

Designing memorable XR under stress: An empirically grounded criterion for long-term memory binding^{*}

Alvaro Pastor^{1,2}, Pierre Bourdin-Kreitz^{1,2}

¹*Dpt of Computer Science, Multimedia and Telecommunication, Open University of Catalonia (UOC), 08018 Barcelona, Spain*

²*XR-Lab, Interdisciplinary Research and Innovation Hub, UOC, Rambla Poblenou 156, 08018 Barcelona, Spain*

Abstract

Designing extended reality (XR) for high-stakes, stressful environments requires understanding how stress reshapes cognition. This position paper translates an empirical finding into a concrete neurocognitive design criterion for resilient XR. A controlled virtual reality study refined the classical emotion-memory model by inducing a sustained high-arousal negative state during incidental encoding of neutral targets. Results revealed a dissociation governed by task relevance: memory for bindings between discrete, task-relevant elements was enhanced via unitisation, while relational binding to spatial context was impaired. This establishes a core criterion: under high-arousal stress, cognitive resources unitise task-critical information into a static memory blend, attenuating its links to fragile contextual relations. Consequently, goal-critical information in XR must be structured to be unitisable. This criterion mandates memory-aware XR information architectures and is fundamental for achieving resilience, trust, and situated explainability. The paper concludes by outlining a research agenda on the limits of unitisation, binding as an evaluation benchmark, and the ethics of adaptive cognitive scaffolding.

Keywords

Virtual reality, Episodic memory, Binding, Face-Name, Face-Context, Emotion

1. Introduction

Extended Reality (XR), an umbrella term for virtual reality (VR), augmented reality, and mixed reality technologies [1], is increasingly being deployed in high-stakes, stressful environments like emergency response and surgery [2].

However, its use in this particular context requires understanding how these states reshape cognition, particularly memory, a function on which operational performance critically depends. Indeed, a core challenge lies in designing for memory: an operator must reliably encode critical information (e.g., procedures, asset locations, and their associations) amidst high arousal and negative-valence emotional states. Although literature has shown that stress fundamentally alters memory processes [3, 4], current XR design paradigms often lack an empirically grounded model of how stress alters fundamental memory processes [5], creating a gap between design and cognitive reality.

This gap is critical for episodic memory binding, the neurocognitive system that integrates disparate elements of an experience (the ‘What’, ‘Where’, and ‘When’) into a coherent memory episode [6]. Indeed, binding is supported by two distinct mechanisms: relational binding, which creates flexible associations between separate items and their context (e.g., linking a tool to its location), and conjunctive binding, which fuses related elements into a single, static representation (e.g., visual features on a control panel). The former is particularly fragile under high-arousal and negative-valence emotional states [7].

Decades of research have established a predictable memory trade-off under negative arousal: enhanced memory for central details at the expense of context [8]. Critically, this prioritisation is governed by task relevance, not merely emotional salience [9]. Neutral information can become central if it is critical to the operator’s goal. This adaptive reallocation has direct, yet unexplored, consequences for

CHI 2026 workshop: "XR for Challenging Environments – Enabling Human Performance and Agency under Stress (XR4CE), April 13–17, 2026, Barcelona, ES

✉ alvaropastor@uoc.edu (A. Pastor); pbourdin@uoc.edu (P. Bourdin-Kreitz)

🆔 0000-0002-0467-8260 (A. Pastor); 0000-0002-8745-3787 (P. Bourdin-Kreitz)



© 2026 Copyright for this paper by its authors. Use permitted under Creative Commons License Attribution 4.0 International (CC BY 4.0).

designing XR information architectures, where all displayed elements may be emotionally neutral but are differentially critical to goal success.

This paper addresses that gap by translating an empirical investigation into a concrete neurocognitive design criterion for XR. A controlled VR study was conducted to isolate how a sustained, high-arousal negative state, operationalised as a realistic stressor, modulates the memory binding of task-relevant neutral information. The study tested a dissociation between two binding types: face-name binding, a challenging association that normally relies on conjunctive and relational processes, and face-context binding, a main relational process [10]. These findings refine the classical emotion-memory trade-off and provide an empirical foundation for resilient, memory-aware XR design.

2. Empirical grounding: The adaptive memory trade-off

To investigate the modulation of memory binding under a state analogous to operational stress, a controlled VR paradigm was designed and implemented. Its objective was to induce a sustained high-arousal negative state (EI) versus a neutral state (EN) during incidental encoding of emotionally neutral targets [11]. This design was based on incidental encoding, controlling for deliberate mnemonic strategies, and the effects of intrinsic emotional properties of the memoranda themselves (e.g., gesturing faces).

A mixed between-subjects design was employed. Forty-four participants completed an incidental encoding VR task as a gate agent, sequentially encountering twenty-four full-body animated characters [12], each endowed with a unique photorealistic face and a name. It should be noted that the characters were artificially created for the occasion, guaranteeing that no participant could have seen their faces before. Presentations were evenly split between two visually distinct virtual locations, serving as the spatial contexts for memory binding. The emotional induction was embedded multi-modally across the experience: a pre-encoding narrative video established the emotional state, while continuous, emotionally congruent auditory and visual elements were presented during encoding. The emotional condition used a video of an aircraft crash-landing followed by pre-validated images (e.g., spiders, snakes) and a soundscape of alarms and crying. This methodology elicited a genuine, sustained emotional response rather than a transient reaction.

Memory performance was assessed via forced-choice recognition tests administered immediately after encoding and after a 24-hour delay. This study measured face-name binding, which requires associating a complex visual percept with an arbitrary lexical label, and face-context binding, which requires associating the same face with its correct spatial location. The 24-hour delay was included to examine effects after memory consolidation.

The effectiveness of the emotional induction was confirmed by self-report measures, with participants in the EI condition reporting significantly higher negative-valence and arousal immediately post-encoding compared to those in the EN condition, before returning to baseline levels after the 24-hour delay.

The central memory findings revealed a clear and theoretically meaningful dissociation. Confirming the adaptive model, high-arousal negative emotion did not cause uniform impairment. Instead, a selective reallocation was observed. Memory for face-name associations was significantly enhanced in the EI condition compared to the EN condition, an advantage that was more pronounced after the 24-hour delay. Concurrently, memory for face-context associations was impaired in the EI condition, with a significant deficit emerging clearly in the delayed test. This indicates the stressful state prioritised binding of task-relevant elements, likely through a unitisation mechanism that fuses them into a single, static representation. In parallel, it attenuated the relational binding of those same elements to their spatial context. The magnification of both effects after consolidation underscores their relevance for the design of systems where information encoded under stress must be retained and retrieved over time.

Thus, the empirical results demonstrate that under sustained high-arousal stress, the long-term fate of information in memory is predictably governed by its task relevance, and that separate, task-relevant elements of high complexity (e.g., faces and names) can survive high-stress encoding experiences via

conjunctive binding, while relational binding fails. This dissociation provides the direct evidence for the neurocognitive design criterion advanced in the following section.

3. The neurocognitive design criterion and its implications

The empirical dissociation establishes a core neurocognitive criterion for the design of XR systems destined for stressful operations. Under high-arousal negative emotion, cognitive resources are allocated to unitise task-critical information, actively attenuating its links to contextual peripheral information, which becomes predictably fragile in memory. Unitisation refers to the process by which two or more discrete items are encoded such that they are remembered as a single ensemble, rather than as separate elements requiring relational binding. This reallocation is a fundamental constraint that must shape XR information architecture for high-stakes use. Critically, this prioritisation is governed by task relevance, not by the intrinsic emotional properties of stimuli or their mere spatial proximity to an emotional source. Neutral, spatially distributed elements that are goal-critical will be drawn into a unified, central representation.

A primary implication is that context cannot be a reliable medium. In a high-stakes XR environment, an interface that communicates vital procedural data through contextual association, for example, by instructing a technician to “use the tool placed near the malfunctioning panel passed-by earlier” builds upon a relational binding that the operator’s cognitive system is actively deprioritising. Consequently, design must ensure that essential data is spatially and conceptually fused with its immediate referent. This could manifest as persistent, world-locked tags on key objects, or by designing procedural steps as unified perceptual chunks where an action icon is inherently linked to a tool model, rather than associating them through layout or narrative.

This necessitates a shift from designing for salience to designing for unitisation. The goal is to structure presented information so that discrete elements are encoded as a single, static representation. For instance, an interface alert should not merely appear near a schematic diagram; it should be visually integrated such that the alert and the relevant component are processed as a unified entity. This design principle challenges prevailing notions of “seamless” or contextual interfaces, arguing instead for deliberate fusion of critical data organisation to withstand the cognitive reallocations induced by stress. This complements existing heuristic approaches to mitigating interface-induced stress [13] by providing a predictive, mechanistic account of why certain information architectures fail under high arousal.

A crucial caveat accompanies this benefit of unitisation. The resulting memory representation is a static blend. While this ensures the survival of the fused information, it simultaneously impairs the flexible, relational processing required to decompose the blend and recombine its elements with new, incoming information during recall or problem-solving. The very process that protects the data under stress also renders it cognitively less adaptable. Therefore, XR systems must not only create unitised representations but also provide external, real-time support for any required relational manipulation of that information, as the user’s capacity for such mental operation is compromised.

Finally, this rule poses a significant challenge to concurrent goals of trust and situated explainability in human–Artificial Intelligence (AI) teaming. An AI explanation that is contextually situated, such as an overlay stating “close this valve because the pressure in the adjacent pipe is critical”, relies on the user’s intact memory for the relational link between the valve and the earlier pressure reading. Under stress, this link is precisely what is impaired. Therefore, achieving trustworthy and explainable AI collaboration under these conditions may require a new paradigm for explanations, unitising the rationale with the actionable object in real-time. This reframes the workshop’s call for situated explainability, suggesting it must be pursued in tandem with conjunctive unitisation strategies for long-term resilience.

In synthesis, these implications converge on the need for memory-aware XR information architectures. Such architectures would treat the user’s binding resources as a key system variable, prioritising unitisable presentation for core task elements while minimising reliance on fragile contextual associations, and providing external scaffolding for the relational reasoning that stress inhibits.

4. Conclusion and future directions

In summary, XR design for high-stress environments cannot ignore neurocognitive reality. The empirical finding of a selective memory trade-off, where stress enhances unitisation of task-critical elements while impairing relational context binding, provides a fundamental criterion for resilient design. This leads to a core proposition: resilient XR must be memory-aware, and memory-aware XR treats the user's binding resources as a key system variable, prioritising unitisable presentation for core task elements while reliance on fragile relational associations is minimised and external scaffolding is provided. Integrating this perspective raises three central questions.

First, the spatiotemporal parameters for conjunctive binding. The study outlined here demonstrates that conjunctive binding can fuse separate, task-relevant elements presented in proximity (e.g., a face and an adjacent name). For XR designers, this raises a critical, unexplored parameter: what is the threshold of spatial and temporal separation beyond which two goal-critical elements will not be fused, and instead require fragile relational binding? Defining this unitisation boundary is essential for creating interfaces that reliably promote resilient memory under stress.

Second, memory binding as a benchmark: Given that stress impairs memory selectively by binding type, should binding-type accuracy become a standard cognitive benchmark for evaluating XR assistance in stressful scenarios? How can assays for unitised versus relational binding, inspired by paradigms like the incidental encoding task used here, be integrated into the iterative design and testing cycles of high-stakes XR systems?

Third, the ethics of adaptive cognitive scaffolding. A memory-aware XR system implements the neurocognitive criterion by making autonomous adjustments to information architecture based on an implicit cognitive diagnosis (e.g., "user is stressed, relational binding impaired"), a technical capability increasingly explored through embedded sensor data [14]. This shifts the system's role towards a cognitive prosthesis and raises critical questions for trustworthy XR. These can be framed as two interconnected challenges: the design and engineering challenge of creating robust, comprehensible systems that maintain user agency, and the human-systems challenge of ensuring such adaptations avoid dangerous skill atrophy and operate with appropriate consent and transparency. Defining these boundaries is a prerequisite for responsible innovation in high-stakes XR.

Addressing these questions is a necessary step in moving beyond designing for an idealised, stable user. It implies engineering for a dynamic symbiosis with the cognitively adaptive human in the loop, an endeavour that requires cognitive science as its grounding discipline. It is therefore recommended that the XR community treat the operator's cognitive architecture as a core system variable to be measured, modelled, and ethically engaged.

Acknowledgments

This work is part of the R+D+i project "Inhabiting the Hybrid: Contributions from Artistic Research in Interactive and Immersive Media", ref. PID2021-128875NA-I00, funded by MCIN/AEI/10.13039/501100011033/"ERDF A way of making Europe".

Declaration on Generative AI

The authors have not employed any Generative AI tools.

References

- [1] D. Burke, H. Crompton, C. Nickel, The use of extended reality (xr) in higher education: A systematic review, *TechTrends* 69 (2025) 998–1011. doi:10.1007/s11528-025-01092-y.

- [2] A. Abdulmajeed, J. Bail, F. AlSuwailem, A. Mostafa, A. Abdel-Azeem, Use of virtual reality exercises in disaster preparedness training: A scoping review, *Sage Open Medicine* 12 (2024) 20503121241241936. doi:10.1177/20503121241241936.
- [3] G. S. Shields, M. A. Sazma, A. P. Yonelinas, The effects of acute stress on core executive functions: A meta-analysis and comparison with cortisol, *Neuroscience & Biobehavioral Reviews* 68 (2016) 651–668. doi:10.1016/j.neubiorev.2016.06.038.
- [4] S. Vogel, L. Schwabe, Learning and memory under stress: implications for the classroom, *npj Science of Learning* 1 (2016) 16011. doi:10.1038/npjscilearn.2016.11.
- [5] L. Colusso, R. Jones, S. A. Munson, G. Hsieh, A translational science model for hci, in: *Proceedings of the 2019 CHI Conference on Human Factors in Computing Systems*, 2019, pp. 1–13. doi:10.1145/3290605.3300231.
- [6] H. Eichenbaum, *Memory binding in hippocampal relational networks*, in: *Handbook of Binding and Memory: Perspectives from Cognitive Neuroscience*, Oxford University Press, Oxford, England, 2006.
- [7] J. Bisby, A. Horner, L. Hørlyck, N. Burgess, Opposing effects of negative emotion on amygdalar and hippocampal memory for items and associations, *Social Cognitive and Affective Neuroscience* 11 (2016) 981–990. URL: <https://doi.org/10.1093/scan/nsw028>. doi:10.1093/scan/nsw028.
- [8] M. Mather, M. Sutherland, Arousal-biased competition in perception and memory, *Perspectives on Psychological Science* 6 (2011) 114–133. doi:10.1177/1745691611400234.
- [9] Y. C. Chiu, F. Dolcos, B. D. Gonsalves, N. J. Cohen, On opposing effects of emotion on contextual or relational memory, *Frontiers in Psychology* 4 (2013). URL: <https://www.frontiersin.org/journals/psychology/articles/10.3389/fpsyg.2013.00103>. doi:10.3389/fpsyg.2013.00103.
- [10] S. N. Moses, J. D. Ryan, A comparison and evaluation of the predictions of relational and conjunctive accounts of hippocampal function, *Hippocampus* 16 (2006) 43–65. URL: <https://doi.org/10.1002/hipo.20131>. doi:10.1002/hipo.20131.
- [11] A. Pastor, P. Bourdin-Kreitz, Comparing episodic memory binding outcomes after emotion induction in virtual reality, *Virtual Reality* (2026). doi:10.1007/s10055-026-01364-9, in press.
- [12] A. Pastor, P. Bourdin-Kreitz, Synthetic generation of photo-realistic 3D characters for face-based episodic memory measures, *Pre Print - Open Science Foundation* (2025). doi:https://doi.org/10.31219/osf.io/4j2g_v1.
- [13] N. Moraveji, C. Soesanto, Towards stress-less user interfaces: 10 design heuristics based on the psychophysiology of stress, in: *CHI '12 Extended Abstracts on Human Factors in Computing Systems*, CHI EA '12, Association for Computing Machinery, New York, NY, USA, 2012, p. 1643–1648. URL: <https://doi.org/10.1145/2212776.2223686>. doi:10.1145/2212776.2223686.
- [14] S. Davies, T. Owen, S. P. Walton, Exploring human behaviours of stress through virtual reality interactions using accessible, embedded sensors and user-personalised machine learning, in: *Proceedings of the Extended Abstracts of the CHI Conference on Human Factors in Computing Systems*, CHI EA '25, Association for Computing Machinery, New York, NY, USA, 2025. URL: <https://doi.org/10.1145/3706599.3719732>. doi:10.1145/3706599.3719732.