



Wildlife Seasonal Modelling

Explorations of interpretable seasonal models and the ecological
structures that emerge from them

Version 1.0

© David Walker 2026

This work is licensed under a Creative Commons
Attribution 4.0 International License.

<https://creativecommons.org/licenses/by/4.0/>

You are free to share and adapt the material provided appropriate attribution is
given, a link to the licence is included, and any changes are indicated.

Contents

<i>Licence and Use</i>	3
<i>Wildlife Seasonal Modelling</i>	4
<i>Resident Detectability Model</i>	6
<i>Blackbird</i>	10
<i>Seasonal Presence Model</i>	13
<i>Bluebell</i>	17
<i>Winter Visitor Model</i>	20
<i>Redwing</i>	23
<i>Worked Example – Bluebell (Hyacinthoides non-scripta)</i>	26
<i>Species Similarity and Clustering</i>	31
<i>Ecological Neighbourhoods</i>	32
<i>Seasonal Ecological Calendars</i>	35
<i>Glossary</i>	37

Licence and Use

Principle

Unless otherwise stated, all original material in this booklet is licensed under the Creative Commons licence specified for this work.

The applicable licence is stated on the cover and applies to the booklet as a whole.

Images

Images in this booklet are licensed according to their content:

- **Identifiable individuals** — All rights reserved
- **Non-identifiable individuals** — CC BY-ND 4.0
- **No people present** — CC BY 4.0

The licence for each image is stated in its caption and should be treated as authoritative.

Where no specific licence is given, the default licence is CC BY 4.0.

Use of Material

Unless otherwise stated, you are free to:

- **Share** — copy and redistribute the material
- **Adapt** — remix and build upon the material (*where permitted by the licence*)
- **Use commercially**

These permissions are granted provided that:

- Attribution is given
- The licence is included
- Changes are indicated

A typical attribution should include:

David Walker, *Field Notes Journal*

Exceptions

All material in this booklet is original unless otherwise stated.

Any third-party material will be clearly identified and should be used according to the terms provided alongside it.

Wildlife Seasonal Modelling

The models described in this booklet explore what simple processes might give rise to the patterns seen in seasonal observations. The observed curves used throughout are derived from long-term monthly aggregation of wildlife observation records collected within the study area.

Each model begins with a small set of assumptions — about presence, detectability, and seasonal change — and asks whether these are sufficient to reproduce the curves observed in the data.

Detectability here refers to the likelihood of observing or recording a species, rather than its absolute abundance.

Across species, three distinct ways of occupying the year emerge.

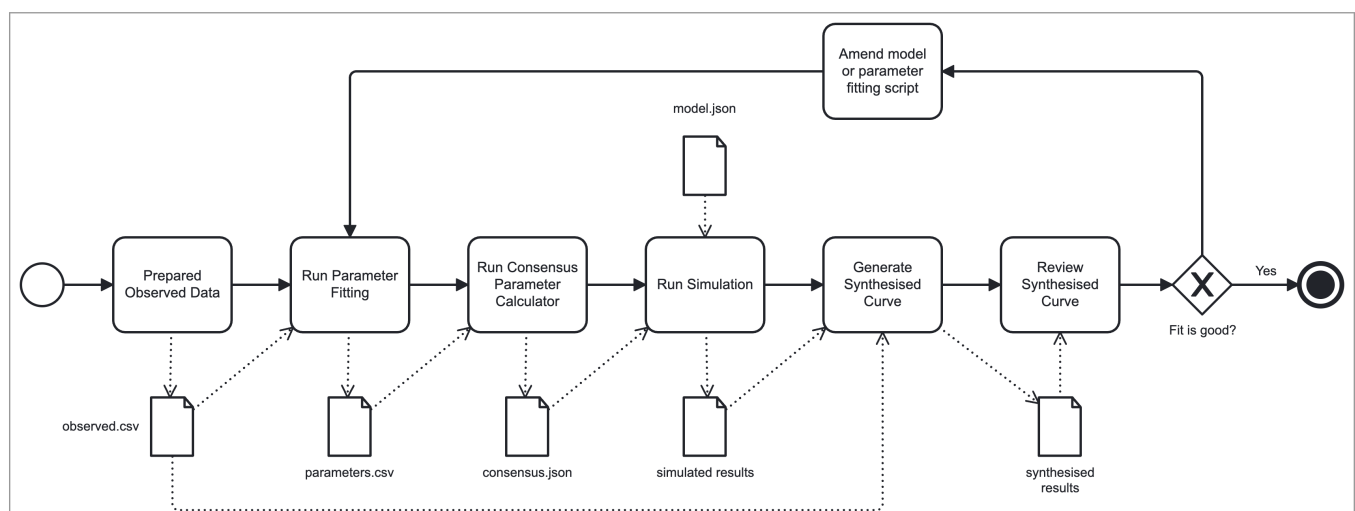
Some species are present only for part of the year, appearing within a defined seasonal window before disappearing again. Others are present throughout the year, but vary in how readily they are observed, their detectability rising and falling with behaviour and conditions. A third group — winter visitors — spans the year boundary, arriving in autumn, peaking in winter, and departing in spring.

These differences suggest three complementary models:

- **Seasonal presence**, where activity is confined to a bounded window
- **Resident detectability**, where presence is continuous but visibility varies
- **Winter presence**, where activity wraps around the year boundary with a winter peak

Each model is deliberately simple. They do not attempt to capture ecological processes in detail, nor are they intended to predict future observations. Instead, they act as a way of testing whether the broad patterns seen in the observed data can arise from straightforward mechanisms.

The modelling workflow is illustrated below:



Wildlife Seasonal Modelling Workflow

A parameter fitting process identifies a search space, defined in terms of the model parameters, and then repeatedly runs the model with parameters selected randomly from within that search space, scoring each run in terms of closeness to the observed data.

On completion, a consensus parameter set is derived from the individual sets that produced the closest matches and this is used to generate a synthesised seasonal curve, one that follows the shape of the simulated output but is scaled onto the observed data scale to allow easy comparison.

This allows each species to be described not only by its pattern, but by a small set of parameters — the consensus set — that together form a simple seasonal “signature”.

The features in that signature describe aspects of seasonal behaviour such as:

- Peak timing
- Seasonal width
- Detectability persistence
- Seasonal symmetry and asymmetry
- Occupancy characteristics
- Derived ecological traits

In this way, the models sit alongside the observations. The observations describe how species occupy the year; the models ask how those patterns might come to be, and provide a consistent way of comparing them.

Resident Detectability Model

This model represents species that are **always present but variably detectable**, describing a continuous presence in which detectability rises and falls through the year without ever reaching zero.

It provides a minimal explanation for patterns seen in the seasonal analysis of observations, showing that variation in observation does not necessarily imply absence, but can arise from:

- Behavioural change
- Seasonal activity patterns
- Variation in visibility

The model is deliberately simple in structure, while still allowing a range of realistic seasonal behaviours to emerge from a relatively small number of interacting processes. It does not attempt to describe detailed ecological mechanisms, but tests whether the observed patterns can arise from a small number of underlying processes.

Concept

This model describes species that are present throughout the year but vary in how readily they are observed.

It answers the question:

How detectable is the species through the year?

Unlike the seasonal and winter visitor models, presence is continuous. What changes is not whether the species is present, but how visible, active, or detectable it is.

The model defines a seasonal target, representing the expected detectability at each point in the year. The observed signal then adjusts towards this target over time, but not instantaneously. Detectability may persist, lag behind seasonal conditions, or decline at different rates through the year.

The target combines:

- A baseline level, representing continuous presence
- A winter peak, reflecting increased visibility or activity
- An optional autumn component
- A summer dip, representing reduced detectability
- Optional spring carry-over behaviour, allowing detectability to remain elevated into late spring and early summer before entering the summer decline

Additional mechanisms allow the model to represent:

- Delayed onset of summer decline
- Prolonged spring persistence
- Asymmetric seasonal rise and fall
- Sharper late-summer reductions in detectability

Together, these produce a continuous annual cycle with no enforced absence.

Model Parameters

A small number of parameters control the behaviour of the model:

Parameter	Purpose
INITIAL_Y	Sets the starting value of the modelled detectability signal
BASELINE	Sets the persistent year-round level of detectability
WINTER_WEIGHT	Controls the strength of the winter / early-spring peak
AUTUMN_WEIGHT	Controls the strength of the autumn / early-winter recovery
SUMMER_DIP	Controls the strength of the summer reduction in detectability
WINTER_PEAK	Sets the timing of the winter / early-spring peak
AUTUMN_PEAK	Sets the timing of the autumn / early-winter recovery
SUMMER_LOW	Sets the timing of the lowest summer detectability
WINTER_WIDTH	Controls the breadth of the winter / early-spring peak
AUTUMN_WIDTH	Controls the breadth of the autumn / early-winter recovery
SUMMER_WIDTH	Controls the breadth of the summer dip
GROWTH_RATE	Controls how quickly detectability rises towards the seasonal target
DECAY_RATE	Controls how quickly detectability falls towards the seasonal target

Together, these parameters define the level, timing, strength, breadth, and responsiveness of the annual detectability cycle.

The peak and low-point parameters are expressed in months on a circular 12-month scale. Width parameters control how broad or concentrated each seasonal feature is. The growth and decay rates allow the model to respond asymmetrically, so that detectability can rise and fall at different speeds.

Extended Seasonal Dynamics

Some species require additional seasonal persistence behaviour.

Optional parameters allow the model to represent:

Behaviour	Purpose
Spring carry-over	Retains elevated detectability into spring and early summer
Delayed summer decline	Prevents summer suppression from beginning too early
Summer decay boost	Allows sharper late-summer reduction
Reduced pre-summer decay	Slows spring decline before the summer low

These mechanisms are mainly important for species with broad spring plateaus or delayed seasonal decline.

Mathematical Form

The model is a first-order system:

$$dy/dt = \text{rate} \times (\text{target}(t) - y)$$

Where:

- $y(t)$ is a relative, dimensionless measure of observable activity
- $\text{target}(t)$ is the seasonal detectability target
- rate controls how quickly the system responds

The target function is constructed from smooth annual components:

$$\text{target}(t) = \text{BASELINE} + \text{winter}(t) + \text{autumn}(t) - \text{summer}(t)$$

Each component is a smooth periodic function over a 12-month cycle. Time is treated as circular, allowing continuous seasonal variation without a defined start or end.

Model Behaviour

When applied over a full year, the model produces a smooth, continuous cycle:

- Detectability rises into a winter or early-spring peak
- Declines through the summer months
- Recovers again towards autumn and winter

Unlike seasonal presence models, the signal does not collapse to zero. Instead, it fluctuates around a persistent baseline, reflecting continuous presence.

The exact shape depends on:

- The timing and strength of seasonal components
- The depth and duration of the summer dip
- The rate at which the system responds to change

Seasonal Persistence

Some resident species do not simply track seasonal conditions directly. Instead, detectability may persist for some time before declining.

For example:

- Some species retain high visibility through spring before entering a relatively abrupt late-summer reduction
- Others follow the seasonal cycle more closely, with a smoother decline into summer

The model allows these behaviours through mechanisms controlling:

- Delayed summer suppression
- Retained spring detectability
- Asymmetric decay rates

This allows the same general framework to describe both:

- Broad, persistent annual patterns
- More rapidly responding resident species

without introducing seasonal absence.

Normalisation

Model outputs are expressed as a relative measure of activity.

To allow comparison across species, results are normalised so that:

- 1.0 → peak activity
- 0.5 → half of peak activity
- 0.0 → effectively zero

This focuses attention on the timing and shape of seasonal variation rather than absolute magnitude.

Parameter Interpretation

After parameter fitting, the parameters are broadly interpretable as follows:

- **WINTER_PEAK** → timing of highest detectability
- **SUMMER_LOW** → timing of lowest detectability
- **WIDTH parameters** → how concentrated or diffuse seasonal features are
- **WEIGHTS** → relative strength of seasonal effects

Together, they describe the *shape* of the species' seasonal behaviour.

As with all simple models:

- Parameters should be treated as estimates rather than exact dates
- Different combinations may produce similar curves
- Interpretation is most reliable when considered alongside the fitted curve itself

In practice, each species can be described by both:

- Its fitted parameters
- The shape of its simulated seasonal curve

Together, these form a compact description of seasonal presence.

An example showing observed data, model output and identified species signature for the blackbird (*Turdus merula*) is shown on the next page.

Blackbird

Seasonal Analysis and Species Classification

Model Family : Resident detectability

Summary

Blackbird

Resident with spring persistence and summer suppression

Blackbird is classified as resident with spring persistence and summer suppression. The fitted resident detectability target peaks around March and reaches its lowest point around September. The model indicates strong baseline presence, strong spring carry-over, strong pre-summer retention, strong summer suppression, and strong summer decay acceleration.

Confidence	Medium
Fit score	0.227
Peak detectability	March
Lowest detectability	September

Traits

resident detectability pattern

strong baseline presence

spring detectability peak

autumn detectability trough

strong spring carryover

strong summer suppression

strong summer decay acceleration


strong pre summer retention

weak autumn component

meaningful year end component

rapid decline biased response dynamics

Seasonal wheel



Calendar strip

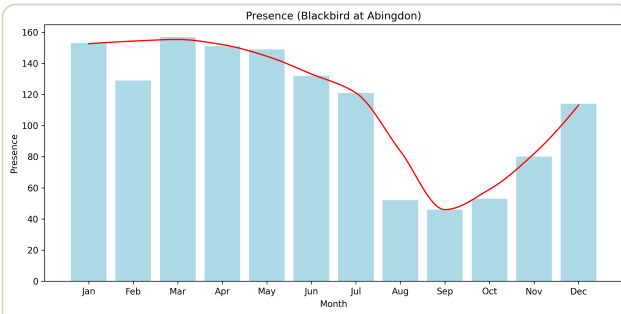


Highlighted months indicate stronger modelled presence or detectability. A ring marks the fitted peak; a hollow mark indicates the trough where available.

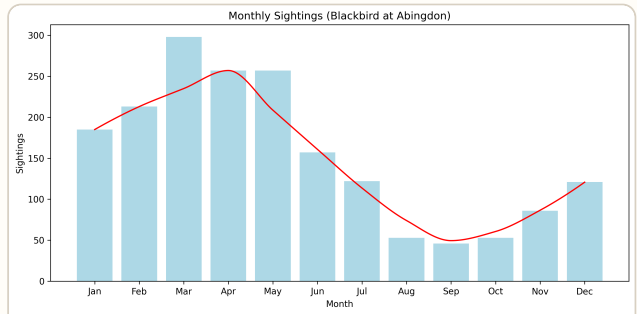
Classification evidence

Target peak month	3
Target peak label	March
Target trough month	9
Target trough label	September
Target peak value	0.963
Target trough value	0.187
Target mean value	0.576
Target amplitude	0.776
Baseline to peak ratio	0.383
Autumn to winter weight ratio	0.063
Year end to winter weight ratio	0.576
Decay to growth ratio	1.214

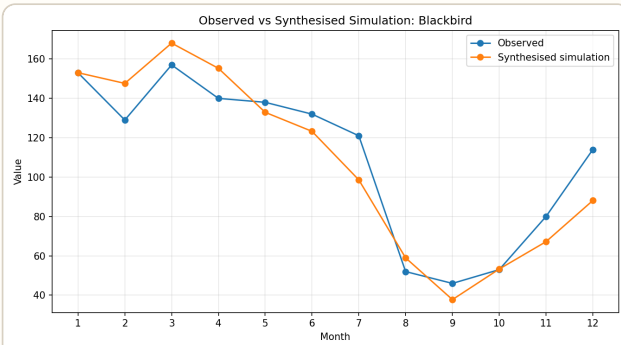
Observed and Simulated Seasonal Patterns



Blackbird Observed Presence, Abingdon



Blackbird Observed Totals, Abingdon



Blackbird Simulated Presence, Abingdon

Seasonal Presence Model

This model is intended to represent species whose presence is **seasonally constrained** — those that are only observable during a particular phase of their life cycle or migration - and describes their rise into activity, the peak of the season, and the eventual collapse back toward absence.

It is therefore intended for species whose observable presence is strongly constrained in time, such as:

- Spring flowers
- Migratory birds
- Butterflies with constrained annual flight periods

The model is deliberately simple in structure - closer to a minimal representation than a description of detailed ecological mechanisms — and is intended to explore whether the observed patterns can arise from a small number of underlying processes, not to predict observations.

It provides a minimal explanation for patterns seen in the seasonal analysis of observations, showing that a small number of simple processes are sufficient to produce:

- Sharply bounded flowering periods
- Migration-driven appearances
- Other forms of seasonal presence

The model does not attempt to describe the underlying biological mechanisms in detail. Instead, it offers a way of understanding how the observed patterns might arise from the interaction of seasonal forcing, constrained availability, persistence, and active seasonal suppression.

Concept

This model describes species that are only present for part of the year.

It answers the question:

When is the species present?

Unlike the resident model, presence is not continuous. Instead, activity is confined to a defined seasonal window.

The model combines four simple elements:

- A **seasonal driver**, representing environmental change through the year
- A **seasonal window**, constraining when presence is biologically possible
- A **baseline decay process**, limiting persistence over time
- A **post-season suppression phase**, accelerating decline once the active season has passed

Together, these produce a dominant seasonal pulse with realistic asymmetry — allowing activity to rise gradually, peak, and then decline more abruptly after the season ends.

Model Parameters

A small number of parameters control the behaviour of the model:

Parameter	Purpose
GROWTH	Controls how strongly seasonal conditions drive the appearance of the species
DECAY	Controls how quickly activity declines during the active period
OOS_DECAY	Increases the rate of decline outside the seasonal window
POST_PEAK_DECAY	Controls how strongly activity is suppressed after the season
POST_PEAK_SHARPNESS	Controls how abruptly post-season suppression activates
SEASON_START	Sets the onset of the active period
SEASON_END	Sets the end of the active period
SHARPNESS	Controls how abruptly the season begins and ends
FORCING_PEAK	Sets the timing of the seasonal driver peak

Together, these parameters define:

- **When** the species appears (SEASON_START / END)
- **How sharply** it appears and disappears (SHARPNESS, OOS_DECAY)
- **How strongly** it responds to seasonal conditions (GROWTH, FORCING_PEAK)
- **How long it persists** once present (DECAY)

The start and end parameters are expressed in months on a continuous scale, allowing the season to be positioned precisely within the year.

Mathematical form

The governing equation combines:

- Seasonal growth
- Seasonal availability
- Time-dependent decline

The effective decay rate increases:

- Outside the seasonal window
- After the active season has passed

This allows the model to reproduce species that decline rapidly after peak activity, reducing unrealistically persistent late-season tails.

Model behaviour

When applied over a full year, the model produces a characteristic seasonal pulse:

- Activity begins as environmental conditions become favourable
- Presence increases within the active window
- Activity reaches a peak during the core season
- Once seasonal conditions deteriorate, activity collapses more rapidly toward absence

The resulting curves are often asymmetric, with relatively sharp post-peak decline. This better reflects many biological systems, where activity does not merely fade gradually, but undergoes active seasonal shutdown through senescence, migration, mortality, or behavioural change.

- **SEASON_START / END** → timing of arrival and departure
- **FORCING_PEAK** → timing of peak activity
- **SHARPNESS** → how abruptly the season begins and ends
- **OOS_DECAY** → how quickly activity falls away outside the season

Normalisation

The model outputs are expressed as a relative, dimensionless measure of activity. As a result, different models — and different parameter choices — can produce values on different scales.

To allow comparison, the outputs are normalised so that the maximum value of y is set to 1, with all other values scaled proportionally.

This produces a simple index:

- 1.0 → peak modelled activity
- 0.5 → half of peak activity
- 0.0 → absence (or effectively zero)

Normalisation is applied after simulation, and does not affect the underlying model behaviour. It allows comparison of the *shape* and timing of seasonal patterns, rather than absolute magnitude.

Parameter Interpretation

After parameter fitting, the parameters are broadly interpretable as follows:

- **SEASON_START / END** → approximate timing of arrival and disappearance
- **FORCING_PEAK** → timing of strongest seasonal forcing
- **SHARPNESS** → abruptness of seasonal onset and termination
- **GROWTH / DECAY** → persistence and responsiveness during the active season
- **OOS_DECAY** → suppression strength outside the seasonal window
- **POST_PEAK_DECAY** → strength of late-season collapse
- **POST_PEAK_SHARPNESS** → abruptness of post-season shutdown

Together, they describe the *shape* of the species' seasonal behaviour.

As with all simple models:

- Parameters should be treated as estimates rather than exact dates
- Different combinations may produce similar curves
- Interpretation is most reliable when considered alongside the fitted curve itself

In practice, each species can be described by both:

- Its fitted parameters
- The shape of its simulated seasonal curve

Together, these form a compact description of seasonal presence.

An example showing observed data, model output and identified species signature for the bluebell (*Hyacinthoides non-scripta*) is shown on the next page.

Bluebell

Seasonal Analysis and Species Classification

Model Family : Seasonal presence

Summary

Bluebell

Narrow spring seasonal presence

Bluebell is classified as narrow spring seasonal presence. The fitted seasonal window runs from about April to June, with a spring peak around May. The season is narrow, with a sharp active window, moderate post-peak decline, and strong off-season suppression.

Confidence	High
Fit score	0.047
Peak	May
Season	April–June

Traits

spring peak narrow season

sharp seasonal window

moderate post peak decline


strong offseason suppression

central peak alignment

Seasonal wheel



Calendar strip




Highlighted months indicate stronger modelled presence or detectability. A ring marks the fitted peak; a hollow mark indicates the trough where available.

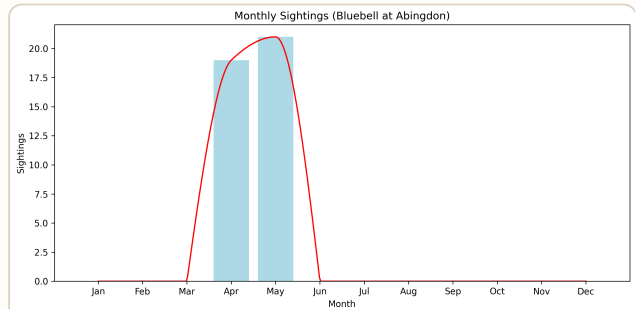
Classification evidence

Season start month	4.185
Season end month	5.595
Forcing peak month	4.88
Season width months	1.41
Season midpoint month	4.89
Season start label	April
Season end label	June
Forcing peak label	May

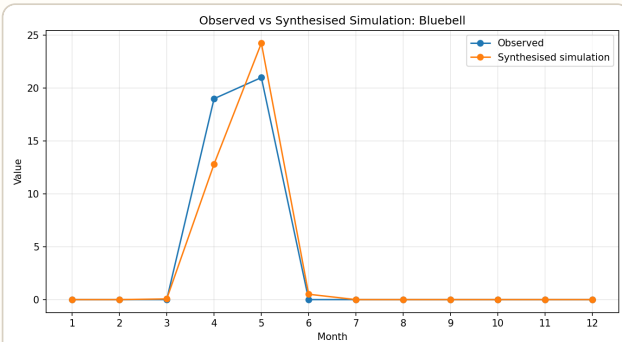
Observed and Simulated Seasonal Patterns



Bluebell Observed Presence, Abingdon



Bluebell Observed Totals, Abingdon



Bluebell Simulated Presence, Abingdon

Winter Visitor Model

This model represents species that are **present only during the winter period**, with a winter peak, activity spanning the year boundary, and near-absence through spring and summer.

It provides a minimal explanation for patterns seen in the seasonal analysis of observations, showing that a small number of simple processes can produce:

- Winter-centred presence
- Distinct arrival phases
- Extended absence through summer

The model does not attempt to describe detailed ecological mechanisms. Instead, it offers a way of understanding how seasonal structure and timing combine to produce the observed patterns.

Concept

This model describes species that are present during the winter months, with activity spanning the end and beginning of the year.

It answers the question:

When is the species present?

Like the seasonal model, presence is limited to part of the year. Unlike it, the active period crosses the year boundary.

The model defines a **seasonal target**, representing expected activity through the year. The observed signal then adjusts towards this target over time.

The target combines:

- A **winter component**, representing peak presence
- An optional **autumn component**, representing arrival
- A **summer suppression**, reducing activity during the off-season

Together, these produce a cycle that rises through autumn, peaks in winter, and falls away into spring.

Model Parameters

A small number of parameters control the behaviour of the model:

Parameter	Purpose
INITIAL_Y	Sets the starting value of the modelled signal
BASELINE	Sets any persistent background level (typically near zero for winter visitors)
WINTER_WEIGHT	Controls the strength of the winter peak
AUTUMN_WEIGHT	Controls the strength of the autumn arrival phase
SUMMER_DIP	Controls the strength of the summer suppression
WINTER_PEAK	Sets the timing of peak winter presence
AUTUMN_PEAK	Sets the timing of autumn arrival
SUMMER_LOW	Sets the timing of lowest summer activity
WINTER_WIDTH	Controls how concentrated the winter peak is
AUTUMN_WIDTH	Controls the breadth of the arrival phase
SUMMER_WIDTH	Controls the breadth of the summer low
GROWTH_RATE	Controls how quickly activity rises towards the seasonal target
DECAY_RATE	Controls how quickly activity declines

Together, these parameters define:

- **When** the species arrives and peaks
- **How concentrated** the winter period is
- **How gradually or abruptly** the species appears and disappears
- **How strongly** the species is absent through summer

All timing parameters are expressed in months on a circular 12-month scale.

Mathematical Form

The model is a first-order system:

$$dy/dt = \text{rate} \times (\text{target}(t) - y)$$

Where:

- $y(t)$ is a relative, dimensionless measure of observable activity
- $\text{target}(t)$ is the seasonal activity target
- rate is selected using `GROWTH_RATE` or `DECAY_RATE` depending on whether the signal is rising or falling

The target function is constructed from smooth periodic components:

$$\text{target}(t) = \text{winter}(t) + \text{autumn}(t) - \text{summer}(t) + \text{BASELINE}$$

Each component is a smooth function over a 12-month cycle, allowing continuous variation without discontinuities.

Model Behaviour

When applied over a full year, the model produces a winter-centred cycle:

- Activity rises through autumn as the species arrives

- Peaks in mid-winter
- Declines through late winter and early spring
- Remains close to zero through summer

The shape depends on:

- The timing and strength of the winter and autumn components
- The breadth of the winter peak
- The strength of summer suppression
- The rate at which the system responds to change

Unlike the seasonal presence model, the season is not bounded within a single part of the calendar year. Instead, it wraps across the year boundary.

Normalisation

Model outputs are expressed as a relative measure of activity.

To allow comparison across species, results are normalised so that:

- 1.0 → peak activity
- 0.5 → half of peak activity
- 0.0 → effectively zero

This focuses attention on the timing and shape of seasonal variation.

Parameter Interpretation

After parameter fitting, the parameters are broadly interpretable as follows:

- **WINTER_PEAK** → timing of peak presence
- **AUTUMN_PEAK** → timing of arrival
- **WINTER_WIDTH** → concentration of winter activity
- **AUTUMN_WIDTH** → spread of the arrival phase
- **SUMMER_DIP / SUMMER_LOW** → strength and timing of absence

As with all simple models:

- Parameters should be treated as estimates rather than exact dates
- Different combinations may produce similar curves
- Interpretation is most reliable when considered alongside the fitted curve itself

In practice, each species can be described by both:

- Its fitted parameters
- The shape of its simulated seasonal curve

Together, these form a compact description of seasonal presence.

An example, the redwing (*Turdus iliacus*), is shown on the next page.

Redwing

Seasonal Analysis and Species Classification

Model Family : Winter presence

Summary

Redwing

Winter visitor with autumn arrival component

Redwing is classified as winter visitor with autumn arrival component. The fitted winter component peaks around February, with a moderate autumn component centred around December. The model has low baseline presence and moderate summer suppression centred around July. The fitted response dynamics suggest slow arrival fast departure.

Confidence	Medium
Fit score	0.109
Winter peak	February
Lowest detectability	May

Traits

year wrapping winter presence

core winter winter peak

moderate autumn component


moderate summer suppression

low baseline presence

moderate winter bump

slow arrival fast departure response dynamics

Seasonal wheel



Calendar strip

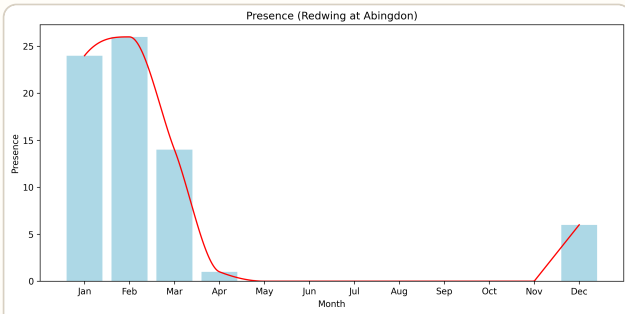


Highlighted months indicate stronger modelled presence or detectability. A ring marks the fitted peak; a hollow mark indicates the trough where available.

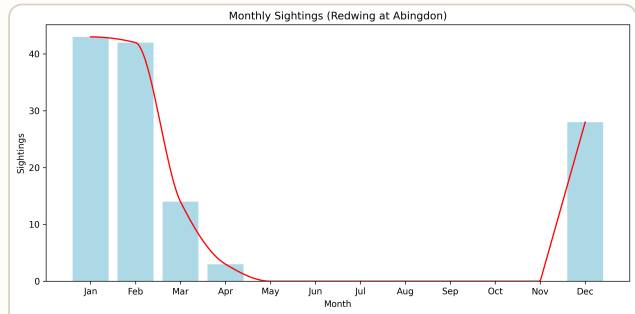
Classification evidence

Winter peak month	1.53
Winter peak label	February
Autumn peak month	11.67
Autumn peak label	December
Summer low month	6.55
Summer low label	July
Autumn to winter weight ratio	0.266
Decay to growth ratio	3.723
Target peak month	1
Target peak label	January
Target trough month	5
Target trough label	May

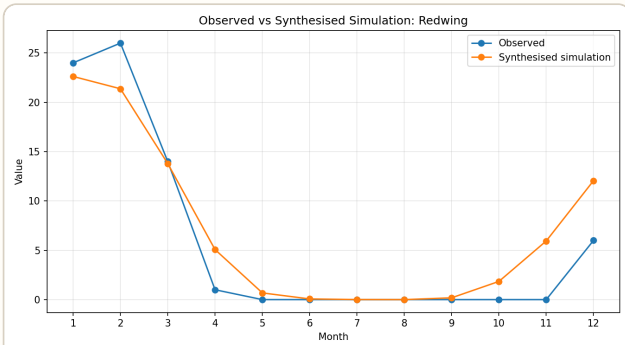
Observed and Simulated Seasonal Patterns



Redwing Observed Presence, Abingdon



Redwing Observed Totals, Abingdon



Redwing Simulated Presence, Abingdon

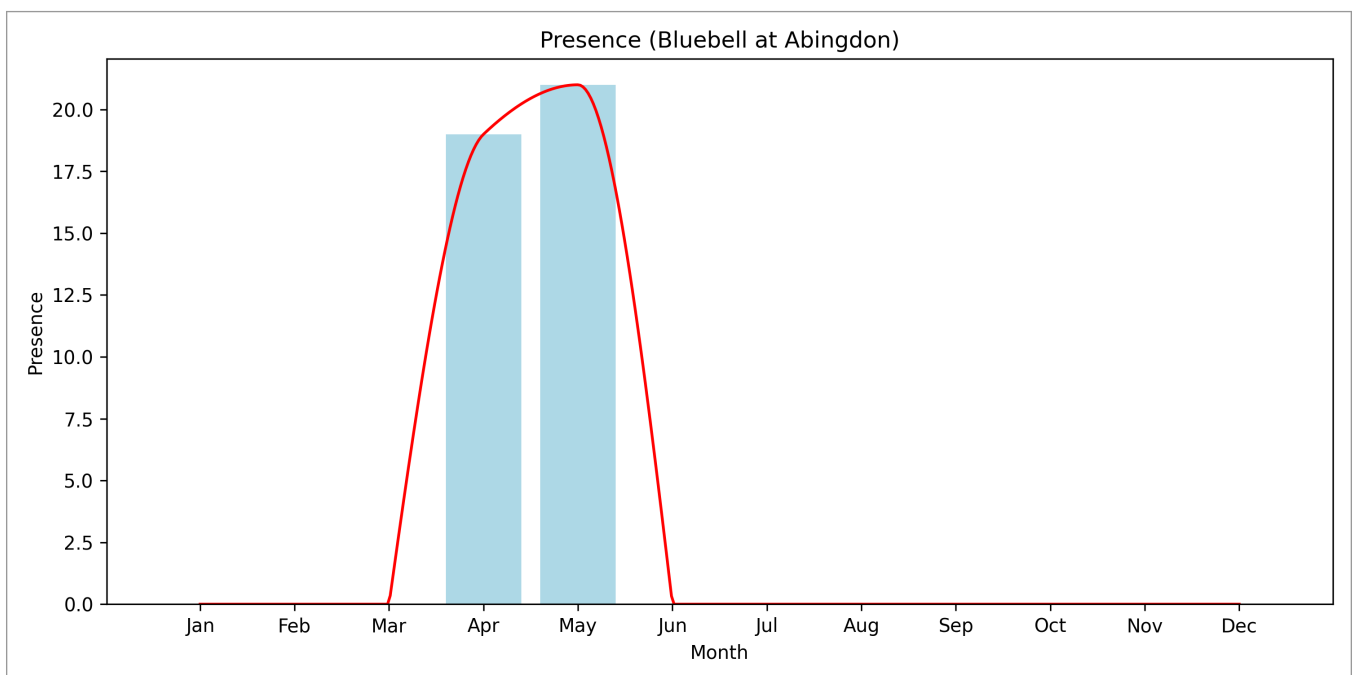
Worked Example — Bluebell (*Hyacinthoides non-scripta*)

The following example illustrates how a single species moves through the modelling workflow, from observed seasonal data to fitted model, classification, and ecological neighbourhood analysis.

The purpose is not to reproduce the technical implementation in detail, but to show how the different stages of the workflow connect together interpretively.

Observed Seasonal Pattern

The starting point is the observed seasonal curve derived from long-term monthly observations.



Observed monthly seasonal pattern for bluebell

The curve shows:

- Near-absence through autumn and winter
- Rapid emergence during spring
- A sharply bounded flowering peak
- Rapid decline after the flowering period ends

This immediately suggests a strongly seasonal presence structure rather than a continuously detectable resident pattern.

Initial Classification

The observed curve is first examined using a rule-based classification stage.

The bluebell pattern shows several strongly seasonal characteristics:

- A narrow active period
- Extended absence outside the flowering season
- Strong asymmetry between emergence and decline
- Low year-round occupancy

These features place the species within the **Seasonal Presence Model** family.

Parameter Fitting

Once a model family has been selected, the parameter fitting process begins.

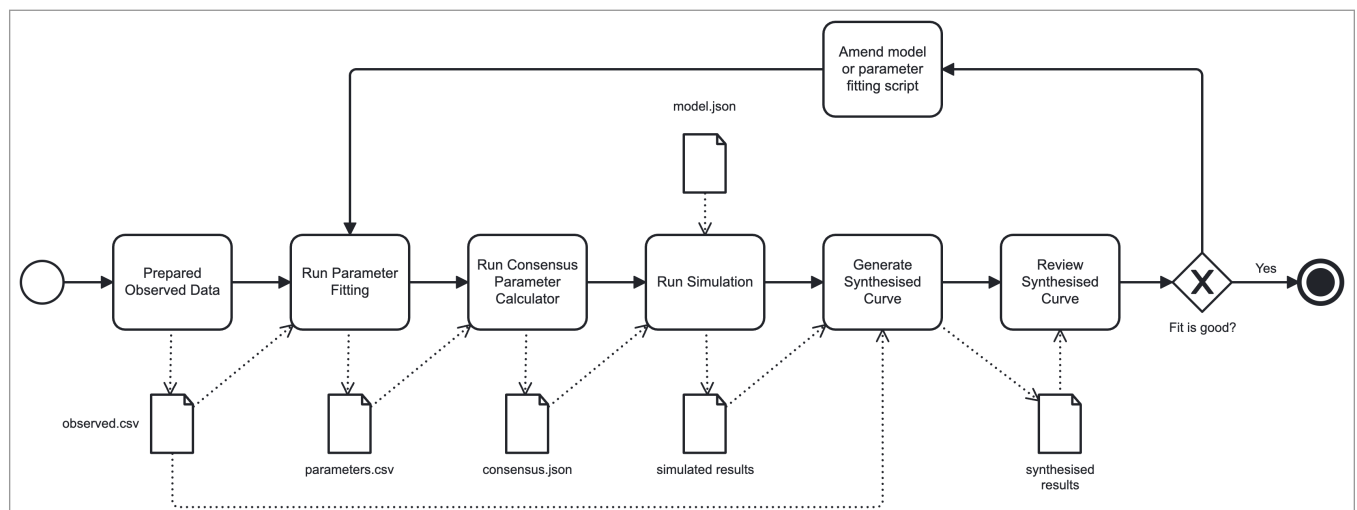
A search space is constructed describing plausible ranges for:

- Seasonal onset and end
- Peak timing
- Seasonal sharpness
- Growth and decay behaviour
- Post-season suppression strength

The model is then repeatedly simulated using randomly selected parameter combinations.

Each simulation is compared against the observed curve and scored according to:

- Overall curve similarity
- Peak timing agreement
- Seasonal width agreement
- Agreement in low-activity months



Parameter fitting and model comparison process

The highest-scoring parameter sets are retained for further analysis.

Consensus Derivation

Rather than selecting a single “best” parameter set, the workflow derives a consensus seasonal description from multiple high-scoring simulations.

This reduces the influence of unstable individual fits and produces a more representative description of the species’ seasonal structure.


The resulting consensus parameters describe features such as:

- Approximate seasonal onset
- Seasonal duration
- Timing of peak activity
- Abruptness of post-season decline

These parameters together form a compact seasonal signature for the species.

Fitted Seasonal Curve

Using the consensus parameter set, a synthesised seasonal curve can be generated and compared directly against the observed data.



Observed and fitted seasonal curves for bluebell

The fitted curve reproduces the main structural characteristics of the observed seasonal pattern:

- Rapid spring emergence
- A concentrated flowering peak
- Strong post-season collapse
- Extended seasonal absence

The aim is not exact prediction, but structural agreement with the observed ecological pattern.

Feature Extraction

Once fitted, the species can be represented within a broader ecological feature space.

The fitted model and observed seasonal characteristics are converted into features describing:

- Seasonal timing
- Width and concentration
- Persistence behaviour
- Occupancy structure
- Seasonal asymmetry

These features allow comparison with species from other model families using a shared seasonal representation.

Ecological Neighbourhoods

Within the resulting ecological space, bluebell clusters with species sharing similar seasonal timing and structural behaviour.

These may include:

- Other spring-flowering plants
- Species with sharply bounded emergence periods
- Species exhibiting strong seasonal concentration

The neighbourhood therefore reflects similarity of seasonal ecological structure rather than taxonomic relationship.

Interpretation

This example illustrates the broader aim of the modelling workflow.

The system does not attempt to reconstruct detailed ecological mechanisms directly. Instead, it asks whether relatively simple seasonal processes are sufficient to reproduce the large-scale structures seen in long-term observational data.

In doing so, the workflow allows species to be described not only individually, but also as components of a larger seasonal ecological landscape.

Further Exploration

The modelling workflow described in this booklet is implemented using a collection of Python-based tools developed as part of the wider Field Notes Journal project.

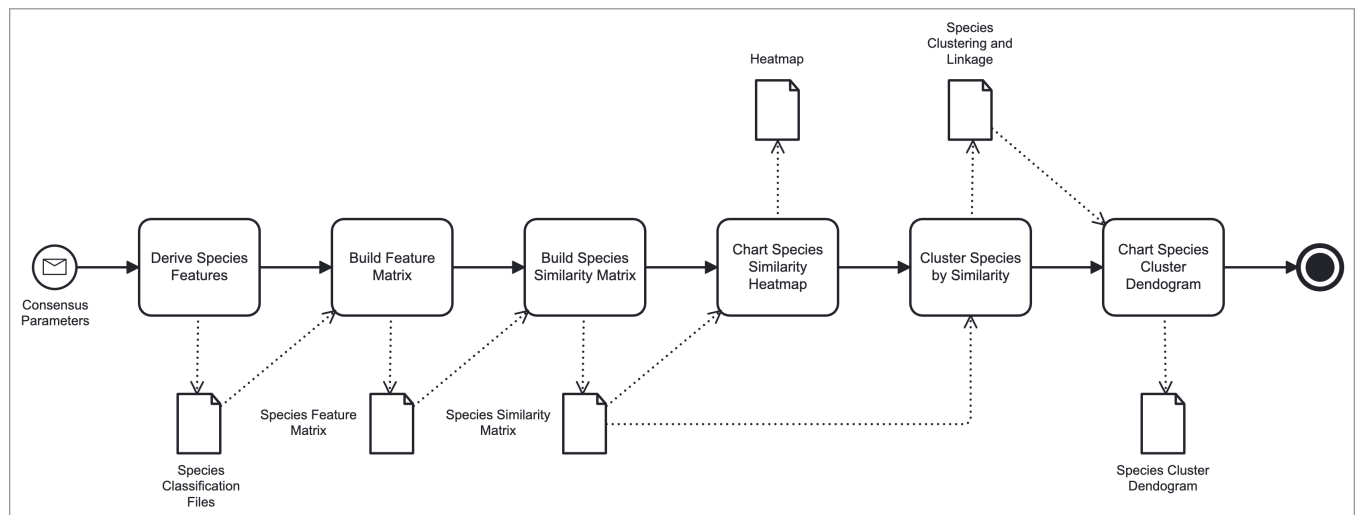
The underlying ODE solver, fitting framework, and analysis workflows are available online for readers interested in exploring or adapting the system for their own observations and seasonal datasets.

- ODE Solver Repository - <https://github.com/davewalker5/OdeSolver>

Species Similarity and Clustering

Once species have been fitted and classified, the resulting parameter sets and observed seasonal characteristics can be converted into a common ecological feature space that allows species from different model families to be compared using a shared representation of seasonal ecological structure.

The workflow is illustrated, below:



Feature Extraction, Species Similarity and Clustering Workflow

The classifications for individual species are collated to create a *feature matrix* that is analysed to produce a *similarity matrix*, comparing species using weighted ecological similarity metrics.

The aim is not to identify taxonomic similarity, but similarity of seasonal ecological signal:

- Shared timing structure
- Similar persistence behaviour
- Overlapping flowering or migration periods
- Comparable detectability dynamics
- Broad phenological synchrony

This allows relationships to emerge not only within groups, but also across the ecosystem as a whole. Butterfly flight periods, flowering seasons, migratory arrival windows, and resident detectability patterns may all align within the same seasonal structure.

In this sense, the system becomes less a species comparison tool and more an exploration of the seasonal organisation of ecological activity.

Ecological Neighbourhoods

Once pairwise similarities have been calculated, the resulting structures can be explored using clustering, dendrograms, and similarity heatmaps.

Rather than grouping species taxonomically, these approaches group species according to similarity of seasonal ecological structure. Species occupying nearby regions of seasonal ecological space form what may be thought of as ecological neighbourhoods — groups sharing broadly similar timing, persistence, detectability dynamics, or seasonal occupancy patterns.

The workflow therefore doesn't ask:

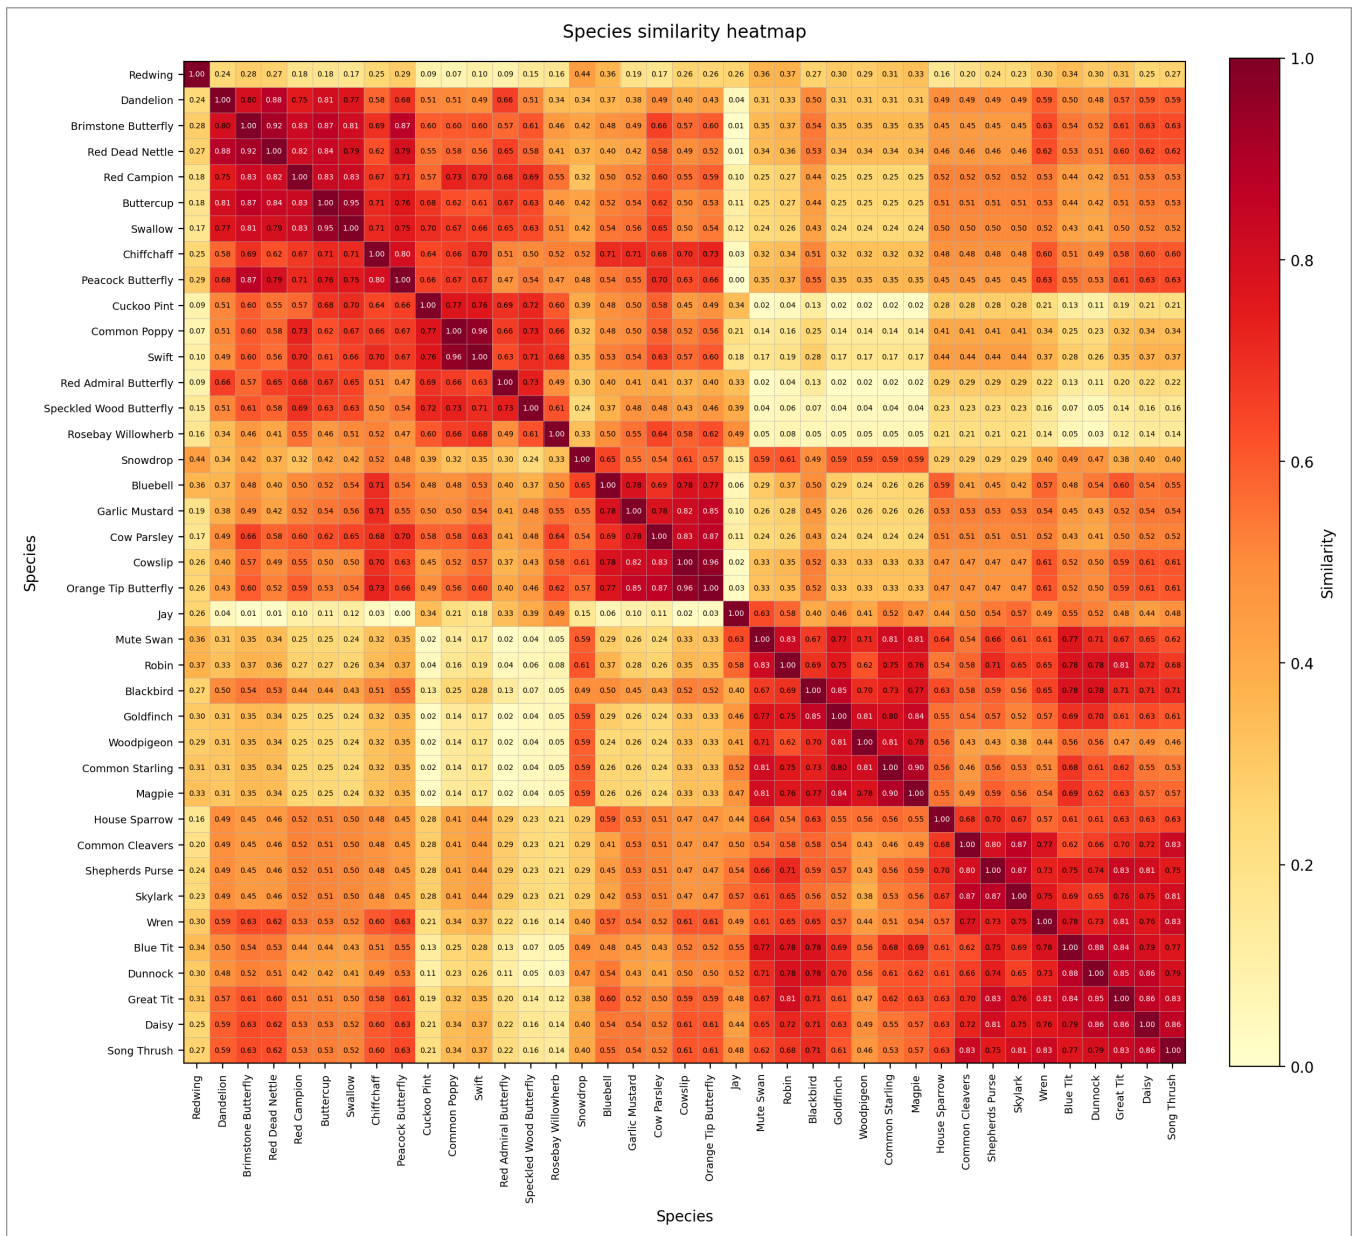
Which species are biologically related?

but, rather:

Which species occupy similar positions within the ecological year?

This allows relationships to emerge across very different groups of organisms. Flowering plants, migratory birds, butterflies, and resident species may all cluster together if their seasonal ecological behaviour follows similar temporal structure.


The resulting similarity matrix can be visualised as a heatmap:



Species Similarity Heatmap

In the heatmap, species are ordered according to hierarchical clustering, causing similar seasonal structures to group together visually. Darker regions indicate stronger similarity within seasonal ecological space.

The same similarity structure can also be represented as a dendrogram:



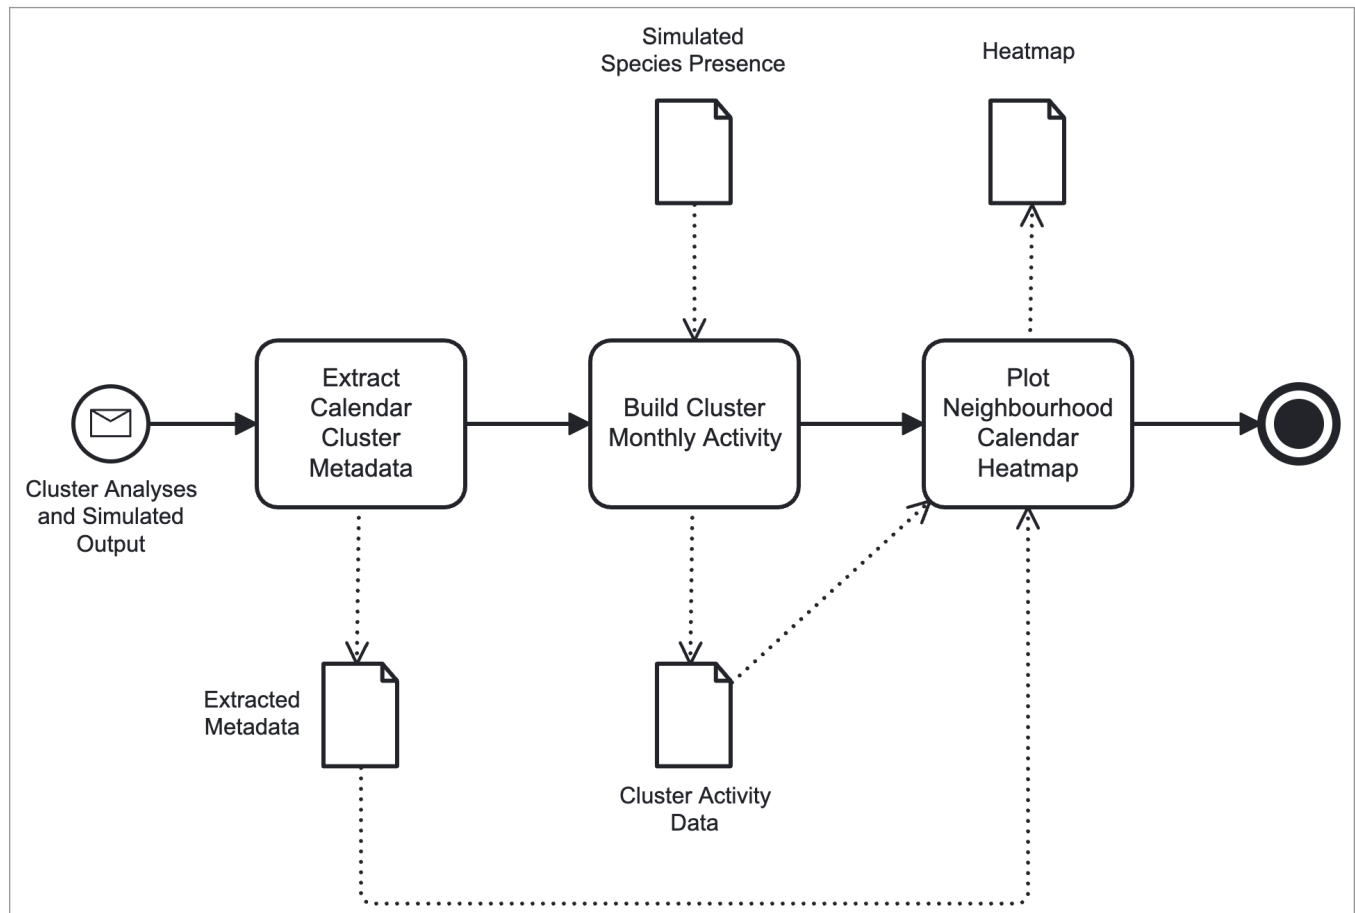
Species Clustering Dendrogram

The dendrogram exposes nested neighbourhood structure within the seasonal ecological system. Species joined by short branch distances share more similar seasonal behaviour, while larger branch separations indicate increasingly distinct ecological timing and structure.

These neighbourhoods should be interpreted as exploratory seasonal assemblages rather than strict ecological or taxonomic categories. They provide a way of examining how seasonal activity is organised across the ecological year, exposing recurring timing structures, overlapping seasonal regimes, and broader phenological organisation within the observed system.

Seasonal Ecological Calendars

The neighbourhood structures can also be aggregated temporally to produce seasonal ecological calendars using the following workflow:



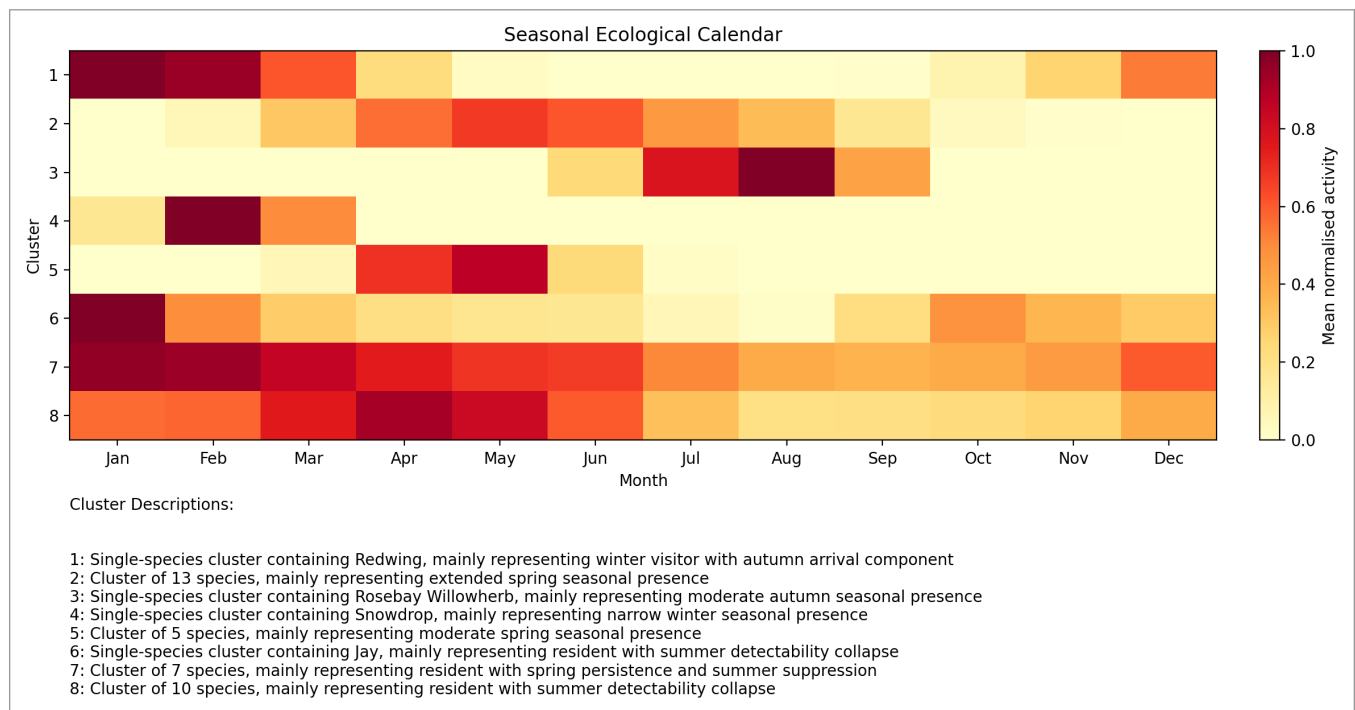
Seasonal Neighbourhood Analysis Workflow

These calendars summarise the mean normalised activity of ecological neighbourhoods through the year, allowing broader community-scale seasonal structure to be visualised directly.

Rather than focusing on individual taxa, they attempt to expose larger seasonal ecological modes, including:

- Winter visitor activity
- Spring flowering and emergence periods
- Resident detectability dynamics
- Extended summer assemblages
- Autumn transitional structure

The resulting heatmaps provide a view of how different regions of seasonal ecological space become active, overlap, and decline through the ecological year:



Cluster Seasonal Activity Heatmap

Toward a Seasonal Ecology of Place

These models do not attempt to explain ecosystems in full. Instead, they provide a way of describing seasonal ecological structure using a common mathematical language. In doing so, they allow long-term observational records to be explored not only as individual species histories, but as components of a larger seasonal ecological system.

Glossary

The following terms are used throughout this booklet in a specific modelling and ecological sense.

Activity Curve

A curve describing how the relative observed activity or detectability of a species changes through the year.

Baseline

The persistent background level of activity or detectability present throughout the year.

Circular Month Distance

A measure of separation between two points in the annual cycle that treats the year as circular, so that December and January are adjacent rather than distant.

Circular Time

A representation of the year in which December and January are treated as adjacent, allowing seasonal cycles to wrap continuously across the year boundary.

Clustering

A method used to group species with similar seasonal ecological characteristics into broader neighbourhood structures.

Consensus Derivation

The process of combining multiple well-fitting parameter sets to produce a representative seasonal description for a species.

Consensus Parameters

A representative parameter set derived from multiple well-fitting model runs, intended to capture the broad seasonal structure of a species.

Cosine Forcing

A smooth seasonal driver based on a cosine-shaped annual cycle, used to represent gradual environmental change through the year.

Detectability

The likelihood of observing or recording a species, rather than a direct measure of absolute abundance.

Ecological Feature Space

A shared mathematical representation in which species are positioned according to their seasonal ecological characteristics.

Ecological Neighbourhood

A group of species occupying nearby regions of seasonal ecological space and sharing broadly similar seasonal structure.

Feature Matrix

A table describing species using fitted parameters and derived seasonal characteristics, allowing comparison across model families.

Fitted Curve

The simulated seasonal curve produced by the model after parameter fitting.

Fitting Method

The overall procedure used to compare model simulations against observed seasonal data and identify parameter sets producing the closest agreement.

Model Family

One of the three broad seasonal model types used in the booklet:

- Seasonal presence
- Resident detectability
- Winter presence

Normalisation

The process of scaling model outputs so that the peak value equals 1.0, allowing comparison of seasonal shape independent of magnitude.

Normalisation Method

The procedure used to scale model outputs onto a common relative range, allowing seasonal patterns to be compared independent of absolute magnitude.

Occupancy

The extent to which a species is present or observable across the annual cycle.

ODE Integration

The numerical process used to simulate how model activity changes continuously through time according to the governing equations of the system.

Parameter Fitting

The process of repeatedly running a model with different parameter combinations to identify those producing the closest agreement with observed data.

Persistence

The tendency for activity or detectability to remain elevated after seasonal conditions begin to change.

Random Search

A parameter fitting approach in which model parameters are repeatedly selected randomly from within a defined search space and evaluated against observed data.

Search Space

The range of allowable parameter values explored during the fitting process.

Seasonal Forcing

A smooth annual driver representing seasonal environmental change within the model.

Seasonal Forcing Function

The mathematical function used to represent cyclical seasonal environmental influence within the model.

Seasonal Suppression

A reduction in modelled activity outside the active seasonal period.

Seasonal Window

The part of the year during which presence or activity is biologically possible within the model.

Similarity Matrix

A pairwise comparison of species based on their fitted seasonal ecological characteristics.

Synthesised Curve

A reconstructed seasonal curve generated from the fitted model and scaled to match the observed data range for visual comparison.

Weighted Ecological Distance

A similarity measure in which some seasonal ecological characteristics contribute more strongly than others when comparing species.

Winter Visitor

A species whose primary period of activity or presence occurs during the winter months, typically spanning the year boundary.