud University.

Edward T. Hall. 1976. *Beyond Culture*. Garden City.

Stephen Harris. 1984. Aboriginal learning styles and formal schooling. *The Australian Journal of Indigenous Education*, 12(4):3–23.

Thomas Headland, Kenneth Pike, and Marvin Harris, editors. 1990. *Emics and Etics: The Insider/Outsider Debate*. Sage Publications.

Paula Helm, Gábor Bella, Gertraud Koch, and Fausto Giunchiglia. 2024. Diversity and language technology: How language modeling bias causes epistemic injustice. *Ethics and Information Technology*, 26(1):8.

Robin Wall Kimmerer. 2013. *Braiding Sweetgrass: Indigenous Wisdom, Scientific Knowledge and the Teachings of Plants*. Penguin.

Bruno Latour. 2007. *Reassembling the Social: An Introduction to Actor-Network-Theory*. Oxford University Press.

Patti Tamara Lenard. 2020. Culture. In Edward N. Zalta, editor, *The Stanford Encyclopedia of Philosophy*.

Jason Edward Lewis, Hēmi Whaanga, and Ceyda Yolgörmez. 2024. Abundant intelligences: Placing AI within Indigenous knowledge frameworks. *AI and Society*.

Keoni Mahelona, Gianna Leoni, Suzanne Duncan, and Miles Thompson. 2023. OpenAI's whisper is another case study in colonisation. *Papa Reo*. https://blog.papareo.nz/whisper-is-another-case-study-in-colonisation/, Accessed 20250215.

Nina Markl, Lauren Hall-Lew, and Catherine Lai. 2024. Language technologies as if people mattered: Centering communities in language technology development. In *Proceedings of the 2024 Joint International Conference on Computational Linguistics, Language Resources and Evaluation*, pages 10085–10099.

Donald Metzler, Yi Tay, Dara Bahri, and Marc Najork. 2021. Rethinking search: Making domain experts out of dilettantes. In *SIGIR Forum*, volume 55, pages 1–27. ACM.

Benjamin Mitchell and William Fleischman. 2020. Responsibility in the age of irresponsible speech. In Mario Arias-Oliva, Jorge Pelegrín-Borondo, Kiyoshi Murata, and Ana María Lara Palma, editors, *Societal Challenges in the Smart Society*, pages 587–592. Universidad de La Rioja.

Sjur Nørstebø Moshagen, Flammie Pirinen, Lene Antonsen, Børre Gaup, Inga Mikkelsen, Trond Trosterud, Linda Wiechetek, and Katri Hiovain-Asikainen. 2023. The GiellaLT infrastructure: A multilingual infrastructure for rule-based NLP. In Arvi Hurskainen, Kimmo Koskenniemi, and Tommi Pirinen, editors, *Rule-Based Language Technology*, pages 70–94. Northern European Association for Language Technology.

Elinor Ochs. 1979. Transcription as theory. *Developmental Pragmatics*, 10:43–72.

Juhyun Oh, Inha Cha, Michael Saxon, Hyunseung Lim, Shaily Bhatt, and Alice Oh. 2025. Culture is everywhere: A call for intentionally cultural evaluation. In *Findings of the Association for Computational Linguistics: EMNLP 2025*, pages 19156–19168, Suzhou, China. Association for Computational Linguistics.

Siddhesh Pawar, Junyeong Park, Jiho Jin, Arnav Arora, Junho Myung, Srishti Yadav, Faiz Ghifari Haznitrama, Inhwa Song, Alice Oh, and Isabelle Augenstein. 2025. Survey of cultural awareness in language models: Text and beyond. *Computational Linguistics*, 51(3):907–1004.

Andrew J. Peterson. 2025. AI and the problem of knowledge collapse. *AI and Society*, pages 1–21.

Kenneth L Pike. 1967. Etic and emic standpoints for the description of behavior.

Aidan Pine, Erica Cooper, David Guzmán, Eric Joanis, Anna Kazantseva, Ross Krekoski, Roland Kuhn, Samuel Larkin, Patrick Littell, Delaney Lothian, Akwiratékha' Martin, Korin Richmond, Marc Tessier, Cassia Valentini-Botinhao, Dan Wells, and Junichi Yamagishi. 2025. Speech generation for indigenous language education. *Computer Speech and Language*, 90:101723.

Karl Popper. 1935. *The Logic of Scientific Discovery*. Routledge.

Abhishek Raj, Harsh Singh, and Anshul Pachouri. 2025. Harnessing the potential of AI for Indian agriculture: Using "Bhashini" as a tool to deploy responsible AI and increase the uptake of AI applications among farmers. In Balaraman Ravindran and Abhishek Singh, editors, *Advancing Responsible AI in Public Sector Application*, pages 119–128. CRC Press.

Edward W. Said. 1978. *Orientalism*. Penguin Random House.

Daniel Schacter, Daniel Gilbert, and Daniel Wegner. 2009. *Psychology*. Macmillan.

Chirag Shah and Emily M. Bender. 2022. Situating search. In *Proceedings of the 2022 Conference on Human Information Interaction and Retrieval*, pages 221–232. ACM.

Ramesh Srinivasan. 2017. *Whose Global Village?: Rethinking how Technology Shapes Our World*. NYU Press.

Larry Sullivan. 2009. Emic and etic perspectives. In *The SAGE Glossary of the Social and Behavioral Sciences*, pages 174–175. Sage Publications.

Sharifa Sultana, Renwen Zhang, Hajin Lim, and Maria Antoniak. 2022. Narrative datasets through the lenses of NLP and HCI. In *Proceedings of the Second Workshop on Bridging Human–Computer Interaction and Natural Language Processing*, pages 47–54.

Paul Ten Have. 1989. The consultation as a genre. *Text and talk as social practice*, pages 115–135.

Maggie Walter and Chris Andersen. 2013. *Indigenous Statistics: A Quantitative Research Methodology*. Taylor and Francis.

Joseph Weizenbaum. 1966. ELIZA: A computer program for the study of natural language communication between man and machine. *Communications of the ACM*, 9:36–45.

Janie Wilson, Rebecca Ryan, and James Pugh. 2010. Professor–student rapport scale predicts student outcomes. *Teaching of Psychology*, 37(4):246–251.

William A. Woods. 1973. Progress in natural language understanding: an application to lunar geology. In *Proceedings of the National Computer Conference and Exposition*, pages 441–450. Association for Computing Machinery.

Naitian Zhou, David Bamman, and Isaac L Bleaman. 2025. Culture is not trivia: Sociocultural theory for cultural NLP. In *Proceedings of the 63rd Annual Meeting of the Association for Computational Linguistics*, pages 25869–25886.

Shoshana Zuboff. 2022. Surveillance capitalism or democracy? The death match of institutional orders and the politics of knowledge in our information civilization. *Organization Theory*, 3(3):1–79.

A Literature Survey

We surveyed the recent NLP literature on culture with the aim to obtain a **representative** picture of the treatment of culture in recent NLP research. Our literature is systematic, but not exhaustive.

We started by obtaining all papers published between 2022 and April 2026 in the main or findings track of a top-tier NLP conference (*ACL, EMNLP) which mention 'culture' or 'cultural' in their title. [1] To keep the scope for manual paper inspection manageable, we removed multimodal approaches, and papers whose primary focus is not culture. We reviewed a total of 107 papers (**full list in the Appendix**). We manually reviewed these papers and extracted the following information:

- **Primary contribution** (resource, method, analysis, survey)
- **Data creation** (predominantly manual, at least partially synthetic, based on a broad background corpus, . . .). Only for empirical papers (i.e., excluding surveys)
- **Task**. Assessment/improvement of cultural knowledge in LLMs (upstream or intrinsic) vs. task-oriented papers proposing culturally sensitive applications (downstream or extrinsic). Only for empirical papers (i.e., excluding surveys).

Summary statistics underpinning our analyses in Section 2 are provided in tables 2 – 4.

Category	N	%
Resource	50	0.467289719626168
Survey	8	0.0747663551401869
Analysis	15	0.14018691588785
Method	31	0.289719626168224

Table 2: Distribution over main contributions in papers in our survey.

Category	N	%
Assess / augment cultural knowledge in LLMs	78	0.728971962616822

Table 3: Fraction of paper with a focus on assessing or augmenting cultural knowledge in LLMs.

Category	N	%
culture annotators	19	0.177570093457944
broad data	41	0.383177570093458
synthetic	33	0.308411214953271
translation	3	0.0280373831775701
– / LLM output analysis	4	0.0373831775700935
– / survey	8	0.0747663551401869

Table 4: Distribution over data set creation approaches in the papers in our survey. Papers of the last two categories were not included, as they did not contribute data sets or investigated culture in LLM output (rather than human artifacts).

[1] We concede that our approach will not capture all works on culture – many won't mention the term in the title and/or are published in other venues – but we argue that this approach gives a representative picture of current foci and approaches.