RESEARCH ARTICLE · TURKISH JOURNAL OF MARITIME AND MARINE SCIENCES

Radar Image-Based Object Detection and Tracking for Autonomous Surface Vehicles

A modular, lightweight HSV-segmentation and motion-aware tracking framework for real-time maritime situational awareness

Sercan KÜLCÜ

Giresun University, Faculty of Engineering, Computer Engineering, Giresun, Türkiye

Volume 12 · Issue 1 (2026) 35-47 | DOI: 10.52998/trjmms.1743670

OVERVIEW

The Study at a Glance

A modular, real-time framework detects and tracks navigational targets on commercial marine radar screens, using colour-space segmentation for detection and a motion-aware identity strategy for tracking — evaluated on 7,200 real radar frames.

98.7%	-59%	-75%	240 ms
True positive detection rate	False positive rate reduction	ID switches vs. baseline tracker	Average processing time per frame

KEYWORDS

- Autonomous surface vehicles
- Radar image processing
- HSV colour-space segmentation
- Multi-target tracking
- Computer vision
- Embedded systems

Why Radar-Based Perception Matters

Autonomous surface vehicles (ASVs) need reliable, real-time detection of surrounding targets. Radar remains essential where optical sensors fail — in low light, fog, or storms — but radar imagery brings its own challenges.

Low resolution & noise

Radar frames show substantial noise and lower resolution than optical images, distorting target shapes.

Fluctuating echo area

Reflections vary with ship orientation, movement, and distance, so static area thresholds often fail.

Fragile identities

Occlusion and sudden manoeuvres cause tracked objects to lose or swap their assigned IDs.

Two Limitations, One Framework

Existing systems typically rely on rigid area-based filtering and fragile ID management. This study addresses both limitations directly.

1

Orientation-aware area filtering

Targets are validated not only by instantaneous pixel area but by consistency with their prior trajectory and heading — preserving objects with temporarily distorted reflections.

2

Motion-predictive ID recovery

Lost targets are kept in an inactive buffer; if a new detection matches the predicted position and trajectory, the original ID is reassigned instead of creating a new one.

System Pipeline Overview

Capture

1200×1024 px
radar frame

→

Preprocess

Gaussian smoothing +
morphological filtering

→

Detect

HSV segmentation +
orientation-aware filtering

→

Track

Euclidean matching +
trajectory-aware ID recovery

Built in Python with OpenCV and NumPy — lightweight enough for real-time embedded deployment, no GPU required.

2.2 HSV-BASED DETECTION

Finding Targets by Colour

Commercial marine radar displays encode targets in yellow tones, with the own-ship marker in white. Converting each frame to HSV separates colour (hue) from brightness (value), making detection robust to lighting changes.

- Threshold yellow hues within configurable HSV bounds to isolate navigational targets
- Morphological open/close operations remove noise and bridge small gaps
- Spatial moments (Eq. 1–3) yield each contour’s centroid and area
- Only contours within a valid area range are kept — unless their centroid still matches the predicted trajectory

TARGET vs. OWN-SHIP ENCODING

Navigational targets

HSV ∈ [10–40]° hue, high sat/val (yellow)

Own-ship marker

Narrow white range: high value, low saturation

2.3 DYNAMIC TARGET ASSOCIATION

Keeping Track of Every Target

Each frame's detections are matched to active tracks by Euclidean distance (Eq. 4). Unmatched detections start new tracks; targets missing for too long move to an inactive buffer rather than disappearing outright.

New Detection → **Distance < threshold?** → **Update Track (reset age)** → **Age > max_lost?** → **Recycle ID or assign new**

Why it works: a frame-indexed deque of recent centroids lets the system distinguish newly appearing objects from re-identifiable ones, while combining spatial proximity with trajectory similarity reduces false reassignments in high-density scenes.

2.4 RELATIVE VELOCITY ESTIMATION

Is a Target Approaching or Receding?

Knowing raw position is not enough — an ASV needs to know how targets move relative to itself. The system compares each target's displacement to the own ship's own motion.

Target velocity

$$v = \Delta position / \Delta t$$

Computed from centroid history in a deque buffer

Relative velocity

$$v_{rel} = v_{target} - v_{ship}$$

Subtracts the ship's own motion (Eq. 6–7)

Smoothing

$$\tilde{v}_t = \alpha \cdot v_t + (1-\alpha) \cdot \tilde{v}_{t-1}$$

Exponential filter stabilises noisy speed readings (Eq. 8)

This distinction — approaching, receding, or holding distance — gives more meaningful situational awareness than absolute motion alone.

Data, Configuration, and Ground Truth

DATA SOURCE

2 hours of screen recording from a Wartsilla ntpro 5000® Bridge Simulator (BridgeMaster radar), Fatsa Faculty of Marine Sciences, Ordu University.

1200×1024 px	Frame resolution
7,200	Radar frames evaluated
labelImg	Ground-truth annotation

KEY CONFIGURATION PARAMETERS

target_hsv_min / max	**[10,50,50] – [40,255,255]**
target_area_min / max	**32 – 196 px²**
distance_threshold	**10 px**
max_lost_frames	**5**
speed_multiplier	**23.47 kts·s/px**

Orientation-Aware Filtering Improves Detection

READING THE RESULT

+4.4 pts

true positive rate over the area-based baseline

−59%

reduction in false positive rate (2.7% → 1.1%)

Combining area with heading/orientation cues lets the system tell real targets apart from noise and clutter far more reliably.

Fewer Lost Identities, Longer Continuity

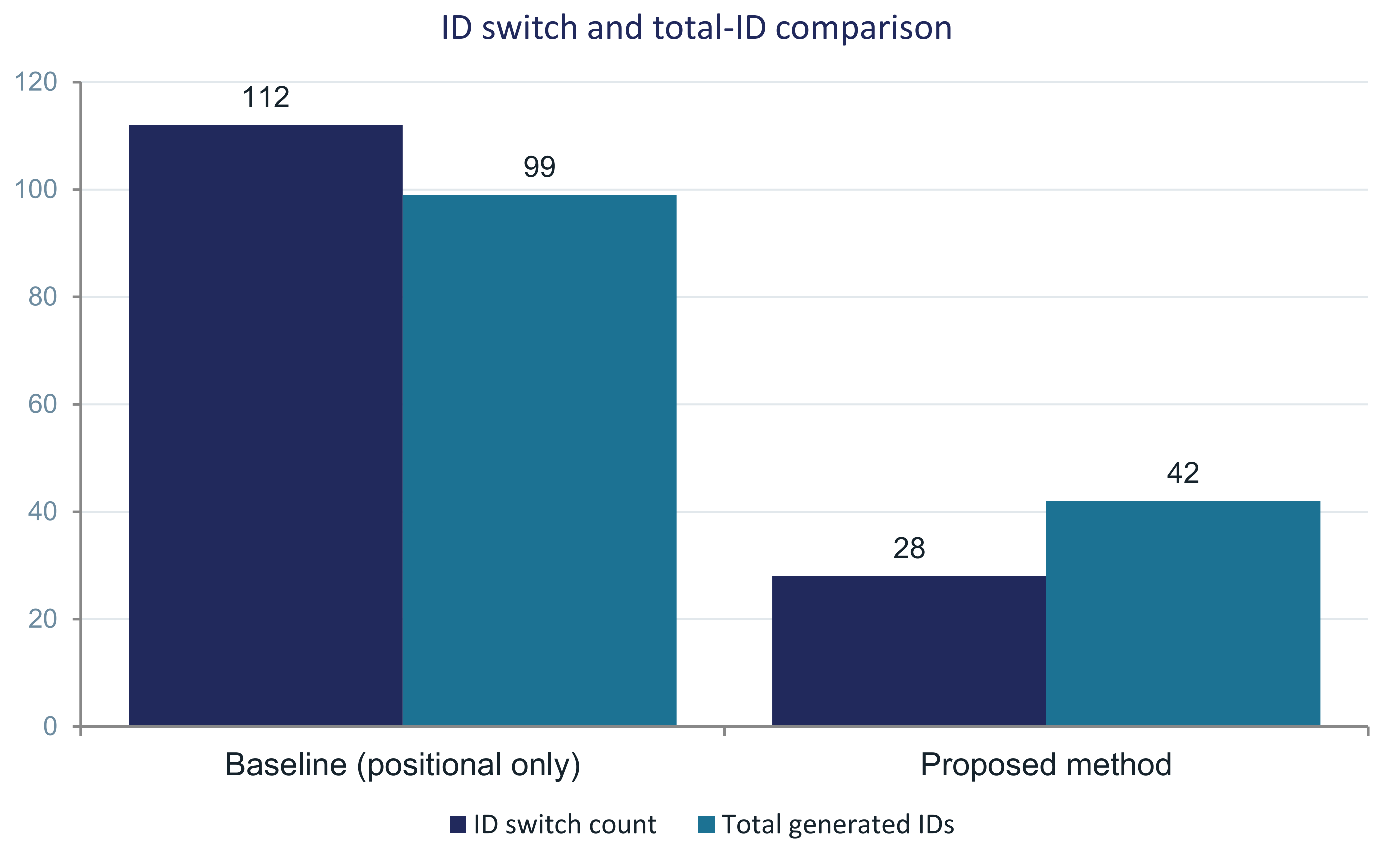

WHY IDENTITIES BREAK

ID switches typically happen when an object is briefly lost and recovered with a new ID — often after abrupt trajectory changes or close interaction with other targets.

By integrating motion features — speed and direction — into the ID-recovery logic, the proposed method preserves consistent trajectories over extended periods, even under occlusion.

Efficient Enough for Embedded Deployment

240 ± 17 ms

Average latency per frame

~4 FPS

Sufficient for a 360° scan every 2–3 s

10.6% ± 2.08

Average CPU usage (peak 15.7%)

111.2 ± 1.9 MB

Average memory usage

Achieved on standard desktop hardware without GPU acceleration — a key requirement for resource-constrained, embedded ASV platforms.

Strengths and Limitations

STRENGTHS

- High detection accuracy (98.7% TPR) and strong tracking continuity (75% fewer ID switches)
- Lightweight and interpretable — no training data or GPU required
- Modular design supports sensor fusion and different radar colour encodings
- Real-time capable on standard embedded hardware

LIMITATIONS & FUTURE WORK

- Relies on manually tuned HSV thresholds and fixed area constraints
- Overlapping targets can still occasionally cause ID swaps
- Future: adaptive thresholding (histogram-based, k-means, ML-driven)
- Future: multi-modal sensor fusion and validation on new radar data

5. CONCLUSION

A Practical Framework for Maritime Perception

By integrating HSV colour segmentation, orientation-aware filtering, and motion-predicted ID reassignment, this framework achieves a 98.7% true positive detection rate across 7,200 radar frames, with an average processing latency of 240 milliseconds — demonstrating clear suitability for resource-constrained embedded systems and reliable real-time situational awareness in ASV navigation and surveillance.

Lightweight | **Real-time** | **GPU-free** | **Modular** | **Extensible**

Sercan KÜLCÜ · Giresun University

sercan.kulcu@giresun.edu.tr

Thank You

Questions & Discussion

Sercan KÜLCÜ · Giresun University, Faculty of Engineering

sercan.kulcu@giresun.edu.tr