WORLD WIDE WEB
FOUNDATION

HOW ARE GOVERNMENTS IN LATIN AMERICA USING ARTIFICIAL INTELLIGENCE?

A proposal for effective and legitimate implementations of AI systems in the public sector

Governments across the world are exploring how to leverage artificial intelligence (AI) and algorithms for decision making and service delivery. While this interest extends to Latin America, the use of AI in the region is still very limited.

The nascent use of this technology presents an opportunity to develop clear policies and practices that will empower public officials interested in implementing AI to ensure that their deployments are both effective and legitimate, maximising the benefits of new technologies while minimising the potential risks to the population.

www.webfoundation.org

A look at Argentina & Uruguay

Based on four case studies from Argentina and Uruguay, two of which are presented below, we observe that although a specific official may have the power to implement an AI system, he or she does not manage or control all the components that will determine its impact. Thus, to understand the effectiveness and legitimacy of an AI system's implementation, collaboration across government units and with other stakeholders is needed.

Key questions public officials should answer for AI systems

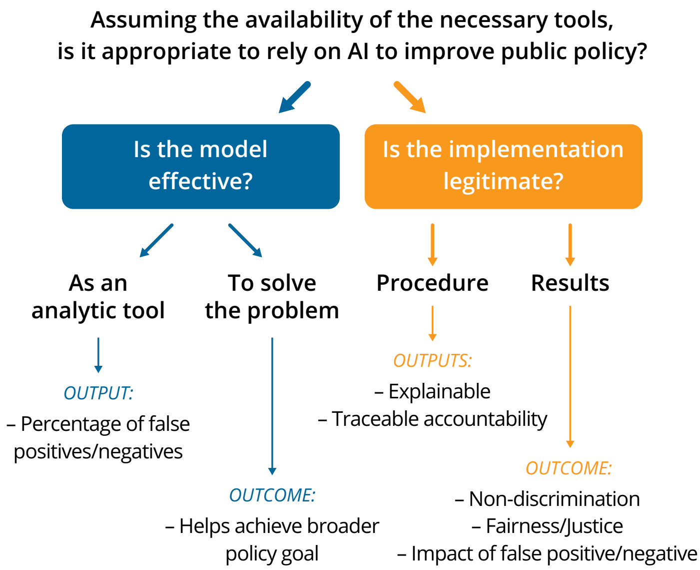

*We suggest that public officials consider **four key areas** to assess the effectiveness and legitimacy of an AI system's implementation:*

1. The process of dataset creation, e.g.:

- Who determines what data to collect?
- Who is included within the data?

2. The setup and design of AI tools, e.g.:

- What variables were included?
- Do they trigger risk of undue discrimination?
- Are outputs explainable? To whom? How?
- How do the outputs compare to human prediction or other non-AI systems?

3. The administrative protocols that surround the tool's output, e.g.:

- Does the tool inform human decisions, or execute policies directly?

4. Interaction with broader social and legal norms target populations are subject to, e.g.:

- Are there mechanisms of appeal for citizens who are impacted by decisions made by AI tools?
- What other safety-nets are available to those who are denied a service?
- How will the community treat a person who the AI classified in a certain way?

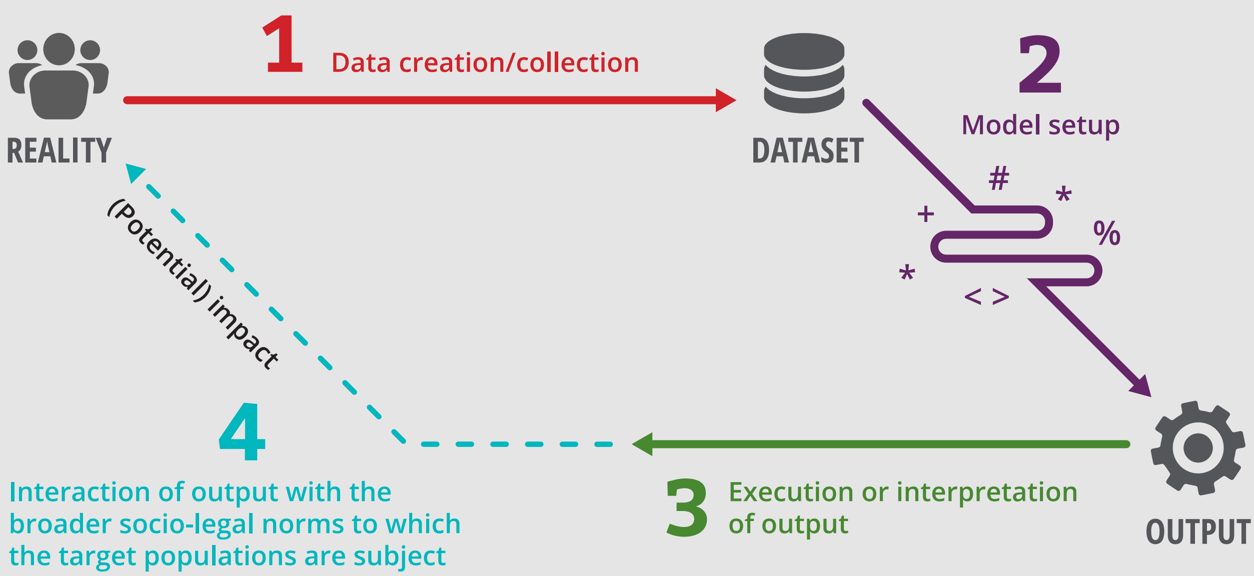

Applying the framework:

A SNAPSHOT OF TWO CASE STUDIES

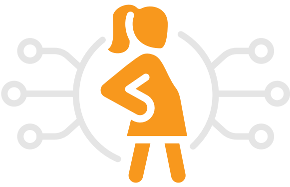

Objective:

PREDICT ADOLESCENT PREGNANCY

Responsible unit: Ministry of Early Childhood, Salta, Argentina

Development: Private (Microsoft) / State

Source of data: Surveys of the Ministry of Early Childhood

Public availability of data: No

Selection of variables: Up to 78 variables (depending on age and sex)

Variables capable of triggering undue discrimination: Yes (nationality)

Model intelligibility: Black box

Output: Assigns a probability of pregnancy to each sampled woman (15-19 years olds), and identifies the subset with highest probability.

Reported error rate: 15% false positives

Executes or assist decisions: Assists

Consequence: Ministry coordinates actions with other Ministries

Impact (reported): Model identified 250 adolescent women with a +70% probability of pregnancy. Impact of subsequent interventions unknown.

SUMMARY OF FINDINGS

Effectiveness of the model: The reported accuracy is high, but independent specialists have questioned the methodology. More openness is necessary to resolve these questions and ensure public trust.

Effectiveness of implementation: The Ministry has not consolidated information about the impact, nor does it have a publicly facing protocol regarding the actions that are taken when a pregnancy is detected.

Legitimacy (process): The software informs the system operator of the key variables driving the percentages presented as output (potential for explainability). Yet the databases are closed, and the design of the model is not available for public scrutiny, thus limiting traceability and effectiveness of independent audits.

Legitimacy (results): No public document consolidates information on the impact of these tools, making judgements on fairness difficult. Local gender specialists have argued it violates the privacy of minors while failing to acknowledge the underlying structural inequality that limit women's ability to exercise their sexual and reproductive rights effectively.

Objective:

PREDICT WHERE A CRIME WILL OCCUR

Responsible unit: Ministry of the Interior of Uruguay

Development: Private (PredPol)

Source of data: Ministry of the Interior

Public availability of data: No

Selection of variables: Type of crime, location, date and time

Variables capable of triggering undue discrimination: Indirectly (location)

Model intelligibility: Black Box

Output: Creates $150m^2$ sections on a map signalling areas with high probabilities of crime to occur within a given time period.

Reported error ratio: Unknown

Executes or assist decisions: Assists

Consequence: Deployment of officers to the area

Impact (reported): Crime was not reduced in absolute terms. There were [small] reductions in areas where it was implemented.

SUMMARY OF FINDINGS:

Effectiveness of the model: There is no available information about PredPol's predictive accuracy. Its replacement by an in-house developed tool suggests it was lower than expected.

Effectiveness of implementation: A reduction of crimes was detected, but only in areas in which it was deployed, not overall. Some argue these tools only displace crime scenes rather than reduce crime rates.

Legitimacy (process): Data is not public, and proprietary nature of software suggests not even the Ministry had access to the black box within which PredPol operated, making the outputs inexplicable.

Legitimacy (results): Local and international organisations have argued that tools like PredPol tend to replicate the biases of the databases used for their training, and are relied on to legitimise unfair and discriminatory practices.

Recommendations for policy-makers in Latin America

1. Develop local infrastructure, expertise and normative frameworks

In Uruguay, black-box proprietary software was used to deploy security forces until in-house experts were brought in. In Argentina, sensitive government data is still being stored abroad because the country lacks secure servers. To ensure projects are sustainable, scalable, and sensitive to the local context, states must foster a local ecosystem, investing in people, infrastructure and normative frameworks. Its nourishment will require training and greater coordination between specialists from a variety of fields, including public officials, regulatory and oversight bodies, citizens, civil society organisations, developers and enterprises.

2. Define risk-assessment criteria for models and implementations

In Latin America, a growing interest in AI tools, paired with an undeveloped local ecosystem for AI could lead to the rapid implementation of prefabricated models, designed in and for other contexts. Given the risks these systems hold, government officials need criteria to help them distinguish models that are potentially very problematic from those that are safer.

3. Promote transparency, public engagement and accountability practices in each of the four key stages

AI has the potential to perpetuate and reinforce existing inequities and biases. The case studies suggest public officials often focus on getting an AI system running, but pay little attention to ensuring the legitimacy of its implementation. Beyond active collaboration across government units that manage each of the four key areas, we recommend implementing practices of transparency, public engagement and accountability at each of the four key stages. Transparency practices, such as publishing metadata on collected data can help ensure the AI model is set up appropriately and that potential biases are detected. Public engagement, such as ensuring that members of affected populations participate in defining the variables used to build the model is key to ensuring that the public trusts these tools. Accountability practices in the following phases, such as town halls in which impact is regularly discussed with the target populations, is key to monitoring unexpected effects and possible feedback loops.

Interested in our research? Get in touch.

contact@webfoundation.org

@webfoundation

This snapshot comes from a report that was written by Juan Ortiz Freuler and Carlos Iglesias, with country-specific research contributions from ILDA (Iniciativa Latinoamericana por los Datos Abiertos).

Supported by a grant from the Omidyar Network.

The Web Foundation was established in 2009 by Sir Tim Berners-Lee, inventor of the World Wide Web. Our mission is to establish the open web as a public good and a basic right.

WORLD WIDE WEB FOUNDATION

World Wide Web Foundation
1110 Vermont Ave NW, Suite 500
Washington, DC 20005, USA

www.webfoundation.org