

Individual Minke Whale Recognition Using Deep Learning Convolutional Neural Networks

Dmitry A. Konovalov^{1*}, Suzanne Hillcoat^{1,2}, Genevieve Williams^{2,3},
R. Alastair Birtles^{1,2}, Naomi Gardiner², Matthew I. Curnock³

¹College of Business, Law & Governance, James Cook University, Townsville, Australia

²College of Science & Engineering, James Cook University, Townsville, Australia

³CSIRO Land and Water, James Cook University, Townsville, Australia

Email: *dmitry.konovalov@jcu.edu.au

How to cite this paper: Konovalov, D.A., Hillcoat, S., Williams, G., Birtles, R.A., Gardiner, N. and Curnock, M.I. (2018) Individual Minke Whale Recognition Using Deep Learning Convolutional Neural Networks. *Journal of Geoscience and Environment Protection*, 6, 25-36.
<https://doi.org/10.4236/gep.2018.65003>

Received: December 19, 2017

Accepted: May 20, 2018

Published: May 23, 2018

Abstract

The only known predictable aggregation of dwarf minke whales (*Balaenoptera acutorostrata* subsp.) occurs in the Australian offshore waters of the northern Great Barrier Reef in May-August each year. The identification of individual whales is required for research on the whales' population characteristics and for monitoring the potential impacts of tourism activities, including commercial swims with the whales. At present, it is not cost-effective for researchers to manually process and analyze the tens of thousands of underwater images collated after each observation/tourist season, and a large data base of historical non-identified imagery exists. This study reports the first proof of concept for recognizing individual dwarf minke whales using the Deep Learning Convolutional Neural Networks (CNN). The "off-the-shelf" Image net-trained VGG16 CNN was used as the feature-encoder of the per-pixel semantic segmentation Automatic Minke Whale Recognizer (AMWR). The most frequently photographed whale in a sample of 76 individual whales (MW1020) was identified in 179 images out of the total 1320 images provided. Training and image augmentation procedures were developed to compensate for the small number of available images. The trained AMWR achieved 93% prediction accuracy on the testing subset of 36 positive/MW1020 and 228 negative/not-MW1020 images, where each negative image contained at least one of the other 75 whales. Furthermore on the test subset, AMWR achieved 74% precision, 80% recall, and 4% false-positive rate, making the presented approach comparable or better to other state-of-the-art individual animal recognition results.

Keywords

Dwarf Minke Whales, Photo-Identification, Population Biology, Convolutional

1. Introduction

The dwarf minke whale (*Balaenoptera acutorostrata* subsp.) is the second smallest baleen whale, born at approximately 2m in length and growing to a maximum measured length of 7.8 m [1]. Dwarf minke whales are distributed throughout the southern hemisphere, including Antarctica, and were first acknowledged as a distinct form of minke in 1985 [1]. The only known predictable aggregation of dwarf minke whales occurs in the Australian offshore waters of the northern Great Barrier Reef (GBR) each year throughout the Australian winter months [3]. This aggregation supports a local swim-with-whales tourism industry [2] [3]. The predictable nature of this aggregation has also enabled dedicated research of dwarf minke whales, which has contributed to seminal work on dwarf minke whale biology [4], behavior [5], and assessment and management of swim-with-whales activities [2]. Outputs from this work have informed and shaped management policies and expanded knowledge of both the subspecies in general and, specifically, the interactions with the tourism industry. The uniqueness of this aggregation presents an opportunity to conduct research and improve the knowledge base for a poorly understood oceanic rorqual whale, as well as a responsibility to ensure that tourism activities are managed sustainably [2] [3] [5].

The identification of individual whales underpins much of the scientific research on dwarf minke whales and the monitoring of tourism activities. While in the GBR, these whales are highly inquisitive, readily approaching vessels and divers and often maintaining contact for prolonged periods [3] [5]. This behavior provides good opportunities for passengers aboard the swim-with tourism vessels to photograph dwarf minke whales. The whales' color patterns have been shown to remain stable over many years, and are sufficiently complex to allow for unequivocal identification of individuals [3] [6] [7]. The stability of these patterns and the regular, in-water access provided to researchers by tourism vessels has made the dwarf minke whale an ideal species for photo-identification (photo-ID) [6] [8].

Photo-ID is a simple, non-invasive technique widely used to study a range of biological and behavioral characteristics of wild animal populations. Ideal candidates for photo-ID are those with stable color patterns and/or other markings that are unique to each individual, so that individuals can be easily distinguished from each other and their identifiable markings remain the same over time. The automation of the photo-ID process is often highly specific to the required species, e.g. fin contour of great white sharks [9]. Due to its fundamental research role, photo-ID is an active research area for many species, e.g. green sea turtles [10], gorillas [11], and dolphins [12].

For minke whales, photo-ID has typically involved visual comparison of large numbers of photographs by trained researchers; thus, the process is time-intensive. Much of the imagery used for photo-identification of dwarf minke whales in recent years has come from tourists and crew aboard swim-with whales dive tourism vessels [8]. The quantity of this donated imagery has increased dramatically with the availability of low-cost digital underwater cameras and the resultant rise in popularity of these items among tourists [8]. Researchers are now obtaining tens of thousands of photographs and video clips each season. Consequently, it is no longer cost-effective for researchers to manually process and analyze such quantities of images, and a large database of historical non-identified imagery exists. In order to utilize the increasing quantity of imagery to address key biological and ecological knowledge gaps about these whales, automatic computer-vision based recognition software is required, and was the main focus of this study.

Over the last few years the Deep Learning Convolutional Neural Networks (CNNs) revolutionized the field of computer-vision image recognition [13]. For example, the Alex Net image classification CNN [14] won the Imagenet Large Scale Visual Recognition Challenge (ILSVRC) [15] in 2012, and since then all the ILSVRC13-ILSVRC17 winners used CNNs of various architectural configurations as their key features, e.g. [16]. It is customary to refer to such CNNs as been *trained-on-Imagenet*.

A typical Imagenet-trained CNN is setup to classify as many as 1000 different types of objects. Therefore, it is plausible to expect that such a CNN could distinguish at least 1000 different individual dwarf minke whales if it is trained or re-trained appropriately. This direct approach, however, has a number of limiting factors. First, millions of images are available in the Imagenet for training CNNs, which is presently not feasible for dwarf minke whales, where the number of images available for an individual whale may vary between one and several thousand. Second, typical Imagenet object categories are very different, e.g. differences in images for dogs and people, whereas all minke whales fit essentially the same category for the Imagenet (*i.e.* near-identical body shape, proportions and general color). Third, the output of a classification CNN is a single probability number for each available class, where category and class are used as equivalent terms in this study. Such probability prediction has limited value to a marine biologist, as it does not explain why/how CNN arrived at its prediction. This is known as the *black-box* perception and/or criticism of the classification CNNs. The black-box CNN prediction is unavoidable in studies where animals are identified by their “faces”, e.g. for gorillas [11], and identification uses facial geometrical proportions and is essentially the full face. Fortunately in the case of dwarf minke whales, they are currently identified by finely detailed color patterns and scars (**Figure 1**), which could be recognized and localized by CNN, and then confirmed by a trained researcher.

The *black-box* limitation of the classification CNNs has a natural solution



Figure 1. Example of individual minke whale distinct fin color pattern and scars.

when the CNNs are configured to perform semantic segmentation of images, where an image is segmented into per-pixel categories [17]. The output of segmentation CNNs is a per-pixel *heat-map* (also known as the probability or activation map) for each class. Therefore, a researcher could easily verify the CNN prediction by viewing the heat-map corresponding to the recognized individual whale (**Figure 2**). This approach was successfully validated in this proof of concept study by training a segmentation CNN to recognize a single whale within 1320 images of 76 different whales.

2. Materials and Methods

2.1. Dataset

The underwater imagery dataset used in this study consisted of 1320 digital photographs of dwarf minke whales (*Balaenoptera acutorostrata* subsp.). All images were sorted according to unique individual animals. In some cases only left or right sides of a whale was identified, without knowing if corresponding images belonged to the same whale or not. Where it was possible to match the left and right sides to the same whale, the related imagery was labeled accordingly and placed together in the same folder. As a result, the dataset identified 76 different whales. The identification process was extremely time consuming even for trained researchers as it required recording and cataloguing the color patterns and scars of 76 different whales, and/or reviewing any new image against at least 76 other whale images thus relying on researchers' memory to identify matches with any efficiency. The number of available images varied greatly between individuals; the MW1020 individual had the largest number of images (179), and several whales had only one image per individual.

2.2. Segmentation Neural Network

As described in the introduction, this study used a segmentation CNN rather than a classification CNN to recognize an individual minke whale and localize the recognized unique features. Specifically, the most accurate segmentation FCN-8s model from the Fully Convolutional Networks (FCN) [17] was selected due to the following considerations.

First, the FCN-8s model is based on the VGG16 CNN model [16], which was one of the top performers in the ILSVRC14 [15].

Second, this study used the Deep Learning python framework Keras [18] with Tensor Flow [19] as the processing backend. The Imagenet pre-trained VGG16 model was available within Keras [18], and the FCN-8s model had a number of publically available Keras-based implementations, e.g. [20]. For this study, FCN-8s version was recreated in Keras directly from the original Caffe source code of the FCN-8s model [21], and released to public domain [22].

Third, at the time of writing, the FCN-8s publication [17] had the largest numbers of citations among segmentation CNNs making it a widely accepted base-line model for semantic segmentation. Adopting this well-known FCN-8s model for this study was intended to make the presented method be reproduced and/or replicated more easily for additional/different minke whale images or for other animal species recognition studies.

In terms of the actual implementation, the FCN-8s model was built by reusing all VGG16 convolutional layers, which were loaded with the Imagenet-trained VGG16 weights available in Keras [18]. Such reuse of CNN weights is often referred to as the *knowledge transfer* [23]. VGG16 was designed to recognize 1000 classes of objects. Since this study was dealing with the maximum of 76 individual whales, the original VGG16/FCN-8s 4096 neurons were reduced to 1024 neurons when the last two dense (non-convolutional) VGG16 layers *fc1* and *fc2* were converted to their convolutional equivalents as per the FCN-8s model. This reduced the total FCN-8s size to approximately 160 MB when stored on disk, comparing to 540 MB for the original FCN-8s model with 4096 neurons in the *fc1* and *fc2* layers. The non-VGG16 convolutional layers were initialized by the uniform distribution as per [24]. *Sigmoid* activation [25] function was used in the last (*i.e.* prediction) layer.

2.3. Data Augmentation and Training Workflow

The adopted FCN-8s [17] segmentation model was a very high capacity neural network, which could overfit if it was presented with the same unchanged training images repeatedly. Therefore, the training images had to be augmented to prevent the FCN-8s model from memorizing the relatively small number of training images, and/or the trivial transient features such as ambient color hue or brightness. Furthermore and for the same reason of regularizing to avoid overfitting, the Imagenet-trained VGG16 convolutional weights were *frozen*, *i.e.* excluded from training.

Two image processing protocols were used. First, all available images were standardized by the following *imagescaling procedure* (ISP640). If a given image had H and W as height and width, respectively, then $L = \min(H, W)$ is the minimum of H and W , and the image was resized by scale $S = 640 / L$. This step scaled all images to have shortest sides be 640 pixels long, hence the abbreviation ISP640.

The second or *training augmentation protocol* (TAP480) was applied to the ISP640 processed images, where each image was:

- Randomly rotated in the range of $[-45, +45]$ degrees, where the input image was reflected to fill pixels outside the original boundary as required;
- Randomly resized in the scale range of $[0.75, 1.25]$, or by up to 25% zooming in or out;
- Randomly shifted in each color channel in the $[-25.5, 25.5]$ range, where 25.5 was the 10% of maximum color values 255;
- Randomly gamma shifted in the $[-25.5, 25.5]$ range, where all color channels values were shifted together;
- Randomly cropped to retain 480×480 pixels;
- Imagenet color mean values were subtracted as commonly done when working with the Imagenet-trained VGG16 model.

The following training workflow was adopted for this study. All available images were sequentially numbered and split into five approximately equal subsets. The first three subsets were used as a single *training* set, *i.e.* 60% of all available images. The fourth and the fifth subsets became the *validation* and *testing* sets, respectively. More precisely, the i^{th} image was allocated to *validation* or *test* if $(i+1)$ or i were multiple of 5, respectively, where all remaining images were assigned to the *training* set.

The training of FCN-8s was done in up to 100 cycles. In each cycle, TAP480 was further applied to the already ISP640-processed images. The training images were loaded into memory as a $X(N_i, M, M, C)$ tensor or a multidimensional matrix, where $N_i = 200$ was the number of images, $M = 480$ was the TAP480 cropping length, and where $C = 3$ was due to the three available color channels. The corresponding to the loaded training images were the *ground-truth* binary per-pixel masks, which were loaded as a *one-hot* encoded $Y(N_i, M, M, K)$ tensor, where $Y(i, m, l, k) = 1$ if the (m, l) pixel belonged to the k^{th} class in the i^{th} image and zero otherwise. The required number of classes K was $K = 1$ for the automatic whale locator and a single whale classifier, as described later on in this paper. The validation X_v and Y_v tensors were constructed in similar fashion.

The per-pixel binary cross-entropy loss function, e.g. p.231 of [25], was averaged as required and used as the training loss metric. Due to the available Graphical Processing Unit (GPU) memory limits, training was done in batches of only four images. Up to 16 training epochs were allowed per cycle, where one feed-forward and one back-propagation passes through all N_i -loaded image-mask pairs were considered to be one epoch. Training for a given cycle was aborted if the validation loss metric did not decrease after two epochs, this is commonly known as *early stopping*. Note that the early stopping was the only place where the validation images were used in training. In order to prevent the indirect overfitting of the validation images, they were augmented by TAP480 before each training cycle similar to the training set.

2.4. Minke Whale Locator

Being a segmentation model, the FCN-8s model required the ground-truth

per-pixel binary mask for each of the training and validation images. Therefore, the auxiliary goal of this study was to design the required workflow to be as scalable as possible for future larger training datasets. Creating the ground-truth per-pixel binary masks was clearly the least scalable component of this study, and required a scalable solution. This was solved by training an instance of FCN-8s to be the Minke Whale Locator (MWL).

To train MWL, 100 images were segmented by hand (including 50 of the MW1020 individual) to produce binary per-pixel ground-truth mask Y for each of the 100 images. Then MWL was trained as per preceding Section 2.2 with the following modifications. In addition to TAP480, images were flipped horizontally with 0.5 probability. The available 100 images were split 70 for training, and 30 for validation, where the rest of the not-segmented images were considered to be the testing set. The Keras version of the RMS prop optimizer was used with 10^{-4} learning rate, and 10^{-3} learning rate decay after each weights update, where RMS prop “divides the learning rate for a weight by a running average of the magnitudes of recent gradients for that weight” [26]. Once the per-pixel validation accuracy stopped improving (usually at around 95%), the Stochastic Gradient Descent (SGD) optimizer was used with 10^{-4} learning rate, 10^{-3} learning rate decay, 0.9 momentum, and enabled Nesterov momentum.

Trained MWL was applied to all available images to automatically generate one largest *rectangular* binary mask per ISP640 pre-processed image. Note that since MWL was fully convolutional, it was rebuilt to accommodate any required image dimensions, where one side was always 640 (due to ISP640) but the other side was varied. The mask generation was done as follows. For each image, the per-pixel prediction heat-map $Y_p(i, j)$ was converted to binary mask B via,

$$B(i, j) = 1, \quad Y_p(i, j) \geq 0.8, \quad (1)$$

where i and j were the row and column pixel location indices, respectively, and where the remaining mask values were set to zero, *i.e.* $B(i, j) = 0$, $Y_p(i, j) < 0.8$. The largest connected non-zero area was filled to complete its minimum-enclosing rectangle, and saved as the only non-zero values of the final binary mask.

2.5. Automatic Minke Whale Recognition

Similar to the preceding MWL model, an instance of the FCN-8s model was created for a required number of K individual whales to be the Automatic Minke Whale Recognition (AMWR) model. To train AMWR, the automatically created (by MWL) masks for the K whales were reviewed for correctness. Specifically, each MWL-generated rectangular mask was checked to make sure it enclosed correct whale if multiple whales were present in an image. Also, if the mask did not enclose the whole whale, the mask was verified to enclose all whales' features, which a biologist could use to identify that whale, *i.e.* fin coloration patterns and distinct scars. Note that in this study, the MWL model was nothing more than a convenience tool to automate ground-truth mask creation. There-

fore where available, the manually segmented masks were used instead of the corresponding MWL masks. MWL produced acceptable bounding boxes in more than 90% cases confirming it to be a viable tool for this project.

The AMWR was trained as per preceding Section 2.2 with the following modifications. For the K selected whales the *positive* ground-truth masks (manually or automatically MWL-segmented) were used. The training masks for the remaining $(76 - K)$ whales were automatically generated as *negative* or all-zeros masks, *i.e.* any of the K selected whales were missing in the remaining images. Then the training proceeded as per MWL but with added regularization weight decay set to 10^{-4} .

3. Results and Discussion

The largest number (179) of images was available for the individual whale MW1020 so it was used as the benchmark of possible accuracy for the utilized dataset and the AMWR model with $K = 1$. As per preceding Sections 2.3 and 2.4, 50 masks were segmented manually, and the rest of available MW1020 images (129) were segmented by MWL and quality-checked visually. The MW1020 training, validation and test sets contained 107, 36, and 36 images, respectively. The rest of other whale images (1141) were automatically labeled as *negative*, and split 60%-training, 20%-validation, and 20%-test. Because there were many more negative labels than positive, for each training cycle an equal number of images (100) were randomly selected from both negative and positive/MW1020 training images. Similarly, all available 36 MW1020 validation images were used with 36 randomly selected negative validation images, where a new random selection of 36 negative images was done before each training cycle. Also due to the highly unbalanced number of positive and negative examples, AMWR classifier was assessed via *precision*, *recall*, *fprate* (false-positive), in addition to the standard *accuracy* [9] [11] [27],

$$\textit{precision} = TP / (TP + FP), \quad \textit{recall} = TP / P, \quad \textit{fp rate} = FP / N \quad (2)$$

$$\textit{accuracy} = (TP + TN) / (P + N), \quad (3)$$

where TP , TN , FP and FN were the numbers of true-positive, true-negative, false-positive and false-negative predictions, respectively, and where P and N were the total numbers of positive (MW1020) and negative (non-MW1020-whale) images.

The main distinct advantage of a *per-pixel* classifier (rather than *per-image*) such as the presented AMWR, is the full control over how “conservative” or “liberal” [27] it could be configured. The highly conservative version was configured by accepting the prediction heat-map values only above 0.99, where the binary per-pixel predictions were set as $B(i, j) = 1$, $Y_p(i, j) \geq 0.99$ and zero otherwise. Furthermore, the largest connected prediction area was only accepted as a positive detection if its area was at least $64 \times 64 = 4096$ pixels, see example in **Figure 2**.



Figure 2. Example of AMWR per-pixel prediction for MW1020 individual. The pixels with the prediction heat-map values above 0.99 were illustrated by amplifying the corresponding image pixel intensities by factor of 1.5.

Table 1. Identification results for MW1020.

Prediction Score	Datasets		
	Train	Validation	Test
Accuracy	0.984	0.924	0.935
Precision	0.935	0.735	0.743
Recall	0.953	0.694	0.805
Fp rate	0.01	0.04	0.04

On the *test* subset, AMWR achieved 4% false-positive rate (**Table 1**). Low *fp rate* was viewed as essential to support a workflow where many thousands of unsorted images could be scanned for the known whales, and the number of “false-alarm” instances would remain feasible to be classified manually. AMWR’s *test* precision (74%) and recall (80%) results (last column of **Table 1**) were better than the corresponding state-of-the-art gorilla identification results [11] of approximately 60%. The AMWR’s test accuracy (93%) and precision (74%) were comparable to the 81% average precision achieved in the state-of-the-art great white shark identification results [9]. The *validation* and *test* prediction metrics were comparable (third and fourth columns in **Table 1**) supporting the achieved test values to be the expected benchmark/baseline values of the AMWR model in future similar circumstances/studies.

4. Conclusion

Due to the increasing abundance of underwater digital imagery, the manual identification of individual dwarf minke whales from images and videos has become cost-ineffective. It has become excessively time-consuming to manually check if an unsorted image contains a new whale or a known whale, e.g. from the 76 labeled whales of this study’s dataset. Considering that photo-identification of

dwarf minke whales represents one of the few methods available to address key knowledge gaps for this species' biology and life history, the application of automated recognition tools can potentially provide new scientific insights that would otherwise be inaccessible to scientists. The quantity of images for individual whales presented a theoretically challenging problem, where the number of available labeled images was too large for further manual labeling, but not large enough to apply Deep Learning classification CNNs. This study demonstrated how the Deep Learning per-pixel segmentation FCN-8s [17] CNN could be trained for an individual minke whale recognition from only 179 positive images. As much as possible the off-the-shelf pre-trained VGG16 [16] CNN was used to assist adoption and reproducibility of the results.

Acknowledgements

The authors are profoundly grateful for the contributions of passengers, crew and owners of the permitted swim-with-whales tourism vessels in the Great Barrier Reef who have helped to provide many of the minke whale images used in this study. We are also deeply grateful to the many Mike Whale Project Volunteers who have helped to sort our minke images. We are particularly indebted to our research colleagues associated with the Minke Whale Project who have facilitated our photo-identification work including especially Dr Susan Soltzick (who developed our main MWP Catalogue), Chrystie Watson, Tara Stephens, Liz Forrest, A/Prof Trina Myers, Dr Dianna Hardy, Prof Ian Atkinson and Kent Adams.

References

- [1] Best, P.B. (1985) External Characters of Southern Minke Whales and the Existence of a Diminutive Form. *Sci. Rep. Whales Res. Inst.*, **36**, 1-33. <http://www.icrwhale.org/pdf/SC0361-33.pdf>
- [2] Curnock, M.I., Birtles, R.A. and Valentine, P.S. (2013) Increased Use Levels, Effort, and Spatial Distribution of Tourists Swimming with Dwarf Minke Whales at the Great Barrier Reef. *Tourism in Marine Environments*, **9**, 5-17. <https://doi.org/10.3727/154427313X13659574649867>
- [3] Birtles, R.A., Arnold, P.W. and Dunstan, A. (2002) Commercial Swim Programs with Dwarf Minke Whales on the Northern Great Barrier Reef, Australia: Some Characteristics of the Encounters with Management Implications. *Australian Mammalogy*, **24**, 23-38. <https://doi.org/10.1071/AM02023>
- [4] Gedamke, J., Costa, D.P. and Dunstan, A. (2001) Localization and Visual Verification of a Complex Minkewhale Vocalization. *J. Acoust. Soc. Am.*, **109**, 3038-3047. <https://doi.org/10.1121/1.1371763>
- [5] Mangott, A.H., Birtles, R.A. and Marsh, H. (2011) Attraction of Dwarf Minke Whales (*Balaenoptera acutorostrata*) to Vessels and Swimmers in the Great Barrier Reef World Heritage Area—The Management Challenges of an Inquisitive Whale. *Journal of Ecotourism*, **10**, 64-76. <https://doi.org/10.1080/14724041003690468>
- [6] Arnold, P.W., Birtles, R.A., Dunstan, A., Lukoschek, V. and Matthews, M. (2005) Colour Patterns of the Dwarf Minke Whale *Balaenoptera acutorostrata* sensual to: Description, Cladistic Analysis and Taxonomic Implications. *Memoirs of the*

- Queensland Museum*, **51**, 277-307.
https://researchonline.jcu.edu.au/4935/1/4935_Arnold_et_al...2005.pdf
- [7] Arnold, P., Marsh, H. and Heinsohn, G. (1987) The Occurrence of Two Forms of Minke Whales in East Australian Waters with Description of External Characters and Skeleton of the Diminutive Form. *Sci. Rep. Whales Res. Inst.*, **38**, 1-46.
<http://www.icrwhale.org/pdf/SC0381-46.pdf>
- [8] Sobtzick, S. (2010) Dwarf Minke Whales in the Northern Great Barrier Reef and Implications for the Sustainable Management of the Swim-With Whales Industry. PhD Thesis, James Cook University.
<https://researchonline.jcu.edu.au/28199/1/28199-sobtzick-2010-thesis.pdf>
- [9] Hughes, B. and Burghardt, T. (2017) Automated Visual Fin Identification of Individual Great White Sharks. *International Journal of Computer Vision (IJCV)*, **122**, 542-557. <https://doi.org/10.1007/s11263-016-0961-y>
- [10] Carpentier, A.S., Jean, C., Barret, M., Chassagneux, A. and Ciccione, S. (2016) Stability of Facial Scale Patterns on Green Sea Turtles *Chelonia mydas* over Time: A Validation for the Use of a Photo-Identification Method. *Journal of Experimental Marine Biology and Ecology*, **476**, 15-21.
<https://doi.org/10.1016/j.jembe.2015.12.003>
- [11] Brust, C., Burghardt, T., Groeneweg, M., Kading, C., Kuhl, H.S., Manguette, M.L. and Denzler, J. (2017) Towards Automated Visual Monitoring of Individual Gorillas in the Wild. *The IEEE International Conference on Computer Vision (ICCV)*, 2820-2830.
http://openaccess.thecvf.com/content_ICCV_2017_workshops/papers/w41/Brust_Towards_Automated_Visual_ICCV_2017_paper.pdf
<https://doi.org/10.1109/ICCVW.2017.333>
- [12] Genov, T., Centrih, T., Wright, A.J. and Wu, G.-M. (2017) Novel Method for Identifying Individual Cetaceans Using Facial Features and Symmetry: A Test Case Using Dolphins. *Marine Mammal Science*. (In press)
<https://doi.org/10.1111/mms.12451>
- [13] LeCun, Y., Bengio, Y. and Hinton, G. (2015) Deep learning. *Nature*, **521**, 436-444.
<https://doi.org/10.1038/nature14539>
- [14] Krizhevsky, A., Sutskever, I. and Hinton, G.E. (2012) Imagenet Classification with Deep Convolutional Neural Networks. In: Pereira, F., Burges, C.J.C., Bottou, L. and Weinberger, K.Q., Eds., *Advances in Neural Information Processing Systems*, Vol. 25, Curran Associates, Inc., 1097-1105.
- [15] Russakovsky, O., Deng, J., Su, H., Krause, J., Satheesh, S., Ma, S., Huang, Z., Karpathy, A., Khosla, A., Bernstein, M., Berg, A.C. and Fei-Fei, L. (2015) Image Net Large Scale Visual Recognition Challenge. *International Journal of Computer Vision*, **115**, 211-252. <https://doi.org/10.1007/s11263-015-0816-y>
- [16] Simonyan, K. and Zisserman, A. (2014) Very Deep Convolutional Networks for Large-Scale Image Recognition. *CoRR*, abs/1409.1556.
<http://arxiv.org/abs/1409.1556>
- [17] Shelhamer, E., Long, J. and Darrell, T. (2017) Fully Convolutional Networks for Semantic Segmentation. *IEEE Transactions on Pattern Analysis and Machine Intelligence*, **39**, 640-651. <https://doi.org/10.1109/TPAMI.2016.2572683>
- [18] Chollet, F. (2015) Keras. <https://github.com/fchollet/keras>
- [19] Abadi, M., Agarwal, A., Barham, P., Brevdo, E., Chen, Z., Citro, C., Corrado, G.S., Davis, A., Dean, J., Devin, M., Ghemawat, S., Goodfellow, I., Harp, A., Irving, G., Isard, M., Jia, Y., Jozefowicz, R., Kaiser, L., Kudlur, M., Levenberg, J., Mane, D.,

- Monga, R., Moore, S., Murray, D., Olah, C., Schuster, M., Shlens, J., Steiner, B., Sutskever, I., Talwar, K., Tucker, P., Vanhoucke, V., Vasudevan, V., Viégas, F., Vinyals, O., Warden, P., Wattenberg, M., Wicke, M., Yu, Y. and Zheng, X. (2015) Tensor Flow: Large-Scale Machine Learning on Heterogeneous Systems. <http://tensorflow.org>
- [20] Ju, J. (2017) Keras-FCN. <https://github.com/jihongju/keras-fcn>
- [21] Shelhamer, E., Long, J. and Darrell, T. (2017) FCN. <https://github.com/shelhamer/fcn.berkeleyvision.org>
- [22] Konovalov, D.A. (2017) Keras-FCN-8s. https://github.com/dmitryako/keras_fcn_8s
- [23] Oquab, M., Bottou, L., Laptev, I. and Sivic, J. (2014) Learning and Transferring Mid-Level Image Representations Using Convolutional Neural Networks. *IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*. <https://doi.org/10.1109/CVPR.2014.222>
- [24] Glorot, X. and Bengio, Y. (2010) Understanding the Difficulty of Training Deep Feedforward Neural Networks. In: Teh, Y.W. and Titterton, M., Eds., *Proceedings of the Thirteenth International Conference on Artificial Intelligence and Statistics*, Proceedings of Machine Learning Research (PMLR), Chia Laguna Resort, Sardinia, Italy, Vol. 9, 249-256. <http://proceedings.mlr.press/v9/glorot10a.html>
- [25] Bishop, C.M. (1995) *Neural Networks for Pattern Recognition*. Oxford University Press.
- [26] Hinton, G., Srivastava, N. and Swersky, K. Overview of Mini-Batch Gradient Descent. http://www.cs.toronto.edu/~tijmen/csc321/slides/lecture_slides_lec6.pdf
- [27] Fawcett, T. (2006) An Introduction to ROC Analysis. *Pattern Recognition Letters*, 27, 861-874. <https://doi.org/10.1016/j.patrec.2005.10.010>