A collaborative co-intelligence draft open to peer-review critiques:

Dispute, disagree, correct and refine.

Slow Thinking and Deep Inquiry: A Framework of Purposeful AI Integration to Emancipate Learning

Executive Summary

The central challenge of AI in education today is not technological — it is attentional. The velocity of AI tool adoption routinely seduced organizations, educators, and learners, generating mountains of curated content while producing shallow engagement and minimal durable learning. The antidote is not less AI, but *more deliberate human thinking before, during, and after AI use*. Drawing on Kahneman's System 1/System 2 framework, cognitive load theory, Guided Inquiry Design, Bloom's Taxonomy, and the emerging science of scrap learning, this report proposes a practical framework for using slow thinking and deep inquiry to design and assess truly efficient, high-impact AI integration in learning.

The Core Problem: Fast Adoption, Shallow Impact

When educators and organizations rush to adopt AI tools, they often operate in what Daniel Kahneman calls **System 1** mode — fast, automatic, pattern-driven, and reactive. The result is a flood of AI-generated content, curated playlists, and repurposed materials that look productive but deliver little lasting change. Researchers describe this waste as **"Scrap Learning"** — learning that is delivered but never applied back on the job or in practice. Studies suggest that a substantial portion of all delivered training falls into this category, representing enormous waste of time, money, and attention.[1][2]

The parallel problem in AI adoption specifically is **Shiny Object Syndrome (SOS)**: the compulsion to continuously adopt new tools in response to hype rather than need. Organizations in its grip display tell-tale symptoms — constantly switching tools, lacking depth of mastery in any of them, measuring activity (content produced, tools tested) rather than outcomes (learning transferred, behavior changed). As one practitioner puts it bluntly: "Depth beats breadth. Consistency beats novelty. Mastery beats dabbling."[3][4]

The stakes extend beyond wasted time. Research published in 2025 found that AI, when used passively, quietly undermines the joy and effort of learning — students and adult learners offload cognition to machines and lose the productive struggle that builds genuine understanding. Seventy percent of teachers now report worrying that AI weakens critical thinking and research skills. The technology itself is not the problem; the absence of intentional design is.[5][6]

Slow Thinking as a Design Posture

What System 2 Thinking Looks Like Applied to AI Use

Kahneman's **System 2 thinking** is slow, deliberate, and effortful — the mode that scrutinizes assumptions, interrogates outcomes, and resists cognitive shortcuts. Applied to AI integration in learning, System 2 thinking means pausing to ask not *"Can AI do this?"* but *"Should AI do this here, and what will the learner be left with?"*[^1]

The Cognitive Synergy Framework articulates this as *"Slow Down to Speed Up"* — a principle that intentional pauses to evaluate which tasks need human attention and which can be delegated to technology are not inefficient, but strategic. Neuroscience supports this: reducing extraneous cognitive load keeps the prefrontal cortex engaged for creativity and problem-solving, instead of triggering stress responses associated with cognitive overwhelm. The paradox is that slowing down tool adoption accelerates meaningful learning outcomes.[^7]

Practically, this means building **pause protocols** into AI workflows: before deploying any AI tool, explicitly asking what cognitive work belongs to the learner, what belongs to the AI, and what the evidence standard for success will be. OpenAI's o1-series models illustrate this at the machine level — they demonstrate that longer deliberation (more "thinking time") produces measurably more accurate outputs, a principle that maps directly onto human learning design.[8][9][^7]

The Reflective AI Practitioner

Recent academic work (2026, ACM) on "Reflective AI" as a slow technology approach found that when learners were required to reflect on their own creative practice and design thinking through AI interaction, the quality of engagement — and outcomes — improved significantly compared to passive use. This points to a design principle: **reflective structure must be built into AI-assisted workflows, not added as an afterthought.**[^10]

The Three-Stage Metacognitive Reflection Framework operationalizes this by requiring learners to engage *before, during, and after* every AI interaction:[^11]

- **Before**: What do I already think? What is my goal? What cognitive work am I about to delegate, and why?
- **During**: Is this AI output helping me think more deeply or bypassing that thinking? Am I evaluating, or just accepting?
- **After**: What did I learn? What did I think for myself? Where did the AI compress or distort something I should understand?

This structure keeps the human as the architect of learning, using AI as a tool rather than a replacement for cognition.[^12]

Deep Inquiry as the Engine of Purposeful AI Design

Starting with Problem, Not Tool

One of the most consistent findings across frameworks is the imperative to **begin with purpose, not with technology**. Guided Inquiry Design (GID), developed by Kuhlthau, Maniotes, and Caspari, provides a research-backed sequence that maps directly onto AI-era learning: Open → Immerse → Explore → Identify → Gather → Create → Share → Evaluate. Each phase demands genuine curiosity, productive uncertainty, and reflective documentation — qualities that resist AI-generated shortcutting.[13][14]

Deep inquiry begins with what GID calls "authentic wondering" — real questions learners care about, not artificial prompts. When that foundation exists, AI becomes a genuine inquiry tool: it can accelerate exploration, surface diverse perspectives, scaffold evidence gathering, and support synthesis. When it is absent, AI merely generates the appearance of inquiry without the cognitive development that makes it valuable.[15][13]

The implication for design: **identify the genuine learning problem first, then ask which AI tools — if any — help learners work through it more deeply**. This reverses the typical pattern of acquiring tools and then searching for applications.[^14]

Bloom's Taxonomy as an AI Deployment Map

Bloom's Taxonomy provides a powerful framework for deciding *where* in a learning sequence AI should and should not be deployed. The key insight from recent research is that **AI is highly capable**

at lower-order tasks (remember, understand, apply) but that human learning and growth require engagement with higher-order tasks (analyze, evaluate, create).[9][16][^17]

Bloom's Level	What AI Can Do	What Humans Must Do
Remember	Generate flashcards, summarize texts, recall definitions[^17]	Verify accuracy, connect to prior knowledge, express in own words
Understand	Explain concepts, offer analogies, rephrase passages[^17]	Evaluate clarity, apply to new situations, test comprehension
Apply	Generate case studies, propose solutions, create problem scenarios[^17]	Adapt solutions to constraints, justify reasoning, debug outcomes
Analyze	Identify patterns, categorize data, compare perspectives[^17]	Critique AI's analysis, identify bias, form independent conclusions
Evaluate	Generate counter-arguments, compare options, create rubrics[^17]	Justify personal judgments, assess source quality, make decisions
Create	Brainstorm ideas, draft outlines, propose structures[^17]	Inject personal voice, iterate with purpose, own the final synthesis

Effective AI integration means **using AI to accelerate lower-order tasks so human energy can be invested in higher-order ones** — not substituting AI for the higher-order work itself. As the University of Texas's responsible AI adoption framework states: the learner must always be in the loop, maintaining agency over their intellectual output and learning decisions.[18][19][^9]

A Framework for Designing AI-Integrated Learning with Impact

The PIER Design Protocol (Purpose → Inquiry → Efficiency → Review)

Drawing from multiple frameworks, the following four-phase design protocol offers a slow-thinking approach to AI integration in learning contexts:

Phase 1: PURPOSE — Define the Learning Problem Before Any Tool

- What specific learning outcome matters, and for whom?
- What evidence would demonstrate that outcome has been achieved?

- What cognitive work *must* the learner do — and cannot be delegated?
- Is AI even the right tool, or is this a relational, mentored, or experiential problem?[19][14]

Phase 2: INQUIRY — Design for Depth, Not Coverage

- Structure learning around authentic questions, not content delivery[15][13]
- Map AI support to Bloom's lower levels; protect Bloom's upper levels for learners[^9]
- Build in structured reflection: before/during/after AI use[11][12]
- Require process transparency — inquiry logs, AI interaction documentation[^13]

Phase 3: EFFICIENCY — Deploy AI Where It Genuinely Helps

- Use AI to manage *intrinsic cognitive load* (adjusting difficulty, scaffolding concepts)[^7]
- Use AI to *reduce extraneous load* (clarifying instructions, improving accessibility)[^7]
- Use AI to *maximize germane load* — the productive mental effort that builds connections and meaning[^7]
- Limit AI tool adoption to a "core stack" of 2-3 deeply mastered tools rather than chasing novelty[^4]

Phase 4: REVIEW — Measure Impact, Not Activity

- Shift metrics from volume and satisfaction to learning effectiveness and transfer[20][2]
- Measure *scrap learning*: what percentage of delivered content is actually applied?[21][2]
- Use longitudinal study design: measure, observe, measure again — to see whether changes persist[^22]
- Collect evidence across four dimensions: **performance** (scores), **process** (how learners reached answers), **persistence** (continued engagement), **transfer** (applying concepts in new situations)[^22]

Cultivating Collaborative Intelligence: The Three-Mode Cognitive Ecology

Beyond Dual Process: The Tri-System Framework

Until recently, the dominant model of human thinking was Kahneman's dual-process framework: System 1 (fast, intuitive, automatic) and System 2 (slow, deliberate, analytical). A landmark 2026 Wharton study by Shaw and Nave has now extended this model into **Tri-System Theory** —

introducing **System 3** as artificial cognition: external, automated, data-driven, and capable of either supplementing or suppressing human reasoning. Their experiments with nearly 1,300 participants found that in 80% of instances where individuals consulted an AI, they accepted outputs without pausing to critically evaluate them — a phenomenon the researchers name **"cognitive surrender"**.[23][24][25][26][^27]

This is the central risk of fast AI adoption: not that the tools are poor, but that human cognition defaults to uncritical deference. The implication for collaborative intelligence design is profound — simply *using* AI does not activate System 2 thinking; it often bypasses it entirely. Designing for collaborative intelligence therefore means deliberately orchestrating all three cognitive modes — not just deploying the tool.[28][29]

The three modes, reconceived for learning design, can be understood as:

- **Mindless / Automatic (System 1 + System 3 in autopilot)**: Habitual, effortless AI use for routine tasks where cognitive load must be minimized — formatting reports, transcribing notes, scheduling, generating first drafts. When well-designed, "mindless computing" leverages the fast and automatic mental processes to reduce friction and free capacity for deeper work. The risk is that without explicit boundaries, this mode bleeds into work that *should* demand human judgment.[^30]
- **Mindful / Slow (System 2 deliberation)**: Reflective, intentional human thinking — evaluating AI outputs, asking whether assumptions hold, questioning outputs, noticing what is absent. This is the mode most at risk from cognitive surrender. It must be *structurally protected*, not left to willpower.[^24]
- **Overflowing / Fast + Creative (Beginner's Mind)**: The Zen concept of *shoshin* — beginner's mind — represents a third, generative mode: rapid, open, non-judgmental exploration in which expertise temporarily sets aside its own conclusions to see what AI, others, or unexpected data surfaces freshly. Shunryu Suzuki articulated this as "In the beginner's mind there are many possibilities, but in the expert's mind there are few". This mode is where innovative prompting, lateral hypothesis generation, and creative recombination happen.[31][32]

The art of collaborative intelligence is knowing *which mode to activate when* — and designing workflows that make that navigation explicit rather than accidental.[33][23]

The Cognitive Ecology Matrix

Collaborative intelligence uses all three modes in rhythm, not competition:

Mode	Cognitive Character	AI Role	Human Role	Risk if Misapplied
Mindless / Automatic	Habitual, frictionless	Executes, formats, retrieves, drafts[^30]	Sets up systems, reviews outputs	Cognitive surrender; deskilling[^29]
Mindful / Slow	Deliberate, evaluative	Surfaces options, reflects back, checks logic	Critiques, decides, integrates, judges[^23]	Paralysis; over-verification of low-stakes outputs
Overflowing / Beginner's	Open, generative, curious	Expands possibility space, offers unexpected angles[^31]	Explores freely, suspends judgment, recombines	Staying in exploration without shifting to synthesis

BCG's organizational learning research confirms that the highest-performing AI-integrated organizations are not those that simply adopt AI, but those that deliberately design the **mutual learning loop** — where human feedback actively shapes AI outputs and AI outputs actively develop human judgment — over time. Organizations combining strong organizational learning with AI-specific learning are **twice as prepared** to manage talent-related disruptions and 60-80% more effective at managing a volatile business landscape.[34][35]

Applying Collaborative Intelligence to High-Impact Use Cases

ROI Reporting and Impact Measurement

ROI measurement is where slow thinking most visibly outperforms fast adoption. Organizations treating AI as a *measured investment* achieve ROI rates of **55%** on their most advanced initiatives, compared to just **5.9%** for those taking an ad hoc approach. Yet an MIT study of 300 public AI implementations found that **95% of organizations** have yet to see any measurable financial return from their AI investments, and 42% of companies abandoned most of their AI projects in 2025 citing unclear value — a dramatic increase from 17% the previous year.[^36]

Collaborative intelligence applied to ROI reporting requires all three cognitive modes in sequence:

- *Mindless*: AI auto-generates baseline dashboards, data visualizations, and draft narrative summaries from structured data[37][38]

- *Mindful*: Humans interrogate the dashboards — asking what the numbers don't show, whether baselines were correctly set, and whether changes are causally linked to the intervention or coincidental[[39]][36]
- *Beginner's*: Strategic teams revisit the original theory of change with fresh eyes — asking whether the right things were measured at all, and what a genuinely useful outcome would look like for the audience[^20]

The critical slow-thinking question in any AI-assisted ROI report: *"Are we measuring what we care about, or what is easy to count?"*[^20]

Marketing Strategy Impact

AI-driven personalization in marketing demonstrates both the power and the trap of the automatic mode. Companies using AI-driven personalization strategies report average increases in consumer spending of **38%**, with 80% of businesses reporting increased consumer engagement when experiences are personalized. Yet this performance collapses when the human layer responsible for *strategic coherence* — brand values, ethical boundaries, contextual judgment — is removed in favor of pure algorithmic optimization.[[40]][36]

A collaborative intelligence approach to marketing means using AI to execute personalization and A/B testing at scale (mindless), while slow-thinking humans maintain strategic oversight of the narrative, the equity implications of targeting patterns, and the long-term trust equation with audiences (mindful). The overflowing/beginner's mode is invoked at campaign conception — setting aside past performance data and asking what genuinely resonates with the human beings being served.[[41]][31][^40]

Changing Mindsets: Metanoia vs. Paranoia

One of the deepest distinctions in applying AI to learning and organizational change is the difference between two categories of mindset shift:

Metanoia — the Greek word for a fundamental *change of mind and heart*, a transformative reorientation of how one sees — is the positive goal of deep learning design. Metanoia is not achieved through content delivery or AI-generated summaries. It requires productive struggle, personal confrontation with dissonance, and the kind of reflective practice that Metanoia SC describes as: *"bringing ourselves to the work every day open to change and to also be changed"*. AI can be an extraordinary catalyst for metanoia when it acts as a **cognitive mirror** — surfacing contradictions, challenging assumptions, offering perspectives that would not otherwise be encountered. But this requires mindful human engagement with those outputs, not passive consumption.[^42]

Paranoia — in this context — represents the overcorrection: the fear-driven, hyper-vigilant, suspicious relationship with AI and its outputs that also impedes learning and action. Researchers identify cognitive inflexibility and strong volatility-sensitivity as the computational markers of paranoid reasoning — a pattern where the learner or organization is so focused on what could go wrong with AI that it defaults to avoidance, excessive verification, or rejection of legitimate AI value. The antidote to paranoia is not naïve trust but **calibrated engagement**: learning to know when to trust the machine, when to verify, and when to override — through practice, feedback, and structural transparency.[43][44][^29]

Designing for metanoia (not paranoia) through AI means:

- Making uncertainty visible in AI outputs — good interfaces prompt critical evaluation rather than presenting binary, falsely confident answers[^29]
- Creating psychological safety where questioning the algorithm is treated as expertise, not inefficiency[^29]
- Building reflection into the workflow so that mindset shifts are documented, not just experienced and forgotten[^11]
- Rewarding the *detection of AI errors* as much as efficiency gains — this structurally resists both cognitive surrender and paranoid avoidance[^29]

Minimum Viable Products (MVPs) and Innovation

The MVP framework — build the minimum version needed to test a core assumption, release it to real conditions, and iterate based on what you learn — is a natural home for all three cognitive modes. AI dramatically accelerates the build phase (mindless/automatic), but the discipline of an MVP is fundamentally a slow-thinking exercise: *What is the one assumption we most need to test? What is the minimum we need to build to test it honestly?*[45][46]

Collaborative intelligence in MVP development looks like this:

1. **Slow thinking**: Define the core learning question and the minimum features needed to test it — resist AI's tendency to generate comprehensive feature lists[^46]
2. **Beginner's mind**: Use AI to stress-test assumptions by generating counter-use-cases, edge-case personas, and alternative framings you haven't considered[^47]
3. **Fast/AI execution**: Use AI to generate prototypes, code scaffolding, user story drafts, and test scripts at speed[^48]

4. **Mindful review**: Humans evaluate real user feedback without letting AI summarize away the qualitative nuance — the uncomfortable observations, the unexpected use cases, the emotional responses[^49]
5. **Metanoia check**: Periodically step back and ask whether the original problem definition was correct — or whether the evidence is pointing toward a deeper reorientation of the product/learning strategy[^42]

AI-accelerated MVP development can reduce time-to-market dramatically, but only if the human team retains genuine epistemic sovereignty over what the product is *for*. Without that, AI produces faster iteration cycles toward the wrong destination.[^48]

The Paranoia-Metanoia Spectrum as an Assessment Tool

One of the most practical applications of the three-mode framework is using the paranoia-metanoia spectrum as a *diagnostic* of where an individual, team, or organization sits in its AI relationship:

Position on Spectrum	Cognitive Signature	What It Looks Like in Practice	Intervention
Deep Paranoia	Hyper-vigilant, refuses AI use, excessive verification	Blocks AI tools, rewrites every AI output, sees AI as threat[^43]	Beginner's mind exercises; low-stakes AI experimentation with documented safety
Productive Skepticism	Calibrated, asks "why should I trust this?"	Verifies selectively, documents AI errors, maintains judgment[^29]	*Optimal zone* — reinforce and reward
Passive Acceptance	Mildly uncritical, accepts outputs with surface review	Edits lightly, rarely questions premises, satisfied with adequate[^27]	Slow-thinking protocols; metacognitive reflection requirements[^11]
Cognitive Surrender	Fully uncritical, defers all judgment to AI	Submits AI drafts unchanged, stops reasoning independently[24][25]	Structural interventions — deliberate error injection, feedback loops, incentive redesign[^29]
Metanoia	Transformed understanding; uses AI as change catalyst	Questions are deeper; learning generates new frameworks, not just new content[^42]	*Aspirational zone* — design for and celebrate

---

Assessment: Moving from Vanity Metrics to Impact Evidence

The Scrap Learning Problem in AI-Curated Content

The concept of *scrap learning* is particularly urgent in the AI context. AI makes it trivially easy to generate and curate enormous volumes of content — reading lists, resource collections, summarized materials, generated lessons. But volume of content has no correlation with depth of learning. High-performing organizations measure learning efficiency, effectiveness, and business outcomes — not volume, cost, or satisfaction.[[2]][50]

A rapid AI-assisted ROI evaluation approach involves analyzing training content *before it is built* against learner profile data to predict what percentage of material is likely to be applied in practice — and redesigning to approach "zero redundancy training time". This scrap learning analysis can be run using AI itself (Claude, GPT-4o) with strong predictive accuracy, turning the tool against the problem it tends to create.[^21]

OpenAI's Learning Outcomes Measurement Suite (LOMS)

OpenAI, in partnership with the University of Tartu and Stanford's SCALE Initiative, has developed the **Learning Outcomes Measurement Suite (LOMS)** — a framework for longitudinal studies measuring how AI use affects learning over time. Its practical steps offer a replicable blueprint:[^22]

1. **Pick a narrow, high-value use case** — start with one curriculum area or cohort[^22]
2. **Establish a baseline** — agree on what "progress" means before measuring AI impact[^22]
3. **Define what counts as learning evidence** — balance performance, process, persistence, and transfer[^22]
4. **Set governance and privacy expectations** — transparent communication, de-identification, human oversight[^22]
5. **Report on a cadence that supports decisions** — monthly pilots with early indicators, learning checks, and qualitative feedback[^22]

Qualitative Signals of Shallow vs. Deep AI Integration

Beyond quantitative metrics, the following contrasts serve as qualitative assessment signals:

Shallow / Reactive AI Use	Deep / Intentional AI Use
AI generates content; learners passively consume it[^18]	AI accelerates inquiry; learners actively evaluate and extend it[^18]
Metrics measure activity (content produced, tools used)[^2]	Metrics measure impact (behavior changed, knowledge transferred)[^20]
New tools adopted at every AI news cycle[^4]	Core stack of 2-3 mastered tools, others evaluated rigorously[^4]
Learning design starts with tools[^14]	Learning design starts with problem and learner[^14]
No reflection protocol; AI replaces thinking[^11]	Before/during/after reflection embedded in every AI interaction[^11]
Lower-order tasks remain with learners; AI handles analysis[^9]	Lower-order tasks offloaded to AI; humans engage in higher-order work[^9]
Impact not measured; satisfaction scores used as proxy[^2]	Scrap learning tracked; transfer and application evidence required[^22]

Avoiding Shiny Object Syndrome: A Governance Protocol

The pull of AI novelty is real and will not diminish. A structured governance protocol operationalizes slow thinking at the institutional or team level:

Before adopting any new AI tool, apply three filters:[^4]

1. Does this tool replace or improve a specific piece of our existing learning stack? If yes, test it. If no, it is content, not infrastructure.
2. Can we evaluate it using a real learning task — not a demo — in under 90 minutes? If it cannot demonstrate improvement on an actual learning challenge, move on.
3. Does it simplify the learning workflow or complicate it? Extra logins, new interfaces, and additional cognitive overhead without meaningful learning return are disqualifying factors.

Establish clear learning objectives before any AI deployment:[[51]][14]

- What is the specific learning problem this AI addresses?
- How will we know it worked? (Not: "Did learners use it?" but "Did they learn differently and better?")

- What is the cost of failure — wasted time, reduced critical thinking, dependency formation?

Build a Culture of Cognitive Stewardship:[5][19]

- Model slow thinking publicly: narrate decisions about when to use and not use AI
- Celebrate the productive struggle — design assessments that reward the *process* of wrestling with complexity, not just the polished output[13][5]
- Regularly audit the AI-human boundary: what cognitive work are learners doing less of because AI is doing it for them, and is that the right trade?[9][7]

The Positive Vision: AI as Cognitive Mirror, Not Oracle

The most generative reframing for AI in learning comes from the "Cognitive Mirror" paradigm — repositioning AI from an oracle that provides answers to a reflective surface that sharpens the quality of a learner's thinking. In this model, AI is deliberately designed to expose the *limits* of a learner's explanation, prompting deeper engagement rather than substituting for it.[^52]

This aligns with Pearson's 2026 data showing that active AI use (iterate, question, evaluate) versus passive AI use (read, accept, submit) produces dramatically better learning outcomes — with repeat active users 24x more likely to demonstrate deep learning behaviors. The variable is not the tool; it is the *posture* with which the learner engages it.[^53]

Slow thinking and deep inquiry are not resistances to AI — they are the conditions under which AI becomes genuinely transformative. When the human remains the architect of learning, when questions precede tools, when reflection is built into every AI interaction, and when impact (not activity) is measured, AI accelerates learning that matters. When those conditions are absent, it accelerates the production of content that doesn't.

Conclusion: The Design Choices That Matter

Purposeful AI integration in learning is, at its core, a series of design choices made before and after the technology is used — not during. The research is convergent: begin with purpose, protect higher-order cognitive work for learners, embed reflective structure around every AI interaction, measure scrap learning and transfer rather than volume and satisfaction, and govern AI tool adoption with the same rigor applied to curriculum design.

The question is not whether to use AI, but whether the humans directing its use are thinking slowly enough to ensure it serves learning rather than substituting for it. That slow thinking — that deep inquiry — is itself the most important learning outcome to cultivate.

References

1. System 1 and System 2 Thinking - The Decision Lab - System 1 thinking is a near-instantaneous thinking process while System 2 thinking is slower and req...

2. [PDF] Confronting Scrap Learning - OSU Community Development - A substantial portion of the learning that is delivered is never applied back on the job and ultimat...

3. How to Handle Shiny Object Syndrome in Data and AI - Hypercube - In this article, we'll explore the nature of shiny object syndrome and examine ways to overcome it, ...

4. Shiny Object Syndrome In AI (And How To Avoid Losing Your Focus) - If you spend ten minutes on LinkedIn, X (Twitter), or YouTube, you are immediately pulled into a vor...

5. How AI quietly undermines the joy and effort of learning - PMC - Policies may limit AI usage, but they cannot force learners to engage deeply or reclaim the joy of l...

6. Rising Use of AI in Schools Comes With Big Downsides for Students - One of the negative consequences AI is having on students is that it is hurting their ability to dev...

7. Slow Down to Speed Up: Designing Learning That Lets Students ... - Dr. Gloria Y. Niles shares why slowing down in the age of AI leads to deeper learning and better stu...

8. System 2 Thinking - by Jurgen Gravestein - Summary: AI labs are eager to solve System 2 thinking in AI in the hope of alleviating some of the i...

9. Using Bloom's Taxonomy to understand AI Adoption in Higher ... - To address this challenge, educators can use Bloom's taxonomy as a framework to determine what types...

10. Reflective AI: A Slow Technology Approach for Design Education - Our analytic approach addressed two areas of inquiry: (1) how students reflected on their own creati...

11. Three-Stage Metacognitive Reflection Framework for AI Engagement - These reflections help students navigate the balance between human and machine agency, build awarene...

12. Teaching Students AI Strategies to Enhance Metacognitive Processing - To foster deeper learning, students can develop skills in the three main areas of metacognition with...
13. Why Guided Inquiry Design® is the Essential Framework for AI-Era ... - Rather than viewing AI as a threat to authentic learning, educators using GID can position it as a p...
14. AI in Education: Toward Purposeful, Human-Centered Transformation - " When teachers maintain their agency and see AI as enhancement rather than replacement, adoption be...
15. When to Slow Down Learning: Why Seeking Meaning Matters More ... - As educators and librarians, we can slow down the learning process, not to do less, but to go deeper...
16. Unlocking Educational Innovation: Integrating Bloom's Taxonomy ... - Higher-order thinking refers to cognitive processes operating at the top tiers of Bloom's taxonomy, ...
17. The AI-Powered Classroom: Navigating Bloom's Taxonomy for ... - Master an AI-powered classroom! Learn how Bloom's Taxonomy guides deeper learning & critical thinkin...
18. How to Shift from Passive AI Consumption to Active AI Learning - For many educators, the current relationship with AI is passive. Students use it to generate answers...
19. Responsible Adoption of AI Tools for Teaching and Learning - The adoption of AI tools should be intentional, empowering all students and educators to take full a...
20. Beyond ROI: Are We Using the Wrong Metric in Measuring AI ... - MIT's recent report claiming that 95% of generative AI projects fail to deliver measurable return on...
21. Rapid AI-Powered RoI Evaluation - by Dr Philippa Hardman - A framework & power-prompt to analyse the likely return on investment of training *before* you build...
22. Learning Outcomes Measurement: How OpenAI Tracks AI Impact - OpenAI's Learning Outcomes Measurement Suite helps schools measure how AI affects learning over time...
23. Thinking Fast and Slow in Human and Machine Intelligence - In this article, we propose an approach inspired by human cognition, specifically drawing on cogniti...
24. We Trust AI Over Our Own Brains, Research Finds - AI has such a profound effect on cognition, Wharton researchers say, that it's time to change a half...

25. How AI is Reshaping Human Reasoning and the Rise of Cognitive ... - A key prediction of the theory is“cognitive surrender”—adopting AI outputs with minimal scrutiny, ov...

26. AI and the danger of cognitive surrender - The Economist - Artificial intelligence can lead to "cognitive surrender" where people stop thinking critically and ...

27. Cognitive Surrender: How AI Is Quietly Reshaping the Way We Think - It adds a third layer, System 3, which represents the external algorithmic cognition that AI provide...

28. "Cognitive surrender" leads AI users to abandon logical thinking ... - “Cognitive surrender” leads AI users to abandon logical thinking, research finds ... Experiments sho...

29. System 3 - Navigating Cognitive Surrender in the AI Era - YouTube - This research explores the concept of System 3, a framework describing how artificial intelligence a...

30. Mindless Computing: Designing Technologies to Subtly Influence ... - Persuasive technologies aim to influence user’s behaviors. In order to be effective, many of the per...

31. Mindful Machines: Zen Buddhism and Ethical AI Development - Zen Buddhism, with its emphasis on deep understanding through simplicity and mindfulness, provides v...

32. Beginner's mind in UX design. Shunryu Suzuki was a practitioner of… - This is known in Zen as beginner's mind (shoshin) and it is useful outside of Zen practice. In user ...

33. Thinking Fast and Slow in AI - Thinking fast and slow in AI. This research project aims to build AI-supported machines that can 1) ...

34. Manage Uncertainty with Organizational Learning and AI | BCG - Discover how AI is enhancing organizational learning, boosting financial results, & improving the ab...

35. Expanding A.I's Impact With Organizational Learning - Our research shows that these companies intentionally change processes, broadly and deeply, to facil...

36. How to measure ROI on generative AI investments: A practical guide - According to the IBM Institute for Business Value, organizations treating AI as a measured investmen...

37. AI ROI calculator: From generative to agentic AI success in 2025 - Learn how to calculate AI ROI with our proven 4-part framework. Get real case studies showing 333% R...

38. Proving the ROI of AI Adoption: Metrics and Dashboards ... - Worklytics - Build the ROI case for AI adoption using proven metrics and dashboards. Covers what to measure, how ...

39. The AI ROI Measurement Framework: From Vibe-Based Spending to ... - Comprehensive AI ROI measurement requires a framework that combines quantitative metrics, qualitativ...

40. Marketing Measurement Framework: A Complete Guide for 2025 - Explore a 2025-ready marketing measurement framework that blends KPIs, attribution models, and first...

41. From AI Adoption to Organizational Capability: Why Mindset Is the ... - Rather than focusing on tools, AI Talent Shift aligns human, technical, and cultural dimensions to a...

42. BUILDING PRACTICES THAT REINFORCE TRAINING - Metanoia - Fundamentally, we bring ourselves to the work every day open to change and to also be changed. Our T...

43. Relationships between cognitive biases, decision-making, and ... - ... paranoia-relevant decision-making behavior in schizophrenia. Finally ... Severity of paranoid an...

44. Relationships between cognitive biases, decision-making ... - Multiple measures of decision-making under uncertainty (e.g. jumping to conclusions (JTC), bias agai...

45. Rapid MVP Iteration in 2024: Agile, Lean, and Beyond - Payline Data - Minimum Viable Product iteration refers to the process of rapidly developing, testing and enhancing ...

46. Leverage the Power of the Minimum Viable Product (MVP) - By employing an MVP, companies can gather valuable user feedback early on, allowing them to iterate,...

47. The Intersection of Design Thinking and AI: Enhancing Innovation - This guide explores how AI and design thinking intersect to drive innovation, the benefits and chall...

48. AI-accelerated MVP: faster, smarter product innovation - SoftDesign - An AI-accelerated MVP is a Minimum Viable Product developed with the support of Artificial Intellige...

49. How to Collaborate with AI: A Comprehensive Review - LinkedIn - The framework defines five types of human-AI collaboration: (1) Humans as optional tools, (2) Consen...

50. Scrap Learning: Wasting Training Dollars on Ineffective Training - “Scrap Learning is a term that describes the gap between delivered training and what is applied back...

51. Managing Shiny Object Syndrome in AI | Heinz Marketing - Shiny Object Syndrome can distract marketing teams with the constant influx of new AI tools, leading...
52. The cognitive mirror: a framework for AI-powered metacognition and ... - As of 2025, chat tools provided by OpenAI and Google offer an "education mode" (OpenAI, 2025; Google...
53. New Data Shows AI Study Tools Turn Passive Reading Into Active ... - Nearly 80 million student interactions with Pearson digital materials show responsible AI use drives...