SC/68C/PH/04

Sub-committees/working group name: PH

Artificial intelligence for right whales: Multi feature photo identification

Khan C, Holmberg J, Hamilton P, Blount D, Parham J, Charlton C, Rowntree V, Vermeulen E, Rayment W, Dawson S, Johnston D, Groch K

Artificial Intelligence for Right Whales: Multi Feature Photo Identification

Christin Khan[1], Jason Holmberg[2], Philip Hamilton[3], Drew Blount[2], Jason Parham[2,5], Claire Charlton[6], Victoria Rowntree[7,8], Els Vermeulen[9], Will Rayment[10], Steve Dawson[10], David Johnston[10], Karina Groch[11]

[1]NOAA Northeast Fisheries Science Center, Woods Hole, Massachusetts, United States
[2]Wild Me, Portland, Oregon, United States
[3]Anderson Cabot Center for Ocean Life at the New England Aquarium, Boston, Massachusetts
[4]deepsense.ai, Warsaw, Poland
[5]Rensselaer Polytechnic Institute, Troy, NY, United States
[6]Centre for Marine Science and Technology, Curtin University, Bentley, Western Australia, Australia
[7]Department of Biology, University of Utah, Salt Lake City, Utah, United States
[8]Instituto de Conservacio´n de Ballenas, Capital Federal, Buenos Aires, Argentina
[9]Mammal Research Institute Whale Unit, University of Pretoria, Hatfield, South Africa
[10]Marine Science Department, Te Tari Putaiao Taimoana, University of Otago, New Zealand
[11]Southern Right Whale Project / Instituto Australis, Imbituba, Santa Catarina, Brazil

Summary

This report provides an update on the use of artificial intelligence (AI) for right whale photo identification presented to the International Whaling Commission (IWC) Science Committee Photo Identification sub-committee in 2020 (SC/68B/PH/03). The NOAA led collaborative project with the New England Aquarium, Wild Me, Deepsense.ai, Kaggle, and right whale researchers around the world has progressed substantially towards the overall aim to use AI for photo identification of right whales. In 2020 at IWCSC68B, we reported that the system was operational on Wild Me's Flukebook platform for matching the heads of right whales taken from an aerial perspective for North Atlantic right whales using the Kaggle-winning Deepsense algorithms. In 2021, the system was further expanded to include Southern right whales, and Flukebook successfully implemented multi-feature matching and new AI techniques. Multi-feature matching allows right whales to be matched by aerial photos of their heads (Deepsense), lateral photos of their heads (Pose Invariant Embeddings), flukes (new CurvRank v2), and peduncle scarring (HotSpotter). This capability to apply new forms of AI and match an individual North Atlantic right whale from multiple poses and marks has then been successfully cross-applied to Southern right whales. Right whale researchers have been slow to adopt these new methods, and we continue to explore ways to streamline the platform and create time saving workflows to encourage broader adoption. This AI photo-identification platform applies to species across multiple IWC sub-committees (Southern Hemisphere, Northern Hemisphere, Conservation Management Plans, and Photo-identification). Additionally this work supports wintering grounds for all Southern right whales which are a priority species for the IWC and should streamline data processing for assessing life history data.

Background

Photo identification is an important tool for estimating abundance and monitoring population trends over time. However, manually matching photographs to known individuals is time-consuming. Motivated by developments in image recognition, NOAA Fisheries hosted a data science challenge on the crowdsourcing platform Kaggle in 2015 to automate the identification of endangered North Atlantic right whales (Figure 1).

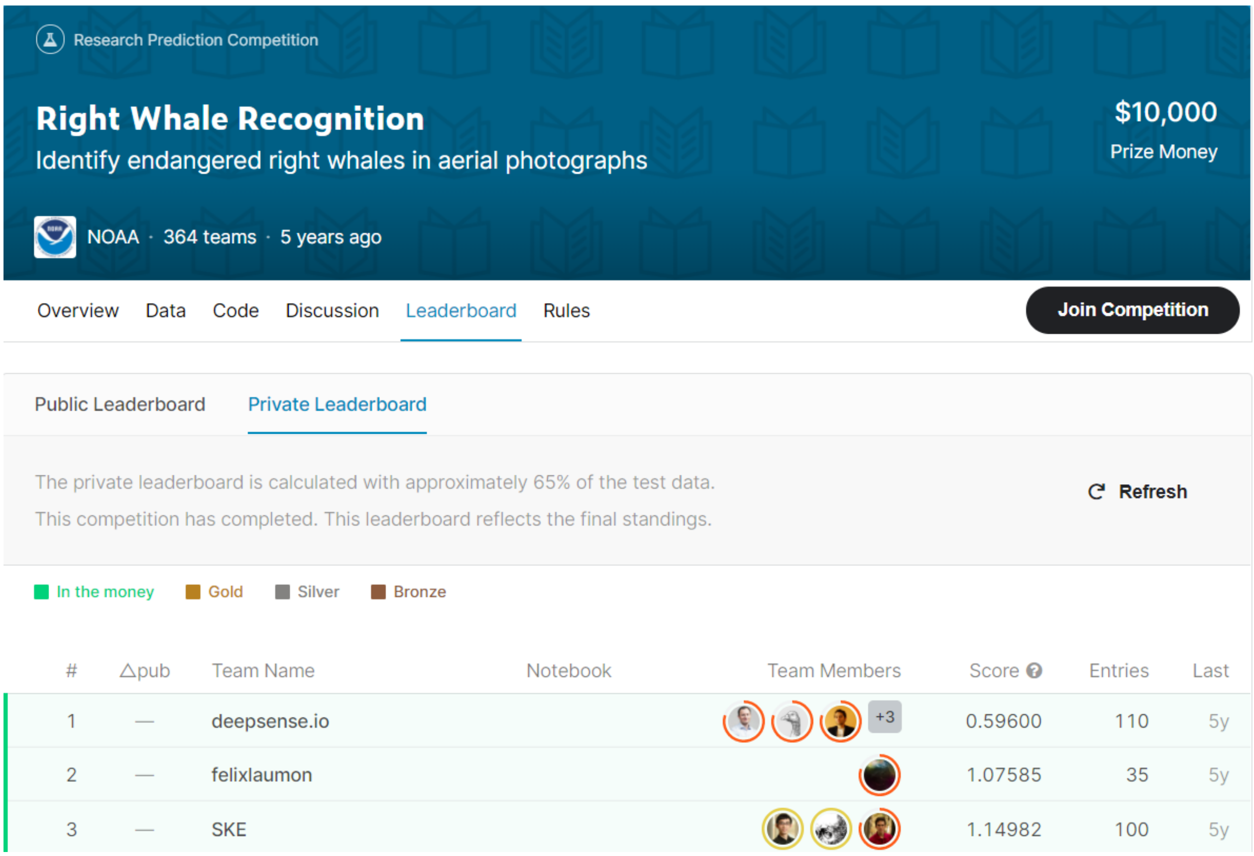

Figure 1. Leaderboard for the NOAA "Right Whale Recognition" Kaggle competition in 2015.

The winning solution by Deepsense.ai automatically identified individual whales with 87% accuracy with a series of convolutional neural networks to identify the region of interest on an image, rotate, crop, and create standardized photographs of uniform size and orientation and then identify the correct individual whale from these passport-like photographs (Figure 2). Recent advances in deep learning coupled with this fully automated workflow have yielded impressive results and have the potential to revolutionize traditional methods for the collection of data on the abundance and distribution of wild populations (Bogucki et al 2019).

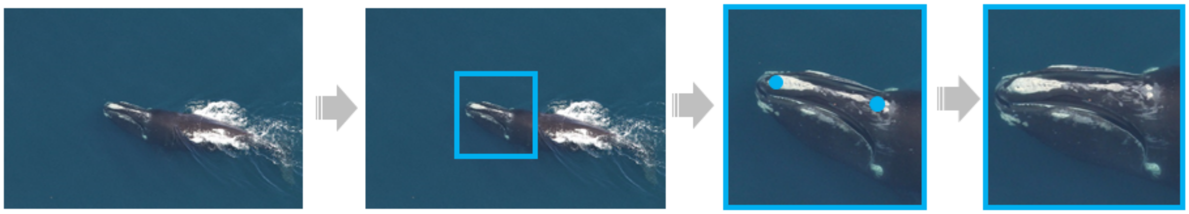

Figure 2. Convolutional neural network pipeline developed by Deepsense.ai as the winning solution in the NOAA "Right Whale Recognition" Kaggle Competition.

To make these algorithms broadly accessible to the research community, NOAA partnered with the Wild Me, a 501(c)(3) nonprofit organization that builds open software and artificial intelligence for the conservation research community. The machine learning experts and software developers at Wild Me retrained the Deepsense algorithm on an updated dataset from the North Atlantic Right Whale Consortium and hosted the system live on the Flukebook platform (Figure 3) in 2019 for the semi-automated matching of aerial photographs of right whale head callosity patterns.

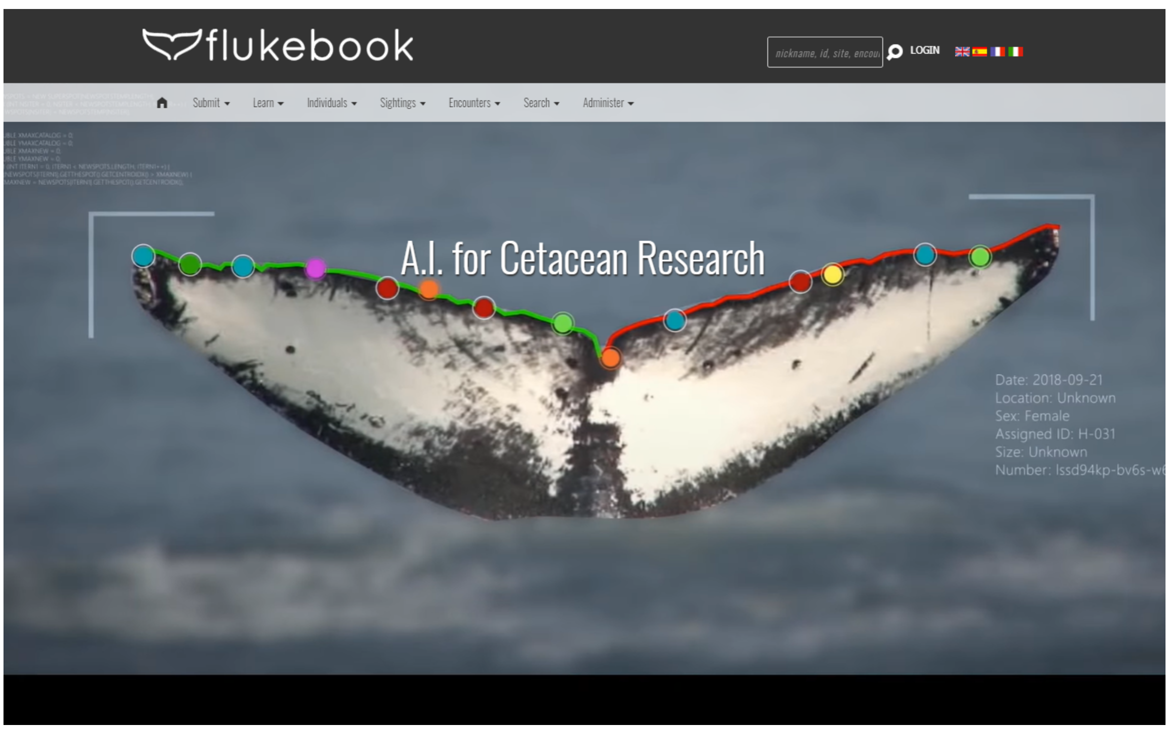

Figure 3. The landing page for the FlukeBook platform (https://www.flukebook.org) which applies computer vision algorithms and deep learning to identify and track individual cetaceans.

The Flukebook interface was integrated into the North Atlantic Right Whale Catalog - the public facing website that contains the comprehensive list of known North Atlantic right whales along with associated metadata such as their date of birth, sex, genetic lineage, and sighting history (Figure 4). Photographers can choose to match the identity of the whale in their images by scrolling through all cataloged whales, searching by matching features, or by using Flukebook's automated matching system. When the Flukebook platform is used to identify the individual whale in the photograph, the algorithm results contain the 12 most likely individuals along with a link to that whale's sighting page (Figure 5) on the North Atlantic Right Whale Catalog website to compare photographs and confirm the id.

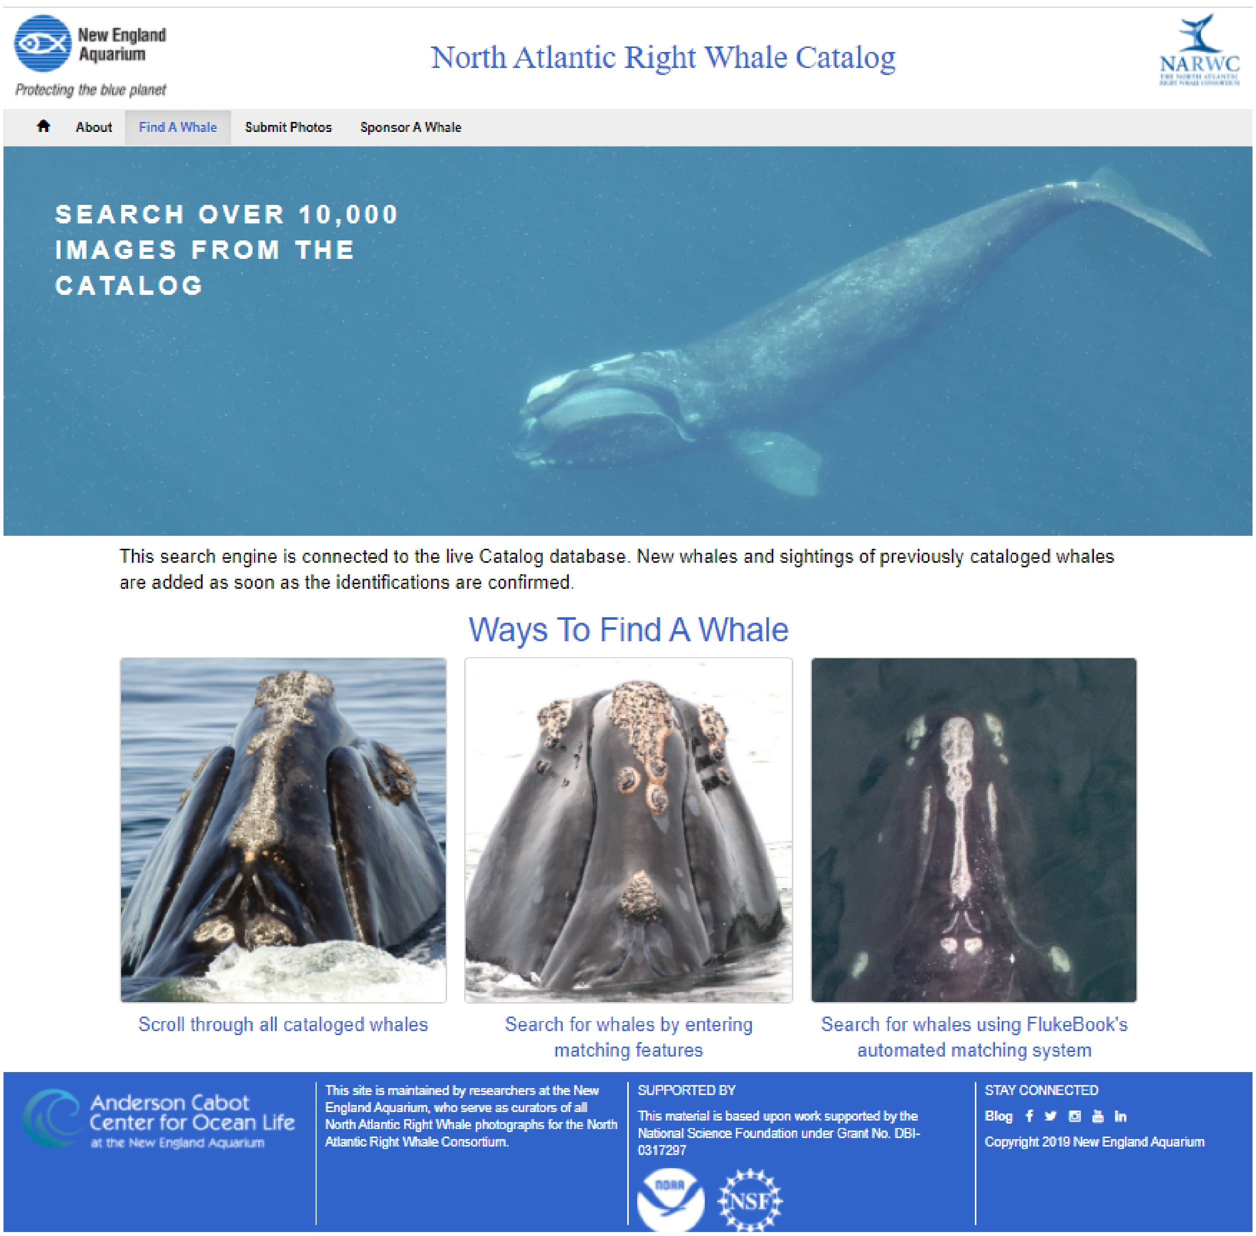

Figure 4. The North Atlantic Right Whale Catalog (http://rwcatalog.neaq.org) which includes the option to 'Search for whales using FlukeBook's automated matching system".

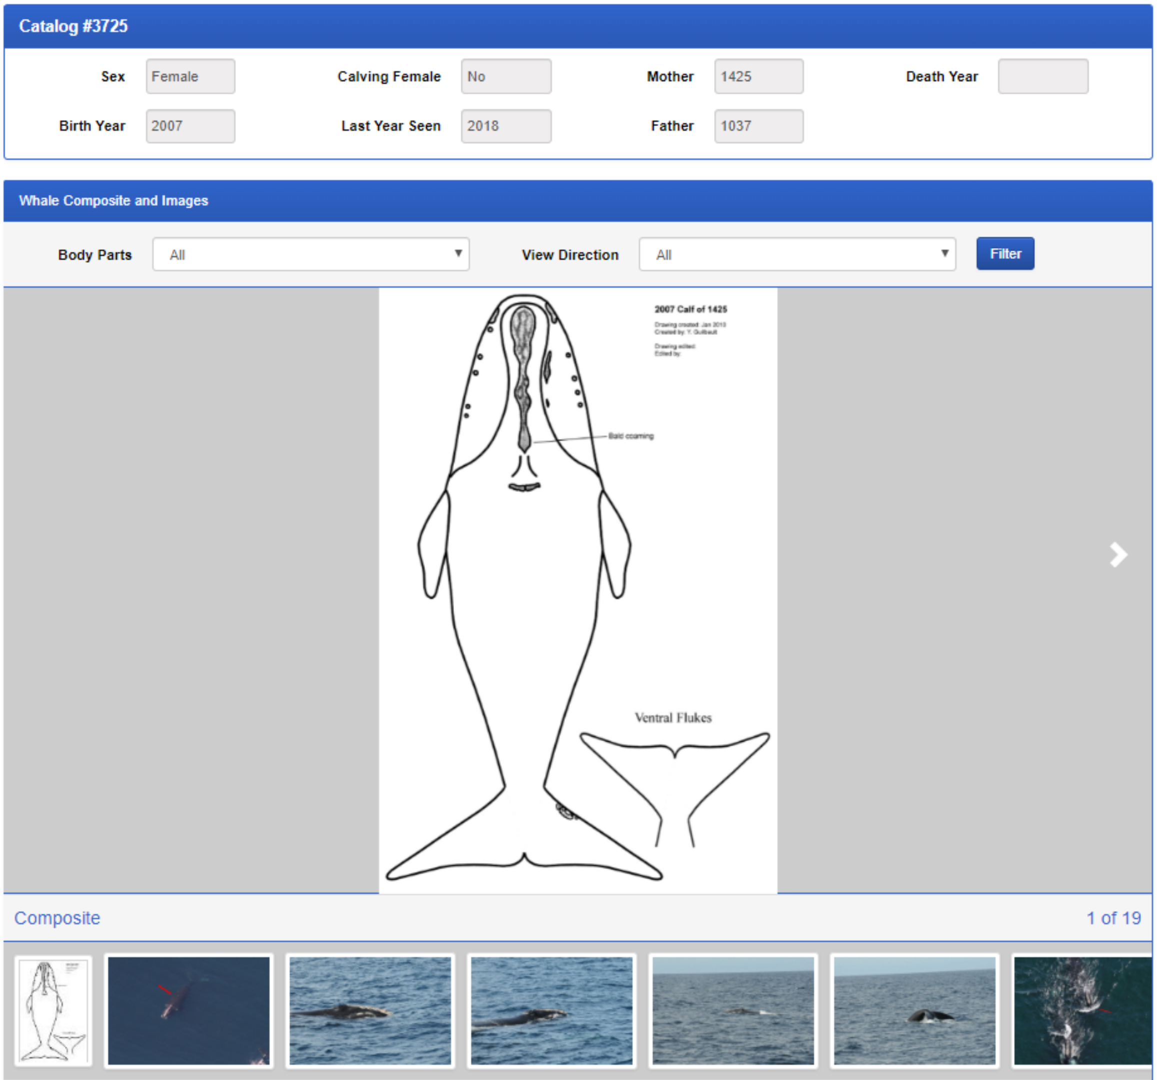

Figure 5. The North Atlantic Right Whale Catalog (http://rwcatalog.neaq.org) page for whale number 3725 including composite drawing and exemplar photos. Here we learn that whale 3725 is a female born in 2007 to the mother 1425 and father 1037.

Southern right whales

The AI pipelines and infrastructure built for North Atlantic right whales was expanded to include Southern right whales (Figure 6): 12,311 images from Australia from Curtin University; 8,461 images from South Africa from the University of Pretoria; 8,952 images from Argentina from the University of Utah; 5473 images from Brazil from Instituto Australis; and 2,913 images from New Zealand from the University of Otago. After filtering for known individuals with at least 2 sightings, the total Southern right whale dataset included 10,451 photographs. The average resight-rate of individual Southern right whales was only 4 (compared to 88 in the North Atlantic) which resulted in significantly less accurate models. Further funding and research is anticipated to make the algorithm more generalizable so that the Southern right whale model can more closely approach the North Atlantic model in accuracy. This effort to bring researchers from around the world together for collaboration and efficiency is tremendously valuable and a model for other conservation efforts.

Figure 6. A Southern right whale demonstrating morphological similarity to the North Atlantic right whale. Photo by Claire Charlton.

Multi-Feature Matching

NOAA recently funded the expansion of the semi-automated photo-identification system for right whales to allow identification via multiple features (head, peduncle, fluke) and viewpoints (overhead from airplane or lateral from vessel). The machine learning and software development team at Flukebook greatly expanded the capacity of the platform to be able to support multi-feature, multi-algorithm matching with a new image intake platform (Figure 7) that assigns annotations (viewpoint, body part) and passes them to one or more appropriate algorithms. This multi-modal, multi-feature, and multi-species machine learning pipeline ("Wildbook Image Analysis") has subsequently enabled new advancements in photo-identification for other species, including simultaneous saddle patch and dorsal matching from a single photo using PIE and CurvRank v2 in orcas.

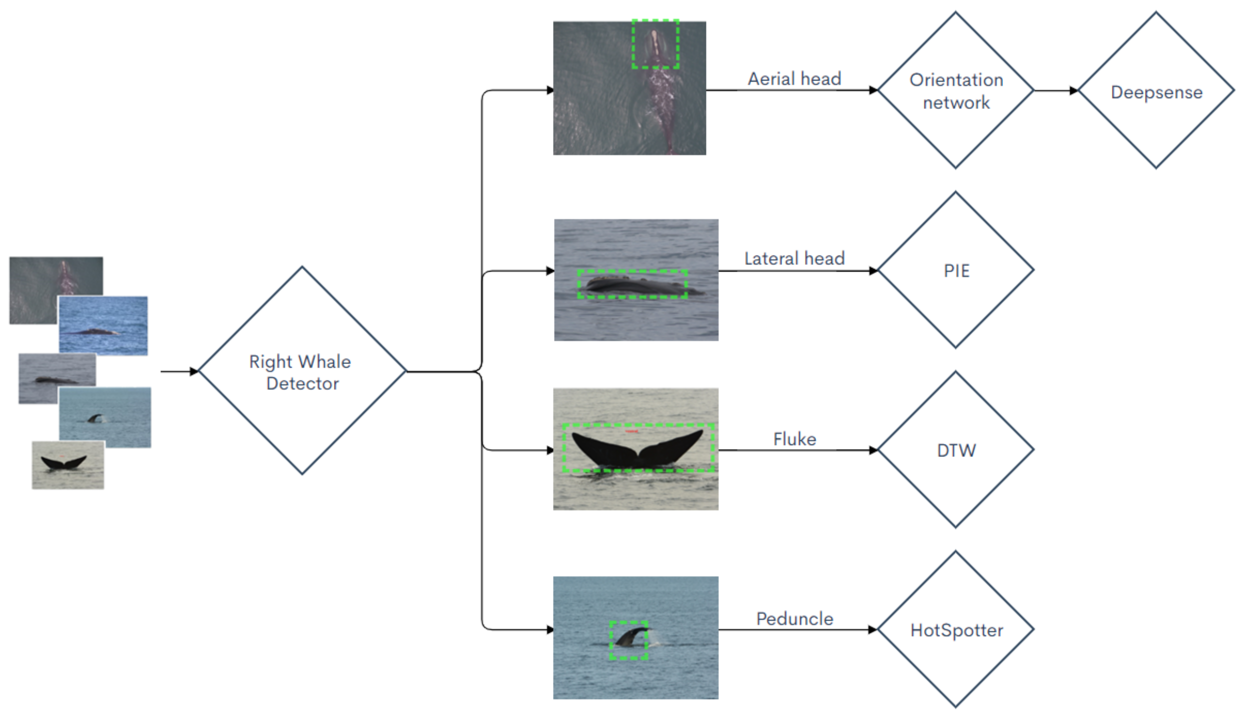

Figure 7. Pipeline for the annotation and identification of multi-feature matching of right whales in the Flukebook platform (https://www.flukebook.org/).

Aerial Photos of the Head: Deepsense

The Deepsense algorithm that won the Kaggle data science competition has since been successfully retrained for both the North Atlantic and the Southern right whale catalogs by the machine learning development team at Wild Me (Figure 8). The Deepsense algorithm is a deep learning convolutional neural network that is both fast and accurate. Neural networks are an established family of machine-learning models inspired by the human brain, able to learn and perform all sorts of tasks by adjusting a huge network of small, flexible units. Convolutional neural networks (CNN) have become the most popular approach in modern machine learning and can be trained to iteratively improve to perform a task as more training data is added. As the number of layers in the state-of-the-art convolutional neural networks increased, the term “deep learning” was coined as a phrase denoting training a neural network with many layers. Each layer within a network receives an input image, performs a transformation on the image, and outputs the results to the subsequent layer. Adding new individual whales to this algorithm requires training a new model and since it is optimized for right whales, it cannot be cross applied to other species.

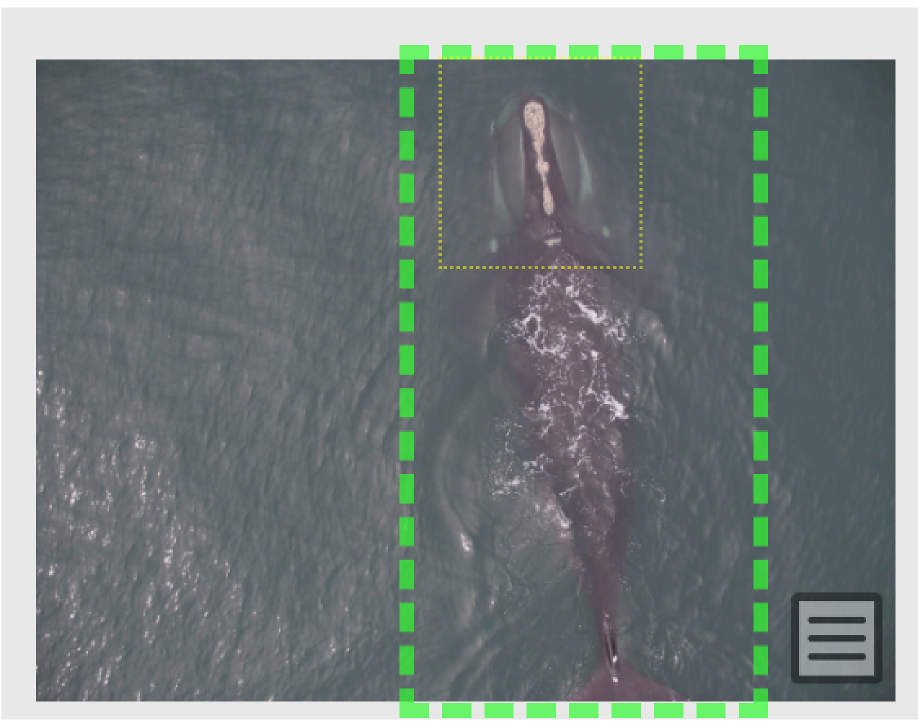

Figure 8. An annotated image of an aerial photograph of a North Atlantic right whale in preparation for matching the head (smaller yellow bounding box) with the Deepsense deep learning algorithm on the FlukeBook platform (https://www.flukebook.org/).

Lateral Photos of the Head: PIE

The Pose Invariant Embeddings (PIE) algorithm was developed by computer science doctoral student Olga Moskvyak at Queensland University of Technology (Moskvyak et al 2019; Figure 9). Unlike Deepsense, the PIE neural network is not trained to classify images into bins of individuals. Instead, its deep neural network is trained to extract embeddings from images. Give an image to PIE, and it returns a list of 256 numbers between 0 and 1. An embedding is an abstract, numerical representation of an image. PIE is trained with the simple concept that images of the same individual should produce similar embeddings, and images of different individuals should produce different embeddings; the distance between two images in embedding space (or "triplet loss") corresponds to the similarity between those images for the purpose of individual ID. An advantage of this approach is that PIE can gracefully add new individuals to its catalog without being retrained: it learns the general task of mapping images into embeddings that represent individuals, rather than the specific task of sorting images into a fixed number of IDs. PIE strikes a balance between a flexible general-purpose identifier and one that can be trained and refined on a given problem. Moskvyak's publicly available PIE algorithm was trained and tuned on lateral photos of right whale heads by Drew Blount at Wild Me, resulting in a model with a top-12 accuracy of 90% and top-5 accuracy of 81%. This difficulty reflects both the challenging nature of these photographs and that this is the first ever automatic system for matching boat-based photos of right whales. Application of PIE to North Atlantic right whales has also subsequently enabled new advancements in gray whale and orca lateral matching in Flukebook.org.

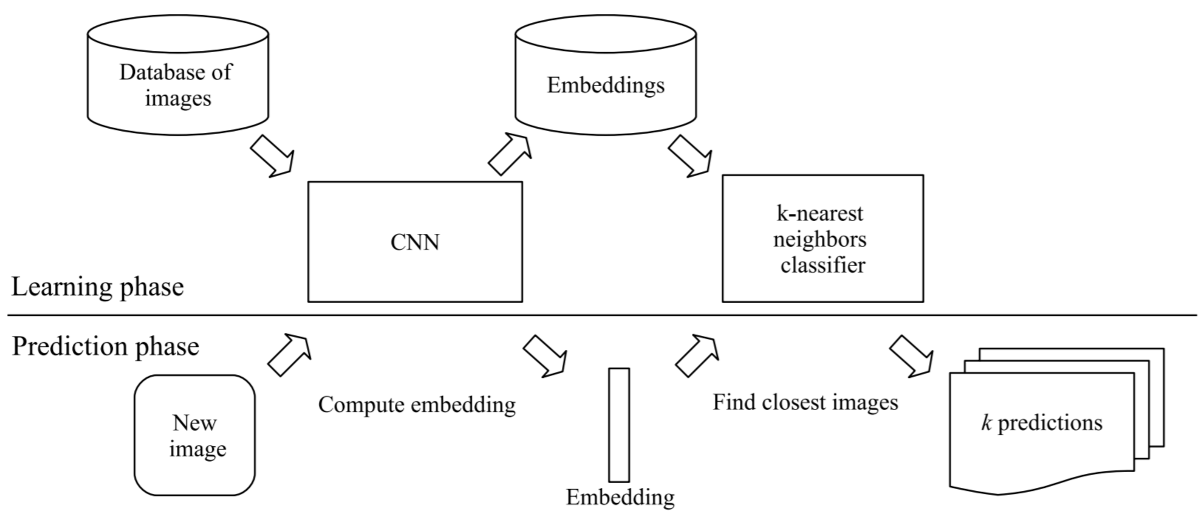

Figure 9. Pose Invariant Embeddings (PIE) algorithm workflow for matching lateral photographs of right whale heads taken from a vessel on the Flukebook platform. From Olga Moskvyak's open source GitHub repository: https://github.com/olgamoskvyak/reid-manta.

Fluke Photos: CurvRank v2

For fluke matching, the Wild Me team implemented the machine learning-based CurvRank v2 algorithm, which has been recently advanced (Charles V. Stewart pers. comm; Alex Mankowski pers. comm.)[1] at Rensselaer Polytechnic Institute over the originally published implementation (Weideman et al 2017; Figure 10). CurvRank v2 uses AI to learn how to extract the trailing edge of a fluke and then searches for matches to other edges by measuring the differences between them, learning to focus on sections that contain more individually identifiable information. The CurvRank algorithm is now available for fluke photographs of both North Atlantic and Southern right whales on the Flukebook platform.

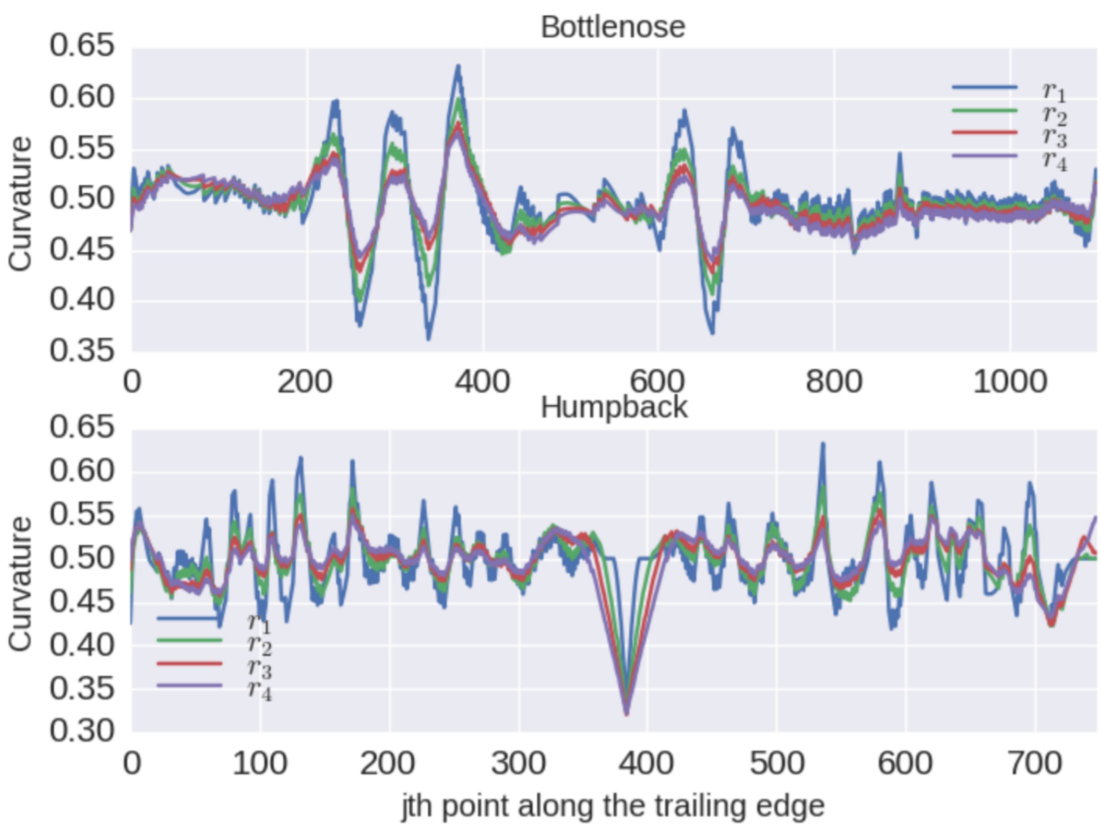

Figure 10. Curvature representation of a dorsal fin from the Bottlenose dataset (top) and fluke from the Humpback dataset (bottom) in CurvRank.

[1] Portions of this work were supported under the Bureau of Ocean Energy Management contract M17AC00013.

Peduncle Scar Photos: HotSpotter

Hotspotter is a SIFT-based computer vision algorithm initially applied to Grevy's zebras, plains zebras, giraffes, leopards, and lionfish (Crall et al 2013). HotSpotter analyzes images for distinct patterns or "hotspots", compares those against other images in the Flukebook database (including recent uploads), and produces a ranked list of potential matches based on "hotspot" similarity. One advantage of this approach is that new individuals can be identified without the need for network retraining. HotSpotter was originally applied to the identification of right whale heads from aerial photographs but given the superior performance of the Deepsense algorithm for this task, it was turned off to maximize server time efficiency. However, the HotSpotter algorithm has been effective in opportunistically matching peduncle scar patterns of right whales and is available for this application.

Figure 11. Example of a correct peduncle insertion scar match using the HotSpotter algorithm in the Flukebook platform (https://www.flukebook.org/). Photo by the New England Aquarium.

Evaluating Model Performance

Top-1 and top-5 accuracy are commonly used evaluation criteria in machine learning algorithms. Top-1 accuracy is the number of times that the model outputs the correct whale identification when allowed to output only one result. Top-5 accuracy is the number of times the model outputs the correct whale identification when allowed to output 5 possible whale identifications.

	top-1	**top-5**
Aerial Heads: Deepsense	89%	98%
Lateral Heads: PIE	55%	81%
Flukes: CurvRank v2	85%	89%
Peduncle Scars: HotSpotter	HotSpotter use is opportunistic, looking for localized scarring and not used exhaustively for individual ID.	HotSpotter use is opportunistic, looking for localized scarring and not used exhaustively for individual ID.

Table 1. Model performance for the different matching algorithms applied to North Atlantic right whales (results not yet available for Southern right whales). The 'top-1' accuracy is the number of times the model suggests the correct individual when only allowed to output one suggestion. The 'top-5' accuracy is the number of times that the correct individual can be found in the top five suggested matches returned.

Conclusions

Building on the work that began with the Kaggle competition and the winning Deepsense algorithm, right whale photo-identification using AI continues to advance. Flukebook successfully implemented multi-feature matching and new AI techniques. Multi-feature matching allows right whales to be matched by aerial photos of their heads (Deepsense), lateral photos of their heads (Pose Invariant Embeddings), flukes (new CurvRank v2), and peduncle scarring (HotSpotter). This capability to apply new forms of AI and match an individual North Atlantic right whale from multiple poses and marks has then been successfully cross-applied to Southern right whales and orcas. Right whale researchers have been slow to adopt these new methods, and we continue to explore ways to streamline the platform and create time saving workflows to encourage broader adoption. Flukebook now allows scientists to bulk import a directory of photographs and associated metadata in an Excel spreadsheet (date, location, species, etc.). Additionally, Flukebook is receiving a new front end interface in early 2022 which should streamline machine learning workflows and simplify the platform for end users. We welcome feedback for further improvements. This progress using AI for photo-identification applies to multiple IWC sub-committees and priority species.

Code Repositories

Deepsense algorithm GitLab repo: https://gitlab.com/deepsense.ai/whales/

Wildbook Image Analysis: https://github.com/WildMeOrg/wildbook-ia
- PIE algorithm Github repo: https://github.com/WildbookOrg/wbia-plugin-pie
- CurvRank v2 Github repo: https://github.com/WildMeOrg/wbia-tpl-curvrank-v2

Literature Cited

Crall JP, Stewart CV, Berger-Wolf TY, Rubenstein DI and Sundaresan SR. 2013. HotSpotter — Patterned species instance recognition. IEEE Workshop on Applications of Computer Vision (WACV), Clearwater Beach, FL, USA, 2013, pp. 230-237, doi: 10.1109/WACV.2013.6475023.

Bogucki R, Cygan M, Khan CB, Klimek M, Milczek JK, and Mucha M. 2019. Applying deep learning to right whale photo identification. Conservation Biology, 33: 676-684. https://doi.org/10.1111/cobi.13226.

Moskvyak O, Maire F, Armstrong AO, Dayoub F, and Baktashmotlagh M. 2019. Robust re-identification of manta rays from natural markings by learning pose invariant embeddings. arXiv:1902.10847.

Weideman HJ, Jablons ZM, Holmberg J, Flynn K, Calambokidis J, Tyson RB, Allen JB, Wells RS, Hupman K, Urian K, and Stewart CV. 2017. Integral Curvature Representation and Matching Algorithms for Identification of Dolphins and Whales. IEEE International Conference on Computer Vision Workshops (ICCVW), Venice, Italy, pp. 2831-2839, doi: 10.1109/ICCVW.2017.334.