

On Structural Depth

A Measure of Authenticity in Physical and Created Systems

Hardik Zayne

Independent Research

February 2026

Abstract

This paper introduces the Structural Depth Function, a formally defined measure that quantifies the number of internally coherent process layers recoverable from any physical surface or created artifact. The central observation is that every real object carries a hierarchy of process signatures organized by temporal scale, from nanosecond-level photon statistics to geological-scale material formation. These layers are not independent; they form a coupled stack in which each layer constrains the statistical properties of those above and below it. A measurable quantity—structural depth—captures the total number of such coherent layers. The paper derives a fundamental asymmetry: the energetic cost of fabricating N coherent process layers grows exponentially with N , while the cost of detecting the absence of any single layer grows only logarithmically with the signal-to-noise ratio at that layer. This asymmetry follows from the Second Law of Thermodynamics and the source coding theorem. The function applies across domains—visual perception, material science, architecture, music, biology—and provides a substrate-independent metric for authenticity that does not depend on the medium, the observer, or the tools of creation.

1. Introduction

A woodworker runs a hand over a tabletop and knows whether the grain was machine-sanded or hand-planed. A musician hears the opening bars of a recording and knows whether it was performed live or synthesized. A geologist picks up a stone and reads its thermal and pressure history in the crystalline structure. A photographer glances at an image and knows whether it was captured through a lens or generated by software.

These judgments feel intuitive, but they are not arbitrary. In each case, the observer is extracting the same kind of information: the depth of recoverable process history encoded in the object. A hand-planed surface carries evidence of the tool angle, the grain resistance, the pressure variation from stroke to stroke, the micro-tearout where the blade caught a reversing grain. A machine-sanded surface carries evidence of uniform orbital motion. Both carry process evidence. One carries more layers of it.

This paper formalizes that quantity. The goal is not to produce a framework for evaluating images or detecting forgeries, but to identify a root-level structural property that applies across every domain where authenticity is meaningful. The claim is direct: there exists a measurable function—here called Structural Depth—that counts the number of internally coherent process layers an artifact encodes. This function predicts perceived authenticity across media, scales, and observer populations. And there is a provable asymmetry at its core:

encoding depth is exponentially expensive, while detecting absent depth is logarithmically cheap. This asymmetry is a consequence of thermodynamics.

2. Definitions

2.1 Process Signature

A process signature is a statistical regularity in a physical surface or signal that is attributable to a specific physical process operating at a specific temporal or spatial scale. A wooden beam in an old barn carries process signatures at multiple scales: the cellular structure from biological growth (decades), the saw marks from milling (seconds), the oxidation patina from atmospheric exposure (years), the load deflection from sustained weight (decades), and the moisture-gradient cracking from seasonal humidity cycling (years). Each is measurable, attributable to a known process, and statistically distinct from the others.

Formally, a process signature P_k at temporal scale τ_k is a statistical distribution $S_k(x)$ over observable features x , where S_k is determined by the physical laws governing the process at that scale. Oxidation follows Arrhenius kinetics. Creep follows power-law viscoelasticity. Grain growth follows thermally activated diffusion. Each process imposes its own statistical signature, and that signature is verifiable against known physics.

2.2 Layer Coherence

Two adjacent process layers L_i and L_{i+1} are coherent if their joint statistics are consistent with a shared physical history. Layer coherence is defined as the conditional mutual information between adjacent layers, given the physical laws governing both:

$$C(L_i, L_{i+1}) = I(S_i ; S_{i+1} \mid \Phi)$$

where Φ represents the physical laws connecting the two processes. For a coherent stack, this quantity is positive and bounded: the layers share information because they share a history. For an incoherent pair—where one layer was fabricated independently of the other—the conditional mutual information drops toward zero, or becomes negative when the layers contradict each other.

Coherence between layers is expensive to fabricate because it requires knowledge of the physical coupling between processes at both scales. Adding realistic-looking grain to a surface is cheap. Making that grain consistent with the oxidation pattern, the load history, and the moisture gradient requires modeling all four processes and their interactions simultaneously.

2.3 Structural Depth

The Structural Depth $D(S)$ of a surface or artifact S is the total number of internally coherent process layers recoverable from its observable properties:

$$D(S) = \max N \text{ such that } C(L_i, L_{i+1}) > \varepsilon \text{ for all } i \in \{1, \dots, N-1\}$$

where ε is a threshold set by the observer's detection sensitivity—the minimum coherence distinguishable from noise. A purely synthetic surface with a single layer of detail has $D = 1$. A photograph of a real scene inherits whatever depth existed in the scene, typically $D = 5$ – 12 depending on the subject. A Vermeer has an

estimated D of 8–10. A quick digital render has $D = 1\text{--}2$. The observer does not count layers consciously, but the perceptual system evaluates them in parallel within the first few hundred milliseconds of exposure.

3. The Temporal Scale Hierarchy

Process signatures are organized by the temporal scale of the process that produced them. For a typical physical surface, the hierarchy spans roughly eighteen orders of magnitude:

Scale	Temporal Range	Process	Observable Signature
Quantum	$10^{-15} - 10^{-9}$ s	Photon arrival statistics	Shot noise distribution
Electronic	$10^{-9} - 10^{-6}$ s	Sensor readout	Read noise, pixel response non-uniformity
Optical	$10^{-6} - 10^{-3}$ s	Light transport through medium	Aberration, vignetting, dispersion
Mechanical	$10^{-3} - 10^0$ s	Hand tremor, tool contact	Motion blur, tremor spectrum
Atmospheric	$10^0 - 10^3$ s	Weather, scattering	Turbidity, wavelength-dependent haze
Biological	$10^3 - 10^7$ s	Growth, healing, wear	Asymmetry, scar topology, patina
Climatic	$10^7 - 10^{10}$ s	Erosion, oxidation	Differential weathering, stress cracking
Geological	$10^{10} - 10^{17}$ s	Material formation	Crystalline structure, isotope ratios

Table 1. The temporal scale hierarchy of process signatures.

The hierarchy is not exhaustive, and sub-layers exist within each band. The point is structural: real surfaces carry process evidence across many temporal scales simultaneously, and these layers are coupled by physics. The oxidation rate depends on the atmospheric conditions at one scale above. The atmospheric scattering depends on particle composition at a scale above that. The material formation process determines the crystalline structure, which determines the erosion rate, which determines the surface texture, which determines the optical reflectance. Each layer constrains every other through physical law.

4. The Markov Property of Depth

A central structural property of the temporal scale hierarchy is approximate conditional independence between non-adjacent layers, given the layer between them:

$$I(S_i ; S_{i+2} \mid S_{i+1}) \approx 0$$

Once you know the process signature at scale $i+1$, the signature at scale i tells you very little additional about the signature at scale $i+2$. The geological formation process determines the crystalline structure, which determines the erosion behavior. But the formation process does not directly predict the erosion pattern except through the crystalline structure. The relationship is mediated.

This Markov property has a practical consequence: to fabricate a coherent stack of N layers, it is sufficient to ensure pairwise coherence between adjacent layers. But this remains an exponentially growing task, because each pairwise coherence condition involves the physics governing the coupling between those two specific scales. There are $N-1$ couplings, each requiring knowledge of a different physical domain. A fabricator

simulating all N layers must solve $N-1$ distinct physics problems correctly and simultaneously.

The Markov structure also explains why partial depth is detectable. If layers 1 through 5 are coherent but layer 6 is absent, the observer's internal model predicts what layer 6 should look like given layers 4 and 5. When the prediction fails, the system registers an error at that specific scale. The error is localized to the missing layer, not distributed across the whole surface.

5. The Creation-Detection Asymmetry

The central result of this paper concerns the asymmetric relationship between the cost of creating structural depth and the cost of detecting its absence.

5.1 Fabrication Cost

To fabricate a coherent process layer, the creator must reduce the entropy of the output at that scale to match the target statistical distribution. By the thermodynamic cost of information erasure, each bit of entropy reduction costs at minimum $k_B T \ln(2)$ joules. The number of bits required to specify a coherent layer at scale τ_k is the description length of the statistical coupling between layer k and layer $k-1$:

$$b_k = I(S_k ; S_{k-1} \mid \Phi_k)$$

For each additional layer, the fabricator must also model the physical law Φ_k governing the coupling. The complexity of Φ_k tends to increase with the temporal separation between scales, because longer-timescale processes involve more interacting variables. The total fabrication cost for N layers is:

$$Q_{\text{fab}}(N) = \sum_{k=1..N} [b_k \cdot k_B T \ln(2) + Q_{\text{model}}(\Phi_k)]$$

where $Q_{\text{model}}(\Phi_k)$ is the computational cost of simulating the physics at layer k . For thermally activated processes—oxidation, creep, diffusion—the governing rate equations depend exponentially on activation energy. Correctly simulating the statistics of these processes requires solving coupled differential equations whose complexity grows with the number of interacting variables. In materials science and atmospheric physics, the computational cost of simulating processes at each additional temporal scale typically grows by roughly one order of magnitude:

$$Q_{\text{fab}}(N) \sim O(e^{\alpha N}) \quad \text{where } \alpha \approx \ln(10) \approx 2.3$$

5.2 Detection Cost

An observer detecting the absence of a process layer does not need to simulate that layer. It need only verify whether the observed statistics match the expected statistics. By the source coding theorem, this verification requires on the order of $\log_2(1 + \text{SNR})$ bits per sample, where SNR is the signal-to-noise ratio at the relevant scale. The total detection cost for checking layer k :

$$Q_{\text{det}}(k) = n_{\text{samples}} \cdot \log_2(1 + \text{SNR}_k) \cdot k_B T \ln(2)$$

This grows logarithmically with SNR and linearly with sample count. It does not grow exponentially with depth. The observer can check each layer approximately independently, because the Markov property means that each layer's expected statistics depend primarily on its immediate neighbors. The total cost of checking all N layers:

$$Q_{\text{det}}(N) = \sum_{k=1..N} n_k \cdot \log_2(1 + \text{SNR}_k) \cdot k_B T \ln(2) \sim O(N \cdot \log \text{SNR})$$

5.3 The Asymmetry

Combining the fabrication and detection costs:

$$Q_{\text{fab}}(N) / Q_{\text{det}}(N) \sim e^{\alpha N} / (N \cdot \log \text{SNR})$$

This ratio diverges exponentially with depth. At $N = 3$, fabrication is roughly 10 times more costly than detection. At $N = 6$, roughly 10,000 times. At $N = 10$, the ratio exceeds 10^9 . This is why a trained observer detects forgery faster than a forger can produce it. It is why biological mimicry targets specific predator sensory channels rather than achieving full-depth replication. It is why recognizing a poor restoration requires only a glance, while producing a convincing one requires understanding the full construction and weathering history.

The asymmetry is thermodynamic, not practical. Creating order at any scale costs entropy. Creating coherent order across multiple scales costs entropy at every scale plus the entropy of the coupling between scales. Detecting disorder at a single scale costs only the information-theoretic minimum for distinguishing signal from noise. The Second Law ensures that creation is always more expensive than detection.

6. The Critical Depth Threshold

Given a fabricator with energy budget E_{fab} and a detector with channel capacity C_{det} , there exists a critical depth N^* beyond which fabrication becomes intractable while detection remains feasible:

$$N^* = (1/\alpha) \cdot \ln(E_{\text{fab}} / (k_B T \cdot \ln 2))$$

This is the maximum structural depth a given fabricator can achieve. Beyond N^* , the fabricator lacks the energy or computational resources to maintain coherence across layers. The detector, however, continues to operate below its capacity ceiling, because detection cost grows only linearly with depth.

Consider current generative image systems as an example. Their training data encodes perhaps 3–5 layers of process evidence: composition, lighting, color statistics, coarse texture, some high-frequency structure. The systems can reproduce statistics at these layers. But the coherent coupling between atmospheric scattering and surface micro-texture, or between material aging kinetics and stress-concentration geometry, or between subsurface light transport and tissue absorption spectra—these require physical simulation that a learned statistical prior does not encode. The human visual system, calibrated by a lifetime of exposure to the real world, detects the absence of these deeper layers within hundreds of milliseconds.

N^* is not fixed. As fabrication technology improves and energy budgets increase, N^* shifts upward. But the exponential growth of fabrication cost means that each additional layer of depth requires roughly an order of magnitude more resources than the last. Progress in fabrication is linear or polynomial; the depth barrier is exponential. This is why the perception of authenticity has remained structurally stable across centuries of improving technology: each generation of tools adds a layer or two of depth, but the gap between achievable depth and the depth of physical reality remains.

7. The Zero-Tolerance Property

Structural depth exhibits a property that distinguishes it from most quality metrics: it is zero-tolerant. A surface with 8 coherent layers out of 10 does not score 80%. The missing layers produce specific, localized prediction errors in the observer's internal model that cannot be compensated by the quality of the layers present.

This follows from the hierarchical organization of predictive processing in biological perception. The visual cortex generates top-down predictions at every level of the hierarchy simultaneously. When the prediction at a particular level is violated—when an expected process signature is absent—a prediction error propagates upward regardless of how well other levels are satisfied. The error is not averaged with the successes. It is processed independently and triggers attentional recruitment at the violated level.

A single inconsistency in an otherwise careful fabrication can produce rejection. A well-rendered face with biologically implausible light transport at the earlobe. A convincingly weathered building with uniform erosion at stress concentrators. A high-fidelity piano recording with no pedal mechanism noise. In each case, the absent layer is more salient than all the present layers combined, because the perceptual system is tuned for anomaly detection, not quality averaging.

Formally, perceived authenticity is not the mean coherence across layers but the minimum:

$$A(S) = \min_k C(L_k, L_{k-1})$$

One weak link defines the chain. This minimum-dominance property follows from Bayesian inference: the observer's posterior probability of “real” is the product of the likelihoods at each layer. A single layer with near-zero likelihood drives the product toward zero regardless of the others.

8. Applications Across Domains

8.1 Visual Perception and Image Generation

The human visual system evaluates structural depth through a hierarchical predictive architecture operating across multiple spatial frequency bands, temporal integration windows, and statistical orders simultaneously. A natural photograph of a weathered stone wall carries process evidence at the quantum level (photon shot noise), the sensor level (pixel response pattern, demosaicing trace), the optical level (lens aberration, vignetting), the atmospheric level (wavelength-dependent scattering), the material level (differential erosion at stress concentrators, anisotropic cracking along crystal planes), and the geological level (mineral composition, thermal history). Each layer constrains the next. A generative system that has learned the marginal statistics at each layer but not the coupling between layers produces images whose individual features look correct but whose cross-layer coherence is absent. This absence is what the visual system detects.

8.2 Architecture and Built Environments

A building occupied for a century carries process evidence at the structural level (load-induced deflection following creep compliance curves), the material level (moisture-gradient cracking with species-specific swelling ratios), the surface level (differential patina following thermally activated kinetics with position-dependent rates), and the use level (wear patterns at contact points following abrasion mechanics). New construction that attempts to simulate age through artificial distressing achieves $D = 2-3$ at best: the visual layer and perhaps the material layer. The structural, chemical, and use layers remain absent. An observer with domain experience detects this quickly—not because the distressing is poorly done, but because

cross-layer coherence is missing.

8.3 Music and Performance

A live musical performance carries process evidence from the performer's motor system (signal-dependent noise with variance proportional to command magnitude), the instrument's physical response (resonance modes, sympathetic coupling, mechanical noise), the acoustic environment (room impulse response with frequency-dependent decay), and the recording chain (microphone coloration, analog noise floor). A synthesized performance may reproduce the note sequence and dynamic contour, but the coupling between motor noise and instrument response—the fact that a harder key strike produces louder sound and more key-bottom noise, more pedal vibration, more string-to-string sympathetic resonance—is typically absent. The listener perceives this as a lack of physicality, even without identifying the specific missing layer.

8.4 Biological Systems

Biological mimicry provides an evolutionary test of structural depth. Mimetic species typically achieve $D = 2-4$: visual pattern, coloration, sometimes movement. Full-depth mimicry—reproducing chemical signatures, behavioral repertoire, and social integration of the model species—is extremely rare, consistent with the exponential cost of each additional layer. Where deep mimicry does occur, it has required millions of years of co-evolutionary pressure—a timescale consistent with the exponential cost of fabrication amortized over evolutionary time.

8.5 Writing and Language

Written text carries process evidence at the lexical level (word choice distributions), the syntactic level (sentence structure), the semantic level (argument coherence), the pragmatic level (audience modeling), and the experiential level (specific sensory details consistent with having witnessed the described events). Machine-generated text has achieved high depth at the lexical and syntactic levels but often lacks coherence at the experiential level—the details are statistically plausible but not mutually constraining. A human writer describing a specific event produces details that cross-validate: the weather affects the clothing, the clothing affects the posture, the posture affects the dialogue. This cross-layer coherence is what careful readers detect as the difference between lived experience and approximation.

9. The Depth of Physical Reality

Physical reality has, in principle, unlimited structural depth. Every surface in the natural world carries process evidence from the quantum level through the cosmological level. The practical limit on detectable depth is the observer's channel capacity—the number of layers the observer's perceptual system can evaluate. For the human visual system, this is estimated at $D = 8-12$ simultaneously evaluable layers for static images and $D = 10-15$ for video, where temporal coherence adds further evaluable layers.

Physical reality therefore always has more depth than any observer can detect. There is always surplus process evidence. This surplus makes reality robust: even if the observer's evaluation at one layer is noisy or imprecise, other layers provide redundant confirmation. Fabricated artifacts lack this surplus. They satisfy the layers they were designed to satisfy and no more. The absence of surplus is itself a detectable signal—a meta-layer that the observer evaluates without being consciously aware of it.

This leads to a counterintuitive prediction: a fabrication that satisfies exactly the layers the observer is known to evaluate, and no others, is more detectable than one that also includes incidental depth at layers the observer is not checking. The residual, unchecked layers serve as a kind of watermark. Their presence signals that the object was produced by a process that did not optimize for the observer's evaluation criteria. Their absence signals that it was.

10. Discussion

10.1 Relationship to Existing Work

The structural depth function draws on several established bodies of work without being reducible to any of them. Process grammar formalized the idea that shape encodes causal history through symmetry-breaking operations. The free energy principle provides the perceptual mechanism by which the brain evaluates process signatures through hierarchical prediction error. The source coding theorem provides the information-theoretic foundation for the detection cost bound. The thermodynamic cost of information erasure provides the foundation for the fabrication cost bound. Allometric scaling laws provide the constraint on the observer's processing capacity. The contribution here is in unifying these into a single measurable function with a derivable asymmetry.

10.2 Limitations

The model is deliberately simplified in several respects. The Markov property of the temporal hierarchy is approximate; some processes create long-range correlations that skip intermediate layers. The exponential scaling of fabrication cost is an empirical observation for currently known processes, not a strict proof for all possible physics. The model treats all layers as equally weighted, whereas the observer's perceptual system applies precision weighting that varies across layers and contexts. A more complete model would incorporate layer-specific precision weights determined by the observer's prior experience and current attentional state.

10.3 Practical Implications

For creators working in any medium, the structural depth function suggests a specific strategy: rather than increasing the quality of individual layers—resolution, detail, polish—increase the number of coherent layers. A rough sketch with five coherent layers (composition, lighting direction, material response, atmospheric effect, maker's hand) will read as more authentic than a polished render with two layers (composition and surface color). This is not because roughness is preferable to polish. It is because depth is more perceptually potent than surface quality.

For authentication systems—forensic analysis, material verification, archival assessment—the depth function provides a principled basis for evaluation: rather than searching for specific artifacts, measure the cross-layer coherence across the temporal hierarchy. Artifacts at any single layer can be individually corrected. Coherence across many layers cannot be simultaneously fabricated without exponential cost.

11. Conclusion

The structural depth function formalizes a simple observation: real things carry more layers of internally coherent process evidence than fabricated things. This is not because fabricators lack skill. It is because the thermodynamic cost of encoding coherent depth grows exponentially, while the cost of detecting absent depth

grows logarithmically. The gap between these two growth rates is the permanent, physics-guaranteed foundation of authenticity.

The function applies wherever authenticity is meaningful: images, materials, buildings, performances, texts, organisms. It does not depend on the medium, the observer's expertise (though expertise raises the detectable threshold), or the tools of creation. It depends on the relationship between the Second Law of Thermodynamics—which makes creation expensive—and the source coding theorem—which makes detection cheap. Neither of these is changing. The structural depth function is therefore as durable as the physics it rests on.

This paper does not assign numerical values to the depth of specific artifacts or calibrate the threshold ϵ for specific observer populations. Those are empirical questions requiring experimental measurement. What is offered here is the theoretical structure: a definition, a scaling law, and an asymmetry result. The rest is measurement.