



Beyond Technostress AI Fatigue as a Cumulative Adaptation Burden in Education and Work

Wael Badawy

Egyptian Russian University, Egypt

ARTICLE INFO

KEYWORDS

AI fatigue; technostress; generative artificial intelligence; cognitive offloading; automation; education; organizational behavior; human-AI interaction; digital well-being; governance

HOW TO CITE

Badawy, W. (2026). Beyond Technostress Ai Fatigue as a Cumulative Adaptation Burden in Education and Work. Emirati Journal of Applied Psychology, 2(2), 33-44.
<https://doi.org/10.54878/bxxewh81>

ABSTRACT

Artificial intelligence has become a persistent layer of work and learning rather than an occasional tool. Users must repeatedly decide when to delegate, how to prompt, whether to trust, how to verify, and how responsibility should be distributed. This integrative review defines AI fatigue as a cumulative reduction in the cognitive, affective, or motivational capacity for adaptive AI engagement when interaction, evaluation, correction, learning, and role-adjustment demands accumulate faster than recovery and support. The review connects technostress, automation vigilance, cognitive load and offloading, trust calibration, digital fatigue, and the Job Demands-Resources perspective. It introduces an Exposure-Adaptation-Depletion-Response framework, a five-domain taxonomy, and an AI Fatigue Load model that separates use volume from uncertainty, task stakes, correction burden, interruption, role disruption, and imposed adoption velocity. The analysis is updated to incorporate a 2026 validated 15-item AI Fatigue Scale; the proposed load model complements that psychophysiological outcome measure by specifying socio-technical antecedents and exposure conditions. Educational and organizational applications show why AI may simultaneously improve output and worsen human sustainability. The article identifies recurring inferential errors, presents a measurement and longitudinal research agenda, and proposes an intervention hierarchy spanning interface design, workflow, curriculum, organizational policy, and governance. AI fatigue is positioned neither as generic screen tiredness nor as a synonym for burnout, but as an adaptation problem located in the coupled human-AI-institution system.



Double-blind peer review by independent reviewers.

Received: 04 Jul 2026, Accepted: 13 Nov 2026, Published: 07 Dec 2026

*Corresponding author: wael@waelbadawy.com

Copyright: © 2026 by the author. Licensee Emirates Scholar Center for Research & Studies, UAE.

Open access under the Creative Commons Attribution (CC BY) license

<https://creativecommons.org/licenses/by/4.0>

1. Introduction

Artificial intelligence has changed category. Earlier digital systems were commonly episodic: a search was made, a spreadsheet opened, or a rule-based system consulted. Contemporary AI is recurrent and increasingly ambient. Generative assistants now appear inside office suites, learning platforms, search tools, development environments, recruitment processes, clinical documentation, and managerial dashboards. The central human task is therefore not simply to learn a tool. It is to make repeated judgments about delegation, prompting, disclosure, verification, correction, and accountability.

The burden created by this transition is described through a fragmented vocabulary. Information-systems research examines technostress, including overload, invasion, complexity, insecurity, and uncertainty [1–3]. Automation scholarship documents complacency, supervisory vigilance, misuse, and disuse [4–6]. Cognitive science addresses load and offloading [7,8], while organizational psychology explains how demands consume resources and how autonomy, support, and competence buffer strain [9]. Education research adds anxiety, dependency, policy churn, and concerns about critical thinking [14–20,23]. These literatures identify parts of the problem but do not by themselves explain why persistent interaction with probabilistic, conversational, and rapidly changing systems may create a distinctive cumulative burden.

AI fatigue has now moved from informal discourse into empirical research. Lau and colleagues developed and validated a 15-item scale with cognitive, emotional, physical, and behavioral dimensions across four studies [24]. That work supplies important evidence that AI fatigue can be measured as a higher-order psychophysiological state distinct from general, digital, and clinical fatigue and from AI-specific technostress. The remaining theoretical need is to explain how exposure becomes fatiguing, why equal use time can produce unequal burden, and how institutions structure the relevant demands and resources.

This review addresses that need. It defines AI fatigue as a cumulative state of reduced capacity or motivation for adaptive AI engagement, arising when recurring interaction, evaluation, correction, learning, and role-adjustment demands exceed available recovery and support. High use is neither necessary nor sufficient. A short, high-stakes verification task can create more load than many routine completions, while extensive use may be sustainable when systems are predictable, autonomy is high, and recovery is protected. Productivity, adoption,

and well-being must therefore be treated as separate outcomes.

2. Review scope and synthesis approach

The review asks four questions: which established theories explain components of AI fatigue; what distinguishes the construct from neighboring forms of strain; which mechanisms recur across education and work; and what measurement and study designs can support cumulative evidence? The unit of analysis is multi-level. Individual cognition matters, but exposure is organized through interfaces, policies, assessment rules, managerial incentives, procurement decisions, and labor-market expectations.

The source survey synthesized work from human-computer interaction, information systems, organizational psychology, education, cognitive science, and automation. Its search logic combined AI terms with fatigue, stress, overload, burnout, anxiety, dependency, offloading, vigilance, resistance, and well-being. Conceptually relevant empirical studies, reviews, and foundational theories were included; purely technical performance studies without a human adaptation outcome were excluded. Evidence was compared by construct definition, exposure operationalization, temporal design, outcome measure, setting, and level of analysis.

This is an integrative rather than a meta-analytic review. Heterogeneity in populations, technologies, measures, and outcomes prevents a meaningful pooled effect. The synthesis therefore uses mechanism mapping: findings are organized around exposure, adaptation work, depletion, response, and contextual resources. Recent peer-reviewed evidence published through August 2026 was checked to correct time-sensitive claims, most importantly the emergence of a validated AI Fatigue Scale [24].

3. Conceptual ancestry and construct boundaries

3.1 Technostress and continual adaptation

Technostress is the closest parent construct. Classical dimensions include overload, invasion, complexity, insecurity, and uncertainty [1–3]. AI intensifies each but also changes its structure. Uncertainty no longer concerns only upgrades or unfamiliar interfaces: generative systems can produce different responses to similar prompts, conceal provenance, and present fluent claims without reliable epistemic cues. Complexity shifts from command syntax toward interaction strategy. Users

must determine what context to supply, how to evaluate a completion, and when correction is more costly than independent work.

Job Demands–Resources theory helps explain why the same AI can support and exhaust simultaneously [9]. AI efficacy, usable interfaces, autonomy, and support can increase engagement, while technostress consumes energy. A three-wave study of 600 employees found resource pathways associated with productivity and positive work outcomes alongside demand pathways associated with exhaustion and work–family conflict [13]. Efficiency is thus not evidence of sustainability.

3.2 Automation vigilance and the offloading paradox

Automation often removes routine execution while leaving people responsible for abnormal or ambiguous cases. Bainbridge called this an irony of automation [4]; later research distinguished appropriate use from misuse, disuse, and abuse [5,6]. Generative AI makes supervisory vigilance continuous because errors can be subtle, contextual, and well written. The user must sustain skepticism toward outputs that appear easier to accept than to inspect.

Cognitive offloading can reduce immediate processing demands [7], and AI can summarize, retrieve, draft, or externalize intermediate reasoning. Yet delegated cognition requires supervision. A student who offloads brainstorming may still need to verify references and defend an argument; a manager who requests alternatives

may receive more plausible options than can be responsibly compared. The first-order task becomes easier while second-order evaluation, integration, and accountability expand. This offloading paradox is central to AI fatigue.

3.3 Trust, resistance, digital fatigue, and burnout

Algorithm aversion demonstrates that observed errors can produce disuse even when an algorithm performs well on average [10]. However, reduced adoption is ambiguous. It may indicate fatigue, distrust, ethical objection, rational recognition of verification costs, or healthy boundary setting. Continued use is equally ambiguous because it may reflect competence, dependency, institutional compulsion, or insecurity. Acceptance measures [11,12] cannot substitute for adaptation quality.

Digital fatigue concerns prolonged digital exposure and can occur without AI. Burnout is a broader and typically more severe response to chronic occupational or role stress. AI fatigue can contribute to both but is defined by an AI-specific demand architecture: prompting, evaluating probabilistic outputs, correcting errors, tracking changing capabilities, negotiating delegation, and adjusting professional or learner identity. AI anxiety may precede substantial exposure, whereas AI fatigue requires repeated demand or sustained exposure.

Table 1. Boundary between AI fatigue and neighboring constructs

Construct	Primary trigger	Relationship to AI fatigue
Technostress	Technology-related demands	Broad parent construct; AI fatigue specifies cumulative AI adaptation and evaluation
Digital fatigue	Prolonged digital exposure	Can occur without AI; screen time is not required for every AI-fatigue pathway
Burnout	Chronic role stress	Potential downstream condition; broader and more severe
Cognitive overload	Task demand exceeds processing resources	Proximal mechanism that may be brief
Algorithm aversion	Error, distrust, or control concerns	Possible response rather than the fatigue state
AI anxiety	Anticipated AI-related threat	May precede repeated use
AI fatigue	Recurrent AI-specific demands	Reduced capacity or motivation for adaptive AI engagement

4. Exposure–Adaptation–Depletion–Response framework

The proposed framework locates AI fatigue in a process rather than a single behavior. Exposure creates adaptation demands; adaptation consumes resources;

incomplete recovery produces persistent depletion; and depletion yields heterogeneous responses. Resources and feedback modify every stage. This structure avoids defining fatigue through use frequency or avoidance alone.




Figure 1. Exposure–Adaptation–Depletion–Response framework.

4.1 Exposure is multidimensional

Exposure includes interaction volume, output uncertainty, task stakes, interruption, correction requirements, role disruption, and the velocity of imposed change. A routine low-stakes assistant can be used frequently with little burden. A single medical, legal, assessment, or personnel decision may demand intensive verification. Exposure measurement should therefore distinguish minutes and prompts from consequence, ambiguity, and responsibility.

Change velocity is especially important. Models, interfaces, policies, and acceptable-use norms evolve faster than many institutions can stabilize training. When a new system arrives before competence with the previous one consolidates, adaptation debt accumulates. Users repeatedly invest in learning that may soon become obsolete, while organizations underestimate the time required for experimentation, error correction, and workflow redesign.

4.2 Adaptation work and resource depletion

Adaptation is active work. It includes learning tool boundaries, crafting instructions, interpreting outputs, checking provenance, reconciling contradictions, editing, documenting disclosure, and deciding whether delegation is appropriate. Much of this labor is invisible in

conventional productivity accounting. Time saved during generation may be recovered by verification or absorbed by higher output expectations.

Depletion is dynamic rather than categorical. Attention, metacognitive control, self-efficacy, emotional energy, and tolerance for uncertainty fluctuate with task design and recovery. Early signs may include slower checking, irritability, repetitive prompting, shallow acceptance, or deliberate avoidance. Severe exhaustion is not required. In safety-sensitive settings, verification attenuation—the gradual decline in checking quality—may be a more consequential early outcome than self-reported tiredness.

4.3 Responses and feedback

Behavioral responses range from careful selective use to disengagement, shallow compliance, oscillation between overuse and avoidance, or dependency. None is diagnostic by itself. Avoidance may be protective; sustained use may conceal exhaustion; reliance may reflect rational task allocation or impaired calibration. Responses also alter future exposure. A user who stops checking may reduce immediate effort but increase error risk; an institution that responds to failures with additional controls may further increase documentation and fatigue.

4.4 Recovery, resources, and adaptation debt

Recovery is more than time away from a screen. It includes a period in which users are not required to

monitor model changes, justify AI use, revise prompts, or anticipate new policies. Micro-breaks may restore attention during a task, but they do not eliminate adaptation debt created by recurrent organizational change. Stable guidance, protected learning time, access to knowledgeable support, and permission to use non-AI alternatives can reduce that debt by allowing competence to consolidate.

Resources operate at several levels. Individual AI literacy improves recognition of limitations but may also increase awareness of risks and therefore verification work. Team resources include shared prompt libraries, peer review, agreed escalation practices, and psychological safety to admit uncertainty. Institutional resources include realistic workload, coherent procurement, accessibility, privacy protection, and stable governance. Interface resources include provenance, uncertainty cues, persistent context,

and easy correction. A model that measures only personal coping would misattribute institutionally produced strain to the user.

The balance between demand and resource is also task dependent. High autonomy can protect a creative worker who uses AI selectively but may provide little protection to a compliance analyst who remains legally responsible for every generated claim. Training may help a novice but create additional burden when delivered during peak workload or when tools change immediately afterward. Resources should therefore be evaluated by whether they reduce exposure, improve adaptation efficiency, or restore depleted capacity—not merely by whether an institution offered them.

5. Five domains of AI fatigue




Figure 2. Five-domain taxonomy of AI fatigue.

Cognitive-evaluative fatigue arises from sustained judgment about accuracy, relevance, bias, provenance, and appropriate reliance. It is likely when surface fluency makes weak outputs difficult to detect or when users lack subject expertise. Interaction fatigue concerns the conversational labor of clarifying prompts, repairing misunderstandings, repeating context, and navigating anthropomorphic or overly verbose responses.

Learning-and-change fatigue reflects continuous adaptation to models, features, policies, and prompting conventions. It is amplified when training is tool-specific, adoption is compulsory, or updates arrive before practice stabilizes. Identity-and-role fatigue occurs when AI changes what counts as competence, authorship, expertise, or legitimate work. Professionals may spend

energy reconciling efficiency expectations with accountability and values; students may struggle to determine which cognitive work remains their own.

Organizational-acceleration fatigue results when institutions convert AI capability into faster cycles, more output, or broader availability without preserving recovery. Saved time becomes new demand. Tool proliferation adds orchestration work: selecting among assistants, moving context, learning inconsistent interfaces, and complying with multiple policies. These five domains overlap, but separating them supports diagnosis and targeted intervention.

The domains can also reinforce one another. Repeated interaction failures increase cognitive evaluation and may reduce confidence. Low confidence motivates additional learning, while constant training highlights role

uncertainty. Organizations may respond to uneven performance by mandating more tools or monitoring, increasing acceleration fatigue. The cumulative experience can become self-reinforcing even if no single interaction is unusually difficult.

Observable signs differ by domain. Cognitive–evaluative fatigue may appear as shorter verification, missed contradictions, or delayed decisions. Interaction fatigue may appear as terse prompting, abandonment after repair attempts, or preference for conventional search and human contact. Learning fatigue may produce avoidance of updates or superficial completion of training. Identity fatigue may emerge as defensiveness, reduced professional meaning, or concern that expertise is becoming invisible. Organizational acceleration fatigue may appear in widespread cynicism, shadow use, policy noncompliance, or turnover intention. These indicators are hypotheses for assessment, not diagnostic rules.

Physical manifestations require particular caution. Headache, eye strain, sleep disturbance, and bodily exhaustion can accompany AI use, and the validated AI Fatigue Scale includes a physical dimension [24]. Yet such

symptoms may arise from general screen exposure, workload, ergonomics, or health conditions. AI specificity should be supported by temporal or task evidence—for example, symptoms increasing during high-correction AI work while comparable non-AI screen work remains stable.

6. AI Fatigue Load: measuring exposure conditions

The AI Fatigue Load model complements the validated 15-item outcome scale [24]. The scale measures cognitive, emotional, physical, and behavioral manifestations; the load model specifies antecedent conditions that may produce those manifestations. A simple research formulation can represent load as a weighted function of volume, uncertainty, stakes, interruption, correction burden, learning demand, role disruption, and adoption velocity, reduced by autonomy, competence, support, interface quality, and recovery. The weights should be estimated rather than assumed.

Table 2. Proposed dimensions of AI Fatigue Load

Dimension	Operational indicator	Interpretive role
Interaction volume	Prompts, sessions, or minutes	Necessary but insufficient
Uncertainty	Variability and unverifiable output	Raises vigilance
Task stakes	Cost of an undetected error	Raises accountability
Correction burden	Editing time / generation time	Captures hidden labor
Change velocity	Rate of tool and policy change	Creates adaptation debt
Role disruption	Change in authorship or responsibility	Creates identity work
Autonomy	Meaningful choice over use	Protective resource
Recovery and support	Time, training, escalation	Restores capacity

Three derived indicators deserve testing. The correction-to-generation ratio captures how much review and repair are required for each unit of output. Calibration entropy represents instability in the user's trust state when system performance is difficult to predict: users repeatedly shift between acceptance and skepticism. Adaptation half-life is the time required for a user or team to regain stable performance and confidence after a material tool or policy change. Together these measures move exposure research beyond simple frequency.

7. Educational overexposure

7.1 Delegation and the learning dilemma

AI can scaffold explanation, feedback, language support, and idea generation. It can also allow learners to bypass the effort through which durable knowledge is built. The relevant question is not whether a student used AI, but which cognitive operations were delegated and whether verification required understanding. Immediate effort reduction can support learning when prompts elicit explanation and feedback; it can undermine learning when answers substitute for retrieval, reasoning, or revision.

Educational delegation and verification loop




Figure 3. Educational delegation and verification loop.

7.2 Assessment, policy, and emotional costs

Educational institutions have often responded to generative AI with rapid and inconsistent policy change. Students navigate different rules across courses, while teachers repeatedly redesign assessment and detection practices. This assessment churn creates learning-and-change fatigue independently of chatbot use. Mandatory platforms can also transform resistance into an adaptive response when tools are unreliable, intrusive, or poorly integrated [17].

Empirical studies link technostress, digital fatigue, dependency, anxiety, and burnout in educational settings [14–16,19,20,23]. These studies are largely cross-sectional and should not be treated as proof of a single causal syndrome. They nevertheless show that AI exposure occurs within broader learning conditions. Workload, language anxiety, assessment pressure, self-efficacy, and institutional support shape outcomes. Education studies should measure task design, verification quality, skill retention, and choice rather than relying on use intention alone.

8. Organizational overexposure

Organizations frequently justify AI adoption through time savings, throughput, and service availability. These gains may be genuine, yet efficiency without recovery creates a rebound effect. Faster drafting leads to more drafts; automated analysis leads to more decisions; constant availability raises response expectations. Employees may deliver more while experiencing greater exhaustion [13].

Tool proliferation creates orchestration burden. Workers choose among assistants, transfer context, resolve conflicting outputs, and comply with fragmented security rules. Accountability is often asymmetric: the organization mandates or encourages AI, but the employee remains responsible for errors, disclosure, and quality. Surveillance and job insecurity add affective load, particularly where AI use is monitored or framed as a benchmark of employability [21].

Table 3. Organizational failure modes that convert AI capability into fatigue load

Failure mode	Mechanism	Likely consequence
Productivity rebound	Saved time converted into more output	No recovery despite efficiency
Tool proliferation	Multiple assistants and policies	Selection and context-switching burden
Verification asymmetry	User liable for machine output	Hidden review and documentation work
Mandatory adoption	Limited meaningful choice	Resistance and loss of autonomy
Surveillance / insecurity	Monitoring and replacement concern	Affective and identity strain

Update churn	Frequent model or policy change	Adaptation debt
--------------	---------------------------------	-----------------

9. Contradictions and recurring inferential errors

The evidence contains apparent contradictions that a process model can reconcile. AI can reduce cognitive load by externalizing work and increase load through supervision. It can increase productivity and exhaustion simultaneously. High use can indicate competence or dependency; low use can indicate aversion, ethical objection, or effective boundary setting. Self-efficacy can buffer strain, yet highly capable users may be assigned more AI-mediated work and conceal overextension.

Three inferential errors recur. Exposure reductionism equates frequency with intensity and ignores stakes, uncertainty, and correction. Outcome substitution treats productivity or continued adoption as evidence of well-being. Attitude essentialism interprets resistance as a stable personal deficit instead of a possible response to accumulated demand, poor design, or compulsory adoption. Research should specify the AI-specific demand, demonstrate temporal accumulation, and measure an adaptive consequence.

A fourth error is construct substitution: applying an AI-fatigue label to any negative outcome observed after AI introduction. Emotional exhaustion following restructuring may be driven by job insecurity; student stress may reflect assessment pressure; physical tiredness may reflect prolonged device use. Studies need a plausible AI-specific pathway and should control adjacent demands. The validated scale improves outcome measurement, but causal attribution still requires exposure and temporal evidence.

A fifth error is level confusion. Individual surveys are often used to make claims about organizational adoption, although organizational practices may vary sharply among teams. Conversely, a university-wide policy may be treated as equal exposure even when some disciplines use AI daily and others rarely. Multi-level sampling should measure both personal encounters and shared adoption regimes. Team norms can normalize constant checking, encourage unsafe shortcuts, or provide mutual support.

Finally, researchers should resist a simple benefit-harm score. AI changes several outcomes at once: speed, quality, learning, autonomy, equity, well-being, and accountability. A system can be beneficial on one dimension and costly on another. Sustainable adoption is a constrained optimization problem, not a search for a universally positive attitude. Reporting should show trade-offs and distributional effects rather than collapse them into an average satisfaction measure.

10. Measurement and research agenda

10.1 Integrating outcome and exposure measures

The validated AI Fatigue Scale provides a major foundation [24]. Future work should replicate its factor structure across occupations, educational levels, languages, cultures, disability groups, and AI modalities. The socio-technical load dimensions proposed here should be studied as antecedents rather than folded uncritically into the fatigue score. Separating exposure from outcome reduces criterion contamination and enables tests of mediation and moderation.

Validation sequence for AI-fatigue measurement



Use self-report together with logs, correction behavior, task quality, recovery, and longitudinal exposure.

Figure 4. Recommended sequence for validating AI-fatigue measures and mechanisms.

10.2 Designs capable of identifying accumulation

Cross-sectional surveys cannot establish whether fatigue accumulates, whether fatigue changes use, or whether

both reflect a third condition. Repeated-task experiments can manipulate uncertainty, correction burden, stakes, and interruption while measuring checking accuracy and recovery. Longitudinal deployment studies can collect baseline measures before adoption and follow users

through training, stabilization, and updates. Experience sampling can capture within-person fluctuations, while objective logs can measure prompts, corrections, model switching, and work outside scheduled hours.

Multi-level models should distinguish person, task, team, and institution effects. Latent-growth and change-point models can test nonlinear trajectories: early enthusiasm

may precede depletion, or initial effort may decline as competence develops. Cross-lagged models can examine reciprocal relations between fatigue and use. Intervention studies should use ablation logic—removing one burden at a time, such as unnecessary alerts, compulsory tools, or unprotected verification—to identify active mechanisms.

Table 4. Reproducible designs for AI-fatigue research

Design	Core manipulation / contrast	Primary outcome
Repeated-task experiment	Uncertainty × stakes × correction	Verification quality and recovery
Longitudinal deployment	Before, during, after adoption	Adaptation trajectories
Experience sampling	Daily demand and affect	Within-person accumulation
Team study	Shared tools and norms	Coordination and contagion
Policy natural experiment	Pacing, choice, or tool consolidation	Institutional causal leverage
Mixed-method audit	Logs + interviews + quality review	Hidden work and mechanism explanation

11. Intervention hierarchy

Individual literacy and recovery practices matter, but they cannot compensate for compulsory exposure, poor interfaces, or unrealistic workload. Intervention should begin at the highest level that creates the demand. Governance determines choice, accountability, procurement, pacing, and permissible uses. Organizational and educational policy determines

workload, tool consolidation, support, and recovery. Workflow design protects verification time and establishes escalation paths. Interface design can provide provenance, uncertainty cues, editable context, and correction support. Individual practices then operate within these conditions.



Figure 5. Intervention hierarchy for sustainable AI use.

Interface interventions should reduce unnecessary conversational turns, preserve context transparently, expose sources, and make uncertainty visible without creating explainability overload. Workflow interventions should specify when human review is mandatory, allocate

time for it, and avoid using AI time savings as an automatic basis for increased volume. Educational interventions should align AI permissions with learning objectives, teach verification, and stabilize policies long enough for competence to develop.

Governance interventions should require impact assessment before mandatory deployment, document accountability across vendors and users, and offer meaningful alternatives where feasible. Procurement should evaluate correction burden, data governance, accessibility, and update practices rather than purchase price and benchmark performance alone. Monitoring should use aggregated, minimally intrusive indicators; a fatigue program that expands surveillance may reproduce the burden it seeks to solve.

12. Ethical, equity, and policy implications

Adaptation resources are unequally distributed. Experienced users, secure employees, well-resourced schools, and workers with greater autonomy can experiment and recover. Students, contingent staff, multilingual users, disabled users, and workers subject to monitoring may face greater exposure with fewer buffers. Mandatory AI therefore raises a question of meaningful choice: whether people can decline, appeal, or use an alternative without penalty.

Fatigue measurement also carries risk. Individual scores could be used to screen workers, pathologize justified resistance, or shift responsibility away from institutional design. Measures should support system improvement, not surveillance. Reporting should use minimum group sizes, protect confidentiality, and combine self-report with process evidence. Ethical governance recognizes that reduced engagement can be a rational boundary response and that human sustainability is a legitimate performance criterion.

12.1 Practical relevance for productivity and safety

Conventional business cases count licenses, active users, output volume, and time saved. A human-sustainability account adds verification time, correction-to-generation ratio, error escape rate, after-hours adaptation, update frequency, and recovery. These measures can reveal a productivity illusion in which generation becomes faster but total responsible completion time changes little. They can also identify where AI creates genuine capacity because correction is low and saved time is not immediately converted into additional demand.

Safety analysis should focus on verification attenuation. In high-stakes work, users may initially inspect every output carefully, then reduce checking as volume grows or surface quality becomes familiar. Organizations should sample decision quality over time, test performance when AI is unavailable, and monitor whether escalation occurs appropriately. Random audits, paired human review, and

protected exception-handling time can provide stronger evidence than annual self-report alone.

Workforce development should distinguish operational fluency from conceptual competence. Knowing how to obtain a polished response does not demonstrate understanding of training data, uncertainty, bias, or appropriate delegation [18]. Training should include failure recognition, independent task practice, disclosure, and accountability. It should also preserve expertise by periodically requiring unaided performance where skill retention is safety relevant.

13. Discussion

13.1 Limitations of the evidence base

The evidence supporting AI fatigue remains uneven despite recent measurement progress. Many studies use convenience samples, self-reported exposure, and cross-sectional designs. These methods are useful for identifying associations but cannot determine whether demanding AI use causes fatigue, whether fatigued people perceive tools more negatively, or whether workload and organizational change produce both. Rapid technological change also creates cohort effects: findings tied to one model generation, interface, or policy environment may age quickly.

Common-method variance is a particular concern when the same respondent reports AI use, technostress, exhaustion, and intention to disengage in one survey. Objective use logs can help but are incomplete because prompt counts do not reveal task stakes, cognitive effort, or unrecorded checking outside the platform. Vendor logs may also create privacy and power concerns. The strongest designs will combine carefully consented behavioral data with self-report, task-quality assessment, interviews, and institutional measures.

The review itself is integrative rather than systematic in the strict PRISMA sense and does not estimate pooled effects. Its cross-disciplinary scope creates terminological inconsistency and risks selective interpretation. Several cited 2025–2026 studies are early contributions whose results require independent replication. The proposed taxonomy and load dimensions are theoretical and have not yet been validated as a unified antecedent model. They should be tested against simpler explanations and revised when evidence does not support their distinctions.

Finally, the construct may not generalize uniformly across cultures or settings. Expectations about authority, technology, disclosure, emotional expression, and acceptable workload influence both exposure and reporting. Measurement invariance is therefore a

prerequisite for group comparison, not a technical afterthought. A responsible research program should seek conditions under which AI fatigue is absent as actively as it seeks confirming evidence.

13.2 Integrated interpretation

The central contribution of this synthesis is to reposition AI fatigue as a socio-technical adaptation problem. The validated scale confirms that fatigue can manifest cognitively, emotionally, physically, and behaviorally [24]. The EADR framework explains how those manifestations may arise from exposure conditions and institutional arrangements. The five-domain taxonomy differentiates sources of burden, while AI Fatigue Load provides measurable antecedents that are not reducible to frequency.

This account also clarifies why AI fatigue is not merely technostress with a fashionable label. The construct remains nested within broader technology strain, but probabilistic output, conversational repair, continuous trust calibration, delegation decisions, and shifting responsibility create a recurrent configuration. Its scientific value depends on discriminant evidence and incremental prediction. If AI fatigue fails to explain outcomes beyond general fatigue, digital fatigue, technostress, workload, and anxiety, the narrower construct should be revised or rejected.

The framework has limits. Some AI interactions are brief, deterministic, and low burden. Expert users may develop stable routines that prevent depletion. Highly supportive institutions may convert the same technology into a resource. Physical fatigue captured by the validated scale may sometimes reflect screen exposure rather than AI-specific demands. The model should therefore be treated as a testable architecture, not a universal diagnosis.

Three propositions summarize the research program. First, use frequency should become a weak predictor after uncertainty, stakes, correction burden, and change velocity are measured. Second, verification attenuation should emerge as an early safety-relevant consequence of cognitive-evaluative depletion. Third, organizational adoption velocity should predict fatigue partly through adaptation debt, with autonomy, tool consolidation, governance clarity, and protected recovery acting as buffers.

14. Conclusion

AI can reduce effort, expand access, and improve output. It can also create a hidden layer of prompting, vigilance, correction, learning, and identity work. AI fatigue arises when these demands recur faster than people and

institutions can adapt and recover. The construct should not absorb every discomfort associated with technology, nor should it be inferred from low adoption or tiredness alone.

A cumulative science of AI fatigue now has two complementary foundations: an empirically validated outcome scale and a socio-technical account of exposure, adaptation, resources, and governance. Research must connect them through longitudinal and behavioral evidence. Practice must recognize that sustainable AI adoption is not achieved by asking individuals to become endlessly resilient. It requires systems that pace change, preserve meaningful choice, protect verification and recovery, clarify accountability, and evaluate human sustainability alongside productivity.

Declarations

Review type: Integrative conceptual review; no new human-participant data were collected.

Data availability: No original dataset was generated for this review.

Ethics: Ethics approval and consent were not applicable because the article reports no original participant research.

Funding: Not stated in the supplied source.

Conflict of interest: Not stated in the supplied source.

Author information: Names, affiliations, ORCID identifiers, and corresponding-author details were not present in the anonymous source and must be inserted before publication.

References

- [1] Brod, C. (1984). *Technostress: The human cost of the computer revolution*. Addison-Wesley.
- [2] Tarafdar, M., et al. (2007). The impact of technostress on role stress and productivity. *Journal of Management Information Systems*, 24(1), 301–328. <https://doi.org/10.2753/MIS0742-1222240109>
- [3] Tarafdar, M., Pullins, E. B., & Ragu-Nathan, T. S. (2015). Technostress: Negative effect on performance and possible mitigations. *Information Systems Journal*, 25(2), 103–132. <https://doi.org/10.1111/isj.12042>
- [4] Bainbridge, L. (1983). Ironies of automation. *Automatica*, 19(6), 775–779. [https://doi.org/10.1016/0005-1098\(83\)90046-8](https://doi.org/10.1016/0005-1098(83)90046-8)
- [5] Parasuraman, R., & Riley, V. (1997). Humans and automation: Use, misuse, disuse, abuse. *Human*

Factors, 39(2), 230–253.
<https://doi.org/10.1518/00187209778543886>

[6] Parasuraman, R., & Manzey, D. H. (2010). Complacency and bias in human use of automation. *Human Factors*, 52(3), 381–410.
<https://doi.org/10.1177/0018720810376055>

[7] Risko, E. F., & Gilbert, S. J. (2016). Cognitive offloading. *Trends in Cognitive Sciences*, 20(9), 676–688. <https://doi.org/10.1016/j.tics.2016.07.002>

[8] Sweller, J. (1988). Cognitive load during problem solving. *Cognitive Science*, 12(2), 257–285.
https://doi.org/10.1207/s15516709cog1202_4

[9] Bakker, A. B., & Demerouti, E. (2007). The Job Demands–Resources model: State of the art. *Journal of Managerial Psychology*, 22(3), 309–328.
<https://doi.org/10.1108/02683940710733115>

[10] Dietvorst, B. J., Simmons, J. P., & Massey, C. (2015). Algorithm aversion. *Journal of Experimental Psychology: General*, 144(1), 114–126.
<https://doi.org/10.1037/xge0000033>

[11] Davis, F. D. (1989). Perceived usefulness, perceived ease of use, and user acceptance of information technology. *MIS Quarterly*, 13(3), 319–340.
<https://doi.org/10.2307/249008>

[12] Venkatesh, V., et al. (2003). User acceptance of information technology: Toward a unified view. *MIS Quarterly*, 27(3), 425–478.
<https://doi.org/10.2307/30036540>

[13] Chuang, Y.-T., Chiang, H.-L., & Lin, A.-P. (2025). Insights from the Job Demands–Resources Model: AI’s dual impact on employees’ work and life well-being. *International Journal of Information Management*, 83, 102887.

[14] Wang, H., Khoso, A. K., & Althubiani, A. R. (2026). Technostress, digital fatigue, and AI dependency as antecedents of burnout and SDG-4 achievement in EFL classrooms. *Scientific Reports*, 16, 15412. <https://doi.org/10.1038/s41598-026-45402-7>

[15] Zhang, H., et al. (2025). The impact of AI-induced educational anxiety on teacher well-being. *BMC Public Health*. <https://doi.org/10.1186/s12889-025-25372-7>

[16] Verano-Tacoronte, D., et al. (2025). Are university teachers ready for generative artificial intelligence? *Education and Information Technologies*.
<https://doi.org/10.1007/s10639-025-13585-7>

[17] Mohd Daud, N., et al. (2026). Resisting the machine: Modeling emotional friction and academic decline in enforced AI learning environments. *Smart*

Learning Environments. <https://doi.org/10.1186/s40561-026-00435-3>

[18] van Zyl, L. E. (2026). The competence paradox: When psychologists overestimate their understanding of artificial intelligence. *AI & Society*.
<https://doi.org/10.1007/s00146-025-02814-9>

[19] Zhang, J., et al. (2026). A SEM–ANN analysis to examine the impact of AI overreliance. *BMC Psychology*. <https://doi.org/10.1186/s40359-026-04246-6>

[20] Carter, J. W., et al. (2025). The emotional cost of AI chatbots in education. *Computers and Education: Artificial Intelligence*.

[21] Zheng, J., et al. (2025). Serial mediation of job insecurity and work-related mechanisms in the association between AI awareness and emotional exhaustion. *Behavioral Sciences*, 15(4), 401.
<https://doi.org/10.3390/bs15040401>

[22] Korosec-Serfaty, M., et al. (2026). Critical thinking in AI-assisted deontologically governed professions. *Information Systems Frontiers*.
<https://doi.org/10.1007/s10796-025-10691-2>

[23] Qiu, J., et al. (2026). Technostress and generative AI in EFL metacognitive reading and digital classroom burnout. *Scientific Reports*.
<https://doi.org/10.1038/s41598-026-47683-4>

[24] Lau, G. R., Kasturiratna, K. T. A. S., Goh, A. Y. H., Tong, E. M. W., & Hartanto, A. (2026). AI fatigue in human–AI interaction: Conceptual framework, scale development and validation, and associations with AI engagement. *Computers in Human Behavior Reports*, 23, 101186.
<https://doi.org/10.1016/j.chbr.2026.101186>