AI for Conservation: Learning to Track Birds with Radar

How deep neural networks can process millions of weather radar data points to help researchers monitor continental-scale bird migration.

Author: Zezhou Cheng, Subhransu Maji, Daniel Sheldon

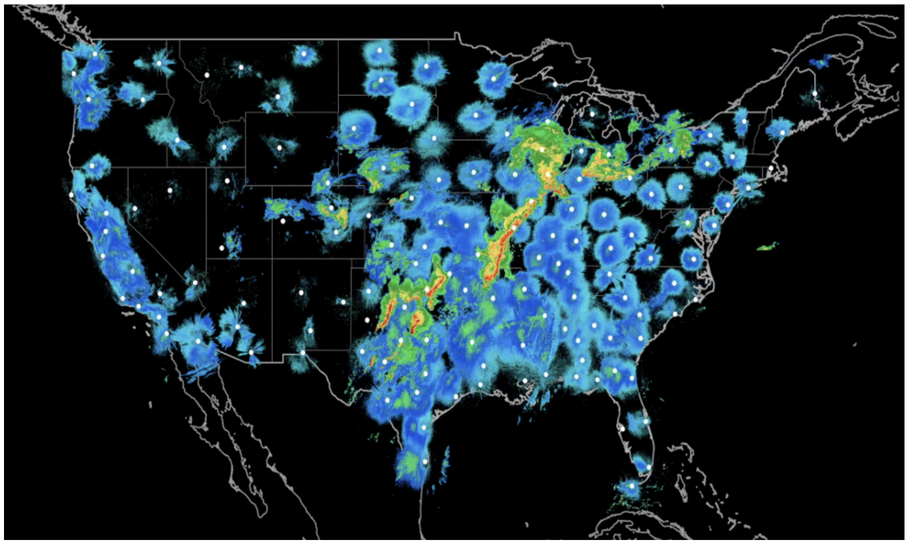

Figure 1: The U.S. Next Generation Weather Radar (NEXRAD) system. The white dots indicate radar stations and the blue blobs surrounding them are detected migratory birds. The green and yellow splotches are storms.

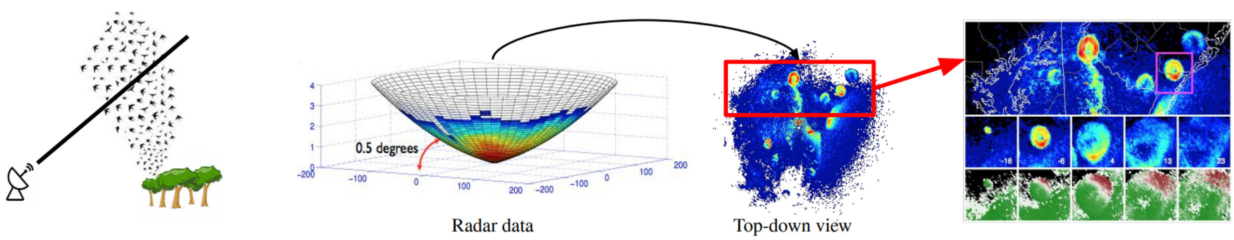

(a) Communal roost exodus

(b) Radar geometry and roost appearance

Figure 2: Roost exodus and appearance in radar data. (a) Illustration of roost exodus. (b) A radar traces out cone-shaped slices of the atmosphere (left), which are rendered as top-down images (center). This image from the Dover, Delaware radar station at 6:52 a.m. on October 2nd, 2010, shows at least eight roosts. Several are shown in more detail to the right, together with crops of one roost from five consecutive reflectivity and radial velocity images over a period of 39 minutes. These show the distinctive expanding ring and “red–white–green” diverging velocity patterns.

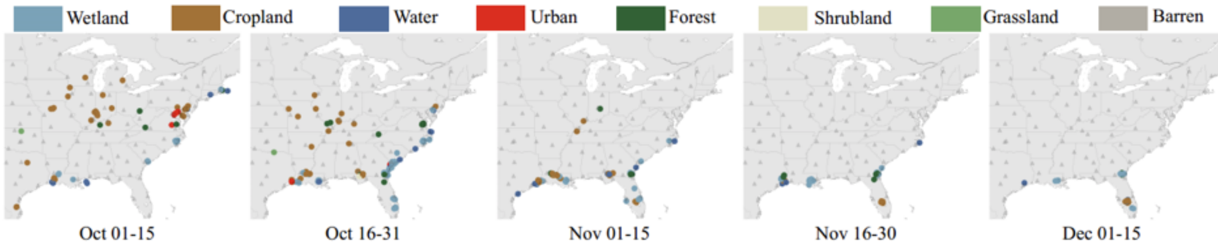

Figure 3: Tree swallow fall migration in 2013. The colored circles show detected roost locations within each half-month period. The color of each circle indicates the habitat type.

The decline of wildlife populations is one of the biggest environmental challenges of this century. As a case in point, the North American bird population has declined by 2.9 billion birds since 1970—a drop of 29 percent [1]. Information about the full life cycles of migratory birds is thus urgently needed to understand and reverse the causes of their decline. A first step is to gather historical information about bird behaviors such as migration and habitat usage. However, these behaviors are incredibly hard to measure. While some individual birds can be tracked with digital devices, technical and logistical challenges prevent us from tracking a large number of birds over long distances to gain a comprehensive understanding of continental-scale behaviors such as migration.

Fortunately, the U.S. weather radar network offers an unprecedented opportunity to study dynamic, continental-scale movements of animals over a long time period. The National Weather Service operates a network of 143 radars in the contiguous U.S. These radars were designed to study weather phenomena such as precipitation and severe storms, but are very sensitive and turn out to also detect flying animals, including birds, bats, and insects. The data archive dates from the early 1990s and provides comprehensive views of a number of significant biological phenomena. These include broad-scale patterns of bird migration, such as the density and velocity of migrating birds flying over the entire country, as shown in Figure 1. Radars can also detect phenomena that can be matched to individual species, such as insect hatches and departures of large flocks of birds or bats from roosting locations. For this reason, radar data has been used in ecological research for decades. However, the scope of ecological research using weather radar has historically been limited due to the difficulty of accessing the biological information in the data. Researchers have only recently begun to apply artificial intelligence (AI) to unlock truly broad-scale and long-term information from this rich data source [2].

Tracking Birds with Artificial Intelligence

We developed an AI system to automatically track the roosting and migration behaviors of several species of songbirds. These birds gather in large communal roosts at night during portions of the year, and each roost may contain thousands to millions of birds. The roosts are packed so densely that the birds are visible on radar when they depart in the morning. Roosts occur all over the U.S., but few are documented, and most are never witnessed by humans since birds enter and leave the roosts during twilight. Those who do witness the swarming behavior of huge flocks at roosts report an awe-inspiring spectacle[1]. We are particularly interested in swallow species, which, due to their unique behavior, create a distinctive expanding ring pattern in radar (see Figure 2). During fall and winter, the tree swallow is the only swallow species present in North America, so roost signatures provide nearly unambiguous information about a single species at continental scale and a daily resolution. This is highly significant from an ecological perspective. Swallows are aerial insectivores, a population that is rapidly declining in North America, possibly due to widespread reductions in insect populations amongst other causes. Precisely documenting migratory patterns is a research priority for understanding the loss of these vulnerable bird populations [3].

The dramatic successes of deep learning for recognition tasks make it an excellent candidate for detecting and tracking roosts using radar data. Our overall approach consists of four steps. First, the radar scans are rendered as multi-channel images and a single-frame roost detector system is built based on Faster R-CNN, an established object detection framework that uses convolutional neural networks [4]. Interestingly, although radar data is visually different from natural images, we found that pretraining the network on a large natural image dataset (in particular ImageNet [5]) is quite useful; without pretraining, the networks took significantly longer to converge and resulted in 15 percent lower performance. Then, a tracking system based on the “detect-then-track” paradigm is developed. The system assembles roost detections into sequences to compute biologically meaningful measures and improve detector performance. Specifically, we employ a greedy heuristic to assemble detections from individual frames into tracks, starting with high-scoring detections and incrementally adding unmatched detections with high overlap in nearby frames. After associating detections, we apply a Kalman smoother to each track using a linear dynamical system model for the bounding box center and radius. Our error analysis suggests that most false positives from our detection and tracking pipeline were due to precipitation, which appear as highly complex and variable patterns in radar images—it is thus common to find small image patches that share the general shape and velocity pattern of roosts. The second leading source of false positives was wind farms. Surprisingly, the wind farms share several features of roosts: they appear as small high-reflectivity “blobs” and have a diverse velocity field due to spinning turbine blades. We exploit auxiliary information about precipitation and wind farms (e.g., a wind turbine database) to reduce these false positives. In the final system, 93.5 percent of the top 521 tracks on test data were correctly identified and tracked.

A significant challenge in this work was the presence of systematic differences in labeling style. We obtained a data set of manually annotated roosts collected for prior ecological research, believed to be nearly all tree swallow roosts. Each label records the position and radius of a circle within a radar image that best approximates the roost. We collected scans starting from 30 minutes before sunrise to 90 minutes after, leading to a large data set of 63,691 labeled roosts in 88,972 radar scans. However, this training data was annotated by different researchers and naturalists using a public tool developed for prior research studies. Since the annotators had different backgrounds and goals, and roost appearance in radar is poorly understood, there was considerable variability in the labels, much of which was specific to individual users. This

[1] See https://people.cs.umass.edu/~zezhoucheng/roosts for an example.

variation makes evaluation using held-out data very difficult and inhibits learning due to inconsistent supervision. To solve this issue, we developed a novel approach to jointly learn a detector together with user-specific models of labeling style. Specifically, we modeled the true label as a latent variable and introduced a probabilistic user model for the observed label conditioned on the image, true label, and user. We developed a variational expectation–maximization learning algorithm (for more details, see [6]) that permits learning with only black-box access to an existing object detection model, such as Faster R-CNN. Our results showed that accounting for user-specific labeling bias significantly improves evaluation and learning.

Insights into a Migratory Season

Finally, we conducted a case study where we used our detector and tracker to synthesize knowledge about continental-scale movement patterns of swallows. We applied our pipeline to 419,000 radar scans collected from 86 radar stations in the eastern U.S. from October 2013 through March 2014. During these months, tree swallows are the only swallow species in the U.S. in fall and winter, and are the predominant species in early spring. They are thus responsible for the vast majority of radar roost signatures. Widespread statistics of roost locations and habitat usage throughout a migratory season have not previously been documented. In this case study, our AI system enabled the first ever comprehensive measurements of a single species of bird across its range on a daily basis. The detector has excellent precision, detects previously unknown roosts, and demonstrates the ability to generate urgently needed continental-scale information about the non-breeding season movements and habitat use of an aerial insectivore. A summary of results from our case study is depicted in Figure 3. In early October 2013, tree swallows left their breeding territories and formed migratory and pre-migratory roosts throughout their breeding range across the northern U.S. By early November, only a few roosts lingered near major water bodies in the central U.S.; some birds have left the U.S. entirely to points farther south, while some remain in staging areas along the Gulf Coast. By December, activity in the Gulf Coast has diminished, and roosts were more concentrated in Florida, where a population of tree swallows will spend the entire winter. This case study highlights the importance of the eastern seaboard and the Mississippi valley as migration corridors, with different patterns of habitat use (wetland versus agricultural) in each. The strong association with agricultural habitats during the harvest season suggests interesting potential interactions between humans and the migration strategy of swallows. These results provide a starting point to better understand and conserve the bird populations.

Conclusion

Looking into the future, our system will help scientists to collect fine-grained measurements of birds—across the continent, at a daily time scale—from the entire 20-year radar archive. The findings can be used to study trends in population size and migration behavior in relation to changes in climate and habitat availability, and we will hopefully gain some of the first insights into where and when birds die during their annual life cycle.

References

[1] Rosenberg, K. V., et al. Decline of the North American avifauna. *Science* 366, 6461 (2019), 120–24.
[2] Lin, T. Y., et al. MistNet: Measuring historical bird migration in the US using archived weather radar data and convolutional neural networks. *Methods in Ecology and Evolution* 10 (2019), 1908–22.
[3] Spiller, K. J., and Dettmers, R. Evidence for multiple drivers of aerial insectivore declines in North America. *The Condor* 121, 2 (2019), duz010.
[4] Ren, S., et al. Faster R-CNN: Towards real-time object detection with region proposal networks. IEEE Transactions on Pattern Analysis and Machine Intelligence 39, 6 (2017), 1137–49
[5] Deng, J., et al. ImageNet: A large-scale hierarchical image database. In *Proceedings of the 2009 IEEE Conference on Computer Vision and Pattern Recognition (CVPR '09)*. 2009, 248–55.
[6] Cheng, Z., et al. Detecting and tracking communal bird roosts in weather radar data. In *Proceedings of the AAAI Conference on Artificial Intelligence (AAAI '20)*. 2020, 378–85.

Biographies

Zezhou Cheng is a Ph.D. student at the University of Massachusetts Amherst under the supervision of Prof. Subhransu Maji and Prof. Daniel Sheldon. His research interests include computer vision, machine learning, and computational sustainability.

Subhransu Maji is an Associate Professor in the College of Information and Computer Sciences at the University of Massachusetts Amherst. He obtained his Ph.D. from the University of California, Berkeley in 2011, and a B.Tech. in Computer Science and Engineering from the Indian Institute of Technology Kanpur in 2006.

Daniel Sheldon is an Associate Professor in the College of Information and Computer Sciences at the University of Massachusetts Amherst. He obtained his Ph.D. from Cornell University in 2009, and an A.B. in Mathematics from Dartmouth College in 1999.