Estimating a seal population using automated counts of drone photographs

Samuel Hawkins, Allister Lundberg, Aiden Lundberg, Guy Oliver, Richard Condit

Samuel Hawkins, Allister Lundberg, Aiden Lundberg, Guy Oliver, Richard Condit, "Estimating a seal population using automated counts of drone photographs," Proc. SPIE 12675, Applications of Machine Learning 2023, 1267505 (4 October 2023); doi: 10.1117/12.2676277

Event: SPIE Optical Engineering + Applications, 2023, San Diego, California, United States

Estimating a Seal Population Using Automated Counts of Drone Photographs

Samuel Hawkins*[a], Allister Lundberg*[b], Aiden Lundberg[c], Guy Oliver[d], and Richard Condit[d]

[a]Bradley University, Peoria, Illinois, United States
[b]LivRobo, Los Altos Hills, California, United States
[c]Consultant, Los Altos Hills, California, United States
[d]Univ. of California, Santa Cruz, Santa Cruz, United States

ABSTRACT

The preservation of Northern elephant seals, with a current population exceeding 250,000, has been due to successful conservation efforts. Down to as few as 100 seals in the 1890s, accurate population monitoring remains crucial. Counting seals from the ground, especially on remote islands where most breed, is difficult and dangerous, and manually counting from aerial photos is time-consuming and error-prone. This research proposes an automated method of counting elephant seals using machine learning. Drone images were collected from Año Nuevo Reserve, California, US, during the 2022 and 2023 winter breeding seasons. The system automatically created orthophotos from drone images, made predictions using a single-stage object detection model from tiles, detected and removed duplicate predictions, classified the seals into males, females, and pups, and mapped the predictions back to the orthophotos as labeled bounding boxes. An optional active learning component also allowed human reviewers to make corrections in a UI and edits could be automatically turned into new training data to improve future surveys. In an examination of the largest aggregation on the Mainland, the model found 99.4% of females, 97.8% of males, and 97.0% of pups. The whole pipeline, including model training, can be run on a laptop, and it can be utilized in remote field sites where there is no internet access.

Keywords: Classification, Counting, Machine Learning, Object Detection

1. INTRODUCTION

Accurately counting elephant seals is important. Northern elephant seals were nearly exterminated by 19th-century hunting seeking oil in their blubber. During the 1890s, there were probably fewer than 100 remaining [1]. A variety of conservation efforts enabled their population to rebound to over 250,000 individuals, filling their former range. Demonstrating this recovery has required repeated counts of breeding females and pups across their entire distribution in Mexico and the US. Some counts are done from the ground, but many of the colonies are best accessed by aerial flights. The flights are expensive, and reaching every colony yearly is not feasible. For example, the most recent global assessment omitted Mexico and relied on rough estimates to fill in the entire population [2]. Our goal is to improve methods for counting females during the breeding season and thus foster understanding of the entire population in case future changes cause population decline.

One tool that could make counting more effective is automated software that recognizes, counts, and maps individual elephant seals from large numbers of aerial photographs. Unmanned aerial drones are in use in surveys of ungulates in Africa, seabirds in Alaska, and penguins in Antarctica, and software that allows rapid counts is in use [3–5]. Our goal here is to develop a software system that can 1) count seals from aerial images using object detection, 2) classify those seals as male, female, or pup, and 3) provide the location of each seal for verification. This paper describes the current system and tests carried out at the Año Nuevo colony in January 2022 and 2023.

*Samuel Hawkins and Allister Lundberg contributed equally to this paper.
Send correspondence to Samuel Hawkins
E-mail: shawkins2@fsmail.bradley.edu

Applications of Machine Learning 2023, edited by Michael E. Zelinski, Tarek M. Taha,
Barath Narayanan Narayanan, Proc. of SPIE Vol. 12675, 1267505

1.1 Background

Elephant seals, both the northern and southern species (*Mirounga angustirostris, M. leonina*) are like most pinnipeds and breed on land while feeding at sea. While breeding, females and pups are readily observable on coastal beaches, where they collect in a few large aggregations. More than 98% of the northern species' population is found at only seven locations along the Pacific coast of Mexico and the United States [2,6]. Moreover, 87% of the breeding females are onshore at the same time at the peak of the winter breeding season in late January [7].

Most elephant seal breeding colonies are on islands where predation pressure is typically low, but access for biologists is difficult and expensive. For the northern species, these include Guadalupe Island in Mexico and three of the Channel Islands off southern California [2]. Since the 1970s, however, mainland sites have been colonized, including beaches at Año Nuevo. This site has been the basis of most research on northern elephant seals, due to easy access and proximity to the University of California at Santa Cruz [8–11]. The research has included regular counts of the number of elephant seals onshore, all carried out by standing on small dunes near the female groups to complete a full head count. A single count of females in late January can be used to estimate the number of pups born in a season, and a single count of pups in late February or early March leads to an estimate of the number surviving [7]. There is a long-term record of the size of the colony and the survival rate of pups [12]. Since the establishment of the colony in 1961, it grew to 2,500 pups born annually by the mid-1990s, then declined slightly to 1,800 each year over the past decade [12].

Population counts of the large island colonies of elephant seals off southern California and Baja California are much more difficult because access is difficult and aggregations of over 1,000 females cannot be counted from the ground. They must be counted from aerial photographs. The last complete aerial survey in 2014 [2] estimated the entire 2010 population in California and Mexico at 210,000, including 48,000 pups produced.

Breeding elephant seals are sexually dimorphic: males are larger (length 4-5 m, mass 1300-2270 kg vs. 2.5-3.5 m, mass 400-770 kg in females) and have darker pelage compared to females. Pups are much smaller (length 0.9-1.4 m) and even darker than males (see Fig. 1). Pups at birth weigh 30 kg and weigh up to 130 kg when weaned. A human observer can readily differentiate adult males, females, and pups, in aerial photographs. Outside the breeding season, juveniles are much harder to separate, but the most important count is pup production during the winter breeding season.

Figure 1. A male elephant seal, 10 female seals, and 7 pups demonstrate the sexual dimorphism found between males and females and show the rapid growth of pups differing in age by only a few weeks. The shape of pups in photographs can vary from semicircular to linear, whereas adults are typically linear in form.

2. METHODS

2.1 Study Site

The elephant seal colony at Año Nuevo (37°06 N, 122°20 W) is composed of an island and a mainland component (Figure 2). The island is 0.8 km offshore and has a single beach with breeding seals. The mainland portion is spread along 2.6 km of beaches, but the seals are concentrated at several locations. The island group and the largest mainland group were used to develop and test an automated system for counting seals. Both groups are on flat sand 1-2 m above the high tide line; the mainland group has 800-900 females and their pups each January, and the island group 400-600.

Figure 2. Año Nuevo Island and Mainland. The elephant seals' largest aggregations are on the triangular beach facing the mainland on the Island and at Año Nuevo Point, which is directly opposite the Island on the mainland. Photo by Daniel P. Costa with permission.

2.2 Limitation of ground counts

Elephant seals can be counted from elevated dunes or cliffs near the animals, and this has been the principal source of information since the 1960s [7, 12, 13]. The greatest limitation is the access to sites near the seals that are high enough above the flat beaches to allow visibility. The large mainland group and the island group at Año Nuevo both have poor visibility. Since the 1970s, aerial photographs have been used to verify ground counts [2], but airplanes are costly and seldom available.

2.3 Drone flights

Flights above the seal groups were flown using a DJI Mavic 2 Zoom quadcopter equipped with a 12 mp sensor and a 24-48 mm zoom lens (SZ DJI Technology Co., Ltd., Shenzhen, China; cost $1,000 USD). Standardized flight plans were designed using the Litchi app (VC Technology Ltd. London, England) for surveys of both the Island and the mainland.

Prior to gaining approval to use drones to monitor the seals within Año Nuevo Reserve, a series of experiments were conducted to determine whether drone flights caused a disturbance to the seals. When the drone was 38 or 45 m above the beaches, seals did not react to their presence. Also, no response was observed from the bird colonies on Año Nuevo Island.

Drones were flown in the early morning when the wind is lightest and before the Park opens to the public. Flight speed was generally 20 km/h and photographs were taken every 2 seconds with the lens set at 24 mm, giving a 74^o field of view. The flight path and sampling frequency were designed to ensure considerable overlap of the photographic images.

2.4 Focus on female count

For this study, the goal was an accurate count of adult females because the female count in late January provides the best estimate of the number of pups born that year, and pup production is the best proxy for overall population size and health [7]. Accuracy of the pup count in late January is less important because not all pups are born yet and others die earlier. Crucial to a counting system, however, is that pups are not miscounted as females.

2.5 Image data

Flights over the Año Nuevo colony have been carried out regularly since 2017. For the current study, 35 images from the mainland and 25 from the island were collected in January 2022 and used for training, and 47 images from the mainland and 23 from the island were collected in January 2023 and used for testing.

2.6 Overview of the Eight-Step Prediction Pipeline

The task of counting elephant seals from drone images can be broken down into eight steps, each of which is an interesting avenue for research in its own right. First, the images from the drone were stitched together into an orthophoto. The orthophoto helps ensure the seals are not counted twice as there is considerable overlap between the adjacent raw drone images. Second, the orthophoto was cut into tiles which were used for predictions. This step was performed for computational efficiency. The tiles were cut with enough overlap to guarantee each seal was fully visible in at least one tile and not touching any of the borders of the tile. Third, object detection was performed on each tile. Fourth, predictions that touch the edges were removed to ensure no partial seals were counted. Fifth, each tiles' predictions were mapped back to the orthophoto. Sixth, overlapping duplicate predictions on the orthophoto were removed. Seventh, predictions were classified as male, female, or pup. Eighth, further duplicate predictions were removed by identifying any in the same category and which had an overlapping center of mass. See Table 2 for software libraries used.

2.7 Details of the Eight-Step Prediction Pipeline

2.7.1 Step 1: Merge raw drone images into an orthophoto

Different stitching methods were compared. A high percentage of overlapping images supported better fusing and improved the final orthophoto quality. OpenDroneMap produced the highest quality orthophoto using SIFT [14]. As an alternative approach, OpenCV with ORB [15] was tested. Similar problems of stitching arise when counting other closely packed animals such as penguins [16]. See Table 3 for details of the settings used for OpenDroneMap.

2.7.2 Step 2: Cut overlapping tiles from the orthophoto

The orthophoto was cut into tiles of 640x640 pixels. This size was chosen to be large enough to include up to a few dozen seals but small enough to allow fast model predictions without down-sampling before feeding the image to the single-stage detection model. A 300-pixel overlap between each tile, nearly 50% of each tile's size, was used. The reason behind the wide overlap was to accommodate instances where a large male seal spanned the boundaries of tiles. This overlap ensured that these seals will always be fully encapsulated within at least one tile without contacting any of its boundaries, thereby ensuring at least one prediction captured the full size of every seal.

2.7.3 Step 3: Predict seals on each tile

Following the tiling process, predictions were made on each tile using a single-stage detection model, YOLO-NAS large [17]. A confidence threshold of .4 worked best to predict as many of the pups as possible as they were the most difficult to detect. However, the low threshold also created multiple predictions on the same seal so steps 4, 6, and 8 were used to remove duplicate predictions.

2.7.4 Step 4: Remove predictions at tile edges

To enhance the reliability of predictions, all predictions located along the edges of each tile were eliminated, retaining only seals fully depicted within the image. This step eliminated problems caused by predictions of seals along the edges. The removal process worked by identifying predictions whose bounding boxes were within 20 pixels of the tile edges; these predictions were removed. The 20-pixel margin was determined empirically to ensure all predictions of seals along the edges were removed, including instances where one seal obscured another near the edge.

2.7.5 Step 5: Map predictions from tiles back to the orthophoto

The individual predictions from the tiles were then mapped back onto the orthophoto. As part of the tiling process, each tile's filename incorporated information on its exact location in the orthophoto. This information allowed for precise remapping of predictions from the tile back to the orthophoto without homography.

2.7.6 Step 6: Remove duplicate predictions

Following the mapping process, a de-duplication process was used to address instances where seals were predicted multiple times due to the overlap in tiles. Duplicates were defined as instances where the intersection of two bounding boxes divided by the union was greater than 80%, a threshold determined empirically. One from each pair of duplicate predictions was removed.

2.7.7 Step 7: Classify predictions to male, female, or pup

After de-duplication process, each seal on the orthophoto was predicted once. The next phase involved classifying seals into one of three categories: male, female, or pup. A weighted Gaussian Mixture Model (GMM) was used to take advantage of the size disparity between the categories. The diagonal distance of the predicted bounding box produced the most reliable seal classifications. Initial weighting of the mixture model approximated a typical distribution of the seals: about 5% males, 55% females, and 40% pups.

However, the classifications from the GMM could not be used directly due to the wide variance in pup sizes. Sometimes the GMM produced an estimate of variance for the pups where females were incorrectly classified as pups. To address this issue, hard thresholds were derived from the GMM predictions. Once the derived thresholds were determined for a specific orthophoto based on the GMM, anything less than the pup threshold was classified as a pup, anything between the pup and male threshold was classified as a female, and anything exceeding the male threshold was classified as a male. The GMM was used to dynamically find the derived thresholds for each orthophoto since the drone height varied between flights.

2.7.8 Step 8: Remove duplicates based on the center of mass on the orthophoto

Post-classification, the predictions were condensed into smaller bounding boxes of 10x10 pixels centered on the middle of each prediction bounding box. This step identified the center of mass of each seal and consequently helped to identify duplicate predictions of the same seal missed in step 6 (see Figure 3 for an example). The low confidence score of 0.4 sometimes resulted in a male or a female being predicted twice but with less than an 80% overlap (intersection over union). By finding the center of mass for the 10x10 bounding box, additional duplicates were removed.

Figure 3. 1. Raw predictions mapped to the orthophoto after IoU greater than 80% removed 2. Predictions mapped to the center of mass 10x10 bounding box 3. Predictions are removed if they overlap on the center of mass bounding box and belong to the same category (both male or both female)

2.8 Training the Object Detection Model

A pre-trained YOLO-NAS large model was fine-tuned on seal images. Tiles of 640x640 pixels with 300-pixel overlap created from orthophotos from the year 2022 resulted in hundreds of training tile images. Orthophotos from the year 2023 were then used to test the pipeline on unseen data.

As part of training process, several image augmentation steps were tested, altering orientation, size, hue, saturation, affine and mosaic transformations. Hyperparameters for each were tuned (see Table 4 for specific image augmentation and model hyperparameters). The training was conducted both on a lower-end GPU, T4 16GB, and also on a 32GB 2019 MacBook Pro laptop.

Alternative software for object detection might offer another way to improve predictions of the pups. GLIP [18] was also tested before settling on YOLO-NAS. GLIP located 10% of seals in orthophotos without any fine-tuning, unlike the pre-trained YOLO V5 or YOLO-NAS models, which did not find seals until after fine-tuning. Due to the initial positive results with GLIP, three different versions of GLIP were tried but these failed to generalize as well as YOLO-NAS on unseen orthophoto images after fine-tuning. GLIP had lower accuracy than YOLO-NAS when making predictions on blurry images of seals. To get around this limitation, experiments were made with GLIP to make predictions directly on the drone images and then map these back onto the orthophoto using homography, but this resulted in misplaced and duplicate predictions. Practically speaking, YOLO-NAS was easier to work with, trained faster, and generalized better after fine-tuning than GLIP.

2.9 Optional Active Learning: Training Data from Human Review

Predictions from the prediction pipeline for 2023 images were saved to a JSON file and then visualized in the LabelMe app. Figure 4 shows predictions in the user interface of the LabelMe software. A human reviewer added missing labels, removed false positives, and corrected classifications as needed, and the corrections were used to automatically create new training data, enabling the model to learn after each survey and adapt when there is a shift in the data. This step is not reflected in the presented results but illustrates how a prediction can be improved and aid in future counts.

Figure 4. Using LabelMe as a UI enables human reviewers to view and optionally edit predictions of males, females, and pups. Males are marked in green, females in red and pups in yellow. Edits are automatically turned into new training data. Images are of a zoomed out view on the mainland and zoomed in view on the island.

2.10 Human Reviewer Accuracy Validation

The same drone images that were analyzed by the computer models were independently counted by direct observation. The stitched images were loaded in the software ImageJ and each female, pup, and male were marked on separate layers (known as a region of interest, or ROI, in ImageJ). One observer completed a count, then a second observer reviewed and corrected until every animal was marked. ImageJ was then used to export the coordinates (in pixels) of all the marked animals. Cases were noted where there was doubt on whether an animal was male or female, and whether or not a dark mark was a pup. Adults were not confused with pups. Human counts were then overlaid on computer counts using pixel coordinates identified in imageJ and the centroid of bounding boxes predicted by the computer model. Discrepancies were counted and checked against the original images. The total number of discrepancies was used as an estimate of the error in the human count. Both the mainland and island images were fully counted before the results from the computer model were checked so that human and computer counts were fully independent.

3. RESULTS

Table 1 summarizes the results. The model counts are presented without active learning (further human adjustment and retraining of the model) to illustrate the expected performance when counts are not known. Human uncertainty is described in section 2.10. For both island and mainland images, the model count of adult females was nearly identical to the independent count made by human observers. Discrepancies amounted to fewer than 2% of the total number, and the computer's estimate was within the range of human uncertainty. Figure 5 shows an image of the results, Figure 5A shows the orthophoto generated by stitching and Figure 5B shows the human counts in blue and computer counts in red.

There were far fewer males in the seal groups, and the computer estimate was close to the human count on the mainland (Table 1). For pup counts, the computer count on the mainland was 3% below the human count but was much lower on the island, where the computer only counted 55% of the pups. Human observers detected large numbers of animals which were unquestionably pups that the computer did not identify.

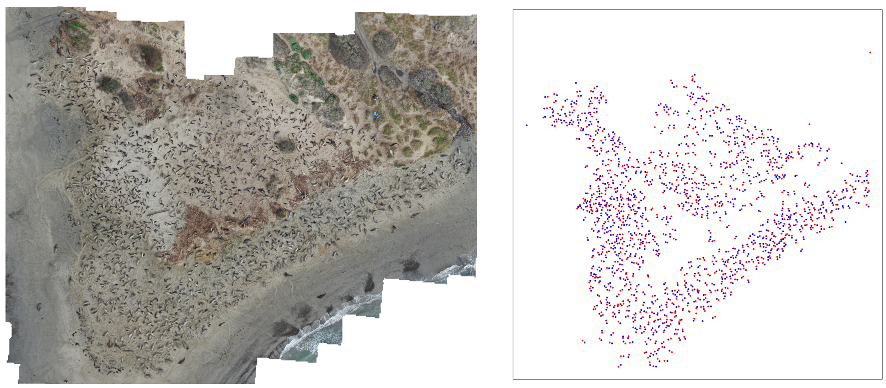

Figure 5. A) Stitched image of drone photographs of the largest aggregation of seals at Año Nuevo Point on the mainland. B) Map showing locations of all females in the mainland image as identified by the computer model (red) and by the independent human count (blue). There are 905 blue points and 903 red, with 900 of those corresponding closely; 3 red points do not have a matching blue point, and 5 blue points have no matching red point. Those differences are within counting error, meaning that approximately 10 of the females were indistinct and could not be conclusively identified. The human count was done by marking seals' heads, while the computer drew a bounding box around each, and red points are the centroid of those boxes; thus, matching red and blue points are separated by the distance from animals' heads to mid-body.

Table 1. Accuracy for Males, Females, and Pups Per Location

Location	Category	Human Count	Human Uncertainty	Model Count	Model Accuracy
Mainland	Female	908	5	903	99.4
Island	Female	458	5	453	98.9
Mainland	Pup	762	10	739	97.0
Island	Pup	242	10	133	55.0
Mainland	Male	44	3	45	97.8
Island	Male	23	3	15	65.2

Counts of three categories of elephant seals at two locations of the Año Nuevo colony, California, were compared between our computer model and independent human counts of the same images. Accuracy is the model count relative to the human count.

4. DISCUSSION

The model's predictions were highly effective, estimating female numbers closely and not mistaking pups for females. In its current state, the model thus is useful, and will greatly reduce the manual effort required in future surveys.

One of the biggest limitations of the model was undercounting pups on the island. Pups on the island image were often hidden in sharp shadows caused by the early morning hour. Also, animal density was higher on the island compared to the mainland. Future surveys with further reviewer feedback and active learning should improve the pup predictions. An avenue for future work would be to use the location of the females to increase the likelihood of a pup being detected in that area. Missing pups were often next to, or partially obscured by their mothers, so it should be possible to base pup predictions on female locations.

Aerial surveys of animals and plants from drone images have many applications. Stitching many raw images into a single orthophoto and then making predictions on the tiles is a well-established and effective method. This research adds to this approach by finding hard-to-detect objects by utilizing a low confidence threshold when making predictions and using specific heuristics to remove duplicate and erroneous predictions that result from using such a low confidence threshold. As a result, the model was able to find female seals on the mainland and the island with high precision.

This research offers two main contributions. First, a pipeline that can accurately find hard-to-detect animals such as seals from aerial photos. It is designed to be easily modifiable and could be used to survey other animals and can run with limited computing resources for training and prediction. Second, using active learning as an optional step. The model can continue to learn and adjust automatically after each survey using the surveyors' edits without requiring any machine learning expertise.

The pipeline developed is expected to greatly reduce the manual effort when doing future surveys at Año Nuevo State Park. It can also potentially be used to survey other locations of larger seal populations, such as the Channel Islands.

This pipeline can also be easily applied to detect other large animals, both terrestrial and marine, trees, or man-made objects. Moreover, it would be a straightforward step to produce precise maps of the objects detected and counted.

APPENDIX A. SUPPLEMENTARY TABLES

Table 2. Software Libraries Used

Library	Usage
OpenDroneMap	Orthophoto creation
Python Imaging Library (PIL), NumPy	Image manipulation
YOLO-NAS	Object detection
Hyperopt	Hyperparameter tuning

Table 3. OpenDroneMap settings

Parameter	Value
orthophoto-resolution	.75
cropping	0
feature-type	sift
feature-quality	high
matcher-type	flann

Table 4. Important Training and Prediction Hyperparameters

Parameter	Value
batch_size	8
warmup_initial_lr	.000001
lr_warmup_epochs	3
initial_lr	.000084
optimizer	Adam
optimizer weight_decay	.00021
nms_threshold	.49
training optimization metric	mAP@0.50
prediction confidence threshold	.4

ACKNOWLEDGMENTS

Thank you Patrick Robinson for being the drone pilot and photo archivist. Thank you Jasmine Espinoza for marking all seals in ImageJ. Thank you to Bradley University students Daniel Ferenc, Matthew Gaby, Paris Geis, Alec Hartman, Jackson Smith, Evan Stuit, Ryan Weil, Zechariah Burrus, Broc Edson, Himal Patel, Josh Patino, George Tang, Claire Tesluk, and Bridget Whitacre for your work on previous versions of the software. Thank you UC Santa Cruz's Año Nuevo Reserve and Professor Dan Costa. The research was approved by the UC Santa Cruz Institutional Animal Care and Use Committee #Costd2009 and the National Marine Fisheries Service marine mammal research permit #23188. Drone flights were authorized by the University of California Center of Excellence for Unmanned Aircraft Systems Safety, the Monterey Bay National Marine Sanctuary, and Anño Nuevo State Park. Funding for the drone flights was provided by the Packard Ocean Science and Technology Endowment and UC Santa Cruz's Año Nuevo Reserve.

REFERENCES

[1] Bartholomew, G. A. and Hubbs, C. L., "Winter population of pinnipeds about Guadalupe, San Benito, and Cedros Islands, Baja California," *Journal of Mammalogy* **33**, 160–171 (1952).

[2] Lowry, M. S., Condit, R., Hatfield, B., Allen, S. G., Berger, R., Laake, J., Morris, P. A., Boeuf, B. J. L., and Reiter, J., "Abundance, distribution, and growth of the northern elephant seal, *Mirounga angustirostris*, in the United States in 2010," *Aquatic Mammals* **40**(1), 20–31 (2014).

[3] Weinstein, B. G., "A computer vision for animal ecology," *Journal of Animal Ecology* **87**(3), 533–545 (2018).

[4] Kellenberger, B., Veen, T., Folmer, E., and Tuia, D., "21 000 birds in 4.5 h: efficient large-scale seabird detection with machine learning," *Remote Sensing in Ecology and Conservation* **7**(3), 445–460 (2021).

[5] Eikelboom, J. A., Wind, J., van de Ven, E., Kenana, L. M., Schroder, B., de Knegt, H. J., van Langevelde, F., and Prins, H. H., "Improving the precision and accuracy of animal population estimates with aerial image object detection," *Methods in Ecology and Evolution* **10**(11), 1875–1887 (2019).

[6] Stewart, B. S., Yochem, P. K., Le Boeuf, B. J., Huber, H. R., DeLong, R. L., Jameson, R. J., Sydeman, W. J., and Allen, S. G., "Population recovery and status of the northern elephant seal, *Mirounga angustirostris*," in [*Elephant Seals: Population Ecology, Behavior, and Physiology*], Le Boeuf, B. J. and Laws, R. M., eds., ch. 2, 29–48, University of California Press (1994).

[7] Condit, R., Le Boeuf, B. J., Morris, P. A., and Sylvan, M., "Estimating population size in asynchronous aggregations: a Bayesian approach and test with elephant seal censuses," *Marine Mammal Science* **23**(4), 834–855 (2007).

[8] Le Boeuf, B. J. and Peterson, R. S., "Social status and mating activity in elephant seals," *Science* **163**, 91–93 (Jan. 1969).

[9] Le Boeuf, B. J. and Panken, K. J., "Elephant seals breeding on the mainland in California," *Proceedings of the California Academy of Sciences* **41**, 267–280 (1977).

[10] Beltran, R. S., Yuen, A. L., Condit, R., Robinson, P. W., Czapanskiy, M. F., Crocker, D. E., and Costa, D. P., "Elephant seals time their long-distance migrations using a map sense," *Current Biology* **32**, R141–R157 (2022).

[11] Costa, D. P., Le Boeuf, B. J., Huntley, A. C., and Ortiz, C. L., "The energetics of lactation in the northern elephant seal, *Mirounga angustirostris*," *Journal of Zoology* **209**, 21–33 (1986).

[12] Le Boeuf, B., Condit, R., Morris, P., and Reiter, J., "The northern elephant seal rookery at Año Nuevo: a case study in colonization," *Aquatic Mammals* **37**, 486–501 (2011).

[13] Condit, R., Allen, S. G., Costa, D. P., Codde, S., Goley, P. D., Le Boeuf, B. J., Lowry, M. S., and Morris, P., "Estimating population size when individuals are asynchronous: a model illustrated with northern elephant seal breeding colonies," *PLoS ONE* **17**(1), e0262214 (2022).

[14] Lowe, D. G., "Distinctive image features from scale-invariant keypoints," *International Journal of Computer Vision* **60**, 91–110 (2004).

[15] Rublee, E., Rabaud, V., Konolige, K., and Bradski, G., "Orb: An efficient alternative to sift or surf," *2011 International conference on computer vision* , 2564–2571 (2011).

[16] Liu, Y., Shah, V., Borowicz, A., Wethington, M., Strycker, N., Forrest, S., Lynch, H., and Singh, H., "Towards efficient machine learning methods for penguin counting in unmanned aerial system imagery," in [*2020 IEEE/OES Autonomous Underwater Vehicles Symposium (AUV)*], 1–7, IEEE (2020).

[17] Aharon, S., Louis-Dupont, Ofri Masad, Yurkova, K., Lotem Fridman, Lkdci, Khvedchenya, E., Rubin, R., Bagrov, N., Tymchenko, B., Keren, T., Zhilko, A., and Eran-Deci, "Super-gradients." https://zenodo.org/record/7789328 (2021).

[18] Li, L. H., Zhang, P., Zhang, H., Yang, J., Li, C., Zhong, Y., Wang, L., Yuan, L., Zhang, L., Hwang, J.-N., et al., "Grounded language-image pre-training," in [*Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*], 10965–10975 (2022).