

Whale ID AI

An AI Platform for Individual Whale Identification and Conservation Intelligence

Executive Summary

Purpose and Significance

Accurate, scalable individual identification is foundational to whale research and conservation. It enables longitudinal tracking of survival, reproduction, injury recovery, habitat use, and movement across years and jurisdictions. Yet most photo-identification programs remain constrained by manual matching workflows that cannot scale with the rapidly increasing volume of imagery from field teams, drones, and partner networks. These constraints slow analysis, limit dataset growth, and delay conservation action.

Whale ID AI addresses this bottleneck through a precision-first AI architecture designed to transform a single submitted image into an identification outcome—high-confidence match, review-ready candidate set, or new individual—while generating auditable intermediate artifacts that support expert verification. The platform architecture supports expansion across whale species as additional reference catalogs and partnerships are incorporated.

What Whale ID AI Is

Whale ID AI is an end-to-end pipeline architecture encompassing:

- Image ingestion (API/webhook-ready design)
- Automated whale detection and localization
- Segmentation and background removal
- Orientation normalization
- Multi-model feature extraction (appearance, keypoints, shape, texture)
- Candidate retrieval via vector search
- Evidence fusion and conservative decisioning
- Database upsert and audit trail

The architecture addresses real field variability—lighting, underwater color cast, occlusion, viewpoint differences, and multi-year appearance drift—while prioritizing catalog integrity through conservative decisioning that escalates uncertainty to human review.

Technical Pipeline

The pipeline integrates proven models and techniques across five stages, each contributing complementary signals to the final identification decision.

Stage 1 — Detection (Localization)

Roboflow YOLO (whale-images/1) detects the whale and outputs a bounding box and confidence score. Reliable localization isolates the region of interest before downstream inference.

Stage 2 — Segmentation (Precise Silhouette + Region Isolation)

SAM2 (Hiera Large) produces a binary mask using detection-guided prompting. Morphological cleanup improves silhouette quality for downstream geometry- and texture-based features. Outputs include a clean mask and segmented whale image.

Stage 3 — Orientation Normalization

An OpenCV PCA-based orientation module estimates the dominant axis from the eroded mask and corrects rotation when required. Normalization reduces pose variance and improves match stability.

Stage 4 — Multi-Model Feature Extraction

Whale ID AI extracts multiple feature families to improve robustness across capture conditions:

A) Global Appearance Embeddings

- MIEW-ID v3 embedding (2152-D) with test-time augmentation for robustness
- MegaDescriptor-L-224 embedding (1536-D) as a complementary global descriptor

B) Local Correspondence Evidence

ALIKED keypoint detection with LightGlue matching (up to 512 keypoints, 128-D descriptors) supplies explicit local correspondence evidence, especially valuable when global embeddings are uncertain.

C) Shape + Texture Descriptors

NeuralWhale structured descriptors using 40 contour and 160 interior keypoints capture morphology and local texture. Enhanced signals include pooled SIFT texture embedding (128-D) and curvature signature (64-D) derived from the segmented contour geometry.

Stage 5 — Candidate Retrieval and Conservative Decisioning

Features are stored and indexed for vector search, enabling efficient retrieval as the catalog grows. Candidate evidence is combined through a conservative decision layer that includes quality gating, compatibility checks, and review escalation. The system produces tiered outcomes:

- **Auto-merge** (very high confidence)
- **Known individual** (confident match with evidence)
- **Needs review** (candidate set + diagnostics for experts)
- **New individual** (no match found)

Current Status

The core pipeline architecture has been designed and key components have been implemented and validated individually. The system produces intermediate artifacts—detection overlays, segmentation masks, keypoint visualizations—suitable for auditing and reviewer workflows. Grant funding will enable integration, optimization, and deployment at conservation scale.

Development Objectives Enabled by Funding

Funding will accelerate the development required to transition from validated components to conservation-scale deployment:

Pipeline Integration and Deployment

- End-to-end integration of pipeline stages into a unified, operational system
- Decision layer calibration and threshold optimization against ground-truth catalogs
- Infrastructure deployment for API access and partner ingestion

Audit-Grade Evidence Retention and Review Workflow

- Durable storage and indexing of pipeline artifacts (masks, overlays, keypoints)
- Reviewer tooling for ambiguous cases with decision traceability

Validation and Partner Deployment

- Systematic validation across years, regions, and capture conditions
- Continuous performance monitoring (precision/recall proxies, review rate, latency, failure modes)

- Partner onboarding pathways for secure ingestion and governance

Species Expansion Roadmap

- Extension across whale species using shared pipeline components
- Species-specific tuning where required, without redesigning the core system

Single-Image 3D Reconstruction (Research Track)

- Methods to approximate 3D whale models from single images using segmentation-derived geometry, anatomical symmetries, and learned priors from accumulated sightings
- Evaluation for conservation-relevant downstream use (improved cross-view matching and consistent morphology assessment)