Beyond Localisation: Cultural Defaults and User Agency in Conversational AI

Anonymous Author(s)

Anonymous Institution

Abstract. This extended abstract proposes a social computing framework for analysing cultural defaults in conversational AI. It distinguishes under-alignment from over-alignment and separates cultural responsiveness from relational positioning.

Keywords: conversational AI · cultural alignment · user agency · social computing

1 Motivation

Conversational AI systems increasingly operate where social interpretation matters, including emotional support, health communication, advice, companionship, and migration support. Cultural mismatch in these settings is not only an accuracy problem but a relational one: systems may treat direct self-expression, individual autonomy, secular reasoning, or low-hierarchy directness as neutral defaults.

This problem has two directions. Under-alignment ignores cultural context, producing alienation, mistrust, and failed disclosure in high-stakes settings [1]. Over-alignment is less discussed: a system may infer that a user belongs to a cultural category and prescribe responses accordingly. This is especially risky for migrant, diasporic, and multi-sited users whose cultural position is negotiated across multiple worlds.

Existing responses, including retrieval, instruction tuning, culturally adapted benchmarks, and local models, improve coverage and benchmark performance [2]. Yet benchmark success is not relational competence: a system may know cultural facts while failing to manage hierarchy, face, or disclosure in live interaction.

2 Framework and Cases

We propose a two-axis framework. The first axis is *depth of cultural responsiveness*: *surface localisation* adapts visible forms; *structural localisation* adapts knowledge infrastructure; and *deep cultural attunement* adapts interactional stance around hierarchy, indirectness, intimacy, stigma, obligation, and user agency. Their risks are tokenism, essentialism, and paternalism.

The second axis is *relational positioning*. A *cultural ally* supports the user's standpoint; a *culture interpreter* explains another cultural logic; and a *cross-cultural mediator* helps users navigate between multiple logics without treating one as universally authoritative. Mediation is the most demanding role because it requires holding value tensions open while preserving user agency [3].

We developed the framework through an abductive synthesis of six recent cases spanning Korean retrieval, Southeast Asian multilingual models, Arab commonsense transfer, Japanese politeness modulation, Saudi emotional support, and Cantonese-speaking immigrant health communication. The resulting 3×3 matrix locates current systems mainly in structural-localisation and ally/interpreter cells, while cross-cultural mediation remains sparsely evidenced.

Across cases, three patterns recur: linguistic coverage does not equal relational competence; cultural knowledge does not equal mediation capacity; and deep attunement within one cultural logic does not equal cross-cultural mediation. A system may support users sensitively within one frame while failing users who must move between frames.

3 Implications

The framework suggests that cultural alignment should be evaluated not only by whether a system knows more cultures, but by whether cultural framing remains visible, contestable, and user-led. The cross-cultural mediator cell remains largely empty, making it a design and evaluation agenda rather than an achieved capability.

This extended abstract contributes a compact diagnostic vocabulary for early-stage evaluation of cultural alignment in conversational AI. It identifies under-alignment and over-alignment as paired risks, separates knowledge coverage from situated judgement, and highlights cross-cultural mediation as a design target for users navigating multiple cultural worlds.

References

1. Aldaweesh, S., Alelsheikh, G., Alhamed, F., Van Kleek, M., Shadbolt, N.: "It Hasn't Lived in Our Society": Investigating cultural sensitivity in LLM chatbots for emotional support. In: Proceedings of the 2026 CHI Conference on Human Factors in Computing Systems. pp. 1–20 (2026), [DOI]
2. Liu, C.C., Gurevych, I., Korhonen, A.: Culturally aware and adapted NLP: A taxonomy and a survey of the state of the art. Transactions of the Association for Computational Linguistics **13**, 652–689 (2025), [DOI]
3. Suchman, L.A.: Plans and Situated Actions: The Problem of Human-Machine Communication. Cambridge University Press (1987)