Prudent AI Integration in Technical Decision-Making: A Case Study in Heat Pump Electrical Design

Robert Lee Gass, Purdue Polytechnic Institute, Purdue University – West Lafayette

Robert ”Bob” Gass is a certified project management executive and a Lean Six Sigma professional with a talent for empowering senior leadership teams struggling to bring order to projects to achieve their strategic goals. Bob cut his teeth on process, planning, and execution of highly complex projects during his career in the US Navy, working his way up through the ranks to becoming a decorated naval officer. Following his retirement from the Navy, Bob began environmental work with the Massachusetts Office of Coastal Zone Management, where he worked on compliance projects from dock and pier construction to the first dam removal in the state. At the Naval Undersea Warfare Center, Bob helped draft federally regulated documentation for the first controlled sinking of the aircraft carrier America, providing invaluable data to aid in constructing next-generation aircraft carriers. Throughout his career, Bob has had extensive cross-functional experience managing projects for high-profile clients, including Nordstrom, Target, Macy's, Morgan Stanley, and the U.S. Postal Service. With this expertise to his credit, Bob was recruited by Telaid to manage projects for its extensive, high-end, global client base. He was continually rewarded for his revenue-generating accomplishments with progressive promotions and enhanced scope of leadership responsibility. Bob's influence within Telaid was profound. He helped grow gross revenue from $30M to $100M+ during his tenure. He authored a white paper that served as the roadmap to Telaid's growth plan trajectory to $100M and instituted the company's first standardized project management process, along with the implementation of a robust risk management program. As Director of Project Planning, Bob implemented his vision of creating ”one standard PMO” of three independent project management offices. Telaid will continue to reap benefits from his efforts, such as synergies, collaboration, and cost savings. Known as the operational ”front man” for the Telaid sales team, Bob's philosophy to his team is ”Sales brings them to us. We keep them coming back.” Bob was most recently the Program Director of the Global Multisite Delivery Team at Black Box Corporation, overseeing a $50M+ portfolio and a team of project management professionals. He introduced best practices for the triple constraints, up-to-the-minute status reporting on all projects, and project planning protocols based on personally developed PTMAPP (plan to make a project plan), paving the way for world-class project performance. Bob has a Master of Marine Affairs from the University of Rhode Island and a Bachelor of Science in Physical Science from the United States Naval Academy. Bob is also a doctoral student at Purdue University, pursuing a Doctor of Technology degree focusing on Project Management. He combines his passion for project management with the discipline he acquired during his naval service to deliver flawless projects and unparalleled customer service that has become the hallmark of his brand.

Dr. Anne M Lucietto, Purdue Polytechnic Institute, Purdue University – West Lafayette

Dr. Lucietto has focused her research on engineering technology education and the understanding of engineering technology students. She teaches in an active learning style that engages and develops practical skills in the students. Currently, she is exploring a variety of topics related to her students' interests.

Prudent AI Integration in Technical Decision-Making: A Case Study in Heat Pump Electrical Design

Abstract

Artificial intelligence tools increasingly appear in engineering design practice, including tasks governed by safety codes and manufacturer requirements. In these settings, AI can generate plausible recommendations, but those recommendations often rest on incomplete assumptions and require careful human review. This paper analyzes how design reasoning develops when AI-generated suggestions are examined within a code-constrained electrical design task rather than accepted at face value.

The study examines an actual residential dual-zone heat pump installation subject to National Electrical Code requirements and manufacturer limitations. Using documented AI recommendations, subsequent revisions, and final validated configurations, the analysis examines how reasoning evolves as constraints are introduced, assumptions are challenged, and compliance is verified. Patterns of collaborative reasoning emerge, showing that design quality improves through iterative human oversight rather than through automated correctness checks.

From an educational perspective, the paper contributes a replicable instructional case for Engineering Technology programs. The case models how students and practitioners can engage AI as a reasoning partner while maintaining responsibility for safety-critical decisions. By foregrounding code interpretation, constraint-based judgment, and documented justification, the work supports ET instruction that reflects professional engineering practice in AI-supported design environments.

Introduction

Artificial intelligence tools now show up routinely in engineering work. Engineers use them to check calculations, suggest component sizes, interpret specifications, and accelerate early design decisions. In practice, these tools are often helpful, but they are also limited. Their outputs tend to reflect generalized design patterns and assumed conditions rather than the specific constraints that govern real installations. In safety-critical domains such as electrical design, those gaps matter. Small errors in conductor sizing, overcurrent protection, or code interpretation can create real hazards. Prior research in engineering decision making and organizational practice shows that users often fail to question automated recommendations closely enough, especially when the output appears confident or technically plausible [1], [2].

Electrical design governed by the National Electrical Code (NEC) highlights this challenge clearly. NEC compliance requires more than numerical thresholds; it depends on how equipment is installed, how circuits are configured, how loads interact, and how manufacturer specifications limit acceptable design choices. AI systems can suggest reasonable starting points based on common residential or commercial practice, but they do not inherently recognize when additional constraints apply unless those constraints are made explicit. Research on explainable and governed AI repeatedly emphasizes that automated systems struggle to align their internal reasoning with regulatory frameworks and safety standards without human supervision [3]–[6]. As a result, reliable electrical design still depends on human judgment, code knowledge, and accountability.

This paper examines these issues through an actual residential HVAC electrical installation in which AI is consulted during the design process. The project involves a dual-zone heat pump system that requires branch circuit design, conductor selection, and overcurrent protection in accordance with NEC Article 440 and manufacturer requirements. During the work, an AI tool provides initial recommendations that appear reasonable but omit key considerations, including minimum circuit ampacity, maximum overcurrent protection, conductor ampacity limits, voltage configuration, and installation environment. The human designer challenges these recommendations, supplies missing context, and verifies all final decisions directly against NEC language and equipment documentation. The case documents how compliant and safe design outcomes emerge through iterative human reasoning rather than through automated output alone.

The intent is not to evaluate AI accuracy or compare models. Instead, the case focuses on how reasoning develops when AI suggestions are treated as provisional and subject to challenge. Research on human-in-the-loop oversight and hybrid intelligence consistently shows that AI is most effective when its outputs are interrogated, justified, and reconciled with authoritative constraints [1], [2], [9]–[21]. In this work, explainability does not reside inside the AI system itself; it emerges through dialogue. When the designer asks the AI to explain a recommendation or justify it against code language, hidden assumptions surface and incorrect generalizations become visible.

This reasoning process has direct implications for Engineering Technology (ET) education. ET programs emphasize applied problem solving, interpretation of codes and standards, and hands-on understanding of real systems. As AI tools become common in professional practice, ET students need experience using them carefully rather than uncritically. Prior work in STEM and Engineering Technology education shows that inquiry-based and experiential approaches are especially effective for building judgment, contextual reasoning, and ethical awareness [7], [8]. The case presented here serves as a structured instructional model that allows students to practice NEC-based reasoning, evaluate AI-generated recommendations, and justify technical decisions in ways that align with professional responsibility.

The paper adopts a case-based analytic approach grounded in documented design artifacts, reconstructed AI-human dialogue, and formal NEC validation. Figures and tables function as evidence rather than illustration, capturing how recommendations change as constraints are introduced and reasoning becomes more precise. By tracing these changes, the paper shows how AI can support learned judgment without displacing it. The contribution is a replicable example that ET programs can integrate into courses on electrical systems, HVAC design, codes and standards, or capstone projects to prepare students for responsible use of AI in safety-critical engineering work.

Conceptual and Methodological Framing

This paper examines AI use in a specific type of engineering problem. The focus is safety-critical electrical design constrained by codes, standards, and manufacturer requirements. In this setting, design quality does not depend on creativity or optimization alone. It depends on careful interpretation of authoritative rules and on the engineer's ability to recognize when additional constraints apply. AI tools increasingly assist with this kind of work, yet their recommendations often rely on generalized assumptions rather than installation-specific conditions. The central concern of this study is not whether AI can generate plausible recommendations, but how engineering reasoning unfolds when those recommendations are examined, challenged, and corrected within a regulated design environment.

The conceptual focus of the paper centers on *collaborative reasoning*. The analysis treats AI output as a provisional contribution to the design process rather than as a solution. Attention is directed to how assumptions surface and where gaps appear. The analysis also examines how reasoning changes once contextual information is introduced, including voltage configuration, environmental exposure, NEC article requirements, and manufacturer limits. This perspective aligns with work on human-in-the-loop oversight and hybrid intelligence, which emphasizes that reliable outcomes emerge when automated reasoning is subjected to structured questioning and human accountability rather than accepted at face value [1], [2], [9]–[21]. In this framing, explainability is not viewed as an internal feature of the AI system but as a property of the interaction between the engineer, the AI output, and the governing standards.

The case is examined using an engineering case analysis approach grounded in documented work artifacts. The paper analyzes a single residential dual-zone heat pump installation, focusing on the electrical design tasks required to supply and protect the equipment in compliance with NEC Article 440 and manufacturer specifications. The evidence base for the analysis consists of AI-generated recommendations, successive revisions prompted by human questioning, explicit NEC citations, and the final validated design configuration. This approach intentionally avoids statistical generalization. It focuses instead on analytic generalization by identifying reasoning patterns typical of safety-critical electrical design and relevant to Engineering Technology education.

The case analysis proceeds by reconstructing the sequence of design decisions and examining how each recommendation holds up when evaluated against authoritative constraints. AI outputs are reviewed for completeness and compliance, and points of divergence between automated suggestion and code-aligned practice are documented. When discrepancies occur, the analysis traces how additional information reshapes the design. This method makes it possible to distinguish between recommendations that are technically acceptable, those that are partially correct, and those that fail once installation conditions are considered. The emphasis remains on reasoning moves rather than on the performance of the AI itself.

Figures and tables serve a central analytic role in this process. Rather than illustrating finished outcomes, they capture intermediate and final design states that emerge through human-AI interaction. Equipment specification figures establish the constraints that govern subsequent decisions. Tables compare AI-suggested configurations with revised and validated designs, highlighting how compliance is achieved through iterative reasoning. These artifacts function as evidence of decision-making under constraint and provide a concrete basis for interpreting how collaborative reasoning evolves over the course of the case.

From an educational standpoint, this framing positions the case as an instructional model rather than as an anecdote. Engineering Technology programs routinely teach students to apply codes, interpret specifications, and justify design choices in real installations. The reasoning patterns surfaced in this analysis—questioning assumptions, verifying automated output, and documenting compliance—map directly onto those instructional goals. The paper foregrounds the analytic lens, the evidence model, and the role of artifacts before introducing technical detail. This structure prepares the reader to engage with the case as a study of applied engineering judgment rather than as a narrative description of events.

Case Study Development and Project Context

This case examines the electrical design of a residential dual-zone heat pump installation, focusing specifically on branch circuit configuration, conductor sizing, and overcurrent protection. The work takes place within a regulated design environment governed by the National Electrical Code (NEC) and manufacturer electrical specifications. The case does not address thermal performance, refrigerant systems, or mechanical layout. Its scope is limited intentionally to the electrical decisions required to supply and protect the equipment safely and in full compliance with applicable standards.

The system consists of two outdoor condenser units and two indoor air-handling units serving separate zones within a single residence. Each outdoor unit requires a dedicated branch circuit sized according to the manufacturer's published minimum circuit ampacity (MCA) and maximum overcurrent protection (MOCP). The outdoor equipment also supplies power to its corresponding indoor unit, creating a combined electrical load that must be evaluated as a single installation scenario rather than as independent components. NEC Article 440 governs these configurations and establishes specific requirements for conductor sizing, overcurrent protection, and equipment disconnects. Compliance requires careful interpretation of both code language and manufacturer constraints.

The design task centers on determining whether AI-generated recommendations align with these requirements once installation context is accounted for. AI is used during the project to generate preliminary suggestions for conductor gauge, breaker rating, and general circuit configuration based on provided MCA and MOCP values. These outputs are treated as starting points, not decisions. They carry no authority and are not accepted without verification. Each recommendation is evaluated against NEC provisions, environmental conditions, voltage configuration, and manufacturer documentation before any design choice is finalized.

The role of the human designer in this process is to identify missing assumptions and introduce constraints that the AI does not initially consider. These include the requirement for two-pole protection on 240-V split-phase circuits, the distinction between indoor and outdoor wiring methods, wet-location conductor ratings, and the application of 125 percent continuous load rules where required. When discrepancies arise between AI output and code-compliant practice, the reasoning behind those discrepancies becomes the focus of analysis. The case therefore emphasizes how design quality improves through questioning and correction rather than through automated generation.

The evidence used in this analysis consists of documented AI recommendations, subsequent revisions prompted by human inquiry, explicit NEC citations, and the final validated electrical configurations. Rather than reporting results as conclusions alone, the case traces how reasoning evolves as additional information is introduced. This approach makes it possible to examine not only what the correct design becomes, but how that design emerges through interaction with automated suggestions under regulatory constraint.

Figures and tables play a central role in supporting this analysis. Equipment specification figures establish the electrical constraints imposed by the manufacturer and define the permissible design space. Configuration diagrams show how system architecture affects circuit sizing and protection decisions. Tables capture how AI-generated recommendations change as NEC requirements and installation conditions are applied, allowing direct comparison between initial suggestions, revised outputs, and final validated choices. These artifacts function as analytic evidence of decision-making under constraint rather than as summaries of finished work.

Although the case focuses on a single installation, it addresses a problem type that is common in Engineering Technology practice: interpreting code requirements, evaluating partial or generic recommendations, and documenting design justification in safety-critical contexts. The structure of the case reflects how similar tasks are encountered in the field and how they can be examined systematically. This framing supports later discussion of instructional use without treating the case as anecdotal or exceptional.

Technical Case Analysis

System Overview and Governing Constraints

The case begins with the definition of the electrical constraints imposed by the installed equipment. The system consists of two independent outdoor heat pump condenser units, each supplying a corresponding indoor air-handling unit. Electrical service to each outdoor unit is provided by a dedicated branch circuit, with the indoor unit powered through feed-through terminals internal to the outdoor equipment. This configuration requires evaluating each zone as a combined electrical load rather than as isolated components.

Figure 1. Zone 1 – 3 Ton System Electrical Specifications (Manufacturer Label/Data Plate)

Figure 1 presents the manufacturer-provided electrical specifications for the Zone 1 (3-ton) system. These specifications establish the permissible design boundaries and serve as the authoritative reference for subsequent decisions. The minimum circuit ampacity (MCA) is listed as 22 A, and the maximum overcurrent protection (MOCP) is listed as 30 A. These values immediately constrain both conductor sizing and breaker selection. Under NEC Article 440, branch-circuit conductors must be sized to carry not less than 125 percent of the equipment load, and overcurrent protection must not exceed the manufacturer-specified maximum. These requirements define the design space within which both AI-generated recommendations and human reasoning operate.

Although the indoor unit draws a smaller current, it is electrically dependent on the outdoor unit. The combined loading condition means that conductor sizing, breaker rating, and wiring method must account for the cumulative demand and the physical installation environment. This interaction between equipment architecture and code requirements becomes a critical point where generalized recommendations often fall short.

Secondary Zone Constraints and System Variation

Figure 2. Zone 2 – 2 Ton System Electrical Specifications (Manufacturer Label/Data Plate)

Figure 2 presents the corresponding manufacturer specifications for the Zone 2 (2-ton) system. While the overall configuration mirrors Zone 1, the electrical values differ in ways that materially affect design decisions. The MCA is lower, and the MOCP limit constrains breaker selection more tightly. These differences prevent direct reuse of Zone 1 decisions and require separate evaluation, despite the similarity in system topology.

The presence of two zones with different electrical characteristics introduces a realistic design challenge. The engineer must determine whether consistency in wiring practice is appropriate or whether each zone requires a unique solution. This tension between uniformity and compliance is a recurring issue in residential electrical design and is not something that automated tools consistently recognize without explicit instruction.

At this stage, the figures function as constraint-setting artifacts. They do not yet show decisions; they define what is legally and technically permissible. All subsequent reasoning unfolds within these boundaries.

Initial AI Recommendations and Identified Gaps

With the governing constraints established, the AI tool is queried for preliminary recommendations related to breaker size, conductor gauge, and general circuit configuration. The AI produces suggestions that align with common residential HVAC practice, including 12 AWG copper conductors and breaker ratings near the MCA values. On first inspection, these outputs appear plausible and familiar to anyone working in residential electrical systems.

Closer examination reveals several gaps. The AI does not consistently account for the 125 percent continuous load requirement when translating MCA values into conductor size. It does not recognize that 12 AWG copper is generally limited to 20 A in this application, nor does it reliably enforce the MOCP limits specified by the manufacturer. In some instances, the AI suggests single-pole breakers, overlooking the requirement for two-pole common-trip protection on 240-V split-phase equipment. Environmental considerations, such as wet-location conductors for outdoor equipment, are also absent from the initial output.

These omissions do not reflect incorrect arithmetic but incomplete reasoning. The recommendations assume default conditions rather than the specific installation context defined by the figures. This distinction becomes central to the analysis: the AI supplies a starting point, but not a complete or compliant solution.

Iterative Reasoning and Code-Based Correction

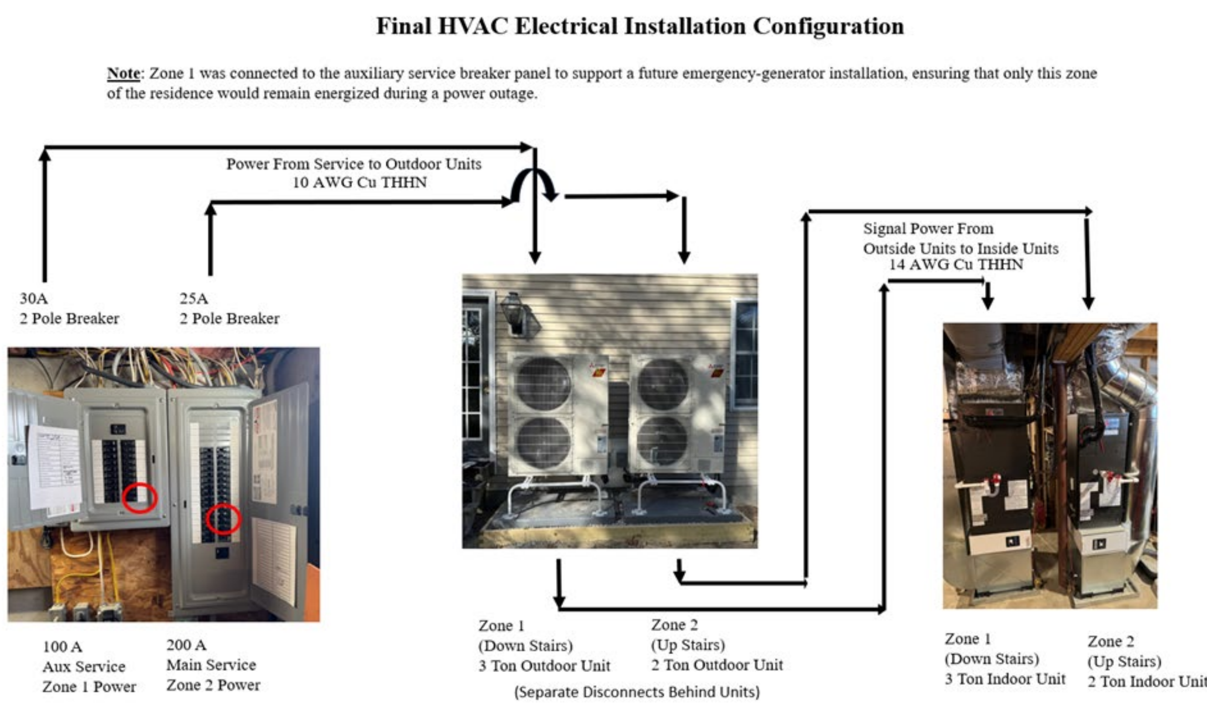

Figure 3. Final HVAC Electrical Installation Configuration

Human intervention begins by introducing the missing constraints one at a time. When the AI is asked to justify its conductor sizing against NEC ampacity tables, it revises its recommendations upward for Zone 1 to 10 AWG copper conductors to support a 30 A breaker compliant with the MOCP. When voltage configuration is clarified, the AI acknowledges the need for two-pole breakers. When environmental exposure is introduced, the wiring method is corrected from generic "copper cable" to THHN/THWN-2 conductors in conduit suitable for outdoor installation.

Figure 3 shows the final validated electrical configuration that results from this iterative process. The figure integrates manufacturer specifications, NEC requirements, and installation context into a single, code-compliant design. Importantly, this outcome does not emerge from a single corrected calculation but from a sequence of targeted challenges. Each correction narrows the solution space and removes assumptions that were not initially visible in the AI output.

The Zone 2 configuration follows the same reasoning process but produces a slightly different outcome due to the tighter MOCP limit. While the AI initially suggests stepping up to a 30 A breaker, NEC and manufacturer constraints limit protection to 25 A. Although 12 AWG conductors technically meet ampacity requirements for this zone, 10 AWG conductors are selected to maintain consistency across installations and reduce ambiguity during future inspection or modification.

E. Comparative Analysis of AI and Validated Design Outputs

Component	AI Initial	AI Revised	Final (Human-Validated)
Unit A Breaker	20 A	25 A	25 A (per MOCP)
Unit A Conductor	12 AWG	10 AWG	10 AWG Cu (per 125% MCA)
Unit B Breaker	25 A	30 A	30 A (per MOCP)
Unit B Conductor	12 AWG	10 AWG	10 AWG Cu (per 125% MCA)
All Conductors	NM-B (Romex)	THHN/THWN-2	THHN/THWN-2

Table 1. AI-Suggested, Revised, and Human-Validated Electrical Design Parameters

Table 1 captures the evolution of design decisions across both zones. It compares the initial AI recommendations, AI revisions prompted by targeted questions, and the final human-validated configurations. The table makes visible how compliance is achieved through successive refinement rather than through wholesale rejection of AI input.

Several patterns emerge. AI recommendations tend to align with common practice but under-constrain the problem. Revisions occur only after specific code language or installation conditions are introduced. Final validated designs consistently reflect conservative interpretation of NEC requirements, prioritizing compliance and safety over minimal material use. These patterns highlight the complementary roles of speed and judgment in collaborative reasoning.

Findings

The technical analysis demonstrates that AI tools can contribute meaningfully to early-stage electrical design by proposing plausible starting points. However, reliable outcomes depend on human oversight that introduces contextual information, verifies assumptions, and enforces code compliance. Figures and tables serve not as illustrations of success but as records of reasoning under constraint. The final design reflects not the authority of automation, but the discipline of iterative, code-based engineering judgment.

Interpretive Analysis of Collaborative Reasoning Patterns

This section interprets the technical case by examining how reasoning changes as automated recommendations are subjected to regulatory constraints and human judgment. Rather than recounting events, the analysis identifies recurring reasoning patterns that emerge when AI-generated suggestions interact with code-governed electrical design practice. These patterns become visible through the comparison of initial AI outputs, subsequent revisions, and final validated configurations documented in the figures and Table 1.

Pattern 1: Plausible Defaults Without Contextual Awareness

Across both zones, AI tools generate initial recommendations that align with common residential practice. Breaker sizes and conductor gauges match what might be expected in a typical installation, and the outputs appear internally consistent. This plausibility is precisely what makes these recommendations easy to accept without further scrutiny. However, the suggestions rely on assumed conditions rather than on the specific constraints imposed by manufacturer limits, NEC provisions, and site conditions.

This pattern shows that automated reasoning tends to operate on generalized defaults unless explicitly constrained. The AI does not reject compliance; it simply does not recognize when additional constraints apply. As a result, initial outputs fall within a broad range of acceptable practice but outside the narrower bounds required for full code compliance.

Pattern 2: Constraint Introduction as the Primary Catalyst for Correction

Meaningful revisions occur only when specific constraints are introduced into the reasoning process. When NEC ampacity rules, MOCP limits, voltage configuration, or environmental exposure are made explicit, the AI adjusts its recommendations accordingly. These corrections do not stem from improved calculation methods but from the addition of missing information that narrows the solution space.

This pattern highlights the role of constraints as catalysts rather than afterthoughts. The figures and Table 1 show that once constraints are articulated, design decisions converge quickly toward compliant solutions. Without those constraints, reasoning remains underspecified. Compliance therefore emerges not through optimization, but through disciplined boundary setting.

Pattern 3: Explainability Emerges Through Challenge, Not Transparency

In this case, explainability develops through dialogue rather than through inspection of internal model logic. When the AI is asked to justify a recommendation against NEC requirements or manufacturer limits, the assumptions underlying that recommendation become visible. Incorrect generalizations are uncovered not because the AI reveals its reasoning unprompted, but because it is forced to reconcile its output with external authority.

This pattern reframes explainability as an interactional process. The human designer functions as an active interrogator who demands justification in code-based terms. The resulting clarity does not reside in the AI system alone; it arises through the exchange that exposes assumptions and requires alignment with standards.

Pattern 4: Conservative Code Alignment as a Human-Anchored Decision

Final validated configurations consistently reflect conservative interpretations of NEC and manufacturer requirements. In situations where multiple compliant options exist, decisions favor clarity, consistency, and ease of inspection rather than minimal material use. This tendency appears most clearly in conductor selection, where larger gauges are retained even when smaller conductors might technically suffice.

This pattern illustrates how human judgment anchors design decisions in safety and accountability rather than procedural sufficiency. The AI does not resist conservative alignment, but it does not initiate it. Responsibility for that alignment remains firmly with the human designer, reinforcing the distinction between recommendation generation and accountable decision-making.

Synthesis of Patterns

Taken together, these patterns show that collaborative reasoning between humans and AI is neither linear nor autonomous. Automated suggestions accelerate early exploration but remain incomplete until subjected to targeted challenge. Design quality improves as assumptions are exposed, constraints are introduced, and reasoning is reconciled with authoritative standards. The figures and Table 1 function as evidence of this progression, documenting how compliance is achieved through structured interrogation rather than through automated correctness.

These identified reasoning patterns provide the analytic bridge between the technical case and its broader implications. They define what is learned from the case without extending into instructional prescription. The following section builds on these patterns to examine their relevance for Engineering Technology education and classroom implementation.

Educational Implications and Contribution to Engineering Technology Education

This paper contributes an instructional model for Engineering Technology (ET) education that centers on disciplined engineering judgment in AI-supported design work. Rather than treating AI as a tool to be mastered or evaluated, the case positions AI as a fallible participant in a code-governed reasoning process. The educational contribution lies in showing how students can be taught to interrogate automated recommendations, apply authoritative constraints, and document design decisions in ways that mirror professional electrical practice.

The reasoning patterns identified in the case align closely with competencies that ET programs already emphasize. ET curricula focus on applied problem solving, interpretation of codes and standards, and the ability to justify decisions under real constraints. The case illustrates how AI-generated recommendations often reflect plausible defaults and how those defaults must be challenged through deliberate constraint introduction. For students, this creates a structured opportunity to practice recognizing incomplete reasoning, identifying missing assumptions, and aligning technical decisions with NEC requirements and manufacturer specifications. These activities reinforce code literacy, contextual judgment, and professional accountability. These skills are central to ET identity.

The case also supports inquiry-based instructional approaches common in STEM and ET education. Instead of presenting students with a fully resolved design, the case invites them to work through successive iterations of reasoning. AI outputs provide an initial point of engagement, but learning occurs as students question those outputs, reconcile them with authoritative sources, and revise decisions accordingly. Prior work in STEM education shows that inquiry-driven and experiential learning environments are effective for developing critical thinking and judgment, particularly when students are required to explain and defend their reasoning rather than select correct answers [7], [8]. The structure of this case naturally supports that kind of engagement without requiring artificial errors or contrived scenarios.

From an instructional standpoint, the case fits readily into existing ET courses without major restructuring. It can be used in courses on electrical systems, HVAC design, codes and standards, or integrated into laboratory exercises and capstone projects. In a classroom setting, instructors can ask students to evaluate AI-generated recommendations, identify which NEC articles apply, justify conductor and breaker selections, and document why certain options are rejected. The figures and tables function as shared reference points that anchor discussion in concrete design artifacts rather than abstract principles.

Assessment in this context focuses on reasoning quality rather than outcome correctness alone. Instructors can evaluate how students articulate assumptions, cite code provisions, and justify design choices in writing or discussion. Comparison between AI-suggested configurations and validated designs makes student thinking visible and assessable without requiring new testing instruments. This approach aligns with broader efforts in ET education to emphasize professional reasoning, communication, and ethical responsibility alongside technical competence.

More broadly, the case helps prepare students for environments in which AI tools are present but cannot replace human accountability. By treating AI as a reasoning partner that must be questioned rather than trusted implicitly, students learn to maintain professional responsibility for safety-critical decisions. This perspective reinforces the role of the engineer or technologist as the final authority, even when automated assistance is available. The instructional value of the case lies not in teaching students how to obtain AI recommendations, but in teaching them how to evaluate those recommendations responsibly.

In sum, this work contributes to Engineering Technology education by providing a replicable, artifact-based case that models how disciplined reasoning develops in AI-supported electrical design. The contribution is practical rather than theoretical and aligns with established ET priorities: applied judgment, code compliance, and accountable engineering practice.

Conclusion

The growing presence of AI in engineering design raises a fundamental challenge in safety-critical, code-governed work: automated recommendations can accelerate analysis, but they do not replace the need for disciplined human judgment. Electrical design remains bounded by authoritative standards, installation context, and professional accountability, all of which require reasoning that extends beyond plausible defaults.

This case demonstrates how design quality improves when AI output is treated as provisional and subjected to structured challenge. As constraints are introduced and assumptions are tested against code language and manufacturer requirements, reasoning becomes more precise and outcomes converge toward conservative, compliant solutions. The analysis shows that compliance does not emerge from automation alone, but from the iterative reconciliation of automated suggestions with human oversight grounded in standards.

From an educational perspective, the case provides a concrete instructional model for Engineering Technology programs. It illustrates how students can be taught to work with AI as a reasoning partner rather than an authority, using codes and specifications to evaluate, refine, and justify design decisions. By foregrounding judgment, accountability, and constraint-based reasoning, the case supports ET instruction that reflects real engineering practice and prepares students to engage responsibly with AI-supported workflows.

References

[1] Y. Shrestha, S. Ben-Menahem, and G. von Krogh, "Organizational decision-making structures in the age of artificial intelligence," *California Management Review*, vol. 61, no. 4, pp. 66–83, 2019.

[2] A. T. Mohan, B. Seligmann, R. Van Coile, and D. Lange, "The assumptions trap: Implications for fire safety engineers," *Journal of Physics: Conference Series*, vol. 3121, no. 1, 2025.

[3] H. Devarajan, "Explainable AI for cloud-based machine learning: Interpretable models and transparency in decision making," *Tuijin Jishu / Journal of Propulsion Technology*, vol. 45, no. 2, pp. 2886–2894, 2024.

[4] M. Karim *et al.*, "Explainable artificial intelligence for bioinformatics: Methods, tools, and applications," *Briefings in Bioinformatics*, vol. 24, no. 5, 2023.

[5] I. Nejadgholi, S. Kiritchenko, K. Fraser, and E. Balkır, "Concept-based explanations to test for false causal relationships learned by abusive language classifiers," in *Proc. ACL Workshop on Online Abuse and Harms*, Jul. 2023.

[6] A. Das and P. Rad, "Opportunities and challenges in explainable artificial intelligence (XAI): A survey," *IEEE Access*, vol. 8, pp. 22456–22463, 2020.

[7] A. M. Lucietto and D. Peters, "Engineering technology and engineering: Incorporating the humanities into the classroom," *Frontiers in Education*, vol. 9, Art. no. 1534407, 2024.

[8] A. M. Lucietto, L. Russell, and E. Schott, "STEM educators: How diverse disciplines teach," *Journal of STEM Education*, vol. 19, no. 3, pp. 40–46, 2018

[9] S. Morandini *et al.*, "Assessing the transparency and explainability of AI algorithms in planning and scheduling tools," in *Human Interaction and Emerging Technologies*, vol. 111, pp. 609–619, 2023.

[10] N. Alberto, "XAI and sustainability: Unifying regulatory standards and solutions for the future of environmental management," EarthArXiv, Nov. 27, 2024. [Online]. Available: https://eartharxiv.org/repository/view/8125/

[11] F. Olinmah *et al.*, "Designing regulatory risk reporting frameworks using automation tools," *International Journal of Scientific Research in Science, Engineering and Technology*, vol. 11, no. 5, 2024.

[12] N. Khan *et al.*, “Explainable and robust artificial intelligence for trustworthy resource management,” *IEEE Communications Magazine*, vol. 62, no. 4, pp. 50–56, 2024.