

CRITICAL COMMENTARY

Governing AI Literacy through a Data-Governance Lens

A critical commentary on the
“Compliance Without Capacity” Policy Brief

Author: Lauren Pachaly

MSt in AI Ethics and Society

University of Cambridge

July 6th, 2026

1. Introduction

The policy brief is addressed to Brazil's Ministry of Education and its National Education Council. It examines a recent decision in which the federal government made adopting the computing complement to the national curriculum — which carries the school system's AI-literacy competences — a condition for a results-based top-up within its main school-funding system. Its concern is that the condition applies everywhere while the means to meet it exist in only some places. The recommendations share one aim, to strengthen the instrument so it builds the capacity that implementing the complement requires rather than assume it already exists.

Where the brief argues over the design of the chosen lever, this commentary steps back to the governance questions behind how a state governs AI literacy at all, several of which stand out. Who decides what AI literacy is, and how it is put into effect? By what rules, and through which instrument, should it be governed? How should the literacy itself be measured? Who is accountable for building the capacity it requires? And what kind of citizen does it form? This commentary sets out to examine two of these, on accountability and citizenship, concepts Brazil's own framework for AI in education places at its centre, where it frames educating for AI as “intellectual autonomy, critical discernment, and the full exercise of democratic citizenship”, and warns that, unless disparities are addressed, technology tends to deepen inequality rather than reduce it (Brazil. Ministry of Education, 2026, p. 9). How far a compliance instrument can deliver on that purpose is not obvious, and it is that distance the commentary explores.

These questions are examined through the literature on data governance and data justice. This choice is deliberate, and less a matter of preference than of what the subject is made of. For Long and Magerko (2020), AI literacy is a set of competencies that enables individuals to critically evaluate, collaborate with, and use AI systems. Built as they are by learning from data, what a model can do, and to whom, is in large part set by the data it is trained on (Jordan and Mitchell, 2015) and the relations that data encodes (Viljoen, 2021, pp. 577–578). Read together, these suggest that a literacy about AI is, above all, a literacy about data and the systems built on it. The same may hold one level up, in governance, since a growing body of work argues that governing AI is largely a matter of governing that data and the relations it produces (Taylor et

al., 2026). If that is right, the broader question this commentary asks, how a state governs the AI literacy its citizens receive, begins to look like a data-governance question at root, and the mature field of data governance and data justice becomes a natural place to look for the concepts an emerging field still lacks, since AI literacy and governance are only beginning to be established as foundations for ethical AI (Karimov and Saarela, 2025). This connection is deliberately compressed, offered as the commentary's analytical lens rather than the only account of how these fields relate.

The sections that follow situate the case within the broader debates it touches, then examine those two questions more closely and consider the limits and challenges of governing AI literacy.

2. Placing the case within broader AI-governance debates

This section places the Brazilian case within the wider debates in AI governance. It offers not an exhaustive map but a first reading of what the lens of data governance can illuminate about governing AI literacy.

The first of these is data justice, which asks not only whether a system works but whom it makes visible and whom it leaves out (Taylor, 2017), and which treats the protection of public infrastructure and public goods as a test of sound governance (Taylor et al., 2026). On that test it is worth asking whether a literacy that presumes connectivity, devices, and prepared teachers distributes a public good evenly or merely tracks one already there, and whether the same instrument might advance literacy in aggregate while widening the gap beneath the average.

The second is the geography of AI. Lehdonvirta et al. (2024) divide the world by access to the compute AI is built on, into a Compute North that hosts the infrastructure to build it, a Compute South with far less, and a Compute Desert with none. The divide, they argue, is not only material but political, since states without compute are less able to set the terms and standards under which AI operates. Brazil sits on the southern side of that line and adopts a framework written in good part on the northern one, from UNESCO to the OECD (Brazil.

Ministry of Education, 2026, p. 117), which the decolonial literature reads as a familiar pattern, the export of one way of knowing as though it were universal (Ricaurte, 2019). One need not go that far to make the modest point that a literacy imported into an under-resourced system, and taught to students who will mostly meet AI built elsewhere, tends to equip them to use such systems while leaving the relations of dependence around them unexamined.

The third is the question of power itself. Read through power, the case raises several questions — who defines the literacy, who benefits, and who bears the cost of its absence. The one this commentary takes as its example, following (Kalluri, 2020), is not whether an instrument is fair but whether it redistributes power or leaves it in place, since a system can be fair and still leave power exactly where it was. A curriculum that teaches students to use AI without the capacity to question it does little to shift power, while one that builds the capacity to contest can begin to change it. This is where Brazil's own framing earns some credit. By placing AI literacy under digital culture rather than narrow technical skills, the curriculum leaves room, in principle, for the broader and more civic conception the critical model calls for (CNE/CEB, 2022b).

Whether that broader conception is realised, though, turns on a fourth and more practical debate, about assessment. The funding condition pays out only once a network can show it has adopted the curriculum, so a civic competence has to be turned into a checkable signal, and in Brazil that signal is the information the network files rather than a window onto the classroom (CIF, 2025). As Ananny and Crawford (2018) argue, seeing a record is not the same as knowing what a system does, and the broader the conception of literacy, the harder it is to verify, which creates a quiet pressure to narrow it back to the checkable, and so to certify the user while believing one is forming the citizen.

Taken together, these debates begin to reframe AI literacy from a curricular question into a governance one. Against that backdrop, the sections that follow take up two governance questions more closely.

3. What kind of citizen the arrangement forms?

As Carmi et al. (2020) argue, literacy is never only a set of skills but a social practice shaped by context and power. They turn that insight into a distinction between a neutral, universal view of literacy — a set of individual skills that can be rolled out everywhere in the same form — and a critical model they call *data citizenship*, oriented toward enabling people to question and act on the systems around them. Their framework sets out three dimensions. *Data thinking* is the critical understanding of how data and the systems around it work. *Data doing* is the everyday practice of using and managing data. And *data participation* is the proactive, collective engagement through which people act with others on the systems that affect them. The three give a test to run any curriculum through, whether it builds understanding and everyday practice and also the capacity to act with others, or stops at the first two.

Data citizenship does not reject competences but builds on them. The mainstream definitions supply the baseline, treating AI literacy as a set of competencies everyone needs rather than a specialist skill (Long and Magerko, 2020). A critical data literacy starts from that baseline and reaches further, toward the capacity to evaluate data systems and their social effects (Pangrazio and Selwyn, 2019). The question is which of the two, the functional competence or the critical capacity, a compliance instrument tends to reward. Read against the critical model, Brazil's curriculum is not empty. Its computing component is organised around computational thinking, the digital world, and digital culture, with data literacy and the ethics of use running through it (CNE/CEB, 2022a), so it speaks, at least on paper, to data thinking and data doing, and gestures at participation through its language of digital citizenship.

Where it falls short is in what Carmi et al. (2020) call networked and contextual literacies. For them a literacy cannot be delivered as a property of the individual learner in isolation, because people engage with data through the people and places around them, so they situate it in citizens' networks of literacy — their families, communities, and neighbourhoods — and reject the idea that one common programme can serve everyone, since people with different backgrounds need different literacy programmes. A curriculum applied in the same form across the country, and certified the same way everywhere, is the universal model they warn against,

and it is the contextual and networked dimension, not the list of competences, that it leaves underdeveloped.

There is a further reason an individual-skills model misfits its subject. Viljoen, (2021, pp. 573–574, 577–578) argues that data is better understood as a social relation than an individualist claim about a single person, since data drawn from one person routinely supports inferences about many others, so its value and harms are produced at the level of populations rather than individuals. A literacy assembled as a personal checklist tends to speak only to the part of the problem visible to the individual, so the networked and collective dimension is not an optional add-on but the part that makes an AI literacy relevant.

In Brazil, the contextual and the collective become a single problem, because the networks and contexts Carmi and colleagues describe are also the populations Viljoen points to — a landscape so varied and unequal that mapping it would be a study in itself. To give one example, access to devices and connectivity is deeply uneven, so the same lesson presumes resources many schools lack. Only 28% of rural schools have computers and internet for students, against 76% of urban ones, and 29% of students in the North reach the internet at school through an institution's computer against 87% in the South (CETIC.br, 2025, pp. 21, 23). A literacy designed for a generic classroom is likely to form one kind of citizen in the school it imagines and another in the schools many students actually attend. Reduced to universal skills, it tends to form the thinner of the two, the digital resignation Carmi and Nakou (2025) warn a functional literacy can produce.

None of this is inevitable and an example of a critical literacy model exists inside the same system. Piauí, a north-eastern state of Brazil, made AI a subject of its own curriculum, now taught across 540 public schools to around 90,000 students a year with over 680 teachers trained to deliver it (UNESCO, 2025). The programme was designed so that teachers grasp AI critically rather than only operationally, and its most cited classroom result — students building an AI tool to catalogue and distribute seeds to local family farmers — is data citizenship in practice, designed from the local context outward rather than imported into it (Porvir, 2025). Set against a nationally uniform requirement, it shows the difference between forming users

and forming citizens, and that the critical model is achievable within the same system, provided the instrument is built to support it.

4. Who is accountable for building the capacity?

Whether a literacy forms citizens or users depends, in good part, on the capacity to teach it, and that returns the analysis to a second governance question, who answers for building it. In Brazil the requirement, as it stands, assumes a capacity to teach that many networks do not yet have. For AI literacy that capacity is not only infrastructural but human and institutional. This is the harder kind to build, and the kind Brazil can least take for granted, because the requirement is layered onto a system still working to secure foundational learning. In PISA 2022, Brazilian fifteen-year-olds scored well below the OECD average across mathematics, reading, and science, with around half below the baseline level of proficiency in reading and roughly three-quarters in mathematics (OECD, 2023, Figures I.3.1 and I.3.4). The workforce that would carry a new and demanding subject is itself stretched. In the final years of primary school many classes are taught by teachers not qualified in the subject, and in much of the North and Northeast fewer than half are (INEP, 2025, Figure 1). Preparation for the digital part of the job is thin and, in places, thinning, as teachers' participation in ongoing professional development on digital technologies fell from 65% in 2021 to 54% in 2024, with the lack of it cited as a barrier by 86% of municipal-school teachers (CETIC.br, 2025, p. 91). This is also where the accountability the instrument sets up comes under strain, since it holds each network answerable for the literacy, and liable to consequences (Bovens, 2007), while the capacity most decisive for it, a prepared teaching workforce, is one many networks simply do not have. To require results before that capacity is in place is to hold a network to account for what it lacks the means to do.

In that setting the capacity to teach AI cannot be assumed, since in many networks it would have to be built from a low base, and this is why teachers are a decisive variable. As the literature on policy enactment argues, teachers do not simply implement a written curriculum but interpret and translate it into what students actually receive (Braun et al., 2011). So where they are unprepared even a well-designed curriculum tends to collapse into its most checkable

form, the operating skills that can be demonstrated rather than the critical capacities that cannot. Reading through the citizenship lens, the implication is sharper still, because a teacher who can only operate AI is unlikely to form students who can question it, which makes the critical, networked literacy depend first on the formation of the people expected to teach it. (Karimov and Saarela, 2025) place professional development among the main channels through which governments build AI literacy, and note that uneven resources and teacher training can hinder adoption, as they observe for the EU.

If that is right, teacher formation is less one investment among several than the step zero on which the rest depends. The thing a capacity-first instrument would fund, and reach the least-resourced networks with, before it asks them to comply. That opens questions this commentary can only raise but not settle. Where should responsibility for forming a teaching workforce for AI sit? With the federal level that sets the standard, the states and municipalities that employ the teachers, or some shared arrangement? And what division of that labour would reach the networks furthest behind? It may be unfair to expect one funding rule to carry all of this. What the Brazilian case points to instead is that a combination of instruments and policies — shared but adaptable curricula, partnerships with universities and support bodies, and the AI-literacy centres some systems are building — is more likely to have lasting effect when the governance roles behind them are made explicit.

5. The limits and challenges of governing AI Literacy

Governing AI literacy in Brazil, a continental country still marked by deep structural gaps in its population's foundational literacy, is a complex governance challenge in itself. This section brings some of that complexity into view by holding the brief's own recommendations up to scrutiny.

The case for the instrument's original design is not weak. In a decentralised system, a binding funding condition may be one of the few levers capable of steering the networks toward a shared goal, and without it AI literacy could remain a marginal concern precisely in the places least able to prioritise it. The brief's critique — that this lever rewards capacity rather than

builds it and so falls hardest on those who lack the means — is persuasive, but its own fixes, despite protecting something real, carry tensions of their own. Turning the condition into a capacity ladder does not remove administrative discretion so much as relocate it, since someone must still set each network's baseline and timeline, and those judgements reopen the influence of capacity, politics, and measurement. Allowing networks to meet the requirement through device-free, unplugged AI literacy protects access where infrastructure is missing, yet, if teaching about AI on paper were treated as equivalent to teaching with it, it would entrench the very divide it addresses and ease the pressure to close the infrastructure gap. Funding capacity before requiring it, on its own, does not settle the matter, as what most decides whether the capacity is actually built — from a prepared teaching workforce to the institutional arrangements for delivering it — reaches beyond financial investment.

Beneath the recommendations also lies a tension between the common national standard Brazil has adopted and the contextual literacy set out earlier. One of the strongest cases for the standard is that a floor guarantees every network owes its students at least a baseline, which in a deeply unequal system is itself a form of equity. The point, though, is not to reject the floor but to deny that it is the ceiling. A baseline of functional competence is a precondition for the critical kind, not a substitute for it (Pangrazio and Selwyn, 2019). An instrument that can only see the floor certifies compliance without guaranteeing the critical capacity it cannot observe (Ananny and Crawford, 2018). Even a fully funded capacity, then, is not the same as citizenship, since a network could meet every requirement and still form users rather than citizens (Carmi et al., 2020). No funding design resolves that trade-off on its own, and it is perhaps at the point where the two questions meet, who builds the capacity and what citizen the literacy forms, that the future of this instrument might most usefully be shaped.

6. Conclusion

Brazil's decision to tie AI literacy to a results-based funding condition is sound in its ambition, and the question it raises is not whether AI literacy should be part of the school curriculum, which is already settled, but how the state governs it and to what end. Seen through

established theories of data governance, several questions stand out, reframing AI literacy from a curricular matter into a governance one. This commentary examines two of the governance questions Brazil's own framework places at its centre — what kind of citizen the arrangement forms and who is accountable for building the capacity it requires — and steps back to weigh the brief's own recommendations.

What this commentary reveals is that AI literacy cannot be governed as curriculum or compliance alone — a task harder still where the ground beneath it is uneven, in systems still working to secure the foundational learning. By bringing the more developed questions and theories of data governance to bear, this commentary aims to give this emerging field a vocabulary and a set of governance questions to prioritise. The questions examined here, and the others left open, are less a conclusion than a starting point. An attempt to equip researchers and policymakers with considerations grounded in a real and unfolding case, in the hope that the angles left unexplored serve as a springboard for further research and public debate on how societies choose to govern the AI literacies their citizens receive.

Bibliography

Ananny, M., and Crawford, K. (2018). Seeing without knowing: Limitations of the transparency ideal and its application to algorithmic accountability. *New Media and Society*, 20(3), 973–989. <https://doi.org/10.1177/1461444816676645>

Bovens, M. (2007). Analysing and Assessing Accountability: A Conceptual Framework¹. *European Law Journal*, 13(4), 447–468. <https://doi.org/10.1111/j.1468-0386.2007.00378.x>

Braun, A., Ball, S. J., Maguire, M., and Hoskins, K. (2011). Taking context seriously: Towards explaining policy enactments in the secondary school. *Discourse: Studies in the Cultural Politics of Education*, 32(4), 585–596. <https://doi.org/10.1080/01596306.2011.601555>

Brazil. Ministry of Education. (2026). *Framework for the Development and Responsible Use of Artificial Intelligence in Education [Referencial para o Desenvolvimento e Uso Responsável da Inteligência Artificial na Educação]*. Ministry of Education (Brazil).

<https://www.gov.br/mec/pt-br/media/segape/referencial-oficial-en.pdf>

Carmi, E., and Nakou, P. (2025). Building data citizenship and learning to resist in the datafied society. *Learning, Media and Technology*, 1–14.

<https://doi.org/10.1080/17439884.2025.2505553>

Carmi, E., Yates, S. J., Lockley, E., and Pawluczuk, A. (2020). Data citizenship: Rethinking data literacy in the age of disinformation, misinformation, and malinformation. *Internet Policy Review*, 9(2). <https://doi.org/10.14763/2020.2.1481>

CETIC.br. (2025). *Survey on the use of information and communication technologies in Brazilian schools: ICT in Education 2024*. Brazilian Internet Steering Committee (CGI.br).

https://www.cetic.br/media/docs/publicacoes/2/en-us/20251217163959/tic_educacao_2024_livro_completo_ingles.pdf

CIF. (2025). *Resolução CIF nº 15, de 12 de junho de 2025 [Resolution CIF No. 15 of 12 June 2025: Methodology for assessing the management-improvement conditionalities for distributing the VAAR per-student top-up in 2026]* [Resolution]. Ministério da Educação, Secretaria de Educação Básica.

<https://www.gov.br/fnde/pt-br/acesso-a-informacao/acoes-e-programas/financiamento/fundeb/legislacao/2025/resolucao-cif-15-2025-fundeb-condicionalidades-vaar-2026.pdf/view>

CNE/CEB. (2022a). *Computação: Complemento à Base Nacional Comum Curricular (BNCC)*

[Computing: Complement to the National Common Curricular Base (BNCC)]. Ministério

da Educação, Conselho Nacional de Educação.

<https://www.gov.br/mec/pt-br/escolas-conectadas/arquivos/BNCCComputaoCompletoDiagramado.pdf>

CNE/CEB. (2022b). *Parecer CNE/CEB nº 2/2022: Normas sobre Computação na Educação*

Básica—Complemento à Base Nacional Comum Curricular (BNCC) [Opinion No. 2/2022:

Standards on computing in basic education—Complement to the National Common

Curricular Base (BNCC)] [Opinion]. Ministério da Educação.

https://www.gov.br/mec/pt-br/cne/pdf/pareceres-do-cne/ceb/2022/pceb002_22.pdf

INEP. (2025). *Censo Escolar da Educação Básica 2024: Notas Estatísticas [Basic Education School*

Census 2024: Statistical Notes] [Statistical Notes]. Instituto Nacional de Estudos e

Pesquisas Educacionais Anísio Teixeira (INEP).

https://download.inep.gov.br/publicacoes/institucionais/estatisticas_e_indicadores/notas_estatisticas_censo_da_educacao_basica_2024.pdf

Jordan, M. I., and Mitchell, T. M. (2015). Machine learning: Trends, perspectives, and prospects.

Science, 349(6245), 255–260. <https://doi.org/10.1126/science.aaa8415>

Kalluri, P. (2020). Don't ask if artificial intelligence is good or fair, ask how it shifts power. *Nature*,

583(7815), 169–169. <https://doi.org/10.1038/d41586-020-02003-2>

Karimov, A., and Saarela, M. (2025). AI Literacy and Governance as Foundations for Ethical AI: A

Cross-National Review of Government Strategies. *2025 IEEE Symposium on Trustworthy,*

- Explainable and Responsible Computational Intelligence (CITREx)*, 1–7.
<https://doi.org/10.1109/CITREx64975.2025.10974932>
- Lehdonvirta, V., Wú, B., and Hawkins, Z. (2024). *Compute North vs. Compute South: The Uneven Possibilities of Compute-based AI Governance Around the Globe*.
- Long, D., and Magerko, B. (2020). What is AI Literacy? Competencies and Design Considerations. *Proceedings of the 2020 CHI Conference on Human Factors in Computing Systems*, 1–16.
<https://doi.org/10.1145/3313831.3376727>
- OECD. (2023). *PISA 2022 Results (Volume I): The State of Learning and Equity in Education*. OECD Publishing. <https://doi.org/10.1787/53f23881-en>
- Pangrazio, L., and Selwyn, N. (2019). ‘Personal data literacies’: A critical literacies approach to enhancing understandings of personal digital data. *New Media and Society*, 21(2), 419–437. <https://doi.org/10.1177/1461444818799523>
- Porvir. (2025). Programa de ensino de IA nas escolas do Piauí ganha prêmio da Unesco [AI teaching programme in Piauí schools wins the UNESCO prize]. *PORVIR*.
<https://porvir.org/piaui-ia-escolas-premio-unesco/>
- Ricaurte, P. (2019). Data Epistemologies, The Coloniality of Power, and Resistance. *Television and New Media*, 20(4), 350–365. <https://doi.org/10.1177/1527476419831640>
- Taylor, L. (2017). What is data justice? The case for connecting digital rights and freedoms globally. *Big Data and Society*, 4(2), 205395171773633.
<https://doi.org/10.1177/2053951717736335>

Taylor, L., De Souza, S. P., Martin, A., and López Solano, J. (2026). Governing artificial intelligence means governing data: (Re)setting the agenda for data justice. *Dialogues on Digital Society*, 2(1), 5–22. <https://doi.org/10.1177/29768640241306800>

UNESCO. (2025). *UNESCO recognizes four initiatives promoting the responsible use of AI in education*.
<https://www.unesco.org/en/articles/unesco-recognizes-four-initiatives-promoting-responsible-use-ai-education>

Viljoen, S. (2021). A Relational Theory of Data Governance. *The Yale Law Journal*, 131(2), 573–654.