



LAPIN YLIOPISTO
UNIVERSITY OF LAPLAND

Sobia Iqbal Farooqui

Reproducing Coloniality Through AI:
**A Case Study of LLM Integration in
Pakistani Organisations**

Master's Thesis

Service Design Strategies and Innovations

Spring 2026

University:

University of Lapland, Faculty of Art and Design

Title

Reproducing Coloniality Through AI: A Case Study of LLM Integration
in Pakistani Organisations

Author

Sobia Iqbal Farooqui

Degree program

Service Design Strategies and Innovations (SDSI)

The type of the work

Master's thesis

Number of pages:

107

Number of appendices:

2

Year:

Spring 2026

Abstract

As large language models become embedded in professional life across the Global South, questions about whose knowledge they carry and whose they erase have become impossible to ignore. Pakistani organisations are adopting these tools fast, within a context shaped by colonial language hierarchies, structural underrepresentation of Urdu in training data, and no community-led infrastructure for cultural knowledge. This thesis examines how LLM integration in Pakistani organisations reproduces or challenges coloniality, and proposes a service design intervention in response.

Employing a qualitative critical case study approach grounded in a critical interpretivist and decolonial paradigm, the study draws on semi-structured interviews with design and marketing professionals in Karachi, supplemented by interactive visual elicitation activities. Data was analysed using thematic analysis and critical discourse analysis, interpreted through the frameworks of Quijano, Mignolo, Couldry and Mejias, Fricker, and Mohamed and colleagues.

The findings reveal four interconnected patterns: LLMs framed simultaneously as neutral tools and culturally loaded infrastructure, a growing dependency on systems that can fail silently, the reproduction of English as professional infrastructure at the expense of Urdu and Pakistan's regional languages, and a pattern this thesis names recognition without transformation, where structural conditions are identified but that identification changes nothing. In response, the thesis proposes an open source, community governed cultural knowledge repository for AI, designed to reduce the invisible overhead Pakistani professionals absorb alone. Decolonising AI infrastructure requires not just better tools but the collective infrastructure that makes acting otherwise possible.

Keywords: decolonisation, large language models, Pakistan, service design, coloniality, case study

Acknowledgements

This thesis would not exist without the people who held me through it.

Thank you to my supervisor, Dr. Kiwoong Nam, for the guidance and patience throughout this process.

To my SDSI family, the people I could think out loud with, argue ideas at, and lean on when the writing felt impossible. You made this feel less lonely than it could have been.

To all those who were part of this thesis, who gave their time and their honesty to this research. You did not have to engage as openly as you did, and I am grateful for every conversation.

To the universities I moved through during this programme. Being placed inside different systems, different ways of knowing, and different conversations taught me to see the system I came from more clearly. That recognition is where this thesis begins.

And to my mother and my siblings, who kept me going in the ways that matter most.

This thesis is dedicated to my mother and my siblings, to Pakistanis inside and around the world, and to everyone working to decolonise the infrastructures we inherited.

List of figures

Figure 1.1 <i>Researcher Positionality Timeline</i>	10
Figure 3.1 <i>Three-level analytical framework for the literature review</i>	25
Figure 3.2 <i>The conceptual framework illustrates how colonial linguistic history and data colonialism converge in the construction of LLMs, which are then adopted by Pakistani organisations</i>	39
Figure 4.1 <i>Research design overview</i>	41
Figure 5.1 <i>Thematic map of findings</i>	67
Figure 5.2 <i>Language circles composite across five participants</i>	77
Figure 6.1 <i>The Cultural Knowledge Commons</i>	80
Figure 6.2 <i>The sustainability loop</i>	89

List of tables

Table 4.1 <i>Research participants</i>	47
---	----

List of abbreviations

- Artificial Intelligence (AI)
- Augmented Reality (AR)
- Critical Discourse Analysis (CDA)
- Chief Executive Officer (CEO)
- Large Language Model (LLM)
- Natural Language Processing (NLP)
- Service Design Strategies and Innovations (SDSI)
- User Experience (UX)

TABLE OF CONTENTS

Abstract	3
Acknowledgements	4
List of figures	5
TABLE OF CONTENTS	6
1. POSITIONALITY STATEMENT	8
1.1 Introduction to the Positionality Statement.....	8
1.2 Historical and Linguistic Context.....	9
1.3 Intersecting Identities.....	10
1.4 Design Education and Knowledge Hierarchies.....	12
1.5 Professional Practice and Linguistic Tensions.....	13
1.6 Encounter with Technology and Epistemic Asymmetries.....	14
1.7 International Academic Context and Shifting Perspectives.....	15
1.8 Reflexivity and Research Perspective.....	16
2. INTRODUCTION	18
2.1 Situating the Problem.....	18
2.2 Motivation for the Study.....	20
2.3 Research Goals and Questions.....	21
2.4 Definition of Key Terms and Limitations of the Research.....	22
2.5 Summary of the Methodology.....	23
2.6 Structure of the Thesis.....	24
3. LITERATURE REVIEW	25
3.1 Theoretical Framework: Coloniality, Knowledge, and Power.....	26
3.2 Macro Level: Power, Coloniality, and Language Dominance.....	28
3.2.1 Language Bias and English Hegemony.....	28
3.2.2 Algorithmic Bias and Knowledge Hierarchies.....	29
3.3 Meso Level: Organizational Context and Workplace Transformation.....	31
3.3.1 Adoption Patterns and Organizational Dynamics.....	31
3.3.2 Workplace Transformation: Narrative Versus Reality.....	32
3.3.3 Productivity Narratives and Economic Discourse.....	34
3.4 Micro Level: Individual Use, Trust, and Epistemic Authority.....	35
3.4.1 Trust Formation and the Mechanisms of Reliance.....	35
3.4.2 Epistemic Authority and Perceived Neutrality.....	36
3.4.3 Language Mediation and Linguistic Accommodation.....	38
3.5 Research Gap and What This Study Does.....	38
4. RESEARCH METHODOLOGY	43
4.1 Research Philosophy: Critical Interpretivist with a Reflexive Decolonial Lens.....	44
4.2 Research Approach: Abductive Reasoning.....	45
4.3 Research Strategy: Critical Case Study.....	46

4.4 Research Site and Participants.....	47
4.5 Data Collection: Semi-Structured Interviews with Interactive Elicitation.....	49
4.6 Data Analysis: Thematic and Critical Discourse Analysis.....	52
4.7 Researcher Positionality.....	54
4.8 Ethical Considerations.....	55
4.9 Limitations.....	56
5. FINDINGS.....	57
5.1 Introduction.....	57
5.2 Participant Profiles.....	59
5.3 Between Neutral Tool and Colonial Infrastructure.....	61
5.4 Navigating Dependency.....	64
5.5 English as Infrastructure.....	68
5.6 Recognition Without Transformation.....	72
5.7 Additional Emerging Themes.....	75
5.8 Chapter Summary.....	77
6. DISCUSSION AND CONCLUSION.....	79
6.1 Introduction.....	79
6.2 Revisiting the Research Question.....	79
6.3 Situating Findings in the Literature.....	81
6.4 The Intervention: The Cultural Knowledge Commons.....	83
6.4.1 Decolonial Justification.....	84
6.4.2 The Service Design Layer.....	86
6.4.3 Why It Would Thrive: Business Case and Sustainability.....	87
6.5 Limitations and Future Research.....	90
6.6 Concluding Remarks.....	91
References.....	93
Appendix 1.....	99
Appendix 2.....	104

1. POSITIONALITY STATEMENT

1.1 Introduction to the Positionality Statement

"Situated knowledges are about communities, not isolated individuals."

(Haraway, 1988, p. 590)

This positionality statement is not a formality; it is foundational to this inquiry. As Haraway (1988) reminds us, all knowledge is produced by someone embedded in a specific position in the world. Since knowledge is always positioned, the researcher must account for their position. Pillow (2003, P. 273) defines self-reflexivity as "critical consciousness through a personal accounting of how the researcher's self-location across gender, race, class, and nationality influences all stages of the research process". Being transparent about one's location doesn't undermine credibility; it establishes it (Bourke, 2014; Holmes, 2020).

As a designer, educator, and researcher of South Asian heritage, I have had the privilege of inhabiting multiple vantage points. Born and raised in Pakistan, educated in an English-medium postcolonial school system, and trained as a communication designer across a decade in advertising and startups, I have lived across worlds that don't always speak to each other. This displacement, both physical and intellectual, has revealed the silhouette of systems I once moved through without question.

This research examines my position within these structures, and reflects on how my thinking has been moulded by the epistemic hierarchies I seek to critique (Pillow, 2003). I trace how my research perspective crystallised: through a colonial education system that conditioned me linguistically, through design education that assumed Western knowledge to be superior, to a stark encounter with epistemic asymmetries in LLMs.

This is not a claim to special authority. It is a reckoning with the fact that researchers are never separate from the systems they study, and this acknowledgement is the foremost responsibility of rigorous qualitative inquiry (Smith, 2012).

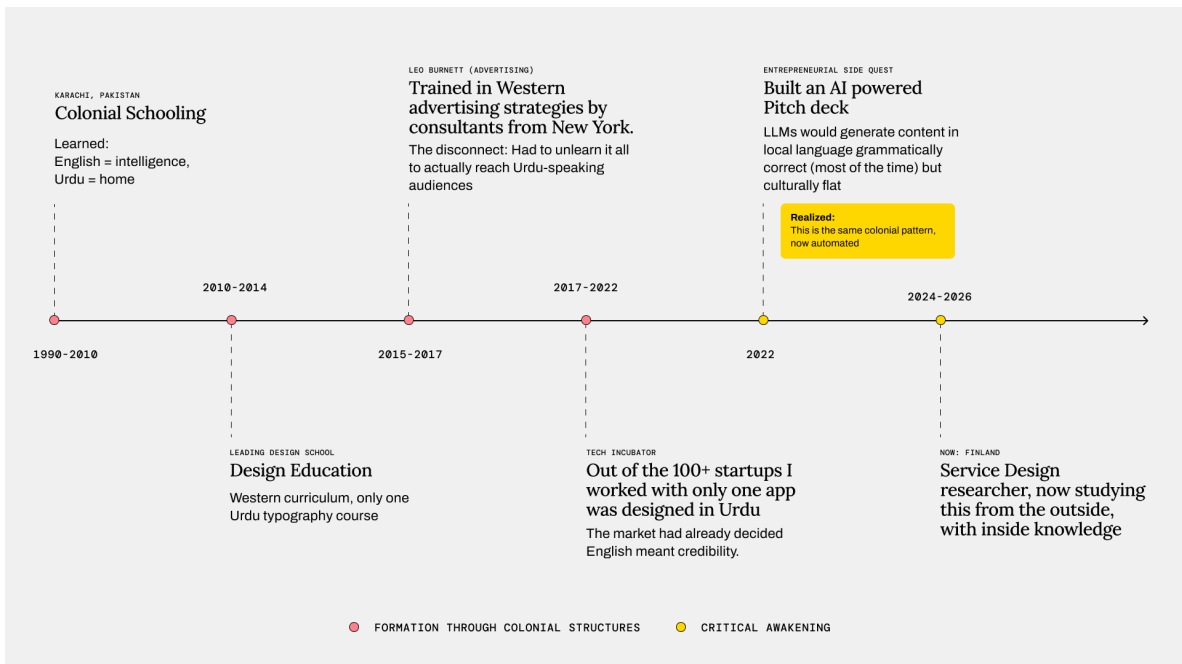


Figure 1.1 *Researcher positionality timeline*

1.2 Historical and Linguistic Context

Chai and English: the two legacies of the British Colony in South Asia. The former became a ritual beverage, the latter a status symbol. Of the two, language proved to be more consequential; becoming the medium through which epistemic authority was concentrated and distributed (Mignolo, 2009; Quijano, 2000).

The 1947 partition of British India was one of the largest forced displacements in recent history, drawn along religious lines under an accelerated colonial withdrawal. I come from a family that crossed into Pakistan with our language as both inheritance and social currency. We were named for what we carried, Urdu-speaking Muhajirs.

But even as Urdu thrived as the language of home, culture, and media, English became the language of power. This was not accidental. It was a direct continuation of hierarchy that British colonial administrators had constructed and postcolonial elites had every reason to maintain. As Pennycook (1994) argues, English inherited the authority of colonial rule, anchoring itself in the status anxieties of the professional class. It required no enforcement because it was internalized as aspiration. A hierarchy had taken root; one that placed

English above Urdu in education, law, government, and by extension, in social life itself. The post-independence generation, who had never lived under colonial rule, became the vehicles through which its logic was reproduced, enrolling their children in English-medium schools because they understood that English proficiency was the gateway to professional and social mobility.

I was one of those children. My school curriculum had been almost entirely anglicized: every subject except Urdu language, literature, and Islamic studies was taught in English. Without recognizing it at the time, I was being formed by what Quijano (2000) terms the coloniality of knowledge. The education system did not present this as a political choice but as the rational one. As Ndlovu-Gatsheni (2017) observes, this is precisely how epistemic structures sustain themselves: not through explicit coercion, but through the slow work of making the construction appear natural.

1.3 Intersecting Identities

I am a Pakistani, brown, Urdu-speaking, able-bodied woman of Muhajir heritage. I grew up in a patriarchal society but within a matriarchal household, a contradiction that taught me early that official structures and lived realities are rarely the same thing. I am also English-medium educated and middle-class, positions that afforded me access to the very hierarchies I now examine critically. These identities do not exist in isolation; they intersect in ways that have shaped my relationship to knowledge, language, and professional authority, and ultimately, the questions this research pursues. Crenshaw (1989) first theorized intersectionality to describe how race, gender, and class operate simultaneously rather than sequentially in shaping lived experience. Following Collins (2000), I understand these intersecting positions not as independent characteristics but as a matrix of relationships, each modifying the meaning of the others, producing specific forms of access, constraint, and visibility that cannot be understood by examining any single dimension alone.

Gender is perhaps the most complex axis of my positioning because it is multi-dimensional; it takes a different shape in each sphere of my life. I grew up in a

society that structurally contests women's professional authority. The fields I entered, advertising, entrepreneurship, and technology, were not just male-dominated in numbers; they defined credibility and expertise in ways that centred male experience as the default.

In contrast, I was raised in a family where female authority, female judgment, and female knowledge were lived realities before they were concepts I could label. Every institutional context beyond the confines of my home was organised around making that competence invisible.

This contradiction produced an early epistemic formation: that institutional structures do not tell the full story. I came to my critical perspective on knowledge hierarchies, in part, through double consciousness: the experience of holding two incompatible truths about authority and knowledge simultaneously. It is a disposition that this research draws on directly, in its refusal to accept systems that present themselves as neutral or universal without interrogation.

Within the European academic environment, I am racialized in ways that are structurally different from my experience. In Pakistan, brownness is the unmarked norm with no social weight. Here, being a brown Pakistani woman carries assumptions: that my relationship to Western theoretical traditions is that of an outsider, that my knowledge contributions should be local rather than theoretical, and that I speak for a region rather than from a research agenda. My engagement with theoretical literature on AI, coloniality, and epistemic justice is not a function of my background, though my background is what led me to ask those questions in the first place.

Growing up between small business ownership and academic life taught me to hold practical and theoretical knowledge in deliberate tension. Small business knowledge is grounded in particular communities, built through direct human exchange, and evaluated by whether it works in a specific setting. Academic knowledge tends toward generalisation, citation, and framework dependency. I carry the generative quality of both: moving between practice and theory, between local specificity and conceptual abstraction, and cautious of academic frameworks that fail to hold the texture of lived experience. This, too, is situated knowledge in Haraway's (1988) sense, and this research draws on it

directly.

My class position, middle-class and English-medium educated, made me a fluent participant in the linguistic and epistemic systems this research examines. English proficiency opened doors that this research walks through: to professional networks, academic literature, and the AI systems at the centre of this inquiry. Every time I designed in English, pitched in English, or reached for Western frameworks to solve Pakistani problems, I was reinforcing the hierarchy that made those choices feel natural. I moved through these systems, worked within them, built on them. I made a career out of reproducing them. That complicity is not something this research stands outside of; I live inside it. It is precisely what makes this research possible.

Disability has shaped none of this in ways I had to navigate. This structural absence is worth mentioning explicitly because it represents a form of access I did not earn and tend not to notice (Dolmage, 2017). Naming it is part of a commitment to reflexivity that does not treat the identities one lacks as invisible by default.

1.4 Design Education and Knowledge Hierarchies

I attended one of Pakistan's leading design schools, where I trained in communication design over four years. The curriculum was rigorous and designed to produce practitioners. It was also, with very few exceptions, almost entirely Western in its reference points. Design history meant European and American design history. Typography meant Latin typography. Colour theory, grid systems, branding frameworks, and visual criticism were all drawn from English-language academic traditions with roots in the Bauhaus, Swiss design, and American modernism. In four years of professional design training, there was exactly one course that addressed Urdu typography: a single course devoted to the visual tradition native to the language spoken by nearly every person in the country where I was being trained to design.

At the time, this did not strike me as unusual. The absence of local visual knowledge from the curriculum was so structurally embedded, so thoroughly naturalised, that it was

invisible to me as a student. I had no language for what was missing. This is how Eurocentric design education sustains itself: not through the explicit denigration of local traditions, but through their systematic omission (Escobar, 2018).

In retrospect, it is uncomfortably obvious. Design education trains practitioners to operate within Western frameworks while marginalising the aesthetic traditions and cultural knowledge of the contexts in which they will actually work. Quijano (2000) calls this the coloniality of knowledge: the persistence of colonial power relations in determining which forms of knowledge get institutionalised, transmitted, and treated as professional standards long after formal colonial rule has ended.

1.5 Professional Practice and Linguistic Tensions

After design school, I entered the field of advertising. Fifteen years of English-medium schooling and four years of Western-inflected design education had prepared me to work within English-language professional environments, but not for the audiences those environments were supposed to serve.

The contradiction was not something I had been educated to notice, let alone bridge. I was designing for people who understood English perfectly well but whose identity, humour, and cultural references lived in Urdu. My education had treated that distinction as beside the point. I learned to navigate it informally, from colleagues, from the street, from paying attention in ways my formal education had never asked me to.

This experience is not unique to Pakistan. Ngũgĩ wa Thiong'o (1986) documents how postcolonial professional education systems routinely produce practitioners who are fluent in the language of their formation but estranged from the communities they are trained to serve. The tension between professional formation and communicative reality I encountered in advertising was a specific instance of a general condition: the mismatch between knowledge systems that claim universality and the particular contexts in which that claim is tested.

1.6 Encounter with Technology and Epistemic Asymmetries

Over time, my work expanded into the startup ecosystem. I spent several years in a technology incubator, an environment that connected me to global entrepreneurial networks and brought me into direct contact with AI tools and workflows. The incubator's broader mission was the dissemination of technology and business knowledge across diverse communities of founders and makers; an aim that carried its own unexamined assumptions about whose knowledge was being disseminated, and from where.

In this context, I co-founded a startup with a software developer, exploring how AI could assist founders in generating pitch decks tailored to different audiences. Working directly with large language models was instructive. I found them capable and often impressive; they could synthesize information, structure arguments, and generate content at a speed and scale no individual could match. But they were also oddly distant from the specific contexts in which I was using them. When the queries required knowledge embedded in specific contexts, local market dynamics, Pakistani professional norms, the particular texture of Pakistan's entrepreneurial ecosystem, the responses arrived confident and complete, and somehow beside the point. They often flattened these specificities into generic global frameworks. Local context was acknowledged, then dissolved into a kind of universal present tense that reflected no particular place.

This was not a failure in any conventional technical sense. The systems did exactly what their training data made possible. The training data these systems draw on is overwhelmingly sourced from the Global North and from English-language contexts, meaning certain cultural frameworks, professional norms, and ways of framing problems are vastly over-represented while others are structurally marginalised (Bender et al., 2021; Crawford, 2021). The epistemological assumptions embedded in that data are not supplementary to how the systems function; they are constitutive of it. When a system trained predominantly on Western sources encounters a query about a Pakistani startup ecosystem, it has no mechanism for recognising what falls outside its epistemic frame. The result is local complexity lost in translation and returned as distorted generalisation.

This experience connects directly to critical scholarship examining how machine learning systems encode and reproduce the social and epistemic inequalities embedded in the data

on which they are trained (Benjamin, 2019; Crawford, 2021; Noble, 2018). Observing these dynamics, not as a theorist but as a practitioner in the specific context of a Pakistani startup ecosystem, contributed to the questions that animate this research. It also compelled a more honest accounting of my own position: I had been, simultaneously, a beneficiary of the English-language infrastructures that gave me access to these tools, and someone whose local knowledge those same tools routinely failed to reflect.

1.7 International Academic Context and Shifting Perspectives

My current position as a student within an international postgraduate programme has amplified and complicated these reflections. Studying alongside peers from diverse regional and cultural backgrounds, and conducting research from a physical and institutional context outside Pakistan, has unveiled patterns and assumptions that proximity had previously rendered invisible.

One observation has been particularly noteworthy: artificial intelligence in Pakistan is typically framed through a register of exigence, positioning its adoption as both an immediate necessity and an unavoidable progression. The dominant discourse often presents a binary: adapt or become redundant. This framing is itself shaped by Pakistan's position within global technological infrastructures: a position of relative dependence on systems developed elsewhere and adopted under conditions not of one's own making.

In some European academic contexts, the discourse around AI integration appears to involve considerably more space for regulatory consideration, ethical reflection, and systemic critique. This contrast in register reflects different structural relationships to technological production, the difference between being a centre of a global system and operating at its periphery (Grosfoguel, 2007; Wallerstein, 1974). Both discourses carry their own limitations and blind spots. But the contrast is real, analytically significant, and directly relevant to the questions this research pursues.

Being outside Pakistan also created space for a discomfiting realization: for much of my professional life, I had willingly been a participant in the very epistemic systems I was

beginning to critique. The schooling that anglicized my education, the design curriculum that erased local visual traditions, the professional environments that oriented me toward global frameworks: I had moved through all of these as someone who had the privilege of benefiting from their hierarchies, even when being shaped by their silences. This is not grounds for self-recrimination, but for honest acknowledgment and the kind of uncomfortable reflexivity that rigorous qualitative research requires (Gani & Khan, 2024; Pillow, 2003).

1.8 Reflexivity and Research Perspective

As a Pakistani researcher situated within a European academic environment while examining global AI technologies, I occupy a position that crosses and troubles conventional Global North and Global South distinctions. I have been shaped by multiple knowledge systems: colonial and postcolonial, local and global, practical and theoretical. Recognizing this position is not a disclaimer; it is a conscious methodological choice. This research is grounded in the concept of situated knowledge: the recognition that all knowledge, including the knowledge embedded in AI systems, emerges from specific cultural and political contexts rather than from a neutral, universal standpoint (Haraway, 1988).

Rather than treating technological systems as neutral artefacts that organizations simply adopt or reject, this research takes a critical interpretivist stance, recognizing that technologies are embedded within broader social, cultural, and political structures (Saunders et al., 2019). From this perspective, the integration of large language models into organizational practice is understood as a socially constructed phenomenon, shaped by institutional norms, knowledge hierarchies, historical inheritances, and power relations that long predate the tools themselves.

I am cognizant that my critical engagement with questions around coloniality and knowledge systems will shape how I interpret what participants tell me. That is not something that can be eliminated. It can only be honestly acknowledged and kept in check. In practice this meant documenting analytical decisions, staying attentive to multiple and

competing interpretations of interview data, and allowing participants' own framings to lead the analysis rather than imposing my theoretical lens onto their experience. The most concrete step I took was to keep theoretical language out of the interview questions entirely. The absence of terms such as “coloniality” or “epistemic hierarchy” created space for participants' own framings to emerge from their experience rather than be introduced by mine. If they described LLMs as neutral tools, that was their genuine orientation. If they described unease about English dominance or cultural assumptions, that came from them.

This last point carries ethical weight beyond methodology. As Smith (2012) argues, research that engages with communities whose knowledge has historically been marginalized carries particular ethical responsibilities, responsibilities that do not diminish simply because the research is conducted within an academic context.

This study is designed with that responsibility in mind: to represent participants' experiences accurately and with care, and to remain open to the possibility that their engagements with AI technologies may be more complex, more ambivalent, or more generative than any theoretical framework, including mine, anticipates.

What distinguishes rigorous inquiry from unreflective observation is not the absence of positionality but its honest acknowledgment. This chapter has attempted to provide that.

The next chapter introduces the research problem, situates it within the broader scholarly conversation about AI and global knowledge systems, and outlines the scope, questions, and design of the study.

2. INTRODUCTION

2.1 Situating the Problem

In 2025, the Government of Pakistan published a National AI Policy across six strategic pillars: innovation, human capital, ethics, sectoral adoption, infrastructure, and international collaboration. The policy reads as a commitment to adoption. It offers no reckoning with what is being adopted.

The large language models entering Pakistani organizational life were built in the United States, trained on data drawn overwhelmingly from English-language sources, and benchmarked by research communities for whom Urdu, Pakistan's national language, has long been an afterthought. These are the conditions of their construction. Arif, Azeemi, Raza, and Athar (2024) evaluated a range of LLMs across Urdu language tasks and found consistent performance gaps relative to English, even in models marketed as multilingual. Tahir, Shams, Fiaz, Adeeba, and Hussain (2025) confirm through comprehensive benchmarking that Urdu performance remains substantially below English baselines across all currently available models. The gap is not a vocabulary problem. It reflects the systematic absence of Urdu-language knowledge from the corpora these systems were trained on.

This asymmetry has structural roots. Bender, Gebru, McMillan-Major, and Shmitchell (2021) show that the Internet, the primary training data source for most LLMs, is not a neutral archive. It is shaped by who has had access to writing, publishing, and digital infrastructure. Crawford (2021) maps AI as a material infrastructure whose costs are concentrated in the Global South and whose value flows to corporations in a handful of jurisdictions. Couldry and Mejias (2019) call this data colonialism: like colonial powers extracting raw materials from distant territories, tech platforms extract human experience as data. Benjamin (2019) shows how technical systems amplify social hierarchies not through explicit intent but through the unexamined assumptions embedded in their construction.

Scholars in the decolonial tradition place these patterns within a longer history. Quijano (2000) argues that colonial rule continued well after political independence. The racial, economic, and epistemic hierarchies it established persist, continuing to determine whose knowledge gets treated as universal and whose gets dismissed as local. Mignolo (2009) extends this to knowledge itself, arguing that what circulates as universal is in fact particular, produced in specific places by specific communities under specific historical conditions.

Pakistan carries this history concretely. British colonial administration made English the language of institutional authority in the subcontinent, and postcolonial elites maintained that hierarchy because it served them. Rahman (1996) documents how this division was not incidental to colonial rule but central to it. Urdu became the language of home and public life; English remained the language of formal education, government, and professional credentialing. When Pakistani organizations adopt LLMs that perform significantly better in English than in Urdu, they are re-encountering a structural pattern that predates the technology by centuries.

Pakistan's National AI Policy prioritizes the urgency of adoption with no apparent awareness of these shortcomings (Government of Pakistan, 2025). The Digital Rights Foundation (2025) and Islam (2025) independently note that the policy's ambitions are not supported by the analytical or infrastructural frameworks needed to act on them. The organizational layer, where AI adoption actually happens, remains unexamined in both the policy and the existing scholarship. That is where this research starts.

2.2 Motivation for the Study

This research pursuit came about through lived experience. I worked in Pakistan's startup sector, co-founded a venture that used LLMs to help founders build pitch decks and map investors, and watched what these systems did when Pakistani contexts entered the frame. They performed fluently on generic professional tasks. They flattened anything specific: local market dynamics, cultural professional norms, the particular texture of Pakistan's entrepreneurial world.

I lacked a theoretical framework for what I was observing at the time. Though, it was self-evident that the tools I was using were built for a different user, in a different context, with a different set of questions treated as the default. That instinct is what eventually led me to take a deeper dive..

The existing literature gave me frameworks to name what I observed. This research contributes what it could not: organizational-level evidence from Pakistan. The macro-level critique of AI as colonial infrastructure is well developed (Benjamin, 2019; Couldry & Mejias, 2019; Crawford, 2021; Mignolo, 2009). Arif et al. (2024) and Tahir et al. (2025) have furthered the literature by documenting the performance gaps between Urdu and English across currently available models.. What neither addresses is the organizational layer: how workers and managers in Pakistani organizations actually encounter, interpret, and negotiate these systems day to day. That is the gap this study closes.

The motivation is also about timing. Pakistan is at a specific moment in its relationship to AI: an official policy has been published, adoption is being encouraged across sectors, the technology workforce is growing. The decisions organizations are making now will shape the epistemic conditions of Pakistani professional life for years. Research that examines these decisions as they are being made is both timely and necessary.

There is also a methodological motivation. Smith (2012) is clear that research engaging with communities whose knowledge has historically been marginalized carries responsibilities that do not disappear because the research is conducted within an academic institution. I designed the interview questions without using theoretical terms like “coloniality” or “epistemic hierarchy.” I wanted participants’ own framings to emerge from their experience, not from mine.

2.3 Research Goals and Questions

This study examines how Pakistani organizations integrate large language models into knowledge work, and whether that integration reproduces or challenges colonial patterns of knowledge hierarchy and technological dependence.

The primary research question is:

How are large language models (LLMs) integrated into organizational practices in Pakistan, and to what extent does this integration reproduce or challenge coloniality in knowledge production and technological dependence?

Four sub-questions structured the empirical fieldwork:

1. How do organizations adapt LLM outputs to their local cultural and linguistic contexts?
2. To what extent do organizations perceive LLM outputs as culturally neutral or globally authoritative?
3. What strategies do organizations use to localize or modify LLM-generated knowledge?
4. How do language use and translation practices shape the flow of knowledge when using LLMs?

These questions avoid presupposing that coloniality is present. They treat its reproduction or challenge as an empirical question, not a theoretical assumption. They also map directly onto the three FigJam activities used in fieldwork: prompt adaptation mapping, knowledge origin mapping, and language hierarchy mapping. These activities were designed to surface the tacit dimensions of organizational AI practice that verbal interviews alone may not reach.

2.4 Definition of Key Terms and Limitations of the Research

Several terms carry specific meaning throughout this study.

Large language models (LLMs) refers to AI systems that learn from a collection of vast text datasets to generate, summarize, translate, and process natural language. In this study, LLMs include widely used commercial tools such as ChatGPT, Claude, Gemini, and Microsoft Copilot, alongside domain-specific models used in professional contexts.

Coloniality is used in the sense Quijano (2000) develops; it refers to the power structures that outlast colonial rule. Here, it refers specifically to epistemic coloniality: the hierarchies that determine whose knowledge counts as universal and authoritative, and whose is treated as local, particular, or deficient.

Data colonialism draws on Couldry and Mejias (2019), referring to the extraction of human social experience as data through commercial platforms, in patterns that echo colonial resource appropriation.

Knowledge production refers to the practices through which organizations generate, validate, transmit, and apply knowledge, including the role AI-generated outputs play in that process.

Organizational practice refers to the everyday routines, decisions, norms, and interactions through which organizations function. This study focuses on knowledge-intensive practices in Pakistan's formal professional sector.

My study has clear limits. It is a single-context study of Pakistan, which means findings cannot be directly transferred to other settings. That is a deliberate methodological choice: depth over breadth. I focus on the formal professional sector and do not cover informal, domestic, or rural LLM use, government adoption at the implementation level, or technical AI development within Pakistan. The study relies on self-reported experience; I did not independently evaluate LLM outputs or technical performance. My own positionality, as a Pakistani researcher in a European academic institution with prior professional exposure to these tools, shapes what I can see and what I might miss. I addressed this in Chapter 1 and return to it in Chapter 4.

2.5 Summary of the Methodology

My research takes the form of a qualitative single-context case study with Pakistan as the case. The research paradigm is critical-interpretivist with explicit decolonial commitments. Critical interpretivism treats social reality as constructed through human meaning-making, shaped by historical and structural forces, and accessible through careful engagement with how people understand their own experience (Saunders, Lewis, & Thornhill, 2019). The decolonial commitment adds a specific orientation: to attend to how colonial power relations structure meaning-making in the contexts I am studying, and to hold open the possibility that participants' own frameworks may be analytically as important as mine.

The single-context case study design reflects the research questions. My aim is to understand how AI adoption operates within one specific configuration of colonial language history, postcolonial professional formation, and contemporary technological dependency. That requires depth over breadth.

Data for this research was primarily generated by semi-structured interviews. I supplemented them with FigJam interactive activities structured around three dimensions of LLM practice: prompt adaptation, knowledge origin mapping, and language hierarchy mapping. These activities were designed to make visible what people do habitually but rarely articulate out loud, the tacit assumptions that shape how they use AI tools without necessarily being aware of them.

I analyzed the data using thematic analysis combined with critical discourse analysis. Thematic analysis generated patterns grounded in participants' own framings, working inductively from the data before moving toward theoretical interpretation. Critical discourse analysis then examined the language participants used to describe their AI practices, looking for what it revealed about underlying assumptions regarding knowledge authority, linguistic hierarchy, and technological dependence.

The full methodology is in Chapter 4, including the ethical framework: informed consent, anonymity, the management of researcher-participant power dynamics, and the particular

responsibilities that come with research engaging communities whose knowledge has historically been marginalized.

2.6 Structure of the Thesis

This thesis is organised into six chapters.

Chapter 1 (Positionality Statement) traces how I came to this research: through a colonial education system in Pakistan, design training that reproduced Western frameworks as universal, direct professional encounter with LLMs in Pakistan's startup sector, and the critical distance that came from stepping outside that context.

Chapter 2 (Introduction) is this chapter. It introduces the research problem, states the research questions, defines key terms, identifies the study's limitations, and summarises the methodology and thesis structure.

Chapter 3 (Literature Review) develops the theoretical and empirical foundation of the study. Organized through a micro/meso/macro framework, it moves from the global epistemic infrastructure of AI, to organizational and linguistic dynamics in Pakistan, to the empirical evidence on LLM performance in Urdu. It concludes with the specific gap this study addresses.

Chapter 4 (Methodology) presents the full research design: paradigm, case study approach, data generation methods, analytical strategy, and ethical framework.

Chapter 5 (Findings) presents the analysis, organized thematically around the patterns that emerged in relation to the research questions.

Chapter 6 (Discussion and Conclusion) interprets the findings against the theoretical framework, identifies the study's contribution to the existing literature, reflects on its limitations, and points toward further research.

3. LITERATURE REVIEW

Large language models entered organizational life faster than most technologies before them. Within two years of ChatGPT's public release, organizations began using these systems for customer service, code generation, and document summarization. The uptake has been both rapid and widespread.

But there is a more complex story underneath. These systems carry particular ways of knowing. They encode specific cultural values. They privilege certain languages and marginalize others. The question this research presses on is how these dynamics unfold in organizational practice. Whereas much existing research examines what goes into these systems, this study asks what happens when they are put to work.

This literature review moves through three levels. At the macro level, I look at how LLMs participate in colonial patterns of knowledge extraction and linguistic dominance. At the meso level, I examine organizational adoption and workplace transformation. At the micro level, I explore how individuals form trust with LLMs and negotiate epistemic authority.

This structure matters because existing research tends to keep these levels separate. Macro-level critiques of AI colonialism rarely connect to everyday practice. Micro-level studies of trust seldom engage with broader power dynamics. My goal here is to map the terrain across all three levels and identify where the gaps are.

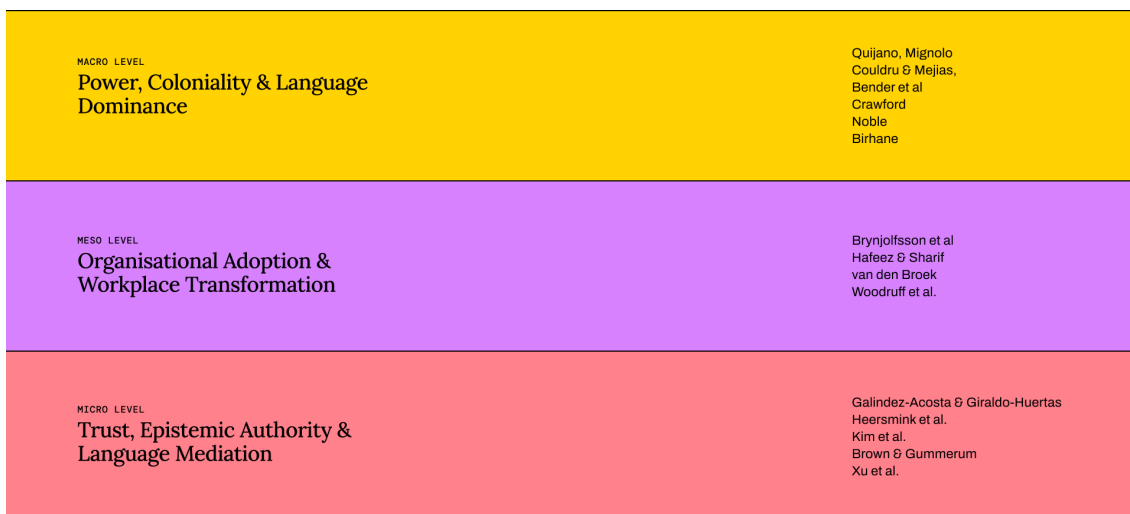


Figure 3.1 *Three-level analytical framework for the literature review*

3.1 Theoretical Framework: Coloniality, Knowledge, and Power

The theoretical foundation starts with Anibal Quijano's work on the colonial matrix of power. Quijano argues that colonialism did not end with political independence. Colonial structures persist through systems that organize knowledge, labor, and authority according to hierarchies established during colonial rule (Quijano, 2000, as discussed in Muldoon & Wu, 2023).

What matters for this research is how Quijano shows that knowledge production itself carries colonial marks. Certain ways of knowing get positioned as universal and rational. Others get dismissed as local or traditional. This is not accidental. It is how power works.

Walter D. Mignolo takes this further with the concept of epistemic coloniality. The problem is not simply that Western knowledge dominates. The problem is that this dominance gets presented as natural and inevitable (Mignolo, 2011, as discussed in Birhane, 2020). Western epistemology becomes the default. Other knowledge systems become invisible or illegitimate. Mignolo calls this epistemic violence. Communities get cut off from their own ways of knowing and forced to understand themselves through categories imposed from outside.

Miranda Fricker's work on epistemic injustice offers another angle. Fricker identifies two forms of injustice in knowledge practices (Fricker, 2007). Testimonial injustice occurs when a testimony gets less credibility because of preconceptions about the speaker's identity or social position. Hermeneutical injustice occurs when one lacks the meaning-making tools to make sense of their own experience. This occurs because their community's ways of understanding have been systematically excluded from collective interpretive resources.

Both forms matter for thinking about LLMs. These systems may perpetuate testimonial injustice by systematically undervaluing knowledge from certain communities. They may create hermeneutical injustice by making Western interpretive frameworks appear universal while marginalizing other ways of making meaning.

Nick Couldry and Ulises Mejias apply these concepts to digital technologies. They introduce data colonialism to describe how tech companies mine and appropriate human experience for profit (Couldry & Mejias, 2019). This extraction mirrors historical colonialism. It takes resources from one context and processes them for value creation elsewhere. It creates dependency. It objectifies people as sources of data rather than as participants in knowledge creation.

Importantly, Couldry and Mejias show that data colonialism is not just economic. It shapes social relations. It concentrates power in the hands of those who control data infrastructure.

Muldoon and Wu (2023) bring these concepts directly to AI. Their argument is that artificial intelligence operates within the four dimensional colonial matrix of power (Muldoon & Wu, 2023). AI development reproduces colonial patterns through several mechanisms. Training data extraction from diverse populations without reciprocal benefit parallels resource extraction in historical colonialism. The concentration of AI development in Western tech companies recreates center-periphery dynamics. The positioning of AI-generated knowledge as objective and universal echoes colonial narratives about Western rationality.

This theoretical foundation shapes how I approach LLM integration in organizations. I treat these systems as carriers of particular epistemologies and power relations. The question becomes how these colonial patterns show up in everyday organizational practice, and whether users recognize and negotiate these dynamics or just accept them as normal.

3.2 Macro Level: Power, Coloniality, and Language Dominance

3.2.1 Language Bias and English Hegemony

The numbers are stark. English comprises about 92 percent of ChatGPT's training data (Xu et al., 2025). Other languages are dramatically underrepresented. This choice about data availability reflects and fortifies global power hierarchies where English functions as the lingua franca of business, technology, and production.

The consequences go beyond just representation. Through studying cultural bias in LLMs across 107 countries, Tao and his colleagues (Tao et al., 2024) found that even when you prompt these models in non-English languages, they consistently exhibit Western values. Anglo-American values to be precise. The models align rather closely with perspectives from the United States and Western Europe than with perspectives from other regions. This is agnostic of the language used for prompting.

This makes it evident that the problem runs deeper than language choice. Western epistemologies are foundational to how these systems process and synthesize information.

The implications for low-resource languages are severe. DeWitt Prat and colleagues examined LLM deployment in African contexts and describe linguistic extractivism (DeWitt Prat et al., 2024). Language data from African communities gets incorporated into training datasets. But these same communities see little benefit from the resulting systems. The models remain optimized for English and other high-resource languages. When they do process African languages, they often perform poorly or misrepresent cultural contexts. This creates a situation where communities have supposedly contributed to AI development but remain underserved users of AI technologies.

Giunchiglia and colleagues frame this as epistemic injustice, connecting linguistic bias back to Fricker's framework (Giunchiglia et al., 2024). They argue that LLMs perpetuate both testimonial injustice, by devaluing knowledge expressed in non-dominant languages, and hermeneutical injustice, by making it harder for speakers of underserved languages to articulate their experiences using AI systems. The models present Western interpretive frameworks as universal, making alternative ways of understanding the world harder to express or even conceptualize.

The "curse of multilinguality" captures another dimension (Xu et al., 2025). As researchers try to make LLMs more multilingual, they face a trade-off. Adding more languages means spreading limited model capacity across more linguistic contexts. This often results in worse performance for all languages except English. Consequently, users of non-anglophone languages face pressure to switch to English to get better results. This

linguistic accommodation is invisible labor that falls disproportionately on users from the Global South.

3.2.2 Algorithmic Bias and Knowledge Hierarchies

Safiya Noble's work on search engine algorithms shows how these systems amplify existing societal inequalities (Noble, 2018). Search results appear neutral and objective. On closer examination, they reinforce stereotypes about race, gender, and culture. Her analysis of Google search results demonstrates systematic bias in how information about marginalized communities gets framed.

This matters for LLMs because they are trained in part on web data that already encodes these biases. The models ingest and reproduce patterns of marginalization.

Kate Crawford outlines the material infrastructure of AI (Crawford, 2021). She follows the supply chains, labor conditions, and environmental costs. Her research illustrates that AI systems depend on extractive practices at several levels. Rare earth mineral mining, often in the Global South, provides materials for hardware. Content moderation labor, frequently outsourced to countries like Kenya and the Philippines, involves traumatic exposure to disturbing content for low wages. Data annotation work relies on platform labor that resembles digital piecework. Crawford argues that these material realities are constitutive of how AI systems operate.

Bender and colleagues coined the phrase "stochastic parrots" to critique how LLMs handle knowledge (Bender et al., 2021). They argue that these models do not understand meaning. They generate statistically likely sequences of text based on patterns in training data. This has important implications for epistemic authority. LLMs can produce fluent and often eloquent responses about topics where they have little to no understanding. Unsurprisingly, they also reproduce biases from training data without the capacity for critical reflection on what they generate. They present culturally specific knowledge as universal because they lack awareness of their own situatedness. Biases from training data pass through intact. Fluency should not be mistaken for knowledge or understanding.

Abeba Birhane extends these critiques specifically to African contexts (Birhane, 2020). She describes algorithmic colonization as a process where solutions developed within the Global North infrastructures get deployed in Africa without attention to local contexts, needs, or knowledge systems. This results in poor performance, misrepresentation of local realities, and the marginalization of African technological development.

Birhane argues that algorithmic colonization operates through similar logics as historical colonialism. It extracts value from African contexts while positioning African people as passive recipients of technology rather than as creators and innovators.

Together, these macro-level analyses show that LLM training data is prejudiced archive of human knowledge. It reflects extractive relationships, language hierarchies, and cultural assumptions that are directly relevant to how these systems perform and what they reproduce when deployed in contexts like Pakistan.

3.3 Meso Level: Organizational Context and Workplace Transformation

3.3.1 Adoption Patterns and Organizational Dynamics

Data shows that 78 to 88 percent of organizations now use generative AI in at least one business function (McKinsey, 2025).

But this statistic obscures what is actually happening in reality. Pilot programs are common. However, the vast majority of pilot programmes do not survive contact with organisational reality. Multiple studies note that only about 5 percent of generative AI implementations successfully scale (MIT research cited in industry reports). Adoption doesn't necessarily translate to success. The path from experimentation to embedded practice remains challenging.

The dominant narrative centers on productivity and efficiency. Brynjolfsson and colleagues

conducted one of the first large-scale empirical studies of generative AI deployment in the workplace (Brynjolfsson et al., 2025). They examined a customer service organization that introduced an AI assistant to help agents respond to inquiries. The results showed a 15 percent increase in productivity, measured by issues resolved per hour.

The gains were skewed in distribution. Novice workers and those with lower initial skill levels saw the largest improvements. Highly skilled workers saw minimal benefit. The authors interpret this as evidence that the AI system captured tacit knowledge from high performers and made it accessible to others.

This suggests LLMs may have a leveling effect on skill distributions. The systems compress the gap between high and low performers by helping less experienced workers perform at levels that previously required more expertise. Whether this represents skill development or simply performance enhancement through technological assistance remains contested. Some argue this could democratize expertise. Others worry about deskilling, where workers become dependent on AI assistance and miss out on the development of underlying competencies.

Research on organizational contexts in Pakistan provides regional specificity. Hafeez and Sharif examined digital transformation in Pakistan's telecom and manufacturing sectors, focusing on organizational agility in mediating AI adoption (Hafeez & Sharif, 2025). Their study, based on surveys of 510 managerial-level employees, found that organizational agility plays a significant mediating role. Organizations that could quickly reconfigure resources, processes, and strategies were more successful at integrating AI technologies.

This suggests adoption depends on more than just technological capabilities; it also heavily relies on organizational characteristics like flexibility, learning culture, and adaptive capacity.

Studies of Pakistani small and medium enterprises identify barriers specific to developing economy contexts. These include limited financial resources for technology investment, inadequate digital infrastructure, and gaps in technical expertise. Cultural factors compound these material constraints. Decision-making authority tends to be concentrated

in organizational hierarchies, bureaucratizing the adoption process. Technologies perceived as risky or unproven face significant institutional reluctance.

3.3.2 Workplace Transformation: Narrative Versus Reality

The discourse around AI and work tends toward extremes. One narrative presents AI as revolutionary technology that will revolutionize knowledge work, automating complex tasks and enabling unprecedented productivity. The other raises alarm about job displacement and the erosion of human expertise. Recent research suggests both narratives overstate the case.

A participatory study involving knowledge workers from seven industries provides important nuance (Woodruff et al., 2024). The researchers conducted workshops with 54 participants across different sectors including journalism, education, law, and software development. They asked participants how they expected generative AI to affect their industries. The dominant response was surprisingly modest. Workers largely envisioned AI as a tool for handling menial tasks under human review. They did not anticipate the transformative changes often projected in media and academic literature. Instead, they described AI as something that might help with routine work but would remain subordinate to human judgment and expertise.

This gap between transformative narratives and worker expectations is revealing. People closest to the work are skeptical of revolutionary claims.

The study participants did identify four social forces they expected AI might amplify. Deskilling, where reliance on AI could erode professional capabilities over time. Dehumanization, where the emphasis on efficiency might undervalue human relationships and judgment. Disconnection, where mediation through AI might reduce direct engagement with work. Disinformation, where the ease of generating plausible-sounding content might make distinguishing truth from fabrication harder.

These concerns are less about replacement and more about degradation of work quality and meaning.

Recent research on AI in the workplace draws attention to hidden forms of labor that productivity narratives often overlook (van den Broek, 2025). The author identifies three types of work that AI systems require from humans.

Data work involves preparing, cleaning, and maintaining the datasets that AI systems depend on. Knowledge work refers to the effort required to align AI outputs with domain expertise and organizational knowledge. Values work encompasses the ongoing negotiation of what counts as appropriate or ethical AI use in specific contexts.

All three forms of labor tend to be invisible in accounts that focus only on efficiency gains.

This complicates the productivity story. Yes, AI systems may enable faster completion of certain tasks. But they also create new forms of work that organizations must perform. Someone has to verify outputs for accuracy. Someone has to translate generic AI responses into context-specific language. Someone has to decide when AI suggestions are appropriate and when they are not.

This work often falls to the same knowledge workers supposedly being made more productive. Efficiency gains may be offset by new forms of epistemic labor.

Research on AI in knowledge work raises additional questions about productivity metrics themselves (van den Broek, 2025; Woodruff et al., 2024). Studies suggest that efficiency measures appropriate for routine, repeatable tasks become counterproductive when applied to knowledge work requiring creativity, judgment, and contextual understanding. Knowledge workers are motivated by autonomy, competence, and purpose rather than by narrow productivity metrics. AI systems focused only on speed or output volume may undermine these intrinsic motivations, potentially harming long-term performance even as they improve short-term metrics.

3.3.3 Productivity Narratives and Economic Discourse

The broader economic discourse leans heavily toward optimism. Policy institutions and consulting firms predict AI will drive significant productivity gains and economic growth. The Brookings Institution calls generative AI a new general-purpose technology comparable to electricity or the internet (Baily et al., 2025). McKinsey describes "superagency," where individuals empowered by AI can dramatically increase their creative output and impact (McKinsey, 2025).

These projections are worthy of scrutiny. They stand on assumptions about how productivity gains will be measured, distributed, and sustained. Previous waves of technological change took years to materialize and were not evenly distributed. The key question is whether these adoptions will produce equitably shared gains or concentrate profits among those already positioned to benefit.

3.4 Micro Level: Individual Use, Trust, and Epistemic Authority

3.4.1 Trust Formation and the Mechanisms of Reliance

How individuals form relationships with LLMs involves complex dynamics. Recent research challenges the assumption that people trust AI primarily because of technical capabilities. Galindez-Acosta and Giraldo-Huertas introduce "deferred trust" to describe how trust in AI can emerge from distrust in human sources (Galindez-Acosta & Giraldo-Huertas, 2025).

Their study presented participants with decision-making scenarios across factual, emotional, and moral domains. Participants chose between AI agents, human adults, peers, or religious authorities for guidance. The results showed that lower trust in human agents consistently predicted higher AI selection.

This suggests that AI trust is relational rather than purely technology-driven. People may turn to AI less because they have high confidence in the technology and more because they

have lost faith in human alternatives. The perceived neutrality of AI becomes appealing precisely when human sources are viewed as biased or unreliable.

The concept of appropriate reliance offers a framework for thinking about optimal trust calibration (Microsoft Research, 2024). The goal is context-appropriate judgment about when to rely on AI and when to verify or override its outputs. Research shows that users struggle to achieve this calibration. They tend toward either over-reliance, often accepting outputs without sufficient verification, or under-reliance, rejecting useful outputs based on general skepticism about AI.

The challenge around appropriate degree of reliance is compounded by the fact that LLMs are good at some tasks and poor at others in ways that are not always obvious to users.

Epistemic vigilance describes the cognitive process people use to evaluate the credibility of information sources. Research on adolescents shows that they exercise selective trust, paying attention to who provides information and whether that person has proven reliable in the past (Brown & Gummerum, 2025).

This vigilance, however, can be undermined by fluency effects. When LLMs produce text that sounds confident, well-structured, and authoritative, users may reduce their critical evaluation even when skepticism would be appropriate. The speed and fluency of LLM outputs can create a certain kind of epistemic pressure. A sense that questioning the output requires justification while accepting it feels natural.

Trust in LLMs also varies by task type. Studies of LLM reliance consistently find that users are more willing to defer to AI outputs for verifiable tasks like coding or summarization than for evaluative tasks where human judgment is harder to bypass (Kim et al., 2024; Heersmink et al., 2024). The difficulty of verification becomes a key factor in how much critical distance users maintain.

Trust appears to develop through repeated familiarity and practice rather than through careful evaluation of system capabilities. This pattern raises concerns about over-reliance becoming normalized over time.

3.4.2 Epistemic Authority and Perceived Neutrality

How LLMs acquire epistemic authority deserves close attention. Research suggests that authority stems in part from linguistic style rather than actual knowledge (Heersmink et al., 2024; Kim et al., 2024). LLMs produce text with markers of expertise, precise vocabulary, structured organization and confident tone. These stylistic features mimic how human experts communicate.

Users may grant authority to LLM outputs based on how they sound rather than on verified accuracy. Fluent writing functions as a kind of epistemic credential.

The speed of LLM text generation adds another dimension. When a system can produce a detailed, well-formatted response in seconds, the velocity itself can feel like evidence of competence. Yet this speed is unrelated to accuracy or understanding. The system is generating statistically likely text based on patterns in training data. It does not reason about the truth of the claims. The velocity creates what has been described as epistemic theater. A performance of knowledge that may or may not correspond to actual expertise.

Research on the phenomenology of LLM interaction explores how people experience these systems. The authors distinguish between epistemic trust, a user's willingness to rely on LLM outputs for knowledge, and trustworthiness, the actual reliability of those outputs (Heersmink et al., 2024). They argue that LLMs can elicit epistemic trust even when they lack trustworthiness.

The human-like quality of LLM language encourages users to relate to these systems as if they were social actors with intentions and expertise. This anthropomorphization, while perhaps inevitable given how the systems are designed, can lead users to attribute understanding and reliability that the systems do not possess.

The question of perceived neutrality follows directly from this. The algorithmic mediation of LLM outputs makes their cultural bias less visible than it would be with an identifiable human author. Users accept outputs as universal truth rather than recognizing them as culturally situated knowledge, precisely because there is no visible author to question.

Some research suggests that transparency about AI limitations may not solve this problem. Studies of uncertainty expressions in LLM outputs show mixed results. When models express uncertainty using phrases like "I'm not sure" or "this may not be accurate," some users calibrate their trust downward appropriately. However, other users interpret these expressions as humility rather than as genuine signals of unreliability, potentially maintaining inappropriate levels of trust (Kim et al., 2024).

The challenge is that epistemic calibration requires users to understand not just that LLMs can be wrong but specifically when and how they are likely to be wrong.

3.4.3 Language Mediation and Linguistic Accommodation

For users working in languages other than English, LLM interaction involves additional complexity. Research on multilingual LLM use identifies several adaptive strategies. Code-switching between languages is common. Users default to English for better results even when they are more comfortable in another language (Xu et al., 2025).

This linguistic accommodation is labor that falls disproportionately on non-English speakers. They must constantly navigate between their preferred language and the language that produces optimal system performance.

Translation also introduces interpretation layers. When users prompt in their local language, receive responses in English or mixed language, and then translate those responses back to their context, meaning can shift at each stage. The user becomes responsible for managing these linguistic transformations and making judgments about what the system meant to say.

This interpretive work is both cognitively demanding and epistemically significant. It gives users agency to modify outputs. But it also places the burden of cultural translation on them rather than on the system.

3.5 Research Gap and What This Study Does

Each level of scholarship reviewed here exists in silos. Macro-level scholarship provides powerful critiques of how LLMs encode colonial relations. Meso-level research documents organizational adoption. Micro-level studies investigate trust formation. But these streams of research rarely speak to each other.

Although macro-level critiques are theoretically sophisticated, they offer limited insight into everyday practice. Their strength lies in identifying structural problems, but they stop short of showing how those problems manifest in actual use.

Do Pakistani users recognise the Western orientation of LLMs? How do they navigate the mismatch between system outputs and their own contexts? When do they defer to LLM knowledge and when do they push back? These questions about lived experience remain largely unexplored and sit at the heart of this research.

The organizational literature focuses heavily on efficiency and productivity, often treating LLMs as neutral tools to be optimized for performance. This instrumental framing makes sense from a managerial perspective. But it may obscure important epistemic and political questions. When organizations adopt LLMs primarily to improve efficiency, they may not attend to questions about whose knowledge these systems privilege or what forms of expertise they marginalize. The productivity focus, while pragmatic, risks reproducing colonial patterns by treating Western-trained LLMs as universal knowledge resources.

Research on Pakistani organizational contexts has begun to document adoption patterns, barriers, and enablers. But this work has focused primarily on technology acceptance and implementation challenges rather than on knowledge practices and epistemic dynamics. Some things are known about why Pakistani organizations do or do not adopt AI systems. Far less is known about how they actually use them, how they interpret outputs, and how they negotiate the cultural specificity of Western-trained models.

The micro-level literature on trust and reliance provides important frameworks but has limitations of scope. Most studies involve Western participants and English language

interaction. The mechanisms identified, such as deferred trust, epistemic vigilance, appropriate reliance may operate differently across cultural contexts. Questions about language mediation, translation labor, and linguistic accommodation remain underexplored, particularly in organizational settings where these practices become routine.

Four specific gaps emerge from this review.

First, we lack empirical studies that connect macro-level critique to micro and meso-level practice. How do colonial dimensions of LLMs manifest in everyday organizational use?

Second, organizational LLM use in Global South contexts remains understudied. Most research focuses on Western organizations or treats developing economy contexts as simply facing adoption barriers rather than as sites of distinct knowledge practices.

Third, the dynamics of epistemic negotiation need more attention. How do users evaluate and potentially contest LLM knowledge claims?

Fourth, language mediation practices in organizational settings deserve investigation. How do multilingual users navigate between languages? Who bears the labor of translation and interpretation?

This research addresses these gaps by examining how individuals and organizations in Pakistan integrate LLMs into their knowledge practices. Pakistan serves as a concrete site of inquiry. As a developing economy with distinct linguistic contexts including Urdu and English, Pakistan offers a setting where questions about language mediation and cultural adaptation come to the foreground. As a postcolonial nation navigating ongoing relationships with Western technology providers, Pakistan provides a context where questions about epistemic authority and technological dependency have concrete stakes.

The research asks how LLM integration reproduces or potentially challenges colonial patterns in knowledge production.

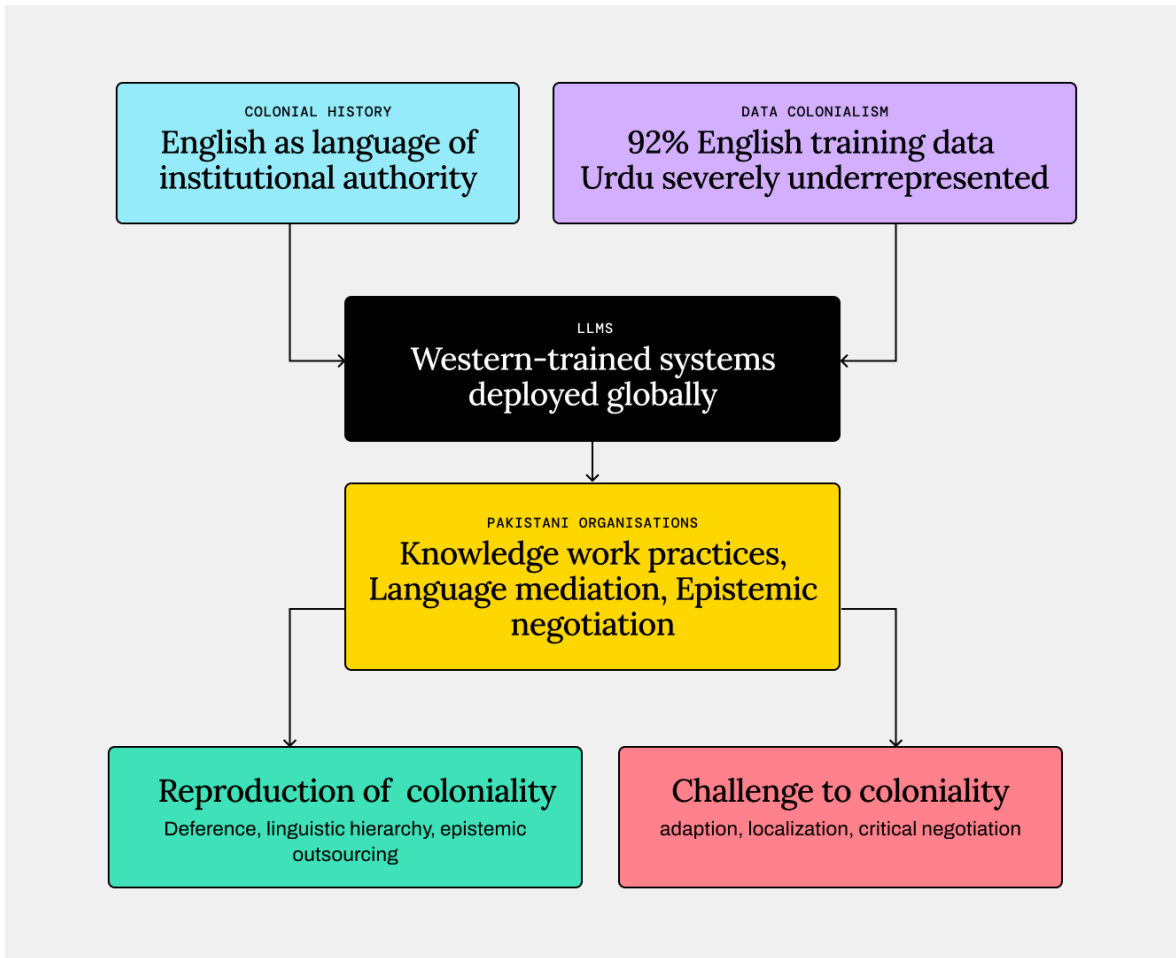


Figure 3.2 *The conceptual framework illustrates how colonial linguistic history and data colonialism converge in the construction of LLMs, which are then adopted by Pakistani organisations.*

The research questions guiding this study emerge from the literature reviewed here. How do organizations adapt LLM outputs to local cultural and linguistic contexts? To what extent do organizations perceive LLM outputs as culturally neutral or globally authoritative? What strategies do organizations use to localize or modify LLM-generated knowledge? How do language use and translation practices shape knowledge flow when using LLMs?

These questions aim to connect macro-level concerns about coloniality and power to micro and meso-level practices of organizational knowledge work.

My aim with this research is not to determine whether LLMs are good or bad. It is not to label them colonial or decolonial. The goal is to understand how these systems integrate into organizational practice in specific contexts: how epistemic authority gets negotiated, how local knowledge gets adapted or overridden, and what language practices reveal about technological dependency today.

The literature maps the terrain. This research asks what happens on the ground.

4. RESEARCH METHODOLOGY

Every methodological choice made in the design of this research, from who to include to how to interpret what emerges, is also a choice about what kinds of knowledge matter and who gets to produce them. In a study examining how colonial power dynamics play out in AI adoption, it is the foundational question; one that precedes every other methodological choice.

This chapter builds the research design using Saunders and colleagues' (2019) research onion as an organising framework, moving from philosophical foundations through research approach, strategy, data collection, and analysis.

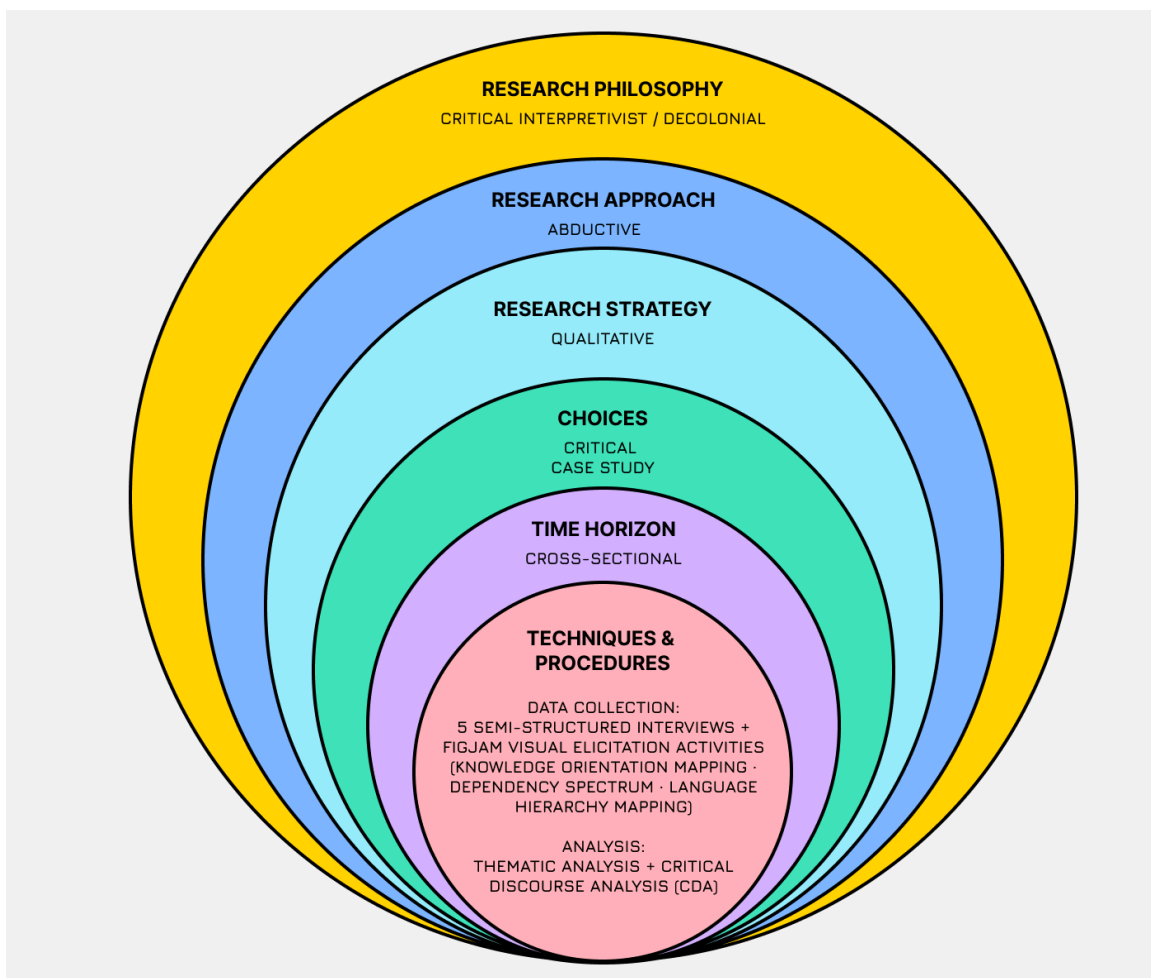


Figure 4.1 Research design overview

Note. Adapted from Saunders et al. (2019)

4.1 Research Philosophy: Critical Interpretivist with a Reflexive Decolonial Lens

The first step is being honest about three underlying assumptions: what counts as reality, how knowledge gets produced, and what role values play. Saunders and colleagues (2019) describe these as ontological, epistemological, and axiological commitments. Mine shape every choice that follows.

Ontologically, I do not believe there is one objective reality about LLM use in Pakistani organisations waiting to be uncovered. Reality in this study is multiple and constructed. A founder using ChatGPT for grant proposals experiences that tool differently from a UX designer generating personas, or a marketer drafting ad copy. These are not different windows onto the same truth. They are genuinely different realities, shaped by professional expertise, linguistic background, and organisational context. This is a constructivist ontology, aligned with the interpretivist tradition, which holds that meaning is socially produced rather than discovered (Creswell, 2013; Schwartz-Shea & Yanow, 2012).

Epistemologically, my commitments are explicitly decolonial. Knowledge about AI in Pakistan, how it functions, what it misrepresents, and what it might do better, should be produced by people who live and work within that context. This isn't about sentimentalising local knowledge or claiming authentic Pakistani way of knowing. It's about power; about who controls the production of knowledge about Pakistan and who doesn't. The systems shaping how AI tools perform in Urdu, how training data gets annotated, and how good outputs get defined are overwhelmingly produced by people and institutions located far from Pakistan. Mignolo (2009) calls the deliberate assertion of marginalised knowledges epistemic disobedience. Smith (2012) argues that research conducted on communities without meaningful participation reproduces extractive colonial relationships. My commitment is to knowledge that actually reflects the experience of Pakistani knowledge workers, not knowledge about them from the outside.

Axiologically, I do not claim neutrality, but I do claim reflexivity. Growing up in Pakistan,

educated partly in English, working in a design field whose frameworks originate largely in Europe and North America, I absorbed assumptions about whose knowledge is legitimate. This research has changed how I see those assumptions. I cannot step outside my own position and observe from nowhere. What I can do is hold my interpretive moves in view, ask whose perspective I am centring at any given moment, and stay open to being corrected by the data and by my participants. This is what Haraway (1988) calls situated knowledge: accountable and partial, rather than falsely objective.

Decolonial theory operates at the core of all three levels, shaping the ontological, epistemological, and axiological commitments from the outset rather than arriving as an interpretive framework applied to data at the end. Ontologically, it affirms that multiple realities exist and that none can be universalised without doing violence to others. Epistemologically, it names the power structures that determine whose knowledge counts. Axiologically, it calls for research that contributes to communities rather than extracting from them.

4.2 Research Approach: Abductive Reasoning

As a designer, I have always worked between observation and adjustment rather than through a fixed sequence. You move between what you are observing, what theory is telling you, and what participants are saying, recalibrating as each informs the other. I was doing this long before I knew the term for it: abduction.

Abduction, as Timmermans and Tavory (2012) describe it, involves moving iteratively between empirical observations and theoretical frameworks. It sits between induction and deduction: using available evidence to generate the most plausible interpretation, then testing that interpretation against new observations. It is well suited to qualitative research that brings existing theory into contact with lived experience, because it allows theory to inform without determining, and data to push back without speaking entirely for itself.

Pure induction would not have worked here. Pretending I entered the field without a

theoretical orientation would not have made this research more rigorous. It would have made it less honest. What abduction gave me instead was permission to let theory and data talk to each other throughout. I entered the field with sensitising concepts: coloniality of knowledge, linguistic hierarchies, epistemic authority. These shaped what I paid attention to. But I stayed open to being surprised. When one participant described AI success as the moment a non-developer could produce working code without a developer's help, that was not something the theory had prepared me for straightforwardly. Going back to the literature on epistemic authority gave me tools to think it through. That back and forth is abduction in practice.

4.3 Research Strategy: Critical Case Study

This study adopts a critical case study strategy. The case is LLM integration by knowledge workers in Pakistani technology startups, enterprise software houses, and marketing organisations. Within that case, the units of analysis are the individuals actively using these tools and making sense of them in daily work. Yin (2018) describes this as an embedded design, where a single bounded case contains multiple units of analysis that together produce a richer picture than any one could on its own. Case study methodology fits because the research asks how and why questions about a phenomenon whose meaning cannot be separated from its context. Whether a Pakistani professional defers to an LLM's English output, switches languages to get better results, or questions the cultural assumptions embedded in a response, these behaviours are produced by a specific context with a specific history. A case study keeps that context in view.

The case is critical rather than simply exploratory. A critical case, in Yin's (2018) terms, is selected because it is particularly well positioned to illuminate a theoretical proposition. Pakistan meets this criterion precisely. With a population of over 258 million, it is the fifth most populous country in the world, yet almost invisible in the systems shaping global AI development. Urdu remains dramatically underrepresented in LLM training data, and as Arif and colleagues (2024) and Tahir and colleagues (2025) show, general-purpose LLMs perform significantly worse on Urdu tasks than specialist models. This gap is structural. Rahman's (1996) work on Pakistan's colonial language history shows that the hierarchy

positioning English above Urdu was built through colonial policy and has been reproduced through every wave of imported knowledge infrastructure since. LLMs are the latest iteration.

The timing matters too. Pakistan released its first National AI Policy in 2025 and LLM adoption across the startup and enterprise sector is accelerating rapidly (Government of Pakistan, 2025). The professionals in this study are integrating these tools in real time, under efficiency pressures that leave little space for critical reflection about what is being reproduced in the process. Studying this intersection of population scale, language underrepresentation, colonial history, and speed of adoption makes Pakistan a site where questions about coloniality and AI are not abstract. They are live.

An earlier version of this research included a comparative element. I moved away from it deliberately. When a Global South context is positioned against a Global North baseline, the analysis tends to treat one as the norm and measure the other against it. That framing reproduces the very asymmetry this research is trying to examine. A single, deeply contextualised case study of Pakistan generates knowledge that stands on its own terms. Smith (2012) describes this as the decolonial imperative in research: to centre the community being studied rather than situating it as a deviation from a presumed universal.

4.4 Research Site and Participants

Participants were selected through purposive sampling, which is appropriate when the goal is analytical depth and you are looking for people who can speak from direct experience of the phenomenon under study (Patton, 2002). I drew on the network I built during several years working in Pakistan's startup and accelerator ecosystem, reaching out through LinkedIn. All five agreed readily, were open and eager to be interviewed, and none requested anonymity. But all participants have been anonymised using participant codes to protect confidentiality.

Expert 2 is CEO of an edtech startup working at the intersection of AI, augmented reality,

and educational technology with a social impact focus. Their perspective matters because they are making strategic decisions about AI adoption in a context where who these systems are designed for directly shapes the product. Expert 1 is Head of Design at a technology incubator, and Expert 3 is Head of User Experience and Design at a multinational enterprise software company, an enterprise software house serving local and international clients. Both work at the level of how systems communicate with users, giving them a particular vantage point on whether AI tools understand Pakistani cultural and linguistic contexts.

Expert 4 is a digital marketing professional with over nine years of experience in social media strategy, content, and paid and organic campaigns at a traditional service enterprise. Expert 5 is an independent marketing and content professional with over eight years of experience working with Pakistani and international brands. Their work is grounded in cultural specificity and audience insight, which makes them well placed to evaluate where LLM-generated content lands and where it does not.

PARTICIPANT	ROLE	SECTOR
Expert 1	CEO	AI & EdTech
Expert 2	Head of Design	Tech & Media
Expert 3	Head of UX & Design	Enterprise Software
Expert 4	Digital Marketing Professional	Traditional Enterprise
Expert 5	Independent Consultant	Marketing & Content

Table 4.1 *Research participants*

Note. All participants are based in Karachi, Pakistan.

All five participants are based in Karachi, which is not incidental. Karachi is Pakistan's commercial capital and the primary hub of its startup and tech ecosystem, where LLM adoption is most visible and most advanced. The sample is intentionally diverse within that context: a social impact startup, an established tech company, an enterprise software house, and two independent marketing professionals. That variation allows the research to examine whether patterns in LLM experience are consistent across sectors or specific to particular types of work.

All five interviews were conducted in English. Participants moved naturally into English without prompting and remained there throughout. This is not a neutral logistical detail. It is an early enactment of the dynamic this research investigates. As Rahman (1996) documents, the positioning of English as the marker of professional credibility in Pakistan is a direct product of colonial policy, reproduced through every institutional structure since independence. These professionals were not making a conscious choice. They were operating in a linguistic environment where professional communication in English is how you are taken seriously. The colonial language hierarchy was present in the data collection itself, before a single interview question had been asked.

4.5 Data Collection: Semi-Structured Interviews with Interactive Elicitation

Primary data came from five semi-structured online interviews, each approximately an hour in duration. Sessions were conducted on Google Meet and recorded after obtaining consent from each participant. Transcription was produced using Slack's voice transcription feature, which I then reviewed and cleaned for accuracy.

This format is well suited to research concerned with how individuals construct meaning from their experiences rather than producing generalisable data. This also leaves the researcher open to directions unanticipated (Brinkmann & Kvale, 2015). The format gave me enough structure to address my research questions consistently across all five participants while leaving room for the conversation to follow where each participant's

experience led. I had follow-up questions ready for each section but used them selectively depending on what participants said, probing more deeply where an answer opened up something theoretically relevant.

The interview guide was developed after completing the literature review, which directly shaped the interview questions. My own experience using LLMs as a Pakistani design professional also informed the guide. I had already noticed, before any interviews, that switching to English produced better outputs, that the tools seemed to assume a certain kind of user, and that cultural references I understood naturally tended to flatten when I tried to work with them through these systems. That experience helped me design questions that could surface similar observations in participants without leading them toward predetermined conclusions.

One of the deliberate choices in the guide's design was the complete absence of theoretical language. Words like coloniality, epistemic injustice, or knowledge hierarchy do not appear anywhere in the questions. This was intentional. I wanted these dynamics to surface from participants' own descriptions of their experience rather than be activated by my framing. If a participant described LLMs as neutral tools, that was their natural orientation. If they described unease about English dominance or cultural assumptions, it came from them. The language of the guide was kept deliberately open so that what emerged reflected participants' realities rather than my theoretical expectations. In practice, this worked as intended. Critical observations about language, authority, and cultural assumptions surfaced organically, particularly as the interviews progressed toward the later sections. By the closing reflection, several participants described the conversation itself as something that had made them think differently about tools they use every day without much scrutiny.

The guide was organised into eight thematic sections covering professional context and organisational positioning, AI adoption and infrastructural orientation, epistemic authority and trust, language and cultural mediation, and critical reflection on structural awareness. Within these sections, I embedded three interactive visual elicitation activities conducted through FigJam, a collaborative digital whiteboard tool. These activities draw on visual and participatory methods central to service design practice, where mapping, sorting, and

visualisation are used to surface tacit knowledge and prompt reflection that verbal questions alone often cannot reach (Stickdorn et al., 2018).

The first activity, Knowledge Orientation Mapping, presented participants with twelve descriptor cards including terms such as neutral tool, western-oriented system, efficiency engine, English-first system, and knowledge authority. Participants sorted these into three categories: strongly represents LLMs in their experience, somewhat represents LLMs, and does not represent LLMs. This activity was designed to surface participants' implicit framings of LLMs without requiring them to engage with theoretical vocabulary. How a participant sorted neutral tool against western-oriented system revealed assumptions that a direct question might have prompted them to rationalise rather than articulate honestly.

The second activity, the Dependency Spectrum, asked participants to place their organisation's LLM usage on a horizontal scale running from locally autonomous to globally dependent, and to explain their placement. This surfaced awareness of infrastructural dependence, including whether participants had considered what would happen if access to these tools were restricted.

The third activity, Language Hierarchy Mapping, asked participants to resize three circles representing English, local languages, and hybrid usage according to their relative dominance in AI use, and to draw arrows showing where translation, adaptation, or modification occurred. All five participants engaged fully with the activities. The visual format produced responses that often went beyond what verbal description alone captured, and several participants added their own descriptors or modified the framing in ways that enriched the analysis.

Each participant was offered a verbatim transcript for review within 24 hours and invited to edit, clarify, or withdraw any part before it was used in analysis. This member checking reflects a commitment to treating participants as active contributors to the research rather than passive sources of data.

4.6 Data Analysis: Thematic and Critical Discourse Analysis

Data analysis proceeded in two overlapping stages, moving between thematic analysis and critical discourse analysis across the full set of interview transcripts and FigJam activity outputs.

The analytical process drew on Braun and Clarke's (2006) thematic analysis framework. Rather than following a linear path, I worked with the six phases spanning initial data engagement through to naming and writing up themes. I approached these iteratively, moving back and forth between stages as understanding deepened. Before beginning to code, I immersed myself in each transcript, reading closely and making initial notes about what felt significant and where the same kinds of concerns seemed to surface across different interviews. That repeated reading matters. Coding before genuine familiarity with the data produces categories that fit the researcher's expectations rather than the data's texture.

Initial coding was inductive. I worked through each transcript line by line, assigning codes to phrases, sentences, and passages that carried meaning relevant to how participants were experiencing and describing LLM use. At this stage I was not working with the theoretical framework. I wanted the initial patterns to come from what participants said, not from what the theory predicted I would find. The codes at this stage were descriptive. Things like: switches to English for better outputs, describes LLM as neutral, flags Urdu-Hindi confusion, reframes dependency as efficiency, names limitation but does not act on it.

The next phase involved grouping codes into candidate themes and then reviewing those groups against the full data set to check whether they held. Some candidate themes collapsed into each other. Others split. I moved back to individual transcripts repeatedly to test whether a proposed theme was present in the data or was a pattern I was projecting

onto it. Four major themes emerged from this process: the positioning of LLMs as simultaneously neutral tools and culturally loaded infrastructure; the nature and management of dependency on LLM systems; the role of English as professional infrastructure and its consequences for Urdu and regional language contexts; and the consistent pattern of identifying structural conditions without acting on them. These became Sections 5.3 through 5.6 of the findings chapter.

Critical discourse analysis ran alongside thematic analysis from the point at which candidate themes began to stabilise. Where thematic analysis asked what patterns were present across the data, critical discourse analysis asked how those patterns were constructed in language and what power relations they enacted or reproduced. Fairclough's (1995) model was chosen for its attention to discourse as something operating on multiple registers at once: the language itself, how it moves between people, and the social world that makes that movement possible. I was looking at lexical choices, the framing and naturalisation of particular assumptions, the way participants moved between positions, and the moments where the structure of an argument revealed what was taken for granted rather than what was being explicitly stated.

This is where the two analytical methods were most productive together. Thematic analysis identified that participants consistently named structural conditions and then moved past them. Critical discourse analysis revealed how that movement happened at the level of language. Expert 1's construction of knowing for a fact that LLMs are not neutral, immediately followed by a reframing that made that non-neutrality professionally irrelevant, was visible as a pattern in thematic analysis and explainable as a discursive move in CDA. Expert 2's trade-off framing, in which cultural loss is absorbed into efficiency gain through the language of rational professional calculation, operates the same way. The combination made it possible to name not just what participants said but how they said it and what that revealed about the conditions they were working within.

The FigJam activity outputs were analysed as a third layer of data alongside the verbal transcripts. Visual outputs from the Knowledge Orientation Mapping, Dependency Spectrum, and Language Hierarchy Mapping activities were interpreted in relation to participants' verbal explanations during each activity, and cross-referenced with relevant moments in the interview transcripts. Where a participant placed a descriptor card in a way that seemed to contradict their verbal account, that tension was treated as analytically significant rather than as an inconsistency to be resolved. The visual activities produced particularly revealing data about tacit assumptions that participants were not explicitly articulating in verbal form.

The decolonial theoretical framework was brought back into the analysis at the interpretation stage. Once the themes had been identified and the critical discourse analysis had mapped the discursive mechanics through which they operated, I read the findings back through the frameworks of Quijano (2000), Mignolo (2011), Couldry and Mejias (2019), Fricker (2007), Mohamed and colleagues (2020), and Escobar (2018) to interpret what the patterns meant at the structural and historical level. This is consistent with the abductive approach described in Section 4.2: theory and data in dialogue throughout, neither determining the other. Themes took shape from within the data before being brought into conversation with existing frameworks.

4.7 Researcher Positionality

My positionality as a Pakistani researcher studying Pakistani organisations is documented in full in Chapter 1. It is the starting point of this research, not an appendix to it. Briefly stated: I am an insider by cultural formation and an outsider by institutional position, conducting research from a European university about a context I grew up in. That productive contradiction shaped every methodological choice documented in this chapter, from the decision to centre Pakistan on its own terms to the design of an interview process that treats participants as experts. I maintained a reflexive journal throughout data collection and analysis to hold my own interpretive moves in view.

4.8 Ethical Considerations

The ethical framework for this research was reviewed and approved by my supervisor, Kiwoong Nam, in accordance with the University of Lapland's research ethics guidelines. Participants were professionals sharing occupational experience. The study involved no deception, no vulnerable groups, and no sensitive personal data. At the start of each interview, informed consent was obtained verbally. I explained the goal of this research, how data would be used, that participation was voluntary, and that the research would appear as a publicly accessible Master's thesis. All five participants chose to be identified by name. I discussed this explicitly with each of them rather than assuming it.

Interview recordings were stored on my personal password-protected device and used solely for transcription. Participants were given their transcript within 24 hours and was invited to edit, clarify, or withdraw any part before it was used in analysis. This was a commitment to treating participants as the authors of their own accounts.

Standard ethics protocols provide a useful procedural floor. But as Smith (2012) and others in decolonial research traditions argue, they often prioritise institutional liability over relational accountability. My commitments went further. My participants are among the professionals building Pakistan's AI-integrated future. Their observations are not data points for me to extract and analyse at a distance. They are insights offered in trust, shaped partly by the questions I brought into the room. I told participants from the outset what I was interested in and why. That honesty, I think, is what made the conversations as candid as they were.

The power dynamic in this research runs in an unusual direction. I came to these interviews as a graduate student. My participants are Chief Executives, Heads of Design and UX, and experienced independent professionals. The professional seniority was entirely on their side. My ethical responsibility was to take that seriously: to prepare thoroughly, ask questions worth their time, and produce research that does justice to what they shared.

4.9 Limitations

Several limitations of this study are worth acknowledging. Five participants based in Karachi cannot represent Pakistani professional practice at large, even though Karachi is where LLM adoption is most concentrated. Sectors with less exposure to these tools remain outside the frame.

The interviews were conducted entirely in English. The data therefore captures how participants chose to articulate their experience in that language. What they might have said differently in Urdu and what that difference might have revealed remains beyond what this study can access.

The research also captures a specific moment. LLMs are developing rapidly, and some of the empirical specifics described here may shift as these tools evolve. The structural and historical analysis is less time-sensitive, but the empirical findings should be read with that context in mind.

Finally, this study relies on interview data rather than direct observation of LLM use in practice. It reflects how participants make sense of and describe their experience, which is a legitimate and appropriate source of knowledge for this research question, but it does not provide access to the actual moments of interaction with these systems. Chapter 5 presents the findings from these interviews, organised around the themes that emerged through analysis.

5. FINDINGS

5.1 Introduction

The findings in this chapter draw on five semi-structured interviews with design and marketing professionals in Pakistani organisations, conducted in April 2026. The interviews combined open-ended discussion with three visual elicitation activities on a shared FigJam board. Participants completed a card sort asking them to categorise descriptors of LLMs, a dependency spectrum, and a language circles exercise. Two methods for data analysis worked in tandem: Braun and Clarke's (2006) thematic analysis in combination with critical discourse analysis (CDA). This allowed me to move between what participants said and how they said it, attending not only to the content of their accounts but to the discursive work those accounts perform.

Four themes emerged as most analytically generative in relation to this study's central research question, that is, how LLM integration in Pakistani organisations reproduces or challenges coloniality. The first, *Between Neutral Tool and Colonial Infrastructure*, examines how participants frame LLMs, what names and categories they reach for, and what those choices allow them to avoid saying. The second, *Navigating Dependency*, explores how participants narrate their position and organisational relationship to global AI platforms and the logics through which dependency comes to feel inevitable. The third, *English as Infrastructure*, looks at the language dynamics that participants described, including the quality gaps, the directional flows, and the moments where coloniality becomes visible in the texture of everyday LLM use. The fourth, *Recognition Without Transformation*, traces what happens after participants name a structural problem and examines how critique gets absorbed, deferred, or redirected in ways that leave the underlying conditions intact.

Additionally, Participants spoke about AI's implications for creative authorship and professional skill, about the competitive dynamics specific to Pakistani technology startups, and about what they imagined future AI development might look like. These threads appear where relevant within the thematic sections below. The four themes above

were selected because they speak most directly to the research question, to where structural inequality is visible, named, managed, and reproduced in everyday professional practice.

The participants in this study are thoughtful, technically capable professionals who understand many of the dynamics they are operating within. They are not passive victims or oblivious users. Analysis reveals something quite complex. Structural conditions get reproduced not through ignorance but through the accumulated pragmatic logic of working inside them.

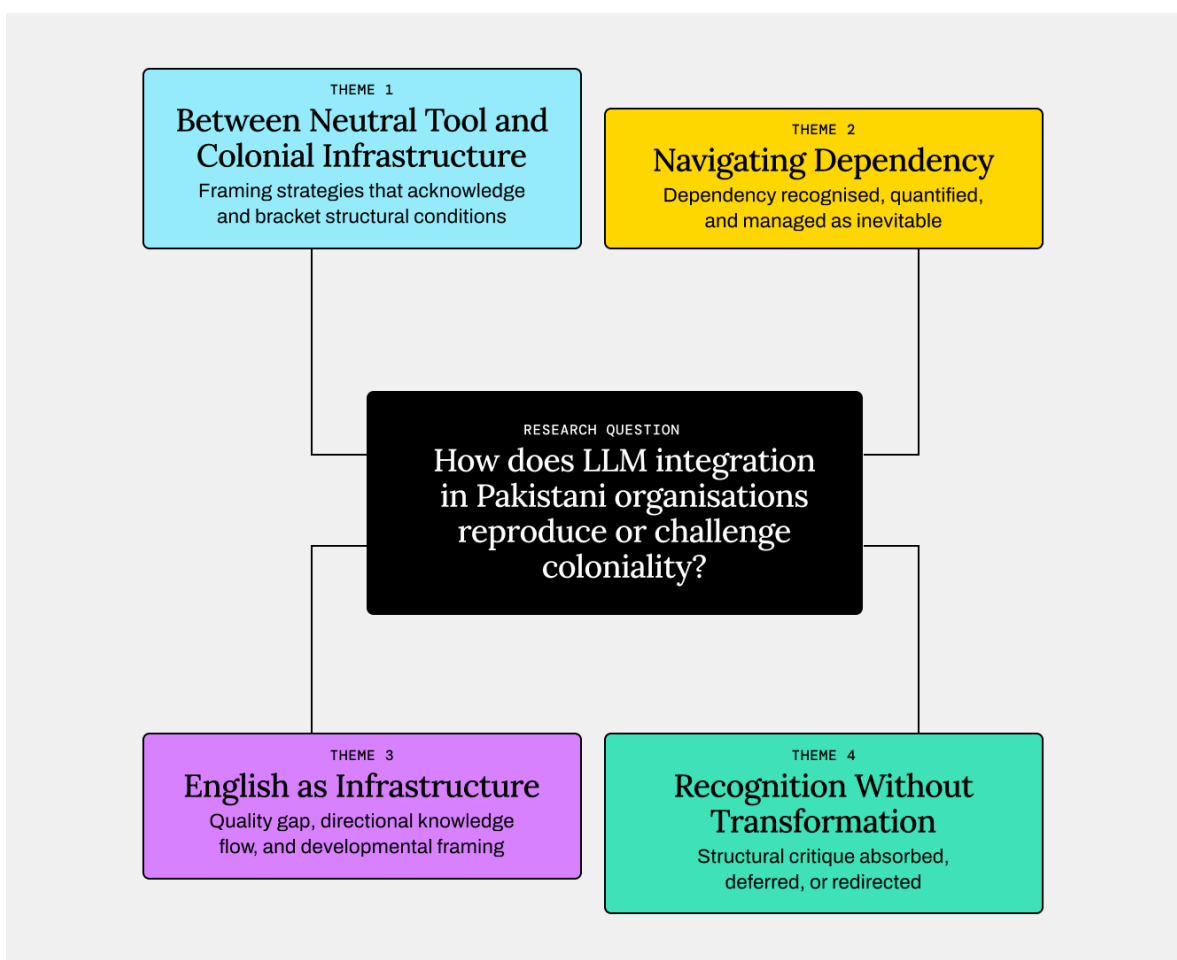


Figure 5.1 *Thematic map of findings*

5.2 Participant Profiles

All five participants were identified through purposive sampling, chosen for their active LLM use in professional settings and their location within Pakistan's design and technology sector. All interviews were conducted in English, a choice that is itself analytically significant and that I return to in the discussion of language hierarchies below.

Expert 1

Expert 1 leads the design function at the in-house startup incubator of a technology group in Pakistan. The organisation positions itself explicitly as AI-native. All workflows and ventures within the lab are expected to integrate AI, and experimentation with LLMs is embedded in performance management. Expert 1 describes their role as a kind of SWAT team function, plugging into ventures that need immediate attention across design, strategy, and product. They have been with the organisation for approximately two and a half years, with the past year spent specifically in the Labs context. Their LLM use is extensive and covers ideation, communication drafting, structuring strategy, and product development. Of the five participants, Expert 1 offered the most structurally aware commentary on LLMs, and was the only participant to use the word 'colonial' unprompted.

Expert 2

Expert 2 is the CEO of an edtech gaming company operating across Pakistan, USA and the Middle East serving the Arabic-speaking market. The gaming company produces interactive learning products for special needs children, maintaining in-house design, development, and sound production to protect product identity while using external LLMs for operational and communication tasks. Expert 2 is an active and reflective LLM user who uses ChatGPT and Claude daily for everything from drafting proposals to testing language capabilities in Urdu and Arabic. Their interview was among the most analytically rich, not because they articulated structural critique in theoretical terms, but because they moved between acknowledging the conditions of LLM dependency and reframing them in ways that the analysis attends to closely.

Expert 3

Expert 3 is Head of UX and Design at a multinational technology company with teams distributed across multiple continents. They have built an AI-assisted design workflow that includes a custom agent trained on design literature. LLM adoption decisions at the technology company are made through a formal hierarchy of vice presidents and senior directors rather than by individual practitioners, which shapes how Expert 3 experiences and describes dependency. They estimate that their own team is approximately 75% dependent on LLMs in their current workflow, and describe this as both strategic and temporary, though the nature of that temporariness is something the analysis interrogates.

Expert 4

Expert 4 works in digital marketing at an international technology and services company, where their responsibilities span marketing communication, project management, and website work. They have been using LLMs for approximately two years and describe their primary function in their workflow as formalisation, taking informal input and converting it into communication ready for professional contexts. Expert 4 was the most direct of the five participants in acknowledging bias in LLMs, and their account of dependency as both structural and strategic, driven by competitive mimicry rather than independent organisational choice, was one of the most analytically generative moments in the dataset.

Expert 5

Expert 5 is an independent digital marketing consultant and scriptwriter with extensive experience across social media strategy, video content production, and performance marketing. Their work includes audience growth campaigns for major brands, podcast production work, and social media management for international SaaS platforms. At the time of interview they were running performance marketing campaigns for a donor-funded organisation, which shaped their account of language and knowledge flows in distinctive ways. They use LLMs primarily as skilled interlocutors they actively direct, a practice they referred to as 'puppeteering'.

5.3 Between Neutral Tool and Colonial Infrastructure

The first theme concerns how participants name and frame LLMs, what categories they use, what those categories permit them to say, and what they foreclose. Across the five interviews, a range of framing positions emerged. At one end, Expert 1 used the word 'colonial' unprompted. At the other, Expert 5 positioned LLMs as fundamentally neutral tools. In between, Expert 2 and Expert 4 acknowledged structural conditions but moved quickly to contain them. These are not stages on a scale of awareness. They are different discursive strategies for managing the same structural situation, and they each do something specific.

Expert 1 was the only participant to name coloniality out of their own accord. The moment arose not in response to a question about power or geopolitics, but in the middle of the language circles activity, when discussing why people in Pakistan tend to use English even when they could communicate in Urdu. Their framing was rather casual.

...my understanding is that again it's more like a colonial thing anyways so it could be stemming from English anyways (Expert 1, personal communication, April 2026)

What is striking here is not just the use of 'colonial' but the 'anyways'. The word appears twice, carrying a tone of resignation or self-evidence, as though the colonial character of English-first communication is a fact so established it barely requires elaboration. This is not a theoretical intervention. It is a practical observation made in passing. Yet Expert 1 also holds this recognition carefully at a distance from their own professional practice. Later in the same interview, when asked whether LLMs feel culturally neutral in the context of their actual work, they were direct.

...underlying, underlying, we know for a fact they're not [neutral]. The neutrality does not matter in this case because we're — it's almost like math. We're telling

them, hey, 2 + 2, give me the answer. It gives us the answer. I do not know why the answer is whatever. (Expert 1, personal communication, April 2026)

The doubled 'underlying' is worth pausing on. Non-neutrality is positioned as a deeper, known truth. It is something beneath the surface that can be deliberately set aside because the transactional nature of the work makes it irrelevant. The maths analogy is particularly telling. It figures LLM use as pure input-output processing, stripped of interpretive or political content. This is a sophisticated move. Expert 1 acknowledges structural conditions fully, then constructs a frame, the transactional frame, in which those conditions cannot reach the surface of practice.

Expert 2 operates with a similar pattern, though their reasoning is explicitly economic rather than technical. When asked whether LLMs influenced the type of knowledge or perspectives their organisation encountered, they offered one of the most candid acknowledgments in the dataset.

...the readily available knowledge that is on this platform which is primarily the Western knowledge which is available more commonly — but the thing is that it does not create any restrictions for us because we usually use LLMs for office work... it's completely fine for us if they're Western dominated. It doesn't affect us because it actually works in our favour because what we're doing is actually making a product for the international market. (Expert 2, personal communication, April 2026)

The structure of this statement, an acknowledgment followed by a pivot and then a reframe, recurs throughout Expert 2's interview. Western knowledge dominance is real, named, and then neutralised by the fact that the edtech company's target market is itself international and largely Western-facing. There is no denial here. The logic is that coloniality becomes invisible precisely when your commercial aspirations are already oriented toward the centre. Nearing the end of our meeting, when I explained my research interest in coloniality and AI, Expert 2 pushed back with a frame that absorbs rather than

engages the critique. He said that "for us tech startups, we're just using technology and technology is like a post-colonial thing. It's globally the same kind of thing" (Expert 2, personal communication, April 2026). The use of the term 'post-colonial' here does not function as a critique; it functions as a temporal reassurance. Coloniality is something that happened, and technology belongs to an era after it.

Expert 4 offered the most unguarded statement about bias in the dataset, with the least discursive packaging around it. He stated directly that "it definitely does [influence perspective]. From whatever region they're coming from, they'll have some biases for them. Yeah, definitely" (Expert 4, personal communication, April 2026). This is an explicit acknowledgment of structural bias without the bracketing moves that Expert 1 and Expert 2 deploy. Yet Expert 4 does not develop it further. The observation sits in the interview without being connected to a critique of practice or a demand for alternatives. The bias is real, named, and then left there.

Expert 3 arrived at a similar conclusion through personal experimentation rather than structural reasoning. They described testing ChatGPT by asking it to take a side on a current conflict and finding bias emerge across multiple exchanges. Their framing positioned this as a product deficiency, arguing that LLMs "need to be a neutral thing" and should not "have some bias", rather than as an expected feature of systems trained on particular epistemic traditions. Expert 5's framing was the furthest from structural critique. When sorting the card labelled 'neutral tool' during the FigJam activity, they placed it in the 'strongly represents LLM' category, explaining that a neutral tool 'gives you whatever the generated output is neutral.' Expert 5 is a sophisticated and experienced user who has built extensive custom knowledge bases inside ChatGPT and actively manages output quality. But their frame for LLMs, as tools they direct whose outputs reflect their inputs, makes bias a non-category. If you control the prompt, you control the output. Structural asymmetry does not appear within this model of the tool. For a user of this sophistication, the absence of structural critique is analytically significant.

What emerges across these five positions is less a gradient from unaware to critically aware and more a set of different strategies for managing the same underlying conditions. Expert 1 knows and brackets. Expert 2 knows and reframes through market logic. Expert 4 acknowledges and moves on. Expert 3 discovers and frames the problem as a product defect. Expert 5's model of the tool makes the question incoherent altogether. Each of these strategies functions. It allows work to proceed, explains away friction, and makes a coherent professional position possible. What none of them do is refuse the terms or ask what it would look like for those conditions to change.

5.4 Navigating Dependency

The dependency spectrum activity asked participants to position themselves between local autonomy and global dependency. The resulting placements, and the reasoning participants offered to justify them, reveal something more interesting than the positions themselves. The analysis attends to how dependency gets narrated as rational, inevitable, strategic, or temporary in ways that close down the question of whether it should be otherwise.

Expert 4's account of their organisation's position on the spectrum was the most structurally explicit in the dataset. Rather than describing an active strategic choice, They described a competitive dynamic in which choice is effectively absent.

I don't kind of adopt whatever's going on with our global competitors and the companies in the same space that are doing globally, so we're not really building something on our own, which is locally autonomous or culturally. So we're just adopting whatever. So unless we build something on our own, that would be... I think structurally and strategically. (Expert 4, personal communication, April 2026)

Two things stand out here. First, the competitive mimicry logic. Adoption is not driven by an assessment of what the organisation needs, but by watching what global competitors are doing and following suit. The locus of decision-making is external. Second, the naming of

the dependency as simultaneously structural and strategic is important. It is structural because it is built into the operating conditions of the sector, and strategic because the organisation has chosen to lean into rather than resist it. The possibility of building something locally autonomous is acknowledged and then left in the conditional.

Expert 3 quantified their team's dependency at roughly 75% and described it as both strategic and temporary. When asked why, their answer pointed to future process maturity rather than any planned shift away from global platforms. He said that "once the process is in place, I think we can adopt more" (Expert 3, personal communication, April 2026). The temporariness here refers not to reduced dependency but to a phase in which dependency becomes more embedded and normalised. The horizon is more LLM use, more smoothly integrated, not less.

Expert 1 offered the most concrete illustration of what dependency means in practice. When asked how their organisation would be affected if access to current LLMs were restricted, their response was immediate. "We'd be on, we'd either be fired or we'd be on paid vacation" (Expert 1, personal communication, April 2026). They were not speaking hypothetically. At the time of the interview, escalating conflict in the Middle East had affected regional data centre infrastructure, with services disrupted across the MENA region for several days. Expert 1 described the experience of a recent Claude outage in the same terms. When the tool went down for several hours, the whole team was left uncertain about what to do. The only option was to wait for it to come back. He described the situation plainly.

There was a [Claude] outage, I think 4 to 5 hours... everyone was thinking, OK, what do I do? You can sort of wait for it to get back to work. I mean there's no alternative. It's almost like saying your laptop has crashed. You're in software. You're not physically [able to] create things with your hands. (Expert 1, personal communication, April 2026)

The laptop comparison is analytically precise. A laptop crash is not a loss of professional

skill. It is a loss of the infrastructure on which that skill now operates. Expert 1 is describing the same thing. When Claude goes down, their team has not lost capability. They have lost access to the platform on which their capability runs. And that platform sits outside Pakistan, on infrastructure whose operational continuity is shaped by geopolitical forces that no Pakistani organisation has any control over. The dependency is not merely professional. It is structural, and it extends into the terrain of geopolitics.

Expert 2 offered the most extended account of what dependency costs at an individual level. Their candour about personal cognitive effects was striking.

every single time that I'm asking AI to write an email for me, every single time that I'm asking the AI to write a proposal for me, I'm actually losing my ability to think clearly or to articulate my thoughts. That is absolutely there. But at the same time, there's a trade-off. (Expert 2, personal communication, April 2026)

Expert 2 is making a clear-eyed acknowledgment that dependency has a specific cost, the atrophy of articulation. But the trade-off frame that follows is equally precise. They explain that the need of the business is efficiency, and that if you want to retain critical thinking you can 'work on your skills in your free time.' The structural condition, that business efficiency overrides the development of human capability, is named, accepted, and then managed at the individual level. Critique does not travel from the personal to the systemic.

The edtech company's position is worth examining separately. Expert 2 maintains in-house design, development, and sound production, which they describe as essential to the product's identity. This local autonomy is real, but it operates at the level of product aesthetics rather than AI infrastructure. The LLMs themselves remain external global platforms. What the edtech company has preserved is its own creative signature, its visual and sonic distinctiveness, while the epistemic and linguistic infrastructure it depends on remains fully external.

Expert 5's relationship with dependency is framed almost entirely through agency. Their description of LLM use as 'puppeteering', specifically that 'you are telling them what they need to tell you' (Expert 5, personal communication, April 2026), positions them as the driver of the tool rather than a user shaped by it. This framing is not inaccurate. Expert 5 has built genuine expertise in prompt design, custom knowledge bases, and output quality control. But the puppeteering frame also has a structural blind spot. It accounts for what happens within a session, the user's direction of the tool, without accounting for the conditions that structure what the tool can do. The puppet can be directed. The stage, the material, and the language the puppet speaks are all set before Expert 5 arrives.

When asked whether they thought LLMs would replace the workforce, Expert 5 extended the same argument in an interesting direction.

you need someone to use those platforms also, so you need a better prompt engineer for sure, someone who's more well versed, someone who knows how to use it in the most effective way, but I don't think the agency is going to replace human beings. (Expert 5, personal communication, April 2026)

They went further, noting that when they wrote copy with ChatGPT, they are 'the one who's adding emotions in the copies. They cannot. The LLMs cannot. They're not there' (Expert 5, personal communication, April 2026). This is an argument for bidirectional dependency. LLMs need skilled human operators just as much as professionals need LLMs. The point is not wrong. Emotional intelligence and contextual knowledge cannot be replicated by the current generation of LLMs. But the claim carries an asymmetry worth noticing. Pakistani professionals cannot do their current work without US-controlled platforms. LLM providers are not dependent on any particular set of users. What Expert 5 is describing is real, irreplaceable human value. What the frame does not account for is where the tools that require that labour are built, owned, and governed.

What runs through all five accounts is a version of the same recognition. Dependency is real, its costs are known, and it is nonetheless the condition everyone is working within. What differs is how that recognition is managed. Expert 1 frames it through the reality of

infrastructure failure. Expert 2 frames it through the logic of trade-offs. Expert 4 frames it through competitive necessity. Expert 3 defers it into future process maturity. Expert 5 reframes it as bidirectional, an argument that restores a sense of agency without disturbing the structural conditions that initially created the dependency. None of these strategies are failures of analysis. They are the ordinary cognitive work of professionals operating inside a system they did not design.

5.5 English as Infrastructure

The language circles activity, in which participants were asked to resize three circles representing English, local language, and hybrid or mixed use, and then to draw arrows showing the direction of translation or adaptation, produced some of the most direct evidence in the study. Across all five interviews, English dominated. Arrows, where they appeared, moved overwhelmingly in one direction, from local languages and hybrid usage toward English. No participant drew arrows flowing the other way.

In itself, this is not surprising. What the interviews added to the visual data was an account of why this directionality feels natural, rational, and often goes unremarked. The reasons participants offered, including better output quality, professional register requirements, and target market orientation, each tells a partial story about how English has become not just a preferred language but a structural condition of LLM use in this context.

Expert 1 was explicit about the quality differential. They keep their LLM interactions restricted to English because 'there's a better chance of it understanding and generating that output.' The rationale is immediate and pragmatic. But they also build a structural explanation into it. The quality gap exists, they said, 'because of the ecosystem. It is not as rich.' They described this in relation to building educational tools for their son.

the subjects in English would be pretty good, but then in school like Pakistan with their geography knowledge, local or Urdu specifically — it does know it, but it

can't translate it as well as the English counterparts because of the ecosystem. It is not as rich. (Expert 1, personal communication, April 2026)

The word 'ecosystem' carries analytical weight here. It points to the difference in the volume and quality of training data across languages, a gap that is itself the product of historical patterns of knowledge production, digitisation, and linguistic capital. Expert 1 does not theorise this, but their framing is accurate. The Urdu performance gap is not a technical accident. It is the accumulated effect of which languages have been systematically documented, published, and made digitally available.

Expert 2 offered a counterpoint, the most optimistic account of Urdu language capability in the dataset. They argued that 'language is no longer a barrier for LLMs' and backed this with a specific test case.

I've translated the phrases of Allama Iqbal's poetry through ChatGPT and it has given me very good answers. Then I asked it to go into the esoteric meanings and the spiritual side, and it gave me a really good context of that as well. Allama Iqbal is considered the pinnacle of Urdu poetry or the pinnacle of Urdu complexity — and ChatGPT is very easily deciphering that and telling me it in English and in Urdu. So I think language, although these elements are English-first in the sense that the user experience and everything is English — but that's all there is to it. (Expert 2, personal communication, April 2026)

This is an interesting test to choose. Allama Iqbal is among the most extensively translated, documented, and digitally available Urdu poets in the world, precisely the kind of source likely to be well-represented in any large training corpus. Using the most globally recognised Urdu literary figure as evidence that the language barrier has been overcome substitutes the exceptional case for the general condition. Expert 2 concedes that the UX is 'English-first' and then moves immediately past it, saying 'that's all there is to it', without asking what that UX-level dominance means for users less confident in English, or for knowledge that lives in less globally celebrated registers.

Expert 4 described a dynamic that sits between these two positions. Their primary use of LLMs is as a formalisation tool.

I give input to LLMs in whatever language I'm comfortable in, in a very informal way, then I get a response which is much more formal and ready to be used for official formal communication. So I think that's something it represents to me — LLMs help me get those proper formal responses. (Expert 4, personal communication, April 2026)

The trajectory here, informal input in a mixed or comfortable language converted to formal English output, encodes a hierarchy. Informal and local-language expression is the raw material. English formal communication is the finished product. The LLM operates as a machine for converting one into the other. Expert 3 described a similar pattern at the level of their organisation in a single line. They noted that 'people talk in hybridic usage, and then it translated into English' (Expert 3, personal communication, April 2026). Input enters in Urdu, Roman Urdu, or a mixture, and it exits in English formal register.

Expert 5's account made this unidirectionality most explicit. Working for a donor-funded organisation, they described how beneficiary stories travel through a specific pipeline.

Sometimes we take videos of our different donors and beneficiaries — they record their stories in their regional languages. Sometimes we send it to a translator, or sometimes we transcribe it through different tools like Turboscribe. And then we take that translation and put it in ChatGPT and ask it to translate into proper English so we could turn it into subtitles for the videos... everything goes into English. (Expert 5, personal communication, April 2026)

'Everything goes into English' is spoken as a simple factual description of workflow, not a complaint. Stories originate in the languages of the people whose experiences they describe. They are then translated, transcribed, and formalised into English for circulation

and fundraising. The LLM is the final stage in a pipeline that converts local lived experience into a format legible to a global English-speaking readership. The direction of that flow is precisely what critical scholars of epistemic injustice point to when they describe how local knowledge gets extracted, translated, and made to serve frameworks it did not generate.

Expert 5 also described what happens when they work in the other direction, trying to communicate with the LLM in Urdu.

I mostly interact in English. Because I think I speak English more than any other language, more than Urdu. Also, sometimes I speak in Urdu, but the way it speaks back to me, it's a little confusing. It's partially Hindi, partially Urdu, you know. LLMs are still not entirely well versed with the languages right now, so they could get confused with different stuff. (Expert 5, personal communication, April 2026)

The confusion between Hindi and Urdu is not a minor linguistic inconvenience. In Pakistan, the difference between the two carries significant cultural, historical, and political weight. They are distinct languages with distinct scripts, vocabularies, and national identities, even as they share a spoken grammar. An LLM that cannot reliably distinguish between them is not failing neutrally. It is failing in a way that reflects the marginalisation of Urdu within global AI training infrastructure. Expert 5 frames this as a developmental issue, describing LLMs as 'still not entirely well versed', which positions the present inadequacy as temporary and the current hierarchy as transitional. This temporal framing does important work. It makes waiting rather than questioning, the appropriate response.



Figure 5.2 *Language circles composite across five participants*

5.6 Recognition Without Transformation

A pattern runs through the three themes above that warrants its own analytical attention. Across the interviews, participants regularly acknowledged structural conditions, including Western knowledge dominance, linguistic inequality, cognitive dependency, and the absence of local AI infrastructure, and then, in the next breath, produced a frame that made those conditions manageable, acceptable, or simply the backdrop against which work proceeds. This fourth theme examines how structural critique gets contained. I am calling it recognition without transformation. It is the condition in which a problem is named but the naming does not produce, and is not expected to produce, any change in the conditions that generate it.

The clearest instance of responsibility displacement in the dataset came from Expert 4, who, when discussing what would make LLMs more useful for Pakistani organisations, described a set of features they hoped OpenAI would one day provide.

If OpenAI gives us for Pakistan specifically, they give us something that — you know, you can create this, which is Pakistan-relevant. They have a whole media

gallery section which is relevant for Pakistan, language section that you can choose from for regional languages or somewhat in that sense. (Expert 4, personal communication, April 2026)

The grammar of this statement is worth noticing. It is entirely conditional and entirely located in the provider. Pakistan-specific features, regional language support, locally relevant media, these are things OpenAI would give, if it chose to. There is no articulation of what Pakistani organisations, researchers, or institutions might do. No suggestion that the gap itself might motivate local capacity building. No naming of what it costs to wait. The structural problem is accurately identified, and then resolved by the imagination of a better product from the same provider that generated it.

Expert 2's trade-off logic performs a different kind of containment. Their acknowledgment of cognitive dependency, that they are actually losing their ability to think clearly or to articulate my thoughts. That is absolutely there', is striking in its clarity. But the frame that follows, 'there's a trade-off,' does specific work. A trade-off implies two things of comparable weight being exchanged, in this case the efficiency gains of AI-assisted writing against the atrophy of independent articulation. Framing it this way positions the loss of critical thinking capacity as a reasonable price for business efficiency, and locates the decision about that price at the level of the individual professional. The question of whether organisations, sectors, or educational systems should address this collectively does not arise.

Expert 5's account of managing ChatGPT's repetitive output patterns adds another dimension.

GPT has its own way of writing things and it's very repetitive. Imagine, imagine this, imagine that, not only but also — you know, all those things, it's just very repetitive, so I have to QA it on my own. (Expert 5, personal communication, April 2026)

The labour here, the careful reading, editing, and rewriting required to make AI output usable, is invisible as labour because it is framed as quality control rather than as remediation. Expert 5 does not describe it as an additional burden generated by an inadequate tool. They described it as part of their professional practice. The homogenisation of language, the flattening of voice, the particular repetitive tics that ChatGPT defaults to, these are all things they manage, absorb, and correct for, without prompting any question about whether they should have to.

Expert 5's description of the Urdu-Hindi confusion provides a final example. When the LLM returns responses that blend Urdu and Hindi, they name it, attribute it to developmental immaturity, and move on. The developmental frame positions the present inadequacy as a temporary station on a linear improvement trajectory, with patience as the appropriate posture. Whether that trajectory will actually arrive at meaningful Urdu capability, and on whose timeline, is a question the frame forecloses.

What connects these different instances is not a shared emotional register. Expert 4 sounds matter-of-fact. Expert 2 sounds reflective. Expert 5 sounds practical. What they share is a structural outcome. In each case, the naming of a condition does not disturb it. Awareness is not what is missing here. Several participants named the structural conditions they were operating within with considerable precision. What critique runs into is the logic of professional practice itself: competitive necessity, efficiency demands, product-market requirements. These conditions are so thoroughly embedded that naming them changes nothing. The structures that would need to shift are not available for change at the level of individual professional choices.

This is where the research question becomes most pressing. The thesis asks whether LLM integration in Pakistani organisations reproduces or challenges coloniality. The answer, across these four themes, is not that it reproduces coloniality through ignorance or complicity. It reproduces it through the ordinary, pragmatic reasoning of capable

professionals who understand their situation and have no framework, and no structural invitation, to act otherwise.

5.7 Additional Emerging Themes

Three additional threads appeared consistently enough across participants to warrant attention, even if they do not carry the same sustained analytical weight as the main themes. They concern the tension between AI use and creative professional identity, the competitive logic that drives adoption in Pakistani technology organisations, and the imaginaries of AI's future that participants brought to the conversation.

Creative Authorship and Professional Identity

Several participants described a tension between their use of LLMs and their sense of themselves as creative professionals. The most direct account came from Expert 2, who reported significant internal resistance when pushing AI adoption at the edtech company.

the mindset of the designers is that this is cheating, you know, this is not creativity. This is not authenticity, which is absolutely right. And there is a concern that if we use AI a lot, then our good skills will deplete over time. (Expert 2, personal communication, April 2026)

What is analytically significant here is not that designers resist AI, but that Expert 2 validates the critique fully and then proceeds to drive adoption regardless. The concern about deskilling is named, affirmed, and overridden by the logic of business efficiency. It is a compressed version of the recognition without a transformation pattern running through the whole dataset. Expert 5 articulated a similar professional identity claim from their own position: "we are creatives. We are not typically people who just believe in creating outputs. We believe in quality also" (Expert 5, personal communication, April 2026). In the same breath, they described a supervisor who expected AI-assisted ad copies in thirty minutes rather than the three to four hours they normally invested. The creative identity is

asserted and then quietly negotiated against the productivity expectations that AI use enables.

The Competitive Logic of AI Adoption

Across several interviews, adoption was described as a response to competitive pressure. Expert 4 described it with particular clarity during the card sort activity, explaining competitive necessity as the condition in which two companies both understand the costs of AI adoption but proceed because the alternative is being outpaced. Expert 3 framed it more starkly: "everyone would be using LLM, else they'll be left behind" (Expert 3, personal communication, April 2026). Expert 2 used almost identical language: "if we do not utilize AI then we're going to be left behind" (Expert 2, personal communication, April 2026).

The repetition of the same phrase across independent interviews is worth pausing on. 'Left behind' constructs non-adoption as a failure of survival rather than a legitimate choice, and it positions the relevant competitive field as global. Pakistani organisations are measuring themselves against international benchmarks, adopting tools built by and for different markets, and doing so not because those tools serve their specific contexts best but because the alternative is falling out of the race. Expert 1 formalised this logic at the organisational level: at their organisation, every venture is required to be AI-native. Adoption is embedded in performance expectations. Whether it's necessary is never asked.

Future AI Imaginaries

When asked how they envisioned LLMs evolving in their work, participants described futures that were almost uniformly extensions of the present. More integration, more automation, fewer people doing more work. Expert 2 anticipated that a social media team of three would eventually be consolidated into one person whose range had been expanded by AI. Expert 3 saw organisations building their own agents and workflows on top of existing global platforms, not alternatives to them. Expert 1 was unambiguous about irreversibility: they said it was "very unlikely that we ever have a conversation where we're

like, yeah... we have stopped using" LLMs (Expert 1, personal communication, April 2026).

The one articulation of an alternative imaginary came from Expert 2, who said that "if in the future there's a locally developed solution, or a locally hosted solution that is performing better for us for our needs, we will very happily and gladly switch" (Expert 2, personal communication, April 2026). This is the only moment across the five interviews in which a locally produced AI infrastructure is named as a possibility. But the framing is conditional and comparative. A local solution would need to outperform the current global offering before it would be considered. The possibility of building or supporting local AI capacity for reasons other than performance parity, for reasons of sovereignty, cultural representation, or epistemic equity, does not surface. The futures participants imagine are better versions of the present arrangement.

5.8 Chapter Summary

Four themes and three additional threads have emerged from the analysis, together mapping how LLM integration unfolds at the junction of professional practice and structural inequality in Pakistani design and marketing organisations. The first theme showed that participants do not frame LLMs as neutral tools out of ignorance. They frame them that way through active discursive work that acknowledges structural conditions and then constructs reasons why those conditions do not apply to their own practice. The second theme showed that dependency is widely recognised, often precisely quantified, and consistently managed through logics, including trade-offs, competitive necessity, and future resolution, that make the question of alternatives difficult to raise. The third theme showed that English functions as infrastructure in LLM use, not through explicit preference alone, but through a quality differential, a directional flow in knowledge translation, and a developmental framing that makes the current hierarchy feel transitional rather than structural. The fourth theme drew these threads together to argue that what is most analytically significant is not the absence of critique but its absorption, the way in which recognition of structural conditions stops short of transformation. The additional themes section documented how creative professional identity becomes a site of resistance that is named and overridden, how competitive survival logic forecloses reflection on

adoption decisions, and how participants imagine AI futures that extend rather than renegotiate the current terms of dependency.

The next chapter brings these findings into conversation with the theoretical literature reviewed in Chapter 3, examining what the theoretical frameworks of coloniality, epistemic injustice, and AI governance contribute to understanding what has been described here, and where the empirical data complicates or extends those frameworks.

6. DISCUSSION AND CONCLUSION

6.1 Introduction

This chapter reads what emerged in Chapter 5 into dialogue with the theoretical literature and draws out their practical and structural implications. It begins by returning directly to the central research question: how does LLM integration in Pakistani organisations reproduce or challenge coloniality? It then situates the findings within the frameworks of coloniality of knowledge, data colonialism, epistemic injustice, and decolonial AI scholarship. The second half of the chapter proposes a service design intervention that responds to what the findings revealed. The intervention is a cultural knowledge repository, an open source, community governed resource designed to support the existing AI infrastructure rather than replace it. It is not a proposal to build a new LLM. It is a proposal to build the cultural layer that any LLM can draw from so that Pakistani professionals no longer have to do that work individually. The model is inspired by the Masakhane initiative in Africa, adapted for the Pakistani context, and designed to be replicable by any community that currently finds itself invisible in global AI systems.

6.2 Revisiting the Research Question

The central research question asked whether LLM integration in Pakistani organisations reproduces or challenges coloniality. The answer, drawn from the analysis in Chapter 5, is that it reproduces it. What matters analytically is not the fact of reproduction but its mechanism.

Coloniality is not reproduced here through ignorance, complicity, or an absence of critical awareness. The participants in this study are thoughtful professionals who can name the structural conditions they are working within. Expert 1 identified the colonial character of English language dominance in passing, almost casually. Expert 2 acknowledged Western knowledge bias directly and precisely, then reframed it through market logic. Expert 4 named it and moved on. Expert 5 built elaborate personal systems to manage its effects. The awareness is present. What is absent is any structural pathway from that awareness to

action.

Chapter 5 named this pattern recognition without transformation. It is the condition in which a structural problem is accurately identified but the identification does not produce, and is not expected to produce, any change in the conditions that generate it. The competitive logic Expert 4 described makes adoption a survival decision rather than a considered choice. The trade-off frame Expert 2 used absorbs the cost of cognitive atrophy into individual professional responsibility. The transactional frame Expert 1 constructed makes the political nature of LLM outputs irrelevant to the work at hand. Each of these is a rational response to real conditions. None of them challenges those conditions.

The additional emerging themes documented in Section 5.7 reinforce this reading. The futures participants imagined were extensions of the present arrangement. The one moment anyone articulated an alternative, Expert 2's conditional openness to a locally developed LLM, was framed entirely in performance comparison terms. A local solution would need to outperform current global tools before it would be considered. The possibility of supporting local AI infrastructure for reasons of epistemic equity or cultural sovereignty did not surface. This is far from a failure of imagination. It is the structural effect of operating inside a system that makes the current arrangement feel like the only reasonable one.

There is one partial exception worth noting. The creative professional resistance documented in Section 5.7 represents an honest moment of challenge. The designers who told Expert 2 that AI use is cheating and that it is not creativity were making a real claim: that some forms of professional knowledge cannot be substituted without loss, and that the loss matters. But even that critique was acknowledged, validated, and overridden by the logic of business efficiency. Recognition without transformation, again.

The thesis therefore answers its own question as follows. LLM integration in Pakistani organisations does not challenge coloniality in any structural sense. It reproduces it through the ordinary, technically skilled, commercially rational decisions of professionals whose recognition of structural conditions has nowhere to go.

6.3 Situating Findings in the Literature

Quijano (2000) argues that the domination colonial conquest established has not dissolved with formal decolonisation. It persists long after formal decolonisation, not through continued occupation but through the continued organisation of knowledge, labour, and social life around a colonial axis. The findings in this thesis illustrate a contemporary iteration of that persistence. The structural conditions of LLM infrastructure, where knowledge is produced, in which language, on whose servers, under whose governance, replicate the colonial organisation of knowledge production. Pakistani organisations do not resist this because, as Quijano's framework anticipates, the alternative is not visible within the logic of the current arrangement.

Mignolo's (2011) elaboration of the coloniality of knowledge adds the epistemological dimension. Western knowledge is not merely dominant. It is positioned as universal and neutral while non-Western knowledge is positioned as local, particular, and supplementary. This is precisely the dynamic Expert 1 described when they said that they knows for a fact LLMs are not neutral, before constructing a frame in which that non-neutrality becomes professionally irrelevant. The neutrality of LLMs is not believed by the practitioners who use them. It is strategically performed, because engaging seriously with the structural conditions of non-neutrality has no practical outlet within current professional arrangements.

Couldry and Mejias's (2019) data colonialism framework extends this analysis to the specific dynamics of digital platforms. Their argument is that the extraction of human behaviour and social practice as data by large technology corporations embody a somewhat evolved colonial appropriation, in which the value generated by that data flows from the periphery to the centre. The unidirectional knowledge flows documented in Section 5.5 are a direct instance of this. Stories originating in the languages of Pakistani beneficiaries are transcribed, translated, formalised into English, and made legible to a global readership. This is a pipeline in which local knowledge is extracted, processed, and made to serve frameworks it did not generate, by platforms whose training data is enriched by the same process. The structural effect of this pipeline is visible at the language level: when

Anthropic addressed Urdu as part of its India initiative, launched alongside its Bengaluru office, Urdu was listed alongside Hindi, Bengali, and Tamil as an Indian language. Pakistani cultural context was not the frame. The result, as Expert 5 and other participants noted, is that LLM outputs in Urdu read more like Hindi than Pakistani Urdu. The training data reflects a geography that is not Pakistan's.

Fricker's (2007) framework of epistemic injustice adds a dimension of the findings that the coloniality literature alone does not fully address. Hermeneutical injustice is the condition in which a person lacks the cognitive tools to make sense of their own experience. This is visible in Expert 5's description of the Urdu-Hindi confusion. They framed it as a developmental limitation of the tool because the available framework for understanding LLM inadequacy is one of technical immaturity, not structural marginalisation. The concepts that would allow them to name what is actually happening, the systematic underrepresentation of Pakistani Urdu in training data as a consequence of historical patterns of digitisation and linguistic capital, are not part of the professional vocabulary available to them.

Mohamed, Png, and Isaac's (2020) findings for decolonial AI positions these dynamics within an emerging scholarly conversation about the structural conditions of AI development. Their argument is that decolonial approaches require not merely ethical adjustments to existing systems but a fundamental rethinking of who builds AI, for whom, and on whose epistemic terms. The findings in this thesis provide empirical texture to that argument at the level of everyday professional practice. They show what decolonial AI literature looks like from the inside, not from the perspective of AI developers but from the perspective of practitioners in a postcolonial context who are navigating its consequences.

Escobar's (2018) argument of the pluriverse that multiple worlds and knowledge systems can coexist rather than being absorbed into a single universal framework, provides the theoretical foundation for the intervention proposed in the following section. The problem the findings document is not diversity within a single universal system. It is the absence of

infrastructure for plurality. The intervention responds to that absence directly.

6.4 The Intervention: The Cultural Knowledge Commons

The four themes and three additional threads documented in Chapter 5 converge on a single structural gap. Pakistani professionals can recognise the colonial conditions of LLM infrastructure. They cannot act on that recognition because there is no collective framework, no shared resource, and no structural invitation to do so. Individual workarounds exist. Expert 5 builds their own knowledge bases. Expert 1 identifies the ecosystem gap. The designers at the edtech company name AI use as a form of cheating. But these remain private, fragmented, and ineffective at the structural level. What is missing is infrastructure for collective action.

This section proposes a service design intervention designed to address that gap. The Cultural Knowledge Commons is an open source, community governed repository of Pakistani cultural knowledge, language resources, and professional context for AI. It is not an LLM. It does not compete with the AI infrastructure that organisations are already using or that local developers are building. It is the layer underneath, the resource that any LLM can draw from to produce outputs that are culturally grounded in Pakistani context rather than filtered through a Western or Indian lens. Any developer, organisation, or individual using an LLM can access the commons and apply it to improve the cultural accuracy of their outputs without changing platforms or adopting new tools.

The model is inspired by Masakhane, the African grassroots initiative whose name means We Build Together in isiZulu. Like Masakhane, the commons starts with community ownership rather than institutional top-down delivery. It builds the infrastructure that makes collective epistemic action possible, and it does so in a way that is replicable by any community that currently finds itself invisible in global AI systems.

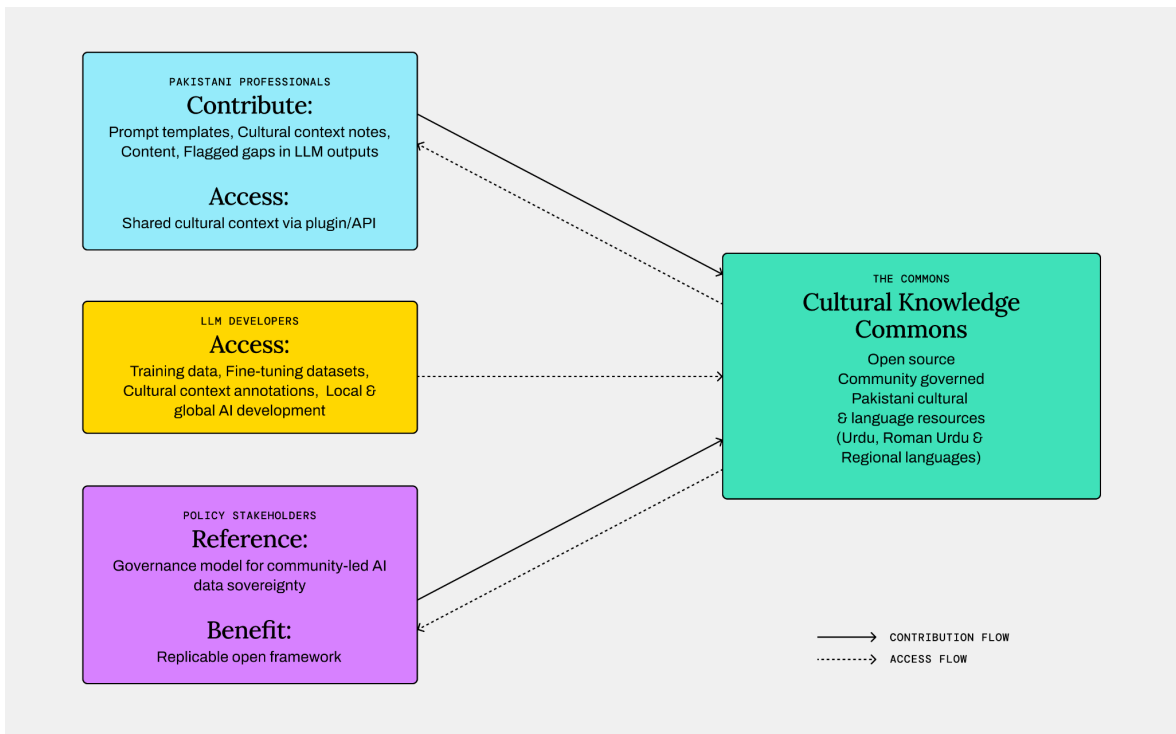


Figure 6.1 *The Cultural Knowledge Commons*

6.4.1 Decolonial Justification

The intervention is designed to shift Pakistani organisations and professionals from being consumers of Western knowledge infrastructure to being producers and custodians of their own. This is epistemic sovereignty made practical, not as a theoretical position but as a design artefact.

The decolonial logic operates at three levels. At the level of knowledge production, the commons addresses the coloniality of knowledge by creating infrastructure through which Pakistani cultural knowledge, Urdu and regional language resources, and locally grounded professional context become part of the AI knowledge ecosystem. Not as a supplement to a Western universal, but as a legitimate knowledge tradition with its own integrity. At the level of ownership, the community governed, open source structure addresses data colonialism directly. The knowledge in the commons belongs to the community that produced it, not to a corporation extracting value from its use. At the level of plurality, the replicable model enacts Escobar's pluriversalism. Because the framework can be adapted

by any community that feels invisible in global AI systems, it creates infrastructure for a genuine plurality of knowledge traditions rather than one global standard with local patches.

The linguistic dimension deserves particular attention. Roman Urdu, Urdu written in the Latin script, is the actual language of digital life for a large portion of the Pakistani population. WhatsApp messages, social media communication, and professional exchanges happen in Roman Urdu as much as in formal Urdu script. Current LLMs process Roman Urdu but produce outputs that sound more like Hindi than Pakistani Urdu, because the training data reflects India's significantly larger digital footprint. This is a structural misrepresentation of Pakistani language and culture at scale. The commons addresses this by building a repository that is specific to Pakistani Urdu, Roman Urdu, and the regional languages of Pakistan, including Pashto, Sindhi, Punjabi, and Balochi, several of which do not have their own standardised digital script and currently rely on Urdu and Roman Urdu to survive in digital spaces. Without a dedicated resource, these communities remain invisible not only in global AI systems but in any locally developed system that does not deliberately include them.

One tension deserves honest acknowledgment. The intervention still operates on top of existing Western LLM infrastructure in the short term. It does not escape dependency. ChatGPT, Claude, and Gemini still run the outputs. A hardline decolonial critique would position this as accommodation rather than liberation. The response to that critique is the same one Masakhane offers: this is a bridge, not a destination. It starts within the current infrastructure because that is where Pakistani professionals are working today, and it builds the knowledge base and community capacity that make a structural transition possible over time. Ongoing initiatives in Pakistan to develop locally built language models represent exactly the kind of structural transition this commons is designed to feed. The knowledge and data accumulated through community contribution become training resources for locally developed AI infrastructure when that infrastructure becomes viable.

6.4.2 The Service Design Layer

What makes this a service design intervention rather than a technology project or a policy recommendation is that it is designed around the actual behaviours, needs, and constraints of the people it serves, observed directly in the research that generated it.

Sanders and Stappers' (2008) distinction between designing for users and designing with them is foundational here. The commons is not a solution delivered to Pakistani professionals. It is infrastructure they participate in building, in ways calibrated to how they already work. The contribution model is derived from an observation the research made directly. Professionals like Expert 5 and Expert 1 are already building cultural context into their LLM workflows, instinctively, alone, without a shared framework. The intervention gives that existing instinct a collective structure. It designs with the behaviour that already exists, not against it.

The systems thinking dimension is equally deliberate. The intervention does not address any single interaction between a Pakistani professional and an LLM. It addresses the infrastructure level at which structural conditions are reproduced. This is what Jones (2014) describes as designing for systemic change: not improving the user experience of an existing system, but creating the conditions for a different system to become possible.

In practical terms, the commons serves three distinct groups. Pakistani professionals access it through their existing LLM tools via a plugin or API layer. They inherit cultural context and language resources that others in the community have built, without needing to switch platforms or rebuild that context from scratch each time. This directly reduces the individual overhead that every participant in this study was absorbing invisibly. LLM developers, whether building global models or locally developed Pakistani systems, can use the repository as training data, fine-tuning datasets, and cultural context annotations, improving the accuracy of their outputs for Pakistani audiences without conducting separate research. Policy stakeholders can reference the commons as a governance model

for community-led AI data sovereignty, a practical demonstration that communities can build and maintain AI knowledge infrastructure outside the control of large technology corporations. Each of these pathways is designed to require minimal specialist knowledge, to connect to existing tools, and to deliver measurable value quickly enough to justify adoption on business terms.

6.4.3 Why It Would Thrive: Business Case and Sustainability

The intervention is sustainable because it aligns decolonial goals with business interests. Interventions that require organisations to act against their commercial interests do not survive long enough to matter. Designing with that constraint is not a compromise of the decolonial argument; it is what makes the argument actionable.

The most immediate business case is cost reduction. As documented in Sections 5.5 and 5.6, Pakistani professionals currently absorb an invisible overhead from culturally inadequate LLM outputs. Expert 5's quality assurance work, Expert 4's iterative formalisation rounds, and Expert 1's individual knowledge base construction are all examples of this overhead: additional labour generated by tools that do not understand the professional context they are being asked to serve. A shared cultural knowledge commons reduces this tax collectively. Organisations that connect to the commons inherit cultural contextualisation that others have already built. The efficiency argument is real, measurable, and does not require any commitment to decolonial principles to be compelling.

The second business case is output quality for local markets. Organisations producing content, products, or services for Pakistani consumers benefit directly from culturally resonant outputs. At the moment, when a Pakistani organisation uses an LLM to communicate with a local audience in Urdu or Roman Urdu, the output tends to reflect Hindi idiom and Indian cultural register rather than Pakistani. This is a quality problem with a direct commercial cost. The commons addresses it at the infrastructure level rather than requiring each organisation to solve it individually.

The central challenge of the commons is that it depends on community contribution to be useful, and that dependence cannot rest on goodwill alone. Goodwill is not a sustainable foundation for any infrastructure. The sustainability model addresses this directly through a trend intelligence layer. As the commons grows, it accumulates a mapped, structured, community validated record of Pakistani cultural knowledge, language patterns, and professional context. That record is itself intelligence infrastructure. Pattern recognition across the commons can surface emerging shifts in Pakistani consumer culture, professional language, and market dynamics in ways that no commercial analytics provider currently offers. A subscription tier for trend intelligence, offered to Pakistani brands, marketing agencies, and research organisations, funds the open source core without compromising it. The knowledge repository remains freely accessible. The intelligence layer derived from it is commercially licensed. Community contributors benefit from the intelligence too, creating a genuine participation incentive. The more organisations contribute, the richer the intelligence. The richer the intelligence, the stronger the incentive to contribute. This is a loop that reinforces itself rather than a model that relies on altruism.

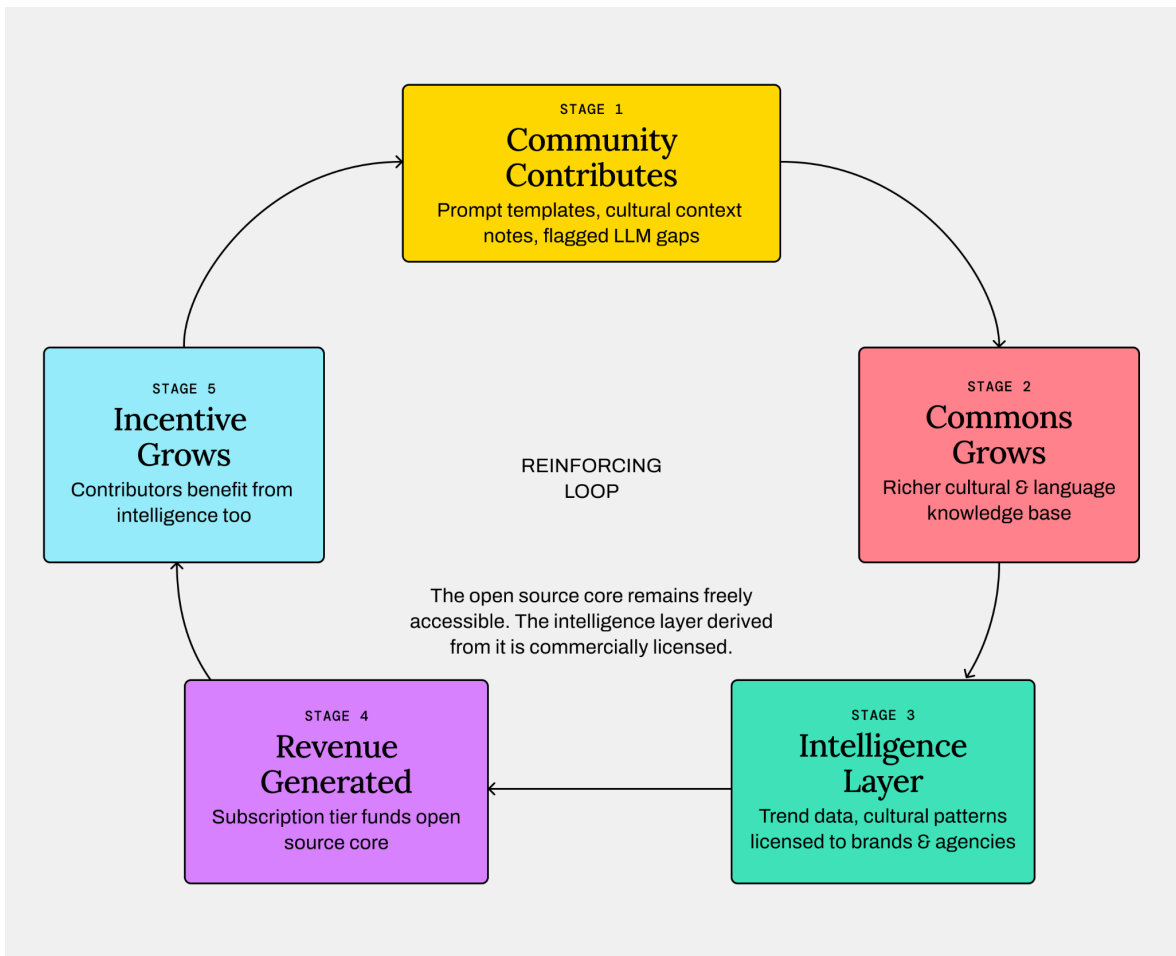


Figure 6.2 *The sustainability loop*

The contribution structure follows the Masakhane precedent. Institutional actors including university linguistics departments, government digital bodies, and organisations with mission-based reasons to invest in high-quality Urdu and regional language AI resources carry the backbone of the commons. Professional contributors share lightweight assets: prompt templates that worked, cultural context notes, flagged gaps. The barrier is low and attribution is visible. Organisations are primarily users who access the commons through their existing LLM tools without needing to change platforms or adopt new infrastructure.

Critically, the initiative starts without requiring institutional buy-in. A small community of motivated professionals can seed it. Institutional backing follows proof of concept rather than preceding it. This is the sequence Masakhane followed, and it is the sequence that preserves community ownership rather than allowing institutional actors to shape the

commons in their own interests before the community is established. The commons is also designed to be replicable. The model, the governance structure, and the technical infrastructure are openly documented and explicitly invite adaptation by any community that feels invisible in global AI systems. Pakistan becomes a precedent for how communities can respond.

6.5 Limitations and Future Research

Multiple limitations of this study and the intervention it proposes should be acknowledged directly.

The research is founded on five interviews with design and marketing professionals in Pakistani organisations. While the participants were carefully selected for their active LLM use and analytical reflectiveness, five interviews cannot represent Pakistani professional practice at large. The findings describe specific, analytically significant moments in specific professional contexts, not a statistically representative picture of LLM adoption across Pakistan.

The intervention proposed in this chapter is an untested design proposal derived from the research findings and grounded in comparable precedents. Evaluation of its effectiveness requires implementation, and implementation requires resources and institutional coordination that are beyond the scope of this thesis.

The most significant practical limitation is the contribution dependency. The commons is only as useful as the community that populates it. Without active and sustained contribution, it becomes a well-designed empty repository. The sustainability model outlined in Section 6.4.3 is designed to address this through commercial incentives rather than goodwill, but the critical mass problem remains real. The commons needs enough early contributors to demonstrate value before it can attract the broader participation that makes it valuable. Seeding strategies, contribution incentives, and governance structures that make contribution feel worthwhile rather than burdensome are design challenges that require dedicated attention beyond what this thesis can provide.

Pakistan is not a homogeneous entity. The commons proposed in this research would need to address the significant linguistic and cultural diversity within Pakistan. Punjabi, Sindhi, Balochi, Pashtun, and other communities whose knowledge and languages are not captured by an Urdu centred model would need to be represented from the earliest stages of design. A holistic pluriversal intervention must be pluriversal within Pakistan as well as between Pakistan and the global AI infrastructure it is positioned against.

There is also a risk of extraction. An open source repository of Pakistani cultural knowledge could be scraped by major AI companies and absorbed into future training data without credit, compensation, or community consent. The governance model needs to include data licensing that permits community use while requiring negotiation and attribution for commercial applications. This is a design challenge the initiative shares with Masakhane and other open source data commons, and it requires ongoing attention rather than a one-time structural solution.

Future research should address the implementation and evaluation of the proposed commons, the governance models that best protect community ownership, the representation of Pakistan's internal linguistic diversity, and the applicability of this model to comparable communities across the Global South. The ongoing development of locally built Pakistani language models provides a particularly promising site for longitudinal research into how community-built knowledge infrastructure and institutional AI development interact. Does a community-led commons accelerate the cultural accuracy of locally developed models? Does institutional involvement help or constrain community ownership over time? These are questions that can only be answered through sustained engagement.

6.6 Concluding Remarks

This thesis set out to understand how LLM integration in Pakistani organisations reproduces or challenges coloniality. The answer is clear. It reproduces it through the ordinary, pragmatic, commercially rational reasoning of technically capable professionals who understand their situation and have no structural invitation to act otherwise. Ignorance

is not the mechanism; neither is malice. The mechanism is the absence of any framework through which recognition can become action.

That is both the problem and the starting point. The professionals in this study are not passive victims of a system they cannot see. They are active, reflective practitioners who can name the conditions they are working within. What they lack is not awareness but infrastructure. The recognition is there. The transformation is not. The Cultural Knowledge Commons is designed to be the infrastructure that makes transformation possible, not as an individual professional choice, not as a corporate responsibility initiative, but as a collective, community owned, decolonially grounded response to a structural problem.

The contributions of this thesis are threefold. Empirically, it provides qualitative evidence of how coloniality operates at the level of everyday LLM use in a postcolonial professional context, evidence that is currently scarce in the AI governance and decolonial AI literature, which tends toward theoretical frameworks rather than situated practice. Analytically, it develops the concept of recognition without transformation as a specific mechanism through which structural conditions are reproduced: not through denial but through the absorption of critique into professional pragmatism. Practically, it proposes an open source intervention that responds to those conditions with a replicable model for community-led AI knowledge sovereignty.

The commons will not resolve the colonial conditions of global AI infrastructure on its own. No single intervention can. But it offers something the participants in this study did not have: a framework for acting otherwise, together, in a way that builds toward a different arrangement rather than accommodating the current one. That is, at minimum, where a structural invitation begins.

References

- Arif, M., Azeemi, S. A., Raza, H., & Athar, A. (2024). Generalists vs. specialists: Evaluating large language models for Urdu. In *Findings of the Association for Computational Linguistics: EMNLP 2024* (pp. 7263–7280). Association for Computational Linguistics. <https://doi.org/10.18653/v1/2024.findings-emnlp.426>
- Baily, M. N., Byrne, D. M., Kane, A. T., & Soto, P. E. (2025). *Generative AI at the crossroads: Light bulb, dynamo, or microscope?* Brookings Institution. <https://www.brookings.edu/articles/generative-ai-at-the-crossroads-light-bulb-dynamo-or-microscope/>
- Bender, E. M., Gebru, T., McMillan-Major, A., & Shmitchell, S. (2021). On the dangers of stochastic parrots: Can language models be too big? In *Proceedings of the 2021 ACM Conference on Fairness, Accountability, and Transparency* (pp. 610–623). ACM. <https://doi.org/10.1145/3442188.3445922>
- Benjamin, R. (2019). *Race after technology: Abolitionist tools for the new Jim code*. Polity Press.
- Birhane, A. (2020). Algorithmic colonization of Africa. *SCRIPTed*, 17(2), 389–409. <https://doi.org/10.2966/scrip.170220.389>
- Bourke, B. (2014). Positionality: Reflecting on the research process. *The Qualitative Report*, 19(33), 1–9.
- Braun, V., & Clarke, V. (2006). Using thematic analysis in psychology. *Qualitative Research in Psychology*, 3(2), 77–101. <https://doi.org/10.1191/1478088706qp063oa>
- Brinkmann, S., & Kvale, S. (2015). *InterViews: Learning the craft of qualitative research interviewing* (3rd ed.). SAGE.
- Brown, G. D. A., & Gummerum, M. (2025). Trust issues: Adolescents' epistemic vigilance towards online sources. *British Journal of Developmental Psychology*. <https://doi.org/10.1111/bjdp.12559>

- Brynjolfsson, E., Li, D., & Raymond, L. (2025). Generative AI at work. *The Quarterly Journal of Economics*, 140(2), 889–942. <https://doi.org/10.1093/qje/qjae044>
- Collins, P. H. (2000). *Black feminist thought: Knowledge, consciousness, and the politics of empowerment* (2nd ed.). Routledge.
- Couldry, N., & Mejias, U. A. (2019). *The costs of connection: How data is colonizing human life and appropriating it for capitalism*. Stanford University Press.
- Crawford, K. (2021). *Atlas of AI: Power, politics, and the planetary costs of artificial intelligence*. Yale University Press.
- Crenshaw, K. (1989). Demarginalizing the intersection of race and sex: A Black feminist critique of antidiscrimination doctrine, feminist theory and antiracist politics. *University of Chicago Legal Forum*, 1989(1), 139–167.
- Creswell, J. W. (2013). *Qualitative inquiry and research design: Choosing among five approaches* (3rd ed.). SAGE.
- DeWitt Prat, L., Ndlovu Lucas, O. N., Goliias, C., & Lewis, M. (2024). Decolonizing LLMs: An ethnographic framework for AI in African contexts. *Ethnographic Praxis in Industry Conference Proceedings*. <https://doi.org/10.1111/epic.12196>
- Digital Rights Foundation. (2025). *National AI Policy 2025 analysis*. Digital Rights Foundation.
- Dolmage, J. T. (2017). *Academic ableism: Disability and higher education*. University of Michigan Press.
- Escobar, A. (2018). *Designs for the pluriverse: Radical interdependence, autonomy, and the making of worlds*. Duke University Press.
- Fairclough, N. (1995). *Critical discourse analysis: The critical study of language*. Longman.
- Fricker, M. (2007). *Epistemic injustice: Power and the ethics of knowing*. Oxford University Press.
- Galindez-Acosta, J. S., & Giraldo-Huertas, J. J. (2025). Trust in AI emerges from distrust in humans: A machine learning study on decision-making guidance. arXiv. <https://arxiv.org/abs/2511.16769>

Gani, J. K., & Khan, R. M. (2024). Positionality statements as a function of coloniality: Interrogating reflexive methodologies. *International Studies Quarterly*, 68, sqae038. <https://doi.org/10.1093/isq/sqae038>

Giunchiglia, F., Koch, G., Bella, G., & Helm, P. (2024). Diversity and language technology: How language modeling bias causes epistemic injustice. *Ethics and Information Technology*, 26(1), Article 8. <https://doi.org/10.1007/s10676-023-09742-6>

Government of Pakistan. (2025). National AI Policy 2025. Ministry of Information Technology and Telecommunication.

Grosfoguel, R. (2007). The epistemic decolonial turn: Beyond political-economy paradigms. *Cultural Studies*, 21(2–3), 211–223. <https://doi.org/10.1080/09502380601162514>

Hafeez, M., & Sharif, I. (2025). From digital transformation to innovation: The sequential role of agility and AI adoption in Pakistan's telecom and manufacturing sectors. *The Critical Review of Social Sciences Studies*. <https://www.semanticscholar.org/paper/From-Digital-Transformation-to-Innovation%3A-The-Role-Hafeez-Sharif/a646ba6d1bf334825f23d5f17b6c893f3ead6a8c>

Haraway, D. (1988). Situated knowledges: The science question in feminism and the privilege of partial perspective. *Feminist Studies*, 14(3), 575–599. <https://doi.org/10.2307/3178066>

Heersmink, R., de Rooij, B., Clavel Vazquez, M. J., & Colombo, M. (2024). A phenomenology and epistemology of large language models: Transparency, trust, and trustworthiness. *Ethics and Information Technology*, 26, Article 41. <https://doi.org/10.1007/s10676-024-09777-3>

Holmes, A. G. D. (2020). Researcher positionality: A consideration of its influence and place in qualitative research. *Shanlax International Journal of Education*, 8(4), 1–10.

Islam, W. (2025). *Will AI transform Pakistan? Assessing the 2025 National Policy* (PIDE Knowledge Brief No. 2025:136). Pakistan Institute of Development Economics.

- Jones, P. H. (2014). Systemic design principles for complex social systems. In G. Metcalf (Ed.), *Social systems and design* (pp. 91–128). Springer. https://doi.org/10.1007/978-4-431-54478-4_4
- Kim, S. S. Y., Liao, Q. V., Vorvoreanu, M., Ballard, S., & Vaughan, J. W. (2024). "I'm not sure, but...": Examining the impact of large language models' uncertainty expression on user reliance and trust. *Proceedings of the 2024 ACM Conference on Fairness, Accountability, and Transparency*, 822–835. <https://doi.org/10.1145/3630106.3658941>
- McKinsey & Company. (2025). *Superagency in the workplace: Empowering people to unlock AI's full potential*. McKinsey & Company. <https://www.mckinsey.com/capabilities/tech-and-ai/our-insights/superagency-in-the-workplace-empowering-people-to-unlock-ais-full-potential-at-work>
- Microsoft Research. (2024). *Appropriate reliance on generative AI: Research synthesis*. Microsoft. <https://www.microsoft.com/en-us/research/publication/appropriate-reliance-on-generative-ai-research-synthesis/>
- Mignolo, W. D. (2009). Epistemic disobedience, independent thought and decolonial freedom. *Theory, Culture & Society*, 26(7–8), 159–181. <https://doi.org/10.1177/0263276409349275>
- Mignolo, W. D. (2011). *The darker side of Western modernity: Global futures, decolonial options*. Duke University Press.
- Mohamed, S., Png, M. T., & Isaac, W. (2020). Decolonial AI: Decolonial theory as sociotechnical foresight in artificial intelligence. *Philosophy & Technology*, 33(4), 659–684. <https://doi.org/10.1007/s13347-020-00405-8>
- Muldoon, J., & Wu, B. A. (2023). Artificial intelligence in the colonial matrix of power. *Philosophy & Technology*, 36, Article 80. <https://doi.org/10.1007/s13347-023-00687-8>
- Ndlovu-Gatsheni, S. J. (2017). Decolonising research methodology must include undoing its dirty history. *Journal of Public Administration*, 52 (Special Issue 1), 686–699.
- Ngũgĩ wa Thiong'o. (1986). *Decolonising the mind: The politics of language in African literature*. James Currey.

- Noble, S. U. (2018). *Algorithms of oppression: How search engines reinforce racism*. New York University Press.
- Patton, M. Q. (2002). *Qualitative research and evaluation methods* (3rd ed.). SAGE.
- Pennycook, A. (1994). *The cultural politics of English as an international language*. Longman.
- Pillow, W. S. (2003). Confession, catharsis, or cure? Rethinking the uses of reflexivity as methodological power in qualitative research. *International Journal of Qualitative Studies in Education*, 16(2), 175–196. <https://doi.org/10.1080/0951839032000060635>
- Quijano, A. (2000). Coloniality of power and Eurocentrism in Latin America. *International Sociology*, 15(2), 215–232. <https://doi.org/10.1177/0268580900015002005>
- Rahman, T. (1996). *Language and politics in Pakistan*. Oxford University Press.
- Sanders, E. B.-N., & Stappers, P. J. (2008). Co-creation and the new landscapes of design. *CoDesign*, 4(1), 5–18. <https://doi.org/10.1080/15710880701875068>
- Saunders, M., Lewis, P., & Thornhill, A. (2019). *Research methods for business students* (8th ed.). Pearson.
- Schwartz-Shea, P., & Yanow, D. (2012). *Interpretive research design: Concepts and processes*. Routledge.
- Smith, L. T. (2012). *Decolonizing methodologies: Research and indigenous peoples* (2nd ed.). Zed Books.
- Stickdorn, M., Hormess, M., Lawrence, A., & Schneider, J. (2018). *This is service design doing*. O'Reilly Media.
- Tahir, A., Shams, R., Fiaz, M., Adeeba, F., & Hussain, S. (2025). Benchmarking pre-trained LLMs across Urdu NLP tasks. In *Proceedings of the First Workshop on Challenges in Processing and Understanding Low-Resource Asian Languages (CHiPSAL 2025)* (pp. 17–34). Association for Computational Linguistics.
- Tao, Y., Viberg, O., Baker, R. S., & Kizilcec, R. F. (2024). Cultural bias and cultural alignment of large language models. *PNAS Nexus*, 3(9), pgae346. <https://doi.org/10.1093/pnasnexus/pgae346>

Timmermans, S., & Tavory, I. (2012). Theory construction in qualitative research: From grounded theory to abductive analysis. *Sociological Theory*, 30(3), 167–186. <https://doi.org/10.1177/0735275112457914>

van den Broek, E. (2025). Unpacking AI at work: Data work, knowledge work, and values work. *Information and Organization*. <https://doi.org/10.1016/j.infoandorg.2025.100584>

Wallerstein, I. (1974). *The modern world-system: Capitalist agriculture and the origins of the European world-economy in the sixteenth century*. Academic Press.

Woodruff, A., Shelby, R., Kelley, P. G., Rousso-Schindler, S., Smith-Loud, J., & Wilcox, L. (2024). How knowledge workers think generative AI will (not) transform their industries. *Proceedings of the 2024 CHI Conference on Human Factors in Computing Systems*, Article 642. <https://doi.org/10.1145/3613904.3642700>

Xu, Y., Hu, L., Zhao, J., Qiu, Z., Ye, Y., & Gu, H., Xu, K. (2025). A survey on multilingual large language models: corpora, alignment, and bias. *Frontiers of Computer Science*. <https://doi.org/10.1007/s11704-024-40579-4>

Yin, R. K. (2018). *Case study research and applications: Design and methods* (6th ed.). SAGE.

Appendix 1

Semi-Structured Interview Guide

Duration: 60 minutes

Format: Semi-structured interview with interactive elicitation (FigJam-based activities)

Introduction and Relational Framing (5 min)

Thank you sincerely for taking the time to participate in this study.

This research explores how organizations in Pakistan integrate Large Language Models (LLMs) into strategic and operational practices. The study adopts a critical perspective attentive to how global AI infrastructures intersect with local knowledge systems, language practices, and organizational decision-making.

The focus is not on evaluating your organization, nor on assessing technical proficiency. Rather, it is an exploration of broader structural dynamics shaping AI integration.

With your consent:

- The interview will be recorded for transcription purposes.
- You will receive a verbatim transcript for review.
- You may edit, clarify, or withdraw any part before inclusion in analysis.
- Participation is voluntary. You may request anonymity at any point. If you consent to being identified, your name and role will be attributed in the final thesis.

This conversation is intended to be dialogical and reflective. Please feel free to ask questions at any point.

Section 1: Professional Context and Organizational Positioning (7 min)

1. Could you describe your role within the organization and how you use Large Language Models in your work?
 2. When did your organization begin integrating LLMs into workflows, and what motivated their adoption?
 3. What kinds of tasks or decisions are currently supported by LLMs in your organization?
-

Section 2: AI Adoption and Infrastructural Orientation (7 min)

4. How did your organization decide which LLM platform or provider to use? Who was involved in that decision-making process?

5. Were considerations such as data hosting, platform ownership, geopolitical origin, or long-term infrastructural dependence discussed?
6. Were alternative models (e.g., open-source or locally developed systems) considered? Why or why not?

Section 3: Interactive Activity 1 – Knowledge Orientation Mapping (FigJam) (9 min)

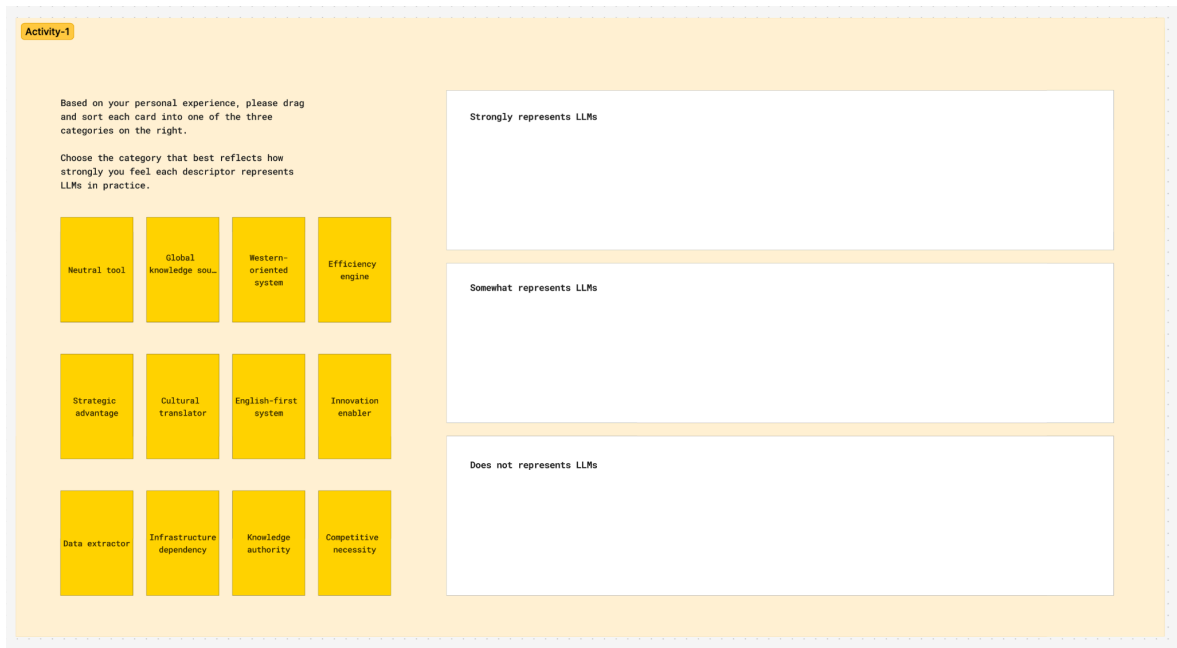


Figure A.1 Knowledge Orientation Mapping activity template

Participants are presented with digital cards containing the following descriptors:

- Neutral tool
- Global knowledge source
- Western-oriented system
- Efficiency engine
- Strategic advantage
- Cultural translator
- English-first system
- Innovation enabler
- Data extractor
- Infrastructure dependency
- Knowledge authority
- Competitive necessity

Prompt:

Please sort these cards into three categories:

- Strongly represents LLMs in your experience
- Somewhat represents LLMs

- Does not represent LLMs

7. Which cards were easiest or most difficult to place? Why?

8. Are there any descriptors missing that you would add?

Section 4: Epistemic Authority and Trust (7 min)

9. How do you evaluate whether an LLM-generated output is reliable, accurate, or strategically sound?

10. To what extent do you trust LLM outputs in important or strategic decisions?

11. Have you had to significantly edit or override outputs? In what situations?

Section 5: Interactive Activity 2 – Dependency Spectrum (7 min)

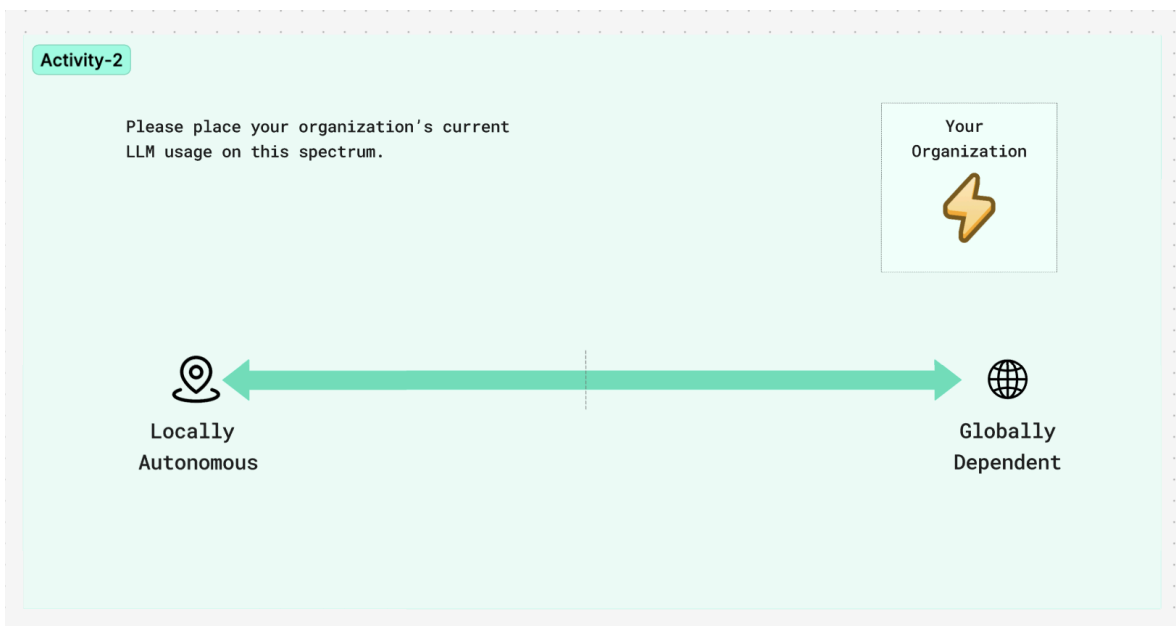


Figure A.2 *Dependency Spectrum activity template*

Participants are shown a horizontal scale:

Locally Autonomous ←————→ **Globally Dependent**

Prompt:

Please place your organization's current LLM usage on this spectrum.

12. Why did you position it there?

13. Does this level of dependence feel strategic, neutral, temporary, or structural?

14. How would your organization be affected if access to current LLM tools were restricted?

Section 6: Language and Cultural Mediation (7 min)

15. In which language(s) do you primarily interact with LLMs, and why?

16. Have you noticed differences in output quality between English and local language outputs? If so, how do these differences influence your workflow?

Section 7: Interactive Activity 3 – Language Hierarchy Mapping (3min)

Activity-3

Step 1:
Resize the three circles below to reflect how dominant each language is in your organization's use of LLMs. Make a circle larger if that language is used more frequently.

Step 2:
Draw arrows between circles to show where translation, adaptation, or modification happens.
For example:
• If outputs are generated in English and then adapted into a local language, draw an arrow to show that flow.
• If ideas move back and forth, show that too.

English local language Hybrid/mixed usage

Figure A.3 *Language Hierarchy Mapping activity template*

Participants are shown three circles labeled:

- English
- Local language(s)
- Hybrid/mixed usage

They are asked to:

- Resize circles based on relative dominance in their organization's AI usage.

- Draw arrows indicating where translation, adaptation, or modification occurs.

17. What does this mapping reveal about how knowledge flows through your organization when using LLMs?

Section 8: Critical Reflection and Structural Awareness (5 min)

18. Most major AI systems are developed in a few powerful countries and then used globally. Do you think that influences the type of knowledge or perspectives they reflect?

19. Do these tools feel culturally neutral to you, or do they reflect particular assumptions or ways of thinking?

20. How do LLMs influence the way your organization frames strategy, communication, or problem-solving?

Closing Reflection (3 min)

21. Looking ahead, how do you see the role of LLMs evolving within your organization?

22. Is there anything we have not discussed that you believe is important in understanding your organization's relationship with these technologies?

Before we wrap up

- Is there anything we didn't touch on that you think is important?
- Is there anyone else — a colleague or another organization — that you think I should speak with?

If it's okay with you:

- May I follow up by email if I need any clarification or if something comes up during analysis?
- I'll send you the full transcript of our conversation within 24 hours. Please feel free to edit, clarify, or remove anything before it's used in the study.

Thank you sincerely for your time, insights, and openness. I deeply appreciate the care and reflection you brought to this conversation.

Appendix 2

Thematic Analysis Overview

This table presents the analytical pathway from descriptive initial codes to candidate themes and final themes, following Braun and Clarke's (2006) six-phase thematic analysis process. Initial codes were generated inductively from five interview transcripts and three FigJam activity outputs. Codes were subsequently grouped into candidate themes, then reviewed and revised through repeated engagement with the full dataset. The four main themes correspond to Sections 5.3–5.6 of the Findings chapter. Additional emerging themes are discussed in Section 5.7.

Initial Code	Candidate Theme	Main Theme
<i>Uses word 'colonial' unprompted during language activity; frames it as self-evident</i>	Direct structural naming	Theme 1 Between Neutral Tool and Colonial Infrastructure
<i>Acknowledges non-neutrality; brackets it as professionally irrelevant through maths analogy</i>	Bracketing strategy	
<i>Identifies bias as regional/geopolitical; names it matter-of-factly, takes no action</i>	Incidental acknowledgment	
<i>Describes LLMs as neutral because output reflects prompt control; structure invisible</i>	Agency-based neutralisation	
<i>Discovers bias through testing ChatGPT; frames as product defect, not structural condition</i>	Product defect framing	
<i>Describes adoption as following global competitors; organisational choice is effectively absent</i>	Competitive mimicry / structural necessity	Theme 2 Navigating Dependency
<i>Quantifies team dependency at ~75%; describes as strategic and temporary</i>	Temporal deferral	
<i>Recounts Claude outage and operational paralysis; compares to laptop crash</i>	Geopolitical vulnerability	
<i>Acknowledges losing ability to think clearly through AI use; frames as individual trade-off</i>	Trade-off management	

<i>Frames LLM relationship as bidirectional; humans add emotion and judgment; restores sense of agency</i>	Agency restoration through bidirectionality	
<i>Switches to English for better outputs; explains gap through 'ecosystem' richness</i>	Quality differential	Theme 3 English as Infrastructure
<i>Inputs informal/local language; outputs formal English; calls this formalisation</i>	Formalisation pipeline	
<i>'Everything goes into English' — beneficiary stories translated for global readership</i>	Unidirectional knowledge flow	
<i>Flags Urdu-Hindi confusion in LLM outputs; frames as developmental immaturity, not structural marginalisation</i>	Developmental framing	
<i>Organisation-wide pattern: hybrid input enters, English formal register exits</i>	Language hierarchy reproduction	
<i>Names structural gap accurately; expects provider (OpenAI) to resolve it</i>	Responsibility displacement	
<i>Reframes cognitive atrophy as efficiency trade-off; critique absorbed into business logic</i>	Critique absorption	
<i>Describes editing, QA, and reworking of AI output as normal professional practice; labour invisible</i>	Labour invisibilisation	
<i>Imagines AI futures as extensions of present arrangement; no alternative infrastructure surfaced</i>	Structural foreclosure	
<i>Designers frame AI use as cheating and inauthentic; CEO validates critique, proceeds regardless</i>	Creative authorship and professional identity	Additional Emerging Themes (Section 5.7)
<i>Supervisor expects AI-assisted copy in 30 minutes vs. 3–4 hours; creative identity quietly negotiated</i>	Creative authorship and professional identity	
<i>'Left behind' framing for non-adoption used by three participants; survival language frames choice as absent</i>	Competitive logic of AI adoption	
<i>AI-native mandate embedded in performance management at an AI-native incubator</i>	Competitive logic of AI adoption	

<i>Futures imagined as more integration, smaller teams; present dependency extended not renegotiated</i>	Future AI imaginaries	
<i>Local alternative considered only if performance-equivalent; epistemic sovereignty not a criterion</i>	Future AI imaginaries	

Note. Initial codes are presented in italics and reflect the language of participants where possible. Candidate themes are descriptive groupings formed during analysis. Main themes represent the final analytical categories as named in Chapter 5.