Computer Vision at Sea: Automated Fish Tracking for Sustainable Fishing

Vienna Saccomanno
Senior Scientist
The Nature Conservancy

Alicia Schandy
ML Engineer
tryo·labs

tryo·labs

Introduction

The conservation challenge in industrial fisheries

- **Context:**
  Billions rely on seafood for essential protein. But > one third of fish stocks are fished at unsustainable levels.
- **Problem**:
  Illegal, unreported, and unregulated (IUU) fishing threatens the sustainability of global fish stocks.
- **Challenge:**
  How can we verify — and incentivize — sustainable industrial fishing with improved monitoring?

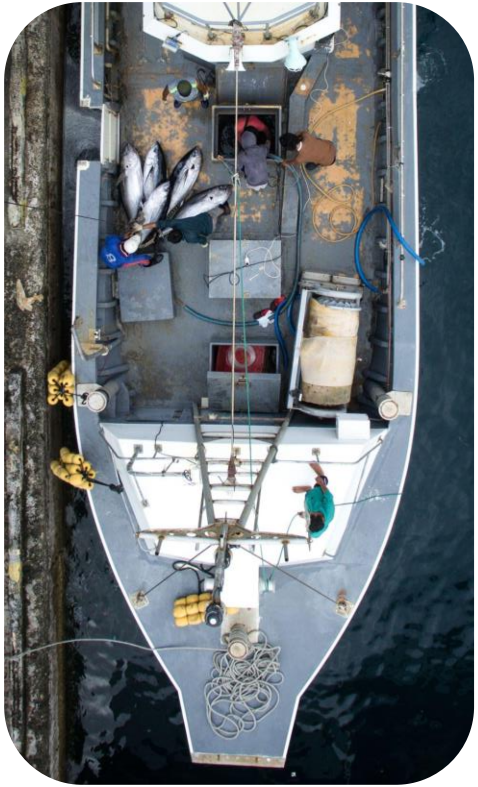

tryo·labs | The Nature Conservancy

Electronic monitoring (EM) of fishing activity

- EM = computers, cameras, gear sensors, and GPS to record and transmit data about fishing activity.
- EM can verify vessel reporting.
- But, EM produces massive amounts of data and footage review can be delayed by months.
- Verified catch in near-real time is needed for IUU detection.

The embedded vision solution

- EM is well suited to CV application, especially longline fishing gear.

- **Goals:**
  1) Transform the status quo of the EM footage review process into a strategic workflow.

  2) Verify catch and identify IUU before products enter the supply chain.

The evolution of fishing monitoring

2025: AI-Powered real-time monitoring

Solution

AI + Edge Computing → Automated real-time tracking & alerts.

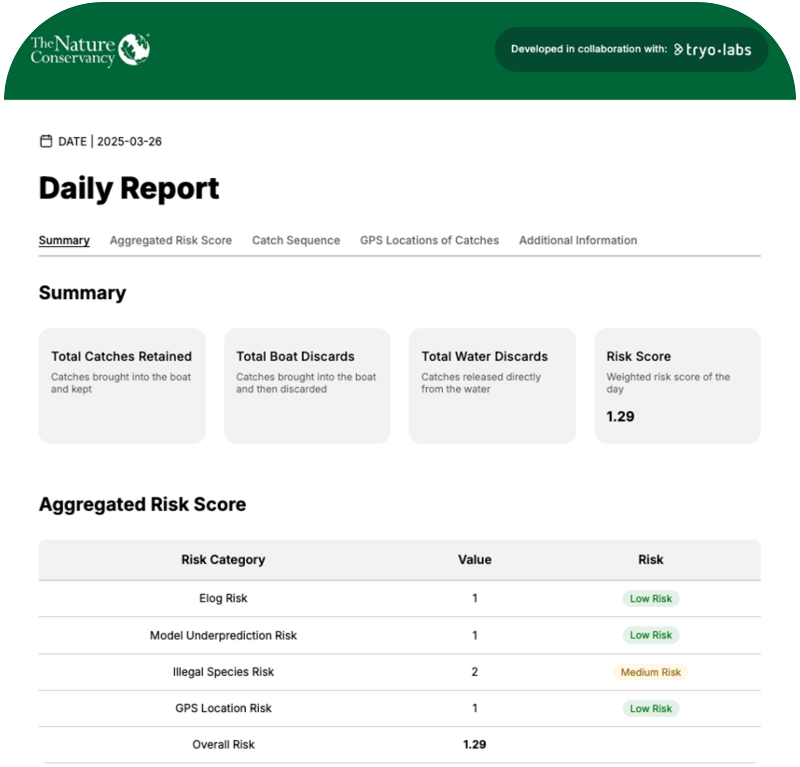

The evolution of fishing monitoring

2025: AI-Powered real-time monitoring

Advantages

- AI counts and classifies fish directly onboard.
- Metadata integration (GPS, vessel speed, logs).
- Alerts suspicious activities within hours, not months.

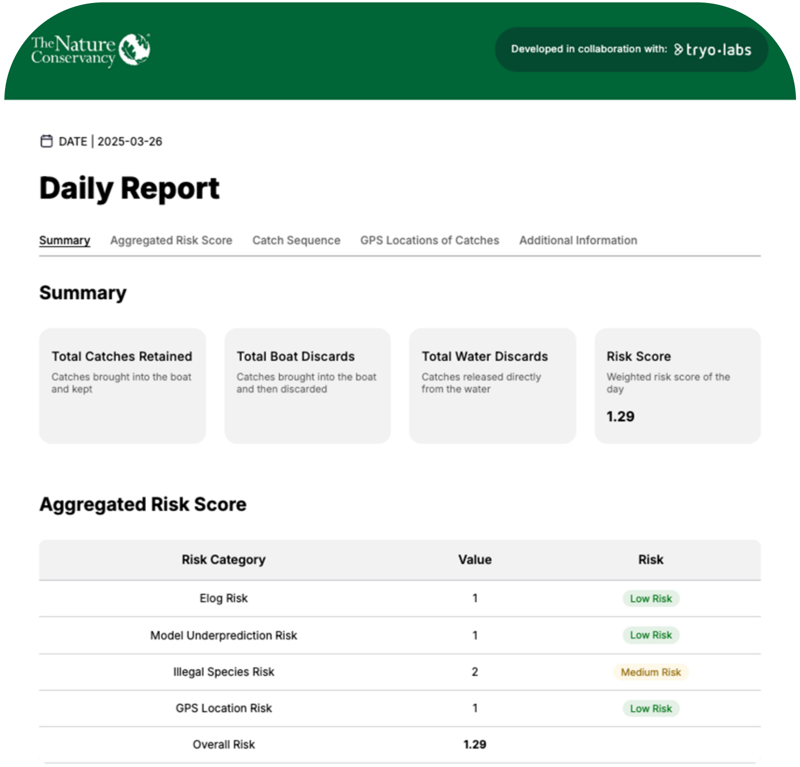

Risk Category	Value	Risk
Elog Risk	1	Low Risk
Model Underprediction Risk	1	Low Risk
Illegal Species Risk	2	Medium Risk
GPS Location Risk	1	Low Risk
Overall Risk	1.29	

The machine learning problem

The machine learning problem

Traditional Object detection

[In]
Image

[Out]
Bounding boxes

Counting problem

[In]
Video

⚠

Tracking is harder than object detection because it requires maintaining identity across frames

[Out]
Bounding boxes with unique IDs

Dataset & model architecture

Dataset characteristics

187 haulings

5,254 individual fishes

63 species

16,104 hours of video

Model architecture

Object detection: YOLO11 medium

Tracker: BoT-SORT

Object counter: Rule based

** All based on Ultralytics implementation*

Technical challenges

Why is tracking fish so hard?

- **Species look alike →** Experts generate training labels
- **Annotating fish tracks is time consuming →** Good tools and patience
- **Massive volume & privacy concerns** → Outsource time-consuming labeling
- **Fish remain in frame after being counted** → Introduce a “fishpile”
- **Severe occlusions** → Well defined labeling guidelines
- **Returned fish must be excluded** from count → Custom counting logic

Species look alike

Species look alike: expensive experts

Only experts ($$$) can **correctly identify** fish species.

There is no way around this step.

We contracted a **company of experts** that **reviews** the footage and **records** all marine species encountered.

They don't generate bounding boxes — not their job and too expensive.

We programmatically extract **expert labels** and add them as bounding boxes at fixed locations for labelers to move around.

Indo-pacific Sailfish

Swordfish

Labeling fish tracks

Labeling fish tracks is time consuming

- CVAT was our labeling tool of choice
- Built-in interpolation updates boxes between keyframes automatically
- But, abrupt movements are very hard to track

Massive volume
& privacy concerns

Massive volume & privacy concerns

- 1 hour at 12 fps → 43,200 frames to label. At 1s per frame it takes 12x the length!
- Volume is too large to label in-house, but preserving fishers anonymity is key
- Large labeling firms won't sign (custom) SoWs for “low” volume transactions
- Labeled the first haulings in-house to avoid delaying the project
- Finally found an external partner (CVAT) to label data

Introducing the fishpile

The fishpile

Problem

Fish remain visible on deck long after being counted — dismembered or piled up

Solution

Introduced “fishpile” as a separate class

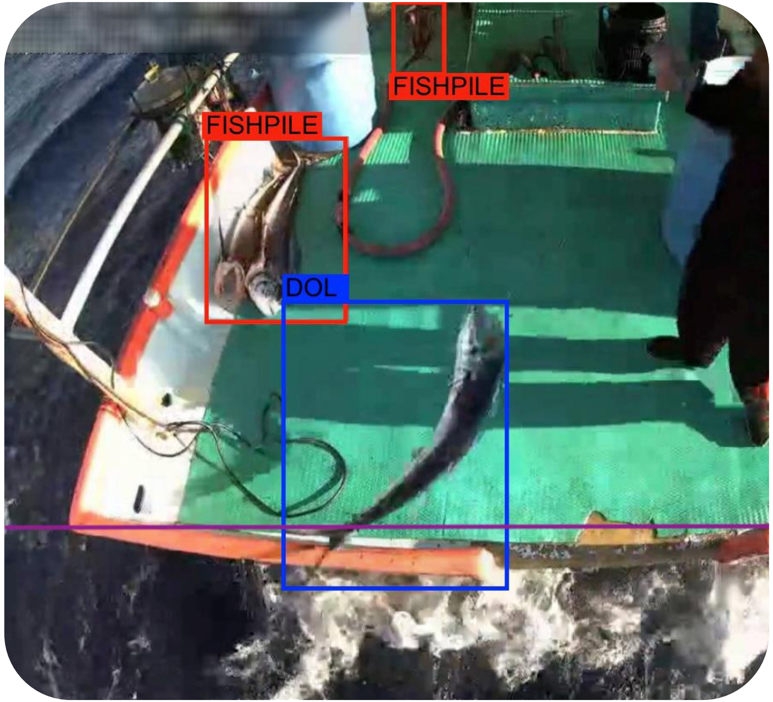

The fishpile

Pros

- Easier & faster annotation
- No need to track individual fish in piles

Cons

- Possible confusion with actual fish class during detection
- Wasted annotation effort for a class we don't actually care about

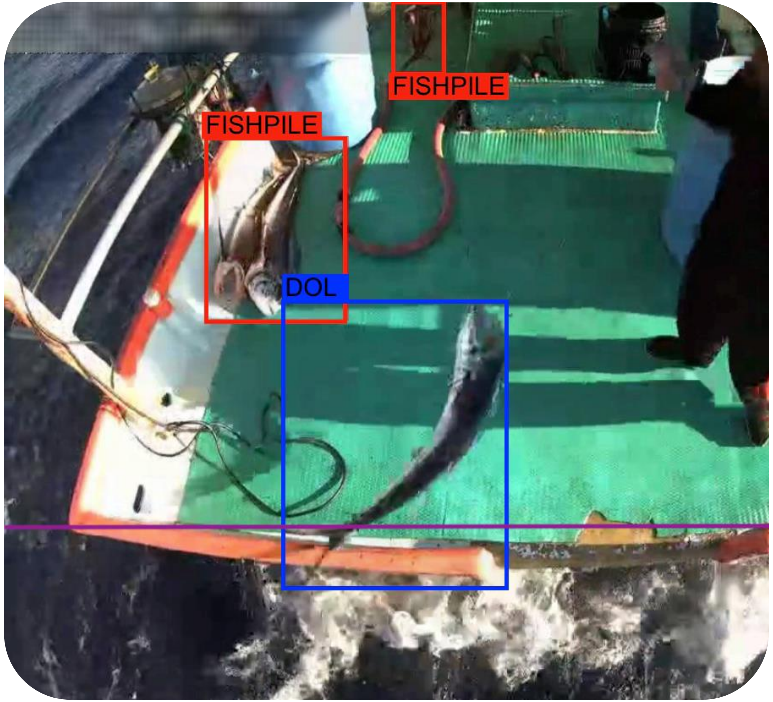

Handling occlusions with good guidelines

Consistent guidelines are key

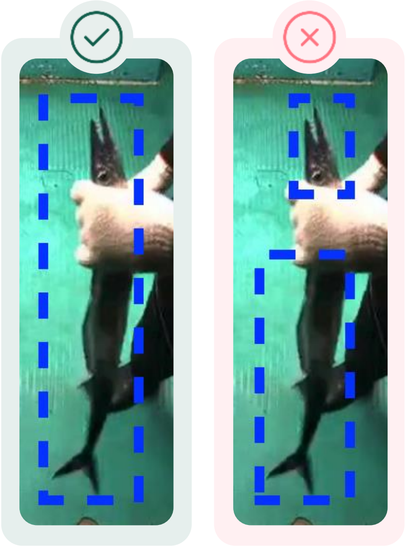

If parts of the fish are visible on both sides of the obstruction, **draw a single bounding box** that covers the **entire fish**

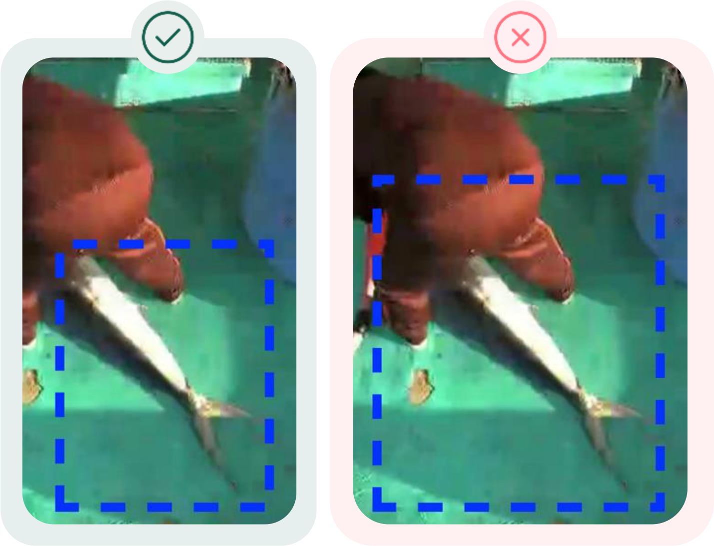

If only part of the fish is visible, **label only the visible portion**

Custom logic for counting

Counting retained and discarded species

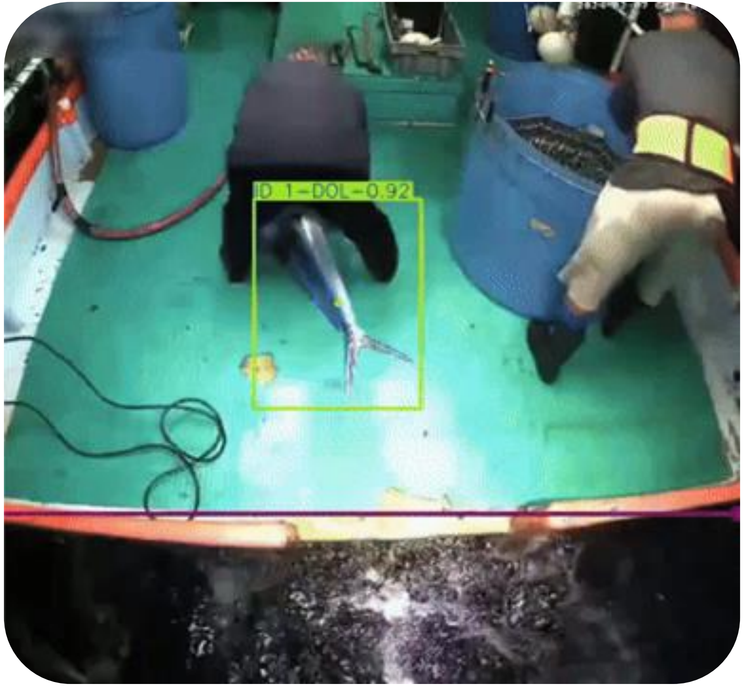

- A **virtual counting line** is drawn along the edge of the deck.
- Each fish is tracked frame by frame, and its movement history is recorded.
- When a fish crosses the line, the event is **logged**—but **not counted yet**
- The system **waits** until a few frames after crossing to gather clearer views before finalizing the count.

Counting retained and discarded species

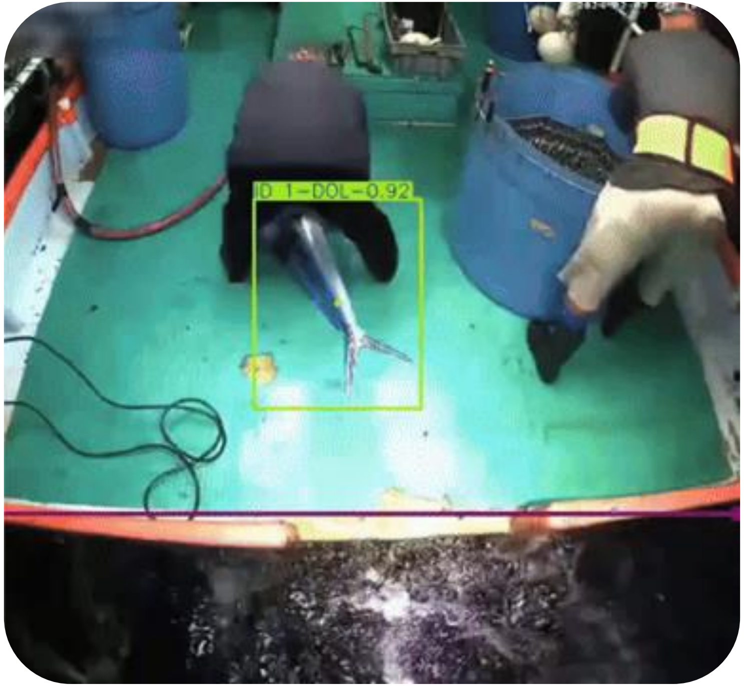

- The direction of movement determines the type of count:
  - From below to above the line → count IN
  - From above to below the line → count OUT
- The most frequent species label seen during the fish's track is assigned to the final count

Results thus far

The 2025 solution

EMBEDDED VISION SUMMIT

KPI 1 Summary Table

Icon	Class	Event	Count					Precision	Recall	f1-score	Error
			GT	Pred	TP	FP	FN				
	CATCH	IN	28	27	27	0	1	100%	96%	98%	3%
	CATCH	OUT	4	3	3	0	1	100%	75%	86%	25%
	CATCH	Retained	24	24							0%

Model: v0.0.8 | Test set: Crafted Test Dataset | Icons: Tryolabs - M&D Team

Wrapping Up

Conclusions

- EM + embedded vision technology significantly reduce the latency of footage review.
- AI is improving and initial results suggest it will be reliable in an operational setting
- Test for scalability
- Open-source solution for equitable access

Thank you!

This work is made possible thanks to an incredible network of partners and advisors, and our collaboration with the Patrick J. McGovern Foundation. Thank you!

- **Project partners:** Deckhand, Thalos, Bureau Veritas, INCOPESCA, CNIP, volunteer vessel owners, captains, and crew
- **Project advisors:** Sanjay Nichani, Rakshit Agrawal, and Jeff Bier

About Tryolabs

tryo·labs

AI, Data & Analytics
Consulting services

15 years helping companies **accelerate the adoption of AI** to drive business

Join us at Booth #806

Trusted by leaders worldwide

The Nature Conservancy · The RealReal · ULTA Beauty · unicef · PACCAR · NVIDIA · Allianz Global Investors · Hyundai Motor Group European Technical Center · LATAM · SES · HALLIBURTON · GRUBHUB

Our partners

Google Cloud · aws · NVIDIA · databricks

