Integrating Variability Models and Ontologies for LLM-Based Multi-Agent Systems in Academic Domains

Noelia Lopez-Duran[1,*,†], Jesús Moreno-León[2,†] and Antonio Balderas[1,†]

[1]Universidad de Cádiz, Escuela Superior de Ingeniería, Avda. República Saharaui, s/n, 11510 Puerto Real, Cádiz, Spain
[2]Universidad de Sevilla, Escuela Técnica Superior de Ingeniería Informática, Avda. de Reina Mercedes, s/n, 41012 Sevilla, Spain

Abstract

This doctoral research investigates how different knowledge representation models can be combined and exploited as a shared knowledge base for multi-agent systems (MAS) in educational contexts. It focuses on the integration of feature models, expressed in Universal Variability Language (UVL), and ontologies to represent entities such as academic programs, courses, roles (e.g., students, professors, administrative staff), and their relationships. While UVL enables the explicit modelling of variability, constraints, and configurations, ontologies provide richer semantic descriptions and relationships between concepts. The research examines how the knowledge resulting from their integration can be adapted into representations that can be effectively interpreted and used by LLM-based agents. Preliminary work compares the use of UVL models and equivalent natural language prompts as knowledge bases for a single large language model, analysing their impact on reasoning and constraint handling. The goal is to support agents in reasoning over structured, semantically grounded, and configurable educational knowledge, providing a robust foundation for Learning Analytics by ensuring that agent-driven interactions are aligned with institutional data and pedagogical constraints.

Keywords

Multi-agent systems, LLMs, Ontology, OWL, Feature model, Universal Variability Language, Knowledge representation

1. Introduction and Motivation

The integration of Artificial Intelligence (AI) and Large Language Models (LLMs) into educational environments is reshaping how students interact with digital systems. These technologies enable new forms of interaction, including conversational interfaces, automated feedback, and on-demand support [1, 2, 3]. As a result, these technologies bring new opportunities to enhance how learners interact with educational systems, enabling more adaptive and data-informed interactions.

Despite this potential, the adoption of LLM-based systems in education requires not only access to domain knowledge but also an understanding of how that knowledge can be interpreted and used in context. In practice, these systems primarily rely on implicit knowledge acquired during training or conveyed through prompts, which may not sufficiently represent the structural and organizational characteristics of educational environments. This lack of explicit structure hinders the ability to perform meaningful Learning Analytics (LA), as there is no formal trace of how the system's pedagogical decisions relate to the actual institutional constraints and learner data.

Educational environments are complex and heterogeneous, involving diverse institutional structures, organizational constraints, contextual dependencies, and user roles. General-purpose AI systems often lack explicit awareness of these aspects, which can result in inconsistent behaviour, limited interpretability, and difficulties in aligning system outputs with the reality present in educational ecosystems.

To address these challenges, structured knowledge representations have been widely used in other domains. Feature Models (FMs), for example, are employed in Software Product Line Engineering

LASI Spain 26: Learning Analytics Summer Institute Spain 2026, Zamora, May 07–08, 2026
*Corresponding author.
†These authors contributed equally.
noelia.lopezduran@uca.es (N. Lopez-Duran); jmorenol@us.es (J. Moreno-León); antonio.balderas@uca.es (A. Balderas)
0009-0002-9783-9779 (N. Lopez-Duran); 0000-0002-3821-5707 (J. Moreno-León); 0000-0003-0026-7410 (A. Balderas)

to capture variability and constraints in complex systems. Textual formats such as the Universal Variability Language (UVL) [4] may enable these models to be interpreted and used by LLM-based systems, facilitating more explicit contextualization, enabling a more interpretable and controllable form of contextualization.

Ontologies are also central to this proposal because they make it possible to formally describe the semantic relationships between educational entities, such as academic programmes, courses, competences, prerequisites, roles, services, and institutional procedures. While FMs are useful for representing configurable alternatives, hierarchy and constraints, ontologies can capture meaning, and complex relationships between concepts. Therefore, this PhD project investigates how both types of representation can be combined to provide LLM-based systems with structured and semantically grounded knowledge about academic domains.

In parallel, Multi-Agent Systems (MAS) provide a paradigm for coordinating multiple specialized agents [5]. This approach distributes responsibilities, supporting modular and scalable solutions compared to monolithic approaches [6]. In educational contexts, this can facilitate the integration of contextual knowledge into targeted tasks by distributing knowledge across agents that interact to generate responses grounded in richer context.

Building on these ideas, our proposal explores how structured knowledge representations can support the contextualization of LLM-based systems in educational scenarios. We investigate the use of formal representations of institutional knowledge and consider their integration within multi-agent architectures, with the aim of enabling more context-aware, reliable, and adaptable interactions for learners, while facilitating the extraction of actionable insights for LA.

2. Research Problem and Objectives

Despite the growing adoption of LLM-based systems in educational contexts [7], their ability to operate consistently across different educational scenarios remains limited. A key issue is the absence of explicit mechanisms to represent and incorporate contextual knowledge into the training and deployment process. As a result, system behaviour depends heavily on implicit knowledge and prompt design, which may lead to biases, variability in outputs and reduced interpretability.

This limitation is particularly relevant in educational environments where structural constraints and relationships between entities play a central role. Existing approaches do not provide a systematic way to represent such variability in a form that can be effectively leveraged by LLM-based systems. Consequently, there is a gap between the complexity of educational contexts and the mechanisms currently used to model them in generative systems.

In this context, our proposal focuses on the integration of structured knowledge representations with LLM-based approaches to support contextualization in LA. The objective is to investigate how explicit models of variability and constraints can be incorporated into the generation process, and how this integration affects the consistency and interpretability of system outputs. Additionally, this work explores the use of multi-agent architectures as a means to distribute and operationalize contextual knowledge across different components.

To guide this investigation, the following research questions are defined:

RQ1: How can the variability and constraints of educational environments be formally represented and transformed into LLM-interpretable formats?

RQ2: How can contextual knowledge be distributed and operationalized across multiple agents within a multi-agent architecture?

RQ3: How does the integration of structured knowledge models affect the quality of generated interactions and the traceability of data for Learning Analytics purposes?

The goal of this proposal is to define a methodology for the design of context-aware multi-agent systems in educational environments, enabling institutions to develop solutions that support the different actors and processes involved in academic settings.

Table 1
Alignment between research questions, methodological focus, and current status.

RQ	Focus	Methodological strategy	Status
RQ1	Formal representation of variability, constraints, and semantic relationships in academic environments.	Compare natural language, UVL-based FMs, ontology-based representations, and hybrid representations in controlled LLM tasks.	Initial FM vs. natural language experiment completed.
RQ2	Distribution and operationalization of contextual knowledge in a MAS.	Define agent profiles from the structured models and evaluate coordination mechanisms for answering academic queries.	Conceptual architecture defined
RQ3	How can Learning Analytics support the monitoring and adaptation of LLM-based educational support? Learning interaction traces, learner activity indicators, and system usage data	Deployment and evaluation in a real educational test case scenario, considering learner interaction patterns, pedagogical usefulness, and support for instructional decision-making	Defined as part of the project design; empirical validation will be conducted through a real test case scenario

Although the three research questions define the conceptual scope of the dissertation, each of them is associated with a distinct methodological strategy and a different level of maturity within the current stage of the project. For this reason, Table 1 summarises the alignment between the research questions, the corresponding methodological strategy, and the present status of development, clarifying how the different components of the research are progressively articulated.

3. Background and Related Work

Numerous studies have explored the use of knowledge modelling techniques to represent characteristics of academic environments. In particular, ontologies have been widely adopted to formalize educational concepts, relationships, and entities, often with a strong emphasis on learner modelling [8, 9, 10]. These approaches have contributed to the representation of profiles, competencies, preferences, and learning trajectories, supporting personalization and adaptive learning processes. However, much of this work has focused on modelling specific dimensions of the educational context rather than capturing the broader structure, variability, and constraints of academic environments as a whole.

In parallel, variability modelling has been investigated in educational technology through approaches inspired by Software Product Line Engineering. These works have shown the usefulness of FM representations for describing configurable educational systems, especially Learning Management Systems and related platforms. However, their main focus has typically been placed on the configuration and customization of technological infrastructures, rather than on representing educational environments [1, 11, 12].

At the same time, Multi-Agent Systems have been explored as a paradigm for building modular and distributed solutions in complex domains, including education [5, 3]. Their ability to decompose functionality across specialized agents makes them particularly suitable for scenarios where different sources of knowledge and responsibilities need to be coordinated. Recent developments combining LLMs and agent-based architectures suggest new opportunities for distributing reasoning processes and contextual information across multiple interacting components.

Despite these advances, the intersection between structured knowledge modelling, variability representation, and multi-agent LLM-based systems remains underexplored. Existing ontology-based approaches provide formal representations of limited and sometimes context-specific educational concepts, while SPL-based approaches capture aspects of platform variability, but neither line of work fully addresses the problem of representing the entire academic environments in a form that can be operationalized by LLM-based multi-agent systems. In this context, our proposal is centred at the

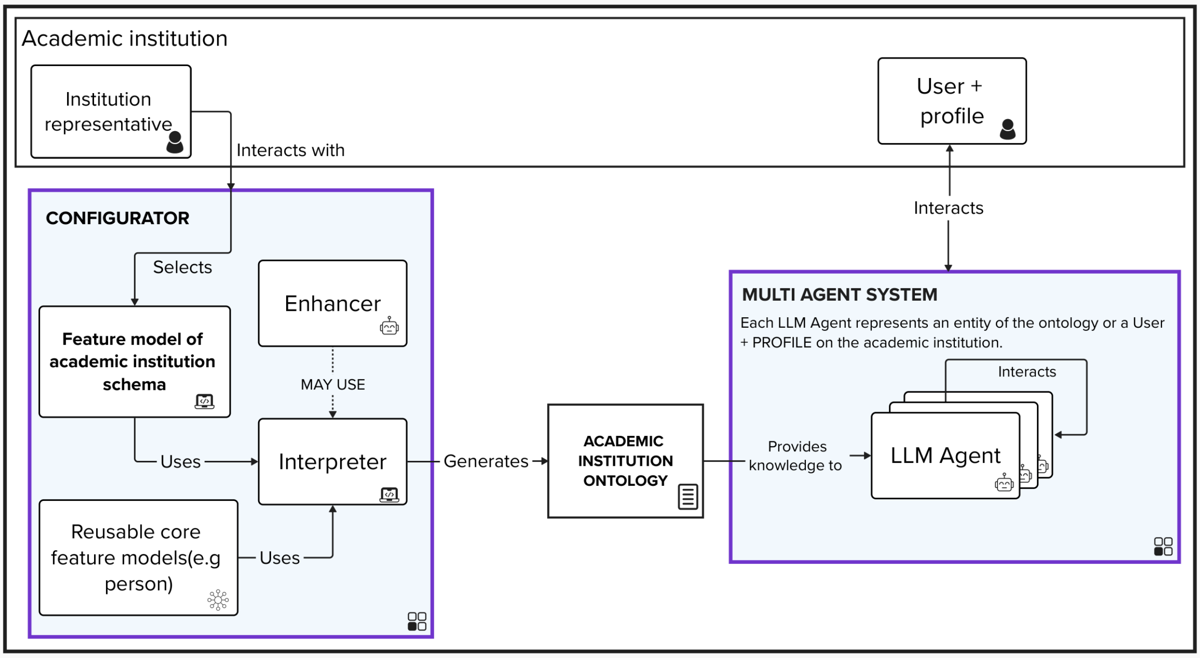

Figure 1: Conceptual representation of the proposed architecture, showing the interaction between configuration components, knowledge representations, and reasoning processes within Multi Agent System.

intersection of these areas, with the objective of investigating how structured models of educational environments can support the design of contextualized and usable multi-agent systems for the academic ecosystem.

4. Proposed Research Approach

As illustrated in Figure 1, the proposal is based on a hybrid model that combines the variability modelling capabilities of FMs with the semantic expressiveness of ontologies. FMs are used to capture the structure and variability of the academic environment, while ontologies provide a formal representation of entities and their relationships. The resulting model is transformed into a structured knowledge base that supports the behaviour and interactions of a multi-agent system.

The approach is organised as a pipeline that begins with the definition of the academic environment through a configurable schema. This schema allows institutional representatives to describe the relevant elements of their context by selecting and adapting features that reflect their specific organisational structure, academic rules, and support needs. The identification of these relevant entities and contextual constraints is informed by continuous consultation with actors from the target educational setting, including teaching staff, so that the resulting representations are aligned with authentic institutional requirements. In addition to institution-specific configurations, reusable core models are incorporated to represent common structural patterns across educational environments, enabling consistency and reuse across different scenarios.

Once defined, the feature-based representation can be used as one possible basis for deriving a corresponding ontology that formalises selected elements and their relationships. This transformation is not assumed to be the only path toward machine-interpretable contextual knowledge, but serves in this research as an operational mechanism for examining how variability-oriented structures can be complemented with semantic formalisation. The use of feature models and ontologies thus responds to their complementary representational roles: feature models provide an explicit mechanism for encoding variability, optional configurations, and institutional constraints, while ontologies contribute a semantic layer capable of expressing conceptual hierarchies, dependencies, and shared domain meaning.

The project therefore does not assume that either formalism is sufficient in isolation, but explores whether their hybrid combination can jointly support a richer and more adequate machine-interpretable representation of educational environments than relying on each of them independently.

The generated ontology is then used as the basis for the MAS. Within this architecture, each agent is associated with a specific entity or profile derived from the ontology, allowing the system to distribute contextual knowledge across multiple specialised components. These agents interact with each other to address user requests, combining their respective knowledge to produce responses that are grounded in the defined academic context.

This design enables the system to adapt to different institutional configurations while maintaining a consistent interaction mechanism. By separating the formal definition of the academic environment from its implementation within the MAS, the approach supports flexibility and reuse, allowing institutions to adjust the represented contextual knowledge to their specific needs without modifying the underlying orchestration logic. In this sense, the MAS functions as the computational layer responsible for operationalising, coordinating, and delivering distributed contextual knowledge through specialised academic support agents.

Methodologically, the research will be developed in four phases. First, the academic domain will be modelled using UVL feature models and ontologies, identifying which elements are better represented as variability constraints and which require semantic relations. Second, these representations will be transformed into LLM-interpretable artefacts, such as structured prompts, serialized model fragments, retrieval units, or agent-specific knowledge profiles. Third, the resulting knowledge base will be operationalized in a MAS, where different agents use selected parts of the model to answer academic queries, validate constraints, and coordinate responses. Finally, the approach will be evaluated in educational scenarios by analysing both the quality of the generated interactions and the traceability of the information used to produce them. In this context, quality refers to the correctness, contextual adequacy, completeness, and pedagogical usefulness of the generated responses, whereas traceability concerns the possibility of identifying which formal knowledge elements, institutional constraints, or contextual fragments contribute to the final answer.

This evaluation stage is intended to examine not only the technical performance of the proposed LLM-based agents, but also their educational impact when situated in authentic learner interactions. By deploying the system in a real educational test case scenario, it will be possible to analyse how students engage with the agent, what types of questions and support requests emerge, how learners react to responses generated from different knowledge representations, and to what extent these interactions influence understanding, engagement, and learning behaviour. In this sense, LA is incorporated as a complementary layer for observing and interpreting learner-system interaction patterns, providing empirical evidence about the pedagogical implications of the proposed approach.

5. Preliminary Work and Current Status

The PhD project is currently in its initial stage. The completed work includes the definition of the research problem, the design of the conceptual architecture, and a first exploratory experiment focused on RQ1. This experiment does not yet address the complete hybrid representation proposed in the thesis; instead, it examines whether an explicit feature-model representation can improve LLM performance when compared with an equivalent natural language description.

The next stages will address the extension of the representation layer with ontologies, the transformation of structured information to LLM-interpretable format, the implementation of the MAS described in RQ2 and the design of LA indicators and evaluation mechanisms for RQ3.

The prototype implements two alternative approaches for representing the academic environment. The first approach relies on structured representations based on FMs, expressed using UVL to capture variability and constraints. The second approach uses natural language descriptions to provide contextual information to the model. Both formulations were constructed from the same institutional context, combining information gathered during the early modelling phase with academic staff and publicly

```
features
    University
        mandatory
            Course
                optional
                    Lab
                    Project
                    CourseType
                        or
                            Theory
                            Practical
constraints
    Lab => Practical
```

The university offers different courses. A course may include a lab or a project. Courses can be theoretical or practical, but labs are only offered in practical courses.

Figure 2: Side-by-side comparison of a concept described in UVL (left) and natural language (right).

available institutional sources, such as official university webpages and academic documentation, in order to reflect realistic organisational structures and frequent support needs. Rather than being jointly employed, these two representations are independently used to guide the responses of an LLM in comparable tasks. Figure 2 illustrates a simplified example of both representations, highlighting the differences between structured and natural language formulations.

The evaluation was conducted through controlled interactions in which the model was asked to answer a set of five questions requiring reasoning over the provided context. These questions focused on constraint validation and dependency understanding. For example:

- Q1: Where can I apply for the issuance of my academic degree certificate?
- Q2: I am Student A, and I would like to choose third-year courses together with my friends. Which courses can I take alongside my friends Student B and Student C, given my academic record?

To evaluate the responses, we defined question-specific rubrics following a common structure. Each rubric assesses (i) the correctness and completeness of the required elements in the answer, and (ii) the inclusion of additional relevant information, while penalizing unsupported or hallucinated content. Each expected element is evaluated independently using a binary score (1 if satisfied, 0 otherwise) and additional criteria are applied when relevant, such as verifying the correctness of referenced entities (e.g., course names).

For the two questions discussed in this section, the rubrics can be summarized in terms of the key elements expected in each response:

- **Q1 - Administrative query:**
  - Identification of the appropriate administrative unit (e.g., Secretariat)
  - Consideration of institutional context (e.g., campus differences)
  - Reference to official university resources (e.g., student service webpages)
  - Description of the application process (e.g., forms or procedures)
- **Q2 - Constraint-based reasoning:**
  - Identification of valid third-year courses
  - Verification of enrolment constraints (completion of previous years)
  - Correct interpretation of each student's academic status
  - Determination of feasible course selections for all students
  - Consistency of the final recommendation across students

The results suggest that the effectiveness of each representation appears to depend on the type of task. For the administrative query (Q1), the natural language representation produces a more complete and contextually appropriate answer. This is because it captures procedural and descriptive information that is not explicitly modelled in the structured representation, such as official processes.

In contrast, for the constraint-based reasoning task (Q2), the structured UVL representation leads to more accurate and consistent responses. The explicit modelling of dependencies and constraints enables the LLM to better interpret the conditions required for course selection and to reason over the academic status of multiple students. The natural language representation, while partially correct, tends to omit or misinterpret some of these constraints, resulting in less reliable answers.

These findings suggest that neither representation is sufficient on its own. While structured models improve consistency and control, they may lack the expressiveness required to capture more complex or implicit aspects of the context. Conversely, natural language representations facilitate richer descriptions but do not provide explicit mechanisms to enforce constraints. This observation motivates the exploration of hybrid approaches that combine both forms of representation.

The preliminary work provides initial evidence supporting the central hypothesis of this proposal, namely that the explicit modelling of context can influence the behaviour of LLM-based systems.Although the present experiment is intentionally limited in scope, it establishes a first empirical basis for the continuation of the project in two directions: first, by extending the comparison toward ontology-based and hybrid representations; and second, by integrating these structured models within the proposed multi-agent architecture for more complex educational assistance scenarios.

6. Future Work and Research Plan

This doctoral research is framed within the GENIALLE project (GENerative Intelligent Agents for digitaL Learning Environments), a framework that leverages LLM-based agents to extend digital platforms with social and proactive capabilities. Building on the preliminary work and as indicated in Table 1, the next stage of this research focuses on extending the current prototype towards a structured architecture for the contextualization of LLM-based MAS in educational environments within the GENIALLE ecosystem. This involves the following aspects: advancing the representation model, integrating multi-agent architectures, and evaluating the approach in realistic scenarios.

A first line of work consists in extending the current representation approach beyond feature-based models by incorporating semantic representations. In particular, ontologies will be integrated to complement the variability modelling capabilities of FMs, enabling a richer and more expressive description of entities and their relationships. This integration aims to support a hybrid representation that combines structural and variability constraints with semantic information, addressing the limitations identified in the preliminary study.

A second line of work focuses on defining and comparing the mechanisms through which these hybrid representations can be translated into LLM-usable contextual artefacts. Rather than assuming that the existence of a formal model is sufficient by itself, this stage of the research will investigate how the encoded knowledge can be exposed to the language model through different strategies, such as structured prompt serialization, retrieval-oriented contextual fragments, or intermediate machine-readable representations. The objective is to determine which forms of knowledge injection allow the LLM to better exploit explicit academic constraints and semantic relations while preserving interpretability and controllability in the generation process.

In parallel, the research will explore the design and implementation of a MAS that distributes contextual knowledge across specialised agents. This involves defining how different agents are associated with specific parts of the structured representation, as well as how they interact to address complex queries. The goal is to enable a modular and scalable system in which contextual reasoning is achieved through the collaboration of multiple agents.

Finally, the proposed approach will be evaluated in real educational scenarios to assess its applicability and impact. This includes analysing how the integration of structured knowledge representations and

MAS affects how users interact and work to improve or resolve problems of the academic ecosystem. The evaluation is expected to provide insights into the feasibility of adopting this approach in heterogeneous academic environments.

The doctoral research follows an iterative methodological approach that combines technical design, controlled experimentation, and authentic educational validation. In an initial phase, formal knowledge representations and agent coordination mechanisms are designed and tested through controlled LLM-based tasks. In a second phase, the resulting architecture will be deployed in a real educational context in order to analyse not only response quality, but also patterns of learner interaction, perceived usefulness, and pedagogical implications. This progression allows the project to move from technical feasibility toward educational relevance, incorporating empirical observations from end users and institutional actors.

Acknowledgments

This work is part of the R&D&i Project PID2023-149674OB-I00, funded by MICIU/AEI/10.13039/501100011033 and ERDF, EU.

Declaration on Generative AI

During the preparation of this work, the author(s) used Gemini (Google) in order to: grammar and spelling check, paraphrase and reword. After using this tool/service, the author(s) reviewed and edited the content as needed and take(s) full responsibility for the publication's content.

References

[1] P. Sánchez, D. García-Saiz, M. Zorrilla, Software product line engineering for e-learning applications: A case study, in: 2012 International Symposium on Computers in Education (SIIE), 2012, pp. 1–6. URL: https://ieeexplore.ieee.org/abstract/document/6403212.

[2] Z. Chu, S. Wang, J. Xie, T. Zhu, Y. Yan, J. Ye, A. Zhong, X. Hu, J. Liang, P. S. Yu, Q. Wen, Llm agents for education: Advances and applications, 2026. URL: https://arxiv.org/abs/2503.11733. arXiv:2503.11733.

[3] O. Zawacki-Richter, V. I. Marín, M. Bond, F. Gouverneur, Systematic review of research on artificial intelligence applications in higher education – where are the educators?, International Journal of Educational Technology in Higher Education 16 (2019) 39. doi:10.1186/s41239-019-0171-0.

[4] D. Benavides, C. Sundermann, K. Feichtinger, J. A. Galindo, R. Rabiser, T. Thüm, Uvl: Feature modelling with the universal variability language, Journal of Systems and Software 225 (2025) 112326. doi:10.1016/j.jss.2024.112326.

[5] A. Matilla-Molina, J. M. Dodero, A. Muñoz, Integration of multi-agent systems and large language models for the creation of personalized and collaborative digital educational environments, in: Proceedings of the Learning Analytics Summer Institute Spain (LASI Spain 2025), volume 4148 of *CEUR Workshop Proceedings*, 2025. URL: https://ceur-ws.org/Vol-4148/Paper11.pdf.

[6] P. G. Balaji, D. Srinivasan, An Introduction to Multi-Agent Systems, 2010, pp. 1–27. doi:10.1007/978-3-642-14435-6_1.

[7] W. Holmes, M. Bialik, C. Fadel, Artificial Intelligence in Education. Promise and Implications for Teaching and Learning., 2019.

[8] A.-B. M. Salem, A. Y. Nikitaeva, Knowledge engineering paradigms for smart education and learning systems, in: 2019 42nd International Convention on Information and Communication Technology, Electronics and Microelectronics (MIPRO), IEEE, 2019, pp. 1571–1574. doi:10.23919/MIPRO.2019.8756685.

[9] G. Kyzylkaya, E. D. Torun, P. Askar, Restructuring e-learning with ontologies, in: 2007 International Conference on Computational Science and its Applications (ICCSA 2007), IEEE, 2007, pp. 161–164. doi:10.1109/ICCSA.2007.76.

[10] K. Klarin, I. Nazor, S. Celar, Ontology literature review as guidelines for improving croatian qualification framework, in: 2019 42nd International Convention on Information and Communication Technology, Electronics and Microelectronics (MIPRO), IEEE, 2019, pp. 1402–1407. doi:10.23919/MIPRO.2019.8757055.

[11] S. Chimalakonda, K. V. Nori, A family of software product lines in educational technologies, Computing 102 (2020) 1765–1792. doi:10.1007/s00607-019-00772-x.

[12] S. Chimalakonda, K. V. Nori, What makes it hard to apply software product lines to educational technologies?, in: 2013 4th International Workshop on Product LinE Approaches in Software Engineering (PLEASE), IEEE, 2013, pp. 17–20. doi:10.1109/PLEASE.2013.6608657.