

Deep Learning for Accurate Vision-based Catch Composition in Tropical Tuna Purse Seiners

Xabier Lekunberri^{1,✉}, Ahmad Kamal¹, Izaro Goienetxea², Jon Ruiz¹, Iñaki Quincoces¹, Jaime Valls Miro^{1,3}, Ignacio Arganda-Carreras^{2,3,4,5}, and Jose A. Fernandes-Salvador¹

¹AZTI, Marine Research, Basque Research and Technology Alliance (BRTA), Txatxarramendi Ugarteia z/g, Sukarrieta (Bizkaia), 48395, Spain

²University of the Basque Country (UPV/EHU), San Sebastian, Spain

³IKERBASQUE, Basque Foundation for Science, Bilbao, Spain

⁴Donostia International Physics Center (DIPC), San Sebastian, Spain

⁵Biofisika Institute (CSIC, UPV/EHU), Leioa, Spain

Purse seiners play a crucial role in tuna fishing, as approximately 69% of the world's tropical tuna is caught using this gear. All tuna Regional Fisheries Management Organizations have established minimum standards to use electronic monitoring (EM) in fisheries in addition to traditional observers. The EM systems produce a massive amount of video data that human analysts must process. Integrating artificial intelligence (AI) into their workflow can decrease that workload and improve the accuracy of the reports. However, species identification still poses significant challenges for AI, as achieving balanced performance across all species requires appropriate training data. Here, we quantify the difficulty experts face to distinguish bigeye tuna (BET, *Thunnus Obesus*) from yellowfin tuna (YFT, *Thunnus Albacares*) using images captured by EM systems. We found inter-expert agreements of $42.9\% \pm 35.6\%$ for BET and $57.1\% \pm 35.6\%$ for YFT. We then present a multi-stage pipeline to estimate the species composition of the catches using a reliable ground-truth dataset based on identifications made by observers on board. Three segmentation approaches are compared: Mask R-CNN, a combination of DINOv2 with SAM2, and an integration of YOLOv9 with SAM2. We found that the last performs the best, with a validation mean average precision of 0.66 ± 0.03 and a recall of 0.88 ± 0.03 . Segmented individuals are tracked using ByteTrack. For classification, we evaluate a standard multiclass classification model and a hierarchical approach, finding a superior generalization by the hierarchical. All our models were cross-validated during training and tested on fishing operations with fully known catch composition. Combining YOLOv9-SAM2 with the hierarchical classification produced the best estimations, with 84.8% of the individuals being segmented and classified with a mean average error of 4.5%.

Computer vision | Tropical tuna identification | Electronic monitoring | Artificial intelligence | Purse seiners | Catch composition estimation

Correspondence: xlekunberri@azti.es

1. Introduction

Tropical tuna species such as skipjack (SKJ, *Katsuwonus pelamis*), yellowfin (YFT, *Thunnus albacares*), and, to a lesser extent, bigeye (BET, *Thunnus obesus*), are among the most heavily fished species worldwide [1]. Purse seine stands out over other fisheries as it is responsible for approximately 69% of the world's total catch of tropical tuna [2]. These three species play a significant role in pelagic ecosystems, acting as predators and prey, contributing to the balance and dynamics of marine food webs [3]. Given the economic value and the key role these species play in their ecosystems, ensuring sustainable fishing with healthy stocks is crucial. In this regard, access to robust data and reliable statistics on catches is essential to ensure effective and sustainable fisheries management. The total weight of the catches made during a purse seine trip can be accurately estimated at landing, but the species composition remains problematic [4]. The catch composition is manually estimated during the fishing operations and recorded in the logbook. However, bias in logbooks has been evidenced since the beginning of the fishery, due to the difficulty in differentiating some species and the large number of individuals in each capture. This bias primarily concerns juvenile YFT and BET individuals, preventing the use of these records as precise estimators [5].

Traditional fisheries monitoring implies that an onboard observer must supervise the fisher's actions directly at sea. The electronic monitoring (EM) systems, already in use in purse seiners, have been proposed as complement to onboard observers. These EM systems consist of several sensors installed onboard the fishing vessel, such

as GPS receivers or CCTV cameras, and they are responsible for recording and storing data generated while the vessel is at sea [6]. The EM can complement, or even replace, external observers, overcoming challenges such as the high cost and safety concerns of bringing additional personnel to the vessel [7, 8]. However, manual review of EM footage is still tedious, as human analysts must watch hours of video and cross-reference it with sensor data to document fishing events.

Several studies have evaluated the effectiveness and improvement over time of EM technology for catch quantification, species identification, and length/weight measurement [9–16]. These works revealed that EM footage analysis can accurately determine key metrics, such as the total tuna catch per set and fishing effort, but requires improvements regarding the accurate collection of species data. Ruiz et al. [17] demonstrated that, even if sampling is done onboard, obtaining precise estimates of species composition per fishing operation is often impractical without extensive and time-consuming sampling because of three main factors: (1) insufficient minimum sample size, (2) difficulty distinguishing juvenile YFT from BET, and (3) challenges in ensuring random sample selection. Such intensive sampling is incompatible with commercial operations, as minimizing the time between capture and freezing is crucial for preserving tuna quality [18]. Recent advances in artificial intelligence (AI) have led to the development of new methods that may overcome the challenges that traditional monitoring and EM still pose, providing a way to monitor the catches accurately.

AI-based systems have already been developed in other fisheries to improve or facilitate species identification. For example, Palmer et al. [19] estimated both the number and size of common dolphinfish (*Coryphaena hippurus*) at a central fish auction. Using Mask R-CNN [20] for segmentation, they achieved a mean average precision (mAP) of 79.8% and a root mean squared error (RMSE) of 2.4 cm in length estimation. Other studies have opted for computationally faster architectures such as YOLO for detection and segmentation tasks. For instance, Sokolova et al. [21] modified YOLOv5 to assign species and measure weights in discarded catches from demersal trawlers. They used a 5-fold cross-validation strategy, achieving a macro F1-score of 94.1% and a weight prediction error of 29.74 grams.

Building on this, Córdova et al. [22] introduced a multi-stage approach that separated detection, classification, and weight estimation into distinct components, further improving accuracy.

Beyond task-specific models, one of the most significant advancements in recent years has been the emergence of foundation models [23]. Trained on massive, diverse datasets, these models provide general-purpose visual representations that can be adapted for a wide range of downstream applications. Despite their versatility, a recent survey indicates that foundation models still require domain-specific fine-tuning to achieve optimal performance [24]. Notable examples include DINOv2 [25], pretrained on a dataset of 142 million images for tasks such as depth estimation or semantic segmentation, and the Segment Anything Model 2 (SAM 2) [26], trained on over a billion masks for image and video segmentation. SAM2 enables zero-shot segmentation of arbitrary objects in photos or videos through simple prompts (e.g., points or bounding boxes). Early applications on underwater fish detection show promising results, but also reveal that their performance is highly dependent on the type of prompt [27, 28].

Despite the remarkable accuracy in some studies, a critical gap remains between performance on controlled environments and the robustness required to implement the pipelines in real-world environments [29]. In our previous work [30] we achieved limited species identification in tuna purse seiners Using historical EM data. Several key points were identified that needed to be improved to develop a robust automated system: (1) the installation of a task-specific cameras providing a top-down view of the conveyor belt, (2) the availability of improved datasets for model training, and (3) accurate ground-truth data for operational testing, beyond statistical validation alone, to mitigate the impact of biased training samples. This present work directly addresses these challenges by leveraging state-of-the-art deep learning methods to improve both segmentation and classification performance of our previous work. To our knowledge, this is the first study to process EM data collected onboard tuna purse seiners, where both train and test datasets derive from real fishing operations. This unique setting enables direct comparison of our computer vision-based pipeline against actual catch composition, granting a precise picture of the system’s perfor-

mance.

2. Material and methods

This section outlines the complete pipeline developed for estimating catches [Figure 1](#). The workflow consists of four main stages: (1) data acquisition, (2) image labeling, (3) model training and validation, and (4) ground-truth testing of the full pipeline.

Initial evaluation of historical EM footage was conducted in collaboration with nine taxonomy experts. The low agreement observed among experts regarding the identification of two of the three target species motivated the acquisition of new ground-truth data directly onboard tuna purse seiners. We selected videos in which the catch composition was perfectly known and only a single species appeared on the conveyor belt to train the models. On the other hand, videos with multiple species were used for the testing.

Training frames were extracted and uploaded to Computer Vision Annotation Tool (CVAT) for annotation [\[31\]](#). Once annotated, segmentation and classification models were trained in parallel, as detailed in [subsection 2.3](#) and [subsection 2.5](#), respectively. Two approaches were developed for both segmentation and classification, each trained and validated using repeated 5-fold cross-validation and subsequently tested on ground-truth data to simulate real-world performance. To assess the accuracy of the final predictions, results from the test phase were compared against the actual species distributions using a Wilcoxon signed-rank test.

2.1 Data acquisition and preprocessing.

Identifying species via images is not always feasible, as the traits that make species unique may not always be visible [\[32\]](#). Two of the target species of this fishery, BET and YFT, are visually very similar, making it difficult or even impossible for analysts to identify them correctly. To quantify this difficulty, we conducted an experiment in which we brought together nine different experts in tuna identification to examine 40 different images. In that experiment, each one of the experts annotated all the tunas they could identify with a high degree of confidence. The dataset was restricted to BET and YFT individuals, since SKJ specimens are easy to identify. The low agreement among experts highlighted the difficulty of distinguishing BET from YFT in EM

footage and led us to discard previously annotated data in favor of generating a new ground-truth dataset through an alternative approach. The findings will be discussed in the Results and Discussion sections.

The raw video data used in this work was recorded aboard vessels equipped with one or two conveyor belts in the fish processing area. Vessels with a single belt carried two EM cameras positioned near the bow and stern, while vessels with a second belt also included an additional camera on that belt., at the midpoint of the boat. All cameras recorded fishing operations at a 1920x1080 pixel resolution and 30 frames per second (FPS). During the trials at sea, our onboard personnel installed additional cameras to record extra footage. We used the ZED 2i camera for the first trip and upgraded to the ZED X series for the second. Both models are from Stereolabs and employ stereoscopic technology to capture the depth information in addition to standard RGB footage. Although 3D reconstruction was not explored in the present study, combining depth data with segmentation results would enable three-dimensional reconstruction of individual fish, potentially offering new opportunities for species classification as well as better size and weight estimation.

The onboard scientists assisting with data collection had vast experience working on tuna purse seiners and were able to reliably identify target species in situ due to their background. Their knowledge was essential, as the data they collected was the basis for the pipeline described here.

During each fishing operation, the procedure was as follows: the conveyor belt was stopped, and all individuals from a given section were transferred into buckets. This process was repeated several times per operation. The number of individuals sampled varied depending on the total catch and the number of belt stoppages. Once the operation concluded, all sampled fish were individually identified to species, measured, and sorted to make up artificial fishing operations (AFOs) used for training and testing the models.

To do so, the identified tunas were placed back on the conveyor belt and recorded with the EM cameras under standard conditions. This procedure allowed us to simulate smaller-scale real fishing operations with fully known species composition and size distributions. For training, the individuals were separated into batches by species be-




Figure 1. Workflow that outlines the pipeline for fish image analysis. It begins with the collection and manual annotation of images, followed by model training using cross-validation to select the best-performing architecture. The selected model is then used to automatically segment and classify fish in images. Finally, the pipeline is tested on real-world data. This process ensures robust and accurate fish identification.

fore being returned to the conveyor belt, ensuring that all individuals selected for manual annotation were correctly labelled. In contrast, in AFOs used for testing purposes, the individuals were returned to the conveyor belt in a single mixed batch, emulating the conditions of real fishing operations.

Image preprocessing followed the methodology described in Lekunberri et al. [30]. The key step in this pipeline was the selection of the conveyor belt region through which the catches pass, correcting the perspective when necessary. This correction was not required when cameras were installed in a top-down position that fully covered the conveyor belt within the field of view and equipped with a lens that corrects optical distortions. Also, a CLAHE (Contrast Limited Adaptive Histogram Equalization) filter was applied to the images to homogenize the histograms and improve overall contrast.

2.2 Manual image annotation. All images were annotated using the Computer Vision Annotation Tool (CVAT), chosen for its maturity, team collaboration features, and the annotators' prior experience with the software. Labels were organized hierarchically, with species-level annotations prioritized. Target species (BET, SKJ, and YFT) were labeled according to the 3-alpha code of the ASFIS List of Species for Fishery Statistics Purposes [33].

In total, 1,395 images were annotated, yielding 24,078 individual fish segments. Some fishing operations from historical data were reused because of their exceptional quality and because the work of labeling them had already been done. In those cases, BET and YFT labels were changed for an additional joint label (BET_OR_YFT). That label and others generated by combining multiple categories were employed exclusively for the segmentation (subsection 2.3) and the hierarchical classification approach (subsection 2.5).

For non-target species captured as bycatch, all individuals were merged into a single NO_TARGET label for classification, although species-level annotations were still collected to support future model improvements. A generic FISH label was added to maximize the number of annotated individuals for the segmentation, even though it did not carry species-specific information. The distribution of annotated segments across labels is shown in Table 1.



Figure 2. Manually annotated image. This image contains a masks for each fish on the conveyor belt that is visible from the point of view of the camera. This frame belongs to a monospecific SKJ sample. It was used during the training of the models.

2.3 Automatic image segmentation. The conventional approach to segmentation and classification addresses both tasks in a single stage, assigning distinct labels to each species. In our case, however, we adopted a two-stage strategy. Since all target species share a highly similar morphology, we replaced the species-specific labels with a single class during the segmentation model's training. This approach allowed us to concentrate annotation efforts on a single segmentation class and to address classification separately using various techniques, as detailed in subsection 2.5. The collected images and manual annotations (Figure 2) were used to train and compare three segmentation approaches:

1. Mask R-CNN [20],
2. YOLOv9 [34] combined with SAM2 [26], and
3. The pipeline described by González-Marfil et al. [35], which performs zero-shot image segmentation using DINOv2 embeddings [36] to compute similarity maps and employs SAM2 to generate segmentation masks of semantically similar objects. This last approach required no training or fine-tuning.

To expand the training set, we applied data augmentation. Image brightness, contrast, and saturation were randomly adjusted to simulate different lightning conditions. Random horizontal and vertical flips were applied to the images and masks to reduce orientation bias. Table A.4 summarizes the computing infrastructure, hyperparameter search spaces, and data augmentation parameters used in model training.

Table 1. Number of segments used to train our models. The segments are also broken down according to their label in a hierarchical manner. The top-level labels are TARGET and NO_TARGET, corresponding to species targeted by this fishery (BET, SKJ, and YFT) and those not (all others). Target species are divided into SKJ and BET_OR_YFT depending on their visual similarity. Finally, BET_OR_YFT is divided into the corresponding species. The generic FISH label does not contain species-specific information.

	TARGET 18.232		NO_TARGET	FISH
SKJ	BET_OR_YFT 10.792		1.204	4.642
7.470	BET 3.592	YFT 6.991		

Mask R-CNN, used in the first approach, is a two-stage instance segmentation algorithm that builds upon Faster R-CNN [37] by incorporating a branch for predicting segmentation masks on each Region of Interest (RoI). The process consists of two steps: the first generates proposals for regions that might contain objects, and the second refines these proposals and predicts masks. In our previous work [30], we used the “Mask R-CNN Inception ResNet V2 1024x1024” implementation from TensorFlow’s Object Detection API, from which we obtained a mean average precision (mAP) of 0.66. Since that API has been deprecated, we migrated to PyTorch and the Torchvision implementation.

In the second and third approaches, segmentation relied on SAM2, a state-of-the-art foundation model capable of segmenting images and videos across multiple domains. Without fine-tuning, SAM2 successfully segmented nearly every fish in our dataset. Its main limitation is the need for an initial prompt (bounding box or point) to locate a fish. The approaches differ in how these prompts were provided. In the second approach, we used the output of a fine-tuned YOLOv9 model as prompts for SAM2. In the third, we exploited DINOv2 embeddings to compute similarity maps based on a reference set of annotated fish, with SAM2 then generating masks for objects detected as semantically similar across other frames.

Segmentation models were evaluated through repeated cross-validation, with performance reported as the mean and standard deviation of both mAP and recall. A predicted segment was considered valid when its intersection over union (IoU) with the ground truth exceeded a predefined threshold. The IoU is measured as the area of the predicted bounding box that overlaps with the ground truth bounding box divided by the total area of both bounding boxes. The mAP is calcu-

lated as the area under the precision-recall curve across IoU thresholds from 0.05 to 0.95 in steps of 0.05.

2.4 Individual fish tracking along the conveyor belt.

Since the fish can move freely on the conveyor belt and the belt occasionally stops, some individuals are captured in more frames than others. The videos are processed frame by frame, resulting in independent segmentations for each frame. Consequently, some individuals may be overrepresented, potentially biasing the final species estimations. Tracking each fish ensures that it is counted only once.

In addition, having multiple images of the same fish might enhance the classification process, as it captures variations in lighting, pose and orientation across the camera’s field of view. This variability improves the robustness of the classification by providing complementary visual information that would otherwise be omitted if only a single image were used.

The original SAM2 implementation is capable of single-object tracking, and some extensions have added support for multiple-object tracking. However, we implemented our own tracking solution to ensure consistency across all segmentation approaches. Different algorithms for multi-object tracking, such as SORT [38], DeepSort [39], or ByteTrack [40], were considered. Studies indicate that no single algorithm consistently outperforms the others, as a trade-off between computational efficiency and tracking precision must be considered [41–43].

We ultimately selected ByteTrack as the individuals in our images are moving in the same direction, and even if a segment’s score was low due to partial occlusion, it was still used to try to match active trackers. Furthermore, ByteTrack’s high computational efficiency was also considered, as any delay introduced by tracking directly affects

the overall processing time of the pipeline. We also included optical flow estimation to determine when operators stop or reverse the conveyor belt's direction. This allows us to pause tracking when the belt stops moving in the right direction and resume tracking when it starts moving again, which ensures that individuals that have left the camera's field of view are not tracked when they reappear.

2.5 Fish classification by species. We evaluated two approaches for species classification:

1. a standard multiclass classification, and
2. a hierarchical classification, where the final label is reached through multiple sequential decisions.

In both approaches, we used the RegNetX400 architecture [44]. As in the segmentation step, random brightness, contrast, and saturation modifications were used as data augmentation techniques, in addition to random horizontal and vertical flips. See Table A.4 for the computing infrastructure, hyperparameter search spaces, and data augmentation parameters used during training.

The first approach follows a conventional multiclass scheme similar to our previous work [30], where a pretrained ResNet50V2 model from TensorFlow's Object Detection API was used. That model achieved approximately 75% of correct classifications. Since the API was deprecated, we migrated to PyTorch and adopted a newer architecture for improved performance and long-term support.

The second approach follows a hierarchical classification strategy, in which labels are grouped in pairs and progressively refined. At each level, one class is final (requiring no further classification), while the other proceeds to the next decision level. This hierarchy includes three classification stages:

1. Non-target species vs target species.
2. Skipjack tuna vs other tunas.
3. Bigeye tuna vs yellowfin tuna.

This structure was designed to increase classification complexity gradually. Although hierarchical classification requires multiple inferences per sample—thus increasing computational time—it enables finer discrimination between species. Similar hierarchical strategies have been

successfully applied to tuna classification in simpler contexts, where individuals do not overlap and are photographed separately [45]. In both approaches, multiple predictions of the same fish were averaged to produce a final class prediction. As the same individual appears in several frames under varying lighting conditions and orientations, this aggregation increases robustness and reduces sensitivity to visual variability.

Finally, all classification strategies were tested independently and in combination. The final system may integrate any subset of these methods, depending on their performance trade-offs. Although hierarchical and weighted schemes introduce a computational cost, they can substantially improve classification accuracy, making the trade-off advantageous in practice.


2.6 Statistical validation and ground truth testing.

All models were validated using robust, stratified, and repeated cross-validation. Except for YOLOv9, which was limited to five repetitions due to training time constraints, all segmentation and classification models underwent 10-times 5-fold repeated cross-validation. Model performance was evaluated using task-appropriate metrics: segmentation models were assessed with **standard COCO metrics**, including recall and mean average precision (mAP) across multiple segment sizes and intersection over union (IoU) thresholds, while classification models were evaluated through confusion matrices and accuracy measures. After validation, all models were tested on out-of-sample data. We used only the portion of the AFOs collected during sea trials that had not been used for training. These samples were collected exclusively for this work, ensuring that ground truth and modeled data corresponded to the same individuals.


The AFOs were processed through the entire pipeline (segmentation, tracking, and classification) with every possible approach combination to assess the overall performance. Finally, Wilcoxon signed-rank tests were employed to determine if statistically significant differences existed between the actual species distribution and the predictions generated by our pipeline.

3. Results

This section presents the outcomes of: (1) the expert-based experiment assessing the ability to distinguish between BET and YFT using EM im-



(a). Number of individuals annotated by each of the nine experts.



(b). Number of individuals, split by species, that have been labeled at least by X experts.

Figure 3. Different statistics for the analysis based on expert identification.

ages, (2) the training and validation of all models, and (3) the comparison of model predictions against the ground-truth data collected onboard different vessels.

3.1 BET and YFT classification based only on images. Nine experts in tuna taxonomy independently identified BET and YFT individuals in a set of 40 EM images. Experts were instructed to label only those fish they could identify with high confidence. Three experts identified substantially more tunas than the others, resulting in a total of 257 unique individuals being annotated (Figure 3). Refer to Table D.7 in the Supplementary Materials to view all the labels.

To calculate the inter-expert agreements, individuals identified by at least four (more than a third) experts were selected. We found that the agreement rate was $42.9\% \pm 35.6\%$ for BET and $57.1\% \pm 35.6\%$ for YFT. There was only consensus among the nine experts for one individual per species. These results indicate that ex-

perts cannot agree on identifying BET and YFT when only electronic monitoring images are available. Consequently, we concluded that remote annotation could not provide sufficient ground truth accuracy. To overcome this limitation, new data were acquired *in situ* onboard tuna purse seiners, where experts could visually confirm the species of each fish during the fishing operation. This ensured that the training footage for our models contained only one species at a time, eliminating ambiguity in subsequent image annotation.

3.2 Segmentation of the images. The combination of YOLOv9 and SAM2 yielded the best performance in terms of both mAP and recall, with statistically significant improvements on the test AFOs, as will be shown later. The cross-validated Mask R-CNN achieved a mAP of 0.63 ± 0.01 for bounding boxes and 0.67 ± 0.01 for segmentation masks. YOLOv9, on the other hand, obtained a mAP of 0.66 ± 0.03 . In terms of recall, Mask R-CNN reached 0.69 ± 0.01 , while YOLOv9 achieved 0.85 ± 0.03 .

No segmentation scores are reported for the second approach, as SAM2 weights were used without fine-tuning. Similarly, we do not provide segmentation scores for the third approach, where DINO was used to compute similarity maps, since this method also does not involve model fine-tuning. Figure 4 shows a raw frame from one of the AFOs together with the corresponding segmentations obtained by each approach. Following the method described by González-Marfil et al. [35], we chose not to include this method in further comparisons, as generating a single frame took up to ten minutes on a laptop GPU. Integrating such a time-consuming process into a future onboard application running on an embedded device would not be feasible, as near-real-time performance is required.

3.3 Classification of the fishes by species.

Four classification models were trained: one using a standard multiclass approach and three using a hierarchical scheme. As with segmentation, all models were cross-validated and tested on out-of-sample data. Figure 5 shows the confusion matrices and corresponding mean accuracies.

Both approaches achieved comparable overall validation scores, but the main difference lies in the NO_TARGET class, where the non-hierarchical classifier obtained a significantly bet-




Figure 4. Comparison of all the segmentation approaches. It can be noted that the right part of the camera has a small water droplet, making the fish in that area blurry and more difficult to segment.

ter accuracy (0.90 ± 0.03 vs 0.69 ± 0.11). This may be explained by (i) the heterogeneous nature of the NO_TARGET label, which groups multiple non-targeted species, including tuna-like fish visually similar to target tunas, and (ii) the limited number of training images for these classes. This imbalance is much more noticeable in the hierarchical classifier, as a one-vs-rest strategy binarizes each step. It should be noted that bycatch species were not targeted during the data gathering step, and that the label has not been prioritized in the development of this work.

For the targeted species, SKJ showed consistently high accuracy in both approaches (0.92 ± 0.02 and 0.95 ± 0.01), while BET and YFT exhibited the lowest performance, as expected given their visual similarity. Even within the hierarchical framework, notable confusion persisted. Approximately one out of every four BET specimens is misclassified as yellowfin tuna YFT. Although in the hierarchical model the error is somewhat higher for this class, the proportion of bycatch in the total catch rarely exceeds 5%. Therefore, in absolute terms of catch, the error is not very significant.

3.4 Testing with artificial fishing operations.

We evaluated the complete pipeline using 21 artificial fishing operations (AFOs) containing multiple species, from two fishing trips conducted on the same vessel (14 and 7 AFOs, respectively). Both segmentation methods segmented more than half of the caught tunas, but YOLOv9 + SAM2 clearly outperformed Mask R-CNN (68.5% vs 52.1%; Table 2). This superiority was consistent across trips (60.3% vs 46.4% and 84.8% vs

True label	BET	SKJ	YFT	NO_TARGET
	0.79 ± 0.03	0.034 ± 0.01	0.17 ± 0.03	0.011 ± 0.0
	0.013 ± 0.0	0.92 ± 0.02	0.044 ± 0.01	0.02 ± 0.01
	0.12 ± 0.03	0.065 ± 0.01	0.78 ± 0.03	0.031 ± 0.01
NO_TARGET	0.018 ± 0.01	0.047 ± 0.02	0.039 ± 0.02	0.9 ± 0.03
	Predicted label			

(a). Standard classification.

True label	TARGET	NO_TARGET
	0.95 ± 0.02	0.046 ± 0.02
NO_TARGET	0.31 ± 0.11	0.69 ± 0.11
Predicted label		

(b). Hierarchical classification, step 1/3.

True label	SKJ	BET_OR_YFT
	0.95 ± 0.01	0.046 ± 0.01
BET_OR_YFT	0.074 ± 0.01	0.93 ± 0.01
Predicted label		

(c). Hierarchical classification, step 2/3.

True label	BET	YFT
	0.76 ± 0.04	0.24 ± 0.04
YFT	0.12 ± 0.03	0.88 ± 0.03
Predicted label		

(d). Hierarchical classification, step 3/3.

Figure 5. Confusion matrices for the four classification models. These matrices are built using the validation data.

63.5%) and statistically significant in all cases (Wilcoxon signed-rank, $p < 0.01$). Therefore, we selected YOLOv9 + SAM2 as the preferred segmentation approach.

Using this segmentation, we compared the stan-

Table 2. Mean and the standard deviation of the percentage of the catches segmented by each segmentation approach. Several Wilcoxon signed-rank tests were run for $p < 0.10$ (*), $p < 0.05$ (†), and $p < 0.01$ (§), with the null hypothesis being that there are no differences in the amount of individuals segmented by the two approaches.

	First trip (§)	Second trip (§)	Both trips (§)
Mask R-CNN	46.4% $\pm$ 6.4%	63.5% $\pm$ 11.5%	52.1% $\pm$ 11.7%
YOLOv9 + SAM2	60.3% $\pm$ 9.6%	84.8% $\pm$ 13.5%	68.5% $\pm$ 16.0%

Table 3. Mean average error (MAE) for classifying individuals segmented by the best approach (i.e., YOLOv9 + SAM2). The best classification approach (i.e., standard or hierarchical) was selected for each trip. Several Wilcoxon signed-rank tests were run for $p < 0.10$ (*), $p < 0.05$ (†), and $p < 0.01$ (§) with the null hypothesis being that there are no differences in the fish counted by the two methods (i.e., the ground truth and the evaluated one).

	First trip	Second trip	Both trips
BET	2.8% $\pm$ 2.4%	2.1% $\pm$ 2.1%	2.3% $\pm$ 2.3% (*)
SKJ	22.9% $\pm$ 11.4% (§)	5.0% $\pm$ 4.7%	19.8% $\pm$ 11.0% (*)
YFT	20.7% $\pm$ 10.8% (§)	5.6% $\pm$ 8.0%	16.1% $\pm$ 11.7% (§)
NO_TARGET	2.8% $\pm$ 1.8% (§)	5.4% $\pm$ 2.8%	3.8% $\pm$ 4.5%

dard and hierarchical classification schemes. Here, we calculate the error for a given species as the absolute difference between the predicted abundance and the ground truth. The standard classifier achieved a lower mean absolute error (MAE) overall (10.5% vs 17.8%), corresponding to roughly two-thirds of the hierarchical model’s error. The results varied by trip. The standard approach performed better on data collected during the first trip (12.3% vs 24.5%), while the hierarchical classification slightly improved results on data collected from the second trip (4.5% vs 6.9%). The differences between the ground truth and the predictions of our model also differ, as the errors for the second trip are much smaller for both standard and hierarchical approaches. Species-wise, using our best segmentation and classification approaches, we got a MAE of 2.1% $\pm$ 2.1% for BET, 5.0% $\pm$ 4.7% for SKJ, 5.6% $\pm$ 8.0% for YFT, and 5.4% $\pm$ 2.8% for NO_TARGET (Table 3). Distribution tests confirmed significant differences only for the first trip. Detailed per-set results are provided in Table B.5 and Table C.6 of the Supplementary Material.

4. Discussion

This work advances the current state of the art in automatic catch identification for tuna purse seiners by generating a reliable, ad hoc ground truth dataset, testing more effective segmentation and classification methods, and validating results against independent ground truth datasets. Comparative analysis identified YOLOv9 + SAM2 as the most effective segmentation approach, while a hierarchical binary classification strategy en-

hanced species discrimination. Tracking individual fish along the conveyor belt further stabilized predictions by averaging outputs across frames.

Data quality has been critical to achieving these results, as expert-labeled datasets underpin every stage of the pipeline. However, species identification from EM imagery remains challenging. Key diagnostic traits are not always visible [32], and even expert judgements can vary widely. The issue of observer bias and reliability in visual identification has also been extensively studied in citizen science contexts. Platforms like *iNaturalist* rely on large-scale, non-expert labeling, and several studies have shown that accuracy strongly depends on participants’ familiarity with the species [46, 47]. Although expert identifications are generally expected to be more reliable, evidence shows they can also be inconsistent. Culverhouse et al. [48] reported self-consistency as low as 67% and agreement between experts of only 43% when identifying plankton species. Our experiment revealed similar variability: inter-expert agreements of 42.9% $\pm$ 35.6% for BET and 57.1% $\pm$ 35.6% for YFT were found. Dropping to only one individual of each species being identified by all nine experts as the same species. We concluded that the agreements were not high enough to rely on those identifications as our ground truth. A posterior analysis by Culverhouse [49], Austen et al. [50] analyzed human factors that may play a role in species identification. Both conclude that expert accuracy in identification is often negatively affected by boredom and fatigue, regardless of their experience. Given the repetitive nature of their job, these two conditions can be very prevalent among workers whose job is to review EM

footage. This does not imply experts cannot distinguish these species, but rather that EM imagery alone is insufficient for reliable identification. In our study, onboard experts could handle the fish directly and confirm diagnostic traits, making our ground truth—and therefore our models—more robust. Although our pipeline’s performance is not perfect, its transparent and unbiased design increases confidence for both authorities and industry users.

The rapid shift in AI research towards large-scale foundation models [51] also affects fisheries monitoring. We evaluated DINOv2 and SAM2, following the approach of González-Marfil et al. [35], in which DINOv2 is used to calculate similarity maps in the image and SAM2 to generate segmentation masks. While this zero-shot method showed promise—segmenting multiple tunas from a few user clicks as seen in Figure 4—it was computationally unfeasible, taking up to ten minutes per frame. Moreover, such approaches still require validated ground truth. As foundation models and hardware continue to improve, real-time applications of this kind may become viable. AI-driven monitoring pipelines are increasingly common across marine research, from aquaculture optimization to ecosystem surveys. Public datasets such as Fishnet [52] and FDWE [21] have accelerated these developments. Córdova et al. [22], van Essen et al. [53], Xu et al. [54] are some prominent works in that field where even better validation scores than ours are obtained, but given that different species and environments are processed, a direct one-to-one comparison is not possible. Our results, based on a custom-made dataset, achieved a validation mAP of 0.66 ± 0.03 and a recall of 0.85 ± 0.03 , an improvement over our previous iteration, segmenting nearly all visible tunas, whilst acknowledging that no model can segment individuals that are not visible in the lower layers of the conveyor belt. To achieve a higher proportion of visible and segmented fish, improvements will depend less on the model itself and more on the way fish are handled and loaded into the wells of these vessels, that is, on fish handling practices. Classification accuracy also improved substantially, 20% for SKJ and 12% for YFT. On the other hand, validation accuracy dropped by 6% for BET. It should be noted that the total number of individuals for both training and validation has increased significantly in this study. Specif-

ically, BET has a 32-fold increase (from 122 to 3,592 individuals), which makes the results much more robust. Nonetheless, standard deep learning metrics should not be the sole evaluation criteria; as Eickholt et al. [29] emphasized, performance must be assessed in relation to management objectives and real-world usability. Our tests, conducted under near-operational conditions, provide robust validation despite the limited number of sampled fish.

Results differed between the two fishing trips used for testing. Both followed identical procedures, but the first used a rolling shutter camera and the second a global shutter model. The latter avoids motion-induced distortions caused by fast-moving fish, resulting in a dramatic reduction in MAE—from 12.3% to 4.5%—. Statistical analyses confirmed significant differences only for the rolling shutter data. Thus, global shutter cameras clearly enhance accuracy and should be preferred in future deployments.

Although these results represent a significant step forward, several aspects could further improve the pipeline’s operational deployment. One promising line is integrating depth (RGB-D) information. Our 3D cameras provide depth maps aligned with RGB images, enabling direct estimation of fish size and even weight through volumetric reconstruction [55]. Similar approaches have improved semantic segmentation [56] and could enhance species-level classification, as morphometric traits differ among species [57]. Incorporating environmental variables that can be used to forecast fishing grounds [58] could also help distinguish visually similar species like BET and YFT, which occupy different habitats [59].

While this study focuses on tropical tuna purse seiners, the pipeline is adaptable to other fisheries. For instance, minimal fine-tuning could make it suitable for albacore tuna (*Thunnus Alalunga*) fisheries in the Cantabrian Sea, where catches are homogeneous and individually retrieved. In contrast, more heterogeneous fisheries, such as demersal bottom trawling, would require major adaptations due to their diverse species composition and debris-filled imagery.

Finally, we plan to generate additional AFOs to improve model robustness, particularly for minority species like BET and bycatch. As fewer individuals are included in these AFOs, error rates on less common species may be overestimated. Although our current seven validated operations

demonstrate that the approach works, a even bigger testing effort is essential to warranty that a system like that could be applied to the whole fishery. Nevertheless, the present results clearly show that we are on the right track.

5. Conclusions

National administrations and Regional Fisheries Management Organizations (RFMOs) are increasingly adopting standards for the implementation of EM systems. In the specific case of tuna RFMOs, all have recently adopted these standards, which has encouraged a growing number of fishing companies to choose this option to comply with monitoring requirements. This is leading to a substantial increase in the volume of data that needs to be analyzed. In the specific context addressed here, there is currently no effective method for accurately estimating species-specific catch data from the tropical purse seine fleet. Traditional approaches, such as onboard observer sampling or visual interpretation of EM-derived imagery, have yielded imprecise estimates. This novel methodology has the potential to fill that gap, providing more reliable data that are critical for effective fisheries management, including stock assessment and regulatory enforcement. Although further work is needed, the progress made can be summarized as follows:

1. Given the similarity between both species, experts disagree on identifying bigeye tuna (BET, *Thunnus Obesus*) and yellowfin tuna (YFT, *Thunnus Albacares*) when only EM footage is available. The observed agreement among experts was $42.9\% \pm 35.6\%$ for BET and $57.1\% \pm 35.6\%$ for YFT, which we considered insufficient for ground-truth annotation. This variability highlights the need to account for observer bias when training and validating AI systems for species classification.
2. SAM2 can segment the catches of this fishery without requiring fine-tuning. When prompted with YOLOv9's bounding boxes, the combined approach achieved near-complete segmentation of all fish visible to the camera.
3. While standard multiclass and hierarchical classification models yielded similar validation metrics, hierarchical classification

demonstrated superior performance in real fishing operations. This finding supports the idea that the evaluation of AI systems should extend beyond statistical validation and involve real-world performance evaluation.

4. Camera characteristics strongly influence accuracy. In tuna purse seiners, the rapid motion of fish on the conveyor belt causes motion blur when using rolling-shutter sensors, leading to a MAE of 12.3% between our predictions and the actual species composition. Switching to global-shutter sensors reduced this error to 4.5%.
5. If a vessel's EM system is equipped with a camera with a global shutter sensor and captures images in a top-down perspective, the pipeline developed in this work can accurately estimate the composition of the catches with a MAE of $4.2\% \pm 4.9\%$ for target species (BET, SKJ, and YFT) and $5.4\% \pm 2.8\%$ for the bycatch species.

ACKNOWLEDGEMENTS

Xabier Lekunberri has received a PhD grant from the IKERTALENT Programme of the Department of Food, Rural Development, Agriculture, and Fisheries of the Basque Government. Open Access funding provided by the University of the Basque Country (EHU). We want to thank the skippers and crew members of the vessels that facilitated our data collection, the observers who accompanied us during both trips (Arnaitz Mugerza and Gorka Ocio), and the observers who worked tirelessly to annotate the images used to train the models (Arnaitz Mugerza, Lluís Horrach, Xuban Zumeta, and Iñaki Oyarzabal). The authors would also like to thank all the experts who generously participated in the classification test of yellowfin tuna (YFT) and bigeye tuna (BET) based on EM images. This paper is contribution no. XXX from AZTI, Marine Research, Basque Research and Technology Alliance (BRTA). The authors used DeepL Write for language improvement during manuscript preparation; all scientific content and interpretations are the authors' own. This work is partially supported by grant GIU23/022 funded by the University of the Basque Country (UPV/EHU), and grant PID2021-126701OB-I00, funded by the Ministerio de Ciencia, Innovación y Universidades, AEI, MCIN/AEI/10.13039/501100011033, and by "ERDF A way of making Europe". This work has been funded by the European Union through the Horizon Europe program projects EVERY-FISH (Grant agreement number 101059892) and OptiFish (Grant Agreement number 101136674) as well as BioBoost+ project within the Biodiversa+ European Biodiversity Partnership program. Views and opinions expressed are however those of the authors only and not necessarily reflect those of the European Union or Fundación Biodiversidad.

Bibliography

1. FAO. The State of World Fisheries and Aquaculture 2024 – Blue Transformation in action. June 2024. doi:[10.4060/cd0683en](https://doi.org/10.4060/cd0683en).
2. Ana Justel-Rubio. Snapshot of the Large-Scale Tropical Tuna Purse Seine Fishing Fleets as of June 2024. *International Seafood Sustainability Foundation*, June 2024.
3. Shane P. Griffiths, Valerie Allain, Simon D. Hoyle, Tim A. Lawson, and Simon J. Nicol. Just a FAD? Ecosystem impacts of tuna purse-seine fishing associated with fish aggregating devices in the western Pacific Warm Pool Province. *Fisheries Oceanography*, 28(1):94–112, January 2019. ISSN 1054-6006, 1365-2419. doi:[10.1111/fog.12389](https://doi.org/10.1111/fog.12389).
4. A. Duparc, Pascal Cauquil, Mathieu Depetris, Patrice Dewals, Daniel Gaertner, A. Hervé, Julien Lebranchu, Francis Marsac, P. Bach, and Pascal Bach. Assessment of accuracy in processing purse seine tropical tuna catches with the T3 methodology using French fleet data. page 1, 2018.
5. Antoine Duparc, V. Aragno, Mathieu Depetris, Laurent Floch, Pascal Cauquil, Julien Lebranchu, Daniel Gaertner, Francis Marsac, Pascal Bach, Mathieu Depetris, Laurent Floch, Pascal Cauquil, Julien Lebranchu, Daniel Gaertner, Francis

- Marsac, and Pascal Bach. Assessment Of The Species Composition Of Major Tropical Tunas In Purse Seine Catches : A New Modelling Approach For The Tropical Tuna Treatment Processing (Case Of The French Fleet In Atlantic Ocean). *ICCAT Recueil de Documents Scientifiques / Collective Volume of Scientific Paper*, 76(6):951–982, January 2020. ISSN 1021-5212. Publisher: ICCAT.
6. Howard McElderry. At-Sea Observing Using Video-Based Electronic Monitoring. January 2008.
 7. Rod Fujita, Christopher Cusack, Rachel Karasik, and Helen Takade-Heumacher. Designing and Implementing Electronic Monitoring Systems for Fisheries: A Supplement to the Catch Share Design Manual. Technical report, Environmental Defense Fund, San Francisco, 2018.
 8. Aloysius T. M. Van Helmond, Lars O. Mortensen, Kristian S. Plet-Hansen, Clara Ulrich, Coby L. Needle, Daniel Oesterwind, Lotte Kindt-Larsen, Thomas Catchpole, Stephen Mangi, Christopher Zimmermann, Hans Jakob Olesen, Nick Bailey, Heidrikur Bergsson, Jørgen Dalskov, Jon Elson, Malo Hosken, Lisa Peterson, Howard McElderry, Jon Ruiz, Johanna P. Pierre, Claude Dykstra, and Jan Jaap Poos. Electronic monitoring in fisheries: Lessons from global experiences and future opportunities. *Fish and Fisheries*, 21(1):162–189, 2020. ISSN 1467-2979. doi:[10.1111/faf.12425](https://doi.org/10.1111/faf.12425).
 9. K Briand, A Bonnieux, W Le Dantec, S Le Couls, P Bach, A Maufroy, A Relot-Stirremann, and P Sabarros. Comparing electronic monitoring system with observer data for estimating non-target species and discards on french tropical tuna purse seine vessels. (Col Vol Sci Pap ICCAT 74):3813–3831, 2018.
 10. Karine Briand, Alexandra Maufroy, Philippe S. Sabarros, Gwenaëlle Wain, Antoine Bonnieux, Sarah Le Couls, Romain Godefroy, Patrick Moëlo, Tiphaine Bettali, Michel Goujon, and Julien Lebranchu. The feasibility and challenges of collecting Electronic Monitoring System (EMS) data on French and associated purse seiners in relation to IOTC minimum standards. 2023.
 11. Eric Gilman, Gonzalo Legorburu, Andrew Fedoruk, Craig Heberer, Mark Zimring, and Amos Barkai. Increasing the functionalities and accuracy of fisheries electronic monitoring systems. *Aquatic Conservation: Marine and Freshwater Ecosystems*, 29(6):901–926, 2019. ISSN 1099-0755. doi:[10.1002/aqc.3086](https://doi.org/10.1002/aqc.3086).
 12. J P Monteagudo, G Legorburu, A Justel-Rubio, and V Restrepo. Preliminary study about the suitability of an electronic monitoring system to record scientific and other information from the tropical tuna purse seine fishery. page 20, 2014.
 13. H. Murua, M. Herrera, J. Morón, F. Abascal, G. Legoburu, M. Roman, G. Moreno, M. Hosken, and V. Restrepo. Comparing Electronic Monitoring and human observer collected fishery data in the tropical tuna purse seine operating in the Pacific Ocean. *IATTC SAC-11 INF-G*, 2020.
 14. Hilario Murua, Jon Ruiz, Ana Justel-Rubio, and Victor Restrepo. Minimum Standards for Electronic Monitoring Systems in Tropical Tuna Purse Seine and Longline Fisheries. Technical report, International Seafood Sustainability Foundation, Washington, D.C., USA, September 2022.
 15. J. Ruiz, A. Batty, P. Chavance, H. McElderry, V. Restrepo, P. Sharples, J. Santos, and A. Urtizberea. Electronic monitoring trials on in the tropical tuna purse-seine fishery. *ICES Journal of Marine Science*, 72(4):1201–1213, May 2015. ISSN 1054-3139. doi:[10.1093/icesjms/fsu224](https://doi.org/10.1093/icesjms/fsu224).
 16. J. Ruiz, I. Krug, O. Gonzalez, and G. Hammann. E-EYE PLUS: Electronic monitoring trial on the tropical tuna purse seine fleet. *ICCAT, Madrid, Spain*, 3, 2016.
 17. Jon Ruiz, Antoine Duparc, Francisco Abascal, Pascal Bach, Jose Carlos Baez, Pedro Pascual, Daniel Gaertner, Francis Marsac, and Josu Santiago. Comments on the assessment of catch by species in the tropical purse seine fishery. 2021.
 18. Zhenze Duanmu, Yajin Zhang, Feng Li, Yao Yao, Hu Shi, Fanbin Kong, and Yang Jiao. Fast freezing, slow tempering: The key to tuna (*Thunnus albacares*) quality retention. *International Journal of Food Science and Technology*, 59(5):3031–3044, May 2024. ISSN 0950-5423. doi:[10.1111/ijfs.17034](https://doi.org/10.1111/ijfs.17034).
 19. Miquel Palmer, Amaya Álvarez Ellacuría, Vicenç Moltó, and Ignacio A. Catalán. Automatic, operational, high-resolution monitoring of fish length and catch numbers from landings using deep learning. *Fisheries Research*, 246:106166, February 2022. ISSN 0165-7836. doi:[10.1016/j.fishres.2021.106166](https://doi.org/10.1016/j.fishres.2021.106166).
 20. Kaiming He, Georgia Gkioxari, Piotr Dollár, and Ross Girshick. Mask R-CNN. *arXiv:1703.06870 [cs.CV]*, January 2018. doi:[10.48550/arxiv.1703.06870](https://doi.org/10.48550/arxiv.1703.06870).
 21. Maria Sokolova, Manuel Cordova, Henk Nap, Aloysius van Helmond, Michiel Mans, Arjan Vroegop, Angelo Mencarelli, and Gert Kootstra. An integrated end-to-end deep neural network for automated detection of discarded fish species and their weight estimation. *ICES Journal of Marine Science*, 80(7):1911–1922, September 2023. ISSN 1054-3139. doi:[10.1093/icesjms/fsad118](https://doi.org/10.1093/icesjms/fsad118).
 22. Manuel Córdova, Maria Sokolova, Aloysius van Helmond, Angelo Mencarelli, and Gert Kootstra. Multi-stage image-based approach for fish detection and weight estimation. *Biosystems Engineering*, 257:104239, September 2025. ISSN 1537-5110. doi:[10.1016/j.biosystemseng.2025.104239](https://doi.org/10.1016/j.biosystemseng.2025.104239).
 23. Rishi Bommasani, Drew A. Hudson, Ehsan Adeli, Russ Altman, Simran Arora, Sydney von Arx, Michael S. Bernstein, Jeannette Bohg, Antoine Bosselut, Emma Brunskill, Erik Brynjolfsson, Shyamal Buch, Dallas Card, Rodrigo Castellon, Niladri Chatterji, Annie Chen, Kathleen Creel, Jared Quincy Davis, Dora Demszky, Chris Donahue, Moussa Doumbouya, Esin Durmus, Stefano Ermon, John Etchemendy, Kavin Ethayarajh, Li Fei-Fei, Chelsea Finn, Trevor Gale, Lauren Gillespie, Karan Goel, Noah Goodman, Shelby Grossman, Neel Guha, Tatsunori Hashimoto, Peter Henderson, John Hewitt, Daniel E. Ho, Jenny Hong, Kyle Hsu, Jing Huang, Thomas Icard, Saahil Jain, Dan Jurafsky, Pratyusha Kalluri, Siddharth Karamcheti, Geoff Keeling, Fereshte Khani, Omar Khattab, Pang Wei Koh, Mark Krass, Ranjay Krishna, Rohit Kudipudi, Ananya Kumar, Faisal Ladhak, Mina Lee, Tony Lee, Jure Leskovec, Isabelle Levent, Xiang Lisa Li, Xuechen Li, Tengyu Ma, Ali Malik, Christopher D. Manning, Suvir Mirchandani, Eric Mitchell, Zanele Munyikwa, Suraj Nair, Avanika Narayan, Deepak Narayanan, Ben Newman, Allen Nie, Juan Carlos Niebles, Hamed Nilforoshan, Julian Nyarko, Gireay Ogut, Laurel Orr, Isabel Papadimitriou, Joon Sung Park, Chris Piech, Eva Portelance, Christopher Potts, Aditi Raghunathan, Rob Reich, Hongyu Ren,
 - Frieda Rong, Yusuf Roohani, Camilo Ruiz, Jack Ryan, Christopher Ré, Dorsa Sadigh, Shiori Sagawa, Keshav Santhanam, Andy Shih, Krishnan Srinivasan, Alex Tamkin, Rohan Taori, Armin W. Thomas, Florian Tramèr, Rose E. Wang, William Wang, Bohan Wu, Jiajun Wu, Yuhuai Wu, Sang Michael Xie, Michihiro Yasunaga, Jiaxuan You, Matei Zaharia, Michael Zhang, Tianyi Zhang, Xikun Zhang, Yuhui Zhang, Lucia Zheng, Kaitlyn Zhou, and Percy Liang. On the Opportunities and Risks of Foundation Models, July 2022. arXiv:2108.07258 [cs].
 24. Zhang Jiaxing and Tang Hao. SAM2 for Image and Video Segmentation: A Comprehensive Survey, March 2025. arXiv:2503.12781 [cs].
 25. Maxime Oquab, Timothée Darcet, Théo Moutakanni, Huy Vo, Marc Szafraniec, Vasil Khalidov, Pierre Fernandez, Daniel Haziza, Francisco Massa, Alaaeldin El-Nouby, Mahmoud Assran, Nicolas Ballas, Wojciech Galuba, Russell Howes, Po-Yao Huang, Shang-Wen Li, Ishan Misra, Michael Rabbat, Vasu Sharma, Gabriel Synnaeve, Hu Xu, Hervé Jegou, Julien Mairal, Patrick Labatut, Armand Joulin, and Piotr Bojanowski. DINOv2: Learning Robust Visual Features without Supervision, February 2024. arXiv:2304.07193 [cs].
 26. Nikhila Ravi, Valentin Gabeur, Yuan-Ting Hu, Ronghang Hu, Chaitanya Ryali, Tengyu Ma, Haiman Khedr, Roman Rädle, Chloe Rolland, Laura Gustafson, Eric Mintun, Junting Pan, Kalyan Vasudev Alwala, Nicolas Carion, Chao-Yuan Wu, Ross Girshick, Piotr Dollár, and Christoph Feichtenhofer. SAM 2: Segment Anything in Images and Videos, October 2024. arXiv:2408.00714 [cs].
 27. Isaak Kavasidis, Amelia Sorrenti, Orazio Tomarchio, Daniela Giordano, Marco Milazzo, Gabriele Turco, Carlo Cattano, and Concetto Spampinato. Training-free fish species population monitoring in unconstrained underwater videos. *Procedia Computer Science*, 257:809–816, January 2025. ISSN 1877-0509. doi:[10.1016/j.procs.2025.03.104](https://doi.org/10.1016/j.procs.2025.03.104).
 28. Shijie Lian and Hua Li. Evaluation of Segment Anything Model 2: The Role of SAM2 in the Underwater Environment, August 2024. arXiv:2408.02924 [cs].
 29. Jesse Eickholt, Jonathan Gregory, and Kavya Vemuri. Advancing Fisheries Research and Management with Computer Vision: A Survey of Recent Developments and Pending Challenges. *Fishes*, 10(2):74, February 2025. ISSN 2410-3888. doi:[10.3390/fishes10020074](https://doi.org/10.3390/fishes10020074). Number: 2 Publisher: Multidisciplinary Digital Publishing Institute.
 30. Xabier Lekunberri, Jon Ruiz, Iñaki Quincoces, Fadi Dornaika, Ignacio Arganda-Carreras, and Jose A. Fernandes. Identification and measurement of tropical tuna species in purse seiner catches using computer vision and deep learning. *Ecological Informatics*, 67:101495, March 2022. ISSN 15749541. doi:[10.1016/j.ecoinf.2021.101495](https://doi.org/10.1016/j.ecoinf.2021.101495).
 31. CVAT. ai Corporation. Computer Vision Annotation Tool (CVAT), July 2024.
 32. Tammy Horton, Leigh Marsh, Brian J. Bett, Andrew R. Gates, Daniel O. B. Jones, Noëlie M. A. Benoist, Simone Pfeifer, Erik Simon-Lledó, Jennifer M. Durden, Leen Vandepitte, and Ward Appeltans. Recommendations for the Standardisation of Open Taxonomic Nomenclature for Image-Based Identifications. *Frontiers in Marine Science*, 8, February 2021. ISSN 2296-7745. doi:[10.3389/fmars.2021.620702](https://doi.org/10.3389/fmars.2021.620702). Publisher: Frontiers.
 33. FAO. ASFIS List of Species for Fishery Statistics Purposes. In: Fisheries and Aquaculture., 2024.
 34. Chien-Yao Wang, I-Hau Yeh, and Hong-Yuan Mark Liao. YOLOv9: Learning What You Want to Learn Using Programmable Gradient Information, February 2024. arXiv:2402.13616 [cs].
 35. Aitor González-Marfil, Estibaliz Gómez-de Mariscal, and Ignacio Arganda-Carreras. DINOsim: Zero-Shot Object Detection and Semantic Segmentation on Electron Microscopy Images, March 2025. Pages: 2025.03.09.642092 Section: New Results.
 36. Hao Zhang, Feng Li, Shilong Liu, Lei Zhang, Hang Su, Jun Zhu, Lionel M. Ni, and Heung-Yeung Shum. DINO: DETR with Improved DeNoising Anchor Boxes for End-to-End Object Detection, July 2022. arXiv:2203.03605 [cs].
 37. Shaoqing Ren, Kaiming He, Ross Girshick, and Jian Sun. Faster R-CNN: Towards Real-Time Object Detection with Region Proposal Networks. *IEEE Transactions on Pattern Analysis and Machine Intelligence*, 39(6):1137–1149, June 2017. ISSN 0162-8828, 2160-9292. doi:[10.1109/TPAMI.2016.2577031](https://doi.org/10.1109/TPAMI.2016.2577031).
 38. Alex Bewley, Zongyuan Ge, Lionel Ott, Fabio Ramos, and Ben Uppcroft. Simple Online and Realtime Tracking, July 2017. arXiv:1602.00763.
 39. Nicolai Wojke, Alex Bewley, and Dietrich Paulus. Simple Online and Realtime Tracking with a Deep Association Metric, March 2017. arXiv:1703.07402.
 40. Yifu Zhang, Peize Sun, Yi Jiang, Dongdong Yu, Fucheng Weng, Zehuan Yuan, Ping Luo, Wenyu Liu, and Xinggang Wang. ByteTrack: Multi-object Tracking by Associating Every Detection Box. In Shai Avidan, Gabriel Brostow, Moustapha Cissé, Giovanni Maria Farinella, and Tal Hassner, editors, *Computer Vision – ECCV 2022*, pages 1–21, Cham, 2022. Springer Nature Switzerland. ISBN 978-3-031-20047-2. doi:[10.1007/978-3-031-20047-2_1](https://doi.org/10.1007/978-3-031-20047-2_1).
 41. Mahmoud Abouelyazid. Comparative Evaluation of SORT, DeepSORT, and ByteTrack for Multiple Object Tracking in Highway Videos. *International Journal of Sustainable Infrastructure for Cities and Societies*, 8(11):42–52, November 2023. Number: 11.
 42. Jumabek Alikhanov and Hakil Kim. Online Action Detection in Surveillance Scenarios: A Comprehensive Review and Comparative Study of State-of-the-Art Multi-Object Tracking Methods. *IEEE Access*, 11:68079–68092, 2023. ISSN 2169-3536. doi:[10.1109/ACCESS.2023.3292539](https://doi.org/10.1109/ACCESS.2023.3292539). Conference Name: IEEE Access.
 43. Diogo Ferreira and Meysam Basiri. Dynamic Target Tracking and Following with UAVs Using Multi-Target Information: Leveraging YOLOv8 and MOT Algorithms. *DRONES*, 8(9):488, September 2024. ISSN 2504-446X. doi:[10.3390/drones8090488](https://doi.org/10.3390/drones8090488). Num Pages: 23 Place: Basel Publisher: MDPI Web of Science ID: WOS:001324028000001.
 44. Iljia Radosavovic, Raj Prateek Kosaraju, Ross Girshick, Kaiming He, and Piotr Dollár. Designing Network Design Spaces, March 2020. arXiv:2003.13678 [cs].
 45. Dena F. Muftaba and Nihar R. Mahapatra. Hierarchical Deep Learning Models

- for Identification of Fish Species. *2022 International Conference on Computational Science and Computational Intelligence (CSCI)*, pages 1558–1562, December 2022. doi:[10.1109/CSCI58124.2022.00278](https://doi.org/10.1109/CSCI58124.2022.00278). Conference Name: 2022 International Conference on Computational Science and Computational Intelligence (CSCI) ISBN: 9798350320282 Place: Las Vegas, NV, USA Publisher: IEEE.
46. Sarah J. Ackland, David M. Richardson, and Tamara B. Robinson. A Method for Conveying Confidence in iNaturalist Observations: A Case Study Using Non-Native Marine Species. *Ecology and Evolution*, 14(10): e70376, 2024. ISSN 2045-7758. doi:[10.1002/ece3.70376](https://doi.org/10.1002/ece3.70376). _eprint: <https://onlinelibrary.wiley.com/doi/pdf/10.1002/ece3.70376>.
 47. Thomas Mesaglio, Kelly A. Shepherd, Juliet A. Wege, Russell L. Barrett, Hervé Sauquet, and Will K. Cornwell. Expert identification blitz: A rapid high value approach for assessing and improving iNaturalist identification accuracy and data precision and confidence. *PLANTS, PEOPLE, PLANET*, 7 (5):1469–1484, 2025. ISSN 2572-2611. doi:[10.1002/ppp3.70005](https://doi.org/10.1002/ppp3.70005). _eprint: <https://nph.onlinelibrary.wiley.com/doi/pdf/10.1002/ppp3.70005>.
 48. Phil F. Culverhouse, Robert Williams, Beatriz Reguera, Vincent Herry, and Sonsoles González-Gil. Do experts make mistakes? A comparison of human and machine identification of dinoflagellates. *Marine Ecology Progress Series*, 247: 17–25, February 2003. ISSN 0171-8630, 1616-1599. doi:[10.3354/meps247017](https://doi.org/10.3354/meps247017).
 49. Phil F. Culverhouse. Human and machine factors in algae monitoring performance. *Ecological Informatics*, 2(4):361–366, December 2007. ISSN 1574-9541. doi:[10.1016/j.ecoinf.2007.07.001](https://doi.org/10.1016/j.ecoinf.2007.07.001).
 50. Gail E. Austen, Markus Bindemann, Richard A. Griffiths, and David L. Roberts. Species identification by conservation practitioners using online images: accuracy and agreement between experts. *PeerJ*, 6:e4157, January 2018. ISSN 2167-8359. doi:[10.7717/peerj.4157](https://doi.org/10.7717/peerj.4157). Publisher: PeerJ Inc.
 51. Johannes Schneider, Christian Meske, and Pauline Kuss. Foundation Models. *Business & Information Systems Engineering*, 66(2):221–231, April 2024. ISSN 1867-0202. doi:[10.1007/s12599-024-00851-0](https://doi.org/10.1007/s12599-024-00851-0).
 52. Justin Kay and Matt Merrifield. The Fishnet Open Images Database: A Dataset for Fish Detection and Fine-Grained Categorization in Fisheries, June 2021. arXiv:2106.09178 [cs].
 53. Rick van Essen, Angelo Mencarelli, Aloysius van Helmond, Linh Nguyen, Jürgen Batsleer, Jan-Jaap Poos, and Gert Kootstra. Automatic discard registration in cluttered environments using deep learning and object tracking: class imbalance, occlusion, and a comparison to human review. *ICES Journal of Marine Science*, 78(10):3834–3846, December 2021. ISSN 1054-3139. doi:[10.1093/icesjms/fsab233](https://doi.org/10.1093/icesjms/fsab233).
 54. Wenjie Xu, Hui Fang, Shengchi Yu, Shenglong Yang, Haodong Yang, Yujia Xie, and Yang Dai. RSNC-YOLO: A Deep-Learning-Based Method for Automatic Fine-Grained Tuna Recognition in Complex Environments. *Applied Sciences*, 14 (22):10732, January 2024. ISSN 2076-3417. doi:[10.3390/app142210732](https://doi.org/10.3390/app142210732). Publisher: Multidisciplinary Digital Publishing Institute.
 55. Alexey Ruchay, Vitaly Kober, Konstantin Dorofeev, Vladimir Kolpakov, Alexey Gladkov, and Hao Guo. Live Weight Prediction of Cattle Based on Deep Regression of RGB-D Images. *Agriculture*, 12(11):1794, November 2022. ISSN 2077-0472. doi:[10.3390/agriculture12111794](https://doi.org/10.3390/agriculture12111794). Publisher: Multidisciplinary Digital Publishing Institute.
 56. Changshuo Wang, Chen Wang, Weijun Li, and Haining Wang. A brief survey on RGB-D semantic segmentation using deep learning. *Displays*, 70:102080, December 2021. ISSN 0141-9382. doi:[10.1016/j.displa.2021.102080](https://doi.org/10.1016/j.displa.2021.102080).
 57. ICCAT. *ICCAT Manual*. International Commission for the Conservation of Atlantic Tuna, 2006. ISBN 978-92-990055-0-7.
 58. Nerea Goikoetxea, Izaro Goienetxea, Jose A. Fernandes-Salvador, Nicolas Goñi, Igor Granado, Iñaki Quincoces, Leire Ibaibarriaga, Jon Ruiz, Hilario Murua, and Ainhoa Caballero. Machine-learning aiding sustainable Indian Ocean tuna purse seine fishery. *Ecological Informatics*, 81:102577, July 2024. ISSN 1574-9541. doi:[10.1016/j.ecoinf.2024.102577](https://doi.org/10.1016/j.ecoinf.2024.102577).
 59. Haritz Arrizabalaga, Florence Dufour, Laurence Kell, Gorka Merino, Leire Ibaibarriaga, Guillem Chust, Xabier Irigoien, Josu Santiago, Hilario Murua, Igaratza Fraile, Marina Chifflet, Nerea Goikoetxea, Yolanda Sagarmínaga, Olivier Aumont, Laurent Bopp, Miguel Herrera, Jean Marc Fromentin, and Sylvain Bonhommeau. Global habitat preferences of commercially valuable tuna. *Deep Sea Research Part II: Topical Studies in Oceanography*, 113:102–112, March 2015. ISSN 0967-0645. doi:[10.1016/j.dsr2.2014.07.001](https://doi.org/10.1016/j.dsr2.2014.07.001).

Supplementary Material A. Reproducibility

Table A.4. Summary of computational resources, data augmentation parameters, hyperparameter configurations, and training durations used for segmentation and classification models.

(a). Computing infrastructure used to train the models.

Component	Model	Quantity
CPU	Intel® Xeon® Silver 4509Y	2
GPU	NVIDIA® H100 NVL (94GB)	1
RAM	—	768GiB

(b). Data augmentation used to train segmentation and classification models. $V_{B,C,S}$ represent the image's original brightness, contrast, and saturation values, respectively.

Parameter	Possible values	Probability
Brightness	$[\max(0, 1 - V_B), 1 + V_B]$	1
Contrast	$[\max(0, 1 - V_C), 1 + V_C]$	1
Saturation	$[\max(0, 1 - V_S), 1 + V_S]$	1
Horizontal flip	—	0.5
Vertical flip	—	0.5

(c). Values assigned to the hyperparameters of the trained **segmentation models**.

Hyperparameter	Search space	Best assignment
Epochs	Mask R-CNN [30, 40, 50, 60]	40
	YOLOv9 [500, 600, 700, 800, 900]	700
Batch size	[2, 4]	2
Train / validation / test	[0.7 / 0.2 / 0.1]	0.7 / 0.2 / 0.1
Optimizer	[SGD]	SGD
Learning Rate	[0.01]	0.01
Momentum	[0.95]	0.95
Weight Decay	[0.0005]	0.0005
LR Scheduler	[Exponential(gamma=0.96)]	Exponential(gamma=0.96)

(d). Values assigned to the hyperparameters of the trained **classification models**.

Hyperparameter	Search space	Best assignment
Epochs	[40, 50, 60]	40
Batch size	[2, 4]	4
Train / validation / test	[0.7 / 0.2 / 0.1]	0.7 / 0.2 / 0.1
Optimizer	[SGD]	SGD
Learning Rate	[0.01]	0.01
Momentum	[0.95]	0.95
Weight Decay	[0.0005]	0.0005
LR Scheduler	[StepLR(step_size=7, gamma=0.1)]	StepLR(step_size=7, gamma=0.1)

(e). Time required to train all the models. This includes all folds and repetitions of the repeated cross-validation. In the hierarchical model, the three levels are added together.

Model	Train duration	
Segmentation	Mask R-CNN	21 days
	YOLOv9	25 days
Classification	Standard	10 days
	Hierarchical	21 days

Supplementary Material B. Mask R-CNN segmentation in AFOs

Table B.5. Catch composition estimations for 21 test AFOs using the **Mask R-CNN** approach for the segmentation and different classifications. We show each AFO's total number of fish, the percentage segmented by our models (evaluation metric for the segmentation, bigger is better), and the catch composition by species percentage (evaluation metric for the classification, similar to GT is better). GT stands for ground truth, S for standard classification, and H for hierarchical.

ID _{AFO}	No. individuals		BET			SKJ			YFT			NO_TARGET		
	GT	Model	GT	S	H	GT	S	H	GT	S	H	GT	S	H
1_01	327	48.3%	11.9%	8.2%	6.3%	46.2%	58.9%	27.8%	39.1%	27.8%	59.5%	2.8%	5.1%	6.3%
1_02	232	45.7%	1.3%	4.7%	8.5%	84.5%	64.2%	38.7%	14.2%	30.2%	46.2%	0.0%	0.9%	6.6%
1_03	204	46.6%	1.0%	2.1%	8.4%	90.7%	68.4%	26.3%	8.3%	26.3%	58.9%	0.0%	3.2%	6.3%
1_04	239	49.4%	12.6%	6.8%	5.1%	48.1%	39.0%	12.7%	39.3%	47.5%	80.5%	0.0%	6.8%	1.7%
1_05	284	42.6%	12.0%	9.1%	9.9%	59.5%	50.4%	25.6%	28.5%	38.8%	62.0%	0.0%	1.7%	2.5%
1_06	217	44.7%	0.0%	3.1%	8.2%	87.1%	54.6%	25.8%	12.9%	37.1%	62.9%	0.0%	5.2%	3.1%
1_07	185	49.7%	1.1%	2.2%	1.1%	85.4%	47.8%	28.3%	13.0%	43.5%	67.4%	0.5%	6.5%	3.3%
1_08	478	33.3%	2.5%	1.9%	2.5%	83.9%	49.1%	30.2%	9.4%	40.9%	65.4%	4.2%	8.2%	1.9%
1_09	418	40.4%	4.1%	4.1%	4.1%	65.1%	49.1%	29.0%	27.5%	43.8%	64.5%	3.3%	3.0%	2.4%
1_10	390	43.8%	1.3%	0.0%	4.7%	87.9%	62.0%	35.1%	10.5%	35.7%	56.1%	0.3%	2.3%	4.1%
1_11	455	61.5%	0.4%	1.1%	3.6%	69.0%	57.9%	31.8%	26.2%	35.7%	61.8%	4.4%	5.4%	2.9%
1_12	499	41.1%	4.6%	1.0%	2.4%	75.8%	55.6%	27.3%	15.0%	35.6%	66.3%	4.6%	7.8%	3.9%
1_13	216	53.7%	1.4%	2.6%	0.9%	50.5%	41.4%	23.3%	48.1%	53.4%	74.1%	0.0%	2.65	1.7%
1_14	309	48.5%	4.2%	2.7%	2.7%	42.1%	30.0%	12.7%	53.7%	62.7%	82.7%	0.0%	4.7%	2.0%
2_01	223	79.8%	2.7%	2.2%	5.6%	66.4%	89.3%	80.9%	30.9%	8.4%	7.3%	0.0%	0.0%	6.2%
2_02	312	57.7%	7.1%	2.8%	3.9%	85.3%	93.3%	83.3%	6.4%	2.8%	3.9%	1.3%	1.1%	8.9%
2_03	344	61.9%	1.2%	1.4%	3.3%	87.8%	96.2%	83.6%	10.8%	2.3%	5.2%	0.3%	0.0%	8.0%
2_04	227	63.9%	4.4%	2.8%	4.1%	85.9%	93.8%	87.6%	5.7%	1.4%	4.1%	4.0%	2.1%	4.1%
2_05	460	46.5%	0.2%	0.0%	0.5%	77.6%	95.3%	87.9%	3.9%	4.2%	5.6%	18.3%	0.5%	6.1%
2_06	232	79.7%	0.0%	0.0%	0.0%	92.7%	97.8%	90.3%	0.0%	2.2%	4.9%	7.3%	0.0%	4.9%
2_07	402	55.0%	1.0%	0.0%	0.0%	75.4%	95.9%	83.7%	8.7%	3.6%	10.9%	14.9%	0.5%	5.4%

Supplementary Material C. YOLOv9 + SAM2 segmentation in AFOs

Table C.6. Catch composition estimations for 21 test AFOs using the **YOLOv9 + SAM2** approach for the segmentation and different classifications. We show each AFO's total number of fish, the percentage segmented by our models (evaluation metric for the segmentation, bigger is better), and the catch composition by species percentage (evaluation metric for the classification, similar to GT is better). GT stands for ground truth, S for standard classification, and H for hierarchical.

ID _{AFO}	No. individuals		BET			SKJ			YFT			NO_TARGET		
	GT	Model	GT	S	H	GT	S	H	GT	S	H	GT	S	H
1_01	327	67.0%	11.9%	6.8%	6.4%	46.2%	49.8%	23.3%	39.1%	39.3%	63.9%	2.8%	4.1%	6.4%
1_02	232	60.8%	1.3%	5.7%	7.1%	84.5%	53.2%	31.2%	14.2%	39.7%	56.7%	0.0%	1.4%	5.0%
1_03	204	51.0%	1.0%	2.9%	7.7%	90.7%	61.5%	22.1%	8.3%	32.7%	63.5%	0.0%	2.9%	6.7%
1_04	239	60.7%	12.6%	4.8%	4.1%	48.1%	34.5%	11.7%	39.3%	52.4%	82.1%	0.0%	8.3%	2.1%
1_05	284	58.8%	12.0%	6.0%	6.6%	59.5%	50.3%	21.6%	28.5%	42.5%	69.5%	0.0%	1.2%	2.4%
1_06	217	59.9%	0.0%	4.6%	7.7%	87.1%	49.2%	20.0%	12.9%	43.1%	67.7%	0.0%	3.1%	4.6%
1_07	185	69.2%	1.1%	0.8%	3.9%	85.4%	42.2%	25.0%	13.0%	53.1%	68.8%	0.5%	3.9%	2.3%
1_08	478	46.7%	2.5%	0.9%	1.3%	83.9%	52.5%	26.5%	9.4%	40.8%	70.0%	4.2%	5.8%	2.2%
1_09	418	52.2%	4.1%	5.0%	5.0%	65.1%	41.3%	20.6%	27.5%	50.9%	72.0%	3.3%	2.8%	2.3%
1_10	390	57.7%	1.3%	0.9%	3.1%	87.9%	55.1%	32.9%	10.5%	41.3%	60.4%	0.3%	2.7%	3.6%
1_11	455	85.5%	0.4%	1.3%	4.9%	69.0%	53.0%	28.3%	26.2%	38.0%	64.5%	4.4%	7.7%	2.3%
1_12	499	47.9%	4.6%	0.4%	2.9%	75.8%	51.9%	22.6%	15.0%	40.6%	70.7%	4.6%	7.1%	3.8%
1_13	216	66.7%	1.4%	1.4%	2.1%	50.5%	38.9%	21.5%	48.1%	54.9%	75.0%	0.0%	4.9%	1.4%
1_14	309	60.8%	4.2%	2.7%	3.2%	42.1%	28.7%	12.8%	53.7%	66.0%	82.4%	0.0%	2.7%	1.6%
2_01	223	102.7%	2.7%	3.1%	7.0%	66.4%	89.5%	81.2%	30.9%	7.4%	6.6%	0.0%	0.0%	5.2%
2_02	312	84.3%	7.1%	2.7%	1.1%	85.3%	93.9%	84.0%	6.4%	2.3%	6.8%	1.3%	1.1%	8.0%
2_03	344	82.3%	1.2%	1.1%	3.5%	87.8%	97.2%	84.1%	10.8%	1.4%	3.2%	0.3%	0.4%	9.2%
2_04	227	81.9%	4.4%	2.2%	3.2%	85.9%	93.5%	86.0%	5.7%	1.6%	4.3%	4.0%	2.7%	6.5%
2_05	460	67.6%	0.2%	0.0%	0.3%	77.6%	97.1%	83.9%	3.9%	2.6%	6.1%	18.3%	0.3%	9.6%
2_06	232	105.2%	0.0%	0.0%	0.0%	92.7%	98.8%	91.0%	0.0%	1.2%	2.5%	7.3%	0.0%	6.6%
2_07	402	69.9%	1.0%	0.4%	0.0%	75.4%	95.7%	82.2%	8.7%	3.2%	7.8%	14.9%	0.7%	10.0%

Supplementary Material D. Expert assigned tuna species

Table D.7. Species assigned to each of the fishes by the nine experts.

ID _F	Species assigned by each expert									Total	
	1	2	3	4	5	6	7	8	9	YFT	BET
001	YFT	YFT	YFT	YFT	YFT	YFT	-	-	YFT	7	0
002	-	-	YFT	YFT	YFT	YFT	-	-	YFT	5	0
003	-	YFT	YFT	YFT	-	YFT	YFT	YFT	YFT	7	0
004	-	-	YFT	-	BET	-	-	-	-	1	1
005	-	-	YFT	-	YFT	-	-	-	-	2	0
006	-	-	-	-	YFT	YFT	-	-	-	2	0
007	-	-	-	-	-	-	-	-	BET	0	1
008	-	-	-	-	YFT	-	-	-	-	1	0
009	-	-	-	-	YFT	YFT	-	-	-	2	0
010	-	-	-	-	YFT	-	BET	-	-	1	1
011	-	-	-	BET	-	-	-	-	BET	0	2
012	-	-	-	-	-	-	-	-	YFT	1	0
013	-	-	YFT	-	-	BET	-	-	-	1	1
014	-	-	-	YFT	YFT	YFT	-	-	-	3	0
015	YFT	-	BET	BET	-	BET	-	-	YFT	2	3
016	-	YFT	YFT	YFT	YFT	YFT	-	YFT	YFT	7	0
017	-	-	-	-	YFT	-	-	-	-	1	0
018	-	-	-	-	YFT	-	-	-	-	1	0
019	-	-	-	BET	-	BET	-	BET	BET	0	4
020	YFT	YFT	YFT	YFT	YFT	YFT	-	-	YFT	7	0
021	-	-	-	BET	YFT	BET	-	-	-	1	2
022	-	-	YFT	BET	YFT	-	BET	-	YFT	3	2
023	-	-	-	BET	YFT	YFT	-	-	-	2	1
024	BET	-	-	BET	BET	BET	BET	-	YFT	1	5
025	-	YFT	YFT	YFT	BET	BET	-	YFT	BET	4	3
026	-	-	-	-	YFT	-	-	-	-	1	0
027	-	-	-	-	YFT	-	-	-	-	1	0
028	-	-	-	YFT	YFT	BET	-	-	-	2	1
029	-	YFT	-	YFT	YFT	-	-	-	-	3	0
030	-	-	YFT	YFT	YFT	YFT	-	-	BET	4	1
031	-	-	-	-	YFT	-	-	-	BET	1	1
032	-	-	-	BET	YFT	BET	-	BET	YFT	2	3
033	BET	-	-	BET	-	BET	BET	-	BET	0	5
034	-	YFT	YFT	-	YFT	-	-	-	-	3	0
035	YFT	YFT	YFT	-	YFT	YFT	YFT	-	-	6	0
036	-	-	-	-	-	-	-	-	BET	0	1
037	-	-	-	YFT	YFT	YFT	-	-	BET	3	1
038	-	-	-	-	-	YFT	-	BET	BET	1	2
039	-	-	-	BET	-	-	-	-	-	0	1
040	-	-	-	-	YFT	-	-	-	-	1	0
041	-	-	-	-	YFT	YFT	-	-	-	2	0
042	-	-	-	-	BET	-	-	-	BET	0	2
043	-	-	-	YFT	BET	-	-	-	YFT	2	1

ID _F	Species assigned by each expert									Total	
	1	2	3	4	5	6	7	8	9	YFT	BET
044	-	-	-	YFT	-	YFT	-	-	-	2	0
045	YFT	YFT	YFT	YFT	YFT	YFT	YFT	YFT	YFT	9	0
046	-	-	-	-	YFT	YFT	-	-	BET	2	1
047	-	-	-	-	YFT	-	-	-	YFT	2	0
048	-	-	-	-	-	YFT	-	-	-	1	0
049	-	-	-	-	YFT	-	-	-	YFT	2	0
050	-	-	-	-	YFT	-	-	-	-	1	0
051	-	-	-	YFT	-	YFT	-	-	YFT	3	0
052	-	-	YFT	YFT	YFT	-	-	-	-	3	0
053	-	-	YFT	YFT	YFT	YFT	-	-	BET	4	1
054	-	-	-	-	-	YFT	-	-	BET	1	1
055	-	YFT	-	BET	-	-	-	-	-	1	1
056	-	-	-	YFT	-	YFT	-	-	BET	2	1
057	-	-	YFT	YFT	YFT	YFT	YFT	-	YFT	6	0
058	-	-	-	YFT	-	-	-	-	-	1	0
059	-	-	-	-	BET	-	-	-	YFT	1	1
060	-	-	-	-	-	-	-	YFT	-	1	0
061	-	-	-	-	-	-	-	-	BET	0	1
062	-	-	-	-	-	-	-	-	BET	0	1
063	-	-	-	BET	-	-	-	-	-	0	1
064	-	-	-	-	YFT	-	-	BET	BET	1	2
065	BET	-	-	BET	YFT	YFT	-	-	YFT	3	2
066	-	-	YFT	YFT	YFT	YFT	YFT	-	YFT	6	0
067	-	-	-	-	-	-	-	-	YFT	1	0
068	-	-	-	-	YFT	YFT	-	-	BET	2	1
069	-	-	-	-	YFT	-	-	-	-	1	0
070	-	-	-	BET	BET	YFT	-	-	YFT	2	2
071	-	-	-	YFT	YFT	-	-	-	YFT	3	0
072	BET	-	BET	BET	BET	YFT	BET	BET	YFT	2	6
073	-	-	-	YFT	YFT	-	-	-	-	2	0
074	-	-	-	-	YFT	-	-	-	-	1	0
075	-	YFT	YFT	YFT	YFT	YFT	-	-	YFT	6	0
076	BET	-	-	YFT	YFT	YFT	-	BET	YFT	4	2
077	-	-	-	YFT	-	-	-	-	YFT	2	0
078	-	-	YFT	YFT	-	-	BET	-	YFT	3	1
079	-	-	-	-	YFT	YFT	-	-	-	2	0
080	-	-	-	-	YFT	-	-	-	YFT	2	0
081	-	-	-	-	YFT	-	-	-	YFT	2	0
082	-	-	-	-	-	-	-	-	YFT	1	0
083	-	-	-	-	-	-	-	-	BET	0	1
084	-	-	-	YFT	BET	-	-	-	BET	1	2
085	BET	-	-	YFT	YFT	-	-	-	BET	2	2
086	BET	-	-	YFT	-	YFT	-	BET	BET	2	3
087	BET	-	BET	YFT	-	YFT	BET	-	BET	2	4
088	-	-	-	YFT	YFT	-	-	-	BET	2	1
089	-	-	-	-	BET	-	-	-	-	0	1

ID _F	Species assigned by each expert									Total	
	1	2	3	4	5	6	7	8	9	YFT	BET
090	-	-	-	-	YFT	YFT	-	-	-	2	0
091	-	-	-	BET	-	-	-	-	YFT	1	1
092	BET	-	-	-	-	-	-	-	YFT	1	1
093	-	-	-	YFT	-	YFT	-	BET	BET	2	2
094	-	-	-	YFT	-	-	-	-	BET	1	1
095	-	-	BET	-	BET	-	BET	-	BET	0	4
096	-	YFT	-	YFT	-	-	-	-	BET	2	1
097	-	-	-	-	-	-	-	-	BET	0	1
098	-	-	-	-	BET	-	-	-	-	0	1
099	-	-	-	-	YFT	-	-	-	-	1	0
100	-	-	-	-	BET	-	-	-	-	0	1
101	-	-	-	-	YFT	-	-	-	-	1	0
102	-	-	-	-	BET	-	-	-	-	0	1
103	-	-	-	-	YFT	-	-	-	-	1	0
104	BET	BET	BET	BET	BET	BET	BET	BET	BET	0	9
105	-	-	BET	BET	-	-	-	-	BET	0	3
106	-	-	-	-	-	-	-	-	BET	0	1
107	-	-	-	-	YFT	YFT	-	-	-	2	0
108	-	-	-	YFT	YFT	-	-	-	YFT	3	0
109	-	-	-	BET	BET	-	-	-	YFT	1	2
110	-	YFT	YFT	YFT	YFT	YFT	BET	-	YFT	6	1
111	-	-	-	YFT	-	-	-	YFT	YFT	3	0
112	-	-	-	YFT	-	-	-	-	YFT	2	0
113	-	-	-	-	YFT	YFT	-	-	YFT	3	0
114	-	-	-	-	BET	-	-	-	-	0	1
115	-	-	-	-	YFT	-	-	-	-	1	0
116	-	-	-	YFT	YFT	YFT	YFT	-	YFT	5	0
117	-	-	-	-	YFT	-	-	-	YFT	2	0
118	-	-	-	-	YFT	-	-	-	YFT	2	0
119	-	-	-	-	-	-	-	BET	-	0	1
120	-	-	-	YFT	YFT	YFT	-	-	-	3	0
121	-	-	YFT	YFT	YFT	YFT	-	-	YFT	5	0
122	-	-	-	-	-	-	-	-	BET	0	1
123	-	-	-	-	YFT	-	-	-	BET	1	1
124	-	YFT	-	-	YFT	-	-	-	-	2	0
125	-	-	-	-	YFT	-	-	-	-	1	0
126	-	-	-	-	-	-	-	-	BET	0	1
127	-	-	-	-	YFT	-	-	-	BET	1	1
128	-	-	-	YFT	YFT	-	-	-	BET	2	1
129	-	-	-	YFT	BET	-	-	-	BET	1	2
130	-	-	-	-	-	-	-	-	BET	0	1
131	-	-	-	YFT	YFT	-	-	-	BET	2	1
132	-	-	-	YFT	-	-	-	-	-	1	0
133	-	-	-	YFT	-	-	-	-	BET	1	1
134	-	-	-	YFT	YFT	-	-	-	BET	2	1
135	BET	-	BET	BET	YFT	YFT	BET	BET	BET	2	6

ID _F	Species assigned by each expert									Total	
	1	2	3	4	5	6	7	8	9	YFT	BET
136	YFT	YFT	YFT	YFT	YFT	-	-	-	BET	5	1
137	-	YFT	-	YFT	BET	-	-	-	BET	2	2
138	-	-	-	BET	BET	-	BET	BET	BET	0	5
139	-	-	-	-	YFT	-	-	-	-	1	0
140	BET	BET	-	BET	BET	-	BET	BET	YFT	1	6
141	-	-	-	-	-	-	-	-	YFT	1	0
142	-	-	-	-	BET	-	-	-	YFT	1	1
143	-	-	-	-	BET	-	-	-	YFT	1	1
144	-	-	-	YFT	-	-	-	-	BET	1	1
145	-	-	-	YFT	-	-	-	-	BET	1	1
146	-	-	-	YFT	YFT	-	YFT	YFT	BET	4	1
147	-	-	-	-	YFT	-	-	-	-	1	0
148	-	-	-	-	BET	-	-	-	-	0	1
149	YFT	YFT	YFT	-	YFT	-	YFT	YFT	YFT	7	0
150	-	-	-	-	-	-	-	-	BET	0	1
151	-	-	-	-	-	-	-	-	BET	0	1
152	BET	BET	YFT	BET	BET	-	BET	BET	BET	1	7
153	-	-	-	YFT	-	-	-	-	-	1	0
154	-	-	-	-	-	-	-	-	BET	0	1
155	-	-	-	-	BET	-	-	-	-	0	1
156	-	BET	-	YFT	YFT	-	BET	YFT	YFT	4	2
157	-	-	-	YFT	-	-	-	-	-	1	0
158	-	-	-	YFT	-	-	-	-	BET	1	1
159	BET	BET	-	BET	BET	-	-	-	YFT	1	4
160	-	-	-	YFT	YFT	-	YFT	YFT	YFT	5	0
161	-	YFT	-	YFT	BET	-	-	-	YFT	3	1
162	-	-	-	YFT	YFT	-	-	-	BET	2	1
163	-	-	-	-	-	-	-	-	YFT	1	0
164	-	-	-	-	-	-	-	-	YFT	1	0
165	-	-	-	YFT	BET	-	BET	-	BET	1	3
166	-	-	YFT	YFT	YFT	YFT	-	BET	BET	4	2
167	YFT	YFT	YFT	YFT	YFT	YFT	BET	YFT	BET	7	2
168	-	-	-	YFT	-	-	-	-	-	1	0
169	-	-	-	-	-	-	-	-	YFT	1	0
170	-	-	-	YFT	YFT	-	YFT	YFT	YFT	5	0
171	YFT	-	-	-	YFT	-	-	-	YFT	3	0
172	BET	-	-	BET	-	-	BET	BET	BET	0	5
173	-	-	-	YFT	-	-	-	-	-	1	0
174	-	-	-	-	YFT	-	-	-	YFT	2	0
175	-	-	-	-	YFT	-	-	-	-	1	0
176	-	-	YFT	YFT	BET	YFT	YFT	BET	YFT	5	2
177	BET	-	-	YFT	YFT	-	-	-	-	2	1
178	-	-	-	BET	-	-	-	-	-	0	1

ID _F	Species assigned by each expert									Total	
	1	2	3	4	5	6	7	8	9	YFT	BET
179	-	-	-	YFT	-	-	-	-	-	1	0
180	-	-	-	BET	-	-	-	-	BET	0	2
181	-	BET	-	BET	-	-	-	-	BET	0	3
182	-	-	BET	BET	BET	-	-	-	BET	0	4
183	YFT	BET	BET	BET	BET	YFT	YFT	BET	BET	3	6
184	-	YFT	-	BET	YFT	-	-	-	BET	2	2
185	-	-	-	BET	-	-	-	-	BET	0	2
186	-	-	-	-	YFT	-	-	-	YFT	2	0
187	-	-	-	BET	YFT	YFT	-	-	YFT	3	1
188	-	-	-	BET	BET	-	-	-	-	0	2
189	YFT	BET	BET	BET	BET	-	YFT	BET	BET	2	6
190	-	-	YFT	YFT	BET	-	-	-	YFT	3	1
191	-	-	-	YFT	-	-	-	-	YFT	2	0
192	-	-	-	BET	-	-	-	-	BET	0	2
193	-	-	-	YFT	YFT	-	-	-	-	2	0
194	-	-	-	-	-	-	-	-	YFT	1	0
195	-	-	-	-	-	-	-	-	YFT	1	0
196	-	-	-	-	-	-	-	-	BET	0	1
197	-	-	BET	YFT	BET	-	-	-	BET	1	3
198	-	-	-	BET	BET	-	BET	-	BET	0	4
199	BET	BET	-	BET	BET	-	-	BET	BET	0	6
200	BET	YFT	-	BET	BET	-	-	-	BET	1	4
201	-	-	-	YFT	BET	-	-	-	BET	1	2
202	-	-	-	YFT	YFT	-	-	-	YFT	3	0
203	-	-	BET	YFT	BET	-	-	-	BET	1	3
204	-	-	-	YFT	-	-	-	-	BET	1	1
205	-	-	-	YFT	BET	-	-	-	BET	1	2
206	-	-	-	YFT	-	-	-	-	YFT	2	0
207	-	-	-	-	-	-	-	-	BET	0	1
208	-	-	-	-	-	-	-	-	BET	0	1
209	-	-	-	-	BET	-	-	-	BET	0	2
210	-	-	YFT	YFT	BET	-	BET	BET	YFT	3	3
211	-	-	-	YFT	BET	-	-	-	YFT	2	1
212	-	-	-	YFT	-	-	-	-	-	1	0
213	-	-	-	-	YFT	-	-	-	-	1	0
214	-	-	-	-	YFT	-	-	-	-	1	0
215	-	-	-	YFT	BET	-	-	-	BET	1	2
216	-	BET	BET	BET	BET	-	BET	-	BET	0	6
217	-	-	-	BET	BET	-	-	BET	YFT	1	3
218	-	-	-	BET	BET	-	-	-	BET	0	3
219	-	-	-	-	-	-	-	-	BET	0	1
220	-	-	-	-	-	-	-	-	YFT	1	0
221	-	-	-	-	BET	-	-	-	-	0	1
222	-	YFT	BET	YFT	YFT	YFT	-	YFT	YFT	6	1
223	BET	-	-	YFT	BET	-	-	-	YFT	2	2
224	-	-	-	BET	-	-	-	-	-	0	1

ID _F	Species assigned by each expert									Total	
	1	2	3	4	5	6	7	8	9	YFT	BET
225	-	-	-	YFT	YFT	-	-	-	YFT	3	0
226	-	-	BET	YFT	BET	-	BET	-	YFT	2	3
227	-	BET	-	BET	-	-	-	-	BET	0	3
228	-	-	-	YFT	YFT	-	-	-	-	2	0
229	-	-	-	-	-	-	-	-	YFT	1	0
230	-	-	-	-	-	-	-	-	YFT	1	0
231	-	-	-	-	YFT	-	-	-	-	1	0
232	-	-	-	-	-	-	BET	-	YFT	1	1
233	-	-	-	YFT	-	-	-	-	YFT	2	0
234	-	-	-	YFT	-	-	-	-	YFT	2	0
235	-	-	-	BET	YFT	-	-	-	YFT	2	1
236	-	-	-	YFT	BET	-	-	BET	YFT	2	2
237	BET	-	-	-	-	-	-	-	-	0	1
238	-	-	-	-	-	-	-	-	BET	0	1
239	YFT	YFT	YFT	YFT	YFT	-	YFT	YFT	YFT	8	0
240	-	-	-	YFT	-	-	-	-	YFT	2	0
241	-	-	-	-	-	-	-	-	YFT	1	0
242	YFT	-	YFT	YFT	YFT	-	YFT	YFT	YFT	7	0
243	-	-	-	YFT	-	-	-	-	-	1	0
244	-	BET	-	-	BET	-	-	-	YFT	1	2
245	-	-	-	YFT	-	-	-	-	YFT	2	0
246	-	-	-	YFT	YFT	-	-	-	YFT	3	0
247	-	-	-	-	YFT	-	-	-	-	1	0
248	-	-	-	-	-	-	-	-	YFT	1	0
249	-	-	-	-	-	-	-	-	BET	0	1
250	-	-	YFT	YFT	YFT	-	-	-	YFT	4	0
251	-	-	-	-	-	-	-	-	YFT	1	0
252	-	-	-	-	-	-	BET	BET	BET	0	3
253	YFT	-	-	YFT	-	-	-	-	YFT	3	0
254	-	YFT	-	-	YFT	-	-	-	YFT	3	0
255	-	-	-	YFT	YFT	-	-	-	YFT	3	0
256	YFT	YFT	YFT	-	YFT	YFT	YFT	YFT	YFT	8	0
257	-	YFT	YFT	-	YFT	-	-	-	YFT	4	0