Drone Sea-Surface Imaging Case Study

Oct 9, 2026 · @嚴國維

Seeing Through the Glare: A Drone Sea-Surface Screening Workflow Built with ENVI 6.3 and IDL



Figure 1. The fish aggregating device used in the survey, on deck before deployment off Keelung (23 September 2026). Red floats, black mesh panels and a tether are typical of the objects the workflow must recognize and exclude.

Fish schools near the sea surface can betray themselves as dark patches, ripples or disturbed wave texture. Spotting them from a drone sounds simple. In practice, the sea surface is crowded with sun glint, whitecaps, foam lines, vessels, fish-aggregating devices

(FADs), and their shadows, all of which can look more striking than a school of fish. At the Fisheries Research Institute (FRI), Ministry of Agriculture, Taiwan, Associate Researcher Dr. Kuo-Wei Yen is building an end-to-end workflow in ENVI 6.3 and IDL 9.3 that separates these distractions from the signals that matter.

The challenge

On 23 September 2026, the team flew a camera drone over the waters off Keelung. They collected 394 still photographs and 4K video from several survey points, flying at 20, 60 and 90 m and at speeds from hovering to 15 m/s. The goal was to answer practical questions for future surveys: which altitude, speed, and camera settings produce imagery that can actually be screened for fish schools.

Three constraints shaped the work. First, we had no confirmed fish-school images to train a detector. Second, the project had to run entirely within ENVI and IDL. Third, the deep learning module was not part of the license, so the solution had to rely on the core software and the license-free ENVI Machine Learning module.

An elimination approach

Rather than teaching software what a fish school looks like, the team asked a different question: what can be confidently ruled out? The workflow removes everything that is known, and treats what remains as a candidate for expert review.

- **Known objects** are detected and masked: FADs, their tethers, vessels, wakes, and the shadows these objects cast, predicted from sun azimuth and elevation.
- **The background** is characterized: normal sea texture and periodic waves.
- **Remaining** dark, irregular, and non-systematic anomalies become candidates for review.

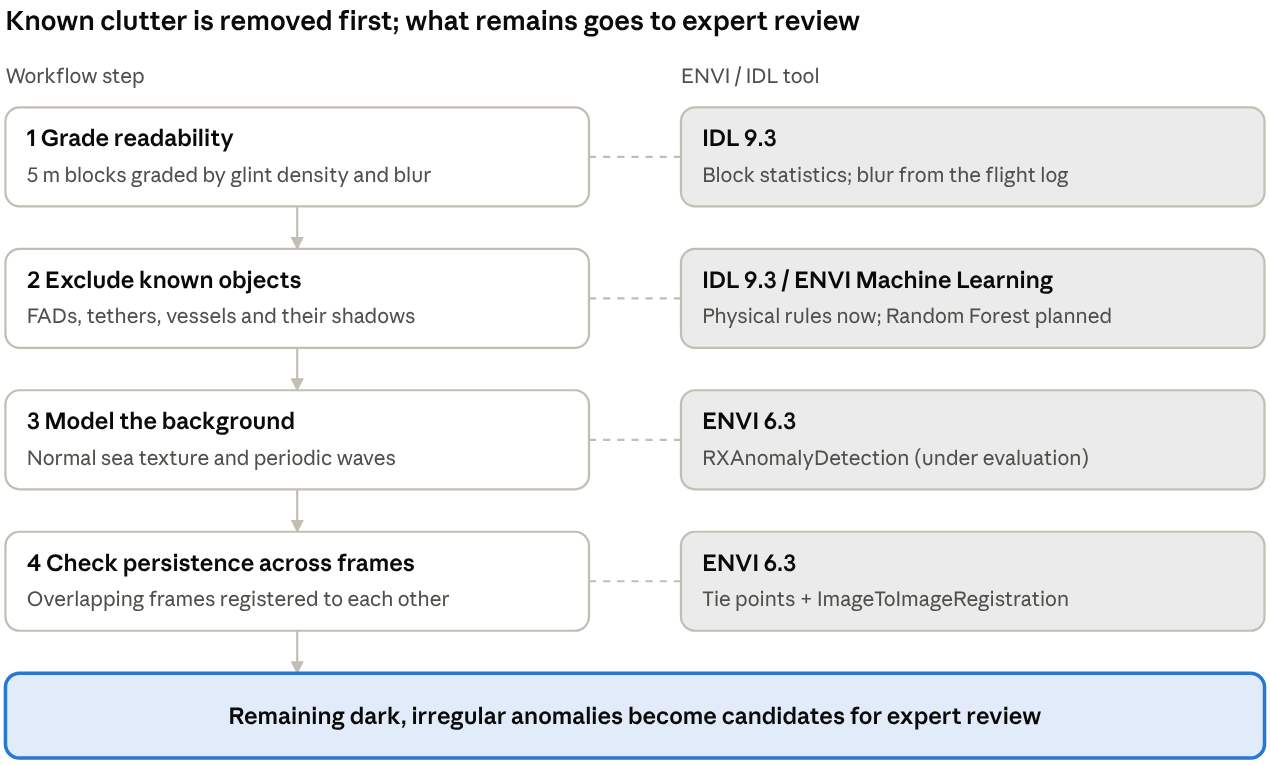


Figure 2 · elimination workflow and the ENVI / IDL tools behind each step

Because FAD designs vary widely, from bamboo rafts to steel floats, the team avoided rules tuned to a single device. Thresholds are expressed in exposure-independent quantities, such as color ratios relative to the surrounding sea, rather than raw pixel values. This lesson came directly from the data: a red-float rule tuned in daylight failed after sunset because the absolute values shifted with the light.



Figure 4. The same 90 m frame before (left) and after (right) known objects are masked: the FAD body, its tether, the vessel and the vessel's shadow. Only the remaining sea is screened for anomalies.

Knowing which pixels to trust

Not every pixel in a drone photo is readable. The workflow grades each 5 m × 5 m block of sea as readable, low-confidence (inside the influence zone of dense sun glint), or unreadable (heavy glint, motion blur, or a known object). An early version mistook scattered foam highlights for glint; switching from single bright pixels to glint density in neighboring blocks fixed this, and visual checks confirmed the grades.

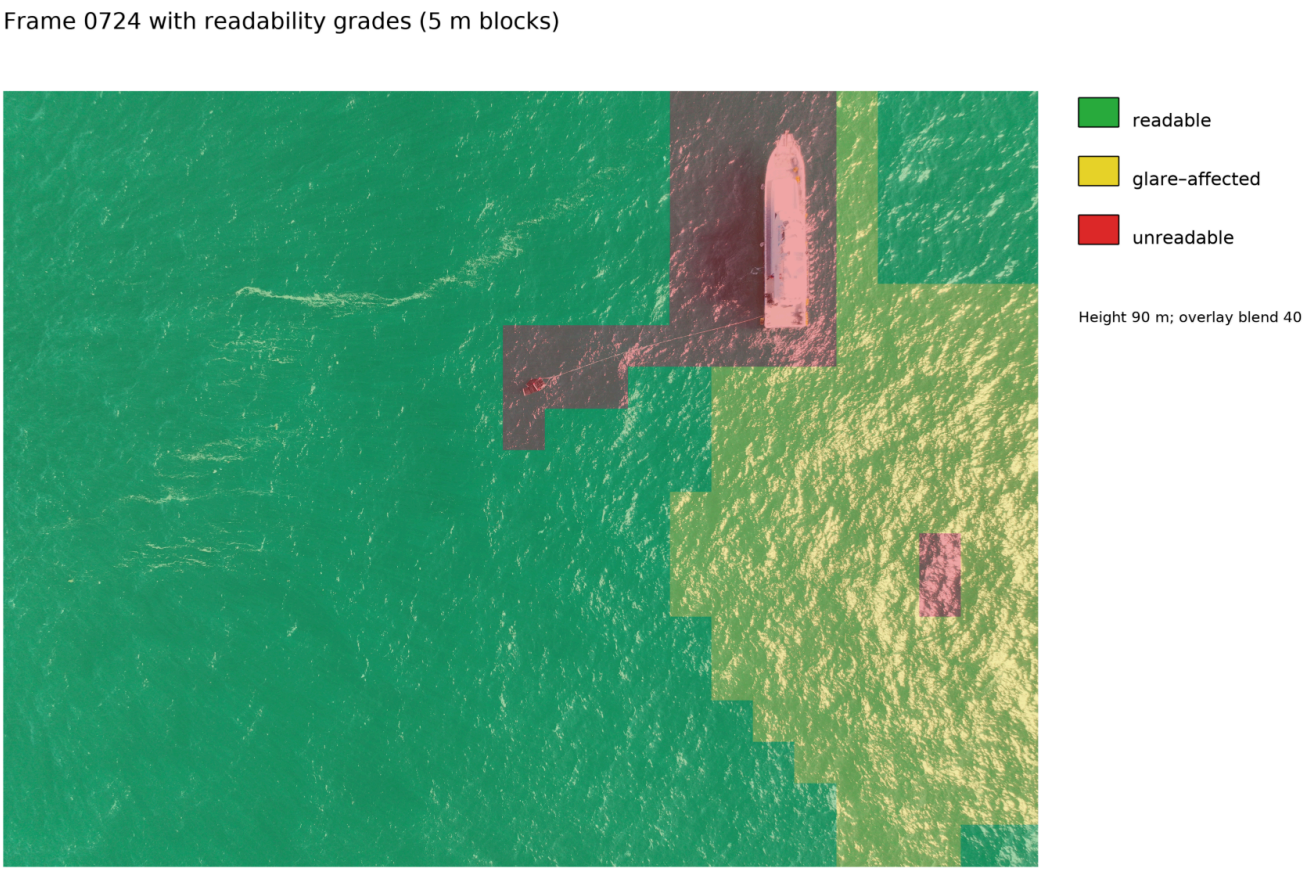


Figure 3. Readability grading of a 90 m frame. Green blocks are readable, yellow blocks lie in the sun-glint influence zone, and red blocks are unreadable or covered by known objects (here the vessel, its shadow and the tethered FAD).

A practical finding: shutter speed, not flight speed

The readability grading produced one of the project's clearest results. Of 394 photographs, 216 (55%) had a predicted motion blur of more than 3 pixels, while sun glint made only 18 unusable. The camera's automatic exposure caused this. At one survey point, it selected a shutter speed of 1/20 s, and every frame from that point was blurred. We predicted the blur using the simple relation flight speed × exposure time ÷ ground sampling distance; we are currently running an image-based check of this prediction. For future surveys, the recommendation is clear: use shutter priority at 1/1000 s or faster and accept a higher ISO.

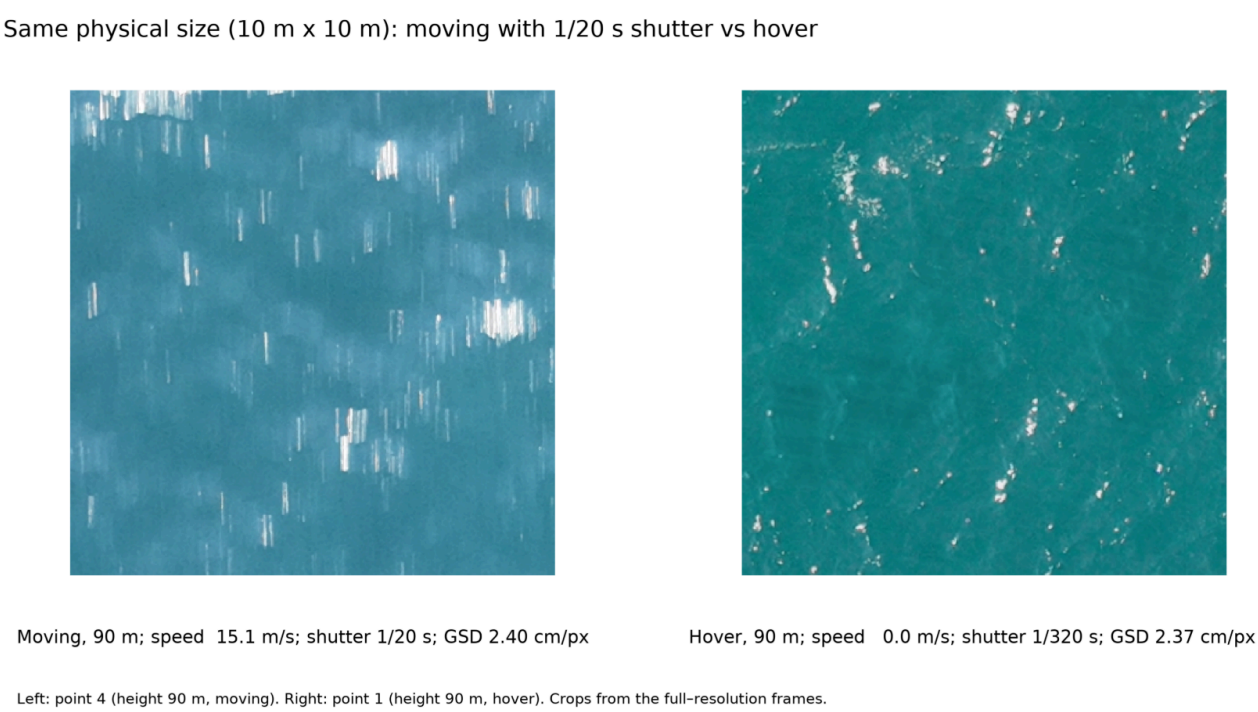


Figure 5. Two 10 m × 10 m crops from frames taken at 90 m. Left: moving at 15 m/s with a 1/20 s shutter selected by automatic exposure, giving a predicted blur of about 31 pixels. Right: hovering with a 1/320 s shutter. The two frames come from different survey points.

Measuring what can and cannot be detected

To set realistic expectations, the team inserted synthetic targets of known size and contrast into real sea imagery and ran the detector over whole frames, holding the false-alarm rate to a fixed budget per hectare. The preliminary results, from twilight imagery, are sobering and useful. Brightness alone finds only large, high-contrast targets, and the trade-off with altitude is stark: covering one hectare takes about 19 frames at 20 m but fewer than one at 90 m. These results point to the next stage: cues that single images cannot provide, such as persistence across overlapping frames.

Making the work reproducible

The workflow runs headless from IDL, with every rule, decision, and assumption versioned in Git. The team holds back validation data from one survey point and keeps it untouched until the methods are frozen. To keep the code aligned with the software, the team uses an internal AI assistant that answers questions only from ENVI and IDL documentation. It also mapped handwritten routines to native ENVI Tasks, such as image-to-image registration, RXAnomalyDetection, and ChangeThresholdClassification, and confirmed that none of these Tasks need an extra license.

Next steps

The team is now testing ENVI's registration chain for cross-frame analysis, RX anomaly detection, and Random Forest classification of known object classes. Hotspot analysis (Getis-Ord Gi*) and an evaluation of super-resolution are planned in IDL.

> "The sea surface is full of things that look more interesting than fish. Building the workflow in ENVI and IDL let us systematically rule those out. It showed us that the camera settings matter as much as the flight plan." — Dr. Kuo-Wei Yen, Associate Researcher, Fisheries Research Institute, Taiwan.

穿透海面耀光：以 ENVI 6.3 與 IDL 建立無人機海面影像篩檢流程



圖1　本次調查使用的人工集魚器，攝於 2026 年 9 月 23 日基隆外海、施放前的調查船甲板。紅色浮球、黑色網片與引繩，正是流程必須辨識並排除的物件類型。

靠近海面的魚群，可能以暗色斑塊、細碎波紋或被擾亂的浪紋露出蹤跡。從無人機上找到它們聽起來不難，實際上海面充滿了太陽耀光、白浪、泡沫條帶、船隻、人工集魚器與它們的陰影，每一樣都比魚群更醒目。農業部水產試驗所嚴國維副研究員正以 ENVI 6.3 與 IDL 9.3 建立一套完整流程，把這些干擾與真正值得關注的訊號分開。

挑戰

2026 年 9 月 23 日，研究團隊在基隆外海以空拍無人機於多個點位拍攝 394 張照片與 4K 影片。飛行高度涵蓋 20、60、90 m，航速從懸停到每秒 15 m。目的是回答未來調查的實務問題：什麼高度、什麼航速、什麼相機設定，拍出的影像才真正能用來篩檢魚群。

這項工作有三個限制。第一，沒有任何確認拍到魚群的影像可供訓練。第二，整套流程必須在 ENVI 與 IDL 內完成。第三，未採購深度學習模組，因此只能依靠核心軟體與不需額外授權的 ENVI Machine Learning 模組。

排除法

團隊不嘗試教軟體「魚群長什麼樣子」，而是反過來問：什麼可以有把握地排除？流程先移除所有已知的東西，剩下的才交給專家審核。

- **已知物件**偵測後遮罩：集魚器、引繩、船隻、尾跡，以及依太陽方位角與高度角推算的物體陰影。
- **背景**加以描述：一般海面紋理與週期性波浪。
- **剩餘的異常**：暗色、不規則、無系統性者，列為候選，交由人工審核。

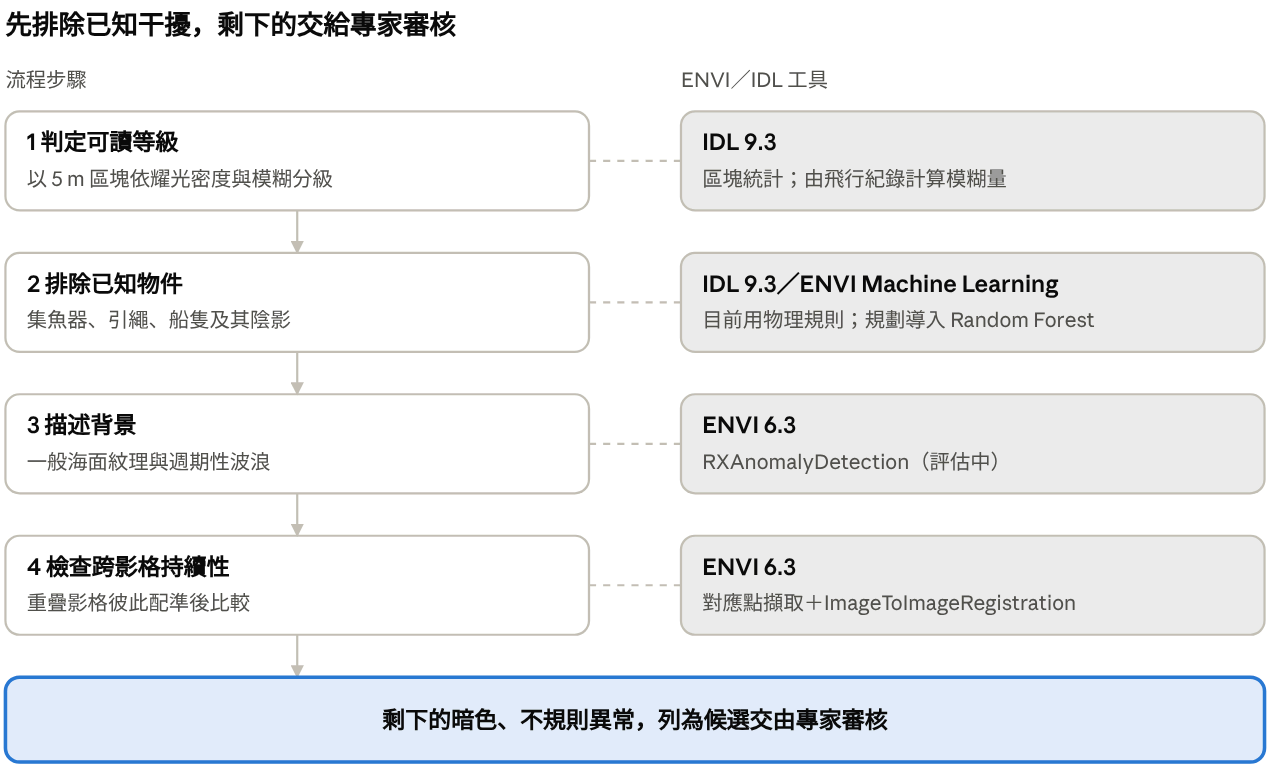


圖 2 · 排除法流程與各步驟使用的 ENVI／IDL 工具

集魚器的樣式差異極大，從竹筏到鋼製浮筒都有，因此團隊避免只針對單一款式調整規則。門檻一律採用不隨曝光改變的量，例如相對於周圍海面的顏色比例，而非相機原始數值。這個教訓直接來自資料：在日間調好的紅色浮球規則，到了日落後就失效，因為絕對數值隨光線改變。



圖 4　同一張 90 m 影像遮罩前（左）與遮罩後（右）：集魚器本體、引繩、船隻與船影都被排除，只有剩下的海面才進行異常篩檢。

哪些像素可以信任

無人機照片中並非每個像素都能判讀。流程將海面切成 5 m × 5 m 的區塊，分為可讀、低信心（位於密集耀光的影響範圍內）與不可讀（強耀光、動態模糊或已知物件）三級。早期版本把零星的泡沫亮點誤判為耀光，改用鄰近區塊的耀光密度判定後問題解決，目視檢查也確認了分級結果。

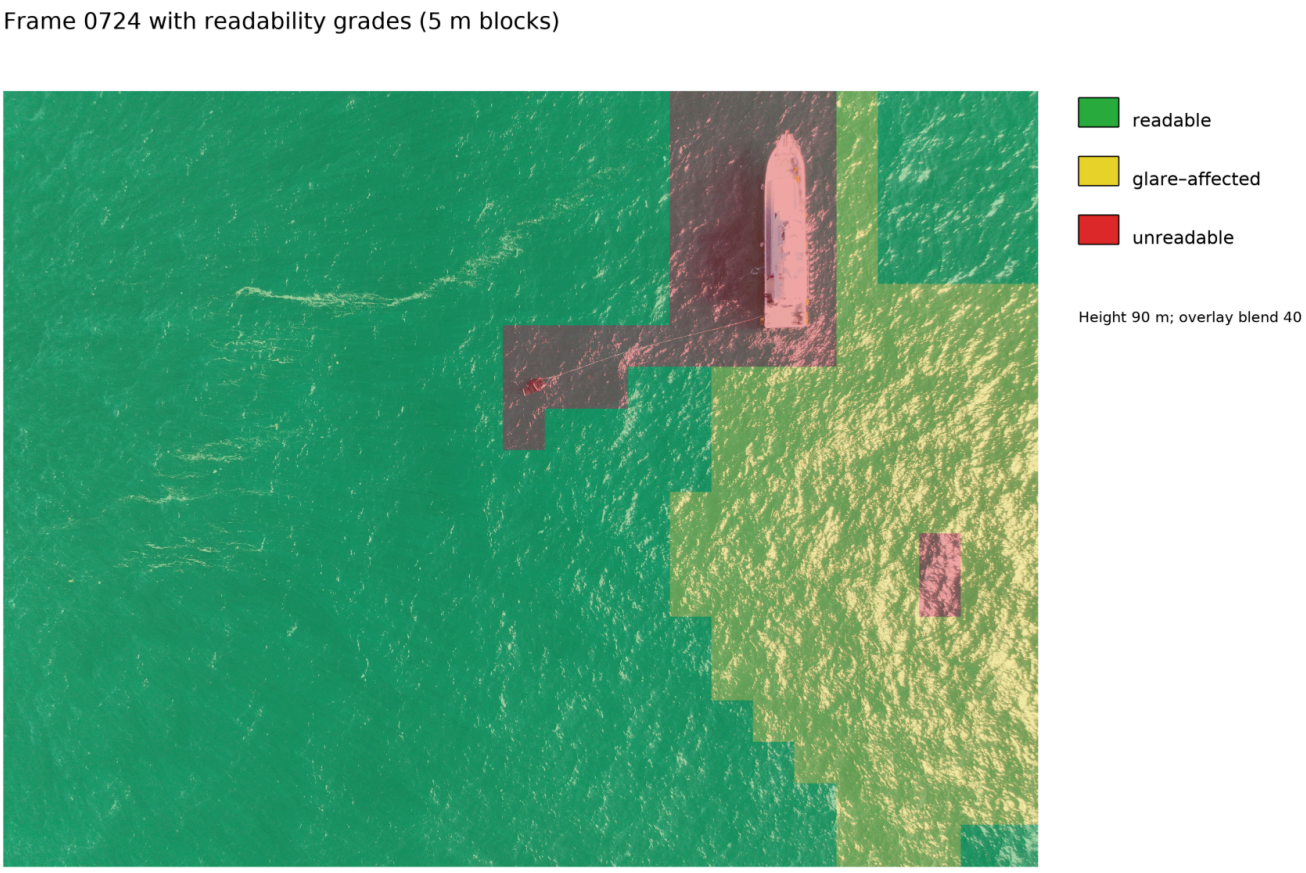


圖 3　一張 90 m 影像的可讀等級。綠色區塊可讀，黃色區塊位於耀光影響範圍，紅色區塊不可讀或被已知物件覆蓋（此例為船隻、船影與繫繩的集魚器）。

實務發現：關鍵是快門，不是航速

可讀等級分析帶出本研究最明確的結果之一。394 張照片中有 216 張（55%）預測模糊超過 3 個像素，因耀光而不可用的只有 18 張。原因是相機的自動曝光：在其中一個點位，相機自動選了 1/20 秒的快門，該點位的影像全部模糊。模糊量以「航速 × 曝光時間 ÷ 地面解析度」的簡單關係預測，以影像實際量測驗證的工作正在進行。對未來調查的建議很明確：使用快門優先模式，設定 1/1000 秒以上，寧可提高 ISO。

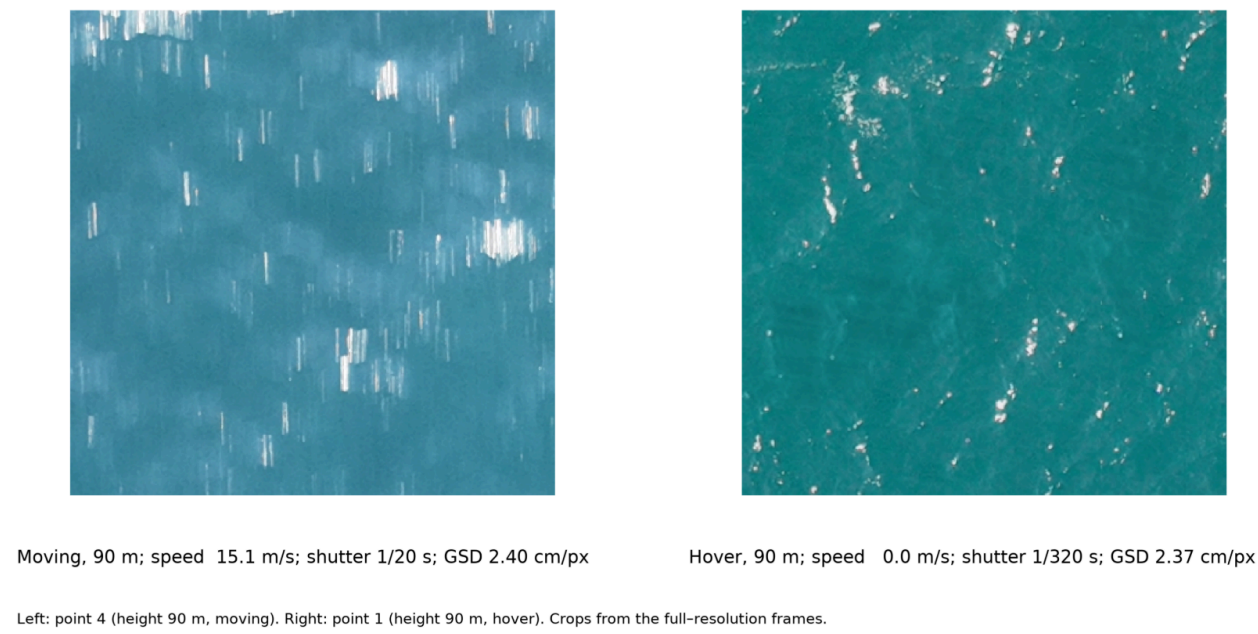


圖 5　兩張 90 m 影像各取 10 m × 10 m 並排。左：以每秒 15 m 航行，自動曝光選了 1/20 秒快門，預測模糊約 31 個像素。右：懸停拍攝，快門 1/320 秒。兩張影像來自不同點位。

量化能偵測與不能偵測的範圍

為了建立合理的期待，團隊在真實海面影像中植入已知大小與對比的合成目標，並在固定的每公頃誤報上限下，對整張影像進行偵測。初步結果來自暮光時段的影像，數字並不樂觀，卻很有參考價值：單靠亮度只能找到大型、高對比的目標；高度的取捨也很明顯，涵蓋 1 公頃海面，20 m 約需 19 張影像，90 m 不到 1 張。這些結果指出下一階段必須借助單張影像無法提供的線索，例如重疊影格之間的持續性。

讓研究可以重現

整套流程以 IDL 在無介面模式下執行，所有規則、決定與假設都以 Git 進行版本管理。其中一個點位的影像保留作為驗證資料，在方法凍結前完全不使用。為了讓程式與軟體保持一致，團隊使用一個只依據 ENVI 與 IDL 官方文件回答的內部 AI 助理。它也協助把手寫的程式對應到 ENVI 原生 Task，例如影像對影像配準、RXAnomalyDetection 與 ChangeThresholdClassification，並確認這些 Task 都不需要額外授權。

下一步

團隊正在測試 ENVI 的配準流程以進行跨影格分析，並評估 RX 異常偵測，以及用 Random Forest 分類已知物件。熱點分析（Getis-Ord Gi*）與超解析度效益評估，也將以 IDL 進行。

> 「海面上比魚群更搶眼的東西太多了。用 ENVI 與 IDL 建立這套流程，讓我們能有系統地把它們排除，也讓我們發現，相機設定和飛行規劃一樣重要。」——農業部水產試驗所副研究員 嚴國維