Research Square

Preprints are preliminary reports that have not undergone peer review.
They should not be considered conclusive, used to inform clinical practice,
or referenced by the media as validated information.

Human-in-the-loop GenAI advising for community college pathway decision-making: a design-based case study of requirement-group modeling and privacy-preserving student-support design

Emery Peck

epeck7@ivytech.edu

Ivy Tech Community College of Indiana

Research Article

Keywords: Generative AI, academic advising, community college, guided pathways, learner agency, privacy-preserving design, self-regulated learning, human-in-the-loop, curriculum modeling

Posted Date: July 17th, 2026

DOI: https://doi.org/10.21203/rs.3.rs-10319737/v1

License: This work is licensed under a Creative Commons Attribution 4.0 International License.
Read Full License

Additional Declarations: No competing interests reported.

Abstract

Generative artificial intelligence (GenAI) research in higher education has focused primarily on tutoring, writing, feedback, and assessment, while comparatively limited attention has been paid to advising and pathway planning as consequential sites of human–GenAI interaction. This gap is especially important in community college contexts, where learners often navigate complex program requirements, transfer uncertainty, and evolving educational and workforce goals. This article reports a design-based case study of a pathway-planning prototype and advising interaction model developed around public curricular data, requirement-group representations, and privacy-preserving architectural separation from student-level records. Drawing on public catalog documents, system design artifacts, workflow analyses, and reflective practitioner documentation, the study asks what interaction-design, governance, and modeling principles are needed for GenAI-supported advising in community college pathway decision-making. The findings are presented as eight design principles: model requirements as flexible groups rather than flat checklists; position GenAI as an explanation and option-framing layer rather than the system of record; separate public program data from protected student records; preserve human advisor review for consequential recommendations; design for learner agency and self-regulated planning; use non-deficit advising language; make uncertainty visible; and align workforce relevance without reducing advising to labor-market demand. The article extends human–GenAI interaction research into advising and student-success operations and offers a privacy-conscious design framework for institutions seeking to use GenAI in educationally consequential but high-governance settings.

Introduction

Generative artificial intelligence (GenAI) now occupies a prominent place in higher education research and practice, with published work examining writing support, feedback generation, assessment reform, tutoring, institutional policy, and student and teacher perceptions (Bond et al., 2024; Crompton & Burke, 2023; Chan & Hu, 2023). Across this literature, scholars increasingly emphasize responsible adoption, human oversight, AI literacy, and the need to align technological innovation with educational purposes rather than simple efficiency gains (Chan, 2023; Walter, 2024; Crompton et al., 2026).

Yet one domain remains comparatively underexamined: advising and pathway planning. This omission matters because advising is not merely an administrative service. In higher education—and especially in community colleges—pathway decisions shape students' opportunity structures, course sequencing, time-to-goal, transfer trajectories, and perceptions of what options remain available to them. These are educationally consequential decisions that affect learner agency and self-regulated planning, even when they occur outside the formal classroom.

The community college setting intensifies this design challenge. Students often navigate highly flexible curricular structures that include fixed requirements, elective categories, general-education menus, concentration choices, wildcard course patterns, and multiple valid routes through a program. Transfer-oriented learners may face additional articulation complexity and information asymmetries across

institutions. When these structures are reduced to flat checklists, advising tools risk misrepresenting how programs actually work, thereby narrowing rather than expanding students' perceived options.

At the same time, recent higher education AI research raises pressing concerns about privacy, trust, and governance. Students and institutions alike express concern about how GenAI systems handle personal information, how much authority they should be given, and how they should be governed in educational settings (An et al., 2025; Wang et al., 2025; Crompton et al., 2026). These concerns are particularly acute in student-success operations, where the stakes of error, overreach, and unauthorized data use are high. In advising contexts, this suggests that GenAI should not be treated as an autonomous recommender or institutional decision-maker, but as a carefully bounded interaction layer embedded within human review and privacy-preserving system design.

This article addresses that gap through a design-based case study of a community-college pathway-planning prototype and associated advising interaction model. The case centers on three intertwined design problems: how to model flexible program requirements faithfully, how to support student-facing what-if advising through GenAI without over-automating consequential decisions, and how to preserve explicit privacy boundaries by separating public curricular data from protected student data.

The study is guided by the following research questions:

RQ1. What interaction-design principles are needed for GenAI-supported advising tools that interpret flexible community college program requirements as requirement groups rather than flat course checklists?

RQ2. How can human-in-the-loop GenAI advising systems support pathway planning while preserving learner agency, advisor accountability, and interpretive transparency?

RQ3. What privacy, governance, and data-separation practices are necessary for GenAI-supported advising in community college student-success contexts, particularly when public curricular data and protected student data must remain distinct?

The article makes four contributions. First, it extends human–GenAI interaction research from classroom-centered use cases into advising and pathway planning. Second, it proposes requirement-group modeling as a more educationally faithful representation of community college pathways than flat checklist logic. Third, it develops a privacy-preserving governance model for GenAI-supported advising grounded in explicit separation between public curricular information and student-level data. Fourth, it offers a set of design principles for institutions seeking to use GenAI in advising without replacing human judgment or making unsupported claims about retention or completion effects.

Background and literature review

Human interaction with GenAI in higher education

Recent higher education literature has shifted from asking whether GenAI matters to asking how it should be integrated and governed. Meta-synthetic and state-of-the-field reviews show rapid growth in research on AI in higher education, but also identify uneven rigor and a need for stronger attention to ethics, collaboration, and educational design (Bond et al., 2024; Crompton & Burke, 2023). Student- and teacher-focused studies likewise suggest that GenAI is experienced not only as a tool of efficiency but as a challenge to existing ideas about authorship, pedagogy, and trustworthy support (Chan & Hu, 2023; Barrett & Pack, 2023; Ellis et al., 2025).

Within IJETHE specifically, recent work has emphasized AI literacy, prompt engineering, institutional policy, and human-AI collaboration as key emergent concerns (Walter, 2024; Lee & Palmer, 2025; Hwang & Lee, 2025; An et al., 2025). However, most of this work remains centered on classroom activities, especially writing, assessment, and feedback.

GenAI, feedback, trust, and learner agency

One important strand of GenAI scholarship concerns feedback. Studies comparing teacher, peer, and GenAI feedback suggest that AI can sometimes produce structured, scalable, and criterion-aligned feedback, but that usefulness depends on context, prompting, student uptake, and interpretive support (Banihashem et al., 2024; Lee & Moore, 2024; Henderson et al., 2025; Farrokhnia et al., 2026). Crucially, higher apparent feedback quality does not necessarily translate into improved learning unless students meaningfully engage with what the feedback says and understand how to apply it (Farrokhnia et al., 2026).

This finding is highly relevant beyond writing tasks. Advising interactions also involve interpretation, uptake, and choice. An advising system that presents options without interpretive support may not strengthen learner agency. By contrast, an advising interaction that helps students understand requirement logic, trade-offs, and alternatives may support forms of self-regulated decision-making analogous to feedback literacy.

Learning analytics, advising, and decision support

Learning analytics and AI-supported advising have long promised more timely and tailored student support, but research has also shown that predictive or dashboard-oriented systems raise difficult questions about interpretability, privacy, and actionability. Scholars have argued that learning analytics should promote student autonomy and informed consent rather than simply intensify surveillance or extractive data use (Slade & Prinsloo, 2019). Ethical-use work in learning analytics similarly emphasizes governance, proportionality, transparency, and actionable interpretation rather than opaque prediction pipelines.

In recent higher education GenAI governance research, privacy, accountability, and human oversight are now treated as core principles rather than optional add-ons (Crompton et al., 2026). This trend is

especially important for advising, where institutions may be tempted to collapse public program information, engagement data, risk labels, and automated messaging into a single system. The present study argues that a better starting point is narrower and more defensible: a public-data advising layer that can provide requirement interpretation and pathway what-if support without depending on protected student records.

Human-centered design and co-design for AI in education

Human-centered AI scholarship in education increasingly argues that AI systems should be designed around the needs, interpretations, and capacities of learners and educators rather than around automation for its own sake. Reviews and policy frameworks emphasize that educational AI tools should support critical engagement, transparency, and responsible collaboration between humans and machines (Markauskaite et al., 2022; Khosravi et al., 2022; Ellis et al., 2025).

Recent IJETHE work likewise points to the importance of framing AI as a collaborative or bounded support layer rather than as an autonomous educational actor (Hwang & Lee, 2025; Crompton et al., 2026). This is especially relevant for advising, where a fully automated system could easily obscure ambiguity, oversimplify pathways, or project false certainty in high-stakes situations.

Privacy, governance, and responsible AI in higher education

Privacy concerns are now central to GenAI adoption in higher education. Institutional policy analyses and student-focused privacy studies show that uncertainty about data use, transparency, and institutional responsibility can reduce trust and complicate adoption (An et al., 2025; Wang et al., 2024; Wang et al., 2025). Governance frameworks increasingly emphasize academic integrity, privacy, equitable access, GenAI literacy, human oversight, and institutional infrastructure as interconnected rather than separable concerns (Crompton et al., 2026).

In educational technology more broadly, explainability and visible uncertainty matter because opaque systems make it harder for users to calibrate trust appropriately (Khosravi et al., 2022). In advising, uncertainty can arise not only from model limitations but from curricular ambiguity itself. Requirement statements are sometimes underspecified, historically layered, or context-dependent. A responsible advising system should therefore make ambiguity visible rather than flattening it into deterministic output.

Community college pathway complexity and workforce alignment

Community colleges are a particularly important but underrepresented context for AI-in-higher-education design. Guided pathways reform has emphasized the value of clear program maps, proactive advising, and momentum-preserving decision structures, but implementation research also shows that pathway

design is institutionally complex and interpretively demanding. In transfer settings, articulation complexity and web-based information asymmetries can further complicate students' attempts to build coherent plans.

At the same time, community colleges are often expected to connect educational pathways to local labor-market opportunities. That expectation can be productive when it helps students interpret curricular options in light of evolving goals. It becomes problematic when workforce alignment collapses into a narrow matching logic that sidelines learner exploration or reduces advising to labor-market sorting. A human-centered GenAI advising design should therefore treat workforce alignment as one input into advising conversations, not as the sole endpoint of decision support.

Summary of the literature gap

Taken together, the literature suggests four unresolved problems. First, GenAI-in-higher-education work remains concentrated on classroom learning, feedback, and policy, with less attention to advising and pathway planning. Second, current scholarship insufficiently addresses how AI systems should interpret flexible curricular requirements rather than flat course lists. Third, although governance frameworks emphasize privacy and oversight, few studies translate those principles into advising-specific architectures. Fourth, community college contexts—where flexibility, transfer, and equity concerns are especially pronounced—remain comparatively underrepresented in human–GenAI interaction research.

Conceptual framework

This study proposes a human-in-the-loop GenAI advising framework organized around seven interacting layers.

The first layer is **public catalog and program data**: catalogs, course descriptions, prerequisites, and published program requirements. These materials are significant because they are public, auditable, and structurally central to pathway understanding.

The second layer is **requirement-group modeling**. Rather than treating each catalog row as an individually mandatory course, the framework represents programs as collections of requirement groups: fixed requirements, option sets, elective buckets, wildcard patterns, and conditional alternatives. This layer is conceptually important because it encodes curricular flexibility directly, rather than treating flexibility as an afterthought.

The third layer is the **GenAI interaction layer**. Here, GenAI serves as an explanatory interface that translates structured requirement logic into plain-language descriptions, scenario-based what-if responses, and learner-facing comparison narratives. In this model, GenAI is not the system of record. It is an interpretive and communicative layer.

The fourth layer is **human advisor judgment**. Advisors remain responsible for exception handling, edge cases, relational interpretation, and final review of consequential recommendations. This reflects both the governance literature's insistence on human oversight and the reality that advising includes contextual and interpersonal judgment that cannot be reduced to template completion.

The fifth layer is **learner agency and self-regulated planning**. In the proposed framework, the purpose of advising is not merely to tell students what to do but to help them understand their options, their constraints, their timing decisions, and their recovery pathways. This layer connects advising to broader educational goals of agency and self-regulation.

The sixth layer is **privacy-preserving governance**. Public curricular data are intentionally separated from protected student-level records. This means the advising model can function in a FERPA-safe mode using public data alone, while any later extension into student-specific planning would require additional governance, permissions, and human review.

The seventh layer is **workforce/curriculum alignment**. Workforce relevance may inform how pathways are explained and compared, but it should not overwrite student preference, exploration, or longer-term educational development.

The framework's central proposition is that when flexible curricular requirements are modeled explicitly as requirement groups and exposed through a bounded GenAI explanation layer, advising systems can improve pathway intelligibility while preserving learner agency, human accountability, and privacy boundaries.

Methods

Research design

This article adopts a design-based case study approach. Design-based research is appropriate when the goal is not simply to test whether an intervention "worked," but to examine how an educational design problem is framed, modeled, iterated, and translated into principled design knowledge. In this case, the focal problem is how to build a GenAI-supported advising system that interprets flexible program requirements accurately while preserving privacy boundaries, visible uncertainty, and human oversight.

A design-based case study is especially appropriate here because the empirical object is an evolving pathway-planning prototype and its associated advising workflow rather than a stabilized intervention with outcome measures. The analysis therefore treats documents, prototype artifacts, architecture decisions, and practitioner memos as evidence of how a human–GenAI advising design was constructed, bounded, and justified. The purpose is to derive transferable design principles from a bounded institutional case, not to claim causal effects on learners.

The case is bounded by the development of a pathway-planning prototype and associated advising workflow model in a community-college context. The study does not report student outcomes, retention effects, or workload reductions. Instead, it analyzes the case as a design problem and develops design principles from a structured corpus of non-student evidence.

Case boundaries and evidence sources

The case is bounded by the development of a community-college pathway-planning prototype and its associated advising workflow model over an eight-month design period. Units of analysis were documents and artifacts rather than people. The evidence base comprises three non-student source classes that together capture the curricular problem space, the prototype's requirement representations, and the governance decisions that bounded GenAI's role in advising.

Evidence source	Description	Role in analysis
Public curricular documents	Published program requirements, course descriptions, prerequisite statements, pathway descriptions, and related catalog materials	Identified recurring forms of requirement-group complexity and ambiguity that flat checklist logic handled poorly
Prototype artifacts	Requirement schemas, pathway logic documents, interface mockups, architecture diagrams, prototype screenshots, workflow maps, prompt templates, and decision rules	Analyzed how curricular flexibility, GenAI explanation behavior, uncertainty handling, and advisor-review checkpoints were modeled
Practitioner design memos	Design notes, architecture rationales, iteration memos, and workflow redesign documentation	Traced why major design decisions were made and how privacy/governance constraints shaped the final interaction model

The corpus is intentionally limited to non-student materials assembled across that eight-month design period. The analyzed materials comprised multiple curricular documents, prototype artifacts, and practitioner memos collected through iterative design work, but this article does not rely on a closed-count archival inventory and therefore reports the corpus by source class and analytic role rather than by exhaustive item totals. Public curricular documents define the problem space; prototype artifacts show how the problem was operationalized; and practitioner documentation provides the analytic bridge between design problem, iteration, and principle derivation. No individual student records, personally identifiable information, or advisor-identifying data are included in the study corpus.

Data analysis

The analysis proceeded in five stages. First, public curricular documents were reviewed to identify recurring forms of structural flexibility in degree requirements, including elective buckets, concentration branches, wildcard patterns, menu-based general-education requirements, and conditional substitutions.

Second, the author assembled an evidence matrix for all analyzed materials, logging source class, artifact function, requirement structure, GenAI role, privacy boundary, human-review point, and notable ambiguity. Third, design artifacts were examined to determine how the prototype and workflow model represented these structures and whether the system treated requirements as flat lists or as grouped, interpretable units. Fourth, architecture and workflow materials were analyzed for their treatment of privacy, human review, and GenAI's system role, with particular attention to whether GenAI functioned as an authoritative planner, an explanation layer, or a hybrid. Fifth, analytic memos were used to synthesize cross-cutting principles through constant comparison across the three source classes: curricular documents, prototype and architecture materials, and workflow and reflective documentation.

The coding logic was primarily deductive-inductive. Deductive attention was given to four sensitizing constructs derived from the literature and research questions: curricular representation, learner agency, human oversight, and privacy/governance. Inductive coding was then used to capture case-specific patterns such as wildcard requirement handling, ambiguity classes, advisor-review checkpoints, and explanation strategies. Principles were retained only when they were supported across more than one source class or when a single artifact documented a design decision with clear governance consequences.

Trustworthiness and study limits

The study's trustworthiness rests on four features: explicit case boundaries, artifact triangulation, a documented evidence matrix and memo trail, and restrained claims. Because the article does not infer student outcomes from design artifacts, it avoids a common overreach in educational technology writing: treating well-designed systems as if their efficacy had already been established. The analysis also preserved ambiguous or weakly specified curricular cases rather than forcing deterministic interpretations, which helped guard against over-cleaning the evidence. At the same time, the study is limited by its focus on one case and by the absence of direct learner-use data in the present article. Its contribution is therefore design knowledge rather than effectiveness evidence.

Use of generative AI in this study

Generative AI tools were used during manuscript preparation to support literature organization, comparative synthesis, outlining, and editorial drafting. The author reviewed, revised, verified, and approved all claims, interpretations, citations, and final wording. No generative AI tool was listed as an author, and no protected student data were entered into a generative AI system.

Findings: design principles for human-in-the-loop GenAI advising

Principle 1. Model academic requirements as flexible requirement groups, not flat checklists

The central design finding is that community college pathways should not be modeled as a sequence of individually mandatory course rows. Published requirements often encode flexibility structurally rather than explicitly. A degree may require one of several mathematics routes, a menu of humanities or social science courses, a concentration branch, or an elective pattern satisfying a wildcard. Treating these as flat checklists can produce false deficits and inaccurate “completion” logic.

Requirement-group modeling makes the logic of curricular choice explicit. It allows a system to represent both what is fixed and what is open, preserving the educational reality that multiple valid pathways often exist within a single credential structure.

Principle 2. Keep GenAI as an explanation and option-framing layer, not the system of record

The reviewed literature suggests that GenAI can be highly useful in explanation-rich tasks, but that its output quality does not automatically justify delegating authority to it (Banihashem et al., 2024; Farrokhnia et al., 2026). In advising, the safest and most educationally appropriate role is therefore interpretive rather than authoritative.

In this case, GenAI is best used to translate requirement logic into plain language, explain why certain options satisfy a group, compare what-if scenarios, and surface likely implications of route changes. This role supports student understanding without converting the model into the institutional system of record.

Principle 3. Separate public catalog data from protected student-level records

A major architectural finding is that a useful advising prototype can be built without ingesting individual student records at all. Public catalog and program data are sufficient for many pathway-explanation, planning, and what-if functions. This separation has two advantages. First, it reduces privacy risk. Second, it supports clearer system governance by distinguishing between publicly interpretable pathway logic and protected student-specific data.

This finding aligns with higher education AI governance literature emphasizing privacy, proportionate data use, and human accountability (An et al., 2025; Crompton et al., 2026; Wang et al., 2025).

Principle 4. Preserve human advisor review for consequential recommendations

The case suggests that advising systems should be designed around reviewability rather than automation closure. When a pathway recommendation could affect time, cost, eligibility, or program

direction, human advisors should remain the accountable reviewers of final guidance. This is not simply a legal or procedural preference. It is also a pedagogical design choice. Advisors contribute relational, contextual, and judgment-based interpretation that a public-data GenAI layer cannot fully replicate.

Principle 5. Design interactions that support learner agency and self-regulated planning

A learner-facing advising system should not only answer “What do I still need?” It should also help students understand why multiple routes exist, what trade-offs different routes create, and how choices made in one term constrain or open later decisions. Such support aligns advising with self-regulated learning by helping students reason about sequence, momentum, timing, and recovery rather than merely consuming instructions.

Principle 6. Use non-deficit advising language

The literature on AI and higher education increasingly warns against reductive framings that weaken student agency or foreground control rather than learning. In advising, this means avoiding deficit-centered or surveillance-heavy language such as “risk score” when the actual design task is to help students understand options. More constructive language includes concepts such as momentum preservation, pathway recovery, requirement completion logic, and scenario comparison.

Principle 7. Make uncertainty visible when rules are ambiguous

Curricular interpretation is not always binary. Some catalog statements are underspecified, locally interpreted, or historically layered. A responsible system should therefore expose uncertainty rather than suppress it. This can be implemented through interface and prompt design that distinguishes between clear requirement matches, likely matches requiring advisor review, and ambiguous cases requiring manual interpretation.

Principle 8. Align advising interaction with workforce pathways without reducing students to employer demand

The case also suggests a cautious but potentially productive role for workforce alignment. Labor-market relevance can help students compare pathways, understand program transferability, and interpret longer-term consequences of educational choices. However, workforce alignment should remain advisory, not determinative. Students’ educational goals, transfer ambitions, and evolving identities may not map cleanly to employer demand categories.

Discussion

This study contributes to human–GenAI interaction research in higher education by shifting attention from classroom-centered tasks to advising and pathway planning. That shift matters because advising is a consequential educational process even when it is institutionally categorized as student support rather than instruction. Pathway decisions shape learners' sense of what is possible, what is recoverable, and what trade-offs they face. In that sense, advising is deeply connected to self-regulated learning, academic agency, and educational persistence, even when those outcomes are not directly measured in this article.

The study's first contribution is conceptual. It argues that requirement interpretation is not a trivial data-parsing problem but a pedagogical design problem. When systems flatten flexible program structures into mandatory lists, they risk narrowing students' understanding of available routes. Requirement-group modeling therefore matters not only computationally but educationally.

The second contribution is architectural. The case demonstrates the feasibility and value of separating public curricular data from protected student-level records. This is especially relevant in an era when many institutions are considering expansive AI integrations without clear data-minimization principles. A public-data advising layer offers a more privacy-conscious entry point for GenAI adoption.

The third contribution is interactional. The case suggests that the most defensible role for GenAI in advising is as an explanation and option-framing layer rather than a decision authority. This aligns with emerging IJETHE scholarship emphasizing human oversight, AI literacy, and bounded forms of human-AI partnership.

The fourth contribution is practical. For community colleges and similarly structured institutions, the study offers a design vocabulary for building advising tools that are more faithful to curricular reality: requirement groups, uncertainty classes, scenario comparison, and advisor-review checkpoints.

Implications for community colleges

Community college advising often operates under constraints of scale, transfer complexity, role differentiation, and privacy sensitivity. These conditions make AI adoption both attractive and risky. Attractive, because GenAI can help translate complex pathway logic at scale; risky, because over-automation can mislead students in high-stakes choices. The design principles proposed here offer a middle path: bounded AI support that increases intelligibility without displacing human review.

Implications for teaching and learning

The article also suggests a broader reframing for educational technology research. Advising should be treated as a learning-adjacent domain in which human–GenAI interaction shapes self-regulation, not merely service efficiency. This reframing may help bridge currently separate literatures on feedback, learner agency, learning analytics, and student support.

Privacy and governance implications

The governance literature increasingly argues that educational AI policy should move beyond plagiarism control toward privacy, access, literacy, accountability, and support infrastructure (Crompton et al., 2026). This case adds a domain-specific extension to that argument: in advising, the strongest early design move may be architectural separation of public curricular logic from protected student data. This helps institutions reduce risk while still learning where GenAI adds value.

Implications for human–GenAI interaction design

This case suggests that interaction quality in GenAI-supported advising should be judged less by conversational fluency than by whether the system expands, clarifies, and appropriately bounds student option awareness. In practical terms, a useful advising interaction should explain requirement logic, surface trade-offs across plausible pathways, distinguish clear cases from ambiguous ones, and route consequential uncertainty to human advisors rather than masking it with confident prose.

These interaction choices matter for learner agency and self-regulated learning. Students need support for comparing routes, anticipating downstream consequences, and revising plans after setbacks; they do not merely need answer-like recommendations. The main trade-off is that a privacy-preserving, review-oriented design may appear less automated and less instantaneous than fully autonomous advising tools. However, in high-governance settings such as pathway planning, slower but reviewable interaction may be educationally and ethically preferable to fast but opaque recommendation.

Limitations

This study has several limitations. First, it is a single-case design-based study and therefore does not claim broad causal generalizability. Second, it does not report learner-use outcomes, advisor workload changes, or retention effects. Third, its design principles are derived from artifacts and workflow analysis rather than from a formal co-design or implementation trial reported here. Fourth, the article focuses on a public-data advising model; future student-specific extensions would require additional governance and empirical validation.

Future research

Three next steps are especially important. First, future studies should test how students and advisors actually use requirement-group-based GenAI pathway explanations in practice. Second, comparative work should examine whether the same design principles transfer across transfer-intensive, baccalaureate, and adult-serving contexts. Third, future work should test governance models for moving from public-data advising layers to carefully permissioned student-specific what-if tools without collapsing privacy boundaries.

Conclusion

GenAI research in higher education has moved quickly, but advising and pathway planning remain comparatively underdeveloped as sites of human–AI interaction. This design-based case study argues that community college pathway decision-making requires a distinctive design response: flexible curricular structures should be modeled as requirement groups rather than flat checklists; GenAI should be used as an explanation and option-framing layer rather than as an authoritative planner; privacy boundaries should be preserved through explicit separation of public curricular data from student-level records; and consequential recommendations should remain subject to human advisor review.

These principles do not demonstrate improved retention, completion, or advising efficiency, and the article does not claim that they do. Instead, the contribution is more foundational: it offers a conceptually and operationally grounded framework for building human-in-the-loop, privacy-conscious GenAI advising systems that respect curricular complexity and preserve learner agency. In a higher education landscape increasingly shaped by GenAI, that may be a necessary precondition for any later efficacy study worth trusting.

Abbreviations

AI: Artificial intelligence

FERPA: Family Educational Rights and Privacy Act

GenAI: Generative artificial intelligence

IJETHE: *International Journal of Educational Technology in Higher Education*

SRL: Self-regulated learning

Declarations

Availability of data and materials

Public curricular materials analyzed in this study are available through the relevant institutional catalog and program websites. Non-student design artifacts and workflow materials analyzed during the current study are not publicly available because they include internal implementation documentation, but non-sensitive materials are available from the corresponding author on reasonable request. No individual student records or personally identifiable student data were used.

Clinical trial number: not applicable.

Consent to Publish declaration: not applicable’

Consent to Participate declaration: not applicable

‘Ethics declaration: not applicable

Competing interests

The author declares that there are no competing interests.

Funding

No external funding was received for this study.

Authors’ contributions

The author conceived the study, conducted the design analysis, synthesized the literature, and drafted the manuscript. The author read and approved the final manuscript.

Acknowledgements

Not applicable for the anonymized review version.

Authors’ information

Removed for double-anonymous review.

GenAI use disclosure

Generative AI tools were used during manuscript preparation to support literature organization, comparative synthesis, outlining, and editorial drafting. The author reviewed, revised, verified, and approved all claims, interpretations, citations, and final wording. No generative AI tool was listed as an author, and no protected student data were entered into a generative AI system.

References

1. An, Y., Yu, J. H., & James, S. (2025). Investigating the higher education institutions’ guidelines and policies regarding the use of generative AI in teaching, learning, research, and administration. *International Journal of Educational Technology in Higher Education, 22*, Article 10. https://doi.org/10.1186/s41239-025-00507-3
2. Barrett, A., & Pack, A. (2023). Not quite eye to AI: Student and teacher perspectives on the use of generative artificial intelligence in the writing process. *International Journal of Educational Technology in Higher Education, 20*, Article 27. https://doi.org/10.1186/s41239-023-00427-0

3. Bond, M., Khosravi, H., De Laat, M., Bergdahl, N., Negrea, V., Oxley, E., Pham, P., Chong, S. W., & Siemens, G. (2024). A meta systematic review of artificial intelligence in higher education: A call for increased ethics, collaboration, and rigour. *International Journal of Educational Technology in Higher Education, 21*(1), 4. https://doi.org/10.1186/s41239-023-00436-z
4. Chan, C. K. Y. (2023). A comprehensive AI policy education framework for university teaching and learning. *International Journal of Educational Technology in Higher Education, 20*, Article 38. https://doi.org/10.1186/s41239-023-00408-3
5. Chan, C. K. Y., & Hu, W. (2023). Students' voices on generative AI: Perceptions, benefits, and challenges in higher education. *International Journal of Educational Technology in Higher Education, 20*(1), 43. https://doi.org/10.1186/s41239-023-00411-8
6. Crompton, H., & Burke, D. (2023). Artificial intelligence in higher education: The state of the field. *International Journal of Educational Technology in Higher Education, 20*, 22. https://doi.org/10.1186/s41239-023-00392-8
7. Crompton, H., Burke, D., Nickel, C., Bozkurt, A., Miao, F., Sharples, M., Greene, J. A., Parsons, D., Gill-Simmen, L., Edmett, A., Pegrum, M., de Waard, I., Bonk, C. J., Garcia, M. B., Curry, J. H., Lindsey, L. A., Yang, M., Marshall, S., Bali, M., et al. (2026). Governing generative AI in higher education: A global Delphi study on policy and practice. *International Journal of Educational Technology in Higher Education, 23*, Article 21. https://doi.org/10.1186/s41239-026-00602-z
8. Dai, W., Tsai, Y. S., Lin, J., Aldino, A., Jin, H., Li, T., Gašević, D., & Chen, G. (2024). Assessing the proficiency of large language models in automatic feedback generation: An evaluation study. *Computers and Education: Artificial Intelligence, 7*, Article 100299. https://doi.org/10.1016/j.caeai.2024.100299
9. Ellis, R., Han, F., & Cook, H. (2025). Qualitatively different teacher experiences of teaching with generative artificial intelligence. *International Journal of Educational Technology in Higher Education, 22*, Article 33. https://doi.org/10.1186/s41239-025-00532-2
10. Er, E., Akçapınar, G., Bayazıt, A., Noroozi, O., & Banihashem, S. K. (2025). Assessing student perceptions and use of instructor versus AI-generated feedback. *British Journal of Educational Technology, 56*(3), 1088–1103. https://doi.org/10.1111/bjet.13558
11. Escalante, J., Pack, A., & Barrett, A. (2023). AI-generated feedback on writing: Insights into efficacy and ENL student preference. *International Journal of Educational Technology in Higher Education, 20*, Article 55. https://doi.org/10.1186/s41239-023-00425-2
12. Farrokhnia, M., Latifi, S., Papadopoulos, P. M., Hogenkamp, L., Gijlers, H., Khosravi, H., & Noroozi, O. (2026). Generative AI offers more, but students revise less: Comparing the effects of teacher and AI feedback on student essay revisions. *International Journal of Educational Technology in Higher Education, 23*, Article 6. https://doi.org/10.1186/s41239-026-00579-9
13. Henderson, M., Bearman, M., Chung, J., Fawns, T., Buckingham Shum, S., Matthews, K. E., & de Mello Heredia, J. (2025). Comparing generative AI and teacher feedback: Student perceptions of

usefulness and trustworthiness. *Assessment & Evaluation in Higher Education*. https://doi.org/10.1080/02602938.2025.2502582

14. Hwang, Y., & Lee, J. H. (2025). Exploring students' experiences and perceptions of human-AI collaboration in digital content making. *International Journal of Educational Technology in Higher Education, 22*(1), 44. https://doi.org/10.1186/s41239-025-00542-0
15. Khosravi, H., Buckingham Shum, S., Chen, G., Conati, C., Tsai, Y. S., Kay, J., & Gašević, D. (2022). Explainable artificial intelligence in education. *Computers and Education: Artificial Intelligence, 3*, Article 100074. https://doi.org/10.1016/j.caeai.2022.100074
16. Labadze, L., Grigolia, M., & Machaidze, L. (2023). Role of AI chatbots in education: Systematic literature review. *International Journal of Educational Technology in Higher Education, 20*(1), 56. https://doi.org/10.1186/s41239-023-00426-1
17. Lee, D., & Palmer, E. (2025). Prompt engineering in higher education: A systematic review to help inform curricula. *International Journal of Educational Technology in Higher Education, 22*(1), Article 7. https://doi.org/10.1186/s41239-025-00503-7
18. Lee, S. S., & Moore, R. L. (2024). Harnessing generative AI (GenAI) for automated feedback in higher education: A systematic review. *Online Learning, 28*(3), 82–106. https://doi.org/10.24059/olj.v28i3.4593
19. Markauskaite, L., Marrone, R., Poquet, O., Knight, S., Martinez-Maldonado, R., Howard, S., Tondeur, J., de Laat, M., Buckingham Shum, S., Gašević, D., & Siemens, G. (2022). Rethinking the entwinement between artificial intelligence and human learning: What capabilities do learners need for a world with AI? *Computers and Education: Artificial Intelligence, 3*, Article 100056. https://doi.org/10.1016/j.caeai.2022.100056
20. Slade, S., & Prinsloo, P. (2019). Learning analytics and higher education: A proposed model for establishing informed consent mechanisms to promote student privacy and autonomy. *International Journal of Educational Technology in Higher Education, 16*, Article 24. https://doi.org/10.1186/s41239-019-0155-0
21. Ursavaş, Ö. F., Yalçın, Y., İslamoğlu, H., Bakır-Yalçın, E., & Cukurova, M. (2025). Rethinking the importance of social norms in generative AI adoption: Investigating the acceptance and use of generative AI among higher education students. *International Journal of Educational Technology in Higher Education, 22*(1), 38. https://doi.org/10.1186/s41239-025-00535-z
22. Walter, Y. (2024). Embracing the future of artificial intelligence in the classroom: The relevance of AI literacy, prompt engineering, and critical thinking in modern education. *International Journal of Educational Technology in Higher Education, 21*, Article 48. https://doi.org/10.1186/s41239-024-00448-3
23. Wang, H., Dang, A., Wu, Z., & Mac, S. (2024). Generative AI in higher education: Seeing ChatGPT through universities' policies, resources, and guidelines. *Computers and Education: Artificial Intelligence, 7*, Article 100326. https://doi.org/10.1016/j.caeai.2024.100326

24. Wang, N., Li, Y., & Cong, F. (2025). University students' privacy concerns towards generative artificial intelligence. *Journal of Academic Ethics*. https://doi.org/10.1007/s10805-025-09658-4