Students' Careers and AI: a decision-making support system for Academia*

Flavio Bertini[1], Alessandro Dal Palù[1,**], Andrea Formisano[2], Andrea Pintus[3], Sara Rainieri[4] and Luana Salvarani[3]

[1]Dept. of Mathematical, Physical and Computer Sciences. Parma University, Italy
[2]Dept. of Mathematics, Computer Science and Physics, Udine University, Italy
[3]Dept. of Humanities, Social Sciences and Cultural Industries, Parma University, Italy
[4]Dept. of Engineering and Architecture, Parma University, Italy

Abstract

In the peculiar realm of higher education, some of the challenges of Public Administration, in terms of quality assurance and data intelligence, can be addressed thanks to the complex ecosystem based on the careers of students and their engagement with the host academia. University governance, ranging from the university Rector and Quality Assurance committee to single heads of degree courses, needs to rely on quantitative and unbiased measures when designing and planning actions. This paper reports on an ongoing project started at Parma University in 2019, that has multiple goals: (1) to collect various sources of students' career-related raw data and to and provide simple access to aggregated analyses through a web portal; (2) to offer an AI based synthesis, in form of automatically generated reports in natural language; (3) to analyze data to detect and predict potential issues (e.g., students drop-out, classes attendance, graduation time estimations, blockages in the career) that can be promptly highlighted, for immediate intervention. As opposed to the majority of academic analytics implementations, particular care is devoted to minimizing ethics and privacy issues and adhering to *explainable AI* principles in the generation of synthetic explanations of charts and reports. The results of lines of research (2) and (3) will be integrated in the portal (1) that is currently deployed at Parma University.

Keywords

learning analytics, explainable Artificial Intelligence, automatic report creation, quality assurance

1. Introduction

In the last decades, learning analytics (LA) received growing attention from educational researchers [1, 2]. Even if there is no general accepted definition of LA, a widely referenced one considers LA as the measurement, analysis and reporting of data about learners, for purposes of understanding and optimizing learning and the environments in which it occurs [3]. A number of benefits arising from LA include the identification of at-risk students, the possibility of developing additional support for coping with academic requirements and expectations [4]. The task of monitoring and improving the quality of the academic experience for students is complex.

In this paper, we present our ongoing work on students' career analytics at Parma University. We focus here on the core aspect of a student's career: exams proficiency through her/his academic life. The whole plethora of services (e.g., housing, libraries, counseling, financial support, sport associations) that contribute to a successful experience have an impact that is less directly measurable and its analysis may result in greater privacy concerns. The design, maintenance and improvement of a degree course require the systematic monitoring of its performance, through gathering and analyzing data about students' careers. The goal is to monitor how the higher education system operates and whether it is reaching, or it will reach with a predictive approach, its objectives and educational targets. To this aim, the central education authority (Ministry) usually defines a set of indicators that are also used to allocate additional resources to universities. Accurate evaluation of these indicators is then essential, and it must include data comparison and effective reporting as well. This scenario involves careful efforts by several actors (quality assurance, degree courses' council and reviewing committees, joint student-teacher committees, single teachers) at different scopes (university, department, degree course, single course). Planned periodical monitoring promotes action planning

Ital-IA 2023: 3rd National Conference on Artificial Intelligence, organized by CINI, May 29–31, 2023, Pisa, Italy

* F.Bertini, A.Dal Palù and A.Formisano are members of the INdAM Research group GNCS. Partially supported by *Interdepartmental Project on AI* (Strategic Plan UniUD–22-25) and by INdAM–GNCS projects CUP E55F22000270001 and CUP E53C22001930001.

** Corresponding author.

flavio.bertini@unipr.it (F. Bertini); alessandro.dalpalu@unipr.it (A. Dal Palù); andrea.formisano@uniud.it (A. Formisano); andrea.pintus@unipr.it (A. Pintus); sara.rainieri@unipr.it (S. Rainieri); luana.salvarani@unipr.it (L. Salvarani)

0000-0001-6925-5712 (F. Bertini); 0000-0003-0353-158X (A. Dal Palù); 0000-0002-6755-9314 (A. Formisano); 0000-0003-3883-4767 (A. Pintus); 0000-0002-5529-2894 (S. Rainieri); 0000-0003-4507-2043 (L. Salvarani)

and evaluation of feedback about previous actions. While projecting students' careers to simple measures (number of freshmen, drop-out/completion rate, average time to pass an exam, etc.) captures only a fraction of the complex dynamics involved, it certainly helps in detecting clear symptoms of potential issues to be further investigated. Our goal is to create a platform that supports academic governance with an effective data-driven pipeline to be integrated into routine activities. The platform enables the structuring of the link between students' outcomes and policy issues. Its design is rather different from typical architectures and it focuses on three building blocks:

1. making information available (Section 2): raw data is processed and aggregated with privacy compliance. Various and rich metrics are introduced in order to capture shades and nuances in career evolutions. A web portal enables browsing information with a Role-Based Access Control approach for visibility;
2. explainable AI-based report creation (Section 3.1): even aggregated analyses require relevant human time to be evaluated. An automated process allows for identifying outliers (i.e., potential issues) in some metrics and synthesizing a discussion about such findings. Moreover, natural language generation controls an unbiased description of charts, for better interpretability and comparison;
3. AI-based predictions (Section 4): university online services usage and its correlation to students' careers can feed predictive models that serve as early detectors of potential issues. Such monitoring can trigger proper actions to be applied while the issues are still developing (and way before the semester is over).

In the literature [5, 6] approaches like ours are classified as Learning Analytics and/or Educational Data Analytics. See [7] for a recent review of data mining techniques used for the prediction of students' drop-out (long-term as well as graduation delays). Interestingly, so far there have been few applications based on Convolutional Neural Networks. There have been specific studies with Italian Universities as case studies: [8], e.g., predicts clusters of students' early drop-outs based on online polls. Various systems have been implemented in the past decade, mainly in USA academia. Predictive Analytics Reporting (PAR) Framework [9] used data analytics to improve student success and retention. In the original formulation analyzed data were: students' backgrounds, GPA and general information about their careers. Since then, such systems expanded and covered many other aspects, including social and financial data. Later, academic analytics examples flourished in Europe as well [10].

Concerns about privacy-related issues in treating personal information have been raised [11]. In [12] another key question is investigated: the implications of educational automation, or, in other words, what kind of responsibility is given to the automated part while being relieved from human activity. Such aspects will be further discussed in Section 3.

1.1. Discussion

Before dwelling into the technical details, let us focus on some key goals and novel opportunities offered by the framework introduced above.

From report creation to report explanation Automated help in raw data processing and issues identification promotes higher quality activities. In particular, it allows one to focus on analyzing possible causes, providing context, explaining dynamics and designing correcting actions to mitigate negative trends, rather than spending time on collecting data, handcrafting charts and writing the report in its descriptive part (which is time-wise predominant). It can be foreseen a shift towards high-level and valuable tasks since manual and mechanic operations are already performed, which translates into increased satisfaction and better use of experts' competencies. This positive impact relies on the usage of an *explainable AI* (or xAI, see Section 3.1). This kind of AI can be trusted, verified and included in high-stakes risk activities.

Unbiased analysis The presence of an xAI-assisted pipeline promotes the reduction of manual errors in data transcription and analysis. Commonly, analyses can be operated by hundreds of people with different roles (e.g., professors, technicians, managers) and backgrounds. A processing that runs on a common baseline smooths out biases in the collection of data. Moreover, natural language descriptions are processed under uniform metrics (e.g., the same modulation of qualitative adjectives) in order to provide unbiased terminology. These standardized metrics allow fair comparisons among different universities, degrees courses and/or the same geographic area.

Improved accuracy vs privacy In common reports and official data of Ministry's reports, quantitative measures are actually limited to macroscopic trends (e.g., number of students/year, amount of credits earned). Finer-level details can be retrieved from raw data exploration and/or from potentially biased investigations (students' feedback about classes, teachers' considerations, etc). Our goal is to include high accuracy (more measures to help uncover small issues that may propagate to large consequences during students' careers) at a limited cost in terms of personal data to be processed. Clearly, anonymized transcripts data (exams proficiency) are at the basis of our analyses. Moreover, we believe that anonymized information about university digital services usage is enough to serve the purpose. In particular, no social nor financial data are included. Quality and care in monitoring data require significant (human) time. Even the availability of an interactive web portal that offers a clear presentation of any analysis does not provide a tool for digesting a fast summary, resulting in a data overload for users. An xAI that sieves relevant facts can handle an increase of metrics and data sources (see Section 4).

Frequency of the analysis An xAI processing of analyses allows one to increase the frequency of monitoring and evaluation of actions' impact. Typically, a monitoring and steering infrastructure meets every 6–12 months (when exams data can be compared). However, it is possible to imagine even real-time monitoring, by moving towards predictive analytics: fresh and available data are needed (by means of a Datalake that collects students' related information) and predictions can be devised. In this case, potential issues can be predicted and tracked earlier and improve the quality assurance impact.

Feedback to students Our data analysis service can be tuned to provide single students a means to track their own progress [13], by comparing individual performance to the one of the associated group (e.g., degree course colleagues, department etc).

2. A web portal for data analysis

Since 2019, Parma University pioneered the deployment of a multi-role internal platform that processes students' data into aggregated analyses. The source comes from views of the Student Management System named Esse3, provided by Cineca (Italian university consortium for research support, IT and HPC services).

The portal is implemented with a combination of Python procedures (for back-end processing), Angular and Echarts (for front-end service). Users are authenticated by the university Shibboleth service and identified according to their role: the head of degree courses can access the analyses carried out for their courses; the head of departments can access degree courses belonging to their department as well as comparisons among those courses; administrators can access to all department data and they can compare all courses in their university.

Weekly, a body of roughly 1.5 million rows relative to students and courses spanning 12 years is processed. Information about each anonymized student's curriculum is processed in order to provide the following main metrics:

- a course has been attended, an exam has been attempted/passed;
- the interval between the last day of lessons and the first exam attempt; the interval between the first attempt and the date of the last attempt or in which the exam was passed;
- number of attempts;
- the grade (in Italian scale between 18 and 30).

The combination of the above presented measures allows the creation of various analyses (around 40 in the current version of the portal), e.g., the load of non-passed exams over time (the main cause of graduation delays); multi-dimensional combinations of proficiency, based on time/mark/number of attempts; delayed exams; patterns in the order of exams. Dashboards and comparative sections highlight distributions and potential outliers from any metric that suggest further investigation.

3. Explainable AI for reports

We believe that simple descriptive analytics, through a web showcase of charts of aggregated data, is only the first step towards a data-driven support system for decision-making. In fact, even at an aggregated level, the amount of information exposed is very large. For example, referred to a large size University as Parma (8K freshmen/year), having roughly 100 degree courses tracked, with an average of 20 exams for 3 years cycle and 12 for 2 years cycles, browsing the rich set of data aggregations and capturing peculiar aspects of careers becomes time-consuming and dispersive.

We envision an additional step devoted to speeding up the task of identifying and reporting issues. Once attention points are clearly defined, it is possible to use AI to generate a higher-level report, written in natural language, that describes relevant charts that contain an issue. Commonly this time-consuming task, namely browsing data, charts generation and text writing is at the basis of any report at any scope level. An automated report can become the basis for the core interpretation, contextualization and decision-making.

The issue with browsing data and their graphic representation is not only about time and practicality. Educational contexts are very sensitive to prejudices, perceptions and assumptions on how different categories of students (slow vs. quick careers, student workers, students with disabilities or learning impairments, non-resident students) can perform and how their pattern of studies will predictably develop towards the degree or the dropout. These assumptions not only can heavily influence the analysis of the data, but also generate professors' attitudes and behaviors that could possibly condition student's learning in order to confirm such assumptions. Using explainable AI for the first level of interpretation can guarantee a bias-free narrative of the most relevant correlation among different phenomena and data sets, providing a powerful instrument for all actors involved and, namely, acting as a continuous professional development tool for professors interested in improving their didactics.

Since the choice of what data to be retained in the report can introduce biases and, in turn, influence political choices, in our opinion, this process must be transparent and trustworthy. We resort to explainable Artificial Intelligence to control the process of selecting the aspects to be described and how to faithfully describe a particular set of data in natural language.

3.1. Explainable Artificial Intelligence

The term *explainable AI* [14, 15] has emerged to capture desirable properties of high-risk systems based on AI. Such systems should ensure transparency, exhibit ethical behavior, and support their results in terms of intelligible descriptions, accountability, security, privacy, and fairness [16]. The adoption of AI systems, especially in Public Administration (PA) contexts, depends on the capability of providing a high-level description of their inner activities. This would promote interpretability and transparency of the inferences that lead to a result.

The urge for explainability in AI applications represents an opportunity for discontinuity with respect to the "traditional" approaches adopted in *sub-symbolic* AI, where the AI system acts as a black box. In other words, such systems, usually relying on Machine Learning (ML), Deep Learning (DL), etc., cannot provide high-level explanations supporting the output of their inferences [17]. The design of an architecture that is both explainable and ML/DL free represents a goal of current research in AI.

As strategic choice, we opt for the use of robust and off-the-shelf technologies of *symbolic* AI, to reach xAI compliance. In this frame of mind, we promote Logic Programming (LP) as the explainable core of an xAI system. LP techniques enable both the representation of knowledge at a higher level of abstraction (ranging from general ontologies to domain-specific knowledge), and the reasoning activity, even by mimicking the human way of thinking. Moreover, in LP-based systems, both the inference steps and the outcome of the reasoning can be immediately justified by singling out which inference rules have been used by the system and how the input knowledge has been processed by these rules. For these reasons, the resulting framework is not only natively explainable, but human users can put themselves in a human-in-the-loop interaction with the xAI system, in order to detect possible flaws in the automated process. This interaction enables a fruitful review of the knowledge base, e.g., to detect incoherent portions of the input, inconsistent inference rules, uncertainty or incompleteness in the sources of the input knowledge.

3.2. Automatic data-to-text

The data-to-text generation (D2T) task consists in automatically generating descriptions from non-linguistic data. Systems able to textually summarize data (e.g, coming from stock prices, healthcare domain etc.), such as time-series, can make data more accessible in cases where the interpretation of visualizations is made difficult or hindered for people with visual impairments or when readers are not expert or have limited cognitive abilities in comprehending and analyzing complex charts.

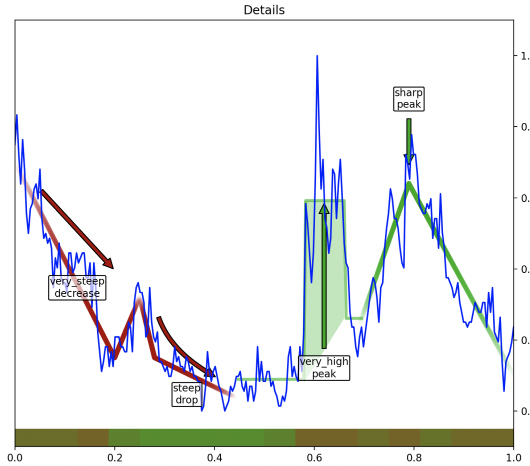

Figure 1: Example of xAI chart visual commentary

The main challenges involved in D2T are the proper identification of *what* to describe —i.e., selecting the key descriptive elements in the input data— and *how* to textually describe such elements in generating the output narration. Our approach, presented in [18], designs an xAI-compliant system integrating Python (to perform raw numerical calculations) and the declarative logic-based framework of Answer Set Programming (ASP) to carry out reasoning. We extract the candidate key descriptors of the series, by applying curve fittings: various of parameterized function prototypes (e.g., lines, polylines, sinusoids, etc.) are matched against fragments of input data. This step, performed by a Python program, produces a collection of candidate descriptions of portions of input data, labeled by a measure of accuracy (e.g., the Root Mean Square Error involved in the approximation). Then, the ASP engine enters into play: The (fragments of) curves are combined to obtain more abstract descriptions of larger portions of the series. For example, a fragment of data described as a decrease followed by a fragment where values increase, are "merged" in a single description of a *valley*. An optimization step identifies the descriptions that better represent the data series. Fig. 1 shows an example where the analysis detects some prominent details in the data series, such as drops and peaks. The last step consists in converting the qualitative descriptions into simple textual narration.

3.3. Automatic report generation

In pursuing transparent processing, an explicit set of attention triggers is defined. For example, a particular distribution is considered to be an outlier if it falls below the 20th percentile. Such filtering among all analyses retrieves those alert cases that will compose the detailed description of the report.

4. AI predictions

In this section we discuss our ideas about predictive analytics in the domain of students' careers. The goal is to predict the evolution of careers so that to anticipate potential disengagement issues ahead of time. As discussed in [19], at the beginning of studies, the compensation of different students' backgrounds is fundamental towards a successful career. Early signs of disengagement dynamics during the first months of academic life should be interpreted as soon as possible by governance, to enact actions like focused tutoring, mentoring and counseling to mitigate those phenomena while they are happening. The typical (manual) revision cycle of courses is performed at the end of the teaching period, possibly enriched by the trend of exams performance. However, any action would be effective only on next year's edition.

Our plan is to develop a fine time granularity analysis that highlights any change in students' engagement at any time, even before measuring the results of the exams. The accuracy and prediction span strongly depend on the type of ingested data. The challenge is to use the *least* amount of information to produce accurate predictions, since privacy issues may arise when intersecting different kinds of personal information. Exams proficiency is also a consequence of students' actions and/or attitudes towards studying. After the spread of the COVID-19 pandemic, teaching methods have adapted to remote emergency teaching [20] and after the release of lockdown restrictions the usage of online resources has been maintained. Digital services can be tracked as fine descriptors of students' engagement with the university. In particular, access to study materials can approximate the amount of relative commitment the student shows in her/his career. Such type of information has different characteristics compared to simple transcripts of the career: it is real-time, accompanies the students throughout their career and can capture different attitudes to studying. It also correlates with exams proficiency, that, however, has a much slower pace (several months). We think that the combination of fine and large-scale phenomena can provide the right balance between accuracy and privacy.

The prediction of a student's career evolution (short and long-term) can be performed with AI tools. It is important to establish whether the amount of opaqueness of a general ML approach (in contrast to xAI) can serve the purpose. It is acceptable to cope with a limited error in the estimation of student drop-outs if the purpose of the prediction is to activate focused support procedures for students. However, when predictions become the basis for other types of interventions, xAI tools should be favored for transparent high-stakes decisions.

Let us present our ongoing research on the AI domain, under the above-mentioned assumptions. In particular, we plan to implement two standard approaches to extract insights from the aggregated careers of students: unsupervised learning and supervised learning. Unsupervised learning models, such as clustering and k-nearest neighbors, discover patterns in untagged data, that is students' careers and use of services in our case. The rationale behind this first line of research is to detect behavioral patterns that are a proxy for student engagement. For example, the use of services can be clustered to identify the group of students who are highly involved in the learning process, the group of those whose involvement wanes as the lessons progress, and finally, the group of students who do not seem to have any interest in the course. The latter activity is particularly challenging because many degree courses have a high percentage of working students who fall into the third group, and anomaly detection techniques may be considered. Overall, this makes it possible to identify critical issues way before the examination session and to intervene with support activities (second group) or better orientation activities (third group). Unsupervised learning models can also be used to identify difficult-to-pass exams, evaluate the relative study approach, and determine whether study groups, which students tend to form on their own, can be beneficial and, where appropriate, foster them with targeted interventions and resources. Complementary, the supervised learning models solve classification and regression problems where the data consists of labeled samples, one example is the students' careers. We plan to exploit supervised models, such as generalized linear models, logistic regression, support vector machines, decision trees and random forest, to address two main time-related issues, that is the time taken to pass every single examination and the time taken to obtain the final title. In the former case, in addition to a basic descriptive statistics analysis, supervised models can help in identifying supporting activities, to be deployed during the course, that could improve performance in terms of final grades and exam pass time. While in the second case, constant career monitoring supported by a model that can predict the progress of examinations passed each year and the date of graduation can be useful for balancing the study load and supplementary resources (e.g., teaching tutors, laboratory activities, study groups).

5. Conclusions

This paper reviewed an ongoing work on learning analytics at Parma University. The multi-role web portal being deployed presents aggregated analyses that help in detecting potential blockages in the career of students and in comparing them at different levels of aggregation (i.e., time-wise, course, degree course, department, university). The portal is the basis for two AI-based challenges: an explainable AI-based automatic report generation and real-time monitoring and prediction of students' careers. The first one can be safely integrated into governance, with advantages in monitoring several metrics with less human cost in preparing documents. The second one offers a solution for a fast-acting governance in order to contribute to lowering the drop-out ratio (according to the EU goals 2030 [21]). Moreover, in a more integrated didactic perspective, both AI-based instruments can be employed in building typologies of students that relate learning styles and approaches with performance in terms of engagement, career length and final outcomes (degrees or dropping out). Caveats about privacy and ethics for the presented approaches have been discussed.

Even if the academic dynamics are not directly applicable to other Public Administration domains, we believe that this methodology could have a positive impact on many other activities. From our point of view, xAI principles and report creation, for example, will be a strategic asset for lighter and more effective quality assurance.

References

[1] G. Siemens, Learning analytics: The emergence of a discipline, American Behavioral Scientist 57 (2013) 1380–1400.

[2] Y. K. Hui, L. F. Kwok, A review on learning analytics, Int. J. of Innovation and Learning 25 (2019) 197–222.

[3] G. Siemens, P. Long, Penetrating the fog: Analytics in learning and education, Educause rev. 46 (2011).

[4] D. Ifenthaler, J. Y. Yau, Reflections on different learning analytics indicators for supporting study success, Int. J. Learn. Anal. AI Educ. 2 (2020) 4–23.

[5] C. Romero, S. Ventura, Educational data mining: a review of the state of the art, IEEE Trans. Syst. Man Cybern., Part C 40 (2010) 601–618.

[6] L. C. Liñán, Á. A. J. Pérez, Educational data mining and learning analytics: differences, similarities, and time evolution, RUSC. Universities and Knowledge Society Journal 12 (2015) 98–112.

[7] F. Agrusti, G. Bonavolontà, M. Mezzini, University dropout prediction through educational data mining techniques: A systematic review, J. of e-learning and knowledge society 15 (2019) 161–182.

[8] A. Sarra, L. Fontanella, S. Di Zio, Identifying students at risk of academic failure within the educational data mining framework, Social Indicators Research 146 (2019) 41–60.

[9] S. L. Watson, B. Gemin, J. Ryan, Using predictive analytics to identify at-risk students in higher education., Educause rev. 47 (2012) 31–40.

[10] A. Bethencourt-Aguilar, et al., The digital transformation of postgraduate degrees. a study on academic analytics at the University of La Laguna, in: Proc. of JICV'21, 2021, pp. 1–4.

[11] N. Sclater, A. Peasgood, J. Mullan, Learning analytics in higher education, Jisc 8 (2016).

[12] N. Selwyn, et al., Digital technologies and the automation of education—key questions and concerns, Postdigital Science and Education (2023) 15–24.

[13] R. Yanosky, D. C. Brooks, Integrated planning and advising services (IPAS) research, 2013. library.educause.edu/resources/2013/8/integrated-planning-and-advising-services-ipas-research.

[14] A. Barredo Arrieta, et al., Explainable artificial intelligence (XAI): Concepts, taxonomies, opportunities and challenges toward responsible AI, Inf. Fusion 58 (2020) 82–115.

[15] A. Adadi, M. Berrada, Peeking inside the black-box: A survey on explainable artificial intelligence (XAI), IEEE Access 6 (2018) 52138–52160.

[16] European Commission, Proposal for a regulation laying down harmonised rules on artificial intelligence, 2021. digital-strategy.ec.europa.eu/en/library/proposal-regulation-laying-down-harmonised-rules-artificial-intelligence.

[17] C. Rudin, Stop explaining black box machine learning models for high stakes decisions and use interpretable models instead, Nature Machine Intelligence 1 (2019) 206–215.

[18] A. Dal Palù, A. Dovier, A. Formisano, An xAI approach for data-to-text processing with ASP, in: Proc. of ICLP'23 (Tech. Comm.), EPTCS, 2023.

[19] C. Aina, E. Baici, G. Casalone, F. Pastore, The determinants of university dropout: A review of the socio-economic literature, Socio-Economic Planning Sciences 79 (2022) 101102.

[20] B. Olasile, S. Emrah, Covid-19 pandemic and online learning: the challenges and opportunities, Interactive Learning Environments 8 (2020) 54–72.

[21] European Commission, Council resolution on a strategic framework for European cooperation in education and training towards the European education area and beyond (2021-2030), 2021.