drmr: A Bayesian approach to Dynamic Range Models in R

Lucas da Cunha Godoy[12] Alexa Fredston[3] R. M. W. J. Bandara[4]
Mark Morales[1] Malin L. Pinsky[1]

April 13, 2026

[1]Department of Ecology and Evolutionary Biology, University of California Santa Cruz
[2]email: lucas@example.com
[3]Department of Ocean Sciences, University of California Santa Cruz
[4]Department of Ecology and Evolution, Rutgers University

Abstract

Predicting how species distributions will respond to environmental change is a critical challenge. Dynamic Range Models (DRMs) offer a powerful mechanistic approach by explicitly modeling the influence of environmental drivers on demographic processes. However, the widespread adoption of DRMs has been hindered by their inherent complexity and a critical gap in the available software. While many tools can simulate range dynamics, none provide a user-friendly framework to statistically estimate the functional relationships between environmental conditions and demographic rates from spatio-temporal data. To make this approach more accessible, we introduce `drmr`, an open-source R package for building, fitting, evaluating, and forecasting age-structured DRMs within a user-friendly Bayesian framework. From spatio-temporal species observations, the `drmr` package allows users to explicitly estimate how environmental drivers affect demographic processes such as recruitment and survival. We demonstrate the package's utility through case studies on summer flounder (*Paralichthys dentatus*) and red-bellied woodpecker (*Melanerpes carolinus*). In both cases, we fit DRMs with environment-dependent recruitment and survival. Additionally, we show how the package visualizes estimated relationships between environmental conditions and demographic rates. The `drmr` package bridges a critical gap between complex, process-explicit models that are parameterized *a priori*, and statistically accessible correlative models that are easily fit to data. By lowering the barrier to entry, it provides a powerful and accessible tool for ecologists to test mechanistic hypotheses and generate more robust ecological forecasts in the face of global change.

Keywords: Climate change, Spatial prediction, Spatio-temporal modeling, Species distribution

1 Introduction

Predicting how species will respond to accelerating global environmental change remains one of the most critical challenges facing ecologists today (Urban et al., 2016). Knowledge of species' population dynamics is essential, as the demographic processes of reproduction, growth, mortality, and movement fundamentally govern species' distributions (Pagel and Schurr, 2012). Dynamic Range Models (DRMs) have emerged as promising tools for this purpose, explicitly linking demographic rates to environmental drivers within a hierarchical Bayesian framework (Pagel and Schurr, 2012). This approach allows for more robust predictions under novel environmental scenarios where traditional, correlative Species Distribution Models (SDMs) often fail (Fredston et al., 2025). Despite their conceptual appeal, the widespread application of DRMs has been limited by their inherent complexity and the computational challenges of fitting them to data (Evans and Moustakas, 2016; Uribe-Rivera et al., 2023).

The importance of biological mechanism in species distribution modeling inspired the development of several `R` packages. However, existing tools are designed for simulating species dynamics under pre-defined assumptions rather than statistically estimating the key relationships from data. For instance, `rangr` (Markowska et al., 2025), and `metaRange` (Fallert et al., 2025) are powerful virtual species platforms (VSP) but do not fit models to observational data. A few packages, such as `demoniche` (Nenzén et al., 2012), `MigClim` (Engler and Guisan, 2009), and `RangeShiftR` (Malchow et al., 2021; Bocedi et al., 2021) take the output of correlative SDMs to inform deterministic population dynamics or movement/dispersal algorithms. The `poems` package can fit spatially explicity metapopulation models to data through approximate Bayesian computation (ABC) (Fordham et al., 2021; Malchow et al., 2024). However, `poems` is exceptionally flexible and requires a steep learning curve for many users. It also relies on ABC, which is computationally expensive and more sensitive to simulation settings and prior constraints than the Markov Chain Monte Carlo (MCMC)-based inference used in `drmr`. To our knowledge, no existing package provides a flexible framework

Table 1: Comparison of `R` packages for spatial species dynamics and climate impacts. "SDM Packages" represents a wide range of standard options, such as `biomod`, `tinyVAST`, and `dismo`. The table evaluates these packages based on whether they explicitly represent demographic processes (Mechanistic), allow fitting spatially explicit models to data rather than just simulating them (Fits to spatial data), and can quantify uncertainty around parameter estimates and forecasts (Uncertainty). It also notes which packages accommodate environmentally-dependent demographic processes (Env-dem) and distinguishes those that use Markov Chain Monte Carlo rather than Approximate Bayesian Computation (ABC) for Bayesian inference (MCMC). Table symbols indicate if a feature is present (✓), absent (✗), or not applicable (·) to a given package.

Package	Mechanistic	Fits to spatial data	Uncertainty	Env-dem	MCMC
SDM Packages	✗	✓	✓	✗	✓
`NicheMapR`	✓	✗	✗	✗	·
`RangeShiftR`	✓	✓	✗	✗	·
`rangr`	✓	✗	✗	✗	·
`steps`	✓	✗	✗	✗	·
`metaRange`	✓	✗	✗	✗	·
`SpatialDemography`	✓	✗	✗	✗	·
`MigClim`	✓	✓	✗	✗	·
`demoniche`	✓	✓	✗	✗	·
`poems`	✓	✓	✓	✓	✗
`drmr`	✓	✓	✓	✓	✓

to directly estimate the functional relationships between environmental drivers and demographic rates within a single, integrated model. This step is essential for understanding why ranges are shifting and for making credible forecasts. To illustrate the gap in functionality across packages, from only simulation to environmentally-dependent population dynamic models, we compare several packages in Table 1, including a broad category for the many packages available for fitting more traditional SDMs (e.g., Hijmans et al., 2024; Thorson et al., 2025).

To bridge this application gap and make DRMs more accessible, we introduce `drmr`: an open-source `R` (R Core Team, 2024) package for developing and applying DRMs using a Bayesian framework. `drmr` is designed to substantially lower the barrier to applying this class of models to real problems by providing a user-friendly interface for constructing, fitting, and projecting age-structured DRMs. It leverages the power and efficiency of `Stan`

via the `cmdstanr` (Gabry et al., 2024) interface. The package offers flexibility, enabling users to easily relate environmental drivers to specific demographic rates like recruitment and survival. In this paper, we introduce the `drmr` framework and its capabilities. We do not attempt an exhaustive comparison against the myriad of modeling approaches already present in the literature.

2 Methods

We adopt the convention of representing random variables with capital letters and their observed values with lowercase letters. Formally, let $Y_{t,i}$ be a random variable representing the density of a focal species, with $y_{t,i}$ representing its realization, that is, what we actually observe. The sub-indices t and i stand for time and patch (or site), respectively. We define a zero-augmented probability density function (pdf) as follows:

$$f(y_{t,i} \mid \mu_{t,i}, \phi, \rho_{t,i}) = \begin{cases} \rho_{t,i}, & \text{if } y_{t,i} = 0, \\ (1 - \rho_{t,i}) g(y_{t,i} \mid \mu_{t,i}, \phi), & \text{if } y_{t,i} > 0, \end{cases} \tag{1}$$

where $\rho_{t,i} = \Pr(Y_{t,i} = 0)$ is the probability of observing a zero density (which might in turn depend on observation effort or $\mu_{t,i}$); and $g(\cdot \mid \mu_{t,i}, \phi)$ is the pdf of a continuous probability distribution with expected value, or theoretical mean, $\mathbb{E}[Y_{t,i} \mid Y_{t,i} > 0] = \mu_{t,i}$ and additional parameter ϕ. This distribution represents the densities of the species when it is present. Popular choices such as the Log-Normal, Log-Logistic, and Gamma distributions are readily available in the package. We parametrize these three families of distributions in terms of their theoretical mean value to facilitate modeling (Ye et al., 2021). A side effect of such a parametrization is that the additional parameter, denoted ϕ, is not always readily interpretable. While this is a limitation, we prioritized being able to model the mean of a distribution over to easily interpreting its nuisance parameters.

2.1 Dynamic range models

While both DRMs and SDMs can relate environmental variables to zero-augmented densities (Equation (1)), DRMs use demographic processes to constrain the dynamics of density across space and through time. Specifically, a DRM separates this observation model (encoded through the zero-augmented distribution) from an explicit process model. Underlying biological mechanisms that drive species distributions are what inform the latter. This separation enables the powerful integration of knowledge regarding species demography, dispersal, and mechanistic interactions with environmental factors (Pagel and Schurr, 2012).

The assumptions of our DRM are placed on the conditional mean density $\mu_{t,i}$, with the probability of absence remaining the same as in a typical zero-augmented SDM. In fact, those assumptions can be interpreted as a biologically informed non-linear function of the environmental variables. Specifically, we put biological constraints on the unobserved age-structure behind the densities at a given time point and patch. Define $Y_{a,t,i}$ as a random variable representing the unobserved density for individuals of age a, where $a \in \{1, \ldots, A\}$. We assume that if the species is absent (i.e., $Y_{t,i} = 0$), then $Y_{a,t,i} = 0$ for all a. We denote the respective age-specific expected densities as $\lambda_{a,t,i} = \mathbb{E}[Y_{a,t,i} \mid Y_{t,i} > 0]$. Although we do not observe $Y_{a,t,i}$, we assume $Y_{t,i} = \sum_a Y_{a,t,i}$. Consequently, by the linearity property of expected values, we have

$$\mu_{t,i} = \sum_a \lambda_{a,t,i}. \tag{2}$$

Biological processes and additional assumptions are encoded through the expected age-specific densities $\lambda_{a,t,i}$. In addition, we can allow for explicit relationships between the mean density $\mu_{t,i}$ and $\rho_{t,i}$, as higher densities should decrease the probability of observing zeros (Lambert, 1992; Yee, 2014).

Our implementation supports explicitly accounting for three key processes: recruitment, survival, and movement. Each process is thoroughly described in the upcoming Sections. The most basic model (when the user does not specify anything about the processes) does

not allow movement between spatial habitat patches and assumes recruitment and survival rates are constant across time and space. In addition, in this basic version, only the recruitment rate is estimated, while the survival rate is a constant, user-defined value. All model configurations require initialized age-specific densities, which can be derived from the total density, assumed mortality rates, or estimated as free parameters via the `est_init` toggle; for more details, see `vignette("init", "drmr")` in the `R` package.

2.1.1 Recruitment

We begin by defining the expected density for recruits at time t and patch i as:

$$\mathbb{E}[Y_{1,t,i}] = \lambda_{1,t,i} = \exp\{\psi\}, \tag{3}$$

where $\exp\{\psi\}$ represents the overall mean recruitment per unit of area. The assumption in Equation (3) is quite restrictive. We can obtain a more realistic model by replacing ψ with

$$\psi_{t,i} = \boldsymbol{\beta}_r^\top \mathbf{x}_{t,i}^{(r)}, \tag{4}$$

where $\mathbf{x}_{t,i}^{(r)}$ is a vector of environmental drivers associated to recruitment at time t and patch i, $\boldsymbol{\beta}_r$ is a vector of corresponding regression coefficients. This depicts a log-linear relationship between recruitment and the environment. Gaussian relationships between recruitment and environmental conditions can be easily included by adding squared terms in $\mathbf{x}_{t,i}^{(r)}$.

2.1.2 Survival

In the proposed model, the expected density at age a, time t, and patch i evolves over time according to survival rates denoted $s_{a,t,i}$. Formally, we have

$$\mathbb{E}[Y_{a,t,i}] = \lambda_{a,t,i} = \lambda_{a-1,t-1,i} \cdot s_{a-1,t-1,i} \text{ for } a > 1. \tag{5}$$

As with recruitment, the survival rates may vary by patch i and time t:

$$\text{logit}(s_{a,t,i}) = \boldsymbol{\beta}_s^\top \mathbf{x}_{t,i}^{(s)}, \tag{6}$$

where $\mathbf{x}_{t,i}^{(s)}$ represents environmental drivers, which can be different from those affecting recruitment, and $\boldsymbol{\beta}_s$ are associated regression coefficients. Note that, despite the age-specific index a, the survival rates are constant across age-groups and, moreover, the logit link function ensures the survival rates within the $[0, 1]$ interval.

External information can be incorporated into Equation (6) to model age-specific survival rates. For example, data on other mortality sources, such as hunting, may be available from other studies or wildlife management agencies (Urban et al., 2016). Another example occurs in fisheries, which often have age-specific fishing rates (Hilborn and Walters, 2013). Assuming $f_{a,t}$ represents an instantaneous rate of mortality for the age-group a at time t, Equation (6) can be updated as follows:

$$s_{a,t,i} = \text{logit}^{-1}\{\boldsymbol{\beta}_s^\top \mathbf{x}_{t,i}^{(s)}\} \exp\{-f_{a,t}\}, \tag{7}$$

where logit^{-1} denotes the inverse of the logit function. Importantly, when provided, the relationship between the mortality rates $f_{a,t}$ and survival is assumed to be deterministic instead of estimated.

The literature often shows that survival rates vary with age (Healy et al., 2019). For mammals, age-specific mortality is available from published life-tables (Heppell et al., 2000). Senescence, an inevitable and irreversible accumulation of damage with age in ecology, leads to functional loss and thus reduced survival rates (Gaillard and Lemaître, 2020; Monaghan et al., 2008). If age-specific survival rates are available from the literature, they can be easily included in Equation (7) by using the negative log of these rates as $f_{a,t}$. When rates are constant across time, users can simply repeat them for all time points.

2.1.3 Movement

Movement can substantially impact species densities and encodes a basic law of geography that similarity increases with proximity. Our package models this using a basic diffusion routine, which represents undirected movement away from an organism's current location (Thorson et al., 2021). Denote P the total number of patches in the present study. For each age-group, we define a $P \times P$ movement matrix denoted $\mathbf{M}_a$. The element $\{i, j\}$ in $\mathbf{M}_a$ represents the probability that individuals of age a move from patch i at time $t-1$ to patch j at time t. Here, the movement matrix $\mathbf{M}_a$ is derived from adjacency matrix $\mathbf{A}$, which encodes connections between patches. In $\mathbf{A}$, the element $\{i, j\}$ is given by

$$\mathbf{A}_{ij} = \begin{cases} 1/N(i), & \text{if } i \sim j, \\ 0, & \text{otherwise,} \end{cases} \tag{8}$$

where $N(i)$ is the number of neighbors of patch i, and $i \sim j$ indicates that i and j are neighbors (i.e., share borders). In this way, spatial structure can easily be defined as one-dimensional, two-dimensional, or as more complex network structures.

The user must also indicate which age-groups can move. If an age group is not assumed to move, the movement matrix becomes the identity matrix. Given those inputs, we define ζ as the probability of individuals remaining in the same patch between two time points. The movement matrix is then constructed as:

$$\mathbf{M}_a = \zeta \mathbf{I}_P + (1 - \zeta)\mathbf{A}, \tag{9}$$

where $\mathbf{I}_P$ is a $P \times P$ identity matrix. In other words, individuals that move have equal probability of moving to any of the neighboring patches.

Once a movement matrix is defined, we adjust the age-specific expected densities accordingly. Let $\boldsymbol{\lambda}_{a,t,\cdot} = [\lambda_{a,t,1}, \ldots, \lambda_{a,t,P}]^\top$ be a vector of length P (number of patches) representing

the expected densities of age a at time t for every patch after survival and before movement. After movement, the expected density is:

$$\tilde{\boldsymbol{\lambda}}_{a,t,\cdot} = \mathbf{M}_a \boldsymbol{\lambda}_{a-1,t-1,\cdot}. \tag{10}$$

Next, $\lambda_{a-1,t-1,i}$ in Equation (5) is replaced by $\tilde{\lambda}_{a-1,t-1,i}$, where the latter represents the i-th element of $\tilde{\boldsymbol{\lambda}}_{a,t,\cdot}$.

2.1.4 Patch, spatial, and temporal dependence

The models presented thus far assume that, after accounting for measured environmental covariates and dispersal, the densities in different patches and at different times are independent. This assumption is often violated in ecological data, as unmeasured factors can create correlations across space and time. To address this, `drmr` allows for the inclusion of random effects to model this residual dependence. These effects can be incorporated into the sub-models for recruitment (Equation (4)), survival (Equation (6)), or the mean density directly (Equation (2)).

The package supports three distinct random effect structures. The simplest structure uses independent and identically distributed (IID) random effects to introduce a patch-specific intercept for a given process. This accounts for time-invariant local factors not captured by the environmental drivers. These effects are assumed to be independent across all patches, with each patch's effect being drawn from a common distribution, effectively shrinking estimates toward a global mean.

When nearby patches are expected to be more similar than distant ones due to unmeasured, spatially structured variables (e.g., Paradinas et al., 2023), a spatial random effect can be used. This structure assumes that the random effect for a given patch is correlated with the effects in its neighboring patches. `drmr` implements this using an intrinsic Conditional Autoregressive (Besag et al., 1991; Morris et al., 2019, ICAR) model, which shrinks

patch-specific effects towards a local, rather than global, mean.

To account for serial autocorrelation in population dynamics that is common to all patches and not explained by the covariates, a temporal random effect can be added. The package implements this as a first-order autoregressive (AR(1)) process, which models the dependency between observations at time t and time $t-1$. By incorporating these random effects, users can build more realistic models that better capture the complex sources of variation inherent in spatio-temporal ecological data.

2.2 Bayesian inference

To allow Bayesian inference, we specify Normal priors for the regression coefficients (e.g., $\boldsymbol{\beta}_r$ and $\boldsymbol{\beta}_s$), a Gamma prior for the nuisance parameter (ϕ), and a Beta prior for the movement parameter (ζ). When random effects are included, we place Normal priors on the log-transformed standard deviation of those effects. For the autocorrelation parameter, denoted α, in the AR(1) structure, we use a Beta prior because negative temporal autocorrelation is rare in practice. Further details on prior specifications and default hyperparameters are available in a package vignette (`vignette("ppt", "drmr")`). We use the No-U-Turn Sampler (NUTS) MCMC algorithm, implemented in `Stan` (Homan and Gelman, 2014), to sample from the joint posterior distribution. MCMC initialization is performed using either `Stan`'s default, drawing from priors, or the Pathfinder algorithm (Zhang et al., 2022). Forecasts incorporating parameter uncertainty are generated by sampling from the posterior predictive distribution. We report posterior medians as point estimates and symmetric 90% credible intervals (posterior percentiles) for uncertainty quantification unless otherwise stated.

3 The `drmr` `R` package

Given the model structure explained so far, the `drmr` package offers functions to facilitate data analyses based on DRMs. The `fit_drm` has an analogous counterpart for GLM-based

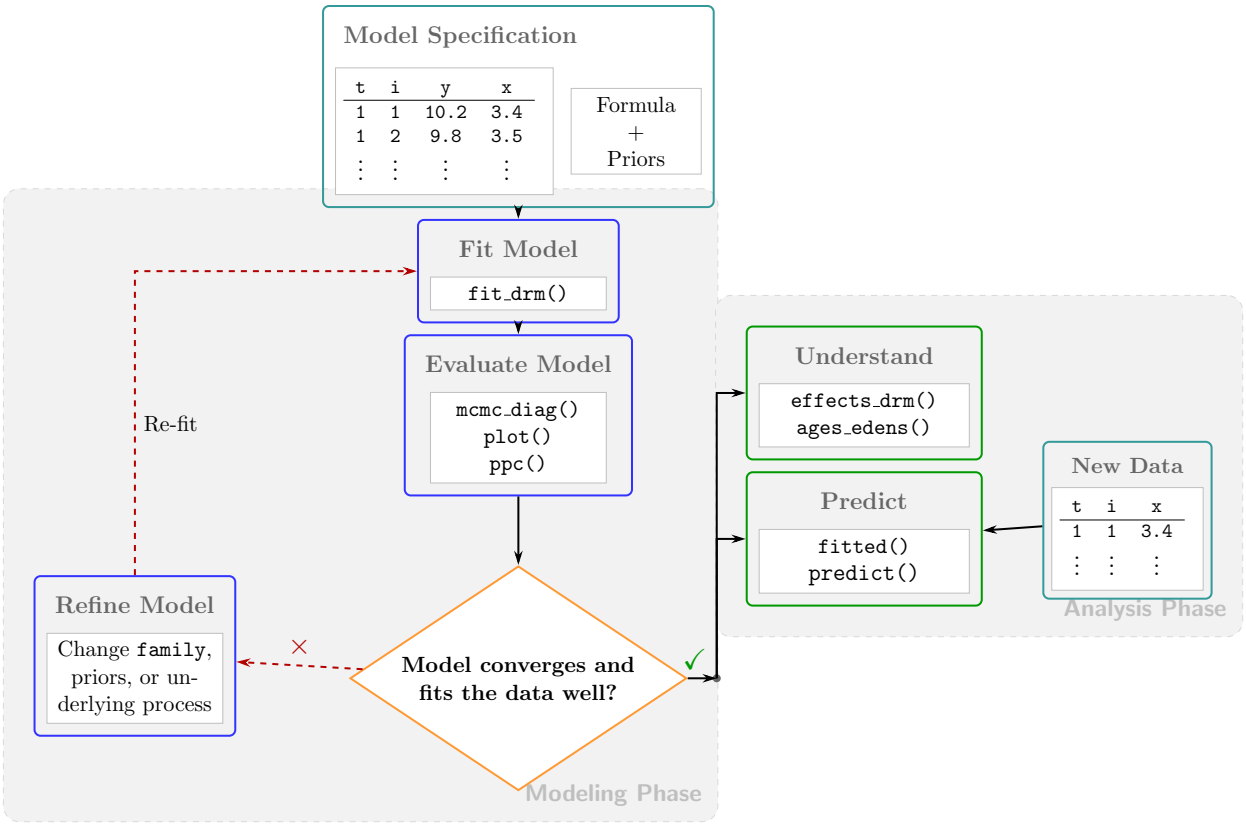

Figure 1: `drmr` modeling workflow. The diagram outlines the iterative analytical process, highlighting key `R` functions and expected input data formats at each stage. In the example datasets presented along the diagram, t represents the time column, i is the patch column, y is the response variable (e.g., density), and x stands for an environmental driver.

SDMs (`fit_sdm`), enabling consistent fitting and fair comparison across model types using the same data. These core functions are `Stan` wrappers that prepare data and run pre-compiled models bundled with the package. Their outputs have their own class (`drmrmodels`) for which we provide `print`, `summary`, `update` and `draws` methods. The latter is used to easily retrieve samples from the parameters' posterior distributions. A suggested workflow is presented in Fig. 1, along with some of the most important functions available in the package.

The following example demonstrates a minimal reproducible workflow using `drmr`'s default priors and assuming no movement between patches. The `summary` method outputs the parameters along with the posterior mean, standard deviation, and percentiles (5%, 50% and 95%, by default) based on the samples from posterior distributions.

```r
## Load the package and example data
library(drmr)
data(sum_fl)
## Creating number of individuals per unit of area to be the response variable
sum_fl <- transform(sum_fl, dens = y / area_km2)
fit <- fit_drm(.data = sum_fl,     ## dataset
               y_col = "dens",     ## column representing response variable
               time_col = "year",  ## column representing time variable
               site_col = "patch", ## column representing patch/site
               seed = 2026,        ## seed for reproducibility
               ## recruitment has an intercept and a coefficient for
               ## surface temperature (stemp)
               formula_rec = ~ 1 + stemp)
summary(fit)
#> Summary of ADRM Fit
#>
#> Estimates:
#> # A tibble: 5 × 7
#>   variable  description                                mean     sd      q5    q50    q95
#>   <chr>     <chr>                                     <dbl>  <dbl>   <dbl>  <dbl>  <dbl>
#> 1 beta_t[1] zero-infl: (Intercept)                   -4.32  1.13   -5.94   -4.48  -2.07
#> 2 beta_r[1] rec: (Intercept)                         -9.03  0.617  -9.86   -9.12  -7.88
#> 3 beta_r[2] rec: stemp                                0.112 0.0301  0.0562  0.116  0.152
#> 4 xi[1]     Relationship between zeros and density -0.578 0.174  -0.821  -0.603 -0.230
#> 5 phi[1]    Dispersion parameter                      0.820 0.0668  0.715   0.817  0.932
```

To use `drmr`, the input data must follow a "long" format where each row corresponds to a specific patch and time point. The dataset must be spatio-temporally complete, containing all patches for every time step across the study period. While missing values are permitted in the response variable and are automatically estimated by the model, information on environmental drivers must be provided for all patch-time combinations. This ensures that the demographic processes underlying the DRM can be modeled continuously throughout

the study duration, even in the absence of species observations.

Running the models provided by the package relies on a small number of critical dependencies. The primary dependency is `CmdStan` (version $\geq$ 2.36). As the computational backbone, `CmdStan` is used via the `instantiate` interface (Landau, 2024) to pre-compile the package's `Stan` models during installation. Please note that `CmdStan` versions prior to 2.36 are unsupported and may lead to unexpected behavior. Another key dependency is the `sf` package (Pebesma, 2018), primarily used to generate the adjacency matrix required for the DRM's movement component. Finally, for availability, the package is hosted on GitHub.

To accommodate varying computational requirements and research stages, `drmr` supports several inference algorithms beyond the default NUTS sampler. Users can utilize the Pathfinder algorithm (Zhang et al., 2022), Laplace approximation, or Automatic Differentiation Variational Inference (Kucukelbir et al., 2017, ADVI) to obtain faster, approximate posterior estimates. These alternatives are implemented in `Stan` and are particularly effective for rapid model prototyping, initial hypothesis screening, or providing robust starting values for a final MCMC run. While these methods offer significant speed advantages, we recommend the default MCMC-based inference for final results to ensure rigorous uncertainty quantification. A comparison of these algorithms and their performance trade-offs is available in the package documentation (`vignette("algos", "drmr")`).

Different assumptions about the DRM can be conveniently modified through a series of "toggles". The toggles allow users to turn some of the features described in Section 2.1 on or off. This is achieved by passing a named list as an argument called `.toggles` to the `fit_drm` function. To turn a feature on, we make the respective list element equal to 1, while to turn it off, we make it equal to 0. A complete list of those toggles and their respective default values is displayed in Table 2. Alternatively, one can consult the default toggles by running `drmr::default_toggles()`. Note that, if nothing is provided, the default values are used. If we provide a list with a single named element, say `list(movement = 1)`, the `movement` feature (e.g., Section 2.1.3) will be turned on, while the other features will be kept at their

Table 2: Toggles associated with specific model features, their descriptions, and default states. For the last three toggles, the options are `"none"` (no random effect), `"rec"` (random effect for recruitment), `"surv"` (random effect for survival), and `"dens"` (random effect for density).

Toggle	Description	Default State
`rho_mu`	Explicitly relates density (μ) and the probability of observing a zero (ρ).	On
`cloglog`	Use the complementary log-log (cloglog) link function for ρ. If off, the logistic link function is used.	Off
`movement`	Enable the movement routine as described in the paper.	Off
`est_surv`	Estimate survival probabilities within the model.	On
`est_init`	Estimate initial values.	Off
`minit`	Initial values driven by mortality rates.	Off
`iid_re`	Include an IID Random effect.	`"none"`
`ar_re`	Include an AR(1) random effect.	`"none"`
`sp_re`	Include an ICAR random effect.	`"none"`

default values.

Complementary tools for diagnosing and understanding the results are also available. The functions `mcmc_diag` provides statistics such as the split-$\hat{R}$, the bulk and tail effective sample size (Vehtari et al., 2021, ESS). Although these statistics alone cannot establish convergence, some rules of thumb are common. For instance, $\hat{R} < 1.01$, bulk-ESS > 400 and tail-ESS > 400 provide reasonable evidence of convergence (Vehtari et al., 2021). In addition, the trace-plots available from calling `plot` on the output of our `fit_*` functions can aid the process of assessing the convergence. Lastly, the `ppc` function provides a posterior predictive check (Gabry et al., 2019, PPC), where the empirical density of the data can be compared to that of the induced posterior predictive distributions. In other words, it shows how well the model fits to the data.

After establishing model convergence, a suite of post-estimation functions allows for interpreting the relationships between environmental drivers and demographic processes. Specifically, the `effects_drm` function extracts the estimated impact of a chosen covariate on

a target process (e.g., recruitment or survival, specified as `"rec"` or `"surv"`). To visualize the estimated population structure, `ages_edens` extracts age-specific latent densities, whereas `fitted` and `predict` return in-sample and out-of-sample predictions, respectively. In the latter two, a `type` parameter allows for chosing between the following types of predictions: (i) the posterior predictive distribution of $Y_{t,i}$ (accounting for data variability and detectability); (ii) its mean, $(1 - \rho_{t,i})\mu_{t,i}$ (accounting for detectability but removing observation error); or (iii) the latent densities, $\mu_{t,i}$ (excluding both variation sources to reflect only the assumed population dynamics). Finally, `log_lik` integrates with the `loo` package to compute the Leave-One-Out Information Criterion (Vehtari et al., 2017, LOOIC), which facilitates model comparison.

4 Case studies

The features of the package are best illustrated through the analysis of real datasets. To this end, we showcase its functionalities using data from two distinct species: the summer flounder (*Paralichthys dentatus*) and the red-bellied woodpecker (*Melanerpes carolinus*). Despite the ecological differences between these species, our interface easily accommodates models with varying assumptions. For both case studies, we fit two distinct DRMs: one with environment-dependent recruitment and another with environment-dependent survival. For the summer flounder, we extend this by fitting a third model that incorporates both environment-dependent recruitment and survival, while also accounting for known fishing mortality. Across all case studies, the DRMs assume a Gamma probability density function for population densities, allow for movement, and include a temporal autoregressive term on recruitment. Finally, we emphasize that these case studies are designed strictly to demonstrate the versatility, workflow, and application of the package. They are not intended to serve as comprehensive ecological evaluations or definitive performance comparisons against existing models.

In each case, we compare the models based on out-of-sample predictions. Specifically, we reserved the last five years of data and assessed point prediction accuracy with root mean square error (RMSE). To assess interval predictions, we used the interval score (Gneiting and Raftery, 2007, IS). The IS penalizes intervals based on their width and whether they contain the true observation. Lower RMSE and IS values indicate better model performance.

4.1 Summer flounder

DRMs have been previously developed and evaluated for summer flounder (Fredston et al., 2025). Here, we demonstrate how to use the package using a somewhat different model structure from Fredston et al. (2025) but the same dataset: scientific bottom trawl surveys in the United States of America (USA) from 1972-2016 (Forrest et al., 2020). The original sampling unit for the survey consists of a "haul", an event during which a fishing net is towed along the seafloor. The haul data was aggregated across 10 patches along the northeast coast, each with 1° latitude in height but varying width. The response variable was the number of fish caught per unit of area. For convenience and reproducibility, this processed dataset is included with the package. Users can view its documentation by running `help("sum_fl", package = "drmr")`.

In our first model, we set the logit of probability of absence as a linear function of the number of hauls (`c_hauls`) and the underlying theoretical density ($\mu_{t,i}$). The latter is controlled by the `rho_mu` toggle. In addition, we establish a quadratic relationship between log recruitment and seawater surface temperature (SST; `stemp`), which approximates a bell-shaped curve once exponentiated. Those functional relationships are easily modified through a formula syntax. We additionally assume an AR(1) effect for recruitment. The syntax to fit this model is as follows:

```
## load package and data
library(drmr)
data(sum_fl)
```

```r
## we omit the code to split the data into train and test here.
## The full script is available on GitHub
drm_rec <-
  fit_drm(.data = dat_train,  ## a data.frame
          y_col = "dens",     ## column for the response
          time_col = "year",  ## column for time
          site_col = "patch", ## column for site
          family = "gamma",   ## non-zero densities' pdf family
          seed = 202505,      ## seed for reproducibility
          ## formula for the probability of absence
          formula_zero = ~ 1 + c_hauls,
          ## recruitment as a function of
          ##   sea surface temperature (stemp)
          formula_rec = ~ 1 + c_stemp + I(c_stemp * c_stemp),
          ## formula for survival (intercept only)
          formula_surv = ~ 1,
          ## f_{a, t} from Equation (6)
          ## In this case, we input it as a matrix with "number of ages" rows
          ## and "number of time points" columns.
          f_mort = f_train,
          n_ages = NROW(f_train),
          ## A matrix from Equations (7) and (8)
          adj_mat = adj_mat,
          ages_movement = c(0, 0,       ## ages not allowed to move
                            rep(1, 12), ## ages allowed to move
                            0, 0),      ## ages not allowed to move
          .toggles = list(rho_mu = 1, ## relates density to probability
                                      ## of absence
                                      ## (0 is the default and turns it off)
                          ar_re = "rec", ## AR(1) term for recruitment
                                         ## ("dens" for density,
                                         ##  "surv" for survival,
                                         ##  "none" is the default includes no AR(1))
```

```
                        movement = 1, ## uses movement routine from Equation (8)
                                      ## (0 is the default)
                        est_surv = 1, ## Estimate the survival rates
                                      ## (0 is the default)
                        minit = 1),   ## Uses mortality to estimate the initial
                                      ## age-classes in the first years of study
                                      ## (0 is the default)
    ## some priors hyperparameters.
    .priors = list(pr_alpha_a = 4.2, pr_alpha_b = 5.8,
                   pr_zeta_a = 7, pr_zeta_b = 3),
    ## algorithm for inference. "nuts" is the default (MCMC)
    ## for other options see ?fit_drm
    algorithm = "nuts",
    ## arguments to be passed to the inference algorithm.
    ## The default for nuts is to run 4 chains sequentially.
    ## here we are specifying to run 4 chains in parallel.
    ## The default arguments can be inspected by
    ##   running: default_nuts()
    algo_args = list(parallel_chains = 4))
```

The priors were set such that prior probability of remaining in the same site is 0.7, while the prior temporal autocorrelation is 0.42. The former is arbitrary but biologically meaningful, while the latter is based on estimates from the fisheries literature (Johnson et al., 2016, e.g.,).

Quick model diagnostics are obtained as follows:

```
mcmc_diag(drm_rec) |>
  print() |>
  summary()
#> # A tibble: 11 × 4
#>    variable     rhat ess_bulk ess_tail
#>    <chr>       <dbl>    <dbl>    <dbl>
#>  1 beta_t[1]    1.00    4253.    3200.
#>  2 beta_t[2]    1.00    5719.    2527.
#>  3 beta_r[1]    1.00    4246.    2852.
```

```
#>  4 beta_r[2]    1.00    5296.    3226.
#>  5 beta_r[3]    1.00    5543.    3196.
#>  6 xi[1]        1.00    4376.    2905.
#>  7 phi[1]       1.00    7562.    3293.
#>  8 zeta[1]      1.00    5280.    2793.
#>  9 beta_s[1,1]  1.00    4144.    3007.
#> 10 alpha[1]     1.00    5448.    3284.
#> 11 sigma_t[1]   1.00    4077.    3355.
#> No issues detected by the MCMC diagnostics.
```

The first column represents the model parameter, with the three remaining columns highlighting three different convergence diagnostics' metrics. The `summary` call quickly scan the rules of thumb to check if there is any evidence of convergence issues. By running `plot(drm_rec)` one can quickly inspect the trace-plots for each model parameter. Since there is no evidence of any convergence issues, we can inspect the parameter estimates using the `summary` function:

```
summary(drm_rec)
#> Summary of ADRM Fit
#>
#> Estimates:
#> # A tibble: 11 × 7
#>    variable    description                                  mean      sd      q5     q50     q95
#>    <chr>       <chr>                                       <dbl>   <dbl>   <dbl>   <dbl>   <dbl>
#>  1 beta_t[1]   zero-infl: (Intercept)                     -5.35   0.541   -6.27   -5.33   -4.50
#>  2 beta_t[2]   zero-infl: c_hauls                        -0.0549 0.0161  -0.0821 -0.0545 -0.0295
#>  3 beta_r[1]   rec: (Intercept)                           -2.57   0.181   -2.85   -2.58   -2.26
#>  4 beta_r[2]   rec: c_stemp                               0.270  0.0276   0.226   0.269   0.316
#>  5 beta_r[3]   rec: I(c_stemp * c_stemp)                 -0.0547 0.00680 -0.0663 -0.0545 -0.0438
#>  6 xi[1]       Relationship between zeros and density    -1.64   0.205   -1.99   -1.63   -1.33
#>  7 phi[1]      Dispersion parameter                       1.45   0.142    1.22    1.45    1.69
#>  8 zeta[1]     Prob. of remaining at the current patch    0.705  0.125    0.477   0.716   0.891
#>  9 beta_s[1,1] surv: (Intercept)                          1.48   0.564    0.604   1.44    2.46
```

```
#> 10 alpha[1]    Temporal autocorrelation              0.477  0.152   0.223   0.478   0.723
#> 11 sigma_t[1]  Temporal random effect std. dev.      0.171  0.0457  0.108   0.166   0.253
```

In this case, the standard output contains the parameters (`variables`), followed by their description, and some summary statistics from their posterior distribution.

An interesting competing hypothesis is that seawater bottom temperature (SBT; `btemp`) drives the species' survival rates. We can update the model to take that into account using the `update` method:

```
drm_surv <-
  update(drm_rec,
         formula_rec = ~ 1,
         ## survival as a function of
         ##   sea bottom temperature
         formula_surv = ~ 1 + c_btemp + I(c_btemp * c_btemp))
```

When using the `update` method, only the provided arguments are modified from the model fit being updated (`drm_rec`, in this case). We can use this method to fit a model under the hypothesis that both demographic processes are influenced by the environment:

```
drm_rs <-
  update(drm_rec,
         ## survival as a function of
         ##   sea bottom temperature
         formula_surv = ~ 1 + c_btemp + I(c_btemp * c_btemp))
```

Notice that, in this case, we retain the original relationship between recruitment and SST.

The `predict` method enables us to seamlessly obtain the out-of-sample predictions. This function always needs at least three arguments, the output of a `fit_drm` call, a dataset with the patches, time-points and environmental variables where we seek to obtain predictions, and a seed for random number generation. When we estimate survival, we need to pass the data from the last year used for model fitting to the function as well. The code below

illustrates how to use this function with the outputs of the DRMs fitted in this Section. Importantly, the argument `f_test` (which represents $f_{a,t}$ associated to `new_data`) is optional and should only be provided when external mortality information is available (see Equation (7)).

```
proj_rec <- predict(drm = drm_rec,
                    new_data = dat_test,
                    past_data = filter(dat_train,
                                       year == max(year)),
                    ## `f_test` is similar to `f_train`. The difference
                    ## is that the number of columns coincides with the
                    ## number of time points where we wish to make predictions.
                    f_test = f_test,
                    seed = 125,
                    cores = 4)


proj_surv <- predict(drm = drm_surv,
                     new_data = dat_test,
                     past_data = filter(dat_train,
                                        year == max(year)),
                     seed = 125,
                     f_test = f_test,
                     cores = 4)


proj_rs <- predict(drm = drm_rs,
                   new_data = dat_test,
                   past_data = filter(dat_train,
                                      year == max(year)),
                   seed = 125,
                   f_test = f_test,
                   cores = 4)
```

The density estimates from the DRM appear to capture more of the population growth in the northern patches, both in the fitted and the out-of-sample forecasts (Fig. 2). In Table 3, we present the predictive skill metrics for each model, sorted by the RMSE. The

Table 3: Predictive skill for the summer flounder models. Out-of-sample indicates whether the predictions are in- or out-of-sample. Model distinguish between the fitted model by highlighting which demographic processes (Rec, Surv) depend on the environment. RMSE is the root mean square error of prediction, and IS is the interval score.

Out-of-sample	Model	RMSE	IS (90%)
No	Rec + Surv	0.02	0.36
	Rec	0.02	0.41
	Surv	0.02	0.47
Yes	Rec + Surv	0.11	1.61
	Rec	0.12	1.66
	Surv	0.12	1.97

results suggest that the DRM (rec-surv) does better both for in-sample and out-of-sample predictions. In terms of interval predictions (assessed using the IS), all models do better in-sample than out-of-sample, as expected.

The value from DRMs extends beyond simply better predictive performance. These models also allow researchers to examine how environmental conditions affect demographic processes. For example, the panel in Fig. 3 depicts the estimated relationship between recruitment and SST and between survival and SST from the DRM that relates both recruitment and survival to environmental variables. Those curves, can be easily obtained by running:

```
drm_effects(drm_rs,                    ## output form fit_drm
            process = "rec",           ## the demographic process we wish to analyze
            variable = "c_stemp") |> ## an environmental variable
  plot()


drm_effects(drm_rs,
            process = "surv",
            variable = "c_btemp") |>
  plot()
```

In addition to the curves, the DRM (rec) model suggests that recruitment is maximized with a SST of 21.3°C (90% CI: 20.6 to 22.5). On the other hand, the DRM (surv) models suggests that the optimal SBT for survival is near 13.7°C (90% CI: 13.0 to 14.5).

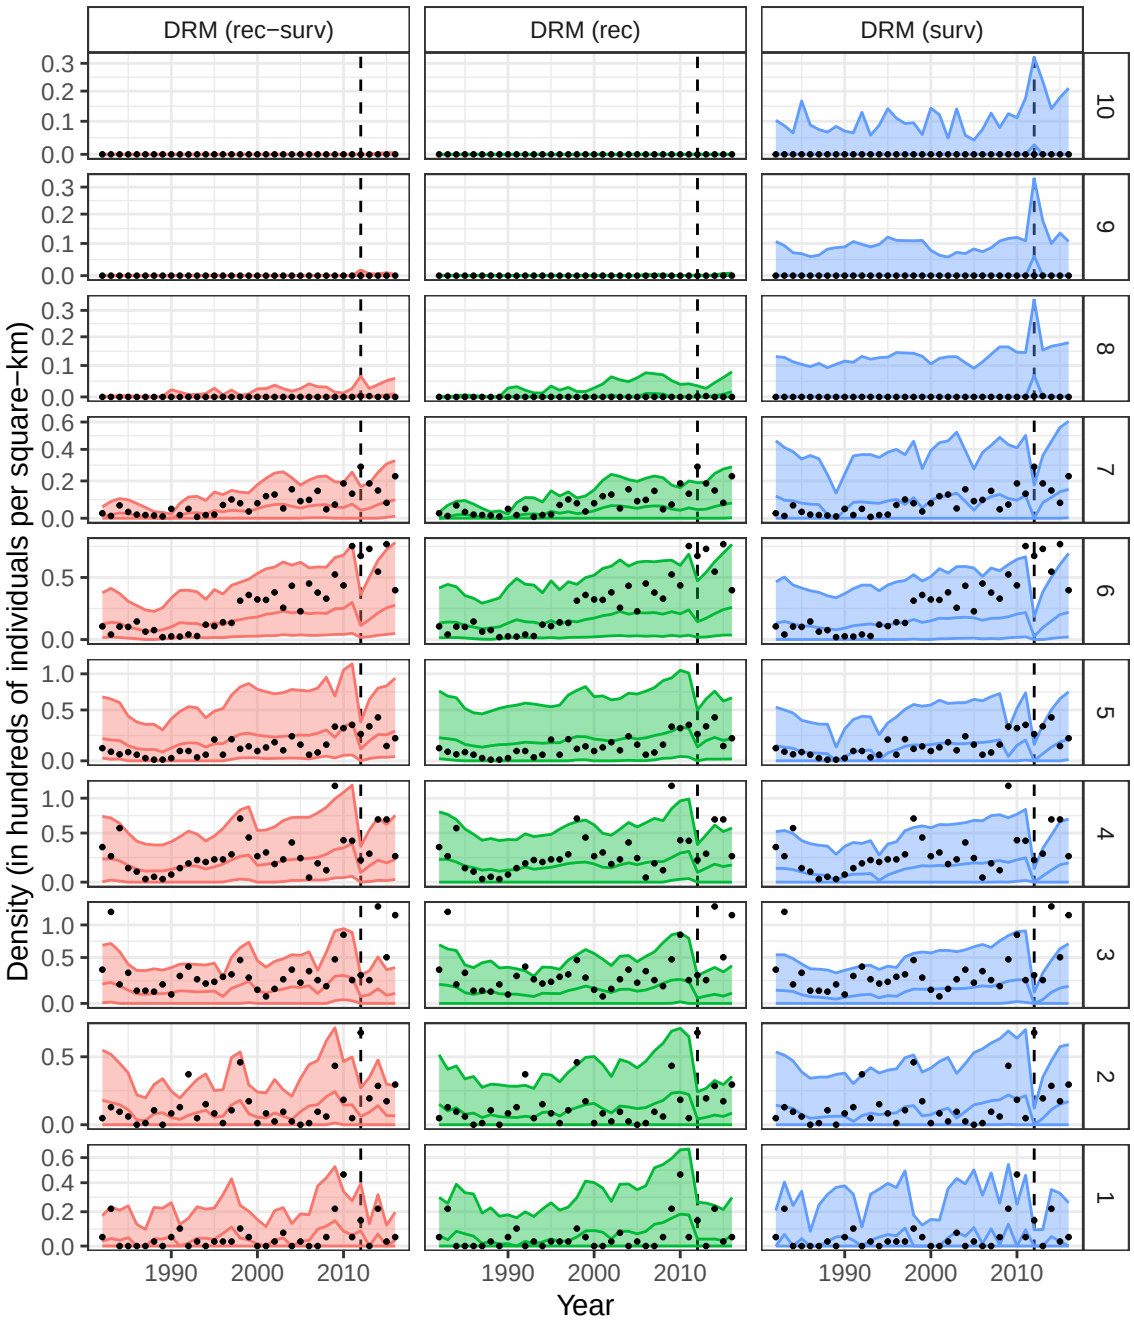

Figure 2: Predicted summer flounder densities per patch in the training years (left of the dashed line) and in the out-of-sample testing years (right of the dashed line). Colored lines show the posterior median, shaded area shows the 90% prediction interval (PI), and black dots show the observations. Patches are arranged from 10 (northernmost) to 1 (southernmost). The columns show results from a DRM with environmentally linked recruitment and survival (rec-surv), a DRM with environmentally linked recruitment (rec), and a DRM with environmentally linked survival (surv).

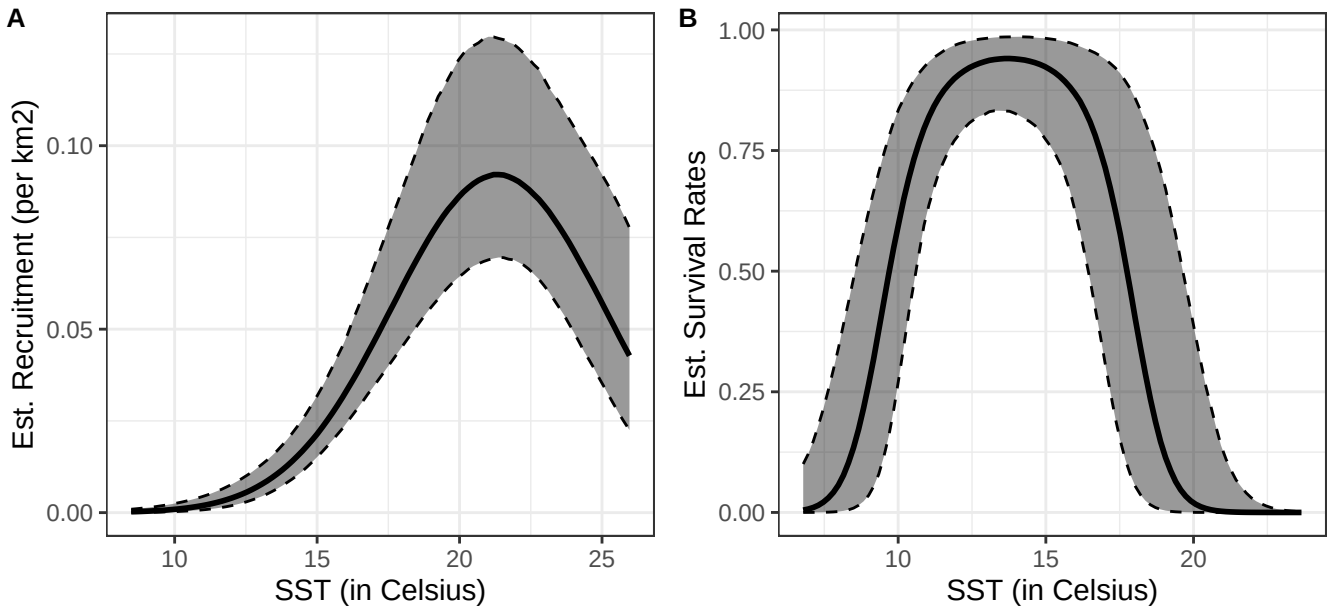

Figure 3: Panel containing the estimated relationship for summer flounder between recruitment and seawater surface temperature (SST) for the DRM (rec) model on the left hand side, and the relationship between survival rates and seawater bottom temperature (SBT) in the DRM (surv) model on the right hand side. The plots show the posterior median and 90% credible interval (CI).

4.2 Red-bellied woodpecker

To demonstrate the package's flexibility, we apply the same models to the red-bellied woodpecker, which is a resident bird species distributed throughout the eastern United States (Shackelford et al., 2000). We obtained abundance data from the North American Breeding Bird Survey (Ziolkowski Jr et al., 2024, BBS) via the `bbsBayes2` package (Edwards et al., 2024). This survey's data are collected along thousands of randomly established roadside routes throughout the continent. For the analysis, we used data from 1980 to 2019 aggregated into 273 $\times$ 273 km grids. The grids were then clipped to the continental USA, resulting in variable cell areas, particularly along the coast. Air temperature data from CHELSA (Karger et al., 2017; Brun et al., 2022) were aggregated to the same spatial resolution.

The code used to fit the models and generate predictions is essentially identical to the previous case study. Thus, it is omitted here but fully available on GitHub. The environment-

dependent recruitment and survival versions of the model were set as a quadratic function of air temperature. For both models, the probability of absence was set as a function of the number of routes and theoretical density ($\mu_{t,i}$). Taking advantage of the larger dataset in this example, we also incorporated an ICAR random effect into the linear predictor of survival (Equation (6)). Movement was allowed among adjacent grid cells.

Table 4: Predictive skill for the red bellied woodpecker models. Out-of-sample indicates whether the predictions are in- or out-of-sample. Model distinguish between the fitted model by highlighting which demographic processes (Rec, Surv) depend on the environment. RMSE is the root mean square error of prediction, and IS is the interval score.

Out-of-sample	Model	RMSE	IS (90%)
No	Rec	0.03	0.62
	Surv	0.03	0.59
Yes	Rec	0.08	0.79
	Surv	0.11	1.26

The model results (Table 4) suggest the environment-dependent recruitment model delivers a better out-of-sampel predictive performance. The model considering environment-dependent survival performed well regarding in-sample predictions, but its the environment-dependent recruitment counterpart provides better out-of-sample predictions (both in terms of RMSE and IS). The results indicated that recruitment peaked at an air temperature of 16.6°C (90% CI: 16.3 to 17.0), while maximum survival occurred at a similar temperature of 15.2°C (90% CI: 15.1 to 15.3).

5 Discussion

The `drmr` package offers an accessible framework for implementing DRMs, addressing the limitations of traditional correlative SDMs and the complexity barrier of DRMs. It provides a user-friendly interface for building, fitting, evaluating, and projecting age-structured DRMs, allowing researchers to consider movement and to explicitly link environmental drivers to demographic processes, like recruitment and survival, that drive shifts in species geographic distributions. Our fish and bird case studies demonstrated `drmr`'s utility in comparing dif-

ferent model structures. These examples highlight the package's value for generating more robust ecological forecasts of species on the move. Key package features include functions for fitting models, generating predictions that account for parameter uncertainty, and visualizing estimated process-environment relationships.

`drmr` distinguishes itself within ecological modeling software by focusing on environmentally-driven and spatially explicit population demography inferred from spatio-temporal observations. This contrasts with other tools addressing related questions. `NicheMapR` excels at mechanistic modeling of individual organismal performance based on microclimate and metabolic theory, and it is parameterized extensively based on prior knowledge (Kearney et al., 2021). Virtual species platforms like `RangeShifter` (Bocedi et al., 2021; Malchow et al., 2021) and `rangr` (Markowska et al., 2025) simulate the spatial, demographic, and/or evolutionary dynamics of an imaginary species across space based on user-specified assumptions about environmental responses and other parameters. By estimating the parameters of mechanistic, spatial population dynamics models from species observations, `drmr` provides an important tool for species that lack the detailed prior knowledge that could otherwise parameterize a virtual species model (while recognizing that virtual species models can be calibrated against data in some cases, e.g., Malchow et al. (2024) and `poems`). In this way, `drmr` combines the ease of fitting SDMs with the mechanistic detail of virtual species models, bridging an important gap between the two approaches. By allowing alternative models to be compared against data, `drmr` also opens the door to testing mechanistic hypotheses about the processes and environmental factors driving range shifts.

The applications of `drmr` are broad, extending to any species where spatio-temporal abundance or density data are available and where understanding the demographic mechanisms of environmentally driven geographic distribution changes are important. Given the large fraction of Earth's biodiversity now shifting to new locations, we expect the application opportunities are many (Pecl et al., 2017). Birds, trees, mammals, insects, and fishes are among the taxonomic groups with particularly extensive observations (Sullivan et al., 2009;

Sauer et al., 2003; Conn et al., 2015; Maureaud et al., 2024; Serra-Diaz et al., 2016). Beyond fundamental knowledge about ecological responses to environmental change, improved understanding and projection of range shifts can provide insights relevant for spatial conservation planning, natural resource management, evaluation of alternative management strategies, species reintroductions, governance planning for transboundary range shifts (Pinsky et al., 2018), and other areas where population-level viability across landscapes is important (Urban et al., 2016).

Despite its utility, `drmr` also has limitations. Currently, it models continuous densities rather than discrete counts, mainly because `Stan` does not support latent discrete variables previously used in this class of models (e.g., Pagel and Schurr, 2012). The implemented movement routine is simplified, assuming diffusion among neighbors, although the package structure allows for future expansion. In addition, survival rates are modeled as constant across age classes unless external age-specific data (e.g., harvest rate-at-age) are provided. The current implementation of temporal random effects is limited to an AR(1) process. Furthermore, increasing model complexity, such as adding movement or more age classes, naturally increases computational time.

In future developments, we aim to enhance `drmr`'s flexibility and realism. Planned extensions include incorporating more sophisticated movement routines that vary with time, patch, age, or environmental conditions (e.g., Thorson et al., 2021); allowing more flexible process-environment relationships using basis functions (as in generalized additive models); and adding alternative population dynamics like negative density-dependent recruitment via Ricker or Beverton-Holt stock-recruit functions (Hilborn and Walters, 2013). Providing support for age- or length-composition data and exploring options for handling count data directly are also key goals. These enhancements will continue to broaden the scope for applying mechanistic DRMs across diverse ecological systems.

Code and Data Availability

The `drmr` package is open-source and available under the GPL-3 license. The source code, documentation, and all data used in the case studies will be released on GitHub upon acceptance.

References

Besag, J., J. York, and A. Mollié (1991). Bayesian image restoration, with two applications in spatial statistics. *Annals of the Institute of Statistical Mathematics 43*(1), 1–20.

Bocedi, G., S. C. F. Palmer, A.-K. Malchow, D. Zurell, K. Watts, and J. M. J. Travis (2021). Rangeshifter 2.0: an extended and enhanced platform for modelling spatial eco-evolutionary dynamics and species' responses to environmental changes. *Ecography 44*(10), 1453–1462.

Brun, P., N. E. Zimmermann, C. Hari, L. Pellissier, and D. N. Karger (2022). CHELSA-BIOCLIM+ a novel set of global climate-related predictors at kilometre-resolution.

Conn, P. B., D. S. Johnson, J. M. V. Hoef, M. B. Hooten, J. M. London, and P. L. Boveng (2015). Using spatiotemporal statistical models to estimate animal abundance and infer ecological dynamics from survey counts. *Ecological Monographs 85*(2), 235–252.

Edwards, B. P. M., A. C. Smith, and S. LaZerte (2024). bbsBayes2: Hierarchical Bayesian analysis of North American BBS data. version 1.1.2.1.

Engler, R. and A. Guisan (2009). MigClim: Predicting plant distribution and dispersal in a changing climate. *Diversity and distributions 15*(4), 590–601.

Evans, M. R. and A. Moustakas (2016). A comparison between data requirements and availability for calibrating predictive ecological models for lowland uk woodlands: learning new tricks from old trees. *Ecology and Evolution 6*(14), 4812–4822.

Fallert, S., L. Li, and J. S. Cabral (2025). metaRange: A framework to build mechanistic range models. *Methods in Ecology and Evolution 16*(1), 49–56.

Fordham, D. A., S. Haythorne, S. C. Brown, J. C. Buettel, and B. W. Brook (2021). poems: R package for simulating species' range dynamics using pattern-oriented validation. *Methods in Ecology and Evolution 12*(12), 2364–2371.

Forrest, D., M. Stuart, R. Batt, and M. Pinsky (2020, 06). pinskylab/oceanadapt: Update 2020.

Fredston, A., D. Ovando, L. da Cunha Godoy, J. Kong, B. Muffley, J. T. Thorson, and M. Pinsky (2025). Dynamic range models improve the near-term forecast for a marine species on the move. *Ecology Letters 00*, (under review).

Gabry, J., D. Simpson, A. Vehtari, M. Betancourt, and A. Gelman (2019). Visualization in bayesian workflow. *Journal of the Royal Statistical Society: Series A (Statistics in Society) 182*(2), 389–402.

Gabry, J., R. Češnovar, A. Johnson, and S. Bronder (2024). *cmdstanr: R Interface to CmdStan*. R package version 0.8.1.9000, commit f2e152b88fde5c2cde01ff078d5715b3b6248628.

Gaillard, J.-M. and J.-F. Lemaître (2020). An integrative view of senescence in nature. *Functional Ecology 34*(1), 4–16.

Gneiting, T. and A. E. Raftery (2007). Strictly proper scoring rules, prediction, and estimation. *Journal of the American statistical Association 102*(477), 359–378.

Healy, K., T. H. G. Ezard, O. R. Jones, R. Salguero-Gómez, and Y. M. Buckley (2019, Aug). Animal life history is shaped by the pace of life and the distribution of age-specific mortality and reproduction. *Nature Ecology & Evolution 3*(8), 1217–1224.

Heppell, S. S., H. Caswell, and L. B. Crowder (2000). Life histories and elasticity patterns: perturbation analysis for species with minimal demographic data. *Ecology 81*(3), 654–665.

Hijmans, R. J., S. Phillips, J. Leathwick, and J. Elith (2024). *dismo: Species Distribution Modeling*. R package version 1.3-16.

Hilborn, R. and C. J. Walters (2013). *Quantitative fisheries stock assessment: choice, dynamics and uncertainty*. Springer Science & Business Media.

Homan, M. D. and A. Gelman (2014, jan). The No-U-turn sampler: Adaptively setting path lengths in hamiltonian Monte Carlo. *Journal of Machine Learning Research 15*(1), 1593–1623.

Johnson, K. F., C. Elizabeth, J. T. Thorson, E. Brooks, R. D. Methot, and A. E. Punt (2016). Can autocorrelated recruitment be estimated using integrated assessment models and how does it affect population forecasts? *Fisheries Research 183*, 222–232.

Karger, D. N., O. Conrad, J. Böhner, T. Kawohl, H. Kreft, R. W. Soria-Auza, N. E. Zimmermann, H. P. Linder, and M. Kessler (2017, Sep). Climatologies at high resolution for the earth's land surface areas. *Scientific Data 4*(1), 170122.

Kearney, M. R., N. J. Briscoe, P. D. Mathewson, and W. P. Porter (2021). Nichemapr–an r package for biophysical modelling: the endotherm model. *Ecography 44*(11), 1595–1605.

Kucukelbir, A., D. Tran, R. Ranganath, A. Gelman, and D. M. Blei (2017). Automatic differentiation variational inference. *Journal of machine learning research 18*(14), 1–45.

Lambert, D. (1992). Zero-inflated Poisson regression, with an application to defects in manufacturing. *Technometrics 34*(1), 1–14.

Landau, W. M. (2024). *instantiate: Pre-Compiled 'CmdStan' Models in R Packages*. R package version 0.2.3.

Malchow, A.-K., G. Bocedi, S. C. F. Palmer, J. M. J. Travis, and D. Zurell (2021). Rangeshiftr: an r package for individual-based simulation of spatial eco-evolutionary dynamics and species' responses to environmental changes. *Ecography 44*(10), 1443–1452.

Malchow, A.-K., G. Fandos, U. G. Kormann, M. U. Grüebler, M. Kéry, F. Hartig, and D. Zurell (2024). Fitting individual-based models of spatial population dynamics to long-term monitoring data. *Ecological Applications 34*(4), e2966.

Markowska, K., K. Malinowska, and L. Kuczyński (2025). rangr: An R package for mechanistic, spatially explicit simulation of species range dynamics. *Methods in Ecology and Evolution 16*(3), 468–476.

Maureaud, A. A., J. Palacios-Abrantes, Z. Kitchel, L. Mannocci, M. L. Pinsky, A. Fredston, E. Beukhof, D. L. Forrest, R. Frelat, M. L. D. Palomares, L. Pecuchet, J. T. Thorson, P. D. Van Denderen, and B. Mérigot (2024, January). FISHGLOB_data: an integrated dataset of fish biodiversity sampled with scientific bottom-trawl surveys. *Scientific Data 11*(1), 24.

Monaghan, P., A. Charmantier, D. H. Nussey, and R. E. Ricklefs (2008, 03). The evolutionary ecology of senescence. *Functional Ecology 22*(3), 371–378.

Morris, M., K. Wheeler-Martin, D. Simpson, S. J. Mooney, A. Gelman, and C. DiMaggio (2019). Bayesian hierarchical spatial models: Implementing the Besag York Mollié model in Stan. *Spatial and Spatio-temporal Epidemiology 31*, 100301.

Nenzén, H. K., R. M. Swab, D. A. Keith, and M. B. Araújo (2012). demoniche–an R-package for simulating spatially-explicit population dynamics. *Ecography 35*(7), 577–580.

Pagel, J. and F. M. Schurr (2012). Forecasting species ranges by statistical estimation of ecological niches and spatial population dynamics. *Global Ecology and Biogeography 21*(2), 293–304.

Paradinas, I., J. Illian, and S. Smout (2023, 05). Understanding spatial effects in species distribution models. *PLOS ONE 18*(5), 1–6.

Pebesma, E. (2018). Simple features for R: standardized support for spatial vector data. *The R Journal 10*(1), 439–446.

Pecl, G. T., M. B. Araújo, J. D. Bell, J. Blanchard, T. C. Bonebrake, I.-C. Chen, T. D. Clark, R. K. Colwell, F. Danielsen, B. Evengård, L. Falconi, S. Ferrier, S. Frusher, R. A. Garcia, R. B. Griffis, A. J. Hobday, C. Janion-Scheepers, M. A. Jarzyna, S. Jennings, J. Lenoir, H. I. Linnetved, V. Y. Martin, P. C. McCormack, J. McDonald, N. J. Mitchell, T. Mustonen, J. M. Pandolfi, N. Pettorelli, E. Popova, S. A. Robinson, B. R. Scheffers, J. D. Shaw, C. J. B. Sorte, J. M. Strugnell, J. M. Sunday, M.-N. Tuanmu, A. Vergés, C. Villanueva, T. Wernberg, E. Wapstra, and S. E. Williams (2017). Biodiversity redistribution under climate change: Impacts on ecosystems and human well-being. *Science 355*(6332), eaai9214.

Pinsky, M. L., G. Reygondeau, R. Caddell, J. Palacios-Abrantes, J. Spijkers, and W. W. L. Cheung (2018). Preparing ocean governance for species on the move. *Science 360*(6394), 1189–1191.

R Core Team (2024). *R: A Language and Environment for Statistical Computing*. Vienna, Austria: R Foundation for Statistical Computing. version: 4.4.1.

Sauer, J. R., J. E. Fallon, and R. Johnson (2003). Use of north american breeding bird survey data to estimate population change for bird conservation regions. *The Journal of Wildlife Management 67*(2), 372–389.

Serra-Diaz, J. M., J. Franklin, W. W. Dillon, A. D. Syphard, F. W. Davis, and R. K. Meentemeyer (2016). California forests show early indications of both range shifts and local persistence under climate change. *Global Ecology and Biogeography 25*(2), 164–175.

Shackelford, C. E., R. E. Brown, and R. N. Conner (2000, January). Red-bellied woodpecker (melanerpes carolinus). *The Birds of North America Online*.

Sullivan, B. L., C. L. Wood, M. J. Iliff, R. E. Bonney, D. Fink, and S. Kelling (2009). eBird: A citizen-based bird observation network in the biological sciences. *Biological Conservation 142*(10), 2282–2292.

Thorson, J. T., S. C. Anderson, P. Goddard, and C. N. Rooper (2025). tinyVAST: R package with an expressive interface to specify lagged and simultaneous effects in multivariate spatio-temporal models. *Global Ecology and Biogeography 34*(4), e70035.

Thorson, J. T., S. J. Barbeaux, D. R. Goethel, K. A. Kearney, E. A. Laman, J. K. Nielsen, M. R. Siskey, K. Siwicke, and G. G. Thompson (2021). Estimating fine-scale movement rates and habitat preferences using multiple data sources. *Fish and Fisheries 22*(6), 1359–1376.

Urban, M. C., G. Bocedi, A. P. Hendry, J.-B. Mihoub, G. Pe'er, A. Singer, J. R. Bridle, L. G. Crozier, L. D. Meester, W. Godsoe, A. Gonzalez, J. J. Hellmann, R. D. Holt, A. Huth, K. Johst, C. B. Krug, P. W. Leadley, S. C. F. Palmer, J. H. Pantel, A. Schmitz, P. A. Zollner, and J. M. J. Travis (2016). Improving the forecast for biodiversity under climate change. *Science 353*(6304), aad8466.

Uribe-Rivera, D. E., G. Guillera-Arroita, S. M. Windecker, P. Pliscoff, and B. A. Wintle (2023). The predictive performance of process-explicit range change models remains largely untested. *Ecography 2023*(4), e06048.

Vehtari, A., A. Gelman, and J. Gabry (2017). Practical Bayesian model evaluation using leave-one-out cross-validation and WAIC. *Statistics and computing 27*, 1413–1432.

Vehtari, A., A. Gelman, D. Simpson, B. Carpenter, and P.-C. Bürkner (2021). Rank-normalization, folding, and localization: An improved $\hat{R}$ for assessing convergence of MCMC (with discussion). *Bayesian Analysis 16*(2), 667–718.

Ye, T., V. H. Lachos, X. Wang, and D. K. Dey (2021). Comparisons of zero-augmented

591 continuous regression models from a Bayesian perspective. *Statistics in Medicine 40*(5),
592 1073–1100.

593 Yee, T. W. (2014). Reduced-rank vector generalized linear models with two linear predictors.
594 *Computational Statistics & Data Analysis 71*, 889–902.

595 Zhang, L., B. Carpenter, A. Gelman, and A. Vehtari (2022). Pathfinder: Parallel quasi-
596 newton variational inference. *Journal of Machine Learning Research 23*(306), 1–49.

597 Ziolkowski Jr, D., M. Lutmerding, V. Aponte, and M. Hudson (2024). North American
598 breeding bird survey dataset 1966–2023. *US Geological Survey data release. https://doi.*
599 *org/10.5066/P97WAZE5*.