
Latin America Governance & Data Safety Dashboard¹

Arthur Thiago Marina
Lyra Dantas Gomes
Mirand Sousa Barbos
a de a
Indepe Azeve Wuha
ndent do n
researchIndepe Univer
her ndent sity
research
her

With
Apart Research

Abstract

Across Latin America, the rapid deployment of AI systems in public services is outpacing the normative frameworks designed to govern them. Existing global tools — such as the OECD.AI Policy Navigator and the IAPP Legislative Tracker — catalog policies at scale but do not produce standardized, verifiable scores that enable direct comparison across countries or distinguish between enacted law and aspirational policy. This gap is particularly consequential for the Global South, where regulatory ambition and enforcement capacity are systematically misaligned.

This paper presents a structured scoring instrument that evaluates eight Latin American countries — Brazil, Argentina, Chile, Colombia, Mexico, Uruguay, Peru, and Bolivia — across three dimensions: Algorithmic Governance (0–4), Data Sovereignty (0–3), and Infrastructure (0–2). Every indicator is grounded in primary official sources and scored through a transparent binary

¹ Research conducted at the [The Global South AI Safety Hackathon](#), June 2026. Track 1: Latin America. Sub-track: AI Governance.

methodology. The tracker is complemented by an AI-powered document analyzer that extracts normative content from uploaded legal texts, classifies it against the indicator framework via a pre-prompted language model, and writes scores directly to the tracker database — reducing what would otherwise require weeks of expert review to a near-instantaneous pipeline.

Key findings reveal a structural decoupling between AI governance ambition and enforcement infrastructure: Peru and Colombia lead on Algorithmic Governance yet lack sovereign cloud policies, while Bolivia imposes data localization requirements without any data protection law. The central takeaway is that high policy scores in the region reflect intent more than capacity — and that closing this gap is a prerequisite for safe AI deployment at scale.

1. Introduction

The rapid evolution of artificial intelligence (AI) and its integration into daily life have triggered urgent regulatory responses from governments worldwide. However, much like the technology itself, these regulatory frameworks have not developed linearly or equitably across the globe. Driven by disparities in digital infrastructure, public funding, and strategic priorities, different regions adopt different approaches toward governance.

One of the most critical dimensions of this divergence is the regulatory asymmetry between the Global North and the Global South. While the Global North leverages substantial structural capacity, technical expertise, and capital to set global standards that often prioritize market leadership or fundamental rights, the Global South faces a distinct reality. For these nations, AI regulation is intrinsically tied to socio-economic development, digital sovereignty, and mitigating the risks of data colonialism. Consequently, regulatory approaches between these two blocks are inherently diverse, reflecting a sociogeographical trend where a country's regulatory strategy follows its position in the global digital economy.

Our project was born from the need for researchers from other fields interested in AI and regulation to find a complete tool that compiles and displays data in an educational way. There are many AI governance and regulatory trackers available online. However, there remains a critical lack of structured data focusing on Latin America that provides the profound benchmarking necessary to orient regional policymaking or academic research. We provide the first quantitative AI governance tracker focused exclusively on Latin America with AI Safety scoring.

Prior tools like CELE Legislative Observatory, OECD.AI Policy Observatory, and ILIA 2025 Index either do not provide benchmarking scores, or cover only general AI capacity, or lack AI Safety metrics. To address this gap, our first contribution is providing an AI governance tracker with quantitative scoring. We created the first regulatory readiness database covering 8 Latin

American countries (Brazil, Chile, Argentina, Colombia, Mexico, Peru, Uruguay and Bolivia) with an explicit 0–9 score and 0–3 AI Safety Score (transparency + auditability + human oversight). The benchmarking system was developed by using previous international frameworks and research to enable AI governance evaluation.

2. Related Work

The most prominent reference on AI and digital governance tracking is the OECD.AI Policy Navigator (oecd.ai), a live database covering 900+ national AI policies across 60+ countries, organized by policy type and implementation stage. Similarly, the OECD.AI Index (2024–2025) offers a composite scoring framework measuring national AI capabilities across dimensions such as R&D, enabling infrastructure, and policy environment. At the regulatory layer, tools like the IAPP Global AI Legislation Tracker and White & Case AI Watch monitor the status of AI bills and laws by jurisdiction, while DLA Piper's Data Protection Laws of the World provides country-level legal analysis of data governance frameworks. For Latin America specifically, the Center for Cybersecurity Policy and Law's AI Governance in Latin America whitepaper (2024) and Brookings Institution analyses have mapped national AI strategies across ten regional countries, identifying governance gaps and the slow pace of regulatory implementation relative to global peers.

While these tools have a comprehensive scope, they share a fundamental limitation: they are descriptive and manually curated, requiring a team to read, classify, and update policy entries, a process and often translated into standardized, comparable scores. The OECD Navigator, for instance, catalogs what policies exist but does not produce a normalized score that allows direct cross-country comparison on specific sub-dimensions such as data localization, algorithmic transparency, or sovereign cloud infrastructure. Existing indices also tend to be global in scope and technically broad, diluting the nuance needed to evaluate Latin American countries, where the gap between enacted norms and operational policy is particularly significant.

Our work addresses this gap in two complementary ways. First, the tracker operationalizes a structured, multi-dimensional scoring framework, covering Algorithmic Governance, Data Sovereignty, and Infrastructure, calibrated specifically for eight Latin American jurisdictions and grounded exclusively in primary normative sources. Second, the integrated AI-powered document analyzer introduces a novel capability absent from all comparable tools: users can upload a legal or policy document, and a pre-prompted language model automatically extracts the relevant normative content, maps it to the indicator framework, and writes the resulting scores directly to the tracker database. This transforms what is typically a weeks-long expert review process into a near-instantaneous, reproducible, and auditable pipeline that jumps the process straight to human review.

3. Methods

Our project introduces a novel composite scoring system that measures AI governance maturity across three layered dimensions, each weighted by regulatory importance:

Category 1: Algorithmic Governance (0–4 points): The most critical dimension, measuring whether countries regulate core AI Safety elements. We evaluate four regulatory indicators: (1) transparency requirements (obligation to disclose AI system use and explainability), (2) auditability/impact assessment (mandatory third-party auditing or impact assessments for high-risk AI), (3) human oversight (requirements for human-in-the-loop decision-making or supervisory mechanisms), and (4) technical standards (existence of AI-specific technical regulations or certification frameworks). Each indicator receives 1 point if present in approved laws, 0 otherwise. This layer captures the AI Safety Score (0–3) as a sub-score (transparency + auditability + human oversight), explicitly identifying Safety gaps.

Category 2: Data Sovereignty (0–3 points): Measures data governance capacity. Indicators include: (1) comprehensive data protection law (e.g., LGPD, Ley 25.326), (2) active data protection authority (e.g., ANPD, AAIP), and (3) regulated international data transfer mechanisms. Each indicator receives 1 point if present. This layer captures regulatory infrastructure for data privacy and cross-border flows.

Category 3: Infrastructure (0–2 points): Measures technical capacity support. Indicators include: (1) cloud sovereignty policy (e.g., Nuvem Soberana, Government Cloud) and (2) data localization requirements. Each indicator receives 1 point if present. This layer captures infrastructure readiness for AI deployment.

Total Score (0–9): Sum of all layers ($4 + 3 + 2$). This weighted structure reflects that algorithmic governance is the most critical dimension ($4/9 = 44\%$ of total score), while data sovereignty ($3/9 = 33\%$) and infrastructure ($2/9 = 22\%$) are supportive but secondary.

The choice of Latin American countries for this project was also not random. We decided to mix countries with high, medium, and low regulatory maturity to enable contrast analysis, provide regional representation, and ensure all 8 chosen countries have publicly available legislative sources, are covered by recognized international organization databases, and have accessible data for validation. The classification of the group analyzed is the following:

Group A: Regulatory Leaders (5 countries): Brazil (highest regulatory maturity in LatAm, LGPD + PL 2338/2023), Chile (first with data center policy, AI law following EU model), Argentina (AI law 27.691 focused on innovation but weak oversight), Colombia (Ethical AI Framework 2021 + sectoral policies), and Mexico (LGPDPPSO + active regulatory debate). These countries represent the upper end of regulatory maturity (scores 6–9/9).

Group B: Intermediate Countries (2 countries): Uruguay (strong data governance via Ley 18.331, but less visible AI regulation) and Peru (AI Law 31814 + implemented regulation). These represent mid-tier maturity (scores 5–7/9).

Group C: Regulatory Gap Country (1 country): Bolivia (near-zero AI governance, no approved AI laws). This represents the lower end (score 0–2/9).

This scoring system was created based on previous AI governance frameworks but advances score construction with a dual-approach on governance and safety. The Algorithmic Governance

category is based on the European High-Level Expert Group on AI framework on evaluating AI systems, drawing its indicators from their four core regulatory elements: transparency, auditability, human oversight, and technical robustness. The Data Sovereignty and Infrastructure categories seek to operationalize the OECD AI Policy Observatory's five policy areas for AI governance: innovation, international cooperation, ethical values, safe and secure systems, and accountability.

We noticed a gap in previous AI governance trackers: while they provide information related to the number of laws and coverage of topics, there is no specific assessment about AI safety. This created a major oversight in the evaluation because a country may have high regulatory maturity and yet no basic safety principles, such as transparency, auditability, and human-in-the-loop.

In particular, Latin America seems to suffer from this issue. One of the systemic barriers for safe AI deployment in the Global South is regulatory gaps. Latin America has only 3 countries with approved laws on AI, and none have explicit AI Safety requirements, reinforcing the need for regional tracking tools with quantitative scores for benchmarking. Our dual-score approach seeks to overcome this major issue.

The data obtained from the “Hackathon Data” database were initially organized in an Excel spreadsheet according to the metrics defined for the analysis. Subsequently, these data were modeled and visualized using Power BI, resulting in the creation of classification tables for governance levels and Artificial Intelligence maturity scores.

To provide greater flexibility in data exploration and enable more detailed analyses, three interactive filters were implemented in the dashboard: Country, Category, and Indicator. These filters allow users to segment the information displayed in charts and maps, facilitating the identification of patterns and comparisons across different Artificial Intelligence (AI) governance and data security scenarios. The Country filter enables the selection of one or more countries for analysis. When applied, all dashboard visualizations are automatically updated to display only the data associated with the selected countries. This functionality supports cross-country comparisons and allows for the individual assessment of each country's performance according to the evaluated criteria.

The Category filter is designed to narrow the analysis to a specific evaluation dimension. Categories represent key aspects of AI governance and data management, such as governance frameworks, data sovereignty, and technological infrastructure. By selecting a category, users can focus on a particular area of interest and gain deeper insights into the corresponding results. The Indicator filter allows users to select specific indicators within each category. This feature provides a more granular level of analysis, enabling the evaluation of country performance according to particular criteria and helping identify strengths, weaknesses, and areas requiring further improvement or investment.

The combined use of these filters supports multidimensional data analysis, enhancing the accuracy and depth of result interpretation. Furthermore, the interactivity provided by Power BI contributes to a more comprehensive understanding of AI maturity levels, governance practices, and data security capabilities across the countries analyzed.



Figure 1: Proposed implementation roadmap

Figure presents the proposed implementation roadmap for the solution, illustrating the sequence of stages that will be developed throughout the project. The workflow represents the planned architecture for document processing, data extraction, validation, storage, and visualization, serving as a guide for the system's progressive implementation.

The implementation is expected to begin with the document upload functionality (Step 1), allowing users to submit files in formats such as PDF, DOCX, PNG, and JPG. These files will be received and managed by the web application (Step 2), which will act as the entry point to the processing pipeline.

As a security and privacy requirement, the solution will include a temporary file handling mechanism (Step 3). The uploaded document will be maintained only during processing and will not be permanently stored. Once the necessary information has been extracted, the original file will be discarded, ensuring compliance with data protection and privacy principles.

The next stage of development will focus on OCR and text extraction capabilities (Step 4), enabling the system to convert document contents into machine-readable text. The extracted text will then be made available temporarily in memory (Step 5) for subsequent processing.

Following text extraction, a Generative AI module (Step 6) will be integrated into the architecture. This component will receive the extracted text along with predefined prompts and will

be responsible for identifying and extracting relevant information from the document. The resulting process, referred to as Specific Data Extraction (Step 7), will transform unstructured content into structured business information.

The extracted information will then be formatted as a JSON response (Step 8), creating a standardized data structure suitable for integration with databases and analytical systems.

To ensure data quality and consistency, the implementation roadmap includes a Data Validation layer (Step 9), where business rules, verification procedures, and quality checks will be applied before any information is persisted.

After validation, only the structured and approved information will be stored in the database (Step 10). The original document will not be retained, reinforcing the security-oriented design of the solution.

Once the data storage layer is established, an ETL/ELT process (Step 11) will be implemented to transform, clean, and prepare the extracted information for analytical use. This stage will support the generation of reliable datasets for reporting and decision-making.

The final stage of the roadmap consists of developing a Business Intelligence (BI) Dashboard (Step 12), where the processed information will be visualized through interactive reports and dashboards. Technologies such as Power BI, Tableau, or Metabase may be employed to provide insights and support data-driven decision-making.

Therefore, the flowchart should be understood as a proposed implementation roadmap rather than a fully completed architecture. Each stage represents a component that will be progressively developed and integrated, culminating in a comprehensive solution that combines OCR, Generative AI, data validation, and business intelligence capabilities.

4. Results

The Digital Sovereignty Tracker was applied to eight Latin American countries (Brazil, Argentina, Chile, Colombia, Mexico, Uruguay, Peru, and Bolivia) across three dimensions: Algorithmic Governance (0–4), Data Sovereignty (0–3), and Infrastructure (0–2). Scores were derived exclusively from primary normative sources and validated against international reference databases (DLA Piper, IAPP, OECD.AI, USTR NTE 2026).

The most salient finding is a structural divergence between Algorithmic Governance scores and Data Sovereignty and Infrastructure scores, a pattern that holds consistently across the sample and is not an artifact of any single country's outlier performance. Peru and Colombia achieved the highest Algorithmic Governance scores (4/4 each), driven respectively by enacted legislation with binding regulations (Peru's Law 31,814/2023 and DS 115-2025-PCM) and a nationally approved policy framework (Colombia's CONPES 4144/2025). Brazil scored 3/4, anchored by the LGPD's provisions on automated decision-making rather than AI-specific law. By contrast, Bolivia, Mexico, Uruguay, Argentina, and Chile each scored 1/4 or 0/4 in this dimension, confirming that most of the region lacks enforceable AI governance norms. All eight countries, however, scored identically or near-identically on Data Sovereignty (6 out of 7 have both a data protection law and an active DPA), and Infrastructure scores were low across the board, with only Brazil and Uruguay presenting consolidated sovereign cloud policies.

This divergence is the central interpretive finding of the study. Countries with higher Algorithmic Governance scores are not the same countries with stronger data protection infrastructure, the two dimensions are decoupled. Peru, which leads on AI governance, has no sovereign cloud policy and only recently updated its data transfer regulations. Colombia, the co-leader on AI governance, has no enacted AI law and its cloud sovereignty rests on a single contractual instrument from December 2025. This suggests that Latin American countries are developing AI governance ambitions faster than the underlying digital infrastructure and legal data protection ecosystem required to make those ambitions operational. High AI governance scores, in this sample, reflect policy intent more than enforcement capacity, a distinction the document analyzer makes explicit by flagging whether the normative basis is an enacted law, a regulatory decree, or a policy instrument, each carrying different levels of legal bindingness.




Figure 2: Overall Scores and Category Breakdown per Country

The robustness of these findings is supported by three features of the methodology. First, every score is traceable to a named primary source with a direct URL, making the evidence independently verifiable and the scoring contestable on transparent grounds. Second, the binary indicator structure (0 or 1 per indicator) minimizes measurement ambiguity: a country either has an enacted norm meeting the criterion or does not, there is no subjective weighting.

5. Discussion and Limitations

The results carry implications that extend beyond regulatory mapping. The decoupling observed between Algorithmic Governance ambition and foundational infrastructure readiness suggests that Latin America faces a compounding AI safety risk: countries are signaling governance intent through policy frameworks and AI strategies precisely as AI systems become more widely

deployed in public services without the enforcement infrastructure to make those frameworks binding. A governance score built on a CONPES document or a national AI strategy is not equivalent to one built on enacted law with a functioning regulatory authority. The region's high policy output relative to low enforcement capacity creates a window of institutional vulnerability during which AI systems can operate in a normative grey zone: formally governed on paper, practically unregulated in practice.

This pattern is especially relevant for AI safety. Algorithmic transparency, human oversight, and impact assessments, the indicators where most countries scored zero, are precisely the mechanisms that safety researchers identify as foundational to detecting and correcting harmful AI behavior before deployment at scale. Their absence is not a secondary concern. It means that in five of the eight countries studied, no legal mechanism currently compels a government agency or private actor to explain an automated decision, audit a high-risk system, or ensure a human can override it. Bolivia scores zero across all four Algorithmic Governance indicators while simultaneously having data localization requirements, a configuration that prioritizes territorial control over data without any corresponding accountability for what systems process that data.

Limitations

Several constraints bound the interpretation of these results. First, the binary scoring structure, while ensuring transparency and replicability, necessarily flattens distinctions between norms of different legal weight. A constitutional provision, an enacted law, a regulatory decree, and a policy document all receive the same score of 1 if they satisfy the indicator criterion, yet their enforceability, sanctions, and institutional backing differ substantially.

Second, the study is limited to eight countries and three dimensions, it does not capture sectoral regulations (financial, health, electoral) that may impose additional AI governance requirements outside the general normative framework examined. Third, all scores reflect the normative landscape as of mid-2026, a rapidly evolving environment where a bill approved after the data collection date would not be reflected.

Most critically, the AI document analyzer component of the tool was not fully developed or tested within the hackathon timeframe. The pipeline architecture, document upload, text extraction, pre-prompted LLM classification, and database write, was designed and partially implemented, but end-to-end validation against the normative corpus was not completed. The scoring results presented in this paper were therefore produced through manual expert review, not through the automated system. This means the analyzer's accuracy, consistency across document types (PDFs, scanned laws, legislative bills), and robustness to ambiguous normative language remain unverified assumptions. If the LLM classification proves unreliable on edge cases, such as provisions that partially satisfy an indicator, the automation pipeline could introduce systematic scoring errors that the current manual approach avoids.

Future Work

The most immediate next step is completing and stress-testing the document analyzer against the full normative corpus, measuring classification accuracy relative to the manually produced ground truth. Beyond that, the framework can be extended in three directions: (i) expanding country coverage to all of Latin America and the Caribbean; (ii) adding a temporal dimension that tracks score evolution year-over-year, enabling trend analysis rather than a static snapshot; (iii) and deepening the indicator granularity to distinguish between norm type (law vs. policy), enforcement authority (dedicated regulator vs. general jurisdiction), and sanction regime (administrative fines vs. criminal penalties). Integration with real-time legislative monitoring APIs could enable the tracker to flag normative changes automatically, making it a genuinely live instrument rather than a periodically updated database.

6. Conclusion

This paper presented the Latin America Governance & Data Safety Dashboard, a structured scoring tool designed to assess the normative readiness of eight Latin American countries across three dimensions of digital sovereignty: Algorithmic Governance, Data Sovereignty, and Infrastructure. By grounding every indicator in primary official sources and applying a transparent binary scoring methodology, the tracker produces comparable, verifiable, and contestable country profiles, a level of traceability absent from most existing regional assessments. The central empirical finding is a systematic decoupling between AI governance ambition and enforcement infrastructure: countries like Peru and Colombia lead on algorithmic governance scores while lacking the sovereign cloud policies, operational data protection ecosystems, and enforcement mechanisms needed to translate normative intent into binding accountability. Bolivia presents the most paradoxical profile, data localization requirements in the absence of any data protection law or AI governance norm, while Brazil emerges as the most structurally coherent case, with layered normative instruments spanning data protection, classified data localization, sovereign cloud, and AI strategy.

The broader implication is that the Latin American AI governance landscape is at a critical inflection point. The region is producing policy frameworks and national AI strategies faster than it is building the institutional and legal infrastructure to enforce them, creating a window of genuine AI safety risk during which high-risk automated systems may operate in public services with no binding transparency, oversight, or audit requirements. The Digital Sovereignty Tracker, and particularly its AI-powered document analyzer when fully operationalized, offers a scalable instrument to monitor this gap continuously, flag normative changes in real time, and support policymakers, civil society, and researchers in holding governments accountable to the governance commitments they have already made on paper.

Code and Data

The dataset was manually built based on research from academic articles and official sources, and then structured in Excel for use in Power BI.

Data/Datasets: Power BI link:

https://drive.google.com/drive/folders/1eCt1r5_K_vGPtEF7L9SIVT8Omxg4D587?usp=sharing

References

CELE (Centro de Estudios en Libertad de Expresión y Acceso a la Información) (2025) Legislative Observatory on AI in Latin America. Buenos Aires: Universidad de Palermo. Available at: <https://cele.palermo.edu.ar> [Accessed: June 2026].

Center for Cybersecurity Policy and Law (2024) AI Governance in Latin America. Washington, D.C.: Center for Cybersecurity Policy and Law. Available at: <https://www.centerforcybersecuritypolicy.org/insights-and-research/ai-governance-in-latin-america> [Accessed: June 2026].

IAPP (International Association of Privacy Professionals) (2025) Global AI Law and Policy Tracker. Portsmouth, NH: IAPP. Available at: <https://iapp.org/resources/article/global-ai-legislation-tracker/> [Accessed: June 2026].

IAPP (International Association of Privacy Professionals) (2025) Global Privacy Law and DPA Directory. Portsmouth, NH: IAPP. Available at: <https://iapp.org/resources/global-privacy-directory/> [Accessed: June 2026].

ILIA (Índice Latinoamericano de Inteligencia Artificial) (2025) ILIA 2025: Latin American Artificial Intelligence Index. Santiago: CENIA / CEPAL. Available at: <https://cenia.cl/indice-ia-latam/> [Accessed: June 2026].

OECD (2026) The OECD.AI Index: Technical Paper. Paris: OECD Publishing. Available at: <https://oecd.ai/en/ai-publications/the-oecd-ai-index-technical-paper> [Accessed: June 2026].

OECD.AI (2025) OECD.AI Policy Navigator. Paris: OECD. Available at: <https://oecd.ai/en/dashboards/overview> [Accessed: June 2026].

White & Case LLP (2026) AI Watch: Global Regulatory Tracker — OECD. New York: White & Case. Available at: <https://www.whitecase.com/insight-our-thinking/ai-watch-global-regulatory-tracker-oecd> [Accessed: June 2026].

DLA Piper (2025) Data Protection Laws of the World. [Online database]. Available at: <https://www.dlapiperdataprotection.com> [Accessed: June 2026].

IAPP (International Association of Privacy Professionals) (2025) Global AI Law and Policy Tracker. Portsmouth, NH: IAPP. Available at: <https://iapp.org/resources/article/global-ai-legislation-tracker/> [Accessed: June 2026].

IAPP (International Association of Privacy Professionals) (2025) Global Privacy Law and DPA Directory. Portsmouth, NH: IAPP. Available at: <https://iapp.org/resources/global-privacy-directory/> [Accessed: June 2026].

ILIA (Índice Latinoamericano de Inteligencia Artificial) (2025) ILIA 2025: Latin American Artificial Intelligence Index. Santiago: CENIA / CEPAL. Available at: <https://cenia.cl/indice-ia-latam/> [Accessed: June 2026].

OECD (2026) The OECD.AI Index: Technical Paper. Paris: OECD Publishing. Available at: <https://oecd.ai/en/ai-publications/the-oecd-ai-index-technical-paper> [Accessed: June 2026].

OECD.AI (2025) OECD.AI Policy Navigator. Paris: OECD. Available at: <https://oecd.ai/en/dashboards/overview> [Accessed: June 2026].

White & Case LLP (2026) AI Watch: Global Regulatory Tracker — OECD. New York: White & Case. Available at: <https://www.whitecase.com/insight-our-thinking/ai-watch-global-regulatory-tracker-oecd> [Accessed: June 2026].

LLM Usage Statement

ChatGPT was used to generate images in this document. Claude was used as a support to structure and correct the grammar and flow in this document. All results and claims were independently verified.