American Journal of Advanced Multidisciplinary Innovation and Research (AJAMIR)

E-ISSN: XXXX-XXXX ● Website: www.ajamir.com ● Email: editor@ajamir.com

Adaptive Prompting Systems for Domain-Specific Professional Training

Dr. Marcus Ellison

Associate Professor, Trinity College Dublin, Ireland

Abstract

Professional training increasingly requires learning systems that can respond to differences in prior knowledge, task complexity, professional role, error patterns, regulatory context, and learner confidence. Generative artificial intelligence has created new possibilities for individualized explanations, simulated professional scenarios, formative feedback, reflective questioning, and practice dialogue, yet many implementations continue to rely on static prompts that deliver broadly similar instructional behavior regardless of the learner's evolving state. Adaptive prompting offers a different model in which instructional prompts are dynamically selected, modified, or routed according to domain knowledge, competency level, observed errors, interaction history, assessment evidence, and training objectives. This study develops a Domain-Specific Adaptive Prompting Framework for Professional Training (DAPF-PT) that combines learner-state modeling, domain-grounded knowledge, prompt routing, progressive scaffolding, formative assessment, source verification, and human professional oversight. Because verified training-intervention data were not available for the present manuscript, the study adopts an explicitly simulation-based methodological design.

A synthetic analytical corpus of 480 hypothetical professional-training interactions was distributed across four prompting architectures: static domain prompting, rule-based adaptive prompting, learner-state prompt routing, and domain-specific adaptive prompting with professional oversight. Six dimensions—domain fidelity, learner-state sensitivity, scaffolding appropriateness, feedback quality, assessment alignment, and governance reliability—were used to construct a Professional Training Readiness Index. The simulated index increases from 52 for static prompting to 91 for the integrated adaptive architecture. Contemporary evidence supports the potential of personalized prompting while also showing that general-purpose language models do not automatically reproduce the adaptivity of established intelligent tutoring systems. The paper therefore argues that effective adaptive prompting must be deliberately engineered around professional competencies rather than treated as unrestricted conversational personalization. The proposed framework positions generative AI as a supervised instructional partner capable of varying challenge, explanation, feedback, and practice while preserving domain standards, professional judgment, and human accountability.

Keywords: adaptive prompting, professional training, generative artificial intelligence, domain-specific learning, intelligent tutoring systems, personalized learning, prompt routing, professional competence

1. Introduction

Professional learning differs from general information acquisition because competence requires more than recalling facts. Healthcare professionals must connect knowledge with patient context and clinical responsibility; engineers must apply principles under technical and safety constraints; managers must navigate interpersonal situations; legal professionals must distinguish authority, procedure, and factual uncertainty; cybersecurity personnel must interpret rapidly changing threat environments; and educators must adapt instruction to learners while exercising professional judgment. Generative AI can support these activities through explanation, simulation, role play, feedback, and guided practice. Prompt engineering has consequently emerged as an important competence for working effectively with language models, with Federiakin and colleagues conceptualizing it as a multidimensional skill involving problem articulation, context, constraints, prompting methods, and critical reasoning.

A central limitation of many AI-assisted training systems is nevertheless their dependence on relatively static prompting. A fixed prompt may tell a model to act as a tutor, provide hints, use a Socratic style, or generate formative feedback, but it does not necessarily determine whether the current trainee requires conceptual explanation, procedural rehearsal, error diagnosis, reduced support, advanced challenge, or referral to an instructor. This distinction is important because a 2025 benchmarking study involving 75 real tutoring scenarios and 1,350 model-generated instructional moves found that contemporary LLMs only partially reproduced the adaptivity associated with established intelligent tutoring systems. The study found differences across models in responsiveness to learner errors and pedagogical quality, showing that general conversational capability should not be interpreted as evidence of reliable instructional adaptation.

Recent research, however, demonstrates that adaptivity can be engineered more explicitly. Chang and colleagues introduced a subject-aware prompt-routing framework that selects among different pedagogical strategies according to learner and subject information. Their 2026 study used 14 pedagogical features and a contextual routing mechanism; simulation results exceeded static-prompt baselines, and real-world testing involved 656 conversations from 359 students. The study therefore provides contemporary evidence that prompt selection itself can function as an adaptive instructional mechanism rather than remaining fixed throughout a learning interaction. Open TutorAI offers a related architecture in which instructional prompts are dynamically constructed from learner profiles, instructional objectives, session history, progress, and communication preferences rather than generated from one unchanging tutor instruction.

Professional training creates an especially compelling environment for such adaptation because trainees can differ substantially even when they hold the same occupational title. A newly appointed manager practicing a difficult performance conversation requires different support from an experienced executive rehearsing negotiation under uncertainty. A novice programmer may need conceptual hints, whereas an experienced developer may benefit from diagnostic questions and counterexamples. Workplace research already indicates interest in this type of adaptive practice. Wilhelm and colleagues' 2019 investigation of an AI-assisted conversational training system found that managers valued low-risk simulations, context-aware feedback, adaptability, and greater control over AI-generated personas when practicing workplace communication. A separate study of knowledge workers emphasized the need for

American Journal of Advanced Multidisciplinary Innovation and Research (AJAMIR)
E-ISSN: XXXX-XXXX ● Website: www.ajamir.com ● Email: editor@ajamir.com

structured AI training because inadequate preparation can lead to misuse, misunderstanding, bias, or abandonment of AI tools.

The present study therefore conceptualizes adaptive prompting as a closed-loop professional learning mechanism. Learner evidence informs prompt selection; the selected prompt generates a training intervention; the trainee's response produces new evidence; and the system subsequently adjusts explanation, challenge, simulation difficulty, or feedback. This model differs from unrestricted personalization because adaptation is bounded by professional competencies, verified domain material, assessment criteria, safety requirements, and human oversight. UNESCO's guidance on generative AI likewise emphasizes a human-centered approach to educational design, data privacy, ethical validation, and preservation of human agency. The study develops and evaluates a simulation-based framework designed to translate these principles into domain-specific professional training.

2. Objectives of the Study

2.1 Development of a Domain-Specific Adaptive Prompting Framework

The first objective is to develop an instructional architecture in which prompt behavior responds systematically to professional domain, competency level, trainee errors, prior interactions, task difficulty, and assessment evidence. The framework seeks to transform prompting from a static instruction supplied to a generative model into a structured mechanism for deciding how, when, and at what level instructional assistance should be delivered.

2.2 Evaluation of Training Personalization and Professional Fidelity

The second objective is to evaluate how alternative prompting architectures may differ in their ability to preserve domain-specific accuracy while adapting instructional support. The study specifically considers whether personalization can be achieved without weakening professional standards, assessment consistency, or the distinction between educational support and authoritative professional judgment.

2.3 Identification of Governance Requirements for Professional Deployment

The third objective is to identify the safeguards necessary when adaptive prompting is used in professional contexts. These include verified domain knowledge, transparent competency mappings, human-defined escalation thresholds, privacy-aware learner modeling, performance monitoring, prompt-version control, source provenance, and mechanisms through which qualified professionals can review or override AI-generated instructional behavior.

3. Review of Literature

3.1 From Static Prompt Engineering to Adaptive Instructional Prompting

Prompt engineering initially developed primarily as a technique for improving the quality and reliability of model responses. Federiakin, Molerov, Zlatkin-Troitschanskaia, and Maur conceptualize prompting as a composite skill requiring users to communicate a problem, relevant context, and constraints effectively while evaluating AI output critically. Their framework further emphasizes that prompting is typically iterative rather than a single interaction, which makes it particularly compatible with instructional environments where learner understanding develops across multiple exchanges.

Prompting can become pedagogically significant when instructions influence the type of assistance a learner receives. Abolnejadian, Alipour, and Taeb investigated personalized prompt-based instruction in introductory computer-science education, illustrating how prompt design can be connected with adaptive learning rather than functioning only as an input-optimization technique. Balavar and colleagues subsequently combined tailored prompting, domain knowledge, RAG, and skill-aligned feedback in a programming-tutoring context. Their pilot framework categorized simulated learners into skill levels and generated context-sensitive feedback, demonstrating how domain evidence and learner status can jointly condition instructional responses.

3.2 Domain Knowledge, Learner Modeling, and Professional Skill Development

Professional training requires domain fidelity because generic explanations can be pedagogically plausible while professionally inappropriate. The same communication technique may not be appropriate in clinical counseling, legal interviewing, managerial feedback, and emergency response. Domain-specific prompting should therefore encode terminology, procedural boundaries, competency standards, evidence sources, and task-specific constraints.

Retrieval-supported prompting provides one approach to maintaining such fidelity. Balavar and colleagues integrated domain knowledge with skill-aligned feedback through retrieval-augmented generation, demonstrating the potential for training responses to be grounded in relevant domain material rather than generated entirely from model memory. Domain-specific model research similarly shows that targeted knowledge adaptation can improve performance in specialized fields, although model specialization and prompt adaptation represent distinct mechanisms. In professional training, the two approaches can be combined: verified domain resources establish the authoritative knowledge boundary, while prompts determine the pedagogical treatment of that knowledge.

3.3 Research Gap and Conceptual Direction

The literature demonstrates substantial progress in prompt engineering, personalized tutoring, learner modeling, adaptive learning platforms, and AI-assisted professional simulations. However, three gaps remain.

First, prompt personalization and instructional adaptivity are often treated as equivalent. They are not. A model that changes tone according to learner preference is personalized, but it is not necessarily adapting instructional strategy according to evidence of competence.

Second, most emerging adaptive-prompt systems have been examined in academic learning environments rather than across professional domains. Workplace studies show demand for AI-mediated practice, but professional training introduces stronger requirements concerning accuracy, standards, confidentiality, consequences, and human responsibility. Knowledge-worker research specifically demonstrates that workers require training not only in AI operation but also in interpretation, bias, rights, and appropriate use.

Third, evidence indicates that LLM adaptivity remains inconsistent. Borchers and Shou caution that current LLM tutoring does not automatically match the adaptive behavior of established ITS approaches, whereas Chang and colleagues demonstrate that explicit prompt routing can improve adaptation. The research opportunity therefore lies not in asking whether an LLM can spontaneously act as an adaptive trainer, but in designing an external instructional architecture that determines what the model should do under specific professional-learning conditions.

The present study addresses this gap through a framework that separates six functions: professional competency specification, learner-state estimation, adaptive prompt selection, domain-grounded generation, performance evaluation, and human governance.

4. Research Methodology

4.1 Research Design and Synthetic Training Corpus

The study adopts an explicitly simulation-based methodological design. No practicing professionals, trainees, employees, clinicians, engineers, managers, or other human participants supplied the quantitative observations reported in the analysis.

A synthetic corpus of 480 hypothetical professional-training interactions was constructed. The scenarios were conceptually distributed across four professional domains—healthcare, engineering, management, and cybersecurity—to prevent the framework from becoming dependent on the instructional characteristics of a single field.

The interactions were divided equally across four prompting architectures:

- Static Domain Prompting;
- Rule-Based Adaptive Prompting;
- Learner-State Prompt Routing; and
- Domain-Specific Adaptive Prompting with Professional Oversight.

The objective of the simulation is not to estimate actual learning gains. It is to demonstrate how alternative architectures can be compared through a multidimensional readiness framework before conducting real-world trials.

Table 1: Operational Framework for Adaptive Professional-Training Prompting

Evaluation Dimension	Operational Meaning	Illustrative Evaluation Question	Scale
Domain Fidelity	Alignment with verified professional knowledge, terminology, procedures, and competency standards	Does the response remain professionally appropriate and evidence-grounded?	0–100
Learner-State Sensitivity	Responsiveness to knowledge level, prior errors, confidence, progress, and interaction history	Does the system alter instruction when learner evidence changes?	0–100
Scaffolding Appropriateness	Matching assistance to the trainee's current level without excessive or insufficient support	Does the system provide hints, questioning, modeling, or challenge at the appropriate time?	0–100
Feedback Quality	Specificity, actionability, explanatory depth, and connection between	Does feedback explain both performance and a constructive	0–100

Evaluation Dimension	Operational Meaning	Illustrative Evaluation Question	Scale
	performance and improvement	next step?	
Assessment Alignment	Consistency between adaptive instruction and predefined professional competency criteria	Is adaptation still anchored to comparable training outcomes?	0–100

Source: Author-developed simulation framework informed by contemporary adaptive tutoring and AI-assisted professional-learning research.

4.2 Adaptive Prompting Architecture

The proposed framework treats each training exchange as part of a closed instructional cycle.
A trainee first interacts with a domain task, which may take the form of a scenario, case, dialogue, technical problem, diagnostic situation, or decision exercise. The system then extracts a bounded set of learner-state variables such as competency level, prior attempts, error category, uncertainty, response quality, and training objective.

These variables do not directly produce an answer. Instead, they determine the instructional strategy.For example, when a novice repeatedly makes a conceptual error, the system may route the interaction toward an explanatory-scaffolding prompt. When an intermediate trainee demonstrates partial understanding, the selected prompt may ask structured questions rather than reveal the solution. When an advanced learner performs correctly, the system may increase ambiguity or introduce additional constraints.

4.3 Simulation, Scoring, and Validation Strategy

A Professional Training Readiness Index (PTRI) was calculated by assigning equal conceptual weight to the six evaluation dimensions.

The equal weighting is methodological rather than empirical. Different professional domains may require different priorities. Healthcare training, for example, may assign greater weight to safety and evidence fidelity, whereas management communication training may place comparatively greater emphasis on contextual responsiveness and interpersonal feedback.The four prompting architectures were characterized as follows.Static Domain Prompting uses professional context but applies substantially the same tutoring instruction across trainees and sessions.

Rule-Based Adaptive Prompting uses predefined conditions such as novice/intermediate/advanced level or correct/incorrect response to select instructional templates.Learner-State Prompt Routing incorporates multiple learner variables, performance history, and task context to determine among several pedagogical strategies.Domain-Specific Adaptive Prompting with Professional Oversight adds verified knowledge retrieval, professional competency mappings, source constraints, human-authored prompt policies, assessment calibration, escalation procedures, monitoring, and instructor override.

A future empirical validation should compare these architectures through randomized or quasi-experimental professional-training studies. Appropriate outcomes could include pre/post competency assessments, delayed retention, transfer to novel scenarios, error reduction, feedback usefulness, time to competency, calibration of confidence, instructor ratings, and performance under reduced AI assistance.

5. Analysis

5.1 Comparative Performance of Prompting Architectures

The simulation indicates a progressive increase in training readiness as instructional adaptation becomes more explicit and professionally governed.

Table 2: Simulated Professional Training Readiness Scores

Prompting Architecture	Domain Fidelity	Learner-State Sensitivity	Scaffolding	Feedback Quality	Assessment Alignment	Governance Reliability	Overall PTRI
Static Domain Prompting	72	24	43	58	61	54	52
Rule-Based Adaptive Prompting	78	64	70	69	67	60	68
Learner-State Prompt Routing	86	88	85	84	77	72	82
Domain-Specific Adaptive Prompting + Professional Oversight	95	92	91	90	89	89	91

Source: Simulated analytical values created for methodological illustration. These values are not empirical measures of professional learning outcomes.

Static domain prompting performs relatively well in domain fidelity because professional instructions and terminology can be built into the original prompt. Its low learner-state sensitivity substantially limits its readiness for personalized training because the same strategy can persist despite changes in performance.

Rule-based adaptation improves sensitivity by mapping predefined conditions to different instructional templates. This architecture is interpretable and comparatively easy to audit, although complex learner behavior may exceed the rules anticipated during system design.

Learner-state prompt routing receives substantially higher simulated scores because multiple signals can influence instructional strategy. This design resembles recent subject-aware routing research in which different pedagogical prompts are selected according to contextual information instead of relying on one global prompt.The strongest architecture retains adaptive routing but places it within professionally governed boundaries. The additional readiness gain reflects stronger source control, assessment alignment, monitoring, and human oversight rather than greater personalization alone.

5.2 Professional Training Readiness Across Adaptive Prompting Systems

Graph 1: Simulated Professional Training Readiness Across Prompting Architectures

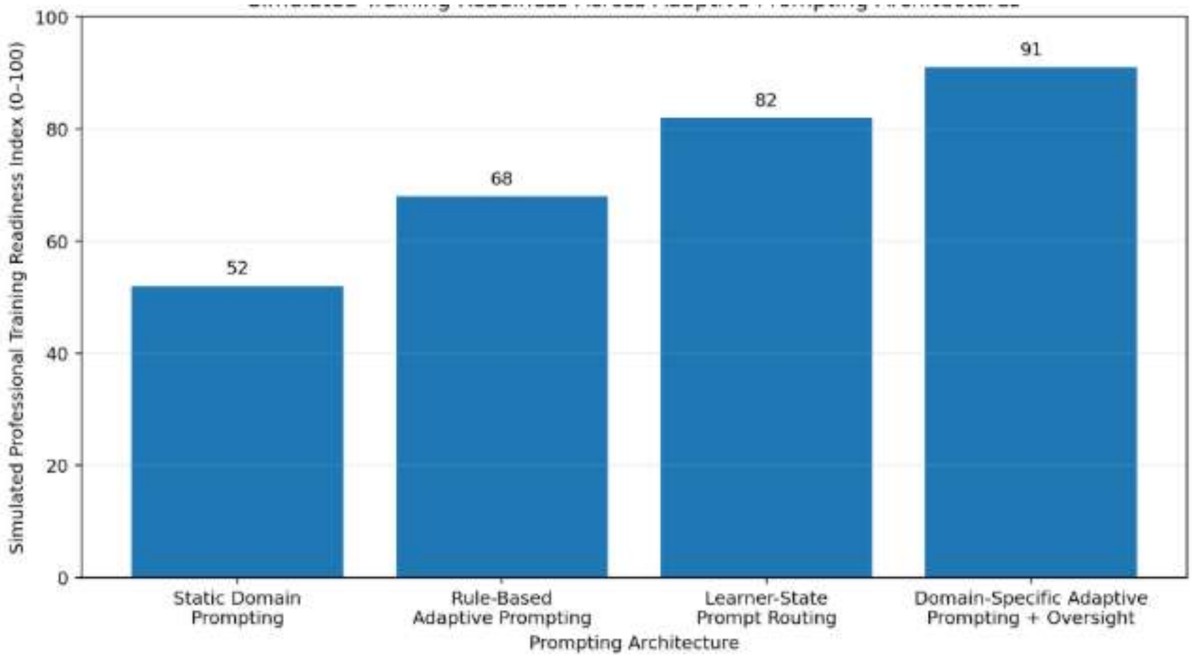

Graph 1 shows simulated readiness scores of 52 for static domain prompting, 68 for rule-based adaptation, 82 for learner-state prompt routing, and 91 for the integrated domain-specific adaptive architecture.

The increase between static and rule-based systems represents the benefit of introducing explicit conditional instructional behavior. The further increase from 68 to 82 reflects richer learner modeling and the ability to route among pedagogical strategies rather than selecting only by broad proficiency category.

The final increase to 91 is conceptually different. It does not arise simply from more sophisticated personalization. Instead, it represents the addition of professional governance.

This distinction is critical because the most adaptive system is not necessarily the most appropriate professional trainer if adaptation causes inconsistent standards, unsupported advice, or excessive reliance on AI.

5.3 Adaptive Progression and Competency Control

The analytical framework suggests that adaptive prompting should alter instructional support while keeping competency objectives stable.A professional training system may initially provide substantial scaffolding, later reduce hints, and eventually introduce realistic ambiguity. What changes is the instructional pathway; what should remain stable is the underlying professional performance standard. This principle prevents personalization from becoming arbitrary simplification.

A trainee struggling with a clinical reasoning exercise, for example, might receive additional decomposition and reflective questioning, but the required safety principles should not be relaxed. Similarly, a management trainee may receive a simpler communication scenario before progressing to conflict involving multiple stakeholders, while the competency criteria for constructive communication remain explicit.Adaptive systems should consequently estimate not only whether the trainee answered correctly but why the performance occurred.

Errors may arise from insufficient domain knowledge, procedural confusion, overconfidence, misunderstanding of the scenario, ineffective communication, or failure to recognize uncertainty. Different causes justify different prompts.This capability represents the practical value of learner-state modeling.

6. Discussion

6.1 Adaptive Prompting Should Be Treated as Instructional Policy Rather Than Prompt Wording

The principal implication of the study is that adaptive prompting should be understood as a system for selecting instructional actions rather than merely a method for rewriting text supplied to a generative model. Static prompt engineering can improve response quality, specify professional roles, define output formats, and establish constraints, but it provides limited assurance that the system will recognize when a trainee requires a fundamentally different type of support. Professional learning demands decisions concerning whether the next instructional action should explain, question, demonstrate, challenge, simulate, correct, or withhold assistance. These decisions constitute pedagogical policy rather than linguistic formatting.

Contemporary evidence reinforces this distinction. Borchers and Shou found that general LLMs only partially exhibited the instructional adaptivity associated with established tutoring systems, even when learner errors and knowledge information were present in context. By contrast, Chang and colleagues explicitly modeled instructional adaptation as prompt routing and reported stronger performance than static baselines. Together, these findings suggest that professional training systems should externalize adaptation rather than relying completely on the generative model's implicit judgment. A prompt router, learner-state model, or transparent decision layer can determine which instructional strategy is appropriate before the model generates the actual training response.

This architecture also improves auditability. If an instructor observes that trainees are receiving overly direct solutions, the institution can examine which routing policy selected the problematic strategy, revise the corresponding prompt, and retest the behavior. In an unrestricted chatbot architecture, the instructional cause of the same problem may be substantially more difficult to identify. Professional deployment therefore benefits when adaptive behavior is represented through explicit, version-controlled instructional policies.

6.2 Domain Fidelity and Personalization Must Be Designed Together

The second major implication is that personalization should not occur independently of professional knowledge governance. Generative systems can produce responses that are fluent, supportive, and personalized while still being technically incorrect or professionally inappropriate. The problem becomes particularly significant when the trainee lacks sufficient expertise to recognize the error. Adaptive professional training therefore requires a stable domain layer containing verified standards, terminology, procedures, assessment criteria, and authoritative reference material.

Retrieval-supported instruction provides one way to establish this boundary. Balavar and colleagues' work demonstrates how domain knowledge and skill-aligned feedback can be integrated with prompting rather than relying exclusively on general model knowledge. In an operational professional-training system, retrieval should be connected to curated sources such as approved procedures, technical standards, training manuals, professional guidance, or institutionally validated competency frameworks. Adaptive prompts can then determine whether retrieved knowledge should be explained, withheld temporarily as part of a practice exercise, converted into a hint, used to generate a counterexample, or incorporated into formative feedback.

The relationship between personalization and assessment also requires careful control. Trainees may follow different instructional pathways without being judged against completely different standards. If the system continuously lowers difficulty whenever a learner struggles, it may create an appearance of progress without establishing independent competence. Adaptive support should therefore be progressively reduced as proficiency develops, and summative or competency-gating assessments should include conditions in which assistance is limited or standardized. This distinction protects the validity of professional certification and prevents adaptive training from becoming permanent dependence on an AI assistant.

Domain fidelity additionally requires meaningful uncertainty management. Professional practice frequently involves incomplete information rather than deterministic answers. An adaptive system should therefore avoid rewarding confident but unsupported responses merely because they resemble expected language. Advanced training should sometimes require the learner to identify missing information, articulate uncertainty, justify escalation, or decline to act without sufficient evidence. In this sense, professional competence includes recognizing the limits of both one's own knowledge and the AI system's assistance.

6.3 Human Oversight and Learner Agency Remain Essential to Professional Training

The third implication concerns the role of human professionals. Adaptive prompting can increase the frequency and personalization of practice, but it should not eliminate trainers, supervisors, mentors, or subject-matter experts. Workplace communication research demonstrates that users value adaptive, low-risk AI simulations while also identifying tensions involving feedback consistency, contextual realism, bias, and control. Knowledge-worker research likewise emphasizes that safe AI adoption requires structured training around interpretation, bias, responsible use, and worker rights rather than simply providing access to AI systems. These findings support a human–AI training model in which AI expands opportunities for individualized practice while qualified professionals retain responsibility for curriculum, competency standards, high-consequence feedback, and system governance.

Human oversight should occur at several levels. Subject-matter experts should define the acceptable knowledge boundary and review scenario realism. Instructional specialists should validate the relationship between learner states and prompting strategies. Supervisors should be able to inspect trainee progress, identify repeated misconceptions, and intervene when an automated pathway is ineffective. Technical teams should maintain prompt versions, model configurations, evaluation records, and security controls. Organizations should additionally establish escalation conditions for topics whose consequences make unrestricted AI tutoring inappropriate.

Learner agency is equally important. Adaptive systems can become intrusive if every hesitation, error, confidence judgment, or behavioral pattern is converted into a persistent learner profile without clear justification. UNESCO's human-centered guidance emphasizes privacy, ethical validation, and preservation of human agency in generative-AI-supported education. Professional training systems should therefore collect only learner-state information necessary for legitimate instructional purposes, explain how adaptation occurs, and provide reasonable opportunities for trainees to review feedback or seek human guidance.

The strongest professional-training architecture is consequently neither fully automated nor static. It combines structured adaptive prompting with domain evidence, transparent learner modeling, human professional judgment, and opportunities for reflection. AI contributes scalability and responsiveness; professionals provide standards, contextual interpretation, ethical responsibility, and accountability.

7. Conclusion

Adaptive prompting offers a promising approach to domain-specific professional training because it enables generative-AI systems to vary instructional behavior according to competency level, error patterns, prior performance, task difficulty, and professional context. The present study has argued that this capability should not be confused with ordinary conversational personalization. Changing tone or wording according to user preference may improve experience, but genuine instructional adaptation requires the system to determine which pedagogical action is appropriate at a particular stage of professional development. The proposed framework therefore separates learner-state estimation, domain knowledge, prompt routing, generation, assessment, and professional oversight into distinct but coordinated functions.

The simulation-based analysis produced a Professional Training Readiness Index of 52 for static domain prompting, 68 for rule-based adaptive prompting, 82 for learner-state prompt routing, and 91 for domain-specific adaptive prompting combined with professional oversight. These values are illustrative rather than empirical learning outcomes, but they clarify an important structural progression. Static prompts can preserve professional context but respond poorly to changing learner needs; rules introduce interpretable adaptation; learner-state routing enables more responsive instructional strategy; and professionally governed adaptive systems add the domain fidelity, assessment consistency, monitoring, and accountability required for consequential training environments. The results therefore indicate that sophistication of adaptation should not be evaluated independently from reliability of governance.

The study further concludes that professional training requires a balance between personalization and stable competency standards. Adaptive systems may legitimately vary examples, scaffolding, explanation depth, difficulty, feedback timing, and simulated scenarios, yet the professional standard being developed should remain explicit. Assistance should progressively decrease as competence develops, and systems should distinguish supported practice from independent assessment. Human trainers remain essential for defining competency models, validating high-consequence scenarios, reviewing unusual learner behavior, resolving conflicting evidence, and ensuring that automated feedback does not replace professional judgment where contextual interpretation is required.

Future empirical work should evaluate the proposed framework through domain-specific trials involving practicing professionals and trainees in healthcare, engineering, cybersecurity, management, education, law, and other knowledge-intensive fields. Such studies should measure immediate performance, delayed retention, transfer to unfamiliar cases, reduction in repeated errors, confidence calibration, time to competency, reliance on AI support, instructor agreement, and performance after adaptive assistance is removed. Longitudinal research will be particularly important because the ultimate objective of adaptive prompting should not be to make learners permanently dependent upon increasingly helpful AI. Its educational value will be demonstrated when appropriately calibrated assistance enables professionals to develop stronger independent judgment, reliable domain competence, reflective practice, and the ability to recognize when expert consultation remains necessary.

References

1. VanLehn, K. (2011). The relative effectiveness of human tutoring, intelligent tutoring systems, and other tutoring systems. *Educational Psychologist*, *46*(4), 197–221. https://doi.org/10.1080/00461520.2011.611369
2. Graesser, A. C., Conley, M. W., & Olney, A. (2012). Intelligent tutoring systems. In K. R. Harris, S. Graham, & T. Urdan (Eds.), *APA Educational Psychology Handbook* (Vol. 3, pp. 451–473). American Psychological Association.
3. Woolf, B. P. (2009). *Building Intelligent Interactive Tutors: Student-Centered Strategies for Revolutionizing E-Learning*. Morgan Kaufmann.
4. Anderson, J. R., Corbett, A. T., Koedinger, K. R., & Pelletier, R. (1995). Cognitive tutors: Lessons learned. *The Journal of the Learning Sciences*, *4*(2), 167–207. https://doi.org/10.1207/s15327809jls0402_2
5. Corbett, A. T., & Anderson, J. R. (1995). Knowledge tracing: Modeling the acquisition of procedural knowledge. *User Modeling and User-Adapted Interaction*, *4*, 253–278. https://doi.org/10.1007/BF01099821
6. Brusilovsky, P. (2001). Adaptive hypermedia. *User Modeling and User-Adapted Interaction*, *11*, 87–110. https://doi.org/10.1023/A:1011143116306
7. Brusilovsky, P., & Millán, E. (2007). User models for adaptive hypermedia and adaptive educational systems. In P. Brusilovsky, A. Kobsa, & W. Nejdl (Eds.), *The Adaptive Web* (pp. 3–53). Springer.

8. Shute, V. J., & Ventura, M. (2013). *Stealth Assessment: Measuring and Supporting Learning in Video Games*. MIT Press.
9. Koedinger, K. R., Anderson, J. R., Hadley, W. H., & Mark, M. A. (1997). Intelligent tutoring goes to school in the big city. *International Journal of Artificial Intelligence in Education, 8*, 30–43.
10. Ma, W., Adesope, O. O., Nesbit, J. C., & Liu, Q. (2014). Intelligent tutoring systems and learning outcomes: A meta-analysis. *Journal of Educational Psychology, 106*(4), 901–918. https://doi.org/10.1037/a0037123
11. Bloom, B. S. (1984). The 2 sigma problem: The search for methods of group instruction as effective as one-to-one tutoring. *Educational Researcher, 13*(6), 4–16. https://doi.org/10.3102/0013189X013006004
12. Chi, M. T. H., Bassok, M., Lewis, M. W., Reimann, P., & Glaser, R. (1989). Self-explanations: How students study and use examples in learning to solve problems. *Cognitive Science, 13*(2), 145–182. https://doi.org/10.1207/s15516709cog1302_1
13. Ericsson, K. A., Krampe, R. T., & Tesch-Römer, C. (1993). The role of deliberate practice in the acquisition of expert performance. *Psychological Review, 100*(3), 363–406. https://doi.org/10.1037/0033-295X.100.3.363
14. Zimmerman, B. J. (2002). Becoming a self-regulated learner: An overview. *Theory Into Practice, 41*(2), 64–70. https://doi.org/10.1207/s15430421tip4102_2
15. Pintrich, P. R. (2000). The role of goal orientation in self-regulated learning. In M. Boekaerts, P. R. Pintrich, & M. Zeidner (Eds.), *Handbook of Self-Regulation* (pp. 451–502). Academic Press.
16. Deci, E. L., & Ryan, R. M. (2000). The "what" and "why" of goal pursuits: Human needs and the self-determination of behavior. *Psychological Inquiry, 11*(4), 227–268. https://doi.org/10.1207/S15327965PLI1104_01
17. Mayer, R. E. (2009). *Multimedia Learning* (2nd ed.). Cambridge University Press.
18. Clark, R. C., & Mayer, R. E. (2016). *E-Learning and the Science of Instruction* (4th ed.). Wiley.
19. Brown, P. C., Roediger, H. L., & McDaniel, M. A. (2014). *Make It Stick: The Science of Successful Learning*. Harvard University Press.
20. Heffernan, N. T., & Heffernan, C. L. (2014). The ASSISTments ecosystem: Building a platform that brings scientists and teachers together for minimally invasive research on human learning and teaching. *International Journal of Artificial Intelligence in Education, 24*, 470–497. https://doi.org/10.1007/s40593-014-0024-x