Automated Deep Learning Analysis of Purple Martin Videos Depicting Incubation and Provisioning

Heather M. Williams
Evolution, Ecology & Behavior
State University of New York at Buffalo
Buffalo, New York
hw49@buffalo.edu

L. Shawn Matott
Center for Computational Research
State University of New York at Buffalo
Buffalo, New York
lsmatott@buffalo.edu

Robert L. DeLeon
Center for Computational Research
State University of New York at Buffalo
Buffalo, New York
rldeleon@buffalo.edu

ABSTRACT

Deep learning models have been developed to automatically analyze video clips of purple martin nesting behavior. Two separate models have been constructed, one for incubation and one for provisioning. The incubation model is a simple two class model that analyzes the videos to determine if an adult purple martin is incubating the eggs/young nestlings or not. The model is a Keras/Tensor Flow convolutional neural network (CNN) trained with 12 thousand still images and achieves a validation data set accuracy of 99.5% percent on still images. A comparison of the results of the automated video analysis with sample validation videos viewed manually shows good agreement; the model approaches human accuracy. Some conclusions from the incubation analysis will be discussed. The provisioning analysis requires a much more complex 3 class model which must distinguish between zero, one parent or both parents on the nest. With training sets including 26 thousand images the CNN model demonstrates a validation set accuracy of 99% on the still images. However, the actual the video analysis presents difficulties. Several different CNN models have been tried but results were similar. The best results to date on analyzing the videos for provisioning events have been 88% accuracy with 10% false positives. A discussion of the conclusions from the provisioning model and model analysis will be presented.

CCS CONCEPTS

• **Computing methodologies → Artificial intelligence**; **Computer vision**;

KEYWORDS

Artificial Intelligence, convolutional neural networks, Purple Martin, Incubation, Provisioning

ACM Reference Format:

Heather M. Williams, L. Shawn Matott, and Robert L. DeLeon. 2019. Automated Deep Learning Analysis of Purple Martin Videos Depicting Incubation and Provisioning. In *Proceedings of Practice and Experience in Advanced Research Computing (PEARC '19).* ACM, New York, NY, USA, 7 pages. https://doi.org/10.1145/1122445.1122456

PEARC '19 , July 28- Aug 1, 2019, Chicago, IL, USA

ACM ISBN 978-1-4503-9999-9/18/06…$15.00
https://doi.org/10.1145/1122445.1122456

1 INTRODUCTION

Purple martins (*Progne subis*) are insectivorous migratory birds in the swallow family. The species spends its winter months in South America, principally Brazil, where it roosts in forest and forested river islands [10], but its breeding range expands throughout much of the Eastern United States, with small pockets of breeding populations in the Pacific Northwest and Southwestern deserts. In common with other aerial insectivores, purple martins are in decline across their eastern breeding range [20]. It is currently thought that the root of this decline lies in the breeding range (as opposed to in South America) [12] where the species is experiencing both phenological mismatch (whereby its northward migration to the breeding range is no longer timed to coincide with peak food availability due to climate change) [11] and competition for nest sites [18]. The eastern subspecies of purple martin (*Progne subis subis*) is highly unusual in that it relies entirely on man-made structures (principally multi-compartment nest boxes and artificial gourds) for nest locations [3].

Population size and stability is determined by the relative levels of fecundity and mortality. Collecting reliable data on factors determining the altricial nestling survival rates is thus essential to pinpointing the source of decline in purple martins. Hatching success is largely determined by incubation attentiveness (i.e. the proportion of time that the adults spend brooding eggs between laying and hatching) [7] and nestling survival to fledging is largely determined by provisioning rate (i.e. how much food is brought to the nestlings by their parents) [29]. However, collecting data on these important life stages has historically been very challenging. Direct observation of active nests is not only time consuming and technically difficult, but also carries the risk of disturbing the very activities researchers wish to document [1]. Researchers have therefore often turned to in-nest devices such as temperature or pressure sensors which can record proxies of incubation or provisioning behavior (e.g. [2, 14]), but which have unquantifiable misclassification errors [22]. In recent years, the use of continuously recording in-nest video cameras has become common place [5], which has reduced disturbance to birds and reliance on proxy data, but has created a new problem—how to review and classify all of the data. Most studies rely on subsampling the data for analysis [6], but this approach can miss important biological variation. Using deep learning techniques to automate the classification of nest camera footage has the potential to allow researchers to analyze the full dataset in a time and cost effective manner.

Here we fitted 20 purple martin nest boxes over two years (2017 and 2018) with high definition nest cameras at Iroquois National Wildlife Refuge in Western New York State. Cameras were set to record from 6am until 9pm (insufficient light level prevented overnight recording) from the start of nest building until the last nestling fledged. Given the average incubation period of 16 days and nestling period of 28 days in the purple martin, this generated close to 13,000 hours of footage. Video clips were categorized by nest stage (incubation; hatch day-day 5; day 5-10; day 11-15 and day 16-fledging) for analysis. We aimed to use deep learning to reliably estimate nest attentiveness during the incubation period and provisioning rate during the nestling period.

The results of the incubation nest attentiveness model will be used in a study of the intrinsic and extrinsic factors affecting nest attentiveness in the purple martin. The results of the provisioning rate model will be used to test the 'parental compensation' hypothesis, which predicts that parents of nestlings afflicted with a high load of nest parasites (fleas, blood-feeding mites and blowfly larvae in this case) will increase their food supply to their brood to negate any fitness cost to the nestlings [25]. To manipulate parasite load, half of the study nests had their nest materials (with associated parasites) replaced with clean nesting materials after day 5 of the nestlings' lives.

2 RELATED WORK

The recent growth in large image datasets in the ecological sciences [8, 26], and the need to reduce the laborious manual classification of them, has led to the emergence of deep learning techniques being used in the field [15, 27]. Studies to date have generally worked with the still images generated by camera traps, and have attempted species classification and abundance problems where "empty" images must be correctly discarded and "occupied" images must have species correctly identified and individuals counted [16, 24, 28]. Many of these studies utilize projects available through the Zooniverse platform (https://www.zooniverse.org/) which allows researchers to share their data projects and invite citizen scientists to volunteer to analyze them. Pre-analyzed datasets such as these provide an in-built validation set for deep learning projects and allow comparison between accuracy achieved by experts vs. citizen scientists vs. deep learning models [28]. In particular, the Snapshot Serengeti dataset [23] has been used to develop deep learning models for mammalian species identification with accuracy approaching human levels (up to 97%) [24] and with an almost entirely automated classification pipeline which represents a huge reduction in labor [16]. So far, however, deep learning techniques for classifying image data in ecology are mostly limited to the analysis of these emerging "standard" datasets. We are not aware of any study so far which has used video nest camera data and attempted to use deep learning to classify avian behavior during breeding.

3 METHOD

3.1 Convolutional Neural Network Structure and Parameters

The model chosen to analyze the purple martin incubation and provisioning videos is a convolutional neural network based on the Keras API for Tensor flow [4] which is based on Python. The calculations were done using NVIDIA Tesla P100 or V100 GPUs. The video frame processing was done with Python cv2. The original high resolution video output image size of 1920X1080 pixels was reduced to 0.1 times this size to speed processing. Note that even with the ten times reduction there was no noticeable pixilation of the images. In addition, also to speed processing, the 30 frame per second videos clips were reduced down to 1 frame per second during the production runs on both the incubation analysis and the provisioning analysis. The majority of the analysis was done with a six convolutional layer CNN. This CNN was chosen because it did a good job fitting the CIFAR-10 data set [13]. The architecture of the CNN network is shown in Table 1. The first layer is convolutional with a rectified linear unit (ReLU) activation; this is followed by a dropout layer set to 0.2 for randomization. The second convolutional layer, also activated by ReLU, is followed by a pooling layer. This structure of convolutional layer, dropout, convolutional layer, pooling layer is repeated two more times. The image is flattened followed by another dropout layer, then a fully connected layer and another dropout layer. The final fully connected output layer has a softmax activation. The optimizer was RMSprop which is similar to gradient descent with momentum. The filter size was empirically varied from 3X3 to 32X32 while monitoring the quality of the fit. The learning rate was varied from 10^{-3} to 10^{-6}. It was found that the CNNs trained the best at 10^{-5} for both the incubation model and the provisioning model. At this learning rate the model would train stably in 20-40 epochs without large oscillations, for example see Figure 1. Example training runs are discussed for each model below. It was found that setting random seeds was critical in obtaining the reproducible outputs required to conduct model optimization; otherwise the random nature of the network initialization lead to very non-reproducible results and made model optimization very difficult and time consuming. The specifics of constructing the incubation model and the provisioning models are discussed below.

3.2 Training and Validation Datasets

Video clips were split into still images sampling at a rate of 1 frame per 3 seconds using FFmpeg [9]. Training images were chosen to ensure that every nest (8 for incubation and 20 for provisioning) was represented approximately equally. We also balanced the training datasets by hour of day (6am-9pm) to ensure all lighting conditions were adequately represented. For the incubation training set we balanced the dataset by day of incubation (0-16). The provisioning models are split into two training sets depending on nestling age (0-5 days and 6-10 days). In both sets, the datasets initially represented each of the 20 nests equally with every nestling age and hour of the day also in balance. Due to the rarity of 2-bird images (see "Deep Learning Model of Provisioning") we subsequently supplemented the training set with additional 2-bird images from select nests. The final incubation training set included 12,144 images, the

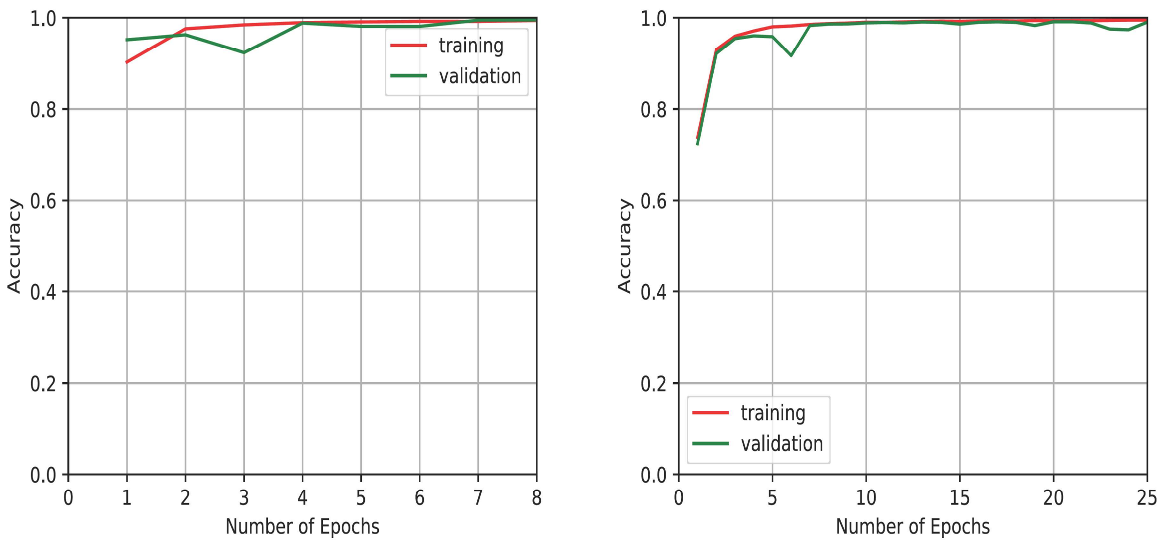

Figure 1: Training the incubation CNN model (left) and the provisioning CNN model for day 6-day 10 (right). The red line is the model accuracy on the training data. The green line is the model accuracy on the validation data; it represents the model generalization accuracy.

Table 1: CNN Architecture

Layer	Output Shape			Parameters
$conv2d_1$	192	108	32	7808
$dropout_1$	192	108	32	0
$conv2d_2$	192	108	32	82,976
$maxpooling2d_1$	96	54	32	0
$conv2d_3$	96	54	64	165,952
$dropout_2$	96	54	64	0
$conv2d_4$	96	54	64	331,840
$maxpooling2d_2$	48	27	64	0
$conv2d_5$	48	27	128	663,680
$dropout_3$	48	27	128	0
$conv2d_6$	48	27	128	1,327,232
$maxpooling2d_3$	24	13	128	0
flatten	39,936			0
$dropout_4$	39,936			0
$dense_1$	1096			43,770,952
$dropout_5$	1096			0
$dense_2$	2			2194

day 0-5 provisioning set included 30,285 images and the day 6-10 provisioning set included 40,825 images.

Validation videos were selected to also balance representation of each nest, day and hour of day. None of these validation videos are included in the training datasets. Validation videos were manually analysed to provide a comparison with the deep learning model outputs.

4 DEEP LEARNING MODEL OF INCUBATION

Conceptually the incubation model is very simple. There are only two classes, they are incubating or not incubating. If one or both parents is on the nest, the model should classify this as incubating; if neither parent is on the nest the model should classify this as not incubating. The model chosen to analyze the incubation videos is the six convolutional layer convolutional neural network (CNN) based on the Keras API for Tensor flow as described above. Another incubation CNN model also was tried but results were similar. The final filter size chosen was 9X9 (as per the provisioning model, see below) and the learning rate was 10^{-5}. Training the model on 12,144 static video images produced an excellent result; see the training curve shown in Figure 1. Training on 80% of the images and reserving 20% for validation yielded a model that achieved 99.5% accuracy on the validation data and produced the Confusion Matrix shown in Table 2. In contrast to changing the layers of the CNN model, altering the training image set was found to be the key to generating a good model. For example, it was found that occasionally the adult purple martin would cover the eggs with leaves before leaving the nest box: see Figure 2. During the development of the model, a preliminary version of the model evaluated this case as incubating, undoubtedly because the eggs were covered, even though it should be not incubating. We added training images that included images of the covered eggs appropriately labelled as not incubating. After training, the model was then able to correctly evaluate this case. The lesson is that the model can only be relied upon to correctly evaluate those situations which are covered by the training images. The final model analysis was compared with manually viewed videos for validation. Based on 43 validation videos mean absolute differences in nest attentiveness is less than 0.02 nest attentiveness

Table 2: Incubation Model Confusion Matrix

Actual	Model zero birds	Model one bird
zero birds	688	8
one bird	4	791

units. This approaches the accuracy which a human can attain and is more than sufficient to use these results for incubation analysis and to combine this data with other data for a more sophisticated analysis of incubation. The incubation videos have been analyzed in a production mode. To speed the analysis only one frame per second is analyzed. Although there is a simple linear relationship between nest attentiveness and the number of images classified as occupied, we checked that analyzing one frame per second produced the same result as 30 frames per second; the differences were out in the fourth decimal place. Using this approach, a 33.5 minute video can be analyzed in less than 4 minutes, that is, almost 10 times faster than real time. Incubation production runs have been completed and the incubation data is in the process of analysis.

5 DEEP LEARNING MODEL OF PROVISIONING

In contrast to the incubation model, the provisioning model is much more complex. The provisioning model has three classes: zero, one or two birds on the nest. The goal is to analyze how many provisioning events (that is when an adult provides food to the nestlings) occur during each video. A provisioning event can be an adult bird entering a nest alone to feed the nestlings or a second parent bird that enters the nest box when it is already occupied by their partner. In preliminary manual analysis it was found extremely rare that an adult would enter the nest and not provision. The way the model works is that a transition from zero to one or from one to two birds is defined to be a provisioning event. The difficulty is that any error can simulate a false positive event. A 6 layer CNN as described above similar to the incubation model was constructed in Keras/Tensorflow. Due to the difficulty in modelling provisioning, the filter size was carefully optimized; 9X9 filters were found to be the best. Other CNN architectures were also tried, notably the 16 CNN layer VGG16 model [21]. Little difference was noted between the CNN models given the same training data. Two different provisioning models were constructed, one for nestlings of age 0 to 5 days (d0-d5) and the other for nestlings of age 6 to 10 days (d6-d10). The models were trained with images specific to their age group. The only exception is that some of the two bird images from the d0-d5 training set were also used in the d6-d10. This is because two bird images were relatively rare and hard to find. When two birds are on the nest the nestlings were not visible so there should be no difference between d0-d5 images and d6-d10 images. An example training plot from the d6-d10 model is shown in Figure 1.

Even though the best models correctly classify better than 99% of the still images, see the Confusion Matricies shown in Tables 3 and 4, a few misidentifications among the many thousands of images on a single video clip can produce a poor analysis. Figure 3 gives an idea of some of the difficulties in properly classifying the provisioning images. One difficulty is in determining exactly when a bird enters the nest. For example, images B and C of Figure 3 show how subtle the difference between on the nest and off the nest can be. The problem is not in knowing exactly when the provisioning event starts. The problem is that often birds will linger and move about in the nest entrance. Another problem is determining whether there is one or two birds on the nest. For example images D and E of Figure 3 show a single bird spread out (D) vs two birds clustered together (E). In image F the bird completely blocks the camera view of the nest making it impossible to determine whether there are one or two birds on the nest. When a human is viewing the video he or she can look at a sequence of images to resolve these difficulties. However, the CNN must classify the images based exclusively on each separate image. In order to correct this we employ post classification filtering. Two assumptions are made to guide the filtering. The first is that provisioning events must last for several seconds while the bird climbs into the nest, feeds the nestlings and leaves. The second is that the provisioning event cannot be repeated within a given period in order to give the bird time to leave, catch more food and return. The filter takes the form of averaging over a number of seconds and eliminating events spaced closer than a specified time. In the final analysis 6 seconds was chosen for the averaging and 60 seconds for the minimum provisioning event separation. For an illustration of how this filtering works see Figure 4. The blue line shows the raw unfiltered analysis and the red line is the corrected, filtered analysis. The red line correctly predicts 6 provisioning events. The combined classification and filtering is far from perfect. A set of 40 videos was used for model validation. At this time the best performance achieved on the day 0 to day 5 hatchlings is 99% compared to manual viewing. (88% accuracy, that is 88% of all events are detected, with 11% false positives). The day 6 to day 10 hatchling data where the provisioning rate is greater is slightly poorer with 93% compared to manual viewing (77% correct and 15% false positives). Although this is considerably poorer than a human can do, nevertheless this accuracy is adequate for performing automated provisioning detection and drawing meaningful conclusions about purple martin provisioning behavior. To check this, we estimated the variance from the comparison of the observed and calculated provisioning validation data set. The data appeared to be normally distributed so we assumed a normal distribution. We also incorporated a bias term linearly dependent on the provisioning rate. That is, the higher the provisioning rate the more likely we are to miss some provisioning events that are close together. With these parameters we ran Monte Carlo simulations. For provisioning rate differences of 10% or greater between the control group and the experimental group the inaccuracy of the model was irrelevant. The model inaccuracy only was important for differences on the order of 5% or less. Variations between different sets of parents for different nests is quite likely to be much larger and is expected to be the dominant source of variation rather than the model itself.

6 CONCLUSIONS-SUMMARY

Using deep learning models for automated analysis of videos of purple martin nesting behavior has been accomplished. Incubation

Figure 2: Three purple martin incubation images are shown. In image A There are no birds on the nest and the 6 eggs are readily observed. This is a not incubation class image. In image B a female purple martin is incubating the eggs which are therefore obscured. This is an incubation class image. In image C no birds are on the nest, however, the eggs are obscured by leaves. This image should be classified as not incubation. Initial models classified this as incubation, based on the fact that the eggs were obscured rather than recognizing that no birds were present. Adding labeled training images which included leaf covered eggs as not incubation produced a model that correctly classified such images.

Figure 3: Six purple martin provisioning images are shown. Image A shows a bird on the nest provisioning one of the nestlings. When this bird entered the empty nest, that is a nest with nestlings but no adult bird, we would record this as a provisioning event. Images B and C show a bird entering the nest. In C the tail of the bird is barely visible but the bird is still considered to be outside the nest. In image B the bird has barely entered the nest. Distinguishing exactly when a bird enters the nest and therefore the start of a provisioning event can be difficult. This is exacerbated by the fact that birds often perch just outside the nest and move in the vicinity of the entrance. Images D, E and F illustrate another difficulty in detecting provisioning events. In image D a single bird is spread out over the nest. In image F two birds are huddled closely together on the nest. In image F a bird is blocking the camera view. In image F there are actually two birds in the nest but there is no way to determine this solely viewing this single image. When a human is viewing the video she or he can look at a sequence of images to resolve these difficulties. However, the model must classify the images frame by frame.

models which are dependent only on getting the great majority of the classifications correct on whether or not an adult bird is incubating or not incubating can be been accurately constructed. The analysis has been found to be largely independent of the actual structure of the CNN model but is very dependent on the training data images which must represent all possible important events. Nest attentiveness accuracy from automated video analysis approaches that which a human can produce manually. Automated provisioning analysis is a much more difficult problem to solve and it also relies on post classification filtering. The provisioning model developed is much less accurate but, based on Monte Carlo simulations, it should be sufficient for analysis of the provisioning videos.

7 FUTURE WORK

The results of the incubation nest attentiveness model will be used in a random forest model of factors affecting incubation effort in the purple martin. Specifically, we aim to analyze the extent to which purple martins adjust their attentiveness in response to their environment (namely, ambient temperature, wind speed and

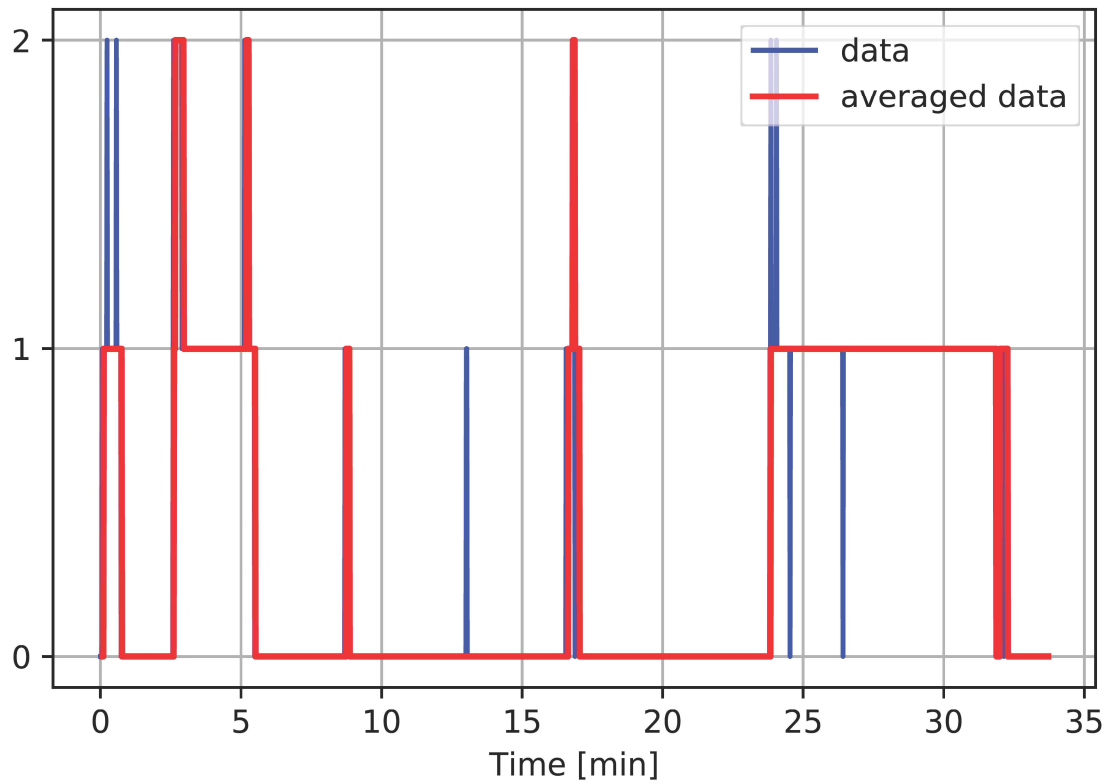

Figure 4: Automated provisioning event analysis of a typical video. The y-axis is the number of adult birds in the nest (0 for zero adult birds, 1 for one adult bird and 2 for 2 adult birds). Seven provisioning events are detected, see the red line. The provisioning events occur at 0.1, 2.6, 5.2, 8.8, 16.6, 23.9 and 32.0 minutes. The blue line shows false positives that have been eliminated by filtering.

Table 3: Provision Model Confusion Matrix Day0-Day5

Actual	Model zero birds	Model one bird	Model two birds
zero birds	1830	6	0
one bird	17	5227	17
two birds	0	48	1020

Table 4: Provision Model Confusion Matrix Day6-Day10

Actual	Model zero birds	Model one bird	Model two birds
zero birds	2820	11	0
one bird	7	2116	13
two birds	0	26	1064

rainfall) and how much variation there is in attentiveness with intrinsic parameters such as clutch size and adult age.

When results from the provisioning rate analyses are available, we aim to test whether parasite load has a positive [25], negative [17] or insignificant effect [19] on provisioning rate in the purple martin. We will test for differences in provisioning rate both between nests which have had experimental parasite reductions and control nests, and for differences in provisioning rate within nests before and after parasite reductions.

Increasing our understanding of incubation and provisioning behaviors in the purple martin helps us determine what affects nestling survival and may help us to refine conservation efforts for this declining long distance migrant.

ACKNOWLEDGMENTS

This work was sponsored by the National Science Foundation (NSF) under award ACI 1445806 for the XD Metrics Service (XMS).

The authors would like to thank: The staff at Center for Computational Research especially Cynthia Cornelius and the former director Dr. Thomas R. Furlani for assisting in this work and providing the computational facilities. Yashas J. Shamaraju for his work on an internship. Dr. Joseph P. White for assistance in the preparation of the manuscript.

H.W. is grateful to the Purple Martin Conservation Association, the Mark Diamond Research Fund and the Friends of Iroquois National Wildlife Refuge for funds to purchase and install nest camera equipment.

REFERENCES

[1] Alexis L. Blackmer, Josh Ackerman, and Gabrielle A. Nevitt. 2004. Effects of investigator disturbance on hatching success and nest-site fidelity in a long-lived seabird, Leach's storm-petrel. *Biological Conservation* 116, 1 (2004), 141–148. https://doi.org/10.1016/S0006-3207(03)00185-X

[2] Grace E. Bottitta, H. Grant Gilchrist, A. Kift, and M. G. Meredith. 2002. A Pressure-Sensitive Wireless Device for Continuously Monitoring Avian Nest Attendance. *Wildlife Society Bulletin* 30, 4 (2002), 1033–1038. http://www.jstor.org/stable/3784269

[3] C.R. Brown and S.A. Tarof. 2017. Purple martin (Progne subis). Account 287 in A. Poole, editor. The birds of North America online. Cornell Lab of Ornithology. http://bna.birds.cornell.edu/bna/species/287

[4] François Chollet et al. 2015. Keras. https://github.com/fchollet/keras.

[5] W.A. Cox, M.S. Pruett, T.J. Bensen, S.J. Chiavacci, and F.R. Thompson. 2012. Development of camera technology for monitoring nests. In *Video surveillance of nesting birds. Studies in Avian Biology*, Thompson F.R. Ribic, C.A. and P.J Pietz (Eds.). University of California Press, Berkley, CA.

[6] S.K. Davis and T.G Holmes. 2012. Spragues Pipit incubation behavior. In *Video surveillance of nesting birds. Studies in Avian Biology*, Thompson F.R. Ribic, C.A. and P.J Pietz (Eds.). University of California Press, Berkley, CA.

[7] Denis Charles Deeming. 2015. *Nests, eggs, and incubation: new ideas about avian reproduction* (first ed.). Oxford University Press, Oxford.

[8] Anthony I. Dell, John A. Bender, Kristin Branson, Iain D. Couzin, Gonzalo G. de Polavieja, Lucas P.J.J. Noldus, Alfonso Perez-Escudero, Pietro Perona, Andrew D. Straw, Martin Wikelski, and Ulrich Brose. 2014. Automated image-based tracking and its application in ecology. *Trends in Ecology Evolution* 29, 7 (2014), 417 – 428. https://doi.org/10.1016/j.tree.2014.05.004

[9] FFmpeg Developers. 2016. ffmpeg tool (Version be1d324). http://ffmpeg.org/.

[10] K. C. Fraser, A. Shave, A. Savage, A. Ritchie, K. Bell, J. Siegrist, J. D. Ray, K. Applegate, and M. Pearman. 2017. Determining fine-scale migratory connectivity and habitat selection for a migratory songbird by using new GPS technology. *Journal of Avian Biology* 48, 3 (2017), 339–345. https://doi.org/10.1111/jav.01091 arXiv:https://onlinelibrary.wiley.com/doi/pdf/10.1111/jav.01091

[11] Kevin C. Fraser, Cassandra Silverio, Patrick Kramer, Nanette Mickle, Robert Aeppli, and Bridget J. M. Stutchbury. 2013. A Trans-Hemispheric Migratory Songbird Does Not Advance Spring Schedules or Increase Migration Rate in Response to Record-Setting Temperatures at Breeding Sites. *PLOS ONE* 8, 5 (05 2013), 1–5. https://doi.org/10.1371/journal.pone.0064587

[12] Kevin C. Fraser, Bridget J. M. Stutchbury, Cassandra Silverio, Patrick M. Kramer, John Barrow, David Newstead, Nanette Mickle, Bruce F. Cousens, J. Charlene Lee, Danielle M. Morrison, Tim Shaheen, Paul Mammenga, Kelly Applegate, and John Tautin. 2012. Continent-wide tracking to determine migratory connectivity and tropical habitat associations of a declining aerial insectivore. *Proceedings of the Royal Society B: Biological Sciences* 279, 1749 (2012), 4901–4906. https://doi.org/10.1098/rspb.2012.2207 arXiv:https://royalsocietypublishing.org/doi/pdf/10.1098/rspb.2012.2207

[13] Magnus Hvass. 2018. Tensor Flow Tutorials. https://github.com/Hvass-Labs/TensorFlow-Tutorials.

[14] S. Charles Kendeigh. 1952. *Parental care and its evolution in birds* (second printing, with minor corrections and an addendum, february, 1955. ed.). Vol. 22, nos 1-3. University of Illinois Press, Urbana.

[15] Roberta Kwok. 2019. AI empowers conservation biology. *Nature* 567, 7746 (March 2019), 133–134. https://doi.org/10.1038/d41586-019-00746-1

[16] Mohammad Sadegh Norouzzadeh, Anh Nguyen, Margaret Kosmala, Alexandra Swanson, Meredith S. Palmer, Craig Packer, and Jeff Clune. 2018. Automatically identifying, counting, and describing wild animals in camera-trap images with deep learning. *Proceedings of the National Academy of Sciences* 115, 25 (2018), E5716–E5725. https://doi.org/10.1073/pnas.1719367115 arXiv:https://www.pnas.org/content/115/25/E5716.full.pdf

[17] Anne Oppliger, Heinz Richner, and Philippe Christie. 1994. Effect of an ectoparasite on lay date, nest-site choice, desertion, and hatching success in the great tit (Parus major). *Behavioral Ecology* 5, 2 (1994), 130–134. https://doi.org//10.1093/beheco/5.2.130

[18] Daniel Raleigh, James D. Ray, Blake A. Grisham, Joe Siegrist, and Daniel U. Greene. 2019. Nest survival data confirm managed housing is an important component to the conservation of the eastern purple martin. *Wildlife Society Bulletin* 43, 1 (2019), 93–101. https://doi.org/10.1002/wsb.941 arXiv:https://wildlife.onlinelibrary.wiley.com/doi/pdf/10.1002/wsb.941

[19] C. A. Rogers, R. J. Robertson, and B. J. Stutchbury. 1991. Parasitism of tree swallow nestlings. In *Bird-parasite interactions: Ecology, Evolution and Behaviour*, J. E. Loye and M. Zuk (Eds.). Oxford University Press, New York.

[20] D. K. Niven J. E. Hines D. J. Ziolkowski Jr K. L. Pardieck J. E. Fallon Sauer, J. R. and W. A. Link. 2017. The North American Breeding Bird Survey, Results and Analysis 1966 - 2015. Version 2.07.2017. https://www.mbr-pwrc.usgs.gov/bbs/.

[21] Karen Simonyan and Andrew Zisserman. 2014. Very Deep Convolutional Networks for Large-Scale Image Recognition. *CoRR* abs/1409.1556 (2014). arXiv:1409.1556 http://arxiv.org/abs/1409.1556

[22] J.A. Smith, C.B. Cooper, and S.J. Reynolds. 2015. Advances in techniques to study incubation. In *Nests, eggs and incubation: new ideas about avian reproduction*, D.C. Deeming and S.J. Reynolds (Eds.). Oxford University Press, New York.

[23] Alexandra Swanson, Margaret Kosmala, Chris Lintott, Robert Simpson, Arfon Smith, and Craig Packer. 2015. Snapshot Serengeti, high-frequency annotated camera trap images of 40 mammalian species in an African savanna. *Scientific Data* 2, 150026 (2015). https://doi.org/doi.org/10.1038/sdata.2015.26

[24] Michael A. Tabak, Mohammad S. Norouzzadeh, David W. Wolfson, Steven J. Sweeney, Kurt C. Vercauteren, Nathan P. Snow, Joseph M. Halseth, Paul A. Di Salvo, Jesse S. Lewis, Michael D. White, Ben Teton, James C. Beasley, Peter E. Schlichting, Raoul K. Boughton, Bethany Wight, Eric S. Newkirk, Jacob S. Ivan, Eric A. Odell, Ryan K. Brook, Paul M. Lukacs, Anna K. Moeller, Elizabeth G. Mandeville, Jeff Clune, and Ryan S. Miller. 2018. Machine learning to classify animal species in camera trap images: Applications in ecology. *Methods in Ecology and Evolution* 10, 4 (2018), 585–590. https://doi.org/10.1111/2041-210X.13120 arXiv:https://besjournals.onlinelibrary.wiley.com/doi/pdf/10.1111/2041-210X.13120

[25] Frederic Tripet and Heinz Richner. 1997. Host Responses to Ectoparasites: Food Compensation by Parent Blue Tits. *Oikos* 78, 3 (1997), 557–561. http://www.jstor.org/stable/3545617

[26] Jana Waldchen and Patrick Mader. 2018. Machine learning for image based species identification. *Methods in Ecology and Evolution* 9, 11 (2018), 2216–2225. https://doi.org/10.1111/2041-210X.13075 arXiv:https://besjournals.onlinelibrary.wiley.com/doi/pdf/10.1111/2041-210X.13075

[27] Ben G. Weinstein. 2018. A computer vision for animal ecology. *Journal of Animal Ecology* 87, 3 (2018), 533–545. https://doi.org/10.1111/1365-2656.12780 arXiv:https://besjournals.onlinelibrary.wiley.com/doi/pdf/10.1111/1365-2656.12780

[28] Marco Willi, Ross T. Pitman, Anabelle W. Cardoso, Christina Locke, Alexandra Swanson, Amy Boyer, Marten Veldthuis, and Lucy Fortson. 2018. Identifying animal species in camera trap images using deep learning and citizen science. *Methods in Ecology and Evolution* 10, 1 (2018), 80–91. https://doi.org/10.1111/2041-210X.13099 arXiv:https://besjournals.onlinelibrary.wiley.com/doi/pdf/10.1111/2041-210X.13099

[29] Ken Yasukawa, Judy L. Mcclure, Rebecca A. Boley, and Juliann Zanocco. 1990. Provisioning of nestlings by male and female red-winged blackbirds, Agelaius phoeniceus. *Animal Behaviour* 40, 1 (1990), 153 – 166. https://doi.org/10.1016/S0003-3472(05)80675-X