

TECHNICAL NOTE

Open Artificial Intelligence in the Educational Context

Initiative:

ProFuturo | Fundação Telefônica vivo | Fundação "la Caixa"

Execution: IA.EDU

Achievement: unesco Chair

Achievement
UNESCO NETWORK OF UNPLUGGED ARTIFICIAL INTELLIGENCE IN EDUCATION - UFAL

Chairs
Ig Ibert Bittencourt
Seiji Isotani

Rector of the Federal University of Alagoas - UFAL
Josealdo Tonholo

Operations Coordinator
Marcela Lorenzoni R. Rabin

Production

Founders
Maria Alice Carraturi
Ig Ibert Bittencourt
Seiji Isotani

Executive Director
Thomaz Edson Veloso da Silva

Authors
Sabrina Gonçalves Raimundo
Tarcízio Silva
Camila Wasserman
Ig Ibert Bittencourt
Seij Isotani

Technical-Editorial Supervision
Eloi Capucho Ferreira
Camila Wasserman

Graphic Design and Layout
Eloi Capucho Ferreira
Fernando André de S. Vasques

Translation and Proofreading
Karina Villarinho

Technical Review
Camila Wasserman

Initiative

General Director
Lia Glaz

Research, Partnerships, and Institutional Mobilization Management
Catherine Rojas Merchan

Research, Partnerships, and Institutional Mobilization Team
Marina Alves Queiroz Maluta
Kelly Cristine de Oliveira Marchetti

ProFuturo

ProFuturo's General Director
María Dolores Martínez Bernabéu

Area Manager – Innovation and Educational Model
Milada Tonarelli Gonçalves

Technology, Data, and Artificial Intelligence Manager
Manuel José Ruiz Garcia

Head of Evaluation and Learning
Virginia Guadalupe Soto Sira

Innovation and Educational Model Specialist
Cleide da Silva

Digital and Online Education Specialist
Ana María Moreno Garduño

ÍNDICE

Executive Summary 4

1. Open AI: Conceptual Framework and Collaborative Ecosystem 6
1.1. What open Artificial Intelligence means 7
1.2. From total openness to proprietary AI 8
1.3. Open AI Models and its components 9
1.3.1. Model Weights 10
1.3.2. Code 10
1.3.3. Training Data 10
1.3.4. Documentation 11
1.4 What is shared and what is built around open models 11
1.5 What open AI does not mean 12
1.5.1. It does not mean free of charge 12
1.5.2 1.5.2 It does not mean licenses with paid features (freemium model) 12
1.5.3. It does not mean without risks 13
1.5.4. It does not mean without governance 13

2. Global Trends for Open AI in Education 15
2.1 Patterns and trends observed globally 15
2.2. Shared public infrastructure and technological sovereignty 16
2.3 Open AI as a public good 18

3. Possibilities and Safeguards for the use of Open AI in Education 20
3.1. Open AI in Education 20
3.2. Open AI Movements in Education 21
3.3. Possibilities for Open AI in Education 23
3.3.1. Uses for Teaching 24
3.3.2. Uses for Learning 22
3.3.3. Uses for Management 25
3.4. Safeguards for the use of Open AI in Education 25
3.4.1. Challenges and risks of Open AI in Education 25
3.4.2 Literacy and critical thinking regarding Open AI 27

4. Strategic Recommendations for Education 31
4.1 Guiding questions for decision-making 31
4.2 Recommendations by stakeholder of the educational ecosystem 32
4.2.2. Development companies and EdTechs 33
4.2.3. Governments and States 33
4.3 Decision-Making Criteria Matrix 35

5. Synthesis 36

References 37

Initiative: ProFuturo | Fundação Telefônica vivo | Fundação "la Caixa"

Execution: IA.EDU

Achievement: unesco Chair

EXECUTIVE SUMMARY

This Technical Note is a collaborative effort, based on data, evidence and international consensus, to support the design, adoption and governance of open Artificial Intelligence (AI) solutions, aligned with the public interest in the educational context.

Written for specialists, institutional leaders, and public policymakers within the education ecosystem, as well as for the general public interested in the topic, it coordinates a survey of academic research and practical experiences to offer analytical frameworks, decision-making criteria, and strategic guidelines on open AI.

We start from the premise that a large portion of educational innovations today relies on proprietary models, which are often associated with high costs, limited transparency, and low adaptability to diverse contexts, challenges that are relevant to Brazil. In our context, open AI emerges as a strong strategic alternative by offering greater transparency, interoperability, and possibilities for local adaptation, aligning itself with the agendas of digital public goods, technological sovereignty, and AI ethics.

The central premise is the realization that openness is not a binary condition, but rather a spectrum of components, licenses, and practices that must be minimally understood by different stakeholders. Its adoption imposes technical, pedagogical, ethical, institutional, and infrastructural challenges that vary according to context, capacities, and regulatory frameworks.

The document, developed in partnership between the Institute of Artificial Intelligence in Education (Instituto de Inteligência Artificial na Educação - IA.Edu), ProFuturo, and Telefônica Vivo Foundation (Fundação Telefônica Vivo - FTV), within the scope of the UNESCO Chair on Unplugged Artificial Intelligence, offers a rigorous, applied, and strategic decision-oriented overview.

The Note answers the following key questions: how open AI is defined in education; which opportunities it offers for teaching, learning, and education; what its risks and limitations are; under what conditions it presents an advantage over proprietary solutions; and what the role of each stakeholder is in promoting the collaborative open AI ecosystem.

SECTION 1 – CONCEPTUAL FRAMEWORK: it establishes the shared framework, positions open AI as a spectrum with a component-by-component analysis — model weights, code, data, and documentation — and dismantles recurring misconceptions (that openness implies free-of-charge access, freemium models, or the absence of risks and governance).

SECTION 2 – GLOBAL TRENDS: comparative analysis of strategies across different countries and regions, highlighting openness as a strategically value searched globally.

SECTION 3 – POSSIBILITIES AND CAUTIONS: examines practical uses for teaching, learning, and management, contrasting opportunities with risks — curation, child and adolescent data protection, cultural appropriateness — and reinforcing critical literacy as a condition for responsible use.

SECTION 4 – STRATEGIC RECOMMENDATIONS: converts learnings into actionable guidance, featuring guiding questions, recommendations by stakeholder and decision criteria among models with different levels of openness.

SECTION 5 – SYNTHESIS: revisits the main themes, outlines future lines of investigation and cooperation, and projects the challenges of an open AI ecosystem grounded in transparency, equity, and protagonism of those who teach and learn.

1.0

Open AI: Conceptual Framework and Collaborative Ecosystem

1. Open AI: Conceptual Framework and Collaborative Ecosystem

Technical advancements and the widespread availability of generative Artificial Intelligence (AI) systems have created new demands and opportunities for this technology in educational institutions. What once seemed distant from daily life has become accessible in just a few clicks, thanks to global investment and competition surrounding technology, becoming a staple in the daily routines of both public and private systems. This Technical Note discusses a topic whose relevance demands the attention of managers, pedagogical teams, and policymakers:

what is open AI, and what are its advantages and disadvantages

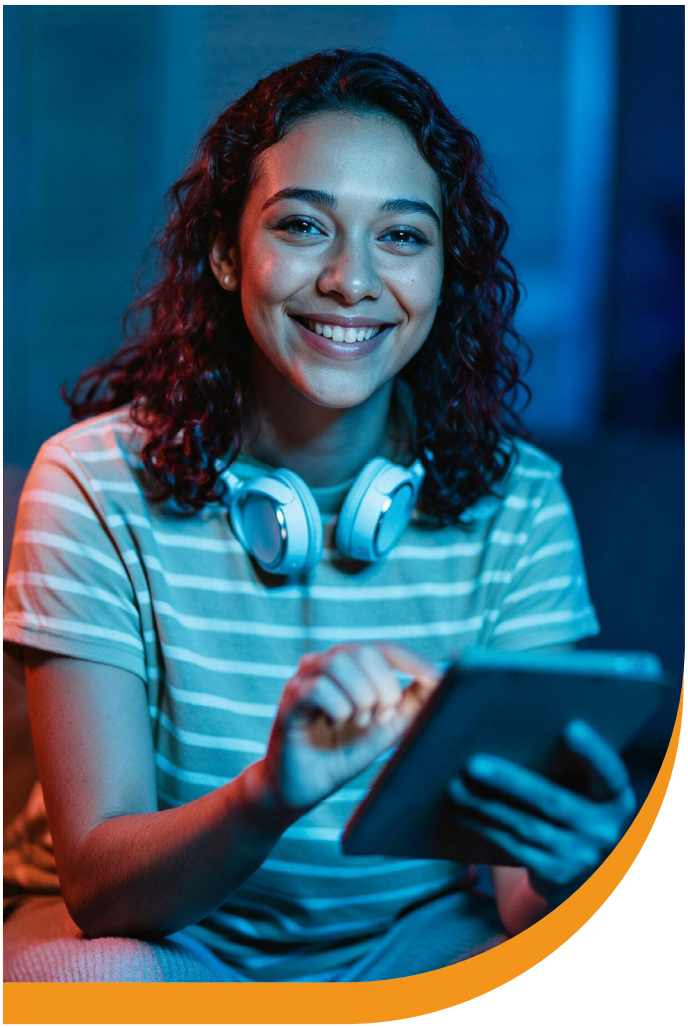
At first glance, the question seems technical, but its implications span budgetary, legal, and pedagogical layers, as well as disputes over digital sovereignty that must be considered before making decisions in a field filled with so many concepts, terms, and complexities.

This section of the document presents a conceptual framework that establishes the foundation for the following sections by introducing:

a) notions of AI based on gradient-of-openness approaches;

b) technical and operational variables underpinning the collaborative open AI ecosystem;

and c) considerations regarding misconceptions and fallacies on the use and characteristics of open AI.



Source: Photo from the Magnific website

1.1. What open Artificial Intelligence means

In the field of computing, the notion of openness far predates contemporary forms of AI. Recalling this origin helps us understand that the term is highly broad and contested, and at times ambiguously understood. In the early decades of commercial computing, software source code circulated relatively freely among developers in universities and companies, creating a culture of sharing that was largely based on the concepts of collective scientific progress, continuously improved and peer-reviewed.

From the 1980s onward, commercial computing became even more relevant in the global scenario with heavy investments from financial capital and state policies that considered its relevance for the technological and scientific leadership of countries with greater production capacity. In response to this growing closure, developer communities consolidated the free software movement, which defines four fundamental freedoms for users:

to run the program for any purpose, to study how it works, to modify it, and to redistribute it.

Subsequently, the notion of open source was established as an alternative in commercial practice, preserving some of these freedoms while being integrated into business models. The binary approach, whether a system was simply "open" or "closed", was still the most common notion.

With the popularization of contemporary AI (especially given the advancements connected to large foundation models trained on massive, multimodal sources of training data), inherited terms have become more complex. Current AI models are not just programs in the traditional sense of the past century. They also simultaneously incorporate parameters learned from long and costly processes using datasets from diverse sources, computational infrastructure with unprecedented capacity, and also multiple layers of application and reappropriation, which include connection and integration layers across many different systems before reaching the end user.

The binary notion – yes or no – to address open systems in artificial intelligence could oversimplify rich possibilities (and variables) at stake when we discuss the topic today. In this regard, to fully leverage the potential of the "open" aspect of AI systems, we must look at which factors differentiate the degrees of openness and which of these are relevant to each specific application or set of applications.

1.2. From total openness to proprietary AI

The international debate has been moving away from a rigid opposition between open and closed, in favor of a more gradual notion of openness. Each component of an artificial intelligence system can individually be in different positions regarding public availability, types of licensing, transparency about training, or usage restrictions.

This notion of differential openness encompasses the premise that two systems labeled as open can actually offer highly distinct possibilities, advantages or disadvantages for those who use them. Thus, there is a spectrum of practices and decisions, and the material meaning of openness varies according to which specific components are present in the system's development lifecycle. Acknowledging this spectrum is a starting point for the institutional decisions on the matter.

At one end of the spectrum, we have absolutely proprietary models, in which the developing organization keeps parameters, code, data, documentation, and licenses under strict control. Access, when it exists, occurs either through end-use in final implementations or by contracting services through third-party interfaces that do not allow inspection, modification, or local hosting. For many, this design protects business secrets and inventions but prevents external analysis and creates technological dependency for those who adopt the system. Therefore, any change is centralized on the developer.

In an intermediate position there are open-weights models, in which the learned parameters are publicly available for download and reuse, enabling local execution, fine-tuning, and contextual adaptation. However, in this design, training data and training code often remain unavailable, and licenses sometimes restrict commercial or competitive uses.

Finally, in a more advanced position, there are models with expanded openness, which also include, in addition to the model weights, the inference and training code, technical reports, and model cards. In full openness, all components, including datasets, evaluation metrics, and reproducible pipelines, are published under licenses that ensure use, study, modification, and redistribution. These typical configurations, however, rarely appear in their full state: most systems adopted today combine characteristics, which makes component-by-component reading even more essential.

AI Systems are not simply open or closed!

The binary open/closed opposition is replaced by a gradual spectrum: each component varies in availability. They range from proprietary models to open weights and full openness, but real systems usually combine the openness of different elements and their characteristics.



Source: Photo from the Magnific website

1.3. Open AI Models and its components

To make this component-by-component reading operational, let us review the elements that can be examined in any system presented as open. These elements can be organized as a progression of depth — a scale ranging from the most frequently released components to the most sensitive ones, culminating in full openness when all of them are published under permissive licenses. The following diagram summarizes this progression.

MAIN COMPONENTS FOR OPEN AI

TRAINING DATA
Training data, frequently considered the most important factor, is usually the least transparent layer. It comprises web content and digital databases used to teach the model, making it a central theme of current debates in the field.

MODEL WEIGHTS
Model weights are numerical parameters that encode the learning from training. When open, they allow local use by anyone, but they do not reveal how they were built. They benefit users, developers, and large companies through distributed learning

DOCUMENTATION
Detailed records of decisions throughout the development and implementation of the system. This includes model and data cards, impact assessments, incident records, and usage policies.

CODE
The code itself: divided into inference (architecture and input processing) and training (how the model learns). Sharing them enables scientific reproducibility: replicating, expanding, or methodologically critiquing what was done.

It is important to emphasize that deciding among any systems described as open involves thinking about the categories of openness and how much each one can influence artificial intelligence-based projects in the long term. The openness of each component can expand what can be done in practice and define the options for adaptation, auditing, and sovereignty. It is a decision-making exercise, especially when it involves goals of scale, as in the case of public education. For this reason, the debate is so intense and rich internationally, engaging technology, software, and artificial intelligence governance entities.

1.3.1. Model Weights

The most explicit component on the scale is the model weights — that is, the learned parameters. These are the numerical parameters with billions or even trillions of values, which encoded everything the model learned during training from massive databases of content, data, and information.

When open, it means that anyone can use the weights and run them locally for specific purposes. However, separately, they do not allow an understanding of how the model was built, its eventual limitations, potential capabilities, or biases.

Today, contemporary models with open weights are even released by major technology companies and research consortia and are easily publicly available. The strategy of publishing model weights openly benefits the entire field by becoming an accessible resource for many users and individual developers — and for the major development companies, which begin to learn about practical applications and implementations in a distributed manner, like a global collective intelligence.

1.3.2. Code

The second level is the code, which in turn can be composed of a) inference code, which is responsible for the architecture and input processing, and b) training code, which dictates how the model learns from the data.

Sharing this level makes it possible to advance toward notions of scientific reproducibility, which means that other people can follow the exact same steps on the same data to replicate the procedure in order to be able to expand or criticize through scientific methodology what was accomplished.

1.3.3. Training Data

Training data, often seen as the less important factor in the age of information, paradoxically, is usually the least transparent layer, even in systems that present themselves as open.

This category is involved in controversies that are still being defined in the field, such as the fact that a significant portion of large language models were trained on generalized extractions of web content and digital databases (including datasets viewed as pirated), which has opened the door for class-action lawsuits related to intellectual property and copyright.



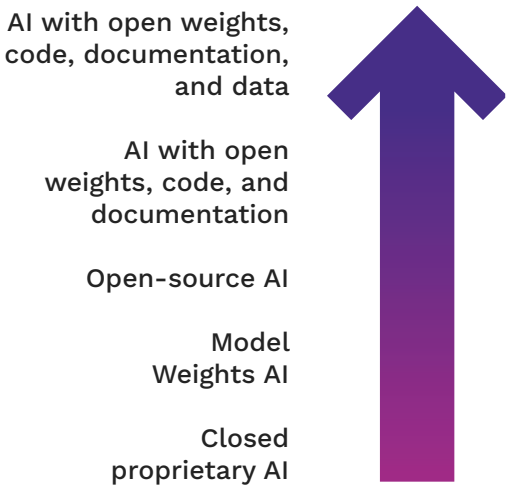
Source: AI image from the Magnific website

1.3.4. Documentation

The fourth layer is documentation, which can include detailed procedures for recording decisions made throughout the development and implementation lifecycle of an artificial intelligence system. This includes model cards, training data cards, impact assessment reports, incident records, and usage policies.

Thus, the four layers, including substantive documentation, could guide open artificial intelligence in a way that converges with those classic freedoms of the open-source movement: to study, modify, and share.

AI models do not present a fixed binary structure, as they can integrate various openness variables. Consequently, each tier of openness reflects distinct levels of autonomy and transparency, as illustrated in the following diagram. It is fundamental to point out, however, that an increase in autonomy proportionally demands greater infrastructure and technical competence to sustain it.



1.4. What is shared and what is built around open models

Around open options for algorithmic systems, which can then be understood as being composed of layers and modules in constant interaction and negotiation, a heterogeneous and complex ecosystem is articulated. This ecosystem includes hosting and cloud computing platforms, development communities, physical infrastructure, data repositories, standardization tools, application layers, and evaluation methodologies as well.

The practical consequences of openness occur in relation to this ecosystem, rather than in isolation. An essential advantage of open AI is that **it creates a two-way street:** it empowers the distributed communities around it and feeds on the collective knowledge that sustains the projects. Multilingual databases with the participation of hundreds of volunteers across dozens of countries show how decentralized large-scale coordination can work. Maintainers of shared resources play social roles that go beyond software development, also including the motivation of collaboration networks, contribution and documentation revision. Sometimes, these initiatives evolve to include governance structures, such as boards, foundations, cooperatives, and review methodologies—which indicate long-term potential.

As for the field of education, within the application layers, initiatives are still consolidating but are gaining traction. Adapting systems from large-scale foundation models is a customization challenge, but a necessary effort to adjust and adapt to the specific demands of education at each level, context, and territory.

1.5. What open AI does not mean

As a concept under dispute and construction, it is possible to identify some common misconceptions about what the openness of algorithmic systems means. Debating such confusion is important because, when adopted uncritically, this concept can lead to poorly calibrated decisions and bad results in the medium and long term. We will name some of these misunderstandings below, beginning with the frequent notion that openness means that it is free of charge.

1.5.1. It does not mean free of charge

The most common confusion is equating openness with being free of charge. Yes, it is true that many model components can be acquired at no cost for use or licensing. However, to run the models, especially when the goal is to make them available to many people, such as for educational systems, one cannot forget the necessary computational infrastructure. Even when running a simplified model on a single computer, it is also important to remember that hardware specifications are indispensable for proper use. Such factors must be kept in mind so that the application of open AI does not become a vector of exclusion: specialized processors, local or cloud storage, electricity, meaningful connectivity, and specialized human resources. All of these are factors that involve costs for a good transformation into relevant implementations.

In countries of continental dimensions and profound inequalities, such as Brazil, this geography of the distribution of opportunities and potential capabilities can reflect the history of regional and socioeconomic disparities. Thus, the adoption of open AI as an educational public policy also requires intelligent budget planning that is adapted to different realities.

1.5.2 It does not mean licenses with paid features (freemium model)

Beyond the apparent free access, we find the most confusing category of supposedly open systems. It is frequent for artificial intelligence services developed by private companies, but with open components, to offer basic versions of their products for free, yet with limitations, while advanced versions (with more parameters, more effective performance, extra capabilities, or technical support) are only available on commercial terms. This practice is typically part of the commercial strategies adopted for these models.

This design, also called freemium, is a legitimate business model, but it can lead to confusion when users or managers mistake it for examples of openness or, worse, as complete solutions for educational purposes. Such a freemium approach, as seductive as it may initially be for popularizing services and experimenting with skills, can introduce vulnerabilities regarding use, limitations in capacity and number of users, or even present greater instability in terms of service maintenance. After all, if there are no well-structured contracts or terms of use, large-scale application in educational contexts can become dependent on fluctuations in the business tactics of major artificial intelligence system developers.

1.5.3. It does not mean without risks

It is common to think that a system described as open is necessarily auditable or transparent and, consequently, more secure. In practice, it is more complex than that. Indeed, the openness of the components makes them more suitable for analysis and investigation, allowing independent developers and researchers to observe and test their features, check for biases, and report vulnerabilities.

This is a real and existing benefit.On the other hand, however, this same openness also opens breaches for harmful and/or malicious uses. Models with open weights can be modified to bypass embedded ethical safeguards or be influenced to generate disinformation or strategic bias, such as political targeting.

In the educational context, the use of open systems can mean a loss of cultural and pedagogical sensitivities and specificities. Systems that have been trained and developed in certain regions can behave inappropriately in other contexts, and this inadequacy is difficult to measure without detailed curation and audits — which demand deeper knowledge of the subject. In the Brazilian context, this means that open AI should only be adopted when accompanied by critical literacy, responsible use policies, mechanisms for monitoring discriminatory effects, as well as continuous training and the strengthening of resources and working conditions for teach

1.5.4. It does not mean without governance

Another myth involves believing that, since the code is available, the weights can be downloaded, and the documentation is public, institutional governance arrangements would be unnecessary.

On the contrary, open models require more governance, not less, because the multiplicity of actors involved increases the complexity of the arrangements. It is necessary to define who maintains the repositories, how vulnerabilities are identified and reported, who records the versioning of releases and how, and how conflicts among contributors are managed. For the educational field, this topic is especially sensitive: the adoption of open AI in public school systems cannot dispense with formal, multi-stakeholder governance bodies that bring together secretariats, schools, universities, developers, EdTechs, organized civil society, and, fundamentally, representatives of the communities directly affected by algorithmic decisions.

Finally, it is worth mentioning a notion that has appeared in communities dedicated to open source: openwashing – the purely rhetorical use of the 'open' label by systems and companies that actually retain control over essential components.

Since, as we will see in the following section, several countries have been adopting AI strategies and regulations that offer advantages to systems with free and open-source models, many developers try to frame their products as open, even though they make no effort to actually allow expanded networks to use, study, modify, and share their components.

2.0

Global Trends for Open AI in Education

2. Global Trends for Open AI in Education

In the previous section, we addressed open AI based on four main component groups: model weights, code, training data, and documentation; and we discussed the most common concepts and misconceptions surrounding the term. Now, we will examine how these ideas manifest on the international landscape through three interconnected trends: the recurring patterns in national strategies and multilateral formulations; the advancement of shared public infrastructure as a technological sovereignty agenda; and the consolidation of open AI as a public good in a field still marked by concentration.



Source: AI image from the Magnific website

2.1. Patterns and trends observed globally

All over the world, strategies toward AI on education are organized around international frameworks produced longitudinally, supported by organizations such as UNESCO and the OECD. There are common aspects articulated around recurring themes, including: educational management and implementation; teaching practices; learning assessment; skills for life and work in the era of AI; and lifelong learning.

Multilateral systems defend the expansion of open infrastructures and strong investment in the production and maintenance of the so called Digital Public Goods (DPG), a concept which encompasses free software, open AI models, open data, open standards, and more. Proposals in the field of digital governance in the United Nations highlight these possibilities as healthy and cooperative alternatives to the concentration of infrastructure in a few large tech companies. Openness, initially seen as a technical matter within developer communities, is now considered an important geopolitical variable.

Reviews from OECD indicated that, in 2025, open-weight models accounted for about 55% of the global commercial supply. The geography of the system providers, however, remains concentrated: the United States accounts for more than half of foundational model providers, followed by China, the United Kingdom, France, and Germany. As we have explored in the previous chapter, the openness of model weights by itself does not dissolve the centrality of these poles in other layers of the value chain, ranging from chip manufacturing to service provision to energy, training data, evaluation infrastructure, and technical training of human resources.

Four trends have been shaping AI strategies across the world:

- The first one is the commitment to large-scale technical and scientific training — in Brazil, with plans to create 5,000 undergraduate places in AI-related courses and the training of 230,000 professionals over five years (Brazil, 2025); in India, with support for more than 500 doctoral students, 8 thousand undergraduates (India, 2025), and 5,000 postgraduates; in Japan, with the integration of AI into primary and secondary education (Japan, 2026).

- The second one is the investment in local linguistic models — such as Portuguese and indigenous languages from Brazil, multilingual platforms in India —, in response to the assessment that global models were trained predominantly in English.

- The third one is the incorporation of cultural frameworks, such as perspectives centered in Kenya in Africa, and attention to the diversity in the Latin American context, also mentioning different groups such as women, rural populations, and indigenous peoples in Colombia.

- The fourth one is the acceleration of legal governance, with intense debates on national and regional regulatory frameworks addressing safeguards that include the education sector, which is sometimes identified as a high-risk

To summarize it, the world has been looking at:

- technical and scientific training in massive proportions;
- allocation of resources for regionalized linguistic systems;
- integration of cultural principles, plurality, and diversity;
- the advancement of regulatory governance frameworks.

2.2. Shared public infrastructure and technological sovereignty

The theme of technological sovereignty appears as a common category in national AI strategies given the concentration of computing capacity, datasets, and models in a few corporate stakeholders, particularly those linked to the world's two major economic leaders: the USA and China.

In a globalized world, this direction does not tend toward isolation or nationalism, but rather toward building domestic capabilities throughout the development lifecycle of artificial intelligence systems – including hardware, models, data, and applications – with discussions on whether the states should have national control over critical infrastructure. Among the most prevalent variables, there are issues related to the dependence on foreign providers and data centers, as well as the limits of jurisdiction over sensitive data and information seen as strategic.

In several national plans, the formulation of sovereignty appears connected to three common fronts: the expansion of computing infrastructure, energy sustainability, and the establishment of national data infrastructures. Analyses of data sovereignty converge on challenging the dependence on foreign data centers as a risk factor for digital independence, coupling this with the mobilization of local energy matrices as a comparative advantage.

A small fraction of dominant development companies controls not only most state-of-the-art AI models, but also the foundational infrastructure that shapes the field — including GPUs for training and inference, as well as cloud systems that host these chips. This is the background from which the date for open AI emerges as an alternative for sovereignty.

In Kenya, for example, the 2025-2030 strategic cycle integrates geothermal energy and the national fiber optic network as bases for a regional processing model. In another illustrative case, in Spain, aiming to position the country at the forefront of the efficient use of computing infrastructure and the development of sustainable AI, initiatives are proposed across three axes: sustainable data processing centers; planning and organization in the installation of data processing centers; and a certification seal as well as an ecosystem around sustainable AI.

In Brazil, the Brazilian Artificial Intelligence Plan 2024-2028 plans for the expansion of national high-performance computing infrastructure, with investments in the National Laboratory for Scientific Computing (Laboratório Nacional de Computação Científica – LNCC) and the National Research Network (Rede Nacional de Pesquisa – RNP), seeking to reduce the dependence on processing in foreign clouds and expand sovereign capacity for the development and execution of AI models within the national territory.

Proposals for regulatory frameworks are beginning to incorporate openness as a legal category, expanding the theme by combining it with community practices and market and competitiveness strategies. In some jurisdictions, general-purpose models released under free and open-source licenses receive a differentiated regulatory regime, with exemptions from certain transparency obligations conditional upon making parameters, architecture, and usage instructions publicly available. In the European AI framework (EU AI Act), this differentiation does not cover implementations identified as carrying systemic risk. Openness is thus actively fostered, both directly and indirectly, by states.

In regions without major reserves of economic power and infrastructure, such as countries of the Global South, fully national and isolated solutions are diagnosed as unfeasible. Regional coordination converges on the recommendation of supranational alliances for the development, acquisition, and public distribution of digital educational resources.

In Latin America, this trend manifests in proposals to share national infrastructure with countries in the region, in the Caribbean, and in Africa — in Brazil's case, with a contribution of around R$ 50 million to expand the regional academic systems (Brazil, 2025). In different areas, the Brazilian experience with federated networks, such as the RNP, demonstrates successful examples of shared models. The Ipê Network, an infrastructure connecting the country's leading research centers and universities, both with each other and with foreign academic networks, is the backbone which integrates more than five hundred institutions and has points of presence in all twenty-seven states of the federation.

Federated models like this allow the high costs of acquisition, operation, and maintenance to be shared, avoiding the duplication of investments and distributing computing capacity beyond major urban centers, while keeping strategic connectivity, computing, and data assets under national governance. By combining resource sharing, territorial capillarity, and institutional cooperation, such designs align with the recommendation for regional coordination and the technological sovereignty agenda.

Finally, the acquisition of software with open components by the public sector resurfaces as a direction for digital sovereignty. Its adoption by governments allows full ownership of customized versions, reduction of licensing costs, and independence from external providers, provided that it is accompanied by technical frameworks for security and sustainability, as well as a sustained long-term investment in infrastructure, education, and human resources.

2.3. Open AI as a public good

The notion of Digital Public Goods has brought the discussion on open AI closer to the debates over shared infrastructure, which have been present on the Open Educational Resources agenda since the 2000s. Multilateral formulations treat open source code, open models, open data, and open standards as layers of an infrastructure oriented toward the public interest, shifting the discussion on AI from a strictly commercial framework to one of essential public value.

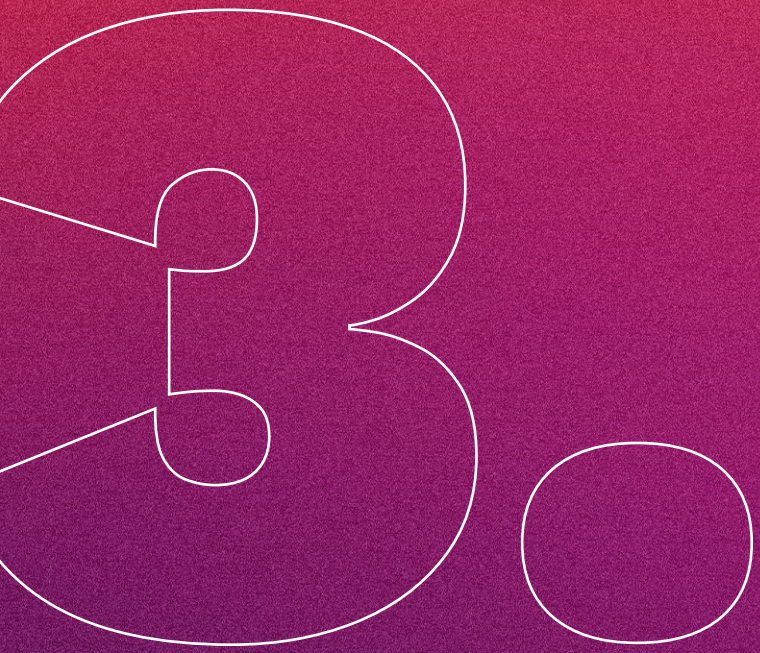
Returning to a theme that we highlighted in the previous section: the transposition of the open-source concept to AI is not straightforward, it is something more complex and systemic. What we generically call 'source code' can refer to inference code, training code, data, or model weighs — components that can be made available independently. Cooperative efforts dedicated to defining open AI converge on the requirement of freedoms to use, study, modify, and share systems, which are different when we think of systems that are necessarily more complex than isolated software and potentially embedded within networks and flows of data, information, and implementations.

The arguments in favor of open AI as a public good leverage technical cooperation, reduction of adoption costs, opposition to the market power of big tech companies, expansion of external research, and community-driven quality reviews of systems. In the educational field, open models enable adaptive tutoring, formative assessment, applications in underrepresented languages, and contextualization and alignment with local territory demands. As we have discussed in the previous chapter, the literature also records adversarial uses: the generation of abusive material, copyright infringement, and large-scale fraud. Once released, a model can hardly be recalled—release decisions are not entirely reversible.

The history of digital common goods offers a relevant conceptual framework. The collaborative production of digital encyclopedias and repositories of open data or components serves as a constant reference for peer production under collective governance outside market pricing. Globally, we have many examples of what has worked and survived for decades, such as the wiki movement, despite growing technological disruptions. Applied to AI in education, this logic suggests decentralized frameworks for collaboration and shared management. It is not enough to advocate for open AI in generic terms: it is necessary to specify which aspects of the technology are accessible, distributable, and reusable by society, and under what governance arrangements —which directly revisits the component-by-component analysis presented in the previous chapter.

Given the already known risks, concerns regarding regulation, data security, and the protection of children and adolescents occupy a central position in the global debate. Therefore, it becomes essential to also enable public participation in decisions concerning the use of AI — particularly at the policy level — in order to ensure sovereignty, strengthen democracy, and mitigate negative impacts. In this sense, the notion of openness connects directly with the promotion of access and interaction within the political, economic, and scientific spheres.

In this context, open AI as a public good designates a collective infrastructure arrangement in which educational systems, multilateral organizations, academic communities, and social movements are coordinated so that technical advancement does not operate only toward corporate concentration. The challenge for the coming years is to convert this repertoire into concrete reconfigurations, with investments in open infrastructure, contextual policies, and international cooperation capable of redistributing the capacities for AI production, and not just consumption — in a sustainable manner given the accumulating indicators of the climate crisis. In the next chapter, we will move forward on how these elements can be organized into concrete recommendations for the Brazilian educational system.



Possibilities and Safeguards for the use of Open AI in Education

3. Possibilidades e Cuidados do uso da IA Aberta na Educação

In the previous sections, we have examined the complexity of the concept of openness applied to AI, noting how it distances itself from conventional definitions related to the understanding of open source, establishing a new paradigm of access and availability.

In 2015, UNESCO reinforced that free access to data, research, and educational resources is an essential pillar of knowledge societies to reduce disparities and promote universal access to information. It is in this direction that initiatives such as free software, open science, public data, open educational resources (recursos educacionais - REA), and, currently, the movement for open AI come together to offer solutions that drive collaboration, social participation, and access.

In this section, we will analyze the materialization of this openness within the educational sphere, presenting examples and practical case studies to explore alternative uses and opportunities for the field. We will also discuss curation criteria to support the choice of open AI models, evaluating their potential capabilities and risks given the tension between security and innovation.



Source: AI image from the Magnific website

3.1. Open AI in Education

The lack of understanding regarding what comprises an open AI system or a standardized and legally consistent definition results in immediate practical consequences. For instance, licensing remains ambiguous and the declared benefits of openness do not necessarily translate into effective transparency for users and institutions. The urgency of this definition intensifies in the educational field, especially when it comes to the use and handling of data from younger students, for whom privacy risks are considerably greater.

Global surveys by the Digital Education Council in 2024 indicate that 86% of higher education students use AI in their studies; among the most frequently used are ChatGPT, followed by Grammarly and Microsoft Copilot, while others such as Gemini and Claude are mentioned in smaller percentages. All of these are proprietary solutions with free access — either partial or total. Critics argue that the adoption of proprietary AI — even under a free license or a freemium model — can deepen inequalities, serve market interests, and impose specific worldviews, without educators or institutions having the ability to audit or challenge the systems they use.

This hegemony of proprietary models and free LLMs in educational environments reflects and perpetuates historical dependencies, in which many territories act as mere recipients of technologies developed in global production centers — as in the case of the Global South. Thus, within primary education, a sharp disparity is observed: private institutions have more than double the access to AI tools compared to public ones. This inequality is even more critical in rural and peripheral areas, where the lack of infrastructure and equipment prevents the effective use of AI systems, proving that the offer of free licenses is insufficient to overcome structural barriers.

Considering this, ensuring that AI applied to education is safe, auditable, and a public good at the service of everyone requires leveraging the potential of openness in its most concrete components (open data and transparent algorithms) as well as its infrastructure. At this level of access, the development of responsible AI systems, the continuous improvement by independent researchers and developers, the traceability of knowledge, and the protection of the integrity of students' rights become viable.

Openness and Security: what changes

The lack of a standardized definition for "open' AI creates legal ambiguity and risks; free proprietary solutions deepen educational inequalities

Besides pedagogical implications, the openness of AI systems can work as a catalyst for higher levels of research and development. Interest in the field is significant in both academia and the private sector, driven by the technology's potential to generate around 13 trillion dollars in global economic activity by 2030, equivalent to an annual increase of 1.2% in global GDP. Promoting openness in research accelerates the circulation of new knowledge and allows a broader audience to develop applications based on state-of-the-art techniques, which can already be observed through the dynamics of accelerated growth in scientific publications on the subject over the last decade.

Therefore, the movement toward open AI models in education is not merely a technical matter of licensing; it is, ultimately, about defining with safety and transparency how and what is learned, and which values are embedded through the technologies that mediate the educational experience of millions of people.

3.2. Open AI Movements in Education

Several initiatives based on varied levels of openness illustrate robust alternatives for using open AI models that can bring concrete benefits to pedagogical planning, institutional management, educational innovation, and, ultimately, learning. AI models and tools based on open components, such as Llama 3, BLOOM, OLMo, Mistral 7B, Falcon, Qwen3, and IndicBERT, are gaining ground in educational environments, driven by communities and collaborative spaces like Hugging Face, a repository of French origin that provides hundreds of thousands of public models for global experimentation.

The application of these models in education frequently occurs through shared or open infrastructures. A striking example in higher education and research is JupyterHub — a multi-user solution for managing computing environments that provides individual server instances for single users or groups. It has become popular for teaching AI and programming, supporting students from institutions such as UC Berkely and the University of São Paulo to move beyond passive consumption. With this tool, students can inspect, modify, and fine-tune open models directly from their browsers.

Another possibility for using open-generative AI in education lies within the field of learning management. For example, Moodle has introduced an agnostic, provider-independent AI API in version 4.5 (2024), allowing public institutions to run automated tutoring and performance analytics on their own servers without exposing student data to third parties. In addition to these models, these management platforms have also been the target of open machine learning algorithms aimed at optimizing and sophisticatedly analyzing data generated through user interactions with the platforms.

However, open platforms and models alone do not solve the issues related to the demands for inclusion and access in a constantly transforming society. Part of the path also involves training those who will use and develop them. In this direction, several impactful initiatives around the world can be mentioned: AI4ALL (Stanford University, USA), which focuses on providing free training for underrepresented university students; Khipu and LatinX in AI, Latin American academic initiatives for strengthening regional research in machine learning; AI4Bharat (IIT Madras, India), which develops open models and training data for the 22 constitutionally recognized Indian languages; and IA.Edu (Federal University of Alagoas and University of São Paulo, Brazil, and Harvard, USA), which promotes the concept of Unplugged Artificial Intelligence in Education — teaching about AI without depending on connectivity, which is already present in thousands of Brazilian municipalities.

This effort is distributed across all regions, although in different ways. In Asia, YUVAi (from the Indian government and the NASSCOM Foundation) engages high school students in AI projects for local development. In Africa, the Deep Learning Indaba, founded by researchers from the University of the Witwatersrand and Stellenbosch University, has held events in 47 countries through its IndabaX program, while Masakhane — a volunteer community linked to researchers from the University of Pretoria and the University of Johannesburg — builds natural language systems for African languages through participatory research. In Europe, AI4InclusiveEducation, developed by the Belgian organization La Scientothèque within the framework of Erasmus+, offers open modules on AI citizenship education for youth from disadvantaged communities.

What these cases collectively indicate is that the adoption of open AI in education is neither an isolated phenomenon nor is it restricted to high-income contexts, but rather a global trend that manifests itself in diverse and complementary ways. The challenge that persists, however, is ensuring that this technical openness translates into an effective openness of opportunities, which requires, beyond available code and models, investments in training, infrastructure, and public policies capable of sustaining what communities and researchers have been building, often with scarce resources.

3.3. Possibilities for Open AI in Education

Preparing lessons, creating assessments, grading essays, and providing individualized feedback are tasks that consume a significant portion of educators' and administrators' time and workload. Artificial Intelligence (AI) systems have a real capacity to support the teaching profession, with the potential to “improve both the productivity and the quality of teaching” by taking over repetitive tasks and freeing teachers up for complex pedagogical interactions that depend strictly on human judgment. However, the success of this integration demands that the use of technology always be guided by a clear pedagogical purpose, rather than by the mere convenience of automation (OECD, 2026).

In this scenario, open AI models emerge as a strategic alternative for educational systems. This is because they enable institutional auditing, safeguards for student data privacy, and the feasibility of fine-tuning algorithms for local realities — ranging from supporting linguistic diversity with large-scale multilingual models to highly specialized platforms. This flexibility is essential in educational contexts, since the capacity to adapt a model to the curriculum, the language, the available infrastructure, and local contexts can be more critical to the educational impact than the overall size or performance of the system.

However, it is worth noting that the most relevant openness for educational equity is not always found in the largest language models or complex infrastructure, but rather in systems tailored to the local context, the curriculum, and the concrete reality of the learner — a dimension that is especially critical in countries with limited infrastructure, where thousands of public schools operate with intermittent or nonexistent connectivity.

Open AI as a Strategic Alternative

Open AI models benefit education by enabling institutional auditing, data privacy, and local adaptation to the curriculum and the infrastructure.

The MathAide tutor, developed by the Center of Excellence in Social Technologies within the framework of the UNESCO Chair of Unplugged Artificial Intelligence in Education, illustrates this need: its architecture is based on the demands of the national curriculum and the logic of simplifying access to democratize the availability of AI. The solution allows teachers to photograph handwritten essays and process them whenever internet access becomes available.

This technical choice is also a pedagogical one: by identifying errors without rewriting or paraphrasing them, it preserves exactly what the student produced, an indispensable condition when the goal is to assess, rather than rewrite on behalf of the learner. This case demonstrates that AI solutions in education need to be flexible enough to function in diverse contexts — and not just where the conditions are ideal.

Seeking to explore the broad possibilities of AI open models, we present below a series of practical use cases, divided into three areas of application:

- USES FOR TEACHING: support for lesson planning, curriculum alignment, and learning assessment.
- USES FOR LEARNING: support for personalized tutoring, cognitive development, and feedback, teacher education.
- USES FOR MANAGEMENT: support for management processes, prediction, and anticipation of problems, and the advancement of teacher education.

3.3.1. Uses for Teaching

Teaching practice encompasses tasks such as planning, developing activities, assessments, and student monitoring. In many cases, these tasks can be automated and supported by AI-based systems.

Regarding planning, research approaches such as LessonPlanLM propose incorporating pedagogical knowledge into the generation process, producing structured lesson plans aligned with curricular objectives in stages that meet requirements for structural integrity, logical coherence, and accuracy (Jia et al., 2025). The result is a lesson plan generation system based on LLMs enriched with curricular knowledge — the full name is Knowledge-enhanced Automatic Lesson Plan Generation — which combines large language models with a specific knowledge base on lesson plans (LPKB, Lesson Plan Knowledge Base).

This principle of contextualized adaptation can also be observed in Twiga, developed by Tanzania AI Community on the Llama model. The solution is trained with national curriculum data and made available via WhatsApp, allowing teachers to create lesson plans, assessments and materials aligned with the Tanzanian curriculum directly from their phones, addressing a concrete shortage of teaching materials in low-infrastructure contexts.

The same potential extends to assessment. An experience conducted at the University of Ljubljana (2024/2025) for the course of Introduction to Bioinformatics demonstrated, in a blind study with over 100 students, that open models used for grading assessments achieved accuracy and feedback quality comparable to those of human evaluators, while preserving student data privacy on the institution's own servers (Rebelo Marcolino et al., 2025).

Open AI has also been incorporated into new categories of teacher support tools. AI-based Teaching Assistants, such as JeepyTA, developed by Ryan Baker and his team at the University of Pennsylvania, integrate directly into university forums to answer questions 24 hours a day, summarize weekly discussions, offer code debugging tips, and provide rapid feedback on essays. Students using this solution rated this type of tool as comparable to human assistants in clarity and accuracy, although weaker in motivation and guidance for development. (OECD, 2026)

The use of open AI can be observed in initiatives across various educational fields and stages. For instance, the use of model MedExamLLM for the alignment of educational resources in medical training contexts and the use of BLOOM (176 billion parameters) for the equitable processing of underrepresented languages reinforce that technical openness also operates as a mechanism for linguistic and curricular inclusion.

When combined, these experiences evidence that the viability of adopting open AI not only expands the repertoire of tools available to educators but does so in a contextually responsive manner: from curricular planning and content production to automated assessment, as well as continuous support via teaching assistants.

3.3.2. Uses for Learning

Although contemporary studies indicate that the interaction between humans and machines may favor learning, this integration demands a strategic approach. One of the central risks mapped in the literature refers to the impact on metacognitive development: specifically, the phenomenon known as “cognitive laziness”, in which tools are used as mental shortcuts, compromising intellectual development instead of strengthening it.

It is precisely at this point that solutions developed with a focus on educational demands distinguish themselves from generic market offerings; by incorporating pedagogical principles into their architecture, they tend to be more specialized and effective regarding the central purpose of educational institutions: learning.

A representative example is the development of personalized tutors that combine large language models, such as Llama 3, with existing intelligent tutoring systems. In this integration, the language model does not operate separately: it accesses the context of the student’s interaction with the tutoring system in real time and, based on this, suggests personalized and pedagogically oriented messages to the carer. In practice, instead of the carer having to improvise how to support the child in a math task, the system offers concrete guidance, such as asking the student to explain the reasoning they used out loud, a strategy that is recognized for fostering metacognition.

Research on this model has indicated that this combination produces messages that are more useful and aligned with good pedagogical practices than the use of language models without anchoring in the instructional context. The fact that the solution is based on open architecture further adds advantages of transparency, control, and the possibility of local execution—making it auditable and adaptable to different institutional contexts.

This potential is already materializing in concrete tools. Students in different parts of the world interact with AI tutors based on Socratic dialogue — such as Khanmigo, Socratic Playground, Alex (aimed toward teaching English in the Czech Republic), and Rori, a mathematics tutor via WhatsApp in Ghana. As a result, there is a growing demand in the literature for specialized models — particularly in STEM components, where there is already evidence of the effectiveness of tutoring through AI agents with greater autonomy in reasoning and intervention (OECD, 2026). However, the technical architecture of most solutions used in schools is anchored in closed-source, proprietary models, which limits their auditability and contextual adaptation.

3.3.3. Uses for Management

The possibilities for using open AI applied to educational management are vast, spanning enrollment processes, dropout monitoring, communication with families, performance reports, indicators, and follow-up. We already have embedding models (systems that transform information such as text, images, and audio into numerical vectors) being used, for example, to map equivalencies between courses and guide admissions and curriculum analysis (Pardos e Borchers, 2026).

Different AI models designed for dropout monitoring offer potential benefits for anticipating problems and applying preventive interventions. Through algorithms trained with logs from platforms such as Moodle, it is possible to predict with high accuracy the risk of academic failure before the situation becomes definitive. A practical case occurred in TTK University of Applied Sciences, where an open-source system analyzed 99,104 records from 154 students to identify risks early on, demonstrating that it is a reproducible solution for institutions with their own training databases (Ovtsarenko, 2026). This possibility reinforces the role of technical openness: tools that guide critical decisions regarding a student's academic path must be fully auditable by those who adopt them, a requirement that proprietary systems, by their very nature, cannot fulfill.

For learning and management:

For learning, open AI can enable individualized, specialized, and adaptable tutoring by the institutions. For management, open models can be safeguard alternatives for caution in educational decisions.

Besides monitoring, there are also real opportunities to integrate course platforms and expand the offer of continuing education and professional development. By combining open infrastructures with open AI models, it is possible to enhance MOOCs (Massive Open Online Courses), strengthening more accessible and scalable learning ecosystems.

For example, Open edX, developed by MIT and Harvard and adopted by the governments of Israel, Ethiopia, Jordan, Saudi Arabia, and India for large-scale training, integrates in its Sumac version open models fine-tuned to the curricular data of the institutions themselves, enabling tutoring systems tailored to the local curriculum for the continuing education of networks.

3.4. Safeguards for the use of Open AI in Education

The possibilities opened by AI models in education are neither realized automatically nor are they free of risks. Responsible adoption requires administrators, teachers, and policymakers to develop a critical understanding of what these technologies do, how they do it, and for whom they are developed; especially when the systems are open and, therefore, modifiable by any agent with the technical capacity to do so, as discussed in the initial sections.

3.4.1. Challenges and risks of Open AI in Education

Challenges and risks of Open AI in Education
The adoption of open AI models in educational contexts cannot be treated as a merely technical decision. The unrestricted availability of models does not equate to safety, nor to a lack of costs, but rather to educational, technical, and pedagogical adequacy. Administrators, pedagogical coordinators, and teachers need to develop a capacity for critical curation. That is: the ability to systematically evaluate whether a specific model meets pedagogical, legal, ethical, and technical criteria before its adoption — considering safety, ethics, transparency, pedagogical intentionality, and more.

This curation is especially demanding in the case of open models because the absence of a central provider accountable for updates, bias corrections, and content adequacy transfers the responsibility for these verifications to the institution itself. Recognizing when an open model poses a risk, and not just when it offers a possibility, is essential for the responsible use of open AI in education.

The conscious use of open AI in education depends on the sharp perception that such models present both potential capabilities and vulnerabilities. Without intending to exhaust the subject or oversimplify reality, the following synthesis provides examples of practical applications across the pillars of teaching, learning, and management, connecting the advantages offered to the structural requirements of each implementation. It is essential to note, however, that success in any of these scenarios invariably demands robustness in infrastructure, institutional governance, and technical competence to ensure operational sustainability. Furthermore, each use may imply more or fewer necessary conditions depending on the purpose for which it is designed.

SOME USES AND REQUIREMENTS OF OPEN AI IN EDUCATION:

Type of Use	Benefit from the use of Open AI	Necessary Conditions
Teaching: Planning and materials	Creating lesson plans, assessments, and materials aligned with the local curriculum; adapting to low-infrastructure contexts.	National curriculum data incorporated into the model; pedagogical alignment, minimum connectivity.
Teaching: Assessment and feedback	Accuracy comparable to those of human evaluators; preserving data privacy on the institution's own servers.	Curriculum and pedagogical alignment, institutional servers for local execution.
Teaching: Teaching assistants	24h availability for questions, summaries, and rapid feedback; frees up teacher's time for complex pedagogical interactions.	Integration with institutional platforms (LMS, forums); teacher supervision to ensure quality and motivation.
Learning: Individualized learning	Responsiveness to student errors, a historical record if individual development that is auditable and adjustable by the institutions themselves, and scalability without dependence on a provider.	Technical capacity for fine-tuning; deliberate pedagogical design to prevent cognitive dependence.
Leaning: Socratic dialogue and active practice	Pacing and level personalization; access via basic devices (WhatsApp, SMS) in low-income contexts.	Specialized architecture: general-purpose models tend to provide ready-made answers instead of guiding the reasoning process.

Managemen: Dropout monitoring	Early risk prediction with high accuracy; a model that can be reproduced by any institution using its own data; total auditability of decisions.	Quality data (logs from platforms such as Moodle); institutional governance for the ethical use of trajectory data.
Management: Curriculum mapping and admissions	Course equivalence and admissions guidance using auditable embedding models.	Structured data infrastructure; strategic selection criteria, and specialized architecture.
Continuing education and MOOCs	Fine-tuning on local curricular data; adapted tutoring aligned with the local curriculum at scale.	Open platforms (Open edX); models fine-tuned with institutional data; long-term sustainability processes.

It is worth noting that, across all types of use, adopting open AI shifts the responsibility for curation, updates, and bias correction to the institution, requiring technical, pedagogical, and governance capabilities that proprietary solutions centralize within the provider. Choosing an open model is, therefore, not just a technical decision: it is an institutional decision regarding the level of autonomy and responsibility the organization is prepared to assume.

3.4.2 Literacy and critical thinking regarding Open AI

Beyond the careful selection of models, consolidating critical Artificial Intelligence literacy requires, among other competencies, recognizing the intricate geopolitical and economic dimensions that underpin contemporary technological development. In the absence of this literacy, students tend do develop a dependence on AI for text production and problem-solving, neglecting the development of their own cognitive capacities, whether analytical or argumentative, which results in a dissociation between task performance and genuine learning (OECD, 2026).

This risk is not exclusive to proprietary models. Favoring access without institutional intermediation can amplify thoughtless use, which means there is a demand for caution with open models as well; any institution, developer, or student can integrate an open model into an educational workflow without any guarantees of pedagogical curation, curricular adequacy, or teacher supervision. AI literacy, for those who learn as well as for those who teach, is therefore a condition of institutional responsibility to ensure that the use of models is responsible and safe under any circumstances, for both open models and proprietary ones, rather than a given premise.

To this end, various entities have been developing support frameworks for curriculum development, resources for AI literacy, and training initiatives across different scales. UNESCO, for instance, adopts a human-centered approach, emphasizing that technology should support human intellectual development and decision-making, rather than replace them.

In conclusion, the pedagogical possibilities offered by open AI depend directly on the implementation model and the intentionality of its use. Therefore, when choosing a specific model and level of openness, it is essential to weigh the benefits of each level of openness, the technical requirements involved, and the intended educational objectives. Below, we present a table with practical uses for each level of openness, considering the necessary legal, technical, and pedagogical criteria.

Level of Openness	Examples of use[1]
N1 **Closed proprietary** *Total control of the provider.* It favors agility and simplicity in operation; however, it creates dependence on the provider and lack of transparency regarding the system's functioning and training.	**Administrative and productivity tasks** • Creating announcements and newsletters for the school community. • Summarizing meetings and meeting minutes. • Preparing initial drafts of tenders, forms, and administrative documents. • Supporting the development of lesson plans, activities, and initial projects. **Why this level?** At this level, tasks that do not involve sensitive data can be carried out because they do not require model auditing.
N2 **Open Weights or Open Source Code** *Model weights are publicly available, partial control of the provider.* It offers more freedom to customize programming and response behavior, despite requiring specific infrastructure and technical expertise.	**Systems that process sensitive data locally** • AI agent running on local servers to support students and teachers. • Pedagogical chatbot customized to the curricula and guidelines designed for the school system or educational institution. • Internal LMS / SMS platforms connected to tutoring systems, diagnostics, and automated learning pathways (e.g., Moodle). • Creation of predictive platforms for educational analysis, formative follow-up, and student performance feedback. • Customized natural language processing tools. • Creation of interactive simulators and complex real-word problem-solving. **Why this level?** At this level, we address tasks that require local privacy for customization to the local and institutional context, but do not require auditing of the training process, so that the basic mathematical architecture of the AI is public and auditable, allowing local training on protected servers. **Practical example:** Twiga is a chatbot for Tanzanian teachers, developed by 'Tanzania AI Community and open-source contributors. Using RAG (retrieval-augmented generation) on the curriculum and textbooks from the Tanzania Institute of Education, Twiga creates assessments, lesson plans and materials aligned with the national curriculum, delivered via WhatsApp.

[1] Regardless of the technology's level of openness, establishing governance frameworks for the use of AI is essential. This includes the critical analysis of generative model outputs, as well as maintaining human supervision and decision-making authority at all stages.

N3
Broader Openness
Model weights, source code, and documentation are available, ensuring greater autonomy for developers.

It ensures greater independence in the creation process, architecture, and control over system reactions. However, its implementation demands a robust technological infrastructure, specialized knowledge, and more rigorous technical management.

Research, assessment and regional scale systems

- Automated essay grading with pedagogical auditing
- Bias analysis in large-scale diagnostic assessments
- Adaptive tutoring systems auditable by researchers
- Teacher training platforms with open metrics

Why this level?
This level addresses demands that require auditable methods capable of peer reproduction. The central objective is to ensure the transparency of processing workflows, providing credibility and trust to the responses predicted or generated by the systems.

Practical example: Global research demonstrates that predictive models, powered by deep learning and machine learning techniques, are viable tools for predicting student dropout. A relevant example is the use of entirely open-source systems that, by analyzing activity logs from Moodle and other learning platforms, can identify students at risk at an early stage. In this regard, the study conducted by Rebelo Marcolino et al. (2025), published in Scientific Reports, used multi-objective optimization and data balancing on these logs to enhance the accuracy of retention predictions.

N4
Total Openness
All components, including data training, are publicly available.

It ensures total independence in the creation process, architecture, and control over system reactions and databases. However, its implementation demands a robust technological infrastructure, specialized knowledge, and more rigorous technical management.

Digital sovereignty, public policies, and minority languages

- Language models for regional and local languages.
- National educational AI infrastructures auditable by the State.
- Open educational assessment training data for public research.
- AI systems for rural education with intermittent connectivity.
- Optimization of student allocation in schools based on transparency and equity criteria.

Why this level?
This level is recommended for scenarios that require sovereignty over data, workflows, and systemic responses. It ensures full autonomy and total control over processes, demanding competencies for the auditing, modification, and distribution of all constituent elements.

Practical example: The project AI4Bharat — IIT Madras, India, develops open models and databases for 22 constitutionally recognized Indian languages. The entire pipeline (data, code, weights, and documentation) is public. It allows Indian states to train educational models in local languages (such as Tamil and Telugu) without relying on external developers.

It is also essential to highlight that the different combinations within the gradation of openness components may vary depending on the developers and the institutional designs of the ecosystem. The examples and cases described here demonstrate, without exhausting the diverse possibilities, the potential that emerges from different contexts, economic conditions, and specific infrastructures.

4.0

Strategic Recommendations for Education

4. Strategic Recommendations for Education

In the previous sections, we covered the points that differentiate the levels of open AI, observed global trends, and identified possibilities and precautions for the use of these systems in education.

Now, we will revisit the lessons learned by breaking them down into actionable strategic recommendations that can be applied to different contexts. We present a set of guiding questions and recommendations categorized by ecosystem stakeholder, as the choices surrounding the adoption of open AI systems are numerous, and the following section can help navigate such complexity.



Source: Photo from the Magnific website

4.1. Guiding questions for decision-making

Managers and pedagogical teams can make decisions about open AI based on different criteria, which are always contextual. We have identified four layers that materialize the best practices and nuances identified throughout the debate on the subject, allowing for choices that preserve educational potential, transparency, equity, privacy, accountability, and meaningful supervision.

The guiding questions presented below serve as a flexible roadmap, subject to adjustments, adaptations, or expansions according to the particularities of each scenario. The lack of viable answers to most of these questions should be interpreted as a warning sign. Given that the implementation, use, and management of artificial intelligence technologies require prudence and ethical commitment, the absence of or difficulty in obtaining crucial data must be seen as a risk factor.

Pedagogical layer	**Concrete problem:** Which learning gap or teaching practice does the solution intend to address? **Evidence of potential:** Is there evidence — even if preliminary — that this technology delivers tangible benefits in this context? **Autonomy or substitution:** Does this tool enhance the judgement of the educator and/or the student, or does it bypass it? **Curriculum alignment:** Is the solution aligned with the curriculum and the learning competencies expected for this stage of education?
Technical layer	**Level of openness and licensing:** What is actually open, under which license, and who needs to understand these terms? **Available infrastructure:** Does the institution have the equipment and connectivity required for the proposed use? **Governance standards:** Does the solution meet security, interoperability, and sustainability requirements for public use? **Bias and performance evaluation:** Are there documented bias assessments, risk evaluations, and benchmark results relevant to the educational context?

Institutional and financial layer	**Motivation for adoption:** Is openness valued for its transparency and adaptability, or merely as a way to reduce costs? **Impact on teacher workload:** Does using the tool add, eliminate, or reorganize tasks in the daily work of its users (teachers, managers, and students)? **Risks of free or freemium models:** What externalities and vulnerabilities are embedded in solutions apparently without costs? **Risk of discontinuation:** What happens if the provider or maintainer of the open component of the open component ceases its development?
Ethical and legal layer	**Children's data protection:** Does the tool comply with the applicable legal framework for data protection, children's rights, and other relevant legislation? **Transparency of training data:** Are the sources of the data and any matters regarding intellectual property clearly disclosed? **Anonymization and data retention:** Does the tool establish clear policies for data processing, anonymization, and data deletion? Community participation: **Do families, carers, and civil society have meaningful** opportunities to participate in the selection of the adopted solutions?

4.2 Recommendations by stakeholder of the educational ecosystem

The collaborative artificial intelligence ecosystem will reach its full potential to the extent that different stakeholders converge on optimized decisions, investments, and practices. The following recommendations are aimed at the key segments regarding open AI in education, but their full effect depends on coordination and a multi-sectoral approach.

4.2.1. Educational institutions

Adoption by educational institutions must be responsible, progressive, transparent, and centered on political-pedagogical projects. Recommended best practices include:

Creating committees for curation and decision-making. The complexity of AI decisions should not rest on a single manager, teacher, or administrative sector. It is recommended to establish committees featuring managers, teachers, technical staff, and student representatives, tasked with evaluating, deciding, and reviewing options considering previously agreed-upon pedagogical, technical, institutional, and ethical criteria.

Considering context and local territories. Disruptive technologies amplify the strengths and can also emphasize the weaknesses of an educational system. Thus, importing a solution designed for another reality — another country, language, curriculum, or infrastructure — can scale problems instead of solving them. It is recommended to prioritize contextualized solutions, even if technically more modest, to generate a more measurable and controllable impact, rather than automatically adopting the largest or most visible model

Promoting AI literacy. Responsible and production adoption depends on empowering those who teach and those who learn, bridging conceptual knowledge, practical application, and ethical awareness. The objective is not only to operate the tools, but also to consider the meticulous process of selection and implementation as a learning opportunity for the people involved.

Conducting continuous documentation, recording, and evaluation. It is recommended to establish routines that register the options and decisions throughout the process, recording what was adopted, for what purpose, with what outcomes, with which metrics and indicators, and with what incidents. This accumulated knowledge is the variable that can distinguish responsible adoption from improvised consumption, allowing course correction before an error solidifies at scale. Additionally, it ensures that long-term success does not depend on specific individuals.

4.2.2. Development companies and EdTechs

The sector holds a strategic position in defining technical advancements and offering technology features. It is essential to establish processes for listening and assessment, rather than just product supply.

Promoting and embracing a culture of impact assessment. Throughout the development lifecycle, prioritizing decisions that enable and facilitate independent evaluation of the developed solutions. Ideally, it should consider the different viewpoints of those who will use or observe the tools: technical implementers, teachers, educators, and managers. Evidence and data must be prioritized over scalability

Moving toward total openness. Even though not every solution must necessarily be totally open, companies can make commitments to seek the maximum level of openness possible and, above all, be precise and transparent about the exact level they achieve. Terminological precision benefits the entire technology market, fosters trust among other stakeholders, and enhances the perception of the sector's integrity.

Seeking interoperability. Modular solutions that integrate with open platforms and standards and do not lock institutions into proprietary ecosystems, preserve institutional autonomy, and support long-term sustainability.

Establishing a commitment to non-negotiable values. Values that represent a consensus in AI governance, such as transparency regarding the system's operation and limitations, non-discrimination, data protection, and human supervision, benefit everyone and support the sustainable consolidation of a technology that is still seen as disruptive. When it comes to children, it is also essential to remember that the absolute priority of their comprehensive protection is legally mandated in the country.

4.2.3. Governments and States

Public institutions have the advantage of establishing regulatory processes and acting as a bridge between the different relevant sectors. Among the common recommendations commonly promoted by international organizations are:

Recognizing and promoting AI as a public good. Recognizing open AI as a strategic digital public good involves shifting the approach to technology beyond decisions on solutions and providers, toward considering it an infrastructure of collective interest. Its integration into educational, science and technology, and digital sovereignty policies can be achieved through a multi-sectoral approach under state coordination.

Strengthening structured and participatory procurement criteria. Decision-making for the funding and public procurement of AI requires prior needs assessment, competitive processes, and neutral functional specifications regarding the provider, preventing decisions from being reduced merely to the lowest price or the most well-known brand. The potential for scalability and the complexity of switching systems after implementation reinforce that specifications gain quality when they incorporate input from educators, school communities, and independent experts, alongside traditional evaluations by technical and legal teams.

Developing regulatory frameworks. Promoting guidelines that distinguish between low- and high-risk uses in the educational context, with specific safeguards for children and adolescents, providing educational networks with legal certainty while prioritizing precautionary approaches.

Advancing international collaboration. Since the frontier of AI is global, cooperation —especially South-South, in sharing infrastructure, language models, and pedagogical experiences — is a concrete strategy to redistribute capabilities for AI production, evaluation, and development in education.

To support strategic choices and critical analysis, we present below a synthesis of reflections structured by degrees of openness: **“Decision-Making Criteria Matrix”.** This framework aims to simplify the operationalization of decision-making processes, providing visual and clarifying support for the complex landscape in which this technical note is situated.

4.3. Decision-Making Criteria Matrix

Level of Openness	Open Components	What is gained	What is necessary	Risks and limits
Closed proprietary Explainability and transparency: **low** Popular examples: ChatGPT, Gemini, Microsoft Copilot, Claude, Manus.	~~Model weights~~ ~~Code~~ ~~Data~~ ~~Documentation~~	• Immediate use without configuration. • Technical support from the provider. • Automatic updates.	• Accepting the terms of use. • Maintaining a license or subscription plan. • Continuous connection to the service.	• Total technological dependence. • Student, strategic, and sensitive data stored on external servers. • Pedagogical and algorithmic opacity. • Vulnerable to changes by the provider. • Risk of openwashing.
Model weights Explainability and transparency: **intermediate** Popular examples: Llama 3, Mistral 7B, Falcon, Qwen3.	**Model weights** ~~Code~~ ~~Data~~ ~~Documentation~~	• Local execution — data remains on the local server. • Adjustments for the local context. • Independence from external API.	• Hardware infrastructure (GPUs). • Technical team for implementation. • Version and update management.	• Unknown training data. • Non-auditable biases. • High cost of infrastructure. • Risk of openwashing.
Broader Openness Explainability and transparency: **high** Popular examples: OLMo, BLOOM.	**Model weights** **Code** ~~Data~~ **Documentation**	• Development autonomy. • Possibility of independent auditing. • Scientific reproducibility. • Local curriculum adaptation. • Greater privacy control.	• Technical capacity to conduct audits. • Formal institutional governance. • Continuous code maintenance.	• Training data remains opaque. • Complex maintenance. • Continuous bias monitoring required. • Risk of openwashing.
Total Openness Explainability and transparency: **very high** Popular examples: ROOTS corpus, Projetos acadêmicos nacionais.	**Model weights** **Code** **Data** **Documentation**	• Full technological sovereignty. • Training data auditing. • Maximum contextual adaptability. • Maximum potential for transparency.	• Robust and sustainable infrastructure. • Multidisciplinary team. • Technical capacity to conduct audits. • Formal multi-stakeholder governance. • Long-term maintenance	• Risk of safeguards being removed. • High institutional cost. • Risk of irreversible decisions. Risk of openwashing.

5. Synthesis

In this document, we have presented a path from conceptual understanding to decision-making, equipping different education stakeholders with analytical and evaluative frameworks for examining open artificial intelligence. We seek to demonstrate that this ideal of openness is far from having a single or dogmatic answer; rather, it encompasses multiple nuances that must be understood in relation to their technical, budgetary, ethical, and pedagogical implications. Thus, the evidence, experiences, and international agreements already identified serve as a background and source of inspiration for concrete discussions.

There is no immediate technical or political threshold that defines and separates open AI from closed AI; what we have, in fact, is a gradient of components — most notably model weights, code, training data, and documentation — each with its own variables and nuances that establish its levels of transparency, usage restrictions, and adaptability. Thus, establishing the conditions and skills to component-by-component reading is the ideal horizon for institutional decision-making, since systems that are seemingly similar and labeled as open can offer vastly different autonomy and possibilities.

Open artificial intelligence does not mean that it is free of charge, without risks, or without governance requirements. Deploying any type of artificial intelligence-based digital technology requires the coordinated availability of essential resources, including funding, infrastructure, energy, technical expertise, and specialized human resources. Ignoring these requirements may turn technology adoption into a driver of exclusion, discrimination, and inequality. The comprehensive protection of children and adolescents, in accordance with data protection legislation and child protection frameworks in online environments, remains a non-negotiable priority. Attention must also be paid to openwashing: the merely rhetorical use of the label “open.”

Internationally, it is evident worldwide that not only global economic powers but also countries across the globe, including the Global South, are undertaking economic, scientific, and social efforts to promote open artificial intelligence as a strategic and geopolitical variable. By allowing advanced levels of transparency, interoperability, and local adaptation, open AI has the potential to enable educational advancements that respond to technological disruption in a way that reduces dependence on proprietary solutions.

To incorporate open AI with pedagogical intentionality and technical responsibility, the challenges are complex yet full of potential. We have mapped reflections and concrete examples for teaching, learning, and management, which include planning aligned with the curriculum framework and the development of competencies and skills, personalized tutors, and linguistic accessibility. These potential benefits, however, do not materialize automatically; they require careful curation and critical literacy as prerequisites. With a chain of accountability over the components and developers' responsibilities, institutions can manage implementation, its adaptations, and mitigate risks or setbacks.

On an operational level, we have converted these insights into recommendations for the various sectors involved. For educational institutions, we propose curation committees, decisions anchored in the political-pedagogical project, and continuous documentation. For developers and EdTechs, a culture of impact assessment, terminological precision regarding the actual level of openness, and a commitment to interoperability. For government institutions, the recognition of open AI as a public good, participatory and vendor-neutral procurement, risk-sensitive regulatory frameworks, and international cooperation. Finally, we have presented a matrix linking levels of openness to use cases, benefits, requirements, and limitations.

The advantages and disadvantages can inform choices so that decision-makers, before adopting any solution advertised as 'open,' apply guiding questions organized into layers (pedagogical, technical, institutional, financial, ethical, and legal) and verify what can actually be achieved. The lack of answers to many of these questions raises a red flag. Such precautions can also minimize the risk that the speed of our digital era and the ease of access amplify thoughtless use and metacognitive disengagement, in which dependence on the tool weakens genuine learning.

Ultimately, the conclusion is that it is possible to pursue and support a collaborative ecosystem anchored in transparency, equity, and the empowerment of both those who teach and those who learn. Openness, when well-understood and well-governed, is a pathway to placing technology at the service of comprehensive training and public interest, cutting across educational policies.

Referências

AMIEL, T.; COX, G.; DE LA HIGUERA, C. Open education principles: resisting metrics of AI black boxes. UNESCO, Paris, out. 2025. Disponível em: https://www.unesco.org/en/articles/open-education-principles-resisting-metrics-ai-black-boxes. Acesso em: maio 2026.

BAKER, R. S. et al. Generative AI as a teaching assistant. In: OECD. OECD Digital Education Outlook 2026: Exploring Effective Uses of Generative AI in Education. Paris: OECD Publishing, 2026. cap. 9. Disponível em: https://doi.org/10.1787/062a7394-en.

BAUER, E. et al. Looking beyond the hype: understanding the effects of AI on learning. Educational Psychology Review, v. 37, n. 45, 2025. Disponível em: https://doi.org/10.1007/s10648-025-10020-8.

BERNAL, M. P. et al. Aquisição de software de código aberto. Nota Técnica - Banco Interamericano do Desenvolvimento, 2025.

BITTENCOURT, I. I. et al. Positive Artificial Intelligence in Education (P-AIED): a roadmap. International Journal of Artificial Intelligence in Education, p. 1–61, 2023. Disponível em: https://doi.org/10.1007/s40593-023-00357-y.

BRASIL. IA para o bem de todos: Plano Brasileiro de Inteligência Artificial. Brasília, DF: Ministério da Ciência, Tecnologia e Inovação. Centro de Gestão e Estudos Estratégicos, 2025.

CENTRO DE INOVAÇÃO PARA A EDUCAÇÃO BRASILEIRA (CIEB). CIEB Notas Técnicas #5: orientações para seleção e avaliação de conteúdos e recursos digitais. São Paulo: CIEB, 2016.

CENTRO DE INOVAÇÃO PARA A EDUCAÇÃO BRASILEIRA (CIEB). Inovação Aberta em Educação. (Estudos 2). São Paulo: CIEB, 2016.

CENTRO DE INOVAÇÃO PARA A EDUCAÇÃO BRASILEIRA (CIEB). Modelos de curadoria de recursos educacionais digitais. (Estudos 5). São Paulo: CIEB, 2017.

CENTRO DE INOVAÇÃO PARA A EDUCAÇÃO BRASILEIRA (CIEB). CIEB Notas Técnicas #21: inteligência artificial na educação básica: novas aplicações e tendências para o futuro. São Paulo: CIEB, 2024.

DEIMANN, M.; LUKÁCS, B. From niche to regime: rethinking openness in the age of artificial intelligence. Journal of Interactive Media in Education, 2026.

DEMSZKY, D. AI tools for lesson planning and teacher support. In: OECD. OECD Digital Education Outlook 2026: Exploring Effective Uses of Generative AI in Education. Paris: OECD Publishing, 2026. cap. 8. Disponível em: https://doi.org/10.1787/062a7394-en.

DIGITAL EDUCATION COUNCIL. Digital Education Council Global AI Student Survey 2024. AI or Not AI: What Students Want. Disponível em https://www.digitaleducationcouncil.com/resource-library-items/digital-education-council-global-ai-student-survey-2024

DIGITAL Public Goods Alliance. Core Considerations for Exploring AI Systems as Digital Public Goods. Community of Practice Discussion Paper, 2019.

ESPAÑA. Estrategia de Inteligencia Artificial 2024. Ministerio para la Transformación Digital y de la Función Pública. Gobierno de España, 2024.

FACER, K.; SELWYN, N. Digital technology and the futures of education: towards 'non-stupid' optimism. Paris: UNESCO, 2021. 19 p. (Futures of Education: Learning to Become). Disponível em: https://unesdoc.unesco.org/ark:/48223/pf0000377071. Acesso em: maio 2026.

FREITAS, E. et al. Learning Analytics Desconectada: um estudo de caso em análise de produções textuais. In: WORKSHOP DE APLICAÇÕES PRÁTICAS DE LEARNING ANALYTICS EM INSTITUIÇÕES DE ENSINO NO BRASIL (WAPLA), 1., 2022. Anais. [S.l.]: SBC, 2022. Disponível em: https://doi.org/10.5753/wapla.2022.226823.

Referências

GEBRU, Timnit et al. Datasheets for datasets. Communications of the ACM, v. 64, n. 12, p. 86-92, 2021.
HU, B.; ZHU, J.; PEI, Y.; GU, X. Exploring the potential of LLM to enhance teaching plans through teaching simulation. npj Science of Learning, v. 10, art. 7, 6 fev. 2025. East China Normal University; Shanghai Institute of Artificial Intelligence for Education. Disponível em: https://doi.org/10.1038/s41539-025-00300-x. Acesso em: maio 2026.

INDIA. India AI Governance Guidelines: Enabling Safe and Trusted AI Innovation. Ministry of Electronics and Information Technology, Government of India, 2025.

INSTITUTO DE INTELIGÊNCIA ARTIFICIAL NA EDUCAÇÃO (IA.Edu). Escolas FigitAIs: IA a serviço da aprendizagem plena para todos. 1. ed. Maceió: Instituto IA.Edu, 2026.

ISOTANI, S. et al. AIED Unplugged: leapfrogging the digital divide to reach the underserved. In: INTERNATIONAL CONFERENCE ON ARTIFICIAL INTELLIGENCE IN EDUCATION (AIED), 2023. Communications in Computer and Information Science, v. 1831. Cham: Springer, 2023. p. 772–779. Disponível em: https://doi.org/10.1007/978-3-031-36336-8_118.

JAPÃO. Artificial Intelligence Basic Plan, 2026. Disponível em https://www8.cao.go.jp/cstp/stmain/20260619ai/aiplan_2601_draft_en.pdf.

JIA, L. et al. Fine-tuning large language models for educational support: leveraging Gagné's Nine Events of Instruction for lesson planning. arXiv preprint, 2025. Disponível em: https://arxiv.org/pdf/2503.09276.

LIESENFELD, A.; DINGEMANSE, M. Rethinking open source generative AI: open-washing and the EU AI Act. In: ACM CONFERENCE ON FAIRNESS, ACCOUNTABILITY, AND TRANSPARENCY (FAccT), 2024. Proceedings. [S.l.]: ACM, 2024. Disponível em: https://doi.org/10.1145/3630106.3658958.

MANYIKA, J.; BUGHIN, J. "As promessas e os desafios da era da inteligência artificial". McKinsey & Company, 15 Out. 2018, disponível em https://www.mckinsey.com/featured-insights/artificial-intelligence/the-promise-and-challenge-of-the-age-of-artificial-intelligence/pt-BR

MAZZUCATO, Mariana; EAVES, David; VASCONCELLOS, Beatriz. Digital public infrastructure and public value: What is 'public'about DPI?. UCL Institute for Innovation and Public Purpose's (UCL IIPP) publication series, 2024.

NOBEL, P.; ROZENSHTEIN, A. Z.; SHARMA, C. Untangling AI openness. Wisconsin Law Review, v. 2026, p. 169, 2026. Stanford University; University of Minnesota Law School; Fordham University School of Law. Disponível em: https://ssrn.com/abstract=5407422. Acesso em: maio 2026.

OECD. OECD Digital Education Outlook 2026: exploring effective uses of generative AI in education. Paris: OECD Publishing, 2026. Disponível em: https://doi.org/10.1787/062a7394-en.

OECD. Empowering learners for the age of AI: an AI literacy framework for primary and secondary education. Paris: OECD Publishing, 2025.

OPENEDX. Open edX Sumac Release. Versão de Dezembro de 2024. Disponível em https://docs.openedx.org/en/latest/community/release_notes/sumac.html

OSBORNE, Cailean. The business of open-sourcing artificial intelligence: commercial interests and collaboration dynamics in machine learning developer communities. Tese de Doutorado. University of Oxford, 2025.

OVTŠARENKO, O. et al. Using AI to forecast student dropout risk in technical education using a learning analytics approach. Scientific Reports, 2026. Disponível em: https://www.nature.com/articles/s41598-026-44919-1.

Referências

PARDOS, Z.; BORCHERS, C. AI in institutional workflows. In: OECD. OECD Digital Education Outlook 2026: Exploring Effective Uses of Generative AI in Education. Paris: OECD Publishing, 2026. cap. 11. Disponível em: https://doi.org/10.1787/062a7394-en.

PARIS, Tamara. From openness to openwashing: Framing a debated concept in AI. Dissertação apresentada ao Mestrado em Ciências da McGill University, 2025.

PINTO BERNAL, M. et al. Aquisição de software de código aberto. Nota Técnica IDB-TN-3090. Washington: Banco Interamericano de Desenvolvimento, 2025.

PORTELA, C. et al. AI in Education Unplugged: support equity between rural and urban areas in Brazil. In: INTERNATIONAL CONFERENCE ON INFORMATION AND COMMUNICATION TECHNOLOGIES AND

DEVELOPMENT (ICTD), 13., 2024. Proceedings. Nova York: ACM, 2024. p. 143–154. Disponível em: https://doi.org/10.1145/3700794.3700810.

RAMÍREZ-MONTOYA, María Soledad et al. Catalysing Open-Source Collaboration for AI Fluency for Educators. In: Open Science and the Future of Education. Routledge. p. 79-92.

REBELO MARCOLINO, M. et al. Student dropout prediction through machine learning optimization: insights from Moodle log data. Scientific Reports, 2025. Disponível em: https://www.nature.com/articles/s41598-025-93918-1.

REVISTA TRIÂNGULO. Relações entre os Cursos Online Abertos e Massivos (MOOCs) e a Inteligência Artificial (IA): uma revisão sistemática da literatura. Revista Triângulo, v. 18, e025001, 2025.

RIBEIRO, L. F. R. et al. Generating summaries with controllable readability levels. In: CONFERENCE ON EMPIRICAL METHODS IN NATURAL LANGUAGE PROCESSING (EMNLP), 2023. Proceedings. [S.l.]: ACL, 2023.

RIMA, S. et al. Systematic literature review on open educational resources recommender systems. International Journal of Interactive Mobile Technologies, 2022.

RIVAS, A.; BUCHBINDER, N.; BARRENECHEA, I. The future of Artificial Intelligence in education in Latin America. Fundación ProFuturo and the Organisation of Ibero-American States for Education, Science and Culture (OEI), 2024.

RODRIGUES, L. et al. MathAIde: a qualitative study of teachers' perceptions of an ITS Unplugged for underserved regions. International Journal of Artificial Intelligence in Education, 2024. Disponível em: https://doi.org/10.1007/s40593-024-00397-y.

SCARLATOS, A. et al. Training LLM-based tutors to improve student learning outcomes in dialogues. arXiv preprint, 2025. Disponível em: https://arxiv.org/abs/2503.06424.

SCHREPEL, Thibault; POTTS, Jason. Measuring the openness of AI foundation models: competition and policy implications. Information & Communications Technology Law, v. 34, n. 3, p. 279-304, 2025.

SURESH, H. et al. Participation in the age of foundation models. In: ACM CONFERENCE ON FAIRNESS, ACCOUNTABILITY, AND TRANSPARENCY (FAccT), 2024. Proceedings. Rio de Janeiro: ACM, 2024. p. 1609–1621.

SURMAN et al. Accelerating Progress Toward Trustworthy AI, Version 0.9 - For Public Input. Mozilla Foundation, 2024.

UNESCO. Steering AI and Advanced ICTs in Knowledge Societies: a rights, openness, access and multi-stakeholder perspective. Paris: UNESCO, 2021.

Referências

UNESCO. Draft Dubai Declaration on Open Educational Resources (OER): digital public goods and emerging technologies for equitable and inclusive access to knowledge. Paris: UNESCO, 2024.

UNESCO. AI Competency Frameworks for Students and Teachers. Paris: UNESCO, 2024. Disponível em: https://www.unesco.org/en/articles/what-you-need-know-about-unescos-new-ai-competency-frameworks-students-and-teachers.

UNESCO. AI and Education: Protecting the Rights of Learners. Paris: UNESCO, 2025. Disponível em: https://www.unesco.org/en/articles/ai-and-education-protecting-rights-learners.

UNIÃO Europeia. Parlamento Europeu e do Conselho. Regulamento (UE) 2024/1689 do Parlamento Europeu e do Conselho, de 13 de junho de 2024, que cria regras harmonizadas em matéria de inteligência artificial. Estrasburgo: Jornal Oficial da União Europeia, 2024.

VENUGOPALAN et al. Combining large language models with tutoring system intelligence: a case study in caregiver homework support. In: LEARNING ANALYTICS AND KNOWLEDGE CONFERENCE (LAK), 2025, Dublin. Proceedings. [S.l.]: ACM, 2025.

WANG, R. et al. Tutor CoPilot: a human-AI approach for scaling real-time expertise. Annenberg Institute at Brown University, 2025. Disponível em: https://doi.org/10.26300/81nh-8262.

WASSERMAN, C. et al. Educar na Era da Inteligência Artificial: caminhos para a BNCC Computação. São Paulo: Fundação Telefônica Vivo; Instituto IA.Edu, 2025.

XING, W.; NIXON, N.; CROSSLEY, S.; DENNY, P.; LAN, A.; STAMPER, J.; YU, Z. The use of large language models in education. International Journal of Artificial Intelligence in Education, v. 35, p. 439–443, 2025. Disponível em: https://doi.org/10.1007/s40593-025-00457-x. Acesso em: maio 2026.

ZHENG, Y.; HUANG, S.; ZENG, X.; HUANG, Y.; LIU, Z.; LUO, W. Knowledge-enhanced large language models for automatic lesson plan generation. Humanities and Social Sciences Communications, v. 12, n. 1, p. 1–14, nov. 2025. Jinan University; TAL Education Group. Disponível em: https://doi.org/10.1057/s41599-025-06004-2. Acesso em: maio 2026.