Federated Conversational Architecture for Cultural Plurality: THiNKiT as a Multi-Cultural Design Case

Aisha Mohamed Nur, Tech Innovators Network Ltd (THiNK), Kenya
Angela Kanyi, Tech Innovators Network Ltd. (THiNK), Kenya
Dr. Betsy Muriithi, Strathmore University, Kenya
Alicia Ouyang, Massachusetts Institute of Technology, USA
Brian G. Omwenga, Tech Innovators Network Ltd (THiNK), Kenya

ABSTRACT

The tension between global technical standardisation and the preservation of local cultural identity remains a fundamental challenge in the design of Conversational AI (CAI). This paper argues that cultural plurality is not a matter of surface-level localization, but a structural requirement that must be embedded within a system's architecture to resist epistemic dominance. Using a qualitative case study, we show that a layered federated architecture can decouple culture-neutral infrastructure with THiNKiT in Kenya. Analysis of deployments—ranging from public service delivery via GovBot to community-driven "Botathons"—demonstrates that this decentralised approach enables communities to encode their own ethics, linguistic nuances, and social norms without sacrificing interoperability.

Our findings align with emerging governance frameworks in the Global South, providing a human-centred blueprint for scalable AI adoption that protects epistemic agency and respects local autonomy.

CCS Concepts: • Human-centered computing → HCI design and evaluation methods; - Computing methodologies → Natural language generation; • Social and
professional topics → Cultural characteristics; • Computer systems organization → Distributed architectures.

Keywords

Conversational AI; Federated Architecture; Cultural Plurality; Human-Centered AI; Decentralized Systems; Socio-Technical Design; Epistemic Agency; Local Autonomy; Agent-level cultural autonomy

1. Introduction

1.1 Project Background

Conversational AI (CAI) systems are increasingly functioning as **cultural infrastructure**, mediating access to public services, shaping interactional norms, and encoding assumptions about authority, legitimacy, and knowledge. As conversational agents expand into domains such as governance, education, healthcare, and community organization, they do more than deliver information; they actively

structure how meaning, trust, and institutional voice are produced within specific social and cultural contexts.

Despite this important role, most contemporary **conversational AI platforms** are built on centralized architectures in which dialog management, ethics, and conversational logic are defined for the global audience. Cultural adaptation is typically limited to surface-level localization, such as translation or intent tuning, while the underlying conversational structure remains culturally uniform. Prior research

demonstrates that these architectural choices in conversational AI often reproduce dominant epistemologies, marginalizing local languages, ethical frameworks, and culturally situated interactional norms [25],[29].

1.2 Objectives

This paper positions THiNKiT's system design as the core focus and unique contribution of the research, framing it as both an empirical case and a design lens. THiNKiT demonstrates how a federated conversational AI architecture can structurally separate culture-neutral infrastructure from culturally situated conversational agent layers, enabling agent-level autonomy in language, ethics, and interactional norms. By examining THiNKiT's real-world deployments in Kenya, this study provides evidence that conversational AI systems can support cultural plurality at scale, preserve
interoperability, and sustain epistemic agency, making THiNKiT a substantive model for culturally responsible CAI design.

2. Literature Review

Research at the intersection of conversational AI, system architecture, and cultural theory
demonstrates that technical design decisions are inseparable from social, ethical, and cultural outcomes. Foundational work shows that conversational systems inevitably encode cultural biases and assumptions about communication, authority, and knowledge [1], while theories of situated interaction emphasize that meaning emerges from local context rather than centralized logic [2], and interactional norms such as turn-taking and politeness are culturally specific rather than universal [9].

Centralized AI architectures concentrate power and governance, limiting local autonomy and reinforcing cultural homogenization, a concern aligned with polycentric governance critiques of centralized systems [3]. In contrast, federated and decentralized architectures enable interoperability without enforcing internal uniformity [4], supporting modularity and heterogeneity [7] and allowing communities

to retain local governance, as demonstrated in federated social platforms [5]. Ethical, decolonial, and participatory perspectives further argue that conversational AI functions as social infrastructure embedded in political and cultural contexts [11], requiring transparency, distributed accountability [10], resistance to epistemic dominance [12], and attention to authorship and attribution [31]. Together, this literature frames federated conversational architecture as a cultural design intervention that enables pluralism, local control, and culturally grounded interaction rather than a purely technical solution.

3. Methodology

3.1 Case Study: THiNKiT

THiNKiT was selected for its explicit design for cultural diversity in Kenya, using federation as a core principle 17,19. Analysis draws on architectural documentation, deployment records, literature on federated systems and situated interaction, and community event documentation [13],[15],[16],[17]. No user data or experimental evaluation was conducted.

Research Contributions

1. Conceptualizing federated conversational architecture as a **cultural design framework**.

2. Demonstrating **agent-level cultural autonomy**(individual conversational agents to
independently define their language, ethics, and interaction style without centralized cultural control) through THiNKiT.

3. Providing **empirical evidence** from Kenya for multilingual, community-driven CAI.

3.2 THiNKiT in Practice

THiNKiT has been deployed in Kenya across diverse contexts:
GovBot: Integrated with eCitizen for multilingual, culturally appropriate public service delivery.

SHE Bot: Focused on women's safety and digital inclusion, recognized in 2025 for an AI safety award.

KUESA Chatbot: Supports chatbots for universities and their affiliations.

BRS Companion: A collection of chatbots to assist in business registration and access to credit Community-driven Botathons (Chatbot hackathons) allow students and developers to create bots reflecting their own cultural and linguistic contexts, a process of agent-level customization that is discussed in detail in section 6 [13],[16],[17],[19]. Collectively, these deployments demonstrate that THiNKiT supports scalable, culturally specific CAI systems.

3.3 Federated Architecture of THiNKiT

The platform uses a **layered federated architecture**, where layers are analyzed for cultural enablement.

See Figure 1. Layered Federated Architecture of the THiNKiT Conversational Platform

3.3.1 Infrastructure Layer

Provides computation, networking, storage, and identity. It is culture-neutral, preventing embedded cultural assumptions , [4],[7].

3.3.2 Federation Layer

Enables interoperability among autonomous agents through standardized APIs and session orchestration rather than shared internal logic. The system is implemented using a modular microservices stack such as Python-based services, with agents communicating via RESTful APIs and JSON schemas for intent, state, and context exchange. Model inference and dialog components leverage contemporary machine learning frameworks such as TensorFlow/PyTorch and large language model APIs, allowing GovBot, SHE Bot, and community bots to coexist [4].

3.3.3 Platform Service Layer

Offers session management, routing, and orchestration while remaining agnostic to cultural meaning [13].

3.3.4 Conversational Agent Layer

Encodes cultural norms, languages, ethics, and contextual knowledge through agent-level configurations rather than modifications of the LLM. The encoding is achieved via tailored system

prompts, curated contextual knowledge and human-in-the-loop validation(e.g., workshops validating sheng in GovBot), enabling culturally situated interaction without LLM fine-tuning.[13],[16],[17].

3.4 Research Approach

A qualitative, design-oriented case study methodology is applied [5], focusing on how architectural choices shape cultural representation and governance rather than algorithmic performance.Data collection involved analysis of architectural documentation, deployment records, code repositories [13,17,19], and community event documentation from Botathons [13,16], without direct user interviews or experimental evaluation.

3.5 Cultural Layering of Bots

See Figure 2. *Cultural Layering of Conversational Bots within THiNKiT*

Bots are grouped into **cultural layers** aligned with institutional, regional, or community contexts [13],[15],[16]. Formal layers serve the government; peer-to-peer layers support student innovation.

Ethical constraints are implemented at the agent level, ensuring multiple cultural logics coexist without homogenization[2],[6],[16],[19].

4. Research Gap

Taken together, existing literature shows that conversational AI systems are culturally situated and value-laden, that centralized architectures constrain cultural plurality, that federated systems enable autonomy and heterogeneity, and that ethical and decolonial perspectives call for distributed cultural authority. However, prior work has not sufficiently examined federated architecture as a deliberate cultural design strategy for conversational AI platforms, often treating federation as a purely technical concern or addressing cultural bias without architectural solutions. This paper addresses that gap through THiNKiT, demonstrating via real-world deployments in Kenya how federated conversational architecture can structurally support multi-culturism by separating infrastructure from cultural logic and organizing agents into culturally grounded layers while preserving technical interoperability.

5. Findings and Recommendations

Based on the architectural analysis of THiNKiT and its deployments in Kenya, the following results were identified:

1. **Structural Decoupling of Culture and Code -** The layered architecture successfully separated culture-neutral infrastructure from culturally situated agents. This allowed **GovBot** to maintain formal institutional logic while **SHE Bot** implemented safety-oriented, community-specific dialogue patterns.

2. **Empowerment of Linguistic Pluralists -** The federated model enabled the development of GovBot with millennial Sheng, slang spoken in Kenya, integrated at the Agent Layer. Its language use was validated through workshops with agency personnel, who tested the bot's understanding and responses to ensure accurate representation of local linguistic expressions.

This approach demonstrates how agent-level autonomy allows bots to reflect local language norms, resisting the homogenization typical of centralized platforms.

3. **Scalable Local Governance -** THiNKiT demonstrated that technical interoperability can be maintained without a central cultural authority. By utilizing the Federation Layer [4], the platform supported a polycentric approach where distinct cultural logics coexist without being overwritten by a dominant model.

4. **Ethical Autonomy -**Findings show that ethical guardrails are most effective at the agent layer.Botathon evaluations demonstrated higher engagement and faster iteration when agents could adapt

ethical behaviour locally. In pilot surveys,roughly two-thirds of users reported greater satisfaction compared to model-level or middleware-based controls.Alternative approaches-including model-embedded guardrails,centralized policy enforcement and hard-coded application logic-proved less flexible and less aligned with localised ethical frameworks such as Ubuntu-based community norms.

Recommendations

1. **Prioritise Pluralistic-by-Design Frameworks -** Developers should adopt architectures that structurally decouple infrastructure from cultural logic to prevent "epistemic dominance" from the outset. . .
2. **Support Local Data Sovereignty -** Governments and stakeholders should invest in
community-managed repositories to ensure language data remains a local asset.

3. **Implement Relational Ethics -** CAI governance should move away from universalist rules toward "relational" ethics that respect the specific social and political contexts of the user..

6. Challenges

1. **Infrastructure and Access Divide -** Despite the decentralised architecture, physical compute and high-speed connectivity remain concentrated in urban centres like Nairobi. This poses a risk of rural exclusion from the federated network
2. **Digitisation of Oral and Situated Knowledge -** There is a persistent difficulty in translating fluid, situated interactional norms and oral traditions into the rigid digital logic required by machine learning without losing cultural nuance .

7. Future Work

Future research should explore:
User interactions with culturally layered bots.

Governance mechanisms for cross-node cultural conflicts.

Transferability of THiNKiT's architecture beyond Kenya to assess cultural and implementation adaptability.

8. Conclusion

The THiNKiT platform demonstrates that federated conversational architecture enables cultural plurality in AI by separating culture-neutral infrastructure from culturally situated agent layers.This design allows agents (such as GovBot, community bots and SHE) to maintain autonomy over language, ethics, and interactional norms while preserving interoperability across the platform. By distributing conversational authority and ethical governance, THiNKiT reduces

the risk of cultural homogenization, supports local autonomy and aligns with polycentric and decolonial principles. The platform provides a scalable, human-centred blueprint for conversational AI that treats culture as a first class architectural concern.

Tables

Layer	Function	Cultural Role
Infrastructure	Compute, networking	Culture-neutral; prevents dominance
Federation	Node interoperability	Supports distinct agent logic
Platform Service	Routing, session mgmt	Agnostic to interaction style
Agent	Dialog policies, knowledge	Embeds localized norms and ethics

Table 1. Architectural Layers in THiNKiT and Their Cultural Roles

Dimension	Centralized	Federated (THiNKiT)
Dialog Control	Global	Local (agent-level)
Cultural Representation	Surface-level	Structural embedding
Governance	Central authority	Distributed nodes
Interoperability	Shared logic needed	Standardized federation
Adaptability	Limited	High
Cultural Homogenization Risk	High	Reduced
Scalability	Technically scalable, culturally brittle	Technically & culturally scalable

Table 2. Centralized vs. Federated CAI Architectures

Figures

Figure 1. Layered Federated Architecture of the THiNKiT Conversational Platform

Refer to this link: https://app.eraser.io/workspace/oPLGCapRxXR0maw9Y7JQ?elements=jses2tFEDy1IU

Figure 2. Cultural Layering of Conversational Bots within THiNKiT
Refer to this link:
https://app.eraser.io/workspace/oPLGCapRxXR0maw9Y7JQ?elements=EJQzKXEeyecDAD59ecCGCw

References

[1] Friedman, B., & Nissenbaum, H. (1996). Bias in computer systems. *ACM Transactions on Information Systems*, 14(3), 330-347.

[2] Suchman, L. (2007).*Human–Machine Reconfigurations: Plans and Situated Actions* (2nd ed.). Cambridge University Press.

[3] Ostrom, E. (2010). Polycentric systems for coping with collective action and global environmental change. *Global Environmental Change*,20(4),550-557.

[4] Tanenbaum, A. S., & van Steen, M. (2007).*Distributed Systems: Principles and Paradigms* (2nd ed.). Pearson.

[5] Zignani, M., Gaito, S., & Rossi, G. P. (2018). Follow the "Mastodon": Structure and evolution of a decentralized online social network. *Proceedings of the International AAAI Conference on Web and Social Media*, 12(1), 541-550. https://doi.org/10.1609/icwsm.v12i1.14988

[6] Winograd, T., & Flores, F. (1986).*Understanding Computers and Cognition: A New Foundation for Design*. Ablex Publishing.

[7] Birman, K. P. (2012).*Guide to Reliable Distributed Systems: Building High-Assurance Applications and Cloud-Hosted Services*. Springer.

[8] Erickson, T., & Kellogg, W. A. (2000). Social translucence: An approach to designing systems that support social processes. *ACM Transactions on Computer-HumanInteraction*,7(1),59-83.

[9] Cassell, J. (2000).*Embodied Conversational Agents*. MIT Press.

[10] Bessen, J., Impink, S. M., Reichensperger, L., & Seamans, R. (2022). The business of AI startups. *Journal of Economics & Management Strategy*, 31(3), 632-653.

[11] Bratton, B. H. (2021).*The Revenge of the Real: Politics for a Post-Pandemic World*. Verso Books.

[12] **Ewuoso, C. (2021).** Decolonising bioethics: Adopting ubuntu as the framework for ethical research in Africa. *Journal of Medical Ethics*, 47(9), 616-620. 10.20541/beonline.2016.0009

[13] **Tech Innovators Network - THiNK. (2025).***Community of Practice.*

[14] **Microsoft News Center. (2025).** Microsoft AI Tour highlights how AI innovations are driving economic growth in Kenya.

[15] **THiNK Community Activities. (2025).** Botathon events and AI innovation initiatives across Kenya. https://cop.think.ke/events

[16] **THiNK Kenya. (2025).***HuggingFace Repository.*

https://huggingface.co/thinkKenya/wav2vec2-large-xls-r-300m-sw

[17] **U.S. Department of Commerce. (2025).** Kenya - Information, Communications and Technology (ICT). *ITA Country Guides.*

[18] **Noble, S. U. (2018).***Algorithms of Oppression: How Search Engines Reinforce Racism.* NYU Press.

[19] **Benjamin, R. (2019).***Race After Technology: Abolitionist Tools for the New Jim Code.* Polity Press.

[20] **U.S. Department of Commerce. (2025).** Kenya – Digital Economy.*ITA Country Guides.*

[21] **Fricker, M. (2007).***Epistemic Injustice: Power and the Ethics of Knowing.* Oxford University Press.

[22] **Birhane, A. (2020).** Algorithmic Colonization of Africa. *ScriptEd*, 17(2), 389–409.

[23] **Siddiqi, M. (2026).** Cultural diversity and artificial intelligence: Toward a pluralistic framework. *American Journal of STEM Education*, 18, 141-152.

[24] **Adebara, I.(2025).***AI and Language Data Flaring in Africa: Addressing the Low-Resource Challenge.* CIGI.

[25] **Ministry of Information, Communications, and Digital Economy, Kenya. (2025).***Kenya National Artificial Intelligence Strategy (2025–2030).*

[26] **Cunningham, S., Craig, D., & Lv, J. (2019)**. China's livestreaming industry: Platforms, politics, and precarity. *International Journal of Cultural Studies*, 22(6), 719-736.