

What is claimed is:

1. A system for computable content identity, comprising:

- a content encoder configured to derive a unique identifier for a digital content artifact by normalizing the digital content artifact to a canonical scalar field in a modality-specific manner and extracting a multi-axis variance vector from the canonical scalar field, and encoding the multi-axis variance vector as a deterministic, slope-indexed fingerprint that encodes a position within a continuous variance space such that cosine similarity between any two unique identifiers is directly computable from their respective variance vectors without decoding a fixed binary digest, wherein the multi-axis variance vector comprises a first axis encoding energy distribution behavior of the artifact across a plurality of spatial scales, a second axis encoding frequency compaction behavior of the artifact across the plurality of spatial scales, and a third axis encoding structural phase persistence of the artifact based on gradient orientation distribution, and wherein the content encoder is operable on artifacts of one or more modalities including raster images, audio waveforms represented as time-frequency scalar fields, textual documents represented as token frequency scalar fields, video frames, and binary objects represented as normalized scalar fields;
- a plurality of anchor nodes, each anchor node scoped to at least one variance band within a global slope continuum, wherein each anchor node stores and governs unique identifiers whose variance-derived slope vectors fall within the variance band assigned to the anchor node;
- a cache memory associated with each anchor node configured to store semantic lineage graphs, alias registrations, mutation records, and policy constraints for content artifacts anchored within the anchor node's band scope;
- an alias resolution engine configured to map human-readable identifiers to variance-derived unique identifiers under cryptographically signed policy constraints enforced by anchor quorum consensus; and
- a provenance validator configured to construct and traverse multi-root lineage graphs by computing slope vector proximity and mutation deltas between a candidate unique identifier and one or more previously anchored unique identifiers.

2. A computer-implemented method for computable content identity, comprising:
  - normalizing a digital content artifact to a canonical scalar field representation, wherein the normalization is modality-specific and comprises, for raster image artifacts, resizing to a canonical square canvas and converting to a luma-weighted grayscale floating-point field; for audio waveform artifacts, computing a time-frequency spectrogram normalized to a canonical resolution; for textual document artifacts, mapping token frequency distributions to a two-dimensional positional grid; and for video artifacts, extracting per-frame scalar fields and computing a temporal delta vector encoding inter-frame variance change rates;
  - extracting a multi-axis variance vector from the canonical scalar field, the multi-axis variance vector comprising a first axis encoding multi-scale energy distribution computed as cross-scale trends in per-cell variance across a plurality of nested grid resolutions, a second axis encoding multi-scale frequency compaction computed as cross-scale trends in per-cell variance standard deviation, and a third axis encoding structural phase persistence computed from a rotation-canonicalized gradient orientation histogram;
  - deriving a unique identifier by computing a weighted linear combination of variance vector components organized into a plurality of axis-dominant triads and hashing the combined vector using a multi-scale quantization scheme;
  - quantizing a global variance value derived from the canonical scalar field into a slope band representing a bounded range within a global variance space, and registering the unique identifier with one or more anchor nodes governing the slope band;
  - constructing a multi-root lineage graph by computing cosine similarity between the variance vector of the registered unique identifier and variance vectors of one or more previously anchored unique identifiers, linking the registered unique identifier to each prior unique identifier whose cosine similarity exceeds a configured semantic continuity threshold, and weighting each lineage edge proportional to the computed cosine similarity; and
  - enforcing, through anchor quorum consensus, cache propagation policies, alias resolution policies, and cryptographically signed delegation constraints governing

mutation eligibility, propagation scope, and temporal validity for the registered unique identifier,

wherein unique identifier derivation, slope band assignment, lineage graph construction, and policy enforcement are performed without reliance on centralized registries, static network addressing, or persistent external credentials.

3. A rights-grade content admissibility platform, comprising:

a pre-release admissibility engine configured to evaluate a candidate content artifact against one or more cryptographically signed policy objects prior to external commitment of the artifact, wherein each policy object defines one or more of: admissible content categories, restricted content classes, jurisdictional constraints, similarity tolerance thresholds, override authorities, and escalation paths, and wherein the candidate artifact is prevented from becoming a committed artifact unless all applicable policy conditions are satisfied;

a structural similarity evaluator configured to compute, prior to commitment, a similarity score between the candidate content artifact and one or more reference artifacts in a governed corpus by comparing variance-derived unique identifiers of the candidate artifact and each reference artifact, wherein similarity is determined by cosine similarity of multi-axis variance vectors and wherein the candidate artifact is rejected, regenerated, or escalated if similarity to any reference artifact exceeds a policy-declared threshold;

a training corpus governance layer configured to admit digital artifacts to a generative model training corpus only under signed, declared corpus policy objects, the admitted artifacts collectively constituting an admissible corpus, and to record a cryptographically verifiable lineage linking trained model artifacts to the admissible corpus from which they were trained; and

a consultation event logger configured to deterministically record, for each generation event that consults a reference artifact through retrieval or structured neighborhood resolution, a consultation record comprising the variance-derived unique identifier of the consulted artifact, an identifier of the policy object, of the one or more cryptographically signed policy objects, under which the consultation was authorized, and a timestamp,

wherein the platform renders structurally impermissible content non-executable prior to artifact commitment.

4. The system of claim 1, wherein the content encoder further extracts a quadrant decomposition fingerprint by dividing the digital content artifact into four spatial sub-regions following canonical normalization and orientation canonicalization, computing an independent variance vector for each sub-region using a coarser quantization scale than the global hash, sorting the resulting quadrant hashes in lexicographic order to produce a rotation-invariant spatial fingerprint, and combining the sorted quadrant hashes with the global variance hash using a multi-segment hash combiner to produce a 320-bit unique identifier.
5. The system of claim 1, wherein the content encoder further extracts a gradient-only structure signature invariant to mean luminance and background fill by computing a rotation-canonicalized sixteen-bin gradient orientation histogram, edge density values at three gradient magnitude thresholds, and histogram variance and peak statistics, and wherein the content encoder further extracts a constellation signature invariant to translation, rotation, and uniform scale change by identifying spatial anchor points whose saliency score, computed as a weighted sum of intra-cell intensity variance and mean local gradient density, exceeds a configured detection threshold, computing normalized inter-anchor distances and rotation-invariant inter-anchor angles, quantizing distances and angles into discrete bins, hashing per-anchor descriptors, and aggregating the hashes in lexicographic order.
6. The system of claim 1, wherein the content encoder, when processing an audio waveform artifact, derives the canonical scalar field from a short-time Fourier transform or mel-spectrogram of the waveform normalized to a canonical time-frequency grid resolution, and wherein the multi-axis variance vector further encodes spectral centroid distribution as a supplementary third-axis component.
7. The system of claim 1, wherein the content encoder, when processing a video artifact, computes a per-frame unique identifier for each frame of the video artifact using the multi-axis variance vector extraction pipeline, derives a temporal delta vector encoding a sequence of inter-frame cosine similarity scores between consecutive frame unique

identifiers, and derives a clip-level unique identifier from the temporal delta vector and a sequence of per-frame unique identifiers.

8. The system of claim 1, wherein the alias resolution engine enforces alias registration constraints using cryptographically signed policy objects specifying delegation rights, propagation scope, temporal validity, and lineage depth limits, and wherein the alias resolution engine initiates cross-band pathfinding when variance drift causes a unique identifier to migrate to a slope band different from its original registration band, traversing band adjacency edges in order of variance proximity until a governing anchor cluster is identified.

9. The system of claim 1, wherein the cache memory is dynamically pruned, retained, or refreshed based on local variance saturation levels, semantic traffic volume, and policy-declared retention priority, and wherein anchor nodes replicate unique identifier registration data across adjacent slope bands based on variance proximity thresholds and mutation propagation rates under an adaptive consensus protocol.

10. The method of claim 2, wherein normalizing the digital content artifact to the canonical scalar field further comprises applying an orientation canonicalization procedure by computing the dominant gradient orientation of the normalized field and rotating the field to align the dominant orientation with a canonical reference angle, reducing unique identifier drift attributable to rotational variation across format conversions and resolution changes.

11. The method of claim 2, wherein extracting the multi-axis variance vector comprises subdividing the canonical scalar field into a coarse grid of eight-by-eight cells, a medium grid of sixteen-by-sixteen cells, and a fine grid of thirty-two-by-thirty-two cells, computing a variance-based proxy for each cell at each grid resolution, deriving components of the first axis from the slope, curvature, and asymptotic value of cross-scale mean variance, and deriving components of the second axis from the compaction rate, spread factor, and variance floor convergence of cross-scale variance standard deviation.

12. The method of claim 2, wherein deriving the unique identifier comprises forming a 27-dimensional vector as three axis-dominant triads, wherein a first triad weights the first axis at 0.6 and the second and third axes each at 0.2, a second triad weights the second axis at 0.5 and the first and third axes at 0.3 and 0.2 respectively, and a third triad weights

the third axis at 0.5 and the first and second axes at 0.3 and 0.2 respectively, and hashing the 27-dimensional vector under multiple overlapping quantization scales using an FNV-variant hash combiner.

13. The method of claim 2, further comprising, when processing a textual document artifact, representing the document as a two-dimensional token frequency scalar field by mapping each token to a positional grid cell weighted by term frequency-inverse document frequency score, and computing byte-level variance across sliding windows of the document as a supplementary variance signal for the third axis of the multi-axis variance vector.

14. The method of claim 2, further comprising enforcing a pre-release admissibility check prior to external commitment of the content artifact by evaluating the variance-derived unique identifier against a governed corpus of reference unique identifiers, computing cosine similarity between a variance vector of the digital content artifact and each reference variance vector, and blocking commitment if any similarity score exceeds a policy-declared threshold, wherein the admissibility check is deterministic, reproducible from versioned policy objects, and executes without reliance on centralized moderation services.

15. The method of claim 2, wherein constructing the multi-root lineage graph further comprises detecting adversarial recombination by verifying that the variance vector of a declared composite artifact is consistent with the weighted variance combination of its declared parent artifacts within a configured tolerance, and flagging as potentially adversarial any composite artifact whose variance profile exhibits discontinuity from the weighted combination of its declared parents beyond the configured tolerance.

16. The platform of claim 3, wherein the structural similarity evaluator evaluates similarity prior to commitment by computing, for the candidate content artifact, a multi-axis variance vector comprising energy distribution, frequency compaction, and structural phase persistence components, and comparing the multi-axis variance vector against variance vectors of reference artifacts indexed in a slope-band-partitioned anchor network, such that similarity evaluation operates without GPU inference, centralized embedding indexes, or per-query API calls.

17. The platform of claim 3, wherein the training corpus governance layer records, for each artifact admitted to the training corpus, a governance record comprising the artifact's variance-derived unique identifier, the signed corpus policy object under which it was admitted, a timestamp, and a cryptographic hash of the policy object, and wherein the lineage linking trained model artifacts to the admissible corpus is verifiable by comparing a set of governance records against the variance-derived unique identifiers of content artifacts presented as potential training data sources.

18. The platform of claim 3, wherein the consultation event logger records consultation events arising from retrieval-augmented generation operations in which a generative model queries a reference corpus by variance-band-indexed neighborhood resolution, and wherein recorded consultation events are aggregated to compute attribution weights for each consulted reference artifact based on consultation frequency, variance proximity of the consulted reference artifact to the candidate content artifact, and policy-declared compensation schedules.

19. The platform of claim 3, further comprising a forbidden content exclusion layer configured to evaluate each candidate content artifact against a policy-defined set of restricted semantic categories by comparing the variance-derived unique identifier and structural signatures of the candidate artifact against a governed exclusion corpus of variance-indexed forbidden content references, wherein a candidate artifact whose variance vector falls within a configured proximity of any forbidden content reference is rendered non-committable prior to release.

20. The platform of claim 3, wherein each cryptographically signed policy object is versioned, machine-evaluable, and comprises typed category constraints, jurisdictional scopes, override authorities, similarity tolerance thresholds, and escalation paths, and wherein admissibility decisions executed under a given policy object version are reproducible and auditable by re-executing the given policy object version against the variance-derived unique identifier and structural signatures of the candidate content artifact under evaluation, and comparing the re-executed admissibility decision to a recorded original admissibility decision.