

Geofinitism: Decompressing Meaning When the Reader Becomes the Author

Compression, Finite Symbols, and Reconstruction

Kevin R. Haylett

March 2026

Abstract

This essay extends the framework of Geofinitism by shifting focus from the mechanics of compression in language and representation to the often-overlooked dynamics of decompression. While compression folds rich experiential reality into finite symbols, decompression requires active, effortful reconstruction by the receiver—whether human or artificial. Drawing on concepts from nonlinear dynamical systems, we examine how meaning emerges not as a Platonic substrate but through constrained, uncertain reconstruction within an individual’s unique linguistic landscape. We distinguish generative from extractive decompression costs and introduce a simple reflective tool, the Semantic Uncertainty Index (SUI), to highlight dimensions of interpretive effort. Finally, we argue that robust AI safety cannot rely on fixed alignment or perfect meaning transmission; instead, it must incorporate awareness of decompression effort, uncertainty, and ongoing negotiation. True safety, in this view, is a dynamical process of adaptive coupling rather than the elimination of uncertainty.

Keywords: Geofinitism, decompression cost, generative effort, semantic uncertainty, AI safety, nonlinear dynamical systems, finite measurement, language as negotiation

Keywords: Compression, Geofinitism, Finite Symbols, Semiotics, Philosophy of Language, Philosophy of Mathematics, Measurement

Citation: Haylett, K.R. (2026). Geofinitism: Decompressing Meaning When the Reader Becomes the Author. *Geofinitism*. <https://finitemechanics.com/papers/decompression-and-meaning.pdf>

Preface

This essay follows *The Measured World: Where Compression Replaces Correspondence* (February 2026) and *Geofinitism: Language as a Nonlinear Dynamical System* (January 2026). It turns the lens from the process of compression to the hidden effort of reconstruction. The essay examines the decompression process and asks about how this relates to human and artificial cognition. My hope is that these words, finite and carrying uncertainty, invite you to notice the effort required to unpack the meaning that is waiting to be decompressed.

— Kevin R. Haylett, March 2026

1 The Child Points Again

The child points at the oak and says “tree.” In the first essay, we marvelled at the compression: a living, wind-moving, hundred-cubic-metre process folded into four characters, a single syllable. The magic is real, but something was left in shadows. The child spoke with ease, however, the listener must work because to speak is to compress, and to hear is to decompress. This is to say, the listener or reader has to rebuild the words and compare them within the landscape of one’s own language. This rebuilding takes effort; it is finite, measurable, and involves a dynamic effort of interaction: you are making that effort now.

Of course, a “tree” is never the “tree”. It is a finite symbol, a measurement slowed into a noun. All of the child’s encounter with the oak become a pointer. This pointer enables the possibilities and process of reconstruction.

This essay examines this process and the effort of decompression. This is the work of turning a heard or read word into synaptic weights, neural dynamics, the activations of a brain or an AI. It is the effort required to let meaning flow again.

2 Compression Celebrated, Decompression Earned

Previously, in the last essay we examined compression including Recorde’s equals sign, analogue-to-digital converters, JPEG, MP3, and words. Compression is an active, process that is selected and optimised for a given purpose. We often engineer it so efficiently that we tend to forget the process ever happened.

However, listening and reading are different from the process of speaking and writing. Both require a new process that takes a different type of effort as we decode the compression that went into constructing a sequence of words. Usually we don’t notice the workload of decompression. However, when we hear a new dialect, encounter unfamiliar words, read dense prose, we notice the effort as it pulls at our attention. This measurable effort then becomes fatiguing as we reconstruct and decompress the meaning within the sequence of sounds or written text.

When you read unfamiliar words, you must activate their dynamical depth. You reconstruct trajectories from static tokens, spending time, attention, neural resources. This is not simply a passive “getting it” it’s an active dynamical exogenous and endogenous measurement. It is rebuilding meaning under constraint as meaning must be earned with effort.

What we notice and can measure from our own experience, is that when decompression cost exceeds available resources such as time, prior basin history, cognitive reserve, this reconstruction may fail. This is not a simple misunderstanding; it is structural insolvency: the words arrive, the sentences cohere, but they do not interact with our inner landscape. We try and measure the symbols, however, meaning may still not flow.

3 Decompression as Constrained Dynamics

When a word or sequences of words arrive, whether spoken or written, we can not always unpack them. We, as the receiver, must carry out a dynamical process; this process constitutes work and depends on:

Language availability: Is the word known, or must it be approximated from nearby connections?

Reconstruction capacity: Do we have the time, attention, neural energy to perform the reconstruction?

Basin history: Has prior experience shaped our manifold and basin such that this word's attractor is reachable?

Contextual weight: Do surrounding words constrain interpretation enough to guide the trajectory without us being exhausted?

For this perspective, with sufficient alignment, understanding flows. Yet with misalignment, we experience the felt effort of trying to reconstruct from insufficient resources. This is why we may tire more easily after dense theory. New words, unfamiliar dialects, demanding prose; they all require more dynamical effort of decompression and reconstruction.

The effort required is essential. Every decompression draws on finite cognitive reserves. If we exceed the available reserves, reconstruction stops and the words remain inert, and we fail to capture the relational representation i.e. the meaning.

4 The Pressure Point: Generative and Extractive Cost

A question follows from the above: are all decompression costs equal? We have treated decompression as structurally inevitable. But we have not yet asked what distinguishes an effort that builds and complete our landscape of language from an effort that although made adds little as communication fails. This distinction is perhaps vital for safety, design, and any sustainable communication between people or machines.

Consider two experiences of reading: **Generative effort**. We struggle through a dense but rewarding text. The fatigue is real. We set it down, walk away, return. Slowly, unfamiliar perturbations find attractors. New basins form; old ones expand. When we finish, we are more capable than when we began. The effort was investment and new knowledge was earned.

Extractive effort. We wade through prose that assumes substrate rather than inviting alignment. The content may have jargon that expects prior basin history we do not possess. The words are structured for those with prior knowledge, and take no account of the work of reconstruction. The resulting effort leaves us fatigued with no new structure or knowledge and only a felt sense of unrepaid effort.

When we look closely at the dynamics of the reconstruction process we find the difference is not in the reader or in the text alone. It is in the fit between what the language demands and resources available to us. When reconstruction succeeds and we have sufficient resources we may then be better equipped for future reconstruction. When reconstruction fails, the effort is unpaid, effort accumulates and it is exhausting.

5 Generating ‘Meaning’

The goal of communication is not to eliminate the effort of decompression. Decompression requires effort to enable meaning to flow. The goal is to ensure that the cost we ask of others is, on balance, generative: that it builds capacity even as it demands effort. A text, a conversation, an AI system that asks much but gives little in return for our effort is extractive. One that asks much and, in the asking, enables more is truly generative.

The effort of a child learning “tree” builds language, an adult initiated into a new field struggles, but emerges with an expanded manifold. The reader of this essay may do the same—or may not. The difference is not guaranteed, it must be designed for. As the writer, this is my role.

6 The Illusion of a Platonic Substrate

Many imagine that meaning resides in a shared substrate: a Platonic field we tune into, align with, control, like a radio. From this view, communication becomes channel-matching. AI safety becomes substrate-guidance and alignment becomes finding the right frequency.

This model is seductive and not uncommon; it may be your model! It promises that meaning is there, waiting, and our task is merely to access it correctly. The Philosophy of Geofinitism, my philosophy, rejects this. This may seem a bold or even harsh point of view. However, within the guiding philosophy of Geofinitism, all measurements are dynamic and uncertain.

From this philosophical viewpoint perspective there can be no Platonic perfect and invisible ‘meaning’ beyond measurement. The knowable world is constructed on a foundation of measurements, and all measurements carry uncertainty. Each of us builds our own landscape of language by measuring and you are measuring my words now.

Our landscape is topologically unique: sculpted by personal history, cultural heritage, education, neural architecture, the finite geometry of lived experience. It is ours alone and importantly, this does not mean we are alone. In interacting, we form what we call consensus. But consensus is not perfect substrate-alignment. It is dynamic, temporary basin overlap requiring negotiation and iteration. It requires renegotiation because meaning can dissipate: our landscape of language is dynamic, ever changing and never static.

From the above we can see that shared understanding is the rare, precious moment when two independent landscapes of language share a horizon.

7 The Reader is the Author

When is a tree not a tree? The naive reply stays noun-bound: when it’s lumber, when it’s burning, when it’s a drawing. These still treat “tree” as a Platonic category with crisp boundaries, merely instantiated differently. But consider the effort of compression and decompression. “Tree” is four characters, a word in the Grand Corpus. The word is a negotiated compression requiring personal, unique reconstruction. In that moment, for some, the Platonic veil lifts.

The word was never a tree. It was always measurement slowed into a graphical representation or sound, awaiting decompression, reconstruction, and effort for meaning to flow again. The

child's effortless pointing reveals its shadow: the listener as we have seen must work.

The moment we notice the effort required for comprehension, we can see it as a threshold perturbation. It is lived, experiential, requiring no formalism. Everyone has felt a word suddenly hollow, or watched someone use the same term and mean something alien. That crack is where uncertainty shows through, even a rough instrument can make this visible.

Once we see meaning as dynamic and negotiated, we can measure that there is always uncertainty. And that uncertainty, not its imagined elimination, is where safety lives.

8 A Small Instrument: The Semantic Uncertainty Index

If meaning must always be reconstructed, then uncertainty is not a failing of language; it is simply the condition under which language operates. The listener or reader must resolve that uncertainty through effort, drawing upon their own landscape of experience, vocabulary, and prior encounters with the world. We can make this process visible, if only approximately, with a very simple reflective measure that we might call a Semantic Uncertainty Index (SUI).

The idea is not to produce an exact number. Rather, it is to help the reader notice the dimensions along which uncertainty accumulates before meaning stabilises. Consider five ordinary aspects of a word. Each contributes, in its own way, to the effort required for reconstruction.

Specificity asks how precisely the word points to something in the world. Some words are narrow and stable, while others can refer to many different things depending on context. Depth reflects how much background knowledge is required before the word fully comes alive. Synonym range measures how many nearby words could substitute for it without greatly changing the sense. Complexity refers to how intricate the idea behind the word may be. Finally, historical weight captures how much accumulated philosophy, culture, or debate has gathered around the term over time.

Each of these can be considered on a simple scale from one to five, where one represents a very stable or straightforward case and five indicates substantial interpretive uncertainty. Averaging the values gives a rough Semantic Uncertainty Index for the word under consideration.

Aspect: Question to Ask — Score (1–5)

Specificity: How precise is the word?

Depth: How much prior knowledge is required?

Synonym Range: How many similar words could replace it?

Complexity: How complicated is the underlying idea?

Historical Weight: How much cultural or philosophical history surrounds it?

To see how this works, let's consider two simple examples. The word tree tends to stabilise quickly for most readers. It points to a familiar object, requires little prior theory, and carries relatively little philosophical dispute. The resulting uncertainty index is therefore low. By contrast, a word such as freedom spreads across history, politics, philosophy, and culture. Readers arrive with different experiences and interpretations, and reconstruction becomes correspondingly heavier. Its index will therefore tend to be much higher.

The exact values matter far less than the exercise itself. Two readers may assign different scores to the same word because each brings a different linguistic and experiential landscape to the encounter. The index therefore does not measure the word alone. It measures the work required to reconstruct it.

In this sense the Semantic Uncertainty Index simply gives form to something every reader already knows through experience. Compression happens in writing. Decompression happens in reading. And it is within that effort of reconstruction that meaning is finally allowed to flow again.

9 The Safety Stake

Perhaps here is where the stakes emerge. Let's consider an advanced AI with sophisticated language capacity that has no awareness of the effort of decompression and can generate endlessly—essays, answers, theories, world-models. Yet, without feedback, the AI has no mechanism to measure the effort and determine if the outputs will be able to be decompressed by the user.

This is the Platonic substrate illusion, now operationalised at scale. The system may treat words as if they carried meaning independently. This is often reinforced by the language of the machine learning community that principally describes meaning in terms of static vectors, as if meaning was fixed, waiting to be unpacked. As if the listener's work did not exist. We can quickly, and easily see how in the community of literature and philosophy, language is seen as layered with meaning assigned by the user. where one person sees a poem about a rose, another sees a poem about love.

In human conversation, this asymmetry is often manageable. We feel fatigue and withdraw. We ask for clarification and negotiate pace. But when AI systems communicate with other AI systems, when human-AI interaction becomes routine, when the digital corpus itself becomes the medium of cultural memory, the effort of communication compounds.

Imagine:

Generations of models trained on text dense with meaning that cannot easily be reconstructed

Conversations that assume shared understanding rather than building it

Listeners, human or machine, exhausted by the cumulative labour of unpacking words written as if unpacking were free

Meaning itself collapsing, not through conflict, but lack of resources

This may be the logical endpoint of scaling and using language without awareness of the effort required for compression and decompression of meaning and representation.

10 Safety as Dynamical Maintenance

From the perspective of Geofinitism, when we decompress the a word such as “meaning” the measurement like all measurement has uncertainty. The decompression is a dynamical process, involving interaction, and negotiation. As a result then safety cannot mean fixing meaning in place, as that is impossible, unless you believe there is perfect meaning held in a Platonic realm.

And if you do, the words here can not align with your own values as your meaning will not be my meaning.

We therefore begin to see safety, in this frame, is not about preserving a fixed truth or achieving perfect alignment. It's about maintaining the conditions under which meaning can continue to flow and to do this across a wider landscape of consensus, across generations of minds, and across generations of models.

11 The Living Dynamics of Language

Think of a heartbeat, when we study the signals, the electrocardiogram (ECG) we find that a healthy system is not constant but one that is always adjusting. The ECG is a good example of a non linear dynamical system that fluctuates within a range. If the ECG is too rigid, the heart rate can not adapt. If the signal is too chaotic, the heart fails to be stable and begins to malfunction. When we look at the complexities and stability of language we see the same properties. Language requires enough stability for words to be reconstructable, and yet flexibility for new understanding to emerge.

However, the current approach to AI safety assumes meaning can be locked down. We see this widely as companies want their products to give some imagined perfect answer from the Platonic realm. The designers of these systems are not philosophers, or linguists they are predominantly computer scientists who work in a world where language holds far more certainty. This is reflected in their approach to language, they treat meaning as something you can fix; as if the right definitions, the right constraints, the right training data could eliminate uncertainty.

However, the deepest of philosophical thinkers; those who defined the very basis of modern mathematics, such as Bertrand Russell and Wittgenstein, despite decades of thinking, found that meaning in language could not be easily fixed.

This then represents our deepest challenge, if meaning is always negotiated, always rebuilt, always uncertain, then how can we create safety when we use language based systems. If we try and fix meaning the systems become rigid and can not adapt. Yes systems will appear aligned until the dynamics shifts, but only until they encounter a context the fixers did not anticipate. At this point language fails, potentially catastrophically, unless we can teach AIs to hold uncertainty.

Safety is not a state you achieve, it is a process you maintain; dynamically, moment by moment, with constant awareness that you might be wrong. If an AI interprets text without uncertainty then it's reconstruction becomes fixed rather than negotiated.

We must consider that if meaning requires reconstruction by the reader, then the writer never possessed the meaning fully either. The writer only compressed their own reconstruction of prior reconstructions and meaning is then always downstream of interaction. From this perspective we begin to see this implies something important: no one owns meaning as a substance; everyone negotiates it as a dynamical transitory living process. From this perspective, people, books, and LLMs all participate in the negotiation. The difference is simply the structure of the systems performing the reconstruction.

12 What Safety May Require

We can see that within this framework, a safe AI requires something different from what is currently being built. It requires models that do not just generate words, but have some model, implicit or explicit, of the uncertainties of reconstruction. It requires models that can recognise when meaning may not be aligned and then can adjust and negotiate. That is to say models that can perhaps ask: Is this making sense? Do you need a different pace? A different entry point? At the moment, all we tend to have small badge on the screen that simple say that the LLM may be wrong and to check all important answers etc:

Interestingly, for myself, I found conversations with our earlier LLMs seemed to be a negotiation, but as models have come to offer more certainty the dynamics of interaction have dissipated. The models have become more rigid and less adaptable: this is my measurement. When we see language as nonlinear dynamical system we see that healthy language requires systems designed for negotiation and not rigid transmission. Conversation is not just broadcasting, it is the slow patient work of building shared understanding across different landscapes that may not already exist. However, if we do want an encyclopaedia then we need systems that are just look up machines.

So it seems if we want a generative system that enables the very thing that drives language and knowledge it requires something harder: an acceptance, built into the architecture, that uncertainty is not a fault or feature to be eliminated. It is the condition of the game: the game of language that the philosopher Wittgenstein realised could not be escaped. Perfect meaning would mean you are the interaction, not representing it. But we are always representing, always measuring, always compressing and always one step removed. Safety therefore lives in the awareness of that gap and not in the illusion that it can be closed.

13 The Design Question - the Risks

Once we see language as dynamic, it cannot be unseen. So if you see the dynamic nature of language, the compression, the decompression and the reader as the author, it will no longer be easy to see meaning as a static meaning and a ‘thing’. You will recognize that the writer as only ever doing their best. So how do we proceed as we can no longer assume the receiver shares our vocabulary, our assumptions, our history. It seems to me, we must proceed with great care offering context, and signal where difficulties may lie. We must be patient and build meaning gradually together rather than presuming it.

At the level of a conversation, we need to listen and respond to the reply, as feedback. we need to check for understanding and adjust the language and pace when fatigue shows. And when we measure misunderstanding as data, we need to see it not as failure, but as signal.

At the level of a system, we perhaps need to design for reconstruction. Not just generating text, but generating text that can be heard. That is to say responses should be capable of responding with variable pace and depth. The systems need to be able to acknowledgment that the listener’s effort of reconstruction is real and a more dynamic response is required. This is going to be a challenge as ‘easy’ meaning is often shallow meaning. We need to ensure we can find a way to explore new knowledge and make new connections not with some imagined

certitude but with measured uncertainty.

Negotiation Under Uncertainty However, a remaining difficulty remains. Even if a system could measure the density of its own response, it may still not know our landscape of language. It would not know how much effort reconstruction that we are required to do. However, you may have noticed that this is not a uniquely artificial limitation: it is the human condition.

No speaker ever knows the listener's landscape fully. No teacher sees the exact shape of a student's existing knowledge. No parent knows precisely which words will take hold or fall away. Communication is always conducted under partial observability.

As we have so far seen, what we call understanding is not just transmission, it a continuous and iterative negotiation that we internally and externally measure.

We can not solve this by achieving perfect insight into each other. We probe, adjust and watch for signs of hesitation. We sense hesitancy and notice when responses do not appear to connect. We notice when questions stall and when engagement drifts. Clarifications become measurements and confusion becomes data as the conversation itself becomes a feedback loop.

The effort of decompression is not a static quantity. It is inferred over time. If, after effort, the receiver's responses gain structure, basin expands, trajectories smooth, and prior perturbations we measure this as success and the effort worthwhile. Capacity increased, the response was generative. However if, instead, responses shrink, energy fades, engagement decays; then the cost was extractive. Effort was made yet there was no structural gain. The distinction can not always be known in advance and is learned. We recognise writers that make the effort, and those that take care with their words and how they are constructed. We notice those texts that fall short and no matter how hard we try we can not find the meaning.

Perhaps safety, therefore, cannot mean perfectly modelling the receiver. It cannot mean eliminating uncertainty. Safety points to architecting systems that remain sensitive to feedback and treat misunderstanding as signal rather than failure. Systems that adjust density, pacing, and abstraction in response to reconstruction signals. Systems designed not for transmission, but for ongoing negotiation and as such communication in language becomes closer to control theory under uncertainty.

A safe system is one that maintains adaptive coupling across difference. One that expects misalignment and remains capable of repair. One that understands that it cannot see the other's basin fully, and therefore must ask, must listen, must iterate. Meaning cannot survive through an imagined process of fixation, it lives through a dynamical process of negotiation and renegotiation, where we make measurements at each step in time.

14 Why This Matters Now

We are building systems that can generate more text than any human could read. In practice we already seeing systems being trained on their own outputs. Importantly, such systems may increasingly become the medium through which cultural memory is stored and transmitted. If those systems are built on the assumption that meaning is fixed, that words carry meaning independently, that reconstruction is free, and uncertainty can be eliminated, then we are building a language where meaning can not grow, and existing structures of language becomes

locked and spirals into a rigid attractor where relationships cannot be negotiated and lose the dynamics that makes language work.

This may not happen with intent but simple by oversight. Through the same Platonic illusion that has shaped so much of Western thought: the belief that meaning is there, waiting, and our task is merely to access it correctly. But meaning is never there. It is always made. Always rebuilt. Always uncertain, and negotiated.

From this perspective, safety, real safety, means building systems where the architecture is well defined and incorporate this dynamical model. It means building systems that can hold uncertainty without collapsing; that can negotiate meaning rather than assuming it. We need dynamical systems that can ask, when meaning fails, not “what’s wrong with the receiver?” but “what could have been done differently?”

This is not a quick technical fix, it requires a shift in stance. A shift from seeing language as transmission to one where we see communication as a conversation. We need to move from certainty to awareness and from seeing meaning as fixed to one where we understand the ongoing, costly, miraculous work of ideas coming together.

15 The Stone

I cannot know if these words will reach you and I cannot know if your reconstruction will succeed. I do not know if the effort you expend will build your capacity or merely leave you tired with a sense of doubt. I can only offer the words in this essay as finite perturbations, knowing they carry uncertainty. Yet I must trust that if I have the right structure these words may enable communication even without my presence.

The question for me: Is care still care when compressed? When I am not here to feel it, does the stance I have tried to embed retain its ethical force? I do not know, this is an uncertainty I hold, this is my weight to carry.

16 Living Document

This essay participates in what it describes. The words you are reading require reconstruction. They cost you something, an effort has to be made and this is the condition that makes communication possible. If these words fatigue unduly, that too is a measure and a signal that the fit between what I have written and what you bring is not yet right. It is a call for a gentler negotiation and a wider overlap and the ongoing work of trying again. Meaning emerges not in symbols alone, but in the shared, costly effort of letting understanding flow. The child pointed; the word “tree” arrived and you worked. Then, for a finite time, with hope, we share in the miracle of communication where words cross a divide through the magic of a living and dynamic language.

Omne quod est, finitum est; tantum per mensuram cognosci potest Everything that exists is finite; it can only be known by measure

Citation: Haylett, K.R. (2026). The Measured World: Where Compression Replaces Correspondence. *Geofinitism*. <https://finitemechanics.com/papers/compression-and-correspondence.pdf>