Teaching Problem Framing in the Age of Generative AI

A Human Insights × Artificial Intelligence (HI x AI) Pedagogy

Kevin Popovic

Abstract

As generative artificial intelligence becomes embedded in HCI education, students can generate analyses and plausible solutions with unprecedented speed, often before clearly articulating the problem they intend to frame. While AI increases apparent productivity, it can also obscure weak sensemaking, unexamined assumptions, and premature convergence during early stages of inquiry. This paper argues that problem framing must be taught explicitly in AI-enabled learning environments.

We present Human Insights × Artificial Intelligence (HI × AI) as a structured pedagogy that formalizes deliberate alternation between human interpretation and AI-supported expansion. The approach is operationalized through the HI × AI Infinity Loop, a four-phase learning rhythm of Frame, Expand, Interpret, and Refine. Grounded in repeated instructional use across undergraduate, pre-college, and large-enrollment contexts, the framework sequences AI integration to preserve human interpretive responsibility while leveraging computational amplification.

By positioning AI as an amplifier rather than an authority, HI × AI offers a practical instructional structure for cultivating disciplined problem framing before solution development in AI-accelerated classrooms.

Keywords: Generative AI, Problem Framing, Human Insights × Artificial Intelligence, HI × AI, HCI Education, Reflective Practice, Hybrid Intelligence, Instructional Design, AI-Enabled Learning

Introduction

Generative artificial intelligence has altered the conditions under which students learn, design, and innovate. By lowering the effort required to produce ideas, analyses, and plausible solutions, AI compresses early stages of inquiry that once demanded sustained sensemaking. In this environment, students can advance quickly without fully articulating the problem they are attempting to frame.

While AI increases speed and apparent productivity, it can also obscure weak reasoning, unexamined assumptions, and premature convergence. Polished outputs may create a sense of understanding without requiring students to surface assumptions, examine evidence, or determine what meaningfully matters. As a result, early-stage reasoning risks becoming implicit rather than deliberate.

This paper argues that problem framing must be taught explicitly in AI-enabled HCI education. Rather than assuming learners will naturally develop this capability, the HI × AI pedagogy treats framing as a central and metacognitive learning outcome that requires structure, reflection, and deliberate practice.

The HI × AI approach structures early-stage inquiry around a four-part learning rhythm in which human judgment establishes direction and manages evaluation, while artificial intelligence expands, tests, and refines understanding through deliberate alternation. By sequencing when and how AI is introduced, the pedagogy preserves human interpretation while leveraging technological capability.

The sections that follow situate this approach within existing design traditions, describe its conceptual foundation, demonstrate how it is operationalized in the classroom, and present observed instructional effects across multiple contexts.

1. Teachable Moment Summary

This paper presents a structured and repeatable instructional approach for cultivating problem framing as a core learning outcome in AI-enabled HCI education. As generative artificial intelligence becomes embedded in classrooms, studios, and project-based learning environments, students can now generate analyses and plausible solutions with remarkable speed, often before clearly articulating the problem they intend to frame. The instructional risk is that polished AI-generated outputs may substitute for disciplined inquiry, masking weak sensemaking and unexamined assumptions, a concern increasingly noted in scholarship on teaching with generative AI (Gallant and Rettinger 2025).

The pedagogical intervention described here addresses this challenge by explicitly teaching students how to think with AI during early-stage inquiry. Rather than focusing on surface-level prompting skill or accelerating artifact production, the approach sequences AI use so that human judgment establishes direction and AI amplifies understanding. Framing is treated not as an implicit prerequisite to design work but as an explicit, assessable, and metacognitive capability grounded in reflective practice (Schön 1983).

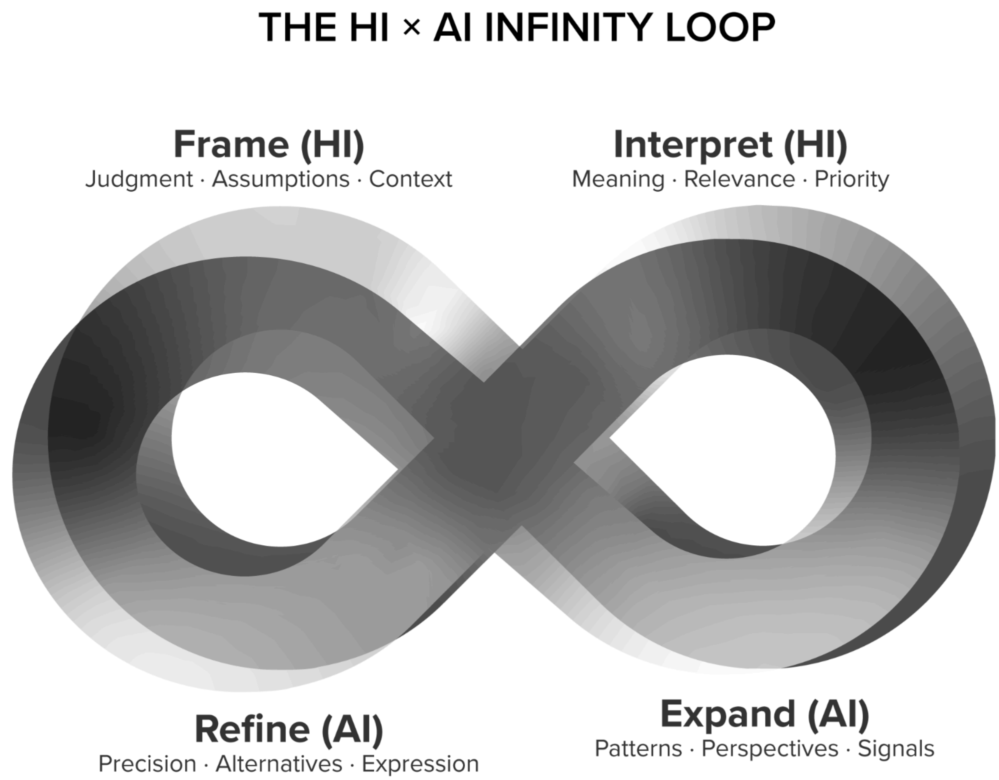

Figure 1 — The HI × AI Infinity Loop

> *The HI × AI Infinity Loop. A four-phase learning rhythm that structures deliberate alternation between human judgment and AI-supported expansion in early-stage inquiry.*

As illustrated in Figure 1, the HI × AI Infinity Loop formalizes this learning rhythm as a four-phase structure that sequences deliberate alternation between Human Insights and Artificial Intelligence. The loop moves through Frame, Expand, Interpret, and Refine before returning to framing in the next iteration. Human judgment establishes context, surfaces assumptions, and determines relevance, while AI expands perspective and supports refinement without assuming interpretive authority. This epistemic ordering aligns with scholarship on hybrid intelligence that emphasizes human-guided machine collaboration (Dellermann et al. 2019).

The contribution of this paper is threefold. First, it formalizes the HI × AI Infinity Loop as a structured method for early-stage inquiry in AI-enabled classrooms. Second, it introduces a sequencing model for AI integration that clarifies when and how generative tools should enter inquiry in order to preserve human epistemic responsibility. Third, it operationalizes this rhythm as an assessable instructional workflow grounded in repeated classroom iteration across undergraduate, pre-college, and large-enrollment innovation contexts.

By sequencing rather than restricting AI, the pedagogy resolves a central instructional tension in contemporary HCI education. AI is neither excluded nor permitted to dominate early reasoning. Instead, it is embedded within a disciplined structure that protects judgment while leveraging computational expansion. In doing so, the approach restores balance between speed and understanding and makes early-stage reasoning visible and accountable within AI-accelerated learning environments.

2. Teaching Context & Origins of the Work

Instructional and Institutional Contexts

The HI × AI pedagogy emerged through sustained teaching practice across multiple educational settings. Rather than originating as a theory-first construct, the framework developed through iterative instructional use with learners at different stages, allowing recurring patterns in early-stage inquiry to surface over time. These contexts included undergraduate innovation workshops, pre-college outreach programs, and a large-enrollment entrepreneurship course.

Case Example 1: Early Opportunity Framing in The Basement

Initial experimentation occurred through Opportunity Framing workshops delivered for The Basement, UC San Diego's interdisciplinary innovation hub. Students from diverse majors were guided through structured exercises designed to move from loosely articulated challenges toward clearer problem statements. Each exercise enacted a version of the HI × AI rhythm, requiring learners to frame the problem independently before expanding perspective through AI-supported exploration.

Despite this structure, many participants initially described surface-level symptoms or gravitated toward solution-oriented language. The deliberate alternation between human articulation and AI expansion made these tendencies visible. These early sessions revealed that the difficulty lay not in creativity, but in disciplined articulation of assumptions, stakeholders, and contextual constraints.

Case Example 2: Pre-College Workshops and Learner-Level Independence

Similar patterns appeared in pre-college workshops conducted through UC San Diego outreach programs. These sessions focused exclusively on problem identification rather than ideation. Students often began with broadly stated challenges lacking stakeholder specificity or explicit assumptions. In the absence of structured progression, inquiry relied primarily on intuition rather than systematic examination.

Many learners also expressed impatience with sustained problem identification, preferring to move toward ideation. However, when guided through a sequenced framing process, students began to recognize the value of each phase. The artifacts generated during these sessions, including written assumptions, revised framings, and AI-expanded perspectives, provided visible evidence of reasoning and supported articulation of how understanding evolved.

The recurrence of these patterns across age groups suggested that early-stage framing difficulties were structural rather than primarily attributable to developmental stage or field of study.

Case Example 3: Refinement at Scale in Essentials of Entrepreneurship

The pedagogy was further refined within a recurring session of a large-enrollment undergraduate course, *Essentials of Entrepreneurship*, offered once per quarter across three academic quarters each year. The author led one of eleven class sessions focused specifically on problem framing. With more than 100 students per offering and repeated delivery over four years, this setting provided sustained opportunities for refinement.

As generative AI became increasingly embedded across successive offerings, many students entered the session struggling to identify a clearly defined problem to address. Some used AI early in their projects to generate structured thinking or exploratory ideas, but these often lacked clearly articulated stakeholders, documented assumptions, or justification for why the problem mattered. Without structured guidance, students frequently moved toward solution concepts before fully identifying the problem space.

The opportunity framing intervention provided a structured pathway for identifying and articulating a problem more deliberately and completely than students typically achieved independently. By structuring inquiry through framing, expansion, interpretation, and refinement, the workflow supported clearer identification of stakeholders, assumptions, and contextual boundaries.

Over successive offerings, the instructional workflow was refined to make the HI × AI alternation more explicit and consistent. The sequencing of framing, expansion, interpretation, and refinement became clearer to students and more tightly integrated into assessment. As a result, patterns in student articulation, assumption identification, and justification of problem boundaries became more consistent across iterations.

From Instructional Iteration to Framework Development

Across these contexts, recurring patterns in early-stage inquiry informed the formalization of the HI × AI framework. What began as the sharing of professional problem-framing practices within instructional settings evolved, through iterative teaching and observation, into a deliberate articulation of the roles between human judgment and technological expansion. The framework presented in this paper is therefore grounded in applied instructional practice rather than abstract design theory.

3. Teaching Challenge: Problem Framing in AI-Accelerated Learning

The widespread availability of generative artificial intelligence has introduced a new instructional challenge for HCI education. Students can now generate analyses and propose solutions with unprecedented speed, often without developing a clear understanding of the problem they are attempting to frame (Gallant and Rettinger 2025). While this acceleration can increase confidence and momentum, it also compresses early stages of inquiry that traditionally required sustained sensemaking. Tasks that once demanded reflection, articulation of assumptions, and examination of underlying causes are increasingly compressed or bypassed, despite the long-standing recognition of such inquiry as central to learning (Argyris & Schön 1978).

In much of the HCI literature on AI design, emphasis is placed on interaction techniques, interface behavior, and system capabilities (Amershi et al. 2019). These perspectives frequently inform how AI is taught in design-oriented courses. Students learn how to prompt models, evaluate outputs, and integrate AI into design workflows. What receives less explicit attention is how learners construct the problem space itself, including how assumptions are surfaced, how relevance is determined, and how ethical responsibility is established prior to design decisions. Yet problem framing has long been recognized as central to reflective practice and design inquiry (Schön 1983; Dorst 2011). When framing is treated as an implicit prerequisite rather than an explicit learning outcome, it becomes especially vulnerable in AI-accelerated environments.

This vulnerability intensifies as AI shifts from being merely available to becoming embedded and normalized within learning environments. When generative tools are introduced early in inquiry, AI-generated language can substitute for student articulation, masking uncertainty and ambiguity that meaningful learning requires. Polished responses may create a sense of understanding without requiring students to engage in the deeper work of examining assumptions, testing interpretations, and determining what meaningfully matters. Learners may converge quickly on plausible framings without being able to explain why a particular problem meaningfully matters, for whom it matters, or how it behaves within a broader system.

From an instructional perspective, this creates a structural tension. Banning AI use is neither practical nor desirable, and doing so risks framing AI as something to be avoided rather than understood. At the same time, unstructured access to AI can undermine the development of judgment, empathy, and critical inquiry capacities central to HCI education. The challenge, therefore, is not whether to include AI in the classroom, but how to structure its use so that it strengthens rather than replaces thinking during early-stage inquiry.

This paper addresses that challenge by positioning problem framing as a central learning outcome and by formalizing deliberate alternation between human judgment and technological expansion. The Human Insights × Artificial Intelligence (HI × AI) framework (Popović 2026) articulates a disciplined learning rhythm that sequences framing, expansion, interpretation, and refinement. By clarifying when and how AI enters inquiry, the framework preserves human epistemic responsibility while leveraging AI's capacity to amplify understanding.

4. Framing in Existing HCI and Design Traditions

The HI × AI framework builds on established traditions within HCI and design education that recognize the importance of framing in creative and professional practice. Research-through-design, as articulated by Zimmerman and colleagues (Zimmerman et al. 2007), emphasizes how understanding evolves through iterative cycles of making and reflection. Within this tradition, framing is embedded in design activity, as designers refine their understanding through engagement with artifacts and inquiry.

Similarly, work on human–AI interaction emphasizes interaction design, transparency, and system behavior once a design direction has been established (Amershi et al. 2019). These approaches provide guidance for interaction design and responsible AI integration.

Both traditions acknowledge the importance of framing, but typically situate it within broader design processes rather than treating it as an explicit and assessable learning outcome. In instructional settings, especially those where generative AI tools are readily accessible, this distinction becomes consequential. When AI can rapidly generate structured artifacts or plausible responses, students may engage in making before developing a sufficiently examined understanding of the problem space.

The contribution of HI × AI lies not in replacing these traditions but in clarifying the sequencing of early-stage inquiry. By making the alternation between human judgment and AI-supported expansion explicit, the framework foregrounds problem framing as a structured component of instruction. It provides educators with a way to preserve reflective interpretation while still leveraging AI's capacity to expand perspective.

In this sense, HI × AI operates as an instructional layer that precedes and strengthens downstream design methods. By structuring when and how AI enters inquiry, it addresses a gap that becomes more visible in AI-accelerated classrooms without altering the core principles of established HCI and design pedagogies.

5. Conceptual Foundation of HI × AI

Across instructional contexts described earlier, recurring patterns in early-stage inquiry revealed a structural imbalance between human judgment and technological expansion. Students frequently relied on AI-generated language before fully articulating their own understanding, or alternatively engaged in intuitive reasoning without structured examination. These tendencies suggested the need for a disciplined alternation between human interpretation and AI-supported exploration.

The HI × AI pedagogy emerged in response to this need. Rather than positioning AI as a tool for accelerating output, the framework clarifies the distinct roles of Human Insights and Artificial Intelligence within early-stage inquiry.

Human Insights (HI) represent the interpretive capacities through which learners make sense of complex situations. These include surfacing assumptions, identifying stakeholders, determining relevance, and exercising judgment about what matters. Such capacities are reflective and situated, requiring deliberate articulation and revision (Schön 1983; Argyris & Schön 1978).

Artificial Intelligence (AI) contributes complementary capabilities, including rapid expansion of perspectives, pattern surfacing, synthesis across information sources, and iterative refinement of language. These capabilities can broaden inquiry beyond the learner's immediate viewpoint. However, AI does not establish meaning, determine priority, or assume ethical responsibility.

The relationship between HI and AI is not symmetrical. Human insight frames and interprets, while artificial intelligence expands and refines. This ordering aligns with research on reflective practice and scholarship on hybrid intelligence emphasizing human-guided machine collaboration (Dellermann et al. 2019).

This epistemic ordering was formalized into the HI × AI Infinity Loop, a four-part learning rhythm:

Frame → Expand → Interpret → Refine

In this rhythm, learners first frame the question using human judgment. AI is then used to expand evidence and perspective. Learners return to interpretation to determine what matters and revise assumptions. Finally, AI supports refinement of articulation and implications. The loop is repeatable rather than linear, enabling iterative strengthening of problem framing before solution development.

The next section demonstrates how this conceptual rhythm becomes operationalized as a structured classroom intervention.

6. Operationalizing HI × AI in the Classroom

The HI × AI framework is enacted in practice through a structured instructional workflow that makes early-stage reasoning explicit and assessable. This section describes how the four-part learning rhythm of Frame, Expand, Interpret, and Refine is implemented within classroom settings.

Teaching Objective

The instructional objective is to develop disciplined problem framing practices before students engage in solution development. Rather than rewarding speed or early ideation, the intervention intentionally slows inquiry and sequences AI use so that it reinforces interpretation rather than replacing it.

Enacting the Four-Part Rhythm

Classroom implementation follows the Infinity Loop explicitly.

Frame (HI)

Students articulate the problem in their own words, identify stakeholders, surface assumptions, and clarify context. AI use is withheld at this stage to ensure students confront ambiguity directly and articulate their understanding independently.

Completion signal: Students can state a concise problem framing, identify the primary stakeholder affected, and name at least one key assumption underlying their interpretation.

Expand (AI)

Once an initial framing exists, students use AI to broaden inquiry by surfacing alternative perspectives, contextual factors, and potential constraints. AI is used to widen the evidentiary field rather than generate solutions.

Completion signal: Students can identify at least one AI-generated insight that meaningfully challenges or deepens their original framing.

Interpret (HI)

Students return to human interpretation, evaluating which expanded insights matter, which assumptions require revision, and how the problem should be bounded or prioritized. This phase requires justification of change.

Completion signal: Students can explain why their revised framing is more meaningful or precise than the original and justify what has been included or excluded.

Refine (AI)

With direction clarified, AI is reintroduced to sharpen articulation and test implications. Students examine alternative phrasings, surface unintended consequences, and improve clarity. AI functions as a clarifier and challenger, not a decision-maker.

Completion signal: Additional AI iterations improve expression without altering the substantive framing.

The loop may be repeated as needed, reinforcing that problem framing is iterative rather than fixed.

Making Learning Visible

Each phase generates artifacts: initial framings, documented assumptions, AI-expanded perspectives, and revised problem statements. These artifacts make reasoning visible for both students and instructors.

First, instructors can assess how students articulate assumptions, revise interpretations, and justify changes in framing.

Second, students can trace how their framing evolved through structured progression rather than intuition or guesswork, strengthening ownership and accountability.

Sequencing Rather Than Restricting AI

A central design principle of this intervention is sequencing rather than prohibition. Students are encouraged to use AI, but only at specific phases of inquiry. This preserves human epistemic responsibility while leveraging AI's capacity for expansion and refinement.

By structuring when AI enters inquiry, the pedagogy resolves the instructional tension described earlier. AI is neither excluded nor allowed to dominate early-stage reasoning. Instead, it is embedded within a rhythm that protects judgment while amplifying understanding.

When Does the Loop Conclude?

The HI × AI Infinity Loop is not intended to run indefinitely. Iteration continues until further expansion no longer materially changes the learner's understanding and the framing reaches sufficient clarity and defensibility. In practice, this occurs when students can articulate their assumptions, explain why the problem meaningfully matters, identify affected stakeholders, and trace how their framing evolved. At that point, the loop transitions into downstream design work.

7. Observed Learning Effects

Evidence Context

The observations presented here are grounded in reflective practitioner practice across undergraduate courses, pre-college workshops, and organizational learning settings. Evidence consists of classroom observation, comparative review of student-produced artifacts across phases of inquiry, and longitudinal refinement across repeated course offerings. These observations do not claim experimental causality. Rather, they document recurring instructional patterns that informed iterative development of the HI × AI framework.

Illustrative Context: Essentials of Entrepreneurship

The most sustained evidence base derives from a recurring session within a large-enrollment undergraduate course, *Essentials of Entrepreneurship*, in which the author led one of eleven class sessions focused exclusively on problem framing. The session was delivered once per quarter across three academic quarters each year for four years, with approximately 100 students per offering.

At the beginning of the session, most students submitted loosely defined problem statements characterized by broad or symptom-focused framing, implicit solution language, unspecified stakeholders, unstated assumptions, and limited justification of why the problem meaningfully mattered. In later offerings, generative AI had often already been used prior to the session. These outputs were frequently polished but lacked documented framing logic or explicit interpretive justification.

Structured Intervention Effects

During the session, students enacted the HI × AI rhythm through sequenced exercises aligned with the Infinity Loop. Each phase generated visible artifacts, including initial framings, explicitly stated assumptions, AI-expanded perspectives, revised framing statements, and final articulated problem statements.

Comparative review of early and later artifacts within the same session revealed recurring instructional shifts. Revised framings demonstrated greater stakeholder specificity, clearer articulation of assumptions, more explicit boundary definition, and justified inclusion and exclusion decisions. Classroom discourse also shifted. Students spent more time interrogating the structure of the problem prior to ideation. Rather than converging rapidly on solution concepts, they revised framings iteratively and articulated reasons for those revisions.

Across repeated offerings, these patterns became more consistent as the sequencing of Frame, Expand, Interpret, and Refine was made increasingly explicit and assessable. The structured alternation between human interpretation and AI-supported expansion appeared to strengthen articulation and accountability during early-stage inquiry.

Micro-Illustrative Example of Framing Revision

The instructional shift can be illustrated through anonymized excerpts from a single session.

Initial submission (Frame phase):

“College students struggle with time management and need better tools to stay organized.”

The statement reflected a broad population and implied a solution orientation. Stakeholders were not differentiated, contextual constraints were unspecified, and assumptions remained implicit.

After completing the Expand and Interpret phases, the revised framing read:

“First-generation college students balancing part-time employment and full course loads experience difficulty prioritizing academic deadlines due to conflicting financial and family obligations.”

In this revision, the student narrowed the stakeholder group, clarified contextual conditions, and removed solution language. During interpretation, the student documented assumptions regarding access to institutional support and variability in work schedules, and justified the decision to bound the problem around competing obligations rather than tool availability.

This progression illustrates how structured alternation supports clearer articulation of stakeholders, assumptions, and boundaries before ideation begins.

Scaffolded Reinforcement: The Role of the Custom GPT

While the classroom intervention made reasoning visible during synchronous instruction, sustaining disciplined alternation required reinforcement beyond the session itself. To extend this structure, students were provided access to a custom GPT configured around the AMPLIFIED! Problem Statements workflow. The GPT did not introduce a new pedagogical model. Rather, it instantiated and enforced the sequencing logic of the HI × AI Infinity Loop as pedagogical infrastructure.

The system implemented phase-gated progression aligned with the four-part rhythm. During Frame, students were required to articulate a concise problem, identify a primary stakeholder, and document at least one explicit assumption before proceeding. During Expand, the GPT prompted contextual broadening and alternative perspective generation while restricting solution-oriented outputs. Progression to Interpret required students to explain how AI-generated insights altered, challenged, or refined their initial framing. Advancement was contingent upon interpretive justification rather than surface revision. During Refine, AI supported articulation clarity and implication testing but did not permit substantive reframing unless explicitly justified.

This configuration operationalized the epistemic ordering of HI × AI. Human articulation preceded AI expansion. Interpretation required justification. Refinement improved precision without displacing judgment.

Importantly, the pedagogical contribution does not depend on the GPT as a proprietary artifact. The HI × AI rhythm can be enacted through instructor-guided worksheets, structured LMS prompts, or facilitated studio exercises. The GPT serves as one implementation modality that automates sequencing and preserves metacognitive accountability, particularly in large-enrollment environments where individualized facilitation is constrained.

In this configuration, generative AI becomes both the subject of instruction and the structural constraint that reinforces disciplined use. Rather than accelerating premature convergence, the system is intentionally architected to slow reasoning at critical junctures and make framing decisions visible.

Together, these observations and scaffolded implementations position HI × AI not as a tool augmentation strategy, but as an instructional structure for preserving interpretive responsibility in AI-accelerated learning environments.

8. Implications for Downstream Design Methods

The HI × AI framework has implications for how existing HCI and design methods are taught in AI-enabled learning environments. Rather than altering established approaches such as research-through-design or UI-for-AI frameworks, HI × AI reshapes the conditions under which students enter those methods.

When problem framing is made explicit and structured prior to ideation and prototyping, downstream design activities become more intentional. Students approach ideation with articulated assumptions, identified stakeholders, and clearer boundaries. Prototyping becomes an exercise in testing defined framings rather than compensating for ambiguity.

In AI-enabled contexts, this sequencing is particularly consequential. Generative tools can accelerate artifact creation, but without disciplined framing they may amplify unexamined assumptions or superficial coherence. By structuring early inquiry through deliberate alternation between human interpretation and AI-supported expansion, instructors help ensure that AI use builds upon examined understanding rather than substituting for it.

This shift also clarifies assessment. Instead of evaluating only the quality of final artifacts, educators can assess the integrity of the framing process itself: how assumptions were surfaced, how evidence was expanded, how interpretations were justified, and how articulation evolved. Early-stage reasoning becomes visible and accountable.

HI × AI therefore functions as an instructional layer that precedes and strengthens downstream design work. By making problem framing explicit and sequenced, it supports more reflective engagement with AI while preserving the interpretive responsibilities central to HCI practice.

9. Conclusion

Generative artificial intelligence has altered the conditions under which students learn, design, and innovate. By lowering the effort required to conceive ideas, analyses, and plausible solutions, AI compresses early stages of inquiry that once required sustained sensemaking. In this environment, students can advance quickly without fully articulating the problem they are attempting to frame.

This paper argues that problem framing must be taught explicitly in AI-enabled HCI education. Rather than assuming learners will naturally develop disciplined framing practices, the HI × AI approach treats framing as a metacognitive capability that requires structured reflection and deliberate sequencing of AI use. By organizing instruction around the HI × AI Infinity Loop, educators make judgment visible and preserve human interpretive responsibility during early-stage inquiry.

Grounded in repeated instructional practice across multiple contexts, the framework demonstrates how deliberate alternation between human interpretation and AI-supported expansion can strengthen articulation of assumptions, deepen engagement with evidence, and clarify why a problem meaningfully matters. Responsible AI use emerges through structure rather than prohibition, aligning technological capability with human judgment.

By foregrounding problem framing as a teachable capability and positioning AI as an amplifier rather than an authority, HI × AI offers a practical instructional structure for helping students develop the discernment necessary to frame meaningful problems before attempting to solve them.

References

Amershi, Saleema, Dan Weld, Mihaela Vorvoreanu, Adam Fourney, Besmira Nushi, Penny Collisson, Jina Suh, Shamsi Iqbal, Paul N. Bennett, Kori Inkpen, Jaime Teevan, Ruth Kikin-Gil, and Eric Horvitz. 2019. Guidelines for human–AI interaction. In *Proceedings of the 2019 CHI Conference on Human Factors in Computing Systems (CHI '19)*, 1–13. https://doi.org/10.1145/3290605.3300233

Argyris, Chris, and Donald A. Schön. 1978. *Organizational Learning: A Theory of Action Perspective*. Addison-Wesley.

Dellermann, Dominik, Philipp Ebel, Matthias Söllner, and Jan Marco Leimeister. 2019. Hybrid intelligence. *Business & Information Systems Engineering* 61, 5, 637–643. https://doi.org/10.1007/s12599-019-00595-2

Dorst, Kees. 2011. The core of “design thinking” and its application. *Design Studies* 32, 6, 521–532. https://doi.org/10.1016/j.destud.2011.07.006

Gallant, Tricia Bertram, and David Rettinger. 2025. Teaching for integrity in the age of generative artificial intelligence. *Change: The Magazine of Higher Learning* 57, 6, 31–37. https://doi.org/10.1080/00091383.2025.2568354

Popović, Kevin. 2026. *AMPLIFIED!: How to Identify the Right Problem and Create the Right Solution with Human Insights × Artificial Intelligence*. CRC Press. In press.

Schön, Donald A. 1983. *The Reflective Practitioner: How Professionals Think in Action*. Basic Books.

Zimmerman, John, Jodi Forlizzi, and Shelley Evenson. 2007. Research through design as a method for interaction design research. In *Proceedings of the SIGCHI Conference on Human Factors in Computing Systems (CHI ’07)*, 493–502. https://doi.org/10.1145/1240624.1240704