Apophatic Randomness: Why a "Truly Random Image" Cannot Be Positively Specified

Douglas C. Youvan

doug@youvan.com

www.youvan.ai

January 4, 2026

In everyday talk, a "random image" means static: pixel-level noise that looks patternless. Yet the moment we try to make that idea precise, we discover a hidden premise. Randomness is never a property of an object alone; it is a property of an object relative to a representation and a probability measure. To call an image "uniformly random," we must first decide what counts as an image (a pixel grid, a basis expansion, a continuous field) and what "uniform" means within that space. Those choices are not neutral: they import structure, privilege a coordinate system, and silently smuggle in a prior. If we demand a stronger notion—randomness "in itself," independent of any basis, distribution, or aesthetic curation—then the standard constructive program collapses. This paper proposes an apophatic alternative: define "truly random image" negatively, by a disciplined refusal of the very specifications that would make it fully describable, reproducible, or compressible. The result is not a recipe but a posture: a framework for speaking honestly about randomness without turning it into an idol of measurement, while still allowing operational approximations and rigorous boundary conditions.

Keywords: apophatic, randomness, measure, prior, representation, basis dependence, entropy, incompressibility, unpredictability, algorithmic information, Kolmogorov complexity, random field, invariance, maximum entropy, operational definitions, epistemology, negative theology, kataphatic scaffolding, selection bias, auditability.

A collaboration with GPT-5.2-Thinking. 51 pages.

Outline

1. Problem Statement: The Phrase “Truly Random Image”
1.1 Ordinary meaning: noise, static, patternlessness
1.2 The hidden premise: randomness requires a distribution
1.3 Thesis: the “truly random image” is ill-posed under positive specification

2. Randomness as a Relational Property
2.1 Object versus description: why randomness is not intrinsic
2.2 Sample space selection: what counts as “an image”
2.3 Measure selection: what counts as “equally likely”
2.4 Observer and test class: randomness relative to detectors of structure

3. Representation Choices as Structure Injection
3.1 Pixels as a basis: discretization is a prior
3.2 Alternative bases: Fourier, wavelets, learned dictionaries
3.3 Non-equivalence of “uniformity” across bases
3.4 The sampling problem: aliasing, correlations, and unintended constraints

4. Why “Maximum Entropy” Is Not an Escape Hatch
4.1 Maximum entropy requires constraints, constraints are structure
4.2 Symmetry principles: translation, rotation, scale and their hidden commitments
4.3 White noise as a choice, not a metaphysical default
4.4 The impossibility of a “neutral” invariance set

5. The Apophatic Turn
5.1 What apophatic means means: truth by negation rather than assertion
5.2 Translating apophasis to randomness: refusing privileged formalisms
5.3 “Image” as a name, not an ontology: the limit-concept move
5.4 Kataphatic scaffolding: permissible approximation without claiming identity

6. An Apophatic Definition of a Truly Random Image
6.1 Not basis-privileged: no coordinate system is canonical
6.2 Not measure-defined: no distribution is final
6.3 Not compressible: refusal of finite capture by fixed description languages
6.4 Not predictably generable: refusal of reproducible recipes
6.5 Not curated: refusal of aesthetic or post-selection filters

7. Operational Approximations Without Metaphysical Overreach
7.1 Physical entropy as a pragmatic anchor: quantum and thermal sources
7.2 Mapping entropy to images: why interpretation is unavoidable

1. Problem Statement: The Phrase "Truly Random Image"

The phrase "truly random image" sounds straightforward because the human mind already carries a convenient exemplar: television static, sensor noise, speckle, snow, grain—visual fields that feel un-authored. In that everyday sense, "random" means "I cannot see a pattern," and "truly" means "not fake." The problem is that this intuition is doing two silent substitutions at once. First, it substitutes a psychological criterion ("patternless to me") for a mathematical one ("drawn from a specified distribution"). Second, it substitutes a provenance criterion ("not generated by a deterministic recipe I distrust") for a metaphysical one ("independent of any describing frame"). Once those substitutions are made explicit, the phrase stops behaving like a simple request ("make me one") and turns into a philosophical claim about what it would even mean for an image to be random in itself.

In other words, the "problem statement" is not merely technical. It is about the hidden commitments embedded in our language: what we think we mean when we say "random," what we implicitly permit as evidence, and what we are unwilling to concede about representation, choice, and interpretation. The section that follows takes the ordinary intuition seriously, then shows why it cannot carry the weight of the word "truly" unless we either (a) accept explicit priors and measures, or (b) adopt an apophatic stance in which "truly random image" is defined primarily by what it refuses to be.

1.1 Ordinary meaning: noise, static, patternlessness

In common speech, an image looks random when it lacks recognizable structure: no edges, no objects, no gradients that read as lighting, no repeated motifs that read as texture. The paradigm case is static: a field that appears locally irregular and globally homogeneous. Importantly, this is not a definition of randomness; it is a report about the observer's pattern detector. Human vision is a compression engine. It is tuned to pick out edges, symmetries, repetitions, and low-dimensional regularities because those are behaviorally useful. So when the visual system fails to compress what it sees—when it cannot summarize the field by "it's a face," "it's a horizon," "it's a checkerboard"—it experiences the field as "noise."

This immediately explains why "random-looking" is not a stable category. A texture can be genuinely stochastic and still appear structured if it happens to contain salient clusters or streaks. Conversely, something can be non-random and still look random to a casual observer: many deterministic processes produce fields that defeat human pattern recognition at a glance. The eye's inability to detect structure is not evidence that structure is absent; it is evidence that the structure is not easily legible to the observer under the observer's informal test suite.

The phrase "patternlessness" is also ambiguous. It can mean "no obvious macro-pattern," while allowing micro-correlations. It can mean "no correlations at any scale," which is far stronger. It

can mean “no compressible description,” which is stronger still and depends on what descriptive languages count. Even the simplest “static” mental image carries an implicit ideal: a kind of visual whiteness—spatially uncorrelated variation with a flat spectrum. But “flat spectrum” is already a statement made in a particular representation (frequency space), and “uncorrelated” is already a statement made relative to a chosen sampling and neighborhood notion. The everyday meaning is thus a bundle of tacit technical commitments masquerading as a purely perceptual judgment.

1.2 The hidden premise: randomness requires a distribution

The hidden premise is blunt: randomness is not a substance; it is a rule for drawing outcomes. To say “X is random” is shorthand for “X was drawn from distribution D over sample space S.” Without S and D, “random” is at best metaphorical. This is why discussions of randomness that begin with “just make it unpredictable” eventually crash into the question, “unpredictable relative to what?”—what knowledge, what model class, what computational power, what prior beliefs, what tests.

When we say “random image,” we have already chosen a sample space, whether we admit it or not. Is the sample space all finite pixel grids of a given size? All possible continuous intensity fields over a square? All outputs of a camera pointed at a dark scene? All arrays of numbers that fit in a file format? Each choice creates a different universe of possible images, and different universes do not share a canonical notion of “uniform.” Even within a finite pixel grid universe, “uniform” depends on encoding decisions: how many intensity levels, what color space, what gamma curve, what channel coupling. If we instead choose the space of Fourier coefficients and declare them “uniform,” we induce a very different distribution on pixel intensities. No contradiction exists here: it is simply that “uniformity” is not invariant under change of coordinates. The distribution is the meaning of “random,” and distributions do not survive arbitrary reparameterization unchanged.

This is exactly where the word “truly” starts doing trouble. If “truly random” means “not pseudorandom,” one can appeal to physical entropy sources and say: the bits arise from phenomena not plausibly predicted by any observer. But if “truly random” means “random in itself,” independent of representation, then the requirement becomes: produce an object whose randomness does not rely on any chosen measure or encoding. And that is not merely hard; it is conceptually incoherent, because “random” without a measure is like “typical” without a population or “fair” without a rule of comparison.

So the hidden premise is not optional; it is what makes the word “random” mean anything at all. When someone insists they want randomness without a distribution, they are often expressing a legitimate discomfort: they do not want hidden priors smuggled in. But the

resolution cannot be to remove priors entirely; it can only be to expose them, or to redefine the aspiration negatively (apophatically) so that the "true" claim is no longer a positive specification.

1.3 Thesis: the "truly random image" is ill-posed under positive specification

The thesis of the paper is that "truly random image," understood as an absolute positive object, is ill-posed. The reason is not merely that we lack good hardware or perfect tests. The reason is that every positive attempt to specify such an image necessarily introduces the very structures the phrase is trying to transcend: a sample space, a distribution, a representation, a mapping, a test class, an acceptance rule. The constructive act is already a commitment to a grammar of description, and that grammar—no matter how carefully chosen—functions as a prior.

If one says, "Make each pixel independent and uniformly distributed," one has privileged pixel coordinates. If one says, "Make Fourier coefficients independent and uniformly distributed," one has privileged frequency coordinates. If one says, "Maximize entropy subject to translation and rotation invariance," one has privileged a symmetry group and a constraint set. If one says, "Use quantum randomness," one has privileged a physical measurement apparatus and a protocol for turning events into bits. None of these are illegitimate; they are what operational work requires. But they do not deliver "randomness without framing." They deliver "randomness relative to a frame," plus a provenance story about why that frame is acceptable.

Therefore, if we keep the adjective "truly" in its strongest sense—meaning "not dependent on any privileged positive description"—then the honest move is to stop offering constructive definitions and to adopt an apophatic one. On an apophatic definition, "truly random image" is not a recipe output; it is a boundary concept that names the refusal of the usual idol-making move: the transformation of uncertainty into an object with a neat positive specification. The paper's wager is that this is not a retreat into vagueness; it is a disciplined clarification. We can still build operational approximations (and must, if we want actual images), but we will no longer confuse the scaffolding (pixels, measures, invariances, entropy sources, tests) with the thing itself. The "truly random image," in the apophatic sense, is precisely what cannot be captured by any final positive statement without ceasing to be what it was meant to name.

2. Randomness as a Relational Property

If the phrase "truly random image" is to mean anything more than an aesthetic impression, we need to relocate randomness from the object to the relation that produces, encodes, and evaluates the object. Randomness is not like "red" or "square," a property that can be read directly off an artifact. It is more like "fair" or "typical": it only becomes meaningful once you specify the rules of the game. This section argues that randomness lives in a four-part relation:

an outcome, a space of possible outcomes, a measure over that space, and an observer (or test class) that tries to compress, predict, or distinguish the outcome from alternatives. When any one of these terms changes, what "random" means changes with it. That is not an inconvenience; it is what prevents the concept from becoming mystical hand-waving. It is also what explains why "randomness-in-itself" is so hard to hold onto: the very act of making a claim about randomness supplies the relational frame that the claim pretends to transcend.

2.1 Object versus description: why randomness is not intrinsic

Consider a single fixed image: a particular finite array of intensity values. By itself, it cannot be random or nonrandom in any absolute sense, because "random" is about how likely it was to appear under some generative rule. The same image could be extremely unlikely under one distribution and entirely typical under another. There is no contradiction in that; the contradiction only appears if we mistakenly believe randomness is stamped onto the artifact like a watermark.

A simple thought experiment shows the dependence on description. Take an image that looks like television static. One person says, "That's random." Another says, "No—someone generated it with a deterministic algorithm using a short seed." Both can be correct relative to their descriptive frame. In the first frame, "random" means "passes my informal perceptual tests." In the second frame, "random" means "incompressible relative to a fixed algorithmic language plus a bounded seed." The dispute is not about the pixels; it is about what counts as an explanatory compression.

This extends beyond provenance into logic. Any finite object can be made "nonrandom" by embedding it in a description that singles it out ("the unique image printed on page 37 of my notebook"). Conversely, any finite object can appear "random" relative to a weak observer who lacks the tools to detect the structure that generated it. So the intrinsic-property view collapses: if randomness were intrinsic, it would not be so easily toggled by changing the descriptive resources or the context of selection.

The deeper lesson is that randomness is not just about unpredictability but about *unpredictability relative to a model class*. To call something random is to say: "Within the language and power I am granting the observer, there is no shorter description, no better-than-chance predictor, no exploitable regularity." That statement is always indexed to a notion of description. And indexing is relational.

2.2 Sample space selection: what counts as "an image"

Before we can even discuss randomness, we must decide what outcomes are allowed—what constitutes "an image." This sounds trivial until you watch it distort every later claim.

One natural sample space is “all pixel grids of size N by M with K intensity levels per channel.” That gives you a finite set. But even here, you have already chosen: resolution, bit depth, color model, and often an ordering convention. If you change any of these, the space changes, and “typical” images change with it. A 64×64 monochrome grid lives in a very different world than a 4096×4096 RGB grid. The intuitive notion “static” is not invariant across these worlds; it is an observer’s projection.

Alternatively, you might define images as continuous functions on a square, perhaps with bounded variation, or square-integrable intensity fields. Now your sample space is infinite-dimensional, and naïve “uniform randomness” immediately becomes problematic: there is no natural uniform distribution over such a space. You can impose one by choosing a basis and choosing distributions for coefficients, but then you have smuggled in a basis as an organizing prior. The sample space choice forces you either into discretization (pixels) or into function spaces where “uniform” is ill-defined without additional structure.

A third option defines the sample space operationally: “the set of all outputs my apparatus can produce under specified conditions.” That can be appealing because it ties the space to physics rather than metaphysics. But it is still a choice. It privileges a device, a measurement protocol, a calibration regime, and a mapping from sensor readings to stored images. The space of “possible images” becomes instrument-relative, not ontology-relative.

So “what counts as an image” is not a neutral preliminary. It is a selection of the universe in which randomness will be judged. The moment you choose a universe, you have already made a claim about relevance—about what variations matter and what variations are excluded.

2.3 Measure selection: what counts as “equally likely”

Even after you pick a sample space, you still do not have randomness. You need a measure: a way to assign probabilities, or at minimum a way to say which outcomes are “typical” versus “rare.”

In finite spaces, people often reach for “uniform.” But uniform with respect to what representation? Uniform over pixel intensities in pixel space induces one distribution over shapes and textures; uniform over coefficients in another basis induces a different distribution over pixel intensities and spatial features. Uniformity is not coordinate-free. It is a property of a measure, and measures are representation-dependent unless there is a canonical symmetry that forces a unique choice.

In continuous or infinite-dimensional spaces, the idea of “uniform measure” becomes even more fragile. There may be no translation-invariant finite measure. Attempts to define “white noise” require choices about covariance structure and topology. Even “maximum entropy” requires constraints, and constraints are effectively priors expressed in polite language. The act

of saying “I will assume translation invariance” is already a substantive commitment about what features matter.

This is where people feel the sense of “flaw” most sharply: to define randomness, you must decide what gets equal weight, and that decision is not derived from randomness; it precedes it. There is no escape from the fact that “equally likely” is a normative phrase. It encodes a judgment about symmetry, ignorance, or fairness—none of which are forced by the mere existence of the object.

Thus, measure selection is not an implementation detail; it is the core of the meaning. Once the measure is chosen, “random” becomes a clear statement: drawn from that measure. But the phrase “truly random” often signals a refusal to accept any measure as final. That refusal is exactly what motivates the apophatic move later: to preserve the word “truly” by refusing the claim that any positive measure exhausts the reality of “randomness-in-itself.”

2.4 Observer and test class: randomness relative to detectors of structure

Finally, even with a sample space and a measure, the claim “this particular image is random” is rarely a direct mathematical statement. In practice, it is a statement made relative to tests: a class of detectors designed to find structure—correlations, periodicities, compressibility, predictability, non-uniformity.

Different observers have different tests. The human visual system is a test class: it detects edges, symmetries, and objects. A simple compression algorithm is another test class: it detects redundancy in a particular coding scheme. A battery of statistical tests is another: it detects specific deviations from independence or uniformity. A sophisticated machine learning model is another: it detects complex regularities learned from data. When someone says “this looks random,” they usually mean “it defeats my current tests.” That can be a strong statement if the test class is strong; it can be a weak statement if the test class is weak. Either way, it is indexed.

This is crucial: randomness is not merely “absence of structure,” but “absence of structure detectable by the granted observer.” That is why pseudorandom generators are useful: they create outputs that are indistinguishable from random to bounded observers. But if the observer is unbounded—if you allow the observer to know the seed or the generating algorithm—then the same output is no longer random in that relative sense.

So we have a layered picture:

- **Ontic randomness** (if you believe in it): outcomes not determined even in principle.
- **Epistemic randomness**: outcomes unpredictable given an observer’s information.

- **Computational randomness**: outcomes indistinguishable from random by observers with bounded resources.
- **Perceptual randomness**: outcomes that defeat human pattern detection.

A “truly random image,” if meant apophatically, refuses to collapse these layers into a single positive certificate. It insists that every test class is a kind of lens—useful, but not final. It asks for humility: label the observer, label the tests, label the assumptions. And where the word “truly” is invoked, it warns that any claim to have *completed* the randomness description is precisely the point at which we have turned a relational property into an idol of specification.

This relational framing is not an evasion. It is the minimum honesty required for the concept to remain coherent. Randomness is not a thing you possess; it is a claim you make under declared constraints. The rest of the paper builds on that honesty to explain why apophasis is not a poetic detour but a disciplined way of preventing the claim from overreaching its warranted frame.

3. Representation Choices as Structure Injection

If randomness is relational, then representation is not a neutral container; it is an active participant in the relation. The moment we decide how an image is represented—what coordinates, what basis, what quantization, what file format assumptions—we do not merely *record* an image. We define what kinds of variation are elementary, what kinds of variation are expensive to describe, and what kinds of regularity are easy to hide or easy to reveal. In that sense, representation “injects structure” not by adding visible patterns, but by imposing a grammar in which some patterns are natural and others are unnatural. This is the heart of the objection you raised earlier: pixel-level talk already carries a metaphysics of discreteness, locality, and adjacency. That metaphysics may be pragmatically unavoidable, but it is not innocent.

This section makes the injection explicit. We start with pixels, then widen to other bases, and then show why the apparently simple phrase “uniform random” is not preserved under change of basis. Finally, we highlight the sampling problem: even if your entropy source is pristine, your representation and sampling choices can create correlations and constraints that are neither intended nor obvious.

3.1 Pixels as a basis: discretization is a prior

A pixel grid is a basis choice: it declares that the “elementary” degrees of freedom of an image are local samples at fixed positions. It also declares adjacency as a primitive notion: pixels are neighbors, and neighborhoods are meaningful. This already biases the kinds of structure that

will be easy to talk about. In pixel space, “uncorrelated noise” means independence across spatial neighbors; “smoothness” means similarity among neighbors; “edges” are local discontinuities. All of these are pixel-native concepts.

Discretization adds additional prior commitments:

- **Finiteness:** An image is not a continuous field but a finite array. This forces any claim of randomness to be about a finite set of possible arrays.
- **Quantization:** Intensities are mapped into bins. This introduces a lattice of possible values, which can impose subtle periodicities in statistics (especially after transformations).
- **Locality:** The representation privileges local interactions. Many global structures (e.g., periodicity, long-range correlation) are not “native” and may be harder to detect without switching to another representation.
- **Geometry:** The grid enforces an implicit Euclidean geometry with a chosen metric. Even the choice of rectangular sampling rather than hexagonal sampling, or uniform sampling rather than adaptive sampling, is a commitment.

None of these commitments makes the outcome “nonrandom” in the operational sense. But they do mean that pixel-level randomness is not randomness-in-itself; it is randomness relative to a discretized local coordinate system. Pixels are not merely a measurement convenience; they define what it means for variation to be elementary.

In an apophatic framing, this is exactly where you would place a warning label: “This representation is scaffolding, not ontology.” The pixel grid is a pragmatic coordinate chart, not the essence of the image. But without such labels, pixel talk quietly becomes metaphysical: it makes locality and discreteness feel like the natural language of reality.

3.2 Alternative bases: Fourier, wavelets, learned dictionaries

Once you recognize pixels as a basis, the door opens: there are many bases, each carrying its own implicit prior about what “simple” and “complex” mean.

- **Fourier basis:** Variation is decomposed into global sinusoidal components. In this representation, periodicity and frequency content become “native.” A single coefficient corresponds to a global pattern extending across the entire image. Smoothness corresponds to dominance of low frequencies; sharp edges correspond to high-frequency energy. Fourier representation privileges global structure and treats locality as emergent from many components.

- **Wavelet bases:** Variation is decomposed into multi-scale localized elements. Wavelets are built to capture both locality and scale: they can represent edges and transients efficiently. A wavelet dictionary implies a prior that images are often structured by sparse features across scales—an assumption that holds well for natural images, which is why wavelets compress them so effectively.
- **Learned dictionaries and latent spaces:** Modern representation learning (whether through autoencoders, sparse coding, or other techniques) constructs bases or coordinate systems adapted to a dataset. These representations embed a strong prior: the "simple" directions are those common in the training distribution. What looks like "noise" in pixel space might look highly structured in a learned latent space, and vice versa. A learned dictionary is not merely a coordinate system; it is a record of what a particular culture, dataset, or model has decided is typical.

These alternatives matter because the phrase "random image" changes meaning across them. An image that is "white" in pixel space may have a very different spectral signature than an image that is "white" in Fourier space. A wavelet-white image will look different still. And a latent-white image might produce outputs that look like weird "natural-ish" textures because the latent coordinates are not aligned with raw pixel statistics.

So representation is not merely where we store the image; it is where we decide what kinds of structure are elementary.

3.3 Non-equivalence of "uniformity" across bases

Here is the conceptual crux: "uniform random" is not invariant under change of basis.

If you choose random values independently and uniformly for pixel intensities, you get what people call "white noise" in pixel space: no pixel is privileged; neighbors have no statistical dependence (in the idealized model). But if you take that same image and look at its Fourier coefficients, they are not independent uniform draws in any straightforward sense; they are constrained by the transform, by conjugate symmetries (for real-valued images), and by the induced distribution of coefficients. The coefficients will have a distribution shaped by the central limit behavior of sums of many pixel variables, and the "uniform" character does not carry over.

Conversely, if you choose Fourier coefficients independently and uniformly (or even with equal expected power across frequencies), the induced distribution in pixel space is not independent uniform pixels. It will exhibit specific correlation structures. The image may look smoother or have global interference patterns because global sinusoids combine into structured spatial fields. Even if you choose coefficients to mimic flat spectral power, pixel-level independence will generally not follow.

Wavelets sharpen the point: choosing wavelet coefficients independently and uniformly typically produces images with conspicuous localized “features,” because wavelet atoms are localized. A “uniform” wavelet draw often looks like a sprinkling of multi-scale edges. Again: not the same object as pixel-wise uniformity.

So uniformity depends on coordinates. You cannot have “uniform random image” without answering: uniform in what representation? And any answer is a prior. There is no basis-free uniformity.

This is not a purely mathematical subtlety; it is the mechanism by which representation injects structure. When you say “I will draw uniformly,” you are not declaring neutrality; you are selecting the coordinate system in which neutrality is defined.

From the apophatic perspective, this is one of the strongest reasons to distrust positive specifications. They masquerade as “no assumptions” while encoding deep assumptions in the coordinate chart.

3.4 The sampling problem: aliasing, correlations, and unintended constraints

Even if we accept representation choices as scaffolding, there is a second trap: the sampling procedure can generate structure that is not present in the entropy source and not intended by the designer.

Sampling problems arise whenever a continuous or high-resolution process is mapped into a discrete representation:

- **Aliasing:** If the underlying signal has variation at scales finer than your sampling grid, those variations can fold into lower-frequency artifacts. In images, this can create moiré patterns, false textures, and apparent correlations. You can get visible structure from purely innocent undersampling.
- **Quantization-induced correlation:** Mapping continuous values into discrete bins can introduce nonuniformities and periodicities, especially if the mapping interacts with hardware noise floors or with transforms applied afterward. The quantizer is not a passive recorder; it is a nonlinear operation that can shape statistics.
- **Device artifacts and coupling:** Real sensors have spatial coupling: neighboring pixels can share electronics, share thermal environments, or share readout paths. Even if the underlying noise source is high-entropy, the readout mechanism can imprint correlations across rows, columns, or blocks.
- **Unintended constraints:** Certain representations enforce constraints automatically. For example, real-valued images impose conjugate symmetry in Fourier space; finite

windows impose boundary effects; certain compression formats impose block structure. A “random” field passed through a lossy or structured pipeline can acquire patterns simply because the pipeline is structured.

- **Post-processing and normalization:** Seemingly harmless steps—clipping values to valid ranges, normalizing brightness, converting between color spaces—can introduce dependencies. If you clip extremes, you create heavier mass at boundaries; if you normalize globally, you couple every pixel to every other pixel through the shared normalization factor.

The key point is that randomness can be destroyed not only by deliberate curation but by innocent engineering assumptions. You can start with an excellent physical entropy stream and still end up with an image whose statistics reveal the fingerprint of your sampling apparatus.

This is why “randomness” in practice is always an audit claim: you must keep track of transformations, quantizations, and mappings, not just the entropy source. And it is also why the apophatic stance is attractive: it discourages the hubris of “I have specified randomness,” and replaces it with “I have built a pipeline; here is what it refuses (curation, privileging, hidden constraints), and here is what it inevitably assumes (representation, sampling, device physics).”

In summary, representation choices do not merely express an image; they define a coordinate metaphysic. Uniformity is not preserved across coordinate changes, so “uniform random” is never basis-free. And even if the entropy source is pure, sampling and mapping can imprint structure. The phrase “truly random image,” if taken seriously, must therefore either accept its scaffolding openly or retreat from positive specification into an apophatic definition that treats every representation as provisional and every “random” claim as indexed and accountable.

4. Why “Maximum Entropy” Is Not an Escape Hatch

When people sense the prior-smuggling problem—“pixels are a choice,” “uniformity is basis-dependent,” “randomness needs a measure”—they often reach for a familiar rescue: maximum entropy. The hope is simple. If priors are unavoidable, then choose the least biased prior: the distribution with maximum entropy subject to what we know. That sounds like a principled way to keep “randomness” from being merely an arbitrary preference.

This section argues that maximum entropy is not an escape hatch from structure injection. It is a disciplined way of *organizing* structure injection. It does not remove priors; it refactors them into an explicit constraint set. And constraint sets are exactly where human commitments hide: symmetry assumptions, boundary conditions, scale choices, and “what counts as knowledge.” Maximum entropy is therefore best understood as a transparency tool rather than a

metaphysical solvent. It tells you what you are assuming by showing you what you had to constrain in order to get a unique distribution at all.

4.1 Maximum entropy requires constraints, constraints are structure

Maximum entropy is often described as "assume nothing beyond the constraints." That sentence contains the entire issue: *the constraints are the assumptions*. Without constraints, maximum entropy either produces a trivial distribution (or fails to exist), and with constraints it produces something precise—but only by accepting those constraints as the structural skeleton of the model.

In the context of images, constraints can come from many places:

- **Range constraints:** intensities lie in a fixed range, channels have fixed bounds.
- **Moment constraints:** fixed mean brightness, fixed variance, fixed covariance structure.
- **Smoothness constraints:** bounded gradients, bounded total variation, limited high-frequency energy.
- **Spectral constraints:** fixed power spectrum, band-limited content.
- **Geometric constraints:** boundary conditions, periodic tiling assumptions, padding conventions.
- **Physical constraints:** sensor noise models, photon statistics, electronics limitations.

Every one of these constraints excludes a vast portion of the sample space. Exclusion is structure: it says, "Images that violate this are not allowed," which is another way of saying, "I have already decided what kind of world we are in." Maximum entropy cannot be "pure ignorance," because pure ignorance over a rich space is not a single distribution; it is either undefined or infinitely non-unique. To get a determinate answer, you must carve the world into "allowed" and "disallowed," and that carving is the prior.

This matters because the phrase "truly random image" often implies something stronger than "as random as possible given some constraints." It implies "as random as possible without making any substantive choices." Maximum entropy cannot satisfy that demand, because its output is exactly a function of the substantive choices encoded as constraints.

So maximum entropy is not wrong. It is honest. It takes the inescapability of priors and says: "Fine. Then put the priors on the table." But it cannot pretend the priors are not there.

4.2 Symmetry principles: translation, rotation, scale and their hidden commitments

A particularly seductive class of constraints are symmetry principles. They feel neutral because they sound like "fairness": don't privilege any location, any orientation, any scale. But symmetry is not a free lunch; it is a commitment about what distinctions do not matter.

- **Translation invariance** says: the statistics of the image do not change if you shift it. This rules out boundaries as meaningful, which is already an assumption. Real images (and real measurement devices) often have boundary artifacts; treating them as irrelevant is a modeling choice. Even the decision of *which* translations count—integer-pixel shifts, continuous shifts, periodic shifts—encodes a geometry.
- **Rotation invariance** says: the statistics do not change under rotation. But which rotations? Continuous rotations or discrete right angles? Around the center or around any point? Do we rotate the square domain and crop? Or do we impose periodic boundaries so rotations are well-defined? These are not technical footnotes; they change the distribution.
- **Scale invariance** says: the statistics look similar under zooming. This is even stronger and even less neutral. Scale invariance picks out a special family of distributions that often exhibit long-range dependencies. It also raises immediate questions about cutoffs: scale invariance cannot hold literally at all scales in a finite pixel image. Where does it break—at one pixel, at the whole image size, or somewhere else? The choice of cutoffs is a structural decision.

Symmetry constraints also conceal cultural and perceptual assumptions. Translation invariance aligns with the idea that "all parts of the image are statistically similar." But natural images are not like that: skies differ from ground; centers differ from edges; subjects tend to be framed. Rotation invariance is often false in human imagery because "up" matters. Scale invariance may or may not align with the world. So symmetry can be a noble aspiration ("don't privilege") while simultaneously being a strong denial of real-world asymmetries.

Thus, symmetry is not neutrality. It is a declaration of what the model refuses to notice.

4.3 White noise as a choice, not a metaphysical default

In many discussions, "maximum entropy image" silently collapses into "white noise": independent, identically distributed pixel values. This is treated as the archetype of randomness, almost as if the universe itself prefers it. But white noise is not metaphysically default; it is one particular selection among many, and it is tightly tied to a representation.

White noise assumes:

- **Independence:** no correlations across space (in the chosen coordinate system).
- **Stationarity:** statistics are constant across the domain.
- **Locality as primitive:** "neighbor" relations are meaningful only because pixels are meaningful.
- **A flat spectrum** in a particular sense (often implied), which again depends on representation.

But consider alternatives. A maximum entropy distribution subject to a fixed power spectrum is not pixel-i.i.d.; it produces correlated textures. A maximum entropy distribution subject to smoothness constraints produces fields that look like clouds rather than static. A maximum entropy distribution subject to scale-invariant constraints produces fractal-like fields. Each can be "as random as possible" relative to its constraints. None is privileged by metaphysics; the privilege arrives via what you constrain.

Even the very idea of "uncorrelated" is not representation-free. Uncorrelated in pixel space does not mean uncorrelated in wavelet space. Independence is not invariant under linear transformations unless the distribution belongs to a special family, and even then the story is delicate. So the notion that white noise is "the" maximum entropy image is a projection of a particular coordinate choice.

So if someone says "maximum entropy gives true randomness," the correct reply is: maximum entropy gives the most spread-out distribution consistent with your declared constraints. If your constraints encode pixel-level independence and equal variance, you get white noise. That's not an escape hatch; it's a consequence of having already decided that pixel-level independence is the neutrality you want.

4.4 The impossibility of a "neutral" invariance set

At this point the deeper obstacle becomes visible: there is no uniquely "neutral" set of invariances. Neutrality itself is a value judgment about what distinctions matter.

Suppose you try to be neutral by adding more invariances:

- Add translation invariance to avoid privileging locations.
- Add rotation invariance to avoid privileging orientations.
- Add reflection invariance to avoid privileging handedness.
- Add scale invariance to avoid privileging size.

- Add intensity-shift invariance to avoid privileging brightness baselines.
- Add contrast invariance to avoid privileging dynamic range.

You quickly collide with impossibility or triviality. Too many invariances can collapse the model into something degenerate or undefined, or force you to introduce cutoffs and boundary conditions that reintroduce the very privileges you tried to remove. In a finite image, scale invariance must break somewhere. In a bounded intensity range, shift invariance must break somewhere. In a finite square domain, rotation invariance must be implemented with conventions that implicitly privilege the center or privilege periodic boundaries.

Conversely, if you use too few invariances, you have not solved the neutrality problem; you have simply hidden it. Why translate but not rotate? Why rotate but not scale? Why scale but not shear? Why Euclidean symmetries but not projective symmetries? Each "why" is a doorway back to priors.

So the impossibility is not a failure of the math; it is a feature of the epistemic situation. "Neutral invariance set" is like "neutral language" or "neutral worldview." You can aim for humility and transparency, but you cannot reach a state where no commitments remain. Even refusing to commit is a commitment: it commits you to leaving distinctions unresolved, which affects what you can conclude.

This is precisely why the apophatic definition becomes attractive. It does not pretend there exists a final, positive, neutrality-guaranteeing distribution. Instead, it treats each invariance set and each maximum-entropy construction as scaffolding: useful for making operational claims, but never entitled to claim metaphysical finality. The apophatic stance is not anti-mathematical. It is anti-idolatry: it refuses to let the convenience of a maximization principle masquerade as a basis-free notion of "true randomness."

In short: maximum entropy is a principled method for building the least-committal distribution given commitments. It cannot remove commitments. Symmetries are powerful, but they are not neutral; they encode what you refuse to distinguish. White noise is not the default of reality; it is a particular coordinate-aligned choice. And no invariance set can be both complete and neutral in any absolute sense. The lesson is not despair; it is clarity: if "truly" is to remain in the phrase "truly random image," it must be guarded apophatically, not granted by a single optimization slogan.

5. The Apophatic Turn

Up to this point the argument has been largely diagnostic: every attempt to specify a "truly random image" positively forces us to choose a space, a measure, a representation, and a test

class. Maximum entropy does not remove these choices; it concentrates them into constraints. Symmetry principles do not erase priors; they encode them in what the model refuses to distinguish. The result is a repeated pattern: the stronger we make the word "truly," the more every constructive definition looks like a category mistake—an attempt to make an object out of what is really a relation plus a declared frame.

The apophatic turn is the proposed resolution. It does not deny the usefulness of positive constructions. It denies their metaphysical entitlement. In apophatic mode, we treat "truly random image" not as a fully specifiable artifact but as a boundary concept that disciplines our speech: it reminds us that the very act of speaking about randomness tends to smuggle in structure, and it offers a way to keep the claim honest by defining the target primarily through negations—what it is not, what we will not pretend, what we refuse to certify as final.

5.1 What apophatic means: truth by negation rather than assertion

In apophatic theology (via negativa), the central motive is not obscurantism but fidelity: God exceeds creaturely concepts, so positive statements risk becoming idols—finite pictures mistaken for the infinite. The apophatic method therefore speaks truthfully by denial: not this, not that. The negations are not empty; they are guardrails. They prevent the mind from taking a conceptually convenient model and treating it as the thing itself.

When translated into epistemology, apophasis becomes a discipline of intellectual humility. It is the recognition that certain targets cannot be fully captured by positive predicates without distortion. The point is not that we cannot say anything; the point is that we must be careful about what kind of saying is appropriate. For apophatic objects, the most accurate statements are often those that refuse to fix a final picture.

This is not a rejection of reason. It is a refusal of premature closure. Apophasis says: your concepts are real tools, but they are also finite lenses. Use them, but do not worship them.

In the context of this paper, the apophatic move is prompted by an analogous risk: the risk that we will confuse a convenient measure (uniform pixels, flat spectrum, maximum entropy under chosen constraints) with "randomness itself." The apophatic method gives us a way to keep "truly" meaningful by refusing the very kind of positive certificate that would make "truly random image" into a complacent slogan.

5.2 Translating apophasis to randomness: refusing privileged formalisms

To translate apophasis into the randomness problem, we identify the analog of idolatry: privileging a formalism and then forgetting we privileged it. In practice this happens whenever we say things like:

- "A random image is i.i.d. pixel noise."

- “A random image has a flat power spectrum.”
- “A random image maximizes entropy.”
- “A random image is what passes these tests.”

Each statement can be operationally useful, but each becomes misleading when treated as final. It silently elevates one coordinate chart, one measure, one constraint set, one test suite into an ontological authority.

An apophatic stance toward randomness therefore begins by refusing the privilege:

- Not pixel-privileged: no single discretization is canon.
- Not basis-privileged: Fourier, wavelets, latent spaces do not carry final authority.
- Not measure-privileged: “uniform” is always “uniform relative to …”
- Not test-privileged: passing a battery of tests is not metaphysical proof.
- Not provenance-privileged as a certificate: even “quantum” does not abolish framing, because the mapping from events to images still encodes choices.

This refusal does not mean we cannot do science or engineering. It means we treat every positive formalism as *indexed*. We stop saying “random” as if it were a stamp and start saying “random relative to these declared commitments.” Where we cannot or will not declare a final commitment, we shift from assertion to negation: we say what we refuse to claim.

In other words, apophatic randomness is not “randomness without math.” It is “randomness without unacknowledged math.”

5.3 “Image” as a name, not an ontology: the limit-concept move

The word “image” is already a metaphysical lure. It suggests a thing: a stable object with intrinsic properties. But if we take seriously the earlier analysis, “image” is a representational act: a way of packaging a field of values into a human-usable artifact. In apophatic mode, we loosen the ontological grip of the word. We treat “image” as a name for a family of encodings, not as a final kind of being.

This is the limit-concept move. A “truly random image” is not something we can point to and say, “This is it, fully and finally.” It functions instead as a boundary term that marks the horizon of positive specification. It names what is approached but never captured by any one representational scheme.

A helpful way to express this is: the apophatic “truly random image” is an “image” only under quotation marks. It is a finite artifact produced by some pipeline, but the word “truly” refuses to

let that artifact be identified with any positive formal certificate. The “image” is the place where entropy meets interpretation; the apophatic stance keeps us from declaring that the meeting has produced an object with basis-free, measure-free randomness “in itself.”

This is not nihilism. It is a discipline for preventing category errors. We can still talk about:

- how an artifact was produced,
- what assumptions were used in mapping entropy to representation,
- which tests it passes,
- what classes of structure it seems to lack.

But we will not talk as if those statements collapse into a single essence called “true randomness.” The apophatic “image” is therefore a conceptual horizon: it preserves the aspiration to purity while denying the claim of possession.

5.4 Kataphatic scaffolding: permissible approximation without claiming identity

Apophatic traditions do not forbid kataphatic speech (positive statements). They subordinate it. Positive language is permitted as scaffolding—useful for prayer, teaching, and navigation—so long as it is not mistaken for exhaustive truth. The same balance is needed here. If we refuse all positive formalisms, we cannot build anything. If we accept a positive formalism as final, we commit the very error apophasis exists to prevent.

So we adopt “kataphatic scaffolding” as an explicit policy:

- We may construct operational images by choosing a representation, a distribution, and a pipeline.
- We may justify those choices pragmatically (instrument limits, intended use, symmetry preferences, testability).
- We may evaluate outputs with test suites and publish the results.
- But we must label the entire construction as scaffolding: “random relative to these declared commitments.”

This yields a workable ethical-epistemic posture:

- **Permissible:** “Here is an image generated from a physical entropy source, mapped to a pixel grid by a declared protocol, and tested against a declared battery; under those declarations, it behaves as random.”

- **Not permissible (apophatically):** “Here is a truly random image, full stop,” where “truly” pretends to mean “free of all representation and measure commitments.”

Kataphatic scaffolding also gives a way to reconcile the paper’s strongest claim (ill-posedness) with practical reality. We are not paralyzed. We can build and test and use. But we keep the word “truly” guarded: it becomes a marker of humility rather than a boast of completion.

In the apophatic turn, then, the phrase “truly random image” is redeemed by being reinterpreted. It no longer claims a final positive object. It becomes a disciplined refusal: a refusal to let any one representation, measure, or test class become an idol. And alongside that refusal, it permits engineering approximations—explicitly labeled as scaffolding—so that we can still act in the world without pretending we have captured randomness in itself.

6. An Apophatic Definition of a Truly Random Image

An apophatic definition does not aim to produce an object with a positive certificate. It aims to prevent a category error: the mistake of treating a convenient formalism as the thing itself. So the definition offered here is intentionally negative. It defines “truly random image” as a disciplined refusal of the main ways we normally “complete” randomness claims—by picking a basis, a measure, a compression language, a generator, and a curation rule—then quietly forgetting we picked them.

This is not meant to be a stunt. It is meant to be a stable posture. In ordinary scientific and engineering practice we inevitably choose representations and distributions, and we inevitably build pipelines that yield actual images. The apophatic definition does not deny that. It says: those are kataphatic approximations—useful scaffolding—but “truly” is reserved for what cannot be exhausted by any one of them. The definition therefore reads like a constitution: a set of prohibitions against false finality.

6.1 Not basis-privileged: no coordinate system is canonical

A truly random image, in the apophatic sense, is not “random in pixel space” as though pixel space were the native language of the image. Nor is it “random in Fourier space,” “random in wavelet space,” or “random in latent space.” It refuses to grant any coordinate system the authority to define what counts as elementary variation.

This is not a denial that coordinate systems exist. It is a denial that any single one is ontologically final. Basis privilege is one of the quietest ways we inject structure. When we call an image “i.i.d. pixels,” we have asserted that locality and adjacency are primitive and that independence should be judged in those coordinates. When we call it “flat spectrum,” we have asserted that global periodic components are primitive and that neutrality should be judged in frequency

coordinates. Each basis makes certain patterns easy to describe and others expensive; each basis makes certain regularities obvious and others invisible.

The apophatic clause “not basis-privileged” therefore means: any claim of randomness that depends on a chosen basis is declared as scaffolding, not essence. The truly random image is not “what looks random in my chosen coordinates,” because that collapses “random” into “coordinate-defined.” Instead, the apophatic image marks the horizon where basis choices are acknowledged but not worshiped.

Operationally, this clause functions as a warning label: whenever you assert randomness, you must also name the basis in which you evaluated it, and you must not treat that naming as a metaphysical completion.

6.2 Not measure-defined: no distribution is final

Randomness, in its standard mathematical form, is always a distributional claim. The apophatic definition does not deny this; it denies that any specific distribution can be treated as the final meaning of “truly random image.”

To say “uniform over pixel values” is to choose a measure. To say “maximum entropy subject to constraints” is to choose a measure. To say “flat spectrum” is to choose a measure (often indirectly via second-order statistics). Each measure yields a different class of “typical” images. The apophatic clause insists that “truly” cannot be attached to any of these as though one of them were the metaphysical ground truth of randomness.

What does it mean, then, to refuse a final measure? It means:

- We may use a measure for a purpose, but we do not reify it as *the* measure.
- We may test an artifact relative to a measure and a test class, but we do not treat passing as an essence.
- We may say “random relative to D,” but we do not say “random simpliciter.”

In apophatic terms, a “truly random image” is not the output of a privileged distribution. It is the refusal to pretend that any distribution closes the case. The word “truly” becomes a guardrail against the seduction of “I have found the correct measure.”

6.3 Not compressible: refusal of finite capture by fixed description languages

Compression is where “randomness” becomes epistemically sharp. A pattern is, in practice, a short description. A random-looking object is one that resists short description. But here again, the trap is subtle: compressibility is always relative to a description language—an agreed set of primitives and rules for summarizing.

The apophatic clause “not compressible” is not the naive claim “it cannot be compressed by any compressor,” which is too strong to certify and too dependent on what compressors exist. It is rather a refusal to let any *fixed* description language become the final judge of randomness. If we pick one language—say, a particular compressor or a particular modeling framework—then “incompressible” means “incompressible to that language.” Another language might compress the same artifact.

So the apophatic stance is: the truly random image is not the kind of object that yields to finite capture within a *privileged* descriptive regime. We do not claim to have enumerated all regimes. We refuse to elevate a particular regime to final authority. We treat compressibility as a diagnostic, not a metaphysical certificate.

This clause matters because positive definitions of randomness often smuggle in a hidden compression language. “Looks like static” is a human compression test. “Flat spectrum” is a frequency-domain compression lens. “Passes NIST tests” is a statistical-lens compression lens. The apophatic clause says: these are lenses. Use them. Do not confuse a lens with the thing.

6.4 Not predictably generable: refusal of reproducible recipes

Here the apophatic definition becomes especially pointed. A reproducible recipe is, by definition, a short description. If I can give you a fixed procedure that reliably produces the image, then I have made the image part of a compressible family. The procedure is a handle; it is structure.

Now, in practice, we can generate images that are unpredictable to observers by using physical entropy sources or cryptographic pseudorandomness. But “truly random image,” in the apophatic sense, refuses to conflate “unpredictable to a bounded observer” with “free of generative handle.” It refuses the boast: “Here is the recipe for true randomness.”

This is not a denial of physical randomness. It is a denial that the existence of a physical entropy stream yields a basis-free, measure-free, recipe-certified object called “the truly random image.” The mapping from physical events to stored pixels is a recipe. Even if the physical events are ontically indeterminate, the *image* is still the result of a pipeline that encodes choices.

So the apophatic clause “not predictably generable” means:

- We will not treat determinism as the only disqualifier.
- We will not treat “quantum” as a magic word that erases representational choices.
- We will not treat the existence of a pipeline as a final certificate of “truly.”

In this frame, the “truly random image” is not something you can promise to regenerate on demand, because “regenerate on demand” is already a claim of control, and control is a kind of

structure. At best you can say: “I can generate new samples from a declared process.” But you cannot claim identity between the declared process and the metaphysical target.

6.5 Not curated: refusal of aesthetic or post-selection filters

Finally, apophasis demands refusal of curation. Curation is where randomness is most commonly betrayed, because it is where human meaning re-enters by selection. Even if the generator is excellent, if you keep only the samples that “look random,” you have conditioned the distribution on a human aesthetic criterion. Conditioning is information. Information is structure.

This includes obvious acts—discarding “boring” outputs, choosing the “most random-looking” frame—but it also includes subtle ones:

- cropping to remove suspicious clusters,
- normalizing contrast “just for display,”
- rejecting outputs with visible streaks that are in fact typical,
- picking the one that best fits a narrative.

The apophatic clause “not curated” therefore functions as an ethical boundary: no post-selection based on appearance, no aesthetic filtration, no silent conditioning. If you must apply transformations for pragmatic reasons (file formats, display constraints), you declare them as scaffolding and acknowledge how they shape the artifact’s statistics.

This clause also disciplines the observer. Humans are pattern engines. They will see meaning in any finite sample. The apophatic posture does not try to suppress that; it refuses to make it authoritative. Curation is precisely the act of letting perception become law.

Synthesis of the apophatic definition

Taken together, these clauses define a “truly random image” apophatically as a boundary concept characterized by refusals:

- no canonical coordinates,
- no final measure,
- no privileged compression language,
- no reproducible essence-by-recipe,
- no aesthetic conditioning.

What remains is not a neat object but a disciplined stance toward claims. “Truly” stops being a boast and becomes a warning: whenever we are tempted to say “this is it,” we remember that our very saying is an act of framing. The apophatic definition keeps that framing visible.

In the next section, we will show how this stance coexists with practical work: how one can build operational approximations—especially using physical entropy—while keeping the scaffolding declared, audited, and prevented from masquerading as metaphysical completion.

7. Operational Approximations Without Metaphysical Overreach

An apophatic definition is a guardrail, not a gag order. It does not forbid building things; it forbids pretending that what we built exhausts what we meant. In practice, people still want random-like images for simulation, cryptography, sampling, art, calibration, and scientific experiments. So we need a disciplined way to proceed: construct operational approximations that are as honest as possible about their scaffolding, while resisting the temptation to convert operational success into metaphysical certainty.

This section outlines a workable posture. We begin with physical entropy as a pragmatic anchor—because it gives a provenance story that is not merely “a deterministic algorithm did it.” Then we confront the unavoidable interpretive step: entropy is not yet an image until we map it into an image representation, and that mapping necessarily imports choices. Then we discuss testing: not as proof, but as falsification and diagnosis. Finally, we describe an audit posture that keeps “construction” and “claim” disentangled, so that one can speak rigorously without claiming a basis-free essence.

7.1 Physical entropy as a pragmatic anchor: quantum and thermal sources

If “randomness” is indexed to a distribution, and distributions require choices, why bother with physical entropy at all? Because provenance matters. A deterministic generator can produce outputs that look random to many observers, but it is still a mechanism whose outputs are, in principle, fixed given internal state. Physical entropy sources provide a different kind of anchor: they yield sequences whose unpredictability is grounded in empirical interaction with the world, not merely in computational opacity.

There are two broad families worth distinguishing conceptually:

- **Quantum-anchored sources:** These rely on phenomena widely treated as fundamentally indeterministic in orthodox quantum mechanics: detection events, decay timing, tunneling fluctuations, vacuum fluctuations. The appeal is not merely “it’s small and noisy” but that, on standard interpretations, the event outcomes are not determined even in principle by prior classical state.

- **Thermal/shot-noise sources:** These rely on microscopic agitation and stochasticity in electronics and photonic processes: Johnson–Nyquist noise, shot noise in diodes, avalanche noise. Here the unpredictability is often treated as emergent from enormous degrees of freedom and practical irreversibility. It may not satisfy every metaphysical appetite for "ontic randomness," but it provides robust unpredictability in operational contexts.

Either way, physical entropy is not used here as a metaphysical trump card ("therefore truly random"). It is used as a pragmatic anchor that improves honesty: it reduces the chance that "randomness" is merely a clever deterministic construction being mistaken for indeterminacy. It also gives a concrete place to audit: the apparatus, its calibration, its shielding, and its failure modes.

However, apophatic discipline demands we say the quiet part out loud: even a perfect physical entropy source does not automatically produce a "truly random image." It produces raw events or raw analog fluctuations. The moment we digitize, threshold, whiten, buffer, packetize, and map those bits into image channels, we have introduced representational and procedural structure. The physical source can anchor unpredictability, but it cannot eliminate the necessity of scaffolding.

So physical entropy is not an escape hatch. It is a more defensible starting point.

7.2 Mapping entropy to images: why interpretation is unavoidable

Entropy is not an image. Entropy is a reservoir of unpredictability—an undirected "difference-maker." To become an image, it must be interpreted through a representation: a grid, a coordinate map, a color model, a bit depth, and an ordering convention. This is where the metaphysical overreach tends to happen, because the mapping step feels like a technicality. In fact it is the moment when "randomness" becomes "random image," and therefore the moment when choices become constitutive.

Several unavoidable interpretive commitments appear here:

- **Defining the target space:** resolution, aspect ratio, channels, quantization. This is the sample space. Choosing it says what outcomes exist.
- **Choosing a coordinate chart:** pixels are a chart; so are Fourier coefficients; so are wavelet coefficients. Any mapping must land in some chart.
- **Choosing an assignment scheme:** how raw bits are grouped into channel values. Grouping is not morally neutral: it can create subtle biases if the raw stream has drift, correlations, or imperfect whitening.

- **Handling boundary conditions:** does the image concept assume wrap-around periodicity or hard edges? Even the act of storing a finite patch implies a boundary policy.
- **Handling constraints of display:** gamma correction, color space transforms, clipping. These operations are interpretive; they can introduce dependencies and distort distributions.

The apophatic point is not "therefore don't map." The point is: mapping is unavoidable, so the honest posture is to treat mapping as scaffolding and to keep it declared. The metaphysical error is to let the mapping disappear and then treat the resulting artifact as if it carried intrinsic randomness.

This is also where the earlier critique of "pixels" becomes constructive. The problem is not that pixels are sinful. The problem is forgetting that "pixelhood" is an interpretive choice. If you choose pixels, you do so knowingly, for operational reasons, and you resist the slide from "chosen chart" to "canonical ontology."

In other words, mapping is where apophasis must be practiced most actively: you perform the interpretation while refusing to identify the interpretation with essence.

7.3 Testing regimes: falsifying non-randomness without "proving" randomness

Once an image is produced, the next temptation is to certify it: "It passed tests, therefore it is random." The apophatic stance rejects that leap. Tests can falsify certain kinds of non-randomness; they cannot prove randomness in the absolute sense the word "truly" tempts us toward.

Testing regimes are best understood as structured interrogations:

- **Do we see obvious biases?** Are some values more frequent than others? Are some regions systematically brighter? Do certain bit patterns recur?
- **Do we see correlations?** Are neighbors more similar than they should be? Are rows or columns coupled by readout electronics? Do blocks share artifacts?
- **Do we see spectral fingerprints?** Are there peaks in frequency content suggestive of periodic interference? Is there unexpected attenuation or amplification in certain bands?
- **Do we see compressibility?** Does a generic compressor shrink the file significantly? (If it does, there is redundancy.) If it doesn't, that's supportive but not decisive: incompressibility to one compressor is not incompressibility in principle.

- **Do we see stationarity failures?** Do statistics drift across the image, suggesting drift in the entropy source or mapping pipeline?

The philosophical discipline is to treat tests as *diagnostics* and *red flags*, not as metaphysical seals. A test suite is a lens. Passing means "no structure detectable by these lenses." It does not mean "no structure exists," and it certainly does not mean "this object is randomness-in-itself."

This is not pessimism; it is normal scientific modesty. We rarely "prove" an empirical generator is ideal. We instead build confidence by repeated attempts to falsify it and by understanding the generator's failure modes.

So the apophatic posture reframes testing as a negative practice: we are more justified in saying what the image is not (not obviously biased, not row-correlated, not spectrally peaked, not trivially compressible) than in saying what it is (the essence of randomness).

7.4 Audit posture: separating "construction" from "claim"

The central operational virtue proposed by the apophatic stance is auditability—not in the sense of revealing "the true randomness," but in the sense of making the scaffolding visible and contestable. This is where the distinction between construction and claim becomes decisive.

- **Construction** is what you did: the apparatus used, the environment, the digitization pipeline, the mapping choices, any whitening or conditioning steps, the file encoding choices, and the display transforms.
- **Claim** is what you assert: random relative to which basis, which measure, which tests; unpredictable relative to which adversary model; acceptable for which use case.

Metaphysical overreach occurs when construction is hidden and claim is inflated. The apophatic posture insists on the opposite: construction is explicit, claims are indexed.

A useful audit posture contains at least these elements:

- **Declared scaffolding:** "We chose this representation, this mapping, these parameters." No pretending the choices were forced by "randomness."
- **Separation of provenance from essence:** "Entropy was sourced from X" is not the same sentence as "therefore the image is truly random." Provenance supports an operational claim; it does not complete a metaphysical one.
- **Declared test class:** "We evaluated with these diagnostics; here is what they can and cannot detect." This prevents test-passing from masquerading as proof.

- **Failure-mode awareness:** “Here are the plausible ways structure could leak in: electronics coupling, drift, quantization artifacts, display transforms.” Naming these is part of honesty.
- **Claim minimalism:** Assert only what your construction plus tests warrant. If you want to say “random,” say “random relative to …” If you want to keep “truly,” reserve it for the apophatic sense: the refusal to claim finality.

In the end, operational approximations are not only compatible with apophasis; they are improved by it. The apophatic stance does not stop you from generating random-like images. It stops you from committing the common conceptual sin: confusing a well-engineered pipeline (with declared assumptions) for an ontological capture of “randomness itself.” It replaces the fantasy of a final certificate with a disciplined separation: build what you can, test what you can, declare what you assumed, and keep the word “truly” as a guardrail against overclaim.

8. Randomness, Meaning, and the Human Pattern Engine

If the earlier sections argued that “truly random image” is ill-posed under positive specification, this section addresses the psychological and cultural fuel that keeps the phrase alive: the human mind is a pattern engine that experiences randomness in a way that is both powerful and unreliable. We do not merely detect patterns; we *need* patterns. Perception and cognition evolved to compress sensory input into actionable summaries. That means we are biased toward finding structure even where none is warranted, and biased toward rejecting real randomness when it fails to conform to our expectations of what “random should look like.”

This is not a side issue. It is one of the main reasons the topic becomes metaphysical. The moment we look at a finite sample, our pattern engine activates, and the artifact becomes entangled with meaning. The apophatic posture, in part, is an ethical response to this entanglement: it refuses to let our projection become an ontological claim. It teaches us to treat the mind’s pattern hunger as a lens that must be acknowledged, not a tribunal that can certify truth.

8.1 Apophenia: why we see structure in any finite sample

Apophenia is the tendency to perceive meaningful patterns in random or meaningless data. In images, it manifests immediately: clusters look like islands; streaks look like motion; gradients look like lighting; two dots look like eyes. Importantly, this is not a glitch. It is the natural output of a system designed to detect faint signals under uncertainty. Missing a true pattern can be costly; seeing a false pattern is often cheaper. So the mind errs on the side of “maybe something is there.”

In a finite random sample, apparent structures are not only possible; they are inevitable. Even if every pixel were generated independently, you will still get clumps, voids, alignments, and near-repetitions. You will get patches that look smoother than average and patches that look noisier than average. You will sometimes get shapes that resemble letters or faces. The mind then does what it always does: it compresses. It says, "This region is special." And the feeling of specialness is experienced as evidence.

But the crucial point is: a finite sample has no internal label that says, "This cluster is meaningful" versus "This cluster is the expected fluctuation of a random field." The mind invents that label by narrative. It treats improbability as message. Yet in a sufficiently large random field, improbable local events are guaranteed to occur somewhere. The surprise arises not from the event itself but from the mind's failure to account for the size of the search space: "If I look long enough, I will find something that looks like something."

This is why "randomness as patternlessness" fails. It asks the mind to do something it cannot do reliably: distinguish "no underlying structure" from "structure too subtle for my detectors" from "structure that is present but coincidental." Apophenia is the reason our intuitive test for randomness is unstable across observers and across contexts.

The apophatic posture treats this honestly: it does not ask perception to certify randomness. It asks perception to confess its bias.

8.2 The "ugliness" of true noise: when randomness looks nonrandom

Many people have an image in their head of what randomness should look like: evenly scattered specks, no big gaps, no strong clusters, no long streaks. That imagined randomness is often not random at all; it is "nicely balanced" disorder—disorder with constraints that prevent conspicuous accidents.

True unconditioned noise violates these expectations. It produces:

- occasional dense clusters,
- occasional conspicuous empty regions,
- accidental lines or arcs,
- repeated motifs by coincidence,
- regions that look "too smooth" or "too rough."

To the human pattern engine, these feel suspicious. They feel like artifacts, like fingerprints, like someone did something. In many cases, nothing happened; the sample simply realized one of the many possible fluctuations that the mind does not intuitively budget for.

This is why true noise can look “ugly.” It does not cooperate with our aesthetic of balanced irregularity. It is indifferent to what we consider acceptable randomness. The “ugliness” is not a property of randomness; it is a collision between randomness and a human expectation that “random” should be decorrelated *and* evenly distributed at every scale. But even distribution at every scale is itself a constraint, and constraints are structure.

In fact, if you reject samples with visible clusters or gaps and keep only those that “look random,” you have performed post-selection. You have turned “true noise” into “noise conditioned on satisfying my aesthetic bias.” You will get a more pleasing texture, but you will have injected information. The thing will no longer be what you claimed.

So the ugliness of true noise is a pedagogical fact: it reveals the difference between randomness and our desire for randomness to look like a certain sanitized ideal.

8.3 Meaning as projection: the boundary between data and interpretation

An image is not only data. For humans, an image is an invitation to meaning. The boundary between data and interpretation is not a line inside the artifact; it is a line in the observer. This is why the same random field can be “nothing” to one person and “a sign” to another. The pixels did not change; the interpretive posture did.

This matters because discussions of “truly random image” often become theological or metaphysical precisely at the point where meaning enters. A person sees a form in noise and asks whether it was “meant.” A person sees a coincidence and asks whether it is “message.” The apophatic stance does not ridicule this. It simply insists on epistemic discipline: meaning may be real as a human experience, but it is not automatically licensed as a property of the data-generating process.

In apophatic terms, the boundary is guarded by refusing to treat projection as ontology. One can say:

- “I experience meaning here,” without also saying,
- “therefore the generator encoded meaning,” or,
- “therefore the artifact is nonrandom.”

This separation is difficult because meaning feels like evidence. But in a high-dimensional space of possible images, meaning-like coincidences are expected to appear occasionally. The apophatic posture does not deny the inner experience. It denies the automatic inference from experience to metaphysical claim.

Thus, the apophatic “truly random image” is also an ethical instruction: do not let your hunger for meaning become a method for certifying claims about the world.

8.4 The temptation to curate: how selection turns entropy into message

Once meaning enters, curation follows. Curation is the act of selecting, editing, or framing an output to serve a narrative. In the context of random images, curation is often subtle and well-intentioned:

- “Let’s pick one that looks more random.”
- “This one has a weird streak—must be a hardware bug.”
- “Let’s crop out that suspicious corner.”
- “Let’s normalize the contrast so viewers can see it better.”
- “Let’s discard the first few outputs while the device warms up.”

Some of these actions may be justified operationally. The problem is not that any adjustment is forbidden; the problem is that adjustments are rarely treated as conditioning events in the probabilistic sense. Selection changes the distribution. It turns entropy into message by imposing a filter that is itself informative.

This is the mechanism by which “randomness” becomes narrative. When you select outputs that conform to your expectation of randomness, you are no longer sampling from the generator’s native distribution; you are sampling from a distribution conditioned on your expectations. The conditioning criterion is a kind of semantic demand: “Give me noise that looks right.” That demand is a form of authorship.

This temptation is especially strong when the random output is meant to persuade. If you are trying to convince someone the generator is random, you will be tempted to show a sample that looks convincingly random. But “convincingly random” is a human category. It is not the same as “random under the generator’s declared assumptions.” The persuasive sample may be the least representative sample.

So the apophatic discipline again functions as an anti-idolatry rule: refuse to turn randomness into a performance. If you must curate for communication (for example, to illustrate typicality), you do it openly: you say what you selected and why, and you avoid implying that the curated output is the essence of randomness.

Synthesis

The human pattern engine ensures that randomness and meaning will always be in tension. Apophenia guarantees that finite samples will display apparent structure. True noise will often look “wrong” by human aesthetic expectations. Meaning will be projected across the boundary

from observer to artifact. And curation will then attempt to stabilize that meaning by selection, quietly transforming entropy into message.

This is not a reason to abandon randomness. It is a reason to handle it apophatically: to reserve "truly" for a posture that refuses to certify essence from appearance, refuses to confuse projection with ontology, and refuses to let curation masquerade as purity.

9. Algorithmic Information and the Apophatic Parallel

The apophatic definition of a "truly random image" might sound theological or poetic until you notice that one of the most rigorous mathematical approaches to randomness already has an apophatic flavor. Algorithmic information theory does not try to define randomness by listing what random objects *are* in a rich positive vocabulary. Instead, it defines randomness primarily by what the object *lacks*: it lacks a short description. It is not compressible. It is not generated by a small program. In other words, it is characterized by a failure of capture.

This is the apophatic parallel: in both cases, the most honest "definition" of randomness is a disciplined refusal of certain claims of knowledge. Yet algorithmic information theory also exposes the limits of this refusal. The moment you attempt to operationalize "no short description," you must fix a description language (a machine, a coding scheme), and the supposedly absolute notion becomes indexed again. Algorithmic randomness therefore provides both a rigorous ally and a rigorous warning. It tells us that negative characterization is powerful, and it tells us that even negative characterization cannot escape structure.

9.1 Incompressibility as a negative characterization

Incompressibility is the cleanest negative way to say "random." Instead of trying to specify the distribution that produced an object, we ask whether the object can be described more briefly than listing it verbatim. If it cannot, then it looks random with respect to that descriptive language.

This aligns with the everyday intuition that patterns are compressions. A checkerboard is compressible: "alternating black and white squares." A gradient is compressible: "linearly increasing intensity." A circle is compressible: "points at fixed radius." These descriptions are far shorter than the raw pixel list. Randomness, in this view, is what refuses such summarization.

The key virtue of incompressibility is that it does not require us to guess what patterns might matter. Instead, it asks a more global question: is there *any* short explanation? This has a natural apophatic feel: the object is defined by the failure of our languages to capture it succinctly.

But incompressibility must be handled carefully. First, it is always relative to a class of descriptions. “Compressible” means “compressible by some allowed method.” Second, in finite objects, incompressibility is not a mystical property; it is a counting fact: most long strings cannot be compressed much, simply because there are not enough short descriptions to cover them all. This is a mathematical way of saying “most objects are unstructured” in the sense of lacking short descriptions.

So incompressibility gives us an honest negative characterization: a random-like image is one that resists description shorter than itself. That is an apophatic-style claim: we do not say what it is; we say what cannot be done with it.

9.2 Kolmogorov complexity as an apophatic metric and its limits

Kolmogorov complexity formalizes incompressibility. Roughly, the Kolmogorov complexity of an object is the length of the shortest program (in some fixed programming language / machine model) that outputs the object and halts. An object is “algorithmically random” if its shortest generating program is about as long as the object itself. That is: there is no hidden recipe.

This is an apophatic metric because it measures *absence*: absence of shorter generative description. It does not define randomness by enumerating features but by ruling out compression.

However, Kolmogorov complexity comes with limits that mirror the paper’s thesis:

- **Language dependence:** Kolmogorov complexity depends on the chosen universal machine (the description language). Different machines change the complexity by at most a constant, but for finite objects that constant can matter. This is the algorithmic version of “no canonical basis.”
- **Uncomputability:** Kolmogorov complexity is not computable in general. There is no universal algorithm that takes an arbitrary object and returns its exact Kolmogorov complexity. This is a profound limit: the apophatic metric cannot be turned into a fully effective certificate. You can sometimes upper-bound it (by giving a program) and lower-bound it in limited ways, but you cannot compute it exactly for arbitrary objects.
- **Finite-sample ambiguity:** For any particular finite image, we can always concoct a short description that singles it out by reference to context (“the image stored in this file at this location”). Algorithmic information theory avoids this by restricting what counts as a valid description language and by treating external references carefully, but the philosophical problem reappears whenever we conflate “a description exists” with “a meaningful intrinsic structure exists.”

So Kolmogorov complexity supports the apophatic approach, but it also warns us: even the most rigorous negative notion of randomness is not absolute in the naive sense. It is indexed to a descriptive regime and cannot be computed into a universal certificate.

This is exactly the paper's posture: negative characterization is powerful, but it does not abolish framing. It makes framing explicit.

9.3 Test classes and "random enough": relativizing the claim honestly

Because Kolmogorov complexity is uncomputable, operational practice replaces "absolute randomness" with "random enough" relative to test classes. This mirrors the earlier discussion of observers and detectors, but algorithmic information theory makes the logic explicit: we choose a class of distinguishers, and we ask whether the output is indistinguishable from the target distribution for that class.

This is the computational analog of apophatic humility:

- We do not claim "truly random."
- We claim "no detectable structure within this declared test class."

The honesty lies in the declaration. If the test class is small (simple statistics, basic correlations), passing is weak evidence. If the test class is broad (many tests, adversarial modeling), passing is stronger evidence. But it is always evidence relative to the chosen lens.

The apophatic parallel is that "random enough" is a negative statement: it asserts a failure of detection, not the presence of an essence. It is essentially a disciplined "not this" claim: not distinguishable by these detectors; not compressible by these compressors; not predictable by these predictors.

Algorithmic information theory therefore supplies a mature vocabulary for what apophasis recommends: make your claims indexed, and resist the temptation to translate "passes my tests" into "is randomness-in-itself."

9.4 The paradox of description: any criterion is itself a structure

Here we meet the deepest convergence between algorithmic information and apophatic theology: the moment you set a criterion for randomness, you have introduced a structure against which objects will be judged. The criterion is a lens. The lens is not neutral. The lens determines what counts as "pattern" and what counts as "noise."

This is the paradox:

- To define randomness, you must define what counts as compressible, predictable, or distinguishable.

- Defining that requires choosing a description language or test class.
- That choice is itself a structure.
- Therefore, the attempt to define randomness introduces exactly the kind of structure it seeks to exclude.

Algorithmic information theory does not resolve the paradox by pretending it is not there. It resolves it by making the dependence explicit and proving robust invariance results (such as "up to a constant" across universal machines) while accepting that absolute, basis-free, criterion-free randomness is not a coherent operational object.

The apophatic stance is, in a sense, the philosophical version of that acceptance. It says: any positive criterion will smuggle structure; therefore the only way to preserve the strongest sense of "truly" is to reserve it for a negative posture that refuses finality. We can still work with criteria—indeed, we must—but we treat them as scaffolding. We do not pretend they are the essence.

Synthesis

Algorithmic information theory shows that the most rigorous talk about randomness already has an apophatic shape: randomness is what refuses compression, refuses short generative description, refuses detection by a declared class. But it also shows that this refusal is never free of framing: it depends on description languages and is blocked by uncomputability from becoming a universal certificate.

So the apophatic parallel is not rhetorical decoration. It is a structural alignment: both domains teach the same discipline. We can speak truthfully about randomness by negation and by indexed claims, while refusing to turn any particular criterion—no matter how elegant—into a final metaphysical declaration of "true randomness."

10. Consequences for AI, Generation, and Epistemic Discipline

The apophatic definition of a "truly random image" is not merely a philosophical ornament. It has direct consequences for how we talk about AI generation, how we distinguish synthetic noise from physical entropy, and how we train ourselves—individually and culturally—to make claims that are both useful and honest. The modern world is saturated with generated artifacts. Many of them look random, and many of them are statistically sophisticated. But the ease with which systems can manufacture "random-like" textures increases the risk of epistemic overreach: mistaking plausibility for truth, indistinguishability for essence, and "passes my test" for "is what I say it is."

This section draws out four consequences. First, generative models cannot produce "truly random" in the apophatic sense because their outputs are always products of learned structure and controlled sampling procedures. Second, synthetic noise and physical entropy differ in provenance even when they can be made visually or statistically similar. Third, the apophatic posture demands an audit culture: label assumptions, choices, and tests. Fourth, it proposes a discipline for language—when to say "random," "pseudorandom," and "unknown"—as a practical ethic of claims.

10.1 Why generative models cannot produce "truly random" in the apophatic sense

Generative models—diffusion models, GANs, autoregressive models, latent-variable models—can produce images that look noisy, texture-like, or "random." They can also produce outputs whose pixel statistics resemble those of a target distribution. But "truly random" in the apophatic sense is not primarily about visual appearance or even about matching a distribution. It is about refusing privileged formalisms, refusing final measures, refusing compressible capture, refusing reproducible recipes, and refusing curation.

Generative models fail this apophatic criterion in multiple ways:

- **They are structured by design.** A generative model is a compressed representation of a training distribution plus an inference procedure. Even if it can emit noise-like outputs, the space of outputs is shaped by learned structure. The model is itself a short description of a large family of images. That is the opposite of apophatic "not compressible" when treated as a claim of essence.
- **They privilege a latent or internal representation.** Whether explicit or implicit, these models define a coordinate system in which sampling is performed. The sampling procedure is coordinate-dependent. That violates "not basis-privileged" if you try to interpret the output as "truly random in itself."
- **They rely on controllable, reproducible recipes.** A model plus a seed plus a sampling algorithm reproduces outputs. Reproducibility is a feature for engineering, but it is also a handle. In apophatic terms, a handle is structure. If you can regenerate the "same random image" by reusing a seed, you have demonstrated that the artifact belongs to a family indexed by a short description.
- **They inevitably import a measure.** Sampling from a model is sampling from a distribution defined by the model and the sampler. That distribution may be complex, but it is still a declared (or at least definable) measure. The apophatic stance refuses to treat any measure as final.

This is not a criticism of generative modeling. It is a boundary condition on metaphysical language. AI can generate outputs that are *random-looking*, *stochastically sampled*, *hard to distinguish from certain noise models*, or *sufficiently unpredictable for some adversary model*. But it cannot, by its nature as a compressed learned mechanism, deliver "truly random" in the apophatic sense—because apophatic "truly random" is defined precisely by refusing the kinds of positive handles that models necessarily provide.

The practical upshot is linguistic: we should stop treating "AI-generated randomness" as an ontological claim. It is an operational claim at best: "noise-like under these tests," "sampled stochastically from this model," "unpredictable to this class of observers." Anything stronger is overreach.

10.2 Synthetic noise versus physical entropy: the provenance distinction

A second consequence is the restoration of provenance as a first-class distinction. In many contexts, synthetic noise and physical entropy can be made to look similar. Both can be whitened. Both can pass statistical tests. Both can be used as dithering, as simulation input, or as aesthetic texture. But they remain different in kind with respect to provenance:

- **Synthetic noise** is produced by a deterministic or algorithmically specified procedure. Even if it is statistically excellent and indistinguishable to bounded observers, its existence is fully accounted for by the algorithm and the seed. It is, in principle, replayable. Its unpredictability is conditional: conditional on ignorance of the seed or on bounded computational resources.
- **Physical entropy** is harvested from interaction with the world. Even if its digitization pipeline introduces scaffolding and even if its unpredictability is not metaphysically settled for all philosophical positions, it is not merely the unfolding of a fixed program. It is, at minimum, an empirical coupling to uncontrolled degrees of freedom, and in many cases it is treated as ontically indeterminate.

Apophatic discipline insists that this distinction not be erased by appearances. If two outputs are visually identical, it does not follow that their epistemic status is identical. Provenance is part of what we mean by "random" in ordinary life ("not faked"), and apophatic honesty preserves that intuition without turning it into metaphysical certainty.

The key is: provenance is not essence either. Physical entropy does not magically produce a basis-free "truly random image." But it does provide a different kind of claim: "this artifact is anchored in physical unpredictability rather than algorithmic replay." In many domains—security, fairness, auditability—that difference matters.

So we should speak in two layers:

- **Statistical layer:** what patterns are detectable under tests.
- **Provenance layer:** what kind of process generated the artifact.

Conflating these layers is one of the main sources of confusion in “randomness” discourse.

10.3 Auditability and humility: labeling what is assumed, chosen, and tested

If we accept that neither AI nor physics yields an effortless certificate of “true randomness,” then the practical ethical response is auditability: make the scaffolding visible and contestable. Apophatic humility is not the refusal to act; it is the refusal to pretend.

Auditability here means that whenever we present a “random image” (especially if we are using it as evidence or as a trust anchor), we should be able to answer four questions cleanly:

- **What was assumed?**
  The target sample space, the constraints, the invariances, the meaning of “random” being used.
- **What was chosen?**
  The representation (pixel grid, basis), the mapping from entropy to image channels, the parameters (resolution, bit depth), any whitening or conditioning steps.
- **What was measured/tested?**
  The test suite: bias checks, correlation checks, spectral checks, compressibility checks, stationarity checks. Also what the tests are blind to.
- **What was withheld or curated?**
  Whether outputs were selected, discarded, cropped, normalized, or otherwise filtered, and why.

Humility means that we treat the answers not as paperwork but as the substance of the claim. In apophatic terms, the audit is the way we prevent the scaffolding from becoming an idol. We can still say “random” in a bounded sense, but we do so with labeled commitments rather than with rhetorical finality.

This has broader implications for AI culture. The same discipline should apply whenever AI outputs are presented as knowledge: the system’s “randomness” (its stochastic sampling) is not the same as truth, and auditability is what keeps the difference visible.

10.4 A discipline for claims: when to say “random,” “pseudorandom,” “unknown”

Finally, the apophatic approach proposes a practical vocabulary discipline—an ethics of naming. The goal is not pedantry; it is preventing category errors that erode trust.

A workable discipline might look like this:

- **Say “random”** when you mean:
  “Random relative to declared assumptions and tests,” and you are willing to state those assumptions and tests. This includes physical-entropy-based outputs and even synthetic ones in contexts where the meaning is explicitly statistical and bounded.

Example posture: “This image is sampled from this declared distribution (or process), and under these diagnostics we find no detectable structure of these kinds.”

- **Say “pseudorandom”** when you mean:
  “Generated by an algorithmic procedure whose outputs are intended to be indistinguishable from random to a declared class of observers, but whose structure is in principle fully determined by seed and algorithm.”

This preserves the crucial provenance difference without implying the outputs are “fake” in a useless sense. Pseudorandom can be excellent. The point is to name the handle.

- **Say “unknown”** when you mean:
  “We do not have sufficient basis to classify this reliably,” either because provenance is unclear, tests are insufficient, or the environment is contaminated by confounds (device artifacts, selection, transformations).

“Unknown” is a virtue word in apophatic epistemology. It prevents the pressure to label from turning into confident guessing.

- **Reserve “truly random”** (if used at all) for the apophatic sense:
  not as a certificate, but as a reminder that any positive claim is indexed and any final claim is suspect. “Truly” becomes a guardrail word, not a triumph word.

This linguistic discipline is the practical heart of the paper. It converts an abstract apophatic thesis into a daily epistemic habit: label your lenses, confess your scaffolding, do not inflate your tests into metaphysical proof, and preserve the provenance distinctions that matter for trust.

Synthesis

For AI, the apophatic framework draws a bright line: generative models can produce stochastic outputs and noise-like textures, but they cannot produce “truly random” in the apophatic sense because they are intrinsically structured and recipe-governed. For practice, the framework

restores provenance as a crucial distinction: synthetic noise and physical entropy are not interchangeable even when they look the same. For culture, it demands auditability: declare assumptions, choices, tests, and curation. And for language, it offers an ethics of naming: use "random," "pseudorandom," and "unknown" carefully, and treat "truly" not as a certificate but as a warning against overclaim.

That combination—provenance clarity, audit posture, and disciplined naming—is the paper's proposed contribution to epistemic hygiene in an age where both humans and machines can fabricate "random-looking" artifacts at will.

11. Objections and Replies

A proposal like "apophatic randomness" invites predictable objections, and that is a good sign. If the idea is doing real work, it will press against habits of speech that people have relied on without noticing the assumptions embedded in them. The replies below are not meant to win rhetorical points; they are meant to keep the central distinction intact: the difference between *operational randomness claims* (which are useful and often rigorous) and *metaphysical completion claims* (which quietly pretend that one representation, one measure, or one test suite exhausts what randomness "truly is").

11.1 "This makes randomness meaningless." Response: it makes the claim honest

The objection usually means: "If you refuse to define randomness positively and finally, then you have removed the concept's power. Everything becomes 'relative,' and nothing can be said."

The reply is that the paper does not remove meaning; it removes false simplicity. Randomness already *is* relational in every rigorous use: it is always "random with respect to a distribution," "unpredictable to a class of observers," "indistinguishable under tests," or "incompressible relative to a description language." The apophatic move does not invent this relativity; it makes it explicit and prevents us from smuggling it in unacknowledged.

In practice, the concept retains plenty of bite:

- You can still say a generator is biased or unbiased relative to declared diagnostics.
- You can still say an artifact shows correlations, periodicities, compressibility, drift.
- You can still compare generators and pipelines.
- You can still make security claims *relative to an adversary model*.
- You can still do science.

What you cannot do—what the apophatic stance refuses—is to take a successful operational story ("passed tests," "maximum entropy," "quantum source") and inflate it into a metaphysical stamp: "therefore this object is truly random in itself." That inflation is not meaning; it is overclaim.

So the concept of randomness becomes *more honest*, not less meaningful. It becomes a disciplined statement that names its frame instead of pretending no frame exists. If anything, that discipline protects the concept from becoming meaningless, because many "randomness" disputes are actually disputes about hidden assumptions. Apophasis forces those assumptions into the open.

11.2 "Physics provides the measure." Response: measurement still encodes choices

This objection says: "You are overcomplicating. The physical world supplies the distribution. If I measure a quantum process, the Born rule gives me probabilities. If I measure thermal noise, physics gives me a noise model. So the measure is not arbitrary; it is given."

The reply is twofold.

First, even if physics supplies a distribution for some *measurement outcomes*, it does not automatically supply a distribution for "images" in any basis-free sense. An "image" is a constructed artifact: you choose a sensor, choose a sampling grid, choose quantization, choose a color encoding, choose a mapping from events to pixel intensities, choose post-processing for display. Those choices induce the distribution over images. Physics may constrain what is possible, but the distribution over *encoded image objects* is shaped by the measurement protocol and representation decisions.

Second, measurement itself is never purely "given." It encodes choices:

- what observable you measure,
- how you threshold and digitize,
- what you treat as noise versus signal,
- how you calibrate and correct drift,
- how you filter and whiten,
- how you define the sample window,
- what you discard as "bad runs."

All of those decisions are the difference between “a physical process happened” and “an artifact was produced and presented.” Even in the quantum case, the mapping from click times to bits to pixel values is not dictated by the Born rule; it is imposed by the experimenter.

So physics can anchor unpredictability and can constrain measures in principled ways, but it does not abolish the representational and procedural scaffolding that the apophatic critique targets. The measure is not simply “given”; it is “given relative to a measurement and encoding protocol.” That is exactly the relational point.

11.3 “All science needs priors.” Response: apophasis is about confessing them

This objection is friendly: “Yes, priors are unavoidable. Science is built on modeling choices. So what are we doing here? Isn’t this just standard practice?”

The reply is: precisely. The paper is not claiming priors are a scandal. It is claiming that priors become dangerous when they are *hidden* and then treated as if they were dictated by the target itself. In the randomness case, the danger is especially acute because “random” often functions rhetorically as “neutral,” “unbiased,” “natural,” “unconstructed.” Those words tempt us to forget that every operational definition of randomness is a choice of representation and measure.

Apophasis, here, is the discipline of confession:

- confess the basis,
- confess the distribution or constraints,
- confess the test class,
- confess the pipeline and selection steps,
- confess the limits of what you can certify.

This is not an attack on science. It is a defense of scientific integrity against a specific rhetorical failure mode: the slide from “this is a useful model” to “this is what the thing is.”

In other words, apophasis is the opposite of anti-scientific mystique. It is the refusal of counterfeit certainty. It keeps modeling honest by keeping scaffolding visible.

11.4 “Apophatic is theology, not math.” Response: it’s also an epistemic method

This objection says: “You are importing theological vocabulary into a technical topic. That’s category confusion.”

The reply is that apophasis is not being used here to do theology *instead of* math. It is being used as a name for a general epistemic method that shows up in many domains: when a target

cannot be captured without distortion by positive predicates, negative characterization and humility about frames can be more truthful than premature closure.

Mathematics itself contains apophatic moments:

- In algorithmic information theory, randomness is characterized by *incompressibility*—a negative property.
- In logic, certain truths are shown to resist complete formal capture (limitations results), and the honest response is to mark boundaries rather than pretend completeness.
- In measurement theory, "no canonical measure" in certain spaces is not a poetic claim; it is a structural fact.
- In statistics, identifiability limits are not moral failures; they are places where the honest statement is "not determinable under these assumptions."

Calling the posture "apophatic" is not a claim that the topic is theological. It is a claim that the *virtue* apophatic theology cultivated—refusing to mistake finite concepts for infinite reality—has a precise analog in epistemology: refuse to mistake a chosen formalism for an intrinsic essence. The vocabulary is borrowed because it names the posture cleanly: negation as a guardrail against idolatry of models.

If the theological resonance is unwelcome, the same idea can be stated in secular terms: "epistemic humility under representational non-uniqueness." But "apophatic" has the advantage of being memorable and morally charged in the right direction: it warns against worshiping our own constructions.

Synthesis

These objections converge on a single point. The apophatic stance does not abolish randomness; it prevents overclaim. It does not deny physics; it denies that measurement removes representation choices. It does not reject priors; it demands their confession. And it does not replace math with theology; it names an epistemic method already implicit in rigorous treatments of randomness and description. The payoff is not metaphysical victory. The payoff is cleaner language, fewer category errors, and a more auditable relationship between what we build and what we claim.

12. Conclusion: The Apophatic Image as a Guardrail

The phrase "truly random image" tempts us because it promises a kind of purity: an artifact free of intention, free of bias, free of framing—an image that simply *is* randomness, as if

randomness were a substance that could be bottled. The argument of this paper has been that this promise cannot be met by positive specification without contradiction. Every constructive route—pixels, Fourier coefficients, maximum entropy, symmetry constraints, test batteries, even physical entropy—requires a frame. The frame is not an accidental bolt-on; it is what makes the phrase coherent at all. The moment we try to say what the "truly random image" is, we have already chosen what counts as an image, how it is represented, what outcomes are allowed, which outcomes are typical, and which kinds of structure we consider disqualifying. Positivity, in other words, does not merely describe randomness; it imports priors and then forgets it imported them.

The apophatic turn is the guardrail against that forgetting. It does not claim that randomness is unreachable or meaningless. It claims that the strongest sense of "truly" cannot be served by a final positive certificate. If we insist on keeping "truly," then "truly random image" must be treated as a limit-concept: a name for what no single privileged formalism can exhaust, and a discipline that prevents our scaffolds from becoming idols.

12.1 Restate the thesis: positivity smuggles priors; apophasis preserves humility

The thesis can now be stated compactly.

First: *positivity smuggles priors.* To define a random image positively is to pick a basis (pixels, waves, latent variables), pick a measure (uniformity somewhere, maximum entropy under constraints), pick an invariance set (translation, rotation, scale), pick a sampling pipeline, and pick a test class. None of these are illegitimate; they are the price of operational clarity. But they are also the sources of hidden commitments. They define what kinds of structure will be visible, what kinds of structure will be invisible, and what "random" will mean in that chosen coordinate chart. When we then speak as if our definition is basis-free or final, we have converted a pragmatic construction into a metaphysical claim.

Second: *apophasis preserves humility.* An apophatic definition does not try to name the essence of the "truly random image" positively. It names what we refuse to pretend: no canonical basis, no final measure, no privileged compression language, no reproducible recipe that counts as essence, no curation that conditions entropy into message. The apophatic image is not a trophy; it is a warning label. It exists to stop the slide from "this is a useful model" to "this is what randomness is."

In that sense, the apophatic posture is a form of rigor. It is not the abdication of precision; it is precision about what precision can legitimately claim.

12.2 Practical takeaway: operational approximations with explicit scaffolding

The practical outcome is not paralysis but posture.

We can still generate random-like images. We can still anchor unpredictability in physical entropy. We can still choose representations and constraints and sampling methods. We can still run tests and report results. But we must speak with explicit scaffolding:

- state the representation and why it was chosen,
- state the mapping from entropy to image values,
- state the distributional assumptions or constraints,
- state the test class and its blind spots,
- state whether any selection, filtering, normalization, or cropping occurred,
- state what claim is warranted (and what is not).

This transforms "randomness" from an incantation into an accountable claim. The key is the separation of construction from claim. Construction can be elaborate and competent; claim should remain indexed and modest. "Random relative to these assumptions and tests" is a strong, useful statement. "Truly random, full stop" is usually a metaphysical overreach disguised as confidence.

In an age of AI generation, this matters even more. Synthetic noise can be made to look convincing. Pseudorandomness can be indistinguishable to many detectors. The apophatic guardrail therefore becomes a broader epistemic discipline: do not let appearance or test-passing become essence; keep provenance, assumptions, and constraints visible.

12.3 Closing synthesis: refusing idols of specification while retaining rigor

The final synthesis is simple: the apophatic image is a way to refuse idols of specification without retreating into vagueness.

The idol is the belief that a clean formula—"i.i.d. pixels," "flat spectrum," "maximum entropy," "quantum source," "passes tests"—can deliver randomness-in-itself. The refusal is the insistence that every such formula is a coordinate chart, a measure, a constraint set, or a test lens. The rigor is the willingness to use these tools while explicitly labeling them as tools, not as ontological completions.

So the "truly random image," defined apophatically, is not an object we possess. It is a guardrail that protects our language from overclaim. It keeps the word "truly" from becoming a boast and turns it into a confession: whatever we build will be scaffolding; whatever we certify will be

indexed; whatever we experience as meaning will be recognized as projection unless independently warranted.

That guardrail is not a retreat from truth. It is a way of remaining faithful to truth in a domain where the temptation to claim completion is strong and the mechanisms of hidden prior-injection are subtle. The result is a disciplined freedom: we can build operational randomness, we can test and compare and deploy it, and we can do so without worshiping our own coordinate choices—retaining rigor while refusing the false comfort of final specification.

References

Bassham, L. E., Rukhin, A. L., Soto, J., Nechvatal, J. R., Smid, M. E., Leigh, S. D., Levenson, M., Vangel, M., Heckert, N. A., & Banks, D. L. (2010). *A Statistical Test Suite for Random and Pseudorandom Number Generators for Cryptographic Applications* (NIST Special Publication 800-22 Rev. 1a). National Institute of Standards and Technology. NIST Computer Security Resource Center+2NIST Publications+2

Chaitin, G. J. (1969). On the length of programs for computing finite binary sequences. *Journal of the ACM*, 16(1), 145–159.

Cover, T. M., & Thomas, J. A. (2006). *Elements of Information Theory* (2nd ed.). Wiley.

Conrad, K. (1958). *Die beginnende Schizophrenie: Versuch einer Gestaltanalyse des Wahns*. Thieme.

Diaconis, P., & Mosteller, F. (1989). Methods for studying coincidences. *Journal of the American Statistical Association*, 84(408), 853–861.

Falk, R., & Konold, C. (1997). Making sense of randomness: Implicit encoding as a basis for judgment. *Psychological Review*, 104(2), 301–318.

Jaynes, E. T. (1957). Information theory and statistical mechanics. *Physical Review*, 106(4), 620–630. APS Link

Knuth, D. E. (1997). *The Art of Computer Programming, Volume 2: Seminumerical Algorithms* (3rd ed.). Addison-Wesley.

Kolmogorov, A. N. (1965). Three approaches to the quantitative definition of information. *Problems of Information Transmission*, 1(1), 1–7.

Li, M., & Vitányi, P. M. B. (2019). *An Introduction to Kolmogorov Complexity and Its Applications* (4th ed.). Springer. Springer Nature Link+1

MacKay, D. J. C. (2003). *Information Theory, Inference, and Learning Algorithms*. Cambridge University Press.

Martin-Löf, P. (1966). The definition of random sequences. *Information and Control*, 9(6), 602–619. ScienceDirect+1

Shannon, C. E. (1948). A mathematical theory of communication. *Bell System Technical Journal*, 27, 379–423, 623–656.

Solomonoff, R. J. (1964). A formal theory of inductive inference. Part I & Part II. *Information and Control*, 7(1), 1–22; 7(2), 224–254.

Tversky, A., & Kahneman, D. (1971). Belief in the law of small numbers. *Psychological Bulletin*, 76(2), 105–110.

Pseudo-Dionysius the Areopagite. (c. late 5th–early 6th century). *The Mystical Theology* (in *The Complete Works*, trans. C. Luibheid). Paulist Press, 1987.

Gregory of Nyssa. (c. 4th century). *The Life of Moses* (trans. A. Malherbe & E. Ferguson). Paulist Press, 1978.

Lossky, V. (1957). *The Mystical Theology of the Eastern Church*. James Clarke.

Turner, D. (1995). *The Darkness of God: Negativity in Christian Mysticism*. Cambridge University Press.