Projeto C2
A melhor decisão

HARBOR: Heading Analysis and Reconstruction from Behavioral Observation and Radar

Joao P. A. Dantas, Paulo F. Silva Filho, Jelton A. Cunha, Gabriel Dietzsch

Institute for Advanced Studies (IEAv) – Brazilian Air Force

AIAA Aviation Forum, 8–12 June 2026

MARITIME SURVEILLANCE — THE OPERATIONAL CONTEXT

Growing Maritime Traffic

Global shipping, fishing fleets, and unregistered vessels have grown significantly.

Persistent Threats

Illegal fishing, smuggling, unauthorized transit through restricted zones, and non-cooperative vessels are constant operational challenges.

Wide-Area Monitoring

Covering large coastal and open-sea regions continuously (independently of vessel cooperation) demands robust sensing.

What situational awareness requires

Where are the vessels?
Position and detection

What type are they?
Classification

Where are they going?
Heading and trajectory

What might they do next?
Prediction

In defense scenarios, this information is often incomplete, unreliable, or deliberately denied

THE PROBLEM

AIS

Automatic Identification System

- Standard system for tracking ships
- Broadcasts identity, position, speed, and heading
- Depends on vessel cooperation
- Can be turned off, lost, or spoofed in defense environments
- Small or non-cooperative vessels often do not transmit at all

SAR

Synthetic Aperture Radar

- Satellite radar imaging sensor
- Works day or night, in any weather
- Covers wide ocean areas independently
- But gives only a single snapshot in time
- Shows where ships are, not where they are going

Can we predict vessel motion from a single SAR image, with no extra data at runtime?

WHY IT IS HARD

Without two images separated in time, speed and direction cannot be measured directly

Approach A

Track ships across multiple SAR passes

Requires frequent satellite revisit

Approach B

Combine SAR with AIS at inference time

Requires cooperative AIS signal

Fails in low-revisit or contested areas

Fails when AIS signal is unavailable

AIS data
(offline calibration only)

HARBOR — single SAR image at inference, no cooperative signal needed

RELATED WORK

Detecting ships in SAR

Modern approaches use threshold-based methods and, more recently, deep learning networks trained on large annotated datasets. Detection itself is a well-studied problem.

Estimating heading and speed

Most methods look for the wake pattern a ship leaves behind to infer its direction and speed. Works well for large ships in calm seas but struggles with smaller vessels or rough conditions.

Combining SAR with AIS

Merging satellite radar images with AIS transmissions improves both detection and classification. But it only works when AIS signals are available, not in denied environments.

Predicting vessel trajectories

Machine learning models trained on historical AIS data can predict where ships are going.

Where HARBOR fits in

We build on these ideas and combine detection, heading estimation, and probabilistic trajectory projection in a single pipeline, designed to work from one SAR image, with AIS used only for offline calibration.

THE CORE IDEA

OFFLINE PHASE

(done once, before operation)

- Large AIS dataset (7M messages)
- Clean and filter records
- Classify vessels by physical size
- Compute median speed and angular spread per category
- Store motion parameters

parameters
passed once

ONLINE PHASE

(at inference time)

- Input: single SAR image
- Detect and classify vessels by size
- Estimate heading from shape and intensity
- Apply calibrated motion parameters
- Output: probabilistic heatmap of future positions

No AIS needed at runtime. No satellite revisit required. One image in, motion map out.

PIPELINE OVERVIEW

1
Clean the image
Preprocessing

2
Find ships
Detection

3
Sort by size
Classification

4
Estimate direction
Motion Estimation

5
Project positions
Trajectory Projection

AIS data feeds step 5 via offline calibration only (not used at runtime)

PREPROCESSING

Raw SAR image

after
processing

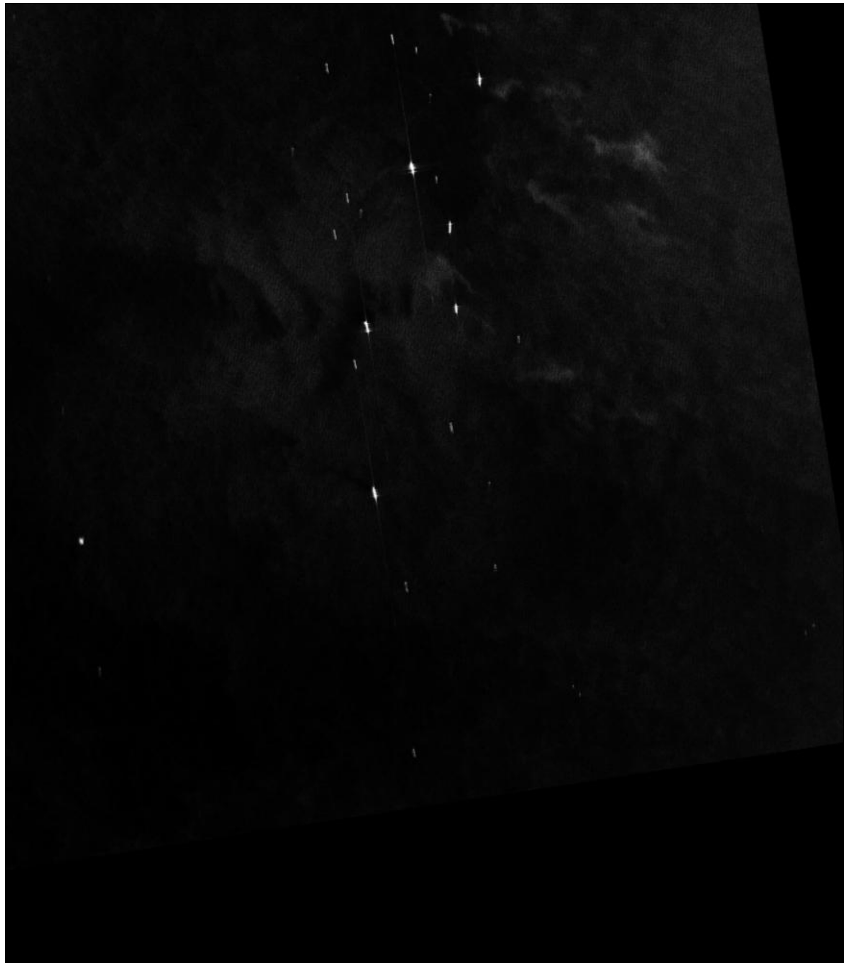

After preprocessing

- The raw image is very noisy (ships are barely visible)
- We apply a series of filters to reduce noise while keeping bright ship-like targets
- We align the image with real geographic coordinates
- We keep only the brightest spots (ships reflect radar signals much more strongly than water)
- Result: a clean image where ships stand out clearly against the dark sea

DETECTION & SIZE CLASSIFICATION

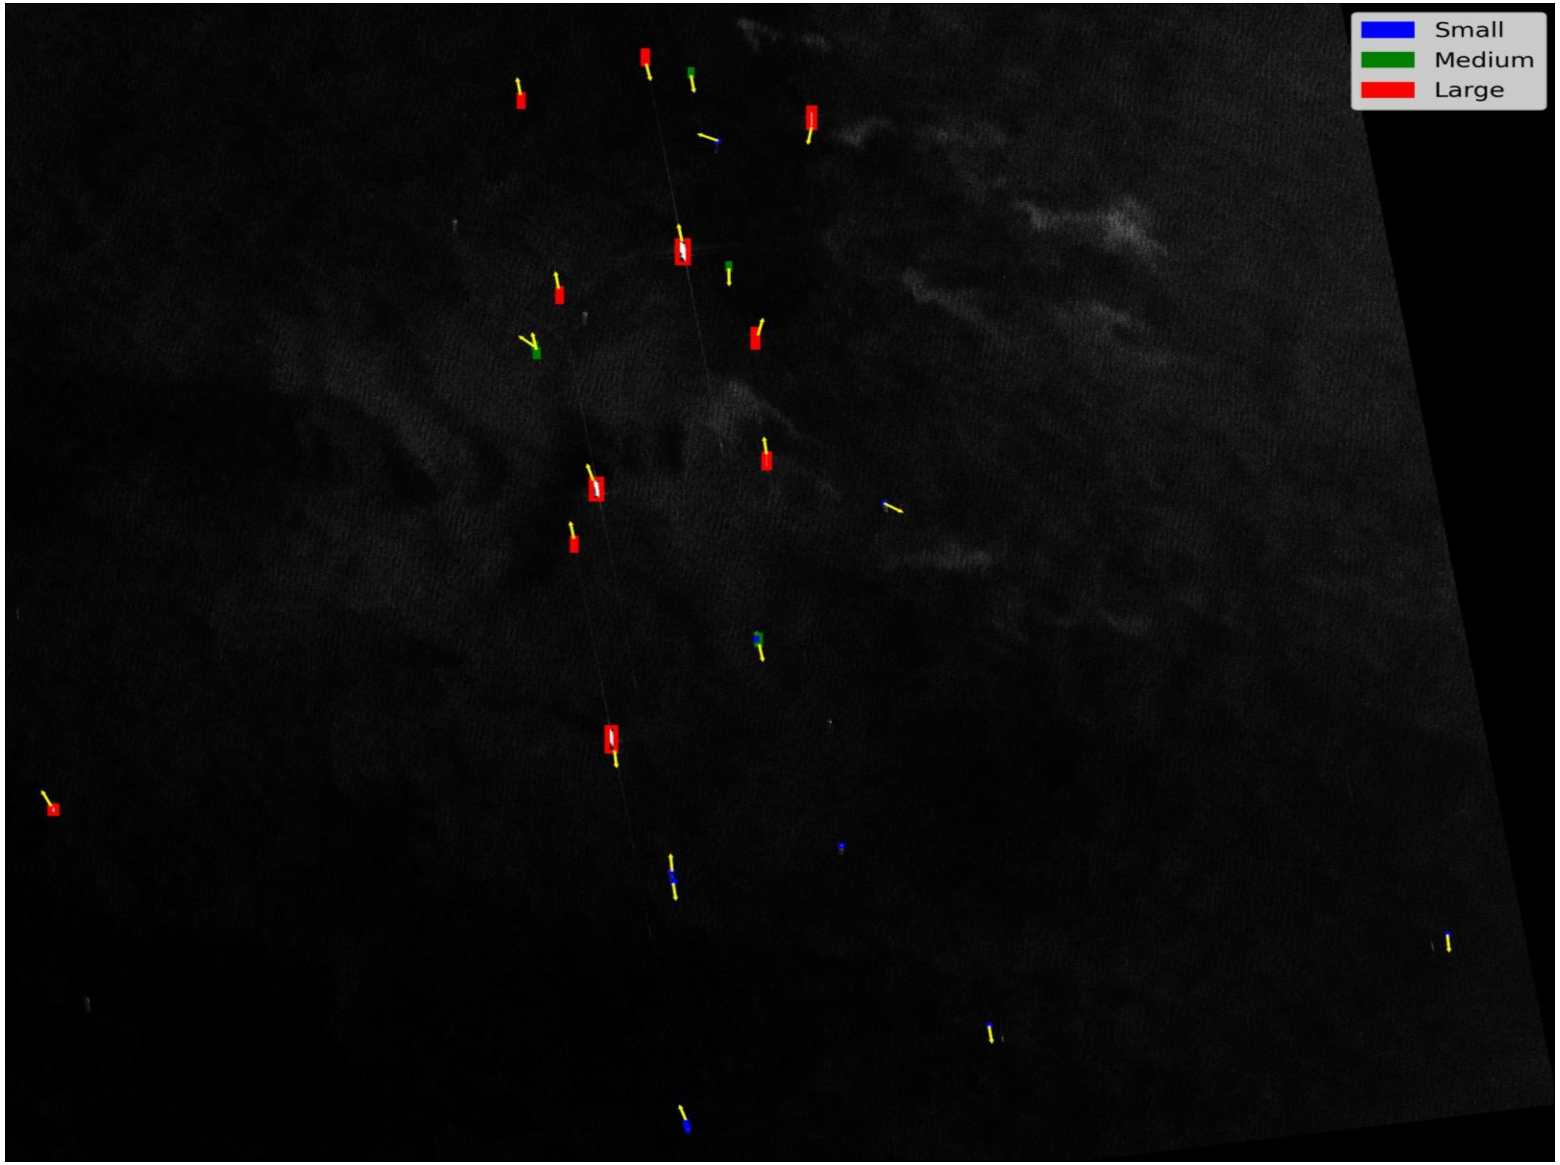

Detected vessels color-coded by size (blue: small green: medium red: large)

How detection works

- Group connected bright pixels into candidate objects
- Discard very small blobs, likely noise, not ships
- No training data needed, fully rule-based

Size classification

Small	*< 1,000 px²*
Medium	*1,000–5,000 px²*
Large	*> 5,000 px²*

Based on bounding box area in pixels

HEADING ESTIMATION

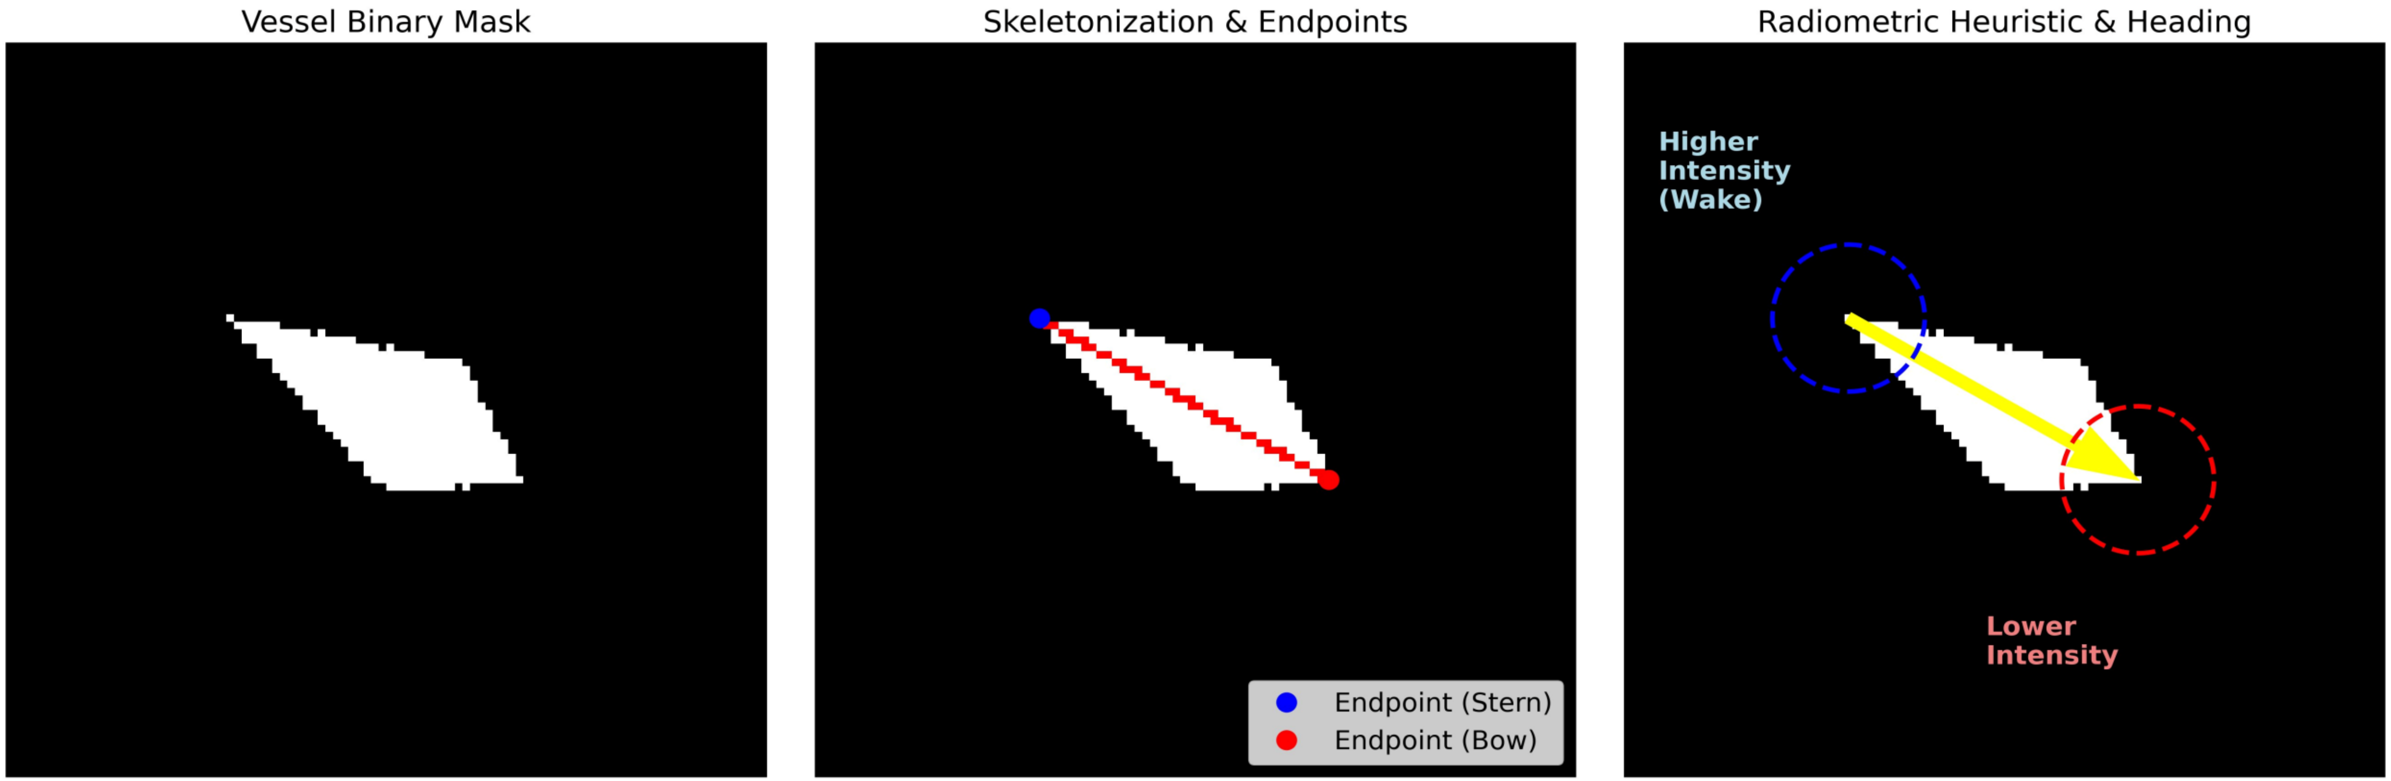

- Reduce the ship shape to a single center line to find its two ends
- One end is the bow (front), the other is the stern (back)
- The stern leaves a wake in the water, which appears brighter in the radar image (we use this to tell the two ends apart)
- The direction arrow points from stern to bow (that is where the ship is going)
- If the two ends look too similar in brightness, we mark the estimate as low confidence

OFFLINE AIS CALIBRATION

- We used a public dataset with about 7 million ship position records
- For each ship, we measured two things: how fast it typically moves, and how much it tends to change direction
- We grouped ships into three size categories (small, medium, and large) based on their physical length
- We computed one representative value per category and stored those as the motion parameters used at runtime

What we learned from the data

Ship size	Typical speed	How much it turns	Ships analyzed
Small	4.5 knots	31° — turns a lot	2,291
Medium	10 knots	5° — mostly straight	623
Large	9.6 knots	2° — very straight	387

Based on 3,301 ship trajectories — small vessels move slower but turn much more than large ones

PROBABILISTIC MOTION MODEL

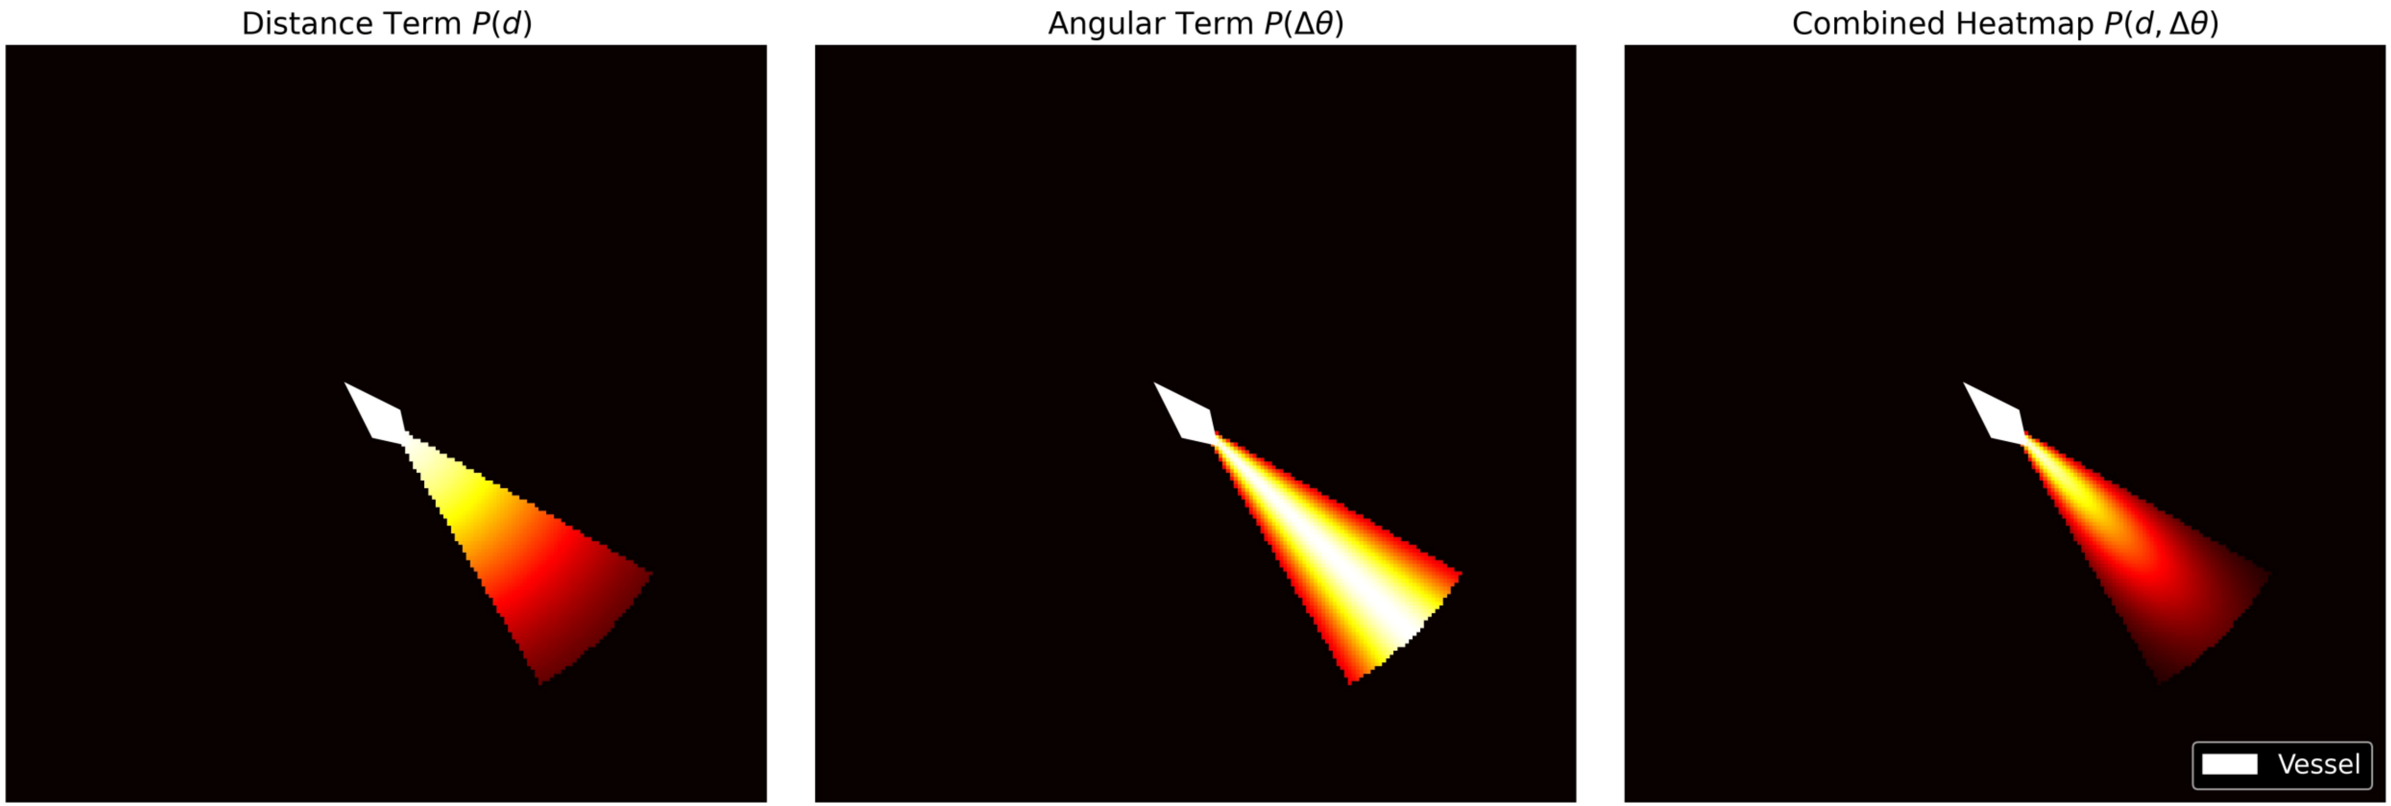

- For each ship, we draw a probability cone pointing in its estimated direction
- The cone gets shorter the further out you go (the ship is less likely to be far away in a short time)
- The cone is wide for small ships (they turn a lot) and narrow for large ships (they go straight)
- All cones are added together (areas with many ships nearby appear brighter)

CASE STUDY — SAR DATA & RESULTS

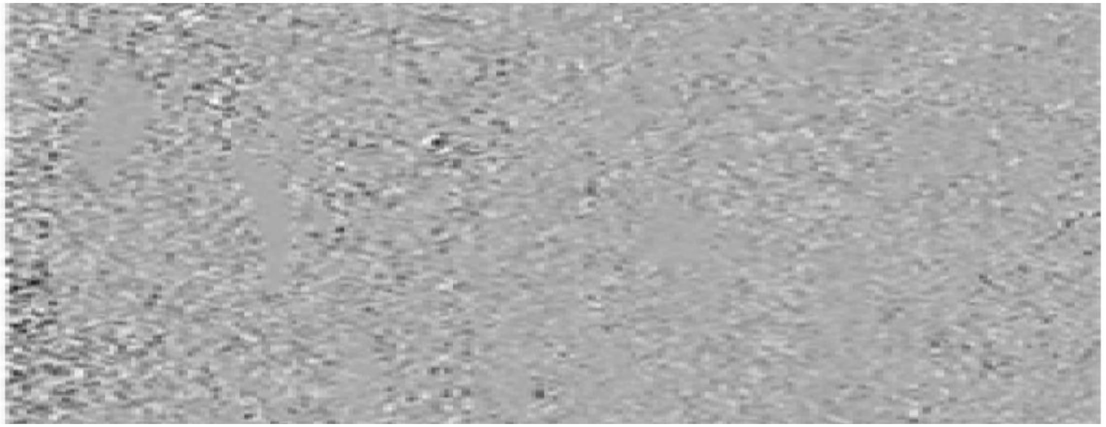

(a) Raw SAR image

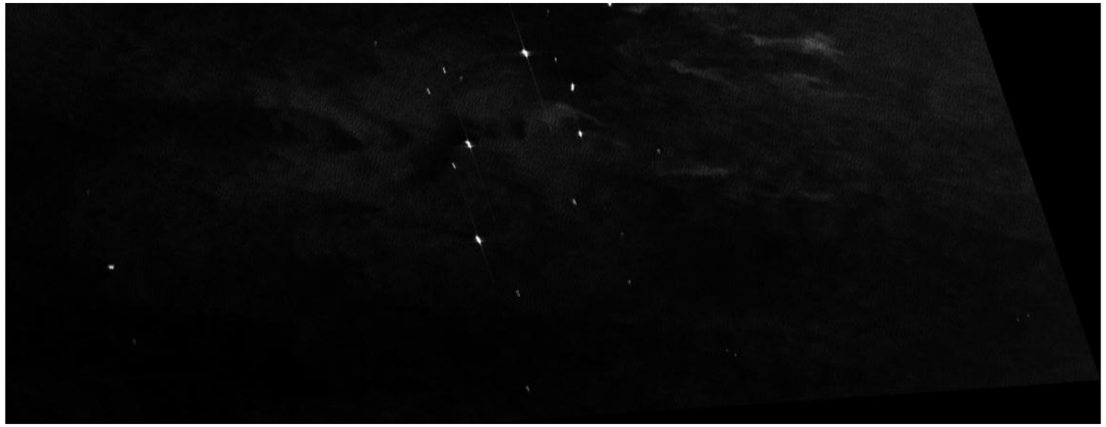

(b) Preprocessed

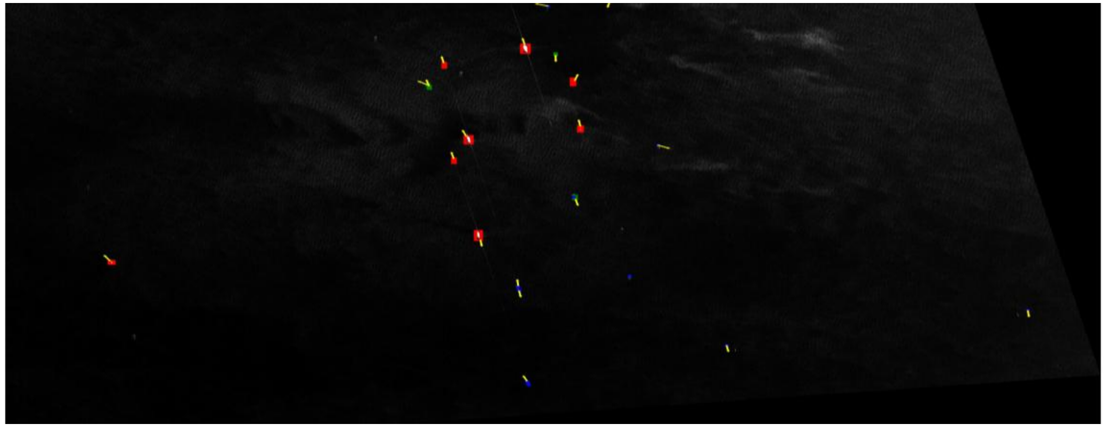

(c) Detected vessels + headings

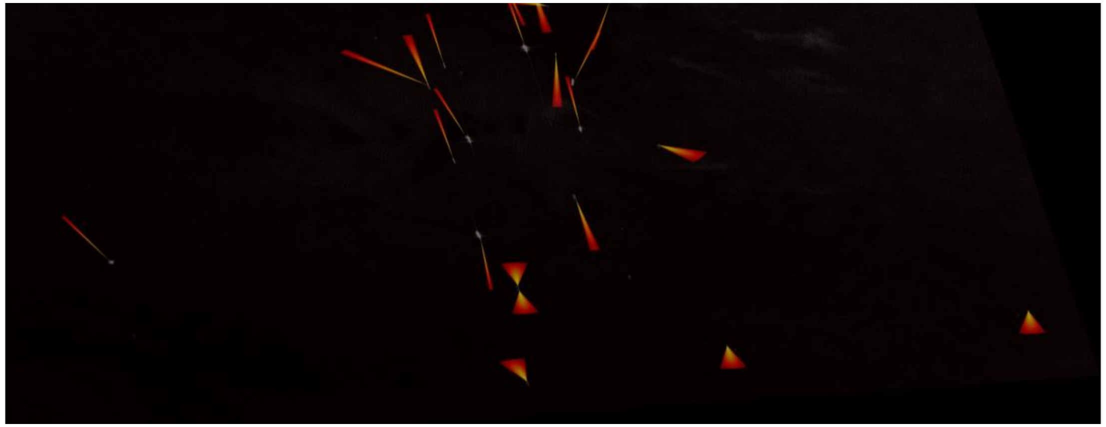

(d) Projected positions — 360 min

27 vessel candidates: 11 small · 5 medium · 11 large | Port of Rio Grande, Brazil

PROJECTION HEATMAP — TIME EVOLUTION

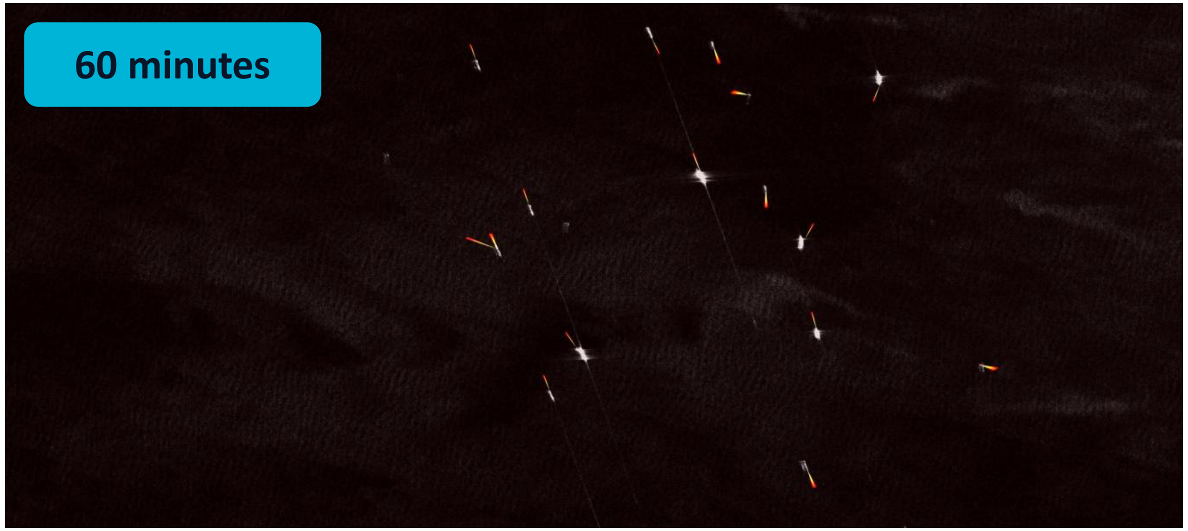

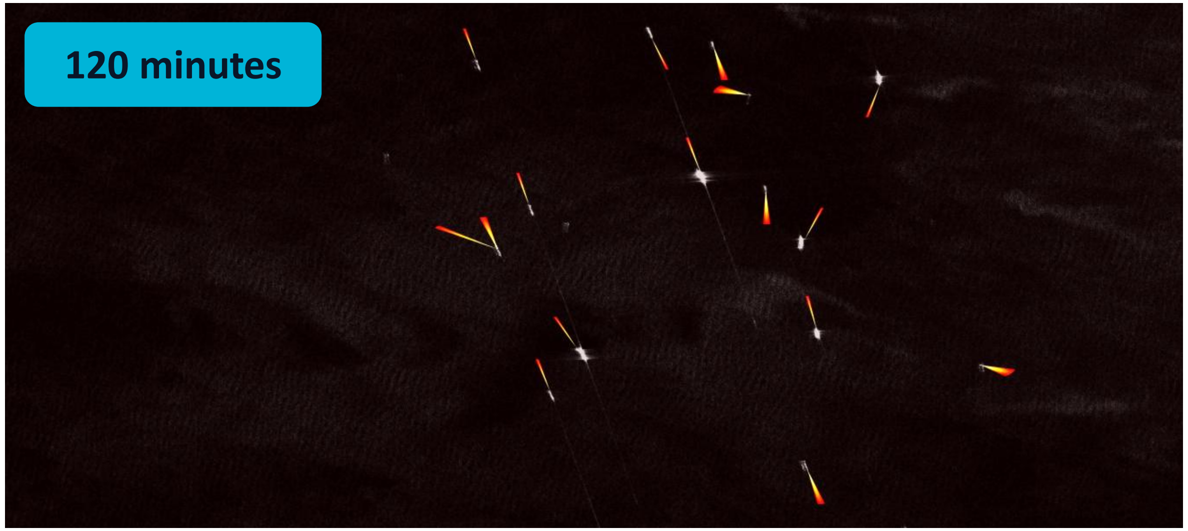

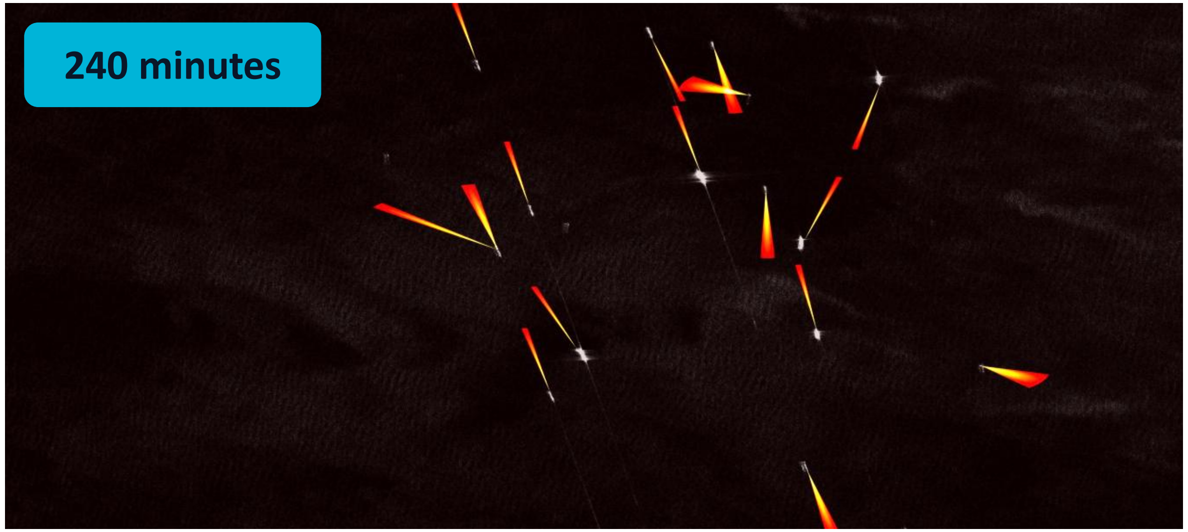

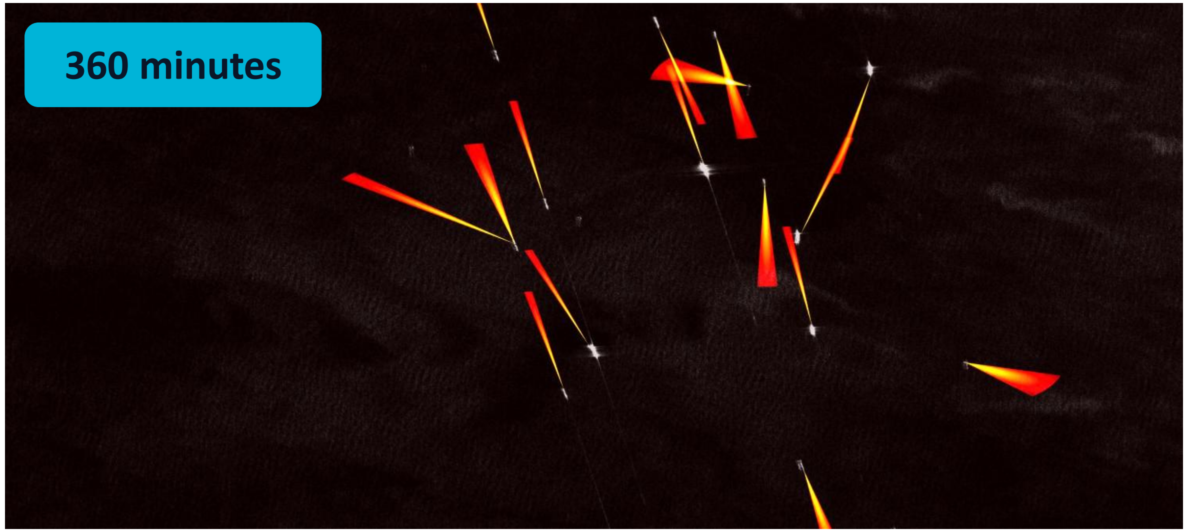

Zoomed in on the same busy region of the scene: as the time horizon grows, the projection cones extend further from each ship

DETAILED VIEW — CONFIDENCE & SIZE CATEGORIES

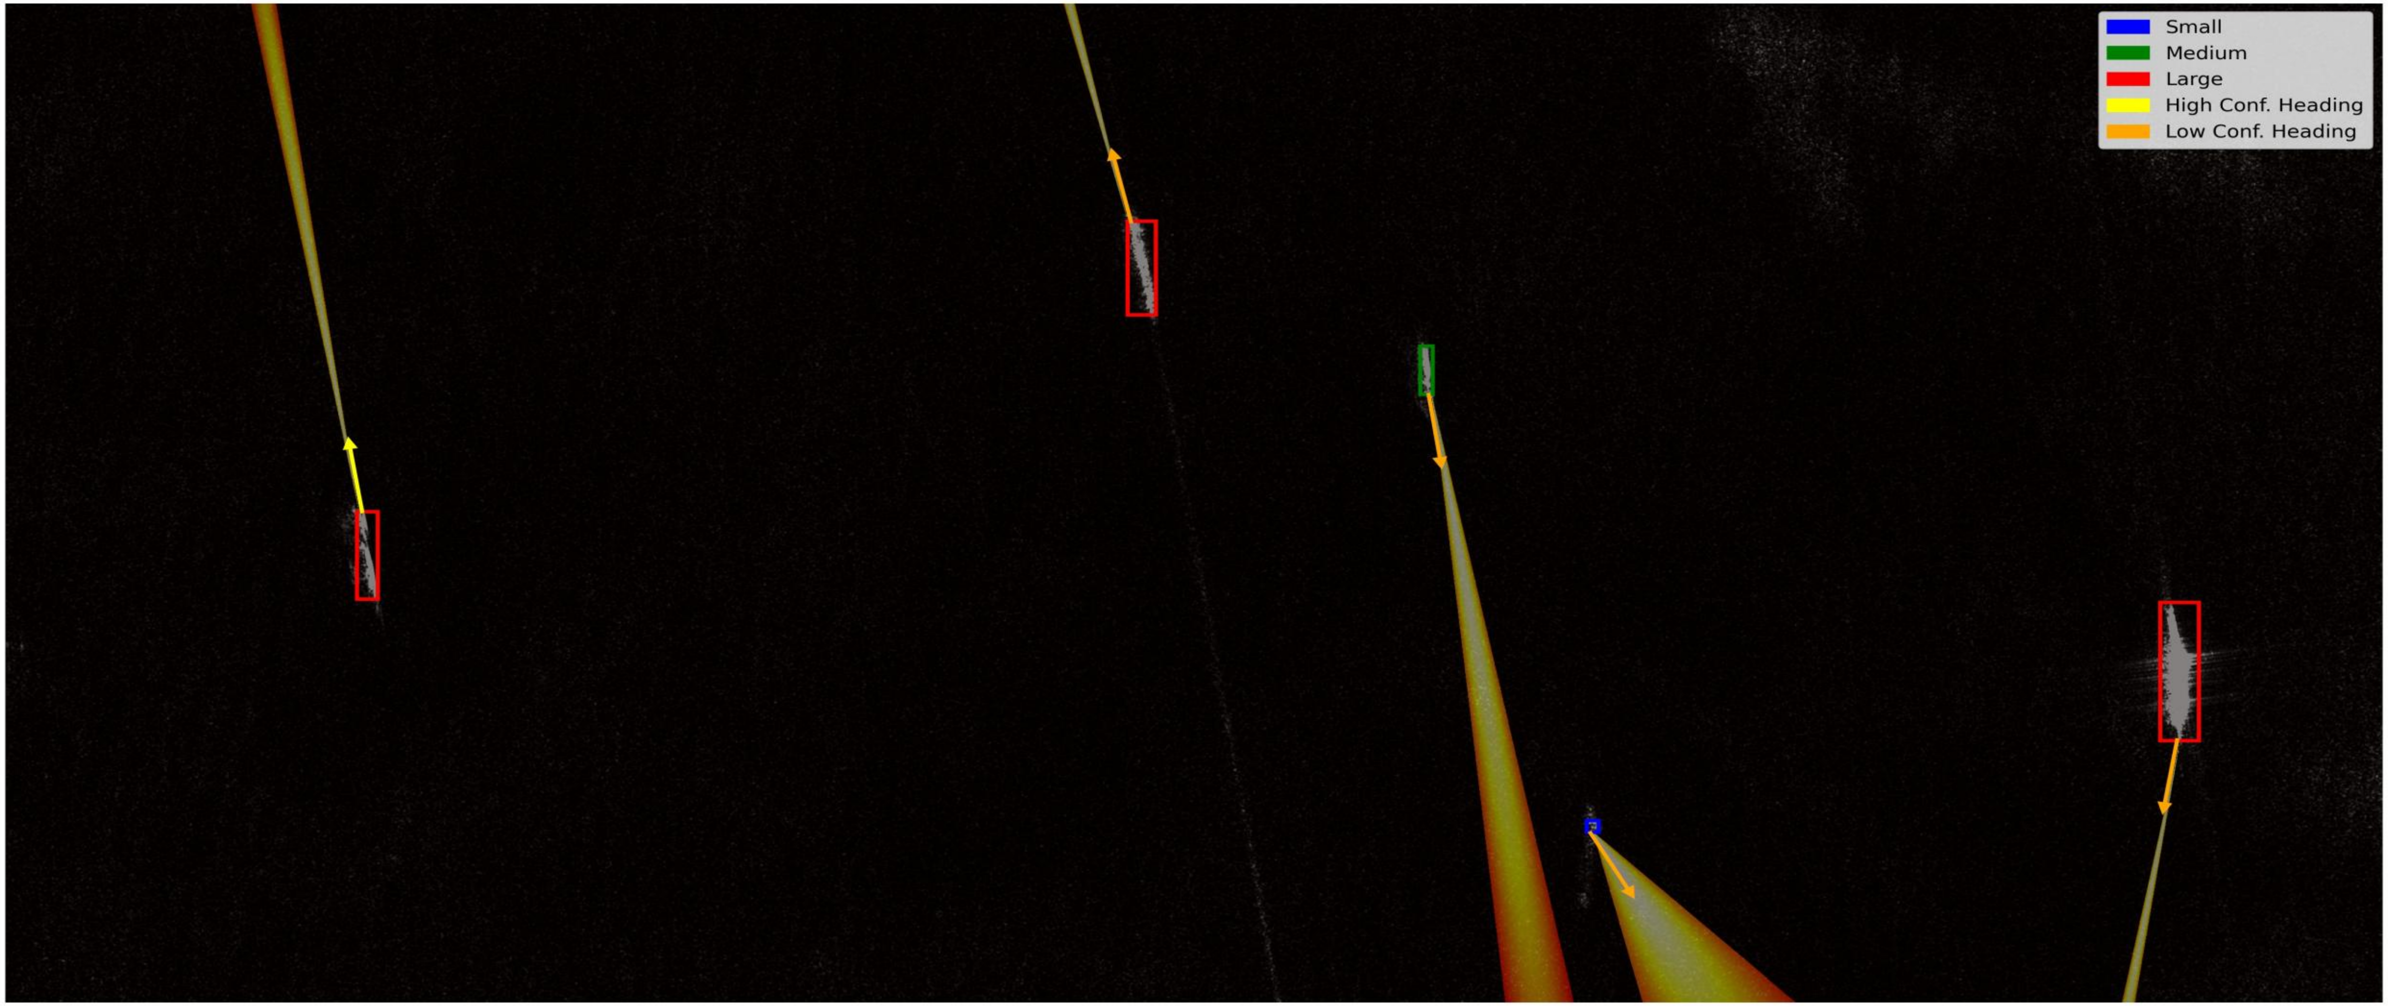

Yellow arrows: high-confidence heading | Orange arrows: low-confidence (<10% intensity difference) | Blue: small – Green: medium – Red: large

CURRENT SCOPE

**This case study proves the pipeline works end to end on real satellite data.
The points below are the natural next steps to mature it further.**

1 Adapting to other sensors

The size categories are based on pixel area. Moving to a different satellite or resolution means recalibrating those threshold.

2 Using richer radar signals

We currently use image brightness. Some satellites also capture additional signal information that could make vessel identification even more reliable.

3 Handling stationary ships

Ships at anchor leave no wake, so the bow/stern signal is unavailable. A complementary indicator for these cases is a natural addition to the pipeline.

4 Expanding the calibration data

Motion statistics were learned from US coastal traffic. Adding data from the region of interest (like Brazilian waters) would refine the projections further.

CONCLUSION & FUTURE WORK

HARBOR transforms a single SAR image into a probabilistic vessel motion map — with no AIS, no repeated passes, and no external data at runtime.

Key contributions

- A complete pipeline: from raw radar image to a map of likely ship movement
- A simple way to estimate which direction a ship is heading, using its shape and wake
- Ship movement patterns learned once, offline — applied instantly to any new image
- Tested on a real satellite image of the Brazilian coast

Next steps

- Top priority: test accuracy against real ship tracking data
- Account for ocean currents and wind in the movement predictions
- Use more visual clues from ships (shape, signature, wake) to improve detection
- Learn movement patterns from more regions and seasons, not just the US coast

github.com/jpadantas/sar-preprocessing-snappy | github.com/jpadantas/harbor

Projeto C2
A melhor decisão

Thank You

Questions?

Joao P. A. Dantas
www.joaopadantas.com