

Eliciting Requirements for Culturally Aware AI-Mediated Social Platforms: Preliminary Findings

Paskalis Andrianus Nani
Department of SSC
Monash University
Melbourne, Australia
paskalis.nani@monash.edu

Chetan Arora
Department of SSC
Monash University
Melbourne, Australia
chetan.arora@monash.edu

John Grundy
Department of SSC
Monash University
Melbourne, Australia
john.grundy@monash.edu

Abstract—Social media platforms, including those with AI-mediation, are largely designed around assumptions drawn from Western, urban and digitally mainstream populations, even as adoption grows in communities governed by distinct kinship structures and customary authority. Recent work has shown that existing research into more culturally aware software engineering has concentrated on macro-level cultural-dimension models and front-end software engineering tasks. This has left the governance of AI-mediated cultural content to be so far largely unaddressed. This paper reports our preliminary findings from an ongoing multi-case field study that is eliciting requirements for a culturally aware AI-mediated social platform for diverse communities. In this paper we use the Malaka regency, Indonesia, a matrilineal community, as our case study community. Drawing on interviews with 12 Malaka stakeholders across six categories, we identify a generational asymmetry in trust and verification of Malaka cultural information, a cross-category consensus that cultural validation authority rests with customary (*adat*) elders, and a conditional, human-in-the-loop acceptance of artificial intelligence in which AI is expected to assist but never authorise cultural validation. We triangulate these findings against gaps in the CAwSE literature and distil three preliminary design principles intended to generalise beyond this case. Data collection in a second case (Southwest Sumba) and a subsequent survey are ongoing to validate these principles across contrasting cultural contexts.

Index Terms—Requirements engineering, Culturally-aware software engineering, Human-AI interaction, Field study

I. INTRODUCTION

Social media platforms are predominantly designed around assumptions drawn from Western, urban, and digitally mainstream user populations, even as much of their growth now comes from communities with distinct kinship structures, customary governance, and norms of communication [1]–[3]. In regions such as Eastern Indonesia, where cultural practice is closely tied to customary *adat*¹ authority and where the appropriateness of sharing, attributing, and validating cultural content is governed by community-specific norms, this mismatch is particularly consequential. A platform that treats all users and all content as interchangeable risks seriously misrepresenting, de-contextualising, or eroding the very cultural knowledge it aims to support. This concern is not merely theoretical: field studies of requirements engineering in

culturally and economically distinct communities have found that stakeholder needs and validation norms diverge substantially from assumptions embedded in mainstream software practice [4], [5].

Recent reviews of work on culturally aware software engineering (CAwSE) [6], [7] have examined how culture has been addressed across software engineering tasks and application domains. Their findings indicate that CAwSE research to date has been concentrated on a small set of macro-level cultural dimension models (most notably Hofstede’s and Schwartz’s frameworks) applied primarily to front-end, user-facing tasks such as requirements elicitation and culturally sensitive design. Later SE tasks, such as testing and deployment, have received only isolated coverage. To date, architecture, implementation, maintenance and DevOps are essentially absent from the literature. In particular, this body of work does not report AI or culturally sensitive intelligent software systems engineering as a distinct theme, despite the growing relevance of AI-mediated content moderation and recommendation on social networks.

This paper is part of a broader, ongoing research program that combines literature review with a two-case exploratory field study and a survey to elicit requirements for a culturally aware AI-mediated social platform designed for communities in Eastern Indonesia. The field study spans two regencies in East Nusa Tenggara with contrasting kinship systems (matrilineal Malaka and patrilineal Southwest Sumba). We hope that this will allow an eventual comparison of which requirements may generalise across cultural contexts, and which are context-specific. Data collection in Malaka is complete, and our data collection in Southwest Sumba and the subsequent survey are ongoing.

In this work we treat culturally sensitive requirements engineering for intelligent social platforms in Malaka as a case study. It makes three main contributions: (a) preliminary findings from interviews with 12 stakeholders across six categories, centred on generational trust in cultural information, the allocation of cultural validation authority, and expectations for AI governance; (b) an initial confrontation of these field findings against gaps identified in the CAwSE literature, from which we distil preliminary design principles intended to generalise beyond Malaka to culturally aware AI-mediated social platforms more broadly; and (c) the design of the

¹a variety of local customary practices and traditions

broader two-case study and its roadmap. The complete results of the two-case analysis and survey will be reported in a separate and subsequent publication.

II. BACKGROUND

Recent reviews [6]–[8] have characterised culture and values in software engineering largely through macro-level cultural dimension models. Within these frameworks, Hofstede’s power distance (PDI), individualism (IDV), uncertainty avoidance (UAI), and masculinity (MAS) dimensions, alongside Schwartz’s values of power, universalism, and self-direction, are the most heavily studied constructs. In terms of SE tasks, culture is incorporated most explicitly in stakeholder-oriented requirements work and culturally sensitive design, while testing and deployment receive only isolated coverage and architecture, implementation, maintenance, and DevOps are essentially absent. In application terms, the literature is split between a large domain-unspecified body of work and a smaller set of recurring domains (particularly education, health, and social systems/social software) suggesting CAwSE research is currently strongest in contexts where interaction, appropriateness, and stakeholder alignment are especially salient.

This work motivates the field study reported in this paper in three specific ways. First, the dominance of macro-level, dimension-based models leaves open how cultural authority and validation actually operate at the level of a specific community (the questions this field study directly elicits from stakeholders rather than inferring from abstract cultural traits). Second, the near-absence of maintenance-related tasks in the literature leaves ongoing, post-publication concerns (such as who may correct or validate cultural content over time) largely unaddressed, despite these concerns emerging as central to participants in this study (Section IV-B). Third, the absence of AI as a theme in the literature leaves open how AI-mediated tools should be governed within culturally sensitive contexts, a gap this study’s field findings (Section IV-C) speak to directly. In addition, AI-mediated cultural validation and authority allocation are not treated explicitly by studies reviewed to date. These gaps are further examined, alongside the corresponding field findings, in the triangulation presented in Section V.

III. STUDY DESIGN

This article is part of a larger and ongoing research program using multi-case study design [9], outlined in Figure 1. Interviews and survey are used to elicit requirements for a culturally aware AI-mediated social platform. We reviewed literature describing how culture has been addressed across software engineering tasks and domains. This revealed a concentration on macro-level cultural dimension models and a near-absence of coverage in later-stage SE tasks. Our reported field study extends this understanding by directly eliciting requirements from communities the target platform is intended to serve.

The field study spans two regencies, the Malaka regency as the first case and the Southwest Sumba regency as the second case in East Nusa Tenggara, Indonesia. These were




Fig. 1. Our three-stage research design

deliberately selected for a sharp contrast in kinship and social organisation: Malaka follows a matrilineal system [10], in which clan lineage passes through women and men marry into their wife’s clan, while Southwest Sumba follows a patrilineal system [11]. Notably, Malaka’s matrilineal structure does not straightforwardly place cultural authority with women: men retain a prominent decision-making role in a wide range of customary (*adat*) matters (from marriage negotiations to other communal *adat* activities) within their clan of origin, even after relocating to their wife’s clan. This distinction between lineage continuity (matrilineal) and customary decision-making authority (exercised prominently by men) makes Malaka a particularly informative case for examining how cultural validation authority is allocated in practice. Because kinship structure shapes who holds cultural authority, how knowledge and property are transmitted across generations, and who is recognised as entitled to speak for or validate cultural practice, the matrilineal-patrilineal contrast between the two cases provides a strong test of whether the governance and validation requirements identified in this study are specific to Malaka’s matrilineal context or generalise across kinship systems. This design supports an eventual cross-case comparison intended to distinguish requirements that generalise across cultural contexts from those that are context-specific, once data collection at both cases is complete.

In this paper we report on our Malaka case, where data collection is complete. Interviews were conducted with 12 participants, purposively sampled across six stakeholder categories with two participants each: Gen-Z youth, *adat* elders, cultural observers, local government representatives, educators, and cultural practitioners. This sampling strategy was intended to capture a wide range of representative perspectives that span both generational position (from young platform users to elder cultural authorities) and institutional role (from grassroots practice to local governance). The interviews were conducted

by the first author in Bahasa Indonesia, following a semi-structured protocol that covers context and cultural experience, cultural transmission practices, digital and social media habits, functional requirements, content sharing and consent norms, attitudes toward AI-assisted tools, and government support. Informed consent was obtained from all participants prior to interview.

The analysis at this stage followed an inductive thematic analysis approach [12], in which the first author read each transcript closely and generated initial codes directly from participants' own language, rather than applying a predefined coding scheme. Related codes were then iteratively clustered into candidate themes through a manual, spreadsheet-based coding process, and reviewed against both the coded extracts and the full data set to check coherence. This process produced the three themes reported in Section IV. Analysis is ongoing, and a second coding pass together with inter-rater checks is planned prior to finalising the complete field study report. Data collection in the second case (Southwest Sumba) and the subsequent survey (targeting more than 100 participants) are both in progress and are described further in Section VII.

IV. PRELIMINARY FINDINGS FROM MALAKA

This section reports our preliminary findings from an initial thematic analysis of our interviews with 12 participants in Kabupaten Malaka, purposively sampled across six stakeholder categories (two participants each): Gen-Z youth, *adat* elders, cultural observers, local government representatives, educators, and cultural practitioners. Given the early stage of analysis and the qualitative and purposive nature of the sample, the findings are organised around major cross-cutting themes and illustrative contrasts rather than frequency counts.

A. Generational Asymmetry in Trust and Verification of Cultural Information

Participants described markedly different strategies for verifying the accuracy of cultural information encountered online or elsewhere. Older and more experienced participants reported that cross-checking cultural information was straightforward, describing it as a matter of "validating it against their own knowledge and experience". Younger participants, by contrast, reported considerable difficulty in the same task, explaining that without a point of comparison, "everything will be considered true".

This asymmetry is notable because it does not appear to be a general digital-literacy gap (participants in all categories were confident and frequent users of social media and search platforms). Rather, the difficulty appears specific to cultural content, where authoritative reference points are tacit and experience-based rather than documented or easily searchable. This suggests that verification support for a culturally aware platform cannot rely solely on generic fact-checking mechanisms, but needs to make elder or expert validation traceable and visible to younger users at the point of consumption. For example, surfacing who validated a piece of content and on

what basis, rather than presenting cultural information as an undifferentiated feed.

B. Cultural Validation Authority a Cross-Category Agreement

Across stakeholder categories, participants consistently pointed to a specific authority structure to validate cultural content before it is shared publicly. When asked who should have the right to verify or correct cultural narratives once published, participants said that "*adat* elders and those who genuinely understand" the subject matter. This was echoed in views on content contribution more broadly: while participants generally agreed that anyone should be able to contribute content, they consistently qualified this position by requiring that content "must be validated before being shared".

The consistency of this position across otherwise very different stakeholder categories (from cultural practitioners to local government representatives) suggests that it reflects a shared normative expectation about cultural authority in this community, rather than an isolated view held only by those who currently hold that authority. This expectation is also consistent with the broader pattern of *adat* authority in Malaka described in Section III: although the lineage in Malaka is matrilineal, decision-making authority in *adat* matters is held prominently by men, and both *adat* elders interviewed in this study were male. Read together, this suggests that the "cultural validation authority" participants stated may be, in practice, a predominantly male authority structure, a nuance that is easily lost if Malaka's matrilineal system is read as implying female-held cultural authority more broadly. For platform design, this points toward a governance model built around designated validator roles tied to recognised cultural standing, rather than crowd-moderation or purely algorithmic content review, an implication that runs counter to the open, crowd-sourced moderation models common in many general-purpose social platforms.

C. Conditional Acceptance of AI: Assistive but Not Authoritative

Participants expressed a consistent and fairly sharp boundary between tasks they considered appropriate for AI assistance and tasks that require human (specifically elder) authority. AI was considered appropriate for lower-stakes mechanical tasks: transcribing interview recordings, translating local languages, organising narratives, answering visitor questions, and classifying photos and videos. However, participants explicitly stated that "cultural validation should not be done by AI, because AI can be wrong." Even in the specific case where an AI system detects a potential error in cultural content, participants indicated that its role should be limited to "providing notification of the error found and suggesting a correction," with final validation remaining with *adat* elders or other trusted individuals.

Trust in AI-provided answers about culture was also explicitly conditioned on traceability: participants indicated that they would trust an AI's answer "if it is accompanied by a valid

source.” Taken together, these findings describe a human-in-the-loop expectation that is not merely a general preference for caution, but a specific allocation of authority: AI as an assistive, source-citation tool, and *adat* elders as the sole final arbiters of cultural correctness.

A further, less expected finding concerns language. AI capable of understanding and speaking the local language was not framed as a threat to linguistic heritage, but as “one way to keep the local language from disappearing”, because participants saw local language AI capability as directly supporting, rather than undermining, preservation goals.

V. DISCUSSION

We contrast the preliminary field findings above with the gaps identified in recent reviews of CAwSE, summarised in Table I. The findings suggest that culturally aware AI requires modelling not only *cultural values*, but *cultural authority*. Existing CAwSE work has largely characterised cultural difference through abstract dimensions. Our participants in Malaka described cultural appropriateness in terms of concrete roles: who may contribute cultural knowledge; who is trusted to verify it; who has final authority to correct it. For AI-mediated social platforms, cultural awareness therefore involves not only adapting system behaviour to cultural values, but also representing and respecting community-specific structures of authority.

The most notable divergence concerns AI. Because the SLR findings does not surface AI as a theme at all, the human-in-the-loop governance expectations described in Section IV-C (including the specific boundary between AI as an assistive, source-citing tool versus *adat* elders as the sole authority for cultural validation) appear to be outside the scope of CAwSE research as currently synthesised. This suggests that as AI-based systems increasingly enter culturally sensitive domains, the largely dimension-based and task-distribution-based framework that has organised CAwSE research to date may need to be extended to explicitly account for AI governance and authority allocation [13], [14], rather than treating AI as an implementation detail within existing task categories such as requirements or design.

These findings from Malaka point toward at least three preliminary design principles that we suggest may generalise beyond this specific platform and case, to culturally aware AI-mediated social platforms operating in communities with strong customary or communal authority structures more broadly: (1) content validation workflows should treat AI as a flagging and suggestion mechanism rather than a decision-maker, with explicit sign-off required from designated cultural authorities prior to publication or correction; (2) governance and access-control model should separate broad, low-friction content contribution from narrow, authority-gated content validation, rather than relying on generic crowd-moderation [15]; and (3) AI assistants embedded in such platforms should be able to speak local-language and should consistently surface their sources when answering cultural questions, directly addressing the verification difficulty reported by younger par-

TABLE I
PRELIMINARY TRIANGULATION OF SLR FINDINGS AND FIELD FINDINGS
FROM MALAKA

Finding from reviews	Field Finding (Malaka)	Relation
CAwSE research is concentrated on a small set of macro-level cultural dimension models (most notably Hofstede (PDI, IDV, UAI, MAS) and Schwartz (power, universalism, self-direction)) used to characterize cultural difference at a national/societal level.	None of the participants’ governance or design expectations were framed in terms of abstract cultural dimensions; instead, requirements were grounded in concrete, community-specific roles (<i>adat</i> elders, event owners) and relationships.	Extends
Culture is incorporated into SE primarily through front-end, user-facing tasks (especially stakeholder-oriented requirements work and culturally sensitive design) while testing and deployment receive only isolated coverage and architecture, implementation, maintenance, and DevOps are essentially absent from the literature.	The most salient governance concern raised by participants (continuous content validation and AI-flagged correction after publication) functions as an ongoing, maintenance-like task, an SE task category the SLR found to be essentially unaddressed in CAwSE research.	New
Applications are split between a large domain-unspecified body of work and a smaller set of recurring domains, particularly education, health, and social systems/social software.	The target platform combines social software with youth-oriented cultural education, sitting at the intersection of two domains the SLR identifies as recurring but does not report being studied together.	Extends
AI is not reported as a distinct theme in the SLR findings; the literature’s treatment of culture in SE is organized entirely around dimension-based models and SE task distribution, without addressing AI-mediated cultural content or AI governance.	A human-in-the-loop boundary for AI (assistive for mechanical tasks, explicitly excluded from cultural validation authority) emerged as one of the strongest and most consistent themes in the field data.	New

ticipants in Section IV-A. For local language adoption in particular, AI tools should not only be accurate and practical but also be designed in close collaboration with the native speaker community [16]. We treat these as preliminary, case study-level principles rather than validated requirements: they are grounded in a single cultural context and are intended to be tested and refined against the contrasting kinship and authority structures of the second case study (Section VII).

VI. THREATS TO VALIDITY

These findings are preliminary and subject to several limitations. The design principles proposed in Section V rest on the theoretical replication logic of this multiple-case study design [9]: Case 1 (Malaka) and Case 2 (Southwest Sumba) were selected for their contrasting kinship systems precisely so that the principles emerging from one case could be tested, rather than merely repeated, in the other. However, at this stage, only Case 1 has been analysed, so this logic remains prospective; the principles proposed in Section V should be read as candidate principles awaiting replication, not as validated or generalised findings. Several further limitations

follow. First, the findings are based on a single case; while Malaka and Southwest Sumba were deliberately selected to differ in kinship structure (matrilineal versus patrilineal), the findings from Case 1 alone cannot be assumed to replicate in Case 2. In particular, the cross-category consensus on elder-held validation authority reported in Section IV-B may be shaped in part by norms of cultural authority specific to Malaka's matrilineal system and should be treated as a hypothesis to be tested against the patrilineal context of Southwest Sumba rather than an established cross-cultural finding. Second, thematic coding at this stage was performed by a single researcher without a second coder or inter-rater agreement check, and analysis has not yet reached saturation; the themes reported here should be read as an initial synthesis rather than a finalised coding scheme. Third, interviews conducted in bahasa Indonesia and reported here in English involved translation, which may introduce shifts in nuance, particularly for culturally specific terms and concepts. Finally, as researchers not native to Malaka, the authors' own positionality may have shaped both the interview dynamic and the interpretation of responses. The authors had a previous working relationship with contacts in Malaka established through earlier collaboration, which facilitated access and initial relationship, but may also have shaped which participants were reachable and how they responded; to help mitigate this, a local research assistant familiar with the community was involved in conducting the interviews.

VII. NEXT STEPS AND EXPECTED CONTRIBUTIONS

The next stages of this research program are to (i) complete data collection and initial analysis for the second case in Southwest Sumba; (ii) conduct a cross-case thematic analysis to identify which requirements generalise across the two cultural contexts and which are context-specific, for instance, examining whether Southwest Sumba's patrilineal kinship system and its associated *belis* (bride-price) customs, involving substantial exchanges of horses and buffalo and entirely absent in Malaka, correspond to a different allocation of cultural validation authority than the one described in Section IV-B; (iii) administer the planned survey to more than 100 participants to quantitatively triangulate themes emerging from both cases; and (iv) synthesise these results into validated design principles for culturally aware AI-mediated social platforms, with a prototype evaluated with participating communities. We will then prototype this platform, evaluate it with end users, and compare different cultural group user feedback on key cultural aspects of this AI-mediated social platform with our requirements. We will use our findings to distil a set of guidelines for improved cultural software engineering practices with diverse communities.

Upon completion, this research program is expected to contribute: to research, one of the first empirically grounded, multi-stakeholder accounts of requirements for culturally aware AI-mediated platforms. To practice, a set of concrete, field-derived design guidance (particularly around AI governance and validation authority) for practitioners building

platforms intended to serve culturally distinct communities in Eastern Indonesia and comparable contexts.

VIII. CONCLUSION

This paper treated Malaka regency, Indonesia, as a case study to obtain candidate design principles for culturally aware AI-mediated social platforms, based on interviews with 12 stakeholders in six categories. Initial findings point to a generational asymmetry in trust and cultural information verification, a cross-category consensus on elder-held validation authority, and a conditional, human-in-the-loop acceptance of AI that the current CAwSE literature does not yet address. These findings motivate ongoing data collection in Case 2 and a subsequent survey toward a validated set of design principles grounded in both the literature and cross-case field evidence.

ACKNOWLEDGEMENT

This research is supported by the Indonesia Endowment Fund for Education (LPDP), Ministry of Finance, Republic of Indonesia.

REFERENCES

- [1] J. Henrich, S. J. Heine, and A. Norenzayan, "The weirdest people in the world?" *Behavioral and Brain Sciences*, vol. 33, no. 2–3, pp. 61–83, Jun. 2010, doi: 10.1017/S0140525X0999152X.
- [2] X. Li, Y. Hu, C. S. Huang, and S. S. Chuang, "Beyond W.E.I.R.D. (western, educated, industrial, rich, democratic)-centric theories and perspectives: masculinity and fathering in chinese societies," *Journal of Family Theory & Review*, vol. 13, no. 3, pp. 317–333, Sep. 2021, doi: 10.1111/jftr.12403.
- [3] A. A. Septiandri, M. Constantinides, M. Tahaei, and D. Quercia, "WEIRD FAccTs: How western, educated, industrialized, rich, and democratic is FAccT?" in *2023 ACM Conference on Fairness, Accountability, and Transparency*. Chicago, IL, USA: ACM, Jun. 2023, pp. 160–171, doi: 10.1145/3593013.3593985.
- [4] T. Kanij, M. Anwar, G. Oliver, and M. K. Hossain, "Developing software for low socio-economic end users: Lessons learned from a case study of fisherfolk communities in Bangladesh," in *2023 IEEE/ACM 45th International Conference on Software Engineering: Software Engineering in Society (ICSE-SEIS)*. Melbourne, Australia: IEEE, May 2023, pp. 96–107, doi: 10.1109/ICSE-SEIS58686.2023.00015.
- [5] C. S. Muzammel, M. Spichkova, and J. Harland, "Cultural influence on RE activities: An extended analysis of state of the art," in *26th International Conference on Mobile Human-Computer Interaction*. Melbourne, VIC, Australia: ACM, Sep. 2024, pp. 1–8, doi: 10.1145/3640471.3680236.
- [6] Anonymous Authors, "Omitted for double-blind reviewing," 2026, omitted for double-blind reviewing.
- [7] T. Alsanoosy, M. Spichkova, and J. Harland, "Cultural influence on requirements engineering activities: a systematic literature review and analysis," *Requirements Engineering*, vol. 25, no. 3, pp. 339–362, 2020.
- [8] M. Shatin, W. Hussain, A. Nurwidyantoro, H. Perera, R. Shams, J. Grundy, and J. Whittle, "Operationalizing human values in software engineering: A survey," *IEEE Access*, vol. 10, pp. 75 269–75 295, 2022.
- [9] P. Runeson and M. Höst, "Guidelines for conducting and reporting case study research in software engineering," *Empirical Software Engineering*, vol. 14, no. 2, pp. 131–164, Apr. 2009, doi: 10.1007/s10664-008-9102-8.
- [10] T. Therik, *Wehali: The Female Land*, ser. Comparative Austronesian Series. Canberra: ANU Press, 2023.
- [11] J. A. C. Vel, *Uma politics: an ethnography of democratization in West Sumba, Indonesia, 1986–2006*, ser. Verhandelingen van het Koninklijk Instituut voor Taal-, Land- en Volkenkunde. Leiden: KITLV Press, 2008, no. 260, doi: 10.1163/9789004253926.
- [12] V. Braun and V. Clarke, "Using thematic analysis in psychology," *Qualitative research in psychology*, vol. 3, no. 2, pp. 77–101, 2006.

- [13] E. Papagiannidis, P. Mikalef, and K. Conboy, "Responsible artificial intelligence governance: A review and research framework," *The Journal of Strategic Information Systems*, vol. 34, no. 2, p. 101885, Jun. 2025, doi: 10.1016/j.jsis.2024.101885.
- [14] E. LaRosa, "Curated information frameworks and the governance of artificial intelligence," *AI and Ethics*, vol. 6, no. 4, p. 407, Aug. 2026, doi: 10.1007/s43681-026-01242-9.
- [15] F. Volz, C. Münch, M. Lohmüller, and C. Küffner, "From data jungle to data governance in digital ecosystems: Empirical evidence from a multiple holistic case study," *Journal of Business Research*, vol. 201, p. 115747, Dec. 2025, doi: 10.1016/j.jbusres.2025.115747.
- [16] M. Eriksson Krutrök and P. Poromaa Isling, "Technolinguistic suspension and the digital futures of minority languages: AI and Meänkieli in Sweden," *New Media & Society*, p. 14614448251385920, Nov. 2025, doi: 10.1177/14614448251385920.