

RESEARCH ARTICLE

# Assessing Rotation-Invariant Feature Classification for Automated Wildebeest Population Counts

Colin J. Torney<sup>1\*</sup>, Andrew P. Dobson<sup>2</sup>, Felix Borner<sup>3</sup>, David J. Lloyd-Jones<sup>4</sup>, David Moyer<sup>5</sup>, Honori T. Maliti<sup>6</sup>, Machoke Mwita<sup>6</sup>, Howard Fredrick<sup>7</sup>, Markus Borner<sup>8</sup>, J. Grant C. Hopcraft<sup>8</sup>

**1** Centre for Mathematics and the Environment, University of Exeter, Penryn Campus, Penryn, Cornwall, United Kingdom, **2** Department of Ecology & Evolutionary Biology, Princeton University, Princeton, New Jersey, United States of America, **3** Frankfurt Zoological Society, Serengeti National Park, Seronera, Tanzania, **4** Independent researcher, P.O. Box 1272, Iringa, Tanzania, **5** Integrated Research Center, The Field Museum of Natural History, 1400 S. Lake Shore Drive, Chicago, United States of America, **6** Tanzania Wildlife Research Institute, P.O.Box 661, Arusha, Tanzania, **7** Tanzania Conservation Resource Centre, Arusha, Tanzania, **8** Institute of Biodiversity, Animal Health and Comparative Medicine; Boyd Orr Centre for Population and Ecosystem Health, University of Glasgow, Glasgow, United Kingdom

\* [colin.j.torney@gmail.com](mailto:colin.j.torney@gmail.com)



OPEN ACCESS

**Citation:** Torney CJ, Dobson AP, Borner F, Lloyd-Jones DJ, Moyer D, Maliti HT, et al. (2016) Assessing Rotation-Invariant Feature Classification for Automated Wildebeest Population Counts. PLoS ONE 11(5): e0156342. doi:10.1371/journal.pone.0156342

**Editor:** Emmanuel Serrano Ferron, Universidade de Aveiro, PORTUGAL

**Received:** February 29, 2016

**Accepted:** May 12, 2016

**Published:** May 26, 2016

**Copyright:** © 2016 Torney et al. This is an open access article distributed under the terms of the [Creative Commons Attribution License](https://creativecommons.org/licenses/by/4.0/), which permits unrestricted use, distribution, and reproduction in any medium, provided the original author and source are credited.

**Data Availability Statement:** Source code is available at <https://github.com/ctorney/wildCount>. Image count data is available as Supporting Information. Images were provided by the Frankfurt Zoological Society.

**Funding:** CJT is supported by a Complex Systems Scholar Award from the James S. McDonnell Foundation. JGCH is supported by a Lord Kelvin Adam Smith Fellowship, funding from the British Ecological Society and the European Union's Horizon 2020 research and innovation programme under grant agreement No 641918 AfricanBioServices. The

## Abstract

Accurate and on-demand animal population counts are the holy grail for wildlife conservation organizations throughout the world because they enable fast and responsive adaptive management policies. While the collection of image data from camera traps, satellites, and manned or unmanned aircraft has advanced significantly, the detection and identification of animals within images remains a major bottleneck since counting is primarily conducted by dedicated enumerators or citizen scientists. Recent developments in the field of computer vision suggest a potential resolution to this issue through the use of rotation-invariant object descriptors combined with machine learning algorithms. Here we implement an algorithm to detect and count wildebeest from aerial images collected in the Serengeti National Park in 2009 as part of the biennial wildebeest count. We find that the per image error rates are greater than, but comparable to, two separate human counts. For the total count, the algorithm is more accurate than both manual counts, suggesting that human counters have a tendency to systematically over or under count images. While the accuracy of the algorithm is not yet at an acceptable level for fully automatic counts, our results show this method is a promising avenue for further research and we highlight specific areas where future research should focus in order to develop fast and accurate enumeration of aerial count data. If combined with a bespoke image collection protocol, this approach may yield a fully automated wildebeest count in the near future.

funders had no role in study design, data collection and analysis, decision to publish, or preparation of the manuscript.

**Competing Interests:** The authors have declared the following interests: This study was partly funded by the British Ecological Society. There are no patents, products in development, or marketed products to declare. This does not alter the authors' adherence to all the PLOS ONE policies on sharing data and materials, as detailed online in the guide for authors.

## Introduction

Aerial surveys, in which the abundance of a population is estimated by flying transects over its habitat and counting the number of animals within a given sampling strip, are an essential tool for assessing wildlife population numbers [1, 2]. Many species are monitored in this way, including birds [3–5], land mammals [6–9], and aquatic fauna [10, 11]. While in-air counts are still used (i.e. animals are enumerated as they are encountered by observers), a common approach, especially with aggregated species living in high densities, is to employ aerial photography and then later count animals within images. The second stage of this process is frequently a labour-intensive procedure [12] that requires highly-skilled counters.

Automating the process of counting animals in images would therefore relieve a significant burden on governmental and non-governmental conservation organizations. Repeated measures of the population size over time allows managers to not only develop accurate estimates of the true population size, but it also enables the estimation of critical parameters about the population such as rates of recruitment, mortality, immigration and emmigration. These diagnostic parameters provide an early warning indicator of a population's health and are core metrics of any adaptive management system. Therefore, increasing the accuracy and the processing speed of a population count enables managers to access critical data and implement preemptive management strategies at an early stage, rather than waiting months for the results to be counted. Furthermore, an automated counting system could increase the frequency between consecutive population counts and thereby increase the temporal resolution of trends.

Achieving automated animal counts has been the subject of extensive research [13–17]. This research forms part of the rapidly evolving field of machine learning and computer vision [18]. Applications of these techniques are diverse and recent advances include the accurate detection of faces [19], facial expressions [20], pedestrians [21], and handwritten text [22]. In the context of ecology and conservation, machine learning has been deployed to classify species based on vocalisations [23], to identify behavioural states [24], and to track and identify moving animals [25]. However the most significant application has been in the automation of animal census methods, either through direct enumeration of animals [13, 14], or through computer-aided mark recapture methods based on automatic identification of individuals [26–28].

In this work, we evaluate the performance of a recently proposed method for the classification of objects [29]. The method is based on the popular histogram of oriented gradients technique [21] but has the distinct advantage of extracting only rotationally invariant features; thus making it suitable for aerial survey images in which animals may be oriented in any direction. We apply the method to the complete set of survey images taken during the 2009 Serengeti National Park wildebeest count. The wildebeest count is performed every 2 to 3 years and involves flying transects at an altitude of 350–400ft above ground. The aircraft travels at a speed of 120–180kph (subject to wind speed and direction), with images taken every 10 seconds from a camera mounted through the floor of the aircraft [30]. The result is approximately 2000 images that take 3 weeks for a single individual to count. To test whether the method proposed by [29] is able to automate the counting of wildebeest we implemented the algorithm, automatically counted the 2009 images, then compared the performance of the method to the manual totals. By testing the method on this dataset we are able to comprehensively evaluate its performance in an applied setting on a task of genuine ecological importance.

## Materials and Methods

### Rotation-invariant image classifiers

For completeness we include here a brief description of the method employed to extract invariant features from images. This is based on [29] and we refer interested parties to that work for a more complete description of their method.

The histogram of oriented gradients (HOG) technique [21] is a popular method that uses the distribution of gradients within regions of images to classify objects. Liu et al. [29] modified this approach so that instead of using a discrete grid, HOG cells are treated as continuous functions that may be approximated using Fourier series. The advantage of this approach is that the extracted features of the image are constant even if the underlying object within the image rotates.

As in [29], to process an image we first construct a matrix of gradients in complex form from the grayscale image  $I$  using a finite-difference scheme

$$\mathcal{G} = \left( \frac{\partial}{\partial x} + i \frac{\partial}{\partial y} \right) I. \tag{1}$$

Hence, each element of  $\mathcal{G}$  denotes the gradient at the corresponding image pixel in the form  $\Delta x + i\Delta y$ . In polar coordinates this may be written as  $re^{i\theta}$ . If we consider each element of  $\mathcal{G}$  as an individual cell [21] then the distribution of gradients is effectively a Dirac delta function, centred at  $\theta$ ,

$$H(\phi) = r\delta(\theta - \phi). \tag{2}$$

Performing a Fourier series expansion of a Dirac delta function leads to

$$H(\phi) = \sum_{m=-\infty}^{\infty} F_m e^{-2\pi i m \phi} = \sum_{m=-\infty}^{\infty} r e^{2\pi i m (\theta - \phi)} \tag{3}$$


Truncating this series at some maximum mode means we are left with a sequence of complex valued coefficients which represent the Fourier transform of the image gradient. The gradient at each pixel is therefore encoded by a sequence of Fourier coefficients and the full transformed image is stored in a 3-dimensional complex array, representing  $x$  and  $y$  coordinates and the modes of the Fourier transform. We denote the 2-d array of mode  $m$  coefficients as  $\mathcal{F}_m$ .

Next we introduce the Fourier basis functions  $U_{j,k}$  shown in Fig 1. By performing a convolution between a basis function  $U_{j,k}$  and a Fourier gradient field  $\mathcal{F}_m$  we obtain a Fourier HOG feature

$$X_{k,m} = U_{j,k} * \mathcal{F}_m \tag{4}$$

which encodes information about the image gradients in the region covered by the basis function. These radially symmetric basis functions act in a manner equivalent to the cells of the original HOG method.

If the original image  $I$  is rotated then each of the complex valued features  $X_{k,m}$  will also be altered, i.e. in this form they are not invariant to rotations. However due to the shift property of Fourier analysis, rotations of the original image can be mapped to multiplications of the Fourier coefficients. A rotation of the original image by an angle  $\alpha$  will result in the movement of pixels to another location and a rotation in the orientation of the gradients. These two effects can be mapped to the HOG feature by firstly rotating the Fourier transform of the gradient field  $\mathcal{F}_m$  by  $\alpha$ , and secondly by rotating the basis functions  $U_{j,k}$  by  $-\alpha$ .



**Fig 1. The basis functions,  $U_{j,k}$ , for performing convolutions are constructed from Fourier modes on concentric circles.** The parameter  $j$  determines the radial distance from the centre of the object, while  $k$  is the wavenumber. The images show the real and imaginary part of the basis function.

doi:10.1371/journal.pone.0156342.g001

If  $X_{k,m}$  is the original HOG feature, and  $X'_{k,m}$  is the corresponding feature calculated after the image has been rotated by  $\alpha$ , then

$$\begin{aligned} X'_{k,m} &= e^{-i\alpha k} U_{j,k} * e^{i\alpha m} \mathcal{F}_m \\ &= e^{i\alpha(m-k)} (U_{j,k} * \mathcal{F}_m) \\ &= e^{i\alpha(m-k)} X_{k,m} \end{aligned} \tag{5}$$

From this equation we can see that if  $m = k$  then image rotations have no impact on the descriptor and it is rotation invariant. Also, by taking the product of two descriptors,  $X_{k_1, m_1}$  and  $X_{k_2, m_2}$ , a composite descriptor is formed,

$$X_{k_1, m_1} X_{k_2, m_2} = e^{i\alpha(m_1 - k_1 + m_2 - k_2)} (U_{j, k_1} * \mathcal{F}_{m_1}) (U_{j, k_2} * \mathcal{F}_{m_2}) \tag{6}$$

Again we note that the composite descriptor remains constant for all angles of rotation angle  $\alpha$  if  $(m_1 - k_1 + m_2 - k_2) = 0$ . We may therefore construct rotational invariant features of the image from the features defined by Eq 4, by firstly using features for which  $m - k = 0$ , and secondly by taking the product of any two features for which  $(m_1 - k_1 + m_2 - k_2) = 0$ .

### Implementation

The 2009 wildebeest count resulted in 2,018 images taken with a Nikon D2X 35mm camera shooting 4288x2848 pixel JPG images. Three separate counts of the aerial images were performed. Firstly, two independent counts were performed simultaneously by two different individuals. A third count was then performed by three individuals for images where there was a discrepancy between initial counts. This final count is taken to be the correct count for our comparison metrics. To evaluate the Fourier HOG method, the full 2009 image set was counted using machine learning software. The adaboost algorithm [31] was employed with a decision tree underlying classifier. Training images were drawn from 100 images taken from the 2012 survey.

The code was written in Python 2.7 ([www.python.org](http://www.python.org)) using OpenCV [32] for image operations and the sci-kit learn package [33] for classification. The classification code was parallelized using PyCUDA [34] and the code was run on an NVIDIA GeForce GT 630 graphics card. All code is based on open source libraries and is available here <https://github.com/ctorney/wildCount>

An iterative process was employed to train the classifier based on the 2012 images. First a set of sample images was generated from the 2012 image set by manually locating wildebeest. Next the classifier was trained on this small training data set and several further images were automatically counted. The results from this count were manually checked and corrected then used to create a larger training data set of 3000 positive samples and 3000 negative samples.

The trained classifier was then applied to the 2009 image set. Images were converted to gray-scale then they were scanned for regions above a threshold level of local contrast. Regions that were uniform were discarded. Next each pixel that was in a non-uniform region was taken to be the centre of an object to be classified and rotation-invariant Fourier HOG features were extracted. Each pixel was then classified either as a wildebeest or not, then contiguous blocks of pixels were grouped and counted as a single individual.

## Results

To assess the accuracy of the method, total wildebeest counts are compared to the multiple counts performed by human counters. When using 3000 training examples for each class (positive or negative) we find good agreement between the automated totals and the manual counts as shown in Table 1. In Fig 2 performance of the algorithm is assessed against the final manual count and the two prior counts are shown for comparison. We note that while the automated total is more accurate than either initial counts, the root mean square error per image is greater. This metric, calculated as  $\sqrt{\frac{1}{N} \sum D_i^2}$  where  $D_i$  is the difference in the count for image  $i$  and the summation is taken over all  $N$  images, reveals that the greater overall accuracy is due to the lack of any systematic bias in the machine learning algorithm. The mean error (defined as  $\frac{1}{N} \sum D_i$ ) in Table 1 shows that each first pass manual count had either a positive or negative bias, whereas the algorithm displayed little systematic bias and was therefore able to obtain a more accurate total count, despite displaying a greater RMS error.

To measure the precision and recall of the method, 100 images were randomly selected and the number of true and false positives, and true and false negatives, were recorded. These results are shown in Table 2.

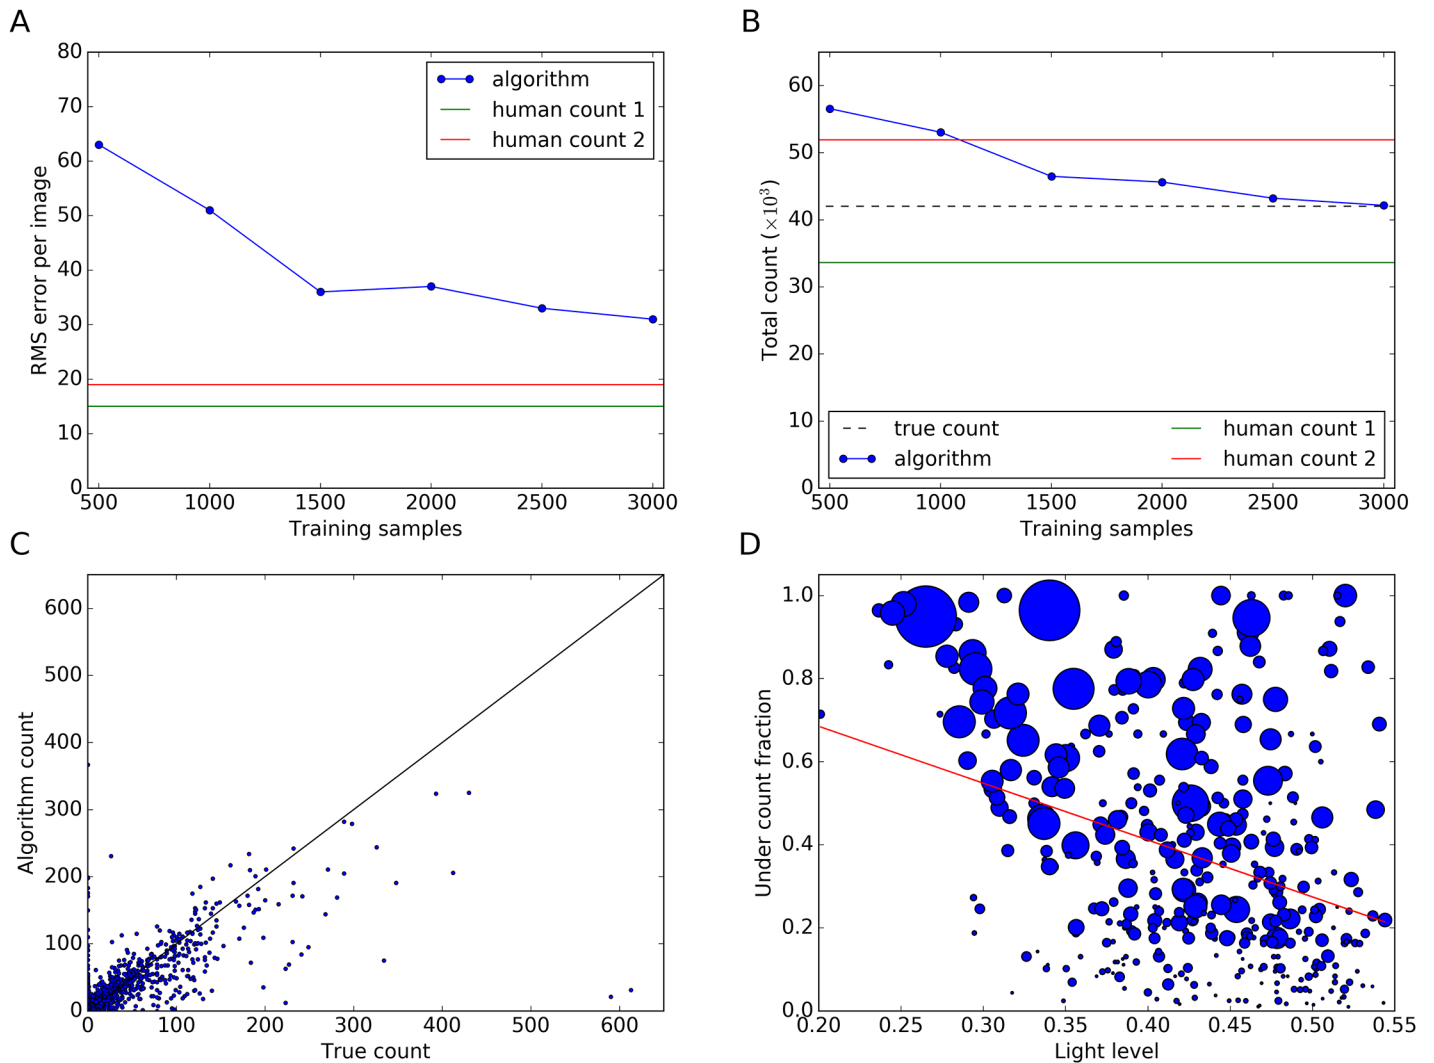
## Discussion

The advance in new technologies such as earth-orbiting satellites [14] or unmanned aerial vehicles [9], has led to a rapid increase in high-resolution, easily accessible image data. To keep pace with this progress, computational tools are required to automate image processing and ensure that these vast amounts of data are transformed into useful information. One area

**Table 1. Comparison of counts between manual and automated methods.**

	Total	Mean (per image)	Coefficient of variation	RMS error	Mean error
First manual count	33644	16.67	2.61	15.64	-4.14
Second manual count	51918	25.73	2.21	19.75	4.91
Final manual count	42007	20.82	2.32	-	-
Automated count	42147	20.89	1.83	31.70	0.07

doi:10.1371/journal.pone.0156342.t001



**Fig 2. Comparing the performance of automated and manual counters.** (A) Root mean square error of counts. The correct count for each image is assumed to be the third and final count. Average per image error is shown for the algorithm (blue line) as a function of the number of training samples from the 2012 survey that were used. For comparison, per image error is shown for each of the first pass human counts (red, green lines). (B) Total wildebeest counted within the image set. The final count is shown by the dashed line. The algorithm (blue line) outperforms both human counters in attaining a closer estimate to the true value. This is because the algorithm exhibits no systematic tendency to over or under count. It should be noted that 3000 was the maximum number of training samples available, and it is plausible that the automated total count will drop below the true count before it asymptotes. (C) Individual image errors. The black line is the  $y = x$  line for reference. While average per image errors are comparable between automated and human counters, the algorithm makes large errors in a small subset of images. Images that contain many false negatives tend to be darker than the training samples, while false positives occur when there is a lot of structure in the landscape. (D) A comparison of image light levels and under counting. A linear regression shows a significant negative relationship between image light level (average of value component of HSV image) and the amount of under counting ( $\beta_1 = -1.37$ ,  $R^2 = 0.12$ ). The under count fraction is calculated as  $1 - \frac{\text{algorithm count}}{\text{true count}}$  and images for which algorithmcount > truecount are excluded. Point sizes are proportional to the absolute value of the under count of wildebeest in the image.

doi:10.1371/journal.pone.0156342.g002

where modern computer vision techniques have the potential to significantly improve current practices is in the automated detection of animals within aerial count images. We have implemented a recent object classification method [29] which uses rotation-invariant features and is therefore suitable for use with these types of images.


By testing the method against multiple manual counts we find that its performance is comparable to a first-pass human count. The algorithm has a greater per image error rate, but is overall more accurate than two individual human counters. This is due to a lack of any

**Table 2. Confusion matrix.** As the accuracy based on the total count does not indicate precision or recall, performance metrics were recorded for a random subset of 100 images. Negative totals are based on the number of non-overlapping regions within each image that are approximately equal in area to a single wildebeest. From these results: precision =  $\frac{TP}{TP+FP} = 74.15\%$ , recall =  $\frac{TP}{TP+FN} = 85.83\%$ .

		Algorithm	
		Positive	Negative
Actual	Positive	1423	235
	Negative	496	1.2m

doi:10.1371/journal.pone.0156342.t002

systematic bias in errors, with landscape features leading to high rates of false positives, and low light conditions leading to false negatives (see Fig 3 for example images). Currently the algorithm is unlikely to outperform multiple human counts, either by trained professionals or through a citizen science approach that averages many counts by non-specialist individuals (such as the snap-shot Serengeti project operated through the zooniverse platform). However a



**Fig 3. Example images.** From top: Correctly detected wildebeest; Pattern and structure in the landscape frequently lead to false positives; The method is able to distinguish between different species; Species such as zebra, that have distinct body shapes are frequently not identified as wildebeest; The ability to distinguish between species is dependent on sufficient training examples, here the algorithm has misidentified a flock of juvenile ostrich as wildebeest.

doi:10.1371/journal.pone.0156342.g003

combination of automated and manual counting would represent an ideal application of the method in its current form, either as a first-pass count or as a method to assess the performance of citizen scientists.

A significant promising aspect of the method is that it appears able to identify and differentiate between animal species. Although we were unable to quantify the performance of the algorithm in this regard with the current data set, these preliminary results show that common species such as zebra may be distinguished from wildebeest by the algorithm. In future we intend to further test this performance with training and testing data sets of multiple species.

A further avenue for future research will involve acquiring 3-dimensional information about the scene. As the features used by the classification method are based on the shape of the object, the error rate will be greatly reduced by obtaining the 3-dimensional structure of the object. This could be achieved through range imaging techniques, such as structure from motion [35] or LIDAR [36]. An alternative approach to increase accuracy would be to include a near-infrared thermal band that could differentiate between endothermic animals and the background. Both thermal and 3-d information could be used in combination with image gradients to enhance the accuracy of the method.

## Supporting Information

**S1 File. Count data.** Human and automated counts for each image for 500, 1000, 1500, 2000, 2500, and 3000 training samples.  
(ZIP)

## Acknowledgments

CJT is supported by a James S. McDonnell Foundation Studying Complex Systems Scholar Award. JGCH is supported by a Lord Kelvin Adam Smith Fellowship, funding from the British Ecological Society and the European Union's Horizon 2020 research and innovation programme under grant agreement No 641918 AfricanBioServices.

## Author Contributions

Conceived and designed the experiments: CJT APD JGCH. Performed the experiments: CJT. Analyzed the data: CJT APD JGCH. Contributed reagents/materials/analysis tools: FB DJLJ DM HTM MM MB HF. Wrote the paper: CJT JGCH.

## References

1. Caughley G, Sinclair R, Scott-Kemmis D. Experiments in aerial survey. *The Journal of Wildlife Management*. 1976;p. 290–300. doi: [10.2307/3800428](https://doi.org/10.2307/3800428)
2. Caughley G. Sampling techniques for aerial censuses. *Aerial surveys of fauna populations Australian National Parks and Wildlife Service, Canberra, Australia*. 1979;p. 15–23.
3. Briggs KT, Tyler WB, Lewis DB. Comparison of ship and aerial surveys of birds at sea. *The Journal of wildlife management*. 1985;p. 405–411. doi: [10.2307/3801542](https://doi.org/10.2307/3801542)
4. Morton SR, Brennan K, Armstrong M. Distribution and abundance of magpie geese, *Anseranas semipalmata*, in the Alligator Rivers Region, Northern Territory. *Australian Journal of Ecology*. 1990; 15 (3):307–320. doi: [10.1111/j.1442-9993.1990.tb01035.x](https://doi.org/10.1111/j.1442-9993.1990.tb01035.x)
5. Mosbech A, Boertmann D. Distribution, abundance and reaction to aerial surveys of post-breeding king eiders (*Somateria spectabilis*) in western Greenland. *Arctic*. 1999;p. 188–203.
6. Goddard J. Aerial census of black rhinoceros using stratified random sampling. *African Journal of Ecology*. 1969; 7(1):105–114. doi: [10.1111/j.1365-2028.1969.tb01199.x](https://doi.org/10.1111/j.1365-2028.1969.tb01199.x)
7. Evans CD, Troyer WA, Lensink CJ. Aerial census of moose by quadrat sampling units. *The Journal of Wildlife Management*. 1966;p. 767–776. doi: [10.2307/3798283](https://doi.org/10.2307/3798283)



8. Norton-Griffiths M. Counting the Serengeti migratory wildebeest using two-stage sampling. *African Journal of Ecology*. 1973; 11(2):135–149. doi: [10.1111/j.1365-2028.1973.tb00079.x](https://doi.org/10.1111/j.1365-2028.1973.tb00079.x)
9. Vermeulen C, Lejeune P, Lisein J, Sawadogo P, Bouché P. Unmanned aerial survey of elephants. *PLoS one*. 2013; 8(2):1–7. doi: [10.1371/journal.pone.0054700](https://doi.org/10.1371/journal.pone.0054700)
10. Marsh H, Sinclair D. Correcting for visibility bias in strip transect aerial surveys of aquatic fauna. *The Journal of Wildlife Management*. 1989;p. 1017–1024. doi: [10.2307/3809604](https://doi.org/10.2307/3809604)
11. Øritsland T, Øien N. Aerial surveys of harp and hooded seal pups in the Greenland Sea pack-ice. *Developments in Marine Biology*. 1995; 4:77–87. doi: [10.1016/S0163-6995\(06\)80012-0](https://doi.org/10.1016/S0163-6995(06)80012-0)
12. Russell J, Couturier S, Sopuck L, Ovaska K. Post-calving photo-census of the Rivière George caribou herd in July 1993. *Rangifer*. 1996; 16(4):319–330. doi: [10.7557/2.16.4.1273](https://doi.org/10.7557/2.16.4.1273)
13. Sirmacek B, Wegmann M, Cross J, Hopcraft P, Dech S. Automatic population counts for improved wildlife management using aerial photography. In: *Managing Resources of a Limited Planet, Sixth Biennial Meeting*, Leipzig, Germany; 2012.
14. Yang Z, Wang T, Skidmore AK, de Leeuw J, Said MY, Freer J. Spotting East African Mammals in Open Savannah from Space. *PLoS one*. 2014; 9(12):e115989. doi: [10.1371/journal.pone.0115989](https://doi.org/10.1371/journal.pone.0115989) PMID: [25551561](https://pubmed.ncbi.nlm.nih.gov/25551561/)
15. McNeill S, Barton K, Lyver P, Pairman D. Semi-automated penguin counting from digital aerial photographs. In: *Geoscience and Remote Sensing Symposium (IGARSS), 2011 IEEE International*. IEEE; 2011. p. 4312–4315.
16. Laliberte AS, Ripple WJ. Automated wildlife counts from remotely sensed imagery. *Wildlife Society Bulletin*. 2003;p. 362–371.
17. Bajzak D, Piatt JF. Computer-aided procedure for counting waterfowl on aerial photographs. *Wildlife Society Bulletin*. 1990;p. 125–129.
18. Chen Ch, Pau LF, Wang PSp. *Handbook of pattern recognition and computer vision*. vol. 27. World Scientific; 2010.
19. Hsu RL, Abdel-Mottaleb M, Jain AK. Face detection in color images. *Pattern Analysis and Machine Intelligence, IEEE Transactions on*. 2002; 24(5):696–706. doi: [10.1109/34.1000242](https://doi.org/10.1109/34.1000242)
20. Michel P, El Kaliouby R. Real time facial expression recognition in video using support vector machines. In: *Proceedings of the 5th international conference on Multimodal interfaces*. ACM; 2003. p. 258–264.
21. Dalal N, Triggs B. Histograms of oriented gradients for human detection. In: *Computer Vision and Pattern Recognition, 2005. CVPR 2005. IEEE Computer Society Conference on*. vol. 1. IEEE; 2005. p. 886–893.
22. Xu L, Krzyżak A, Suen CY. Methods of combining multiple classifiers and their applications to handwriting recognition. *Systems, Man and Cybernetics, IEEE Transactions on*. 1992; 22(3):418–435. doi: [10.1109/21.155943](https://doi.org/10.1109/21.155943)
23. Acevedo MA, Corrada-Bravo CJ, Corrada-Bravo H, Villanueva-Rivera LJ, Aide TM. Automated classification of bird and amphibian calls using machine learning: A comparison of methods. *Ecological Informatics*. 2009; 4(4):206–214. doi: [10.1016/j.ecoinf.2009.06.005](https://doi.org/10.1016/j.ecoinf.2009.06.005)
24. Bidder OR, Campbell HA, Gómez-Laich A, Urgé P, Walker J, Cai Y, et al. Love thy neighbour: automatic animal behavioural classification of acceleration data using the k-nearest neighbour algorithm. *PLoS one*. 2014; 9(2):e88609. doi: [10.1371/journal.pone.0088609](https://doi.org/10.1371/journal.pone.0088609) PMID: [24586354](https://pubmed.ncbi.nlm.nih.gov/24586354/)
25. Pérez-Escudero A, Vicente-Page J, Hinz RC, Arganda S, de Polavieja GG. idTracker: tracking individuals in a group by automatic identification of unmarked animals. *Nature methods*. 2014; 11(7):743–748. doi: [10.1038/nmeth.2994](https://doi.org/10.1038/nmeth.2994) PMID: [24880877](https://pubmed.ncbi.nlm.nih.gov/24880877/)
26. Sherley RB, Burghardt T, Barham PJ, Campbell N, Cuthill IC. Spotting the difference: towards fully-automated population monitoring of African penguins *Spheniscus demersus*. *Endangered Species Research*. 2010; 11(2):101–111. doi: [10.3354/esr00267](https://doi.org/10.3354/esr00267)
27. Bolger DT, Morrison TA, Vance B, Lee D, Farid H. A computer-assisted system for photographic mark-recapture analysis. *Methods in Ecology and Evolution*. 2012; 3(5):813–822. doi: [10.1111/j.2041-210X.2012.00212.x](https://doi.org/10.1111/j.2041-210X.2012.00212.x)
28. Kühl HS, Burghardt T. Animal biometrics: quantifying and detecting phenotypic appearance. *Trends in ecology & evolution*. 2013; 28(7):432–441. doi: [10.1016/j.tree.2013.02.013](https://doi.org/10.1016/j.tree.2013.02.013)
29. Liu K, Skibbe H, Schmidt T, Blein T, Palme K, Brox T, et al. Rotation-invariant HOG descriptors using fourier analysis in polar and spherical coordinates. *International Journal of Computer Vision*. 2014; 106(3):342–364. doi: [10.1007/s11263-013-0634-z](https://doi.org/10.1007/s11263-013-0634-z)
30. Sinclair ARE, Arcese P. *Serengeti II: dynamics, management, and conservation of an ecosystem*. vol. 2. University of Chicago Press; 1995.

31. Freund Y, Schapire RE. A decision-theoretic generalization of on-line learning and an application to boosting. In: Computational learning theory. Springer; 1995. p. 23–37.
32. Bradski G. Dr Dobb's Journal of Software Tools. 2000;.
33. Pedregosa F, Varoquaux G, Gramfort A, Michel V, Thirion B, Grisel O, et al. Scikit-learn: Machine Learning in Python. *Journal of Machine Learning Research*. 2011; 12:2825–2830.
34. Klöckner A, Pinto N, Lee Y, Catanzaro B, Ivanov P, Fasih A. PyCUDA and PyOpenCL: A Scripting-Based Approach to GPU Run-Time Code Generation. *Parallel Computing*. 2012; 38(3):157–174. doi: [10.1016/j.parco.2011.09.001](https://doi.org/10.1016/j.parco.2011.09.001)
35. Ullman S. The interpretation of structure from motion. *Proceedings of the Royal Society of London B: Biological Sciences*. 1979; 203(1153):405–426. doi: [10.1098/rspb.1979.0006](https://doi.org/10.1098/rspb.1979.0006) PMID: [34162](https://pubmed.ncbi.nlm.nih.gov/34162/)
36. Lin Y, Hyyppä J, Jaakkola A. Mini-UAV-borne LIDAR for fine-scale mapping. *Geoscience and Remote Sensing Letters, IEEE*. 2011; 8(3):426–430. doi: [10.1109/LGRS.2010.2079913](https://doi.org/10.1109/LGRS.2010.2079913)