Toward Pedagogically-Guided Large Language Models: the Role of Knowledge Components and SOLO in Enhancing Smart Learning Environments*

Zeynab (Artemis) Mohseni†, Juan I. Asensio-Pérez, Miguel L. Bote-Lorenzo, Eduardo Gómez-Sánchez, Yannis Dimitriadis

GSIC-EMIC Research Group, University of Valladolid, Valladolid, Spain

Abstract

Smart Learning Environments (SLEs) aim to provide adaptive, context-aware support across hybrid learning settings, yet they still face numerous challenges that limit their pedagogical effectiveness. Among these, this paper focuses on three persistent challenges: 1) the limited alignment between Learning Design (LD) and Learning Analytics (LA); 2) student models that rely primarily on structured data rather than conceptual and cognitive evidence; and 3) constrained forms of adaptive support that do not sufficiently respond to learners' reasoning processes or real-world contexts. Recent advances in Large Language Models (LLMs) offer new opportunities to address these limitations, but only when their generative capacity is steered by pedagogically meaningful structures. This paper explains the theoretical rationale for using Knowledge Components (KCs) and the Structure of Observed Learning Outcome (SOLO) taxonomy to enhance LLM interpretation of learning data and generation of personalised and context-aware learning tasks in SLEs. KCs help LLMs identify conceptual targets in learner inputs, while SOLO captures the quality of the learner's reasoning. Used together, they support meaningful links between the pedagogical intentions of teachers (reflected in a learning design), student modelling, and personalised and context-aware reactions in SLEs. Contrary to the usual case of traditional Intelligent Tutoring Systems (ITS), KCs are not used here as elements of an expert-driven, pre-defined knowledge model, but as the result of decomposing teachers' learning objectives into assessable units that can drive the behaviour of LLMs in SLEs. The goal of this paper is not to introduce a new approach, but to articulate why combining KC and SOLO guidance is a promising direction for enhancing SLEs, that make use of LLMs, in theoretically sound, teacher-aligned ways.

Keywords

Learning Analytics, Learning Design, SOLO taxonomy, Knowledge Components, Large Language Model, SLE

1. Introduction

Smart Learning Environments (SLEs) have long aimed to provide continuous support across physical and digital contexts through cycles of sensing, analysing, and reacting to learner behaviour [1]. Although multiple challenges affect their effectiveness, three persistent issues remain central: 1) the weak alignment between Learning Design (LD) and Learning Analytics (LA) [2]; 2) the narrow reliance on structured data that obscures conceptual and cognitive evidence [3]; and 3) the difficulty of generating adaptive, context-sensitive pedagogical reactions grounded in learners' reasoning processes [4]. Recent advances in Generative AI in general, and Large Language Models (LLMs) in particular, have renewed interest in addressing these limitations [4]. LLMs can interpret explanatory or reflective learner responses, classify open-ended reasoning, and support teachers during LD [5, 6]. Nevertheless, when unguided, their outputs often remain pedagogically shallow, inconsistent, or

* *LA4GenAISEL'26: Learning Analytics for Human-Centered Generative AI-Enhanced Smart Learning Environments Workshop, LAK26, April 28, 2026, Bergen, Norway.*

† Corresponding author.

zmohseni@uva.es (Z. Mohseni); juaase@tel.uva.es (J. I. Asensio-Pérez); migbot@tel.uva.es (M. L. Bote-Lorenzo); edugom@tel.uva.es (E. Gómez-Sánchez); yannis@tel.uva.es (Y. Dimitriadis)

0000-0002-3297-0189 (Z. Mohseni); 0000-0002-1114-2819 (J. I. Asensio-Pérez); 0000-0002-8825-0412 (M. L. Bote-Lorenzo); 0000-0003-0062-916X (E. Gómez-Sánchez); 0000-0001-7275-2242 (Y. Dimitriadis)

disconnected from pedagogical intentions [5, 7]. Their linguistic fluency does not guarantee educational coherence.

This motivates the need for pedagogically meaningful structures capable of orienting LLM reasoning. In this paper, two are emphasized: Knowledge Components (KCs) and the Structure of Observed Learning Outcome (SOLO) taxonomy [7-9]. Crucially, the KCs used here differ from the traditional Intelligent Tutoring Systems (ITS) lineage that relies on predefined fine-grained skills or production rules [9-11]. This distinction concerns the purpose they serve rather than any inherent comparison of quality. Whereas predefined KCs in ITSs support domain-specific knowledge tracing, in our context they are used to represent the conceptual targets embedded in teachers' learning objectives so that LLMs can interpret learner inputs in pedagogically aligned ways.

Recent studies show that KCs can be obtained through multiple methods, including LLM-based extraction and data-driven discovery: LLMs can generate KCs that are semantically coherent and interpretable for teachers, though still requiring oversight, and automatically learned KCs may reflect blended or emergent cognitive patterns rather than isolated micro-skills [12-14]. Together, these findings illustrate that KCs being appropriate for hybrid and open-ended learning contexts do not need to mirror the granularity of ITS-style skill decompositions. Further, data-driven approaches reveal that learned KCs often reflect blended or emergent cognitive patterns rather than isolated micro-skills [14], reinforcing that KCs do not need to mirror the granularity of ITS-style skill decompositions. This paper focuses on the theoretical justification for using KC and SOLO guidance to enhance SLEs. By combining concept-level interpretation with reasoning-level insight, we hypothesize that LLM-supported SLEs can better align LA with LD, construct richer student models, and produce adaptive reactions that are pedagogically sound and contextually meaningful.

2. Why Focus on Knowledge Components and SOLO?

2.1. Knowledge Components for Conceptual Clarity

KCs were defined in foundational work by Koedinger et al. [9] as "an acquired unit of cognitive function or structure that can be inferred from performance on a set of related tasks". This definition situates KCs as fine-grained cognitive units that reveal what learners must know or be able to do in order to perform a task successfully. Subsequent work expands this view by noting that KCs may represent facts, concepts, procedures, misconceptions, or skills, thereby encompassing a broad range of conceptual elements relevant for student modelling [15, 16].

In learning sciences research, KCs have traditionally been defined as the inferable cognitive structures activated during task performance [17]. Within ITS, these units typically appear as stable, fine-grained skills or production rules [10, 11]. However, recent findings show that KCs generated by LLMs or data-driven models differ meaningfully from this tradition. For instance, Moore et al. [12] demonstrate that LLM-generated KC labels derived from domain materials are often clearer and more intuitive to human evaluators than expert-authored KCs. Similarly, Sampaio de Alencar et al. [13] show that automatically derived KCs can be meaningfully clustered and reorganized to support teacher interpretation. Shi et al. [14] further demonstrate that KCs can emerge as latent patterns in learner behaviour, capturing blended or composite reasoning tendencies not easily expressible as isolated ITS-style skills. Together, these findings indicate that KCs suited to hybrid, open-ended learning tasks must emphasize conceptual interpretability rather than strict micro-skill decomposition. The KCs used here therefore represent the conceptual targets embedded in learning materials such as ideas, relationships, distinctions, or principles that learners must activate. This makes them pedagogically appropriate anchors for orienting LLMs in environments where learners express understanding through unstructured, situated, or reflective tasks.

Integrating Koedinger et al.'s [9] definition with recent LLM-based work highlights that the role of KCs has broadened: while they remain units of cognitive structure, they are no longer restricted to predefined curricular elements. This expansion reflects not only the variety of sources from which KCs may be obtained but also the different methods used to derive them, ranging from traditional expert specification to matrix-factorization and clustering approaches, as well as LLM-based extraction from instructional materials [12-14]. As explained in previous studies, KCs may also emerge from patterns in learner behaviour [14] or from automated analyses of instructional materials [12, 13]. This expanded view preserves their inferable cognitive nature while allowing them to reflect the conceptual granularity actually present in learning tasks. This evolution positions KCs as the most viable representational layer for aligning pedagogical intentions with LLM-supported interpretation [9, 18]. Because KCs make explicit the conceptual distinctions embedded in teacher-provided materials—including not only instructional documents but also elements of the learning design such as learning objectives and task descriptions—they provide a tractable and pedagogically meaningful structure for LLMs to recognize within learner responses [12]. At the same time, their flexibility accommodates the heterogeneous, sometimes composite cognitive patterns identified in contemporary data-driven research, ensuring that they remain interpretable even as their origins shift from expert specification to semi-automated derivation [13, 14, 19].

2.2. SOLO for Characterizing Reasoning Quality

The SOLO taxonomy classifies learner responses according to the structural organization of ideas, offering a way to evaluate the quality of understanding rather than the type of task or the action verbs associated with learning objectives [8, 20, 21]. SOLO describes a progression from pre-structural responses, in which no meaningful engagement with the task is evident, to uni-structural and multi-structural responses that introduce one or several ideas, and further to relational and extended abstract responses, where those ideas are integrated or generalized. This contrasts with taxonomies such as Bloom's revised taxonomy, Depth of Knowledge, or Marzano's New Taxonomy, which classify cognition based primarily on action types or task complexity rather than the structure of the learner's explanation [22, 23].

The literature emphasizes that many established taxonomies are limited when analysing unstructured learner outputs, because they focus on cognitive actions or task demands rather than on the internal structure of the learner's response [22, 23]. SOLO addresses this limitation by evaluating the quality and organization of the ideas expressed by the learner rather than the task itself [20, 21]. SOLO's developmental sequence (from isolated ideas to integrated and generalized understanding) has been repeatedly described in the literature, and this progression is central to its suitability for analysing learning outcomes [20, 21]. This makes SOLO particularly relevant for SLEs that need to interpret learner-generated content [22, 23]. Recent evidence demonstrates that LLMs can reliably classify learner responses according to SOLO levels, capturing distinctions in reasoning quality beyond correctness or keyword matching [6]. This compatibility with LLM-based analysis strengthens the case for using SOLO as a guiding structure in SLEs, since it provides a taxonomy that can be operationalized for evaluating open-ended or explanatory responses. Overall, SOLO offers a well-established approach for representing the structural complexity of learner understanding in a way that aligns with the needs of SLEs, especially when these environments rely on generative models to interpret learners' explanations or reflections [20-23].

2.3. Complementarity

KCs and SOLO serve different but mutually reinforcing roles. KCs identify the conceptual content present in a learner's response [9, 16, 18], while SOLO captures the structural complexity of the reasoning [20, 21]. Used together, they provide a balanced perspective that avoids the limitations of either approach alone. This dual perspective allows LLMs to interpret open-ended responses in ways that remain precise, pedagogically aligned, and sensitive to developmental progression. Because KCs and SOLO capture different dimensions of performance, bringing them together enables interpretations that are both conceptually specific and cognitively meaningful. This complementarity is reinforced by findings showing that KCs can vary in granularity and may represent blended or emergent patterns rather than predefined skills [13, 14]. SOLO, by focusing on the structural complexity of the response itself, offers a way to interpret such heterogeneous conceptual units without requiring them to conform to a fixed skill hierarchy [20, 21]. Thus, while KCs reveal which conceptual elements appear in a learner's explanation, SOLO provides a lens for interpreting how these elements are expressed and interconnected. The combination is further supported by research showing that learning objectives and competences often aggregate multiple skills, concepts, or knowledge components [1, 24-26]. When KCs and SOLO are used together, they help make these broader intentions analytically tractable: KCs offer insight into which parts of the intended conceptual landscape learners engage with, while SOLO clarifies the developmental sophistication with which these engagements occur. As a result, the two constructs jointly support the kind of pedagogically aligned interpretation that SLEs require but often lack. KCs ensure that interpretations remain tied to the intended conceptual content, while SOLO ensures that learners' cognitive performance is evaluated in terms of structural quality rather than surface correctness. These complementary insights allow LLMs to generate analyses and reactions that more faithfully reflect pedagogical goals, particularly in hybrid and open-ended learning contexts where conceptual variability and reasoning differences are both central to understanding student progress.

3. Why KC and SOLO-guided LLMs Can Enhance SLEs

Analytic models in many SLEs fail to reflect the pedagogical intentions embedded in LD [1, 4]. Learning objectives are typically too broad to be operationalized computationally [27]. In contrast, KCs derived from teacher-provided materials offer conceptual units that can be recognized by LLMs. When paired with SOLO-aligned examples, these units make analytics sensitive not only to the learner's conceptual focus but also to their reasoning quality [20, 21]. This connection helps reduce the longstanding disconnect between design and interpretation.

Traditional student models in SLEs rely heavily on structured traces such as quiz scores or interaction logs [28]. These overlook the conceptual and cognitive richness of learners' open-ended responses. KC-guided interpretation allows LLMs to identify which conceptual targets appear in a student's explanation, while SOLO-guided classification captures developmental differences in how ideas are organized. Combined, they allow student models to represent conceptual mastery patterns and reasoning trajectories that were previously invisible [15, 16, 29, 30]. The feasibility of using LLMs to detect such reasoning patterns has also been demonstrated in recent work using SOLO-aligned classification of student responses [6]. Existing SLE reactions are often rule-based or generic, lacking sensitivity to learners' conceptual understanding or cognitive depth [1, 31]. When LLMs generate adaptive support anchored in KCs and SOLO, reactions can reflect both what the learner is addressing and how deeply it is understood. This enables more diverse and pedagogically meaningful feedback, including reflective prompts, conceptual hints, or follow-up tasks aligned with SOLO levels [8]. Because SLEs frequently integrate contextual data such as time, weather, geolocation, or environmental conditions [26, 32], LLM-generated reactions can be adapted to real-world conditions, supporting feasible and situated hybrid learning experiences. Concerns about AI reducing teacher control are well known [33]. KC and SOLO structures mitigate this risk by providing interpretable intermediaries that teachers can inspect, merge, refine, or reject. Since the KCs used here can originate from teacher materials rather than from predefined skill lists (as is typically the case of ITSs), they can remain more closely aligned with pedagogical intentions while still enabling LLM-assisted personalisation and context-aware adaptation [24, 25].

4. Conclusion

This paper has argued that combining KCs and the SOLO taxonomy offers a theoretically grounded way to guide LLMs in SLEs. The KCs adopted here are not the predefined micro-skills of ITS research; instead, they reflect interpretable conceptual units derived from teacher materials or from LLM-assisted extraction, as shown in studies demonstrating the pedagogical advantages of LLM-generated or data-driven knowledge units [12-14]. SOLO complements these units by capturing the structural quality of learner reasoning, providing essential cognitive insight for interpreting open-ended responses [8, 20, 21]. By making conceptual targets and reasoning structures explicit, KC- and SOLO-guided LLMs offer a promising direction for addressing three persistent SLE challenges: misalignment between LD and LA, limited student models, and insufficiently adaptive reactions [2, 4]. Their combination enables more meaningful personalization that remains transparent, situated, and aligned with teacher intentions.

Acknowledgements

This research was partially supported by grant PID2023-146692OB-C32 funded by MICIU/AEI/10.13039/501100011033 and by ERDF/EU as well as by grant VA176P23 funded by the Regional Government of Castile and Leon together with the ERDF.

Declaration on Generative AI

The authors declare that no generative AI tools were used for the writing, editing, or preparation of this manuscript.

References

[1] García-Tudela, Pedro Antonio, Paz Prendes-Espinosa, and Isabel María Solano-Fernández. "Smart learning environments: a basic research towards the definition of a practical model." Smart Learning Environments 8.1 (2021): 9.

[2] Mangaroska, Katerina, and Michail Giannakos. "Learning analytics for learning design: A systematic literature review of analytics-driven design to enhance learning." IEEE Transactions on Learning Technologies 12.4 (2018): 516-534.

[3] Wu, Siyu, et al. "A comprehensive exploration of personalized learning in smart education: From student modeling to personalized recommendations." arXiv preprint arXiv:2402.01666 (2024).

[4] C. Delgado Kloos, J. I. Asensio Pérez, D. Hernández Leo, P. M. Moreno-Marcos, M. L. Bote Lorenzo, P. Santos, ... and B. Tabuenca, GENIE learn: human-centered generative AI-enhanced smart learning environments, in: Proceedings of the 17th International Conference on Computer Supported Education (CSEDU), V. 2, 2025, pp. 15-26.

[5] Giannakos, Michail, et al. "The promise and challenges of generative AI in education." Behaviour & Information Technology 44.11 (2025): 2518-2544.

[6] S. Zhang, P.S. Meshram, P. Ganapathy Prasad, M. Israel, and S. Bhat, An LLM-Based Framework for Simulating, Classifying, and Correcting Students' Programming Knowledge with the SOLO Taxonomy, in: Proceedings of the 56th ACM Technical Symposium on Computer Science Education V. 2, Pittsburgh PA USA, Feb. 2025, pp. 1681-1682.

[7] R. Niousha, A. O'Neill, E. Chen, V. Malhotra, B. Akram, and N. Norouzi, LLM-KCI: Leveraging Large Language Models to Identify Programming Knowledge Components, in: Proceedings of the 56th ACM Technical Symposium on Computer Science Education V. 2, Feb. 2025, pp. 1557-1558.

[8] Biggs, John B., and Kevin F. Collis. Evaluating the quality of learning: The SOLO taxonomy (Structure of the Observed Learning Outcome). Academic press, 2014.

[9] Koedinger, Kenneth R., Albert T. Corbett, and Charles Perfetti. "The Knowledge-Learning-Instruction framework: Bridging the science-practice chasm to enhance robust student learning." Cognitive Science 36.5 (2012): 757-798.

[10] Corbett, Albert T., and John R. Anderson. "Knowledge tracing: Modeling the acquisition of procedural knowledge." User modeling and user-adapted interaction 4.4 (1994): 253-278.

[11] H. Nguyen, Y. Wang, J. Stamper, and B. M. McLaren, Using Knowledge Component Modeling to Increase Domain Understanding in a Digital Learning Game, in: International Conference on Educational Data Mining (EDM), 12th, Montreal, Canada, International Educational Data Mining Society, Jul 2-5, 2019.

[12] S. Moore, R. Schmucker, T. Mitchell, and J. Stamper, Automated generation and tagging of knowledge components from multiple-choice questions, in: Proceedings of the eleventh ACM conference on learning@ scale, July 2024, pp. 122-133.
[13] R. Sampaio de Alencar, M. A. Demirtas, A. S. Saha, Y. Shi, and P. Brusilovsky, Integrating Expert Knowledge with Automated Knowledge Component Extraction for Student Modeling, in: Proceedings of the 33rd ACM Conference on User Modeling, Adaptation and Personalization, June 2025, pp. 307-312.
[14] Y. Shi, R. Schmucker, M. Chi, T. Barnes, and T. Price, KC-Finder: Automated Knowledge Component Discovery for Programming Problems, in: In M. Feng, T. K¨aser, and P. Talukdar, editors, Proceedings of the 16th International Conference on Educational Data Mining, Bengaluru, India, July 2023, pp. 28–39.
[15] Duan, Zhangqi, et al. "Automated Knowledge Component Generation and Knowledge Tracing for Coding Problems." arXiv preprint arXiv:2502.18632 (2025).
[16] Wei, Yumou, Paulo Carvalho, and John Stamper. "KCluster: An LLM-based Clustering Approach to Knowledge Component Discovery." arXiv preprint arXiv:2505.06469 (2025).
[17] Pelánek, Radek. "Managing items and knowledge components: domain modeling in practice." Educational Technology Research and Development 68.1 (2020): 529-550.
[18] N. Luburić, B. Šarenac, L. Dorić, D. Vidaković, K. G. Grujić, A. Kovačević, and S. Prokić, Challenges of knowledge component modeling: A software engineering case study, in: 8th International Conference on Higher Education Advances (HEAd'22), May 2022, pp. 901-908.
[19] Wang, Canwen, Jionghao Lin, and Kenneth R. Koedinger. "Leveraging Large Language Models for Identifying Knowledge Components." arXiv preprint arXiv:2511.09935 (2025).
[20] Adeniji, Saidat, Penelope Baker, and Martin Schmude. "Structure of the Observed Learning Outcomes (SOLO) model: A mixed-method systematic review of research in mathematics education." EURASIA Journal of Mathematics, Science and Technology Education 18 (2022): 1.
[21] Chan, Charles C., et al. "Applying the structure of the observed learning outcomes (SOLO) taxonomy on student's learning outcomes: An empirical study." Assessment & Evaluation in Higher Education 27.6 (2002): 511-527.
[22] AlAfnan, Mohammad Awad, and Mohammad Al. "Taxonomy of educational objectives: Teaching, learning, and assessing in the information and artificial intelligence era." Journal of Curriculum and Teaching 13.4 (2024): 173-191.
[23] Irvine, Jeff. "Taxonomies in education: Overview, comparison, and future directions." Journal of Education and Development 5.2 (2021): 1.
[24] Heller, Jürgen, et al. "Competence-based knowledge structures for personalised learning." International Journal on E-learning. Vol. 5. No. 1. Association for the Advancement of Computing in Education (AACE), 2006.
[25] Lu, Hsiu-Lien, and Hsiao-Fang Lin. "A concept model of competency tasks in competency-based education." Technology, Pedagogy and Education (2025): 1-19.
[26] Tabuenca, Bernardo, et al. "Greening smart learning environments with Artificial Intelligence of Things." Internet of Things 25 (2024): 101051.
[27] Yang, Mei, and Jiatao Li. "Under the Digital Background Marzano's New Taxonomy of Educational Objectives: Conception of Deep Learning Analysis." Fuzzy Systems and Data Mining IX. IOS Press, 2023. 867-876.
[28] Jiménez Macías, Alberto Alejandro, et al. "Content modeling in smart learning environments: A systematic literature review." (2024).
[29] Moon, Hyeongdon, et al. "Using large multimodal models to extract knowledge components for knowledge tracing from multimedia question information." arXiv preprint arXiv:2409.20167 (2024).
[30] R. Niousha, M. Hoq, B. Akram, and N. Norouzi, Use of large language models for extracting knowledge components in cs1 programming exercises, in: Proceedings of the 55th ACM Technical Symposium on Computer Science Education V. 2, March 2024, pp. 1762-1763.

[31] Chang, Maiga, and Yanyan Li, eds. Smart learning environments. Springer Berlin Heidelberg, 2015.
[32] Mouri, Kousuke, and Hiroaki Ogata. "Ubiquitous learning analytics in the real-world language learning." Smart Learning Environments 2.1 (2015): 15.
[33] Dron, Jon. "Smart learning environments, and not so smart learning environments: a systems view." Smart learning environments 5.1 (2018): 25.