

AEYES - Visual Content Intelligence Platform

Finally, software that actually SEES your visual content

The User Story: From 50,000 Lost Images to Instant Discovery

Sarah, a professional AI artist, has 50,000 Midjourney generations scattered across folders named "MJ_Download_1", "MJ_Download_2"... She remembers creating a perfect cyberpunk character months ago but can't find it among thousands of files. With AEYES, she searches "purple cyberpunk woman with neon eyes" and finds it instantly—along with 12 similar variations she'd forgotten existed. What took hours of manual searching now takes seconds, and she rediscovers creative directions she thought were lost forever.

This is why the world needs AEYES: not another file manager, but AI eyes that actually see and understand your creative work.

Phased Expansion: Local → Cloud → Platform → Enterprise

Development Phases: Controlled Growth Strategy

Phase 1: Local Prosumer Application (MVP)

Target: Professional creators with proper hardware

- **Minimum Requirements:** RTX 3070+, 16GB RAM, modern CPU, fast storage
- **No Compromises:** Smooth 60fps rolling interface, full AI processing power
- Rolling physics interface with subdivision scaling (optimized for high-end hardware)
- Image-only analysis and organization with professional-grade CLIP intelligence
- Freemium model with optional ads (targeting prosumer market)
- Export capabilities and data sovereignty (no vendor lock-in)
- **Marketing:** "Professional-grade visual intelligence for serious creators"

Business Model: \$0-25/month, targeting users who value performance over accessibility **Target Users:** AI artists, digital creators, professionals with modern workstations

Phase 2: Enhanced Local + Essential Video Features (Growth)

Target: Strategic video expansion without scope creep

- **Core video features:** Real-time scene segmentation and basic motion vector analysis only
- **Deferred to Phase 3:** Audio fingerprinting, advanced temporal analysis, complex transitions
- **Multi-GPU workstation support:** Distribute tasks across multiple RTX cards
- **3D asset integration:** Real-time previews, texture analysis, polygon counting

- **Streamlined batch processing:** Queue scene detection for thousands of video files
- Platform integration plugins (Automatic1111, ComfyUI, DaVinci Resolve, Blender)
- **Advanced prompt reconstruction** with confidence scoring and uncertainty visualization
- **Team collaboration features:** Bulk tagging, project management, version control
- **Early cloud beta:** Limited sync (10GB) to test infrastructure and user trust

Business Model: \$15-50/month, plugin licensing, enterprise pilot programs **Target Users:** Video creators, 3D artists, studios with RTX 4080/4090 workstations

Phase 3: AEYES Cloud Platform (Universal Access)

Target: Democratize access via cloud infrastructure while maintaining premium local experience

- **Hybrid sync model:** Sensitive files stay local, others sync to AEYES Cloud
- **Universal device support:** Phones, tablets, low-end laptops via cloud streaming
- **Server-side processing:** High-end GPU farms handle all AI computation
- **Responsive streaming interface:** 60fps rolling experience on any device
- **Seamless migration wizard:** Preserve all metadata, GUIDs, and user preferences
- **On-premise deployment:** Enterprise security and compliance options
- **Advanced collaboration:** Real-time metadata syncing with conflict resolution
- **GPU-accelerated analytics:** Real-time visualization of asset usage trends

Business Model: \$20-150/month storage + compute, enterprise contracts **Target Users:** Mobile creators, enterprises, global teams, security-conscious studios

Phase 4: Creative Content Ecosystem (Expansion)


Target: Full creative platform with additional revenue streams

- **Multi-modal intelligence:** Images, videos, 3D assets, audio
- **Marketplace integration:** Stock licensing, NFT minting, print services
- **White-label solutions:** License technology to other platforms
- **Enterprise media management:** Studio and agency solutions
- **API marketplace:** Third-party developers build on our platform

Business Model: \$50-500/month, transaction fees, licensing deals, marketplace cuts

Video Content Revolution:

Rolling Video Interface:



Video Intelligence Capabilities:


- **Frame-by-frame analysis:** "Find videos with blue skies at 0:30"
- **Motion detection:** "Fast-paced action scenes" vs "slow cinematic shots"
- **Audio analysis:** "Videos with dramatic music" or "dialogue scenes"
- **Scene transitions:** "Jump cuts" vs "smooth fades"
- **Color grading analysis:** "Dark/moody" vs "bright/vibrant" videos
- **Duration intelligence:** "Short clips" vs "long sequences"

Video Search Examples:

- "Show me all drone footage"
- "Find videos with explosions"
- "Sunset timelapses with mountains"
- "Close-up portrait shots"
- "Videos with similar color grading to this one"

Interface Design: Simplicity on Surface, Complexity Under Hood

The AEYES Rolling Gallery Interface (Signature Feature with Success Metrics):



Success Metrics for Rolling Interface (MVP Validation):

- **Time-to-first-find:** <30 seconds for specific image in 10k+ collection
- **User satisfaction scores:** >4.5/5 vs traditional grid layouts
- **Browse engagement:** >3x longer session times vs folder browsing
- **A/B testing:** Rolling vs grid control group during beta
- **Pivot trigger:** If rolling doesn't outperform grid by 25%+ on key metrics, switch to enhanced grid with physics-based momentum scrolling

High-Performance Interface Features:

- **4K/8K display optimization** with smooth 60fps rendering
- **Ultra-wide monitor support** for professional creative setups
- **Customizable physics:** Adjustable momentum, rolling speed, carousel layout
- **VR/AR gallery mode:** Immersive browsing with Oculus/Vision Pro support
- **Multi-GPU acceleration:** Dedicated processing for rendering and AI analysis
- **Real-time previews:** 3D assets, video thumbnails with GPU rendering

Interface Elements:

- **Search Box:** Natural language queries ("cyberpunk cats", "dark portraits")
- **Rolling Image Strip:** Click/drag with momentum physics - images roll naturally
- **Subdivision Controls:** Left arrow = bigger images (fewer), Right arrow = smaller (more)
- **90% screen space** = your art, **10%** = interface

Interaction Model:

- Click and drag image strip → smooth rolling with momentum
- Release with velocity → images continue rolling and gradually stop

- Left/Right arrows or arrow keys → subdivide view (1→4→16→grid)
- Mouse wheel → roll through images
- Click any image → full view with hidden complexity revealed

Hidden Complexity:

- All CLIP analysis, confidence scoring, style classification happens invisibly
- AI organizes and surfaces relevant images in the rolling strip
- Complex features accessible but never cluttering the main interface

Platform Integration Business Model (The Big Prize)

B2B SaaS Opportunity:

Midjourney Integration:

- "ArtSight for Midjourney" - cloud-based addon (\$3-5/month)
- Automatically organizes user's Discord downloads
- Enhanced search using their existing CSV exports
- "Find all my fantasy characters" across thousands of generations
- Style evolution tracking and prompt optimization suggestions

Stable Diffusion Platforms:

- Automatic1111 WebUI plugin (premium tier)
- ComfyUI workflow integration
- Invoke AI library management
- Direct integration with generation metadata

Creative Platforms:

- **Adobe Creative Cloud** plugin for video projects
- **DaVinci Resolve** addon for footage organization
- **Blender** asset library enhancement
- **Unreal Engine** texture/material management

The Business Transformation:

Personal Tool → Platform Integrations → Enterprise SaaS
 \$10/month → ... \$50/month → ... \$500/month

Revenue Potential:

- **Midjourney:** 16M+ users × \$3/month = \$48M+ market
- **Adobe CC:** 26M+ subscribers × \$5/month = \$130M+ market
- **Enterprise Video:** Studios pay \$100-500/month for content management
- **API Licensing:** Platform integration fees + revenue sharing

The Vision: ArtSight becomes the **visual intelligence layer** that every creative platform integrates with - like how Stripe became the payment layer for the internet.

Technical Architecture: High-Performance Multi-Modal Intelligence

AEYES Visual Intelligence Engine (Competitive Differentiator):

- **Faster embedding generation:** Custom-optimized CLIP pipeline, 2x speed vs standard implementations
- **Higher accuracy:** Fine-tuned models for AI-generated content with 95%+ style classification accuracy
- **Prompt-aware metadata:** Reconstructs generation parameters with confidence scoring
- **Local model improvement:** User feedback trains personalized classification locally
- **Multi-GPU acceleration:** Parallel processing across RTX cards for studio workstations

Image Analysis (Phase 1 - GPU Optimized):

- **CUDA/Metal acceleration:** Tensor core utilization for RTX cards, Metal optimization for Apple Silicon
- **4K/8K processing:** Native high-resolution support with real-time analysis
- **Color palette extraction** and style classification with confidence scoring
- **Quality assessment** with user feedback loops and local fine-tuning
- **Performance benchmarks:** RTX 4090: ~4,000 images/hour, M2 Ultra: ~2,500 images/hour (Lite Mode)

Essential Video Analysis (Phase 2 - Focused Scope):

- **Core features only:** Real-time scene segmentation and basic motion vector analysis
- **Deferred complexity:** Audio fingerprinting, advanced temporal analysis moved to Phase 3
- **GPU optimization:** Efficient frame sampling and processing for studio workflows
- **Batch processing:** Queue scene detection for large video collections
- **Integration ready:** Works with DaVinci Resolve, Adobe Premiere metadata

Phase 3 Video Expansion:

- **Audio fingerprinting:** Music genre, dialogue vs effects detection
- **Advanced temporal analysis:** Pacing, rhythm, complex transition recognition

- **Cloud processing:** Server-side analysis for mobile and low-end devices

3D Asset Integration (Phase 2-3):

- **Real-time 3D previews:** GPU-rendered thumbnails in rolling interface
- **Texture analysis:** Resolution, quality assessment, material classification
- **Polygon counting:** Complexity analysis for game/film asset management
- **Material detection:** PBR, procedural, hand-painted classification

Enterprise Analytics (Phase 3-4):

- **GPU-accelerated analytics:** Real-time visualization of asset usage trends
- **Bulk processing:** Handle 100k+ assets with multi-threaded analysis
- **Performance benchmarks:** RTX 4090: ~4,000 images/hour, M2 Ultra: ~3,200 images/hour
- **Memory optimization:** Hierarchical indexing for massive collections

The Universal Value Proposition

For Everyone:

- LITE MODE works on ANY system (Windows, macOS, Linux) TODAY
- FULL MODE: Windows/NVIDIA now, Apple Silicon Metal Q2 2025
- Realistic performance expectations set during onboarding
- Just point it at a folder of images and go
- Virtual collections with migration wizard (never moves your files)

For Power Users:

- Enhanced with CSV metadata when available
- Cross-reference visual findings with reconstructed prompts + confidence scores
- Technical parameter correlation and optimization insights
- Interactive style evolution tracking with gamified discovery
- Collaborative libraries and professional portfolio export tools

Enhanced Search Capabilities

Visual Intelligence Search:

- "Show me all portraits with blue eyes" → Visual detection, no prompts needed
- "Find cyberpunk-style images" → Style classification with 2024+ aesthetic awareness
- "Dark moody lighting" → Lighting analysis from pixels
- "Find images similar to this one" → One-click visual similarity matching

- "Show me my space scenes" → Object/environment detection

Smart Query Building:

- Boolean combinations: "portraits, warm colors, NOT anime"
- Natural language parsing with progressive disclosure of advanced features
- Time-based queries: "Show my work from last month"
- Quality filtering: "My best work" (customizable quality metrics)
- Zen mode toggle for professional environments (hides gamification)

Transparent Prompt Reconstruction:

- AI-powered estimates with visual uncertainty indicators ("CFG 7 ± 1 , 61% confidence")
- Colored rings and $\pm$ icons keep expectations realistic without killing magic
- "How did I make this?" reverse engineering with clear probability boundaries
- Manual correction system that improves accuracy over time
- Complete technique analysis: sampler, CFG, LoRA, negative prompts

Gamified Discovery:

- "Rediscovery Badges" for finding forgotten artwork
- "Style Explorer" rewards for portfolio diversification
- Progress tracking and achievement system
- Community challenges and style trend competitions

The Broader Market

This solves problems for EVERYONE who:

- Downloads images from Midjourney Discord (no metadata)
- Uses local Stable Diffusion (inconsistent exports)
- Has years of AI art scattered across folders
- Tries different platforms and loses track of techniques
- Wants to understand their artistic evolution
- Needs to find specific images from vague visual memories
- Works on Mac, Windows, or Linux systems
- Has modest hardware but wants smart organization

Technical Implementation Priority

Phase 1: Core Rolling Interface + Visual Engine

python

Clean interface with powerful backend

ROLLING_INTERFACE:

- .. - Physics-based image carousel with momentum
- .. - Single search box with natural language processing
- .. - Left/right subdivision controls (keyboard shortcuts)
- .. - 90% screen space dedicated to image display
- .. - Minimal UI elements that never compete with art

VISUAL_ENGINE:

- .. - Adaptive memory management for smooth rolling performance
- .. - Real-time image analysis as user browses
- .. - LITE MODE: Universal cross-platform support (Windows, macOS, Linux)
- .. - FULL MODE: Windows/NVIDIA RTX at launch, Apple Silicon Metal Q2 2025
- .. - GUID storage in sidecar files for migration resilience

Phase 2: Smart Organization + Hidden Complexity

python

Invisible organization with visible results

BACKGROUND_INTELLIGENCE:

- .. - Auto-categorization surfaces in rolling strip order
- .. - File-GUID pairs in sidecar files + PNG/JPEG metadata
- .. - Interactive timeline charts (accessible but not prominent)
- .. - Quality-based curation influences rolling order
- .. - Processing summaries available but not intrusive

HIDDEN_RECONSTRUCTION:

- .. - "Beta reconstruction" available on image click
- .. - Parameter estimates with visual uncertainty indicators
- .. - Crowd-sourced confidence model training
- .. - All complexity hidden behind clean single-image view

Phase 3: Advanced Features (Invisible Integration)

python

Power features that don't clutter interface

SMART_ROLLING:

- AI curates rolling order based on search relevance
- Visual similarity influences adjacent image placement
- Natural language search drives rolling content
- Progressive disclosure of advanced features on demand

HIDDEN_WORKFLOWS:

- Data export capabilities (CSV, JSON, ZIP) **in** settings
- Krita/Blender plugins **for** auto-indexing (background operation)
- API access **for** third-party tools (invisible to end user)
- All workflow integration happens without UI clutter

Phase 4: Enterprise Collaboration + Community Features

python

Professional collaboration with concrete functionality

- Multi-user libraries **with** role-based permissions (Viewer, Editor, Admin)
- Real-time collaborative tagging **with** conflict resolution
- Version control **for** metadata changes **with** audit trails
- Team workspaces **with** project-based organization
- Advanced analytics **and** team productivity insights

Community and scalability features

- Interactive style evolution visualization **with** trend predictions
- Hierarchical indexing **for** 100k+ image collections
- Advanced portfolio export tools (PDF, HTML, interactive galleries)
- Community-driven style trend detection **and** challenges
- Optional cloud-assisted indexing **for** massive collections
- Open-source community edition **with** plugin architecture
- Cross-platform feature parity (Windows, macOS, Linux)

Business Model: Community-Driven Freemium

Business Model Evolution: Staged Risk Management

Phase 1 Revenue (Professional Market):

AEYES Prosumer Application

- **Free:** 2,500 images, contextual ads, basic AI vision features
- **AEYES Pro:** \$19/month - unlimited images, advanced AI, professional features
- **AEYES Studio:** \$39/month - collaboration, exports, multi-GPU support, priority processing

Revenue Target: \$200k/month (5k paying users avg \$40) **Infrastructure Cost:** Minimal (desktop software) **Risk Level:** Low (proven prosumer market pays for quality tools)

Phase 2 Revenue (Professional Expansion):

Enhanced AEYES + Essential Video Features

- **Free:** 2,500 images, contextual ads, basic features
- **AEYES Pro:** \$25/month - unlimited images/videos, multi-GPU support, 3D assets
- **AEYES Studio:** \$49/month - collaboration, bulk processing, advanced integrations
- **AEYES Enterprise:** \$149/month - on-premise deployment, custom integrations, dedicated support, SLA guarantees

Revenue Target: \$500k/month (10k users avg \$50) **Infrastructure Cost:** Low (local processing + limited cloud beta) **Risk Level:** Medium (video complexity managed through focused scope) **Enterprise Support Model:** Dedicated customer success, implementation assistance, custom training

Phase 3 Revenue (Scale Platform):

ArtSight Cloud Hosting

- **Free:** 1GB storage, basic organization
- **Pro:** \$15/month - 100GB, full AI features, cross-device
- **Studio:** \$39/month - 1TB, collaboration, portfolio hosting
- **Enterprise:** \$99+/month - unlimited, white-label, APIs

Revenue Target: \$1M+/month (20k users avg \$50) **Infrastructure Cost:** High (cloud hosting, CDN, processing) **Risk Level:** High (scaling challenges, competition)

Phase 4 Revenue (Ecosystem):

Full Creative Platform

- **Subscription tiers:** \$20-500/month based on usage
- **Transaction fees:** 5-15% on marketplace, prints, NFTs
- **Licensing deals:** \$100k-1M+ for white-label solutions
- **Enterprise contracts:** \$10k-100k/month for studios

Revenue Target: \$10M+/month (platform economics) **Infrastructure Cost:** Very high (global platform) **Risk Level:** Very high (platform competition, regulation)

Technical Requirements

Hardware Philosophy: Professional Standards with Cross-Platform Performance

- **Phase 1 Full Mode:** Premium local application requiring professional hardware:
 - **Minimum:** RTX 3070, 16GB RAM, fast NVMe storage
 - **Recommended:** RTX 4080/4090, 32GB+ RAM, multi-GPU support
 - **Display:** 4K/8K optimization, ultra-wide monitor support
 - **VR Ready:** Oculus, Vision Pro compatibility for immersive galleries
- **Phase 1 Lite Mode (Mac/Linux):** Matches Lightroom/Capture One browsing performance
 - **Apple Silicon:** M2/M3 chips with 16GB+ RAM, Metal acceleration
 - **Performance guarantee:** Smooth browsing even without full physics
 - **Feature parity:** All AI analysis, search, organization (minus advanced physics)
- **No compromises:** 60fps rolling interface on supported hardware, real-time AI processing
- **Phase 3+:** Universal access via hybrid cloud (sensitive files local, processing in cloud)

Secure Update Strategy

Integrity & Trust:

- Signed updates with SHA-256 verification
- Public changelog page documenting every patch change
- Deterministic build scripts (post-launch) for supply-chain verification
- Quarterly LoRA patches for emerging aesthetic trends (<50MB)
- Incremental model improvements without full re-downloads
- User feedback drives local fine-tuning improvements

Platform Roadmap:

- Windows/NVIDIA launch (proven stability)
- Apple Silicon Metal optimization (Q2 priority)
- Linux community edition (open-source driven)
- Cross-platform feature parity by end of year
- Explicit timelines and platform support commitments

Risk Management Strategy

Phase 1: Minimize Risk, Validate Core

What We're Testing:

- Does the rolling interface actually improve user experience?
- Will users pay for AI-powered visual organization?

- Can we build sustainable visual intelligence technology?
- Is the freemium + ads model viable?

Success Metrics:

- 10k+ active users within 6 months
- 20%+ conversion to paid tiers
- 4.5+ app store ratings
- <\$50 customer acquisition cost

Pivot Points:

- If rolling interface doesn't test well → traditional grid interface
- If freemium doesn't work → paid-only model
- If AI accuracy is poor → human-assisted curation
- If market is too small → pivot to specific vertical (photographers, etc.)

Phase 2: Prove Scalability

What We're Testing:

- Can we handle video processing at scale?
- Do users want cloud sync badly enough to pay more?
- Will platforms accept our integration plugins?
- Is there enterprise demand for team features?

Success Metrics:

- 50k+ active users, \$200k+/month revenue
- Video feature adoption >60% of Pro users
- 3+ successful platform integrations
- 10+ enterprise pilot customers

Phase 3: Platform Validation

What We're Testing:

- Will users trust us with hosting their creative work?
- Can we compete with Google Photos/Apple Photos on convenience?
- Is our AI differentiated enough to justify switching costs?
- Can we build sustainable competitive moats?

Success Metrics:

- 100k+ cloud users, \$1M+/month revenue
- 50%+ user retention after 12 months
- Clear competitive advantages in AI accuracy
- Successful API ecosystem with 3rd party developers

Phase 4: Market Leadership

What We're Proving:

- Can we become the standard for creative content management?
- Will enterprise clients pay premium prices for our platform?
- Can we successfully compete with Adobe, Google, etc.?
- Is the marketplace model sustainable and profitable?

Community Engagement:

- Beta testing program with early adopters across all platforms
- Feedback portal with feature voting and model correction tools
- Community challenges and style exploration contests
- User-generated tutorials and technique sharing
- Regular "Rediscovery Spotlights" featuring community finds

Implementation Roadmap

1. **Bulletproof MVP:** Adaptive memory management + honest capability matrix
2. **Cross-Platform Foundation:** macOS Metal + Linux community builds
3. **Early Value:** Beta reconstruction in Phase 2 for immediate parameter estimates
4. **Professional Ready:** Zen mode + workflow integrations + API access
5. **Trust Building:** Confidence scores + secure updates + transparent changelogs
6. **Community Growth:** Open-source core + crowd-sourced model training

Unique Selling Points (Final)

- **Truly Cross-Platform:** Works on Windows, macOS, and Linux with appropriate feature sets
- **Safe Organization:** Virtual collections never move your original files
- **Learning System:** Gets smarter as you use it and provide feedback
- **Style Archaeology:** Discovers your artistic evolution patterns with interactive visualizations
- **Prompt Detective:** Reconstructs techniques from images alone with confidence scoring
- **Platform Agnostic:** Universal compatibility without vendor lock-in

- **Privacy First:** Everything runs locally, your art stays private
- **Professional Grade:** Zen mode and workflow integrations for studio environments

Critical Success Factors

Technical:

- Preemptive memory assessment with forced ultra-small batches on constrained systems
- GUID persistence in sidecar files AND image metadata ensures migration resilience
- Intelligent index pruning with transparent user education about trade-offs
- Visual uncertainty indicators prevent unrealistic reconstruction expectations
- Secure update mechanism with post-launch supply-chain verification capabilities
- Cross-platform native experience with honest timeline commitments

UX:

- Truth-first hardware matrix with concrete performance benchmarks prevents disappointment
- Progressive learning system with feature discovery prevents overwhelming new users
- First-run personalization sets appropriate interface complexity immediately
- Live memory monitoring and detailed processing summaries for complete transparency
- Visual uncertainty communication (colored rings, $\pm$ ranges) manages reconstruction expectations
- Contextual tooltips and gradual feature introduction support ongoing mastery

Market:

- **Sustainable free tier** funded by tasteful, relevant advertising
- **Multiple revenue streams** reduce dependency on any single monetization method
- **Developer/investor appeal** through clear path to profitability and scale
- Full data portability prevents vendor lock-in concerns
- Concrete enterprise collaboration features capture studio market
- Open-source core with deterministic builds for developer trust
- **Network effects** - larger user base improves ad rates and community features
- **Low barrier to entry** makes viral growth possible

AEYES empowers creators to rediscover and organize their visual content through artificial intelligence that actually sees. It's not just organization—it's creative archaeology for the visual generation, powered by AI eyes that understand your work.