Ecological Informatics 89 (2025) 103158

Contents lists available at ScienceDirect

Ecological Informatics

journal homepage: www.elsevier.com/locate/ecolinf

Exploiting facial side similarities to improve AI-driven sea turtle photo-identification systems

Lukáš Adam [a], Kostas Papafitsoros [b,c,*], Claire Jean [d], ALan F. Rees [c], Vojtěch Čermák [e]

[a] University of West Bohemia in Pilsen, FEE, RICE, Pilsen, Czechia
[b] School of Mathematical Sciences, Queen Mary University of London, UK
[c] ARCHELON, the Sea Turtle Protection Society of Greece, Athens, Greece
[d] Kélonia, l'Observatoire des tortues marines, La Réunion, France
[e] Czech Technical University in Prague, FEE, Prague, Czechia

A R T I C L E I N F O

Dataset link: https://github.com/sadda/sides-matching, https://www.kaggle.com/datasets/wildlifedatasets/zakynthosturtles, https://www.kaggle.com/datasets/wildlifedatasets/amvrakikosturtles, https://www.kaggle.com/datasets/wildlifedatasets/reunionturtles

Keywords:
Re-identification
Sea turtles
Artificial neural networks

A B S T R A C T

Animal photo-identification (photo-ID), the process of identifying individual animals from images, has proven to be a valuable tool for various studies on sea turtles, increasing the knowledge of their ecology and informing conservation efforts. Photo-ID in sea turtles is predominantly based on the geometric patterns of the scales of their two head sides, which are unique to every individual and different from side to side. As such, both manual and automated photo-ID techniques are traditionally performed under a side-specific setting. There, an image showing a single profile of an unknown individual is compared only to images showing the same side of previously identified individuals. In this paper, we show for the first time an inherent visual similarity between left and right facial profiles of the same individuals in three sea turtle species. We do so by employing two state-of-the-art automated neural network-based photo-ID methods, one local feature-based and one deep embedding-based, designed to rank profiles based on their similarities. Both methods rank the similarity of the left and right profiles of the same individual higher than those of different individuals. These similarities are detectable even when images are taken years apart under diverse conditions. We further show that the exploitation of this similarity results in improved accuracies when compared to the traditional side-specific photo-ID setting. Our results indicate two concrete guidelines for improving automated sea turtle photo-ID workflows. When trying to match a photo of a given profile, searches should not be restricted only to photos of the same profile. As the first method of choice, a deep embedding model finely-trained using a photo-database of the focal sea turtle population should be used. In the absence of such training database, a neural network-based local feature method is preferable, but in that case searches should be performed with both the original query image and its horizontally flipped version.

1. Introduction

Animal photo-identification (photo-ID; commonly referred to as *re-identification* in the computer vision literature) denotes identifying individual animals from images by exploiting their unique and stable across time external morphological patterns. Photo-ID is particularly widespread in studies of wildlife populations of various taxa and species like cetaceans (Cheeseman et al., 2017), sharks (Araujo et al., 2017), sea turtles (Schofield et al., 2020), rays (Marshall and Pierce, 2012), bears (Anderson et al., 2010), seals (Koivuniemi et al., 2016) giraffes & zebras (Parham et al., 2017) to name a few representative works. It overcomes the logistical difficulties of capturing and tagging the animals using external tags while being minimally invasive, causing minimal or no stress, and it can be performed on large scales when it is combined with active or passive citizen science (Holmberg et al., 2009; Papafitsoros et al., 2021).

Photo-ID is often performed manually by visually comparing newly obtained images of individuals to an existing database which contains previously identified individuals. Due to the ever-increasing number of obtained images, the major challenge in animal photo-ID is that it is often time-consuming. Speeding up the process is possible by employing divide-and-conquer strategies, e.g. by splitting photo-databases into sub-databases of much smaller size, each one corresponding to a particular characteristic, e.g. location, sex or a fine categorisation of the identifying patterns (Schofield et al., 2008; Lloyd et al., 2012; Papafitsoros et al., 2025). However, this only partially solves the scalability problem and is not always applicable.

* Corresponding author at: School of Mathematical Sciences, Queen Mary University of London, UK.
E-mail addresses: lukas.adam.cr@gmail.com (L. Adam), k.papafitsoros@qmul.ac.uk (K. Papafitsoros).

https://doi.org/10.1016/j.ecoinf.2025.103158
Received 13 September 2024; Received in revised form 16 April 2025; Accepted 17 April 2025
Available online 7 May 2025

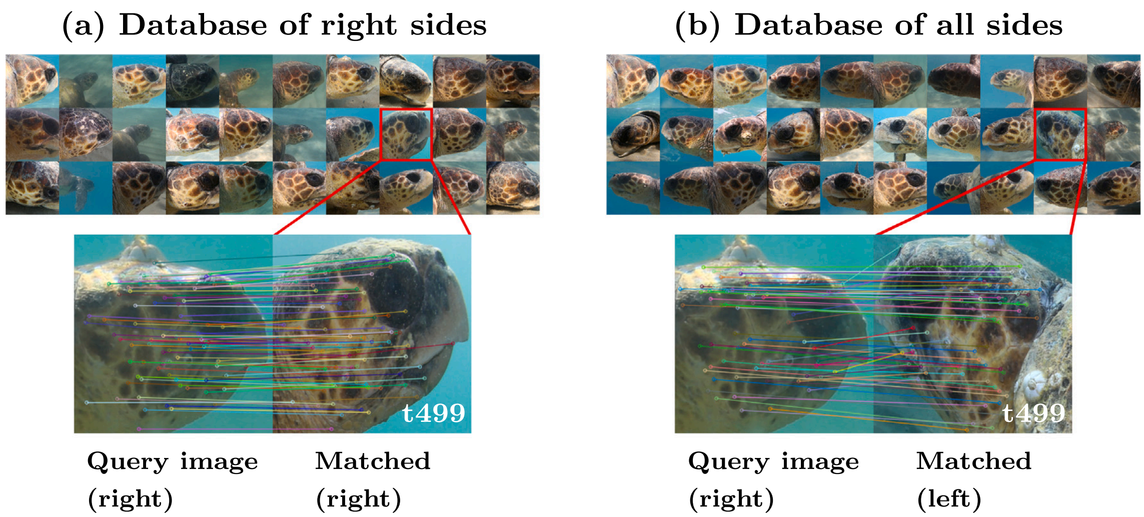

Fig. 1. (a) "Classical" sea turtle photo-ID workflow where a query image (right side of the head) is compared to a database of images of the same side (only right sides). This setting utilises the exact similarity of polygonal patterns, which differ from side to side. (b) Workflow where also opposite sides are compared. The image showing the right side of the individual "t499" is matched to an image showing its left side since this is the most similar one in the database. While so far, photo-ID methods followed approach (a), we show that approach (b) is also feasible and beneficial when using deep methods like MegaDescriptor or ALIKED.

This need to reduce manual labour has driven the development of fast and accurate automated photo-ID methods. Following Čermák et al. (2024b), these methods belong to one of three categories:

- **Local feature-based methods** extract local keypoints (coordinates in the image) and their corresponding descriptors (local image features of salient characteristics, e.g. edges). A classical example of such a method is the Scale Invariant Feature Transform (SIFT) (Lowe, 2004), which, in the animal photo-ID context, extracts descriptors characterising small morphological regions of the imaged animal that do not depend on the overall scale and orientation of the animal within the image. On the other hand, more modern extractors are based on deep neural networks (Goodfellow et al., 2016). They are usually trained on large datasets and then used as pre-trained models. Examples include ALIKED (Zhao et al., 2023), DISK (Tyszkiewicz et al., 2020) or SuperPoint (DeTone et al., 2018). After extracting the keypoints between two images (circles in Fig. 1), the method matches them to determine if the two images show the same individual.
- **Deep embedding-based methods** exploit the power of deep neural networks to represent an image of an individual via an embedding vector of a much lower dimension than the image. The neural network is trained on sufficiently large databases and learns how to efficiently encode characteristics of individual animals to these embedding vectors (Čermák et al., 2024b,a; Shinoda and Shiohara, 2024). These identifying characteristics can be richer than local features and include global geometric patterns, colouration, texture, etc. These methods are often faster than the local feature-based ones because they do not require the expensive matching of keypoints.

  Fine-tuning of deep embedding-based methods: It is well-known that deep embedding-based methods (or any neural network-based method for that matter) are able to identify individuals of a given species with higher success rate when images of these species have been included in their training set (Čermák et al., 2024b; Otarashvili et al., 2024). For example, a model only trained on one sea turtle species, e.g. loggerheads, is not guaranteed to perform well on a different species, e.g. green turtles, since they have slightly different identifying scale patterns (Papafitsoros et al., 2025). In that case, it is advisable to *fine-tune* the model with images of the new species. This typically means starting a new training process using the latter set of new images, having as initial network parameters the ones of the old model (Goodfellow et al., 2016) (transfer learning). As such, the model learns to identify the new species as well without compromising its performance on the species initially trained for. In certain cases, this fine-tuning must be done with images of the same species the model was initially trained with, but taken in different environments. For instance, if a model was initially trained using only underwater images of a turtle species, fine-tuning might still be required if the target is to perform photo-ID on turtles of the same species taken out of the water.
- **Species-specific methods** exploit particular morphological characteristics of the focal species and cannot be used for other species. For example, Anderson et al. (2010), Kelly (2001), and Jean et al. (2010) designed methods that can be used exclusively for polar bears, cheetahs and sea turtles, respectively.

Commonly used automated methods for sea turtle photo-ID have been mainly local feature-based (Buteler et al., 2022) with the most prominent of those being Hotspotter (Crall et al., 2013), which is based on SIFT descriptors, see also Dunbar et al. (2021). On the other hand, the application of deep embedding-based methods to sea turtle photo-ID remains relatively unexplored with some recent emerging publications (Čermák et al., 2024b; Otarashvili et al., 2024).

Sea turtle photo-ID is predominantly based on the animal's facial scales polygonal pattern, which is unique to every individual. It has been shown that these patterns are stable throughout the animals' lives (Carpentier et al., 2016). Except for drone-based sea turtle photo-ID (Comis et al., 2022) which uses the top of the head, or the more rarely used flipper-based photo-ID (Gatto et al., 2018; Pursley, 2020), the identification is predominantly based on the left and right sides of the head (Papafitsoros et al., 2025). The left and right polygonal scale patterns are different in a given individual. As a result, both manual and automated sea turtle photo-ID are "classically" performed under a side-specific image retrieval setting: Any new image (*query image*) depicting a side of the head of an unknown individual is visually or automatically compared to a group of images (the *database*) showing the same side of previously known individuals until a *matching* occurs, see Fig. 1(a). Thus, the individual cannot be identified if only the opposite side of one of the query images is available in the database. However, it has recently been reported (Papafitsoros et al., 2025) that despite the different patterns, there is still some inherent visual similarity between left and right profiles at a given individual, not only in colouration and texture but also geometrically. Yet, this similarity has neither been quantified nor exploited in sea turtle photo-ID and this is precisely the topic of this paper. As a prelude, Fig. 1(b) illustrates our suggested photo-ID workflow, which is based on the fact that state-of-the-art neural network-based photo-ID methods are capable of matching not only the same but also opposite profiles.

In this work, we make the following contributions:

- We quantitatively show the inherent visual similarity between left and right facial profiles of the same individual in three sea turtle species, Loggerheads (*Caretta caretta*), Greens (*Chelonia mydas*) and Hawksbills (*Eretmochelys imbricata*). We do so by leveraging four multi-year databases of these species and by employing three state-of-the-art automated methods, each one being a representative from the three above categories: ALIKED (Zhao et al., 2023), MegaDescriptor (Čermák et al., 2024b) and TORSOOI (Jean et al., 2010).
- We show that allowing comparisons of opposite sides leads to improved accuracies for the deep embedding method (MegaDescriptor), as it can match opposite profiles. Furthermore, the deep embedding method (ALIKED) has the highest accuracy in datasets on which it was trained. On the other hand, the local feature-based method generally outperforms a non-fine-tuned deep embedding method, but only in the case where the same sides are compared. When comparing opposite profiles, simply flipping the query image horizontally to artificially create the same orientation makes the local feature-based method capable of exploiting the similarity of opposite profiles and increases its performance.
- We propose the following photo-ID method selection guidelines: A deep embedding model fine-tuned with images of the focal database should be the first method of choice. Without fine-tuning, a neural network-based local feature method is preferable, but searches should be performed with both the original query image and its horizontally flipped version.
- Finally, to encourage and standardise further research in this novel direction of sea turtle photo-ID we make our datasets and codes publicly available (links provided at the end of the paper).

2. Testing datasets

Before describing our methodology and experimental design, we present the datasets used to perform our experiments. We used four datasets: two of loggerhead sea turtles (one of underwater photographs and one of out-of-water ones), one of green and one of hawksbill sea turtles, both of underwater photographs. We included the following metadata for all images: turtle identity, head orientation (left or right) and timestamp (the year each photograph was taken). We briefly describe the characteristics of these datasets.

2.1. Dataset Zakynthos-Loggerheads

This dataset consists of photographs taken between 2018 and 2024 in Laganas Bay, Zakynthos Island, Greece (37°43′N, 20°52′E), which is a main breeding site for the Mediterranean loggerhead sea turtles (Margaritoulis et al., 2022). All photographs were captured underwater by the same photographer (the second author) during snorkelling surveys from a distance ranging from 7 m to a few centimetres. A Canon 6D full-frame DSLR camera (5472 × 3648 pixels) combined with a Sigma 15 mm fisheye lens was mainly used. A few photographs were captured using a Canon R8 mirrorless camera with either the Sigma 15 mm or the Tokina 10–17 mm lenses. The water depth ranged from 1 to 8 m, with most photographs taken less than 5 m deep.

This dataset consists of images of the same *distribution*, i.e. same capture conditions and sea turtle population, as the SeaTurtleID2022 dataset (Adam et al., 2024a), which was part of the training set of the MegaDescriptor, see Section 3.2. However, it consists of new individual loggerheads which were not included in SeaTurtleID2022 and thus were not seen during the MegaDescriptor training.

Table 1
Overview of the four testing datasets: Number of individuals, total number of images (number of individuals ×4), conditions of photo-capturing and average time-span of photographs in years.

Dataset	Individuals	Images	Conditions	Average span
Zakynthos-Loggerheads	40	160	Underwater	2.5 years
Amvrakikos-Loggerheads	50	200	On a boat	4.4 years
Reunion-Greens	50	200	Underwater	4.7 years
Reunion-Hawksbills	34	136	Underwater	3.4 years

2.2. Dataset Amvrakikos-Loggerheads

This dataset also consists of photographs of Mediterranean loggerhead sea turtles taken at Amvrakikos Gulf, Greece (39°02′N, 21°06′E), which is a well-known foraging site for adult and juvenile turtles. Photographs were collected as part of a long-term capture-mark-recapture project conducted by ARCHELON, the Sea Turtle Protection Society of Greece (Rees et al., 2013). Turtles were captured from a boat using the sea turtle rodeo technique, and among other data collected, photographs of the head sides were taken while the animal was on the boat. In all photographs, the whole side of the head was fully shaded or fully illuminated by the sun. All photographs in this dataset were taken during the summer months (June–August) between 2014 and 2022, using a selection of different digital cameras of varying optical resolution. While it is known that many of the turtles foraging in Amvrakikos Gulf are genetically similar to the ones of Zakynthos (Rees et al., 2017) and thus share common morphological characteristics, the different conditions of photograph collection make it a distinct dataset to the Zakynthos-Loggerheads one.

2.3. Datasets Reunion-Greens & Reunion-Hawksbills

This dataset consists of photographs taken between 2007 and 2024 on Reunion Island, a French territory in the Indian Ocean (21°06′S, 55°30′E), whose shallow waters are known to be a development and foraging ground for green and hawksbill turtles (Chassagneux et al., 2013). Recreational divers took photographs as part of a citizen science programme with no specific associated protocols (no specific viewing angle or distance). Citizens shared their photographs whenever they wished for further analysis with scientists. All photographs are stored in the TORSOOI database; see Section 3.4 for more details.

2.4. Data selection and preprocessing

We briefly summarise the selected photos for each dataset in Table 1. For every individual in each dataset, we used four photographs for our experiments: two (left and right side) taken during a given year and two (left and right side) taken during a different year, see Fig. 2 for sample photos from each dataset. We had a desired quota of 50 individuals per dataset which was achieved for Amvrakikos-Loggeheards and Reunion-Greens. In the other two datasets, Zakynthos-Loggerheads and Reunion-Hawksbills, only 34 and 40 individuals, respectively, met our selection criteria. We note again that for Zakynthos-Loggerheads, we used individuals not included in SeaTurtleID2022 (Adam et al., 2024a), which is a dataset containing images of 400 individuals from Zakynthos Island.

The main reason for including images taken during different years is to eliminate the undesirable phenomenon when an algorithm decides that two images contain the same turtle, not based on the turtle's morphological patterns but on external factors such as similar background or lighting conditions. Images taken during different years do not share the latter two and are only matched due to turtles' morphological characteristics (Adam et al., 2024a). To further reduce any background effect, we applied a bounding box around each turtle's head and all the experiments were performed on the turtles' heads only. We note that for images taken outside of water (Amvrakikos-Loggerheads; see the second row of Fig. 2), there are minimal global lighting differences across the years compared to underwater photos.

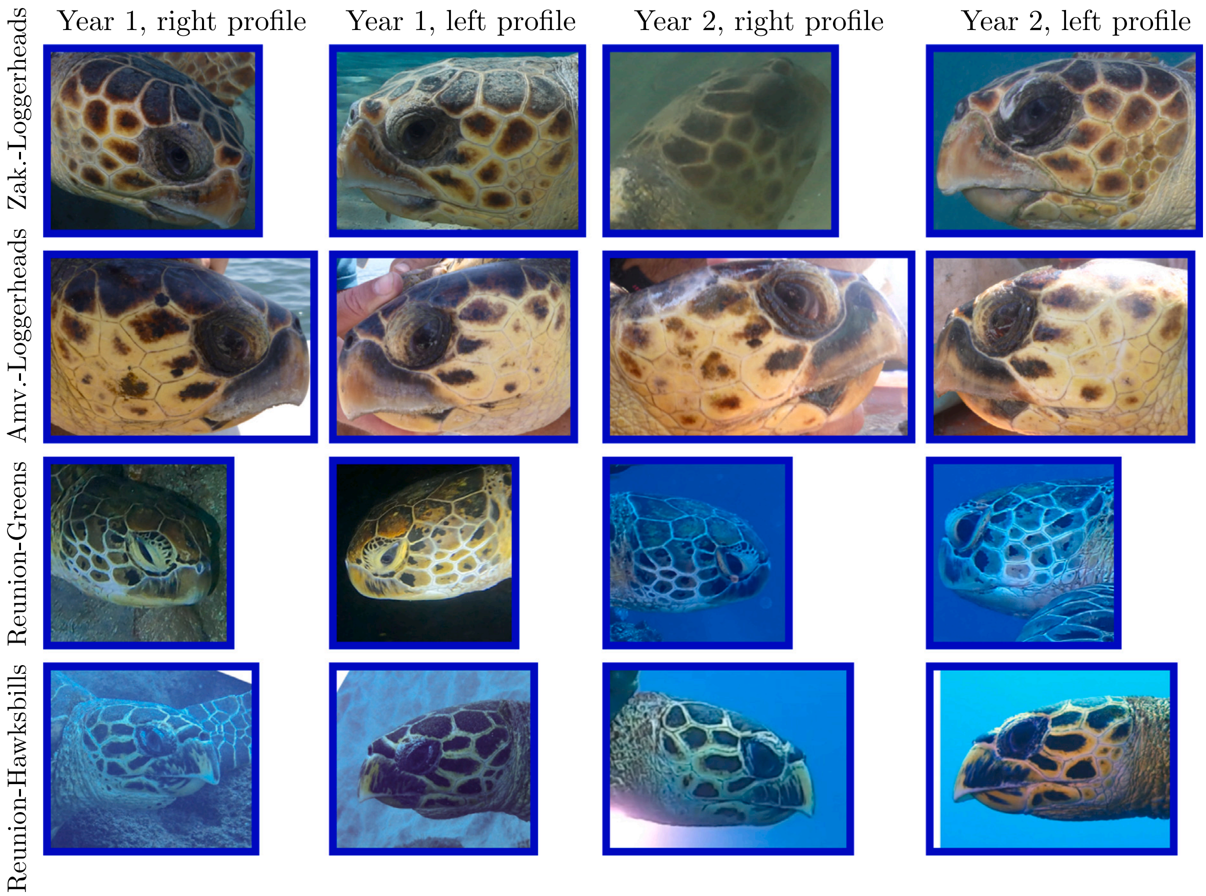

Fig. 2. Sample photos after applying a bounding box around their heads for individuals belonging to each of the four considered datasets (one single individual per row). Four photos were chosen for every individual, taken in two different years, with both profiles represented in both years.

3. Methodology

This section describes the algorithms used and the experimental design. We first describe the general framework of image retrieval based on similarity comparisons between images. We then explain in more detail the representatives of the three method categories mentioned in the introduction: MegaDescriptor (a state-of-the-art deep embedding method), ALIKED (a deep method extracting local features) and TORSOOI (a sea turtle-specific method). We proceed by describing our experimental design, which includes several settings for examining the similarity of sea turtle head profiles and a series of related photo-ID experiments and their evaluation.

3.1. Image retrieval and automated photo-ID methods

Automated photo-ID methods typically operate under an *image retrieval* setting: This assumes that there is a set of images of known individuals, the *database*. For each image depicting a not-yet-recognised individual (*query image*), such a method predicts which image from the database most likely depicts the same individual as the query image. To do so, it first extracts some embedding vectors from all images. These vectors, also called local feature or deep embedding vectors depending on the method, are low-dimensional representations of the identifying areas. The method then compares the extracted embedding vector of the query image with the embedding vectors of all images in the database and selects the closest one based on some *similarity metric*. This selection can further be manually verified. Since the embedding vectors of one individual should form clusters, newly recruited animals (not in the database) should be easily recognised because their embedding vectors should be far from all other embedding vectors in the database. The goal of an automated method is to design the embedding extractor and the similarity metric such that the similarity of two images is high for images of the same individual and low for images of different individuals.

3.2. Deep embedding-based method: MegaDescriptor

is a recently introduced deep embedding vector extractor (Čermák et al., 2024b). It was trained so that images of the same (resp. different) individuals have similar (resp. distinct) deep embedding vectors. MegaDescriptor is freely available from HuggingFace. It is written in Python and designed for easy use, e.g. embedding extraction is performed by one line of code. We briefly sketch how MegaDescriptor is used under an image retrieval setting in Fig. 3. The left side shows the query image, and the right side shows the database. An embedding vector is extracted from each image, and the similarities between embedding vectors are computed using the cosine similarity metric which measures the angle between the embedding vectors. Then, the individual in the query image is predicted to have the same identity as the identity of the database image with the highest similarity score (t086 in the case of Fig. 3).

MegaDescriptor was trained on over 30 datasets of various species, including primates, carnivores, reptiles, whales, mammals and fish. A significant subset of its training dataset is also available separately in an easily accessible form (Adam et al., 2024b). Sea turtles were included in 2 of the above 30 datasets: SeaTurtleID2022 (Adam et al., 2024a), of which Zakynthos-Loggerheads is an extension, as well as ZindiTurtleRecall (Turtle Recall: Conservation Challenge, 2022) which contains green and hawksbills sea turtles ashore. Thus, MegaDescriptor only saw these species in selected environments. However, since it was

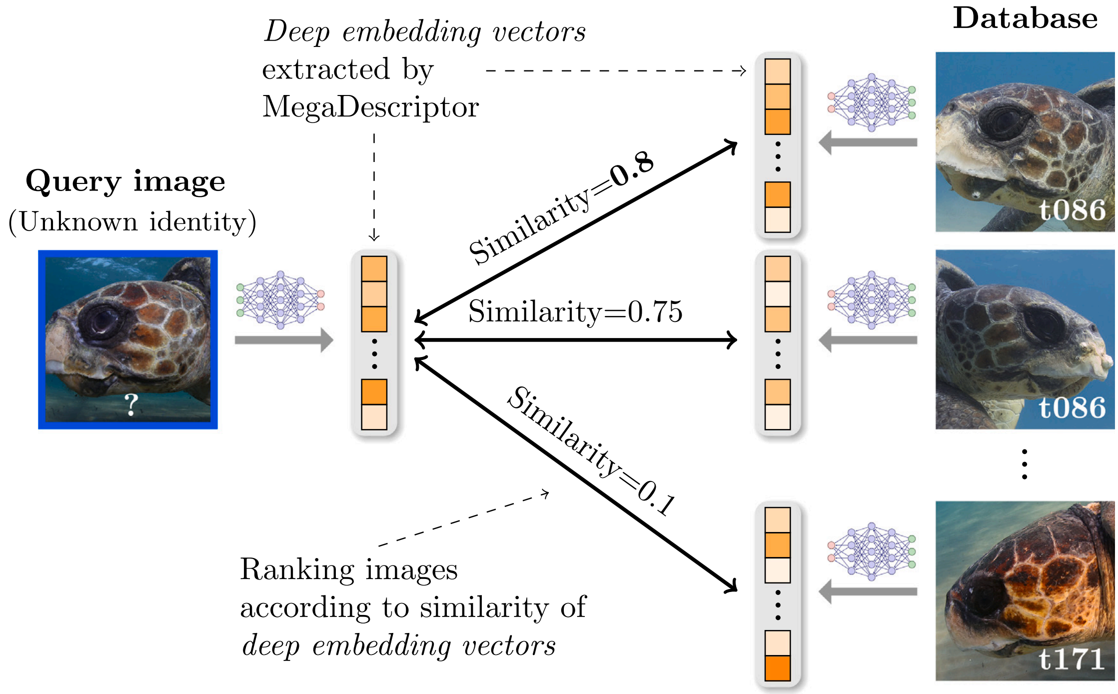

Fig. 3. MegaDescriptor photo-ID workflow: MegaDescriptor extracts *deep embedding vectors* both from the query image and the database images, which can be of any side — left or right. The database embedding vectors are ranked by decreasing similarity to the query embedding vectors, giving a ranking to the database images.

trained on many different species, it is expected to perform well even for species unseen during its training.

3.3. Local feature-based method: ALIKED

is a deep learning-based method for extracting keypoints and descriptors from an image (Zhao et al., 2023). It improves classical methods such as SIFT (Lowe, 2004) by introducing deformable transformations that improve descriptor flexibility and robustness. Keypoints and descriptors extracted from two images are then matched, see Fig. 1, to determine whether the images contain the same individual. The matching is done based on the similarity of descriptors. Since the keypoints are the geometric positions in the image, they are sometimes also used for the matching procedure to verify that the two images could correspond to one object photographed from two different angles. We used LightGlue (Lindenberger et al., 2023) as the matching algorithm.

3.4. Species-specific method: TORSOOI

is a database which is part of a large collaborative sea turtle project, which includes a web application for sea turtle photo-ID. Identifying individual turtles is based on the TORSOOI codes. To generate these codes, the scales on each head profile are divided into columns and represented by a series of numbers denoting the number of edges in each scale in that column (Jean et al., 2010). The similarity between two TORSOOI codes is computed as the number of matching numbers across all scales. Therefore, it computes how often the scales at similar locations have the same number of edges. The predictions are then made using the image retrieval setting as in the two previous methods. We only applied TORSOOI to the Reunion-Greens and Reunion-Hawksbills as the codes were only available for these datasets.

3.5. Experiment 1: Evaluating the similarity of left and right profiles

To properly evaluate the similarities of left and right turtle profiles, we adopt five settings (A)–(E) summarised in Fig. 4. Given an image, each setting defines the set of images to be compared with it. The pairwise comparisons are done among profiles of the same individuals in the settings (A), (B) and (C) and among different individuals in the settings (D) and (E). In the setting (A), images of opposite sides taken in the same year are compared. In the setting (B) (respectively (C)), images of the same (respectively opposite) side taken in a different year are compared. Since by dataset design, there are only four images for each individual, in each of the settings (A), (B), (C), a given image is compared with precisely one other image.

In Experiment 1, we compute the similarity scores given by MegaDescriptor, ALIKED and TORSOOI under all settings (A)–(E). For example, for setting (A), only similarities of pairs of images of the opposite profiles of the same individual, taken in the same year, are considered. The higher the similarity scores, the more visually similar the images are. Therefore, it is reasonable to expect high similarities for setting (B) and low similarities for setting (E).

3.6. Experiment 2: Predicting sea turtle identities under the different settings

Experiment 2 extends Experiment 1 to realistic photo-ID scenarios where sea turtles must be identified. We use the image retrieval setting, where the query image is always one image from the dataset. We control the database, where we always include all images of different individuals but we include only specific images of the same individual according to five different settings, see Table 2. In setting (full), all the remaining three images of the individual are included in the database, while settings (A)–(C) correspond to the settings from Fig. 4. For instance, in setting (A), the only database image that depicts the same individual as the query image is the one of the opposite side taken during the same year; similarly for the settings (B) and (C). Furthermore, we introduce here the setting (B+C), where the database contains both sides of the individual of the query image taken in different years. The setting (B+C) is the most relevant because it corresponds to the common situation when the database contains all the photos from previous years and the query image comes from the current year.

To numerically evaluate every experiment, we use the top-k accuracy ($k = 1, 2, \dots$). Here, the database images are ranked in order of

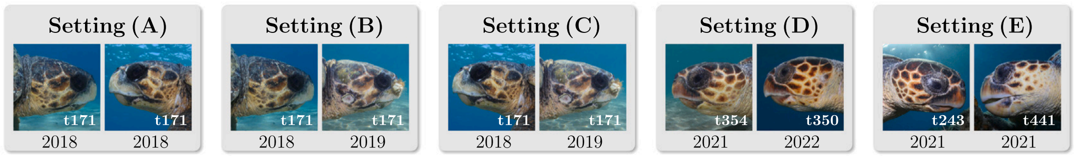

	Setting (A)	Setting (B)	Setting (C)	Setting (D)	Setting (E)
Individual:	Same	Same	Same	Different	Different
Profile:	Opposite	Same	Opposite	Same	Opposite
Year:	Same	Different	Different	Any	Any
Count:	1	1	1	Many	Many

Fig. 4. Settings for comparing similarities of sea turtle head profiles (Experiment 1). "Individual", "Profile" and "Year" refer to whether images in pairwise comparisons under each setting show the same individuals, profiles, or capture years. The last row "Count" refers to the number of images each query image is compared to under each setting.

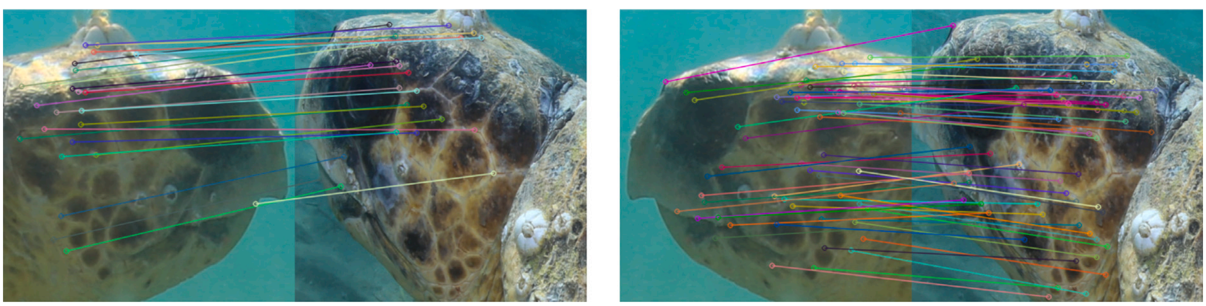

Fig. 5. Motivation for performing experiments after horizontally flipping the query images: While initially ALIKED produces only 24 (mostly wrong) matching keypoints for opposite profiles, when the query image is horizontally flipped, the number of matching keypoints increases to 62 with most of them matching the same head positions.

Table 2
Different settings regarding the series of image retrieval experiments (Experiment 2). "Same individual" refers to which images of the same individual as the query image are included in the database each time, with the last column showing the number of those images. Note that in every setting, all images of all the other individuals are included in the database.

Setting	Same individual	Images of same individual
(full)	All	3
(A)	Opposite side, same year	1
(B)	Same side, different year	1
(C)	Opposite side, different year	1
(B+C)	Both sides, different year	2

decreasing similarity. If at least one of the top-k images shows the same individual as the query image, we deem the whole prediction correct. The overall top-k accuracy for each setting is the ratio of correct predictions over the number of query images. Top-1 accuracy reduces to the standard accuracy. The idea behind multiple ranked predictions is that the user may manually compare only a few images and decide whether the prediction is correct.

3.7. Flipping query images

Besides performing Experiments 1 and 2 as described above, we also performed them with horizontally flipped query images. This is done to systematically examine to what degree artificially "changing" the head orientation makes it simpler for methods to perform matches. We provide a motivation for this in Fig. 5. The left subfigure shows matching keypoints of the two opposite profiles of turtle "t499". There are only 24 matching ALIKED keypoints, which mostly match the wrong positions in the images, e.g. a position of the mouth is erroneously matched to a facial scale. The right subfigure shows matching keypoints after the query image (left) was flipped. Not only did the number of matching keypoints increase to 62, but now they match the same head positions.

This flipping was not performed for the TORSOOI method since flipping does not alter the TORSOOI codes.

4. Results

4.1. Experiment 1: Evaluating the similarity of left and right profiles

In Fig. 6, we show the similarity scores for profile comparisons under all five settings (A)–(E) of Fig. 4, inferred by MegaDescriptor, ALIKED and TORSOOI (three columns) for all four datasets (four rows) for the original (red) and flipped (blue) query images. For every method, higher scores mean higher similarities. However, we note that the similarity values at the horizontal axes among the different methods are incomparable.

MegaDescriptor did not differentiate between the original and flipped images, meaning that it learns inner representations of the individuals independent of the capturing profile. On average, we have the following order for the similarity scores:

MegaDescriptor: $(A) > (B) \approx (C) > (D) \approx (E)$,

TORSOOI: $(B) > (A) \approx (C) > (D) \approx (E)$,

ALIKED original query: $(B) > (D) > (A) \approx (C) \approx (E)$,

ALIKED flipped query: $(A) > (C) > (E) > (B) \approx (D)$.

MegaDescriptor and TORSOOI inferred that the average similarity under (C) (same individual, opposite profile, different year) is larger than those under (D) and (E) (profiles of different individuals). This confirms quantitatively that the similarity between the left and right profiles of the same individual is, on average, higher than the similarity between profiles of different individuals. The same is inferred by ALIKED, which can be seen from (C flipped) > (D original).

Both MegaDescriptor and TORSOOI assigned higher similarities to the same individuals ((A), (B) and (C)) than different individuals ((D) and (E)), which is the desirable property for a photo-ID method. On the other hand, ALIKED (original) assigned higher similarities to images of the same side ((B) and (D)), with (B)>(D) but very low similarities to

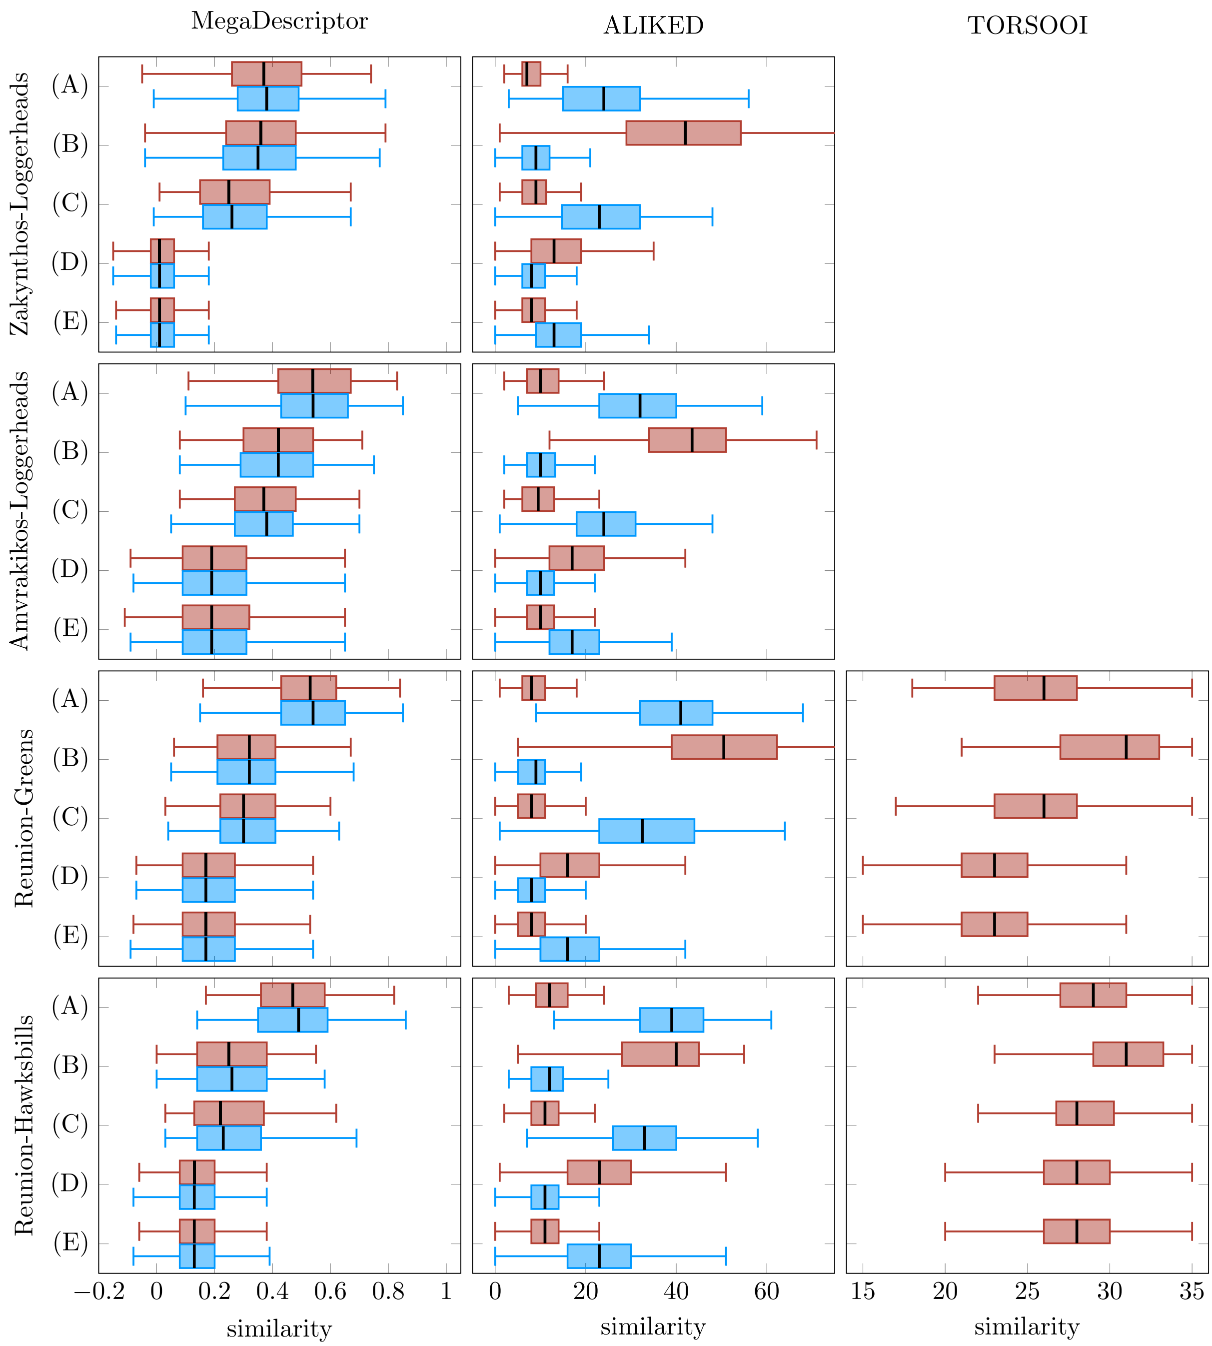

Fig. 6. Similarity scores for profile comparisons under all five settings (A)–(E) of Fig. 4, inferred by MegaDescriptor, ALIKED and TORSOOI (three columns) for all four datasets (four rows). Red and blue bars correspond to experiments with the original and flipped query images, respectively. (For interpretation of the references to colour in this figure legend, the reader is referred to the web version of this article.)

images of the opposite sides ((A), (C) and (E)). This indicates it is a suitable method but only for the classical side-specific photo-ID workflow. However, similarities of opposite sides for ALIKED significantly improved when one of the images was flipped. This can be explained by the fact that, being a local feature-based method, ALIKED produces more matching keypoints for images whose geometrical patterns are more "aligned". This happens only when these images are flipped to artificially make them have the same orientation, see again Fig. 5.

When focusing on similarities between images of the same individuals (A)–(C), MegaDescriptor infers that in all four datasets, the similarity under (A) (same individual, opposite profile, same year) is higher than the similarity under (B) (same individual, same profile, different year). On the other hand, we observe a reverse situation with TORSOOI with (A) < (B) and also with ALIKED, which infers that (A flipped) < (B original) in all datasets. This difference indicates that deep embedding-based methods are more likely to be affected by "global" changes to the animals over the years (e.g. colouration) than by the mere geometrical differences between opposite profiles. On the other hand, local feature-based methods rely more on geometrical features, which are robust (unaffected by time) for the same individual.

In Fig. 7 (top two rows), we have provided examples of photos of setting (A) (same individual, opposite profile, same year) with the largest similarity scores according to MegaDescriptor. One can readily see the reasons why the similarity was high. For example, the left and right profiles from the top Zakynthos-Loggerheads individual have

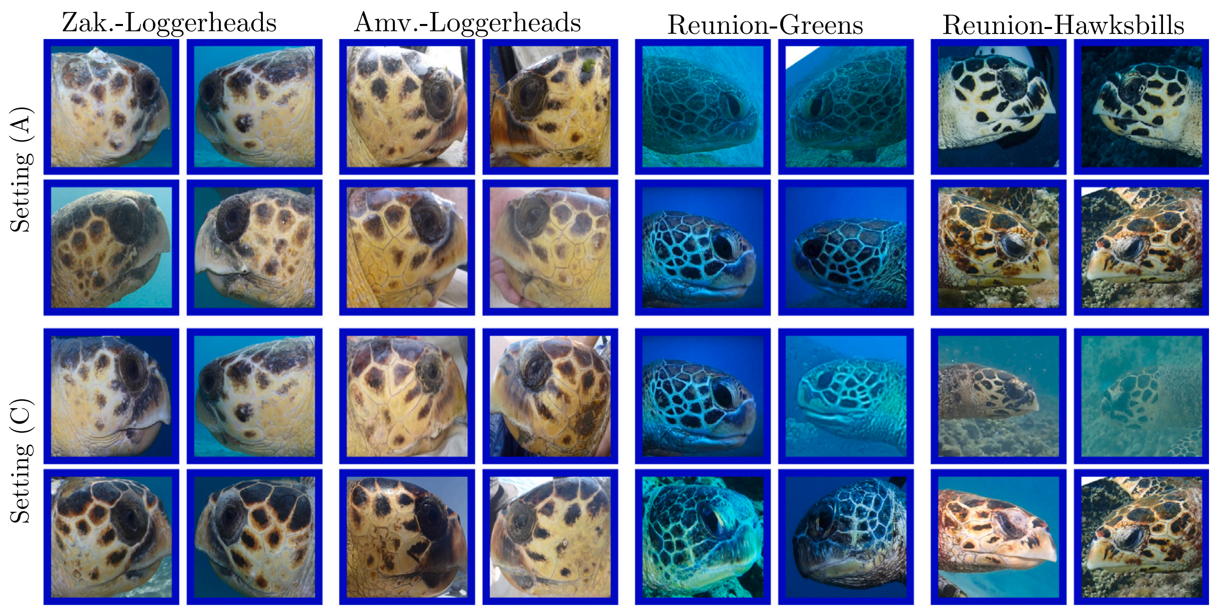

Fig. 7. Top: Photos of opposite profiles of the same individuals taken in the same year (setting (A)) ranked with the two highest similarity scores. Bottom: Photos of opposite profiles of the same individuals taken in different years (setting (C)) ranked with the two highest similarity scores. All rankings are according to MegaDescriptor.

Table 3
Top-5 accuracy for the image retrieval experiments under all settings from Table 2. For each dataset and setting, we have highlighted the best-performing method.

Method	Dataset	Setting				
		(full)	(A)	(B)	(C)	(B+C)
MegaDescriptor	Zakynthos-Loggerheads	0.987	0.919	0.862	0.694	0.912
	Amvrakikos-Loggerheads	0.795	0.730	0.360	0.225	0.420
	Reunion-Greens	0.830	0.765	0.260	0.255	0.350
	Reunion-Hawksbills	0.882	0.838	0.390	0.324	0.463
ALIKED original	Zakynthos-Loggerheads	0.750	0.019	0.737	0.006	0.744
	Amvrakikos-Loggerheads	0.805	0.005	0.805	0.000	0.805
	Reunion-Greens	0.830	0.000	0.830	0.000	0.830
	Reunion-Hawksbills	0.618	0.007	0.610	0.000	0.610
ALIKED flipped	Zakynthos-Loggerheads	0.625	0.387	0.006	0.406	0.412
	Amvrakikos-Loggerheads	0.655	0.495	0.005	0.280	0.280
	Reunion-Greens	0.845	0.720	0.000	0.555	0.555
	Reunion-Hawksbills	0.757	0.618	0.000	0.382	0.382
TORSOOI	Reunion-Greens	0.780	0.315	0.720	0.265	0.765
	Reunion-Hawksbills	0.618	0.184	0.507	0.140	0.559

high similarity due to the common whitish texture and the not well-defined pigmentation within the scales. The profiles of both pairs of the Amvrakikos-Loggerheads dataset look similar regarding colouration and geometry. For Reunion-Greens and Reunion-Hawksbills, the similarity of left–right profiles is also evident (especially the almost identical pattern in the top Hawksbill pair). Fig. 7 (bottom two rows) shows pairs of images with high similarity for setting (C) (same individual, opposite profile, different year). Once again, we notice visual similarities between the left and right profiles with respect to both colouration and geometry. Since photos were taken in different years, the similarities are guaranteed not to be attributed to external factors (background) but only to the animals themselves.

4.2. Experiment 2: Predicting sea turtle identities under the different settings

In Table 3, we provide the results of the image retrieval experiments in terms of top-5 accuracy. For each dataset and setting, we have highlighted the best-performing method. From this table, we have the following results:

- MegaDescriptor outperforms the other methods on Zakynthos-Loggerheads in all settings. Since MegaDescriptor was trained on the dataset SeaTurtleID2022, which is very similar to Zakynthos-Loggerheads, this shows the benefits of fine-tuning deep networks. TORSOOI was never the best method, showing the general advantage of neural network-based methods over hand-crafted ones.
- MegaDescriptor is the best method for settings (full) and (A), i.e. when query images are compared among others with images from the same year. However, ALIKED performs well for the most relevant setting in photo-ID, that is (B+C), where the query images come from different years than the ones in the search database. Notably, it significantly outperformed MegaDescriptor in the three datasets which the latter was not fine-tuned on.
- For (B+C), the difference between ALIKED-original and ALIKED-flipped measures the similarity of opposite profiles. There are two extremes. If the accuracy for ALIKED-flipped is zero, it would suggest no similarity between opposite profiles. If the accuracy for ALIKED-flipped were the same as for ALIKED-original, it would suggest that the opposite profiles are identical. Since the accuracy for ALIKED-flipped is usually more than half for ALIKED-original, it suggests a significant similarity between opposite profiles.

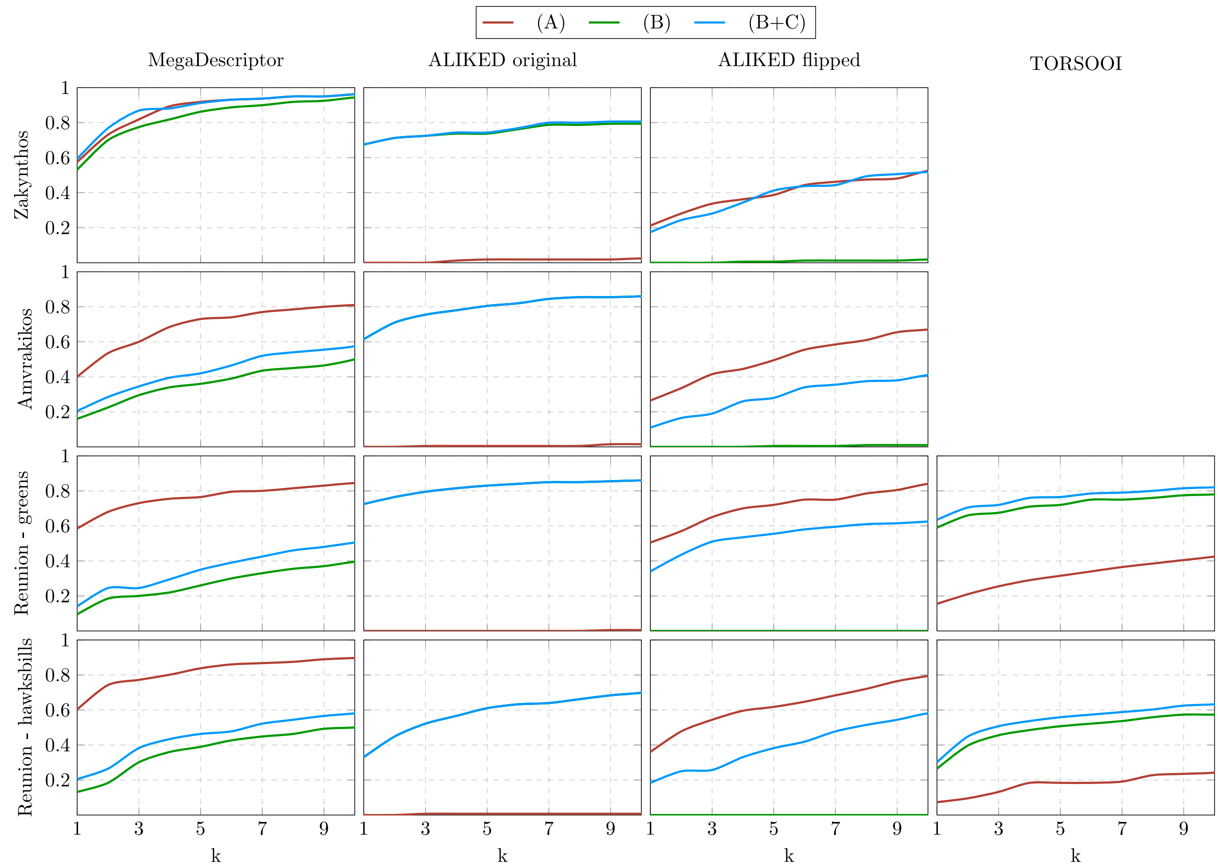

Fig. 8. Top-k accuracies for the series of image retrieval experiments, for various values of k, under selected settings from Table 2.

- ALIKED is unable to match opposite profiles. This manifests in the essentially zero accuracy for settings (A) and (C) for ALIKED-original and (B) for ALIKED-flipped. This is consistent with the very low similarity scores of the same settings from Fig. 6.

We complement Table 3 by Fig. 8, where we provide top-k accuracy for various values of k. For illustrational simplicity, we plot only setting (A), (B) and (B+C). Table 2 can be deduced from Fig. 8 by considering $k = 5$. Upon careful observation, we can deduce that the results stated above hold even for Fig. 8.

5. Discussion

Sea turtle photo-ID has seen remarkable progress over the years with respect to expanding its range of ecological applications and the ever-increasing improvement of automated techniques. Up to now, both the currently used automated photo-ID methods (Dunbar et al., 2014; Calmanovici et al., 2018; Jean et al., 2010; Dunbar et al., 2021; Mills et al., 2023) and those based on manual divide-and-conquer strategies (Schofield et al., 2008; Papafitsoros et al., 2025) are local feature-based or species-specific methods operating under a side-specific image retrieval setting where only comparisons between the same profiles are performed, i.e. left vs. left or right vs. right. Our results imply that methods based on deep neural networks, be it embedding-based (MegaDescriptor) or local feature-based (ALIKED), can assign high similarity scores not only to pairs of the same profiles of the same individual but also to pairs of opposite profiles.

Direct benefit to the sea turtle photo-ID workflows

The capability of these methods to recognise the left–right profile similarities promises to directly benefit sea turtle photo-ID workflows. Firstly, it is particularly relevant to cases where query images of individuals and their available database entries show only opposite profiles due to, i.e. skittish animal behaviour or incomplete citizen science submissions (Papafitsoros et al., 2025). We showed that a fine-tuned deep embedding method and a query-flipped local feature one can match such images. Secondly, we showcased that this capability also results in increased photo-ID accuracy. We observed in all datasets, especially when using MegaDescriptor, that when comparing photos of different years, restricting the searches to the same profiles provides only a slight advantage over doing so based on the opposite profiles only. Most importantly, there is an apparent gain in using both profiles. It is known that the more photos of an individual already exist in the database, the higher the chances that a query photo of that individual will be matched to at least one of those, resulting in a correct identity prediction (Dunbar et al., 2014). Therefore, database curators should still collect as many images of each individual as possible of both sides, but they should do the searches in a non-side-specific manner. Since there are two profiles per individual and only one dorsal side, this hints at the superiority of using the lateral sides over the dorsal side for photo-ID. However, dorsal-based sea turtle photo-ID is still the only option when drones are the primary means for capturing photos (Comis et al., 2022).

Factors contributing to similarity of opposite profiles

Even though the higher similarity of left and right profiles of a given individual compared to the one between different individuals was detected by all methods, we could not quantify to what degree this was due to the similar geometrical patterns of the scales, the inherent colouration and pigmentation of the individual or other characteristics, e.g. the shape of the head (Casale et al., 2017). Colouration, pigmentation and texture are common throughout the head of an individual turtle. As such, it is expected that these factors do drive the detected high similarity. The reasons that skin colouration and pigmentation differ among individuals are not understood well but are likely attributed to genetic factors, differences in foraging habits and sun exposure (Papafitsoros et al., 2025). Furthermore, the fact that ALIKED and TORSOOI take into account spatial information suggests that the high similarity scores were also driven by similar geometrical patterns between left and right profiles. Interestingly, there are essentially no biological theories that could explain the reasons behind this left–right geometric similarity and further research is required, particularly concerning the mechanisms behind scale formation at the embryonic stages (Moustakas-Verho et al., 2014; Zimm, 2019).

While geometrical scale patterns are stable throughout a turtle's life span (Carpentier et al., 2016), pigmentation and colouration on a turtle's facial skin can change due to various factors like ageing, shifts of foraging habits, or seasonal changes in sun exposure (Adam et al., 2024a; Papafitsoros et al., 2025). Factors like algae, injuries and scratches can also change the appearance of a turtle year after year. This could explain the high similarity scores that were assigned by MegaDescriptor to opposite pairs of the same individual taken in the same year (setting (A)) compared to pairs taken in different years (setting (C)). It turned out that for MegaDescriptor, it was easier to recognise individuals based on a comparison of photos of opposite profiles of an individual taken in the same year, compared to comparisons of the same profiles taken in different years (setting (B)). We could not entirely exclude the possibility that common backgrounds and global colouration, which are often similar in photos taken during the same encounter (Adam et al., 2024a,b), could have artificially inflated the similarities in (A). It is also possible that other factors, such as the same degree of focus or use of strobe lights on photos of the same encounter resulting in increased contrast and scale definition, could have influenced these results. This could be the case for the Reunion-Greens and Reunion-Hawksbills datasets, where the relationship (A) > (B) was more pronounced. On the other hand, we note that (A) > (B) holds, albeit to a lesser degree, also for the Zakynthos-Loggerheads and Amvrakikos-Loggerheads datasets even though the photo-capturing conditions were essentially uniform across the years. Thus, we argue that the high similarity scores in (A) compared to (C) is most likely attributed to changes in the turtles' appearance.

Guidelines for automated method selection

The accuracy of MegaDescriptor for the Zakynthos-Loggerheads dataset was much higher than the other three datasets. This is almost certainly because images of individuals from the SeaTurtleID2022 dataset were part of the training set of MegaDescriptor. Even though the individuals of Zakynthos-Loggerheads were not part of SeaTurtleID2022, turtles from both of these datasets belong to the same population, thus sharing similar morphological characteristics and their photos were taken under similar conditions, with similar cameras and by the same photographer. As already mentioned, deep embedding and, in general, neural network-based methods become more accurate when trained on images of the same characteristics (same distribution) as the ones tested on (Goodfellow et al., 2016). Therefore, MegaDescriptor, or generally deep embedding-based methods, should be used whenever there is a sufficient amount of training data in terms of number of images, identities, capture settings and orientations which can be used for fine-tuning them. The order of magnitude of the number of images needed typically depends on many factors, such as the number of network parameters, the size of the turtle population, and the diversity of image capture environments. On the other hand, the local feature method ALIKED outperformed MegaDescriptor for the three other datasets. This implies that such a method should be used whenever training data are not available for fine-tuning a deep embedding-based method. In that case, our results imply that performing the matching for both the original and horizontally flipped query image is advantageous. This is particularly useful when individuals have only a limited amount of photos in the database since by doing two such searches, there is a higher chance that one or more of the top-retrieved images will be a successful match.

The importance of public high-quality datasets

Our work further highlighted the need for multiple well-curated, publicly available multiple-year-spanned sea turtle datasets that can be used both for developing algorithms (training) and their proper evaluation (testing). In particular, it is crucial that timestamps, i.e. capture dates, and orientation labels of the heads, are included in the metadata of each image. For instance, training deep embedding-based methods with photos of individuals spanning several years could force the extracted embedding vectors to encode these identifying characteristics that are stable in time. This is particularly crucial for developing methods to identify individuals over large time scales, e.g. from juveniles to adults. Maturation of sea turtles up to a point where they stop growing can take decades (Baldi et al., 2023), making the assemblage of such datasets particularly challenging. On the other hand, adding timestamps to each photo is simple yet crucial for proper method evaluation (Adam et al., 2024a). Comparing the query images only with those taken in different years allows us to mimic realistic photo-ID matching workflows. It avoids the comparison of photos taken in the same conditions/encounters, which could artificially inflate the accuracy of algorithms due to the shared background or global colouration. We also encourage method evaluations using multiple datasets, with various species, created with different settings. Most sea turtle photo-ID studies typically use a single dataset for a single location to perform experiments. However, we showed how the results of a method can vary on different datasets and how this can be explained based on their characteristics. At the moment, there are only two publicly available sea turtle datasets for algorithmic training and testing: SeaTurtleID2022 (Adam et al., 2024a), and ZindiTurtleRecall (Turtle Recall: Conservation Challenge, 2022), and we strongly recommend researchers to publish also their datasets. Here, we contribute to that need by making all our four datasets publicly available.

6. Conclusion

In this work, we showed that state-of-the-art deep re-identification methods can detect the similarity of left and right profiles in sea turtles. While our current work focuses on sea turtles, it paves the road for exploiting morphological symmetries of opposite sides of other animals as well, when performing photo-ID with deep neural networks. Such detectable symmetry, for instance, has been shown to exist in dolphins (Genov et al., 2018), and it is likely to exist in other species whose identification has been so far performed under the same side setting. We thus anticipate further research towards this direction, which will be facilitated by the constant progress of computer vision and deep learning on the one hand and the ever-increasing availability of public databases of wild animals on the other (Čermák et al., 2024b).

CRediT authorship contribution statement

Lukáš Adam: Writing – original draft, Software, Methodology, Data curation. **Kostas Papafitsoros:** Writing – original draft, Methodology, Data curation, Conceptualization. **Claire Jean:** Writing – review & editing, Data curation. **ALan F. Rees:** Writing – review & editing, Data curation. **Vojtěch Čermák:** Writing – review & editing, Software, Methodology.

Funding

This research has been supported by the Ministry of Education, Youth and Sports of the Czech Republic under project SGS-2024-017 and by the Technology Agency of the Czech Republic, Czechia, project No. SS05010008.

Declaration of competing interest

The authors declare that they have no known competing financial interests or personal relationships that could have appeared to influence the work reported in this paper.

Acknowledgements

We would like to thank the two anonymous reviewers whose comments and suggestions improved the quality of the paper.

Data availability

Code is available at:
https://github.com/sadda/sides-matching
Datasets are available at:
https://www.kaggle.com/datasets/wildlifedatasets/zakynthosturtles
https://www.kaggle.com/datasets/wildlifedatasets/amvrakikosturtles
https://www.kaggle.com/datasets/wildlifedatasets/reunionturtles.

References

Adam, Lukáš, Čermák, Vojtěch, Papafitsoros, Kostas, Picek, Lukas, 2024a. SeaTurtleID2022: A long-span dataset for reliable sea turtle re-identification. In: Proceedings of the IEEE/CVF Winter Conference on Applications of Computer Vision. WACV, pp. 7146–7156.

Adam, Lukáš, Čermák, Vojtěch, Papafitsoros, Kostas, Picek, Lukas, 2024b. WildlifeReID-10k: Wildlife re-identification dataset with 10k individual animals. arXiv preprint arXiv:2406.09211. URL https://arxiv.org/abs/2406.09211.

Anderson, Carlos J.R., Da Vitoria Lobo, Niels, Roth, James D., Waterman, Jane M., 2010. Computer-aided photo-identification system with an application to polar bears based on whisker spot patterns. J. Mammal. 91 (6), 1350–1359, URL https://doi.org/10.1644/09-MAMM-A-425.1.

Araujo, Gonzalo, Snow, Sally, So, Catherine Lee, Labaja, Jessica, Murray, Ryan, Colucci, Anna, Ponzo, Alessandro, 2017. Population structure, residency patterns and movements of whale sharks in southern leyte, Philippines: results from dedicated photo-ID and citizen science. Aquat. Conserv.: Mar. Freshw. Ecosyst. 27 (1), 237–252, URL https://doi.org/10.1002/aqc.2636.

Baldi, Giulia, Furii, Giovanni, Del Vecchio, Massimo, Salvemini, Pasquale, Vallini, Carola, Angelini, Valeria, Pari, Sauro, Lombardi Moraes, Kamyla, Margaritoulis, Dimitris, Çurri, Anxhela, Lazar, Bojan, Casale, Paolo, 2023. Growth rates and age at maturity of mediterranean loggerhead sea turtles estimated from a single-population foraging ground. Mar. Biol. 170 (36), URL https://doi.org/10.1007/s00227-023-04189-1.

Buteler, Candela, Bardier, Cecilia, Cabrera, Mario Roberto, Gonzalez, Yaco, Vélez-Rubio, Gabriela Manuela, 2022. To tag ornot to tag: comparative performance of tagging and photo-identification in a long-term mark-recapture of Juvenile Green Turtles (*Chelonia mydas*). Amphibia- Reptil. 44 (1), 45–58, URL https://doi.org/10.1163/15685381-bja10119.

Calmanovici, Bruna, Waayers, David, Reisser, Julia, Clifton, Julian, Proietti, Maira, 2018. I3S pattern as a mark-recapture tool to identify captured and free-swimming sea turtles: an assessment. Mar. Ecol. Prog. Ser. 589, 263–268, URL https://doi.org/10.3354/meps12483.

Carpentier, Alice S., Jean, Claire, Barret, Mathieu, Chassagneux, Agathe, Ciccione, Stéphane, 2016. Stability of facial scale patterns on green sea turtles chelonia mydas over time: A validation for the use of a photo-identification method. J. Exp. Mar. Biol. Ecol. 476, 15–21, URL https://doi.org/10.1016/j.jembe.2015.12.003.

Casale, Paolo, Freggi, Daniela, Rigoli, Alessandro, Ciccocioppo, Amedeo, Luschi, Paolo, 2017. Geometric morphometrics, scute patterns and biometrics of loggerhead turtles (*Caretta caretta*) in the central mediterranean. Amphibia- Reptil. 38 (2), 145–156, URL https://doi.org/10.1163/15685381-00003096.

Čermák, Vojtěch, Picek, Lukas, Adam, Lukáš, Neumann, Lukáš, Matas, Jiří, 2024a. WildFusion: Individual animal identification with calibrated similarity fusion. arXiv preprint arXiv:2408.12934. URL https://arxiv.org/abs/2408.12934.

Čermák, Vojtěch, Picek, Lukas, Adam, Lukáš, Papafitsoros, Kostas, 2024b. Wildlife-Datasets: An open-source toolkit for animal re-identification. In: Proceedings of the IEEE/CVF Winter Conference on Applications of Computer Vision. WACV, pp. 5953–5963.

Chassagneux, Agathe, Jean, Claire, Bourjea, Jérôme, Ciccione, Stéphane, 2013. Unraveling behavioral patterns of foraging hawksbill and green turtles using photo-identification. Mar. Turt. Newsl. 137.

Cheeseman, Ted, Johnson, Tory, Southerland, Ken, Muldavin, Noah, 2017. Happywhale: Globalizing marine mammal photo identification via a citizen science web platform. Happywhale, St. Cruz URL https://www.happywhale.com/home.

Comis, C., McKenzie, A., Witzmann, J., Vallianos, N., 2022. Sea turtles in the harbour of argostoli and the koutavos lagoon. Wildl. Sense Tech. Rep. URL https://wildlifesense.com/en/wp-content/uploads/2021-Sea-Turtles-in-the-Harbour-of-Argostoli-and-the-Koutavos-Lagoon-2021-Wildlife-Sense-Report.pdf.

Crall, Jonathan P., Stewart, Charles V., Berger-Wolf, Tanya Y., Rubenstein, Daniel I., Sundaresan, Siva R., 2013. Hotspotter–patterned species instance recognition. In: IEEE Workshop on Applications of Computer Vision. WACV, pp. 230–237, URL https://doi.org/10.1109/WACV.2013.6475023.

DeTone, Daniel, Malisiewicz, Tomasz, Rabinovich, Andrew, 2018. Superpoint: Self-supervised interest point detection and description. In: Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition Workshops. pp. 224–236.

Dunbar, Stephen G., Anger, Edward C., Parham, Jason R., Kingen, Colin, Wright, Marsha K., Hayes, Christian T., Safi, Shahnaj, Holmberg, Jason, Salinas, Lidia, Baumbach, Dustin S., 2021. HotSpotter: Using a computer-driven photo-id application to identify sea turtles. J. Exp. Mar. Biol. Ecol. 535, 151490, URL https://doi.org/10.1016/j.jembe.2020.151490.

Dunbar, S.G., Ito, H.E., Bahjri, K., Dehom, S., Salinas, L., 2014. Recognition of juvenile hawksbills *Eretmochelys imbricata* through face scale digitization and automated searching. Endanger. Species Res. 26 (2), 137–146, URL https://doi.org/10.3354/esr00637.

Gatto, Christopher R, Rotger, Andreu, Robinson, Nathan J, Tomillo, Pilar Santidrián, 2018. A novel method for photo-identification of sea turtles using scale patterns on the front flippers. J. Exp. Mar. Biol. Ecol. 506, 18–24, URL https://doi.org/10.1016/j.jembe.2018.05.007.

Genov, Tilen, Centrih, Tina, Wright, Andrew J., Wu, Gi-Mick, 2018. Novel method for identifying individual cetaceans using facial features and symmetry: A test case using dolphins. Mar. Mam. Sci. 34 (2), 514–528, URL https://doi.org/10.1111/mms.12451.

Goodfellow, Ian, Bengio, Yoshua, Courville, Aaron, 2016. Deep Learning. MIT Press.

Holmberg, Jason, Norman, Bradley, Arzoumanian, Zaven, 2009. Estimating population size, structure, and residency time for whale sharks *Rhincodon typus* through collaborative photo-identification. Endanger. Species Res. 7, 39–53, URL https://doi.org/10.3354/esr00186.

Jean, Claire, Ciccione, Stéphane, Talma, Elke, Ballorain, Katia, Bourjea, Jérôme, 2010. Photo-identification method for green and hawksbill turtles - first results from reunion. Indian Ocean. Turt. Newsl. 11.

Kelly, Marcella J., 2001. Computer-aided photograph matching in studies using individual identification: an example from serengeti cheetahs. J. Mammal. 82, 440–449, URL https://doi.org/10.1644/1545-1542(2001)082<0440:CAPMIS>2.0.CO;2.

Koivuniemi, Meeri, Auttila, Miina, Niemi, Marja, Levänen, Riikka, Kunnasranta, Mervi, 2016. Photo-ID as a tool for studying and monitoring the endangered saimaa ringed seal. Endanger. Species Res. 30, 29–36, URL https://doi.org/10.3354/esr00723.

Lindenberger, Philipp, Sarlin, Paul-Edouard, Pollefeys, Marc, 2023. Lightglue: Local feature matching at light speed. In: Proceedings of the IEEE/CVF International Conference on Computer Vision. pp. 17627–17638.

Lloyd, Jane R., Maldonado, Miguel Á., Stafford, Richard, 2012. Methods of developing user-friendly keys to identify green sea turtles (*Chelonia mydas* L.) from photographs. Int. J. Zoöl. 2012 (1), 317568, URL https://doi.org/10.1155/2012/317568.

Lowe, David G., 2004. Distinctive image features from scale-invariant keypoints. Int. J. Comput. Vis. 60, 91–110, URL https://doi.org/10.1023/B:VISI.0000029664.99615.94.

Margaritoulis, Dimitris, Lourenço, Gonçalo, Riggall, Thomas E., Rees, ALan F., 2022. Thirty-eight years of loggerhead turtle nesting in Laganas Bay, Zakynthos, Greece: A review. Chelonian Conserv. Biol. 21 (2), 143–157, URL https://doi.org/10.2744/CCB-1531.1.

Marshall, A.D., Pierce, S.J., 2012. The use and abuse of photographic identification in sharks and rays. J. Fish Biol. 80 (5), 1361–1379, URL https://doi.org/10.1111/j.1095-8649.2012.03244.x.

Mills, Sophie K., Rotger, Andreu, Brooks, Annabelle M.L., Paladino, Frank V., Robinson, Nathan J., 2023. Photo identification for sea turtles: Flipper scales more accurate than head scales using APHIS. J. Exp. Mar. Biol. Ecol. 566, 151923, URL https://doi.org/10.1016/j.jembe.2023.151923.

Moustakas-Verho, Jacqueline E., Zimm, Roland, Cebra-Thomas, Judith, Lempiäinen, Netta K., Kallonen, Aki, Mitchell, Katherine L., Hämäläinen, Keijo, Salazar-Ciudad, Isaac, Jernvall, Jukka, Gilbert, Scott F., 2014. The origin and loss of periodic patterning in the turtle shell. Development 141 (15), 3033–3039, URL https://doi.org/10.1242/dev.109041.

Otarashvili, Lasha, Subramanian, Tamilselvan, Holmberg, Jason, Levenson, JJ, Stewart, Charles V, 2024. Multispecies animal re-ID using a large community-curated dataset. arXiv preprint arXiv:2412.05602. URL https://arxiv.org/abs/2412.05602.

Papafitsoros, Kostas, Hoh, Daphne, Köhnk, Stephanie, Montagna, Micol, Witzmann, Josh, 2025. Sea turtle photo identification: A practical guide. Preprint, URL https://www.researchgate.net/publication/388909753_Sea_Turtle_Photo_Identification_A_Practical_Guide.

Papafitsoros, Kostas, Panagopoulou, Aliki, Schofield, Gail, 2021. Social media reveals consistently disproportionate tourism pressure on a threatened marine vertebrate. Anim. Conserv. 24 (4), 568–579, URL https://doi.org/10.1111/acv.12656.

Parham, Jason, Crall, Jonathan, Stewart, Charles, Berger-Wolf, Tanya, Rubenstein, Daniel, 2017. Animal population censusing at scale with citizen science and photographic identification. In: AAAI Spring Symposium-Technical Report. URL https://cdn.aaai.org/ocs/15245/15245-68194-1-PB.pdf.

Pursley, K., 2020. In-water computer aided photo-ID of juvenile green turtles (*Chelonia mydas*) using flipper scales and affine transformations. Mar. Turt. Newsl. 161, 1–8.

Rees, ALan F., Carreras, Carlos, Broderick, Annette C., Margaritoulis, Dimitris, Stringell, Thomas B., Godley, Brendan J., 2017. Linking loggerhead locations: using multiple methods to determine the origin of sea turtles in feeding grounds. Mar. Biol. 164 (30), URL https://doi.org/10.1007/s00227-016-3055-z.

Rees, ALan F., Margaritoulis, Dimitris, Newman, Robert, Riggall, Thomas E., Tsaros, Paul, Zbinden, Judith A., Godley, Brendan J., 2013. Ecology of loggerhead marine turtles *Caretta caretta* in a neritic foraging habitat: movements, sex ratios and growth rates. Mar. Biol. 160, 519–529, URL https://doi.org/10.1007/s00227-012-2107-2.

Schofield, Gail, Katselidis, Kostas A., Dimopoulos, Panayotis, Pantis, John D., 2008. Investigating the viability of photo-identification as an objective tool to study endangered sea turtle populations. J. Exp. Mar. Biol. Ecol. 360, 103–108, URL https://doi.org/10.1016/j.jembe.2008.04.005.

Schofield, Gail, Klaassen, Marcel, Papafitsoros, Kostas, Lilley, Martin, Katselidis, Kostas A., Hays, Graeme C., 2020. Long-term photo-id and satellite tracking reveal sex-biased survival linked to movements in an endangered species. Ecology 11 (7), e03027, URL https://doi.org/10.1002/ecy.3027.

Shinoda, Risa, Shiohara, Kaede, 2024. PetFace: A large-scale dataset and benchmark for animal identification. In: European Conference on Computer Vision. Springer, pp. 19–36.

2022. Turtle recall: Conservation challenge. URL https://zindi.africa/competitions/turtle-recall-conservation-challenge.

Tyszkiewicz, Michał, Fua, Pascal, Trulls, Eduard, 2020. DISK: Learning local features with policy gradient. Adv. Neural Inf. Process. Syst. 33, 14254–14265.

Zhao, Xiaoming, Wu, Xingming, Chen, Weihai, Chen, Peter CY, Xu, Qingsong, Li, Zhengguo, 2023. Aliked: A lighter keypoint and descriptor extraction network via deformable transformation. IEEE Trans. Instrum. Meas. 72, 1–16.

Zimm, Roland, 2019. On the Development of the Turtle Scute Pattern and the Origins of Its Variation (Ph.D. thesis). University of Helsinki, Finland.