The Internalize-Reference-Delegate (IRD) Framework for AI in Power Systems Education

Hantao Cui
Department of Electrical and Computer Engineering
North Carolina State University
Raleigh, NC 27606, USA
hcui9@ncsu.edu

Kevin Tomsovic
Department of Electrical and Computer Engineering
Clemson University
North Charleston, SC 29405, USA
ktomsov@clemson.edu

***Abstract*—Knowledge work differs fundamentally from physical automation because it requires reasoning to justify outcomes. Large language models provide valuable support in education, but they may also impede the development of engineering judgment, which is critical for power systems engineers. This paper introduces the Internalize-Reference-Delegate (IRD) framework, which classifies learning tasks along two dimensions: the need for internalization and the difficulty of verification. Each dimension has three levels. For internalization, the levels are knowledge that must be internalized, procedural skills that can rely on AI reference, and tasks that can be delegated. For verification difficulty, the levels are easy to verify, require expert knowledge to verify, and hard to verify. These categories form a three-by-three matrix that helps instructors and learners decide how to integrate AI into learning activities. A project on Newton-Raphson power flow illustrates how treating AI generated suggestions as hypotheses that require mathematical verification can strengthen the understanding of fundamental concepts while still benefiting from the efficiency of AI.**

***Index Terms*—Power systems education, large language models, engineering education, AI literacy, educational technology, human-AI collaboration**

I. Introduction

The integration of large language models (LLMs)-based AI into power systems education presents a paradox. These tools surface definitions, worked examples, and code templates in seconds, yet they can also weaken the reasoning training that underpin engineering judgment. Educators and students therefore face a shared challenge: determine where AI assistance improves learning efficiency, and where it disrupts the practice needed to internalize theory. This paper aims to clarify that boundary rather than assuming AI belongs everywhere or nowhere.

Recent empirical evidence reveals the severity of AI-induced skill degradation. Studies show AI assistance accelerates skill decay and hinders skill development without users' awareness and creates dangerous illusions of understanding [1]. In medical diagnostics, physicians using AI tools showed degraded detection abilities when AI became unavailable [2]. These findings carry profound implications for power systems education, where engineers must maintain technical expertise throughout decades-long careers.

The fundamental challenge lies in the nature of knowledge work, including engineering. Physical automation is widely accepted because the outcomes can be readily verified. For instance, a washing machine either cleans the clothes or it does not. Power systems understanding, however, requires building mental models of invisible phenomena. Students must develop intuition for what is reactive power, why oscillations happen, and why new protection schemes are needed for converters. These cognitive structures only develop from working through problems, correcting mistakes, and recognizing patterns. When students bypass this deliberate practice through AI assistance, they lose the capacity to recognize when something "does not look right."

This vulnerability becomes critical when we consider how LLMs work. They are probability-based text predictors that generate likely tokens, rather than reasoning based on first principles. LLMs cannot validation its outputs but only restate patterns from training data that may have assumptions and context. For power systems education, every LLM output is fundamentally a hypothesis awaiting verification. Such verification requires precisely the domain knowledge that students are trying to develop.

These limitations create a joint responsibility of the instructor and the learner. Instructors must distinguish between tasks where AI scaffolding is useful and manual (conceptual) work designed for skill development. Students must practice *intellectual honesty* and recognize that accepting AI-generated explanations without working through actual problems is not a shortcut but a detour. The proposed framework supports this shared accountability by mapping activities to appropriate AI integration levels with explicit verification requirements.

This paper proposes the Internalize-Reference-Delegate (IRD) framework that maps power systems learning activities to proper AI integration levels. Rather than binary prescriptions of AI use or prohibition, an actionable matrix is presented on two dimensions: internalization need and verification difficulty. Our contributions are threefold: (1) a three-by-three decision matrix that categorizes tasks on the two dimensions of internalization need and verification difficulty, (2) verification protocols tailored to power engineering, and (3) a Newton-Raphson power flow project as an implementation example that preserves skill development while benefiting from AI. The framework builds on established educational taxonomies while incorporating unique requirements to power engineering, though empirical validation remains future work.

979-8-3315-8138-1

The remainder of this paper is organized as follows. Section II reviews the cognitive science foundations and the unique challenges in power systems education. Section III presents the IRD framework with detailed task categorization and verification protocols. Section IV discusses implementation strategies. Section V draws the conclusions.

II. Cognitive Foundations and Domain-Specific Requirements

A. Evidence of AI-Induced Skill Degradation

The impact of AI assistance on human expertise follows patterns documented across multiple safety-critical domains. While this research area is relatively new, a 2025 MIT study [3] demonstrated through EEG measurements that students using ChatGPT for writing essays show significantly weaker neural connectivity and accumulated "cognitive debt" compared to those writing manually, with effects persisting even after discontinuing AI use. In medical diagnostics, physicians using AI tools showed degraded ability to detect precancerous issues after just three months of use [2]. These effects manifest severely in domains requiring reasoning, which are precisely the cognitive components central to knowledge work, including power systems engineering.

The mechanism underlying skill degradation relates to cognitive offloading and reduced neural activation during problem-solving. Neuroeducation research demonstrates that pattern recognition and expertise development require deliberate practice to create neural pathways, with synaptic changes occurring only in response to active engagement [4]. When students receive AI-generated solutions without engaging in effortful reasoning, the neural mechanisms essential for learning remain dormant. However, students believe they comprehend material because they can produce correct-looking outputs, yet lack the internal representations needed to reason independently.

B. Prior Work on LLMs in Power Systems Education

Several studies have directly examined LLM applications in power engineering. It is shown that LLMs show varying success rates for power systems tasks, with GPT-4 achieving 78% accuracy on routine unit commitment problems but only 31% on new optimization challenges [5]. Their human-in-the-loop framework proved essential for complex tasks where LLMs lack cutting-edge knowledge.

Another work provides concrete evidence of LLM capabilities and limitations in power systems education, showing that GPT-4 can generate functional OpenDSS circuits for load flow studies when given adequate guidance [6]. The work includes detailed verification protocols for students to check. The published transcripts document specific errors that students must identify, such as incorrect impedance values and unrealistic reactive power assumptions. In addition, broader implications of LLMs for the electric energy sector are discussed, including data governance, security, and human-in-the-loop requirements as critical considerations [7].

C. Limitations of Current Approaches

While existing studies demonstrate LLM potential in specific power systems tasks, they reveal fundamental gaps in addressing educational needs. Current approaches treat AI as either a productivity tool or a knowledge repository, overlooking how it disrupts the cognitive development process itself. In [5], accuracy metrics focus on final answers rather than reasoning pathways. In [6], verification protocols assume students already possess the judgment to identify errors. None address how students *develop* this judgment in the first place when AI provides immediate answers.

The existing literature also fails to distinguish between different types of power systems knowledge. Generating syntactically correct OpenDSS models differs fundamentally from understanding why a particular network parameter might experience low voltage issues. Current frameworks evaluate AI performance on discrete tasks but neglects how tasks contribute to integrated understanding. This gap becomes critical when students use AI to complete assignments without developing the conceptual foundations.

D. The Two-Layer Structure of Power Systems Knowledge

Power systems education operates across two interconnected layers that AI affects differently. The theoretical foundation encompasses the mathematical and physical principles governing electrical networks. This includes circuit analysis, differential equations for dynamics, optimization, and control theory. Students must internalize abstract relationships to understand why systems behave as they do. When AI provides formulas or derivations, students who accept without working through them will miss the development of intuition and are likely to fail in complex tasks where AI hallucinates.

The implementation layer translates theoretical understanding into practical systems through simulation software, programming, and hardware interfaces. This includes writing Python scripts for data analysis, configuring simulation models, and setting up hardware tests. While AI can generate code snippets or configuration files, it cannot always convey the judgment needed to recognize, e.g., numerical issues or physically unrealizable solutions. Students learn these subtleties through debugging their own implementations and seeing how theoretical issues manifest as practical failures.

The critical challenge lies in the coupling between these layers. Real expertise requires moving fluidly between abstract mathematical concepts and their concrete implementations. A protection engineer must understand both fault transients and algorithms in digital relays to approximate these dynamics with discrete samples. Students who use AI to jump between layers without understanding the connections develop fragmented knowledge, on which problematic designs can be built.

E. The Verification Problem in Knowledge Work

The verification challenge in power systems education extends beyond Bainbridge's classical ironies of automation [8]. While Bainbridge identified that operators must maintain skills to take over when automation fails, educational

979-8-3315-8138-1

TABLE I
Decision Matrix for AI Integration in Power Systems Education

Internalization Need	Easy to Verify	Requires Expertise	Hard to Verify
Must Internalize (Core concepts, intuition)	**AI for reference after mastery** 1) Basic formulas, 2) Standard equations, 3) Unit conversions *Example: phasor; per unit conversion; transformer pi-circuit model*	**Manual first, then AI review** 1) Problem-solving methods, 2) Analysis techniques, 2) Design procedures *Example: Power flow equations after deriving them manually*	**Manual practice essential** 1) Physical intuition, 2) Pattern recognition, 3) Failure mode detection *Example: Recognizing unstable equilibrium points in phase plots*
Can Reference (Standards, methods)	**Full AI with spot checks** 1) NERC standards, 2) Typical values, 3) Component specs *Example: Under voltage limit to trigger load shedding*	**AI with detailed verification** 1) Simulation setup, 2) Model configuration, 3) Case studies *Example: Generator subtransient parameters*	**AI with peer review** 1) Technical reports, 2) Design documentation, 3) Literature reviews *Example: Derivation of holomophic embedding*
Can Delegate (Tools, automation)	**AI freely with output checks** 1) Data formatting, 2) Plot generation, 3) Code syntax *Example: Matplotlib visualizations*	**AI with domain validation** 1) Script generation, 2) Batch processing, 3) Data analysis *Example: Load profile data filtering scripts*	**AI with expert oversight** 1) Optimizationg, 2) Transient stability simulation *Example: Coding an optimization model in Python*

contexts face a more fundamental problem. Students must simultaneously develop the expertise needed for verification while using tools that bypass the very experiences that build such expertise.

Consider a student using AI to analyze transient stability. The AI might correctly identify that a fault causes loss of synchronism and suggest increasing generator inertia as a solution. Verifying this recommendation requires understanding swing equations and critical clearing times. But these concepts only become meaningful through repeatedly solving stability problems, observing how different parameters affect outcomes, and developing intuition for system behavior. When AI provides answers directly, students lose these formative experiences.

The delayed feedback inherent in power systems compounds this problem. A mis-configured protection scheme might function correctly for years until a specific fault pattern exposes the error. A poorly tuned controller might seem adequate until system operates inn abnormal conditions. Unlike software development where bugs surface quickly through testing, power systems decisions may not reveal their flaws until rare events occur. Students trained with AI assistance may graduate without experiencing enough failure modes to develop proper verification instincts.

Recent frameworks for AI literacy in education emphasize critical evaluation skills [9], [10], yet remain generic across disciplines. The ALTL framework [9] identifies technical understanding, evaluative skills, practical application, and ethical considerations as core competencies. However, these frameworks provide limited guidance for fields where verification requires years of accumulated pattern recognition. Power systems education needs domain-specific protocols that account for both the theoretical depth and practical judgment required for safe system operation.

III. A Decision Framework for AI Integration in Power Systems Education

The limitations identified in current approaches necessitate a framework that moves beyond binary prescriptions of AI use. Existing studies focus on what AI can do technically but fail to address when its use supports or undermines learning objectives. This section presents a decision matrix based on two critical dimensions: the need for knowledge internalization and the difficulty of verification. Unlike frameworks that simply categorize tasks by cognitive level, our approach recognizes that power systems expertise requires both memorized foundations and the judgment to know when external tools are appropriate.

A. Task Categorization and Decision Matrix

We propose categorizing power systems learning activities along two dimensions. The first dimension, internalization need, distinguishes between knowledge that must become automatic, knowledge that can be referenced when needed, and tasks that can be delegated to tools. The second dimension, verification difficulty, captures how much expertise is required to confirm that an output is correct. This yields nine categories with distinct guidance for AI integration.

The matrix deliberately avoids terms like "prohibited" since enforcement is impractical and counterproductive. Instead, it provides guidance on when manual work is essential for building expertise versus when AI can enhance efficiency without compromising understanding. The key insight is that verification difficulty often determines whether AI assistance helps or hinders learning. Tasks that are easy to verify can leverage AI more freely because students can quickly identify errors. Tasks requiring expertise to verify demand more caution since students may accept incorrect outputs without recognizing problems.

B. Temporal Progression Through the Matrix

While the matrix defines appropriate AI use for different task types, students' readiness for these guidelines evolves across their educational journey. During initial exposure to new concepts, even tasks classified as "Can Delegate" should be performed manually to establish foundational understanding. As students progress through practice and reinforcement phases, they can begin following the matrix guidance more directly. By the mastery and application stage, typically in senior design projects, students can fully leverage AI according to the matrix categories.

This temporal overlay means that a single task might migrate across the matrix during a student's education.

979-8-3315-8138-1

Consider the derivation of power flow equations : they begin as "Must Internalize / Practice Essential" in Power System Analysis I, since students need to develop intuition for nonlinear equations and Jacobian matrix. The same problem transition to "Can Reference / AI with Detailed Verification" for routine cases in graduate-level computational methods once fundamental concepts are internalized. Eventually, practicing engineers place these tasks in "Can Reference / Requires Expertise" as they maintain core understanding while using AI for computational efficiency.

The matrix thus provides not a static prescription but a framework that adapts to students' developmental stages and accumulating expertise. Instructors should communicate explicitly where students are in this progression for each topic area. A statement like "We're treating power flow as 'Must Internalize' this semester, but by next year you'll be able to use AI tools with verification" helps students understand why certain restrictions exist and how their relationship with AI tools will evolve.

C. Verification Protocols by Category

Each cell in the matrix requires specific verification approaches tailored to both the knowledge type and the verification challenge. These protocols ensure that AI assistance enhances rather than replaces understanding.

For tasks in the "Easy to Verify" column, verification can often be done through simple checks that students learn early. When using AI to look up generator reactance values, students compare against typical ranges from textbooks or datasheets. For data formatting tasks, they visually inspect outputs for obvious errors. These verifications require minimal expertise and can be performed quickly, making AI assistance relatively safe even for new learners.

The "Requires Expertise" column demands more sophisticated verification protocols. When AI assists with simulation setup, students must check multiple aspects that require domain knowledge. For a transient stability simulation, this includes verifying the initial conditions to be within limits, the load converted to the correct model, etc. Students must also run benchmark cases with known solutions to validate setup. This level of verification requires understanding both the physical system and the simulation tool that develops through coursework and practice.

Tasks in the "Hard to Verify" column present the greatest challenge and require the most stringent protocols. Consider reactive power planning for a sophisticated AI capable of calling power flow tools. Checking the voltage profile, AI might suggest placing capacitor banks at a low voltage bus. But the capacitance may be large enough to push up the nose point and cause the system to collapse at a high voltage level. Students need to quantitatively analyze the reactive support and verify the PV curves to check for potential overcompensation issues. Such implications are inconspicuous to the untrained eye. It is thus crucial for instructors to provide scaffolded review processes to ensure students develop the expertise even though AI suggestions may appear plausible.

The verification burden also varies by row in the matrix. "Must Internalize" tasks require students to demonstrate understanding without AI assistance before they can use AI for reference. Verification here focuses on conceptual understanding rather than just correct answers. Students might be asked to explain why a particular power flow solution is infeasible without referring to AI outputs, demonstrating internalized knowledge of system constraints. "Can Reference" tasks allow concurrent AI use but require documentation of what was referenced and how it was validated. "Can Delegate" tasks need only output verification, though students should still understand the general approach even if they delegate execution details.

These verification protocols serve a dual purpose. They ensure technical correctness while developing the metacognitive skills students need to work effectively with AI throughout their careers. By explicitly practicing verification at different levels of complexity, students learn to calibrate their trust in AI outputs and recognize when human judgment is essential.

IV. Implementation Recommendations

The framework presented in Section III requires translation from conceptual matrix to classroom practice. This section demonstrates concrete implementation through a detailed course example, provides integration strategies across curricula, and addresses the practical challenges educators will face when adopting this approach.

A. Implementation Example: Power Flow Project

Consider a senior-year power systems analysis course focusing on Newton-Raphson power flow implementation. This semester-long project exemplifies how the framework guides AI integration through different phases of learning. The instructor provides Python code skeletons, and students build a complete power flow solver while learning when AI assistance helps versus hinders understanding.

Week 1-3 introduce Python programming fundamentals, classified as "Can Delegate / Easy to Verify." Students learn numpy operations, object-oriented programming with classes and methods, and file I/O. AI assistance is encouraged for explaining programming concepts and debugging syntax errors since code functionality is immediately testable. The syllabus states: "Use AI freely for programming questions. Every line of code can be verified by running it."

Week 4-8 focus on mathematical modeling, classified as "Must Internalize / Requires Expertise." Students derive power flow equations on paper and implement transmission line π-models, generator PV buses, and load PQ buses without AI assistance. Critical concepts like sign conventions for power injection must be understood manually, since AI frequently hallucinates inconsistent ones. Students implement admittance matrix formation and understand the physical meaning of elements. A typical assignment: "Derive the complex power balance equations for a PQ bus. Implement the model in Python without AI assistance. Explain your sign convention choice and why consistency matters."

Week 9-12 cover Newton-Raphson solver implementation, categorized as "Can Reference / Requires Expertise." Students may use AI suggestions for sparse matrix techniques

979-8-3315-8138-1

and numerical methods, but must treat all suggestions as hypotheses. When debugging convergence issues, AI might suggest checking flat start conditions or adjusting tolerances. Students must verify these suggestions through mathematical reasoning and physical understanding. The rubric weights verification at 40%: "Document each AI suggestion used, explain why it might work, and show mathematical or physical evidence supporting or rejecting it."

Week 13-15 involve testing on IEEE test systems, placed in "Can Reference / Easy to Verify." Students validate their implementation against published results, debug edge cases, and analyze convergence patterns. AI can assist with visualization and data analysis, but students must explain discrepancies through physical reasoning.

B. Curriculum Integration and Immediate Actions

The framework requires coordinated implementation across the four-year curriculum. Years 1-2 focus on foundations with content primarily "Must Internalize," *limiting* AI to basic tasks like plotting. Year 3 introduces strategic AI use as students have internalized fundamentals. Year 4 aligns with industry practice while maintaining manual requirements for safety-critical topics.

For immediate deployment, modify syllabi with explicit matrix mappings. Template: "Module 3 (Symmetrical Components) is classified as Must Internalize / Requires Expertise. Complete manual calculations for all homework problems. AI may be used only to verify final answers after showing full work." Assessment should use the two-part structure: Part A (40% weight) allows AI with mandatory verification documentation, Part B (60% weight) tests understanding without AI. This design ensures students cannot bypass core learning since they face closed-AI evaluations.

Departments should maintain a shared framework mapping updated each semester. When Circuit Analysis I treats phasor analysis as "Must Internalize," Power Systems I must acknowledge this progression before allowing AI assistance. Verification rubrics weight process over results: "Physical Validity Check (15 points): Full credit for identifying and explaining physical constraint violations in AI output."

C. Challenges and Continuous Improvement

Academic integrity concerns diminish when verification becomes central to assessment. The framework reduces cheating incentives since students must defend AI-generated solutions and face closed-AI evaluations. Policy: "Using AI without acknowledgment is academic dishonesty. Using AI with poor verification shows incomplete understanding. Using AI with thorough verification demonstrates competency."

Continuous improvement requires tracking three metrics: correlation between AI usage and manual performance, student core competencies, and employer feedback on graduate preparedness. Annual reviews will adjust matrix classifications as AI capabilities evolve. What requires expertise to verify today may become easily checkable tomorrow. The goal remains preparing students for careers where AI is ubiquitous but engineering judgment irreplaceable.

V. Conclusion

This paper presents a decision framework that categorizes power systems learning tasks by internalization need and verification difficulty, providing educators practical guidance for AI integration. Evidence shows unstructured AI use accelerates skill decay in safety-critical domains [1], [2], yet different knowledge types require different approaches. Tasks demanding internalized understanding require manual practice, while procedural work can leverage AI once students develop verification capabilities. Success will be measured not by efficiency with AI tools, but by whether graduates possess the engineering judgment to recognize issues and intervene when automation fails.

Finally, to advance AI for education, sustained efforts are needed to develop domain-specific applications. General-purpose LLMs are powerful, yet their domain-specific outputs can be superficial. Two directions are open to research. First, problem-specific LLMs can be developed and fine-tuned using curated materials as training data (e.g., models specialized in voltage stability or user-defined model development). Second, LLMs can be coupled with power system analysis subroutines that provide well-defined APIs and informative outputs (for instance, calling ANDES for transient stability simulation, similar to "tool calls" in modern AI-assisted coding tools). Both directions are promising and require community investment for further research.

References

[1] B. N. Macnamara *et al.*, "Does Using Artificial Intelligence Assistance Accelerate Skill Decay and Hinder Skill Development Without Performers' Awareness?," *Cognitive Research: Principles and Implications*, vol. 9, no. 46, 2024, doi: 10.1186/s41235-024-00572-8.

[2] C. Natali, L. Marconi, L. D. Dias Duran, and F. Cabitza, "AI-Induced Deskilling in Medicine: A Mixed-Method Review and Research Agenda for Healthcare and Beyond," *Artificial Intelligence Review*, vol. 58, no. 52, 2025, doi: 10.1007/s10462-025-11352-1.

[3] N. Kosmyna *et al.*, "Your Brain on ChatGPT: Accumulation of Cognitive Debt when Using an AI Assistant for Essay Writing Task," *arXiv preprint arXiv:2506.08872*, 2025.

[4] K. Pradeep, R. S. Anbalagan, A. P. Thangavelu, S. Aswathy, V. G. Jisha, and V. S. Vaisakhi, "Neuroeducation: Understanding Neural Dynamics in Learning and Teaching," *Frontiers in Education*, vol. 9, 2024, doi: 10.3389/feduc.2024.1437418.

[5] R. Li *et al.*, "LLM-based Frameworks for Power Engineering from Routine to Novel Tasks," *arXiv preprint arXiv:2305.11202*, 2023, [Online]. Available: https://arxiv.org/abs/2305.11202

[6] R. S. Bonadia, F. C. L. Trindade, W. Freitas, and B. Venkatesh, "On the Potential of ChatGPT to Generate Distribution Systems for Load Flow Studies Using OpenDSS," *IEEE Transactions on Power Systems*, vol. 38, no. 6, pp. 5965–5968, 2023, doi: 10.1109/TPWRS.2023.3315543.

[7] S. Majumder, L. Dong, F. Doudi, and others, "Exploring the Capabilities and Limitations of Large Language Models in the Electric Energy Sector," *Joule*, vol. 8, no. 6, pp. 1544–1549, 2024, doi: 10.1016/j.joule.2024.05.009.

[8] L. Bainbridge, "Ironies of Automation," *Automatica*, vol. 19, no. 6, pp. 775–779, 1983, doi: 10.1016/0005-1098(83)90046-8.

[9] M. Kassorla, M. Georgieva, and A. Papini, "AI Literacy in Teaching and Learning: A Framework for Higher Education," 2024. [Online]. Available: https://www.educause.edu/content/2024/ai-literacy-in-teaching-and-learning/executive-summary

[10] UNESCO, "AI Competency Framework for Students," 2024. [Online]. Available: https://www.unesco.org/en/articles/ai-competency-framework-students

979-8-3315-8138-1