

Countering Epistemic Centralisation in Frontier AI: Towards Hermeneutical Justice

Hall, M. (2026). Countering Epistemic Centralisation in Frontier AI: Towards Hermeneutical Justice. Groningen Declaration Network, Cape Town, 20–22 October 2026.

“Epistemic centralisation” is a consequence of building Large Language Models dependent on pre-training datasets that disproportionately represent global knowledge sources. An emerging body of research is revealing the extent to which the design of these Frontier AI models is leading towards what has been called “generative hermeneutical erasure”, favouring the “default speaker”, a “voice from nowhere”. In turn, this is prompting the formation of regional networks committed to equitable language representation and representative training datasets.

Epistemic centralisation and the systematic injustices it causes have significant implications for the recording of academic achievement, particularly for the value of credentials recognised and exchanged across national borders, the focus of the Groningen Declaration Network. Ensuring hermeneutical justice in credentialing starts with quality assurance; the benchmarks and standards which underpin trust in teaching and learning. New benchmarks should acknowledge that Artificial Intelligence will soon be ambient, and that assessment must measure the ability to use AI appropriately and critically. The GDN is well positioned to contribute towards these goals through secure forms of digital credentialing and Comprehensive Learner Records, helping create a network of networks that recognises epistemic validity across national borders and global regions.

Countering Epistemic Centralisation in Frontier AI: Towards Hermeneutical Justice

Martin Hall

Allan Grey Centre for Values-Based Leadership, Graduate School of Business, University of Cape Town

Groningen Declaration Network Annual Meeting, Cape Town, 20-22 October 2026

Abstract

“Epistemic centralisation” is a consequence of building Large Language Models on pre-training datasets that disproportionately represent global knowledge sources. An emerging body of research is revealing the extent to which the design of these Frontier AI models is leading towards what has been called “generative hermeneutical erasure”, favouring the “default speaker”, a “voice from nowhere”. In turn, this is prompting the formation of regional networks committed to equitable language representation and representative training datasets. The global reach of a small set of commercially driven AI platforms is also pushing nations towards “Sovereign AI”, although recent surveys show that most sovereign projects depend on partnerships with a small set of US-based Big Tech companies. These trends indicate that creating regional or focused networks to address epistemic centralisation does not, in itself, ensure epistemic justice. As with all technologies, social, economic, and political factors have a determining role, expressing forms of power, grounded in the history of colonialism, that persist long after formal colonial administrations have ended.

The systemic injustices of epistemic centralisation have significant implications for the recording of academic achievement, particularly for the value of credentials recognised and exchanged across national borders, the focus of the Groningen Declaration Network. Ensuring hermeneutical justice in credentialing starts with quality assurance; the benchmarks and standards which underpin trust in teaching and learning. New benchmarks should acknowledge that Artificial Intelligence will soon be ambient, and that assessment must measure the ability to use AI appropriately and critically. The GDN is well positioned to contribute towards these goals through secure forms of digital credentialing and Comprehensive Learner Records, helping create a network of networks that recognises epistemic validity across national borders and global regions.

Epistemic Centralisation

“Epistemic centralisation”, in the context of Artificial Intelligence, is a consequence of building Large Language Models that depend on pre-training datasets that disproportionately represent global knowledge sources (Leslie & Perini, 2024). These “frontier” models - the most advanced general-purpose forms of Artificial Intelligence available at a given time - favour dominant languages and cultures in constructing their outputs and then present these outputs as definitive, authoritative and universal; the “default speaker”, an “unmarked identity that conversational systems perform when they claim to speak from nowhere” (Gaba, 2026). Unlike bias and hallucination, which are episodic, epistemic centralisation is a structural issue that leads towards “epistemicide”; the consequence of one way of knowing erasing and replacing another (De Sousa Santos, 2014).

Epistemic centralisation starts with the language inequality that is embedded in the tokenisation model used by all the major Frontier AI models. Tokenisation divides text into units, “tokens”, which may represent whole words, parts of words, punctuation or other character sequences. Each token is then given a numerical label, “an integer identifier”, which is used to prepare text for processing by the model. During training, the model learns patterns from tokenised data; when responding to a user prompt, it uses these learned patterns to generate a sequence of tokens, which is converted back into readable text as the model’s output. Because tokenization systems use English as a standard, Frontier AI models favour data sources and outputs written in English. As a result, all other languages face a premium that requires proportionately more tokens, resulting in higher costs, lower efficiency, and a significant disincentive to include non-English resources in pre-training data sets. Languages that use non-Latin scripts, such as Bangla, Hindi, Arabic, and Urdu, can increase processing costs fifteen-fold compared with English (Petrov et al., 2023; Lundin, Zhang et al., 2026).

For example, take the greeting “Hello, how are you?” In English, “Sawubona, unjani?”, in isiZulu, one of South Africa’s twelve official languages. OpenAI’s tokenizer algorithm uses six tokens to process this greeting in English, but nine tokens to capture the same message in isiZulu. In effect, this is a tariff of 50% on any textual material processed by OpenAI’s ChatGPT models in isiZulu; “an active design decision that determines who benefits from AI and who bears hidden costs” (Hasse & Pokutta 2026).

Language	Text	Token divisions	Number of Tokens
English	Hello, how are you?	Hello · , · how · are · you · ?	6
isiZulu	Sawubona, unjani?	S · aw · ub · ona · , · un · j · ani · ?	9

An emerging body of research is revealing additional ways in which the design of Frontier AI models is eroding epistemic validity. AfroBench is a set of measures for evaluating the performance of Large Language Models across 64 African languages, 15 tasks and 22 datasets. A study using these criteria found significant performance deficits when uses of these languages were compared with monolingual English. When a language is weakly represented in the underlying data, then everyday knowledge expressed primarily through that language is much less likely to become part of what a Large Language Model can retrieve and incorporate in its outputs (Ojo, Ogundepo et al., 2025). For example, curricula in many African countries incorporate information about agricultural practices that require contextual ecological knowledge, but these essential details are often excluded when curricula are generated by Frontier AI models: “This trend of epistemic erasure of African farming methods is not new and has been historically ushered in via imperialism and colonisation through industrialisation processes, but AI risks further erasing Indigenous knowledge. For example, Western rationality perceives plants as objects, whereas Indigenous Eastern African knowledge sees plants as beings with spirit and agency. Yet these East African values are often absent from AI-generated agricultural curricula” (Maisiri & Musonza, 2025).

The World Values Survey measures people’s values, beliefs and attitudes across more than 120 countries and tracks how they change over time. Questions cover a wide range of subjects, including attitudes towards politics, religion, democracy and political authority, family and marriage, gender equality, work, national identity, ethnic diversity, and subjective wellbeing (World Values Survey, 2026). A study of five Large Language Models found that, when prompted with questions from the World Values Survey, all initially generated outputs aligned with survey responses from English-speaking and Protestant European countries. But in this case, epistemic centralisation was not the consequence of inadequate training data; when provided with area-specific guidance, the models’ cultural alignment improved by between 70% and 80%, showing that Western cultural assumptions serve as the default, while other cultural positions have to be specified as contextual exceptions (Tao, Viberg et al., 2024).

Considerations of epistemic centralisation lead in turn to the principles of epistemic justice. Here, Miranda Fricker’s critical philosophy anticipates questions that are now being raised about the programming of Frontier AI models (Fricker, 2007). Fricker distinguishes between “testimonial injustice”, when a person’s credibility is diminished due to identity prejudice, and “hermeneutical injustice”, which occurs when a person cannot make sense of, or explain, their own life experience because of the lack of the words or concepts to understand it.

Applied in the context of Artificial Intelligence, testimonial injustice occurs when an AI application exhibits episodic bias across gender, ethnicity, and race (Gándara, Anahideh et al., 2024; Gándara and Anahideh, 2025). Hermeneutical injustice arises in contexts such as using AI-generated content to learn about farming, plants, and animals in African languages, or when the Protestant values of English-speaking countries are imposed as the “default speaker” in the formulation of AI outputs. Given the extraordinary reach and impact of Frontier AI models, Mollema has argued that the systematic misrepresentation of marginalised social groups’ experiences will ultimately result in what he calls “generative hermeneutical erasure” (Mollema 2025).

Countering Epistemic Centralisation

The dominance of a small set of Frontier AI models and their global reach is prompting the formation of regional networks committed to equitable language representation and representative training datasets.

LatamGPT is a regional model that brings together fifteen countries to assemble datasets that reflect regional cultures and histories in Spanish, Portuguese, and indigenous languages. The project is coordinated by Chile's National Centre for Artificial Intelligence (CENIA) (Hope, 2026; LatamGPT, 2026): "What makes Latam-GPT remarkable isn't its technical architecture, but its institutional coordination. In a region where cross-border collaboration often stalls, Latam-GPT's design reflects the region's institutional diversity. It provides a shared base that can be fine-tuned for national priorities - Argentine legal documents, Colombian health records, and Mexican curricula, among other use cases. This federated logic acknowledges that common linguistic roots coexist with wide differences in governance and user needs" (Yeyati, 2025).

Another example is SEA-LION (Southeast Asian Languages in One Network), a set of open-source Large Language Models enabled by Singapore's National Research Foundation and hosted by the National University of Singapore (SEA-LION, 2026; Ng, Nguyen et al., 2025). Southeast Asia has more than 1000 languages spoken by some 700 million inhabitants. To counter the consequences of epistemic centralisation, SEA-LION is putting together models based on Southeast Asian linguistic and cultural material. SEA-LION's open-source strategy allows other organisations in the region to offer applications without recreating a foundation model themselves: "epistemic and technological decentralisation within an interconnected global AI ecosystem" (Ng, Nguyen et al., 2025).

AI4Bharat, a project of the Indian Institute of Technology Madras, also uses open-source models as an alternative to proprietary applications and has developed a series of multilingual Large Language Models tailored for India's 22 scheduled languages and hundreds of dialects, brought together in datasets such as IndicCorpora and Sangraha (AI4Bharat, 2026; Khan, Mehta et al., 2024). In contrast to the web-scraping and copyright violations for which Frontier AI models are notorious, AI4Bharat trains its models through the authorised uploading of a large corpus of locked PDF content held in India's institutional repositories, including government reports, educational material, religious texts and literature that are authenticated source materials for the traditions, practices, and historical accounts that are unique to different parts of the country (Khan & Khapra, 2024). AI4Bharat's IndicVoices dataset also includes some 23 000 hours of speech from about 50 000 speakers (Javed, Nawale et al., 2024).

Building networks to counter epistemic centralisation is also in the interests of its perpetrators because of the risk of "model collapse", which increases as Large Language Models come to depend on synthetic and superficial data; sources created by previous and repetitive iterations that are fed back into training data as ultimately destructive feedback loops (Villalobos, Ho et al., 2024; Shumailov, Shumaylov et al., 2024). In response, Frontier AI companies are prioritising pre-training datasets that include "real-world" high-quality content obtained through networks of

organisations that create high-quality real-world data. These include some of the largest for-profit academic publishers (Oderinwale & Kazlauskas, 2025). For example, in mid-2024, Informa, the parent company of the prominent academic publisher Taylor and Francis, signed a major data-sharing deal with Microsoft without academic authors' consent, provoking substantial backlash (Jack, 2024).

Factors such as these come together as a contradiction inherent in Sovereign AI; a country's ability to build, run and govern Artificial Intelligence using local data, hardware infrastructure and software in alignment with domestic values, laws and security needs (McKinsey, 2026; Song, Tsukioka et al., 2026). On the one hand, Sovereign AI enables projects such as AI4Bharat and SEA-LION to counter epistemic centralisation at scale by assembling pre-training data sets marginalised by the first wave of large commercial AI models. But most Sovereign AI initiatives rely on commercial partnerships with Frontier AI providers. CNAS's Sovereign AI Index reports that "in practice, most 'sovereign' AI projects rely heavily on foreign - overwhelmingly American - technology providers. These foreign companies remain deeply embedded across more than 170 national infrastructure and model projects tracked by the index. More than three-fifths of tracked projects disclose at least one foreign partner, and four-fifths of these involve a U.S. company" (Chavez, Chilukuri & Scanlon, 2026).

This indicates that creating regional or focused networks to address epistemic centralisation does not, in itself, ensure epistemic justice. This is anticipated in Fricker's discussion of how injustice manifests as "hermeneutical hotspots" in collective social experiences, resulting in forms of epistemic disadvantage that are coercive rather than the consequence of individual decisions, such as the use of a particular technology. Fricker emphasises that this form of marginalisation is, by definition, an expression of collective powerlessness and exclusion from practices and opportunities that would have had value to participants: "It is only when the collective impoverishment is concretely situated in specific social situations that it comes to be especially and unjustly disadvantageous to some groups but not others. Hermeneutical lacunas are like holes in the ozone - it's the people who live under them that get burned. Fundamentally, then, hermeneutical injustice is a kind of structural discrimination" (Fricker, 2007).

This becomes evident in the case of indigenous languages and their tangible manifestations as cultural property (United Nations, 2007; Kwok, 2026). Here, the risk is that the open source principles that underpin projects such as SEA-LION, AI4Bharat, and Latam-GPT leave indigenous knowledge resources vulnerable to commercially driven data extraction, in which human experiences become a resource to be mined for economic value (Couldry and Mejias 2019).

Indigenous Data Sovereignty (IDS) is a framework that extends the political notion of sovereignty into the realm of information: it asserts that data relating to Indigenous communities should be governed by those communities, according to their values and for their collective benefit (Carroll, Rodriguez-Lonebear & Martinez, 2019). This is set out in the CARE Principles for Indigenous Data Governance. Rather than focusing on making data open and shareable, CARE principles focus on people and purpose, emphasising that data use should align with the

rights and needs of indigenous communities. This approach focuses on who benefits from data, who has authority over it, who is responsible for its use, and whether its use is ethical (Global Indigenous Data Alliance, 2019; Kwok, 2026). An example of the successful application of the CARE principles is Te Hiku Media, a Māori non-profit broadcaster that collects and archives oral recordings of indigenous knowledge and licenses use under the Kaitiakitanga License, based on the Māori concept of community guardianship (kaitiakitang) (Spano & Zhang, 2025; Kwok, 2026; Te Hiku, 2026).

Addressing epistemic centralisation and its consequences, then, cannot be limited to technical solutions for assembling Large Language Models that better represent world languages and diverse data sources; as with all technologies, social, economic, and political factors play a key role. These dimensions come together in the concept of coloniality; the persistence of forms of power that are grounded in the history of colonialism and persist long after formal colonial administrations have ended (Quijano; 2000; Couldry & Mejias, 2019, 2021; Muldoon & Wu, 2023; Cambraia & Pyrrho, 2025).

Epistemic injustice and credentialing academic achievement

When Generative AI applications first became widely available in late 2022 and early 2023, many educational organisations' first instinct was to ban their use. This absolutist approach was soon abandoned, although Manichaeian solutions are still debated in student assessment, with positions ranging from returning to handwritten submissions to fully incorporating AI into new assessment modalities (Desai, 2025; Grove, 2026). Other risks of using Artificial Intelligence applications for educational purposes are well documented. Studies have described “cognitive offloading”, which can occur when reliance on AI applications reduces the efficacy of learning processes, as well as the emotional dependency that can follow from personalised and automated student support (Fan et al., 2025; Namvarpour, Brofsky, et al., 2025).

Meanwhile, the sophistication and scope of AI models have continued to evolve in the exponential-like pattern characteristic of digital innovations. Artificial Intelligence is becoming ambient, ever-present in day-by-day uses of the digital facilities offered by Microsoft, Apple and Google (Microsoft Ignite, 2025; Twarogal & Grębski, 2025; Acallar, 2026; Apple, 2026). In turn, this is enabling fully-automated sequences for core university processes, from student enrollment management (CollegeVine, 2025), running a research laboratory (University of Toronto, 2026), and for the full process of writing and publishing an academic paper, from formulating the research problem to authoring and reviewing the manuscript reporting a project's outcomes (Gottweis, Weng et al., 2026; Nature, 2016).

Less attention has been given to the implications of epistemic centralisation for the quality of the knowledge transmission that underpins the value of a credential, both for those studying within “hermeneutical hotspots” and for everyone else, everywhere, as the “default speaker” flattens variegated forms of knowledge into a more uniform and synthetic landscape. A good place to

start a discussion of what this means for credentialing is quality assurance; the benchmarks and standards which underpin trust in teaching and learning. The succession of guidances issued by Australia's Tertiary Education Quality and Standards Agency provides an apt starting point (TEQSA, 2023, 2025, 2026).

TEQSA's approach is underpinned by "learning assurance" as a quality principle; the requirement that institutions provide evidence of both formative work and appropriate assessment. In turn, this requires that students can understand AI's epistemic role critically and that assessment tasks are "designed to foster responsible and ethical use of AI in ways that are authentic to both the task and the discipline" (TEQSA, 2023). Rather than focusing on AI detection mechanisms, TEQSA encourages institutions to redesign assessment methods to capture authentic evidence of student capabilities for using new technologies appropriately (TEQSA, 2025).

TEQSA's most recent guidance focuses on the importance of "adaptive capabilities"; the distinctively human capacity to apply lived experience, embodied understanding and ethical awareness. "Institutions should consider investing in developing a systematic capacity to collect and analyse data on learning processes. This requires the coordinated development of learning analytics infrastructure and research methodologies, as well as staff technical capabilities to work with data that illuminates how students are engaging with tasks over time". This suggests quality assurance standards that include evidence of "the traces students generate as they plan, monitor, seek feedback, revise work, interact with resources and make decisions, including how they engage with gen AI" (TEQSA, 2026).

Implementing any form of quality standard in teaching and learning provision leads in turn to conventions for setting learning outcomes and signifying learner attainment in the form of a formal record that serves as a credential that is meaningful and can be trusted. This requires what Bernstein called a "pedagogic device"; a mechanism that shows how knowledge and different ways of knowing are translated for the purposes of teaching and learning (Bernstein, 2000, Maton, 2026). Building on TEQSA's guidance, this recontextualisation requires records of the traces that students create as they plan and revise their work, engage with Artificial Intelligence, use resources, and demonstrate, through appropriate forms of assessment, that they have developed the adaptive capabilities that are the distinctive feature of human intelligence (Floridi, 2025; TEQSA, 2026).

The GDN has led in advancing the concept of a "network of networks", strengthening the key mechanism for protecting the diversity of the global knowledge commons. Here, the objective is to facilitate the portability of digital credentials across and between national higher education systems: "Participants in this ecosystem work together on the safe administration, processing, consultation, transfer, exchange, acceptance and recognition of digital student data in addition to appropriately verified paper-based records, using the best standards and technology available, in order to guarantee trust in this digital student data ecosystem" (GDN, 2026). The GDN has also championed the opportunities inherent in the Comprehensive Learning Record as

a mechanism for a secure, immutable and mobile record of learner attainment (IMS Global, 2025).

This positions the GDN to address the consequences of epistemic centralisation through secure forms of digital credentialing that recognise epistemic validity across national borders and global regions; a key element for a future “federated learning” model that joins secure digital credentialing with a digital knowledge commons, diamond open-access publishing and open AI applications.

Federated learning, originally conceptualised by Google, is an architecture in which a Large Language Model is trained on datasets stored on institutionally based servers, without the data ever leaving its original location (Stanford HAI, 2026). In higher education, federated learning has been used successfully in health sciences research, where data protection obligations prevent patient records from being copied to a generic pre-training database (Topaloglu, Morrell et al., 2021). However, given the urgency of developing an alternative to Frontier AI models, federated learning has significant potential to bring universities together as a global network of networks.

To match the scale and scope of commercial Frontier AI models, a federated network of universities will require metadata that indexes an equally large ocean of source materials that are quality-assured and ethically assembled. This system has been prototyped in initiatives such as LA Referencia, a regional network of institutional repositories across Hispanic American countries, including Brazil, Argentina, Mexico, and Colombia (Bueno-de-la-Fuente & Agustín-Lacruz, 2025). The Confederation of Open Access Repositories (COAR) now has more than 140 members and partners across the world: “The real value of repositories lies in their potential to become an interconnected network that can provide unified access to an aggregated set of scholarly and related outputs that machines and researchers can work with in new ways. This vision of networked repositories promised to make research outputs more discoverable, enable connections between related content, and contribute to the development of other value-added services built on repository resources” (Shearer, 2024; COAR, 2026).

A federated digital data commons comprising institutional repositories could licence data sources for commercial use, in line with the CARE Principles for Indigenous Data Governance, ensuring data quality and protecting authors' attribution rights. Such a system would provide diamond open access, free of processing charges and commercial publishing profit margins, as envisaged in the 2021 UNESCO recommendation on Open Science; a network of local diamond open access “capacity centres,” regional “capacity hubs”, and global co-ordination: a “globally distributed, aligned, high-quality, and sustainable scholarly communication infrastructure that is managed and owned by the scholarly community, and where the quality of outputs is derived from the trust in the communities running the journals” (Mounier and Rooryck, 2023).

A federated digital data commons would also provide high-quality “real-world” data for training fully open and transparent alternative AI models such as AI2’s OLMo framework: “The OLMo project, ‘by scientists, for scientists,’ aims to empower academic and research communities to

study language models. Its open framework seeks to drive new scientific inquiry into AI's fundamental questions, including data-model relationships, bias, and risks. This open science approach counters AI's commercialization, proving it can rival proprietary systems while fostering a deeper, more accountable understanding of these powerful technologies" (Mehmood 2025).

Taken together, this set of initiatives has the potential to turn the principles of epistemic and hermeneutical justice into a practical reality across global Higher Education, assuring the continuing value of secure digital credentialing, the GDN's specific area of focus and expertise.

References

Acallar, L. J. (2026). What is Ambient AI? Ultimate Guide for Clinicians. *Heidi*. 25 June 2026. <https://www.heidihealth.com/en-za/blog/ambient-ai>

AI4Bharat. (2026). *Building AI for India*. Retrieved 21 August 2026 from <https://ai4bharat.iitm.ac.in/>

Alvarado, R. (2023). AI as an Epistemic Technology. *Science and Engineering Ethics*, 29(32). <https://doi.org/10.1007/s11948-023-00451-3>

Apple. (2026). Apple introduces Siri AI, a profoundly more capable and personal assistant. *Newsroom*. 8 June 2026. <https://www.apple.com/newsroom/2026/06/apple-introduces-siri-ai-a-profoundly-more-capable-and-personal-assistant/>

Bernstein, B. (2000). *Pedagogy, Symbolic Control and Identity: Theory, Research, Critique*. Taylor and Francis. ISBN:978-0-8476-9576-8

Beyvers, S., Hochmuth, J., Brehm, L., Hansen, M., Goesmann, A., & Förster, F. (2025). Towards FAIR and federated Data Ecosystems for interdisciplinary Research. *arXiv*. <https://arxiv.org/abs/2504.20298>

Bueno-de-la-Fuente, G., & Agustín-Lacruz, M. d. C. (2025). Repositorios institucionales, redes y organización del conocimiento: un estado de la cuestión en Hispanoamérica. In C. Naumis Peña & A. Suárez Sánchez (Eds.), *Organización del conocimiento : bibliotecología, terminología e información* (pp. 163–192). Universidad Nacional Autónoma de México, Instituto de Investigaciones Bibliotecológicas y de la Información.

Cambraia, L., & Pyrrho, M. (2025). Generative artificial intelligence and the risk of techno-digital colonialism. *Frontiers in Political Science*, 7. <https://doi.org/10.3389/fpos.2025.1628139>

Carretero, J., Izquierdo-Moreno, G., Vasile-Cabezas, M., & Garcia-Blas, J. (2018). Federated Identity Architecture of the European eID System. *IEEE Access*, 6, 75302 – 75326. <https://doi.org/10.1109/access.2018.2882870>

Carroll, S. R., Rodriguez-Lonebear, D., & Martinez, A. (2019). Indigenous Data Governance: Strategies from United States Native Nations. *Data Science Journal*, 18(31), 1–15.
<https://doi.org/10.5334/dsj-2019-031>

Chavez, P., Chilukuri, V., & Scanlon, R. (2026). *Sovereign AI Index. Tracking the Global Push for AI Self-Reliance*. CNAS: Center for a New American Security.
<https://interactives.cnas.org/reports/sovereign-ai-index/>

COAR. (2026). *Building a sustainable, global knowledge commons*. Confederation of Open Access Repositories. Retrieved 5 October 2026 from <https://coar-repositories.org/>

CollegeVine. (2025). *AI Recruiter*. Retrieved 10 December 2025 from
<https://www.collegevine.com/agents/recruiter>

Couldry, N., & Mejias, U. A. (2019). *The Costs of Connection. How Data is Colonizing Human Life and Appropriating it for Capitalism*. Stanford University Press. ISBN: 9781503603660

Couldry, N., & Mejias, U. A. (2021). The decolonial turn in data and technology research: what is at stake and where is it heading? *Information, Communication & Society*, 26(4), 1–17.
<https://doi.org/10.1080/1369118X.2021.1986102>

Desai, H. (2025). What's worth measuring? The future of assessment in the AI age. *Futures of Education*. 19 May 2025. <https://www.unesco.org/en/articles/whats-worth-measuring-future-assessment-ai-age>

De Sousa Santos, B. (2014). *Epistemologies of the South: Justice against Epistemicide*. Routledge. ISBN: 978-1612055459

Fan, Y., Tang, L., Le, H., Shen, K., Tan, S., Zhao, Y., Shen, Y., Li, X., & Gašević, D. (2025). Beware of metacognitive laziness: Effects of generative artificial intelligence on learning motivation, processes, and performance. *British Journal of Educational Technology*, 56(2).
<https://doi.org/10.1111/bjet.13544>

Floridi, L. (2025). AI as Agency without Intelligence: On Artificial Intelligence as a New Form of Artificial Agency and the Multiple Realisability of Agency Thesis. *Philosophy and Technology*, 38(30). <https://doi.org/10.1007/s13347-025-00858-9>

Fricke, M. (2007). *Epistemic Injustice. Power and the Ethics of Knowing*. Oxford University Press. <https://doi.org/10.1093/acprof:oso/9780198237907.001.0001>

Gaba, A. (2026). *The Default Speaker: Why Neutral Conversational AI is a Myth Worth Retiring*. CUI '26: Proceedings of the 8th ACM Conference on Conversational User Interfaces. 45,1-4.
<https://doi.org/10.1145/3816046.3816304>

Gándara, D., & Anahideh, H. (2025). Using AI to predict student success in higher education. Brookings, 17 April 2025. <https://www.brookings.edu/articles/using-ai-to-predict-student-success-in-higher-education/>

Gándara, D., Anahideh, H., Ison, M. P., & Picchiarini, L. (2024). Inside the Black Box: Detecting and Mitigating Algorithmic Bias Across Racialized Groups in College Student-Success Prediction. *AERA Open*, 10. <https://doi.org/10.1177/23328584241258741>

GDN. (2026). *Over a Decade of Contributing to the Digital Ecosystem to Support Education Access and Inclusivity*. Groningen Declaration Network. Retrieved 30 August 2026 from <https://groningendeclaration.org/who-we-are/>

Global Indigenous Data Alliance. (2019). *CARE Principles for Indigenous Data Governance*. Retrieved 31 March 2026 from https://static1.squarespace.com/static/5d3799de845604000199cd24/t/6397b363b502ff481fce6baf/1670886246948/CARE%2BPrinciples_One%2BPagers%2BFINAL_Oct_17_2019.pdf

Gottweis, J., Weng, W.-H., Daryin, A., Tao Tu, Petar Sirkovic, Artiom Myaskovsky, Grzegorz Glowaty, Felix Weissenberger, Alessio Orlandi, Dan Popovici, Anil Palepu, Keran Rong, Ryutaro Tanno, Khaled Saab, Fan Zhang, Jacob Blum, Andrew Carroll, Kavita Kulkarni, Nenad Tomašev, ... Natarajan, V. (2026). Accelerating scientific discovery with Co-Scientist. *Nature*, 655, 487–496. <https://doi.org/10.1038/s41586-026-10644-y>

Grove, J. (2026). Faculty Moving Away From Outright Bans on AI, Study Finds. *Inside Higher Ed*. 13 February 2026. <https://www.insidehighered.com/news/faculty/learning-assessment/2026/02/13/faculty-moving-away-outright-bans-ai-study-finds>

Haase, J., & Pokutta, S. (2026). The Hidden Cost of Tokenization: Why (most) Non-English Speakers Pay More for Less. *Zenodo*. <https://doi.org/10.5281/zenodo.18416604>

Hope, G. (2026). Meet Latam-GPT, the New Open Source AI Model for Latin America. *AI Business*, 12 February 2026. <https://aibusiness.com/generative-ai/the-new-open-source-ai-model-for-latin-america#close-modal>

IMS Global. (2025). *Comprehensive Learner Record Standard Version 2.0 Final Release*. <https://www.imsglobal.org/spec/clr/v2p0>

Jack, P. (2024). Academic backlash as publisher lets Microsoft train AI on papers. *Times Higher Education*. 31 July 2024. <https://www.timeshighereducation.com/news/academic-backlash-publisher-lets-microsoft-train-ai-papers>

Javed, Tahir, Janki Nawale, Eldho George, Sakshi Joshi, Kaushal Bhogale, Deovrat Mehendale, Ishvinder Sethi, Aparna Ananthanarayanan, Hafsah Faquih, Pratiti Palit, Sneha

Ravishankar, Saranya Sukumaran, Tripura Panchagnula, Sunjay Murali, Kunal Gandhi, Ambujavalli R, Manickam M, C Vijayanthi, Krishnan Karunganni,...Khapra, M. (2024). IndicVoices: Towards building an Inclusive Multilingual Speech Dataset for Indian Languages. *Findings of the Association for Computational Linguistics: ACL 2024*, 10740–10782. <https://doi.org/10.18653/v1/2024.findings-acl.639>

Khan, M. S. U. R., & Khapra, M. M. (2024). Indic LLM Suite. *AI4Bharat*, 14 March 2024. <https://ai4bharat.iitm.ac.in/blog/indicllm-suite>

Khan, M. S. U. R., Mehta, P., Sankar, A., Kumaravelan, U., Doddapaneni, S., Suriyaprasaad B, Balan, V., Jain, S., Kunchukuttan, A., Kumar, P., Dabre, R., & Khapra, M. M. (2024). IndicLLMSuite: A Blueprint for Creating Pre-training and Fine-Tuning Datasets for Indian Languages. *arXiv*. <https://arxiv.org/html/2403.06350v2>

Kwok, J. C. Y. (2026). Data colonialism and indigenous languages in AI: a critical review of existing initiatives and their struggles with data sovereignty. *AI & Society*, 41, 6719–6229. <https://doi.org/10.1007/s00146-026-03091-w>

LatamGPT. (2026). *LatamGPT - A language model made in Latin America, for Latin America*. LatamGPT. Retrieved 26 August 2026 from <https://www.latamgpt.org/en>

Leslie, D., & Perini, A. M. (2024). Future Shock: Generative AI and the International AI Policy and Governance Crisis. *Harvard Data Science Review, Special Issue 5: Grappling with the Generative AI Revolution*. <https://doi.org/10.1162/99608f92.88b4cc98>

Lundin, J. M., Zhang, A., Karim, N., Louzan, H., Wei, V., Adelani, D. I., & Carroll, C. (2026). *The Token Tax: Systematic Bias in Multilingual Tokenization* Proceedings of the 7th Workshop on African Natural Language Processing, <https://aclanthology.org/2026.africanlp-main.10.pdf>

Maisiri, J., & Musonza, S. (2025). The cultural cost of AI in Africa's education systems. *UNESCO Idea*, 20 July 2025. <https://www.unesco.org/en/articles/cultural-cost-ai-african-education-systems>

Maton, K. (2026). *Sociology of Possibility. An Invitation to Legitimation Code Theory*. Routledge. ISBN:978-0-367-76532-3

McKinsey (2026). *Sovereign AI: Building ecosystems for strategic resilience and impact*. <https://www.mckinsey.com/industries/technology-media-and-telecommunications/our-insights/sovereign-ai-building-ecosystems-for-strategic-resilience-and-impact#/>

Mehmood, M. (2025). Analysis of the OLMo Training Framework: A Commitment to Truly Open Science, 25 June 2026. <https://genaiops.ai/analysis-of-the-olmo-training-framework-a-commitment-to-truly-open-science>

Microsoft Ignite. (2025). *Microsoft 365 Copilot overview*. Microsoft. Retrieved 24 September 2025 from <https://learn.microsoft.com/en-us/copilot/microsoft-365/microsoft-365-copilot-overview>

Mollema, W. J. T. (2025). A taxonomy of epistemic injustice in the context of AI and the case for generative hermeneutical erasure. *AI and Ethics*, 5, 5535–5555. <https://link.springer.com/article/10.1007/s43681-025-00801-w>

Mounier, P., & Rooryck, J. (2023). Towards a federated global community of Diamond Open Access. *The Diamond Papers*, 23 December 2023. <https://thd.hypotheses.org/296>

Muldoon, J., & Wu, B. A. (2023). Artificial Intelligence in the Colonial Matrix of Power. *Philosophy and Technology*, 36(80). <https://doi.org/https://link.springer.com/article/10.1007/s13347-023-00687-8>

Namvarpour, M. M., Brofsky, B., Medina, J. Y., Akter, M., & Razi, A. (2025). Understanding Teen Overreliance on AI Companion Chatbots Through Self-Reported Reddit Narratives. *arXiv*. <https://arxiv.org/pdf/2507.15783v3>

Nature. (2026). AI scientists are changing research — institutions, funders and publishers must respond [25 March 2026]. *Nature*, 651, 853–854. <https://doi.org/10.1038/d41586-026-00934-w>

Ng, R., Nguyen, T. N., Yuli, H., Tai Ngee Chia, Leong Wai Yi, Wei Qi Leong, Xianbin Yong, Jian Gang Ngui, Yosephine Susanto, Nicholas Cheng, Hamsawardhini Rengarajan, Peerat Limkonchotiwat, Adithya Venkatadri Hulagadri, Kok Wai Teng, Yeo Yeow Tong, Bryan Siow, Wei Yi Teo, Tan Choon Meng, Brandon Ong,...Teo, L. (2025). *SEA-LION: Southeast Asian Languages in One Network*. Proceedings of the 14th International Joint Conference on Natural Language Processing and the 4th Conference of the Asia-Pacific Chapter of the Association for Computational Linguistics, Mumbai. <https://aclanthology.org/2025.ijcnlp-long.30/>

Oderinwale, H., & Kazlauskas, A. (2026). The Economics of AI Training Data: A Research Agenda. *arXiv*. <https://arxiv.org/pdf/2510.24990>

Ojo, J., Ogundepo, O., Akintunde Oladipo, Kelechi Ogueji, Jimmy Lin, Pontus Stenetorp, & David Ifeoluwa Adelani. (2025). AfroBench: How Good are Large Language Models on African Languages? *Findings of the Association for Computational Linguistics*, 2025, 19048–19095. <https://aclanthology.org/2025.findings-acl.976.pdf>

Petrov, A., La Malfa, E., Torr, P. H. S., & Bibi, A. (2023). Language model tokenizers introduce unfairness between languages. NIPS '23: Proceedings of the 37th International Conference on Neural Information Processing Systems, New Orleans. <https://dl.acm.org/doi/proceedings/10.5555/3666122>

Quijano, A. (2000). Coloniality of Power, Eurocentrism and Latin America. *Nepantla: Views from the South*, 1(3). <https://doi.org/10.1177/0268580900015002005>

SEA-LION. (2026). *Catalysing AI Innovation for Southeast Asia*. SEA-LION.AI. Retrieved 24 June 2026 from <https://sea-lion.ai/>

Shearer, K. (2024). Bridging the Digital Divide: How Repository Networks Are Shaping the Future of Open Science. 11 December 2024. COAR: Confederation of Open Access Repositories. <https://coar-repositories.org/bridging-the-digital-divide-how-repository-networks-are-shaping-the-future-of-open-science/>

Shumailov, I., Shumaylov, Z., Zhao, Y., Papernot, N., Anderson, R., & Gal, Y. (2024). AI models collapse when trained on recursively generated data. *Nature*, 631, 755–759. <https://doi.org/10.1038/s41586-024-07566-y>

Song, H., Tsukioka, K., Ubaldi, B., Andersen, R., Harvard, J., Ward-Jackson, G., Langengen, T., West, D., & Teo, M. (2026). *Sovereignty in the Age of AI: Strategic Choices, Structural Dependencies and the Long Game Ahead*. Tony Blair Institute for Global Change. <https://institute.global/insights/tech-and-digitalisation/sovereignty-in-the-age-of-ai-strategic-choices-structural-dependencies>

Spano, I., & Zhang, Y. (2025). Indigenous data sovereignty in intangible cultural heritage governance: A complementary approach to public–private partnerships. *International Journal of Cultural Property*, 32(2), 167–193. <https://doi.org/10.1017/S0940739125100064>

Stanford HAI. (2026). *What is Federated Learning?* Stanford Human-Centered Artificial Intelligence. Retrieved 4 October 2026 from <https://hai.stanford.edu/ai-definitions/what-is-federated-learning>

Tao, Y., Viberg, O., Baker, R. S., & Kizilcec, R. F. (2024). Cultural bias and cultural alignment of large language models. *PNAS Nexus*, 2024(3). <https://doi.org/10.1093/pnasnexus/pgae346>

Te Hiku. (2026). *Kuaka Mārangaranga*. Retrieved 1 January 2026 from <https://tehiku.nz/>

TEQSA. (2023). *Assessment reform for the age of artificial intelligence*. Tertiary Education Quality and Standards Agency Australia. <https://www.teqsa.gov.au/guides-resources/resources/corporate-publications/assessment-reform-age-artificial-intelligence>

TEQSA. (2025). *Enacting assessment reform in a time of artificial intelligence*. Tertiary Education Quality and Standards Agency Australia. <https://www.teqsa.gov.au/sites/default/files/2025-09/enacting-assessment-reform-in-a-time-of-artificial-intelligence.pdf>

TEQSA. (2026). *Assuring quality learning in a gen AI-integrated future: The role of adaptive capabilities*. Tertiary Education Quality and Standards Agency Australia.

<https://www.teqsa.gov.au/sites/default/files/2026-06/assuring-quality-learning-in-a-gen-AI-integrated-future.pdf>

Topaloglu, M. Y., Morrell, E. M., Rajendran, S., & Topaloglu, U. (2021). In the Pursuit of Privacy: The Promises and Predicaments of Federated Learning in Healthcare. *Frontiers in Artificial Intelligence*, 4. <https://doi.org/10.3389/frai.2021.746497>

Twarogal, P., & Grębski, R. (2025). Google Gemini vs Microsoft Copilot in Action: From Docs to Code. Neontri, 4 July 2025. <https://neontri.com/blog/gemini-vs-copilot/>

United Nations. (2007). *United Nations Declaration on the Rights of Indigenous Peoples*. https://www.un.org/development/desa/indigenouspeoples/wp-content/uploads/sites/19/2018/11/UNDRIP_E_web.pdf

University of Toronto. (2026). *Accelerating materials discovery for a sustainable future*. Retrieved 6 January 2026 from <https://acceleration.utoronto.ca/vision>

Villalobos, P., Ho, A., Sevilla, J., Besiroglu, T., Heim, L., & Hobbhahn, M. (2024). Will we run out of data? Limits of LLM scaling based on human-generated data. *arXiv*. <https://doi.org/10.48550/arXiv.2211.04325>

World Values Survey. (2026). World Values Survey. Retrieved 26 August 2026 from <https://www.worldvaluessurvey.org/>

Yeyati, E. L. (2025). Latam-GPT and the search for AI sovereignty. *Brookings*, 25 November 2025. <https://www.brookings.edu/articles/latam-gpt-and-the-search-for-ai-sovereignty/>