Can we be unconsciously conscious?

Can "human error" be just a corollary of the complexity of cerebral circuitry?
- David Slater

"Human error" is one of those phrases that claims / appears, to explain everything while often explaining very little. It tells us that a person acted, or failed to act, in a way that contributed to an unwanted outcome. But it does not tell us why that action made sense to the person's body, brain, situation, training, expectation, task design, or surrounding system at the moment it occurred. Too often, the phrase simply closes the inquiry at the point where the inquiry should begin.

The question raised for example by Gareth Lock's Linked in reflections on aborted take-offs, switching things off rather than on, and signals passed at danger is therefore a serious one. Are these actions always consciously committed in the ordinary sense of the word? And what do we mean by conscious? Is there a simple way of illustrating its ontology?

If we imagine a two-dimensional representation of the extents of awareness and agency /autonomy, that we have in deciding on our responses to external stimuli, we can draw a simple management consultancy two by two matrix to highlight the complexity of the issues. The four quadrants define the ontological spaces naturally as the Rumsfeld-esque Figure 1;

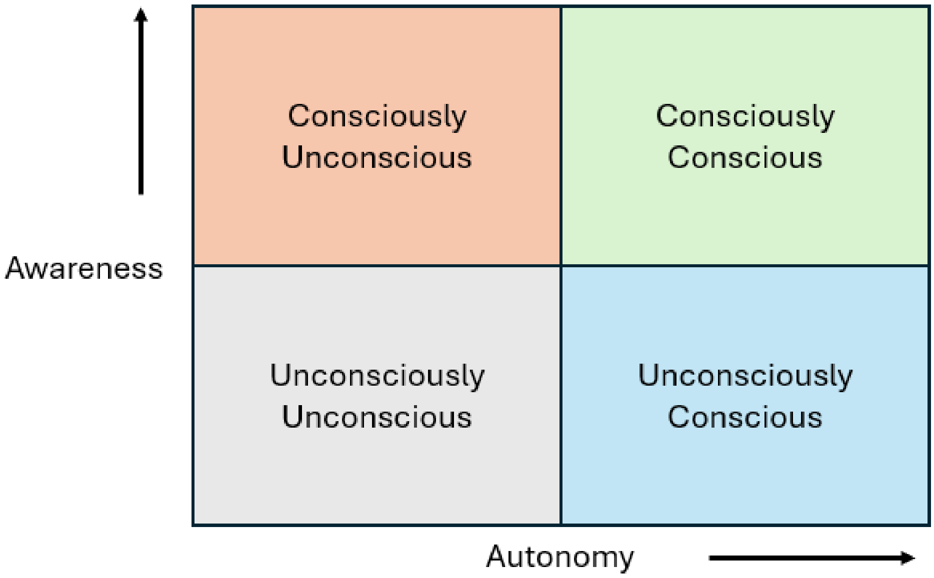

Figure 1 – The Awareness / Agency dimensions of consciousness (after Rumsfeld's unknown unknowns.)

- Unconsciously unconscious, (pure reflexes?),
- consciously unconscious, (breathing, heart rate, sweating?)
- Unconsciously conscious, (System 1, autopilot, in the flow?)
- Consciously conscious, (System 2, in charge, wait a minute!?)

We all instinctively know this and are continually reminded of it by perhaps realising that we had been "daydreaming" on autopilot; and seemingly the more intelligent the individual, the more skilled the artist or athlete, the more prone to such "lapses" of concentration we are. When looking at organisational accident rates the highest numbers always seem to be in the research labs!

Human error, automaticity, and the false comfort of blame

This sounds paradoxical only if we treat consciousness as a single thing. In practice, it is not. A person can be awake, alert, competent, trained and engaged in a task, while much of the perception, orientation and motor control that makes action possible is being handled below the level of reflective awareness. That is not a defect in the human brain. It is one of its most important adaptations.

The world presents too much information for deliberate consciousness to process item by item. Vision, sound, proprioception, expectation, memory, task demand, social context, time pressure and bodily state all arrive together. If every signal had to pass through slow, reflective reasoning before action could occur, ordinary life would become impossible. Skilled action depends on compression. The brain learns regularities, predicts what is likely to happen next, and allows well-practised routines to proceed with limited conscious supervision. Predictive-processing accounts describe the brain not as a passive receiver of information, but as an active prediction system, continuously matching incoming sensory information against expectation and updating behaviour when prediction error becomes significant (Clark, 2013).

Thinking fast and slow

This is why the familiar distinction between “System 1” and “System 2” remains useful, provided it is not taken too literally. Fast, automatic processes generate default responses; slower, more reflective processes intervene when the situation demands working memory, inhibition, hypothetical reasoning or a deliberate change of course. Evans and Stanovich describe Type 1 processing as rapid and autonomous, while Type 2 processing depends heavily on working memory and higher-order reasoning (Evans and Stanovich, 2013).The terms are imperfect, but the practical point is clear: much of what people do competently is not done by continuously thinking it through.

A better way to frame the issue is to distinguish **conscious access** from **executive control**. Conscious access refers to whether information is available for report, reflection and deliberate use. Executive control refers to whether an action is currently being selected, inhibited or modified by slow, intentional control. On that view, the interesting category is not the fully unconscious reflex, nor the fully deliberate decision, but the middle ground: the person is conscious of the task, but not necessarily conscious of the detailed micro-selection of action. They are, in this sense, **unconsciously conscious**.

We all know this state. You intend to drive to the shop on a Saturday morning and find yourself on the road to work. You reach for the wrong switch in a familiar sequence. You open the wrong drawer while thinking about the right object. The action is yours, but it is not experienced as a fully deliberate choice. A routine has been captured by context. A learned sequence has run ahead of reflective supervision. The surprise comes afterwards: “Why did I do that?”

Very “Human” factors

Human factors has long recognised this distinction. The Health and Safety Executive distinguishes unintended errors from deliberate violations and notes that slips and lapses commonly occur during familiar tasks, including highly trained procedures where the person no longer needs to concentrate on each action. HSE is also explicit that such errors cannot simply be trained away; better design can reduce their likelihood and make systems more tolerant when they occur.

James Reason’s work remains important because it separated different kinds of human failure rather than treating them as one moral category. A slip is not the same as a mistake, and neither is the same as a violation. In a slip or lapse, the intention may be correct while the execution, memory or sequencing of action fails. Reason’s later work also emphasised that adverse outcomes usually arise from combinations of active failures and latent conditions, not from isolated personal deficiencies (Reason, 1990; Reason, 2000).

The aviation example of Boeing 777 G-VIIT at Gatwick is unusually useful because it is documented in those terms. During the take-off roll on 28 June 2024, the co-pilot began retarding the thrust levers at V1 instead of removing his hand from them. He then advanced them again and initiated a rejected take-off shortly afterwards. The AAIB report page describes the event as an inadvertent thrust reduction and notes that the report considered the complex nature of the take-off roll, the possible benefit of mental rehearsal of motor actions, and the wider industry challenge of preventing action slips.

That incident matters not because it proves that all such events are unconscious, but because it illustrates the kind of phenomenon we need to take seriously. A highly trained, competent person can perform a wrong action in a safety-critical context without that action being consciously chosen in the ordinary moral sense. The wrong movement can be generated within a familiar motor ecology before the reflective self has fully audited it. The recovery may be immediate, but the action has already entered the world.

The seemingly incorrigible SPADS

Rail provides a parallel. Signals passed at danger are often described publicly as if the explanatory burden is carried by the phrase “the driver passed a red signal”. But that is only a description of the final visible act. RSSB reports that, in one analysis of 125 SPADs, 70% were cited as slip/lapse errors, while also emphasising that most SPAD events are not due to individual error alone and have multiple causes. That is the important safety lesson. The visible action is not the explanation. The explanation lies in the conditions that made that action possible, plausible, insufficiently interrupted and insufficiently recoverable.

Orientation before decisions?

This also reframes OODA. The brain does not simply Observe, then consciously Orient, then Decide, then Act in a neat serial loop. Much of orientation is already happening before conscious deliberation arrives. The brain acts through expectation. It samples the world for what matters. It escalates surprise, conflict and prediction error when the routine no longer fits. Only then does the slower, more effortful mode of reasoning become dominant. The “Decide” function is therefore not always a discrete conscious command; often it is an intervention into an already-moving stream of orientation and action.

Energy conservation

There is also an energy argument, although it should be handled carefully. The brain is metabolically expensive: Raichle and Gusnard note that it accounts for about 20% of the body’s oxygen and calorie consumption. But the stronger point is not simply that thinking uses energy. It is that conscious, controlled processing is limited and must be used selectively. It is needed for conflict, novelty, uncertainty, inhibition and re-planning. Routine action must be delegated, because if it were not, skilled performance would collapse under its own attentional burden.

Implications

The practical consequence is uncomfortable. Switches will be toggled the wrong way. Levers will be moved at the wrong moment. Signals will sometimes be missed. People will sometimes continue along the wrong well-worn path. This is not pessimism about human beings. It is realism about the conditions under which human beings achieve competent performance most of the time.

The wrong response is to imagine that an additional procedure, a stronger instruction, a poster, or another week of training can abolish automaticity. Training matters. Mental rehearsal matters. Briefing matters. Deliberate pause points matter. But they do not remove the underlying architecture of skilled cognition. If anything, expertise depends on that architecture. The more practised the task, the more likely it is that much of it will be executed through automatic routines. The safety question is therefore not, “How do we stop good people ever slipping?” It is, “How do we design systems in which predictable slips are less likely, more visible, more reversible and less consequential?”

That question leads directly to design. Is the wrong action physically easy? Is the right action perceptually obvious? Are incompatible controls sufficiently separated by shape, location, feel, direction of movement and consequence? Does the system provide feedback before commitment? Are high-consequence transitions protected by forcing functions, confirmation, cross-checks or delay? Can a second person, automated monitor or procedural rhythm detect the mismatch before it matters? Does the interface support the intended action under time pressure, or merely document the operator's failure afterwards?

Conclusion

This is where "human error" becomes a poor endpoint, but a useful starting point. The human brain is not a defective component inserted into an otherwise rational system. It is an adaptive, predictive, energy-constrained, habit-forming control system. Its great strength is that it can act without full deliberation. Its corresponding vulnerability is that, under some conditions, it can act before deliberation has caught up.

The implication for safety is profound. We should stop designing systems as if humans are continuously, consciously in charge of every relevant action. They are not, and could not be. We should design systems on the assumption that skilled humans will mostly perform brilliantly, sometimes slip predictably, and always need work environments that make the wrong action difficult, detectable, recoverable and survivable.

To be "unconsciously conscious" is not to be absent, negligent or irrational. It is to be human in the ordinary operating state of skilled work: conscious of the task, but supported by layers of learned, embodied and predictive action that do not wait for permission from reflective thought. Many accidents begin in that middle ground. Better safety begins by admitting that this is not an exception to human performance. It is one of its normal forms.

References

Air Accidents Investigation Branch (AAIB) (2025) 'Boeing 777-236, G-VIIT: Inadvertent thrust reduction during the takeoff roll, London Gatwick Airport, 28 June 2024', *AAIB Bulletin*, 5/2025, AAIB-30154. Available at: GOV.UK (Accessed: 4 May 2026). (GOV.UK Assets)

Clark, A. (2013) 'Whatever next? Predictive brains, situated agents, and the future of cognitive science', *Behavioural and Brain Sciences*, 36(3), pp. 181–204. doi: 10.1017/S0140525X12000477. (PubMed)

Dekker, S. (2014) *The Field Guide to Understanding "Human Error"*. 3rd edn. Boca Raton, FL: CRC Press. doi: 10.1201/9781317031833. (Taylor & Francis)

Evans, J.St.B.T. and Stanovich, K.E. (2013) 'Dual-process theories of higher cognition: Advancing the debate', *Perspectives on Psychological Science*, 8(3), pp. 223–241. doi: 10.1177/1745691612460685. (PubMed)

Health and Safety Executive (HSE) (2024) *Managing human failures: Overview*. Available at: HSE (Accessed: 4 May 2026). (HSE)

Le Coze, J.-C. (2022) 'The "new view" of human error: Origins, ambiguities, successes and critiques', *Safety Science*, 154, 105853. doi: 10.1016/j.ssci.2022.105853. (ScienceDirect)

Lock, G. (2026) ‘Pilot retards throttles during take-off’ (Available at Linked in [https://www.linkedin.com/feed/update/urn:li:activity:7456990778423422976/) -Accessed: 4 May 2026). (LinkedIn)

Mashour, G.A., Roelfsema, P., Changeux, J.-P. and Dehaene, S. (2020) ‘Conscious processing and the global neuronal workspace hypothesis’, *Neuron*, 105(5), pp. 776–798. doi: 10.1016/j.neuron.2020.01.026. (PubMed)

Mylopoulos, M. (2022) ‘Oops! I did it again: The psychology of everyday action slips’, *Topics in Cognitive Science*, 14(2), pp. 282–294. doi: 10.1111/tops.12552. (Wiley Online Library)

Norman, D.A. (1981) ‘Categorization of action slips’, *Psychological Review*, 88(1), pp. 1–15. doi: 10.1037/0033-295X.88.1.1. (ResearchGate)

Norman, D.A. (2013) *The Design of Everyday Things: Revised and Expanded Edition*. New York: Basic Books.

Raichle, M.E. and Gusnard, D.A. (2002) ‘Appraising the brain’s energy budget’, *Proceedings of the National Academy of Sciences of the United States of America*, 99(16), pp. 10237–10239. doi: 10.1073/pnas.172399499. (PNAS)

Reason, J. (1990) *Human Error*. Cambridge: Cambridge University Press.

Reason, J. (2000) ‘Human error: Models and management’, *BMJ*, 320(7237), pp. 768–770. doi: 10.1136/bmj.320.7237.768. (BMJ)

RSSB (2026) *Signals Passed at Danger: RSSB Graphic Insights*. Available at: RSSB (Accessed: 4 May 2026). (rssb.co.uk)

Schneider, W. and Shiffrin, R.M. (1977) ‘Controlled and automatic human information processing: I. Detection, search, and attention’, *Psychological Review*, 84(1), pp. 1–66. doi: 10.1037/0033-295X.84.1.1.

Shiffrin, R.M. and Schneider, W. (1977) ‘Controlled and automatic human information processing: II. Perceptual learning, automatic attending, and a general theory’, *Psychological Review*, 84(2), pp. 127–190. doi: 10.1037/0033-295X.84.2.127.

SKYbrary Aviation Safety (n.d.) *Error Management (OGHFA BN)*. Available at: SKYbrary Aviation Safety (Accessed: 4 May 2026). (Skybrary)

SKYbrary Aviation Safety (n.d.) *Human Error Types*. Available at: SKYbrary Aviation Safety (Accessed: 4 May 2026). (Skybrary)

Toner, J., Montero, B.G. and Moran, A. (2015) ‘The perils of automaticity’, *Review of General Psychology*, 19(4), pp. 431–442. doi: 10.1037/gpr0000054. (journals.sagepub.com)