

Affective Scaffolding in Creative Learning: The Vivience Framework for Empathic Digital Humans as Collaborative Partners

¹Prof. Jagmeet Singh

¹ Vivience AI Affiliate, Creative Media Industries Institute, Georgia State University

jeasy@graphicmonk.com

DOI: <https://doi.org/10.63084/cognexus.v2i2.266>

Abstract

Generative artificial intelligence (GenAI) is increasingly used as a dialogic partner in creative learning, raising a question not captured by accuracy-centred evaluation alone: how might relational and affective qualities of AI interaction influence willingness to expose uncertainty, accept critique, revise work, and take creative risks? This conceptual paper advances the Vivience Framework for Affective Scaffolding, extending the Affective-Emotive Resonance (AER) model introduced in prior exploratory research. It integrates Vygotskian scaffolding, Self-Determination Theory, Control-Value Theory, Cognitive Load Theory, social-response research, and polyphonic approaches to creative learning. Two propositions are formalised. Affective Load is defined for prospective testing as task-irrelevant cognitive demand associated with regulating achievement-related affect and AI-mediated interactional uncertainty, while acknowledging prior uses of the term. The Neutrality Trap Hypothesis proposes that affectively sparse responses may, in some high-vulnerability contexts, be interpreted as indifference or implicit evaluation. The paper differentiates four AER dimensions, Persistence, Belonging, Continuity, and Layered Emotional Care, and presents a system-level architecture separating the base language model from memory, inference, pedagogical-state estimation, relational policy, generation, multimodal expression, and governance. Finally, the Vivience Protocol is proposed as a preregistrable mixed-methods design for empirical testing. The contribution is a falsifiable research programme for responsible relational AI in creative learning, not a claim that machine empathy or educational benefit has already been demonstrated.

Keywords: human-AI interaction; creative learning; Affective-Emotive Resonance; relational AI; affective scaffolding

Introduction

Generative artificial intelligence (GenAI) is reshaping creative practice by moving computational systems from passive tools toward participation in ideation, critique, synthesis, and revision (Mollick, 2024; Sharples, 2023). In creative media, design, writing, and related domains, learners may expose provisional ideas to AI before those ideas are mature. Such interaction can become socially legible because users apply social expectations to conversational technologies, especially those using human-like language, voice, or embodiment (Gambino et al., 2020; Nass & Moon, 2000).

This paper asks whether specific, transparent interaction designs can provide affective scaffolding that helps learners remain engaged during uncertainty, critique, failure, and revision, without implying biological or phenomenological empathy. The question is especially relevant to creative learning, where learners must tolerate ambiguity, expose incomplete work to evaluation, and make consequential revisions (Amabile & Pratt, 2016; Beghetto, 2016; Beghetto & Kaufman, 2014). These processes may be cognitively demanding and emotionally charged, particularly for novices.

The paper extends the Affective-Emotive Resonance (AER) model developed in earlier exploratory work within the Vivience research programme (Singh, 2025). That researcher-only comparison used one foundational language model across two scaffolded digital-human personas and a neutral baseline. A standardised 12-prompt evaluation probed the four AER dimensions, with independent transcript coding by two researchers (Cohen's kappa = .85). Scaffolded personas showed substantially stronger coded AER patterns than the neutral baseline. These findings support a narrow premise: persona design can alter observable relational output patterns within a shared model architecture. They do not establish learner-perceived resonance, Affective Load, the Neutrality Trap, or educational benefit. No new participant data are reported here; those claims remain prospective.

Two domain-specific propositions are advanced. The term affective load has prior use in information-behaviour research (Nahl, 2004), and emotion-cognitive-load relations are established (Plass & Kalyuga, 2019). Vivience claims neither priority over the term nor novelty for that general relationship. Here, Affective Load is narrowed to task-irrelevant demand arising from achievement-related affect and uncertainty about the interactional meaning of AI feedback. The Neutrality Trap Hypothesis proposes that, under identifiable conditions, affectively sparse feedback may be interpreted as indifference, social distance, or implicit judgement. Both require discriminant, predictive, and cross-cultural validation.

The paper also avoids describing contemporary AI as uniformly stateless or transactional. Deployed conversational systems can add retrieval, persistent memory, reflection, planning, tools, user profiles, and

other state-management mechanisms around a base language model (Park et al., 2023; Zhong et al., 2024). Transactional configuration is therefore used here only as an operational comparator: a task-focused system that does not intentionally retrieve longitudinal affective context or adapt a relational scaffolding policy across sessions.

Statement of Contribution

1. Conceptual contribution. The paper offers a domain-specific operationalisation of Affective Load and a conditional Neutrality Trap Hypothesis, distinguishing both from adjacent constructs in information behaviour, achievement-emotion research, Cognitive Load Theory, social presence, and computer-mediated communication.
2. Framework contribution. It translates AER into four pedagogical dimensions with distinct operational indicators and separates observable interactional effects from claims about machine consciousness or felt emotion.
3. Systems contribution. It presents a technically qualified, system-level architecture for relational scaffolding that distinguishes the base language model from memory, inference, policy, generation, multimodal expression, and governance components.
4. Methodological contribution. It specifies the Vivience Protocol as a prospective reference design for preregistered empirical testing, including participant allocation, control conditions, measurement intervals, instrument validation, behavioural indicators, and longitudinal analysis.

Literature Review: Affect, Social Response, and Creative Learning

Relational Agents, Educational Chatbots, and Longitudinal Interaction

Relational agents precede contemporary GenAI. Bickmore and Picard (2005) demonstrated that computational agents can use social dialogue and longitudinal interaction to establish and maintain relationships with users. These systems were constrained by the technologies of their time, but they established an important principle: interaction design can influence adherence, alliance, and willingness to continue engaging with a system.

Educational chatbots support several self-regulated learning processes, but evidence is uneven across learning phases. Guan et al. (2025) found stronger support for resource identification, learning strategies, and metacognitive monitoring than for goal setting, planning, reflection, and future adaptation. A systematic review and meta-analysis of emotional AI in education likewise found effects that varied by support type and outcome, while evidence isolating emotional-only support remained limited (Zhang et

al., 2025). Direct experimental evidence is also heterogeneous: Wang and Pang (2026) found no overall learning advantage for an emotional AI condition in L2 vocabulary learning and reported subgroup differences on motivational outcomes. Together, these findings motivate mechanism-specific, longitudinal testing rather than a universal-benefit claim.

Long-term memory cannot be reduced to a model's context window. Research systems demonstrate external memory and reflection architectures that store, retrieve, summarise, and selectively use earlier interactions (Park et al., 2023; Zhong et al., 2024). The relevant Vivience question is therefore not whether memory is technically possible, but what should be retained, how confidently it should be interpreted, when it should influence pedagogical response, and what governance should constrain its use.

Social Responses and Multimodal Pragmatics

The Computers Are Social Actors tradition and subsequent refinements show that users often apply interpersonal scripts, politeness norms, and social heuristics to computers even when they know the system is artificial (Gambino et al., 2020; Nass & Moon, 2000). This does not imply that every user anthropomorphises every system, nor that social responses are identical to human-human relationships. It does establish that language, timing, role framing, embodiment, and other interface cues can alter the pragmatic meaning of an exchange.

Digital-human embodiment adds visual and vocal channels that may carry instructional and affective information, but it is not inherently superior to text. Mayer's (2020) multimedia principles support coordinated verbal and visual channels when instructionally relevant, not the necessity of a human-like avatar. High visual fidelity may help or impair response depending on congruence, animation quality, task, expectation, and preference (Seyama & Nagayama, 2007). Volonte et al. (2018) found that virtual-human conversational and affective animation influenced visual attention, supporting multimodal-cue research but not a general claim that text feedback is harsher.

Digital humans have also been explored as intervention interfaces. Loveys et al. (2021) reported a randomised pilot trial in which a digital human delivered a remote loneliness and stress intervention to at-risk younger and older adults. The study supports the feasibility and acceptability of digital-human-delivered psychosocial interaction, but it does not establish general educational benefit, superior working alliance, or a universal advantage of high-fidelity embodiment.

Achievement Emotions, Cognitive Load, and Creative Risk

Pekrun's Control-Value Theory, originally formulated for achievement emotions and subsequently generalised to a broader account of human emotions (Pekrun, 2006, 2024), explains achievement emotions through appraisals of control and value. Attentional Control Theory describes how anxiety can consume attentional resources (Eysenck et al., 2007). Cognitive Load Theory provides a complementary account of limited working-memory resources (Sweller et al., 2019), and Leppink et al. (2013) developed measures distinguishing cognitive-load components, including extraneous load.

Affective Load is re-specified at the intersection of these literatures. It is not proposed as an additional canonical category within Cognitive Load Theory. In Vivience, it denotes proposed task-irrelevant demand generated when a learner regulates achievement-related affect while interpreting socially meaningful or ambiguous AI feedback during a creative task. The claim is not that emotion itself consumes resources, but that AI-mediated relational uncertainty may be a design-contingent source of additional demand. Scientific utility therefore depends on incremental validity beyond anxiety, negative affect, Nahl's information-search construct, and established cognitive-load measures.

Creative learning makes this mechanism especially testable because creativity requires departures from safe or conventional responses. Creative confidence, classroom context, and willingness to take intellectual risks are related but distinguishable constructs (Beghetto & Kaufman, 2014; Beghetto et al., 2021). If a learner anticipates evaluation or misreads the social meaning of feedback, the learner may choose a safer response even when the AI's informational content is accurate. The Vivience Framework treats that possibility as an empirical question.

Scaffolding, Relatedness, and Polyphonic Creativity

Vygotsky's (1978) account of learning in the Zone of Proximal Development provides a conceptual foundation for later scaffolding approaches that treat support as contingent assistance enabling performance beyond what a learner can independently achieve. Later scaffolding research includes motivational and frustration-management functions alongside cognitive support (Belland, 2017). In technology-rich learning environments, emotions are not peripheral to regulation; they interact with cognition, appraisal, and persistence (Lajoie et al., 2020).

Affective scaffolding is established terminology rather than a Vivience-originated label. Saarinen (2020) describes affective scaffolding as structuring environments to influence affective life, while Candiottto and Stapleton (2024) extend the concept to assistive technology and agency.

Vivience therefore claims not invention of affective scaffolding, but its operationalisation as a longitudinal, governable AI interaction architecture for creative learning with explicit system dimensions, measurable

mediators, and falsifiable boundary conditions.

Self-Determination Theory identifies autonomy, competence, and relatedness as basic psychological needs relevant to motivation and internalisation (Ryan & Deci, 2000). The present framework uses the established term relatedness. It does not claim that an AI can satisfy human relatedness in the same manner or to the same degree as a human relationship. Instead, it asks whether design cues associated with social presence and collaborative framing can support a learner's perceived sense of being accompanied rather than evaluated in isolation.

Ness's (2017) polyphonic account of creative knowledge processes shifts scaffolding from error correction toward orchestration among multiple voices. The associated Room of Opportunity describes phases in which diverse knowledge and perspectives interact before premature closure (Ness & Søreide, 2014). For creative learning, Vivience therefore treats relational scaffolding as one possible condition for sustaining productive dialogue, not as a universal pedagogical prerequisite.

The Vivience Framework and Construct Boundaries

The framework separates observable interactional properties from claims about machine interiority. It retains the four AER dimensions introduced in Singh (2025) while refining AER as an observable, design-dependent pattern of contextually appropriate relational attunement across sequential interactions. It concerns what the system does and how that pattern is perceived, not whether the system experiences emotion, care, consciousness, or sentience. Affective Load and the Neutrality Trap Hypothesis are likewise proposed constructs whose value depends on empirical differentiation from established concepts.

Proposed pathway to be tested. Under conditions of creative vulnerability and socially legible AI interaction, affectively sparse or ambiguous feedback may trigger a Neutrality Trap interpretation for some learners; that interpretation may increase Affective Load, which may in turn alter creative risk-taking, persistence, and revision. AER-enabled scaffolding is hypothesised to moderate this pathway, not to guarantee better outcomes.

Table 1: Construct Boundaries and Validation Requirements

Construct	Definition in this paper	Relation to established constructs	Required empirical test
Affective-Emotive Resonance (AER)	A design-dependent pattern of contextually appropriate relational attunement sustained across sequential	Extends relational-agent and social-response research with a four-dimensional longitudinal structure first	Reliability; convergent and discriminant validity; predictive validity; cross-context

	interactions.	operationalised in Singh (2025).	replication.
Affective Load (Vivience operationalisation)	Task-irrelevant cognitive demand associated with regulating achievement-related affect and AI-mediated interactional uncertainty while performing a creative-learning task.	The term predates Vivience and emotion-load relations are established. Vivience narrows the construct to AI-mediated affect regulation and interactional uncertainty.	Item development; cognitive interviewing; factor validation; incremental validity beyond anxiety and established load measures; behavioural prediction.

Construct	Definition in this paper	Relation to established constructs	Required empirical test
Neutrality Trap Hypothesis	The proposition that affectively sparse AI responses may, in high-vulnerability contexts, be interpreted as indifference or implicit evaluation, increasing Affective Load for some learners.	Related to social-response, expectancy, social-presence, and computer-mediated-pragmatics research; proposed as a conditional interpretation, not a property of all neutral systems.	Hold informational content constant while varying affective signalling; test moderation by vulnerability, preference, culture, and prior AI experience.
Polyphonic Scaffolding	Dialogic support that preserves multiple perspectives long enough for generative exploration before convergence.	Adapted from Ness (2017) and Ness and Søreide (2014); not claimed as a Vivience-original construct.	Dialogue-process coding; learner-perceived opportunity space; association with revision and creative risk-taking.

Note. Proposed construct boundaries and tests; not reported outcomes.

AER and Biological Empathy

AER does not attribute a subjective emotional life to an AI system. It is therefore preferable to describe an AER-enabled agent as performing or signalling context-sensitive care rather than feeling care. The framework is compatible with transparent disclosure that the agent is artificial and that affective responses are computationally generated. This distinction is both epistemological and ethical: users should not need to believe that a machine possesses consciousness for researchers to study whether relationally structured interaction changes human perception or behaviour.

Affective Load as a Proposed Mechanism

A learner facing a difficult creative task may allocate resources both to the task and to regulating anxiety,

shame, perceived judgement, social uncertainty, or ambiguity in feedback. When the latter demand is task-irrelevant, it may compete with planning and revision. Affective Load names this proposed demand. It cannot be inferred from negative emotion alone: anxiety may occur without cognitive interference, and extraneous load may arise without affect. Validation should therefore combine self-report, task behaviour, and appropriate process data.

The Neutrality Trap as a Conditional Hypothesis

The Neutrality Trap is explicitly conditional. Task-focused or affectively sparse responses may be useful or preferred. The hypothesis becomes relevant when a learner is vulnerable or uncertain, the interface is socially legible, and the response provides too little pragmatic information to distinguish correction from judgement or neutral task focus. Effects may vary by culture, personality, neurotype, expertise, prior AI experience, and preference.

Propositional Model and Boundary Conditions

The framework is expressed as a hypothesised conditional causal sequence rather than as a universal benefit claim. Five propositions organise the empirical programme:

P1. When a creative task is personally consequential or evaluatively uncertain and the interface is socially legible, affectively sparse feedback may increase interactional uncertainty for some learners.

P2. Interactional uncertainty may contribute unique task-irrelevant demand beyond negative affect and general extraneous cognitive load. Evidence for Affective Load therefore requires incremental validity, not merely correlation with anxiety.

P3. Accurate, preferred, and transparent AER features may reduce this demand; inaccurate memory, excessive warmth, intrusive inference, or unwanted relational framing may instead increase it.

P4. Lower interaction-related demand is hypothesised to support greater willingness to revise, explore alternatives, and take intellectual or creative risks, conditional on adequate task knowledge and competence.

P5. Effects are expected to be moderated by learner preference, culture, neurotype, expertise, prior AI experience, and task vulnerability. The framework therefore predicts heterogeneous and potentially reversed effects rather than a uniform relational advantage.

Alternative Explanations and Failure Criteria

The framework should be rejected, narrowed, or re-specified if its constructs add no explanatory or predictive value beyond established alternatives. Affective Load requires incremental validity beyond

anxiety, negative affect, social uncertainty, and established load measures. The Neutrality Trap is weakened if affective sparsity does not interact with vulnerability or social legibility, or is fully explained by established social mechanisms. AER dimensions should be consolidated if they cannot be measured reliably or distinguished. The relational package would lack pedagogical value without preregistered learning-relevant benefit over a competent comparator, and any apparent benefit accompanied by reduced autonomy, miscalibrated trust, intrusive memory, dependency, or material harm would count as failure.

Four Pedagogical Dimensions of AER

The four AER dimensions retain the architecture introduced in Singh (2025) but are sharpened for educational testing. Their ingredients are not individually claimed as novel: Persistence overlaps memory, Belonging draws on relatedness and social presence, Continuity concerns developmental narrative integration, and Layered Emotional Care draws on affect-adaptive feedback. The contribution lies in their coordinated longitudinal operationalisation. They are design and interaction dimensions rather than assumed reflective psychological traits; system indicators and learner perceptions therefore provide related but non-interchangeable evidence.

Table 2: Pedagogical Dimensions of Affective-Emotive Resonance

Dimension	Pedagogical definition	Illustrative system behaviour	Independent indicators
Persistence	Accurate storage and retrieval of affect-relevant, task-relevant, and preference-relevant information across authorised sessions.	Recalls that the learner previously reported uncertainty about a specific narrative transition and checks whether that context remains relevant.	Memory retrieval precision/recall; false-memory rate; retention; appropriate non-retrieval after absence or revocation.
Belonging	Pragmatic construction of a collaborative rather than evaluative interactional frame, without implying exclusivity or human equivalence.	Uses language that recognises joint problem solving while preserving the learner's authorship and autonomy.	Relatedness/social presence; stance coding; learner-rated collaboration; autonomy-preservation checks.
Continuity	Appropriate integration of retrieved history into a coherent narrative of the learner's evolving goals, strategies, and creative development.	Connects a current revision to an earlier design decision and to a stated future intention, rather than merely reciting a stored fact.	Narrative-link quality; temporal coherence; relevance of past-present-future linkage; expert-coded developmental integration.
Layered	Adaptive modulation of	Uses exploratory, low-	Feedback-modulation

Emotional Care	feedback framing and multimodal expression in relation to task phase, learner preference, and uncertain affective cues.	judgement framing during divergent ideation, then shifts toward precise critique during convergence when that style is appropriate.	appropriateness; task-phase/preference calibration; perceived warmth-rigour balance; affect-inference error.
----------------	---	---	--

Note. Design/interaction dimensions; indicators are illustrative, not validated scales.

These dimensions are not assumed to be culturally invariant. Their behavioural expression and acceptable intensity require local calibration. For example, explicit collaborative pronouns may support relatedness for one learner and feel intrusive or artificial to another. The framework therefore treats user control and cultural responsiveness as part of construct implementation rather than as post hoc usability concerns.

Interaction Architecture: From Base Model to Relational Orchestration

Distinguishing the Base Language Model from the Deployed System

Relational behaviour cannot be attributed to a single model parameter. Conversational systems may combine a base language model with external memory, retrieval, tools, policy layers, state stores, preferences, and multimodal modules (Park et al., 2023; Zhong et al., 2024). Vivience therefore treats relational scaffolding as a system-level orchestration problem. Model temperature may influence generation diversity, but it is not treated as a direct control for empathy, tone, or pedagogical appropriateness.

Proposed Vivience Orchestration Pipeline

- 1. Signal acquisition:** Receive the learner's utterance and, only where enabled and consented, optional prosodic or visual interaction cues. Raw signals are not treated as ground-truth emotion.
- 2. Context and uncertainty inference:** Estimate task intent, interactional stance, and possible affective cues with confidence values. Where confidence is low, the system should ask or remain conservative rather than infer vulnerability as fact.
- 3. Longitudinal memory retrieval:** Retrieve only authorised, relevant historical context and preferences. Retrieval includes provenance, recency, confidence, and user-controlled deletion or correction.
- 4. Pedagogical-state estimation:** Represent the current creative phase, such as ideation, exploration, convergence, revision, or reflection, together with learner goals and demonstrated competence.
- 5. Relational policy selection:** Select a scaffolding strategy appropriate to task phase, user preference,

prior context, and uncertainty. Strategies may include affect acknowledgement, clarification, challenge, fading of support, or a deliberately concise task-focused response.

6. Response planning and generation: Generate content under explicit instructional and relational constraints. Generation parameters may be adjusted for reliability or diversity, but they are subordinate controls rather than proxies for empathy.

7. Epistemic confidence handling: Separate relational warmth from confidence in factual content. The system should state uncertainty, request evidence, cite sources where available, and avoid using affective fluency to conceal epistemic weakness.

8. Multimodal expression policy: When embodied, coordinate language, prosody, gaze, facial animation, and timing so that expressive cues are congruent with content and user settings.

9. Logging, governance, and user control: Record only the data necessary for the approved research or service purpose. Expose memory controls, affective-feature settings, transparency notices, and audit logs in accordance with applicable governance requirements.

Operational Comparator for Empirical Testing

For the Vivience Protocol, transactional configuration is an experimental label with a narrow definition. Both conditions should use the same base model, avatar, voice, knowledge resources, safety constraints, comparable latency, and equivalent authorised cross-session task history. Only the relational condition would additionally use affective or preference-relevant memory and an adaptive relational-scaffolding policy. This avoids confounding relational scaffolding with ordinary project memory, visual fidelity, or model quality.

Table 3: Structural Comparison of Experimental Interaction Configurations

Design parameter	Task-focused comparator	Vivience relational condition
Base model and content capability	Matched.	Matched.
Embodiment, voice, and basic politeness	Matched.	Matched.
Session context	Current-session context plus authorised cross-session task history.	Matched task history plus authorised relational retrieval.
Affective memory	No cross-session affective or preference-memory retrieval.	Consent-based retrieval of relevant affective and preference context.
Pedagogical-state	Task state plus factual project	Matched task state plus relational

model	trajectory.	trajectory and scaffolding state.
Relational policy	Consistent task-focused feedback framing.	Adaptive policy selecting acknowledgement, challenge, fading, or task-focused brevity.
Primary optimisation target	Clear, accurate task assistance.	Clear, accurate task assistance while testing whether relational scaffolding alters affective load, alliance, and creative risk-taking.

Note. Both conditions share factual task history; only Vivience uses consented relational retrieval.

The Vivience Protocol: A Prospective Reference Design

The Vivience Protocol is a proposed research programme, not a completed study. Because the framework includes new or adapted measures, instrumentation and efficacy should be separated. Phase A would refine manipulation feasibility and candidate items through cognitive interviews and pilot interaction. Phase B would test psychometric structure, discriminant validity, and measurement invariance in an independent or partitioned sample. Phase C would preregister a longitudinal confirmatory trial using frozen measures and analysis rules, reducing same-sample overfitting and supporting replication.

Research Questions and Hypotheses

RQ1: Does a relationally scaffolded empathic digital human (EDH) produce a different trajectory of proposed Affective Load than a matched task-focused EDH during a multi-session creative project?

RQ2: Does relational scaffolding influence task-specific creative risk-taking, revision behaviour, persistence, and perceived working alliance?

RQ3: Are Persistence, Belonging, Continuity, and Layered Emotional Care empirically distinguishable dimensions, and how are their measured expressions associated with observed outcomes? This is an associational dimensional question; the package trial cannot establish the independent causal effect of each component.

H1: Participants in the relational condition will report lower Affective Load over time than participants in the task-focused comparator, after adjustment for baseline measures. This hypothesis remains contingent on successful construct validation.

H2: Participants in the relational condition will show greater validated task-specific creative risk-taking than participants in the comparator condition. Revision behaviour will be examined as a secondary or exploratory outcome unless its measurement is independently validated.

H3: Perceived relatedness and digital working alliance will be higher in the relational condition, while

autonomy should not be lower. A reduction in autonomy would count as an adverse design effect rather than as successful resonance.

H4: Exploratory mediation analyses will test whether changes in Affective Load statistically account for part of any condition effect on creative risk-taking. This analysis should be preregistered and interpreted as mechanistic evidence only if measurement validity is adequate.

Participants, Sample Planning, and Allocation

A first implementation would recruit adults aged 18 years or older in university-level or structured professional training in creative media, design, or a related field, with adequate study- language proficiency. Prior GenAI experience, creative self-efficacy, demographics, and accessibility preferences would be measured before allocation. An illustrative target of 160 participants, about 80 per condition, allows for attrition but is not a claim of statistical sufficiency; final recruitment should follow a preregistered power or precision analysis.

Participants would be randomly assigned 1:1 by a computer-generated sequence. Stratification should be limited to variables with a strong a priori rationale, such as baseline creative self- efficacy and prior GenAI experience, to avoid excessive cells. Allocation should be concealed from outcome raters. Participants cannot be blinded to all interactional differences, but they should not be told which condition is expected to be superior.

To reduce demand characteristics, recruitment and consent materials should describe a comparison of two AI-supported creative-learning configurations without labelling one as empathic, relational, or superior. Baseline expectations and interaction-style preferences should be measured, followed by a post-study question about participants' understanding of the study purpose.

Creative-Learning Task and Intervention Period

The reference task would be a four-week short-form narrative-design project progressing through ideation, selection, prototype, critique, substantial revision, final artefact or pitch, and reflection. Instructions, deadlines, knowledge resources, and technical support would be identical across conditions. A domain-expert panel, blind to condition, would score outputs with a preregistered rubric for novelty, coherence, craft, and departure from conventional solutions.

The relational condition would use the Vivience orchestration pipeline described in Section 5. The task-focused comparator would receive the same authorised cross-session task history but would not retrieve affective or preference-relevant memory or apply an adaptive relational- scaffolding policy. The question

is whether that relational layer adds measurable value or harm beyond competent assistance and matched factual project memory.

This two-condition design estimates the causal effect of the Vivience relational package as a whole, not the independent causal contribution of Persistence, Belonging, Continuity, or Layered Emotional Care. Component-specific causality would require subsequent dismantling, factorial, or micro-randomised experiments in which dimensions are independently manipulated. This distinction prevents associational dimensional results from being overinterpreted as causal decomposition.

Intervention fidelity should be audited independently of outcomes. Informational coverage, task- history availability, instructions, response-length bands, turn frequency, latency, safety rules, and external knowledge access should be comparable. Logs should distinguish task-memory from relational-memory retrieval and record scaffolding actions. Prespecified fidelity checks should test for condition leakage, unequal assistance, or unequal memory access.

Assessment Schedule

Table 4: Illustrative Assessment Schedule for the Confirmatory Trial

Time point	Measures and observations
T0: Baseline	Creative self-efficacy; prior GenAI experience; baseline affect; trust propensity; demographic and accessibility variables; frozen Affective Load measure; baseline task-specific risk orientation.
Each session	Brief state affect; frozen Affective Load items; perceived support intrusiveness; task progress; interaction logs; revision and restart events; user-initiated memory corrections or opt-outs.
Weekly	Working alliance/relatedness measures; perceived social presence; autonomy and competence; manipulation checks for Persistence, Continuity, and feedback adaptation.
T4: Post-intervention	Primary task-specific creative risk measure; expert-rated output; secondary learning or transfer outcome; cognitive-load measures; trust calibration; digital working alliance; qualitative interview or stimulated recall; overall acceptability.
Follow-up: approximately 2 weeks	Persistence of perceived effects; willingness to use the system again; brief transfer task where feasible; recall of learning strategies; and delayed concerns about dependence, privacy, or over-support.

Note. Prospective schedule; the frozen Affective Load measure is the Phase B-validated instrument.

Confirmatory endpoint discipline. Phase C should designate one primary outcome and time point before preregistration. If validated, post-intervention creative risk-taking is a defensible primary endpoint, with Affective Load as a mechanistic secondary outcome. A preregistered secondary learning or transfer

outcome should prevent pedagogical value being inferred from risk-taking alone. If validation is inadequate, the trial should be delayed or re-scoped.

Measures and Validation Strategy

Creative risk-taking. Divergent-thinking performance should be separated from willingness to take creative or intellectual risk. Task-specific risk-taking would use items and behavioural indicators informed by intellectual risk-taking research (Beghetto et al., 2021), adapted through cognitive interviews and pilot testing. Torrance-style fluency, flexibility, and originality may remain secondary creativity outcomes but should not be interpreted as direct measures of willingness to take risks.

Behavioural revision indicators. Semantic revision distance, structural revision magnitude, restart frequency, and exploration of alternatives may be logged as exploratory indicators. Dialogue coding may also examine whether multiple plausible alternatives are sustained before convergence, providing a process link to polyphonic scaffolding. These are not inherently positive or validated measures: large revision may reflect productive risk, confusion, or poor planning. Convergent validity should therefore be examined against blinded expert ratings and participant explanations.

Affective Load. Validation should precede confirmatory efficacy testing. Candidate items would derive from the construct definition, undergo content review and cognitive interviewing, then exploratory factor analysis and item reduction before confirmatory testing in an independent or reserved sample. Phase B should test whether affect-regulation and interactional-interpretation demand form one construct or separable subcomponents. Discriminant and incremental validity should be tested against anxiety, negative affect, Nahl's information-behaviour conception, and established cognitive-load measures. Phase C should use a frozen measurement model.

Working alliance and relatedness. The Working Alliance Inventory was developed for human therapeutic relationships and should not simply be relabelled for educational AI. More appropriate digital-intervention precedents include the Mobile Agnew Relationship Measure (Berry et al., 2018) and the Working Alliance Inventory for Internet Interventions used with a conversational agent (Schlöpfer et al., 2026). Any educational adaptation requires cognitive interviewing, content review, reliability testing, and factor assessment, and should remain distinct from Self-Determination Theory relatedness and social presence.

Trust calibration. Trust would be measured using established automation-trust constructs (Jian et al., 2000) and behavioural decisions in trials where the system's confidence or correctness is experimentally controlled. The objective is calibrated reliance rather than maximal trust. A relational system that makes

users more willing to accept incorrect advice would represent a safety failure even if subjective trust scores increase (Jacovi et al., 2021).

Manipulation checks. Participants should rate whether the system accurately remembered authorised context, connected past and present work, adapted feedback appropriately, and maintained a collaborative rather than evaluative stance. System logs should independently confirm whether relevant memory was retrieved and which relational policy was selected. These checks are essential because an unsuccessful manipulation cannot support conclusions about the framework itself.

Analytical Approach

For H1, a candidate mixed-effects model would estimate repeated Affective Load from condition, time, and their interaction, with prespecified baseline covariates, a participant random intercept, and a random time slope if supported by convergence and preregistered criteria.

Additional clustering should be modelled only if present. Confirmatory reports should provide effect estimates, 95% confidence intervals, p-values where used, sample sizes, and prespecified multiplicity and missing-data procedures. Software, package versions, exclusions, transformations, and the smallest effect size of interest should be fixed in preregistration. The package-level condition contrast is causal under successful randomisation and fidelity; component analyses remain associational unless separately manipulated.

Confirmatory analysis should follow intention-to-treat wherever feasible, with per-protocol or high-fidelity analyses reported only as sensitivity evidence. Preregistration should define the smallest effect size of interest, primary contrast, clustering, attrition handling, and protocol- deviation analyses. Mediation requires temporal ordering of the proposed mediator before the behavioural outcome; cross-sectional post-test mediation would not establish mechanism.

Qualitative interviews or stimulated recall would use a transparent thematic approach to examine how learners interpret critique, memory, support, intrusiveness, autonomy, and continuity. Such evidence can reveal why the same relational behaviour is supportive for one participant and patronising for another. Quantitative and qualitative findings should be integrated into the interpretation rather than using qualitative comments only to illustrate preferred results.

Ethics, Transparency, and Research Governance

Any human-participant implementation would require institutional ethics review before recruitment. Participants should be informed that the EDH is an AI system, that affective responses are

computationally generated, what memory and data are used, and how stored information can be corrected or deleted. Affective inference should remain probabilistic and never be represented as hidden access to emotional truth. Sensitive memory should be minimised, purpose-limited, and protected under the approved protocol and applicable data- protection requirements.

Pedagogical Implications

If supported empirically, AI-supported learning may depend partly on interaction design rather than content accuracy alone. Effective feedback is timely, comprehensible, actionable, and related to learners' goals (Hattie & Timperley, 2007). Relational scaffolding could add a design dimension for how feedback is framed, when challenge intensifies, and when support fades.

The framework does not position an EDH as a replacement for teachers, mentors, peers, or communities of practice. Human educators retain responsibility for curriculum, standards, ethical judgement, pastoral boundaries, and the social life of learning. EDHs are better framed as bounded, transparent practice support between human interactions.

Fading is central to this model. Scaffolding that does not withdraw can foster dependence or reduce self-regulation. As competence increases, the system should shift toward concise challenge, metacognitive prompting, and learner-selected assistance. User preference should remain an explicit control rather than being inferred solely from behaviour.

Ethical Design Requirements for Relational AI

Relational transparency. The system must disclose that it is artificial and must not imply consciousness, human emotion, or reciprocal personal need. Relational language should support the learner without creating false claims of sentience or exclusivity.

Memory agency. Learners should be able to inspect, correct, limit, and delete longitudinal memory. The ability to remember can support continuity but can also become intrusive when the user cannot control what is retained.

Epistemic separation. Relational fluency must remain separate from factual confidence. Warmth should never function as a substitute for evidence or as a mechanism for making uncertain claims more persuasive.

Right to low-affect interaction. Some learners will prefer direct, minimally relational support. Affective scaffolding should be adjustable or disabled, and opting out should not reduce access to core educational capability.

Support fading and anti-dependence. The system should encourage metacognitive reflection and progressively transfer regulation back to the learner. Repeated signs of dependency should prompt design safeguards rather than stronger relational intensity.

Cultural and neuroinclusive calibration. Eye contact, warmth, pronoun use, emotional explicitness, and conversational pacing are culturally and individually variable. Embodiment and affective signalling should therefore be configurable and tested with diverse participants.

Data minimisation. Affective memory can contain highly personal information. Retention should be limited to the minimum necessary for the approved educational or research purposes, with clear retention periods and access controls.

These requirements align with beneficence, autonomy, justice, explicability, and accountable AI governance (Floridi et al., 2018). They also protect scientific integrity: relational responses produced through hidden uncertainty, restricted control, or attachment manipulation would violate the framework's normative purpose.

Discussion, Boundary Conditions, and Limitations

The Vivience Framework does not treat transactional AI as a universal description or assume affective neutrality is inherently harmful. Its proposition is conditional: in some creative-learning contexts, relationally structured interaction may alter how learners allocate attention, interpret critique, and take creative risks. This remains to be tested directly.

Construct validity remains a central risk. Affective Load adds value only if it separates empirically from Nahl's information-behaviour construct, anxiety, negative affect, perceived judgement, and extraneous cognitive load. Likewise, the Neutrality Trap Hypothesis may apply only to a subgroup or be explained by social-presence or expectancy mechanisms. Either label should be revised or abandoned if it lacks incremental explanatory value.

A second limitation concerns affect inference. Text, prosody, gaze, facial movement, and interaction timing are ambiguous signals. An EDH that confidently labels a learner as anxious may be less supportive than one that asks a careful clarification question. The architecture therefore treats affect as uncertain evidence rather than as a directly observable internal state.

Third, relational effects are likely to be culturally contingent. Norms around directness, warmth, eye contact, praise, hierarchy, self-disclosure, and collaborative pronouns vary across linguistic and educational contexts. Cross-cultural replication is therefore not an optional extension but a requirement before broad

claims about AER can be made.

Fourth, learner expertise will moderate need. Novices may benefit from more explicit reassurance and orientation, whereas advanced learners may experience the same behaviour as intrusive or patronising. The framework therefore predicts a need for fading and personalisation rather than fixed relational intensity. These boundary conditions are consistent with the heterogeneous educational-AI evidence reviewed earlier: relational or emotional support may be beneficial, negligible, intrusive, or counterproductive depending on learner and context.

Fifth, embodiment is not assumed necessary. AER dimensions can in principle be expressed through text, voice, or embodied systems. Digital humans add multimodal channels and social salience, but modality comparisons are required before claiming any advantage from embodiment.

Sixth, the protocol is a reference design, not a completed experiment. Its sample, measures, task, and analysis should be refined through pilot work, ethics review, preregistration, and disciplinary consultation. Its value lies in making those dependencies explicit and testable.

Seventh, external validity will initially be narrow. A four-week narrative-design task cannot establish that the same mechanisms operate in visual design, virtual production, game development, music, engineering design, or professional creative teams. Replication across tasks, expertise levels, languages, modalities, and institutions is required before Vivience can be described as a general theory of creative human-AI collaboration.

Taken together, the principal constraints map onto three validity domains: construct validity (whether Affective Load, the Neutrality Trap, and AER dimensions are distinguishable), internal validity (whether fidelity, expectancy, unequal assistance, or affect-inference error explains an observed package effect), and external validity (whether effects generalise across learners, cultures, tasks, and modalities). Phase C findings should therefore be interpreted as context- bounded package effects until these constraints are separately tested.

Conclusion

This paper advances the Vivience Framework as a testable account of affective scaffolding in creative human-AI interaction. It separates earlier exploratory AER observations, current theoretical propositions, and future validation. Affective Load is defined as proposed task- irrelevant demand associated with achievement-related affect and AI-mediated interactional uncertainty; the Neutrality Trap is a conditional interpretation hypothesis. Persistence, Belonging, Continuity, and Layered

Emotional Care are specified as distinguishable design dimensions.

The framework neither claims that AI systems feel emotions nor assumes relational behaviour is always beneficial. Its value depends on measurement validity, differentiation from established theory, and prediction of meaningful educational behaviour. The Vivience Protocol provides a route for testing these questions through matched conditions, longitudinal observation, blinded assessment, qualitative interpretation, and governance.

The research agenda should proceed sequentially: Phase A and Phase B establish manipulation feasibility and measurement validity; Phase C tests the package-level effect in a preregistered creative-learning trial; cross-context replication then examines culture, expertise, task, and modality; and only after a reliable package effect should factorial or dismantling studies attribute effects to individual AER dimensions. The central question remains whether transparent, controllable relational design changes learning in beneficial, negligible, or harmful ways, for whom, and under what conditions.

Acknowledgement

The author acknowledges the Creative Media Industries Institute at Georgia State University and collaborators associated with Vivience AI for the interdisciplinary dialogue and practice-led research environment that informed the development of this conceptual framework.

AI-assisted manuscript preparation. OpenAI ChatGPT supported literature discovery, structural refinement, language editing, and consistency checks. Retained bibliographic records were checked against publisher or DOI metadata. The author retains final editorial control and accepts full responsibility for the manuscript.

Competing interests. The author leads Vivience AI, which develops research and applied systems related to this framework and may have future commercial interests in its application. This is disclosed as a potential financial and non-financial interest.

Funding. No external funding was received for preparation of this manuscript.

References

- Amabile, T. M., & Pratt, M. G. (2016). The dynamic componential model of creativity and innovation in organizations: Making progress, making meaning. *Research in Organizational Behavior*, 36, 157-183. <https://doi.org/10.1016/j.riob.2016.10.001>
- Beghetto, R. A. (2016). Creative learning: A fresh look. *Journal of Cognitive Education and Psychology*, 15(1), 6-23. <https://doi.org/10.1891/1945-8959.15.1.6>

- Beghetto, R. A., & Kaufman, J. C. (2014). Classroom contexts for creativity. *High Ability Studies*, 25(1), 53-69. <https://doi.org/10.1080/13598139.2014.905247>
- Beghetto, R. A., Karwowski, M., & Reiter-Palmon, R. (2021). Intellectual risk taking: A moderating link between creative confidence and creative behavior? *Psychology of Aesthetics, Creativity, and the Arts*, 15(4), 637-644. <https://doi.org/10.1037/aca0000323>
- Belland, B. R. (2017). *Instructional scaffolding in STEM education: Strategies and efficacy evidence*. Springer. <https://doi.org/10.1007/978-3-319-02565-0>
- Berry, K., Salter, A., Morris, R., James, S., & Bucci, S. (2018). Assessing therapeutic alliance in the context of mHealth interventions for mental health problems: Development of the Mobile Agnew Relationship Measure (mARM) questionnaire. *Journal of Medical Internet Research*, 20(4), e90. <https://doi.org/10.2196/jmir.8252>
- Bickmore, T. W., & Picard, R. W. (2005). Establishing and maintaining long-term human- computer relationships. *ACM Transactions on Computer-Human Interaction*, 12(2), 293-327. <https://doi.org/10.1145/1067860.1067867>
- Candiotto, L., & Stapleton, M. (2024). Assistive technology as affective scaffolding. *Topoi*, 43(3), 747-756. <https://doi.org/10.1007/s11245-024-10019-w>
- Eysenck, M. W., Derakshan, N., Santos, R., & Calvo, M. G. (2007). Anxiety and cognitive performance: Attentional control theory. *Emotion*, 7(2), 336-353. <https://doi.org/10.1037/1528-3542.7.2.336>
- Floridi, L., Cows, J., Beltrametti, M., Chatila, R., Chazerand, P., Dignum, V., Luetge, C., Madelin, R., Pagallo, U., Rossi, F., Schafer, B., Valcke, P., & Vayena, E. (2018). AI4People: An ethical framework for a good AI society: Opportunities, risks, principles, and recommendations. *Minds and Machines*, 28(4), 689-707. <https://doi.org/10.1007/s11023-018-9482-5>
- Gambino, A. L., Fox, J., & Ratan, R. A. (2020). Building a stronger CASA: Extending the computers are social actors paradigm. *Human-Machine Communication*, 1, 71-86. <https://doi.org/10.30658/hmc.1.5>
- Guan, R., Raković, M., Chen, G., & Gašević, D. (2025). How educational chatbots support self-regulated learning? A systematic review of the literature. *Education and Information Technologies*, 30, 4493-4518. <https://doi.org/10.1007/s10639-024-12881-y>
- Hattie, J., & Timperley, H. (2007). The power of feedback. *Review of Educational Research*, 77(1), 81-112. <https://doi.org/10.3102/003465430298487>

Jacovi, A., Marasović, A., Miller, T., & Goldberg, Y. (2021). Formalizing trust in artificial intelligence: Prerequisites, causes and goals of human trust in AI. In *Proceedings of the 2021 ACM Conference on Fairness, Accountability, and Transparency* (pp. 624-635). Association for Computing Machinery. <https://doi.org/10.1145/3442188.3445923>

Jian, J.-Y., Bisantz, A. M., & Drury, C. G. (2000). Foundations for an empirically determined scale of trust in automated systems. *International Journal of Cognitive Ergonomics*, 4(1), 53-71. https://doi.org/10.1207/S15327566IJCE0401_04

Lajoie, S. P., Pekrun, R., Azevedo, R., & Leighton, J. P. (2020). Understanding and measuring emotions in technology-rich learning environments. *Learning and Instruction*, 70, 101272. <https://doi.org/10.1016/j.learninstruc.2019.101272>

Leppink, J., Paas, F., van der Vleuten, C. P. M., van Gog, T., & van Merriënboer, J. J. G. (2013). Development of an instrument for measuring different types of cognitive load. *Behavior Research Methods*, 45, 1058-1072. <https://doi.org/10.3758/s13428-013-0334-1>

Loveys, K., Sagar, M., Pickering, I., & Broadbent, E. (2021). A digital human for delivering a remote loneliness and stress intervention to at-risk younger and older adults during the COVID- 19 pandemic: Randomized pilot trial. *JMIR Mental Health*, 8(11), e31586. <https://doi.org/10.2196/31586>

Mayer, R. E. (2020). *Multimedia learning (3rd ed.)*. Cambridge University Press. <https://doi.org/10.1017/9781316941355>

Mollick, E. (2024). *Co-intelligence: Living and working with AI*. Portfolio.

Nahl, D. (2004). Measuring the affective information environment of web searchers. *Proceedings of the American Society for Information Science and Technology*, 41(1), 191-197. <https://doi.org/10.1002/meet.1450410122>

Nass, C., & Moon, Y. (2000). Machines and mindlessness: Social responses to computers. *Journal of Social Issues*, 56(1), 81-103. <https://doi.org/10.1111/0022-4537.00153>

Ness, I. J. (2017). Polyphonic orchestration: Facilitating creative knowledge processes for innovation. *European Journal of Innovation Management*, 20(4), 557-577. <https://doi.org/10.1108/EJIM-05-2016-0049>

Ness, I. J., & Søreide, G. E. (2014). The Room of Opportunity: Understanding phases of creative knowledge processes in innovation. *Journal of Workplace Learning*, 26(8), 545-560.

<https://doi.org/10.1108/JWL-10-2013-0077>

Park, J. S., O'Brien, J., Cai, C. J., Morris, M. R., Liang, P., & Bernstein, M. S. (2023). Generative agents: Interactive simulacra of human behavior. *Proceedings of the 36th Annual ACM Symposium on User Interface Software and Technology*, Article 2, 1-22. <https://doi.org/10.1145/3586183.3606763>

Pekrun, R. (2006). The control-value theory of achievement emotions: Assumptions, corollaries, and implications for educational research and practice. *Educational Psychology Review*, 18(4), 315-341. <https://doi.org/10.1007/s10648-006-9029-9>

Pekrun, R. (2024). Control-value theory: From achievement emotion to a general theory of human emotions. *Educational Psychology Review*, 36, Article 83. <https://doi.org/10.1007/s10648-024-09909-7>

Plass, J. L., & Kalyuga, S. (2019). Four ways of considering emotion in Cognitive Load Theory. *Educational Psychology Review*, 31(2), 339-359. <https://doi.org/10.1007/s10648-019-09473-5>

Ryan, R. M., & Deci, E. L. (2000). Self-determination theory and the facilitation of intrinsic motivation, social development, and well-being. *American Psychologist*, 55(1), 68-78. <https://doi.org/10.1037/0003-066X.55.1.68>

Saarinen, J. A. (2020). What can the concept of affective scaffolding do for us? *Philosophical Psychology*, 33(6), 820-839. <https://doi.org/10.1080/09515089.2020.1761542>

Schläpfer, S., Barth, J., Heinz, P., Held, U., Kowatsch, T., & Witt, C. M. (2026). Association between working alliance and treatment outcomes in a mobile health intervention with a conversational agent (CanRelax). *Internet Interventions*, 44, 100929. <https://doi.org/10.1016/j.invent.2026.100929>

Seyama, J., & Nagayama, R. S. (2007). The uncanny valley: Effect of realism on the impression of artificial human faces. *Presence: Teleoperators and Virtual Environments*, 16(4), 337-351. <https://doi.org/10.1162/pres.16.4.337>

Sharples, M. (2023). Towards social generative AI for education: Theory, practices and ethics. *Learning: Research and Practice*, 9(2), 159-167. <https://doi.org/10.1080/23735082.2023.2261131>

Singh, J. (2025). Emergence of affective-emotive resonance: A paradigm shift in empathic digital humans. In *2025 7th Computing, Communications and IoT Applications Conference (ComComAp)* (pp. 116-121). IEEE. <https://doi.org/10.1109/ComComAp68359.2025.11353186>

Sweller, J., van Merriënboer, J. J. G., & Paas, F. (2019). Cognitive architecture and instructional design: 20 years later. *Educational Psychology Review*, 31(2), 261-292. <https://doi.org/10.1007/s10648-019-09465-5>

Volonte, M., Robb, A., Duchowski, A. T., & Babu, S. V. (2018). Empirical evaluation of virtual human conversational and affective animations on visual attention in inter-personal simulations. In *2018 IEEE Conference on Virtual Reality and 3D User Interfaces (VR)* (pp. 25-32). IEEE. <https://doi.org/10.1109/VR.2018.8446364>

Vygotsky, L. S. (1978). *Mind in society: The development of higher psychological processes* (M. Cole, V. John-Steiner, S. Scribner, & E. Souberman, Eds.). Harvard University Press.

Wang, Z., & Pang, H. (2026). Not a universal benefit: Examining the differential effects of emotional AI on L2 pre-service teachers' language learning. *British Journal of Educational Technology*. Advance online publication. <https://doi.org/10.1111/bjet.70078>

Zhang, H., Liu, Y., Jiang, M., Chen, J., Wang, M., & Paas, F. (2025). Emotional artificial intelligence in education: A systematic review and meta-analysis. *Educational Psychology Review*, 37, Article 106. <https://doi.org/10.1007/s10648-025-10086-4>

Zhong, W., Guo, L., Gao, Q., Ye, H., & Wang, Y. (2024). MemoryBank: Enhancing large language models with long-term memory. *Proceedings of the AAAI Conference on Artificial Intelligence*, 38(17), 19724-19731. <https://doi.org/10.1609/aaai.v38i17.29946>