Biodiversity and Energy Online Mapping

About the layers

Layer: Modeled Migratory Bird Stopovers
Date of this document: 3 April 2013
Layer developed by: Cornell Laboratory of Ornithology and New York Natural Heritage Program

Short Description: This layer depicts the results models predicting stopover habitat for 28 species of migratory birds. For each species, known stopover locations were obtained from the eBird database and used to predict the probability of occurrence throughout the state using landcover data. The probability of occurrence for each species was rescaled from 0 to 1 and the values for each species were added together. The higher the number in a grid cell, the more species are predicted to stop over within that cell during their species-specific migration window.
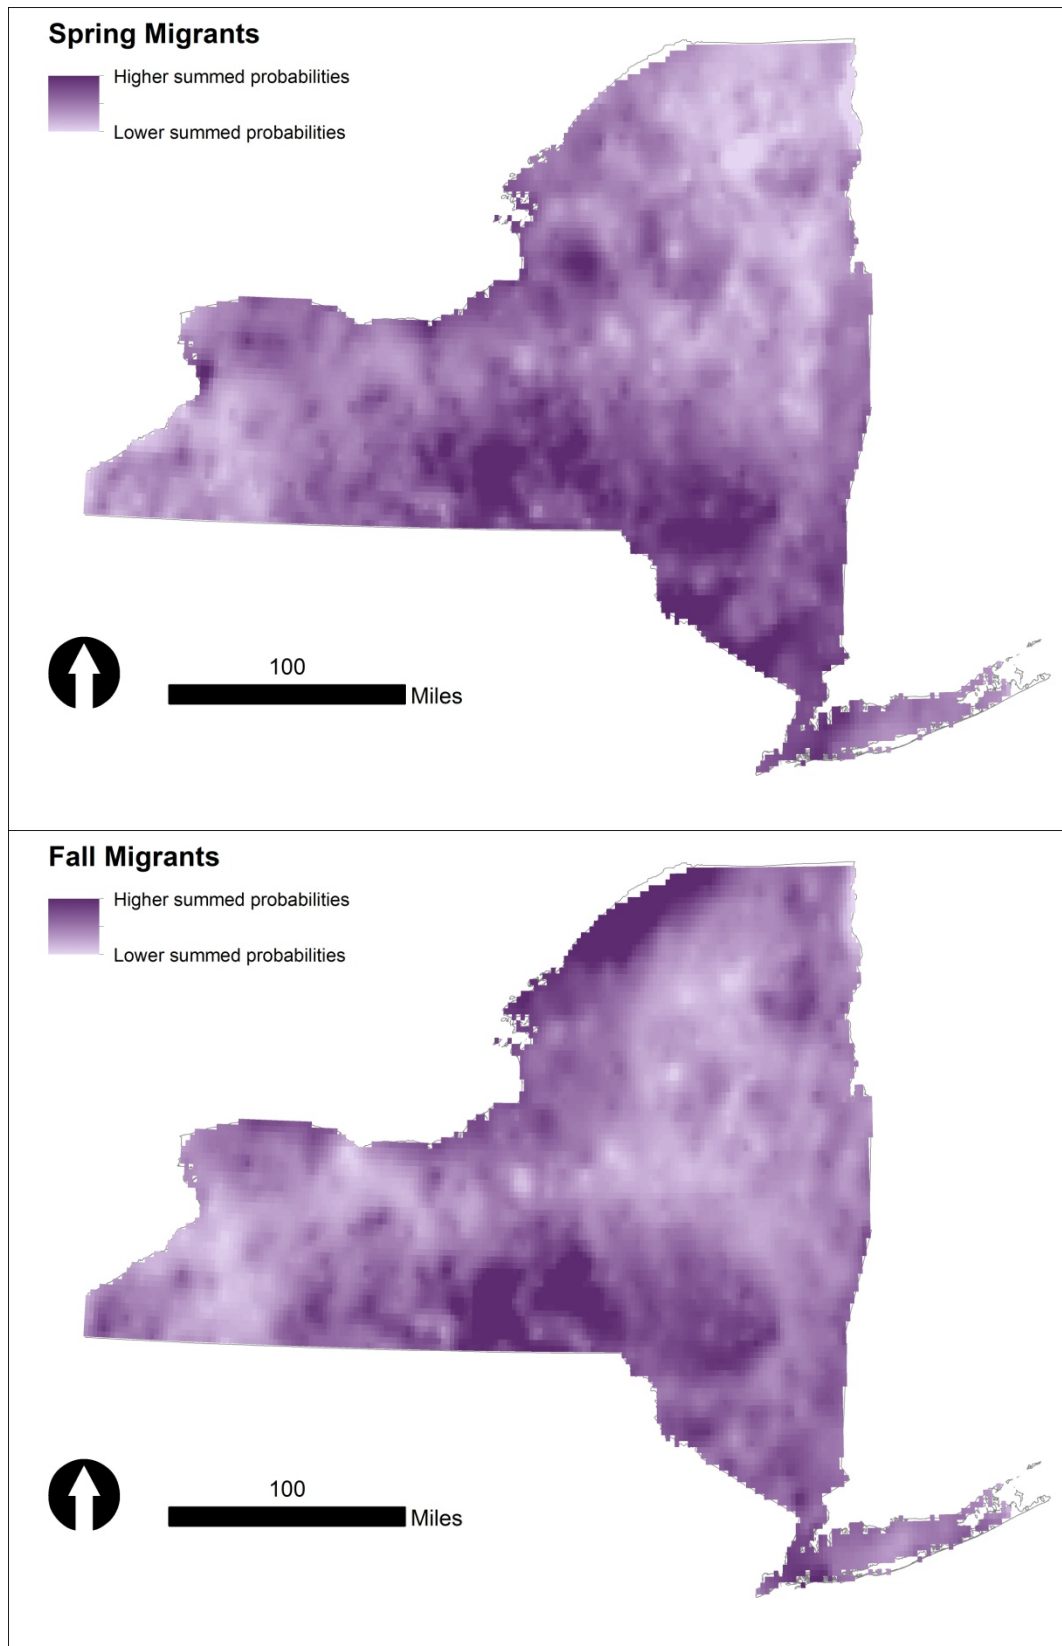
Why this layer matters:
Migratory birds are known to be killed by wind turbines and may be especially susceptible within heavily used flyways and stopover hotspots. Understanding where migratory birds are most likely to fly through and stop over may be important for avoiding impacts of development or for mitigation during certain times of year or in certain weather conditions.

Source: Location data derived from the eBird dataset. Explanatory variables are described in the full report, below.

Processing Overview:
1. Methods for generating continuous probability of occurrence surfaces for each of the 28 species in both spring and fall are described in the report below.
2. We normalized each species by dividing all values by the maximum to rescale each species from 0 to 1.
3. We added the values for each species in each cell to yield a composite value representing the probability of multiple species occurring in each cell during each species' migration window. This range was labeled from low to high.

Spring Migrants

Higher summed probabilities

Lower summed probabilities

100
Miles

Fall Migrants

Higher summed probabilities

Lower summed probabilities

100
Miles

eBird Species Distribution Modeling of Migrant Landbirds for:

Balancing Wind and Wildlife: New Data and Tools
To Improve Wind Project Siting for Biodiversity Conservation

Daniel Fink, Frank A. La Sorte, Andrew Farnsworth,
Steve Kelling, and Ron Rohrbaugh

August 2012
(Revised March 2013)

INTRODUCTION
The primary goal for this portion of the "Balancing Wind" project was to use observational data from the Cornell Lab of Ornithology's eBird project (www.ebird.org) to develop spatio-temporal models of migration stopover locations and habitat in New York State as a means of informing geographical positioning of new wind energy developments. This report includes three components:

1) Spatio-temporal Exploratory Model (STEM) data products, including maps, habitat associations, and model performance statistics for 28 species.
2) A technical report that provides species-level interpretation of the STEM output and methodology describing the modeling process.
3) A PDF of a PowerPoint presentation (Webinar) given to the project "PAC" by Ron Rohrbaugh and Frank LaSorte.

Section 1 of this report describes the conceptual background behind the modeling effort and criteria for selecting species, while sections 2-5 focus on various aspects of STEM methodology and evaluation, and qualitative species-specific model interpretations. The STEM data products, including maps, habitat specific results, and model performance metrics have been provided in a separate package.

SECTION 1: Concept
Most migratory passerine and shorebird species migrate at night, taking flight soon after dusk and landing in migratory stopover habitats a few hours prior to dawn. This nightly cycle is typically repeated for several weeks until birds arrive at their wintering or breeding grounds, depending on the direction of migration. Migratory stopover sites and their associated habitats are ecologically vital for resting and foraging. Little is known about how birds select stopover sites or how habitat and landscape features shape migration pathways. In general, most species appear to use stopover habitats that are very similar in composition, configuration, and age to their breeding habitat. For example, Wood Thrush breed in mature deciduous forest and stopover in the same, while Bobolink breed and stopover in grassland habitats. These stopover sites, however, don't seem to be chosen randomly from the available appropriate habitat, but rather are likely linked to fine-scale migration pathways.

By modeling eBird data in a spatio-temporal framework and visualizing the results at fine scales (3-6 km), we hoped to map important migratory stopover sites and perhaps infer migratory pathways by linking the flight trajectories among sites. We modeled the spring and fall "peak migration" distributions for 28 species. Peak migration period for each species in both spring and fall was determined by developing a frequency distribution of eBird observations within New York State and then selecting the

roughly two-week period that best represented migration activity. It was important to avoid modeling data from the breeding period, as this would simply produce a map of breeding locations.

Section 1.1: Species Selection
We modeled 28 species in four classes: 1) selected migratory passerines that commonly breed in New York State, 2) selected migratory passerines that rarely or do not breed in New York State, but migrate through ("Fly-through Migrants"), 3) selected shorebirds, and 4) species identified by Arnold and Zink (2011) as "super colliders." Breeding birds were grouped by habitat type to reveal migratory stopover sites within forests, shrublands, and grasslands (Table 1). We were limited by sample size and only chose those species that had adequate numbers of observations in eBird to produce reliable model results.

Table 1. Species distribution models were developed for 28 bird species that breed in or migrate through New York State.

Migratory Breeding Birds	"Fly-through" Migrants	Shorebirds	Super Colliders
Forest Associates	American Tree Sparrow	Least Sandpiper	Fox Sparrow
Wood Thrush	White-crowned Sparrow	Lesser Yellowlegs	Golden-winged Warbler
Veery	White-throated Sparrow	Greater Yellowlegs	Black-throated Blue Warbler
Scarlet Tanager	Swainson's Thrush	Dunlin	Bay-breasted Warbler
Ovenbird	Orange-crowned Warbler		
Shrubland/Young Forest Associates	Canada Warbler		
Chestnut-sided Warbler	Tennessee Warbler		
Indigo Bunting	Wilson's Warbler		
Yellow Warbler	Blackpoll Warbler		
Grasslands Associates	Olive-sided Flycatcher		
Bobolink			
Eastern Meadowlark			
Savannah Sparrow			

SECTION 2: Data
In Section 2.1 we discuss the bird observation data obtained from eBird and in section 2.2 we provide a brief discussion of the data sources and pixel sizes for our covariate data.
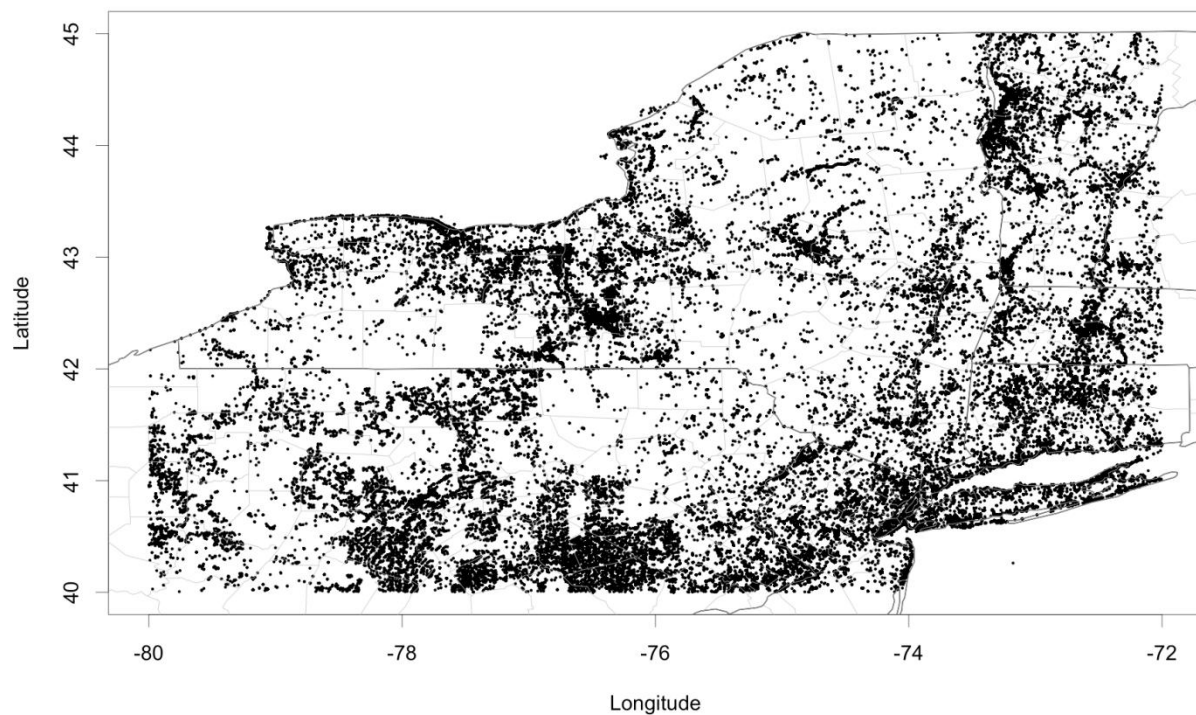
SECTION 2.1: eBird Data
The bird observation data used to develop the species distribution models comes from the Cornell Lab of Ornithology citizen-science project eBird (Sullivan et al. 2009). Specifically, we used eBird Reference Dataset 4.0 (Munson, et al. 2012). eBird is unique among broad-scale bird monitoring projects in that it collects observations made throughout the year. Participants follow a checklist protocol, in which time, location, and counts of birds are all reported in a standardized manner. By asking participants to indicate when they have contributed "complete checklists" of all the species they detected, we can assume that species with no detected individuals conveys absence information for that location. A subset of eBird participants used standardized protocols designed to collect additional information on search effort. Together, the reports of absence and effort add valuable information, allowing the analytical control of variable detection rates when inferring absences.

The analyses described in this report are based on presence-absence data from complete checklists collected under the "traveling count" and "stationary count" protocols from January 1, 2004 - December

31, 2011. Transect distances were limited to 8.1 km (5 miles), start times were restricted to daylight hours between 5AM and 8PM, and the total search time was limited to less than three hours.

We analyzed eBird data for the New York State region. This is defined as the region of the United States falling within the bounding rectangle formed between 40 and 45 degrees latitude and 80 to 72 degrees west longitude. The bounding box was selected to be larger than New York's political boundaries to reduce boundary effects, especially for modeling migrations. Species distribution models for this analysis were trained using 285,343 observations made across 39,059 unique locations (Figure 1). Approximately ten percent of data, 31,571 observations made at 4,192 locations, were held aside for model validation.

Figure 1. eBird training data locations for the New York State Region.

SECTION 2.2: Predictor Data
We include three general classes of predictors in the model.
1) Effort and observation/detection predictors: these are included to account for variation in detection rates, an unavoidable complexity e when making inference about species occupancy.
2) Temporal predictors to account for trends, and
3) Spatial predictors to account for spatial and spatiotemporal patterning.

There are three effort variables included in the model to account for variation in detection rates: the hours spent searching for species, the length of transects traveled during the search, and the number of people in the search party. The time of the day at observation is used to model variation in availability for detection; e.g., diurnal variation in behavior, such as participation in the "dawn chorus" (Diefenbach et al. 2007), may make species more or less conspicuous. An indicator of observations made under the "traveling count" protocol was included to allow the model to capture systematic differences in species

3

detection between the two protocols.

Temporal information comes from the covariate that records the day of the year (1-366) on which the search was conducted. This covariate is used to capture day-to-day changes in occurrence, and, similarly, the year of the observation is included to account for year-to-year differences.

Spatial information is captured by landcover and elevation data. To account for habitat selectivity, each eBird location was linked with a set of remotely sensed landcover variables. The U.S. 2006 National Land Cover Database (NLCD) (Homer et al. 2007) classifies vegetation into one of 16 classes at a 30-m cell resolution. From this information we calculated the percent coverage of each land-cover class in a 1.5 km pixel (225 ha) centred on each location. Elevation measured at 30 arc second (GTOPO30, http://eros.usgs.gov/#/Find_Data/Products_and_Data_Available/gtopo30_info) resolution was also included.

SECTION 3: Species Distribution Modeling with STEM
Species distributions were estimated one species at a time using the SpatioTemporal Exploratory Model (STEM) (Fink et al. 2010). The analysis preformed here was constructed to estimate the daily distribution of species occurrence, averaged across the 8-year study period.

The STEM was fit using a partition with cells of 2 degrees latitude and 2 degrees longitude and temporal window of 40 days. The 40-day window has been found to capture a wide variety of complex avian migration patterns across a diverse set of terrestrial species (e.g., NABCI 2011). The 2-degree size was selected as a practical compromise--small enough to potentially detect spatial "non-stationarity" within the state, but large enough to provide good predictive performance. The quality of predictive performance was judged based on both numerical summaries of validation set predictive performance as well as expert opinion. Validation set predictive performance was evaluated based on the numerical summaries of monthly predictive performance using the method discussed in Fink et al. (2010). Expert review was used to assess the correspondence between New York and national results and assess the quality of New York state estimates.

The STEM ensemble was created as a two-step procedure. In the first step, a bootstrap replicate of the training data was selected and in the second step a randomly located spatiotemporal partition was created across the study area. Four randomly located partitions were generated for each of fifty bootstrap replicates to create the ensemble for each species' analysis. Thus, a maximum of 200 base models, support the prediction at each location. Variation across bootstrap replicates can be used to quantify the variation in predictions, partial dependence estimates, as well as other model summaries.

For each cell in each partition of the ensemble, a boosted regression tree base model was fit to the data within that cell. The minimum data size per cell was 30 observations. The boosted regression tree was fit using the gbm package version 1.6.3.1 (Ridgeway G. 2010 in R version 2.14 (R Development Core Team 2011)) with the bagging fraction equal to 80 percent, shrinkage set to 0.01, the number of trees for each boosted tree ensemble set to 500, each with an interaction depth of 3. These parameters have been found to produce good results across a diverse set of terrestrial species.

SECTION 4: Species' Occurrence Estimation
Variation in detectability associated with observation effort was controlled by assuming that all effort predictors (search time, transect length, time of day, number of observers, and protocol) were constant and additively associated with the true occurrence probability. Thus, the quantity we used to estimate

species distributions was defined as the probability that a typical eBird participant would detect the species on a search from 7-8AM while traveling 1 km on the given day of 2011. This quantity is a relative measure of species occupation, to the degree that variation in detection rate has been corrected.

The occurrence estimates account for variation in detection rates associated with observation effort, as discussed in the above paragraph. Occurrence estimates were developed individually for each base-model in the STEM ensemble. However, this does not address the issue of seasonal variation in detection rates associated with effort. This means that temporal patterns of occurrence may reflect variation in detection rates.

For each species we calculated one daily occurrence estimate per week for all 52 weeks. The distribution surface was estimated based on a series of predictions designed to uniformly represent the study area. Because of the large study area relative to the spatial resolution of model predictors, it is computationally expensive to estimate the distributions with very fine spatial resolution.

To reduce the computational cost, we created two geographically Stratified Random Designs (SRD) to estimate distributions. The "30km SRD" consists of 2,219 locations within New York State generated with 15 locations sampled uniformly from each ~30-km pixel (strata) in a regular grid. The "3km SRD" consists of 15,805 locations within New York State generated with one location sampled uniformly from each 3km pixel (strata). The analysis, including time replicates and bootstrapping, took approximately 150 hours per species.

NOTE: Because SRD locations do not lie on a regular grid, it is necessary to use a pixel size a little larger than the nominal size of the design to "fill in" the complete extent for visualization of distributions. For example, distribution estimates based on the 30-km SRD were plotted with ~ 34-km pixels and 3-km SRD estimates were plotted with ~6-km pixels.

Section 4.1: Kriged County Estimates
To refine the spatial scale at which we can estimate and visualize distributions, we kriged the predictions at the 3-km and 30-km SRD locations using the fields package in R (Fields Development Team 2006). Each map includes an estimate of the Effective Spatial Range (ESR), defined here as $-\log(0.05)\lambda.$, where $\lambda.$ is the estimated range parameter for the exponential covariance model. The ESR is an estimate of the distance at which spatial correlation drops to under 5% of its maximum.

The ESR is often used as a description of the distance between "independent" observations and a measure of the amount of smoothing performed by the kriging operation. Large values mean that the smoothing is taking place across large regions. This is expected when the distribution surface is relatively flat or smooth, such as when a species is uniformly absent or mostly absent (black). Similarly, small values of the ESR mean that the estimated distribution exhibits fine-scale spatial structure. Thus, it is expected that ESR will change seasonally as the distributions of species change.

SECTION 5: Predictive Performance Assessment
Species distribution models of prevalence are commonly used to estimate two related, but distinct quantities. The first quantity is the binary response, presence-absence, often coded as 0-1. For example, we used STEM to predict if a new eBird search will detect or not detect a given species, at a given location, conducted under specified search conditions (e.g. over 1 km from 7- 8am). The second quantity is the expected occurrence rate. This quantity can be thought of as a long-run average frequency of occurrence for a given species, at a given location, conducted under specified search conditions. We

refer to the expected occurrence as the predicted probability of occurrence, a continuous measure from 0 to 1.

There are a number of different metrics used to describe model predictive performance for binary classification for both of these inferential quantities. In an empirical study, Caruana and Niculescu-Mizil (2004) found that most metrics fall into a small number of distinct groups, suggesting that they measure different aspects of performance. As part of the workflow we computed estimates of six different numerical summaries of predictive performance. AUC and RMSE measure the ability of the model to estimate the expected occurrence rates. Sensitivity, Specificity, Accuracy, and the Kappa statistics are calculated to measure the ability of the model to estimate the binary outcome for a new observation. All of these statistics were calculated using an independent validation set, sub-sampled to help control for the uneven spatial distribution of the validation data (Fink et al 2010). We used the spatial extent of the entire study area to compute predictive performance.

Because of the strong variation in seasonal distributions of migrant species, annual summaries are not sensitive enough to be useful diagnostics of model performance. Therefore, we opted to compute monthly predictive performance metrics in addition to annual measures. Annual and monthly predictive performance was computed for all statistics, except RMSE, using the R package PresenceAbsence (Freedman and Moisen, 2008). RMSE was computed only for the monthly assessment.

SECTION 5.1: Assessing Monthly Occurrence Rate Estimates
Many applications require one to accurately rank occurrence rates. The Area Under the Curve (AUC) measures a model's ability to discriminate between positive and negative observations (Fielding and Bell 1997). The AUC is equal to the probability that the model will rank a randomly chosen positive observation higher than a randomly chosen negative one. Thus, AUC depends only on the ranking of the predictions.

The AUC statistic is formally related to several other familiar statistics. It can be shown that the area under the ROC curve is closely related to the Mann–Whitney U, which tests whether positives are ranked higher than negatives. It is also equivalent to the Wilcoxon test of ranks. The AUC is related to the Gini coefficient (G1) by the following formula G1 + 1 = 2AUC. The AUC statistic ranges from 1.0 to zero. When AUC equals 1 it indicates perfect discrimination, a value of 0.5 indicates random discrimination, and values less than 0.5 are worse than random (Figure 2).
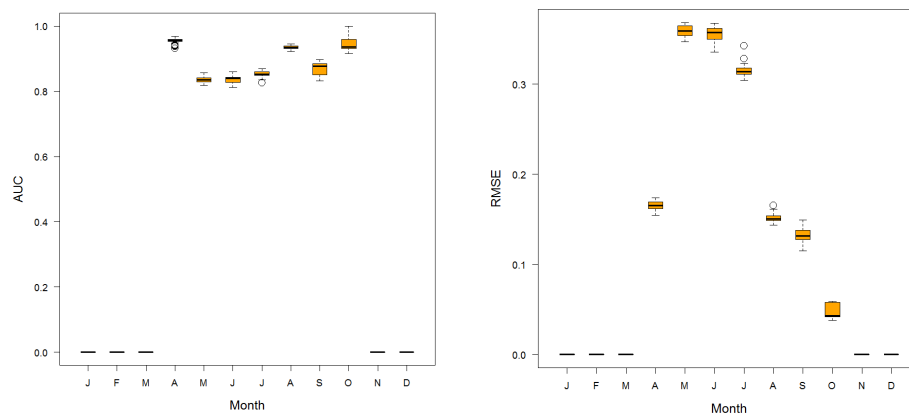
Figure 2. Assessing monthly occurrence rate estimates using 1-AUC and RMSE statistics. Example graphs are shown for Wood Thrush.

The AUC of a model gives no indication of how close the actual predicted probabilities are to the observed outcome. To take this aspect of a model's performance into account, we have to look at the calibration of the models. Calibration refers to the agreement between predicted probabilities and the "true probabilities." Several measures are commonly used to measure model calibration. We computed the Root Mean Squared Error (RMSE) between observed presence/absence coded as 0 and 1, and the predicted probabilities on the validation set. For measuring differences between binary observations and predicted probability scores, the RMSE ranges from 1.0 (worst) to 0.0 for a set of perfect predictions.

SECTION 5.2: Assessing Monthly Presence-Absence Predictions
Predicting the presence or absence of a species at a given location and time provides a basic source of ecological information. It is used to identify suitable habitat for species or to delineate a species' range. Here we use "range" to specifically mean the boundary that distinguishes from contiguous regions occupied by the species during searches from regions where the species is absent.

Given predictions of presence and absence, the binary response, one can summarize model performance by comparing the relative number of correct predictions and erroneous predictions for observations that were truly present and absent. We report Sensitivity, Specificity, and the Percent Correctly Classified (PCC) statistics. Sensitivity, also called recall rate in some fields, measures the proportion of positives that are correctly identified. Specificity measures the proportion of negatives that are correctly identified. PCC is the percentage of both positive and negative correctly classified. Note that these three statistics will vary with the background prevalence of positives in the training data (see Figure 3 for examples).

The Kappa statistic is another commonly used measure of predictive performance designed to take into account chance agreement. Perfect classification produces a Kappa = 1.0 and negative values indicate performance worse than chance alone. Values above 0.0 indicate predictive power better than chance. Despite being designed to take into account background prevalence, Kappa is known to vary with prevalence, with higher Kappa scores as classes become closer to equi-probable (Sim and Wright, 2005). Using predictive performance metrics to compare distributions that have different background prevalence remains a challenge. It is not yet clear how to compare metrics that themselves vary with prevalence.

7
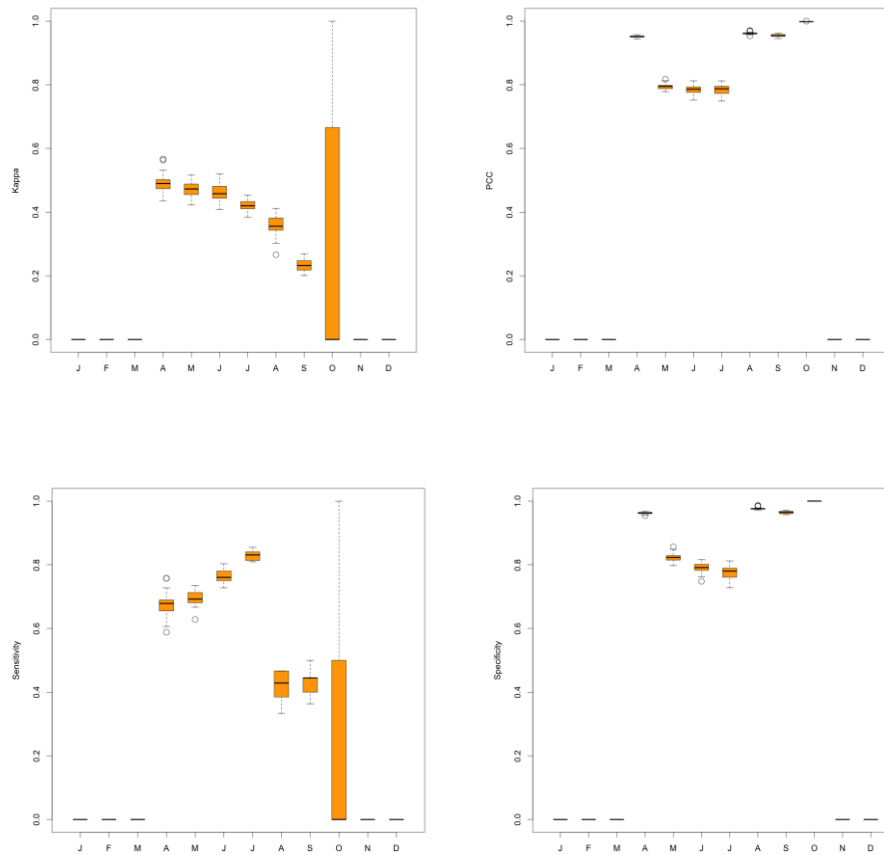
Figure 3. Monthly presence-absence predictions were assessed using Sensitivity, Specificity, Percent Correctly Classified, and Kappa Statistics. Example graphs are shown for Wood Thrush.

SECTION 5.3: Estimating Presence-Absence from Expected Occurrence Rate
To estimate locations where species are present or absent, we computed the threshold for probability of occurrence that maximizes kappa based on the validation data for each month. These thresholds vary with season, providing a way to make inference about occurrence corrected for seasonal variation in detection rates.

STEM estimates of the expected occurrence rate are provided as a value between 0 and 1. These scores can be interpreted as a measure of habitat suitability: points with the greatest likelihood of having suitable vs. non-suitable habitat. To convert these probabilities to a binary 0-1 outcome we used a simple threshold rule: the prediction is said to be present and given a value of 1 when the predicted probability is above the threshold and 0 otherwise. Generally speaking, the threshold divides points where presence is most likely to be correctly classified (greatest probability of the species being observed) and points where absence is most likely to be correctly classified (greatest probability of the species not being observed).

Thresholds can be calculated in various ways; the main difference among methods is in tradeoff between false negatives and false positives. As part of the predictive performance assessment (above), the threshold that maximizes the kappa statistic across the spatial extent of study was estimated for

8

each month of the year. To do this, we calculated kappa across a range of threshold values, and the one that maximized kappa was retained. By taking into account prevalence, Kappa strikes a balance between both types of errors averaged across the spatial extent and evaluation month.

The monthly threshold estimates can be used to estimate thresholds at any time throughout the year by smoothing them. Temporal variation in species prevalence across the study area is expected to be high because of the seasonal variability in occurrence for migratory species, and also because of seasonal variation in detectability. For this reason, we opted to smooth the monthly estimates using the R package akima version 0.5-7 (Akima 1978).

SECTION 6: Partial Dependence
Partial dependence (PD) estimates are low dimensional projections of the predictors. They are calculated by marginalizing for select predictors while integrating out the effects of all other predictors (see Hastie et al. 2009, Section 10.13). Because PD summarizes the property of a single model, the minimum spatiotemporal scale of these statistics is determined by the extent of the model. We calculated PD estimates for the additive effects of all predictors for each base model in the ensemble. These were computed using the plot.gbm() function in the gbm package.

The goal is to estimate PD within the selected region of interest and quantify the variation in this estimate across bootstrap replicates. This uncertainty estimate is important because it can be used to identify predictors with important additive effects and to compare local predictor effects. To estimate the variation in the expected PD estimate requires that we take into account the nested sampling design underlying the STEM estimator. To do this, we have implemented a two-step analysis procedure to take into account the nested effects of random partitions within each bootstrap replicate. In the first step, a Generalized Additive Model (GAM) (Wood 2006) is used to average out partition-level variation by estimating PD for each individual bootstrap replicate based on all partition replicates for that bootstrap sample. In the second step, another GAM is fit to estimate the conditional mean across bootstrap replicates.

In the second step, we computed the prediction intervals to capture variation across bootstrap replicates. The prediction intervals estimate the amount of variation expected from a PD estimate based on a new bootstrap replicate. Thus, the prediction intervals provide a conservative estimate of variation when making inference about the conditional mean estimate. PD plots showing species specific relationships to NLCD land cover types have been provided with the overall package of STEM data products.

Section 7: Qualitative Biological Interpretation of Species-specific Models

Below, based on expert human review, we provide a qualitative biological interpretation of the model results for each species. These interpretations, along with with performance metrics for each species, (see Section 5 and previously submitted package of data products) provide the basis for including/excluding species for "stacked modeling," and ultimately for inclusion in the online mapping tool.

A note on detectability—Because of behavior and plumage patterns, many bird species can be more difficult to detect during fall versus spring migration. In fall, males rarely sing and instead give chip notes that are difficult to identify. Both sexes spend much time foraging in heavy cover and are thus less conspicuous. Warblers, often in non-breeding plumage, form mixed species flocks, making it confusing

9

to make accurate identifications. This difference in detectability is relative within the eBird dataset and does not diminish the overall quality of the data for a given species. Poor detectability during fall can; however, result in underestimating the spatial extent of some species distributions. In the results that follow, we make light of species' fall migration distributions that potentially are impacted by detectability.

Migratory Breeding Birds

Forest Associates

Spring **Wood Thrush** migration shows major activity in the lower Hudson Valley and the Southern Tier region. This is certainly accurate, though results also likely reflect high observer effort in those areas. Fall detection is likely to be biased by poor detection rates. While the Southern Tier and St. Lawrence Valley are clearly important in fall for migrating Wood Thrush, it's likely that other heavily forested areas in the state are important as well.

Veery is likely similar to Wood Thrush in its patterns of spring and fall migration. Easterly bias in spring is likely real, and could easily represent many birds moving up the Appalachians. However, the fall pattern is likely biased by poor detection rates, although there is probably a strong movement in the Appalachians. Numerous accounts and field recordings agree that high densities of thrushes, including Veery, do pass through this region. It's likely that other areas have similarly high fall migration importance, but are not appearing because of detection issues.

Ovenbird spring hotspots in the Shawangunk Mountains and Catskills are probably reflective of breeding rather than stopover. Broad-scale patterns of occurrence in this species suggest that it arrives en masse in many areas and passes through without being constrained by any particular vegetation community structure (given the extent of the passage). Fall patterns are likely driven by challenges in detection, as the widespread but sparse distribution suggests that more individuals are likely present than are being detected.

Scarlet Tanager spring stopover locations seem to mirror very closely the species' forested breeding habitat, suggesting that similar landscapes are important for stopover. Fall hotspots in the Appalachians are probably real, with birds potentially using that area as a passage point for movements south. Eastern Lake Ontario plains contain important locations as well. The dearth of Adirondack occurrences may reflect lack of observers as much as lack of appropriate stopover habitat.

Shrubland/Young Forest Associates

Chestnut-sided Warbler seems to occur en masse in spring and fall migration in many parts of the state. Other than where lacking in some portions of the west, northwest, and Hudson Valley; the species seems to occur in many regions without showing preference for specific hotspots. However, spring standouts do appear in the Tug Hill and in the Southern Tier just north of the Pennsylvania border; in the fall, northeastern areas and some similar pre-Appalachian areas light up.

The explosion of **Indigo Bunting** in spring presumably represents a rapid arrival of this species from coastal and interior locations at once, which is likely a function of the species' passage from widespread origins to the south. The fall pattern suggests a much greater bias to the central portion of the state, mostly west of the Catskills and possibly aligned with movements that focus on birds departing the northeast via an Appalachian route. However, there is also a tendency to follow the Hudson River valley, similar to the pattern we see in Yellow Warbler, though not as extensive.

Yellow Warbler is a common migrant across the state, with concentrations in shrub and forested wetlands, including river valleys like the Hudson River valley. These hotspots are probably real corridors that this species uses for passage, whether for staging or simply for targeting specific habitats. The fall maps highlight the importance of the Hudson River valley, the St. Lawrence River valley, and areas in the Finger Lakes and Lake Erie/Ontario shores. Additionally, some areas in southeastern New York might be important.

Grassland Associates

Bobolink maps suggest that non-coastal, open habitats are used for stopover and probably passage. Areas that light up in the spring are widespread in the center and eastern portions of the state, corresponding to movements around large mountain ranges. Presumably this species expands into the area from Appalachian based movements, with the mountains limiting birds to easterly and westerly routes around them. Similarly, in the fall, it appears that birds might be taking more westerly paths through the state, moving west of the Adirondacks and south toward the Appalachians.

Eastern Meadowlark passage is probably not well represented by the STEM models; rather the maps show areas of annual importance for this species. Passage for this species is nocturnal, and though diurnal flights do occur, most detections are probably focused on non-stopover reports. Regardless, open habitats are critical for this species, and the pattern of occurrence during breeding likely mirrors the pattern of occurrence during migration. Separating the two would be difficult, but is presumably unnecessary.

Savannah Sparrow is a widespread, common and facultative nocturnal migrant. Finding important migratory stopover during spring may be challenging because breeding and migrating individuals are often mixed over the same spatial extent. However, fall movements clearly indicate a substantially greater presence for the species in the coastal area, presumably representing migration activity.

Fly-through Migrants
American Tree Sparrow spring patterns suggest a much broader passage across much of the state, with the Southern Tier and southern Adirondacks appearing to have higher prevalence. This pattern may be driven by rapid departure, thus mirroring more of winter range and prevalence than migratory stopover. In fall the species shows affinity toward lake shore occurrence, which may reflect recent fall arrivals. The St. Lawrence Valley and Lake Ontario shore seem to be important, which would follow with birds moving south and east across Canada and passing over the lake to points just offshore.

In spring **White-crowned Sparrow** shows a pattern of westerly movements, with most wintering birds pulling out quickly and migrating in a primarily lakeshore-driven pattern. In fall this species arrives en masse. Hotspots along the lakeshore and the Hudson Valley are likely accurate, as this region contains open space and the potential for topographically driven concentration points. Additionally, some southeastern hotspots, like the Shawangunk grasslands, appear to be important.

White-throated Sparrow is a common migrant and wintering species, and separating what is stopover habitat from wintering areas as well as breeding areas in some places is challenging. Hotspots in the southern Finger Lakes, Hudson River Valley, Ontario lakeshore, and NYC urban metroplex are likely important areas in both spring and fall. The species is a common migrant, and some of these areas, particularly urban areas, may see huge numbers of this species move on single nights in spring and fall.

Swainson's Thrush exhibits a pattern that is probably driven, in part, by breeding distribution in the Northeastern portion of the state and by potential biases in detection probability. That said, it is likely a real pattern that the terminus of the Appalachians light up, as this is a high traffic area for the species in fall based on acoustic and observer data. Spring patterns may also be biased by observer behavior and a generally low occurrence frequency, but some locations in the western and southern parts of the state are certainly stopover locations. However, the STEM maps probably do not do a comprehensive job of describing where and when the species stops over and the habitat in which they stop over.

Orange-crowned Warbler is a primarily a western species that occurs in New York in the fall. Its occurrence is concentrated mostly in the St. Lawrence Valley, and Lake Erie and Ontario lakeshores. Although data are scant, this pattern of fall occurrence is likely accurate, with Orange-crowned much less common in spring than in fall. Furthermore, habitat types associated with the lakeshore are probably relevant, as this species commonly uses open, shrubby and wet habitats. Note: because of small sample size, no STEM data products are available for Orange-crowned Warbler in spring.

Canada Warbler maps for both seasons are probably biased by poor detection rates, and it appears that there may be a strong bias to detecting stopover locations that are also tightly correlated with breeding locations. However, the Adirondack to Appalachian connection is probably real, with birds using this pathway, though the importance of other areas for stopover is probably missed given low detection rates.

Tennessee Warbler accurately shows a western bias in the spring. Fall hotspots include the northern Appalachians, Catskills, St Lawrence Valley, and Champlain regions. These hotspots are probably correct and represent a more inland route traveled by this species, as individuals return via trans-Gulf movements and circum-Gulf movements to non-breeding areas. Additional locations may be important, as the species is probably underrepresented in these models given the low probabilities of occurrence.

Wilson's Warbler is difficult to detect in the Eastern U.S. during migration. Hotspots depicted for this species represent stopover locations in wet, shrubby and somewhat open habitats. The eastern population is increasingly rare, and although detection is probably an issue, the model appears representative of the species' scattered distribution as it passes through New York. The more western slant to fall migration is likely correct, and the lower likelihood of detection in the spring is likely a function of more westerly movements, typically west of the Appalachians.

Blackpoll Warbler typically migrates over the Atlantic Ocean in the fall, and the pattern we see in New York reveals birds moving east of the Appalachians and departing to the ocean in southern New England and northern mid-Atlantic areas, and birds moving down the Appalachians and departing in the southern mid-Atlantic. The low probability of occurrence in the St. Lawrence is probably a detection issue, but the species prevalence in mountainous areas in the eastern part of the state is accurate, as is the occurrence in the western part of the state. Spring patterns are interesting, though some reflect low detection rates rather than an actual distribution of stopover importance. Adirondack birds are probably breeders. Species passage suggests perhaps that multiple populations of birds may well be passing through the state.

Olive-sided Flycatcher is a long-distance migrant, and the present maps are highly biased toward breeding distribution rather than stopover. Considering the species' use of snags, many habitat types should appear as potentially important stopover locations. However, detection is certainly an issue. It is interesting to note that fall occurrence is higher than background in southwestern New York, potentially

indicating some sort of movement along the Appalachians. However, this is a weak pattern at best, as it is unclear how this species moves to and from South America. In the spring, some additional locations pop out in the southeastern part of New York, but whether this is a function of a different migration strategy in the spring is unclear. It is interesting that Tug Hill pops out more in the spring, something that may be useful in considering whether in fact this is a hotspot for migrants or part of the breeding distribution.

Shorebirds

Least Sandpiper follows a similar pattern to other widespread, non-coastal shorebird migrants. Areas that are wet and have exposed flats and mud appear as hotspots for the species. This is more typical for fall than spring, presumably because of more birds passing through the region after breeding, but also conceivably as a function of distinct spring and fall biases in migration.

Lesser Yellowlegs follows the pattern of non-coastal shorebirds like Least Sandpiper and Greater Yellowlegs. This species appears in wet areas with exposed land or mud for feeding. Similar spring and fall hotspots appear and these focus squarely on specific wet habitats. One important thing to note--it is possible that some shorebirds are using stopover in a very specific way, more for staging, but we are not yet able to tease this apart in our data presentations.

Greater Yellowlegs appears in typical shorebird-friendly habitats in fall, with hotspots appearing throughout wetlands in the Finger Lakes and west, as well as the St. Lawrence Valley and in many marsh habitats in Long Island and coastal southeastern New York. A similar but weaker pattern appears in the spring. This, along with other shorebirds, is a clear case in which stopover habitat basically defines the distribution of the species, and in this case, open wetlands with some exposed mudflats and agricultural areas with open wet space dominate.

For **Dunlin**, coastal Long Island and large wetland areas are important. For all shorebirds, presumably, certain mudflat habitats are going to be critical, but this map does not do a good job of showing all potential places where the species might be, particularly in the Lake Champlain, Lake Ontario, and Finger Lakes regions. That said, the locations that appear in the binary maps are reasonable and likely important, rather than some spurious location. Long Island is certainly the most likely stopover location for this species.

Super Colliders

Fox Sparrow shows a series of hotspots in fall stopover in the Catskills and Tug Hill plateau, which probably correspond to mixed forest and shrubby habitats from the Shawangunks up to the upper Hudson River valley. The Catskills show up strongly in the spring, and less so but still in an obvious manner in the western Lake Ontario and Adirondack regions.

Golden-winged Warbler appears in the expected areas of high breeding density, and given its rarity as a migrant, this is presumably correct. This species probably uses Appalachian to trans-Gulf of Mexico strategies for moving between breeding and nonbreeding areas, and the hotspots that appear are more likely biased toward breeding sites than stopover habitat. The combination of rarity and potential challenges with coverage make for low confidence in anything other than breeding hotspots.

Black-throated Blue Warbler shows hotspots in the northeastern part of New York State, a pattern that is probably real as birds move south from Canada. But this is likely biased a bit by breeding birds that linger, as other locations to the south might be more indicative of actual passage points and stopover.

The Adirondacks and Catskills, however, do seem to provide important stopover locations. In spring, a similar pattern occurs, and both spring and fall suggest that there is a substantial Appalachian component to the stopover and passage patterns for this species. This fits with what we know about the arrival of the species in the U.S. from the Caribbean.

Bay-breasted Warbler patterns in the fall suggest a passage through the state that links the St. Lawrence and Appalachians, which may be a relevant path. This species is probably using the Appalachians and mountains east of there as a passage strategy. However, note that low prevalence is true for this species, suggesting that some other patterns that may exist are not getting detected. Spring patterns mirror the Appalachian-centric view of this species' movements, suggesting that birds enter New York from farther west. This presumably follows with a primarily trans-Gulf flight of birds moving west of and up through the Appalachians as they move north.

Section 7.1: Discussion

As of this writing, the Cornell Lab of Ornithology is working with The Nature Conservancy to develop a standardized approach to evaluating the biological accuracy and statistical integrity of each species' model. From this evaluation, a subset of species will be selected for inclusion in subsequent modeling efforts that are being conducted by The Nature Conservancy.

We are confident that the 28 species distribution models that we have provided can be combined to help understand landscape-scale patterns of bird migration and the use of stopover locations; however, it is beyond the scope of this report (and the deliverables of this sub-contract) to make predictions about how the final stacked models and online mapping tool will inform the siting of new wind energy in New York. That said, the Cornell Lab of Ornithology is eager to see these models used for applied conservation and to that end, we are very open to continued collaboration with The Nature Conservancy and NYSERDA to help develop publications and "end-user" tools that contribute to science and conservation in New York State.

References:
Akima, Hiroshi "A Method of Bivariate Interpolation and Smooth Surface Fitting for Irregularly Distributed Data Points", ACM Transactions on Mathematical Software, Vol. 4, No. 2, June 1978, pp. 148-159.
Caruana, R. and Niculescu-Mizil, A. 2004 "Data Mining in Metric Space." International Conference on Knowledge Discovery and Data Mining (KDD).
Diefenbach, D. R., M. R. Marshall, J. A. Mattice, and D. W.Brauning. 2007. Incorporating availability for detection in estimates of bird abundance. Auk 124:96–106.
Fields Development Team (2006). fields: Tools for Spatial Data. National Center for Atmospheric Research, Boulder, CO. URL http://www.cgd.ucar.edu/Software/Fields.
Fink D, Hochachka W. 2009. Gaussian semiparametric analysis using hierarchical predictive models. Pages 1011–1035 in Thomson D, Cooch E, Conroy M, eds. Modeling Demographic Processes in Marked Populations. Springer. doi:10.1007/978-0-387-78151-8
Freeman, E.A., Moisen, G.2008 PresenceAbsence: An R Package for Presence Absence Analysis. Journal of Statistical Software. Vol. 23, Issue 11
Hastie, T., R. Tibshirani, and J. Friedman. 2009. The elements of statistical learning: data mining, inference, and prediction, 2nd edition. Springer Verlag, New York, New York, USA.

Munson, M. A., K. Webb, Sheldon, D. , Fink, D., Hochachka, W.M., Iliff, M., Riedewald, M., Sorokina, D., Sullivan, B., Wood, C., and Kelling, S. 2010. The eBird Reference Dataset 2.0. (http://www.avianknowledge.net/content/features/archive/eBird_Ref)."

North American Bird Conservation Initiative, (NABCI) U.S. Committee, 2011. The State of the Birds 2011 Report on Public Lands and Waters. U.S. Department of Interior: Washington, DC. 48 pages.

R Development Core Team. 2011. R: a language and environment for statistical computing. R Foundation for Statistical Computing, Vienna, Austria. http://www.R-project.org

Ridgeway G. 2010. gbm: Generalized Boosted Regression Models. Version 1.6-3.1. Available from http://CRAN.R-project.org/package=gbm

Sim J; & Wright, C. C (2005). "The Kappa Statistic in Reliability Studies: Use, Interpretation, and Sample Size Requirements". Physical Therapy 85: 257–268.

Sullivan, B., Wood, C. Iliff, M. J. Bonney, R. E. Fink, D. and Kelling, S. 2009. eBird: A Citizen-based Bird Observation Network in the Biological Sciences: Biological Conservation 142, 2282–2292.

Wood, S.N.(2006) *Generalized additive models: an introduction with R*. Chapman & Hall/CRC