The eBird Reference Dataset, Version 4.0

M. Arthur Munson†, Kevin Webb‡, Daniel Sheldon†,
Daniel Fink‡, Wesley M. Hochachka‡, Marshall Iliff‡,
Mirek Riedewald*, Daria Sorokina**, Brian Sullivan‡,
Christopher Wood‡, Steve Kelling‡

† Cornell University Computer Science Department, Ithaca, NY 14853
‡ Cornell Lab of Ornithology, Ithaca, NY 14850
* Northeastern University, Boston, MA 02115
** Carnegie Mellon University, Pittsburgh, PA 15213

October 1, 2012

Abstract

This document describes the eBird reference data set and the processing steps taken during creation. We hope this data will be a useful resource for studying avian dynamics and for developing new ecological modeling techniques.

1 Usage and Copyright

The eBird reference data is freely available for all usages. The observational data included in the data set have data access level 5 in the Avian Knowledge Network (AKN) data warehouse and are published here in compliance with the AKN data sharing policy.[1] eBird[2] is run by the National Audubon Society[3] and the Cornell Lab of Ornithology[4], and the data is copyrighted by both organizations. A primary goal in publishing this data is to provide a common data resource for studying and comparing ecological models; as such, derivative versions of the eBird reference data set must not be distributed without explicit permission from the copyright holders.

The data set is a snapshot of submitted observations for years prior to 2012 that were submitted to eBird and reviewed by January 1, 2012. Observations with an observation date later than December 31, 2011 were not included.

Published results using this data should cite this document as follows:

> M. Arthur Munson, Kevin Webb, Daniel Sheldon, Daniel Fink, Wesley M. Hochachka, Marshall Iliff, Mirek Riedewald, Daria Sorokina, Brian Sullivan, Christopher Wood, and Steve Kelling. *The eBird Reference Dataset, Version 4.0.* Cornell Lab of Ornithology and National Audubon Society, Ithaca, NY, January 2012.

Please direct any questions to:

> M. Arthur Munson at `mmunson@cs.cornell.edu` and
> Steve Kelling at `stk2@cornell.edu`

---

[1] See `http://www.avianknowledge.net/content/about/akn-data-sharing-policy` for details.
[2] `http://ebird.org`
[3] `http://www.audubon.org`
[4] `http://www.birds.cornell.edu`

2 Introduction

This dataset contains count data for bird species observed by novice and experienced bird observers (a.k.a. *birders*). The data was submitted by volunteers to the eBird Citizen Science Project, run by the Cornell Lab of Ornithology and the National Audubon Society. A record in this dataset corresponds to a *checklist* that a birder uses to mark the number of birds of each species detected; one checklist is submitted per sampling event (i.e. birding session). Each checklist submitted from the 48 states in the contiguous United States is additionally annotated with hundreds of predictor variables (called covariates below) that are derived from the location of the sampling event.

One pragmatic note: Excel is unable to handle the larger data files in this dataset. The data from year 2008 contains more records than Excel supports (rows were truncated around 175,000); previous years contain fewer records. In our experiments with Excel, columns were truncated from the US48 checklist files and the extended covariates files.

Finally, the *stratified_random_design* (SRD) dataset is a companion dataset to the eBird reference dataset with a similar format. The SRD data contains the predictor values for random locations within the contiguous United States, locations were chosen using a stratified random design. As of the 2.1 release of the ERD there are 2 SRD data sets available for use, one set of SRD point sets, roughly based on a 30km grid, contains around 130,000 random locations, and the second data set contains approximately 933,700 random locations roughly based on a 3km grid of the lower 48 states. The SRD data can be combined with an ecological model to generate maps of predicted surfaces. In brief, use the model to predict at each location in the SRD data, and then plot the predicted values overlaid with geographical boundaries. Helpful tools for plotting maps include Matlab and R. The SRD data can be downloaded from the AKN: `http://www.avianknowledge.net/content/download/packages/reference_data`.

3 Dataset Organization

The eBird reference data set actually consists of two separate datasets:

US48 — all checklists from the 48 states in the contiguous US, with detailed location-based covariate information.

AMERICAS — all checklists from the 48 states in the contiguous US. No location-based covariates are provided.

The two datasets are organized in a similar fashion, and observations are *repeated* between them. That is, all checklists included in the US48 dataset are also included in the AMERICAS dataset.

3.1 Column Summaries

The data is organized by checklist, with one row per checklist. The columns are split into three groups:

- Checklist description. 16 columns describing the sampling event (date, time, lat, long, duration, distance traveled, protocol, observer id, etc.), plus variables for all the species in the data set. In the US48 data set there is one column per species, with each column listing the number of birds observed. In the AMERICAS data set the observation count data uses a sparse representation, and only the observed species are listed (all within a single column called SPECIESLIST). See sections 4.2 and 4.3 for more details.

- Core covariates (US48 only). Approximately 46 columns that we believe a priori are most important for most of the species. These are suggested as a starting point for analyses. Details are in section 4.4.

- Extended covariates (US48 only). These include extensive statistics about habitat configuration, fine-grained climate measurements, and observer expertise. See section 4.5 and Tables 3 and 4.

In the US48 dataset, all of the files are row aligned. That is, the covariates for the sampling event in line 2 of `checklists.csv` are in lines 2 of `core_covariates.csv` and `extended_covariates.csv`.

3.2 File Organization and Directory Structure

The checklists are ordered chronologically by year; within a year they are ordered by the observation location. Files are split into years for manageability. Data for a given year are organized in a single directory containing three row-aligned files, each containing one of the column groups described above. For example, the directory structure of the US48 dataset looks like this:

us48/2010:

```
checklists.csv
core_covariates.csv
extended_covariates.csv
```

us48/2011:

```
checklists.csv
core_covariates.csv
extended_covariates.csv
```

The structure of the AMERICAS dataset is similar except it excludes the covariates.

americas/2010:

```
checklists.csv
```

americas/2011:

```
checklists.csv
```

In addition, there is also a `docs` folder containing this document and text files describing the names and types of all data columns:

docs:

```
checklists.names
core_covariates.names       [US48 only]
extended_covariates.names   [US48 only]
taxonomy.csv
speciesfreq.csv
```

See section 4.1 for the names files' format.

The data set uses scientific names (i.e. latin names) to name species. The file `taxonomy.csv` lists all of the scientific names and the corresponding common names in english, as well as taxonomic codes for the species. The first line of the file contains the names of the columns.

Finally, how often each species is observed is given in `speciesfreq.csv`. The numbers are the percentage of checklists containing the species for all years up through December 2011. In US48, these frequencies determine the order of the species' columns in `checklists.csv`.

4 Data Set Description

The data is written in CSV format with one line (data record) per checklist. All data files include the column SAMPLING_EVENT_ID, a unique identifier for each checklist. This can be used to verify that all column groups are correctly aligned. The first line of every file contains the names of the columns.

Two special values are used in the dataset. When the value for a variable is missing / unknown, it is represented as `?`. Second, when a covariate measurement is not applicable to the context of an observation,

it is represented as `NA`. How to deal with these special cases depends on how the data is used and the analysis tools, and is best decided by the data analyst.

The rest of this section describes the names file format and the different variable groups.

4.1 Names File Format

Each names file lists the names and types of all variables in the group. One variable is listed per line, and the order of variables matches the order of data columns in the corresponding CSV file. Variables can be `string`, `continuous`, or `nominal` valued. Example variable descriptions are

```
SAMPLING_EVENT_ID: string.
ELEVATION: continuous.
BCR: 1,2,3,...,35.
```

The variable name precedes the colon, and type information comes after. Continuous variables are real- or integer-valued measurements. Nominal variables are categorical, and can take one of a small set of values; the allowed values are listed after the variable name. String variables are high-arity nominal-valued variables where a legal value is a sequence of letters or numbers (no whitespace). Generally string variables should not be used for predictive modeling, but they are useful for fitting random effects and for data provenance.

4.2 Species Counts

All species reported in complete checklists to eBird are included in the reference dataset. Species reported outside the contiguous US are not included in the US48 dataset. When a birder submits a *complete checklist*, they are reporting on all birds that they were able to identify. By only including complete checklists in the dataset, one can assume that if a species is not reported it was not present.

Casual count checklists, and random count checklists are also included in this release. *Casual counts* are observations made while birding was not the submitter's primary activity. *Random counts* are observations made at a randomly selected location over a period of at least five minutes. The addition of *complete casual counts*, and *random counts* provide a varied set of observation protocols for research.

In the US48 dataset, there is one column per species, containing the count of how many were observed or a 0 if the species was unobserved (the zero-filling relies on the above assumption).

If the species was reported as present without a count, the count is replaced with an `X`. Birders often use present-without-count if they do not have a good idea how many they detected (perhaps because there were many individuals of that species) or if they did not want to bother with counting the species (i.e. if the species is not particularly interesting to the birder). These records contain useful information about the presence of a species, but should probably be discarded when modeling species abundance.

In most cases, a model should be fit using a single species variable as the response variable. Unless the research question involves relationships between bird species, the remaining species variables should be ignored during the modeling.

Researchers wondering which species to choose for modeling may wish to refer to the species matrix table available at `http://www.avianknowledge.net/content/files/MatrixV1.xls`. This table lists a few dozen species and categorizes them along multiple axes (e.g. migration and population trends) based on domain expert opinion. A copy of the file is also included in the `docs/` folder.

4.3 Sampling Event Covariates

Each data record contains information describing when, where, and how the observations were made, as well as a unique identifier for the checklist. Table 1 lists the covariates tied to the sampling event.

Location is highly correlated with many covariates describing the environment. Further, the pair LATITUDE-LONGITUDE is highly correlated with the sampling event. For these reasons, the decision to include latitude and longitude as predictor variables should be made carefully.

Table 1: Summary of sampling event covariates.

Variable Name	Comments
SAMPLING_EVENT_ID	Unique identifier for each data sample / checklist.
LATITUDE	Decimal latitude. Location is tied to starting position of traveling counts. Datum = WGS 84.
LONGITUDE	Decimal longitude. Datum = WGS 84.
COUNT_TYPE	What kind of observation the sample is: stationary (P21), traveling (P22, P34), area (P23, P35), casual (P20), or random (P48). Protocol P34 is a small amount of data contributed from the Rocky Mountain Bird Observatory that we believe is high quality. Protocol P35 data are back-yard area counts made on consecutive days (see `http://www.birds.cornell.edu/MyYardCounts`).
COUNTRY	Full name of country / political unit where observation took place. Useful for extracting sub-portions of the data.
STATE_PROVINCE	Name of the state, province, or region where observation was made. Useful for extracting sub-portions of the data.
YEAR	
MONTH	Month of the year, ranging from 01 through 12. Useful for extracting sub-portions of the data.
DAY	Day of the year, ranging from 1 through 366.
TIME	Time when observation started, ranging over $[0, 24)$. Fractional hours represent minutes (e.g. $13.5 = 1$:30PM). Times are local times (including daylight savings when/where appropriate).
EFFORT_HRS	Duration of observation for the checklist, in hours.
EFFORT_DISTANCE_KM	Distance traveled during observation period, in kilometers. Equals 0 for non-traveling counts.
EFFORT_AREA_HA	Size of survey area for area counts, in hectares. Equals 0 for non-area counts.
OBSERVER_ID	Identifier for the person who submitted the data.
NUMBER_OBSERVERS	Number of observers in the birding party.
GROUP_ID	ID string indicating a group checklist. Group/shared checklists contain summations of the highest counts for each species seen by members of the group.

4.4 Core Covariates

Table 2 summarizes the core covariates that we feel are generally useful for most species. Two elevation variables are included because a) the different resolutions serve complementary purposes, and b) sometimes one will have a measurement while the other's value is missing. See section 4.5 for details about the NLCD2001 and NLCD2006 covariates.

All of these covariates are static, and are tied to the location of the observation. The climate variables use the month when the observation is made to select the appropriate climate value from those listed in Table 3.

4.5 Static Environment Covariates

Based on an observation's location, covariate information about climate and habitat is extracted from GIS databases and joined to the checklist records. These measurements are static, in that they are derived from environmental snapshots tied to a time frame independent of when observations are made.

Table 3 summarizes the static environment covariates. Information about landcover / habitat comes from raster data that can be downloaded from `http://www.mrlc.gov/nlcd.php`.

Table 2: Summary of core covariates.

Variable Name	Comments
SAMPLING_EVENT_ID	Unique identifier for each data sample / checklist.
BCR	Bird conservation region (numeric identifier).
BAILEY_ECOREGION	Ecoregions defined by common climatic and vegetation characteristics. Details at `http://www.nationalatlas.gov/mld/ecoregp.html`
OMERNIK_L3_ECOREGION	Ecoregions defined by vegetation, animal life, geology, soils, water quality, climate, and human land use. Details at `http://www.nationalatlas.gov/mld/ecoomrp.html`
CAUS_PREC†	Mean total precipitation for month in which observation made.
CAUS_SNOW†	Mean snow depth for month in which observation made. Always missing (coded as '?') for observations in May through Sept. since no data available on snow depth from the climate atlas for these months.
CAUS_TEMP_AVG†	Mean daily average temperature for month in which observation made.
CAUS_TEMP_MIN†	Mean daily minimum temperature for month in which observation made.
CAUS_TEMP_MAX†	Mean daily maximum temperature for month in which observation made.
ELEV_GT	Elevation in meters. Horizontal resolution is roughly 1km by 1km. Source: GTOPO30 elevation dataset, acquired from USGS in 2004. Details at `http://eros.usgs.gov/products/elevation/gtopo30.php`
ELEV_NED	Elevation in meters. Horizontal resolution is roughly 30m by 30m. Source: National Elevation Dataset, acquired from USGS. Described at `http://seamless.usgs.gov/website/seamless/products/1arc.asp` and `http://www.usgsquads.com/elevationdata.htm#NED_Info`
HOUSING_DENSITY‡	Number of housing units per square mile (2000 census) for the census blockgroup containing the location.
HOUSING_PERCENT_VACANT‡	Percentage of housing units in census blockgroup vacant in 2000.
NLCDYYYY_FS_cTT_7500_PLAND	Percent of surrounding landscape that is habitat class TT. Data values extracted from source landcover rasters for years YYYY = 2001 and YYYY = 2006 are included. See text and Table 4 for more details.
POP00_SQMI‡	Population per square mile (2000 census) for the census blockgroup containing the location.

† Source: Climate Atlas of the US, v2 (1961–1990), from NOAA-NCDC. Grid cell resolution is 4km by 4km. Note that these are *climate* variables averaged over 30 years, not weather variables for the year the observation was made. Described at `http://www.ncdc.noaa.gov/oa/about/cdrom/climatls2/info/atlasad.html`

‡ Source: US 2000 census; acquired from ESRI summer 2004. `http://www.census.gov/geo/www/tiger/glossary.html`

Table 3: Summary of static environment covariates.

Variable Name	Comments
SAMPLING_EVENT_ID	Unique identifier for each data sample / checklist.
CAUS_PRECMM†	Mean total precipitation for month MM.
CAUS_SNOWMM†	Mean snow depth for month MM.
CAUS_TEMP_AVGMM†	Mean daily average temperature for month MM.
CAUS_TEMP_MINMM†	Mean daily minimum temperature for month MM.
CAUS_TEMP_MAXMM†	Mean daily maximum temperature for month MM.
CAUS_LAST_SPRING_32F_MEAN†	Last 32 F day in spring (mean). Numbers indicate date ranges.
CAUS_LAST_SPRING_32F_MEDIAN†	Last 32 F day in spring (median).
CAUS_LAST_SPRING_32F_EXTREME†	Last 32 F day in spring (extreme dates).
CAUS_FIRST_AUTUMN_32F_MEAN†	First 32 F day in autumn (mean). Numbers indicate date range.
CAUS_FIRST_AUTUMN_32F_MEDIAN†	First 32 F day in autumn (median).
CAUS_FIRST_AUTUMN_32F_EXTREME†	First 32 F day in autumn, (extreme dates).
DIST_FROM_FLOWING_FRESH‡	Distance from flowing fresh water.
DIST_IN_FLOWING_FRESH‡	Distance inside flowing fresh water.
DIST_FROM_STANDING_FRESH‡	Distance from standing fresh water.
DIST_IN_STANDING_FRESH‡	Distance inside standing fresh water.
DIST_FROM_WET_VEG_FRESH‡	Distance from wet vegetation, fresh water.
DIST_IN_WET_VEG_FRESH‡	Distance inside wet vegetation, fresh water.
DIST_FROM_FLOWING_BRACKISH‡	Distance from flowing brackish water.
DIST_IN_FLOWING_BRACKISH‡	Distance inside flowing brackish water.
DIST_FROM_STANDING_BRACKISH‡	Distance from standing brackish water.
DIST_IN_STANDING_BRACKISH‡	Distance inside standing brackish water.
DIST_FROM_WET_VEG_BRACKISH‡	Distance from wet vegetation, brackish water.
DIST_IN_WET_VEG_BRACKISH‡	Distance inside wet vegetation, brackish water.
NLCD01_CANOPYMEAN_RADRR	Mean canopy cover in square neighborhood around location with "radius" $RR \in \{75, 750, 7500\}$, measured in meters.
NLCD01_IMPERVMEAN_RADRR	Mean imperviousness to water in square neighborhood around location with "radius" $RR \in \{75, 750, 7500\}$, measured in meters.
NCLDYYYY_FS_*	Landscape and landcover statistics describing the habitat in a square neighborhood around the location. Computed from the National Land Cover Data from MRLC (www.mrlc.gov/nlcd.php) source data years YYYY = 2001 and YYYY = 2006 using the FRAGSTATS program. See text and Table 4 for more information.
SUBNATIONAL2_CODE	String encoding the state and county of the location. Useful for extracting sub-portions of the data.

† Source: Climate Atlas of the US, v2 (1961–1990), from NOAA-NCDC. Grid cell resolution is 4km by 4km. Month code 13 denotes the annual aggregate statistic. Described at http://www.ncdc.noaa.gov/oa/about/cdrom/climatls2/info/atlasad.html
‡ Source: Shapefiles provided by Idaho National Gap Analysis Program http://www.gap.uidaho.edu/bulletins/10/idaho_gap.htm. Refer to Appendix Table 5 for the description of distance values.

Ecologists generally agree that both the type of habitat and its configuration are important factors in ecological processes.[5] Consequently, we processed the raw 2001 and 2006 NLCD habitat data using the FRAGSTATS program [MCNE] to generate covariates describing landscape configurations. The configuration settings for FRAGSTATS can be found in Appendix A.1. We included the subset of FRAGSTATS statistics we felt were most likely to be informative for a variety of species across the continent.

More specifically, we extracted the landcover information from the NLCD database for a grid centered on each checklist location, creating a landcover matrix. Each landcover matrix was given as input to FRAGSTATS, which returned an array of landscape statistics summarizing the habitat neighborhood for the corresponding checklist. Since the ideal neighborhood size depends on the species under consideration, we repeated this process for three different spatial extents: 2.25 hectares, 225 hectares, and 22,500 hectares. These extents were selected to cover local ecological processes at small, medium, and large ranges. We decided not to include spatial extents large enough to cover entire migration areas because a) the computation costs would be considerable, and b) habitat configurations becomes less distinct as they are averaged over larger and larger areas. The scale of each covariate is indicated by the "radius" of the neighborhood in meters, and appears in the name of each covariate. *Radius* is actually a misnomer. The neighborhoods are square regions centered on the location. The length of the neighborhood square side is twice the "radius". In other words, the radius number is the radius for a circle inscribed within the neighborhood square.

We post-processed the FRAGSTATS covariates to recode most not-applicable (NA) values as numeric values, in almost all cases a recoding to numeric zero. This was done because the NA values actually do have a biological meaning when recoded to numeric values. For example, in a landscape that is entirely composed of grassland, FRAGSTATS would return NA values for metrics describing landcover types such as forest that were not present. However, absence of forest really does mean that, for example, there is no forest edge; zero-values are justified. Brief descriptions of when NA values were recoded are given in Table 4.

Table 4: Summary of habitat statistics.

Variable Name	Comments
Class Level Statistics	
NLCD*YYYY*_FS_C*TT*_*RR*_ED	Edge density for patches of habitat type TT. Ratio of total edge length to landscape area (meters per hectare). NA occurs if habitat TT not in landscape. Recoded as 0 since no edges implies 0 density.
NLCD*YYYY*_FS_C*TT*_*RR*_LPI	Largest patch index. Percentage of landscape area comprised by the largest patch of habitat type TT. NA occurs if habitat TT not in landscape. Recoded as 0.
NLCD*YYYY*_FS_C*TT*_*RR*_PD	Patch density. Number of patches of habitat class TT per 100 hectares in surrounding landscape. NA values occurred if no TT patches in landscape. Recoded as 0.
NLCD*YYYY*_FS_C*TT*_*RR*_PLAND	Percent of surrounding landscape that is habitat class TT. NA values occurred if no TT patches in landscape. Recoded as 0.
Landscape Level Statistics	
NLCD*YYYY*_FS_L_*RR*_ED	Edge density for the landscape. Ratio of sum of all edges between patches over total landscape area (meters per hectare).
NLCD*YYYY*_FS_L_*RR*_LPI	Percentage of landscape area occupied by the largest patch (any habitat type).
NLCD*YYYY*_FS_L_*RR*_PD	Patch density (number of patches per 100 hectares).

* All habitat statistics summarize square neighborhood around location with "radius" RR.
** YYYY = 2001 and 2006

A second class of corner case was caused by the fully automated application of FRAGSTATS. A checklist location near the edge of the NLCD map has unknown values in its landscape matrix, corresponding to the grid cells that extend past the NLCD map edge. We handled this by setting these cells to -9999, and

---

[5]A list of relevant literature can be found in the FRAGSTATS online documentation (background section): `http://www.umass.edu/landeco/research/fragstats/documents/ConceptualBackground/LiteratureCited/LiteratureCited.htm`

configuring FRAGSTATS to treat this value as background that is omitted from computations. Essentially, this truncates the extent of the landscape to the parts of the matrix with known values.

Acknowledgements

References

[MCNE] K. McGarigal, S. A. Cushman, M. C. Neel, and E. Ene. *FRAGSTATS: Spatial Pattern Analysis Program for Categorical Maps.* University of Massachusetts, Amherst. Version 3. Available from: http://www.umass.edu/landeco/research/fragstats/fragstats.html.

A   Covariate Processing Details

Covariate processing depends on whether the raw, unprocessed data are stored as shape files or raster data. Shape files are lists of polygons and their locations, and a covariate's value is constant within a polygon's region. Raster data represent the covariate data as a regular grid superimposed on a spatial area; each grid cell contains the covariate value for the corresponding region.

The census data (covariates HOUSING_DENSITY, HOUSING_PERCENT_VACANT, and POP00_SQMI) and climate data (all covariates named CAUS_) start as shape files that are subsequently processed using version 1.5.2 of the Geospatial Data Abstraction Library (GDAL)[6]. A C++ program calls the library to convert the decimal latitude and decimal longitude for each sampling event into the native coordinate system for a particular shape file and then assigns each event the appropriate covariate value based on the polygon containing each location. Locations that fall outside of all polygons are assigned NODATA values that are later converted into missing values (represented as ? in the dataset). Locations can fall outside all polygons because the location is outside the boundaries of the covariate data source, or because there are quirks with the data source (e.g., the climate data polygons vary in their boundaries from month to month).

The covariates BCR, COUNTRY, STATE_PROVINCE, and SUBNATIONAL2_CODE also started as shape files from ESRI that were then loaded into a spatial database (Oracle Spatial) that directly supports location-based queries.

The raw elevation, canopy, and impervious surface data are stored as raster images. A second C++ program calls the GDAL library to associate these data with sampling event locations. Decimal latitude and longitude are converted from the WGS 84 coordinate system to the coordinate system of the raster data. Next, the location is re-projected to a flat surface and converted to pixel coordinates to facilitate a fast data lookup. For the canopy and impervious surface covariates, a GDAL-based C++ program computes the mean data value for a square neighborhood around the location.

The landcover-based (habitat) covariates are also derived from raster data; these covariate values are statistics computed by the FRAGSTATS program [MCNE] using NLCD-2001 and NLCD-2006 raster data

---

[6]http://www.gdal.org

as input. Version 1.0 of the ERD and SRD datasets contain FRAGSTATS covariate statistics generated from v3.3 of the FRAGSTATS program. Version 2.1 of the ERD and SRD datasets contain FRAGSTATS covariate values generated from v2.0 of the FRAGSTATS program. Source code for version 2.0 of the Fragstats program is publicly available and provides for much higher processing throughput with the drawback of generating a less rich suite of statistics compared to those generated by the 3.3 release. Processing locations using the v2.0 release of the Fragstats code makes use of the GDAL library to extract a square neighborhood mini-raster around the location being processed, and then passing this mini-raster buffer to the Fragstats processing code built into the program. The processing pipeline for the v3.3 Fragstats is described as follows. A python wrapper script takes the decimal latitude and longitude coordinates for a sampling event and prepares a mini-raster data file describing the landscape neighborhood around the location. The script then passes the mini-raster to the FRAGSTATS executable. To create the mini-raster, the script reprojects the coordinates using an Albers Equal Area projection (using the GDAL utility program `cs2cs`) and calls the ArcGIS scripting library to extract the neighborhood around the location into the mini-raster. The mini-raster includes a landscape border around the neighborhood data (1 extra row / column on all sides) as recommended in the FRAGSTATS user manual.[7] The resolution of the mini-raster matches the original landcover data's resolution (30m x 30m).

A.1 FRAGSTATS Configuration

FRAGSTATS is run with the following configuration settings:

- Patch neighbors use the 8 cell rule (i.e., cells with touching corners are adjacent).

- Boundary does not count as edge (since there is a landscape border outside the landscape extent).

- Cells with unknown (missing) values are coded as background (value 9999).

- A similarity index file is used to customize the similarity of the different habitat classes defined in the National Land Cover Database. By default, each habitat class has 0 similarity with the other classes and perfect similarity with itself (1). We use the following custom similarities:

  - deciduous forest (class 41) and evergreen forest (class 42): 0.2
  - deciduous forest and mixed forest (class 43): 0.7
  - evergreen forest and mixed forest: 0.2
  - pasture, hay (class 81) and cultivated crops (class 82): 0.7
  - grassland, herbaceous (class 71) and pasture, hay: 0.3
  - grassland, herbaceous and cultivated crops: 0.3

  We believe these settings are adequate for most bird species.

- Similarly, we use an edge contrast file to soften the edges (dissimilarity) between the above habitat class pairs. Specifically, the edge contrast for the above pairs is set to 1 minus the similarity setting.

---

[7]See the section *Backgrounds, Borders, and Boundaries* in the FRAGSTATS user guidelines.

A.2   Hydrography Covariate Distance Values

Each hydrography feature type listed in Static Covariate Table 3 has a distance-in and distance-from value. If a checklist lat/lon location is inside a particular water feature, the distance-in value is described by the range bins listed in the Hydrography Distance Values Table 5; the matched distance-from value for this location will be represented as a missing value. A lat/lon location can only be inside a single hydrography feature type, so other distance-in variables for this location will be represented as missing values with the remaining distance-from values described by the range bins listed in the Hydrography Distance Values Table 5 as appropriate.

Similarly, if a checklist lat/lon location is outside all water features, the distance-from values are described by the range bins listed in the Hydrography Distance Values Table 5, and all distance-in values for this location will be represented as a missing value.

Table 5: Hydrography Distance Values

Value	Distance Range (meters)
1	0 - 30
2	30 - 60
3	60 - 120
4	120 - 250
5	250 - 500
6	500 - 1,000
7	1,000 - 2,000
8	2,000 - 4,000
9	>4,000

B   Change History

2012/10/1 Dataset release and document update (version 4.0) - K. Webb

2012/6/5 Dataset release and document update (version 3.1) Includes taxonomy updates and corrected core climate variables referenced in section 4.4 - K. Webb

2011/4/19 Dataset release and document update (version 3.0) - K. Webb

2010/12/21 Dataset release and document update (version 2.1) - K. Webb

2010/6/28 Dataset release and document update (version 2.0) - K. Webb

2010/2/4 Updated documentation to describe covariate processing and include SRD dataset.

2009/8/5 Release of covariates for 130,000 random locations in contiguous USA (SRD dataset).

2009/6/8 Initial dataset release (version 1.0).