THE ROAD TO EDUCATIONAL INCLUSION: FOUR STEPS TO DEVELOP SYSTEMS TO PROTECT EDUCATIONAL PATHWAYS

STEP 2

Designing early warning systems:

From systems based on expert knowledge and indicators to artificial intelligence

Elena Arias Ortiz, Cecilia Giambruno, Natalia González Alarcón,
Marcelo Pérez Alfaro, Cristina Pombo and Roberto Sánchez Ávalos

IDB

This document was prepared in the framework of the fAIr LAC initiative of the Inter-American Development Bank (IDB), which seeks to promote the ethical and responsible use of data and AI-based systems in the region, especially in the provision of social services.

Table of contents

1 Introduction

This document is the second of four publications in the series *"The Road to Educational Inclusion: Four Steps to Develop Systems to Protect Educational Pathways"*. It aims to serve as a guide for education ministries and secretariats in the region that are interested in designing and implementing a pathway protection system.

The series consolidates existing knowledge about the protection of students' educational pathways in a context in which the challenge of exclusion is ever more widespread in Latin America and the Caribbean (LAC). Even before the COVID-19 pandemic, educational exclusion rates in LAC were already very worrying. In addition, it is estimated that school closures during the health crisis could have an impact that, in terms of educational exclusion, reaches more than 3 million of the region's children and young people from preschool to tertiary level (UNESCO, 2020), affecting particularly those who are most vulnerable (Acevedo et al., 2020). This calls for a redoubling of efforts to find systemic evidence-based responses, using the new technologies available, to ensure that the region's children and adolescents have a real opportunity to pursue uninterrupted and complete educational pathways, guaranteeing their right to education.

The objective of a system to protect educational pathways is to build the conditions so that the journeys of children and adolescents within the education system are continuous, complete and high-quality, reducing lags and early dropout and generating equality of opportunities for learning and development (UNICEF, 2020: 1). These systems are usually structured around two main components: 1) detection (with early warning systems as the principal tool) and 2) timely interventions and remedial strategies. Both elements are essential to achieve the objective of reducing educational exclusion.

Step 1 in this series, "Educational exclusion in LAC: how systems to protect educational pathways can help", introduced the main notions related to pathway protection systems: conceptualization, objectives, components, evidence and lessons learned. Step 2 now goes on to present the different approaches to designing early warning systems. They range from simple systems that trigger warnings by combining specific indicators through to more sophisticated systems using machine-learning methodologies. Step 2 also addresses the aspects that are key for designing an effective system and sets out essential guidelines for the use of data, the definition of indicators and the potential application of artificial intelligence (AI).

.

International experience has shown that prevention, intervention and guidance from an early stage are key for designing and implementing effective pathway protection systems. If education systems are to abandon their reactive approach in favor of a preventive approach, early warning systems are a vital tool in helping to identify students at risk of exclusion. This, in turn, permits better targeting in activating timely interventions in line with the needs and demands of students.

2 What are early warning systems?

Early warning systems are central to identifying students at risk of exclusion within the framework of a system to protect educational pathways. It is important to bear in mind that dropout from school is not a single event but, rather, the result of a long process of disconnection between the student and the system during which there are clear warning signs (Jimerson et al., 2000; Lamb et al., 2010; Román, 2013). Early reading of these signals permits detection of the risk before dropout occurs (detection component) and, therefore, facilitates the implementation of timely interventions and remedial strategies (intervention component) as a means of avoiding new exclusion processes.

In general, early warning systems apply a "red flags" logic (UNICEF, 2018). In this way, they help the authorities and educational institutions to identify students at risk based on performance patterns and potential contextual factors that may result in the interruption of an educational pathway.[1] In other words, they permit early identification of students with behavior or academic performance that puts them at risk of dropout (Frazelle and Nagel, 2015).

Early warning systems have proven to be a cost-effective tool and many countries have implemented them as part of their strategy to protect educational pathways and reduce educational exclusion (Gutiérrez and Vázquez del Mercado, 2021). One of their main advantages is building on structures that already exist in ministries, school districts or schools themselves. As a result, they do not imply large additional workloads and, at the same time, foster and facilitate coordination.

Different types of early warning systems can be used depending on conditions in each country or education system. Some are simple, using a combination of certain indicators to generate warnings, while others are more sophisticated and use machine-learning methodologies.[2]

1. Definition adopted by the US Department of Education through the National Survey on High School Strategies (HSS) Designed to Help At-Risk Students Graduate (2017).
2. Section 3.3 presents both types of early warning systems.

3 Early warning systems as a tool to support public policy decisions

Early warning systems are, in essence, a tool to support public policy decisions. They generate systematized information for use in decisions about actions or public policy interventions in the framework of a system to protect educational pathways.

Thanks to the rapid technological development of recent decades and the growing generation and collection of data from different sources, systems of this type now tend to use AI methodologies for data processing and/or the application of decision-making rules.[3] In line with the definition adopted by the fAIr LAC initiative,[4] AI-based systems offer a wide range of solutions, from those that use machine-learning algorithms and are able to receive data and learn by themselves without pre-programmed decisions to those based on expert knowledge and rules in which the decisions and criteria are pre-programmed by humans (Pombo et al., 2020).

In the specific case of the protection of educational pathways, systems for making or supporting decisions range from those based on indicators that activate a warning when a predefined threshold is reached, to those that use machine learning and can identify those students at higher risk of dropping out through prediction and reference to a set of data about the student and their environment.

In general, the design and implementation of effective technological tools to support public policy decisions calls for an understanding of the interrelation between the public policy cycle and the AI life cycle (Figure 1) (González, Ortiz and Sánchez, 2020). The public policy cycle is a simplified framework that seeks to represent the phases of policy development: identification of the problem, design or formulation of the intervention, implementation and evaluation. Figure 1 shows how this cycle can be expanded through the use of AI-based systems as public policy tools for making or supporting decisions.

3. The Organisation for Economic Co-operation and Development (OECD) describes artificial intelligence (AI) as "a machine-based system that can, for a given set of human-defined objectives, make predictions, recommendations, or decisions influencing real or virtual environments" (OECD, 2019), with the idea that AI systems are designed to operate with different levels of autonomy

4. fAIr LAC is an alliance between the public and private sectors, civil society and academia that seeks to influence both public policy and the entrepreneurship ecosystem in order to promote responsible and ethical use of AI (Pombo et al., 2020). See https://fairlac.iadb.org/.

FIGURE 1 ▪ **AI-ASSISTED PUBLIC POLICY CYCLE**

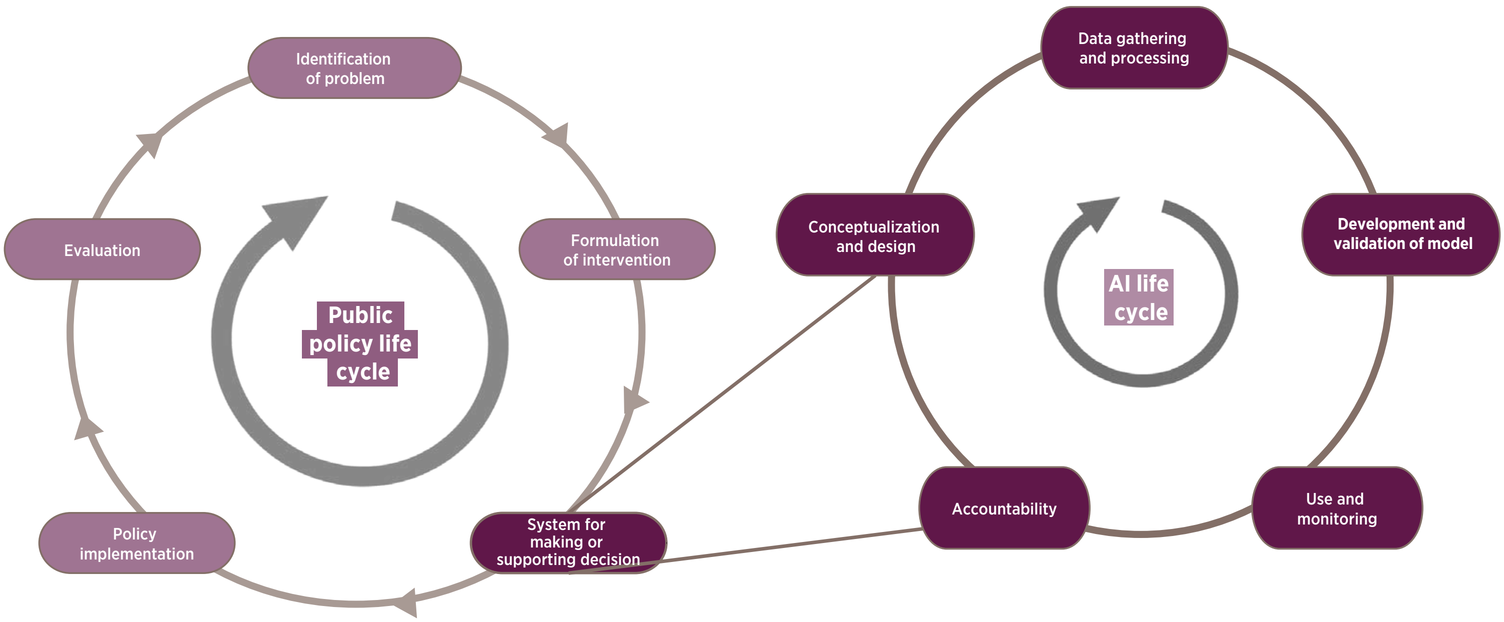

Source: Gonzalez, Ortiz and Sánchez (2020).

Early warning systems are, therefore, a tool for optimizing policy implementation and maximizing the impact of the action or intervention to be implemented (González, Ortiz and Sánchez, 2020).

3.1 Conceptualization and design

As shown in Figure 1, the first step in the public policy cycle is to identify the problem and define the objectives. In the specific case of systems for the protection of educational pathways, the problem may be, for example, a high percentage of students outside the education system. In turn, the public policy objective describes the ideal state to be achieved in terms of mitigating the problem. The objectives must be specific, measurable, achievable and realistic and have a defined timeline.

During the *conceptualization and design* stage, the person responsible for taking the public policy decisions must specify both the target population and the groups and attributes to be protected by the interventions (gender, ethnicity, religion, disability, income level, etc.). Expert knowledge must be used to define the attributes to be protected, identifying where situations of inequity have historically been observed in the social system or where they are most likely to arise. This informa-

tion is essential so that the technical team can subsequently analyze the available data, define its quality, coverage and representativeness and search for inequalities or undesirable states affecting population groups in those intersectionalities.[5]

Once the problem and the objectives have been identified, the next step in the public policy cycle is to formulate an intervention or, in other words, specify the actions to be taken. The policy measures deployed in a pathway protection system seek to take a comprehensive approach to the different causes and factors associated with student dropout. The measures can include personal support (face-to-face meetings with families, feedback among peers, academic tutoring programs), school scholarship programs, flexible formats and the adaptation of study plans. The possible interventions should be defined and specified before starting to develop an early warning system since, as discussed below, there are several ways in which they can condition how the AI-based model should work.

Step 3 of this series of publications will focus on timely interventions and remedial strategies to protect educational pathways.

5. This point is even more important for systems based on machine learning because algorithms learn from patterns in the data. Historical biases can, therefore, be transferred to the models, perpetuating the situation of discrimination. Analysis of these subgroups, during both data collection and the system's training, will help in taking decisions about the system's use or not.

TABLE 1 ▪ **EXAMPLES OF INTERVENTIONS FOR THE PROTECTION OF EDUCATIONAL TRAJECTORIES**

DIMENSION	PROGRAM
Measures to improve infrastructure	Repair of infrastructure
	Construction of canteens
Measures to train teachers	Training workshops for teachers
Measures to improve educational programs	Preparation of personalized study plans
	Flexible formats
Measures of guidance and personal support	In-person meetings between teachers, school directors and parents
	Educational agreements
	Preparation of a community strategy for support during holidays
	Spaces for stimulation of learning
	Feedback among peers
	Academic tutoring

Source: Compiled by authors based on CODICEN (2016) and US Department of Education (2016).

The characteristics of policy interventions under a pathway protection system can condition the ideal approach of the proposed early warning system in very different ways. Before conceptualizing and designing the system for making or supporting decisions, it is, therefore, of the greatest importance to have detailed information about the set of interventions and strategies to be implemented. The main aspects of the interventions that affect the design of the early warning system and must, therefore, be specified include the following:

Granularity: This refers to the level at which the intervention will be implemented. For example, if public policy interventions are envisaged only at the aggregate school or school district level, it may not be necessary to develop a student-level early warning system since the additional information may not improve the system's predictions.[6]

Intervention timeframe: This involves the temporal dimension in which the public policy intervention is expected to be applied and, therefore, also affects the time horizon of the early warning system. Illustrating this, Márquez et al. (2016) describe the implementation of an early warning system for Guatemala and Honduras, while Muñoz Estuardo (2019) analyzes the case of Uruguay. The authors focus on the transition from primary to secondary school when most of the reported school dropout occurs. In this case, the timeframe defines the system's development because, if the interventions are to be implemented in a specific period, the system must be designed to generate prior warnings. A system that generates warnings after the transition, however accurate, would be of no use.

Frequency of intervention: This refers to the frequency with which the intervention is expected to be implemented. It is necessary to define whether it will be monthly, half-yearly or each school year and, on this basis, determine the system's updating cycles.

Both the identification of the problem and the definition of the action or public policy intervention to be implemented are the basis for conceptualizing and delineating an early warning system. Its design will also depend on another key dimension: the availability of data. In this sense, the selection of information and access to relevant data are essential for early warning systems. Indeed, it is vital to have high-quality, accurate, relevant and representative data about the groups to be protected.

Early warning systems can also be used at multiple levels of aggregation of the education system (student, classroom, school, region, etc.) to address different problems in that particular sphere. With systems that generate information at the level of the individual, it is possible to provide targeted support to those students at risk of educational exclusion. When this is not possible or the data lacks sufficient granularity and quality, it will be necessary to develop systems to identify trends in the aggregate behavior of indicators by territorial area, focusing programs and resources on the regions and schools with the highest level of risk or, in other words, where they are most needed (Bruce et al., 2011; Josephson and Jayaram, 2018).

6. If the intervention will not be student-level, it is advisable to minimize the personal data used by the early warning system.

During the *conceptualization and design* stage, it is useful to consult the fAIr LAC initiative's Manual de Formulación de Proyectos and, in particular, its design and feasibility sheet. This is a tool for identifying the key aspects of an AI project in order to assess its viability, determine the suitability of AI as a solution and gather the information required for the project's design (Denis et al., 2021).

3.2 Data collection and handling

Sources of information, data handling and protection

As indicated above, the scope of an early warning system will depend on the availability and quality of the data. This is true both for models based on expert knowledge or indicators and for predictive models based on machine learning. What is needed is an information collection and processing strategy that improves the data's availability for decision-making and does not imply an overload of functions (mainly in schools).

In practice, most early warning systems are based on registers of attendance, grades and disciplinary incidents, including suspensions or expulsions (US Department of Education, 2016). Even though student's socioeconomic characteristics are relevant in educational exclusion (Adelman et al., 2017), observable signs - such as absenteeism, isolation from the group or academic performance - are usually a better predictor of risk of dropping out (Mac Iver and Mac Iver, 2009).

Education Management and Information Systems (SIGEDs) are the foundation for the construction of effective early warning systems since they contain the necessary information for their implementation.[7] SIGEDs provide a comprehensive overview of educational processes. They comprise six key management processes and two structural conditions (Figure 2). Early warning systems are one of the main tools for strategic management (process 6) and draw on information generated by daily management of the different levels of the education system in the other five key management processes.

7. This source of data contains most of the institutional and individual factors of the model of Rumberger (2012), presented in Step 1 of this series of publications.

FIGURE 2 ▪ MANAGEMENT PROCESSES AND STRUCTURAL CONDITIONS IN A SIGED

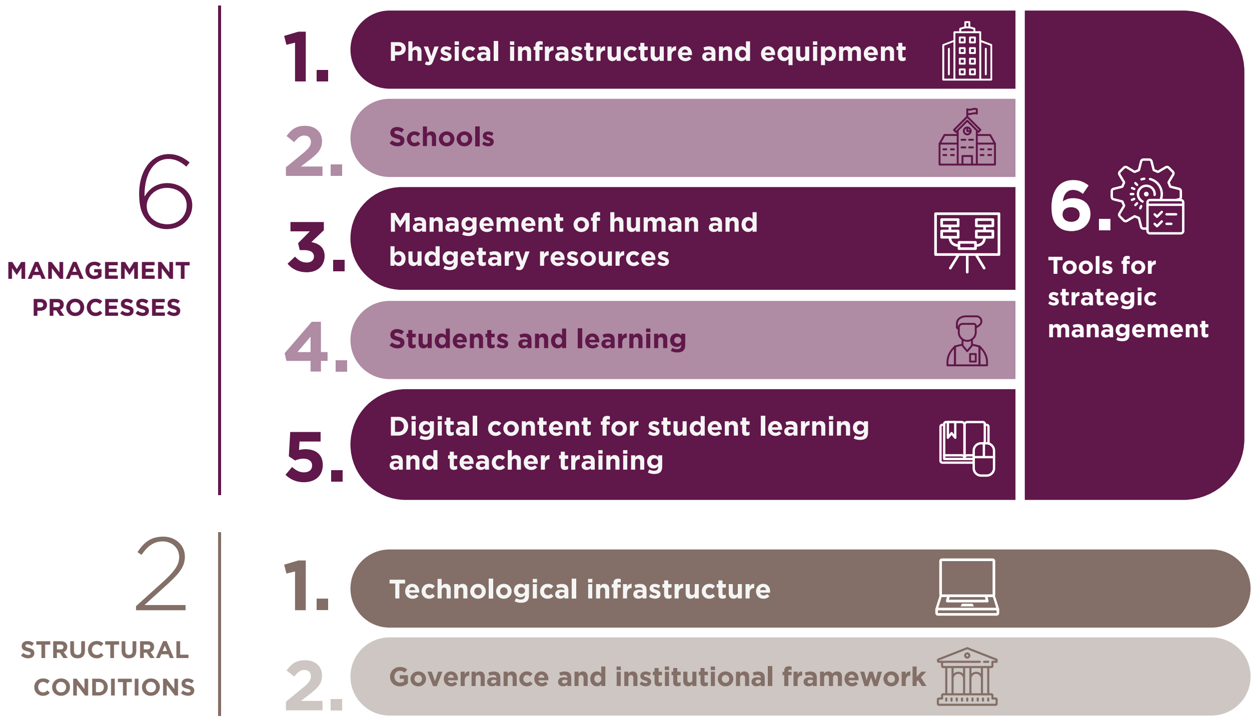

Source: Arias Ortiz et al. (2021).

For the implementation of early warning systems, the students and learning process acquires particular importance since it generates the information through which to monitor students' pathways. An established SIGED has a unique register of students in digital format, making it possible to follow them throughout their school life. It is a repository of personal, academic (digital student file), socio-educational and behavioral information, certificates, permission slips and data about attendance, exams, promotion and the repetition of years[8] (Arias Ortiz et al., 2019).

In addition, the availability of historical series of indicators makes for a more complete approach to risk pathways and, therefore, greater precision in identifying students at risk of exclusion. The

8. In addition to permitting the digitization and automation of management processes through which to monitor educational pathways, SIGEDs facilitate practices such as automatic enrollment, particularly in the case of cycle changes, that can help to reduce dropout. This has been the experience of Uruguay, which introduced digital enrollment for the transition from primary to secondary school and was able to comply with the family's preferences in 90% of cases, practically eliminating dropout in this transition (Arias Ortiz et al., 2019).

digitization of daily educational management processes facilitates access to students' historical information and its update (Rivera Pizzarro, 2020). Interoperability between information systems is particularly important when institutional organization and governance differ between educational cycles. For example, when the level of governance in primary schools differs from that in secondary education, it is vital to strive for coordination and interoperability between systems in order to have a historical record of students across their different educational cycles.

The information needed to implement effective early warning systems will differ depending on the type of model or approach used. In models based on machine learning, information from different levels, about educational pathways and socio-demographic variables and non-academic risk factors, is key for training the models and the accuracy of their predictions. By contrast, systems based on expert knowledge and indicators do not necessarily require this level of information and tend to use observable signs related to students' ties with their school. Table 2 shows the dimensions and variables used to design early warning systems. Information both at the individual level and about the characteristics of the student's school is relevant and complementary.

TABLE 2 ▪ VARIABLES USED FOR EARLY WARNING SYSTEMS THAT CAN BE OBTAINED FROM A SIGED

LEVEL	DIMENSIONS	VARIABLES
SCHOOL	Physical infrastructure and equipment	Geo-referenced location of school attended by the student
		School's infrastructure and utilities (electricity, water, gas, telephone, Internet, etc.)
		Number of students per classroom
	Educational institutions	Syllabus and curriculum
	Human, financial and budgetary resources	Teacher evaluation (internal and external)
		Information about human resources: degrees, training, work background
		Teachers' level of satisfaction
		Relation between teachers, content and students
STUDENT	Educational pathway and tie with the school	Analysis of repetition of years and over-age students
		Attendance (absence and reasons)/chronic absenteeism
		Academic performance
		Qualifications and results in standardized exams
		Scholarships - loans - social assistance
		Attitudes towards studying (discipline, conduct, commitment, etc.)
	Socio-demographic variables	Demographic attributes (sex, age)
		Disabilities
		Ethnic-racial descent
		Place of residence, migratory situation
		Diet and health
		Composition of family
		Domestic violence
		Special education
		Household economic situation
	Non-academic risk factors	Violence at school, bullying
		Adolescent parenthood, other family care responsibilities
		Problematic consumption of alcohol and/or other substances

Source: Compiled by authors.

Note: In the case of demographic attributes, although variables such as sex, age and disabilities do not affect the probability of dropout, it is important to be able to include these attributes when developing the system since it facilitates analysis of the results for bias and inequity.

The maturity of SIGEDs varies widely across the region. For example, not all countries have a nominal register of students that permits unique identification of each student, their pathway and the particular school(s) attended. There are also serious interoperability and fragmentation problems between the different subsystems used for the register, resulting in databases that do not communicate with each other, due to both conceptual and technological differences. This implies lost opportunities and costly inefficiencies (Arias Ortiz et al., 2021). The main dimensions and processes that a SIGED can cover are described in the publication Education Management and Information Systems (SIGEDs) in Latin America and the Caribbean: The Road to the Digital Transformation of Education Management, which also includes the results of a diagnostic study of 16 public education systems in LAC and offers policy recommendations for increasing the efficiency of educational management (Arias Ortiz et al., 2021).

In LAC, limits on data availability mean that it is not always possible to establish early warning systems that permit analysis at the student level. This can be for several reasons: because, for example, it takes time to gather data for the school period into a single database or because digitized records do not exist at that level of disaggregation. In these cases, an alternative is to develop systems at the school level. As shown in Table 2, some variables of dimensions within schools, aggregated at the school level, can be used as input for an early warning system. A simple model that permits identification of the schools attended by high-risk students will help the authorities plan aggregate policies that would not otherwise be available.

In recent years and, particularly, in 2020 in the context of the COVID-19 pandemic, education systems have expanded their use of digital tools for learning through new platforms. This digitization of the learning environment has increased the availability of information - in some cases, this information is already integrated into the SIGED - that can also be incorporated into early warning systems and facilitates the application of analytical and machine-learning models for research and learning evaluation (Cohen, 2017).

Some early warning systems also use information from surveys and questionnaires to complement other data (Sansone, 2019). This can be the case, for example, of indicators about daily hours of study, group study, study habits, level of motivation, personality type, resources for studying, number of siblings, position in the family as the eldest, middle or youngest child, encouragement to study by parents, distance from school, interest in subjects and their level of difficulty (Márquez et al., 2016).

It can also be useful to incorporate data from other government programs and agencies since different ministries compile a number of educational and social indicators as part of their regular activities. Key information of this type includes participation in social programs and, particularly, access to non-contributory transfers as well as data about aspects such as health, finance and social security situation. Coordination and joint work with other areas of the state is also important in

order to avoid the duplication of efforts and increase efficiency (Rivero Pizarro, 2020). In this sense, the use of the identity documents issued by civil registries facilitates the exchange of information with other programs or bodies that record this data about its beneficiaries (Arias Ortiz et al., 2019).

One of the main challenges as regards data handling is the governance and protection of personal data. Some initiatives have used information about students from outside school, such as thematic profiles of their interests and integration and connections in social networks (Berens et al., 2018). This information can be very useful, but it is necessary to ensure the existence of the legal powers and consent required for its use, guaranteeing students' right to privacy, minimizing the personal information that will be used for the early warning system, and confirming the presence of specific governance and security structures and mechanisms.[9] The rapid transition to digital education that occurred during the pandemic has revealed gaps in the knowledge of personnel about the practices that must be implemented so as not to compromise the security of students' data. According to a survey carried out by the Inter-American Development Bank (IDB) and C Minds' Eon Resilience Laboratory in seven LAC countries,[10] 72.3% of teachers do not have training (or it is insufficient) on data privacy issues and the responsible use of digital platforms and tools (Del Pozo, del Campo Alcocer and Róo Rubí, forthcoming). Moreover, individuals do not know about or consent to the use of their personal data for purposes not directly related to the educational institution by which it was collected (Del Pozo, del Campo Alcocer and Róo Rubí, forthcoming).

Biases, undesirable situations and incomplete information about the population

Once the available data has been identified and the ethical, legal and governance considerations for its use validated, the next step consists of an exploratory analysis of the data's quality, relevance and coverage. For this purpose, all the information provided during the *conceptualization and design* stage must be taken into account. The development of AI is affected not only by the understanding of the problem and the system's objectives, but also by the population and attributes to be protected as defined during the public policy cycle.

Given that every social setting has inequities, the person responsible for public policies must be in constant communication with the technical team in charge of developing the early warning system. In this phase, exploratory analysis, based on the historical register of students, is required to cross-reference and count the variables of interest, separated by the groups and attributes to be protected (defined during the *conceptualization and design* stage). This will show whether coverage

9. It is advisable to complete the data profile (González, Ortiz and Sánchez, 2020).
10. Brazil, Colombia, Costa Rica, Mexico, Peru, Panama and Uruguay.

is sufficient to guarantee representativeness of the population subgroups of interest or if there is an over or under-representation of some group in particular that could affect the model's results. By way of illustration, Table 3 shows the results of counting observations of students from the historical register who passed or failed the year in urban and rural schools in Uruguay (ANEP, 2021). There are more observations for urban than rural areas and the proportion of students passing the year differs by geographical area. In addition, the results reveal a lack of information for the second year in rural areas. These differences and the absence of certain values can have numerous causes, ranging from systems' lack of connectivity in rural areas to difficulties in data collection and the diversity of educational approaches and programs. The differences found must be taken into account so as not to generalize a model that, in the long run, could harm minority groups within the target population - in this case, students in rural areas.

TABLE 3 ▪ **EXAMPLE OF GROUPING OF VARIABLE OF INTEREST BY PROTECTED GROUP AND NUMBER OF SECONDARY STUDENTS IN URBAN AND RURAL AREAS, URUGUAY**

PLAN	YEAR	URBAN AREA			RURAL AREA		
		Passed	Failed	TOTAL	Passed	Failed	TOTAL
CES	1	88,987	47,425	136,712	5,364	1,460	6,824
	2	91,090	11,242	102,332	0	0	0

Source: Consultancy documents prepared by ANEP (2021).

Although bias and a lack of representativeness are found in decision-support systems of all types, they are particularly a concern in those based on machine learning because algorithms learn patterns from historical information. Therefore, the system may learn and replicate these patterns if the training database lacks information about a protected group or attribute or has historical biases. For example, if students were surveyed about their interest in different subjects, biases could arise as a result of structures of social inequity that it is not the intention to perpetuate. This is the case of the gender gap in the Science, Technology, Engineering and Mathematics (STEM) sector. Although the gap has been narrowing gradually, there continue to be more men than women (WEF, 2016). The technical team must be very clear as to which attributes and populations must be subject to comparative analysis to control their coverage and possible historical biases in collecting the data and align the model with the public policy objectives.

3.3 Development and validation of model

Early warning models based on expert knowledge or indicators

The first attempts to create early warning systems focused on models based on expert knowledge and indicators. Models of this type function by generating simple indicators of risk, such as low grades and low school attendance. They then define static thresholds of anomalies and establish rules for their aggregation (Bowers et al., 2013). When well implemented, these models have proven effective in identifying students at risk of exclusion.

As its name suggests, this type of model draws on accumulated evidence-based knowledge about the main risk factors associated with educational exclusion. The selection of indicators is crucial for the precise identification of those students who are really at risk (Bruce et al., 2011). It must be based on the results of research that associates the phenomenon of educational exclusion with specific factors. In order to achieve greater precision, it must also take into account the particular characteristics of exclusion in the context in which the system will be used (Balfanz, 2008). Some indicators can be considered "basic" and should be part of any early warning system, but the relevance of others may depend on the characteristics of the country's population and/or education system. One of the most common ways of selecting indicators is to analyze historical records - that is, the information available about students who have dropped out in the past - as a guide to the key indicators for the preventive detection of situations of risk in future generations (UNICEF, 2018).

The attendance, behavior and course performance (ABC) model summarizes the three basic components that the literature has identified as the most powerful predictors of dropout (Bruce et al., 2011). It has been implemented by the National High School Center Early Warning System, which proposes reference thresholds for each dimension: student attendance (absences of 10% or more), course performance (one or more courses failed and a grade average of less than or equal to 2 on a 4-point scale), and behavior (referrals, in-school or out-of-school suspension or behavior grades whose thresholds must be defined locally) (Frazelle and Nagel, 2015). When any of these variables is outside the expected range, the model generates a warning.

John Hopkins University also offers suggestions for defining thresholds to detect students at risk of exclusion. It uses three categories: "off track", "sliding" and "on track to graduation" (Table 4) (Frazelle and Nagel, 2015). In line with the points raised in previous sections, it is also necessary to consider the temporal dimension since, if the system is expected to generate warnings throughout the school year, measurement parameters and thresholds for different periods will be required. It is also recommended that the teams be able to analyze and evaluate cut-off points based on historical data to ensure that the reference thresholds make sense in the specific context of their education system.

TABLE 4 ▪ **EARLY WARNING SYSTEM OF INDICATORS AND THRESHOLDS SUGGESTED BY JOHN HOPKINS UNIVERSITY**

Threshold	Attendance (days missed)		Behavior				Academics	
			Office referrals		Suspensions			
	Quarter	Full year	Quarter	Full year	Quarter	Full year	Math and English language arts [a]	Core courses [b]
Off track	9 days	36 days	2	6	1	2	Report card grade of F	Report card grade of F
Sliding	5-8 days	19-35 days	1	3-5	0	0-1	Report card grade of D	Report card grade of D
On track to graduation	4 days or fewer	18 días	0	0-2	0	0-1	Report card grade of A, B or C	Report card grade of A, B or C

a: Middle school.
b: High school.
Source: John Hopkins University (2012: 10).

A survey by the National High School Center Early Warning System found that, in 2015, 52% of schools in the United States were using early warning systems based on expert knowledge. In most cases, the systems used data on school attendance (92%), course grades (91%), absenteeism and/or chronic absenteeism (82%) and disciplinary incidents, including suspensions or expulsions (79%) or, in other words, the basic predictors of the ABC model. The survey also found that some schools were using complementary data and that schools with the lowest graduation rate and the highest poverty rate tended to include the largest quantity of data additional to that of the ABC model. The main complementary indicators used were involvement with the criminal justice system, lack of housing, adolescent pregnancy/parenthood and over-age students. The National High School Center Early Warning System underscores this finding as positive and considers that, while ABC predictors are central to almost all early warning systems, schools have also exercised discretion in collecting additional data tailored to their contexts and populations (US Department of Education, 2016). When early warning systems are developed at the national or regional level, they must take into account the regional and local dimensions that are relevant as regards educational exclusion. In other words, the design of early warning systems by the central level must ensure the inclusion of local indicators in order to generate systems that are more efficient and sensitive in identifying students at risk.

In LAC, in particular, the most visible additional indicators include adolescent pregnancy and over-age students (Adelman and Székely, 2017). Being over-age - or, in other words, older than the official school-age for the corresponding year - due either to a late start of schooling or the repetition of courses, is one of the main predictors of educational exclusion in the region (Manacorda, 2006; Adelman and Székely, 2017). It is also important to bear in mind that the relative importance of the different indicators changes over a student's life cycle: for example, indicators of the level of violence or crime in the community tend to be more relevant during adolescence (Adelman and Székely, 2017).

In early warning systems based on expert knowledge, the number of indicators selected is key. In their case, as distinct from machine-learning systems, "more" is not always "better". Experts warn that selecting too many indicators may result in an overwhelming quantity of data to analyze and interpret, which is not advisable (University of Chicago Consortium on Chicago School Research, 2014). A good strategy is to start by establishing a basic set of indicators - for example, the ABC model - and then incorporate other complementary indicators and see whether this significantly increases the number of students identified as being at risk of dropout (Frazelle and Nagel, 2015). The sensitivity of the selected indicators is another factor to bear in mind. It is important to avoid systems that, rather than helping target the measures to be taken, identify too large a proportion of students as being at risk and, therefore, hamper timely identification of those who may potentially drop out (UNICEF, 2018). To guard against this, the use of a combination of indicators and their weighting are recommended.

Colombia, Brazil and Uruguay are among the LAC countries that have developed early warning systems based on expert knowledge or indicators. They have selected a set of indicators and rules to detect students at risk of dropout. For example, Colombia uses the Index of Risk of the School Educational Pathway (IRTE), a dashboard with 16 indicators grouped into four sub-indices: student family responsibilities, school climate, socioeconomic context and past educational pathway. Similarly, as part of its System of Protection of Educational Pathways (SPTE), Uruguay has launched an early warning system of repeated absences for use by schools and, particularly, the teams responsible for monitoring and supporting educational pathways (Muñoz Stuardo, 2020). The SPTE was used in pilot form in 2018 and, in May 2019, was introduced for all the Uruguayan educational system (Muñoz Stuardo, 2020). In Brazil, the State of Paraná has been implementing the School Dropout Combat Program since 2013. It is based on a single indicator of school attendance and, when a student has had five consecutive absences or seven non-consecutive absences, the system notifies the Social Protection Network for Children and Adolescents - comprising schools, social assistance centers, community councils and other actors - so that it implements measures and strategies appropriate to each situation. Step 4 of this series of publications describes various experiences of early warning systems, with different levels of automation and data use, that have been implemented in LAC.

Early warning systems based on machine learning

In recent years, with the increasing availability of data and analytical processing methodologies, knowledge-based early warning models have begun to be replaced by machine-learning algorithms. They mainly use approaches of two types: i) unsupervised approaches for the application of group interventions through grouping or clustering algorithms, and ii) supervised approaches for the prediction of variables of interest (for example, school dropout), based on algorithms that range from random forests (Aguiar et al., 2015) to neural networks (Sattar et al., 2016).

Supervised learning algorithms are models trained using historical databases about the behavior of previous generations of students who have already passed through the educational cycle and whose pathway is, therefore, known. These models serve to predict different characteristics or risk factors, that is, target variables or variables of interest. The algorithms learn from the relationship between descriptive characteristics (or predictors) and the target variable to accurately predict the outcome of future observations or arrive at a better understanding of the relationship between the predictors and the target variable (Chung and Lee, 2019).

As described in the *conceptualization and design* section, the target variable must be defined as a function of the type of action or public policy intervention to be carried out. In the literature, the main target variable used in the case of educational exclusion is "school dropout". It is constructed as a dichotomous variable that takes the value of 1 for students with at least one episode of dropout during their school life and 0 otherwise (Sansone, 2019). Based on this, the models predict the probability of dropout for new students over a given time horizon. However, early warning systems can also be trained to identify other educational risks, such as the risk of not graduating on time (Aguiar et al., 2015), poor academic performance or absences. For example, a school might want to identify students who are good candidates for tutoring to improve their grades. In digital contexts and, particularly, in the case of models applied in online courses or massive open online courses (MOOCs), the definition of the target variable can consider the frequency of visits and the duration of sessions in the platform (Prenkaj et al., 2020).

Unsupervised learning algorithms, unlike supervised algorithms, do not aim to predict a result during training; instead, they search for structural patterns in the information in order to create associations between observations and discover similar groups or individuals in the training information (clustering) (Pombo et al., 2020). Indeed, clustering algorithms, such as k-means and hierarchical clustering, have been used to create academic performance profiles for students with similar characteristics (Iam-On and Boongoen, 2017). In this way, student groups or profiles can be devised for group interventions through which to reduce the cost of the strategies to be implemented.

It is important to note that an early warning system does not necessarily have to be based on a single model, but can be designed as a set of models that includes both supervised and unsupervised algorithms. For example, Sansone (2019) uses a supervised model to identify students at risk of educational exclusion and then applies unsupervised algorithms to classify them according to their observable characteristics. With this differentiation, it is also possible to assess how a policy's impacts differ across the different groups. Similarly, Aguiar et al. (2015) combine two types of supervised algorithms, performing first a binary estimation of the probability of dropout and then applying survival analysis, using a Cox-type regression, to create an indicator of "urgency". Survival analysis methodologies are supervised algorithms that focus on predicting the expected time until an event of interest - in this case, school dropout - occurs. With systems of this type, it is possible to model a student's pathway, create a dynamic indicator of dropout risk over time and know when the risk is greatest (Figure 3). Based on this, priority can be given to measures or interventions for students who require more immediate attention. In other words, if two students are classified as at high risk of dropping out of school, priority can be given to intervention for the student whom the model predicts will drop out first.

FIGURE 3 ▪ **SURVIVAL MODEL APPLIED TO SCHOOL DROPOUT**

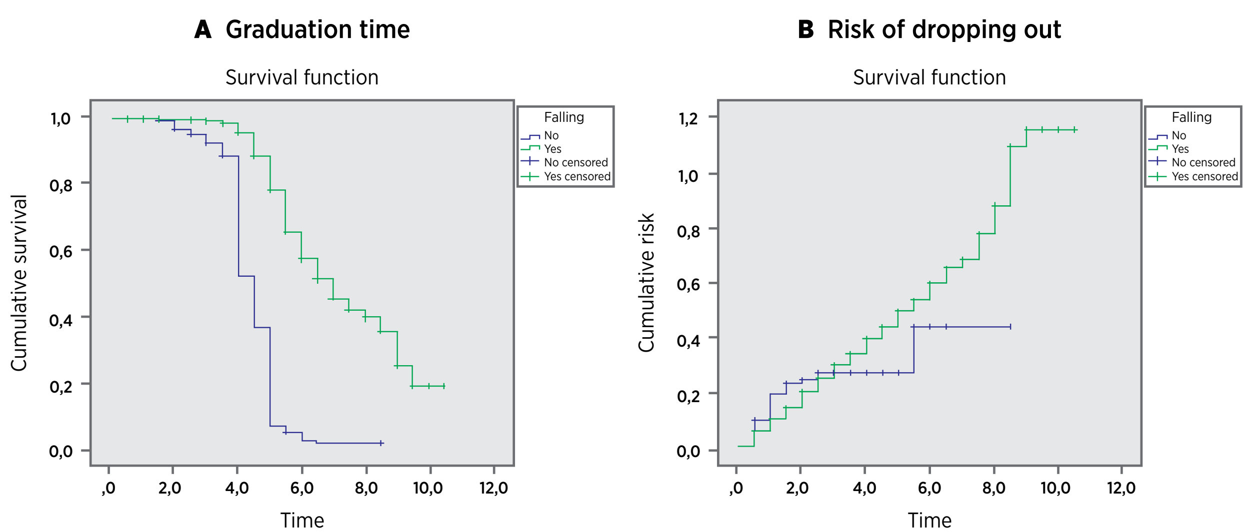

Source: Costa et al. (2018).

As discussed above, the prediction horizon and the frequency of training must be in line with the design and timing of the intervention. Logically, the sooner a reliable prediction can be made, the more time educational teams will have to apply public policy strategies and interventions to reduce

the risk of exclusion. However, the earlier a prediction, the smaller the amount of information that will be available about the student's educational pathway and the greater the system's likely margin of error (Márquez et al., 2016). Indeed, one of the keys to implementing systems of this type lies in striking a balance between obtaining warnings that are early enough to permit intervention and, on the other hand, allowing the time required to gather sufficient information for a high-quality prediction.

It is also possible to design systems that include more than one model over the course of an educational pathway, updating the risk identified per student as more information becomes available and adjusting the interventions accordingly. The definition of cut-offs and intervention horizons will depend mainly on the availability of data and the capacity of the institutions to adjust their public policy measures.

FIGURE 4 ▪ **INTERVENTION TIMEFRAME**

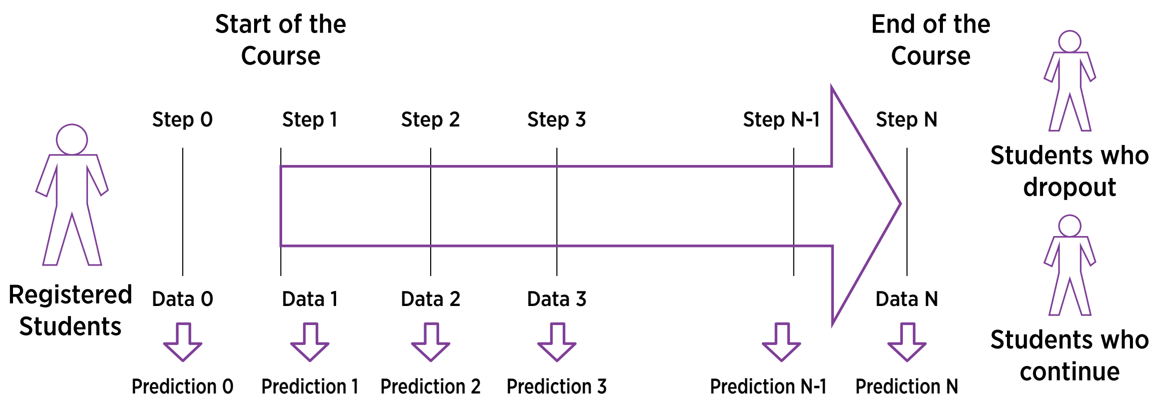

Source: Márquez et al. (2016).

Unbalanced data

One of the main problems in training a machine-learning early warning system for school dropout is the training database's inherent imbalance between those who drop out and those who do not. There tends to be much less historical information about students who have abandoned the education system than about those who have remained in it for the simple reason that the former are fewer in number. In supervised machine-learning systems, training about the minority class (that

is, the one with the fewest number of cases) is more difficult and this could hamper construction of the predictive model. To mitigate phenomena of this type, under and over-sampling techniques such as the Synthetic Minority Over-Sampling Technique (SMOTE) can be used. This helps to balance the classes by increasing the relative importance of the minority group when training the model (Figure 5).[11]

FIGURE 5 ▪ **DISTRIBUTION OF FIRST-YEAR SECONDARY STUDENTS BY GRADUATION, ANEP PROJECT**

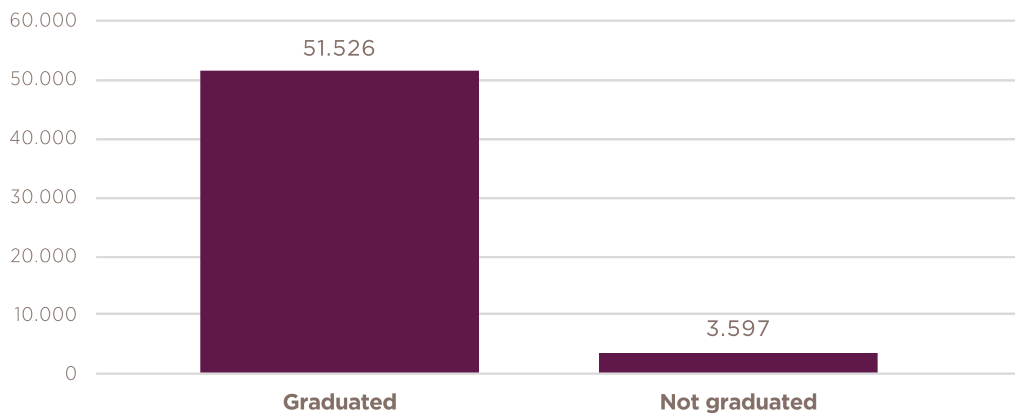

Source: Consultancy documents prepared by ANEP (2021).

Evaluation and validation of models

A key point that must be borne in mind in the case of supervised models trained using unbalanced databases is that analysis of their error may be misleading.[12] This is because metrics such as accuracy - the proportion of observations that the model predicted correctly - analyze the performance of the classification for both categories - those who drop out and those who do not - which can

11. All under or over-sampling techniques must be applied after the database's division into training and evaluation. To do so previously would cause a leakage of information from the training to the validation table (González, Ortiz and Sánchez, 2020).
12. The fAIr LAC initiative's publication, *IA Responsable: Manual Técnico: Ciclo de Vida de la Inteligencia Artificial,* (in Spanish only) describes the main errors and technical risks that should be taken into account in each of the phases of development of a machine-learning model and indicates the measures that can be taken to mitigate them (González, Ortiz and Sánchez, 2020).

lead to misinterpretation and misplaced confidence in the system's performance. For example, if an education system has a 5% dropout rate - and, therefore, a continuity rate of 95% - a trained model that predicts 100% continuity and 0% exclusion will have an accuracy of 95%. In this extreme case, the accuracy metric can lead to false confidence in the system's performance since, although it has an accuracy of 95%, its usefulness for predicting students at risk of dropout is nil. It is, therefore, always necessary to break the error down into false positives (mistakenly identifying a student as at risk of dropping out) and false negatives (mistakenly identifying a student as not being at risk of dropping out) and analyze metrics such as precision, sensitivity and specificity as a whole (Table 5). It is important to clarify that supervised models are evaluated and validated using historical data about students who have already passed through the educational cycle, a process in which these metrics are obtained.

TABLE 5 ▪ **CLASSIFICATION CONFUSION MATRIX**

		REAL	
		POSITIVES	**NEGATIVES**
PREDICTED	**POSITIVE**	True positive (TP)	False positive (FP)
	NEGATIVE	False negative (FN)	True negative (TN)

$$\text{Accuracy} = \frac{TP + TN}{TP + TN + FP + FN} \qquad \text{Precision} = \frac{TP}{TP + FP} \qquad \text{Sensitivity} = \frac{TP}{TP + FN}$$

Source: González, Ortiz and Sánchez (2020).

The choice of an error metric to select the classification probability cutoff points, as well as the choice between different models involves a public policy decision that must be considered as such. In some contexts, given their application space, the person responsible for public policies may want to minimize false positives while, in other cases, it may be more important to minimize false negatives. This decision is a tradeoff between the two objectives since it is not always possible to optimize one without sacrificing the other. For example, a school or education system that wants to cover the largest possible number of students at risk of exclusion will prefer the most sensitive model (minimizing false negatives) even though it increases the cost of the intervention by incorrectly classifying some students as having a high probability of dropout. On the other hand, a school or system that has few resources may prefer to maximize precision (minimizing false positives), thereby reducing the number of students wrongly classified as at risk, even though this may mean failing to identify some students who are actually at risk of dropping out.[13]

As already described for the previous stages, incomplete information, a lack of representativeness and assumptions about its collection as well as errors in the model's development can cause biased results that lead to decisions that are undesirable, unfair or discriminatory for different subgroups (González, Ortiz and Sánchez, 2020). Therefore, the chosen error metrics must not only be evaluated for the population in general, but also differentiated by groups and protected populations.[14] Figure 6 shows an example of analysis of the confusion matrix by groups, indicating in percentages the composition of the false positives and false negatives in the model's results disaggregated by sex. In this case, the percentage of false negatives and false positives is very similar for both of the groups analyzed, indicating that there are not differences in the error metric by sex. This type of analysis should be carried out for all the intersectionalities defined by the public policy decision-maker during the *conceptualization and design* stage.

13. In general, for this type of system, use of the precision recall curve is recommended to evaluate trained classifiers with unbalanced databases. For more details about error metrics and the risks of evaluating them out of context, see González, Ortiz and Sánchez (2020).
14. Completion of the model profile is recommended (González, Ortiz and Sánchez, 2020).

FIGURE 6 ▪ ANALYSIS OF BIAS BY SEX, ANEP PILOT PROGRAM

MALES

FALSE POSITIVES (%)	FALSE NEGATIVES (%)	ACCURACY (%)
13,8	**1,7**	**84,4**

Confusion Matrix

		PREDICTED		
		POSITIVE	NEGATIVE	**TOTAL**
ACTUAL	POSITIVE	84.1% (243)	1.7% (5)	85.8% (248)
	NEGATIVE	13.8% (40)	0.3% (1)	14.2% (41)
	TOTAL	97.9% (283)	2.1% (6)	

FEMALES

FALSE POSITIVES (%)	FALSE NEGATIVES (%)	ACCURACY (%)
13,7	**1,4**	**84,8**

Confusion Matrix

		PREDICTED		
		POSITIVO	NEGATIVO	**TOTAL**
ACTUAL	POSITIVE	84.4% (178)	1.4% (3)	85.8% (181)
	NEGATIVE	13.7% (29)	0.5% (1)	14.2% (30)
	TOTAL	98.1% (207)	1.9% (4)	

Source: Consultancy documents prepared by ANEP (2021).

There are also other risks that must be borne in mind when developing a system based on machine learning. They include risks related to validation errors, information leaks and degradation of performance.[15]

15. It is recommended that the technical team use the publication *"IA Responsable: Manual Técnico: Ciclo de Vida de la Inteligencia Artificial"* (in Spanish only) (González, Ortiz and Sánchez, 2020).

3.4 Use and monitoring

Once the most appropriate model has been developed, validated and selected, the implementation stage begins and the AI system will start to provide results for decision-making - in this particular case, early warnings about the students most at risk of dropping out. In this phase, it is essential to identify two aspects: using the results to activate the public policy intervention and monitoring the model's performance to guarantee that the early warnings and, therefore, the resulting interventions remain relevant over time.

Model appropriation

A key element for the effective functioning of early warning systems is their presentation to teachers, school directors and the authorities. This is important in order to guarantee correct interaction with the system's results when a student is at risk of dropping out. If the end-users are to appropriate the system, the visualization tools must be effective and user-friendly (Frazelle and Nagel, 2015). The challenge here is to position the information provided by the system as a relevant input for the daily work of teachers, school directors and management levels. This implies that communication and active listening about the daily needs of the tool's end-users are fundamental in the appropriation stage.

It is also important to consider the different levels at which the information may be relevant. Therefore, it is useful to have different types of users, for either the report or the data dashboard, to improve use of the information at each level and foster the tool's appropriation throughout the education system. For example, two types of user can be created: one, designed mainly for teachers, with access to information disaggregated by student, and the other, designed for school directors and municipal, regional and national authorities, with access to consolidated information at the school, regional and country level, thereby promoting use of the data for efficient and opportune decision-making.

When the entity responsible for education policy opts to outsource an early warning system's development to a third party (private company, university or another state body), it is essential to guarantee a fluid and transparent transfer and appropriation process that engages the different actors of the education sector. The technical team in charge of the system's development must train users so they are familiar with key aspects of the model, its objective, the logic under which it produces results, how it works and how to interact with it. In line with this, preparation of a user manual is suggested to facilitate appropriation of the new tool.

Uruguay's recent experience serves as an example. In 2020 and the first half of 2021, its National Public Education Administration (ANEP) worked with the country's University of the Republic and the IDB to create a machine-learning early warning system through the development of a model to predict dropout in first and second-year secondary education. The technical team built nine models and, once they had been validated, transferred the code for pre-processing the data and generating the models to ANEP, along with the application programming interfaces (API)[16] for the models' execution, so they could be used as part of its pathway protection system. The transfer and appropriation process also included workshops with different officials related to one or more of the stages of definition of the early warning system implemented.

Monitoring and evaluation

Like most systems for supporting and making decisions, early warning systems must include a follow-up and monitoring process to ensure that they continue to provide relevant results and do not deviate from their intended purpose. It is, therefore, important to define metrics and performance indicators through which to detect the potential deviations or unexpected changes in the system's operation that may occur over time. It is particularly important to establish clear metrics and performance thresholds for systems based on machine learning. Because algorithms continue to learn after they have been developed, they may deviate from the original concept and produce different results, implying that the system must be retrained.

The evaluation of an early warning system consists specifically in measuring the precision and quality of its identification of students at risk of educational exclusion. For this, it is essential to monitor the metrics established in the modeling and the cutoff thresholds defined for systems based on expert knowledge and to review data capture assumptions to ensure they have not changed. In other words, if the coverage, representativeness or quality of the data has changed for any reason, this will change how the system works, making its monitoring, evaluation and update essential.

It is also important to bear in mind that a high-precision system that permits adequate identification of students at risk of exclusion will only be useful if the interventions it triggers are effective in achieving their objective of reducing exclusion. In other words, the warning system's greater or lesser success in correctly identifying students at risk is not synonymous with the success of the resulting policy intervention. Therefore, to evaluate a pathway protection system's impact, it is also necessary to evaluate the impact of the policy intervention and, for example, the extent to which the tutoring and support programs implemented have reduced educational exclusion. Different impact evaluation techniques are available, including natural experiments, randomized controlled

16. APIs are a set of definitions and protocols used to develop and integrate applications' software.

experiments and quasi-experiments (Shadish, Cook and Campbell, 2002). A randomized controlled experiment for mathematics was used to evaluate the effectiveness of a pilot dropout prevention project using early warnings in four Asian countries: Cambodia, India, Tajikistan and Timor-Leste. Each country used personalized predictors to identify at-risk students, accompanied by differentiated interventions/activities in response to the warning, and each conducted an independent impact evaluation. Although the interventions had the same objective - to prevent dropout - they differed, ranging from after-school tutoring programs in Tajikistan to computer literacy programs in Cambodia and structured recreation programs in India and Timor-Leste. The impact evaluation showed that: i) in three of the four countries, attendance measured as the percentage of days of school attendance among students in the intervention group improved by around 1.4 percentage points on average; and ii) in Cambodia, dropout rates fell from 41.1% in the control group to 38.7% and 39.3% in each treatment group (USAID, 2015).

Thanks to their high level of validity, randomized controlled experiments are seen as one of the best ways to measure impact. However, the feasibility of carrying out experiments of this type depends on the availability of data for each project since solid counterfactuals may be difficult to obtain. Step 3 of this series of publications, which looks at timely interventions and remedial strategies in greater depth, discusses key aspects of the evaluation of the measures implemented.

3.5 Accountability

Throughout the implementation of public policy interventions, information must be made available to citizens or the target population in order to guarantee transparency. When education systems implement early warning models, they must be capable of explaining the decision-making processes that led them to allocate resources and interventions to some students - those at risk of dropping out - and not others. This is particularly important in the case of machine-learning systems.

When an early warning system fails to identify a student who is, in fact, at risk of dropping out (a false negative) or when a student not at risk is included in a policy intervention (a false positive), this implies both inefficiency in the allocation of resources and, in the case of false negatives, a potential increase in the risk of interruption of these students' educational pathways. In this context, in order to provide the best interventions and support, teachers need to be able not only to identify the students at risk, but also to understand why a student is at risk (Bowers, 2021).

In the simplest systems, based on a set of indicators, this process is far more direct because their previously defined thresholds, which are known to the educational community, favor accountability and the interpretability of results. By contrast, machine-learning systems tend to perform better (fewer allocation errors), but their design's complexity makes interpretation of the data far more difficult (Sansone, 2019). This tradeoff between performance and interpretability becomes important in adopting and using sophisticated early warning systems in school districts (Knowles, 2015).

This represents a challenge, especially when seeking to explain a prediction at the individual student level. In many cases, there is a legal need to provide individual explanations of certain decisions, particularly when they involve public resources. Because machine-learning systems tend to aggregate different types of student data across multiple areas (behavior, attendance, academic performance, etc.) and apply very advanced statistical methods to reach a risk prediction, teachers do not have explicit information about how the model arrived at the estimate (Coleman, 2021). This challenge is common in decision-making systems based on machine learning. However, there are statistical techniques[17] that analyze how the model works for individual predictions and facilitate the interpretability of early warnings without sacrificing precision, making them conducive to the design of more personalized student interventions. (Coleman, 2021).

17. Further information can be found in *"IA Responsable: Manual Técnico: Ciclo de Vida de la Inteligencia Artificial"* (in Spanish only) (González, Ortiz and Sánchez, 2020) and the references mentioned in this manual.

4 Next steps

This Step 2 publication on the construction of systems to protect educational pathways sets out the different approaches to designing early warning systems and examines key aspects for their effectiveness. In addition, it discusses critical characteristics of the systems as regards the use of data, the definition of indicators and the potential application of AI.

The next two publications in this series continue to look in depth at different topics related to the design and implementation of systems to protect educational pathways. Step 3 focuses on timely interventions, examining the evidence about effective dropout-reduction interventions and good practices for timely interventions of remediation and support. Finally, Step 4 brings together a selection of experiences from pathway protection systems in LAC and summarizes what has been achieved in terms of early warning systems and timely interventions and remedial strategies.

References

Adelman, M., F. Haimovich, A. Ham and E. Vázquez. 2017. Predicting School Dropout with Administrative Data: New Evidence from Guatemala and Honduras. Policy Research Working Paper No. 8142. Washington, D.C.: World Bank.

Adelman, M. A. and M. Székely. 2017. An Overview of School Dropout in Central America: Unresolved Issues and New Challenges for Education Progress. E*uropean Journal of Educational Research,* 6(3): 235-259.

Aguiar, E., H. Lakkaraju, N. Bhanpuri, D. Miller, B. Yuhas, K. Addison and R. Ghani. 2015. Who, When, and Why: A Machine Learning Approach to Prioritizing Students at Risk of not Graduating High School on Time. ACM International Conference Proceeding Series, Vol. 16-20-March-2015: 93-102.

ANEP (Administración Nacional de Educación Pública). 2021. Consultoría sobre estudiantes aprobados y no aprobados en escuelas urbanas y rurales de Uruguay. Montevideo: ANEP.

Arias Ortiz, E., J. Eusebio, M. Pérez Alfaro, M. Vásquez and P. Zoido. 2019. From Paper to the Cloud: Guiding the Digital Transformation of Education Management and Information Systems (SIGEDs). Washington, D.C.: IDB. Available at https://doi.org/10.18235/0001749.

------. 2021. *Education Management and Information Systems (SIGEDs) in Latin America and the Caribbean: the road to the digital transformation of education management.* Washington, D.C.: IDB. Available at https://publications.iadb.org/es/los-sistemas-de-informacion-y-gestion-educativa-siged-de-america-latina-y-el-caribe-la-ruta-hacia.

Balfanz, R. 2008. Three steps to building an early warning and intervention system for potential dropouts. PowerPoint slides. Baltimore, MD: Center for Social Organization of Schools, John Hopkins University. Available at jhu. edu/CSOS/graduation-gap/gradgap. html.

Berens, J., S. Oster, K. Schneider and J. Burghoff. 2018. Early Detection of Students at Risk: Predicting Student Dropouts Using Administrative Student Data and Machine Learning Methods. Schumpeter Discussion Papers, No. 2018-006.

Bowers, A. J., R. Sprott and S. Taff. 2013. Do We Know Who Will Drop Out? A Review of the Predictors of Dropping out of High School: Precision, Sensitivity and Specificity. *The High School Journal.* Available at https://doi.org/10.1353/hsj.2013.0000.

Bowers, A. J. 2021. Early Warning Systems and Indicators of Dropping Out of Upper Secondary School: The Emerging Role of Digital Technologies. En: OCDE, *Smart Data and Digital Technology in Education: Learning Analytics, AI and Beyond*. París: OCDE.

Bruce, M., J. M. Bridgeland, J. H. Fox and R. Balfanz. 2011. On Track for Success: The Use of Early Warning Indicator and Intervention Systems to Build a Grad Nation. Civic Enterprises.

Cohen, A. 2017. Analysis of student activity in web-supported courses as a tool for predicting dropout. *Educational Technology Research & Development*, 65(5): 1285-1304. Available at https://doi.org/10.1007/s11423-017-9524-3.

Coleman, C. 2021. Exploring a Generalizable Machine Learned Solution for Early Prediction of Student At-Risk Status. Thesis presented to Columbia University. Available at https://doi.org/10.7916/d8-5scb-n214.

Costa, F. J., M. S. Bispo and R. F. C Pereira. 2018. Dropout and retention of undergraduate students in management: a study at a Brazilian Federal University. RAUSP Management Journal, 53, 74-85.

Denis, G., M. Hermosilla, C. Aracena, R. Sánchez Ávalos, N. González Alarcón y C. Pombo. 20021. Uso responsable de IA para política pública. Manual de formulación de proyectos. Washington, D.C: BID. Disponible en: https://publications.iadb.org/es/uso-responsable-de-ia-para-politica-publica-manual-de-formulacion-de-proyectos

Departamento de Educación de Estados Unidos. 2016. Issue Brief: Early Warning Systems. Office of Planning, Evaluation and Policy Development. Washington, D.C.: Departamento de Educación de Estados Unidos. Available at https://www2.ed.gov/rschstat/eval/high-school/early-warning-systems-brief.pdf.

Frazelle, S. and A. Nagel. 2015. A practitioner's guide to implementing early warning systems (REL 2015-056). Washington, D.C.: US Department of Education, Institute of Education Sciences, National Center for Education Evaluation and Regional Assistance, Regional Educational Laboratory Northwest. Available at http://ies.ed.gov/ncee/edlabs/projects/project.asp?ProjectID=396.

González, F., T. Ortiz and R. Sánchez. 2020. *Uso responsable de IA para política pública: manual de ciencia de datos*. Washington, D.C.: BID. Available at https://publications.iadb.org/es/ia-responsable-manual-tecnico-ciclo-de-vida-de-la-inteligencia-artificial.

Iam-On, N. and T. Boongoen. 2017. Generating descriptive model for student dropout: a review of clustering approach. En: *Human-centric Computing and Information Sciences*. Available at https://hcis-journal.springeropen.com/articles/10.1186/s13673-016-0083-0.

Jimerson, S., B. Egeland, L. A. Sroufe and B. Carlson. 2000. A Prospective Longitudinal Study of High School Dropouts Examining Multiple Predictors Across Development. *Journal of School Psychology*, 38(6): 525-549. Available at https://doi.org/10.1016/S0022-4405(00)00051-0.

Josephson, K., R. Francis and S. Jayaram. 2018. Políticas para promover la culminación de la educación media en América Latina y el Caribe. Lecciones desde México y Chile. Caracas: CAF. Available at https://scioteca.caf.com/handle/123456789/1246.

Knowles, J. E. 2015. Of Needles and Haystacks: Building an Accurate Statewide Dropout Early Warning System in Wisconsin. *Journal of Educational Data Mining*, 7(3): 18–67.

Lamb, S., E. Markussen, R. Teese, N. Sandberg and J. Polesel. 2010. School dropout and completion: International comparative studies in theory and policy. Berlín: Springer Science & Business Media.

Mac Iver, M. A. and D. J. Mac Iver. 2009. Beyond the Indicators: An Integrated School-Level Approach to Dropout Prevention. Washington, D.C.: George Washington University Center for Equity and Excellence in Education.

Manacorda, M. 2006. Grade Failure, Drop out and Subsequent School Outcomes: Quasi-Experimental Evidence from Uruguayan Administrative Data. Londres: STICERD. (Documento mimeografiado.)

Márquez, V., A. Cano, C. Romero, A. Mohammad, H. Mousa Fardoun and S. Ventura. 2016. Early dropout prediction using data mining: a case study with high school students. Expert Systems. Available at https://onlinelibrary.wiley.com/doi/10.1111/exsy.12135.

Muñoz Stuardo, G. 2020. Sistematización de antecedentes, principales características e implementación del SPTE: fortalezas, debilidades y desafíos para su continuidad y fortalecimiento. Consultoría de Sistematización de la implementación del Sistema de Protección de Trayectorias Educativas (SPTE) del Uruguay.

OECD (Organisation for Economic Co-operation and Development). 2019. Artificial Intelligence in Society. Paris: OECD.

Pombo, C., M. Cabrol, N. González Alarcón and R. Sánchez Ávalos. 2020. fAIr LAC: Adopción ética y responsable de la inteligencia artificial en América Latina y el Caribe. Washington, D.C.: BID. Available at: http://dx.doi.org/10.18235/0002169.

Prenkaj, B., P. Velardi, G. Stilo, D. Distante y S. Faralli. 2020. A Survey of Machine Learning Approaches for Student. ACM Computing Surveys.

Rivera Pizarro, J. 2020. Sistemas de alerta temprana en la región SICA: Revisión de la literatura y de la práctica. San José: SICA. Available at https://ceccsica.info/sites/default/files/inline-files/SISTEMAS%20DE%20ALERTA%20TEMPRANA%20EN%20LA%20REGI%C3%93N%20SICA.pdf.

Román, M. 2013. Factores asociados al abandono y la deserción escolar en América Latina: Una mirada en conjunto. *Revista Iberoamericana sobre Calidad, Eficacia y Cambio en Educación (RE-ICE)*, 11(2): 33-59.

Rumberger, R. W. 2012. Dropping Out: Why Students Drop Out of High School and What Can Be Done About It (Reprint edition). Cambridge, MA: Harvard University Press.

Sansone, D. 2019. Beyond Early Warning Indicators: High School Dropout and Machine Learning. *Oxford Bulletin of Economics & Statistics*. Disponible en: DOI 10.1111/obes.12277.

Sattar, A., F. Mahtab, C. Ratna and R. Chandan. 2016. Survival Analysis based Framework for Early Prediction of Student Dropouts. Proceedings of the 25th ACM International on Conference on Information and Knowledge Management (CIKM), 16: 903-912. Disponible en: https://doi.org/10.1145/2983323.2983351.

Shadish, W. R., T. D. Cook y D. T. Campbell. 2002. *Experimental and quasi-experimental designs for generalized causal inference*. Boston, MA: Houghton, Mifflin and Company.

UNICEF (United Nations Children's Fund). 2018. Early Warning Systems for Students at Risk of Dropping Out: Policy and Practice Pointers for Enrolling All Children and Adolescents in School and Preventing Dropout. UNICEF Series on Education Participation and Dropout Prevention, Vol. 2. New York, NY: UNICEF.

------. 2020. Proteger trayectorias educativas de niñas, niños y adolescentes para mejorar los aprendizajes y la eficiencia del sistema educativo. Nueva York, NY: UNICEF.

University of Chicago Consortium on Chicago School Research. 2014. Selecting effective indicators. College Readiness Indicator Systems Resource Series. Seattle, WA: Fundación Bill &Melinda Gates.

USAID (United States Agency for International Development). 2015. Do Early Warning Systems and Student Engagement Activities Reduce Dropout? Findings from the Four-Country SDPP Evaluation. Washington, D.C.: USAID School Dropout Prevention Summit.

US Department of Education. 2016. Issue Brief: Early Warning Systems. Office of Planning, Evaluation and Policy Development. Washington, D.C.: US Department of Education. Available at https://www2.ed.gov/rschstat/eval/high-school/early-warning-systems-brief.pdf.

IDB
Inter-American
Development Bank