Ogma: An Intelligent Platform for Anticipating Academic Failure and Recommending Recovery Measures

Ricardo Queirós
CRACS, INESC TEC & ESMAD, Polytechnic of Porto, Portugal

Abstract

The growth of class sizes, the increasing complexity of course content, and the multiplicity of digital platforms make it progressively harder for instructors to identify, in a timely manner, students who are at risk of academic failure. This short paper presents the concept of *Ogma*, a *learning analytics* application with artificial intelligence components designed to anticipate signs of academic decline and suggest personalized recovery measures. The proposal combines heterogeneous data – student-produced work, grade records, forum messages, and logs from other educational applications – to build a dynamic risk profile. Based on this profile, the system provides instructors with a dashboard containing graded alerts, explanations of risk factors, and pedagogical recommendations generated by a pipeline that combines *machine learning*, *retrieval-augmented generation*, and large language models. The paper presents the motivation for the problem, synthesizes related work, and describes the design of a reference architecture for Ogma.

2012 ACM Subject Classification Applied computing → Education; Applied computing → E-learning; Human-centered computing → Interactive systems and tools; Computing methodologies → Artificial intelligence; Human-centered computing → HCI design and evaluation methods; Human-centered computing → Collaborative and social computing systems and tools

Keywords and phrases Generative AI, Educational chatbot, Pedagogical agent, Gamification, Moodle, Retrieval-augmented generation, Self-regulated learning

Digital Object Identifier 10.4230/OASIcs.ICPEC.2026.4

Category Short Paper

Funding *Ricardo Queirós*: This work is funded by national funds through FCT – Fundação para a Ciência e a Tecnologia, I.P., under the support UID/50014/2025 (`https://doi.org/10.54499/UID/50014/2025`).

1 Motivation

Early identification of at-risk students is a central problem in higher education. In demanding courses taught to large and heterogeneous classes, instructors rarely have enough granularity to closely monitor the individual progression of each student. Even when abundant data exist, such as learning management system accesses, forum participation, intermediate grades, and assignment deadlines, these signals are typically fragmented across multiple screens and systems and are not presented in an actionable way. The learning analytics literature shows precisely that collecting, measuring, and analyzing student data can support the understanding and optimization of learning and of the contexts in which it occurs [5, 2].

In real settings, the problem is not merely to predict a low final grade. The pedagogical goal is to detect patterns of disengagement before failure materializes: fewer Moodle accesses, repeated delays in assignment submission, reduced participation in forums, signs of frustration in messages, sudden drops in formative assessment performance, or overload signals observable in other educational platforms. Early warning systems have shown usefulness in supporting more timely interventions [1, 3]. However, many solutions still depend on a limited number of variables, remain insufficiently transparent to instructors, or are weakly connected to concrete recovery recommendations.

This is where *Ogma* emerges. Its central idea is to build a platform capable of aggregating multimodal and multi-application evidence to produce two complementary outputs. First, an updated risk score throughout the semester. Second, intervention suggestions contextualized to the course, the student, and the teaching moment. Instead of showing instructors only a predictive value, the system should answer operational questions: *who* is falling behind, *why*, *with what level of urgency*, and *which recovery measure* makes sense to propose now.

This approach is particularly relevant in complex domains such as programming, mathematics, statistics, or engineering, where small early gaps tend to propagate and produce cumulative failure. In such scenarios, delays in detection may make recovery unfeasible. A tool such as Ogma can bridge the gap between the abundance of digital data and the need for fast, sustained, and explainable pedagogical intervention.

2 Related Work

The field of *learning analytics* has become established over the last decade as an intersection of data science, educational technology, and pedagogical decision support. Foundational texts stress that the usefulness of data analysis in education does not depend solely on predictive performance, but also on its translation into concrete actions in instructional design and teaching intervention [5, 2]. This concern is especially visible in the research line on *early warning systems*, whose purpose is to flag at-risk students early enough for institutions and instructors to act.

Among the most widely cited early warning systems is *Course Signals*, developed at Purdue University, which combined academic and engagement indicators to classify students into risk levels and communicate status through a simple color scheme [1]. In a similar vein, the Open Academic Analytics Initiative showed that predictive models can be operationalized at the institutional level to support scalable alerts and interventions [3]. More recent review studies confirm the maturity of the field and show that academic performance prediction increasingly relies on *ensemble learning*, decision trees, *gradient boosting*, and deep models, although challenges of generalization, explainability, and transfer across contexts remain [7, 6].

A second line of work relevant to Ogma concerns the multimodal nature of the data. In many courses, risk signals are not found only in Moodle clicks, but also in late submissions, forum messages, frequency of access to materials, communication patterns, and the textual quality of student productions. Recent reviews show that combining interaction data with textual and socio-emotional data can improve the sensitivity of models to states of disengagement and frustration [9]. This observation justifies including sentiment analysis over forum messages, team channels, or other authorized educational applications in the Ogma proposal.

A third line concerns the presentation of results. Learning analytics dashboards for instructors can improve situational awareness about the class, but they are effective only when they make alerts understandable, actionable, and aligned with real pedagogical decisions [10, 8]. Alerts without explanation tend to generate mistrust; visualizations without prioritization generate overload. For this reason, the Ogma interface should combine visual simplicity with explanations of the main risk factors.

Finally, the emergence of large language models creates an opportunity to move from prediction to recommendation. The *retrieval-augmented generation* (RAG) architecture makes it possible to ground text generation in specific document bases and reduce the probability of generic or decontextualized answers [4]. In the context of Ogma, this means that instructors can upload pedagogical materials, course regulations, and previously defined recovery strategies, creating a knowledge base that supports recommendations better suited to the local context.

3 Design and Implementation

3.1 Architecture overview

The proposed architecture was designed to support the early detection of academic failure through the integration of Learning Analytics, Machine Learning, Explainable Artificial Intelligence (XAI), and Generative AI technologies. The system aims not only to identify students at risk but also to support teachers through contextualized recommendations and intelligent decision-support mechanisms.

The architecture follows a modular and layered approach composed of six main components:

1. Data collection
2. Data processing and feature engineering
3. Predictive analytics
4. Explainability
5. Retrieval-augmented generation
6. Visualization

3.1.1 Data collection

The first layer is responsible for data acquisition from multiple educational platforms. Moodle acts as the primary source of learning analytics data, including login frequency, access patterns, assessment results, assignment submissions, forum participation, and interaction with learning resources. Additional academic and administrative information, such as previous grades, completed credits, repeated courses, and student status, may also be integrated. To ensure interoperability between heterogeneous systems, the architecture incorporates MCP (Model Context Protocol), allowing the platform to communicate with external systems such as academic management software, Microsoft Teams, Discord, and other services.

3.1.2 Data processing and feature engineering

After collection, the data is processed and transformed into features suitable for predictive analysis. This stage includes data cleaning, normalization, aggregation, and feature engineering. Several behavioral and engagement indicators are generated, such as weekly activity consistency, delays in submissions, learning regularity, inactivity periods, and academic performance trends. These features are then stored in a centralized feature repository to support both training and inference processes.

3.1.3 Predictive analytics

The predictive layer is based primarily on supervised machine learning techniques. Since the objective is to predict academic risk using historical data, supervised learning represents the most appropriate approach. Historical datasets from previous academic years are used to train models capable of classifying students according to different risk levels. Logistic

Regression may be used as a baseline model due to its interpretability, while ensemble-based approaches such as Random Forest and XGBoost are expected to provide higher predictive performance for tabular educational data.

3.1.4 Explainability

To increase transparency and trustworthiness, the architecture integrates an Explainable AI layer. Techniques such as SHAP (SHapley Additive exPlanations) – a popular Explainable AI (XAI) framework that uses game theory to explain the output of machine learning models – are used to identify the contribution of each feature to the prediction outcome. This allows teachers to understand why a student was classified as being at risk, highlighting factors such as low engagement, missing assignments, or declining academic performance.

3.1.5 Retrieval-augmented generation

Beyond prediction, the architecture also incorporates a Retrieval-Augmented Generation (RAG) component combined with Large Language Models (LLMs). While the machine learning model is responsible for identifying risk, the LLM layer focuses on generating contextualized pedagogical recommendations. The RAG mechanism retrieves relevant information from institutional regulations, pedagogical guidelines, previous interventions, and scientific literature before sending the context to the language model. Based on this information, the system may generate personalized corrective measures, intervention suggestions, summaries, or teacher-oriented recommendations.

3.1.6 Visualization

Finally, the visualization layer provides an intelligent dashboard for teachers and academic coordinators. The dashboard presents risk indicators, predictive scores, behavioral trends, explainability outputs, and recommended interventions using visual alerts and color-based prioritization mechanisms. This approach supports early intervention and facilitates pedagogical decision-making without replacing the role of the teacher.

3.2 AI Workflow Orchestration

In addition to the architectural layers previously described, the proposed system incorporates an orchestration layer responsible for coordinating the interaction between predictive models, explainability services, retrieval mechanisms, generative AI components, and visualization modules. Since the architecture combines heterogeneous technologies and processing stages, an orchestration mechanism is required to manage the execution flow and ensure interoperability between components. Figure 1 summarizes this workflow.

To support this process, the system adopts LangFlow as the orchestration framework. LangFlow was selected due to its modular and visual workflow capabilities, integration with LangChain-based components, support for Retrieval-Augmented Generation pipelines, and suitability for rapid prototyping and research-oriented environments.

The orchestration layer acts as the runtime coordination engine of the platform. Rather than executing machine learning training processes directly, LangFlow is responsible for managing the inference and decision-support workflows triggered by new educational events or user interactions. Predictive models are trained externally using Python-based machine learning frameworks and deployed as independent services accessible through APIs. The orchestration layer then invokes these services whenever predictions are required.

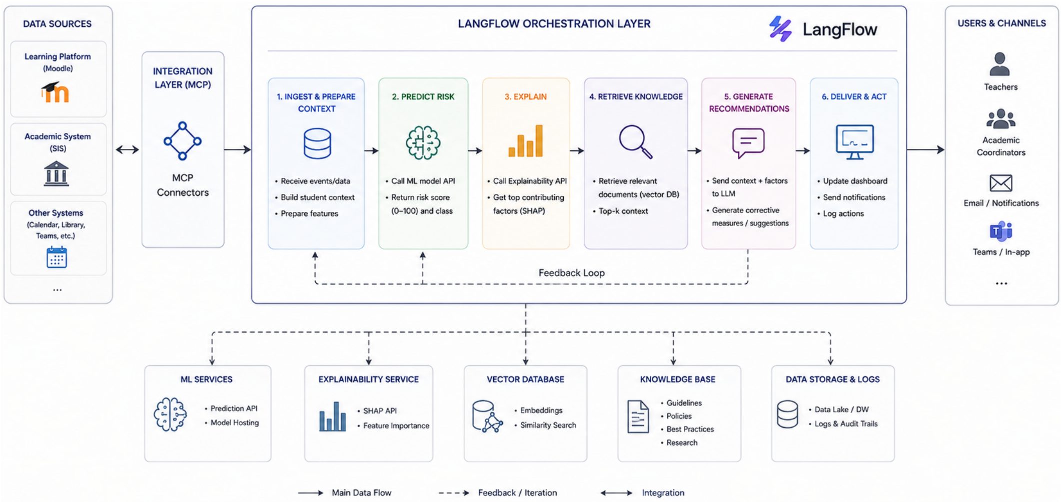

Figure 1 AI Workflow Orchestration of Ogma ecosystem (created by ChatGPT v. 5.3 using as prompt the content of this section).

The execution flow begins with the ingestion of new data from Moodle or other integrated systems through the MCP integration layer. After preprocessing and feature extraction, the student profile is sent to the predictive service, which returns a risk classification score. When the predicted risk exceeds a predefined threshold, the orchestration workflow automatically activates additional components, including explainability modules and retrieval pipelines.

The explainability service generates feature importance information using SHAP values, identifying the main factors contributing to the prediction outcome. Simultaneously, the RAG pipeline retrieves relevant contextual information from the knowledge base, including institutional policies, pedagogical guidelines, previous interventions, and scientific literature. This contextual information is then forwarded to the Large Language Model, which generates personalized recommendations, corrective measures, or intervention suggestions for the teacher.

The orchestration layer also supports event-driven workflows. Instead of operating exclusively through manual requests, the system may react automatically to behavioral changes such as prolonged inactivity, missing submissions, repeated low assessment scores, or sudden decreases in engagement. These events can trigger prediction pipelines, recommendation generation, and notification mechanisms in real time.

Finally, the generated outputs are delivered to the visualization layer, where teachers and academic coordinators can access dashboards containing risk indicators, explanatory insights, and recommended interventions. This orchestration approach enables the integration of predictive analytics and generative AI into a unified educational intelligence workflow capable of supporting proactive pedagogical intervention strategies.

4 Conclusion

In summary, Ogma proposes an evolution of early warning systems toward an intelligent pedagogical platform oriented to intervention. Its main conceptual contribution lies not only in predicting academic failure, but in transforming fragmented data into explainable signals and contextualized recovery measures, supporting instructors in managing complex classes and mitigating academic failure in a timely way.

This paper represents only the starting point of a broader work that will be developed over the next year, in which the goal is to implement and validate this architecture in real higher education contexts.

References

1 Kimberly E. Arnold and Matthew D. Pistilli. Course signals at purdue: Using learning analytics to increase student success. *Proceedings of the 2nd International Conference on Learning Analytics and Knowledge*, pages 267–270, 2012. `doi:10.1145/2330601.2330666`.

2 Rebecca Ferguson. Learning analytics: drivers, developments and challenges. *International Journal of Technology Enhanced Learning*, 4(5–6):304–317, 2012. `doi:10.1504/IJTEL.2012.051816`.

3 S. M. Jayaprakash, E. W. Moody, E. J. M. Lauría, J. R. Regan, and J. D. Baron. Early alert of academically at-risk students: An open source analytics initiative. *Journal of Learning Analytics*, 1(1):6–47, 2014. `doi:10.18608/jla.2014.11.3`.

4 Patrick Lewis, Ethan Perez, Aleksandras Piktus, Fabio Petroni, Vladimir Karpukhin, Naman Goyal, Heinrich Küttler, Mike Lewis, Wen-tau Yih, Tim Rocktäschel, Sebastian Riedel, and Douwe Kiela. Retrieval-augmented generation for knowledge-intensive NLP tasks. In *Advances in Neural Information Processing Systems 33*, pages 9459–9474, 2020.

5 Phil Long and George Siemens. Penetrating the fog: Analytics in learning and education. *EDUCAUSE Review*, 46(5):31–40, 2011.

6 Jairo López-Zambrano, Juan A. Lara, and Cristóbal Romero. Towards portability of models for predicting students at risk of failing and dropping out in university. *Applied Sciences*, 11(23):11363, 2021. `doi:10.3390/app112311363`.

7 Abdullah Namoun and Abdullah Alshanqiti. Predicting student performance using data mining and learning analytics techniques: A systematic literature review. *Applied Sciences*, 11(1):237, 2021. `doi:10.3390/app11010237`.

8 Martin Paulsen and Kasper Hornbæk. Learning analytics dashboards for teachers: A systematic review of design, data, and actionability. *Computers and Education Open*, 6:100176, 2024. `doi:10.1016/j.caeo.2024.100176`.

9 A. Pilicita Garrido and E. Barra. Sentiment analysis with transformers applied to education: Systematic review. *International Journal of Interactive Multimedia and Artificial Intelligence*, 9(2):72–83, 2025. `doi:10.9781/ijimai.2025.02.008`.

10 Beat A. Schwendimann, María Jesús Rodríguez-Triana, Andrii Vozniuk, Luis P. Prieto, Mina S. Boroujeni, Adrian Holzer, Denis Gillet, and Pierre Dillenbourg. Perceiving learning at a glance: A systematic literature review of learning dashboard research. *IEEE Transactions on Learning Technologies*, 10(1):30–41, 2017. `doi:10.1109/TLT.2016.2599522`.