STRUCTURAL CONTENT IDENTITY AND RIGHTS-GRADE ADMISSIBILITY FOR DIGITAL ARTIFACTS

RELATED APPLICATION DATA

[0001] This application claims the benefit of priority of U.S. Provisional Patent Application Serial No. 63/808,372, filed on May 19, 2025, titled "Entropy-Based Content Anchoring System with Slope-Governed Routing, Adaptive Caching, and Policy-Enforced Alias Resolution", which is incorporated by reference herein in its entirety.

FIELD

[0002] The present disclosure generally relates to content identity systems, distributed indexing, and rights governance for digital artifacts. In particular, the present disclosure is directed to structurally-derived content identity and rights-grade admissibility for digital artifacts, and methods thereof.

BACKGROUND

[0003] Conventional content delivery and asset indexing systems reference digital assets by static identifiers such as uniform resource locators, cryptographic hash pointers, or file-system paths derived from storage location or transmission metadata rather than from the internal structure of the content. Such identifiers are invalidated by mutation, format conversion, resolution change, lossy compression, or replication, producing identity fragmentation across versions and derivatives. Content distribution networks compound this limitation by routing on geographic or address-based heuristics that are structurally unaware of payload semantics and provide no mechanism for derivative attribution, remix lineage, or variance-based similarity across stored objects.

[0004] Existing perceptual hashing systems, including difference hash, average hash, and perceptual hash algorithms, produce low-dimensional binary signatures from downsampled image representations. They lack multi-scale structural analysis, directional orientation decomposition, spatial sub-region identity, and a continuously scaled similarity score suitable for slope-based banding or lineage tracing, and their fixed-width binary outputs cannot encode the gradient structure, variance flow, or compositional geometry required for distributed anchor assignment or derivative attribution. Watermarking and metadata tagging approaches embed

identity signals in the content stream or a sidecar record; watermarks are removable through transcoding, cropping, or generative reconstruction, and metadata records are decoupled from content structure and require persistent external storage. Neither approach supports slope-indexed distribution, variance-band routing, or governance-enforced propagation.

[0005] Blockchain-based asset registration systems anchor ownership records to distributed ledgers through hash-based proofs of existence, requiring global consensus, imposing transaction costs, and binding identity to key-pair ownership rather than content structure; they cannot detect semantic similarity between related or derivative content objects and provide no mechanism for alias resolution, composite attribution, or variance-band-scoped caching. Conventional caching frameworks rely on time-to-live heuristics, frequency-of-access counters, or manual invalidation signals that are indifferent to the structural variance of cached content, the mutation distance between cached and current versions, or governance constraints on propagation scope.

[0006] Generative artificial intelligence systems evaluate content admissibility through post-generation moderation filters applied after an output artifact has been produced, which cannot prevent impermissible content from existing as an internal artifact and do not provide reproducible, auditable admissibility decisions verifiable from versioned policy records. No existing system interposes a structural pre-release admissibility evaluation between content generation and external commitment using variance-derived content identity, governs training corpus admission through cryptographically verifiable lineage, or determines whether a generative model output is structurally proximate to specific training-data artifacts without requiring white-box or gray-box access to the model.

[0007] Accordingly, there is a need for systems and methods that address these shortcomings.

SUMMARY OF THE DISCLOSURE

[0008] In an embodiment, a system for computable content identity comprises a content encoder configured to derive a unique identifier for a digital content artifact by extracting a multi-axis variance vector from the internal structure of the artifact, the multi-axis variance vector comprising a first axis encoding cross-scale energy distribution, a second axis encoding cross-scale frequency compaction, and a third axis encoding structural phase persistence based

on gradient orientation distribution, the unique identifier encoding a position within a continuous variance space such that cosine similarity between two unique identifiers is directly computable without decoding a fixed binary digest, the content encoder operable across raster images, audio waveforms, textual documents, video frames, and binary objects represented as normalized scalar fields. The system may further comprise one or more anchor nodes each scoped to at least one variance band within a global slope continuum, a cache memory at each anchor node storing lineage graphs, alias registrations, mutation records, and policy constraints, an alias resolution engine mapping human-readable identifiers to variance-derived unique identifiers under cryptographically signed policy constraints enforced by anchor quorum consensus, and a provenance validator constructing and traversing multi-root lineage graphs by slope vector proximity and mutation delta computation. The system enables decentralized content identity, alias governance, mutation tracking, and provenance verification across variance-partitioned substrates without reliance on centralized registries or static network addressing.

[0009] In an embodiment, a computer-implemented method for computable content identity comprises normalizing a digital content artifact to a canonical scalar field representation in a modality-specific manner; extracting a multi-axis variance vector from the canonical scalar field; deriving a unique identifier by computing a weighted linear combination of variance vector components organized into axis-dominant triads and hashing the combined vector under a multi-scale quantization scheme; quantizing a global variance value into a slope band 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 registered unique identifier and previously anchored unique identifiers and weighting each lineage edge proportional to the computed 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. The method may be performed without reliance on centralized registries, static network addressing, or persistent external credentials.

[0010] In an embodiment, a rights-grade content admissibility platform comprises a pre-release admissibility engine evaluating a candidate content artifact against one or more cryptographically signed policy objects prior to external commitment, each policy object defining admissible content categories, restricted content classes, jurisdictional constraints,

similarity tolerance thresholds, override authorities, and escalation paths; a structural similarity evaluator computing, prior to commitment, cosine similarity between the candidate artifact's multi-axis variance vector and variance vectors of reference artifacts in a governed corpus, with rejection, regeneration, or escalation when similarity exceeds a policy-declared threshold; a training corpus governance layer admitting digital artifacts to a generative model training corpus only under signed corpus policy objects and recording cryptographically verifiable lineage linking trained model artifacts to the admissible corpus; and a consultation event logger deterministically recording each generation event that consults a reference artifact through retrieval or structured neighborhood resolution. The platform renders structurally impermissible content non-executable prior to artifact commitment, and admissibility decisions are reproducible and auditable from versioned policy objects and consultation event logs.

BRIEF DESCRIPTION OF THE DRAWINGS

[0011] For the purpose of illustrating the disclosure, the drawings show aspects of one or more embodiments of the disclosure. However, it should be understood that the present disclosure is not limited to the precise arrangements and instrumentalities shown in the drawings, wherein:

FIG. 1 illustrates the quadrant decomposition and spatial sub-region fingerprinting process by which a 320-bit unique identifier is constructed from rotation-invariantly sorted per-quadrant hashes;

FIG. 2 illustrates the slope band indexing and anchor node distribution topology, including the band adjacency graph and the quorum-based resolution mesh by which anchor governance is exercised within and across slope bands;

FIG. 3 illustrates the multi-root composite lineage graph and mutation versioning model by which derivative artifacts are linked to multiple parent artifacts through weighted attribution edges; and

FIG. 4 illustrates the training-level content governance and curriculum integration architecture, including the closed feedback loop by which per-band validation loss governs subsequent batch composition.

DETAILED DESCRIPTION

1. Overview of the Computable Content Identity Platform

[0012] The present invention discloses a platform for computing a persistent, structurally derived, and policy-governed identity for digital content artifacts. The platform assigns each artifact a unique identifier (UID) derived deterministically from the artifact's internal variance and structural features, rather than from its storage location, file name, cryptographic key, or transmission metadata. UIDs produced under this system encode not merely a binary fingerprint but a multi-dimensional semantic position within a continuous variance space, enabling similarity detection, mutation tracing, derivative attribution, and distributed governance without reliance on centralized registries, static addressing, or persistent credential infrastructure.

[0013] Content objects processed by the platform---including raster images, vector graphics, audio waveforms, video sequences, textual documents, binary blobs, and mixed-media artifacts---are analyzed through a pipeline that extracts variance-reflective signals at multiple structural scales. These signals are encoded into a multi-axis variance vector that characterizes each artifact's energy distribution, frequency compaction behavior, gradient orientation structure, and spatial sub-region composition. The vector is then hashed under multiple overlapping quantization scales to produce a UID that is stable under controlled transformations e.g., format conversion, resolution rescaling, and lossy compression within defined thresholds, while diverging predictably as variance-shifting mutations occur.

[0014] A content encoder receives a digital content artifact and routes it through the multi-axis variance vector extraction stage, which produces a structured numerical representation of the artifact's internal composition. The slope-band assignment module quantizes the resulting variance vector into one or more variance bands within a global slope continuum and routes the artifact's UID to the appropriate cluster within the distributed anchor node network. The alias resolution engine maintains mappings between human-readable symbolic identifiers and variance-derived UIDs, enforcing policy constraints through anchor quorum consensus. The provenance validator constructs and traverses multi-root lineage graphs linking UIDs to their semantic predecessors and derivative descendants.

[0015] A provenance validator receives outputs from the alias resolution engine and verifies the authenticity and chain of custody of the content artifact before final UID generation. The

system produces a UID output representing the final computed unique identifier for the content artifact, derived from the provenance-validated alias resolution.

[0016] In an embodiment, the disclosed system extends a content anchoring layer described in connection with an adaptive network framework, and discloses computational methods, data structures, vector encodings, hashing schemes, spatial decompositions, and comparison operators that together constitute an enabled content identity system operable across image artifacts of varying resolution, format, compression level, and spatial composition. The present disclosure further incorporates by reference U.S. Nonprovisional Patent Application No. 19/230,933, titled "Cognition-Native Semantic Execution Platform for Distributed, Stateful, and Ethically-Constrained Agent Systems," filed June 6, 2025; U.S. Nonprovisional Patent Application Serial No. 19/326,036, titled "Adaptive Network Framework for Modular, Dynamic, and Decentralized Systems," filed September 11, 2025; and U.S. Nonprovisional Patent Application Serial No. 19/366,760, titled "Cognition-Compatible Network Substrate and Memory-Native Protocol Stack," filed October 23, 2025, each of which is incorporated by reference in its entirety for the subject matter of agent execution, memory-native substrate protocols, and adaptive indexing respectively.

[0017] The architecture is substrate-agnostic and operates independently of device-bound session state, transport-layer addressing, or centralized certificate authorities. Content UIDs may be computed on any conforming processing node, stored in memory-native anchor substrates, and resolved through variance-band-routed queries without external registry consultation. This independence enables deployment across centralized data centers, federated institutional clusters, decentralized peer networks, mobile edge nodes, and intermittently connected devices, as further described in Section 10.

2. Multi-Axis Variance Vector Extraction

[0018] In an embodiment, the multi-axis variance vector extraction pipeline accepts a digital content artifact and produces a 9-dimensional variance vector organized into three structural axes, designated X, Y, and Z. Each axis encodes a distinct and complementary aspect of the artifact's internal composition, and together they represent a semantically rich, mutation-sensitive characterization of the artifact suitable for UID derivation, slope-band assignment, and similarity comparison.

[0019] For raster image artifacts, the extraction pipeline first converts the input to a normalized grayscale floating-point representation through the grayscale conversion module. The conversion applies a perceptual luminance weighting of approximately 0.299 for the red channel, 0.587 for the green channel, and 0.114 for the blue channel, yielding intensity values normalized to the range [0, 1]. This luma-weighted conversion is luminance-preserving and suitable for structural analysis of photographic, illustrative, and diagrammatic content. For non-image artifacts e.g., audio waveforms, textual documents, or binary objects, equivalent normalization procedures are applied to extract a comparable scalar field from which multi-scale variance and structural features may be computed, as further described in Section 2.7.

[0020] The multi-scale variance flow analyzer subdivides the normalized scalar field into three nested grid resolutions: a coarse grid of 8x8 cells covering the full artifact extent, a medium grid of 16x16 cells, and a fine grid of 32x32 cells. For each cell at each grid resolution, the analyzer computes a variance-based proxy defined as the variance of pixel intensity values within the cell. The per-cell variance values for each grid resolution are aggregated to produce a mean variance and standard deviation of variance for each scale level, yielding six scalar values collectively characterizing variance distribution across the spatial hierarchy of the artifact.

[0021] The X-axis energy behavior vector encodes the artifact's energy distribution trend across scale. Three components are computed: the slope of mean variance from the coarse to fine grid resolution, representing the rate at which variance concentrates or disperses as spatial resolution increases; the curvature of mean variance at the medium scale, representing deviation from a linear energy trend; and the asymptotic fine-scale energy value, representing the limiting energy concentration of the artifact at maximum examined resolution. These three components together capture whether an artifact is characterized by broad, diffuse energy distribution or by fine, spatially concentrated energy structures.

[0022] The Y-axis frequency compaction vector encodes the behavior of variance dispersion across scales. Three components are computed: the rate of change of variance standard deviation from coarse to fine resolution, indicating whether the artifact's variance distribution becomes more or less spatially uniform at finer scales; the spread factor between maximum and minimum per-scale variance standard deviation, characterizing the range of distributional variability; and the variance floor convergence value, measuring the proximity of scale-level mean variance to

the global variance of the full artifact, indicating how closely multi-scale behavior approximates the global statistical profile. These components are sensitive to frequency content, texture regularity, and compositional heterogeneity.

[0023] The Z-axis structural phase persistence vector encodes orientation and structural stability. The gradient histogram module computes a histogram of gradient magnitudes across eight angular bins spanning zero to pi radians, aggregated over all interior pixels of the normalized scalar field. The orientation canonicalization stage rotates the histogram so that the dominant angular bin is positioned at index zero, producing a rotation-invariant representation of the artifact's edge orientation distribution. Two components of the Z-axis vector are derived from this histogram: the horizontal-vertical orientation bias, computed as the mean weight of horizontal bins minus the mean weight of vertical bins; and the diagonal-axial bias, computed as the mean weight of diagonal bins minus the mean weight of axial bins. The third Z-axis component is a stability coefficient computed as one minus the absolute difference between the artifact's edge density and its normalized global variance, measuring the structural coherence between edge density and statistical variance.

[0024] The edge density module computes the fraction of interior pixels whose gradient magnitude exceeds a threshold of 0.1, providing a normalized measure of structural complexity. The global variance module computes the full-image variance as a proxy for information density. These two scalar values supplement the multi-scale variance flow and gradient histogram in the construction of the Z-axis vector.

[0025] The resulting nine-dimensional vector, comprising three components each along the X, Y, and Z axes, constitutes the primary variance representation of the artifact. This representation is designed to be stable under format conversion, resolution rescaling within a defined canonical size, and lossy compression at moderate quality levels, while varying predictably with semantic-content-altering transformations e.g., object insertion, removal, significant cropping, style transfer, or compositional remixing.

[0026] The nine-dimensional vector is combined with quadrant sub-region vectors and optionally supplemented by structure and constellation signatures, as described in Sections 3 and 4, to produce the final UID. The combined 27-dimensional vector used for hashing is formed as a weighted linear combination of the X, Y, and Z axes organized into three dominant triads: an X-

dominant triad weighting X at 0.6 and Y and Z each at 0.2; a Y-dominant triad weighting Y at 0.5 and X and Z at 0.3 and 0.2 respectively; and a Z-dominant triad weighting Z at 0.5 and X and Y at 0.3 and 0.2 respectively. This weighting scheme ensures that no single axis completely dominates the UID and that the fingerprint is sensitive to structural perturbations along each behavioral dimension.

[0027] For non-image digital artifacts, the extraction pipeline applies modality-appropriate preprocessing to produce a normalized scalar field from which the multi-scale variance, gradient, and density analyses described above may be computed. For audio waveforms, the input signal may be represented as a two-dimensional time-frequency spectrogram normalized to a canonical resolution, whereupon the multi-scale variance flow, gradient histogram, and edge density computations operate over the spectrogram matrix in an identical manner to image processing. For textual documents, the normalized scalar field may be derived from token frequency distributions mapped onto a two-dimensional positional grid, or from byte-level variance across fixed-width content windows. For binary objects and structured data artifacts, any representation that produces a bounded, two-dimensional scalar field of normalized values may serve as input to the extraction pipeline. Such modality-specific normalization is within the scope of the present disclosure.

3. Quadrant Decomposition and 320-Bit UID Construction

[0028] FIG. 1 illustrates the quadrant decomposition and spatial sub-region fingerprinting process. Following global variance vector extraction, the content artifact is processed through a spatial decomposition pipeline that produces independent fingerprints for four non-overlapping quadrant regions, enabling detection of partial similarity, regional mutation, and spatial composition matching that is not captured by the global nine-dimensional vector alone.

[0029] The canonical normalization stage (300) rescales the input artifact to a square canvas of 256 pixels by 256 pixels by computing the ratio of the target dimension to the longest edge of the source artifact and applying uniform scaling along both dimensions, then centering the scaled image on a black background fill. This letterbox-style normalization preserves aspect ratio while ensuring all artifacts occupy a consistent canonical coordinate space for quadrant extraction. Image smoothing is disabled during rescaling to prevent anti-aliasing from introducing artificial variance signals along rescaled edges.

[0030] The orientation canonicalization module (302) computes the dominant gradient orientation of the normalized artifact by extracting the peak bin of an eight-bin gradient histogram over the full canonical image. If the dominant orientation angle exceeds an absolute value of approximately 0.1 radians, corresponding to approximately 5.7 degrees, the artifact is rotated by that angle about the center of the canvas prior to quadrant extraction. This canonicalization step ensures that artifacts with predominant orientation structure are presented in a consistent spatial frame prior to spatial decomposition, reducing inter-format and inter-resolution UID drift attributable to minor rotational variations.

[0031] The four spatial quadrant extraction modules (304a), (304b), (304c), and (304d) extract non-overlapping rectangular sub-images from the canonical image: (304a) extracts the top-left quadrant, (304b) the top-right quadrant, (304c) the bottom-left quadrant, and (304d) the bottom-right quadrant, each comprising one-quarter of the canonical image area. Per-quadrant variance vector computation (306a), (306b), (306c), and (306d) applies the full nine-dimensional extraction pipeline described in Section 2 to each quadrant sub-image, producing an independent X, Y, Z triplet for each spatial region.

[0032] Per-quadrant hashing (308a), (308b), (308c), and (308d) converts each quadrant's nine-dimensional vector into a 256-bit hexadecimal hash using a coarsened quantization scheme in which the X and Y axis components are quantized at a step of 1/32 and Z axis components at a step of 1/8 prior to hashing. This coarser quantization absorbs JPEG compression noise and format conversion artifacts at the sub-image level, where localized compression introduces greater per-pixel variance than at the full-image level. The hash function applies multiple overlapping FNV (Fowler–Noll–Vo)-variant hash functions at two quantization scales (16 and 20) and XORs the resulting 128-bit outputs to produce the final 256-bit quadrant hash.

[0033] The rotation-invariant quadrant sorting module (310) sorts the four quadrant hashes in lexicographic order of their hexadecimal string representations, then assigns the sorted hashes to canonical positions q0 through q3. This sorting step ensures that the assignment of quadrant identifiers is independent of the spatial orientation of the artifact, so that a rotated or mirrored version of an artifact produces the same set of sorted quadrant hashes even if the individual spatial positions of the quadrants differ.

[0034] The combined 320-bit unique identifier construction module (312) applies a multi-segment FNV-64 hash combiner to the global 256-bit hash and the four sorted quadrant hashes, producing five 64-bit hash segments that are concatenated to form a 320-bit UID. The combiner applies five distinct initialization seeds to the same ordered concatenation of hash inputs, yielding five independent 64-bit outputs that together constitute an address space of 2^320 distinct UIDs. This address space is sufficient to ensure negligible collision probability across any foreseeable scale of digital content deployment.

[0035] The first 16 hexadecimal characters of the 320-bit UID, representing 64 bits, serve as a backward-compatible short-form identifier suitable for human-readable display, database indexing, and bandwidth-constrained transmission. The full 320-bit representation is used for anchor node assignment, slope-band routing, and similarity comparison operations. The variance band classification, as described in Section 5, derives from the global variance value and is carried as a supplementary field in the UID record.

4. Structure Signature and Constellation Signature

[0036] These signatures are optional components of the full UID record and serve specialized matching functions: the structure signature supports recognition of logos, icons, and graphically sparse artifacts across background color changes and flat-fill variations; the constellation signature supports matching of artifacts that share distinctive spatial compositions across cropping, scale change, and partial occlusion.

[0037] The gradient-only structure vector module computes a 21-dimensional structure vector derived exclusively from spatial gradient information, suppressing sensitivity to mean luminance and background fill. A 16-bin gradient histogram is computed over the full canonical image, with each bin accumulating gradient magnitude weighted by edge strength across all interior pixels. The histogram is canonicalized by rotating to place the maximum-weight bin at index zero. The multi-threshold edge density stage computes edge density at three gradient magnitude thresholds of 0.05, 0.10, and 0.15, producing three scalar values that characterize the structural content of the image across fine, medium, and coarse edge-strength levels. The gradient histogram moment computation derives the variance and peak of the normalized 16-bin histogram, providing compact statistics of gradient energy distribution. The structure hash output concatenates the 16-bin histogram values, the three edge density scalars, the histogram variance,

and the histogram peak into a 21-dimensional structure vector, which is hashed using the STRUCTURE_COARSE_SCALES scheme at quantization scales of 24, 32, and 40, producing a 256-bit structure hash that is further extended to a 320-bit structure identifier.

[0038] The saliency hotspot detection module identifies up to five spatially distributed high-salience anchor points within the canonical image. Saliency is computed over a 12x12 coarse grid by scoring each cell as the sum of twice its intra-cell intensity variance and its mean local gradient density. Cells are sorted by descending score with a minimum inter-cell separation of one grid unit to prevent multiple hotspots from clustering on a single prominent feature. At least three hotspots are required for constellation signature computation; if fewer than three are detected, the constellation signature is not emitted and the corresponding field in the UID record is null.

[0039] The micro-constellation computation stage builds an independent descriptor for each detected hotspot. For each hotspot serving as a focal anchor, the nearest neighbors among the remaining hotspots are identified by Euclidean distance. Distances are normalized by the maximum pairwise inter-hotspot distance within the image, rendering the constellation scale-invariant. Normalized distances are quantized into eight bins. Angles between the vector from the focal anchor to the nearest neighbor and each subsequent neighbor vector are computed in the range zero to pi radians, making the constellation rotation-invariant. Angular values are quantized into 12 bins of 15 degrees each. A descriptor string encoding the hotspot grid coordinates, quantized distances, and quantized angles is hashed using a 64-bit FNV variant to produce a micro-constellation hash for each focal anchor.

[0040] The rotation-invariant micro-hash aggregation module sorts the set of per-hotspot micro-constellation hashes lexicographically, eliminating any dependence on the order in which hotspots were detected or processed. The sorted hashes are concatenated and hashed using the structure hash function to produce the 256-bit constellation hash, which is extended to a 320-bit constellation identifier by the multi-segment FNV-64 combiner.

[0041] The anchor ID module derives a non-authoritative 128-bit anchor identifier from the 320-bit structure identifier and the 320-bit constellation identifier by applying the multi-segment combiner to their concatenation and retaining the first two 64-bit segments. This anchor ID serves as a compact routing handle for content objects whose structure and geometric

composition are intended to match across background substitution, format conversion, and partial cropping, as commonly encountered in logo recognition, reverse image search, and derivative content attribution workflows.

5. Variance Band Classification and Anchor Node Distribution

[0042] FIG. 2 illustrates the slope band indexing and anchor node distribution topology. The variance band classification module assigns each content UID to one of five variance bands based on the global variance value of the artifact. Band 1 encompasses artifacts with global variance below 0.02, characterizing near-uniform content e.g., solid fills, blank documents, and flat-color graphics. Band 2 encompasses artifacts with global variance between 0.02 and 0.06, characterizing low-detail content with sparse but identifiable structure. Band 3 encompasses artifacts with global variance between 0.06 and 0.12, characterizing moderate-complexity content with visible but regular structure. Band 4 encompasses artifacts with global variance between 0.12 and 0.22, characterizing high-complexity content with varied spatial structure. Band 5 encompasses artifacts with global variance at or above 0.22, characterizing very-high-variance content e.g., natural photographs, dense text documents, and richly textured compositions.

[0043] Each variance band defines a routable index segment within the global slope continuum. Anchor nodes are assigned to one or more variance bands and declare governance responsibility for the UIDs whose global variance values fall within their assigned band or bands. The anchor cluster assignment logic (504) maps each UID to a primary anchor node (506) and one or more secondary anchor nodes (508a), (508b), and (508c) based on the UID's variance band and the current topology of the anchor network. The band adjacency graph (510) records which bands share governance boundaries, enabling cross-band resolution pathfinding for content that drifts between bands under mutation, as described in Section 7.

[0044] Within each variance band, anchor nodes are organized into a quorum-based resolution mesh (512). Each anchor in the mesh independently stores a fragment of the UID index for its assigned bands and participates in an Adaptive Consensus Protocol to validate alias registrations, resolve conflicting UID assignments, and maintain deterministic routing consistency across the distributed cache. In an embodiment, the Adaptive Consensus Protocol comprises (i) trust-weighted asynchronous voting in which each anchor's vote weight is derived

from its declared band scope, historical reliability, and trust-zone authority; (ii) per-mutation quorum thresholds configurable by mutation type, such that routine UID registration requires a configurable minimum quorum and structural mutations e.g., anchor recruitment or band split require a configurable supermajority; and (iii) lineage-preserving commitment in which each accepted mutation records its predecessor state, the participating anchor signatures, and the policy version under which it was evaluated, admitting later replay verification. The Adaptive Consensus Protocol is described in further detail in U.S. Patent Publication No. 20260010525, which is incorporated by reference herein for such disclosure. Quorum thresholds for alias mutation and UID registration are configurable per band and per trust zone, enabling stricter governance for high-value content ranges and lighter-weight coordination for low-variance or low-activity bands.

[0045] Slope binning enables content resolution to be routed without reference to the network address, location, or identity of any specific node. Any processing node that knows the variance value of a target artifact's UID can determine which variance band governs the UID and direct a resolution query to the appropriate anchor cluster without consulting a central directory. This property ensures that UID resolution remains functional in disconnected, asynchronous, and adversarial environments where central registries may be unreachable or untrustworthy.

[0046] In an embodiment, alternative deployments define a finer band granularity such as 10, 20, or 100 bands or a continuous slope spectrum with fuzzy band boundaries, depending on corpus scale, variance diversity, and the routing resolution required. The UID structure supports arbitrary band granularity without modification.

6. Adaptive Cache Governance

[0047] Each anchor node maintains a memory-resident UID cache that stores UID records, alias registrations, lineage graph fragments, and policy constraints for content objects within its governed variance bands. Cache behavior is governed dynamically rather than by static TTL configuration, enabling anchors to self-regulate based on local variance saturation, access frequency, policy constraints, and the semantic relevance of cached objects.

[0048] The variance saturation monitor continuously evaluates the density of active UIDs within each governed variance band. When a band approaches a configured saturation threshold,

the monitor signals the eviction engine to identify and remove low-priority UID entries. Priority is computed as a function of recency-weighted access frequency, slope proximity to recently queried UIDs, and policy-declared retention priority. Entries for stale, low-access, or policy-expired artifacts are evicted first; entries for high-value, frequently accessed, or policy-protected artifacts are retained.

[0049] The access frequency tracker maintains rolling access counts and recency timestamps for each cached UID entry. These statistics inform both eviction priority and replication decisions. UIDs experiencing sustained access volume above a configurable threshold may trigger selective replication to secondary anchor nodes within the same band cluster, as governed by the replication controller. Replication is scoped to the band cluster and constrained by the trust zone policy of the originating anchor, preventing unauthorized propagation to external zones.

[0050] The policy constraint evaluator enforces governance rules embedded within each UID's policy field at the time of cache admission, replication, and eviction. Policy objects may specify zone-local propagation only, prohibiting replication to anchors outside the originating trust zone even when variance proximity would otherwise qualify an adjacent anchor for replication. Policy objects may specify time-bounded retention, after which cached entries are automatically marked for eviction regardless of access frequency. Policy objects may also specify read-only proxy status, permitting the anchor to serve UID metadata and alias lookups without retaining full lineage graphs, as implemented by the read-only proxy cache.

[0051] The cache invalidation event handler processes quorum-initiated invalidation signals, semantic rollback events, and fork adjudication outcomes. When a cached UID is invalidated through anchor quorum decision---for example, because a lineage audit detects variance discontinuity exceeding the configured similarity threshold, or because a duplicate alias claim is resolved against the cached artifact---the handler flags the affected UID entry, notifies downstream subscribers if configured, and removes or quarantines the entry. Invalidation events are recorded in the anchor's event log and may be included in future resolution responses to inform querying nodes of the invalidation status.

[0052] Cache replication coordination among anchors within a slope band cluster operates under the Adaptive Consensus Protocol. When a new UID is registered with a primary anchor,

that anchor may initiate a selective replication proposal to secondary anchors in the cluster. The proposal specifies the UID record, the variance band, the applicable policy constraints, and a replication priority score derived from the global variance of the artifact and the access history of the registering agent. Secondary anchors evaluate the proposal against their local memory availability and band governance configuration and respond with acceptance or rejection. A quorum of accepting anchors is required before the UID is considered durably replicated within the cluster.

7. Semantic Alias Registration and Resolution

[0053] In an embodiment, the alias registration and resolution framework enables human-readable symbolic identifiers to be registered, maintained, and resolved in association with variance-derived UIDs. Aliases allow content creators, institutional publishers, platform operators, and agents to assign persistent symbolic references to content objects without replacing or conflating the deterministic, variance-derived identity of those objects. Aliases are treated as mutable, scoped references governed independently of the UID itself and resolved through slope-band-scoped anchor consensus.

[0054] An alias registration request specifies a human-readable string, the UID of the target content object, the requested scope, and a cryptographically signed policy object. The scope classifier categorizes the registration into one of three scope levels: band-local, restricting the alias to the variance band of the target UID; zone-local, extending the alias to all variance bands governed by anchors within the same trust zone; or global, requiring the alias to be uniquely resolvable across all participating anchor clusters. Global alias registrations are subject to heightened consensus requirements, including proof-of-variance-priority or multi-anchor confirmation, to prevent namespace collision in the global alias space.

[0055] The anchor quorum validator evaluates the alias registration request against the governance rules of the target band cluster. The validator checks that the requesting party possesses mutation authority over the target UID, that the proposed alias does not conflict with an existing registration at the same or broader scope level, and that the policy object satisfies the format and signature requirements of the trust zone. If validation succeeds, the alias record is committed to alias record storage on the registering anchor and propagated to secondary anchors in the cluster in accordance with the Adaptive Consensus Protocol.

[0056] Alias resolution is performed by the multi-stage resolution engine. When a resolution query is received, the variance band identification stage determines the most likely variance band for the alias based on content-type metadata, prior resolution history, or explicit band hints included in the query. The UID association lookup queries the anchor cluster of the identified band for a UID associated with the provided alias string. If a match is found, the policy validation stage verifies that the querying party satisfies the scope constraints and access conditions of the alias record. Upon successful validation, the canonical UID output returns the target UID, any active lineage annotations, and applicable policy constraints to the querying agent.

[0057] If the alias cannot be resolved within the initially identified variance band, adjacent bands in the band adjacency graph (510) may be queried in order of proximity, subject to policy scope limits and anchor resource availability. If an alias is queried without knowledge of the associated slope band, anchors may perform an index pathfinding traversal using deterministic band adjacency rules to locate the governing anchor cluster, as described in Section 7.2. All resolution attempts, including negative results, are recorded in the anchor event log to support auditability and conflict detection.

[0058] The fork adjudication handler is invoked when a resolution query reveals conflicting alias claims across slope-divergent UIDs. Conflict resolution proceeds through anchor quorum review, which evaluates registration timestamps, variance proximity to the canonical UID, policy lineage, and trust zone authority. The quorum may resolve the conflict in favor of one registration, freeze both registrations pending external governance escalation, or quarantine the alias pending variance continuity verification. All adjudication outcomes are recorded and are accessible to querying parties as provenance metadata.

[0059] Aliases may be reassigned or updated by parties with mutation authority over the target UID, provided the alias policy permits reassignment and the updated target UID satisfies the policy-declared slope proximity threshold. All alias mutation events are recorded in the anchor memory field and contribute to the provenance graph of the UID.

8. Composite Attribution and Multi-Root Lineage Graphs

[0060] FIG. 3 illustrates the multi-root composite lineage graph and mutation versioning model. In an embodiment, the platform supports composite content attribution by enabling each UID to participate in a directed lineage graph that may contain more than one parent. Content artifacts generated through combination, transformation, remixing, or derivation from multiple prior artifacts may be registered under a new UID linked to all contributing source UIDs, with lineage edge weights reflecting the degree of variance inheritance from each contributor.

[0061] When content artifact C (804) is registered as a derivative of content artifact A (800) and content artifact B (802), the variance vector comparison module (806) computes the cosine similarity between the variance vector of C and the variance vectors of A and B, respectively. Cosine similarity is computed over the concatenated X, Y, and Z axis components of the nine-dimensional global variance vector as the inner product of the two vectors divided by the product of their magnitudes. The slope delta computation module (808) computes the Euclidean distance between the variance vector of the candidate derivative and each proposed parent, normalized over a configured maximum distance threshold to produce a similarity score in the range [0, 1]. These similarity scores determine whether the observed slope proximity is within the configured semantic continuity threshold required to establish a lineage edge.

[0062] Weighted parent attribution (810) assigns a contribution weight to each confirmed parent UID proportional to its cosine similarity with the derivative UID. These weights are recorded as edge annotations in the directed lineage graph but do not constitute legal determinations of authorship or ownership; they serve as structural signals that may inform licensing, attribution display, policy inheritance, and governance enforcement by downstream systems. The directed lineage edge registration module (812) commits the edges, slope delta values, contribution weights, and anchor endorsement signatures to the lineage graph stored in the governing anchor's cache memory.

[0063] The version history record (814) stores a directed graph of UID transitions ordered by registration timestamp, where each edge encodes the mutation type, the slope delta between predecessor and successor UIDs, and the policy inheritance conditions applicable to the mutation. Versioning is structural and does not depend on author declaration or explicit version numbering: any new UID that falls within the configured slope continuity threshold of a

previously registered UID is automatically considered a candidate for version linkage, subject to anchor quorum confirmation.

[0064] The delegation policy evaluator (816) enforces ownership delegation constraints at the time of UID registration and mutation. Delegation policy objects specify permitted mutation types, temporal validity bounds, UID lineage depth limits, trust zone whitelists, and transferability conditions. If a proposed mutation violates a delegation constraint---for example, if an unauthorized agent attempts to register a derivative UID from a policy-protected source---the mutation rejection handler (818) denies registration and records the attempt in the anchor event log. Delegation constraints may propagate through version lineages in accordance with the inheritance mode specified in the policy object.

[0065] In an embodiment, the comparison framework operates on UID records and produces per-axis cosine similarity scores for the X, Y, and Z axes, an aggregate directional cosine similarity, a Euclidean distance-based similarity score, and per-quadrant similarity scores together with hash match flags. The quadrant-level comparison supports localized mutation detection: a derivative artifact that modifies only one spatial region of a source exhibits quadrant similarity scores near 1.0 for unchanged regions and reduced scores for modified regions, enabling spatially resolved attribution even when the global similarity score remains high.

[0066] A composite content artifact is considered fully auditable. Any party with access to the governing anchor cluster may trace the UID of a composite artifact to its contributing parents, inspect mutation metadata, evaluate policy lineage, verify anchor quorum endorsements, and assess whether the composition satisfies the governance requirements of applicable trust zones. Adversarial recombinations---in which a party attempts to register a derivative that suppresses attribution to one or more source UIDs by manipulating variance features---are detectable because the slope profile of the composite UID will exhibit measurable divergence from the weighted variance combination of its declared parents. Anchors may flag such entries and subject them to heightened governance review.

9. Semantic Routing and Cache Policy Enforcement

[0067] In an embodiment, the platform resolves content queries through variance-band-targeted semantic routing rather than IP address resolution, DNS lookup, or static path traversal.

When an agent or node initiates a UID resolution query, the query is directed to anchor nodes responsible for the variance band corresponding to the target UID. If the exact governing band cannot be determined from the query parameters, anchors in the estimated band evaluate the query and, if unable to resolve, refer the query to adjacent bands through bandwidth-scoped propagation subject to policy limits. Routing may also include alias redirection, fallback to lower-fidelity proxy caches, and meta-policy escalation if propagation constraints are triggered during traversal.

[0068] Cache nodes within an anchor cluster may participate in scoped replication contracts that define when and how a node may replicate a UID, including variance verification requirements, mutation validity thresholds, and TTL parameters where applicable. Full anchors maintain authoritative mutation histories and alias assignment records. Read-only proxy nodes serve UID metadata and alias lookups without retaining full lineage traces, reducing storage and bandwidth requirements for nodes operating at the periphery of an anchor cluster.

[0069] In a disconnected or low-variance execution environment, fallback caching retains content artifacts from prior authenticated queries; cached copies are invalidated, refreshed, or verified upon reconnection. Cached UIDs exceeding the configured staleness threshold are flagged for re-validation before being served, preventing cached content from drifting beyond its permitted propagation scope.

[0070] Cache propagation boundaries are enforced through policy-scoped replication filters derived from governance objects embedded in the UID record, defining whether a UID may be cached locally, forwarded to adjacent variance bands, exported to external zones, or retained beyond a specified mutation threshold. Anchors in federated deployments may instantiate partial caches for UIDs governed by adjacent band clusters in other administrative zones, retaining only essential metadata for resolution fallback without replicating full lineage graphs or policy enforcement state.

10. Substrate Deployment and Interoperability

[0071] In an embodiment, the content identity system is configured for modular deployment across heterogeneous computational environments, including centralized data centers, federated institutional nodes, decentralized peer clusters, mobile edge substrates, and intermittently

connected devices. The system operates independently of conventional transport-layer infrastructure, deriving routing eligibility and cache governance from slope-indexed identity and policy-scoped anchor memory rather than from network addresses, session state, or external certificate services.

[0072] In an embodiment, anchors in centralized deployments serve as persistent authorities for high-volume variance bands, maintaining full-resolution UID registries, mutation histories, and alias graphs. Centralized anchor configurations support deterministic resolution latency, continuous policy enforcement, and zone-scoped governance consistent with institutional mandates or regulatory frameworks. In a further embodiment, a batch processing architecture retrieves unprocessed content records from a persistent store, computes variance vectors and UIDs, and commits sanitized UID records back to the store, providing a reference architecture for high-volume centralized UID computation pipelines.

[0073] In federated deployments, anchors operate semi-autonomously under shared trust frameworks, exchanging variance delta vectors, slope lineage graphs, and alias conflict states across administrative boundaries while preserving local decision authority. Federated anchors may resolve UID queries for content governed by remote band clusters by maintaining low-fidelity partial caches or by initiating inter-anchor referral protocols, as described in Section 7.

[0074] In a further embodiment, decentralized and edge-based substrates host anchors in partial or transient configurations, retaining only scoped index segments or alias route mappings. These anchors operate in quorum-supplemented read-only mode, validating UID slope continuity and deferring final mutation or alias decisions to upstream authorities upon reconnection. In variance-constrained environments e.g., sensor clusters with limited memory, anchors cache recent UID queries and execute only partial slope delta checks, retaining full rehydration capability until complete validation becomes feasible.

[0075] Interoperability with memory-native agent systems is optional. Content UIDs may be referenced, embedded, or governed by semantic agents operating on the cognition-native execution platform described in related applications, or exist as standalone slope-bearing objects without agent encapsulation. Anchor nodes may serve as semantic routing interfaces during agent delegation or alias query events. Trust zones under cognition-native governance models

may define propagation scope, alias authority, and mutation eligibility for content UIDs using the same policy enforcement architecture applicable to agent objects.

[0076] In an embodiment, a record-sanitization procedure addresses substrate compatibility concerns of persistent document stores that prohibit nested array values and require content UID records to be flattened before persistence. The procedure recursively traverses a UID record structure and converts any nested array or nested object value within an array field to a JSON string representation, preserving information content while conforming to the flat-value constraint of a target substrate. The disclosure is not limited to any particular storage format or substrate.

11. Multi-Modal Content Identity: Audio, Text, Video, Binary, and Streaming Modalities

[0077] In an embodiment, the multi-axis variance vector extraction pipeline disclosed in Section 2 operates on any normalized scalar field representation of a digital artifact and is therefore extensible to content modalities beyond raster images. The present section discloses the specific normalization procedures, extraction adaptations, and UID construction methods for audio waveforms, textual documents, and video sequences. These disclosures establish full written description support for multi-modal claims and enable a practitioner of ordinary skill to implement the extraction pipeline for each modality without undue experimentation.

[0078] For audio waveform artifacts, the normalization procedure computes a time-frequency representation of the waveform using a short-time Fourier transform with a Hann window of approximately 2048 samples and a hop length of approximately 512 samples, producing a two-dimensional magnitude spectrogram. The spectrogram is mapped to the mel frequency scale using a filterbank of 128 mel bins spanning the frequency range of the artifact's sample rate, and normalized to a canonical resolution of 256 time frames by 128 frequency bins, with log-magnitude scaling applied to compress dynamic range. This normalized mel-spectrogram serves as the scalar field input to the multi-axis variance vector extraction pipeline. The variance-based proxy computed over this scalar field captures frequency energy distribution across time-frequency cells, the gradient histogram captures transitions between frequency regions and temporal onset patterns, and the edge density metric captures the structural complexity of the spectral profile. The resulting nine-dimensional variance vector encodes audio texture, onset density, harmonic richness, and spectral centroid behavior in a form directly

comparable to image-domain variance vectors through the same cosine similarity operator. The temporal delta vector for audio artifacts encodes the cross-frame cosine similarity between consecutive short-time spectral windows, providing a compact representation of temporal dynamics suitable for clip-level UID derivation.

[0079] For textual document artifacts, the normalization procedure maps the document to a two-dimensional token frequency scalar field. Each distinct token in the document vocabulary is assigned a positional index along one axis and a frequency-weighted salience value along the other, computed as the product of the token's term frequency within the document and the inverse of its document frequency within a reference corpus. The resulting matrix is normalized to a canonical resolution of 256 by 256 cells, with tokens assigned to cells by their positional index and cell values representing aggregated salience scores. Byte-level variance is computed as a supplementary signal across sliding windows of 64 bytes of the document's UTF-8 encoded representation, producing a secondary scalar profile from which variance-based proxies are derived and blended with the token frequency field variance at a configurable weight. The gradient histogram over the token frequency field captures the distributional sharpness of vocabulary usage, the concentration of semantic content, and the regularity of positional token patterns, providing a structural fingerprint sensitive to document genre, authorship style, and compositional density. The resulting variance vector encodes the statistical and structural texture of the document's semantic content in a form directly operable by the slope-band assignment, anchor registration, and lineage graph construction methods.

[0080] For video artifacts, the UID derivation system operates at two levels. At the frame level, each frame of the video is treated as a raster image artifact and processed through the full image extraction pipeline described in Section 2, including canonical normalization, orientation canonicalization, global variance vector extraction, quadrant decomposition, and optionally structure and constellation signature computation. At the clip level, a temporal delta vector is derived by computing the cosine similarity between consecutive frame variance vectors and recording the resulting similarity scores as a one-dimensional temporal profile. The temporal delta vector encodes the rate and magnitude of variance change across the clip, capturing scene transitions, motion intensity, and compositional rhythm. A clip-level UID is derived by applying the multi-axis extraction pipeline to the temporal delta vector treated as a one-dimensional signal, producing a variance vector that encodes the clip's dynamic structure and can be

registered with anchor nodes for clip-level identity, lineage tracing, and similarity matching. This two-level architecture supports both frame-level deduplication and clip-level provenance tracking, enabling detection of partial reuse, remix, and cross-format reformatting across video content ecosystems.

[0081] The modality classifier determines the content type of the input artifact and routes it to the appropriate normalization path. The image normalization path produces a grayscale scalar field. The audio normalization path computes a short-time Fourier transform, applies a mel filterbank, and produces a mel-spectrogram scalar field. The text normalization path applies TF-IDF weighting and positional grid mapping to produce a token frequency scalar field. The video normalization path extracts per-frame scalar fields through the frame extraction module and computes a temporal delta vector through the inter-frame cosine similarity module. All four paths converge on the shared multi-axis variance vector extraction stage, which produces the variance vector output regardless of the source modality.

[0082] For real-time streaming content, including live video broadcasts, audio streams, and continuous sensor data, the UID derivation system operates over a sliding window of the stream rather than over a discrete artifact. A sliding window of configurable duration, e.g., 10 seconds for audio or 30 frames for video, is extracted from the stream, normalized, and processed through the multi-axis variance vector extraction pipeline to produce a window-level UID. Consecutive window UIDs are compared by cosine similarity to detect structural continuity or discontinuity in the stream, and each window-level UID is registered with the anchor network for real-time provenance tracking. When the cosine similarity between consecutive window UIDs falls below a configured continuity threshold, the system records a scene transition event in the anchor's event log. When the cosine similarity between a window-level UID and a registered reference UID exceeds the policy-declared similarity threshold, the system generates a real-time match event that may trigger policy enforcement actions including blocking of unauthorized retransmission, generation of a consultation event record, or invocation of the pre-release admissibility engine. This streaming architecture enables live broadcast monitoring, real-time deepfake detection, and continuous provenance tracking across streaming content platforms without requiring offline batch processing.

[0083] For binary object artifacts, including executable files, compiled archives, container payloads, and source code files, the normalization procedure maps the byte sequence of the artifact to a two-dimensional scalar field by reshaping the byte stream into a square or near-square matrix at a canonical resolution and computing per-cell statistics from sliding-window byte variance, byte-frequency entropy approximated as the variance of byte-frequency counts within the window, and structural-section profile values where the artifact carries a recognized container structure such as a portable executable, executable and linkable format, or archive index. The resulting scalar field is normalized to the canonical resolution and processed by the same multi-axis variance vector extraction pipeline as image and audio modalities, producing a UID that encodes the byte-level structural texture, sectional layout, and statistical regularity of the binary artifact and that supports cosine-similarity-based comparison across recompiled, repacked, or partially patched variants of the same underlying payload.

[0084] For vector graphics artifacts, including scalable vector graphics documents, vector portions of portable document format pages, and machine-readable diagram formats, the normalization procedure rasterizes the vector content at a canonical resolution using a deterministic rasterization profile and processes the resulting raster scalar field through the image extraction pipeline, optionally augmented by a path-density scalar field derived from the cumulative path length and control-point density of the vector primitives within each cell of the canonical grid. For structured tabular artifacts, including comma-separated values, columnar data files, and structured object collections, the normalization procedure projects column-wise statistical descriptors, including per-column variance, cardinality, and inter-column correlation, into a two-dimensional cell grid indexed by column position and statistical descriptor type, producing a scalar field operable by the same extraction pipeline. These normalization recipes are illustrative; alternative scalar-field projections that preserve structural variation across cells fall within the scope of the present disclosure.

[0085] The modality-specific normalization procedures and streaming extension are non-limiting examples. Any normalization that produces a bounded, two-dimensional scalar field of normalized values from a digital artifact of any type may serve as input to the multi-axis variance vector extraction pipeline, enabling the content identity system to operate across arbitrary digital content modalities without modification to the extraction, hashing, slope-band assignment, anchor governance, or lineage construction components.

12. Rights-Grade Content Admissibility and Pre-Release Governance

[0086] The candidate artifact input enters the pre-release admissibility engine, which routes the artifact through two parallel evaluation tracks before reaching the commitment gate. The first track is the policy object evaluator, which receives a versioned signed policy object and produces an admissibility decision. The second track is the structural similarity evaluator, which queries the governed exclusion corpus indexed in the slope-band anchor network and produces a similarity score that feeds into the forbidden content exclusion layer, which compares the score against the policy-declared threshold. If both tracks confirm admissibility, the commitment gate permits a committed artifact output. If either track fails, the rejection handler produces a regeneration signal or an escalation signal. All generation events that consult reference artifacts are recorded by the consultation event logger as consultation records. Admitted training artifacts enter the training corpus governance layer, which appends each artifact to the governance record log.

[0087] In an embodiment, the content identity infrastructure disclosed in the preceding sections provides the technical substrate for a rights-grade content admissibility platform that governs generative and distributable content at the commitment boundary rather than through post-hoc moderation. A commitment is any irreversible or externally visible side effect of a content generation or distribution event, including public release, customer delivery, API return, licensing event, marketplace publication, training data admission, or cross-platform provenance anchor. Rights-grade governance interposes an admissibility evaluation between the generation of a candidate artifact and its commitment, ensuring that structurally impermissible content cannot become a released artifact.

[0088] The pre-release admissibility engine evaluates a candidate content artifact against one or more policy objects prior to commitment. Each policy object is versioned, cryptographically signed, and machine-evaluable, defining typed category constraints, jurisdictional scopes, override authorities, similarity tolerance thresholds, and escalation paths. Admissibility decisions are reproducible and auditable: given the variance-derived UID and structural signatures of the evaluated artifact and the policy object version, any authorized party may verify the determination by replaying the evaluation, in distinction from conventional content moderation that relies on opaque classifiers whose decisions are not independently verifiable.

[0089] The structural similarity evaluator computes, prior to commitment, the cosine similarity between the variance vector of the candidate artifact and the variance vectors of reference artifacts indexed in a governed corpus. The governed corpus is a slope-band-indexed anchor network whose entries are registered under signed corpus policy objects specifying admissibility scope, exclusion classes, and similarity tolerance thresholds. If the similarity score between a candidate artifact and any reference artifact in the exclusion corpus exceeds the policy-declared threshold, the candidate is rejected, regenerated under modified generation constraints, or escalated to an authorized override authority. Because similarity evaluation operates over variance-derived unique identifiers rather than requiring GPU inference or centralized embedding indexes, it can be executed client-side, at generation time, and at the scale of real-time content production without per-query compute costs proportional to corpus size.

[0090] The training corpus governance layer operates at the boundary between data ingestion and model training. Digital artifacts are admitted to the training corpus only under signed, declared corpus policy objects that specify permissible content categories, excluded classes, jurisdictional constraints, and usage rights. Each admitted artifact receives a governance record comprising its variance-derived UID, the governing policy object under which it was admitted, a timestamp, and a cryptographic hash of the policy object. These records are appended to an audit log that constitutes a verifiable lineage from the trained model artifacts back to the admissible corpus. This lineage enables an operator to demonstrate corpus scope, governing policy, and artifact provenance as verifiable execution facts rather than as assertions of responsible sourcing practice, shifting the legal posture of model training from self-declared compliance to structurally verifiable lineage.

[0091] The consultation event logger records, for each generation event that consults a reference artifact through retrieval-augmented generation or structured neighborhood resolution, a deterministic consultation record comprising the variance-derived UID of the consulted artifact, the governing policy object, the variance proximity score between the consulted artifact and the generated output, and a timestamp. These records enable attribution and compensation mechanisms to attach to governed consultation events rather than requiring reverse-engineering of model weights. When generation consults reference artifacts within policy-defined similarity neighborhoods, those consultation events are logged as computable attribution events from which compensation obligations may be derived under policy-declared schedules. This

architecture renders creator payment computable from consultation event logs rather than from approximations of training data influence.

[0092] The forbidden content exclusion layer maintains a governed exclusion corpus of variance-indexed forbidden content references, including content classes defined as structurally impermissible under applicable policy objects. Candidate artifacts are evaluated against the exclusion corpus by comparing their variance vectors and structural signatures against the indexed forbidden content references. A candidate artifact whose variance vector falls within a configured proximity of any exclusion corpus entry is rendered non-committable prior to release. This evaluation occurs before the artifact is externally released, ensuring that impermissible content never becomes a committed artifact and that governance prevents forbidden content from existing as released media rather than filtering it after exposure has already occurred.

[0093] The rights-grade admissibility architecture is structurally distinguished from conventional content moderation and watermarking. Conventional moderation applies filters after artifact creation or release, permitting impermissible content to exist as an artifact before detection. Watermarking attaches identity signals severable by transcoding, cropping, or generative reconstruction. The present system evaluates admissibility at the commitment boundary using structurally derived, embedding-free, registration-free variance vectors computed from the content itself, making admissibility evaluation both pre-release and substrate-independent. Nothing is embedded in the artifact, no enrollment is required, and no central registry is needed.

13. Training-Level Content Governance and Curriculum Integration

[0094] FIG. 4 illustrates the training-level content governance and curriculum integration architecture. The training corpus (1100) is processed by the variance band classifier (1102), which assigns each artifact to one of the variance bands. The variance-governed curriculum sequencer (1104) orders training batches by ascending variance band during early training phases and transitions to dynamic ordering in later phases. The slope-band batch composition module (1106) receives per-band validation loss feedback from the model training loop (1108) and adjusts band sampling weights to preferentially admit artifacts from bands where loss remains elevated. The structural provenance trace module (1110) computes cosine similarity between training artifact variance vectors and generated output variance vectors, producing a

memorization proximity score. The depth-wise content attention adapter (1112) receives the variance band classification of the current batch and uses it to modulate pseudo-query initialization and block boundary placement in the training architecture.

[0095] In an embodiment, the content identity infrastructure extends to the training level of generative model development, enabling variance-derived structural properties of training data to govern training corpus admission, curriculum ordering, batch composition, and provenance tracing at the weight level. This section discloses methods for integrating the multi-axis variance vector extraction pipeline and slope-band indexing architecture with neural network training procedures, producing a governed training regime in which the structural identity of training data is a first-class input to the training process rather than an extrinsic metadata annotation.

[0096] Variance-governed curriculum ordering assigns each training data artifact a variance band classification based on its global variance value and variance vector profile, as described in Section 5. Training batches are composed by sampling artifacts in order of ascending variance band during initial training phases, presenting low-variance, structurally simple artifacts before high-variance, structurally complex artifacts. This variance-ordered curriculum mirrors the empirical observation that models trained on progressively increasing structural complexity exhibit more stable early-phase gradient distributions and more uniform convergence across depth. The variance band classification of each training artifact is derived deterministically from its structural properties and requires no human annotation, enabling fully automated curriculum construction from any corpus whose artifacts have been processed by the multi-axis variance vector extraction pipeline.

[0097] Slope-band batch composition governs the variance profile of each training batch to maintain a target distribution across variance bands throughout training. Rather than sampling training data uniformly or purely by curriculum stage, the batch composition module evaluates the current variance band distribution of the training loss surface and adjusts batch sampling weights to preferentially admit artifacts from bands where the model's current loss is elevated. This dynamic batch composition uses the model's per-band validation loss as a feedback signal for sampling, analogous to priority experience replay in reinforcement learning but operating over the structural variance space of the training corpus rather than over a replay buffer of interaction histories.

[0098] The structural provenance trace provides a method for evaluating whether a trained generative model has memorized specific training data artifacts or generalized from their structural features. For each artifact in the training corpus, the system computes the cosine similarity between the artifact's variance vector and the variance vectors of outputs generated by the trained model when prompted with semantically related queries. A high cosine similarity between a generated output's variance vector and a training data artifact's variance vector, combined with a similarity score exceeding the policy-declared threshold, indicates that the model's output is structurally proximate to the training artifact in variance space. This structural proximity measurement does not require access to model weights or activation patterns; it operates entirely over the variance-derived UIDs of training artifacts and generated outputs, enabling provenance evaluation at inference time without model introspection.

[0099] The depth-wise content attention method integrates structural variance signals from training data with the depth-wise aggregation mechanism of neural network architectures that employ learned, input-dependent weighting of preceding layer representations. In such architectures, the pseudo-query vector associated with each layer governs the weight distribution over preceding layer outputs. The present invention discloses initializing or adapting these pseudo-query vectors using variance-vector-derived features of the training batch, such that layers processing high-variance training artifacts assign different depth-wise aggregation weights than layers processing low-variance artifacts. This integration allows the depth-wise attention mechanism to be sensitive to the structural complexity of the content being processed, enabling adaptive depth allocation as a function of content variance rather than as a fixed architectural property. The variance band of the training batch may further be used to modulate the block boundary placement in block-partitioned depth-wise architectures, dynamically adjusting block granularity based on the structural complexity of the current training distribution.

[0100] The training-level governance methods are enabled by the multi-axis variance vector extraction pipeline and slope-band indexing architecture disclosed in the preceding sections, which together provide the structural identity measurements, variance band classifications, and cosine similarity operators required for variance-governed curriculum construction, dynamic batch composition, structural provenance tracing, and depth-wise content attention integration.

14. Adversarial Robustness, Deepfake Detection, and Screenshot Recapture

[0101] In an embodiment, the candidate artifact input is processed by the variance vector extractor, which produces the multi-axis variance vector. The lineage query module queries the anchor network for registered parent UIDs within the slope continuity threshold and returns a lineage match result. The orphan detector flags artifacts with no registered lineage as structurally unanchored. The screenshot recapture classifier evaluates the Z-axis gradient histogram component for characteristic screen-capture variance signatures and produces a recapture probability score. The synthetic content detector compares the candidate variance vector against a generative model output distribution and produces a synthesis probability score. The composite risk score aggregator combines lineage absence, recapture probability, and synthesis probability into a governance signal that routes to the pre-release admissibility engine.

[0102] The adversarial robustness architecture of the present disclosure exploits a structural property that distinguishes synthetically generated content from photographically or digitally captured content: a generatively synthesized artifact has no structural lineage to any prior registered artifact in the anchor network. When an image, audio clip, or video sequence is generated by a diffusion model, a generative adversarial network, or a language model operating over visual tokens, the resulting artifact's variance vector position in slope space reflects the statistical properties of the generative model's output distribution rather than the variance profile of any specific prior artifact. The lineage query module detects this condition by querying the anchor network for registered parent UIDs within a configured slope continuity radius of the candidate's UID. If no registered parent UID falls within the continuity radius, the orphan detector classifies the artifact as structurally unanchored, meaning it has no provable lineage connection to any content registered in the governed corpus. Structurally unanchored artifacts are not necessarily fraudulent or impermissible, but they cannot be admitted under a policy object that requires verifiable provenance, and they trigger heightened scrutiny under policy objects that govern synthetic content.

[0103] The screenshot recapture detection method exploits a characteristic variance signature introduced when a digital display renders an image and a camera or screen-capture device re-captures the rendered output. Screen rendering introduces a periodic spatial frequency structure in the luminance channel attributable to the sub-pixel geometry of the display, compression and dithering artifacts from the display pipeline, and the optical point-spread function of the capturing lens or sensor. These artifacts manifest in the Z-axis gradient histogram component of

the re-captured artifact as elevated energy in the horizontal and vertical orientation bins relative to the diagonal bins, producing a horizontal-vertical bias score in the Z-axis vector that is systematically elevated compared to the original digital artifact. The screenshot recapture classifier evaluates this Z-axis horizontal-vertical bias against a policy-calibrated threshold and produces a recapture probability score. This detection method requires no reference to the original artifact and operates entirely from the structural features of the candidate artifact itself, enabling recapture detection without corpus lookup.

[0104] The synthetic content detector operates by comparing the candidate artifact's variance vector against a generative model output distribution represented as a slope-band-indexed statistical model of the variance vector profiles of known synthetic content. When a candidate artifact's variance vector falls within the high-probability region of the synthetic content distribution and outside the high-probability region of the authentic content distribution for the relevant content category, the detector produces an elevated synthesis probability score. The synthetic content distribution may be constructed empirically from samples of generative model outputs registered in a governed reference corpus, and may be updated continuously as new generative model architectures produce artifacts with distinct variance signatures. The system thus adapts to evolving generative model outputs without requiring retraining of an inference model or deployment of new classification infrastructure.

15. Client-Side Execution Architecture

[0105] In an embodiment, the content input stage receives a file object or media stream from a standard browser file input or media capture API. The Canvas 2D API normalization module performs canonical resizing, grayscale conversion, and orientation canonicalization using only the standard HTMLCanvasElement and CanvasRenderingContext2D interfaces available in any conforming browser without WebGL, WebAssembly, or server-side dependency. The client-side variance vector computation module performs multi-scale variance analysis, gradient histogram computation, and 27-dimensional vector construction using standard JavaScript arithmetic operations. The client-side hash module produces a 320-bit UID by applying the multi-scale FNV-variant hash combiner in standard JavaScript. The local similarity evaluation module computes cosine similarity against a locally cached exclusion corpus fragment and evaluates the result against a locally stored policy object. The anchor query module transmits only the

computed UID to the anchor network for corpus-scale resolution rather than transmitting the raw artifact.

[0106] In an embodiment, the complete pipeline from content input through UID computation, local similarity evaluation, and anchor query dispatch executes within a standard browser execution context without server-side infrastructure, GPU compute, or per-query API service. The raw content artifact does not leave the client device during the admissibility evaluation phase; only the computed UID and the resulting admissibility decision are transmitted. This architecture admits client-side admissibility evaluation at upload time, conforms to data minimization requirements that restrict transmission of personal media content, and avoids per-query inference costs proportional to content volume.

[0107] The locally cached exclusion corpus fragment is a slope-band-filtered subset of the full governed corpus, pre-fetched from the anchor network for the variance bands most likely to be relevant to the content categories the client is authorized to process. The locally stored policy object is a versioned, signed policy object pre-fetched from the governing authority under which the client operates. Both the corpus fragment and the policy object are verifiable by their cryptographic signatures without requiring a live connection to the anchor network at evaluation time. This enables client-side admissibility evaluation in disconnected or intermittently connected environments, with the assurance that the policy object and corpus fragment are authentic and unmodified.

16. UID Resolution Query Protocol

[0108] In an embodiment, the querying client computes a candidate UID and a variance band from the candidate artifact using the extraction pipeline. A band-targeted query is dispatched to the primary anchor cluster governing the identified variance band. The anchor lookup module searches the band's UID index for slope-proximate entries and returns a ranked candidate match list with cosine similarity scores. The policy validation stage evaluates each candidate match against applicable policy objects and filters the results to policy-permitted entries. If no match is found in the primary band, the cross-band referral module dispatches adjacency-ordered referral queries to neighboring band clusters in order of variance proximity. The result aggregator consolidates ranked matches from primary and referred band queries and returns a final

resolution response comprising matched UIDs, cosine similarity scores, lineage annotations, and policy constraints.

[0109] The UID resolution query protocol is the commercial interface layer through which all platform and application integrations interact with the anchor network. A querying party submits a candidate UID--computed locally from a candidate artifact without transmitting the artifact itself--and receives in return a resolution response containing all policy-permitted matches whose variance vectors fall within the configured proximity radius of the query UID. The protocol defines four resolution modes: identity resolution, which returns exact or near-exact matches indicating the candidate is a known registered artifact; derivative resolution, which returns matches with cosine similarity between the policy-declared continuity threshold and the identity threshold, indicating the candidate is a probable derivative of a known artifact; orphan resolution, which returns an empty match set indicating the candidate has no registered lineage within the governed corpus; and conflict resolution, which returns multiple matches with overlapping lineage claims and routes to the fork adjudication handler of Section 7.

[0110] The resolution protocol is stateless from the querying client's perspective: each query carries the candidate UID, the querying party's policy scope identifier, and a timestamp, and receives a complete resolution response without requiring any prior session establishment or registration by the querying party. Anchor nodes verify the policy scope identifier against the governing policy object for the relevant band and return only those resolution results that fall within the querying party's authorized scope. This stateless, policy-scoped resolution design enables the protocol to be implemented over any transport layer including HTTP, WebSockets, WebRTC, and delay-tolerant mesh protocols, and supports deployment in disconnected environments where queries are batched and resolved upon reconnection to the anchor network.

[0111] The resolution protocol further supports bulk resolution, in which a querying client submits a batch of candidate UIDs from multiple variance bands in a single request. The anchor network routes each UID in the batch to its corresponding band cluster in parallel, aggregates the results, and returns a bulk resolution response whose entries are indexed to the submitted UIDs. Bulk resolution enables platforms processing large volumes of content--e.g., upload pipelines, content moderation systems, and training data ingestion pipelines --to resolve the identity and admissibility status of thousands of artifacts per second without issuing individual queries per

artifact, reducing per-artifact resolution latency to the latency of a single network round-trip amortized over the batch size.

17. Model Output Provenance Fingerprint

[0112] In an embodiment, the generative model output is processed by the output variance vector extractor, which applies the full multi-axis extraction pipeline to the generated artifact without regard to the generative model that produced it. The output UID computation module produces a slope-indexed output UID. The anchor network proximity query queries the training corpus anchor index for training artifacts whose variance vectors fall within a configured proximity radius of the output UID. The proximity match evaluator computes cosine similarity between the output variance vector and each candidate training artifact variance vector and applies the policy-declared memorization threshold. When the threshold is exceeded, the memorization signal generator produces a memorization proximity score and a ranked list of structurally proximate training artifacts. The provenance report module generates an auditable provenance record comprising the output UID, matched training artifact UIDs, cosine similarity scores, and the applicable policy version.

[0113] The model output provenance fingerprint method provides a mechanism for determining whether a generative model's output is structurally proximate to specific training data artifacts without requiring access to model weights, activation patterns, gradient information, or training logs. The method is deployable by any party that possesses the output artifact and access to the training corpus anchor index, including rightsholders evaluating potential infringement by a model they did not train, regulators auditing model output for training data scope compliance, and model operators conducting pre-release provenance checks. The method does not require model introspection, membership inference attacks, or the cooperation of the model operator, because it operates entirely over the structural features of the generated output and the indexed variance vectors of training artifacts.

[0114] The memorization proximity score is computed as a function of the cosine similarity between the output variance vector and the most similar training artifact variance vector, weighted by the variance band of the output and calibrated against the policy-declared memorization threshold for the relevant content category. A score above the threshold does not constitute a legal determination of infringement; it constitutes a structural proximity signal that

may inform further investigation, rights enforcement action, or compensation computation under the creator attribution architecture of Section 18. The provenance record generated by the provenance report module is versioned, cryptographically signed, and appended to the anchor network's event log, making it available as evidence in legal, regulatory, or contractual proceedings.

[0115] The model output provenance fingerprint architecture supports both individual artifact evaluation and batch evaluation over a corpus of generated outputs. Batch evaluation over a corpus of model outputs allows statistical characterization of a model's structural proximity to its training data: the distribution of memorization proximity scores across a representative output corpus indicates whether the model exhibits systematic structural proximity to specific training data regions, suggesting memorization of content in those regions, or whether proximity scores are uniformly distributed, suggesting generalization from structural features without memorization of specific artifacts. This statistical characterization is a novel application of the variance vector infrastructure that enables quantitative, auditable assessment of training data influence on model outputs at scale.

18. Creator Attribution and Compensation Routing

[0116] In an embodiment, a content artifact with a registered alias has an alias record comprising the alias string, the variance-derived UID, a compensation routing field encoding a payment address, a compensation schedule reference, and a jurisdictional scope. The consultation event log stores consultation records comprising the consulted artifact UID, the governing policy version, the variance proximity score, and a timestamp for each generation event that consulted the artifact. The attribution computation module aggregates consultation events per consulted artifact and computes an attribution weight as the product of consultation frequency and mean variance proximity score. The compensation computation module multiplies the attribution weight by the schedule rate from the compensation schedule reference and produces a payment obligation record. The payment routing module credits the computed obligation to the payment address associated with the consulted artifact's alias record under the declared jurisdictional scope. The compensation audit log appends each payment obligation record for regulatory and dispute resolution purposes.

[0117] The creator attribution and compensation routing architecture operationalizes the principle that content creator compensation for generative AI use of training data should be computable from verifiable structural proximity measurements rather than from approximations of training data influence derived from model weight analysis. The architecture requires three pre-conditions: the training data artifact must have a registered UID in the anchor network; the creator must have registered an alias for that UID with a compensation routing field specifying a payment address and compensation schedule; and the generative model must operate within a governed execution environment that logs consultation events when the model queries reference artifacts through retrieval-augmented generation or structured neighborhood resolution during inference.

[0118] The compensation schedule reference is a versioned, machine-evaluable, cryptographically signed document specifying the rate structure under which attribution weights translate to payment obligations. Schedule formats include per-consultation flat rates, proximity-weighted rates in which higher cosine similarity scores produce proportionally higher payments, category-specific rates, and jurisdiction-specific rates. Schedule versioning ensures compensation computations are reproducible and auditable from the schedule version in effect at the consultation event.

[0119] The payment address in the compensation routing field may designate a bank account routing number, a digital payment service identifier, a blockchain wallet address, an internal ledger account within a platform payment system, or any other machine-resolvable payment endpoint. The architecture is payment-infrastructure-agnostic: the routing field specifies the destination and the schedule specifies the amount; the payment routing module resolves the destination format and executes the credit through the appropriate payment interface, ensuring the architecture remains operable as payment infrastructure evolves.

[0120] The compensation audit log maintains an append-only record of payment obligation records, comprising the consulted artifact UID, the computed attribution weight, the schedule version applied, the payment obligation amount, the payment address credited, the jurisdictional scope, and a timestamp. The log is verifiable against the consultation event log: any party may confirm that payment obligations are consistent with the corresponding consultation events, the

compensation schedules in effect at the relevant times, and the computed attribution weights — without requiring access to model weights, training logs, or proprietary platform data.

19. Terminology

[0121] As used herein, "Unique identifier" or "UID" means a deterministically computed digital token derived from the internal variance structure of a content artifact, encoding a multi-axis variance vector and optionally one or more supplementary structural fingerprints, and serving as the primary identity reference for the artifact within the content anchoring platform. A UID is not a static hash of the SHA-256 or MD5 class; it is a slope-bearing identifier whose numerical components encode the artifact's position in variance space and whose similarity to other UIDs is directly computable without decoding a black-box digest.

[0122] As used herein, "Variance vector" means a multi-dimensional numerical vector whose components encode variance-based proxy values extracted from a digital content artifact at multiple spatial scales and across multiple structural dimensions including energy distribution, frequency compaction, and gradient orientation.

[0123] As used herein, "Slope band" or "variance band" means a bounded contiguous range of variance-vector space within a global slope continuum, used as the basis for assigning anchor node governance responsibility to content UIDs whose variance vectors fall within the range. Slope bands are quantized segments of variance space and serve as the routing unit for UID registration, alias resolution, and cache propagation within the anchor network.

[0124] As used herein, "Anchor node" means a processing node within the content identity platform that stores and governs content UIDs, alias registrations, lineage graphs, and policy constraints for artifacts whose variance-derived slope vectors fall within one or more designated slope bands. An anchor node is not a fixed network endpoint; it is a governance entity that may be instantiated and that derives its routing responsibility from its declared band scope.

[0125] As used herein, "Cosine similarity" means the inner product of two vectors divided by the product of their magnitudes, yielding a value in the range [-1, 1] where 1 indicates identical direction, 0 indicates orthogonal structure, and -1 indicates opposite structural orientation. Cosine similarity is used herein as a measure of directional alignment between variance vectors rather than as a measure of magnitude equality.

[0126] As used herein, "Constellation signature" means a geometric fingerprint derived from the relative positions and angular relationships of high-salience spatial anchor points detected within a content artifact, computed in a manner invariant to translation, rotation, and uniform scale change. The constellation signature is distinct from and supplementary to the variance-vector-based UID; it encodes spatial relational geometry rather than distributional variance.

[0127] As used herein, "Structure signature" means a gradient-only fingerprint derived from the orientation distribution and edge density of a content artifact without reference to mean luminance or background fill values. The structure signature is stable across background color changes, flat-fill variations, and format conversions that alter pixel values without altering edge structure.

[0128] As used herein, "Mutation" means any transformation applied to a content artifact that produces a new artifact with a distinct variance vector, including editing, recomposition, cropping, style transfer, compression, resizing, re-encoding, format conversion, remixing, and generative synthesis. Transformations within the quantization tolerance of the hash function produce the same UID and are not treated as distinct mutations for governance purposes.

[0129] As used herein, "Provenance" means the record of the origin, ownership, mutation history, lineage relationships, and alias registrations associated with a content artifact UID, as maintained in anchor node memory and traversable through the directed lineage graph.

[0130] As used herein, "Policy object" means a versioned, machine-evaluable, cryptographically signed structured artifact specifying any of: admissible content categories, restricted content classes, jurisdictional scopes, similarity tolerance thresholds, retention periods, propagation scopes, override authorities, escalation paths, delegation rights, and compensation schedules. Policy objects may be evaluated at admission, replication, eviction, alias registration, alias resolution, and pre-release admissibility events; their evaluation results are reproducible from the policy object version and the artifact's variance-derived UID and structural signatures.

[0131] As used herein, "Trust zone" means a bounded set of anchor nodes operating under a shared governance framework defined by a versioned, cryptographically signed trust-zone policy object that specifies admissible policy objects, anchor recruitment rules, propagation scope, and quorum threshold profiles for the trust zone. Trust zones may be nested and federated, and a single anchor node may participate in multiple trust zones with distinct authorities in each.

[0132] As used herein, "Adaptive Consensus Protocol" means a trust-weighted, asynchronous, lineage-preserving consensus mechanism for anchor coordination, comprising (i) per-anchor vote weights derived from declared band scope, historical reliability, and trust-zone authority; (ii) per-mutation quorum thresholds configurable by mutation type; and (iii) commitment records preserving the predecessor state, participating anchor signatures, and policy version under which each mutation was evaluated, admitting later replay verification.

[0133] As used herein, "Consultation event" means an event in which a generative model queries a reference artifact through retrieval-augmented generation, structured neighborhood resolution, or analogous mechanism, recorded as a deterministic record comprising the variance-derived UID of the consulted artifact, the governing policy object, the variance proximity score between the consulted artifact and the generation output, and a timestamp.

[0134] As used herein, "Commitment" or "commitment boundary" means an irreversible or externally visible side effect of a content generation or distribution event, including public release, customer delivery, API return, licensing event, marketplace publication, training data admission, or cross-platform provenance anchor registration. The pre-release admissibility engine interposes admissibility evaluation at the commitment boundary.

[0135] As used herein, "Governed corpus" means a slope-band-indexed reference corpus of registered UIDs governed by one or more signed corpus policy objects, against which candidate artifacts are evaluated for admissibility, similarity, exclusion, or memorization proximity.

[0136] As used herein, "Exclusion corpus" means a governed corpus comprising variance-indexed forbidden content references, against which candidate artifacts are evaluated for proximity in the pre-release admissibility pipeline.

[0137] As used herein, "Fork adjudication" means a governance procedure invoked when a resolution query reveals conflicting alias claims across slope-divergent UIDs, comprising anchor quorum review of registration timestamps, variance proximity to the canonical UID, policy lineage, and trust-zone authority, with quorum-recorded outcomes accessible to querying parties as provenance metadata.

[0138] As used herein, "Memorization proximity score" means a function of cosine similarity between a generative model output's variance vector and the most similar training

artifact variance vector within the governed corpus, weighted by the variance band of the output and calibrated against a policy-declared memorization threshold for the relevant content category.

[0139] As used herein, "Near real-time" or "real time" describes a process that produces a given result with a slight but acceptable delay between an event (e.g., data acquisition or update) and the result. In the present disclosure, a slight but acceptable delay is in the range of about 250 milliseconds.

[0140] As used herein, "About" with reference to a value or range is used in its plain and ordinary sense as understood by persons of ordinary skill in the art, referring to standard tolerances for the referenced parameter; when standard tolerances are not applicable, a value or range defined with "about" is met when a change does not alter the relevant performance characteristics by more than five percent (5%).

[0141] The disclosed system and methods may be embodied in software on a non-transitory computer-readable medium executed by one or more processors, deployed on a stand-alone computer, an intranet-accessible server, or an Internet-accessible server.