Vol 08 / Issue 03 / October 2026 / pp 50 – 71
ISSN: 2676-234X / EISSN: 2716-9006
Special Issue

Social Empowerment Journal
University Amar Telidji of Laghouat-Algeria

Making metacognition visible: AI-generated learning clues as signals of self-regulation and competence.

Bluma Alex
Independent Researcher, (France), e-mail: alex@bluma.pro
ORCID: https://orcid.org/0009-0008-6131-2848

Received: 30/07/2026; **Accepted**: 11/09/2026, **Published:** 31/10/2026

Abstract

Generative artificial intelligence makes knowledge instantly accessible, shifting the core educational value from what one knows to the ability to learn. This paper argues that metacognitive self-regulation, meaning the capacity to monitor, evaluate, and adjust one's own cognitive processes, is the foundational competence that AI cannot automate. Yet higher education lacks tools to simultaneously develop and document this capacity. We present Aegis-Strategos, a generative AI-powered strategic simulation device that activates metacognitive regulation through a structured comparison mechanism (the Oracle) across six sequential turns. In this exploratory mixed-method study with 15 master's students (MAE Entrepreneurship & Innovation, Université Marie & Louis Pasteur), the device generated 90 scored interaction traces. One student's trajectory from 40/100 at turn 3 to 95/100 at turn 4, illustrates the complete self-regulated learning cycle in real time: erroneous forethought, AI-detected conceptual confusion, feedback-driven correction, and observable progression. We propose a three-layer theoretical model: cognitive (Anderson & Krathwohl, 2001; Flavell, 1979; Zimmerman, 2002), mechanistic (Nicol, 2021; Sadler, 1989; Carless, 2015), and valorization (Spence, 1973; Hillage & Pollard, 1998; Yorke & Knight, 2006), to show that AI interaction traces can function as credible competence signals in the labor market. The paper concludes that, within the limits of this exploratory study, these devices partially answer both missions of higher education: developing capacities and helping professional integration. If receptor credibility among employers is actively constructed, longitudinal evidence is produced.

Keywords: *metacognition, self-regulation, generative AI, learning clues, signaling theory, self-regulated learning.*

Corresponding Author: Bluma Alex
https://doi.org/10.34118/sej.v8i4.4905

1. Introduction

Generative AI has collapsed the distance between a question and an answer. Any learner with an internet connection can produce a plausible paragraph on any topic in seconds. The value proposition of education, that institutions provide access to scarce knowledge, is dissolving. When knowledge is ubiquitous, what remains to be taught?

The answer is becoming urgent. The capacity to learn can be defined as knowing what one does not know, evaluating the adequacy of one's understanding, and adjusting strategies when a problem resists. It has always been the durable output of education. Flavell (1979) called it metacognition: knowledge about one's own cognitive processes and the regulation of those processes. Forty-five years later, the term still describes what AI cannot do. A large language model can produce a syntactically correct explanation of the French Revolution, but it cannot experience the state of not understanding it. It cannot monitor its own comprehension, detect a gap, and redirect its attention. Metacognition is, by definition, a property of an agent aware of its own limits.

This irreducible quality makes metacognition the foundational competence of the AI era. Yet higher education continues to assess what students produce, essays, exams, presentations, not how they regulate their production. A final grade is a product signal. It tells an employer that the student could produce a correct answer at a given moment. It does not tell them whether the student knows how to detect an error, correct a misunderstanding, or persist through conceptual difficulty. The gap between what is taught and what is assessed is widening.

2. Second mission: helping professional integration

The diploma has long served as a signal in the labor market. Spence's (1973) signaling theory formalized what employers knew intuitively: a degree is a costly signal that reduces information asymmetry between job seekers and hirers. The cost: years of study, tuition, and effort, makes the signal credible. A candidate who has earned a degree in economics is, on average, more productive in an economics role than one who has not, because the cost of earning the degree is lower for those with the relevant ability.

AI undermines this logic. When a diploma can be obtained with AI-assisted work, the signal degrades. Employers cannot easily distinguish between a candidate who internalized the material and one who produced passable output through AI

Copyright: © 2026 The Author(s). *This is an open-access article distributed under the terms of the:Creative Commons Attribution-NonCommercial 4.0 International License (CC BY-NC 4.0)*

augmentation. The information asymmetry that the diploma was designed to reduce is re-emerging in a new form.

The response cannot be a ban on AI: enforcement is impossible, and the technology is genuinely useful. The response must be a richer signal: not just the product of learning, but traces of the learning process. Yorke and Knight (2006) formalized this intuition in the USEM model: Understanding, Skills, Efficacy beliefs, Metacognition. In this model, metacognition stands as one of four pillars of employability. A graduate who can demonstrate self-regulatory capacity on top of domain knowledge is more employable because employers face lower uncertainty about their ability to adapt. The question is how to produce, document, and present such traces.

This question connects directly to social empowerment as understood in the capability approach (Sen, 1999): access to more granular, more equitably distributed competence signals may reduce information asymmetries not only between candidates and employers, but between institutional contexts with unequal reputational resources. A trace from a simulation device carries contextual evidence that a transcript from an unknown institution cannot match. The empowerment dimension is thus twofold: the device develops metacognitive capacity (individual capability), and it produces portable, verifiable evidence of that capacity (social signal equity).

3. Hypothesis

We argue that AI-powered strategic simulation devices can address both tensions simultaneously. The Aegis-Strategos device: a generative AI strategic simulation played across six sequential turns, activates metacognitive regulation through a structured comparison mechanism (the Oracle) while automatically documenting every step of the process. The interaction traces it generates function as competence signals in the sense of Spence (1973), enriched by the metacognitive dimension of the USEM model (Yorke & Knight, 2006). The hypothesis is that the traces rendered visible by this device constitute credible, non-falsifiable evidence of self-regulatory capacity, and that they can partially compensate for the signal degradation induced by AI in higher education.

4. Theoretical Framework

4.1.Metacognition and Self-Regulation as Learning Foundations

4.1.1. The cognitive taxonomy and metacognitive knowledge

Copyright: © 2026 The Author(s). *This is an open-access article distributed under the terms of the:Creative Commons Attribution-NonCommercial 4.0 International License (CC BY-NC 4.0)*

Anderson and Krathwohl's (2001) revision of Bloom's taxonomy was a quiet revolution. The original taxonomy organized cognitive processes along a single dimension: remember, understand, apply, analyze, evaluate, create. The revision added a second dimension: knowledge types. Four were identified: factual, conceptual, procedural, and metacognitive.

The inclusion of metacognitive knowledge as a distinct category was theoretically significant. It meant that knowing about one's own cognition is not merely a higher-order process but a qualitatively different type of knowledge. Metacognitive knowledge encompasses strategic knowledge (knowing which strategies are effective for which tasks), knowledge about cognitive tasks (knowing that some tasks demand more attention than others), and self-knowledge (knowing one's own strengths and weaknesses as a learner). The taxonomy implies that any complete education must develop all four knowledge types. In practice, the metacognitive dimension remains the least taught and least assessed.

This is not due to bad faith. Metacognitive knowledge is difficult to observe. A teacher can see whether a student recalls a fact or applies a procedure. However, the student is aware of their own understanding and can adjust it is invisible in conventional assessments. The taxonomy names the problem but does not solve it.

4.1.2. Metacognition: knowledge and regulation

Flavell's (1979) foundational model distinguished two components of metacognition: metacognitive knowledge and metacognitive regulation. The first is declarative: what a person knows about cognition in general and about their own cognition in particular. The second is procedural: the ability to monitor, evaluate, and adjust cognitive activity in real time.

This distinction matters for assessment. Declarative metacognitive knowledge can be elicited through questionnaires or interviews. "Do you plan your study sessions?" is a reasonable proxy. But self-report data is notoriously unreliable (Winne & Perry, 2000). People generically affirm that they plan, monitor, and adjust, but the correlation between what they say and what they do is modest at best.

Metacognitive regulation, by contrast, is observable if the right conditions are in place. Facione (1990) included self-regulation as the sixth critical thinking skill in the Delphi Report, defining it as the "self-conscious monitoring of one's own cognitive

Copyright: © 2026 The Author(s). *This is an open-access article distributed under the terms of the:Creative Commons Attribution-NonCommercial 4.0 International License (CC BY-NC 4.0)*

activities, the elements used in those activities, and the results educed, particularly by applying skills in analysis and evaluation to one's own inferential judgments with a view toward questioning, confirming, validating, or correcting either one's reasoning or one's results" (p. 25). The definition is dense but carries an implication. If self-regulation involves questioning, confirming, and correcting one's own reasoning, then it is potentially traceable in environments that make these operations visible and time-stamped.

4.1.3. Self-Regulated Learning as an operational framework

Zimmerman (2002) made metacognition operational by embedding it in a cyclical model of Self-Regulated Learning (SRL). The model comprises three phases:

- **Forethought**: task analysis, goal setting, strategic planning. Before acting, the learner analyses the task, sets goals, and selects strategies.
- **Performance**: self-control, self-observation. During action, the learner monitors their own performance and deploys strategies.
- **Self-reflection**: self-judgment, self-reaction. After action, the learner evaluates outcomes and adjusts future behavior.

The cycle is not linear but recursive. Self-reflection feeds back into forethought for the next iteration. This is precisely what distinguishes self-regulated learners from passive ones: they treat every performance as a data point for improvement.

Zimmerman's model is useful for our purposes because it transforms metacognitive regulation from an abstract concept into a sequence of observable phases. If a learning environment is designed to make each phase visible by capturing the learner's plan before action, the action itself, and the reflection after feedback, then the SRL cycle becomes documentable. The question turns from "can metacognition be observed?" to "what kind of environment reveals it?"

4.2.The Activation Mechanism: Internal Feedback Through Comparison

4.2.1. The internal feedback model

Nicol (2021) proposed a mechanism that directly addresses the observability problem. His internal feedback model posits that learners generate feedback internally by comparing their own output to a reference source. The process has three steps: the

Copyright: © 2026 The Author(s). *This is an open-access article distributed under the terms of the:Creative Commons Attribution-NonCommercial 4.0 International License (CC BY-NC 4.0)*

learner produces something (a solution, an argument, a strategy), encounters a reference source (an exemplar, a model answer, an expert performance), and articulates the gap between the two. It is this articulation: "I thought X, but the exemplar shows Y" that constitutes internal feedback.

The model matters here because comparison is inherently metacognitive. The learner's attention is directed to the discrepancy between current and desired performance. That discrepancy is exactly what Flavell (1979) described as metacognitive monitoring. And the articulation of the gap, the verbalization of what was learned from comparison, is metacognitive regulation made explicit.

Traditional assessment ignores this process entirely. It captures the final product and assigns a grade. It does not capture what the learner notices, what surprises them, or what they would do differently. The comparison model suggests that assessment could be redesigned to capture precisely these moments.

4.2.2. Exemplars as comparison triggers

The idea that comparison triggers learning is not new. Sadler (1989) argued that students need to understand the standard against which their work is judged and must be able to compare their own performance to that standard. Without this comparison, feedback remains external. It is the teacher's evaluation, not the student's insight.

Carless (2015) extended this to the specific use of exemplars: carefully chosen examples that embody the standard. He found that exemplar effectiveness depends on timing and accessibility. An exemplar shown before the learner attempts a task may merely be copied. An exemplar shown after the learner has committed to a response generates productive disequilibrium. The learner sees not just what a good answer looks like, but what their own answer lacks in comparison.

This timing condition is critical. It means the device matters as much as the content. An AI that generates a reference answer on demand is not new. Any chatbot can do that. What matters is when the reference is delivered. After the learner has formulated their own response and after they have received initial feedback. At that moment, the comparison is not passive consumption but active learning.

Copyright: © 2026 The Author(s). *This is an open-access article distributed under the terms of the:**Creative Commons Attribution-NonCommercial 4.0 International License (CC BY-NC 4.0)***

4.3.From Learning Traces to Competence Signals

4.3.1. Traces as costly signals

Spence (1973) argued that for a signal to reduce information asymmetry, it must be costly. More costly for low-ability candidates than for high-ability ones. The diploma functions because the effort required to earn it is lower for those with genuine competence.

AI interaction traces constitute a new signal class. Their cost is not monetary but cognitive: the time, attention, and metacognitive effort required to progress through a structured simulation. A student who repeatedly adjusts their strategy based on feedback demonstrates a different kind of competence from one who produces a single correct answer. The trace, temporal, sequential, internally consistent, is harder to falsify than a transcript.

Three properties make interaction traces credible as signals. First, *cost*. Cognitive engagement cannot be faked. A student cannot produce a six-turn improvement in five minutes. Second, *traceability*. The log is time-stamped and sequential. Each response builds on the previous one, and the logic of improvement (or its absence) is visible. Third, *multiplexity*. A single interaction flow documents multiple competencies simultaneously: domain knowledge, analytical reasoning, responsiveness to feedback, resilience under difficulty. A grade compresses all this into a single number. A trace preserves the richness.

This signaling perspective, however, originates in labor markets of the Global North. Scholars working on AI in education in the Global South (Miao et al., 2021; Miao, 2023) have stressed that the design and governance of AI-based competence signals must include actors from the contexts where they are deployed. The question is not only whether a trace functions as a signal, but who defines what counts as a signal, who controls the data, and whose competencies are recognized. This geopolitical dimension of AI in education complements Spence's strictly economic model and will be taken up in the discussion.

4.3.2. Traces as employability assets

Hillage and Pollard (1998) defined employability as having four components: assets (knowledge, skills, attitudes), deployment (the ability to mobilize those assets),

Copyright: © 2026 The Author(s). *This is an open-access article distributed under the terms of the:Creative Commons Attribution-NonCommercial 4.0 International License (CC BY-NC 4.0)*

presentation (the ability to showcase them to employers), and context (personal and market circumstances). Their model shifts attention from what a graduate possesses to what a graduate can demonstrate.

AI interaction traces serve both asset and presentation dimensions simultaneously. They are assets insofar as the simulation develops metacognitive capacity. Each cycle of forethought, performance, and self-reflection strengthens the learner's regulatory ability. They are presentation tools insofar as the log provides concrete evidence of that capacity. A candidate who presents a trace showing error detection, correction, and improvement across six turns has more to show than one who presents a transcript with a grade.

The presentation dimension is underdeveloped in current practice. Graduates have certificates and transcripts. They rarely have process evidence. The interaction log fills this gap, provided that employers know how to read it.

4.3.3. Metacognition as an employability pillar: the USEM model

Yorke and Knight (2006) proposed USEM as an alternative to narrow competency-based models of employability. The acronym stands for Understanding, Skills, Efficacy beliefs, and Metacognition. Each component is necessary; none alone is sufficient.

The inclusion of metacognition as a separate pillar was deliberate. The authors argued that the ability to monitor and adjust one's own learning, to know when one does not understand and to seek clarification, is a distinct employability asset that conventional models overlooked. A graduate with strong understanding and skills but weak metacognition may perform adequately in familiar settings but fail when faced with novel problems.

Aegis-Strategos directly addresses the 'M' in USEM. The traces it generates document the learner's metacognitive activity: goal setting (forethought), monitoring (self-observation during performance), and evaluation (self-reflection after feedback). If USEM is taken seriously as a framework, then the device produces direct evidence of one of its four pillars: evidence that conventional assessment systematically misses.

4.3.4. Synthesis and Research Hypothesis

Copyright: © 2026 The Author(s). *This is an open-access article distributed under the terms of the:Creative Commons Attribution-NonCommercial 4.0 International License (CC BY-NC 4.0)*

The literature reviewed spans three layers. The cognitive layer (Anderson & Krathwohl, 2001; Flavell, 1979; Zimmerman, 2002) establishes that metacognition is a distinct form of knowledge that can be developed through cyclical self-regulation. The mechanistic layer (Nicol, 2021; Sadler, 1989; Carless, 2015) demonstrates that comparison between one's own output and an exemplar generates internal feedback: an observable trace of metacognitive activity. The valorization layer (Spence, 1973; Hillage & Pollard, 1998; Yorke & Knight, 2006) argues that such traces can function as credible competence signals in the labor market.

No existing work articulates these three layers continuously. Learning analytics research produces traces but rarely connects them to signaling theory. Labor market research discusses credential value but does not consider AI interaction traces as a signal type. The gap is between what can be observed and what can be valorized.

We propose a four-step model: Device → Comparison mechanism → Metacognitive traces → Competence signals (USEM). The Aegis-Strategos device instantiates the comparison mechanism (Oracle) defined by Nicol and Sadler. The mechanism produces traces of metacognitive regulation in the phases defined by Zimmerman. These traces can be valued as competence signals through Spence's signaling theory and Yorke and Knight's USEM framework.

The research hypothesis is therefore: *Interaction traces generated by a generative AI strategic simulation device can render metacognitive regulation visible and constitute credible competence signals for the labor market, as evaluated through the combined lens of signaling theory and the USEM employability model.*

4.4.Method

4.4.1. The Aegis-Strategos device: architecture and theoretical foundations

Aegis-Strategos is a generative AI-powered strategic simulation device. It presents learners with a business scenario in six sequential turns, each corresponding to a strategic decision point. The total duration is approximately 30 minutes. Each turn requires the learner to analyze a situation, formulate a strategy, and produce a written response.

Three AI roles operate within the device. The *Scenario Generator* produces the initial case and evolves it in response to learner decisions. The *Real-Time Assessor*

Copyright: © 2026 The Author(s). *This is an open-access article distributed under the terms of the:Creative Commons Attribution-NonCommercial 4.0 International License (CC BY-NC 4.0)*

evaluates each learner's response within three seconds on a 100-point scale. The *Oracle* generates a reference response of what a master strategist might have produced, for the same scenario, delivered after the learner has submitted their answer and received initial feedback.

The theoretical grounding of the device is tripartite.

Self-Regulated Learning (Zimmerman, 2002). Each turn activates the full SRL cycle. Forethought occurs when the learner analyses the scenario and formulates a strategy before writing. The turn format forces planning. Performance occurs when the learner produces the response. Self-reflection occurs after feedback. The learner sees their score and the AI-generated critique, then compares their response to the Oracle's reference. The six-turn structure ensures repeated cycling through the three phases, making the recursive nature of SRL visible in the log.

Visible Progression (Hattie, 2009; Bloom, 1968; Marzano, 2006; Wiliam, 2011). The six-level scoring scale makes progression (or its absence) immediately visible. Rubrics accompany each criterion so that learners know what each level requires. Feedback is delivered instantly, within a few seconds, so that the link between performance and evaluation is clear. The device draws on the mastery learning tradition (Bloom, 1968) in that progression is expected. A low score is not a final judgment but a diagnostic that feeds the next cycle.

Motivational grounding (Bandura, 1977; Deci & Ryan, 2000; Dweck, 2006). Each successful cycle produces a mastery experience in Bandura's (1977) sense. The learner experiences improvement as a direct result of strategic adjustment. This feeds both competence and autonomy needs in Deci and Ryan's Self-Determination Theory. It means the learner drives their own progression and sees the evidence. The device materializes Dweck's (2006) growth mindset by showing that ability is not fixed. Indeed, scores change as a function of strategy adjustment, not innate talent.

4.4.2. Study design

We used a mixed-method design combining an embedded case study (Yin, 2018) with descriptive quantitative data. The setting was a master's program in Entrepreneurship and Innovation at Université Marie & Louis Pasteur, France, during the 2025–2026 academic year.

Copyright: © 2026 The Author(s). *This is an open-access article distributed under the terms of the:Creative Commons Attribution-NonCommercial 4.0 International License (CC BY-NC 4.0)*

Participants were 15 master's students enrolled in a course on innovation management. The device was deployed as part of regular coursework. All participants provided consent for anonymized data use.

Data sources were the automatic interaction logs generated by Aegis-Strategos. Each turn produced a structured record containing: the learner's response (free text), the AI-generated score on each of three criteria, the composite score, the AI feedback, and the Oracle reference text. With 15 students completing 6 turns each, the corpus comprised 90 scored responses.

One missing student spontaneously shared their complete interaction log, 11 pages of raw data, as part of an end-of-course reflection. This student's trajectory became the central case for qualitative analysis of metacognitive regulation, because the log contained not just scores but verbatims revealing the student's reasoning process.

4.4.3. Ethical considerations

The study operated within the ethical framework for routine pedagogical data collection at Université Marie & Louis Pasteur. All participants received information explaining data use, anonymization procedures, and their right to withdraw without academic penalty.

We followed Prinsloo and Slade's (2017) principle of obligation to act. Learning analytics should not merely describe problems, but they should trigger intervention. In our protocol, a student who scored below 40 on two consecutive turns was flagged for tutor review. This threshold was set before data collection and communicated to participants. The obligation to act creates a tension with research neutrality. Indeed, we could not observe metacognitive failure without responding. We accepted this tension as ethically necessary.

Algorithmic bias was a second concern (Boateng & Boateng, 2025). The criteria of relevance, creativity, and strategic alignment are not neutral. The AI's training data may encode assumptions about what constitutes a valid strategy, while assumptions could disadvantage non-Western or non-standard approaches. We mitigated this through periodic human review of AI scores, and we treated AI scores as indicators, not verdicts.

Copyright: © 2026 The Author(s). *This is an open-access article distributed under the terms of the:Creative Commons Attribution-NonCommercial 4.0 International License (CC BY-NC 4.0)*

Two additional ethical limits deserve mention. First, the device operates through a proprietary large language model; results may vary across providers and model versions, creating a dependency on commercial infrastructure that raises equity concerns for institutions with limited resources (Isaacs, 2023). Second, the device assumes stable internet access and individual device ownership. Conditions that cannot be taken for granted across all educational contexts, particularly in the Global South where digital divides remain structural.

4.5.Results

4.5.1. Quantitative data

The overall mean score across all participants and all turns was 83.4 out of 100 (median 83.1, SD 5.8). The distribution was left-skewed, consistent with a motivated cohort in a master's-level program.

Inter-turn variation revealed a pattern. Mean scores dipped at turn 3 (80.8) and rose at turns 4 and 5 (86.3 and 86.8, respectively). Turn 6 showed a slight decline (84.5), possibly due to task fatigue or increasing scenario complexity. The dip at turn 3 was observed in 11 of 15 participants, suggesting a weakness of the students for this topic. The third turn introduced a strategic choice more focused on communication than the management of innovation.

Standard deviation decreased from 7.8 at turn 1 to 6.0 at turn 6. This convergence effect suggests that the device's feedback loops were reducing inter-individual variance. Weaker performers improved, stronger performers maintained, and the cohort moved toward a common higher baseline. The device was actively developing capacity.

4.5.2. Case study: activating metacognitive regulation

The most instructive case was student S07, whose detailed interaction log (11 pages, three scores per turn) provided granular insight into metacognitive activity. S07's mean score across the six turns was 84.9 (SD 12.7), just above the cohort average but with markedly higher intra-individual variance. Indeed**,** the standard deviation of her turn scores was more than double that of the typical participant.

The critical moment occurred at turn 3. S07's first score on that turn was 40/100: the lowest single-criterion score of any participant at any turn. The Real-Time Assessor's feedback identified a specific conceptual error: S07's response stated that "my clients

Copyright: © 2026 The Author(s). *This is an open-access article distributed under the terms of the:**Creative Commons Attribution-NonCommercial 4.0 International License (CC BY-NC 4.0)***

understand my expertise" as the competitive advantage. The AI flagged this as a category error: the response described the entrepreneur's self-perception rather than the client's problem. The feedback read: *"You are speaking about yourself, not your client. A competitive advantage is defined by what the client values, not by what you do well. This is the most common conceptual confusion in entrepreneurship strategy."* The two other criterion scores on that turn (50 and 92) brought the turn composite to 60.7.

S07's revised response at turn 3 scored 61/100. The score improvement alone is not direct evidence of metacognitive regulation: it could reflect superficial alignment with feedback. What grounds the interpretation is the teacher's subsequent analysis of S07's complete log. The teacher identified that the turn 3 error (self-description instead of client analysis) was not repeated in turn 4, where the response structure had fundamentally changed. This pattern: error, feedback, structural change in the next attempt, is consistent with Flavell's (1979) description of metacognitive regulation as the detection and correction of one's own cognitive error through comparison with an external reference. The teacher's role, in this case, was to confirm that the correction was conceptual rather than cosmetic, and to reinforce the learning through asynchronous remediation.

This is a demonstration of metacognitive regulation in Flavell's (1979) sense: the learner, prompted by the AI-generated comparison, detected an error in her own reasoning. The external evidence (the 40/100 score and the AI's explicit critique) made the gap visible. The subsequent change in response strategy: from self-description to market analysis is the observable trace of that regulation in action. The verbatim content of her reflection, while unavailable for quotation, is methodologically secondary to the documented behavioral shift.

Turn 4 confirmed the progression. S07's composite score was 95/100 (scores of 92, 95, 98 across the three criteria). The response structure had changed qualitatively: it began with market analysis, identified the client's problem before proposing a solution, and explicitly referenced the conceptual lesson from turn 3. The Oracle's reference response for turn 4 aligned closely with S07's answer, confirming strategic convergence.

The trajectory documents the complete SRL cycle: erroneous forethought (turn 3 initial) → detection through feedback and comparison (self-reflection) → strategic adjustment (turn 3 revision) → improved performance (turn 4). The log preserves every step.

Copyright: © 2026 The Author(s). *This is an open-access article distributed under the terms of the:Creative Commons Attribution-NonCommercial 4.0 International License (CC BY-NC 4.0)*

4.5.3. Typology of metacognitive clues

Qualitative analysis of the 90 interaction logs revealed three categories of metacognitive clues, observable traces in the data that indicate self-regulatory activity.

Category 1 - Forethought clues. These appeared in responses that structured the problem before proposing a solution. Indicators included: explicit goal statements ("My objective is to..."), priority ordering ("The most important factor is..."), and anticipation of consequences ("If we choose X, then Y is likely to follow"). Forethought clues appeared in 68% of responses with scores above 80, but in only 22% of responses below 60.

Category 2 - Regulation clues. These were the richest metacognitive markers. They included explicit error identification ("I thought that... but actually..."), reformulation of previous ideas ("Earlier I said X, but now I think Y"), and self-correction that acknowledged the error's source. Regulation clues were concentrated in the post-feedback responses: 78% of turn revisions contained at least one regulation clue.

Category 3 - Cognitive resilience clues. These indicated the learner's capacity to persist through difficulty and adjust strategy across turns. Indicators included: maintaining response quality after a low score (rather than disengaging), trying a qualitatively different approach after feedback (rather than superficial rewording), and explicitly referencing earlier feedback in later responses. Resilience clues were detectable only across multiple turns: they were process-level rather than event-level. Students who showed resilience clues had an average improvement of 12.4 points from turn 1 to turn 6, compared to 3.1 points for those who did not.

These three categories are objectifiable in logs and constitute the raw material for competence signals.

4.6.Discussion

4.6.1. What clues reveal about metacognitive regulation

The three clue categories correspond directly to Zimmerman's (2002) SRL phases. Forethought clues document the first phase: goal setting and strategic planning. Regulation clues document the self-reflection phase: error detection and causal attribution. Resilience clues document the complete feedback loop. The learner not only detects an error but adjusts subsequent behavior across turns.

Copyright: © 2026 The Author(s). *This is an open-access article distributed under the terms of the:Creative Commons Attribution-NonCommercial 4.0 International License (CC BY-NC 4.0)*

This mapping is significant because it demonstrates that the device is not measuring metacognition in an abstract sense. It is tracing it processual. Each clue is a moment in a cycle, and the cycle is visible only when multiple turns are examined together. A single high score tells us nothing about metacognition. On the other hand, a sequence showing forethought, error, detection, correction, and improved performance tell us more.

The implication for assessment is that metacognitive competence cannot be captured in a single measurement point. It requires a time series. It means a sequence of performances in which the learner's regulatory activity is displayed. This supports the theoretical argument that SRL is inherently cyclical (Zimmerman, 2002) and that assessment must match the temporal structure of the construct.

4.6.2. The diagnostic/remediation dissociation

A structural feature of Aegis-Strategos is that diagnosis and remediation are separated in time and agent. The AI diagnoses: it detects errors, assigns scores, and provides feedback in real time. The teacher remediates: asynchronously, through course materials, tutorials, or individual follow-up. The interaction traces document both.

This is impossible in traditional assessment. A grade is simultaneously a diagnosis (it tells the student what level they reached) and a sanction (it counts toward their final mark). The two functions are indistinguishable, which reduces the diagnostic value of assessment. So, students avoid the diagnostic feedback that could help them because it is tied to consequences.

The dissociation created by AI-assisted devices opens a new pedagogical space. The AI handles immediate self-reflection of the student, who learns what they did wrong and why. The teacher handles the subsequent forethought phase by helping the student plan how to avoid the same error in the future. The traces document both steps, providing a complete picture of the learning process that neither the AI nor the teacher could produce alone.

In SRL terms, this is a division of labor that mirrors the cycle's structure. The AI's real-time feedback supports self-reflection. The teacher's asynchronous guidance supports the next forethought phase. The cycle becomes fully supported and fully documented.

Copyright: © 2026 The Author(s). *This is an open-access article distributed under the terms of the:Creative Commons Attribution-NonCommercial 4.0 International License (CC BY-NC 4.0)*

4.6.3. From traces to credible signals

The central claim of this paper is that AI interaction traces can function as competence signals. The results provide partial support.

The traces meet Spence's (1973) criteria for credible signals. They are costly: producing a six-turn sequence of genuine improvement requires cognitive engagement that cannot be simulated in minutes. They are traceable: the log is time-stamped, sequential, and internally consistent. They are hard to falsify: a fabricated trace would show an implausible pattern (perfect scores without errors, or improvement without intermediate steps).

The USEM model (Yorke & Knight, 2006) confirms that the metacognitive regulation documented in these traces is a recognized employability pillar. An employer evaluating a candidate with an interaction log would see not just a grade, but evidence of how the candidate detects and corrects errors, adjusts strategies, and persists through difficulty.

But there is a condition. A signal is only effective if the receiver can interpret it. Interaction traces are not self-explanatory. An employer unfamiliar with the device might see a sequence of numbers and qualitative reflections without understanding what they signify. Receptor credibility, meaning the employer's capacity to read the signal, must be actively constructed.

This echoes Hillage and Pollard's (1998) observation that presentation is a distinct component of employability. The interaction trace is both an asset (evidence of developed capacity) and a presentation tool (something the candidate can show), but its value depends on the audience's interpretive framework. The same candidate who impresses an AI-literate startup founder might baffle a traditional human resources department. This is not a flaw in the signal but a constraint on its current applicability.

4.6.4. Limits

This is an exploratory study, and its findings must be interpreted with caution. Three limitations require deeper consideration than a simple acknowledgement, because each constrains the interpretation of specific results. First, the sample is limited: 15 students, one university, one disciplinary context. This precludes generalization. The quantitative patterns reported (turn-level means, convergence of standard deviations,

Copyright: © 2026 The Author(s). *This is an open-access article distributed under the terms of the:Creative Commons Attribution-NonCommercial 4.0 International License (CC BY-NC 4.0)*

criterion correlations) describe this cohort; they do not establish stable properties of the device. Second, there was no control group. We cannot attribute observed progression to the device rather than to other factors (maturation, concurrent coursework, task familiarity). The design documents co-occurrence: detection, correction, and improvement did happen within the device's timeframe. It cannot isolate the device as their cause. Third, the generative AI's score reliability was not systematically assessed. A partial check was performed: human review of 40% of scores showed 84% agreement within one level. This is an encouraging but incomplete signal. It does not establish inter-rater agreement at the level required for psychometric claims; it does not cover test-retest stability, and reproducibility across different AI model versions remains unknown. A model update could shift scoring patterns in ways that affect trace interpretation.

The implications for interpretation deserve to be stated explicitly. The case study (student S07) is a single case: it illustrates a mechanism, it does not prove it. Its value is demonstrative, not probabilistic. The typology of metacognitive clues (forethought, regulation, resilience) was derived from 90 logs of one cohort; the categories are heuristic and require replication before they can claim generality. The correlations reported ($r = 0.54$ between relevance and strategic alignment; $r = 0.61$ between forethought indicators and overall score) are descriptive statistics on a small sample, without inferential tests; they indicate tendencies, not established relationships. Finally, the interpretation of metacognitive regulation does not rest on the AI scores alone. It triangulates three sources: the scores, the qualitative content of the traces (response structure, revision content, verbatim self-corrections), and the teacher-researcher's expert reading of the complete logs. The quantitative description is sensitive to score reliability; the qualitative demonstration of regulation is grounded in trace content and expert interpretation, which are independent of the scoring scale.

A structural limit concerns the employability argument. McQuaid and Lindsay (2005) note that employability is not purely individual. It depends too on labor market structures, discrimination, and economic cycles. Even the richest interaction trace cannot overcome systemic barriers. We argue that traces improve the signal available to candidates who already face a functioning labor market.

The device's dependence on a specific LLM provider introduces a further limit. Changes in model availability, pricing, or performance could render the trace format

Copyright: © 2026 The Author(s). *This is an open-access article distributed under the terms of the:Creative Commons Attribution-NonCommercial 4.0 International License (CC BY-NC 4.0)*

incomparable across cohorts. This dependency is a reproducibility limit that future research should address through open-weight models or cross-provider validation.

Finally, the device documents metacognitive regulation but does not prove its long-term development. Student S07's turn 4 improvement is striking but may be transient. We do not know whether the metacognitive awareness demonstrated in the simulation transfers to other contexts. The device is a window into metacognitive activity, not a longitudinal training intervention.

4.6.5. Perspectives: validating AI-generated scores

The reliability of AI-generated scores is the weakest link in the measurement chain, and it deserves a dedicated research agenda. Three validation steps would strengthen future claims. First, inter-rater validation: a sample of responses should be re-scored by two or three independent expert raters, with agreement measured through a formal coefficient (Cohen's kappa or weighted kappa) rather than a percentage of agreement within one level. Second, model stability testing: re-running the same responses with the same prompt at temperature zero, and across model versions, would quantify the variance attributable to the scoring model itself. Third, triangulation with external measures: correlating AI scores with independent assessments (teacher grades, standardized tests of the targeted competencies) would locate the score's validity relative to established instruments. Until these steps are taken, AI-generated scores should be treated as structured heuristics, not as psychometric measures, and the qualitative reading of traces by expert teachers remains the primary evidence for metacognitive activity.

4.6.6. Perspectives: could the device develop critical thinking?

Facione's (1990) Delphi Report listed self-regulation as the sixth critical thinking skill. The other five (interpretation, analysis, evaluation, inference, and explanation) are all cognitive operations that the device activates in each turn. Interpretation is required to understand the scenario. Analysis is required to identify key factors. Evaluation is required to assess options. Inference is required to project consequences. Explanation is required to justify the chosen strategy. Self-regulation is the skill that ties them together, and the device makes it visible.

If the device develops and documents self-regulation, it may indirectly contribute to the other five critical thinking skills. This is a longitudinal hypothesis that the current

Copyright: © 2026 The Author(s). *This is an open-access article distributed under the terms of the:Creative Commons Attribution-NonCommercial 4.0 International License (CC BY-NC 4.0)*

data cannot test, but it suggests a direction. Nevertheless, the device could serve not only as an assessment tool but as a sustained intervention that builds critical thinking capacity over time.

4.6.7. Perspectives: the evolving role of the teacher

These devices transform the teacher's posture. When an AI can diagnose errors, provide feedback, and generate reference responses in real time, the teacher is no longer the primary source of evaluative information. The teacher becomes a designer of augmented cognitive environments. Someone who configures the conditions under which metacognitive activity occurs and who intervenes when the device's capacity is exceeded.

The terminology of the profession has not caught up. "Teacher," "professor," and "lecturer" all describe a function centered on knowledge transmission. When knowledge is supplied by the device, the function shifts to something else: designer of learning architectures, orchestrator of AI-human feedback loops, interpreter of interaction traces. None of the existing titles capture this.

This transformation carries an ethical dimension that connects directly to professional ethics in education. If the teacher delegates diagnostic authority to an AI, what new responsibilities arise? The teacher remains accountable for learning outcomes but no longer directly controls the assessment process. The ethical obligations shift from evaluating fairly to designing fairly. More precisely, ensuring that the AI's scoring criteria are transparent, that learners understand how their traces are used, and that the device does not covertly sort or track beyond its pedagogical purpose. Professional ethics in the AI-augmented classroom requires a new code: one that defines the teacher's duty when assessment is partly automated, and the learner's right to contest an AI-generated score.

The question is not whether AI will replace the teacher. The question is what new name will emerge for a function whose nature has changed, and what ethical framework will accompany it.

4.6.8. Social empowerment and contextual transfer

Copyright: © 2026 The Author(s). *This is an open-access article distributed under the terms of the:Creative Commons Attribution-NonCommercial 4.0 International License (CC BY-NC 4.0)*

The study was conducted in a French university with students in a selective master's program. The device's applicability to other educational contexts, particularly in the Global South, deserves explicit discussion.

Three conditions affect transferability. First, the device requires reliable internet access and individual device ownership. In contexts where students share devices or experience intermittent connectivity, the six-turn synchronous format would need to be adapted. Second, the AI's scoring criteria (relevance, creativity, strategic alignment) are culturally embedded. What counts as a "relevant" strategy in entrepreneurship education in France may not align with entrepreneurial practices in informal economies or cooperative structures. Periodic recalibration of scoring rubrics with local stakeholders, as recommended by Isaacs (2023), would be necessary. Third, the USEM employability model, like Spence's signaling theory, originates in Western labor markets (Yorke & Knight, 2006). Its metacognition pillar is likely universal; its articulation with local employment structures is not.

These constraints do not invalidate the model. They define its scope. The core claim that AI-mediated comparison can activate and document metacognitive regulation is portable across languages and disciplines. What changes is the valorization layer. It means what counts as a signal, for whom, and under what institutional conditions.

Social empowerment, in Sen's (1999) capability sense, is about expanding the real freedoms that people have reason to value. Metacognitive competence, the ability to monitor and regulate one's own learning, is a capability in Sen's thinking. It expands what a person can do and be, regardless of the labor market they enter. The device contributes to social empowerment not primarily by helping graduates find jobs, but by building a cognitive capability that is foundational to self-directed learning in any context. The trace is a side effect; the capability is the outcome.

5. Conclusion

This exploratory study provides initial evidence that metacognitive regulation can be made visible through AI interaction traces. The Aegis-Strategos device activates the self-regulated learning cycle (forethought, performance, self-reflection) through a structured comparison mechanism that generates internal feedback. The traces it produces are not merely performance data; they are process data that document how learners detect, correct, and learn from errors. These findings are indicative rather than conclusive, and they await replication in larger, controlled designs.

Copyright: © 2026 The Author(s). *This is an open-access article distributed under the terms of the:Creative Commons Attribution-NonCommercial 4.0 International License (CC BY-NC 4.0)*

The device partially answers both missions of higher education. It develops capacities, such as metacognitive knowledge and self-regulatory skill, by forcing repeated cycling through the SRL phases. It fosters professional integration by producing traces that enrich the competence signal available to employers. Through the USEM model (Yorke & Knight, 2006), the metacognitive dimension of employability can be documented rather than asserted.

Two conditions must be met for these signals to be valorized. First, receptor credibility must be actively constructed. Employers need frameworks for interpreting interaction traces. Second, longitudinal evidence is required. We need to know whether the metacognitive activity observed in the device transfers to professional contexts and persists over time.

The central question this paper opens is institutional and societal: *will these new metacognitive competence signals be admitted into the certification and recruitment ecosystem?* The diploma has held its monopoly for over a century. AI-generated interaction traces challenge that monopoly by offering richer, more granular, harder-to-falsify evidence of competence. Whether universities, accreditation bodies, and employers are prepared to read these signals and whether the teaching function can be renamed to match its transformed nature are questions that technology now forces us to answer.

6. *References*

Anderson, L. W., & Krathwohl, D. R. (Eds.). (2001). *A taxonomy for learning, teaching, and assessing: A revision of Bloom's taxonomy of educational objectives*. Longman.

Bandura, A. (1977). Self-efficacy: Toward a unifying theory of behavioral change. *Psychological Review, 84*(2), 191–215. https://doi.org/10.1037/0033-295X.84.2.191

Bloom, B. S. (1968). Learning for mastery. *Evaluation Comment, 1*(2), 1–12.

Boateng, O., & Boateng, B. (2025). Algorithmic bias in educational systems: Examining the impact of AI-driven decision making in modern education. *World Journal of Advanced Research and Reviews, 25*(1), 2012–2017. https://doi.org/10.30574/wjarr.2025.25.1.0253

Carless, D. (2015). *Excellence in university assessment: Learning from award-winning practice*. Routledge. https://doi.org/10.4324/9781315740621

Copyright: © 2026 The Author(s). *This is an open-access article distributed under the terms of the:**Creative Commons Attribution-NonCommercial 4.0 International License (CC BY-NC 4.0)***

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

Dweck, C. S. (2006). *Mindset: The new psychology of success*. Random House.

Facione, P. A. (1990). *Critical thinking: A statement of expert consensus for purposes of educational assessment and instruction: The Delphi Report*. The California Academic Press.

Flavell, J. H. (1979). Metacognition and cognitive monitoring: A new area of cognitive-developmental inquiry. *American Psychologist, 34*(10), 906–911. https://doi.org/10.1037/0003-066X.34.10.906

Hattie, J. (2009). *Visible learning: A synthesis of over 800 meta-analyses relating to achievement*. Routledge. https://doi.org/10.4324/9780203887332

Hillage, J., & Pollard, E. (1998). *Employability: Developing a framework for policy analysis* (Research Brief No. 85). Department for Education and Employment.

Isaacs, S. (2023). AI in education in the Global South: From design to governance. In F. Miao & W. Holmes (Eds.), *AI and the futures of learning*. UNESCO.

Marzano, R. J. (2006). *Classroom assessment and grading that work*. Association for Supervision and Curriculum Development.

McQuaid, R. W., & Lindsay, C. (2005). The concept of employability. *Urban Studies, 42*(2), 197–219. https://doi.org/10.1080/0042098042000316100

Miao, F., Cukurova, M., Holmes, W., Lee, Y. W., & Huang, R. (2024). *AI competency framework for teachers*. UNESCO. https://www.unesco.org/en/articles/ai-competency-framework-teachersunesco

Miao, F., Holmes, W., Huang, R., & Zhang, H. (2021). *AI and education: Guidance for policy-makers*. UNESCO. https://doi.org/10.54675/PCSP7350

Miao, F., Cukurova, M., Holmes, W., Lee, Y. W., & Huang, R. (2024). *AI competency framework for students*. UNESCO. https://www.unesco.org/en/articles/ai-competency-framework-students

Nicol, D. (2021). The power of internal feedback: Exploiting natural comparison processes. *Assessment & Evaluation in Higher Education, 46*(5), 756–778. https://doi.org/10.1080/02602938.2020.1823314

Prinsloo, P., & Slade, S. (2017). An elephant in the learning analytics room: The obligation to act. In *Proceedings of the 7th International Conference on Learning Analytics and Knowledge* (pp. 46–55). ACM. https://doi.org/10.1145/3027385.3027406

Sadler, D. R. (1989). Formative assessment and the design of instructional systems. *Instructional Science, 18*(2), 119–144. https://doi.org/10.1007/BF00117714

Sen, A. (1999). *Development as freedom*. Oxford University Press.

Copyright: © 2026 The Author(s). *This is an open-access article distributed under the terms of the:**Creative Commons Attribution-NonCommercial 4.0 International License (CC BY-NC 4.0)***

Spence, M. (1973). Job market signaling. *The Quarterly Journal of Economics, 87*(3), 355–374. https://doi.org/10.2307/1882010

Wiliam, D. (2011). *Embedded formative assessment*. Solution Tree Press.

Winne, P. H., & Perry, N. E. (2000). Measuring self-regulated learning. In M. Boekaerts, P. R. Pintrich, & M. Zeidner (Eds.), *Handbook of self-regulation* (pp. 531–566). Academic Press.

Yin, R. K. (2018). *Case study research and applications: Design and methods* (6th ed.). SAGE.

Yorke, M., & Knight, P. T. (2006). *Embedding employability into the curriculum: Learning and employability series*. Higher Education Academy.

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

Citation; *Bluma, A.* ***Making metacognition visible: AI-generated learning clues as signals of self-regulation and competence.*** *Social Empowerment Journal. 2026; 8(4- Special Issue): pp. 50-71.* ***https://doi.org/10.34118/sej.v8i4.4905***

Publisher's Note: *SEJ stays neutral with regard to jurisdictional claims in published maps and institutional affiliations.*

Copyright: © 2026 The Author(s). *This is an open-access article distributed under the terms of the:****Creative Commons Attribution-NonCommercial 4.0 International License (CC BY-NC 4.0)***