


Deep Learning for Underwater Domain Awareness

Capability overview — passive-sonar acoustic intelligence: detection, classification and vessel re-identification



Output of Oravont's DEMON analysis tool — blade-rate harmonics ($1\times$ – $5\times$ BR) extracted from the cavitation envelope of a vessel recording.

Oravont Systems builds deep-learning systems that listen to the ocean — detecting acoustic events, classifying their sources, and recognising individual vessels from their radiated sound. The practice pairs three decades of naval engineering and submarine acoustic-stealth design authority with applied machine learning: a designer's understanding of how a vessel makes noise, applied to teaching machines to recognise it.

Physics first — the DEMON toolkit


A vessel's propeller amplitude-modulates its own cavitation noise, and that modulation carries the propulsion signature. Oravont's DEMON (Detection of Envelope Modulation on Noise) tool extracts it: bandpass the cavitation-dominated band, take the Hilbert envelope, and transform the envelope into a low-frequency modulation spectrum — revealing shaft-rate and blade-rate lines and their harmonics, as in the plot above. Because blade rate = shaft rate $\times$ blade count, the harmonic comb reads out RPM and an integer hull attribute — the number of blades — directly from the sound. These interpretable, briefing-ready physics features run alongside the learned representations described on the following pages: two views of the same acoustics, one transparent, one propagation-robust.

The engine — SKANN, a raw-waveform network

Conventional pipelines analyse spectrograms computed with a single, hand-set analysis window — yet no single window can resolve both slow machinery tonals and millisecond transients. SKANN (Selective Kernel Attention Neural Network) instead operates directly on the raw waveform: an eight-branch 1-D filterbank applies parallel convolution kernels spanning roughly 0.5 ms to 64 ms — the longest kernels capture the slow temporal structure of vessel-noise envelopes, the shortest capture sharp transients such as snapping shrimp and clicks — and channel-wise selective-kernel attention learns which scales to trust for each input segment. A five-stage 2-D convolutional backbone then distils the fused feature map into a 512-dimensional embedding. No frequency band is hand-assigned; the network learns where the information lives.

SKANN (Selective Kernel Attention Neural Network) Model Architecture

HybridSKEncoderV3 — Hybrid 1D–2D Convolutional Architecture for Raw-Waveform Underwater Acoustic Classification





B = batch size, T = time steps after 1D stage (approx. 224,000) All convolutions use 'same' padding unless otherwise noted.

SKANN / HybridSKEncoderV3 — hybrid 1-D/2-D convolutional architecture for raw-waveform underwater acoustic classification.

PS12 — underwater acoustic detection & classification

Under Problem Statement 12 of the national-security AI Grand Challenge, Oravont designed an end-to-end system that takes raw hydrophone recordings and returns detected, classified acoustic events across seven target classes — large and small vessels, whales, dolphins, other marine life, natural and other man-made sound. Detection adapts to local ambient conditions; classification runs on SKANN with calibrated per-class probabilities; a runtime clustering stage aggregates evidence within each recording. The pipeline was developed entirely on open-source acoustic corpora, runs offline on GPU or CPU, and produced strong, well-calibrated classification across acoustically diverse classes.




Left: class territory map — UMAP projection of the learned embedding space, with Voronoi territories drawn from class centroids (backbone stage, four acoustic superclasses). Sources separate cleanly in the space the network learns. Right: the PS12 inference pipeline.

The same embedding space that separates acoustic classes is the foundation the vessel-identity work builds on — which is where the second engagement begins.


DISC5 Challenge 10 — vessel re-identification from passive sonar

Classification answers *what kind of source is this*. Re-identification answers a harder question: *which specific vessel is this*? As technical partner under DISC5, Oravont is building the deep-learning core that distils each recording into a compact acoustic fingerprint — a 512-number embedding trained with metric-learning objectives (ArcFace) so that recordings of the same hull cluster together while different hulls separate.



Acoustic fingerprints compared by cosine similarity — two recordings of the same hull agree ($\cos 0.92$); a different hull does not ($\cos 0.16$).

Fingerprints make identification a measurable, repeatable operation: every comparison returns a number, so matches can be ranked and thresholded rather than re-judged by eye. A new contact is either matched to a known hull in the gallery or flagged as new — an open-set decision — and enrolling a new vessel is a single forward pass, with no retraining and no new class.



Built for sovereign deployment

- **Frozen model, growing gallery** — the user organisation enrolls and scores on its own data, entirely in-house
- **No data egress, no retraining** — fully offline operation on the user's infrastructure
- **Robust by design** — training augmented with synthetic channel colouring, measured ambient noise, multipath and time shifts, so the fingerprint keys on the vessel, not the recording chain
- **Licence-clean** — developed on open-source acoustic corpora; open-source dependencies; GPU-accelerated with CPU fallback

A research collaboration in formation with IIT Delhi, in association with a DRDO/NSTL programme, extends this work toward propagation-robust re-identification through deep metric learning and physics-based acoustic-propagation modelling.

Capt (Dr) Sunil Tyagi, Indian Navy (Retd) — Founder & Principal

PhD, Applied Machine Learning (DIAT) · Former Head, Acoustic Stealth Group, Directorate of Submarine Design, Naval HQ
oravontsystems.com · suniltyagi.in · styagi@oravontsystems.com

This overview describes methodology at an unclassified, public level. Programme specifics, customer data and evaluation results are not disclosed.