Nick Clark · Published August 16, 2026 · PDF

Mixpeek: Multimodal Retrieval and Structural Content Identity

Multimodal retrieval platforms and structural content identity systems begin from the same practical question: given an enormous pool of video, image, audio, and document material, how does a system find the right item and say what it is related to? Mixpeek, as publicly described, is a multimodal data warehouse and retrieval platform whose public materials describe managed indexing, typed feature extractors, and multi-stage retrievers running over a vector store. The disclosure filed as PCT/US26/28630 describes a different starting point, deriving an artifact identifier from a multi-axis variance vector whose geometry makes cosine similarity computable directly between identifiers, and interposing a pre-release admissibility engine that evaluates a candidate artifact against cryptographically signed policy objects before commitment. Retrieval quality and admissibility governance are separable engineering problems, and an operator running content at scale can reasonably want both.

Two Questions That Look Like One

Anyone who has operated a large media library has met both of the following questions, usually on the same afternoon. One is a retrieval question: out of several million frames, transcripts, and pages, which items answer this query, and which are near

neighbors of the one I already have? The other is a commitment question: this candidate artifact is about to be published, delivered to a customer, or admitted into a training corpus, and someone must decide whether that is permitted, on a record that will still make sense to an auditor a year from now.

They look like one question because both are answered by comparing content against other content. Retrieval optimizes for ranking quality against a user's intent, and its failure mode is a mediocre result list. Commitment optimizes for a defensible decision against a declared rule, and its failure mode is an artifact that has already been released. The filed disclosure frames the second problem in terms of where the evaluation sits. It describes post-generation moderation filters, applied after an output artifact has been produced, as unable to prevent impermissible content from existing as an internal artifact and as not yielding determinations reproducible and auditable from versioned policy records. The architecture it describes therefore requires the evaluation to be interposed between generation and commitment, and requires the resulting decision to be replayable from a versioned record.

How Mixpeek Describes Its Platform

Mixpeek is publicly positioned as a multimodal data warehouse and retrieval platform for video, image, audio, and document search. Named products in its public materials include a vector store supporting dense, sparse, and BM25 search on object storage, and managed indexing that extracts scenes, faces, OCR, transcripts, and embeddings from raw files. That is a substantial piece of engineering: the awkward part of multimodal search is rarely the nearest-neighbor lookup, it is getting from source media to typed, queryable features without hand-building a pipeline for every modality.

Building blocks described publicly include retrievers performing multi-stage filter, join, and rerank; typed feature-extractor pipelines for faces, scenes, transcripts, OCR, and fingerprints; and support for many embedding models. Named extractor models in

public materials include SigLIP for image embeddings, E5-Large for text embeddings, and SCRFD with ArcFace for face identity. Public materials also describe taxonomies built from clustering output and from model-generated labels, and cross-modal joins.

One detail is worth calling out on its own merits. As publicly described, every extracted feature receives a URI encoding its provenance, enabling tracing from a search result back to the source frame, timestamp, or audio segment. Any team that has debugged a result list knows how much work sits behind a clean answer to where a given hit came from. What follows is not a critique of that objective.

What the Filed Architecture Derives and Gates

The subject matter disclosed in PCT/US26/28630 starts one layer beneath retrieval, at what an artifact's identifier is made of. In an embodiment, a content encoder derives a unique identifier for a digital content artifact by extracting a nine-dimensional multi-axis variance vector, across three axes, from the internal structure of the artifact. The X-axis energy behavior vector encodes cross-scale energy distribution, computed across a coarse grid of 8x8 cells, a medium grid of 16x16 cells, and a fine grid of 32x32 cells. The Y-axis frequency compaction vector encodes how variance dispersion behaves across those same scales. The Z-axis structural phase persistence vector derives from an eight-bin gradient orientation histogram, yielding a horizontal-vertical orientation bias, a diagonal-axial bias, and a stability coefficient.

The property that matters architecturally is what happens next. In described embodiments, the identifier encodes a position within a continuous variance space such that cosine similarity between two unique identifiers is directly computable without decoding a fixed binary digest. Identity and similarity become the same object. Spatial resolution comes from the quadrant pipeline of FIG. 1: canonical normalization (300) rescales the artifact to a square canvas, orientation canonicalization (302) rotates it when the dominant gradient angle exceeds a threshold of approximately 0.1 radians, quadrant extraction (304a) through (304d) feeds per-quadrant variance computation

(306a) through (306d) and hashing (308a) through (308d), a rotation-invariant sorting module (310) orders those hashes lexicographically, and a combined construction module (312) applies a multi-segment FNV-64 hash combiner to produce a 320-bit UID.

Placement in the network follows from the same number. In an embodiment, a variance band classification module assigns each content UID to one of five variance bands by global variance, from Band 1 below 0.02 through Band 5 at or above 0.22, and anchor cluster assignment logic (504) maps the UID to a primary anchor node (506) and secondary anchor nodes (508a) through (508c). A band adjacency graph (510) supports querying adjacent bands in order of proximity, and a quorum-based resolution mesh (512) runs the Adaptive Consensus Protocol over alias registrations. Any processing node that knows the variance value of a target UID can determine which band governs it and route a resolution query to the appropriate anchor cluster without consulting a central directory.

The governance layer sits on that substrate. In an embodiment, a pre-release admissibility engine routes a candidate artifact through two parallel evaluation tracks before a commitment gate. A policy object evaluator receives a versioned, cryptographically signed policy object defining typed category constraints, jurisdictional scopes, override authorities, similarity tolerance thresholds, and escalation paths. A structural similarity evaluator computes cosine similarity between the candidate's variance vector and reference artifacts in a governed corpus indexed in the slope-band anchor network, feeding a forbidden content exclusion layer that compares the score against the policy-declared threshold. If both tracks confirm admissibility the gate permits a committed artifact output; if either fails, a rejection handler produces a regeneration or escalation signal. A commitment, in the disclosure's own terms, is any irreversible or externally visible side effect of a content generation or distribution event, including public release, customer delivery, API return, or training data admission.

Two adjacent components complete the picture. A training corpus governance layer admits artifacts only under signed corpus policy objects, and each admitted artifact receives a governance record comprising its variance-derived UID, the governing policy object, a timestamp, and a cryptographic hash of that policy object. A consultation event logger records a deterministic consultation record for each generation event that consults a reference artifact through retrieval-augmented generation or structured neighborhood resolution. Outcomes stay conditioned on declared bounds: where the similarity score between a candidate and an exclusion corpus entry exceeds the policy-declared threshold, the candidate is rejected, regenerated under modified generation constraints, or escalated to an authorized override authority.

Convergent Category, Divergent Architecture

The convergence is genuine. Both architectures index heterogeneous media, both compute similarity between items, and both treat tracing a result back to its source as a first-class concern. Mixpeek's public materials describe per-feature provenance URIs reaching back to a frame, timestamp, or audio segment. The filed disclosure describes multi-root lineage graphs in which a derivative UID (804) links to parent UIDs (800) and (802) through a variance vector comparison module (806), a slope delta computation module (808), and weighted parent attribution (810), with weights recorded as edge annotations that the disclosure states do not constitute legal determinations of authorship or ownership.

The divergence is in what the identifier is for. A retrieval architecture built on embedding models places semantically related items near each other in a learned space, which is what a search product needs. The architecture in the filed disclosure derives its coordinate from the artifact's own multi-scale variance statistics, which yields a different property set: the vector is computed rather than learned, one 320-bit identifier serves as both routing key and similarity operand, and client-side execution is described in a conforming browser using the Canvas 2D API and standard JavaScript arithmetic, with the computed UID transmitted rather than the raw artifact.

A second axis of divergence is where a decision is enforced. The filed architecture requires a signed, versioned policy object to be evaluated before the commitment gate opens, and requires the determination to be verifiable by an authorized party replaying the evaluation from the artifact's variance-derived UID and structural signatures together with the policy object version. That is an evidentiary requirement rather than a ranking one, and it typically originates in an operator's compliance posture.

Coexistence and Open Edges

A plausible deployment runs both. Managed indexing and multi-stage retrievers, as publicly described, turn raw video, image, audio, and document files into typed features and answer queries over them. The identity and admissibility layer sits at the boundaries, at upload and at release. On ingestion, an artifact's variance vector and 320-bit UID are computed and registered to the governing variance band. At the commitment boundary, the pre-release admissibility engine evaluates the candidate against the governed exclusion corpus and the applicable signed policy object. The disclosed resolution protocol further supports bulk resolution, in which the network routes each UID in a submitted batch to its band cluster in parallel and returns a response indexed to the submitted UIDs, the shape an upload pipeline needs. Retrieval answers what an operator should show a user; admissibility answers what the operator is permitted to commit.

It is equally worth naming the edges of the disclosed architecture itself. Its similarity operator is structural rather than semantic: two artifacts with similar multi-scale variance geometry score as proximate whether or not a reader would call them about the same subject. The disclosure describes identity derivation, anchor governance, admissibility, and lineage. Its constellation signature requires at least three saliency hotspots, and where fewer than three are detected the signature is not emitted. Its stability under transformation is described as holding for format conversion, resolution

rescaling, and lossy compression within defined thresholds, a bounded statement rather than an unconditional one. Nothing in the architecture makes a policy decision correct; it makes the decision reproducible against the policy object in force.

Disclosure Scope

This article is a technical description of subject matter disclosed in PCT/US26/28630, "Structural Content Identity and Rights-Grade Admissibility for Digital Artifacts," which claims priority to U.S. Provisional Application Serial No. 63/808,372. Architectural statements above are grounded in that filing and use its own mechanism names and reference numerals. Features are described as embodiments. The applications referenced are pending, and descriptions of behavior are conditioned on the thresholds and signed policy objects the filing specifies.

References to Mixpeek are to public materials and are used for comparison only; no relationship, endorsement, or infringement is asserted. Descriptions of Mixpeek reflect publicly available product and documentation materials reviewed on August 5, 2026. Nothing above asserts that Mixpeek practices, overlaps with, or was influenced by the subject matter disclosed in the cited filing, or says anything about the relative timing of any party's work.

Content Anchoring (/content-anchoring) All 49 steps → (/inventive-steps)

Computable identity for media. Provenance from structural variance.

PCT/US26/28630 (/patents/pct-us26-28630)

PRIMARY TECHNICAL DISCLOSURE

– Content Anchoring: Computable Identity for Media That Changes (/articles/content-anchoring-computable-identity-for-media-that-changes)

SECONDARY TECHNICAL

- Multi-Axis Variance Vector Extraction: Nine Dimensions of Structural Content Identity (/articles/content-anchoring/variance-vector)
- Quadrant Decomposition: Spatial Sub-Region Fingerprinting for Partial Similarity Detection (/articles/content-anchoring/quadrant-decomposition)
- 320-Bit UID Construction: Multi-Segment Hashing for Negligible Collision Probability (/articles/content-anchoring/uid-construction)
- Structure Signature: Background-Invariant Matching Through Gradient-Only Descriptors (/articles/content-anchoring/structure-signature)
- Constellation Signature: Geometry-Invariant Matching Across Crop, Scale, and Occlusion (/articles/content-anchoring/constellation-signature)
- Five-Band Variance Classification: Content Routing by Structural Complexity (/articles/content-anchoring/variance-classification)
- Variance Saturation-Governed Cache Eviction: UID Density Replacing Static TTL (/articles/content-anchoring/cache-eviction)
- Multi-Root Composite Lineage Graphs: Provenance Through Variance Vector Similarity (/articles/content-anchoring/composite-lineage)
- Multi-Modal Content Identity: Unified Pipeline Across Image, Audio, Text, and Video (/articles/content-anchoring/multi-modal-identity)
- Rights-Grade Pre-Release Admissibility: Policy Evaluation Before Content Commitment (/articles/content-anchoring/pre-release-admissibility)
- Training Corpus Governance: Verifiable Lineage From Training Data to Model (/articles/content-anchoring/training-corpus-governance)
- Consultation Event Logging: Deterministic Records of Every Generation Reference (/articles/content-anchoring/consultation-logging)
- Model Output Provenance Fingerprint: Structural Proximity Without Model Access (/articles/content-anchoring/output-provenance)
- Creator Attribution and Compensation Routing: Payment From Consultation Lineage (/articles/content-anchoring/creator-attribution)
- Adversarial Robustness and Deepfake Detection: Content Identity as Detection Substrate (/articles/content-anchoring/adversarial-robustness)
- Client-Side Execution Architecture: Privacy-Preserving Variance Computation on Device (/articles/content-anchoring/client-side-execution)
- UID Resolution Query Protocol: Distributed Lookup Across Anchor Node Networks (/articles/content-anchoring/uid-resolution)
- Orientation Canonicalization: Rotation-Invariant Processing Through Gradient Normalization (/articles/content-anchoring/orientation-canonicalization)

- Cross-Band Resolution Pathfinding: Traversal Between Variance Bands Under Mutation (/articles/content-anchoring/cross-band-resolution)
- Identity by Position: Media as a Third Navigable Space (/articles/content-anchoring/identity-by-position)

APPLICATIONS · GENERAL

- Forbidden-Content Blocking at Upload and Generation Time: Pre-Release Exclusion Against Signed Policy (/articles/content-anchoring/forbidden-content-blocking)
- Structural Provenance for Software Supply Chains: Binary and Firmware Identity Independent of SBOM Metadata (/articles/content-anchoring/software-supply-chain-provenance)
- Rights-Grade Generative AI: How to Pay Creators, Exclude Forbidden Content, and Prevent Infringement Before Release (/articles/content-anchoring/rights-grade-generative-ai)
- Deepfake Detection by Structural Provenance: Verifying Synthetic Media Without Watermarks (/articles/content-anchoring/deepfake-provenance)
- Creator Economy Attribution Without Platform Intermediaries (/articles/content-anchoring/creator-attribution-economy)
- Verifying Source Photos and Video in the Newsroom: Content Anchoring for Journalism (/articles/content-anchoring/journalism-verification)
- Detecting Image Manipulation and Proving Figure Provenance in Research Publications (/articles/content-anchoring/academic-research-integrity)
- Content Anchoring for Legal Evidence Chains (/articles/content-anchoring/legal-evidence-chain)
- Content Anchoring for Insurance Claims Evidence (/articles/content-anchoring/insurance-claims-evidence)
- Content Anchoring for Real Estate Documentation (/articles/content-anchoring/real-estate-documentation)
- Art Authentication and Provenance Verification with Content Anchoring (/articles/content-anchoring/art-authentication)
- Detecting Screenshot and Recapture Fraud in Identity-Document KYC With Structural Content Identity (/articles/content-anchoring/identity-document-kyc-recapture)
- When the Same Work Returns Wearing a New Hash (/articles/content-anchoring/content-anchoring-stakes)

APPLICATIONS · SPECIFIC

- C2PA vs Content Anchoring: Attached Provenance or Content-Intrinsic Identity? (/articles/content-anchoring/c2pa)
- Google SynthID Alternative: Content-Intrinsic Identity Beyond Watermarking (/articles/content-anchoring/google-synthid)

- Beyond Shutterstock: Content-Intrinsic Identity That Survives Re-Encoding and Cropping (/articles/content-anchoring/shutterstock)
- Spotify Alternative for Music Provenance: Structural Content Identity Beyond the ISRC Database (/articles/content-anchoring/spotify)
- Getty Images Alternative for Provenance: Structural Content Identity Beyond Metadata (/articles/content-anchoring/getty-images)
- Adobe Stock vs Structural Content Identity: Licensing Records Are Not Content Identity (/articles/content-anchoring/adobe-stock)
- YouTube Content ID vs Content Anchoring: Matching Against a Database, or Identity in the Content Itself (/articles/content-anchoring/youtube-content-id)
- Audible Magic Alternative: Structural Content Identity Beyond Database-Matched Fingerprinting (/articles/content-anchoring/audible-magic)
- Digimarc vs Structural Content Identity: Watermarks Are Added, Not Intrinsic (/articles/content-anchoring/digimarc)
- Irdeto vs Structural Content Identity: DRM Protects the Channel, Not the Payload (/articles/content-anchoring/irdeto)
- Truepic alternative: capture-time provenance versus structural identity derived from the artifact itself (/articles/content-anchoring/truepic)
- Microsoft PhotoDNA vs structural content identity: hash-matching known images versus screening artifacts before release (/articles/content-anchoring/microsoft-photodna)
- Pex alternative: structural content identity vs enrolled fingerprint matching (/articles/content-anchoring/pex)
- **Mixpeek: Multimodal Retrieval and Structural Content Identity (/articles/content-anchoring/mixpeek)**

HOW-TO GUIDES

- How to Attribute AI-Generated Content to Its Source for Licensing (/articles/guides/attribute-ai-generated-content-for-licensing)
- How to Detect a Screenshotted or Re-Photographed Document (/articles/guides/detect-recaptured-document)
- How to Detect AI-Generated Images Without a Watermark (/articles/guides/detect-ai-images-without-watermark)
- How to Verify Content Authenticity Without a Central Registry (/articles/guides/verify-content-authenticity-without-registry)

TERMINOLOGY

- constellation signature (/glossary#constellation-signature)
- fork adjudication (/glossary#fork-adjudication)
- governed corpus (/glossary#governed-corpus)
- memorization proximity score (/glossary#memorization-proximity-score)
- slope band (/glossary#slope-band)

Content Anchoring overview → (/content-anchoring)