

The Geometry of Boredom

Compression Progress and the Conditions of Sustained Attention

Flyxion

2026

Abstract

Boredom is standardly treated either as a deficit of stimulation or, on a folk-psychological reading, as a response to excessive predictability. This essay argues that neither raw stimulation nor raw predictability is the relevant variable, and gives the correct variable a precise formal home within this corpus's existing history-compression apparatus. Building on *History Before Function's* compression map $C(H) = F$ and on Schmidhuber's independently developed theory of compression progress, boredom is characterized as the condition in which no accessible action offers positive expected compression gain: every available option either has already been fully absorbed into the current operator F , or belongs to a domain no accumulation of history could ever compress. The same formal quantity, evaluated instead in its positive interior, recovers Csikszentmihalyi's flow channel, giving the account an independent empirical anchor. Voluntary handicapping, didactic reframing, and structural generality-seeking are then shown to be three distinct, formally specifiable mechanisms for restoring positive compression gain once a domain has saturated, rather than three unrelated pieces of folk pedagogical advice.

1 The Puzzle, Restated

Two intuitive theories of boredom are each individually inadequate. The first identifies boredom with insufficient stimulation, but repetitive, low-stimulation activities — a long walk, a familiar chore, a well-practiced skill — are frequently not boring, and are sometimes the opposite. The second identifies boredom with excessive predictability, but this cannot be the whole story either: a well-understood domain revisited for the pleasure of fluent execution is not boring merely because its outcomes are foreseeable, and conversely, a domain that is maximally unpredictable — pure noise, a shuffled deck examined card by card — is not interesting merely because nothing about it can be anticipated.

What the two failed theories share is an attempt to locate boredom in a single-variable property of the environment: how much is happening, or how surprising it is. This essay proposes that the relevant quantity is neither, but a comparison between an environment and the current state of an agent's own compressed model of it — specifically, whether engaging further with that environment offers any prospect of improving the compression.

2 Compression Progress

Definition 2.1 (Expected Marginal Compression Gain). *Let H be an agent's accumulated history and $F = C(H)$ its current compressed operator, in the sense of History Before Function. For an available action*

a with (possibly stochastic) outcome $o(a)$, define

$$\text{EMCG}(a) = \mathbb{E}_{o(a)} \left[|C(H)| - |C(H \cup \{a, o(a)\})| \right],$$

the expected reduction in description length achieved by folding a 's outcome into the compressed model, relative to not taking a .

This is a direct application, to this corpus's own H, C, F apparatus, of what Schmidhuber's theory of curiosity and creativity calls compression progress: an agent's intrinsic interest in an experience tracks not its absolute complexity or absolute predictability, but the *rate at which its predictive model is improving* as a result of the experience. The definition immediately separates three cases that a single-variable theory of stimulation or predictability cannot.

Definition 2.2 (Saturation). *An option space is saturated for an agent at time t if $\text{EMCG}(a) \approx 0$ for every action a accessible to the agent.*

Remark (Two kinds of saturation). *Saturation occurs for two structurally distinct reasons, and collapsing them is the error both folk theories make.*

Mastery saturation. *F already compresses the domain fully: every accessible a 's outcome is already implied by F , so folding it in adds nothing. This is the boredom of the solved puzzle, the completed level, the fully memorized route.*

Noise saturation. *No amount of additional history would improve compression, because the domain is not compressible relative to the agent's available model class: a 's outcome is genuinely unpredictable and remains so no matter how much H accumulates. This is the boredom of white noise, of a game of pure chance with no exploitable structure, of static.*

Both present subjectively as boredom. They are opposite conditions: one is a fully learned domain, the other an unlearnable one, and they call for opposite remedies, addressed in §4.

Remark (Recovering the flow channel). *Between the two saturation extremes lies the region $\text{EMCG}(a) \gg 0$ for some accessible a : the domain is not yet compressed by F , but is compressible with further history. This is, on the present account, the formal content of Csikszentmihalyi's flow channel, arrived at independently and from a different empirical tradition: engagement is sustained precisely in the region between mastery (further practice yields no compression gain) and overwhelm (the material is not yet, and may not currently be, compressible at all — a state closer to noise saturation from the agent's present vantage point, even where the domain is not truly incompressible in principle). The three-zone structure — boredom, flow, anxiety — recovers exactly the three cases $\text{EMCG}(a) \approx 0$ (mastery), $\text{EMCG}(a) > 0$ (interior), and $\text{EMCG}(a) \approx 0$ under present incapacity to compress (noise-like from the agent's current position, whether or not objectively noise) that the definition above distinguishes formally.*

3 Boredom Is Not the Absence of Stimulation

The overstimulation-based response to boredom, in which environments are populated with constant novelty and rapid activity switching, addresses the wrong variable on this account. Rapid switching guarantees a stream of nominally new experiences, but if none of them is retained long enough to be folded into a stable, improving F , no compression progress accumulates regardless of how much surface novelty is presented. The learner encounters a sequence of experiences rather than a system of relations: each item is consumed once and discarded, H grows in raw length

without growing in structure, and EMCG for the *next* externally supplied novelty remains just as available as before, but the learner’s own capacity to generate positive EMCG independently — to find or construct compressible structure once the external supply stops — has not been built at all. This is a dependency relationship, not a solution: engagement persists only for as long as external orchestration continues to supply pre-packaged, momentarily uncompressed material, and the underlying saturation reappears, undiminished, the moment it stops.

4 Three Mechanisms for Restoring Positive Compression Gain

Proposition 4.1 (Voluntary Handicapping). *Let F be saturated relative to a domain D (mastery saturation). Let $F' = C_{\mathcal{M}'}(H)$ be the compression of the same history H under a strictly weaker model class $\mathcal{M}' \subsetneq \mathcal{M}$ — fewer permitted operations, fewer available tools, a reduced representational vocabulary. Then generically $\text{EMCG}(a) \gg 0$ for some $a \in D$ relative to F' , even though $\text{EMCG}(a) \approx 0$ for every such a relative to F .*

Remark. *This is immediate from the definitions rather than a deep result, and its value is in what it makes explicit: handicapping does not alter the domain D at all. It alters which model class is doing the compressing. A weaker \mathcal{M}' has not yet exhausted its compression capacity against D , even where the unrestricted \mathcal{M} has, so restricting the tool set does not create difficulty out of nothing — it demotes the agent to an operator for which the difficulty, latent in D all along, becomes exploitable again. Solving arithmetic without a calculator, programming without high-level abstractions, or reading without relying on whole-word shortcuts are all instances of choosing $\mathcal{M}' \subsetneq \mathcal{M}$ deliberately, precisely because \mathcal{M} had already driven EMCG to zero on the material in question.*

Proposition 4.2 (Didactic Reframing). *Fix a single history H and a single completed operator $F = C(H)$ for one predictive target. Define a distinct target-relative compression $F_i = C_i(H)$ for each of several different structural questions $i = 1, \dots, n$ askable of the same H (what generalizes, what fails to generalize, what would break under a small perturbation, and so on). Then saturation of F_1 places no constraint on EMCG relative to F_2, \dots, F_n .*

Remark. *Didactic framing, in these terms, is the practice of never allowing a single H to be exhausted by a single compression target. Treating a repeated situation as a probe — asking not merely what happened but why it had to happen that way, or what would happen if one variable were changed — multiplies the number of simultaneously live F_i against the same underlying H , so that mastery saturation of one does not imply saturation of the whole attentional field.*

Proposition 4.3 (Structural Generality as Meta-Compression). *Let F_1, \dots, F_n be already-compressed operators, individually saturated with respect to their own domains. Define a second-order compression $F^{(2)} = C^{(2)}(F_1, \dots, F_n)$, folding the operators themselves into a higher-order operator capturing what is shared across them.*

Remark. *This is not merely an analogy; it is the same operation as History Before Function’s compression map applied one level up, and it directly instantiates that essay’s Open Problem 3 (higher-order compression of operators into meta-operators). Simultaneous saturation of F_1, \dots, F_n individually is compatible with EMCG remaining strongly positive at the level of $F^{(2)}$: the question of what algebra, a particular programming paradigm, and a particular grammar have in common as instances of a shared constraint structure can remain wide open exactly when each has, on its own, nothing further to teach. This is the formal content behind treating a single mastered domain as a route to structural generality rather than a terminus: what has been called elsewhere comparative generality is meta-compression across already-saturated operators, while structural generality within a single domain is the practice of Proposition 3.2 applied recursively to one’s own developing F_i before it saturates.*

5 The Precondition: Unstructured History

None of the three mechanisms above can operate on a history that does not yet exist. *Recursive Continuation*, §8, requires a recoverable H_t as a standing precondition for repair; the present account requires the same object as a standing precondition for engagement, for an identical reason. An environment saturated with continuous external structuring — every interval scheduled, every difficulty pre-resolved — does not merely fail to teach handicapping, reframing, or meta-compression as techniques; it forecloses the accumulation of the raw H those techniques would need to operate on. A learner who has never been given the chance to sit with an unstructured interval has never needed to construct their own F from scratch, and consequently has no practice generating positive EMCG independently once external supply ceases. Unstructured time, on this reading, is not the absence of a solution to boredom; it is the substrate the solution is built from.

6 Relation to the Corpus

This essay's central quantity, EMCG, is a direct application of *History Before Function's* compression map to a target Schmidhuber's independent theory of curiosity had already identified as the right one: not complexity, not predictability, but the rate of compression's own improvement. The saturation/mastery distinction sharpens that corpus's Open Problem 3 into a specific claim about when meta-compression becomes attractive (exactly upon saturation of the lower-order operators, not before), and the requirement of a pre-existing H before any of this can occur is the same requirement *Recursive Continuation* places on repair, applied here to engagement instead. *The Ecology of Constraints*, not yet written, is a natural next step from this essay's Proposition 3.1: voluntary handicapping is a single agent choosing its own model-class constraint, and an ecology of many interacting constraint-choosers is the natural generalization this essay's single-agent treatment does not attempt.

Open Problems

1. **Measuring compressibility independent of the agent's current model.** EMCG as defined depends on the current \mathcal{M} ; a domain that looks noise-saturated to one model class may be richly compressible to another. Distinguishing genuine incompressibility from currently-unavailable compressibility is not resolved here and matters directly for telling frustration from productive difficulty in practice.
2. **Productive boredom.** Some traditions of creative practice treat boredom itself as a precursor state to insight, rather than purely an aversive one to be escaped. Whether sustained mastery saturation on a narrow domain systematically triggers the meta-compression of Proposition 3.3, or merely disengagement, is an empirical question this essay's formalism poses but does not answer.
3. **Multi-agent compression targets.** This essay treats a single agent's H and F . Shared or competing compression targets across multiple agents — a classroom, a research community, a game with opponents — introduce dynamics (strategic incompleteness, deliberate obfuscation) this single-agent account does not model.

References

1. Jürgen Schmidhuber. "Formal Theory of Creativity, Fun, and Intrinsic Motivation (1990–2010)." *IEEE Transactions on Autonomous Mental Development*, 2(3), 2010.
2. Jürgen Schmidhuber. "Curious Model-Building Control Systems." *Proceedings of the International Joint Conference on Neural Networks*, 1991.
3. Mihaly Csikszentmihalyi. *Flow: The Psychology of Optimal Experience*. Harper & Row, 1990.
4. Daniel E. Berlyne. *Conflict, Arousal, and Curiosity*. McGraw-Hill, 1960.
5. Flyxion. *History Before Function: Compression as the Origin of Operators*. 2026.
6. Flyxion. *Recursive Continuation: Autopoiesis, Symptoiesis, and the Compression of Self-Modification*. 2026.
7. Flyxion. *Never Bored: Voluntary Constraint, Didactic Framing, and the Structural Conditions of Sustained Attention*. Unpublished manuscript.