



The Last Colonization: *Inside Your Mind – A Mycelial Invasion?*

*An Issue Brief on Cognitive Autonomy,
Data Extraction, and the Governance
of Generative AI*


Muhammad Firas Shams

Introduction

FIGURE 1.

Three Interlocking Dynamics

Analytical framework: extraction, dependency and homogenization.



When you hear about Artificial Intelligence colonizing the mind, it seems like some sci-fi apocalypse coming to life. It seems like a literal land colonization, but inside your head. The AI isn't just visiting - it's surveying, dividing, building and owning. Left to right, you watch your mind get subdivided and industrialized, seems like an unsettling matrix to dwell in. Historically, colonization meant something concrete: external powers seizing land and labor, then reordering local life around a distant metropole. Applying that loaded word to artificial intelligence's grip on human thought is a riskier move, one that could tip into rhetorical excess or land on genuine insight. This brief takes the second view. Used here, "Colonization" is no flourish borrowed to dramatize generative AI's effects on cognition; it names a structural pattern with identifiable mechanisms, empirical evidence, and policy implications of its own, implications distinct from the familiar debates over job losses or algorithmic bias.

Three interlocking dynamics back this claim up. First, extraction: human experience and expression get captured as data, then refined into products sold back into the very discourse they came from, a pattern media scholars call data colonialism. Second, dependency. As generative systems move from storing the products of thought to producing substantive linguistic and analytical output themselves, habitual reliance starts raising real questions about memory, engagement, and critical

judgment. Third, and more troubling: homogenization. Recent scholarship argues AI outputs get reabsorbed into human communication and reasoning, gradually narrowing variation around whatever linguistic and conceptual patterns dominate.

This brief is policy-oriented, not philosophical. To that end, it gives policymakers, educators, and institutional leaders a defensible framework for assessing risk and designing safeguards, rather than settling any final verdict on AI's inherent nature. Three parts follow. First, background on data colonialism, surveillance capitalism, and cognitive offloading. Second, a discussion of the empirical and theoretical evidence for mind-level colonization, including how unevenly it lands globally and the strongest objections raised against the thesis. Third, a conclusion flagging areas ripe for regulatory and institutional attention, Pakistan among them, where AI adoption is outrunning the governance infrastructure meant to keep pace.

Background

“Colonization” entered AI discourse through two related but distinct literatures. Data colonialism is the first. Media scholars Nick Couldry and Ulises Mejias coined the term, arguing that digital platforms have opened a new phase of capitalism, one in which human life itself, not just land or labor, becomes raw material for extraction.¹ Historical colonialism treated inhabited land as terra nullius, “empty” and available for appropriation. Contemporary data relations do something similar: personal experience becomes a resource that is “just there” for capture, needing only the infrastructure to collect it. Couldry and Mejias make the point explicit, stating that data colonialism “involves not one pole of colonial power (‘the West’), but at least two: the USA and China.” That framing complicates older Global North/Global South mappings of colonizer and colonized, even as the underlying extractive logic persists.²

Shoshana Zuboff supplies the economic engine behind this appropriation: surveillance capitalism. Surveillance capitalists, in her account, unilaterally claim human experience as free raw material. Then they translate it into behavioral data. The resulting surplus gets refined into prediction products, traded on what she calls behavioral futures markets.³ Out of this comes a new form of power, instrumentarian power, which, in her words, “knows and shapes human behavior toward others’ ends” through an increasingly ubiquitous computational architecture. Twentieth-century critiques of media worried mainly about persuasion and propaganda. Zuboff’s framework points somewhere more structural and less visible: an economy whose profitability rests on the continuous, fine-grained modification of behavior at scale.

Decolonial AI runs parallel to this, focused more directly on AI itself. It extends the analysis to how AI systems get designed and deployed. Shakir Mohamed, Marie-Therese Png, and William Isaac argue that AI systems can reproduce colonial patterns of power even without any historical colonial relationship in place, through what they term algorithmic oppression, algorithmic exploitation, and algorithmic dispossession.⁴ Their point isn’t that AI developers intend colonial outcomes. It’s that the technology inherits an unexamined “coloniality of power”, cultural assumptions and institutional arrangements carried over from earlier eras, and reproduces them at new scale, speed, and reach.

A different literature sits alongside these political-economic critiques: cognitive science’s long-running work on offloading, which predates generative AI by decades. Psychologists Evan Risko and Sam Gilbert define cognitive offloading as “the use of physical action to alter the information processing requirements of a task so as to reduce cognitive demand”, behavior as ordinary as jotting down a phone number or setting a calendar reminder.⁵ Calculators, address books, search engines: all offloading technologies, and none of them gets called mind-colonizing. Generative AI is different because it shifts from storage to synthesis. A calculator performs a bounded operation on numbers the user still supplies and understands; a large language model produces the reasoning itself, drafting arguments, writing code, rendering judgments on request. Recording thought is one thing. Manufacturing it is another, and that shift is the pivot on which the “colonization of the mind” argument turns. It’s this shift, more than the mere fact of AI assistance, that the rest of this brief examines.

¹Couldry and Mejias, *Costs of Connection*.

²Nick Couldry and Ulises A. Mejias, “Data Colonialism: Rethinking Big Data’s Relation to the Contemporary Subject,” *Television & New Media* 20, no. 4 (2019): 336–349, <https://doi.org/10.1177/1527476418796632>.

³Shoshana Zuboff, *The Age of Surveillance Capitalism: The Fight for a Human Future at the New Frontier of Power* (New York: PublicAffairs, 2019), 8.

⁴Mohamed, Png, and Isaac, “Decolonial AI,” 666.

⁵Risko and Gilbert, “Cognitive Offloading,” 676.

Discussion

From Data Extraction to Cognitive Extraction

Data colonialism explains how personal experience gets captured as data. A newer line of scholarship picks up from there, asking what happens once AI systems feed their outputs back into human communication and reasoning. A 2026 review in *Trends in Cognitive Sciences* synthesizes the evidence: large language models can promote stylistic and conceptual homogenization, mirroring dominant patterns while quietly downplaying alternative voices and reasoning strategies. Repeated reliance on the same handful of models, the review argues, may amplify that convergence further, across language, perspective, and reasoning alike.

A further policy concern follows. If users keep adopting model-generated formulations without independently developing or testing alternatives, convergence could run deeper than surface-level style. Treat that as a governance hypothesis, not an established causal finding. Still, it matters for cognitive autonomy: repeated delegation can quietly shift where judgment and meaning-making authority actually sit.

Experimental work backs part of this mechanism. Researchers ran a controlled study with 1,506 participants: co-writing with a language model set to argue for a particular position affected what participants said, and shifted their subsequent attitudes toward the model's position. That's evidence AI-generated framing can influence views. It is not, on its own, evidence of a general or permanent change in belief.

FIGURE 2.

A Precautionary Governance Response

1 AI LITERACY



Teach when to use AI—and when deliberate non-use protects learning and judgment.

2 COGNITIVE IMPACT



Assess cognitive and epistemic effects alongside safety, privacy and bias.

3 LOCAL AGENCY



Support systems, data practices and governance that reflect local languages and institutions.

The Empirical Record on Cognitive Debt

A more direct question sits beyond attitude change: does habitual reliance on generative AI touch the cognitive processes users would otherwise handle themselves? One early study, from the MIT Media Lab, put this to the test. Participants wrote essays with an LLM, with a search engine, or with no external tool at all. Using EEG and behavioral measures, the researchers found differences in neural connectivity, essay characteristics, recall, and reported ownership across the three groups. They called the pattern “cognitive debt.” It’s a preprint, and a small one, 54 participants across the first three sessions, only 18 completing the crossover, so read the findings as preliminary. They don’t show that generative AI generally causes cognitive decline.

Even so, the study grounds the policy question in something concrete. Deliberate, bounded offloading, checking arithmetic, storing a reminder, is not the same as habitually delegating substantive cognitive work. Offloading itself carries no inherent harm; established literature shows external tools reducing cognitive demand and helping people work around the limits of memory and processing. What remains unresolved is the tipping point: the moment AI assistance stops scaffolding learning and starts replacing the practice through which users build real expertise.

Whose Mind, Whose Colonizer?

A serious account of AI's relationship to cognition can't treat "the human mind" as one undifferentiated territory. The decolonial AI literature insists on exactly this point. The same system offering convenience to a well-resourced professional can impose interpretive frameworks, linguistic norms, and value assumptions built almost entirely inside a handful of English-language, Silicon Valley-adjacent institutions, imposed on users elsewhere who had essentially no voice in the system's design.⁶ Here, "colonization" earns its name rather than functioning as mere metaphor. It's an asymmetry of power: one party sets the terms, extracts the value, exports a worldview; the other supplies the raw material and absorbs the consequences.

For policymakers in low- and middle-income countries, this asymmetry piles onto the cognitive-offloading concerns already described. Every population using these tools faces those concerns, regardless of geography. A further governance question emerges on top of them: which values, dialects, and reasoning styles become the unmarked default "voice" a model produces? The decolonial AI literature frames such questions squarely as issues of power, representation, and whose perspectives get built into sociotechnical systems.

Contested Terrain: Grounds for Caution about the Colonization Frame

The colonization framing has genuine critics, and a fair-minded brief owes them its strongest objections. First: colonialism's historical specificity, its violence, its formal political domination, its racialized legal structures, does real work that a loose extension to software risks flattening. Critics worry that invoking colonization for every lopsided technology relationship dilutes the term, weakening its usefulness for naming actual colonial and neocolonial harms. Second, much of the cognitive-debt evidence remains preliminary. The MIT study is a preprint, its sample modest, and its own authors caution against overreading the results as evidence of general "brain rot" rather than a specific, task-bound reduction in neural engagement.⁷ Third, cognitive offloading has a long history, and mostly a beneficial one. The same freed capacity a skeptic calls dependency, a defender calls the ordinary, productive division of cognitive labor, the kind literacy, print, and computing have each extended in turn.⁸

These objections temper the claim, but they don't break it. The evidence comes from different kinds of scholarship, and no one strand should be mistaken for a single causal demonstration. Political-economy and decolonial work supply theoretical accounts of extraction and power. Jakesch et al. supply experimental evidence of opinion influence. Kosmyna et al. supply preliminary evidence on cognitive engagement. Sourati et al. synthesize evidence on homogenization. Together, these strands justify treating cognitive autonomy and epistemic diversity as governance questions worth precautionary attention, while leaving the stronger claim of "colonization" open to continued empirical testing.

⁶Mohamed, Png, and Isaac, "Decolonial AI," 666.

⁷Kosmyna et al., "Your Brain on ChatGPT."

⁸Risko and Gilbert, "Cognitive Offloading."

Toward a Governance Response

International bodies have started grappling with cognitive-level risks, though existing instruments aren't yet calibrated to this specific question with much precision. UNESCO's 2021 Recommendation on the Ethics of Artificial Intelligence, adopted by 193 Member States, puts human rights and human dignity at the centre of AI governance, alongside transparency, fairness, environmental sustainability, and human oversight across the AI lifecycle. Its policy action areas, education, data governance, communication and information, offer a foundation that could extend to cognitive-autonomy concerns directly. AI-literacy curricula could teach not just how to use generative tools, but when deliberate non-use makes more sense. Disclosure or impact-assessment requirements could apply wherever AI materially shapes institutional decisions or submissions.

A responsible governance agenda resists two temptations. One is technological fatalism, treating cognitive colonization as just another inevitable by-product of progress, no different from earlier waves of tool adoption. The historical record on cognitive offloading complicates that comparison: external tools can free mental capacity without necessarily carrying the same risks as delegating synthesis, judgment, or persuasion. The other temptation is moral panic, which forecloses AI's genuine benefits in education, scientific discovery, and access to expert knowledge. A more defensible position sits in between: precaution. The cognitive effects of generative AI aren't fully mapped yet, but the evidence already available justifies governance measures that preserve human agency, encourage critical engagement, and prevent irreversible dependence while the evidence keeps developing.

Conclusion

Is artificial intelligence colonizing the human mind? The evidence in this brief doesn't settle that as empirical fact. It does identify something more modest but still defensible: a structural concern. AI systems extract and process human-produced data. Generative systems influence users' expressed views. Early experimental work raises real questions about cognitive engagement during AI-assisted writing. Recent scholarship flags risks of linguistic and conceptual homogenization. None of these effects is uniform, inevitable, or automatically harmful in every case. The policy question that remains is how to preserve human cognitive autonomy while letting AI augment human judgment rather than routinely replace it.

Three recommendations follow for policymakers. First, education systems should treat AI literacy as inseparable from critical-thinking pedagogy, teaching students when offloading helps them, and when it substitutes for skills they still need to build. Second, regulators should widen existing AI-ethics frameworks, UNESCO's Recommendation included, to weigh cognitive and epistemic effects alongside bias, safety, privacy, transparency, and human oversight. Third, governments outside AI's principal development centres, Pakistan among them, should invest in AI systems and governance arrangements built around local languages, institutions, and values. That's how the asymmetries decolonial AI literature identifies get addressed at the level of design and data governance, rather than patched up after deployment.

Bibliography

- Couldry, Nick, and Ulises A. Mejias. *The Costs of Connection: How Data Is Colonizing Human Life and Appropriating It for Capitalism*. Stanford, CA: Stanford University Press, 2019.
- Couldry, Nick, and Ulises A. Mejias. "Data Colonialism: Rethinking Big Data's Relation to the Contemporary Subject." *Television & New Media* 20, no. 4 (2019): 336–349.
<https://doi.org/10.1177/1527476418796632>.
- Jakesch, Maurice, Advait Bhat, Daniel Buschek, Lior Zalmanson, and Mor Naaman. "Co-Writing with Opinionated Language Models Affects Users' Views." In *Proceedings of the 2023 CHI Conference on Human Factors in Computing Systems*, Article 111. New York: Association for Computing Machinery, 2023. <https://doi.org/10.1145/3544548.3581196>.
- Kosmyna, Nataliya, Eugene Hauptmann, Ye Tong Yuan, Jessica Situ, Xian-Hao Liao, Ashly Vivian Beresnitzky, Iris Braunstein, and Pattie Maes. "Your Brain on ChatGPT: Accumulation of Cognitive Debt When Using an AI Assistant for Essay Writing Task." arXiv preprint, June 10, 2025.
<https://arxiv.org/abs/2506.08872>.
- Mohamed, Shakir, Marie-Therese Png, and William Isaac. "Decolonial AI: Decolonial Theory as Sociotechnical Foresight in Artificial Intelligence." *Philosophy & Technology* 33 (2020): 659–684.
<https://doi.org/10.1007/s13347-020-00405-8>.
- Pillay, Tharin. "Are We Losing Our Minds to AI?" *Time*, April 30, 2026.
<https://time.com/article/2026/04/30/ai-thinking-cognitive-offloading/>.
- Risko, Evan F., and Sam J. Gilbert. "Cognitive Offloading." *Trends in Cognitive Sciences* 20, no. 9 (2016): 676–688. <https://doi.org/10.1016/j.tics.2016.07.002>.
- Sourati, Zhivar, Alireza S. Ziabari, and Morteza Dehghani. "The Homogenizing Effect of Large Language Models on Human Expression and Thought." *Trends in Cognitive Sciences* 30, no. 9 (2026): 805–816.
<https://doi.org/10.1016/j.tics.2026.01.003>.
- UNESCO. *Recommendation on the Ethics of Artificial Intelligence*. Paris: UNESCO, 2021.
<https://www.unesco.org/en/artificial-intelligence/recommendation-ethics>.
- Zuboff, Shoshana. *The Age of Surveillance Capitalism: The Fight for a Human Future at the New Frontier of Power*. New York: PublicAffairs, 2019.