Proceedings of CONF-MLA 2025 Symposium: Intelligent Systems and Automation: AI Models, IoT, and Robotic Algorithms
DOI: 10.54254/2755-2721/2025.LD31407

Generative-AI–Driven Intelligent Socratic Dialogue: A Theoretical Analysis and Model Construction

Qiwen Yu[1], Xiangyu Wang[1*]

[1]Department of Elementary Education, Shandong Normal University, Jinan, China
**Corresponding Author. Email: 2568144086@qq.com*

Abstract. Imagine a classroom where every student gets a personal tutor trained in the Socratic art—not to give answers, but to ask just the right questions. That vision, however, has long run into a practical wall: how do you scale a method built on live, adaptive dialogue to hundreds or thousands of learners? In this article, we explore a generative-AI model designed to do exactly that. Instead of reducing Socratic teaching to a scripted Q&A, we treat the AI as a "midwife" for thinking—one that operates from deliberate ignorance, provides cognitive scaffolding, and dynamically shifts initiative back to the learner. Technically, the system works through a three-layer "Goal–Agent–Conversation" architecture, turning broad instructional aims into living, personalized dialogues. Learners move through a kind of cognitive spiral: they externalize their assumptions, run into contradictions, reflect critically, and gradually rebuild understanding. What we're really trying to address here is an old tension in education—the tug-of-war between standardization and genuine personalization. Could this approach ease that friction? Looking ahead, we see the beginnings of a data-informed learning model that doesn't just deliver knowledge but cultivates the skill of thinking itself. We also sketch where the work goes from here—from technical fine-tuning and new use cases to the ethical questions that will inevitably follow.

Keywords: generative artificial intelligence, Socratic dialogue, cognitive partner, educational paradigm shift

1. Introduction

Socratic dialogue has always represented an educational ideal: learning as a guided act of discovery, where probing questions, not delivered answers, stir a learner's own reasoning into motion [1]. Yet in practice, this ideal is notoriously difficult to mass-produce. Its magic lies in responsive, adaptive exchange—precisely what gets lost in large classrooms or standardized digital platforms.

Here's where generative AI enters the picture. With their emerging grasp of context and reasoning, large language models could, in principle, simulate the questioning art of a human tutor [2]. But let's be honest—most educational AI today still operates like a sophisticated answer engine. It's a one-way street for information, which frankly misses the whole point of Socratic "midwifery." The real goal isn't to transfer knowledge, but to help thinkers uncover it for themselves [3].

DOI: 10.54254/2755-2721/2025.LD31407

So, can we design an AI that stays true to that spirit while actually working at scale? That's the puzzle we take on in this study. We develop a generative-AI model for intelligent Socratic dialogue, built around three guiding questions:

What, theoretically, defines an "intelligent" Socratic dialogue in an age of AI collaboration?

How do we build a system that honors Socrates' philosophical essence without becoming a scripted Q&A machine?

And if such a system works, how does it reshape traditional teacher–student roles—and perhaps the broader educational paradigm?

The discussion that follows moves from theory to design to implication. We start by grounding the model in its philosophical roots and outlining its human–AI collaborative nature. From there, we lay out the design principles and technical architecture that bring it to life. We then consider what learners might gain from this approach—not just in knowledge, but in thinking habits. Finally, we step back to reflect on what this shift might mean for education at large, and where the road leads next, from technical hurdles to ethical considerations.

2. Rethinking socratic dialogue for the AI era

If we want to build an AI that can genuinely think with a learner—not merely answer them—we have to slow down and ask a deceptively simple question: what actually counts as a Socratic dialogue? It isn't just a classroom trick or a clever questioning routine. At its best, it's a way of treating learning as something you do, not something you receive. That leaves us with a translation problem. We have to keep the philosophical spine of the method intact while giving it a new body—one that can live inside a generative AI system. The aim of this section is to sketch a framework for that translation, tracing the thread from ancient practice to something that can operate, credibly, in an algorithmic environment.

2.1. The philosophical core: from midwifery to productive uncertainty

Any serious digital heir to Socrates has to begin where the original method begins: with two ideas that quietly govern everything else.

One is intellectual midwifery [4]. Socrates didn't present himself as a master pouring knowledge into empty minds. He acted more like a midwife—someone who helps bring forth an understanding that is already forming inside the learner, even if it arrives awkwardly, in fragments, and with a bit of pain. The implication is unmistakably constructivist: knowledge isn't "delivered," it is built. For AI design, that pushes us toward a posture that feels almost counterintuitive in the age of instant answers: disciplined restraint. The system needs to practice a kind of deliberate ignorance. Even when it could supply the solution, it holds back—because the point is not to rescue the learner from uncertainty, but to use uncertainty as the engine of reasoning. Its job becomes the careful shaping of questions that trigger the learner's own inferential work.

The second idea is the elenchus, the method of refutation [5]. Here the dialogue sharpens. Through persistent, targeted questioning, people are nudged to examine their own claims, notice inconsistencies, and confront the discomfort of not having a clean answer. And that discomfort matters. Anyone who has watched a student pause mid-sentence and think, "Wait—why doesn't that line up?" knows that this small moment of friction often marks the beginning of real critical thought. So an intelligent system cannot be satisfied with "asking questions" in a generic sense. It has to listen closely to what the learner actually says, detect where the reasoning blurs or where a hidden assumption slips in, and then push—gently but insistently—until the gap becomes visible. In other

DOI: 10.54254/2755-2721/2025.LD31407

words, the system must embody the central Socratic purpose: it should extend thinking, not close it down.

2.2. A new kind of partnership: defining the human–AI cognitive model

Once those principles are on the table, we can define intelligent Socratic dialogue in contemporary terms: a generative-AI-driven interactive system that uses structured questioning to strengthen critical thinking and metacognition. That sounds neat on paper, but the definition only makes sense if we accept three meaningful shifts.

A shift in goal. The target is no longer efficient knowledge transfer. The focus moves to the learner's thinking process—how ideas are formed, tested, revised, and occasionally abandoned. Problems become a stage on which a full cognitive cycle can play out: hypothesizing, bumping into contradictions, reflecting, and rebuilding understanding in a more stable form.

A shift in roles. The AI stops behaving like a static warehouse of information and starts functioning as an inquiry partner—something closer to a persistent, attentive interlocutor [6]. At the same time, the learner is no longer positioned as a receiver. They become the active constructor, the person responsible for doing the intellectual assembly work. In practice, this reconfigures the familiar teacher–student dyad into a three-way relationship: a human–machine co-cognitive partnership.

A shift in method. The model also serves as a bridge between craft and system. Traditional Socratic teaching often depends on tacit expertise—an almost artisanal feel for when to press, when to pause, when to ask a sharper question, and when to let the learner wrestle [7]. A well-designed AI model can translate parts of that skill into a practice that is explicit enough to study, refine, and scale. The larger promise is not that machines replace mentorship, but that high-level thinking instruction no longer remains trapped inside rare one-on-one interactions.

2.3. What makes it tick: core features of the paradigm

So what would an algorithmic version of Socratic dialogue look like when it stops being a slogan and starts being a working system? The paradigm can be understood through three interlocking features—each a contemporary interpretation of Socratic philosophy, not a departure from it.

At the base is the algorithmic enactment of guided interrogation. This is the point where the system differs most sharply from standard educational AI. The system operates under constraints that prevent it from short-circuiting the learner's struggle. Sometimes it must "play dumb," not because it lacks the answer, but because giving the answer would collapse the space in which thinking happens. In effect, it tries to imitate what skilled human tutors do: parse a learner's response, locate the conceptual stumble or logical misstep, and generate a coherent chain of follow-up questions that nudges the learner to do the repair themselves. If the questioning is careless, the whole method becomes theatrical; if it is precise, it becomes transformative.

Two further capabilities follow from that foundation.

The first is the observability of thinking-in-progress. As learners articulate reasons, caveats, and half-formed intuitions in order to respond, their private cognitive work becomes external and traceable. The dialogue logs create something like a map of the learner's "logical footsteps." That map can support metacognitive review in a concrete way: instead of evaluating only the final answer, instructors (and learners themselves) can examine the route taken, notice recurring patterns, and intervene at the level of reasoning rather than outcomes.

DOI: 10.54254/2755-2721/2025.LD31407

The second is a dynamic re-balancing of teaching responsibility. The model invites a practical division of labor: let the AI take on the sustained, scalable work of structured cognitive questioning, while human teachers concentrate on what remains distinctly human—emotional attunement, moral nuance, situational judgment, and personal encouragement. One might summarize the collaboration as "AI for analysis, teacher for counsel," though in real classrooms the boundary will never be perfectly clean. Still, the intent is clear: to ease the old tension between efficiency and individualized attention without pretending the tension can be eliminated.

Taken together, these features make the cognitive process more visible and redistribute instructional roles in a way that is both philosophically grounded and operationally plausible. More importantly, they suggest a realistic answer to an old educational problem: how do we cultivate deep thinking skills—not only for the students lucky enough to receive intensive mentorship, but for many more learners who rarely get that kind of guided intellectual pressure?

3. Architecture and dynamics of the system

Building an AI that can sustain a true Socratic exchange is less about engineering a brilliant oracle and more about designing a thoughtful, restrained partner. The system rests on three core principles that act as guardrails, steering the technology away from becoming just another Q&A interface and keeping it aligned with the spirit of guided discovery.

3.1. The guardrails: core design principles

Principle 1: AI as Deliberate Ignorance.

This is the system's foundational discipline. The AI must actively resist its own capability to provide answers. We treat generative AI not as a knowledge sovereign, but as a guide that knows when to stay quiet. Technically, this is enforced through rigorous prompt engineering—hard-coded rules that forbid declarative statements and mandate that all guidance be framed as a question. The system even monitors its own outputs in real time, rewriting any response that smells too much like a verdict back into an inquisitive form.

Why go to such lengths? It's to prevent a common failure mode: the AI doing the thinking for the student. By maintaining what we might call productive knowledge gaps, the system ensures the learner's mind stays in the driver's seat. Those gaps, after all, are where intellectual independence grows.

Principle 2: Cognitive Scaffolding.

Drawing from Vygotsky's Zone of Proximal Development [8], this principle demands that the system diagnose and adapt in real time. It continuously reads multiple signals—the coherence of a learner's reasoning, the efficiency of their responses, even subtle affective cues in their language—and adjusts its support accordingly.

When a learner is on a roll—reasoning clearly, connecting ideas without much strain—the system can afford to step back and open the room a little. It might nudge the conversation toward more abstract territory, or toss in an open-ended question that invites the learner to generalize, defend, or even rethink their claim. But when the learner starts to wobble—circling a contradiction, leaning on a fuzzy assumption, or simply going quiet—the system should tighten the frame instead. It can break the idea into smaller, more workable pieces, ask a crisp clarifying question, or slip in a counterexample that gently tests the learner's logic.

The goal, in most cases, is to keep cognitive tension in a useful range. Too little pressure and the dialogue turns into comfortable chat; too much and the learner stalls, frustrated or disengaged.

DOI: 10.54254/2755-2721/2025.LD31407

Somewhere between those extremes is the zone where you can almost see thinking change shape—where a student pauses, revises, and comes back with a stronger line of reasoning. Isn't that the real point of the exercise?

Principle 3: Dynamic Initiative.

Dynamic initiative is the system's way of managing the handshake between human and machine. It deals with a very practical tension: we want to scale genuinely personalized dialogue, but we also do not want to sand down the parts of teaching that still depend on human judgment. So the system keeps a quiet watch on the interaction and looks for cues that it should step back—patterns like the learner getting trapped in the same logical loop, frustration rising in the tone or pacing, or the discussion drifting into ethical territory where a teacher's discretion matters more than a model's fluency [9].

Most of the time, during routine cognitive probing, the AI can and should take the lead. It has the stamina to ask the seventh follow-up question without impatience, and it can sustain structured inquiry across many learners at once. Still, the model should also know when to yield. Once certain thresholds are crossed, initiative shifts—ideally without drama—to a human teacher who can read context, respond to emotion, and weigh values in a way that is hard to formalize.

What emerges is less a replacement of teaching than a workable division of labor. The AI carries the scalable load of logic-building and disciplined questioning; teachers focus on the parts that rarely scale well but matter most when they do appear: nuanced guidance, emotional support, and mentorship shaped by human values.

3.2. The blueprint: a layered collaborative architecture

Principles are easy to admire in the abstract. The harder question is what they look like when a real system has to make decisions, sentence by sentence, with an actual learner on the other end. To operationalize the approach, we built a three-layer "Goal–Agent–Conversation" architecture. You can picture it as a chain: broad educational intentions sit at one end, and the live back-and-forth of dialogue sits at the other, with a set of translating mechanisms in between.

The Goal Layer functions as the strategic command center. It defines what "success" means for a given session by aligning the dialogue with curriculum expectations and targeted thinking skills. In practice, this layer keeps the system honest. It prevents the conversation from drifting into cleverness for its own sake by continually pulling the interaction back toward its educational purpose.

The Agent Layer sits in the middle as the system's real-time orchestrator. It receives direction from above and reads signals from below, then decides what to do next: press for deeper reasoning, soften the scaffolding, pivot the questioning strategy, or hand the initiative to a human teacher. This is where the "dynamic initiative" principle stops being a slogan and becomes a control mechanism—the hinge that decides who should lead, and when.

The Conversation Layer is the point of contact, where the interaction actually happens. It is composed of three tightly coupled modules that work almost like gears:

A dialogue-state tracker that maintains a live cognitive profile of the learner as the exchange unfolds.

A questioning-strategy selector that draws from a repertoire of Socratic moves—probing definitions, tracing implications, testing for consistency—based on what the learner has just said and what the system is trying to cultivate.

A natural-language generator that turns the chosen move into questions that sound fluent, varied, and situationally appropriate rather than formulaic.

DOI: 10.54254/2755-2721/2025.LD31407

When these layers stay coherent, the dialogue can feel like a thoughtful conversation with an attentive partner. When they fall out of sync, it collapses into something learners recognize immediately: a stiff, procedural interrogation.

3.3. The dance: the dynamic operating mechanism

The real craft here is not the mere act of questioning. It is the way questions move a learner—how they change what the learner notices, what they doubt, what they revise, and what they can finally articulate. To capture that motion, the system follows a four-stage cognitive spiral intended to mirror (and, in some places, strengthen) the kind of guidance a skilled mentor provides.

Stage 1: Probing and Externalizing.

The opening move is less about correctness than about surfacing the learner's starting point. The system asks questions that draw out unexamined assumptions and half-formed intuitions—those ideas people often carry confidently until someone invites them to spell them out. The prompts stay open-ended on purpose. By nudging the learner to put their thinking into words, the system turns implicit cognition into explicit material that can be examined, revised, and built upon.

Stage 2: Exposing Contradiction.

Once the learner's position has shape, the AI becomes a disciplined interlocutor. It tests the reasoning with careful pressure: it asks for evidence, challenges a premise, follows an implication to its endpoint, or offers a counterexample that quietly asks, "Does your claim survive this case?" The point is not to defeat the learner. The point is to help them reach the edge of their own logic—often marked by that small, revealing pause when things no longer quite add up. That moment of productive confusion is not an error state; it is the signal that deeper inquiry is now necessary.

Stage 3: Supporting Reflection.

After introducing friction, the system changes posture. It stops acting as the challenger and starts acting as the scaffold. It prompts the learner to look back at the assumption they just tripped over, to consider alternatives, and to evaluate the strength of their own argument. Metacognitive prompts play a central role here: the learner is asked not only what they think, but how they arrived there. The conversation becomes less about defending a position and more about inspecting the reasoning that produced it.

Stage 4: Constructing and Ascending.

Reflection creates room for reconstruction. Learners begin to assemble a more robust understanding—one that integrates the earlier tension rather than ignoring it. The system supports this synthesis, helping the learner connect insights into a coherent model they can actually use. And then, crucially, it does not treat that model as the finish line. It becomes the new baseline. The system introduces a more demanding question, and the spiral begins again—ideally each turn starting from a higher level of clarity than the last. Progress, in this view, looks less like a straight sprint toward an answer and more like an iterative climb: practice-driven, occasionally uncomfortable, and steadily more capable.

4. Looking ahead: reimagining practice through partnership

Developing an intelligent Socratic dialogue system is, at its heart, an attempt to let teaching philosophy and technology infuse one another. The immediate benefits show up in smoother interactions and more responsive guidance, but the deeper value lies in how it reframes persistent problems. This model starts from education's stubborn tensions and asks: if we took these

DOI: 10.54254/2755-2721/2025.LD31407

challenges seriously, what would we build? Its potential breakthroughs, therefore, aren't about flashy novelty, but about targeted responses to long-standing constraints.

4.1. Bridging the scale-vs-depth divide

At the heart of modern education sits a familiar contradiction. Traditional models prioritize scale and access, especially in foundational education [10], while deep, Socratic-style mentorship—so effective for higher-order thinking—remains hopelessly resource-intensive. The result is an inequitable reality: high-quality thinking instruction becomes a luxury good, reserved for the few.

Intelligent Socratic dialogue proposes a new path by reallocating cognitive labor. The AI agent can handle the scalable, parallel work of foundational questioning and logical probing—the very tasks that drain teacher bandwidth. This shifts the teacher's role. Freed from the intensity of real-time interrogation, they can redirect energy toward creative lesson design and exploratory pedagogy. In this partnership, a new practice emerges: teachers can analyze the "argumentation chains" recorded by the system—the visible pathways of student reasoning—to diagnose thinking patterns and design far more targeted interventions [11]. They also gain capacity for the one-on-one care that students with special needs require.

From an equity standpoint, widespread adoption could begin to dissolve the barriers of time and geography that currently gatekeep rigorous intellectual practice. Learners from vastly different backgrounds could access the same high-caliber dialogue platform. For remote or under-resourced regions, where educational infrastructure is thin, this isn't just an improvement—it could be a practical entry point for raising instructional quality altogether.

4.2. Toward intelligent quality assurance

In a lecture-centered model, "instructional quality" often rises and falls on factors that are difficult to standardize, let alone evaluate: a teacher's particular strengths, their bandwidth on a given day, the mood of a class after lunch, the small social currents that can either energize discussion or flatten it. That natural variability is not anyone's fault—but it does make fair resource allocation messy, and it makes rigorous, comparable assessment of teaching quality surprisingly hard to pull off.

The intelligent Socratic model offers a technical route toward quality assurance that is more stable without becoming sterile. Its advantage lies in delivering consistent, standards-aligned practice in thinking—questioning that is grounded in cognitive science, logically sequenced, and tied to explicit goals. At the same time, it resists the trap of rigidity. Because the system watches the learner's reasoning as it unfolds, it can adjust pacing and support in real time: widening the space when the learner is ready to generalize, tightening the frame when their argument starts to slip. This is what "standardized yet personalized" looks like in practice—clear instructional norms that still flex around individual cognition rather than forcing everyone into the same tempo.

Assessment changes just as sharply. The system's dialogue logs capture not only what a learner arrived at, but the route they took—the hesitations, the revisions, the points where a premise collapsed and had to be rebuilt. For administrators, that process-level record becomes a usable evidence base. They can look for group-wide patterns, identify where cohorts routinely struggle, and also trace individual trajectories over time with far more granularity than a test score can provide. There is an additional, quieter benefit: the system can function as a kind of backstage mentor for teachers themselves. Novice educators can learn faster by studying the system's questioning heuristics in action, while experienced teachers gain breathing room to innovate—building on a stable baseline instead of repeatedly reinventing core dialogic routines.

DOI: 10.54254/2755-2721/2025.LD31407

4.3. The data-informed personalization pathway

In conventional classrooms, teachers rarely lack commitment; they lack signal. Even excellent instructors often do not have timely, fine-grained insight into each student's current cognitive state. When information is thin, personalization tends to lean on intuition—sometimes brilliantly, but rarely in a way that scales beyond a small number of students.

This model reverses that constraint. As learners interact with the system, it continuously captures trace data of their thinking—what they claim, how they justify it, where they hedge, what kinds of prompts unlock progress, and what kinds reliably trigger confusion. Over time, those traces support a high-resolution cognitive model that maps knowledge structures, reasoning habits, and emerging needs. With that empirical footing, the familiar ideal of teaching within a learner's "zone of proximal development" shifts from an appealing metaphor to something closer to an operational target.

Adaptability is what makes the pathway work. When a student moves quickly, the system can introduce more demanding questions that push them toward higher-order exploration rather than letting them coast on easy wins. When a student needs time, it can slow the tempo and supply calibrated scaffolds—clarifying prompts, smaller steps, sharper comparisons—so the challenge stays constructive instead of demoralizing. The result is a responsive loop of "assess–intervene–reassess" that is almost impossible to sustain manually for more than a handful of learners at once.

Taken together, this approach points to a new phase of personalized learning. Traditional differentiation depends heavily on a teacher's tacit expertise—a precious resource and, by definition, scarce [12]. By lowering the cost of high-fidelity personalization through technology, the model opens a plausible path toward adaptive, thinking-centered education at scale. The broader implication is not merely technical. It suggests a way to make deep, equitable learning opportunities less of an aspirational slogan and more of a workable, system-level reality.

5. Conclusion and paths forward

This exploration of an intelligent Socratic dialogue model proposes a different role for AI in education. It moves beyond the tool that delivers information or automates tasks, positioning AI instead as a collaborator within the learner's own cognitive process. This shift addresses a core tension in mass education: how to nurture distinctive, individual thinking within systems designed for standardized delivery. Our model doesn't claim to resolve this tension completely, but it offers a fresh pathway—one that treats dialogue, not content, as the primary engine for learning.

Of course, a conceptual blueprint is just a beginning. For this paradigm to take root in real classrooms, work must advance along three interconnected frontiers.

Technical Refinement: The system's core intelligence needs deepening. It must get better at discerning not just what a learner says, but what they mean—detecting subtle logical bends and breaks in reasoning. Equally crucial is a sharper sensitivity to the emotional texture of dialogue: signs of frustration, anxiety, or disengagement that signal when a learner is nearing overload. Furthermore, since real inquiry rarely respects disciplinary borders, the model will need more robust "connective" thinking power—the ability to weave concepts across fields without resorting to vague generalities.

Contextual Adaptation: A one-size-fits-all implementation would miss the point. The dialogue's form must adapt to its subject. In STEM contexts, it might emphasize hypothesis testing and stepwise logic; in the humanities, it may live in comparative interpretation and analogical reasoning.

DOI: 10.54254/2755-2721/2025.LD31407

A major task ahead is charting these subject-specific pathways—translating the Socratic spirit into distinct question grammars and progress indicators for different domains.

Ethical Governance: We cannot afford to be reactive. Proactive, comprehensive guidelines for AI in education are needed—frameworks that go beyond platitudes about "responsible use" to concretely address algorithmic bias, data privacy, transparency, and accountability. The goal is clear yet demanding: to ensure technological development remains firmly aligned with the fundamental commitment to educational equity.

As this work progresses, we also need a richer, evidence-based conversation about how such a paradigm might reshape the educational ecosystem. If intelligent Socratic dialogue becomes a widespread cognitive infrastructure, it won't replace teachers. It could, however, powerfully complement what they uniquely provide: emotional resonance, ethical guidance, and the kind of creative spark no system can truly replicate.

In that collaborative space, a compelling possibility emerges: each learner might receive sustained, tailored support in developing their capacity to think—not just to answer. In one sense, this aligns with a natural trajectory for educational technology. In a deeper sense, it feels like a return: a move back toward the enduring idea that education is, at its heart, the cultivation of a mind.

References

[1] Zhao, X.-W., Shen, S.-S., & Zhu, Z.-T. (2024). Digital-intelligent Socrates: Shaping learners' subjectivity through dialogue. Zhongguo Yuancheng Jiaoyu [Distance Education in China], 44(6), 13–24.

[2] Zhou, G.-L., & Huang, Y.-L. (2025). Exploring how large language models empower ideological and political theory teaching in universities. Xuexiao Dangjian Yu Sixiang Jiaoyu [School Party Building and Ideological Education], (13), 70–73.

[3] Li, H., Wang, W., Li, G.-X., et al. (2024). Intelligent midwifery teaching method: A teaching practice case of the "Intelligent Socratic Conversational Robot". Kaifang Jiaoyu Yanjiu [Open Education Research], 30(2), 89–99.

[4] Liu, Z.-F. (2008). Modern applications of the Socratic teaching method. Gaojiao Fazhan Yu Pinggu [Higher Education Development and Evaluation], (2), 98–102, 124.

[5] Sun, Y.-J. (2020). Aristotle and Socrates' method of elenchus (refutation). Zhexue Yanjiu [Philosophical Research], (8), 98–107.

[6] Liu, B.-Y., & Wen, Y.-H. (2021). From propositions to questions: Building a "scaffold" for progressive thinking. Sixiang Zhengzhi Ke Jiaoxue [Ideological and Political Course Teaching], (9), 38–41.

[7] Liu, T.-F. (1998). Classical humanistic education: The direction of education and reflections on modern education (Part I). Jiaoyu Lilun Yu Shijian [Education Theory and Practice], (4), 6–9.

[8] Liu, N., & Yu, S.-Q. (2020). Research on precision teaching based on the zone of proximal development. Dianhua Jiaoyu Yanjiu [E-Education Research], 41(7), 77–85.

[9] Wu, A., Wang, H.-L., Liu, L., et al. (2025). A survey of large-scale models for emotion recognition. Jisuanji Kexue Yu Tansuo [Computer Science and Exploration], 1–32. Retrieved November 23, 2025, from https://link.cnki.net/urlid/11.5602.TP.20251121.1421.004

[10] Wang, J.-H. (2025). On the industrialization and individuation of education. Gaoxiao Jiaoyu Guanli [Journal of Higher Education Management], 19(4), 39–47.

[11] Hao, X.-J., & Gu, X.-Q. (2025). A new understanding of human–machine collaborative learning: Humans and AI co-regulate learning. Zhongguo Dianhua Jiaoyu [China Educational Technology], (6), 9–17, 79.

[12] Wu, N.-Z., Yuan, X.-M., & Hu, Z.-C. (2025). Tracing classroom forms of large-scale personalized instruction and their configurations. Dianhua Jiaoyu Yanjiu [E-Education Research], 46(10), 80–88.