RESEARCH ARTICLE · SENSORS (MDPI), VOL. 26 · 2026

High-Precision Marine Radar Object Detection Using Tiled Training and SAHI Enhanced YOLOv11-OBB

Finding small ships on radar screens with over 95% precision

Sercan Külcü
Computer Engineering Department, Giresun University, Turkey

Sensors 2026, 26, 942 · doi:10.3390/s26030942

SUMMARY

What Did This Study Find?

1 Main Finding

Tiled training + SAHI + YOLOv11-OBB reaches over 0.95 mAP@0.5 on real radar data — beating standard full-image detection.

2 Why OBBs Help

Oriented bounding boxes (OBBs) fit ships more tightly than ordinary boxes or masks, since they follow the ship's true heading.

3 Practical Payoff

The pipeline is lightweight enough to run in near real time on edge hardware — suitable for small and medium-sized vessels.

4 Maritime Impact

Accurate orientation and centre-point estimates support better multi-target tracking and collision avoidance.

Why Is Marine Radar Detection So Hard?

The core challenge

- Radar screens are full of sea clutter, small distant targets, and fixed navigation markers that look like ships.
- Classic methods (like CFAR statistical detectors) work, but struggle when clutter changes quickly or targets are faint.
- Recent deep learning models (YOLO-based, GAN-based, multi-sensor fusion) do better, but most still use plain, axis-aligned boxes.
- A key gap remains: resizing large radar images to fit standard network inputs (e.g. 640×640) makes small ships disappear.

What earlier studies achieved

CFAR + superpixels	84.8% detection rate
Marine-Faster R-CNN	93.65% recall
YOLO-SWFormer	91.8% recall
YOSMR (lightweight)	93.1% recall, 92.0% accuracy
Radar + camera fusion	88.6% mAP@0.5

None of these combine tiled training + slicing + orientation-aware boxes — the gap this paper fills.

Source: Section 1 (Külcü 2026)

CONTRIBUTION

Four Contributions to the Field

Tiled + SAHI + OBB pipeline

A combined pipeline that fixes the resolution mismatch in large radar images, greatly improving small-target detection.

Orientation-aware detection

OBBs model the true heading of each vessel, cutting false positives from clutter and fixed navigation aids.

Semi-automatic annotation

A contour-based labelling method builds three label types — box, oriented box, and mask — straight from raw radar echoes.

New performance benchmark

Over 0.95 mAP@0.5 and sub-10-pixel localization error, while staying fast enough for edge devices.

The DAAN Marine Radar Dataset

Collected by the German Aerospace Center (DLR) in the Baltic Sea

- X-band radar mounted on a moving "own-ship", using standard nautical PPI mode at a 3 nautical-mile range.
- 4 target vessels, tracked for about 13 minutes, with very different speeds, headings, and crossing/overtaking manoeuvres.
- Frames also contain wake trails, sea clutter, and interference from fixed navigation aids.
- AIS data (real GPS-based ship positions) is used as accurate ground truth for evaluating detections.

760 full-resolution radar frames

4 vessels tracked, plus the moving own-ship

1 Hz frame capture rate, ~13 minutes total

Source: Section 2.1, Figs. 1–2 (Külcü 2026)

2.2 · METHOD

Turning Raw Echoes into Training Labels

1. Find Candidates

Convert the image to HSV colour space and threshold the bright yellow used for strong echoes.

→

2. Build Labels

For each surviving shape, compute a box, an oriented box, and a smoothed outline (mask) at the same time.

→

3. Human Check

About 15% of frames needed a manual fix — splitting merged ships, joining broken echoes, or removing false navigation-aid shapes.

Only ~15% of frames needed manual correction — because the radar display uses a fixed colour code for strong echoes, automatic thresholding worked well for most frames. This check happens only once, offline; the trained model runs fully automatically afterwards.

Source: Section 2.2 (Külcü 2026)

Three Label Formats, One Contour

Axis-Aligned Box

The smallest upright rectangle around the shape. Simple, but loose-fitting for tilted ships.

Oriented Box (OBB)

The smallest rotated rectangle. Matches the ship's real heading for a tight, accurate fit.

Segmentation Mask

A smoothed polygon outline (Douglas–Peucker algorithm), capturing the exact silhouette.

Yellow-echo threshold (HSV)

$$\text{hue} \in [20°, 40°],\ \ \text{saturation} \in [100, 255],\ \ \text{value} \in [100, 255]$$

Shapes smaller than 8 pixels are dropped as noise.

Mask smoothing precision

$$\varepsilon = 0.01 \times \text{perimeter}$$

Keeps the outline's shape while removing extra, redundant points.

Source: Section 2.2.1–2.2.2, Eq. (1)–(2), Algorithm 1 (Külcü 2026)

Tiled Training: Keeping Small Ships Visible

Resizing a large radar image down to 640×640 pixels makes small, distant ships shrink to almost nothing. The fix: cut each image into overlapping tiles first.

640×640 px

size of each tile
(20% overlap, 512 px steps)

760 → 3,040

training images
after tiling (4× more)

4,457 → 11,893

labelled ship instances
available for training

Only tiles that contain at least one ship are kept — this avoids wasting training time on empty sea. Labels are automatically re-calculated for each tile's own coordinate space.

Source: Section 2.4.2 (Külcü 2026)

Three YOLOv11 Variants, Enhanced with SAHI

YOLOv11-detect

Standard detection with axis-aligned boxes (center, width, height).

YOLOv11-obb

Predicts a rotated box using four corner points — matches vessel heading.

YOLOv11-seg

Adds pixel-level masks for the exact shape of each echo.

🔍 SAHI — Slicing Aided Hyper Inference

At test time, SAHI cuts each full-resolution radar image into overlapping 640×640 slices, runs the model on each slice separately, and merges the results back together. This matches the resolution the model was trained on — without ever retraining it.

Source: Sections 2.3, 2.5 (Külcü 2026)

Choosing a Backbone: Lightweight YOLO Comparison

Model	Parameters	GFLOPs	mAP@0.5	mAP@0.5:0.95	Inference (ms)
YOLOv5n	2.5 M	7.1	0.981	0.787	2.7
YOLOv8n	3.0 M	8.1	0.987	0.797	2.7
YOLO11n	2.6 M	6.3	0.984	0.816	2.9
YOLO26n	2.4 M	5.2	0.987	0.807	2.8

All variants score above 0.98 mAP@0.5. YOLO11n was chosen: it gives the best mAP@0.5:0.95 (0.816) — the strictest, most demanding accuracy measure.

Source: Table 1, Section 3.1 (Külcü 2026)

Full-Image Resizing vs. Tiled Training

Model	Training Method	mAP@0.5	mAP@0.5:0.95
YOLOv11-detect	Resized	0.921	0.681
YOLOv11-detect	Tiled	0.984	0.816
YOLOv11-obb	Resized	0.955	0.799
YOLOv11-obb	Tiled	0.989	0.882
YOLOv11-seg	Resized	0.922	0.668
YOLOv11-seg	Tiled	0.985	0.810

Tiled YOLOv11-OBB wins overall: 0.989 mAP@0.5 — the highest score in the study, and a clear jump over its resized counterpart (0.955).

Source: Table 2, Section 3.2 (Külcü 2026)

Tuning SAHI: Overlap Ratio and Slice Size

Overlap ratio (0.1 → 0.5)

Overlap	mAP@0.5	mAP@0.5:0.95	mAP@0.75
0.1	0.962	0.487	0.380
0.2	0.962	0.488	0.381
0.4	0.961	0.489	0.392
0.5	0.962	0.490	0.394

Slice size (640 px → 240 px, overlap fixed at 0.2)

Slice	mAP@0.5	mAP@0.5:0.95	mAP@0.75
640 px	0.962	0.488	0.381
512 px	0.973	0.506	0.426
448 px	0.971	0.509	0.447
320 px	0.977	0.512	0.486
240 px	0.976	0.501	0.459

Key takeaways

- A bigger overlap gives a small, steady gain — useful when precision matters more than speed.
- Shrinking the slice size helps a lot — mAP@0.75 nearly doubles (0.381 → 0.486) — but going too small (240 px) starts to hurt, since there are too many boundary artifacts.
- 320 px slices give the best accuracy; 640 px with 0.2 overlap is the best speed/accuracy balance.

Source: Tables 3–4, Section 3.3 (Külcü 2026)

Comparing Variants: Accuracy vs. Precision/Recall

mAP at different strictness levels

Model	mAP@0.5	mAP@0.5:0.95	mAP@0.75
Detect	0.962	0.488	0.381
OBB	0.954	0.374	0.428
Segment	0.948	0.494	0.441

Precision & Recall at IoU = 0.5

Model	Precision	Recall
Detect	0.929	0.958
OBB	0.912	0.955
Segment	0.873	0.955

How to read this: the plain detect model has the best raw precision and recall — but its accuracy drops fastest at strict thresholds (mAP@0.75 = 0.381). OBB and Segment keep much more of their accuracy at mAP@0.75 (0.428 and 0.441), proving that they draw tighter, more reliable boundaries around each ship — even if they occasionally flag a few more false positives.

Source: Tables 5–6, Section 3.3 (Külcü 2026)

How Close Are Predictions to the Real Ship Position?

Mean distance (in pixels) between the predicted centre point and the true AIS-based position:

Vessel	Model	Mean (px)	Median (px)	Std. Dev. (px)
Vessel 2	Detect	6.33	6.05	3.10
Vessel 2	**OBB**	6.30	6.00	3.04
Vessel 2	Segment	6.75	6.30	2.98
Vessel 4	Detect	5.14	5.20	2.99
Vessel 4	**OBB**	5.06	5.10	3.10
Vessel 4	Segment	5.62	5.45	2.93

All three models stay under 10 pixels of mean error — and OBB is consistently the most accurate at finding a ship's true centre, confirming that orientation-aware boxes track vessels more faithfully.

Source: Table 7, Section 3.4 (Külcü 2026)

Speed on Real Hardware, and When It Fails

Computational efficiency

Model size	**~5.5–6.0 MB**
Peak memory use	**< 4.2 GB**
Full-image inference	**18–22 FPS**
Tiled + SAHI inference	**4–6 FPS**

4–6 FPS is a 3–4× slowdown versus full-image inference — but still enough for near-real-time collision avoidance and AIS-linked tracking.

Six typical failure patterns

- Weak echo from signal attenuation → missed target
- Wave clutter hides the vessel's signature
- Slicing (SAHI) splits one ship's echo across two tiles
- Sharp manoeuvres change the echo shape too fast
- One ship's echo splits into two false detections
- Two nearby ships merge into a single, undercounted echo

Source: Sections 3.5, 3.7 (Külcü 2026)

CONCLUSION

Scale-Aware and Orientation-Aware Design Wins

- Combining tiled training, SAHI, and YOLOv11-OBB sets a new benchmark for marine radar detection: over 0.95 mAP@0.5 with sub-10-pixel localization error.
- This solves two problems at once — the resolution mismatch that hides small ships, and the orientation blindness of ordinary bounding boxes.
- The segmentation variant adds further value for tasks like separating overlapping wakes or tracking individual vessels.
- Future work: speed up the pipeline (model pruning, lighter backbones), add sea-state labels to study weather effects, and test on more public radar datasets.

REFERENCE

Source

Külcü, S. (2026).
High-Precision Marine Radar Object Detection Using Tiled Training and SAHI Enhanced YOLOv11-OBB.
Sensors, 26(3), 942. doi:10.3390/s26030942

Dataset: DAAN marine radar repository, German Aerospace Center (DLR) — https://doi.org/10.26090/rwmr

This presentation is an academic summary of the research paper above. For the full reference list (27 citations), please see the original article.

Thank You