A Guide to Pre-Processing High-Throughput Animal Tracking Data

Pratik Rajan Gupte[1, 2] Christine E. Beardsworth[2] Orr Spiegel[3]
Emmanuel Lourie[4] Sivan Toledo[5] Ran Nathan[4]
Allert I. Bijleveld[2]

1 Groningen Institute for Evolutionary Life Sciences, University of Groningen, The Netherlands.

2 Department of Coastal Systems, NIOZ Royal Netherlands Institute for Sea Research, 1790 AB, Den Burg, The Netherlands.

3 School of Zoology, Faculty of Life Sciences, Tel Aviv University, Tel Aviv 69978, Israel.

4 Movement Ecology Lab, Department of Ecology, Evolution, and Behavior, Alexander Silberman Institute of Life Sciences, The Hebrew University of Jerusalem, 91904 Jerusalem, Israel.

5 Blavatnik School of Computer Science, Tel Aviv University, Israel.

Correspondence Author

Pratik R. Gupte

Address

Groningen Institute for Evolutionary Life Sciences,

Nijenborgh 7/5172.0581 9747 AG Groningen

The Netherlands.

Email

pratikgupte16@gmail.com OR p.r.gupte@rug.nl

ORCID

https://orcid.org/0000-0001-5294-7819

Running Headline

Cleaning high-throughput animal tracking data

Abstract

1. Modern, high-throughput animal tracking studies collect increasingly large volumes of data at very fine temporal scales. At these scales, location-error can exceed the animal's step size, leading to mis-estimation of movement metrics such as speed. 'Cleaning' the data to reduce location errors prior to analyses is one of the main ways movement ecologists deal with noisy data, and has the advantage of being more scalable to massive datasets than more complex methods. Though data cleaning is widely recommended, and ecologists routinely consider cleaned data to be the ground-truth, uniform guidance on this crucial step is scarce.

2. A pipeline for cleaning massive high-throughput datasets must balance ease of use, computational efficiency. Data cleaning must also reject location errors without discarding valid animal movements. Another useful feature of a pre-processing pipeline is efficiently segmenting and clustering location data for statistical methods, while also being scalable to large datasets and robust to imperfect sampling. Manual methods being prohibitively time consuming, and to boost reproducibility, a robust pre-processing pipeline must be automated.

3. In this article we introduce a pipeline to pre-process high-throughput animal tracking data in order to prepare it for subsequent analysis. This pipeline, consisting of removing outliers, smoothing the filtered result, and thinning it to a uniform sampling interval, is easily scaled to very large datasets. We demonstrate this pipeline on simulated movement data with location errors. We also show a case study of how large volumes of cleaned data can be quickly segmented-clustered into biologically meaningful 'residence patches'. We use calibration data to show how pre-processing improves its quality, and to verify that the residence patch synthesis accurately captures animal space-use. Finally, turning to tracking data from Egyptian fruit bats (*Rousettus aegyptiacus*), we demonstrate the pre-processing pipeline and residence patch method in a fully worked out example.

4. To help with fast implementation, and to help standardise methods, we developed the `R` package `atlastools`, which we also introduce here. Our pre-processing pipeline and `atlastools` can be used with any high-throughput animal movement data in which the high data-volume combined with knowledge of the tracked individuals' movement capacity can be used to reduce location errors. The `atlastools` function is easy to use for beginners, while providing a template for further development. The use of common pre-processing steps that are simple yet robust promotes standardised methods in the field of movement ecology and leads to better inferences from data.

Keywords high-throughput movement ecology (HTME), biotelemetry, data cleaning, ATLAS tracking, *atlastools*, residence patch

1 Introduction

Tracking individual animals is the methodological mainstay of movement ecology, which seeks to link animal movement with internal causes, environmental factors, and resulting consequences (Holyoak et al., 2008; Nathan et al., 2008). Investigating fine-scale environmental and social drivers and the consequences of movement requires position data from many individuals at high temporal and spatial resolution relative to the phenomenon of interest (high-throughput tracking). High-throughput tracking is possible using GPS tags (see recent examples in Harel et al., 2016; Papageorgiou et al., 2019; Strandburg-Peshkin et al., 2015); additionally, 'reverse-GPS' systems developed to track animals over land (MacCurdy et al., 2009, 2019; Toledo et al., 2014, 2016, 2020; Weiser et al., 2016) and in aquatic systems (Aspillaga et al., 2021; Baktoft et al., 2019, 2017; Hussey et al., 2015; Jung et al., 2015) routinely produce high-throughput location data at much lower costs (Wikelski et al., 2007).

Although high-resolution tracking provides a massive amount of data on the path of a tracked animal, this data presents a two-fold challenge to ecologists. First, the location-error of each position may approach or exceed the true step size of the animal compared to low-resolution tracking with the same measurement error. This biases derived metrics such as speed and tortuosity (see Calenge et al., 2009; Hurford, 2009; Noonan et al., 2019; Ranacher et al., 2016). Additionally, the absolute location-error around a position introduces uncertainty when studying the location of an animal relative to fixed landscape features (e.g. roads) or mobile elements (e.g. other tracked individuals). Users have two main options to improve data quality: making inferences after modelling the system-specific location-error (Aspillaga et al., 2021; Fleming et al., 2014, 2020; Johnson et al., 2008; Jonsen et al., 2005, 2003; Patterson et al., 2008), or pre-processing data to clean it of positions with large location errors (Bjørneraas et al., 2010). The first approach may be limited by the animal movement models that can be fit to the data (Fleming et al., 2014, 2020; Noonan et al., 2019), and may be beyond the computational capacity of common hardware, leading users to prefer data cleaning instead. When attempting data cleaning, the second challenge of high-throughput tracking reveals itself: the large number of observations themselves (Toledo et al., 2020; Weiser et al., 2016). Having large volumes of data to clean makes manual identification and removal of errors from individual tracks prohibitively time consuming, incentivising automation based on a protocol.

Pre-processing steps must reliably discard large location errors, also called outliers, from tracks (analogous to reducing false positives) while avoiding the overzealous rejection of valid animal movements (analogous to reducing false negatives). How well researchers balance these imperatives has consequences for downstream analyses (Stine and Hunsaker, 2001). For instance, small-scale resource selection functions can easily infer spurious preference and avoidance effects when there is uncertainty about an animal's true position and movement (Visscher, 2006). Ecologists recognise that tracking data are imperfect observations of the underlying process of animal movement, yet they implicitly consider cleaned data equivalent to the ground-truth. This assumption

is reflected in popular statistical methods in movement ecology such as Hidden Markov Models (HMMs) (Langrock et al., 2012), stationary-phase identification methods (Patin et al., 2020), or step-selection functions (SSFs) (Avgar et al., 2016; Barnett and Moorcroft, 2008; Signer et al., 2017), which expect low location errors relative to real animal movement (i.e., a high signal-to-noise ratio). This makes the reproducible, standardised removal of location errors crucial to any animal tracking study. While gross errors are often removed by positioning-system algorithms, 'reasonable' errors often remain to confront end users (Fischler and Bolles, 1981; Ranacher et al., 2016; Weiser et al., 2016). Further, as high-throughput tracking is deployed in more regions, standardised pre-processing steps should be general enough to tackle animal movement data recovered from a range of environments.

Despite the importance and ubiquity of reducing location errors in tracking data, movement ecologists lack formal guidance on this crucial step. Pre-processing protocols are not often reported in the literature, or may not be easily tractable for mainstream computing hardware and software. Furthermore, filtering out positions on their location-error estimates may not be appropriate for outliers which represent unrealistic movement but have low error measures (Ranacher et al., 2016; Weiser et al., 2016). This makes identifying and removing biologically implausible locations from a track an important component of recovering true animal movement (Bjørneraas et al., 2010). Even after removing unrealistic movement, a track may comprise of positions that are distributed around the true animal location (Noonan et al., 2019). The large data-volumes of high-throughput tracking allow for a neat solution: tracks can be 'smoothed' to reduce small location errors that have remained undetected. This large data-volume becomes a challenge when users seek to examine animal space-use in relation to relevant environmental covariates, such as the predictors of prolonged residence in an area (see Aarts et al., 2008; Bijleveld et al., 2016; Bracis et al., 2018; Fleming et al., 2014; Harel et al., 2016; Oudman et al., 2018). To avoid issues of non-independent observations due to strong spatio-temporal autocorrelation, and to make datasets computationally tractable for mainstream statistical software, users may want to reduce data-volumes by thinning or clustering.

Here, we present a hands-on guide to pre-processing high-throughput tracking data, and demonstrate a relatively simple pipeline to prepare these data for subsequent analyses (see Fig. 1). We take two important considerations into account, (1) that methods must be computationally efficient, and (2) that the pre-processing steps should be easily understood and reproduced. `R` is the computational environment of choice in movement ecology (Joo et al., 2020) and formalising tools as an `R` package improves portability and reproducibility (Marwick et al., 2018). With these considerations in mind we have developed the `R` package `atlastools` (Gupte, 2020; R Core Team, 2020), which implements important steps of the pre-processing pipeline. `atlastools` functions are easy to use and understand, and rely on the efficiency of the `R` package `data.table` (Dowle and Srinivasan, 2020). Beginning with basic spatio-temporal filtering of positions, we cover filtering outliers from movement tracks, and reducing the effect of location-error with a median smooth. Then, we suggest one solution to issues

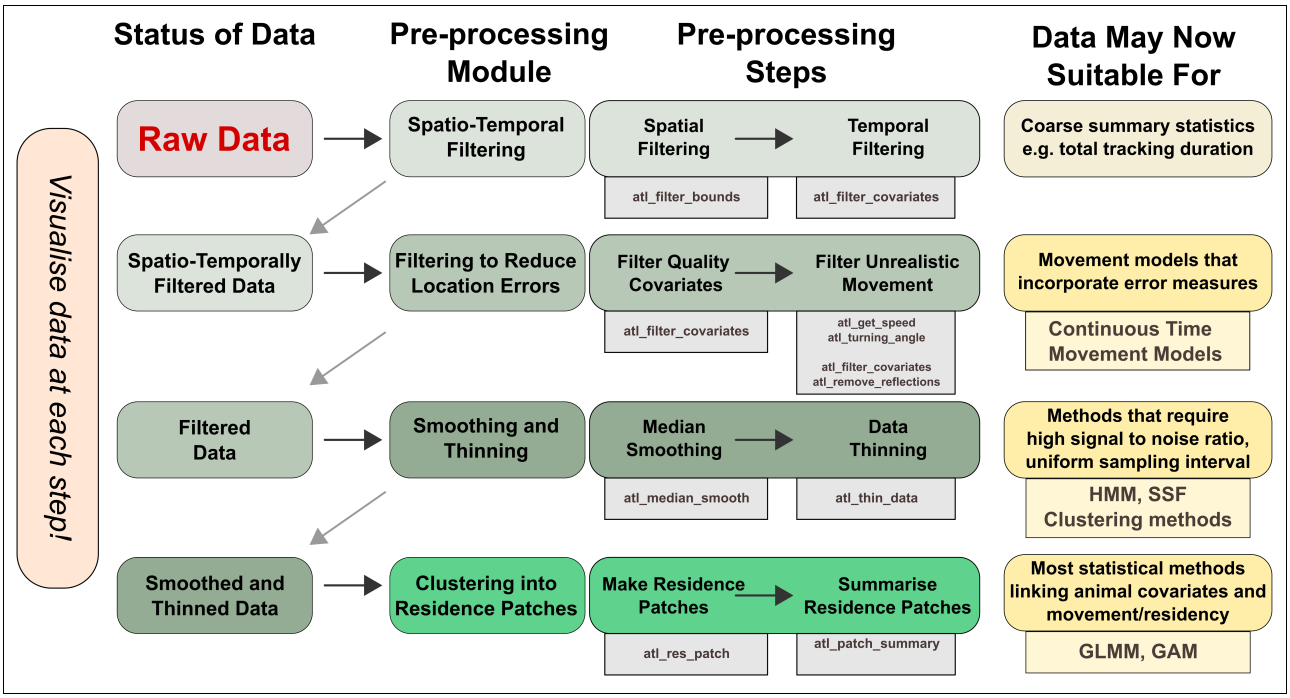

Figure 1. A general, modular pipeline for pre-processing high-throughput tracking data from raw localisations to cleaned data, and optionally into residence patches. Users should apply the appropriate pre-processing modules and the steps therein until the data are suitable for their intended analysis, some of which are suggested here. The `atlastools` function that may be used to implement each pre-processing step is shown in the grey boxes underneath each step. Popular statistical methods are shown underneath possible analyses (yellow boxes). Users are strongly encouraged to visualise their data and scan it for location errors as they work through the pipeline, always asking the question, could the animal plausibly move this way?

of computational tractability, which is to synthesise residence patches from clusters of spatio-temporally proximate positions, and so reveal important aspects of animal space use (*sensu* Barraquand and Benhamou, 2008; Bijleveld et al., 2016; Oudman et al., 2018). Using calibration data from a manually transported ATLAS tag, we demonstrate how the residence patch segmentation-clustering method can be used to accurately identify areas of prolonged residence under real field conditions. Finally, we turn to data from Egyptian fruit bats (*Rousettus aegyptiacus*) tracked in the Hula Valley, Israel, to show a fully worked out example of the pre-processing pipeline and the residence patch method.

2 Pipeline Overview, Usage, and Simulating Data

2.1 The Pipeline and `atlastools`

We lay out a modular pipeline for pre-processing raw high-throughput tracking data using the `R` package `atlastools` (Fig. 1). While the pipeline and package were designed with ATLAS systems in mind, the principles and functions can be used with any high-throughput tracking data. Users may follow the pipeline in full, or implement the module most suitable for their data. Users are advised to visualise their data throughout their

workflow to check for evident location errors (Slingsby and van Loon, 2016).

1. Users can use simple spatio-temporal filters to select positions within a certain time or area.

2. Next, users should reduce gross location errors by removing unreliable positions based on a system-specific error measure, or by the plausibility of associated movement metrics, such as speed and turning-angle (Calenge et al., 2009; Seidel et al., 2018).

3. Users should then reduce small-scale location errors by applying a median smooth.

4. Users who need uniformly thinned data can either aggregate or resample it. At this stage, the data are ready for a number of popular statistical treatments such as Hidden Markov Model-based classification (Langrock et al., 2012; Michelot et al., 2016).

5. Finally, users wishing simple, efficient segmentation-clustering to identify areas of prolonged residence can classify their data into residence patches based on the movement ecology of their study species, after filtering out travelling segments (see SEGMENTING AND CLUSTERING MOVEMENT TRACKS INTO RESIDENCE PATCHES).

2.2 Demonstrating the Pipeline with Simulated Data

To demonstrate the pipeline, we simulated a realistic movement track of 5,000 positions (unbiased correlated velocity model; UCVM) using the `R` package `smoove` (Gurarie et al., 2017, see Fig. 2.a). We added three kinds of error to the simulated track: (1) normally distributed small-scale offsets to the X and Y coordinates independently, (2) normally distributed large-scale offsets to a random subset (0.5 %) of the positions, and (3) large-scale displacement of a continuous sequence of 300 of the 5,000 positions (indices 500 – 800) (Fig. 2.a). To demonstrate the residence patch method, we chose to simulate three independent rotational-advective correlated velocity movement (RACVM) tracks of 500 positions each (ω = 7, initial velocity = 0.1, μ = 0; see Gurarie et al. 2017), and connected them together with a roughly linear path (see Fig. 5.a). RACVM models approximate the tracks of soaring birds which circle on thermals over a relatively small area, and move between thermals ('thermalling'; Gurarie et al., 2017; Harel et al., 2016). This complex track structure provides a suitable challenge for the residence patch method and helps to demonstrate its generality.

3 Spatio-Temporal Filtering

3.1 Spatial Filtering Using Bounding Boxes and Polygons

First, users should exclude positions outside the spatial bounds of a study area by comparing position coordinates with the range of acceptable coordinates (the bounding box), and removing positions outside them (Fig.

```
filtered_data <- atl_filter_bounds(
                        data = data,
                        x = "X", y = "Y",
                        x_range = c(x_min, x_max),
                        y_range = c(y_min, y_max),
                        sf_polygon = your_polygon,
                        remove_inside = FALSE
                    )
```

Listing 1. The `atl_filter_bounds` function filters on an area defined by coordinate ranges, a polygon, or all three; it can remove positions outside (`remove_inside = FALSE`), or within the area (`remove_inside = TRUE`). The arguments `x` and `y` determine the X and Y coordinate columns, `x_range` and `y_range` are the filter bounds in a coordinate reference system in metres, and the data can be filtered by an `sf-(MULTI)POLYGON` can be passed using the `sf_polygon` argument. The output is a `data.table`, which must be saved as an object (here, `filtered_data`).

2.a; Listing 1). A bounding box filter does not require a geospatial representation such as a shapefile, and can help remove unreliable data from a tracking system that is less accurate beyond a certain range (e.g. in ATLAS systems Beardsworth et al., 2021). In some special cases, users may wish to remove positions *inside* a bounding box, either because movement behaviour within an area is not the focus of a study, or because positions recorded within an area are known to be erroneous. An example of the former is studies of transit behaviour between features which can be approximated by their bounding boxes. Instances of the latter are likely to be system specific, but are known from ATLAS systems (Bijleveld et al. *in prep.*). Bounding boxes are by definition rectangular, and users seeking to filter for other geometries, such as a circular or irregularly-shaped study area, need a geometric intersection between their data and a spatial representation of the area of interest (e.g. shapefile, geopackage, or `sf`-object in `R`). The `atlastools` function `atl_filter_bounds` implements both bounding box and explicit spatial filters, and accepts X and Y coordinate ranges, an `sf`-polygon or multi-polygon object (Pebesma, 2018), or any combination of the three to filter the data (Listing 1).

3.2 Temporal and Spatio-temporal Filters

An animal's movement may be non-representative when it has been fitted with a tracker but has not been released, leading to artificial stationary positions, or in the time shortly after release when its movement may be influenced by the stress of capture and handling. Periods of poor tracking quality may result from system malfunctions, and users may also wish to exclude these data. Temporal filtering can exclude positions from intervals when data are expected to be unreliable for ecological inference, either due to abnormal movement behaviour or system-specific issues. Temporal filters can be combined with spatial filters to select specific time-location combinations. For example, users who want to study the response of individuals to the stress of capture can select positions shortly after release and within a certain distance of the release location. Users should apply filters in sequence rather than all at once, and visualise the output after each filtering step ('sanity checks'). The atlastools function `atl_filter_covariates` allows convenient filtering of a dataset by any number of logical

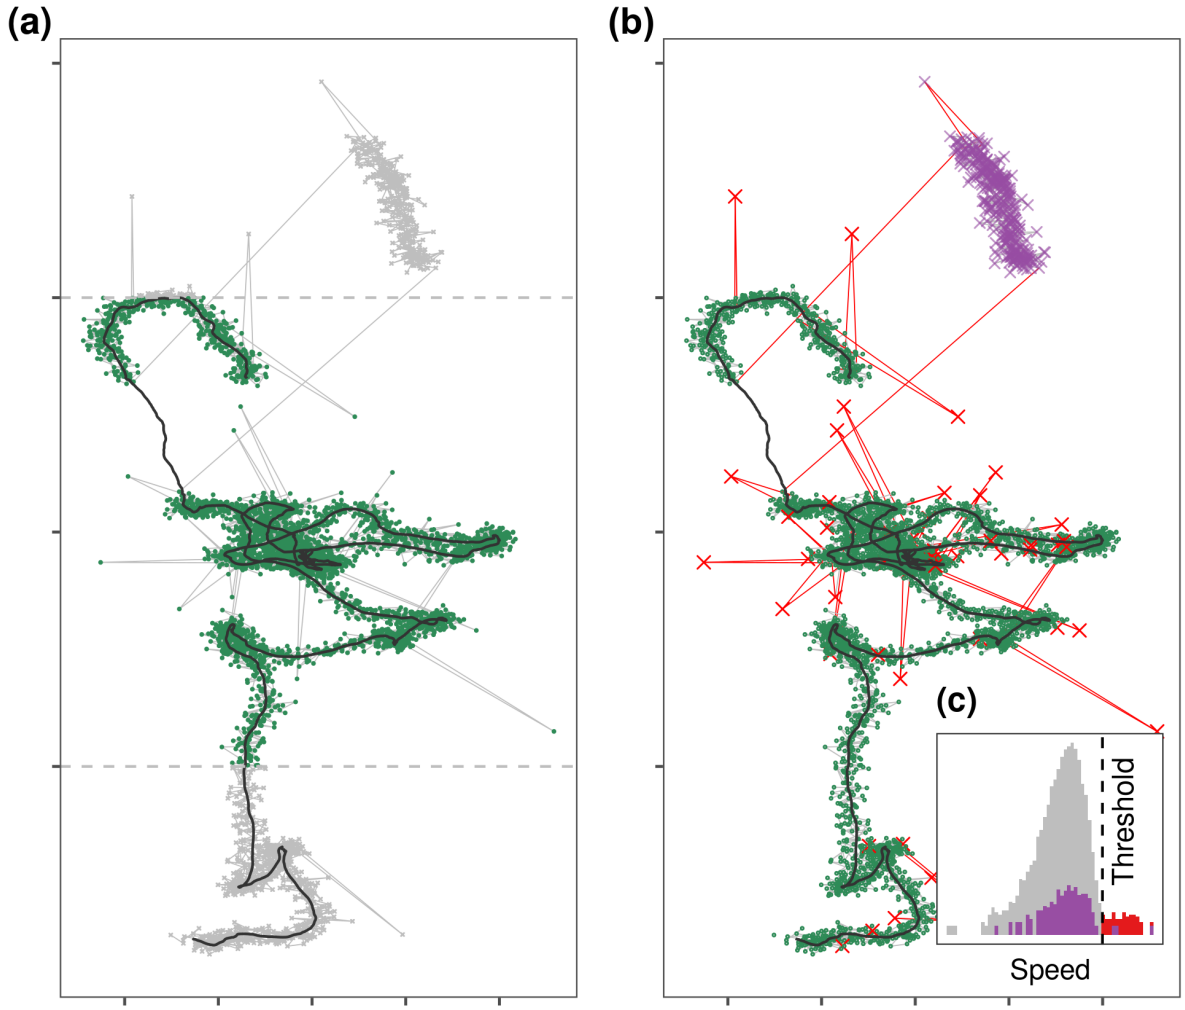

Figure 2. Simulated movement data (black line), with three kinds of artificially added errors: (1) normally distributed small-scale error on each position, (2) large-scale error added to 0.5% of positions, and (3) 300 consecutive positions displaced to simulate a gross distortion affecting a continuous subset of the track. **(a)** Tracks can be quickly filtered by spatial bounds (dashed grey lines) to exclude broad regions (green = retained; grey = removed). **(b)** Location-error may affect single observations resulting in point outliers or 'spikes' (red crosses and track segments), or continuous subsets of a track, called a 'prolonged spike' (purple crosses, top right), and both represent unrealistic movement. **(c)** Histograms of speed for the track (grey = small-scale errors, red = spikes), and the prolonged spike (purple). While spikes can be removed by filtering out positions with high incoming and outgoing speeds and turning angle, prolonged spikes cannot be removed in this way, and should be resolved by conceptualising algorithms that find the bounds of the distortion instead. Users should frequently check the outputs of such algorithms to avoid rejecting valid data.

statements, including querying data within a spatio-temporal range (Listing 2). The function keeps only those data which satisfy each of the filter conditions, and users must ensure that the filtering variables exist in their dataset in order to avoid errors.

```
night_data <- atl_filter_covariates(
                    data = dataset,
                    filters = c(
                        "!inrange(hour, 6, 18)",
                        "between(x, x_min, x_max)"
                    )
                )

filtered_data <- atl_filter_covariates(
                        data = data,
                        filters = c(
                            "NBS > 3",
                            "SD < 100",
                            "between(day, 5, 8)"
                        )
                    )
```

Listing 2. Data can be filtered by a temporal or a spatio-temporal range using `atl_filter_covariates`. Filter conditions are passed to the `filters` argument as a character vector. Only rows in the data satisfying *all* the conditions are retained. Here, the first example shows how nighttime data can be retained using a predicate that determines whether the value of 'hour' is between 6 and 18, and also within a range of X coordinates. The second example retains ATLAS locations calculated using > 3 base stations (`NBS`), with location-error (`SD`) < 100, and data between an arbitrary day 5 and day 8.

4 Filtering to Reduce Location Errors

4.1 Filtering on Quality Covariates

Tracking data covariates can be good indicators of the reliability of positions calculated by a tracking system (Beardsworth et al., 2021). GPS systems provide direct measures of location-error during localisation (Ranacher et al., 2016, Horizontal Dilution of Precision, HDOP in GPS), while reverse-GPS systems provide a similar estimate (called Standard Deviation, SD; MacCurdy et al., 2009, 2019; Ranacher et al., 2016; Weiser et al., 2016). Tracking data can also include indirect indicators of data quality. For instance, GPS systems' location-error may be indicated indirectly by the number of satellites involved in the localisation. In reverse-GPS systems too, the number of base stations involved in each localisation is an indirect indicator of data quality, and positions localised using more receivers are usually more reliable (the minimum required for an ATLAS localisation is 3; see Weiser et al., 2016). Unreliable positions can be removed by filtering on direct or indirect measures of quality using `atl_filter_covariates` (Listing 2).

4.2 Filtering Unrealistic Movement

Filtering on system-generated measures of error may not result in the removal of all erroneous positions, and data may remain which would require biologically implausible movement. Users are encouraged to visualise their tracks before and after filtering point locations, and especially to 'join the dots' and connect consecutive positions with lines (Figure 2.b). Whether the resulting track looks realistic is ultimately a subjective human

judgement, but one which implicitly integrates prior knowledge of the movement ecology of the study species to ask, 'Does the animal move this way?'. Segments which appear to represent unrealistic animal movement are often obvious to researchers with extensive experience of the study system (the non-movement approach; see Bjørneraas et al., 2010). Since it is both difficult and prohibitively time consuming to exactly reproduce expert judgement when dealing with large volumes of tracking data from multiple individuals, some automation is necessary. Users should first manually examine a representative subset of tracks and attempt to visually identify problems — either with individual positions, or with subsets of the track — that persist after filtering on quality covariates. Once such problems are identified, users can conceptualise algorithms that can be applied to their data to resolve them.

An example of a problem with individual positions is that of point outliers or 'spikes' (Bjørneraas et al., 2010), where a single position is displaced far from the track (see Fig. 2.b). Point outliers are characterised by artificially high speeds between the outlier and the positions before and after (called incoming and outgoing speed, respectively; Bjørneraas et al., 2010), lending a 'spiky' appearance to the track. Removing spikes is simple: remove positions with extreme incoming and outgoing speeds. Users must first define plausible upper limits of the study species' speed (Calenge et al., 2009; Seidel et al., 2018). Here, it is important to remember that speed estimates are scale-dependent; high-throughput tracking typically overestimates the speed between positions where the animal is stationary or moving slowly, due to small-scale location errors (Noonan et al., 2019; Ranacher et al., 2016). Even after data with large location errors have been removed, it is advisable to begin with a liberal (high) speed threshold that excludes only the most unlikely of speeds. Estimates of maximum speed may not always be readily obtained for all species, and an alternative is to use a data-driven threshold such as the 95th percentile of speeds from the track. Once a speed threshold S has been chosen, positions with incoming *and* outgoing speeds $> S$ may be identified as spikes and removed.

Some species can realistically achieve speeds $> S$ in fast transit segments when assisted by their environment, such as birds with tailwinds, and a simple filter on incoming and outgoing speeds would exclude this valid data. To avoid removing valid, fast transit segments while still excluding spikes we suggest combining the speed filter with a filter on the turning angles of each position (see Calenge et al., 2009). This combined filter assumes that positions in high-throughput tracking with both high speeds and large turning angles are likely to be due to location errors, since most species are unable to turn sharply at high speed. Users can then remove those positions whose incoming and outgoing speeds are both $> S$, and where $\theta > A$ (sharp, high-speed turns), where θ is the turning angle, and A is the turning angle threshold. Many other track metrics may be used to identify implausible movement, and used to filter data (Seidel et al., 2018). Spike removal can be implemented using the `atl_filter_covariates` function (Listing 3).

Sometimes, entire subsets of the track may be affected by the same large-scale location-error. For instance, multiple consecutive positions may be roughly translated (geometrically) away from the real track and form

```
data$speed_in <- atl_get_speed(
                    data,
                    x = "x", y = "y",
                    time = "time",
                    type = c("in")
                )

data$angle <- atl_turning_angle(
                    data,
                    x = "x", y = "y",
                    time = "time"
                )

filtered_data <- atl_filter_covariates(
                    data = data,
                    filters = c(
                        "(speed_in < S & speed_out < S) | angle < A"
                    )
                )
```

Listing 3. Filtering a movement track on incoming and outgoing speeds, and on turning angle to remove unrealistic movement. The functions `atl_get_speed` and `atl_turning_angle` are used to get the speeds and turning angles before filtering, and assigned to a column in the data (assignment of `speed_out` is not shown). The filter step only retains positions with speeds below the speed threshold S *or* angles above the turning angle threshold θ, i.e., positions where the animal is slow but makes sharp turns, and data where the animal moves quickly in a relatively straight line.

'prolonged spikes', or 'reflections' (see Fig. 2.b). These cannot be corrected by targeted removal of individual positions, as in Bjørneraas and colleagues' approach (2010), since there are no positions with both high incoming and outgoing speeds. Since filtering individual positions will not suffice, algorithms to correct such errors must take a track-level view, and target the displaced sequence overall. Track-subset algorithms are likely to be system-specific, and may be challenging to conceptualise or implement. In the case of prolonged spikes, one relatively simple solution is identifying the bounds of displaced segments, and removing positions between them. Users are strongly encouraged to visualise their data before and after applying such algorithms, We caution that these methods are not foolproof, and data that are heavily distorted by errors affecting entire track-subsets should be used with care when making further inferences.

5 Smoothing and Thinning Data

5.1 Median Smoothing

After filtering out large location errors, the track may still look 'spiky' at small scales, and this is due to smaller location errors. These are caused by errors in sampling the true animal position, and are especially noticeable when the individual is stationary or moving slowly (Noonan et al., 2019). These smaller errors are challenging to remove since their covariates (such as speed and turning angles) are within the expected range of movement behaviour for the study species. The large data-volumes of high-throughput tracking allow users to resolve

```
atl_median_smooth(
    data = track_data,
    x = "x", y = "y",
    time = "time",
    moving_window = 5
)
```

Listing 4. Median smoothing a movement track using the function `atl_median_smooth` function with a moving window $K = 5$. Larger values of K yield smoother tracks, but K should always be some orders of magnitude lower than the number of observations.

this problem by smoothing the positions. A 'smooth' works by approximating the value of an observation based on neighbouring values. For a one-dimensional series of observations, the neighbouring values are the K observations centred on each index value i. The range $i - (K-1)/2 \ldots i + (K-1)/2$ is referred to as the moving window as it shifts with i, and K is called the moving window size. A common smooth is nearest neighbour averaging, in which the value of an observation x_i is the average of the moving window K. The median smooth is a variant of nearest neighbour averaging which uses the median rather than the mean, and it is robust to outliers (Tukey 1977). The median smoothed value of the X coordinate, for instance, is

$$X_i = Median(X_{i-(K-1)/2} \ldots X_{i+(K-1)/2})$$

Users can apply a median smooth with an appropriate K independently to the X and Y coordinates of a movement track to smooth it (see Fig. 3.a – e). The median smooth is robust to even very large temporal and spatial gaps, and does not interpolate between positions when data are missing. Thus it is not necessary to split the data into segments separated by periods of missing observations when applying the filter (see Fig. 3).

Smoothing does not change the number of observations, but does decouple the coordinates from some of their covariates. For instance, smoothing breaks the relationship between a coordinate and the location-error estimate around it (VARX, VARY, and SD in ATLAS systems, or HDOP in GPS tracking). This makes subsequent filtering on covariates of data quality unreliable, and smoothed data are unsuitable for use with methods that model location uncertainty (Calabrese et al., 2016; Fleming et al., 2014, 2020; Noonan et al., 2019). Furthermore, while larger K may result in smoother tracks (Fig. 3.b – e), one drawback of using a large K is that short, quick forays away from the main track are likely to be smoothed away, leading to a loss in detail of the individual's small-scale movement. Users must themselves judge how best to balance large-scale and small-scale accuracy, and choose K accordingly. Median smoothing is provided by the `atlastools` function `atl_median_smooth`, with the only option being the moving window size, which must be an odd integer (Listing 4).

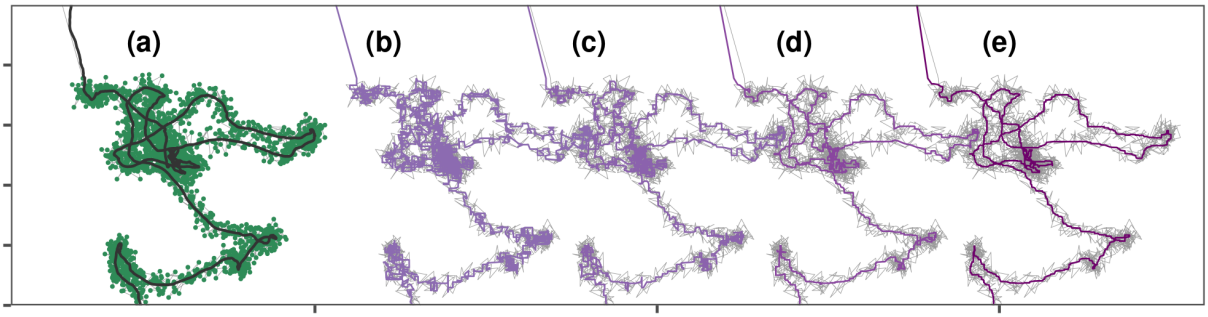

Figure 3. Median smoothing position coordinates reduces small-scale location-error in tracking data. **(a)** The goal of this step is to approximate the simulated canonical track (black line), given positions with small-scale error remaining from previous steps (green points; grey lines). Median smoothing the given position coordinates (green points, in **(a)**) over a moving window (K) of **(b)** 3, **(c)** 5, **(d)** 11, and **(e)** 21 positions respectively, yields approximations (purple lines) of the canonical track. Grey lines show the track without smoothing. Users are cautioned that there is no correct K, and they must subjectively choose a K which most usefully trades small-scale details of the track for large-scale accuracy.

5.2 Thinning Movement Tracks

Most data at this stage is technically 'clean', yet its volume alone may pose challenges for lower-specification or older hardware and software if these are not optimised for efficient computation. Thinning data need not compromise researchers' ability to answer scientific questions with them; for instance, proximity-based social interactions lasting 1 – 2 minutes would still be detected on thinning from a sampling interval of 1 second to 1 minute. Indeed, some analyses require that temporal auto-correlation in the data be broken by subsampling the data to a lower resolution (such as the estimation of home-ranges or step-selection functions; Dupke et al., 2017; Fleming et al., 2014). Added to the requirement of uniform sampling intervals by many methods in the field, evenly reducing data-volumes is a worthwhile endeavour (e.g. Avgar et al., 2016; Fleming et al., 2014; Michelot et al., 2016). Two plausible approaches here are resampling and aggregation, and both approaches begin with identifying time-interval groups (e.g. of 1 minute). Resampling picks one position from each time-interval group while aggregation involves computing the mean values of all covariates for positions within a time-interval group. Both approaches yield one position per time-interval group (Fig. 4.a). Categorical variables, such as the habitat type associated with each position, can be aggregated using a suitable measure such as the mode. We caution users that thinning causes an extensive loss of small-scale detail in the data, and should be used carefully.

Both aggregation and resampling have their relative advantages. The aggregation method is less sensitive to selecting point outliers by chance than resampling. However, when users want to account for location-error with methods such as state-space models (Johnson et al., 2008; Jonsen et al., 2005, 2003), or continuous time movement models (Calabrese et al., 2016; Fleming et al., 2014, 2020; Gurarie et al., 2017; Noonan et al., 2019),

```
thinned_data <- atl_thin_data(
                    data,
                    interval = 60,
                    id_columns = c("animal_id"),
                    method = "aggregate"
                )
```

Listing 5. Code to thin data by aggregation in `atlastools`. The method can be either "aggregate" or "resample". The time interval is specified in seconds, while the `id_columns` allows a character vector of column names to be passed to the function, with these columns used as identity variables. Both methods return a dataset with one rows per time-interval.

correctly propagating the location-error is important, and resampling directly propagates these errors without further processing. In the aggregation method, the location-error around each coordinate provided by either GPS or reverse-GPS system can be propagated to the averaged position as the sum of errors divided by the square of the number of observations contributing to each average (N):

$$Var(X)_{agg} = \left(\sum_{i=1}^{i=N} Var(X)_i \right) / N^2$$

Similarly, the overall location-error estimate for the average of N positions in a time-interval can be calculated by treating it as a variance:

$$SD_{agg} \text{ or } HDOP_{agg} = \sqrt{\left(\sum_{i=1}^{i=N} SD_i^2 \text{ or } HDOP_i^2 \right) / N^2}$$

Users may question why thinning, which can obtain consensus positions over an interval and also reduce data-volumes should not be used directly on the raw data. We caution that thinning prior to excluding unrealistic movement and smoothing (Fig 4.b) can lead to preserving artefacts in the data, and estimates of essential metrics — such as speed — that are substantially different from the true value (see Fig. 4.c; Noonan et al., 2019). In our example for instance, the data with errors included would have to be thinned to $1/30^{\text{th}}$ of its volume for the median speed to be correctly estimated; an undesirable step if the aim is fine-scale tracking (Fig. 4.c). The mis-estimation of track metrics could have knock-on consequences for the implementation of subsequent filters based on detecting unrealistic movement. However, thinning before data-cleaning may have its place as a useful step before exploratory visualisation of the movement track, since reduced data-volumes are easier to handle for plotting software. Thinning is implemented in `atlastools` using the `atl_thin_data` function, with either aggregation or resampling (specified by the `method` argument) over an interval using the `interval` argument. Grouping variable names (such as animal identity) may be passed as a character vector to the `id_columns` argument (Listing 5).

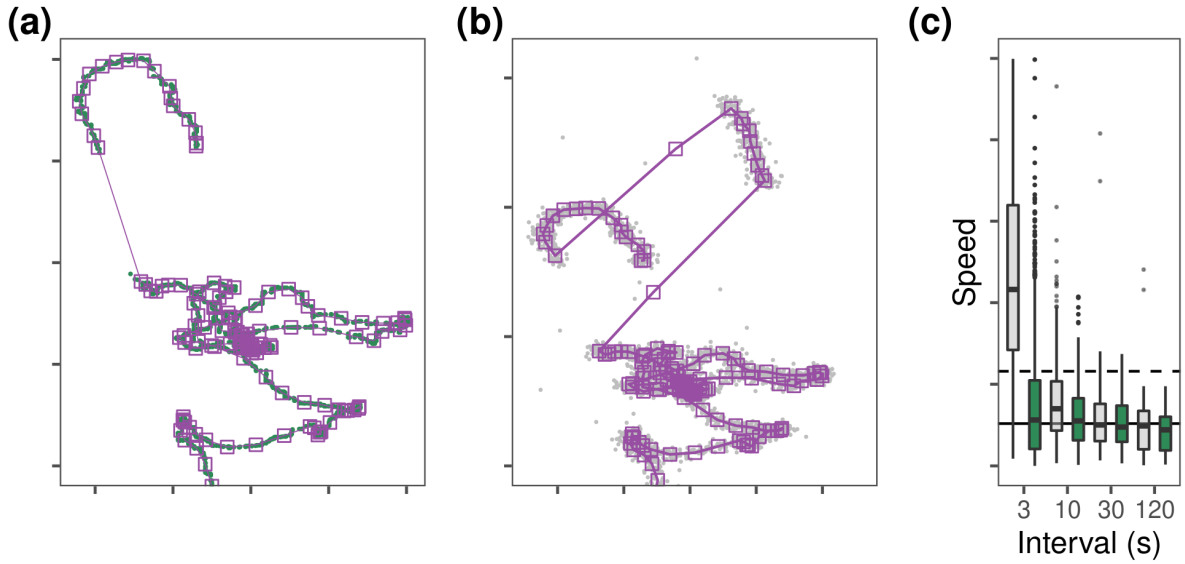

Figure 4. **(a)** Aggregating a filtered and smoothed movement track (green points) preserves track structure while reducing data-volume, but **(b)** aggregating before filtering gross location errors and unrealistic movement (grey points) leads to the persistence of large-scale errors (such as prolonged spikes). The thinned track is shown as purple squares and lines in both panels. **(c)** Thinning before data cleaning can lead to significant misestimations of essential movement metrics such as speed at lower intervals. Boxplots show the median and interquartile ranges for speed estimates of tracks aggregated over intervals of 3, 10, 30, and 120 seconds (grey = thinning without removing location errors; green = thinning after filtering and smoothing). For comparison, the median and 95th percentile of speed of the canonical track are shown as solid and dashed horizontal lines, respectively. The unfiltered data would have to be thinned to $1/30^{th}$ volume to correctly estimate median speed.

6 Segmenting and Clustering Movement Tracks into Residence Patches

6.1 The Residence Patch Algorithm

An important and common analysis with animal tracking data is to link space-use with environmental covariates. This is difficult even with smoothed and thinned high-throughput data, as these may be too large for statistical packages, or have strong autocorrelation. Users aiming for such analyses can benefit from segmenting and clustering the data into spatio-temporally independent 'residence patches', and metrics such as the distance travelled within and between patches can help compare broad-scale individual movement strategies, or be linked with environmental data (Barraquand and Benhamou, 2008; Bijleveld et al., 2016; Oudman et al., 2018). Common methods of segmentation-clustering are either not based in the biology of the tracked individual, or are not scalable to very large, yet often gappy datasets (Patin et al., 2020), making a first-principles approach an essential part of a pre-processing pipeline. Furthermore, making the patch the unit of observation conveniently sidesteps pseudo-replication and reduces computational requirements.

First, users should identify positions representing bouts of stationary behaviour, for instance on their speed or residence time (Bracis et al., 2018). Patch identification assesses whether consecutive stationary positions and the bouts they comprise are spatio-temporally independent, and clusters them together if they are not.

The splitting and lumping of positions into bouts, and bouts into patches is based on simple user specified thresholds — the distance and the time interval between positions (and bouts) beyond which they should be considered independent. Users are encouraged to base these thresholds in the movement habits of their study species. For example, residence patch classification of red knot movement tracks considers consecutive stationary positions independent if they are 20m apart and considers consecutive bouts independent if the distance between them $\geq$ 100m, or the time difference between them $\geq$ 30 minutes (Listing 6). The 20m distance represents a maximum speed of 6.667 m/s between positions, above which the individual is more likely in transit. The 100m and 30 minute thresholds are chosen to account for potentially missing data between bouts; if a track ends abruptly and then reappears $\geq$ 100m away or $\geq$ 30 minutes later, this is more safely considered a new residence patch.

A cleaned movement track can be classified into residence patches using the function `atl_res_patch` (see Fig. 5.b), which requires three parameters: (1) the distance threshold between positions (called `buffer_size`), (2) the distance threshold between clusters of positions (called `lim_spat_indep`), and (3) the time interval between clusters (called `lim_time_indep`). Clusters formed of fewer than a minimum number of positions can be excluded. This residence patch algorithm is capable of correctly identifying clusters of related points from a movement track (Fig. 5). This includes clusters where the animal is relatively stationary (orange and green patches, Fig. 5.b), as well as clusters where the animal is moving slowly (blue patch, Fig. 5.b). This flexibility is especially useful when studying movements that may represent two different modes of the same behaviour, for instance, area-restricted search, as well as slow, searching behaviour with a directional component.

The function `atl_patch_summary` can be used to extract patch-specific summary data such as the median coordinates, the patch duration, the distance travelled within the patch, and the patch area. Position covariates such as speed may also be summarised patch-wise by passing covariate names and summary functions as character vectors to the `summary_variables` and `summary_functions` arguments, respectively. Setting the `which_data` argument to `"spatial"`, returns `sf MULTIPOLYGON` objects, and setting `which_data = "points"` returns the positions in each patch, with patch-specific covariates.

6.2 Validating the Residence Patch Method

We applied the pre-processing pipeline using `atlastools` functions described above to a calibration dataset to verify that the residence patch method could correctly identify known stopping points (see Fig. 6). We collected the calibration data (n = 50,816) by walking and boating with a hand-held WATLAS tag (sampling frequency = 1 Hz) around the island of Griend (53.25°N, 5.25°E) in August 2020 (Beardsworth et al., 2021, Bijleveld et al. *in prep.*). Stops in the calibration track were recorded as waypoints using a handheld GPS device (Garmin Dakota 10) at each stop. We estimated the real duration of each stop as the time difference between the first and last position recorded within 50m of each waypoint, within a 10 minute window before and after the waypoint

```
patches <- atl_res_patch(
                data = track_data,
                buffer_radius = 10,
                lim_spat_indep = 100,
                lim_time_indep = 30,
                min_fixes = 3,
                summary_variables = c("speed"),
                summary_functions = c("mean", "sd")
            )

patch_summary <- atl_patch_summary(
                    patch_data = patches,
                    which_data = "summary",
                    buffer_radius = 10
                )
```

Listing 6. The `atl_res_patch` function can be used to classify a track into residence patches. The arguments `buffer_radius` and `lim_spat_indep` are specified in metres, while the `lim_time_indep` is provided in minutes. In this example, specifying `summary_variables = c("speed")`, and `summary_functions = c("mean", "sd")` will provide the mean and standard deviation of instantaneous speed in each residence patch. The `atl_patch_summary` function is used to access the classified patch in one of three ways, here using the `summary` option which returns a table of patch-wise summary statistics.

timestamp (to avoid biased durations from revisits). Stops had a median duration of 10.28 minutes (range: 1.75 minutes – 20 minutes; see Supplementary Material). We cleaned the data before constructing residence patches by (1) removing a single outlier ($>$ 15 km away), removing unrealistic movement ($\geq$ 15 m/s), smoothing the data (K = 5), and (4) thinning the data by resampling over a 30 second interval. The cleaning steps retained 37,324 positions (74.45%), while thinning reduced these to 1,803 positions (4.8% positions of the smoothed track). Details and code are provided in the Supplementary Material (see VALIDATING THE RESIDENCE PATCH METHOD WITH CALIBRATION DATA).

We identified stationary positions (residence time $\geq$ 5 minutes) using the `recurse` package (n = 837, 46.42 %; radius = 50m Bracis et al., 2018). We clustered these positions into residence patches with a buffer radius of 5m, spatial independence limit of 50m, temporal independence limit of 5 minutes, and a minimum of 3 positions per patch. Inferred residence patches corresponded well to the locations of stops (see Fig. 6.c). However, the residence patch algorithm detected more stops than were logged as waypoints (n = 28, n waypoints = 21). One of these was the field station on Griend where the tag was stored between trips (red crossed-square, Fig. 6.c). The method also did not detect two stops of 105 and 563 seconds (1.75 and 9.4 minutes) since they were data poor and aggregated away in the thinning step (n positions = 6, 15). To determine whether the residence patch method correctly identified the duration of stops in the calibration track, we first extracted the patch attributes using the function `atl_patch_summary`. We then matched the patches to the waypoints by their median coordinates (rounded to 100 metres). We assigned the inferred duration of the stop as the duration of the spatially matched residence patch. We compared the inferred duration with the real duration using a linear model with the inferred duration as the only predictor of the real duration. Inferred duration was a good

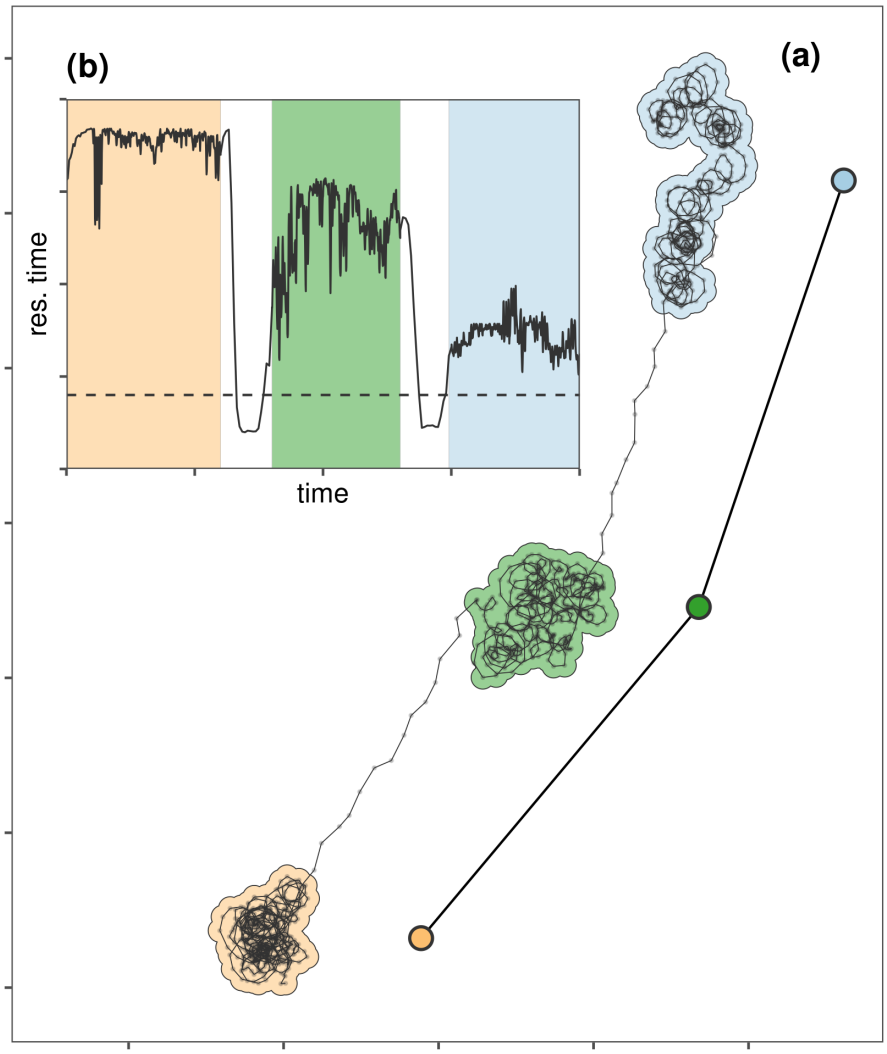

Figure 5. Movement tracks can be classified into residence patches, which are areas of prolonged residence (clusters of points), while leaving out the transit between them. **(a)** The residence patch method correctly identifies clusters of positions where the individual is relatively stationary (orange and green patches), as well as positions where it is moving slowly (blue patch). This is especially useful when studying more complex behaviour such as area-restricted search, which may have a directional component. The residence patch method loses the details of movement between patches, but can efficiently represent the general pattern of space-use (see coloured points representing patch centroids, and lines joining them). **(b)** A plot of residence time against time (solid line; Bracis et al. 2018) shows how the residence patch algorithm segments and clusters positions of prolonged residence. Regions are shaded by the temporal bounds of each residence patch. The arguments passed to `atl_res_patch` determine the clustering, and can be adjusted to fit the study system. Users are cautioned that there are no 'correct' arguments, and the best guide is the biology of the tracked individual.

predictor of the real duration of a stop (linear model estimate = 1.021, t-value = 12.965, $p < 0.0001$, $R^2 = 0.908$; see Supplementary Material Fig. 1.7). This translates to a 2% underestimation of the stop duration at a tracking interval of 30 seconds.

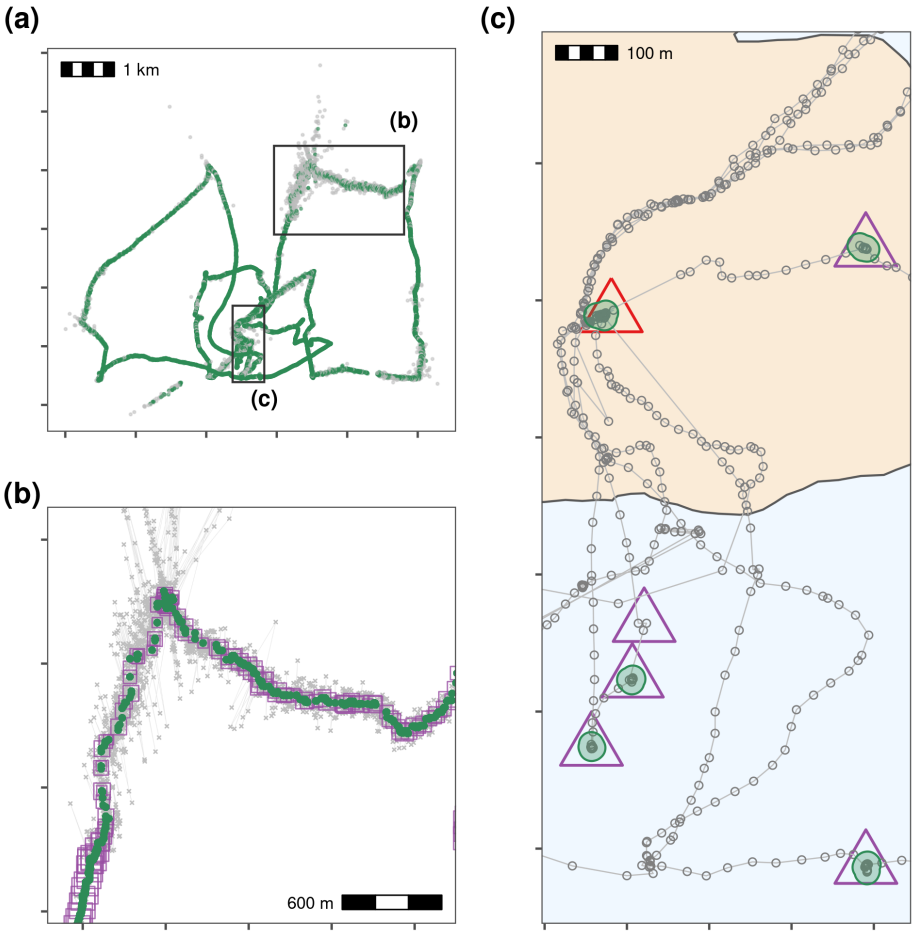

Figure 6. Pre-processing steps for WATLAS calibration data showing filtering on speed, median smoothing and thinning by aggregation, and making residence patches. **(a)** Positions with incoming and outgoing speed > 15 m/s are removed (grey crosses = removed, green points = retained). **(b)** Raw data (grey crosses), median smoothed positions (green circles; moving window $K = 5$), and the smoothed track thinned by aggregation to a 30 second interval (purple squares). Square size corresponds to the number of positions used to calculate the averaged position during thinning. **(c)** Clustering thinned data into residence patches (green polygons) yields robust estimates of the location of known stops (purple triangles). The algorithm identified all areas with prolonged residence, including those which we had not intended to be recorded, such as stops at the field station (n = 12; red triangle). The algorithm also failed to find two stops of 6 and 15 seconds duration, since these were lost in the data thinning step; one of these is shown here (purple triangle without green polygon).

7 Worked-Out Example on Animal Tracking Data

We present a fully worked-out example of our pre-processing pipeline and residence patch method using movement data from three Egyptian fruit bats tracked using the ATLAS system (*Rousettus aegyptiacus*; Toledo et al. (2020)). Code can be found in the Supplementary Material (see PROCESSING EGYPTIAN FRUIT BAT TRACKS). Bats were tracked over three nights (5th, 6th, and 7th May, 2018) in the Hula Valley, Israel (33.1°N, 35.6°E), with an average of 13,370 positions (SD = 2,173; range = 11,195 – 15,542; interval = 8 seconds) per individual. Plotting the tracks showed severe distortions (see Supplementary Material Fig. 2.1). We first reduced location errors by removing observations with ATLAS SD > 20, and observations calculated using fewer than four base stations

(mean positions remaining = 10,447 / individual; 78% of the raw data on average). We removed unrealistic movement represented by positions with incoming and outgoing speeds > 20 m/s leaving 10,337 positions per individual on average (98% of previous step). We median smoothed the data with a moving window K size = 5, and no observations were lost.

We began the construction of residence patches by finding the residence time within 50 metres of each position (Bracis et al., 2018). Bats may repeatedly traverse the same routes, and this could artificially inflate the residence time of positions along these routes. To avoid confusing revisits with residence, we limited the summation of residence times at each position to the period until the first departure of 60 minutes or more. Thus, two nearby locations ($\leq$ 50m apart) each visited for one minute at a time, but separated by an interval of some hours would not have a residence time of two minutes each, but only one minute each. Bats had a mean residence time at locations of 100.54 minutes (SD = 114.7); this measure was strongly biased by time spent at the roost. We opted for a first-principles approach and selected as residence positions any locations with a residence time > 5 minutes, reasoning that a flying animal stopping for > 5 minutes at a location should plausibly indicate resource use or another interesting behaviour. This step retained 7,819 positions per bat on average (75.6%) of the smoothed data, showing that bats are actually relatively stationary at the daily scale, and move only mostly between their roost and foraging sites (see Fig. 7).

We constructed residence patches with a buffer distance of 25m, a spatial independence limit of 100m, a temporal independence limit of 30 minutes, and rejected patches with fewer than three positions. We extracted summary data and spatial polygons from the constructed residence patches. Plotting the bats' residence patches and the linear paths between them showed that though all three bats roosted at the same site, they used distinct areas of the study site (Fig. 7.a). Bats tended to show prolonged residence near known food sources (fruit trees), travelling repeatedly between previously visited areas (Fig. 7.b). However, bats also appeared to spend some time at locations where no fruit trees were recorded, prompting questions about their use of other food sources, or another behaviour entirely (Fig. 7.b, 7.c). Bats occurring close together did not have strongly overlapping residence patches, and their paths to and from area of co-occurrence were different (Fig. 7.a, 7.c). Constructing residence patches for multiple individuals over multiple nights suggests interesting dynamics of within- and between-individual overlap.

8 Discussion

Our guide anticipates that high-throughput animal tracking data will become increasingly common. A uniform pipeline and toolset for data cleaning promotes reproducibility and standardisation across studies, making comparative inferences more robust. The open-source `R` package `atlastools` serves as a starting point for methodological collaboration among movement ecologists. Efficient location-error modelling approaches (Aspillaga et al., 2021; Fleming et al., 2020) may eventually make data-cleaning optional. Yet cleaning tracking data

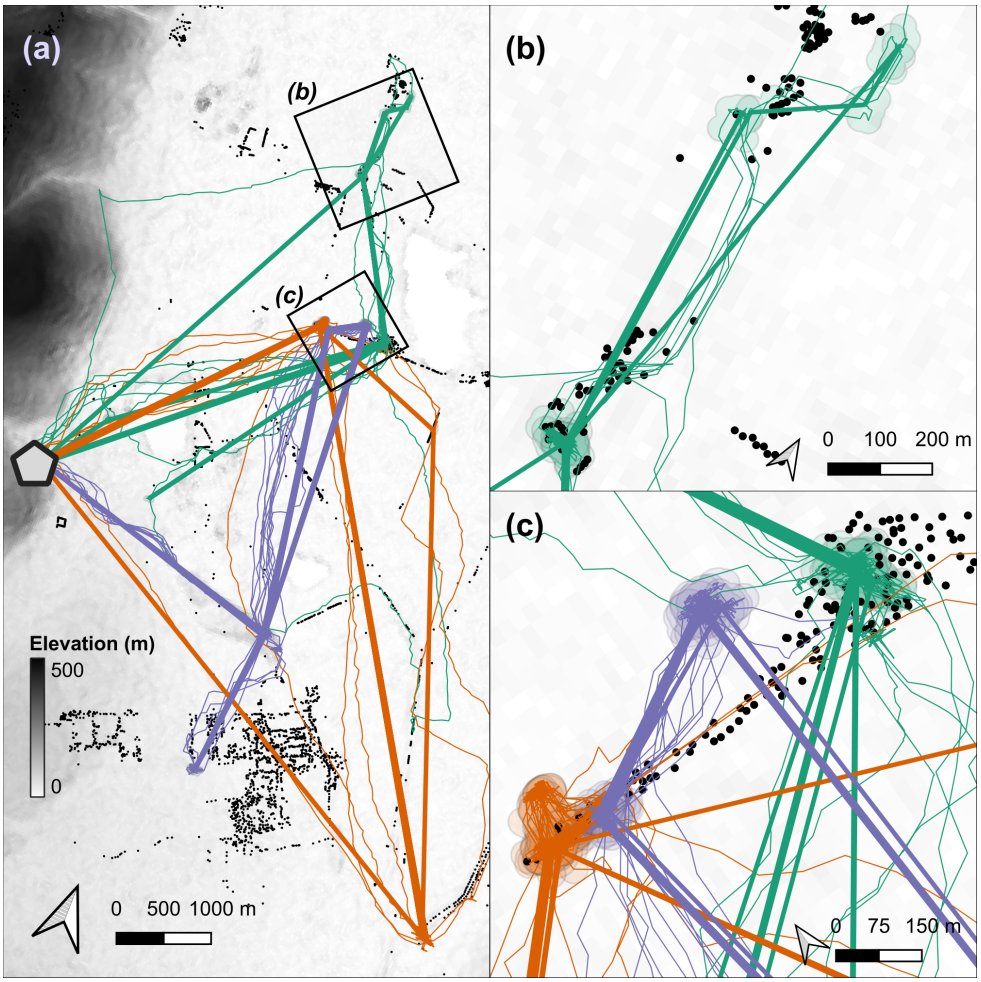

Figure 7. Synthesising animal tracks into residence patches can reveal movement in relation to landscape features, prior exploration, and other individuals. **(a)** Linear approximations of the paths (coloured straight lines) between residence patches (circles) of three Egyptian fruit bats (*Rousettus aegyptiacus*), tracked over three nights in the Hula Valley, Israel. Real bat tracks are are shown as thin lines below the linear approximations, and colours show bat identity. The grey hexagon represents a roost site. Black points represent known fruit trees. Background is shaded by elevation at 30 metre resolution. **(b)** Spatial representations of an individual bat's residence patches (green polygons) can be used to study site-fidelity by examining overlaps between patches, or to study resource selection by inspecting overlaps with known resources such as fruit trees (black circles). In addition, the linear approximation of movement between patches (straight green lines) can be contrasted with the estimated real path between patches (irregular green lines). **(c)** Fine scale tracks (thin coloured lines), large-scale movement (thick lines), residence patch polygons, and fruit tree locations show how high-throughput data can be used to study movement across scales. Patches and lines are coloured by bat identity.

even partially before modelling location-error is faster than error-modelling on the full data, and the removal of large location errors may improve model fits. Thus we see our pipeline as complementary to these approaches (Fleming et al., 2014, 2020). Finally, we recognise that the diversity and complexity of animal movement and data collection techniques often requires system-specific, even bespoke, pre-processing solutions. Though the principles outlined here are readily generalised, users' requirements will eventually exceed the particular tools we provide. We see this as an incentive for more users to be involved in developing methods for their systems. We offer our pipeline and package as a foundation for system-specific tools in the belief that simple, robust

concepts are key to methods development that balances system-specificity and broad applicability.

9 Backmatter

9.1 Competing Interests

The authors declare that they have no competing interests.

9.2 Acknowledgements

PRG would like to thank Pedro M. Santos Neves for introducing PRG to `R` package development, for help with setting up `atlastools`, and for help with archiving it on Zenodo; Aparajitha Ramesh and Franjo Weissing for discussions on error propagation; Geert Aarts, Jacob RL Gismann, Evy Gobbens, and Roos Kentie for feedback that improved the manuscript; members of the Modelling Adaptive Response Mechanisms Group (Weissing Lab), and the Theoretical Biology department at the University of Groningen for helpful discussions on `atlastools` and the manuscript. We thank the many volunteers, students, and NIOZ staff involved in operating the WATLAS tracking system, and most importantly Frank van Maarseveen, Bas Denissen and Anne Dekinga. We also thank the attendees of ATLAS workshops held in May and June 2020 at the Hebrew University of Jerusalem for helpful comments on the pipeline and `atlastools`. This work was partly funded by the Dutch Research Council grant VI.Veni.192.051 awarded to AIB.

9.3 Authors' Contributions

PRG wrote the manuscript and inline code snippets, performed the analyses, prepared the figures, and developed the `R` package `atlastools`. CEB and AIB collected the calibration track, and EL collected the bat movement data and fruit tree locations. RN conceived the idea of writing this manuscript, and PRG, AIB, OS, CEB, ST, and EL contributed to its design, and the design of `atlastools`. All authors contributed to the writing of the manuscript, and the design of figures.

9.4 Data Availability

The data and source code to reproduce the figures and analyses in this article and in the Supplementary Material can be found in the Zenodo repository at https://doi.org/10.5281/zenodo.4287462.

9.5 Supplementary Material

1. Supplementary Material 1: Code for worked out examples on calibration data from the Dutch Wadden Sea, and bat tracking data from the Hula Valley, Israel.

2. Supplementary Material 2: Manual for the `R` package `atlastools`.

References

Aarts, G., M. MacKenzie, B. McConnell, M. Fedak, and J. Matthiopoulos. 2008. Estimating space-use and habitat preference from wildlife telemetry data. Ecography 31:140–160.

Aspillaga, E., R. Arlinghaus, M. Martorell-Barceló, G. Follana-Berná, A. Lana, A. Campos-Candela, and J. Alós. 2021. Performance of a novel system for high-resolution tracking of marine fish societies. Animal Biotelemetry 9:1.

Avgar, T., J. R. Potts, M. A. Lewis, and M. S. Boyce. 2016. Integrated step selection analysis: Bridging the gap between resource selection and animal movement. Methods in Ecology and Evolution 7:619–630.

Baktoft, H., K. Ø. Gjelland, F. Økland, J. S. Rehage, J. R. Rodemann, R. S. Corujo, N. Viadero, and U. H. Thygesen. 2019. Opening the black box of high resolution fish tracking using yaps. bioRxiv page 2019.12.16.877688.

Baktoft, H., K. Ø. Gjelland, F. Økland, and U. H. Thygesen. 2017. Positioning of aquatic animals based on time-of-arrival and random walk models using YAPS (Yet Another Positioning Solver). Scientific Reports 7:14294.

Barnett, A. H., and P. R. Moorcroft. 2008. Analytic steady-state space use patterns and rapid computations in mechanistic home range analysis. Journal of Mathematical Biology 57:139–159.

Barraquand, F., and S. Benhamou. 2008. Animal movements in heterogeneous landscapes: Identifying profitable places and homogeneous movement bouts. Ecology 89:3336–3348.

Beardsworth, C. E., E. Gobbens, F. van Maarseveen, B. Denissen, A. Dekinga, R. Nathan, S. Toledo, and A. I. Bijleveld. 2021. Validating a high-throughput tracking system: ATLAS as a regional-scale alternative to GPS. bioRxiv page 2021.02.09.430514.

Bijleveld, A. I., R. B. MacCurdy, Y.-C. Chan, E. Penning, R. M. Gabrielson, J. Cluderay, E. L. Spaulding, A. Dekinga, S. Holthuijsen, J. ten Horn, M. Brugge, J. A. van Gils, D. W. Winkler, and T. Piersma. 2016. Understanding spatial distributions: Negative density-dependence in prey causes predators to trade-off prey quantity with quality. Proceedings of the Royal Society B: Biological Sciences 283:20151557.

Bjørneraas, K., B. V. Moorter, C. M. Rolandsen, and I. Herfindal. 2010. Screening Global Positioning System Location Data for Errors Using Animal Movement Characteristics. The Journal of Wildlife Management 74:1361–1366.

Bracis, C., K. L. Bildstein, and T. Mueller. 2018. Revisitation analysis uncovers spatio-temporal patterns in animal movement data. Ecography 41:1801–1811.

Calabrese, J. M., C. H. Fleming, and E. Gurarie. 2016. Ctmm: An r package for analyzing animal relocation data as a continuous-time stochastic process. Methods in Ecology and Evolution 7:1124–1132.

Calenge, C., S. Dray, and M. Royer-Carenzi. 2009. The concept of animals' trajectories from a data analysis perspective. Ecological Informatics 4:34–41.

Dowle, M., and A. Srinivasan. 2020. Data.Table: Extension of 'data.Frame'.

Dupke, C., C. Bonenfant, B. Reineking, R. Hable, T. Zeppenfeld, M. Ewald, and M. Heurich. 2017. Habitat selection by a large herbivore at multiple spatial and temporal scales is primarily governed by food resources. Ecography 40:1014–1027.

Fischler, M. A., and R. C. Bolles. 1981. Random sample consensus: A paradigm for model fitting with applications to image analysis and automated cartography. Communications of the ACM 24:381–395.

Fleming, C. H., J. M. Calabrese, T. Mueller, K. A. Olson, P. Leimgruber, and W. F. Fagan. 2014. From Fine-Scale Foraging to Home Ranges: A Semivariance Approach to Identifying Movement Modes across Spatiotemporal Scales. The American Naturalist 183:E154–E167.

Fleming, C. H., J. Drescher-Lehman, M. J. Noonan, T. S. B. Akre, D. J. Brown, M. M. Cochrane, N. Dejid, V. DeNicola, C. S. DePerno, J. N. Dunlop, N. P. Gould, J. Hollins, H. Ishii, Y. Kaneko, R. Kays, S. S. Killen, B. Koeck, S. A. Lambertucci, S. D. LaPoint, E. P. Medici, B.-U. Meyburg, T. A. Miller, R. A. Moen, T. Mueller, T. Pfeiffer, K. N. Pike, A. Roulin, K. Safi, R. Séchaud, A. K. Scharf, J. M. Shephard, J. A. Stabach, K. Stein, C. M. Tonra, K. Yamazaki, W. F. Fagan, and J. M. Calabrese. 2020. A comprehensive framework for handling location error in animal tracking data*. bioRxiv page 2020.06.12.130195.

Gupte, P. R. 2020. Atlastools: Pre-processing Tools for High Frequency Tracking Data. Zenodo.

Gurarie, E., C. H. Fleming, W. F. Fagan, K. L. Laidre, J. Hernández-Pliego, and O. Ovaskainen. 2017. Correlated velocity models as a fundamental unit of animal movement: Synthesis and applications. Movement Ecology 5:13.

Harel, R., N. Horvitz, and R. Nathan. 2016. Adult vultures outperform juveniles in challenging thermal soaring conditions. Scientific Reports 6:27865.

Holyoak, M., R. Casagrandi, R. Nathan, E. Revilla, and O. Spiegel. 2008. Trends and missing parts in the study of movement ecology. Proceedings of the National Academy of Sciences of the United States of America 105:19060–5.

Hurford, A. 2009. GPS Measurement Error Gives Rise to Spurious 180° Turning Angles and Strong Directional Biases in Animal Movement Data. PLOS ONE 4:e5632.

Hussey, N. E., S. T. Kessel, K. Aarestrup, S. J. Cooke, P. D. Cowley, A. T. Fisk, R. G. Harcourt, K. N. Holland, S. J. Iverson, J. F. Kocik, J. E. Mills Flemming, and F. G. Whoriskey. 2015. Aquatic animal telemetry: A panoramic window into the underwater world. Science 348:1255642–1255642.

Johnson, D. S., J. M. London, M.-A. Lea, and J. W. Durban. 2008. Continuous-Time Correlated Random Walk Model for Animal telemetry Data. Ecology 89:1208–1215.

Jonsen, I. D., J. M. Flemming, and R. A. Myers. 2005. Robust State–Space Modeling of Animal Movement Data. Ecology 86:2874–2880.

Jonsen, I. D., R. A. Myers, and J. M. Flemming. 2003. Meta-Analysis of Animal Movement Using State-Space Models. Ecology 84:3055–3063.

Joo, R., S. Picardi, M. E. Boone, T. A. Clay, S. C. Patrick, V. S. Romero-Romero, and M. Basille. 2020. A decade of movement ecology. arXiv:2006.00110 [q-bio] .

Jung, K. W., Z. D. Deng, J. J. Martinez, D. R. Geist, G. A. McMichael, J. R. Stephenson, and P. J. Graf. 2015. Performance of an acoustic telemetry system in a large fishway. Animal Biotelemetry 3:17.

Langrock, R., R. King, J. Matthiopoulos, L. Thomas, D. Fortin, and J. M. Morales. 2012. Flexible and practical modeling of animal telemetry data: Hidden Markov models and extensions. Ecology 93:2336–2342.

MacCurdy, R., R. Gabrielson, E. Spaulding, A. Purgue, K. Cortopassi, and K. Fristrup. 2009. Automatic Animal Tracking Using Matched Filters and Time Difference of Arrival. JCM 4:487–495.

MacCurdy, R. B., A. I. Bijleveld, R. M. Gabrielson, and K. A. Cortopassi. 2019. Automated Wildlife Radio Tracking. Chap. 33, pages 1219–1261 in Handbook of Position Location. John Wiley & Sons, Ltd.

Marwick, B., C. Boettiger, and L. Mullen. 2018. Packaging Data Analytical Work Reproducibly Using R (and Friends). The American Statistician 72:80–88.

Michelot, T., R. Langrock, and T. A. Patterson. 2016. moveHMM: An R package for the statistical modelling of animal movement data using hidden Markov models. Methods in Ecology and Evolution 7:1308–1315.

Nathan, R., W. M. Getz, E. Revilla, M. Holyoak, R. Kadmon, D. Saltz, and P. E. Smouse. 2008. A movement ecology paradigm for unifying organismal movement research. Proceedings of the National Academy of Sciences 105:19052–19059.

Noonan, M. J., C. H. Fleming, T. S. Akre, J. Drescher-Lehman, E. Gurarie, A.-L. Harrison, R. Kays, and J. M. Calabrese. 2019. Scale-insensitive estimation of speed and distance traveled from animal tracking data. Movement Ecology 7:35.

Oudman, T., T. Piersma, M. V. Ahmedou Salem, M. E. Feis, A. Dekinga, S. Holthuijsen, J. ten Horn, J. A. van Gils, and A. I. Bijleveld. 2018. Resource landscapes explain contrasting patterns of aggregation and site fidelity by red knots at two wintering sites. Movement Ecology 6:24–24.

Papageorgiou, D., C. Christensen, G. E. C. Gall, J. A. Klarevas-Irby, B. Nyaguthii, I. D. Couzin, and D. R. Farine. 2019. The multilevel society of a small-brained bird. Current Biology 29:R1120–R1121.

Patin, R., M.-P. Etienne, E. Lebarbier, S. Chamaillé-Jammes, and S. Benhamou. 2020. Identifying stationary phases in multivariate time series for highlighting behavioural modes and home range settlements. Journal of Animal Ecology 89:44–56.

Patterson, T. A., L. Thomas, C. Wilcox, O. Ovaskainen, and J. Matthiopoulos. 2008. State–space models of individual animal movement. Trends in Ecology & Evolution 23:87–94.

Pebesma, E. 2018. Simple Features for R: Standardized Support for Spatial Vector Data. The R Journal 10:439–446.

R Core Team. 2020. R: A Language and Environment for Statistical Computing. R Foundation for Statistical Computing, Vienna, Austria.

Ranacher, P., R. Brunauer, W. Trutschnig, S. V. der Spek, and S. Reich. 2016. Why GPS makes distances bigger than they are. International Journal of Geographical Information Science 30:316–333.

Seidel, D. P., E. Dougherty, C. Carlson, and W. M. Getz. 2018. Ecological metrics and methods for GPS movement data. International Journal of Geographical Information Science 32:2272–2293.

Signer, J., J. Fieberg, and T. Avgar. 2017. Estimating utilization distributions from fitted step-selection functions. Ecosphere 8:e01771.

Slingsby, A., and E. van Loon. 2016. Exploratory Visual Analysis for Animal Movement Ecology. Computer Graphics Forum 35:471–480.

Stine, P. A., and C. T. Hunsaker. 2001. An Introduction to Uncertainty Issues for Spatial Data Used in Ecological Applications. Pages 91–107 in C. T. Hunsaker, M. F. Goodchild, M. A. Friedl, and T. J. Case, eds. Spatial Uncertainty in Ecology: Implications for Remote Sensing and GIS Applications. Springer, New York, NY.

Strandburg-Peshkin, A., D. R. Farine, I. D. Couzin, and M. C. Crofoot. 2015. Shared decision-making drives collective movement in wild baboons. Science 348:1358–1361.

Toledo, S., O. Kishon, Y. Orchan, Y. Bartan, N. Sapir, Y. Vortman, and R. Nathan. 2014. Lightweight low-cost wildlife tracking tags using integrated transceivers. Pages 287–291 in 2014 6th European Embedded Design in Education and Research Conference (EDERC).

Toledo, S., O. Kishon, Y. Orchan, A. Shohat, and R. Nathan. 2016. Lessons and Experiences from the Design, Implementation, and Deployment of a Wildlife Tracking System. Pages 51–60 in 2016 IEEE International Conference on Software Science, Technology and Engineering (SWSTE).

Toledo, S., D. Shohami, I. Schiffner, E. Lourie, Y. Orchan, Y. Bartan, and R. Nathan. 2020. Cognitive map–based navigation in wild bats revealed by a new high-throughput tracking system. Science 369:188–193.

Tukey, J. W. 1977. Exploratory Data Analysis, vol. 2. Reading, MA.

Visscher, D. R. 2006. GPS measurement error and resource selection functions in a fragmented landscape. Ecography 29:458–464.

Weiser, A. W., Y. Orchan, R. Nathan, M. Charter, A. J. Weiss, and S. Toledo. 2016. Characterizing the Accuracy of a Self-Synchronized Reverse-GPS Wildlife Localization System. Pages 1–12 in 2016 15th ACM/IEEE International Conference on Information Processing in Sensor Networks (IPSN).

Wikelski, M., R. W. Kays, N. J. Kasdin, K. Thorup, J. A. Smith, G. W. Swenson, M. Kurita, and Y. Le Maho. 2007. Going wild: What a global small-animal tracking system could do for experimental biologists. Journal of Experimental Biology 210:181–186.