PIT POLICY LAB

Public Interest Technology

AI Adoption Challenges in Mexico

FROM INTEREST TO IMPLEMENTATION

Perspectives from business, civil society, and research

Mark Cherches
September 2026

CONTENTS

Executive summary 3

Introduction 4

Methodology and limitations 5

1. Uneven Starting Points 6

2. Moving from Experimentation to Implementation 7

3. The Advantages and Constraints of Smaller Firms 9

4. Trust and Organizational Acceptance 10

5. Data, Privacy, and Governance 11

6. Sector Differences 13

7. Policy Implications 15

Conclusion 17

Works cited 18

Acknowledgments 19

EXECUTIVE SUMMARY

This report combines desk research with five semi-structured interviews to examine how organizations in Mexico move from interest in AI to sustained implementation. Two interviews concern health startups, one concerns civil society, and two provide research and cross-sector perspectives. The interviews offer evidence about adoption mechanisms and disagreements rather than a representative account of Mexico's productive sector. Across the material reviewed, needs change with digital maturity: basic infrastructure and skills precede implementation, while trust, validation, and governance become more prominent as use expands.

WHAT THE EVIDENCE SUGGESTS

The desk research and interviews indicate that there is no one single barrier preventing adoption across Mexico's productive sector. Organizations begin from very different levels of digital maturity, and the most important constraint changes as adoption progresses. For firms with limited digital foundations, the challenge begins with connectivity, structured data, modern information systems, and basic technical skills. Organizations that have started experimenting with AI are more likely to struggle with identifying useful applications, redesigning workflows, assigning responsibility, and evaluating whether a tool creates enough value to justify its cost. More mature adopters increasingly confront questions of trust, systems integration, validation, cybersecurity, governance, and regulatory uncertainty.

01 **Foundations**

Connectivity, structured data, modern systems, and basic skills.

02 **Implementation**

Useful applications, workflows, responsibility, and evaluation.

03 **Scale and governance**

Trust, integration, validation, cybersecurity, and regulatory clarity.

Author's synthesis of the desk research and five interviews. The sample supports interpretation of mechanisms, not estimates of their prevalence.

INTRODUCTION

Mexico's AI-adoption challenge is increasingly one of implementation rather than awareness. Interest in AI is widespread, investment is growing, and organizations across the productive sector are experimenting with generative AI and other tools. However, relatively few firms have integrated AI into their operations in a sustained or strategic way. Centro México Digital reports AI use among 8.0% of Mexican firms with at least ten employees in 2023, using INEGI's 2024 Economic Census, alongside an OECD average of 20.1% drawn from countries' latest available observations, mostly 2025. These are contextual benchmarks rather than a like-for-like comparison. Within Mexico, adoption ranges from approximately 6% among small firms to 17% among large firms. Significant regional differences also remain, with adoption reaching 12.6% in Nuevo León compared with 5.3% in Chiapas (Centro México Digital, 2026).

The gap is not simply the result of companies being unfamiliar with AI. A 2025 survey found that 41% of surveyed Mexican organizations had general knowledge of AI but no business case, while another 31% wanted to implement AI but lacked certainty about how to proceed. Only 27% reported having a defined AI strategy (KPMG México, 2025). These findings suggest that many organizations have moved beyond asking whether AI is relevant but still struggle to determine where it can create value, how it should be incorporated into existing work, and which internal or external capabilities are needed to implement it responsibly.

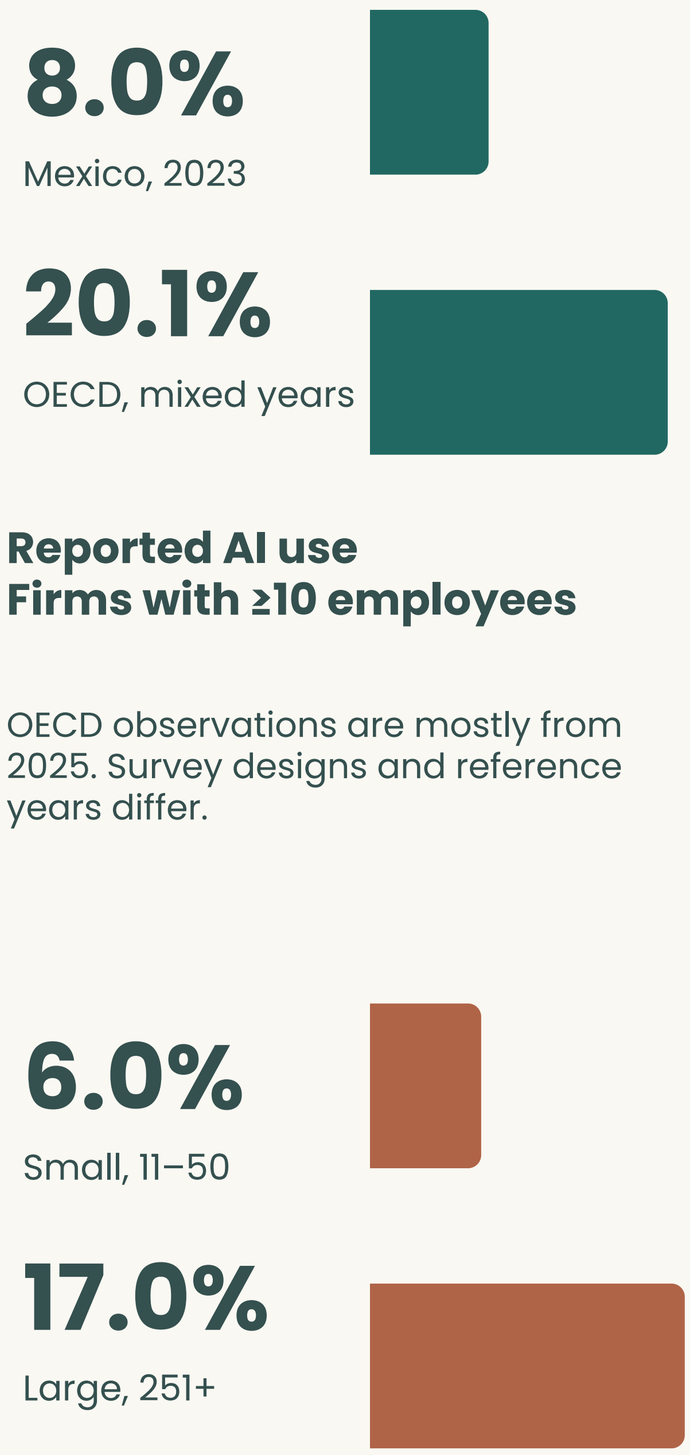

41%
General knowledge of AI, no business case

31%
Wanted AI, unsure how to implement

27%
Defined AI strategy

Sources: Centro México Digital (2026), pp. 6–8, using INEGI 2024 census data for 2023; OECD country data compiled by CMD. Strategy figures: KPMG México (2025).

METHODOLOGY AND LIMITATIONS

This synthesis combines desk research on AI adoption in Mexico and Latin America with five completed semi-structured interviews: Cristina Campero Peredo, founder and CEO of PROSPERiA Health; Jorge Madrigal, co-founder and CEO of Vivanta; Juan Manuel Casanueva, co-founder and director of SocialTIC; Hiram Ponce, director of the Center for Innovation and Technological Development at Universidad Panamericana and president of the Mexican Society for Artificial Intelligence; and Leonardo Martínez Flores, Director of Research at Centro i para la Sociedad del Futuro.

Interviewees were selected purposively for their expertise and direct experience with AI adoption in their fields and organizations. The common interview guide covered current uses, implementation decisions, barriers, trust, data governance, support needs, and policy, with follow-up questions tailored to each participant's role. All five participants agreed to be named and to have their comments attributed in the report.

SCOPE OF THE EVIDENCE

Given that the interview sample is concentrated in health and cross-sector perspectives, it is used primarily to explain mechanisms, identify disagreements, and assess how the broader research appears in practice. Findings about manufacturing, financial services, and agroindustry remain grounded mainly in the desk research and should be tested through additional interviews with companies in those sectors.

The sample does not include government officials, business chambers such as CONCAMIN, CANACINTRA, or COPARMEX, workers, or interviewees selected to represent lagging regions. The regional statistics in the introduction establish context, but the interviews do not explain those regional differences. With two of five interviews drawn from health startups, the sample cannot establish the prevalence of barriers across sectors or regions.

No direct financial-services company was included in the completed interview sample, making this an important area for additional research. Interviews should test how incumbent banks and fintechs make build-versus-buy decisions, which regulatory uncertainties are delaying specific applications, and whether adoption is improving access for underserved customers or primarily increasing efficiency within existing markets.

5 **interviews**

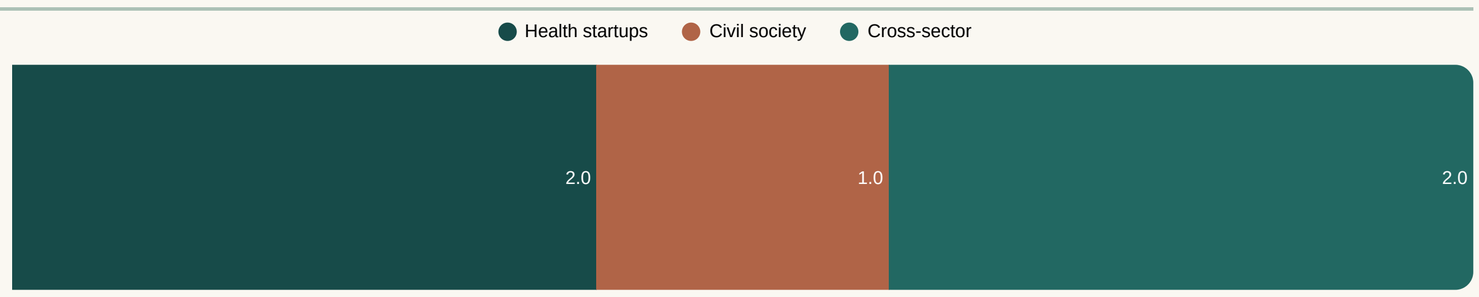

Interview composition described above; counts indicate perspectives, not sector representation.

1. UNEVEN STARTING POINTS

National adoption figures conceal substantial differences in organizations' underlying readiness. Based on INEGI's 2024 Economic Census, with 2023 as the reference year, Centro México Digital reports that 22.4% of Mexican microenterprises had a computer or tablet and 23.6% had internet access. Among medium and large firms, each measure exceeded 90%. These are separate indicators, not a combined measure of access. Even among firms with at least ten employees, large companies adopt AI at nearly three times the rate of small firms. These differences suggest that AI adoption remains closely connected to the earlier process of business digitalization (Centro México Digital, 2026, pp. 6, 23).

Hiram Ponce described a similar divide based on his work with companies in manufacturing, agroindustry, and health. According to Ponce, large companies are beginning to apply computer vision, machine learning, autonomous robotics, and natural-language processing to operational problems, while many small and medium-sized firms remain limited to generative AI tools for office work or employee assistance. He explained that family-owned firms may still rely on manual or highly traditional processes and often do not consistently collect the operational data that more advanced AI applications require. In these organizations, the immediate barrier is not the availability of an AI model, but the absence of the underlying information and systems needed to use one. (Hiram Ponce Interview, 2026).

Juan Manuel Casanueva identified a comparable pattern among the civil-society organizations supported by SocialTIC. Casanueva explained that many organizations, particularly those working in rural or underserved communities, continue to face basic technology and data-literacy constraints. Staff may begin using AI because it has been incorporated into familiar platforms, but this informal use does not necessarily mean that the organization understands how the technology works, has identified appropriate applications, or knows how to protect sensitive information. Casanueva distinguished this natural adoption from a more strategic and critical form of adoption in which an organization deliberately evaluates where AI can be useful and where it should not be used. (Juan Manuel Casanueva Interview, 2026).

Taken together, the national data and interviews show that AI adoption begins well before an organization selects a model or vendor. For many firms and civil-society organizations, the first requirements are still connectivity, structured data, cybersecurity, basic digital literacy, and the consistent use of existing technology. Policies or products that begin with advanced AI implementation may therefore miss the more immediate barriers facing less digitally mature organizations.

22.4% Microenterprises Computer or tablet

23.6% Microenterprises Internet access

Source: Centro México Digital (2026), p. 23, based on INEGI's 2024 Economic Census. Reference year: 2023. Measures are separate.

2. MOVING FROM EXPERIMENTATION TO IMPLEMENTATION

A second group of organizations has sufficient digital capacity to experiment with AI but lacks a clear implementation path. This is the group most clearly captured by the KPMG Survey findings: organizations understand the technology and may want to use it, but have not identified a strong business case or determined how implementation should proceed. Their challenge is less about obtaining AI and more about translating general interest into a specific organizational decision.

Decisions behind implementation

1

Problem

What should improve?

2

Ownership

Who is responsible?

3

Workflow

How does AI fit people and systems?

4

Evaluation

How will success be assessed?

STUDY SAMPLE

122

Mexican workers

López-García et al. (2025)

The broader literature supports this interpretation. The OECD's research on AI adoption in firms finds that successful implementation depends not only on access to technology, but also on complementary capabilities such as skilled workers, usable data, collaboration with external institutions, and access to technology-diffusion support (OECD, BCG, and INSEAD, 2025). Research involving 122 Mexican workers similarly found that organizational support, workforce training, ethical governance, and participatory workplace cultures were at least as important as technical readiness. (Lopez-Garcia et al., 2025)

Leonardo Martínez Flores argued that many organizations confuse AI adoption with the digitalization of existing processes. Martínez explained that a firm may convert an analog process into digital form and add an AI tool without reconsidering how the process itself should change. In his view, meaningful adoption requires a broader transformation involving strategy, culture, responsibilities, and employee practices. Otherwise, organizations may automate existing processes without addressing whether those processes are effective or whether employees understand the role of the new system. (Leonardo Martínez Flores Interview, 2026)

The health interviews provide examples of more integrated adoption. Cristina Campero Peredo explained that PROSPERiA develops its core clinical computer-vision models internally and increasingly uses external generative-AI tools in software development, employee training, internal operations, and patient communication. Potential applications are discussed with the relevant operational teams before being assigned to machine-learning and development staff. One example involved replacing the manual allocation of health promoters' time with an AI-assisted process, beginning with an assessment of which parts of the workflow consumed the most time and which decisions required greater care.

Jorge Madrigal described a similarly broad use of AI at Vivanta. The company uses AI to analyze patient-reported behaviors, forecast health trajectories, determine when human intervention is necessary, and support internal finance and management decisions. Madrigal explained that decisions about new applications typically begin with practical questions about how much value the tool could create, how long implementation would take, and whether the benefit justifies the cost.

Two implementation examples

PROSPERiA	Manual allocation of health promoters' time	→ AI-assisted allocation
Vivanta	Patient-reported behaviors	→ Health forecasts and human-intervention decisions

Source: Campero Peredo and Madrigal interviews (2026). Interview-based examples.

These examples show that productive adoption is tied to a recognizable organizational problem rather than the general objective of "using AI." Organizations must decide what they are trying to improve, who will own the implementation, how AI will interact with existing workers and systems, and how success will be evaluated. The technology may be available, but the organization still needs the capacity to translate it into a functioning workflow.

AUTHOR'S SYNTHESIS

The technology may be available, but the organization still needs the capacity to translate it into a functioning workflow.

3. THE ADVANTAGES AND CONSTRAINTS OF SMALLER FIRMS

The national statistics show that larger firms adopt AI at higher rates, but the interviews suggest that firm size affects different stages of adoption in different ways. Large organizations possess more resources, data, specialized employees, and formal technology teams, but their implementation processes may be slower because decisions involve more departments, compliance reviews, and layers of approval.

Firm size and implementation

Dimension	Smaller AI-capable firms	Larger organizations
Capacity	Limited time and specialist staff	More resources, data, and specialized teams
Decisions	Quick decisions by a small leadership team	More departments, compliance reviews, and approvals

Source: Author's synthesis of interviews with Campero Peredo and Madrigal (2026); firm-size context from Centro México Digital (2026).

Campero explained that PROSPERiA's small size allows it to implement a new application relatively quickly once leadership decides that the use case is valuable. At the same time, the company does not have unlimited staff available to monitor new AI products and evaluate every possible use. If more resources were available, she would dedicate someone specifically to tracking new tools and developments.

Madrigal described the same tension at Vivanta: smaller companies may hear about relevant tools later because they lack dedicated monitoring capacity, but they can move quickly once a small leadership team decides to implement something. Both interviews suggest that the principal constraint for an AI-capable startup may not be technical implementation itself, but the time required to discover, compare, and prioritize possible applications.

This complicates a simple interpretation of the firm-size adoption gap. Smaller firms are generally less likely to possess the digital foundations required for AI, but those that are technically prepared may be more agile than large organizations. Support should therefore distinguish between firms that need foundational digitalization and firms that already possess technical capacity but need help evaluating tools, identifying use cases, or managing implementation.

4. TRUST AND ORGANIZATIONAL ACCEPTANCE

Trust emerged as a consistent barrier across the interviews, but it takes distinct forms for different groups: clinicians, who bear professional liability for AI-assisted decisions; patients, whose data and care are at stake; employees, whose roles may change; and institutions such as pharmaceutical and insurance companies, which face regulatory uncertainty. Desk research on Mexican health care found that only 9% of physicians use AI tools in daily practice, while 71% describe themselves as only somewhat informed and 18% report no knowledge of AI uses in healthcare. Although attitudes are generally positive, physicians continue to raise concerns about privacy, algorithmic bias, implementation costs, and the effect of AI on the doctor-patient relationship. (Culebro, 2025).

Campero identified clinician trust as one of PROSPERiA's most important early barriers. The company deliberately presented its technology as a screening or pre-diagnostic tool rather than as a replacement for ophthalmologists, allowing clinicians to view AI as an additional source of support. Campero also observed that widespread use of tools such as ChatGPT has made AI more familiar to clinicians and patients, reducing some of the ambiguity that previously surrounded the technology.

Madrigal offered a more differentiated perspective. He reported that patients generally care about the quality and accessibility of care rather than whether AI is involved, while clinicians remain more resistant to technological change. His view of patients is a single-source observation that sits in tension with the desk research on privacy concerns, and should be read as one perspective rather than a general finding. Larger pharmaceutical and insurance companies were described as particularly cautious because they were uncertain about the direction of Mexican regulation and the risks of becoming involved with an AI-enabled product. (Madrigal, 2026; Culebro, 2025).

Casanueva described a different form of skepticism among civil-society organizations. Their concerns were connected to privacy, cybersecurity, labor displacement, environmental impacts, and the handling of information about vulnerable communities. SocialTIC helps organizations establish their own internal agreements about AI use, including which information should never be uploaded to a commercial platform and which tasks require human judgment. These discussions may lead an organization to adopt AI selectively rather than broadly, or to seek offline tools that reduce privacy risks.

The interviews show that resistance should not automatically be treated as a lack of understanding. In some cases, skepticism reflects legitimate uncertainty about professional responsibility, sensitive data, employment, or potential harm. Adoption is more likely to be sustained when organizations can explain what a system does, identify its limitations, preserve meaningful human oversight, and respond to the concerns of employees or affected communities.

5. DATA, PRIVACY, AND GOVERNANCE

The ServiceNow Enterprise AI Maturity Index 2026 identifies data quality, systems integration, and governance as recurring constraints on AI adoption. Its global findings emphasize problems with data accuracy, accessibility, and management, alongside legacy systems and regulatory complexity. The report's Mexico discussion highlights greater concern about legacy-system integration than in most countries. These findings support treating data and integration as implementation challenges, while distinguishing the report's global survey results from its country-specific observations (ServiceNow, 2026, pp. 10–11).

Data quality

Accurate, accessible, well-managed data

Integration

Connections to legacy systems

Governance

Controls, oversight, and risk management

Source: ServiceNow, Enterprise AI Maturity Index 2026, pp. 10–11. Qualitative themes from global findings and the Mexico discussion; no country percentages implied.

The importance of data changes as organizations mature. For the smaller firms described by Ponce, the challenge may be that operational information is not collected or stored in a usable form. For more advanced organizations, the issue becomes whether the data are accurate, representative, secure, and appropriate for the intended application.

Campero emphasized that clinical AI must perform equitably across gender, race, and socioeconomic background, particularly if it is intended to reduce gaps in access to healthcare. PROSPERiA already maintains privacy and security practices because it handles sensitive patient information, but the company must also ensure that its models are validated and that results are communicated with appropriate clinical oversight.

Casanueva explained that data governance may lead civil-society organizations to reject particular uses entirely. Organizations working with migrants, journalists, or human-rights defenders may determine that personal or operational information should never be entered into a commercial AI platform, even when the tool could improve efficiency. Once organizations understand how commercial systems process information, decisions about acceptable use often become clearer.

GOVERNANCE IN CONTEXT

These findings show that governance becomes more important as AI moves into consequential work. Organizations need internal rules covering data handling, permitted uses, responsibility for errors, human review, system monitoring, and the process for suspending a tool that performs poorly. As a policy recommendation, adopting organizations should remain accountable for the ways they incorporate AI into their work and for its effects on employees, customers, patients, or communities. Vendors should disclose limitations and support validation, while government should establish minimum protections and credible avenues for accountability. This proposed allocation of responsibilities is the author's synthesis, informed by the OECD AI Principles and UNESCO's Mexico readiness assessment. Legal obligations depend on each actor's role and the applicable data-protection and sector rules (OECD, 2024; UNESCO, 2024).

EXISTING MEXICAN REQUIREMENTS

Existing data-protection and sector requirements already shape some AI uses in Mexico. The following distinctions provide legal context for the recommendations, without implying that every AI application is regulated in the same way.

Personal data

For personal data held by private entities, the LFPDPPP establishes principles including purpose, consent, proportionality, and responsibility, alongside security duties. The current law assigns enforcement functions to the Secretaría Anticorrupción y Buen Gobierno. Responsibilities depend on the actor's role in processing personal data (LFPDPPP, arts. 5, 18, 38–39).

Health software

Health software that qualifies as a medical device falls within COFEPRIS's device framework. NOM-241-SSA1-2025 defines software as a medical device and includes validation requirements. These provisions do not make every use of AI in healthcare a medical device (NOM-241-SSA1-2025, §§3.118, 11.16).

Financial technology

For regulated financial-technology institutions, the Fintech Law assigns supervision to the CNBV and Banco de México within their respective competencies. Sector oversight therefore remains relevant alongside data protection (Ley para Regular las Instituciones de Tecnología Financiera, art. 3).

Legislative proposals

The legislative debate also includes proposals such as the Ley General para el Desarrollo Ético y Tecnológico de la Inteligencia Artificial, introduced in the Chamber of Deputies on 17 September 2025. That initiative is cited here as a proposal, separately from existing law. Martínez's concerns about fragmented legislative efforts should be read in this context (Cámara de Diputados, 2025; Martínez Flores, 2026).

The OECD AI Principles and UNESCO's Mexico readiness assessment provide additional guidance on accountability, human oversight, and institutional preparedness. They are policy frameworks and diagnostics, not substitutes for Mexican legal requirements (OECD, 2024; UNESCO, 2024).

6. SECTOR DIFFERENCES

Manufacturing

Manufacturing is central to Mexico's economy and nearshoring strategy, but the reviewed research does not provide a reliable sector-wide AI-adoption rate. Existing evidence documents applications in predictive maintenance, computer-vision quality inspection, robotics, and process optimization, while also suggesting that advanced adoption remains concentrated among large companies. Industrial SMEs face a combination of legacy systems, limited data collection, shortages of middle-skill automation and data workers, and uncertainty about whether examples drawn from multinational plants apply to their own operations (American Industrial Magazine, 2026).

Ponce's experience reinforces this divide. He described large companies experimenting with sophisticated robotics and computer-vision applications while smaller manufacturers remain at an earlier stage of digitalization. The main challenge is therefore moving from general demonstrations of industrial AI to applications that can be integrated into Mexican mid-market plants. Firms need evidence of returns under local conditions, technical assistance connecting new tools to older equipment and enterprise systems, and training for the workers expected to operate or supervise the technology. (Hiram Ponce Interview, 2026).

Financial Services

Financial services encompass banking, lending, payments, insurance, and investment management. Within this broader sector, fintech companies represent a more mature AI-adoption environment. The evidence reviewed here concerns fintechs rather than incumbent banks, so it should not be generalized to financial services as a whole. The Finnovista Fintech Radar reports that 77% of Mexican fintechs use AI and 27% operate under an AI-first model (Finnovista, 2026). Data-rich applications in fraud detection, credit assessment, customer service, and automation provide comparatively clear business cases.

The sector's challenge is less about demonstrating that AI can create value and more about ensuring that it is secure, explainable, and inclusive. Automated financial decisions may reproduce existing inequalities if the underlying data are incomplete or biased, particularly in a country where access to formal financial services remains geographically and socioeconomically uneven. Clearer model-governance expectations, auditing practices, data protections, and consumer safeguards will become increasingly important as AI is incorporated into consequential decisions.

Agroindustry

Evidence from agroindustry remains case-based rather than representative. Applications such as AI-assisted livestock monitoring show potential value in disease detection, animal welfare, and productivity, including systems that process information locally to accommodate limited rural connectivity. However, the sector faces infrastructure constraints, uncertain affordability at smaller scales, limited access to technical support, and unclear standards for AI-enabled monitoring (Guillermo Solomon, 2025).

Ponce described advanced agricultural companies experimenting with autonomous robots, computer vision, and machine learning, while emphasizing that adoption remains much more difficult for smaller firms. The most promising applications may therefore be those designed around rural operating conditions rather than adapted from large urban firms. Local processing, shared services, producer associations, regional universities, and public-private pilots could help demonstrate value and reduce the cost of adoption for smaller producers. (Ponce, 2026).

Health

Health combines significant potential with particularly demanding trust, data, and validation requirements. Desk research indicates low routine AI use among physicians, fragmented health-information systems, and limited familiarity with clinical applications (Culebro, 2025). The interviews nevertheless identified practical applications in retinal screening, patient monitoring, health-trajectory analysis, triage, and internal operations (Campero Peredo, 2026; Madrigal, 2026).

9%

Physicians using AI in daily practice

Culebro (2025)

Campero presented AI as a way to address shortages of specialists by helping identify patients who need ophthalmological attention. Madrigal similarly described AI as a means of extending chronic-disease monitoring without requiring staff to contact every patient manually. Both cases demonstrate how AI could expand access where Mexico’s healthcare infrastructure is limited. (Campero Peredo, 2026; Madrigal, 2026).

At the same time, the interviews show that technical performance alone will not produce adoption. Broader use depends on local validation, clear responsibility, secure patient-data practices, professional acceptance, and confidence that AI will support rather than replace clinical judgment.

7. POLICY IMPLICATIONS

The evidence does not support a single intervention for every organization or sector. Policies should reflect differences in digital maturity, firm size, sector, and region. Organizations that lack basic digital foundations need infrastructure, data systems, and introductory training, while more mature firms may benefit more from implementation support, validation standards, financing, and regulatory clarity.

Support needs and digital maturity

Basic digital foundations

Infrastructure
Data systems
Introductory training

More mature adopters

Implementation support
Validation standards
Financing and regulatory clarity

Author's synthesis of the policy recommendations in this section.

Education emerged as a particularly important theme. Ponce argued that universities must update curricula across professional fields because AI is becoming relevant not only to programmers, but also to engineers, physicians, lawyers, and managers (Ponce, 2026). Martínez similarly identified the absence of clear educational guidance as a major problem extending from elementary school through higher education and professional research (Martínez Flores, 2026). The desk research supports these concerns: the World Economic Forum reports that Mexican employers expect AI to transform operations and anticipate growth in big-data and information-security roles (World Economic Forum, 2025).

Government, universities, chambers, and technical institutions could also reduce the implementation burden by providing trusted, vendor-neutral assistance. Readiness assessments, shared testing facilities, sector-specific demonstrations, pilot support, and independent validation would be particularly valuable for SMEs that cannot maintain dedicated AI teams. Public incentives could reduce the financial risk of adoption, although support should be tied to continued implementation and measurable value rather than the purchase of a tool or completion of a short pilot.

AUTHOR'S SYNTHESIS

Support should be tied to continued implementation and measurable value rather than the purchase of a tool or completion of a short pilot.

The interviews also revealed disagreement about regulation. Campero argued that Mexico needs clearer health-AI rules, formal certification pathways, and reasonable approval timelines. Although regulatory ambiguity allows experimentation, she explained that it also creates uncertainty and permits lower-quality products to compete without equivalent validation. Madrigal was more skeptical of government intervention and worried that premature regulation could slow useful innovation, even as he acknowledged that regulatory uncertainty discourages pharmaceutical and insurance companies from engaging with AI products.

Different views on regulation

Cristina Campero Peredo

Clearer rules, certification pathways, and reasonable approval timelines.

Jorge Madrigal

Concern that premature regulation could slow useful innovation.

Regulatory uncertainty affects adoption

Source: Author's synthesis of Campero Peredo and Madrigal interviews (2026).

Ponce and Martínez both emphasized the lack of coordination across existing policy efforts. Ponce favored a central institution or national council capable of establishing direction, while Martínez described legislative initiatives developing without a common understanding of the problems being addressed. Casanueva added that coordination would not be sufficient unless civil society and affected communities had a meaningful role alongside government and industry, particularly in discussions involving privacy, labor, environmental impacts, and data-center development.

Taken together, these perspectives suggest that Mexico needs clearer policy direction without a rigid, uniform approach. Risk-based guidance, predictable sector-specific pathways, minimum expectations for data protection and human oversight, and opportunities for supervised experimentation could reduce uncertainty while preserving flexibility. The process should also include organizations and communities affected by AI rather than limiting policy development to government and industry representatives.

Participation in policy design

Government	Industry	Civil society	Affected communities

Author's synthesis of the interview perspectives on inclusion and coordination.

CONCLUSION

Mexico’s AI-adoption gap is not explained by one shortage or one policy failure. The desk research establishes substantial disparities by firm size, region, and sector, while the interviews show how those disparities appear inside organizations and how the barriers change as adoption progresses.

For many smaller firms and civil-society organizations, AI adoption remains constrained by foundational issues such as connectivity, structured data, modern systems, and workforce literacy. Organizations that have crossed that threshold often struggle to identify worthwhile use cases and integrate new tools into existing operations. More mature adopters increasingly confront trust, validation, privacy, governance, and regulation.

The most effective response is therefore not simply to increase access to AI technology. It is to help organizations develop the capabilities required to choose appropriate applications, integrate them into real workflows, evaluate their results, and govern them responsibly. This will require:

- Different forms of support across sectors and stages of maturity.
- Stronger university–industry connections.
- Implementation-oriented services.
- Clearer regulatory pathways.
- Greater attention to the workers and communities affected by adoption.

AUTHOR’S SYNTHESIS

The central challenge for Mexico is no longer only whether organizations are interested in AI, but whether they can turn that interest into sustained implementation that produces measurable value without weakening privacy, accountability, or public trust.

WORKS CITED

Source titles link to the full publications. Personal interviews are identified by name and date.

American Industrial Magazine. "Inteligencia Artificial Industrial en México 2026: Casos Reales, Datos Oficiales y Ventajas para la Manufactura." 18 Feb. 2026.

Cámara de Diputados. Initiative for the Ley General para el Desarrollo Ético y Tecnológico de la Inteligencia Artificial. Gaceta Parlamentaria, 17 Sept. 2025.

Campero Peredo, Cristina. Personal interview. 9 Aug. 2026.

Casanueva, Juan Manuel. Personal interview. 18 Aug. 2026.

Centro México Digital. Inteligencia artificial en México: De la promesa al impacto económico—Adopción empresarial e implicaciones económicas. 2026.

Culebro, Enrique. "From Hesitation to Innovation: AI in Mexico's Medical Landscape." Mexico Business News, 13 Feb. 2025.

Finnovista, Mastercard, and Galileo Financial Technologies. Finnovista Fintech Radar México 2026. Finnosummit, 2026.

IMCO Staff. "Reporte Futuro del Empleo 2025." Instituto Mexicano para la Competitividad, 24 Jan. 2025.

INEGI. Censos Económicos 2024: Resultados definitivos. 2025. Reference year 2023.

KPMG México. Panorama de la inteligencia artificial en México y Centroamérica 2025: Estrategias para acelerar la adopción efectiva. 2025.

López-García, Jessica, et al. "Barriers to AI Implementation in the Mexican Workforce: Social and Organizational Insights." ICAI-TEMS, IEEE, 2025. doi:10.1109/ICAI-TEMS64510.2025.11467701.

Madrigal, Jorge. Personal interview. 9 Aug. 2026.

Martínez Flores, Leonardo. Personal interview. 19 Aug. 2026.

México. Ley Federal de Protección de Datos Personales en Posesión de los Particulares. 20 Mar. 2025; amended 14 Nov. 2025. Arts. 5, 18, 38–39.

México. Ley para Regular las Instituciones de Tecnología Financiera. 9 Mar. 2018. Art. 3.

OECD. OECD AI Principles. Updated 2024.

OECD, Boston Consulting Group, and INSEAD. The Adoption of Artificial Intelligence in Firms: New Evidence for Policymaking. OECD Publishing, 2025. doi:10.1787/f9ef33c3-en.

Ponce, Hiram. Personal interview. 19 Aug. 2026.

Secretaría de Salud. NOM-241-SSA1-2025, Buenas prácticas de fabricación de dispositivos médicos. 2025. §§3.118, 11.16.

ServiceNow. Enterprise AI Maturity Index 2026. Research with ThoughtLab, 2026. pp. 10–11.

Solomon, Guillermo. "Inteligencia Artificial y videovigilancia en la agroindustria: El futuro del monitoreo animal." DPL News, 9 July 2025.

Stubrin, Lilia, et al. Startups x AI: An Overview of Artificial Intelligence Adoption in Latin America and the Caribbean. Inter-American Development Bank, May 2026. doi:10.18235/0014037.

UNESCO. México: Evaluación del estadio de preparación de la inteligencia artificial. 2024.

World Economic Forum. The Future of Jobs Report 2025. Chapter 5, Mexico. 2025.

ACKNOWLEDGMENTS

I would like to express my deepest thanks to Cristina Martínez Pinto, CEO of PIT Policy Lab, for her guidance and support throughout this entire journey. Having no prior experience in AI policy research, I worried that this project would be overwhelming, difficult to navigate, and intimidating. After meeting with Cristina for the first time, however, I knew this would not be the case. Instead, Cristina made both my research and my overall time with PIT Policy Lab exciting, engaging, and rewarding.

Whether designing numerous tools that helped me track my research or helping me source points of contact at numerous organizations, Cristina dedicated herself to supporting my research with her expertise and wisdom. Cristina, thank you so much for creating and supporting what has been an incredible educational experience. I look forward to implementing everything I have learned from you in my future endeavors!

I would also like to thank everyone involved in Stanford's Tech Ethics & Policy Fellowship, which made this research possible. Specifically, I want to thank Ashlyn Jaeger, Alexandra Cambra, and Drew Spence for overseeing the fellowship and the teaching of ETHICSOC85: Introduction to Tech Ethics and Policy Career Pathways. Both the class and fellowship have fundamentally shaped who I am as a student and professional - I am immensely grateful for the knowledge and support.

Mark Cherches

Lead Researcher

Mark Cherches is a rising sophomore at Stanford University from Baltimore, MD, studying Symbolic Systems with a concentration in Artificial Intelligence.

He is deeply passionate about the intersection of technology, economics, and society, with a particular emphasis on applying AI to drive meaningful, real-world impact. With experience across a range of technical and data-driven work, Mark brings both analytical rigor and practical execution to solving complex problems.

