From Practice to Nudge: A Hybrid Intelligence Framework for Instructional Decision Support[⋆]

Alex Liu[1,*], Shawon Sarkar[1], Lief Esbenshade[1], Victor Tian[1], Kevin He[2], Zachary Zhang[2] and Min Sun[1]

[1]*University of Washington*
[2]*Hensun Innovation, LLC*

Abstract

This study introduces a practice-informed AI nudging system designed to support responsible, context-aware instructional decision-making during lesson planning. Grounded in a hybrid intelligence framework, the system integrates multiple sources of human insight: individual teacher preferences, collective behavior patterns, and pedagogical research, with computational support from large language models to generate timely and explainable nudges. A co-occurrence network of instructional strategies, derived from annotated teacher-AI interactions, guides contextual nudge generation using behavioral heuristics. Initial analysis shows strong alignment between system recommendations and teacher' pedagogical goals and values, as well as increased diversity in effective instructional practices. By balancing professional autonomy with evidence-based guidance, this work advances responsible AI in education, contributing an explainable, value-sensitive system that fosters sustainable pedagogical improvement in real-world settings.

Keywords

Digital nudging, hybrid intelligence, context-aware decision support, teacher-AI interaction, value-sensitive design

1. Introduction

Lesson planning is a central activity through which teachers translate pedagogical and content knowledge into classroom instruction [1, 2]. This process involves sequencing and adapting activities to engage diverse learners, drawing on both expertise and instructional judgment [3]. More than routine preparation, it serves as embedded professional learning, requiring integration of content, pedagogy, and learner needs [4]. Without sufficient support, this cognitively demanding task may yield suboptimal instruction, highlighting the need for real-time, context-aware decision support. Traditional professional development often falls short: it is difficult to access, insufficiently personalized, and disconnected from day-to-day planning [5]. Even when training is available, instructional planning remains time-intensive and complex [6]. Recent advances in artificial intelligence (AI) present an opportunity to address these gaps by supporting teachers' instructional decision-making. However, without thoughtful design, AI-generated recommendations risk falling short in both effectiveness and trustworthiness [7].

Nudging, rooted in behavioral science, refers to subtle prompts that influence decisions without limiting choice [8]. In education, teachers, like all individuals, are subject to cognitive constraints from limited time, information, and energy [9, 10]. Nudges can reduce instructional inertia by gently steering teachers toward better practices while preserving autonomy [11]. When embedded in AI systems, nudges can be dynamically tailored to a teacher's context, supporting decisions around content, pedagogy, and classroom management [12]. Their potential lies in leveraging real-time data to generate

HHAI-WS 2025: Workshops at the Fourth International Conference on Hybrid Human-Artificial Intelligence (HHAI), June 9-13, 2025, Pisa, Italy
*Corresponding author.
alexliux@uw.edu (A. Liu); ss288@uw.edu (S. Sarkar); lief@uw.edu (L. Esbenshade); ztian27@uw.edu (V. Tian); kevin@colleague.ai (K. He); zac@colleague.ai (Z. Zhang); misun@uw.edu (M. Sun)
0000-0002-4785-1801 (A. Liu); 0000-0003-2551-211X (S. Sarkar); 0000-0003-1549-9337 (L. Esbenshade); 0009-0002-8263-0753 (V. Tian); 0000-0001-5832-1534 (M. Sun)

timely, personalized prompts [13], making decision support itself a form of embedded professional development.

While research on AI-powered nudging in instructional preparation is limited, studies in classroom orchestration show that real-time feedback can enhance instructional flow and student engagement [14, 15]. These findings underscore nudging's potential in complex, high-stakes environments. However, teachers vary in experience, pedagogical philosophy, and technological fluency, requiring systems that not only model behavior but also interpret goals and preferences. Ignoring these differences risks eroding trust and reinforcing a one-size-fits-all approach.

Equally important are the ethical stakes. Without transparency and alignment with teacher values, nudges risk being perceived as manipulative, raising concerns about agency and power imbalances in AI-mediated decision-making [16, 17, 13]. If teachers feel their expertise is undermined, they may resist adoption, especially in autonomy-centered environments [18]. To be trusted, nudging systems must preserve professional judgment, communicate their rationale clearly, and act as supportive partners rather than prescriptive agents.

This underscores the role of human-centered design (HCD) in building AI that complements rather than constrains teacher decision-making. HCD ensures systems are not only technically effective, but also aligned with user values, workflows, and autonomy [19]. In education, where complex judgment is central, AI systems must respect expertise and avoid overriding intent [20]. Nudges must be transparent and explainable [21], and personalization must go beyond surface tailoring to align with evolving pedagogical needs, classroom contexts, and teacher goals [22].

2. Design Tensions: Adoption vs. Effectiveness

A central challenge in designing AI-powered nudges for educators lies in balancing two potentially competing objectives: *alignment with existing practices and values versus introduction of effective but unfamiliar strategies.* On one hand, nudges must resonate with teachers' current routines and pedagogical orientations to maintain trust, foster engagement, and uphold professional autonomy. On the other hand, to meaningfully support instructional improvement, nudges must also surface research-backed practices that teachers may not typically consider—strategies that go beyond existing habits while remaining contextually appropriate.

To address the tension between respecting teacher autonomy and promoting instructional improvement, we propose an AI-powered nudging system grounded in a **hybrid intelligence framework**. This approach integrates four complementary sources of insight: (1) individual teacher preferences captured through usage patterns, (2) collective behavioral trends derived from educator interactions on the platform, (3) pedagogical best practices from instructional research, and (4) the reasoning and generative capabilities of LLMs. By integrating multiple sources of human and LLM contributions, the system can enhance the contextual relevance and pedagogical validity of its recommendations, encouraging sustained instructional growth without enforcing prescriptive behavior change.

To evaluate this system, we investigate the following research questions:

- RQ1: To what extent do the system-generated nudges align with the pedagogical goals, values, and instructional intent expressed by teachers during interaction with the platform?
- RQ2: Do nudges lead teachers to adopt a more diverse set of high-leverage instructional practices than they would typically select on their own?
- RQ3: Does the hybrid intelligent framework affect the explainability, contextual alignment, and adoption of nudges across varied educational contexts?

3. 3. Hybrid Intelligence Framework

3.1. Data Source

Our analysis draws on de-identified chat history logs from an AI-powered instructional assistant used by K-12 educators across grade levels, subject areas, and various instructional contexts. These interactions reflect authentic teacher goals, instructional challenges, and pedagogical reasoning during lesson planning.

3.2. Identification and Qualitative Annotation

We conducted an extensive literature review on effective classroom instructional approaches. Drawing from established research in teacher education and instructional design frameworks, we developed a codebook of high-leverage instructional strategies, including opportunities for student discourse, scaffolding techniques, the use of meaningful real-world examples, and formative assessment [23].

Using this codebook, we applied automated qualitative coding to teacher-AI dialogues to identify instances of these instructional strategies, along with contextual information such as subject area, educational level, and learner needs (e.g., English language learners, students with special education needs, neurodiverse students). Following automated coding by an LLM, experienced qualitative researchers manually verified and refined the results to ensure accuracy and interpretive validity. This process surfaced both explicit instructional moves and implicit pedagogical values, preferences, and intentions embedded within teacher queries and interactions.

3.3. Practice Network Construction

From the annotated data, we constructed a directed **co-occurrence network** that models how instructional practices tend to be sequenced across similar user contexts. The resulting network captures both collective educator behavior (as it emerges from platform usage) and transitions grounded in pedagogical research. This structure enables us to generate recommendations that *reflect both the normative behaviors of peers and evidence-based instructional trajectories.*

To enhance personalization and contextual fidelity, we also generate a **user-specific preference network** for each teacher with sufficient interaction history on the platform. This individualized layer captures the teacher's own prior behaviors and choices, demonstrating commonly selected instructional strategies and frequent content goals. After implementation of this pipeline, we will further incorporate individuals' responsiveness to previous nudges to this user preference network. These personal usage patterns are modeled using a weighted subgraph derived from the global co-occurrence network, giving additional weight to pathways the user has historically followed. Both networks will be updated periodically to reflect the recent educators behavior on the platform.

During nudge generation, the system balances the collective network (peer norms and research-based best practices) with this personalized subgraph. This **compound network architecture** allows the system to recommend actions that are both familiar and growth-oriented: they reflect the teacher's preferences while gently nudging toward more effective or underused strategies observed in the broader community. However, the optimal weight to balance these two networks remains to be investigate.

3.4. Nudging Mechanism

At runtime, the system incorporates LLM tagging the teacher's current instructional step using established practice codebook and identifying educational context demonstrated from teachers' input. The compound practice network is then queried to predict the most likely and pedagogically meaningful next actions to nudge. Nudges are selected based on contextual similarity, using the teacher's subject, educational levels, and learner needs, and delivered using behavioral heuristics such as status quo, social proof, and peer effect.

Nudges are phrased to be suggestive rather than prescriptive, maintaining the teacher's sense of agency while introducing effective alternatives grounded in both literature and peer activity. For example, if a teacher exploring inquiry-based science methods, the system may suggest a transition to seeking real-world applications, a practice used frequently by other teachers in similar contexts and recommended by previous studies in science education [24].

4. Planned Analysis

Following implementation of the nudging pipeline within the instructional platform, we will evaluate its effectiveness through a mixed-methods strategy. To do so, users will be randomly assigned to one of three groups: a hybrid intelligence group (receiving nudges generated from the practice-informed network), an LLM-inference group (receiving LLM inferred recommendations), and a control group (receiving generic or no nudges). Table 1 summarizes the analysis plan for each research question, integrating controlled experiments like A/B testing [12], platform usage analysis, and qualitative feedback.

Table 1
Overview of the planned analysis strategy

Research Question	Analysis Approach
RQ1: Alignment with Pedagogical Goals and Values	Conduct A/B testing comparing intervention and control groups. Code teacher actions and intentions from usage logs to assess overlap with system-nudged recommendations. Evaluate alignment not only with behavior but with expressed pedagogical goals and guiding frameworks.
RQ2: Diversity of Instructional Practices	Compare distributions of instructional practices between teacher-initiated and system-nudged teacher-AI conversations. Use richness metrics to capture diversity. Analyze how different weighting strategies between co-occurrence network and user preference graphs influence the range of adopted practices.
RQ3: Explainability, Contextual Alignment, & Adoption Potential	Conduct interviews and surveys to assess teachers' perceptions of nudge transparency, peer grounding, and relevance. Analyze adoption behavior quantitatively through click-through rates [?]. Experiment with network configuration (e.g., relative weighting of personal vs. collective behavior) to test effects on both perceived and actual adoption.

4.1. Expected Results

Preliminary retrospective analysis using a held-out dataset, which was not used during network construction, indicates that system-generated recommendations closely align with teachers' actual instructional actions, validating the predictive strength of the co-occurrence network. The system also surfaces a broader range of effective instructional strategies that are pedagogically appropriate for the given context. This suggests the system is capable of nudging teachers toward effective, research-backed practices beyond their habitual repertoire.

We expect this hybrid approach, which combines individual usage patterns, collective educator behavior, and pedagogical research, to promote both adoption and instructional growth. We expect that recommendations align with their peers' behavior are more likely to be perceived as trustworthy and contextually relevant by teachers, while the integration of evidence-based strategies enhances the depth and diversity of instructional planning. We anticipate that the practice-informed network will outperform LLM-only inference in two key areas: (1) contextual precision and teacher acceptance, due to its grounding in real educator behavior, and (2) pedagogical value, by encouraging the uptake of strategies aligned with effective teaching frameworks. These findings will be further validated through controlled experimentation, user behavior analysis, and qualitative teacher feedback. These findings will be further validated through controlled experimentation, user behavior analysis, and qualitative teacher feedback.

5. Conclusion

This study introduces a practice-informed, behaviorally grounded AI nudging pipeline that supports instructional decision-making without compromising teacher autonomy. By combining a co-occurrence network of educator behavior with pedagogical research, the system generates nudges that are explainable, context-aware, and value-sensitive. Our approach exemplifies hybrid intelligence in education, blending human insight with computational support to guide, not prescribe, teacher action. The findings offer practical and theoretical implications for designing responsible AI systems in education and other professional domains.

Acknowledgments

This work is supported by the Institute of Education Sciences of the U.S. Department of Education, through Grant R305C240012 and by several awards from the National Science Foundation (NSF #2043613, #2300291, #2405110) to the University of Washington, and a NSF SBIR/STTR award to Hensun Innovation LLC (#2423365). The opinions expressed are those of the authors and do not represent views of the funders.

Declaration on Generative AI

During the preparation of this work, the author(s) used ChatGPT-4o in order to: Grammar and spelling check. After using these tool(s)/service(s), the author(s) reviewed and edited the content as needed and take(s) full responsibility for the publication's content.

References

[1] D. K. Cohen, S. W. Raudenbush, D. L. Ball, Resources, instruction, and research, Educational Evaluation and Policy Analysis 25 (2003) 119–142.
[2] M. Ding, M. A. Carlson, Elementary teachers' learning to construct high-quality mathematics lesson plans, The Elementary School Journal 113 (2013) 359–385.
[3] J. Choppin, J. M. Amador, C. Callard, C. Carson, R. Gillespie, Synchronous online model for mathematics teachers' professional development, in: Handbook of Research on Online Pedagogical Models for Mathematics Teacher Education, IGI Global, Hershey, PA, 2020, pp. 176–202.
[4] J. Shen, S. Poppink, Y. Cui, G. Fan, Lesson planning: A practice of professional responsibility and development, Educational Horizons 85 (2007) 248–258.
[5] L. Darling-Hammond, M. E. Hyler, M. Gardner, Effective Teacher Professional Development, Learning Policy Institute, 2017.
[6] N. D. Jones, E. M. Camburn, B. Kelcey, E. Quintero, Teachers' time use and affect before and after covid-19 school closures, AERA Open 8 (2022) 23328584211068068. doi:10.1177/23328584211068068.
[7] R. de Brito Duarte, J. Campos, Looking for cognitive bias in ai-assisted decision-making (2024).
[8] R. H. Thaler, C. R. Sunstein, Nudge: Improving Decisions About Health, Wealth, and Happiness, Yale University Press, New Haven, CT, 2008.
[9] W. Samuelson, R. Zeckhauser, Status quo bias in decision making, Journal of Risk and Uncertainty 1 (1988) 7–59.
[10] R. S. Nickerson, Confirmation bias: A ubiquitous phenomenon in many guises, Review of General Psychology 2 (1998) 175–220.

[11] M. E. Rodriguez, A. E. Guerrero-Roldán, D. Baneres, A. Karadeniz, An intelligent nudging system to guide online learners, The International Review of Research in Open and Distributed Learning 23 (2022) 41–62. doi:10.19173/irrodl.v22i4.5407.
[12] M. Weinmann, C. Schneider, J. vom Brocke, Digital nudging, Business & Information Systems Engineering 58 (2016) 433–436.
[13] R. Rebonato, A critical assessment of libertarian paternalism, Journal of Consumer Policy 37 (2014) 357–396.
[14] M. Tissenbaum, J. Slotta, Supporting classroom orchestration with real-time feedback: A role for teacher dashboards and real-time agents, International Journal of Computer-Supported Collaborative Learning 14 (2019) 325–351.
[15] K. Holstein, B. M. McLaren, V. Aleven, Co-Designing a Real-Time Classroom Orchestration Tool to Support Teacher-AI Complementarity, Technical Report, Grantee Submission, 2019.
[16] L. Bovens, The ethics of nudge, in: S. O. Hansson, T. Grüne-Yanoff (Eds.), Preference Change: Approaches from Philosophy, Economics and Psychology, Springer Netherlands, Dordrecht, 2009, pp. 207–219.
[17] H. Farrell, C. R. Shalizi, Pursuing cognitive democracy, in: D. Allen, J. Light (Eds.), From Voice to Influence: Understanding Citizenship in a Digital Age, University of Chicago Press, Chicago, IL, 2015, pp. 211–231.
[18] F. Furedi, On Tolerance: A Defence of Moral Independence, Bloomsbury Publishing, London, 2011.
[19] Y. Dimitriadis, R. Martínez-Maldonado, K. Wiley, Human-centered design principles for actionable learning analytics, in: C. Mouza, N. Lavigne (Eds.), Research on E-Learning and ICT in Education: Technological, Pedagogical and Instructional Perspectives, Springer, Cham, 2021, pp. 277–296.
[20] A. Südkamp, J. Kaiser, J. Möller, Teachers' judgments of students' academic achievement: Results from field and experimental studies, in: S. Krolak-Schwerdt, S. Glock, M. Böhmer (Eds.), Teachers' Professional Development, Brill, Leiden, 2014, pp. 5–25.
[21] B. Shneiderman, Human-centered artificial intelligence: Three fresh ideas, AIS Transactions on Human-Computer Interaction 12 (2020) 109–124.
[22] S. J. B. Shum, R. Luckin, Learning analytics and ai: Politics, pedagogy and practices, British Journal of Educational Technology 50 (2019) 2785–2793.
[23] N. R. C. U. C. on a Conceptual Framework for New K-12 Science Education Standards, A Framework for K-12 Science Education: Practices, Crosscutting Concepts, and Core Ideas, National Academies Press, Washington, DC, 2012.
[24] T. Mirsch, C. Lehrer, R. Jung, Digital Nudging: Altering User Behavior in Digital Environments, Technical Report, 2017.