Agent EDDIE – A Model-Agnostic AI Agent for Cross-LMS Content Migration, OER Integration and Didactic Quality Assurance

Dominik Giel[1], Sebastian Moser[1], Marcia Håkansson Lindqvist[2], Peter Mozelius[2]

[1]Offenburg University of Applied Sciences, Germany
[2] Mittuniversitetet, Sweden

Abstract

The integration of Artificial Intelligence in Higher Education is often limited to standalone tools, neglecting the structural complexity of Learning Management Systems (LMS). This paper introduces "Agent EDDIE" (Educational Didactic Digital Interchange Engine), a collaborative project between Mid Sweden University and Offenburg University of Applied Sciences. EDDIE is designed as an AI-agnostic agent capable of bridging the gap between LMS environments like Moodle and Canvas. Unlike provider-specific solutions, our architecture ensures strategic independence by decoupling the agent's logic from any specific Large Language Model (LLM). At the current stage of development, the agent has successfully demonstrated core functionalities: it autonomously authenticates within Moodle, validates course accessibility via unique IDs, generates a comprehensive inventory of student-visible elements, and utilizes LLM APIs to synthesize rudimentary course descriptions. However, the primary focus of this contribution is the strategic roadmap of the project. We outline the planned evolution of EDDIE through four critical milestones: first, the activity-by-activity technical transfer between Moodle and Canvas; second, the integration of automated linguistic translation; third, the integration of Open Educational Resources (OER) from external repositories to enhance existing courses, and finally, the implementation of a didactic audit tool. This tool will evaluate courses against abstract "studyability" rubrics, providing instructors with feedback based on pedagogical best practices. By presenting this work-in-progress, we aim to discuss the challenges of maintaining identical usability across systems and the importance of academic sovereignty in the age of AI.

Keywords: *AI-Ed, OER, LMS Migration, Digital Sovereignty, Moodle, Canvas*

1. Introduction

In the contemporary discussions on the challenges and opportunities of Artificial intelligence (AI) in education (AIED), the focus has often been on the importance of teacher training [1] [2], and how to address students AI cheating [3] [4]. Less has been published on the opportunities of using AI techniques to facilitate course management and assess the quality of course content. As pointed out by Zawacki-Richter et al. [5] AIED has been around for about more than 30 years, but it is still unclear how educators could make a pedagogical advantage of AI techniques on a broader scale, and how to achieve a meaningful actual impact on teaching and learning in higher education.

The integration of AI in Higher Education is often limited to standalone tools, neglecting the structural complexity of Learning Management Systems (LMS). Moreover, the concept of open educational resources (OER) has become prominent in technology enhanced learning [6]. At the same time, there have been research reporting on the need of quality assurance of both OER [7] [8], and entire courses in LMS environments [9] [10]. As argued in [11], the course quality can be further divided into factors such as course content quality, course relevance, course instructor quality, and course design quality.

To improve the quality of different existing LMS this paper suggests the design and development of Agent EDDIE (Educational Didactic Digital Interchange Engine). Agent Eddie is a model-agnostic AI agent who can assess the quality of LMS-courses according to variable definitions which is up to the user. Moreover, EDDIE should be able to suggest and implement improvements in different types of LMS.

Course review has long been a source of inspiration for producing new courses and developing and advancing existing courses. For teachers in higher education, the course reviews are sometimes conducted through peer review [13]. Faculty peer review of courses and course design can lead to improved and developed teaching practice, stronger teaching cultures, pedagogical exchange, as well as serving as additional information along with student evaluations [15], [16], [18]. Challenges include lack of time, faculty anxiety, unclear purposes, variable feedback quality, and lack of training for reviewers [13], [15]. This course review process can be compared to the peer review process in

publications, where there are similar challenges. Although peer review is seen as important, it can be biased, opaque, slow and harsh, with research calling for fairness, transparency, recognition, and training more open models in peer review [17]. The use of AI, or a hybrid solution, may offer alleviations to these challenges and provide new possibilities [14]. As in the publication process, teachers in higher education may appreciate the possibilities of AI course review as a faster, objective and transparent review than a review done by a colleague.

2. Method and Materials

This paper should be seen as the first part of a Design science research project comprising the five-folded process outlined by Johannesson and Perjons [20]. As presented and discussed in Humble and Mozelius [19], small scale studies, or parts of larger studies in academia can be carried out with selected parts of the full five-folded process. The first iteration of the initial phases described in this paper involves the three first phases of the process depicted below in Figure 1. below. Initially, the problem presented in this paper was identified and formulated by brainstorming and discussions among the authors during a staff training week. Later the problem explication, the defining of requirements and the beginning of the design phase have been conducted individually and with e-mail discussions among the authors. The focus has been on the first two phases, but also involves a first proof-of-concept of the design (prototype EDDIE V0.9, section 4).

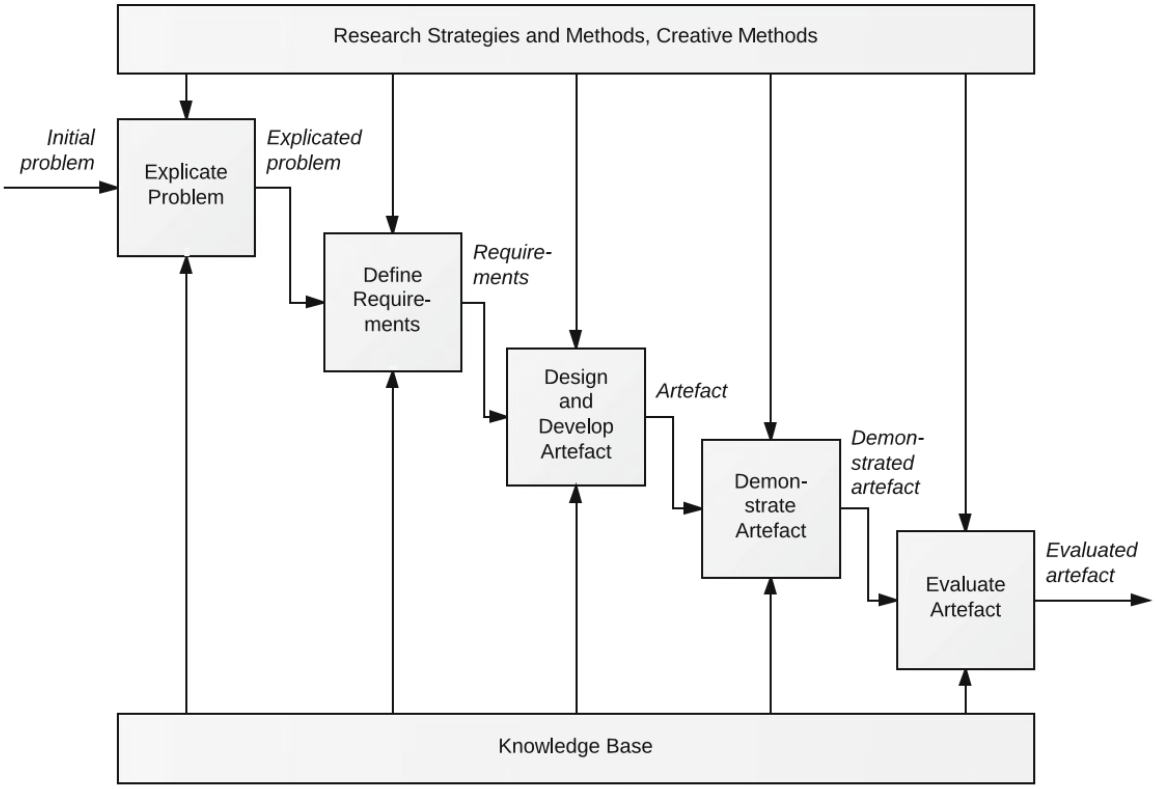

Fig. 1. The design science process as outlined by Johannesson and Perjons [12], p. 82

2.1 Current Situation of AI in Higher Education

Current research into the integration of AIED has been established for over 30 years, yet it remains primarily focused on profiling and prediction, adaptive systems, or intelligent tutoring systems [5]. While systematic reviews identify a vast array of potential applications, there is a noted lack of critical reflection on pedagogical implications and a "trust gap" regarding the delegation of course management to automated systems. The evaluation of didactical content through AI-mediated rubrics is an emerging pattern in the literature, though currently limited to sub-course granularity. Recent studies [26] have demonstrated the use of Natural Language Processing (NLP) to score classroom discussion transcripts against instructional quality rubrics. Similarly, projects have utilized GPT-4 to compare human-crafted

learning objectives against AI-generated ones based on criteria such as "SMART" alignment and course goal consistency. A systematic surveys of LLM Agents for Education [27] reports no existing project that combines course-activity-level rubrics with automated AI evaluation.

3. Concept

Faculty trust for granting access to a live LMS must be earned at the level of the agent's operating envelope, not at the level of its outputs alone. To help building trust to the concept, we have stated basic principles (P1-P5) for EDDIE:

3.1 Non-delete (P1)

EDDIE's first commitment is that no API call ever deletes, overwrites, or unpublishes existing material on the host LMS. Every action of the agent produces a new item, for example: an additional course, section, activity, or file whose existence is independent of the source course it derives from. EDDIE targets two platforms whose API surfaces and instructional models differ. On either of them, the design intent is the classical principle of least privilege of [21] applied to the agent's webservice token. On Moodle the token's scope is restricted to the `core_course_create` family and `core_files_upload`. On Canvas the equivalent OAuth scope is `manage_courses_add`. The corresponding delete and update functions against source-course resource handles are absent from the agent's token profile and therefore unreachable. The most defensive case of additive-only is no writes at all: The prototype (section 4) implements this case and grounds the principle in deployed code.

3.2 Invisible by Default (P2)

Every course that EDDIE creates begins in a non-discoverable state. On Moodle this is the `visible=0` in conjunction with `visibleold=0` flag-pair. On Canvas it is the equivalent `workflow_state="unpublished"` at course creation. The flag is set in the create call itself, not in a follow-up update. This is a deliberate choice that closes the otherwise small but real window in which the course exists in the public catalogue before being hidden, and is a residual risk addressed explicitly in the summary below. The LMS's own access-control layer enforces non-discoverability for invisible courses independently of any agent-side discipline.

3.3 Explicit Promotion (P3)

Switching the visible course pointer from a source course to its EDDIE-improved variant is a deliberate human action, never an agent action. A structured file (uploaded by the agent into the agent-produced course as an ordinary file resource) lists every change since the source: per activity, the agent records the source identifier, the agent-produced identifier, the type of semantic delta and a provenance bundle with the LLM provider identity, the system-prompt class, the timestamp, and the place where the instructor's acknowledgement signature is to be recorded. The manifest serves as the explanation surface aligned with the user-control guidelines and stands in analogy to the learning-to-defer framework of [22], where a learned classifier hands the decision to a human expert rather than committing to its own answer; in EDDIE the analogous boundary is promotion, the single decision that the agent never makes on its own. The provenance fields are modelled on the Content Credentials data model of the C2PA specification [23], with the practical intent that the per-activity audit log can be ingested by provenance-aware tooling outside EDDIE [24]. The same events are mirrored as xAPI statements [25] for analytics pipelines that already consume xAPI streams.

3.4 Reversibility (P4)

The source course is, at every moment in time, exactly the course the instructor authored. No agent-issued API call ever lands on its resource handles, and the principle of least privilege of [21] already invoked for P1 against deletions here serves a complementary role as it isolates the write surface. EDDIE's adapters allocate every agent-produced course as a separate `course_id` on either platform, and the agent never holds a write-handle on the source course. An instructor who deletes an agent-produced course therefore cannot affect the source. The host LMS's standard backup/restore workflow remains the institutional fallback, nothing about EDDIE's operation degrading it.

3.5 Open Source (P5)

The four preceding principles are technical commitments. The fifth is a deliberate boundary on those commitments: The EDDIE codebase will be published under a permissive MIT-class license whose warranty disclaimer is the legal complement of P1 - P4.
This means that institutional deployments are wrapped by the host institution's own risk acceptance and governance regime. Stating P5 here is not a defence against P1 - P4 but is the recognition that an architectural promise cannot replace an institutional contract, and that the two layers do different work.

3.6 Concept Summary

Three weaknesses survive the principles above and are stated explicitly here. First, data flows to the LLM provider as even a read-only agent forwards course content to an external service during course-description and matching steps. Our prototype Agent EDDIE V0.9 (section 4) uses an internal LLM proxy, which keeps the egress traffic inside the institutional perimeter, but data minimisation across providers remains an open work item. Second, the visibility race introduced under P2 is closed by setting the `visible=0` flag in the create API call itself, not in a follow-up update: but the soundness of this mitigation depends on the LMS's own atomicity guarantees, which neither Moodle nor Canvas documents formally. Third, the mis-promotion failure, since an instructor may click acknowledge without reading the manifest. UX guard rails as a required reading delay, a counter-signed-by line or a visualisation of the diff prior to the action are deferred to future work. In summary, defensive operation is necessary but not sufficient and is a necessary condition for an institution's consent to deploy Agent EDDIE.

4. Prototype

As a proof-of concept. a prototype of agent EDDIE with limited functionality has been build. The authors wish to thank Mohammad Dahi (HS Offenburg) for his implementation, EDDIE V0.9. It provides the most defensive instance of P1 as it has read-only access. The prototype observes a Moodle test instance, hands the observation to a controlled LLM endpoint, and renders the result as a per-module traffic-light view as a quality measure.
V0.9 is a four-step pipeline. A `core_course_get_courses` web-service call, authenticated by a single `wstoken`, returns the full course list with each entry's visible flag. The list is combined with a static module list in V0.9 hard-coded into the script, with module-handbook ingestion deferred to the first milestone (M1) discussed in Section 5. Both lists are forwarded to the HSO-internal LLM proxy (llm-proxy.imla.hs-offenburg.de, model gpt-5.1, temperature=0), which returns structured JSON mapping to each module to matched courses and their visibility status. The JSON is finally rendered as a tabular per-module traffic-light view (Fig. 2). The Moodle endpoint accessed by V0.9 is the institutional test instance, not production as the boundary between the two remains an institutional governance question.

	Modul	Kurs	Status
0	Mathematik 1	Mathe Gisela (G. Hillenbrand, Fakultät M, CeLT)	sichtbar
1	Physik Grundlagen	Physik Chatbot - RAG mit eBook-PDF (D. Giel, F...	sichtbar
2	Informatik Einführung	None	fehlt
3	Programmierung 1	None	fehlt
4	Statistik	Statistik HSOh! (Sokratischer Dialog) RAG (ext...	sichtbar
5	Statistik	Statistics HSOh! (Sokratischer Dialog) RAG-min...	sichtbar
6	Statistik	Statistik mit Python (Decker)	sichtbar
7	Software Engineering	Software-Engineering (C. Schmidt, V. Sänger, K...	sichtbar
8	Betriebspraktikum	Betriebspraktikum im WS 2022/23 + SS 2023 (AI,...	unsichtbar

Fig. 2. Sample output of EDDIE V0.9 provided by Mohammad Dahi (HS Offenburg)

EDDIE V0.9 was executed once by the implementer on 2026-04-30 against the HSO Moodle test instance. The inputs were a static list of nine study modules excerpted from a faculty handbook and the live course inventory of the test instance. Notice that there is a triple-match to "Statistik" with three corresponding courses (all with similar, but not identical titles) identified by the LLM's matching.

5. Future Work

EDDIE V0.9 is the first artifact created in order to build a fully-functioning AI agent, a proof-of-concept for the idea of using AI to automate certain (easy) tasks. We derived four milestones on the way which according to their complexity are:

5.1 Inter-System Course Transfer (M1)

At milestone M1, EDDIE is capable of transferring a basic LMS course from Moodle into Canvas and vice versa. Due to the open nature of the LMS (for example, novel plug-ins), the number of transferable items is limited and the transfer thus always partial. As pointed out earlier, EDDIE needs a rubric to evaluate course quality: The number of “transferable items” might be used (yet not improved by EDDIE) at this stage.

5.2 Inter-Language Course Transfer (M2)

At milestone M2, EDDIE can translate the language of LMS courses without altering the course items. As a quality rubric, one could pass part of the translated text to an LLM to evaluate if and how well the language matches a given particular language (for example, technical English). Then EDDIE could compare different system prompts to achieve an optimal for optimum translation quality. Translation itself could also be done using an LLM (not necessarily the same as for the quality rubrics).

5.3 OER Integration (M3)

At milestone M3, EDDIE is capable of analysing a course’s use of external learning resources (books, audio, video). As a quality metrics, it is capable of determining the appropriate abundance of external resources and estimate their respective license information. If material without or with unclear licence status is found, EDDIE offers to add an appropriate OER replacement or license. If OER material with insufficient license information is found, EDDIE offers to add the necessary information.

5.4 Didactic Quality (M4)

At the final milestone M4, EDDIE is capable of evaluating how good courses for a desired outcome, prior knowledge and student work-load look like. EDDIE can conduct an automated activity-level audit on courses running on a live LMS, suggest measures to improve didactic quality [9] and implement them (in hidden courses, according to P2 and P3).

6. Conclusion

Agent EDDIE represents a significant step toward achieving academic sovereignty in the integration of AI within Higher Education. By decoupling the logic of course auditing and migration from specific LLM providers, the project ensures that universities maintain control over their pedagogical data and instructional standards. The current V0.9 prototype successfully demonstrates the feasibility of automated course auditing via a model-agnostic proxy layer. Future milestones are inter-system migration and automated OER integration. EDDIE will provide educators with the tools to bridge the "trust gap" and ensure that digital learning environments remain both high-quality and pedagogically sound in an increasingly automated landscape.

Disclaimer and Acknowledgements

The creation of this paper was supported by the ChallengeEU project. The authors acknowledge that Artificial Intelligence was utilized during the preparation of this paper for the purpose of language improvement and linguistic refinement.

REFERENCES

[1] Fissore C., Floris F., Conte M. M., & Sacchet M., “Teacher training on artificial intelligence in education”, *Smart learning environments in the post pandemic era: Selected papers from the CELDA 2022 conference*, Cham, Springer Nature Switzerland, 2024, pp. 227-244.
[2] Bekdemir Y., “The urgency of AI integration in teacher training: Shaping the future of education”, *Journal of Research in Didactical Sciences*, 3(1), 2024, pp. 37-41.
[3] Xie Y., Wu S., & Chakravarty S., “AI meets AI: Artificial intelligence and academic integrity-A survey on mitigating AI-assisted cheating in computing education”, *Proceedings of the 24th annual conference on information technology education*, 2023, pp. 79-83.
[4] Bubas G., & Cizmesija A., “A Critical Analysis of Students' Cheating in Online Assessment in Higher Education: Post-COVID-19 Issues and Challenges Related to Conversational Artificial Intelligence”, *MIPRO*, 2023, pp. 905-910.
[5] Zawacki-Richter O., et al., “Systematic review of research on artificial intelligence applications in higher education”, *Int J Educ Technol High Educ*, 16, 2019, 39.
[6] Otto D., “Adoption and Diffusion of Open Educational Resources (OER) in Education”, *Open Education: International Perspectives in Higher Education*, 2021.
[7] Lopes A. P., Soares F., Softic S., & Andriichenko Y., “The Art of Quality Assessment for Open Educational Resources (OER)”, *International Conference on Interactive Collaborative Learning*, Cham, Springer Nature Switzerland, 2024, pp. 259-270.
[8] Putra F. P. E., Degeng I. N. S., Ulfa S., & Kamdi W., “The evolution of quality education: Impacts and challenges of using open educational resources (OER) and open educational practices (OEP) in the conceive-design-implement-operate (CDIO) framework”, *TEM Journal*, 13(1), 2024, 386.
[9] Crews T. B., & Wilkinson K., “Online quality course design vs. quality teaching: Aligning quality matters standards to principles for good teaching”, *The Journal of Research in Business Education*, 57(1), 2015, pp. 47-63.
[10] Karanfiloğlu M., & Bulut M. A., “Techno-pedagogical communication, ed-tech and media professionals: crossroads for enhancing instructional quality”, *Interactive Learning Environments*, 33(7), 2025, pp. 4193-4217.
[11] Ucha C. R., “Role of course relevance and course content quality in MOOCs acceptance and use”, *Computers and Education Open*, 5, 2023, 100147.
[12] Johannesson P., & Perjons E., “Design science”, *An introduction to Design Science*, Springer, 2014, 10.
[13] Cutroni L., & Paladino A., “Peer-ing in: A systematic review and framework of peer review of teaching in higher education”, *Teaching and Teacher Education*, 133, 2023, 104302.
[14] Farber S., “Comparing human and AI expertise in the academic peer review process: towards a hybrid approach”, *Higher Education Research & Development*, 44(4), 2025, pp. 871–885.
[15] Harrison R., Meyer L., Rawstorne P., Razee H., Chitkara U., Mears S., & Balasooriya C., “Evaluating and enhancing quality in higher education teaching practice: a meta-review”, *Studies in Higher Education*, 47(1), 2022, pp. 80–96.
[16] Thomas S., Chie Q. T., Abraham M., Jalarajan Raj S., & Beh L. S., “A qualitative review of literature on peer review of teaching in higher education: An application of the SWOT framework”, *Review of Educational Research*, 84(1), 2014, pp. 112-159.
[17] Kadaifci C., Isikli E., & Topcu Y. I., “Fundamental problems in the peer-review process and stakeholders' perceptions of potential suggestions for improvement”, *Learned Publishing*, 38(1), 2025, e1637.
[18] Fleckney P., Thompson J., & Vaz-Serra P., “Designing effective peer assessment processes in higher education: a systematic review”, *Higher Education Research & Development*, 44(2), 2025, pp. 386–401.
[19] Humble N., & Mozelius P., “Design science for small scale studies: Recommendations for undergraduates and junior researchers”, *European Conference on Research Methodology for Business and Management Studies (ECRM 2023)*, ACI Academic Conferences International, 2023, pp. 87-92.
[20] Johannesson P., & Perjons E., “A method framework for design science research”, *An introduction to design science*, Cham, Springer International Publishing, 2021, pp. 77-93.
[21] Saltzer J. H., & Schroeder M. D., “The protection of information in computer systems”, *Proceedings of the IEEE*, 63(9), 1975, pp. 1278-1308.

[22] Mozannar H., & Sontag D., “Consistent estimators for learning to defer to an expert”, *International Conference on Machine Learning (ICML)*, 2020.
[23] C2PA, “C2PA Specification v1.0”, *Coalition for Content Provenance and Authenticity*, 2026.
[24] Bakharia A., Corrin L., de Barba P., Kennedy G., Gašević D., Mulder R., Williams D., Dawson S., & Lockyer L., “A conceptual framework for linking learning design with learning analytics”, *Proceedings of the Sixth International Conference on Learning Analytics & Knowledge*, 2016, pp. 329-333.
[25] van den Berg A., de Wit M., & van der Klis M., “The Dutch xAPI reference profile”, *SURF*, 2016.
[26] Tran, N., Pierce, B., Litman, D., Correnti, R., & Matsumura, L. C. (2023). Utilizing natural language processing for automated assessment of classroom discussion. In International Conference on Artificial Intelligence in Education (pp. 490-496). Cham: Springer Nature Switzerland.
[27] Chu, Zhendong, Shen Wang, Jian Xie, Tinghui Zhu, Yibo Yan, Jinheng Ye, Aoxiao Zhong, Xuming, Hu, Jing Liang, Philip S. Yu, and Qingsong Wen (2025). “LLM Agents for Education: Advances and Applications”, *Findings of the Association for Computational Linguistics: EMNLP 2025* 10.18653/v1/2025.findings-emnlp.743